

Blue Coat[®] Systems ProxySG[™]

Configuration and Management Guide



Contact Information

Blue Coat Systems Inc.
650 Almanor Avenue
Sunnyvale, California 94085

North America (USA) Toll Free: 1.866.362.2628 (866.36.BCOAT)

North America Direct (USA): 1.408.220.2270

Asia Pacific Rim (Japan): 81.3.5425.8492

Europe, Middle East, and Africa (United Kingdom): +44 (0) 1276 854 101

bcs.info@bluecoat.com

support@bluecoat.com

www.bluecoat.com

Copyright© 1999-2005 Blue Coat Systems, Inc. All rights reserved worldwide. No part of this document may be reproduced by any means nor modified, decompiled, disassembled, published or distributed, in whole or in part, or translated to any electronic medium or other means without the written consent of Blue Coat Systems, Inc. The Software may not be modified, reproduced (except to the extent specifically allowed by local law), removed from the product on which it was installed, reverse engineered, decompiled, disassembled, or have its source code extracted. In addition to the above restrictions, the Software, or any part thereof, may not be (i) published, distributed, rented, leased, sold, sublicensed, assigned or otherwise transferred, (ii) used for competitive analysis or used to create derivative works thereof, (iii) used for application development, or translated (iv) used to publish or distribute the results of any benchmark tests run on the Software without the express written permission of Blue Coat Systems, Inc., or (v) removed or obscured of any Blue Coat Systems, Inc. or licensor copyrights, trademarks or other proprietary notices or legends from any portion of the Software or any associated documentation. All right, title and interest in and to the Software and documentation are and shall remain the exclusive property of Blue Coat Systems, Inc. and its licensors. Blue Coat Systems, Inc. specifications and documentation are subject to change without notice. Information contained in this document is believed to be accurate and reliable, however, Blue Coat Systems, Inc. assumes no responsibility for its use. ProxySG™, ProxyAV™, CacheOS™, SGOS™, are trademarks of Blue Coat Systems, Inc. and CacheFlow®, Blue Coat®, Accelerating The Internet®, WinProxy®, AccessNow®, Ositis®, Powering Internet Management®, and The Ultimate Internet Sharing Solution® are registered trademarks of Blue Coat Systems, Inc. All other trademarks contained in this document and in the Software are the property of their respective owners.

BLUE COAT SYSTEMS, INC. DISCLAIMS ALL WARRANTIES, CONDITIONS OR OTHER TERMS, EXPRESS OR IMPLIED, STATUTORY OR OTHERWISE, ON SOFTWARE AND DOCUMENTATION FURNISHED HEREUNDER INCLUDING WITHOUT LIMITATION THE WARRANTIES OF DESIGN, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL BLUE COAT SYSTEMS, INC., ITS SUPPLIERS OR ITS LICENSORS BE LIABLE FOR ANY DAMAGES, WHETHER ARISING IN TORT, CONTRACT OR ANY OTHER LEGAL THEORY EVEN IF BLUE COAT SYSTEMS, INC. HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. The Software and all related technical information, documents and materials are subject to export controls under the U.S. Export Administration Regulations and the export regulations of other countries.

Document Number: 231-02679

Document Revision: Halfdome 004—02/24/2005

Third Party Copyright Notices

Blue Coat Systems, Inc. Security Gateway Operating System (SGOS) version 4 utilizes third party software from various sources. Portions of this software are copyrighted by their respective owners as indicated in the copyright notices below.

The following lists the copyright notices for:

BPF

Copyright (c) 1988, 1989, 1990, 1991, 1992, 1993, 1994, 1995, 1996

The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that: (1) source code distributions retain the above copyright notice and this paragraph in its entirety, (2) distributions including binary code include the above copyright notice and this paragraph in its entirety in the documentation or other materials provided with the distribution, and (3) all advertising materials mentioning features or use of this software display the following acknowledgement:

This product includes software developed by the University of California, Lawrence Berkeley Laboratory and its contributors.

Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

DES

Software DES functions written 12 Dec 1986 by Phil Karn, KA9Q; large sections adapted from the 1977 public-domain program by Jim Gillogly.

EXPAT

Copyright (c) 1998, 1999, 2000 Thai Open Source Software Center Ltd.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Finjan Software

Copyright (c) 2003 Finjan Software, Inc. All rights reserved.

Flowerfire

Copyright (c) 1996-2002 Greg Ferrar

ISODE

ISODE 8.0 NOTICE

Acquisition, use, and distribution of this module and related materials are subject to the restrictions of a license agreement. Consult the Preface in the User's Manual for the full terms of this agreement.

4BSD/ISODE SMP NOTICE

Acquisition, use, and distribution of this module and related materials are subject to the restrictions given in the file SMP-READ-ME.

UNIX is a registered trademark in the US and other countries, licensed exclusively through X/Open Company Ltd.

MD5

RSA Data Security, Inc. MD5 Message-Digest Algorithm

Copyright (c) 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing this software or this function.

License is also granted to make and use derivative works provided that such works are identified as "derived from the RSA Data Security, Inc. MD5 Message-Digest Algorithm" in all material mentioning or referencing the derived work.

RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

THE BEER-WARE LICENSE" (Revision 42):

<phk@FreeBSD.org <mailto:phk@FreeBSD.org>> wrote this file. As long as you retain this notice you can do whatever you want with this stuff. If we meet some day, and you think this stuff is worth it, you can buy me a beer in return. Poul-Henning Kamp

Microsoft Windows Media Streaming

Copyright (c) 2003 Microsoft Corporation. All rights reserved.

OpenLDAP

Copyright (c) 1999-2001 The OpenLDAP Foundation, Redwood City, California, USA. All Rights Reserved. Permission to copy and distribute verbatim copies of this document is granted.

<http://www.openldap.org/software/release/license.html>

The OpenLDAP Public License Version 2.7, 7 September 2001

Redistribution and use of this software and associated documentation ("Software"), with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain copyright statements and notices.
2. Redistributions in binary form must reproduce applicable copyright statements and notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution, and
3. Redistributions must contain a verbatim copy of this document.

The OpenLDAP Foundation may revise this license from time to time. Each revision is distinguished by a version number. You may use this Software under terms of this license revision or under the terms of any subsequent revision of the license.

THIS SOFTWARE IS PROVIDED BY THE OPENLDAP FOUNDATION AND ITS CONTRIBUTORS "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OPENLDAP FOUNDATION, ITS CONTRIBUTORS, OR THE AUTHOR(S) OR OWNER(S) OF THE SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The names of the authors and copyright holders must not be used in advertising or otherwise to promote the sale, use or other dealing in this Software without specific, written prior permission. Title to copyright in this Software shall at all times remain with copyright holders.

OpenLDAP is a registered trademark of the OpenLDAP Foundation.

OpenSSH

Copyright (c) 1995 Tatu Ylonen <ylo@cs.hut.fi>, Espoo, Finland. All rights reserved

This file is part of the OpenSSH software.

The licences which components of this software fall under are as follows. First, we will summarize and say that all components are under a BSD licence, or a licence more free than that.

OpenSSH contains no GPL code.

1) As far as I am concerned, the code I have written for this software can be used freely for any purpose. Any derived versions of this software must be clearly marked as such, and if the derived work is incompatible with the protocol description in the RFC file, it must be called by a name other than "ssh" or "Secure Shell".

[Tatu continues]

However, I am not implying to give any licenses to any patents or copyrights held by third parties, and the software includes parts that are not under my direct control. As far as I know, all included source code is used in accordance with the relevant license agreements and can be used freely for any purpose (the GNU license being the most restrictive); see below for details.

[However, none of that term is relevant at this point in time. All of these restrictively licenced software components which he talks about have been removed from OpenSSH, i.e.,

- RSA is no longer included, found in the OpenSSL library
- IDEA is no longer included, its use is deprecated
- DES is now external, in the OpenSSL library
- GMP is no longer used, and instead we call BN code from OpenSSL
- Zlib is now external, in a library
- The make-ssh-known-hosts script is no longer included
- TSS has been removed
- MD5 is now external, in the OpenSSL library
- RC4 support has been replaced with ARC4 support from OpenSSL
- Blowfish is now external, in the OpenSSL library

[The licence continues]

Note that any information and cryptographic algorithms used in this software are publicly available on the Internet and at any major bookstore, scientific library, and patent office worldwide. More information can be found e.g. at "<http://www.cs.hut.fi/crypto>".

The legal status of this program is some combination of all these permissions and restrictions. Use only at your own responsibility. You will be responsible for any legal consequences yourself; I am not making any claims whether possessing or using this is legal or not in your country, and I am not taking any responsibility on your behalf.

NO WARRANTY

BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION. IN NO EVENT UNLESS REQUIRED BY APPLICABLE

LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

2) The 32-bit CRC compensation attack detector in deattack.c was contributed by CORE SDI S.A. under a BSD-style license.

Cryptographic attack detector for ssh - source code

Copyright (c) 1998 CORE SDI S.A., Buenos Aires, Argentina. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that this copyright notice is retained. THIS SOFTWARE IS PROVIDED ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES ARE DISCLAIMED. IN NO EVENT SHALL CORE SDI S.A. BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY OR CONSEQUENTIAL DAMAGES RESULTING FROM THE USE OR MISUSE OF THIS SOFTWARE.

Ariel Futoransky <futo@core-sdi.com> <<http://www.core-sdi.com>>

3) ssh-keygen was contributed by David Mazieres under a BSD-style license.

Copyright 1995, 1996 by David Mazieres <dm@lcs.mit.edu>. Modification and redistribution in source and binary forms is permitted provided that due credit is given to the author and the OpenBSD project by leaving this copyright notice intact.

4) The Rijndael implementation by Vincent Rijmen, Antoon Bosselaers and Paulo Barreto is in the public domain and distributed with the following license:

@version 3.0 (December 2000)

Optimised ANSI C code for the Rijndael cipher (now AES)

@author Vincent Rijmen <vincent.rijmen@esat.kuleuven.ac.be>

@author Antoon Bosselaers <antoon.bosselaers@esat.kuleuven.ac.be>

@author Paulo Barreto <paulo.barreto@terra.com.br>

This code is hereby placed in the public domain.

THIS SOFTWARE IS PROVIDED BY THE AUTHORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHORS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

5) One component of the ssh source code is under a 3-clause BSD license, held by the University of California, since we pulled these parts from original Berkeley code.

Copyright (c) 1983, 1990, 1992, 1993, 1995

The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

6) Remaining components of the software are provided under a standard 2-term BSD licence with the following names as copyright holders:

Markus Friedl
 Theo de Raadt
 Niels Provos
 Dug Song
 Aaron Campbell
 Damien Miller
 Kevin Steves
 Daniel Kouril
 Wesley Griffin
 Per Allansson
 Nils Nordman

Simon Wilkinson

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

OpenSSL

Copyright (c) 1995-1998 Eric Young (ey@cryptsoft.com). All rights reserved.

<http://www.openssl.org/about/>

<http://www.openssl.org/about/>

OpenSSL is based on the excellent SLeay library developed by Eric A. Young <mailto:ey@cryptsoft.com> and Tim J. Hudson <mailto:tjh@cryptsoft.com>.

The OpenSSL toolkit is licensed under a Apache-style license which basically means that you are free to get and use it for commercial and non-commercial purposes.

This package is an SSL implementation written by Eric Young (ey@cryptsoft.com). The implementation was written so as to conform with Netscapes SSL.

This library is free for commercial and non-commercial use as long as the following conditions are adhered to. The following conditions apply to all code found in this distribution, be it the RC4, RSA, lhash, DES, etc., code; not just the SSL code. The SSL documentation included with this distribution is covered by the same copyright terms except that the holder is Tim Hudson (tjh@cryptsoft.com).

Copyright remains Eric Young's, and as such any Copyright notices in the code are not to be removed. If this package is used in a product, Eric Young should be given attribution as the author of the parts of the library used. This can be in the form of a textual message at program startup or in documentation (online or textual) provided with the package.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement: "This product includes cryptographic software written by Eric Young (ey@cryptsoft.com)" The word 'cryptographic' can be left out if the routines from the library being used are not cryptographic related :-).
4. If you include any Windows specific code (or a derivative thereof) from the apps directory (application code) you must include an acknowledgement: "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"

THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The license and distribution terms for any publicly available version or derivative of this code cannot be changed. i.e. this code cannot simply be copied and put under another distribution license [including the GNU Public License.]

Copyright (c) 1998-2002 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"

4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact openssl-core@openssl.org.
5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.
6. Redistributions of any form whatsoever must retain the following acknowledgment: "This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF

SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

PCRE

Copyright (c) 1997-2001 University of Cambridge

University of Cambridge Computing Service, Cambridge, England. Phone: +44 1223 334714.

Written by: Philip Hazel <ph10@cam.ac.uk>

Permission is granted to anyone to use this software for any purpose on any computer system, and to redistribute it freely, subject to the following restrictions:

1. This software is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

2. Regular expression support is provided by the PCRE library package, which is open source software, written by Philip Hazel, and copyright by the University of Cambridge, England.

ftp://ftp.csx.cam.ac.uk/pub/software/programming/pcre/

PHAOS SSLava and SSLavaThin

Copyright (c) 1996-2003 Phaos Technology Corporation. All Rights Reserved.

The software contains commercially valuable proprietary products of Phaos which have been secretly developed by Phaos, the design and development of which have involved expenditure of substantial amounts of money and the use of skilled development experts over substantial periods of time. The software and any portions or copies thereof shall at all times remain the property of Phaos.

PHAOS MAKES NO WARRANTIES, EXPRESS OR IMPLIED, INCLUDING WITHOUT LIMITATION THE IMPLIED WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, REGARDING THE SOFTWARE, OR ITS USE AND OPERATION ALONE OR IN COMBINATION WITH ANY OTHER SOFTWARE.

PHAOS SHALL NOT BE LIABLE TO THE OTHER OR ANY OTHER PERSON CLAIMING DAMAGES AS A RESULT OF THE USE OF ANY PRODUCT OR SOFTWARE FOR ANY DAMAGES WHATSOEVER. IN NO EVENT WILL PHAOS BE LIABLE FOR SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

RealSystem

The RealNetworks® RealProxy™ Server is included under license from RealNetworks, Inc. Copyright 1996-1999, RealNetworks, Inc. All rights reserved.

SNMP

Copyright (C) 1992-2001 by SNMP Research, Incorporated.

This software is furnished under a license and may be used and copied only in accordance with the terms of such license and with the inclusion of the above copyright notice. This software or any other copies thereof may not be provided or otherwise made available to any other person. No title to and ownership of the software is hereby transferred. The information in this software is subject to change without notice and should not be construed as a commitment by SNMP Research, Incorporated.

Restricted Rights Legend:

Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013; subparagraphs (c)(4) and (d) of the Commercial Computer Software-Restricted Rights Clause, FAR 52.227-19; and in similar clauses in the NASA FAR Supplement and other corresponding governmental regulations.

PROPRIETARY NOTICE

This software is an unpublished work subject to a confidentiality agreement and is protected by copyright and trade secret law. Unauthorized copying, redistribution or other use of this work is prohibited. The above notice of copyright on this source code product does not indicate any actual or intended publication of such source code.

STLport

Copyright (c) 1999, 2000 Boris Fomitchev

This material is provided "as is", with absolutely no warranty expressed or implied. Any use is at your own risk.

Permission to use or copy this software for any purpose is hereby granted without fee, provided the above notices are retained on all copies.

Permission to modify the code and to distribute modified code is granted, provided the above notices are retained, and a notice that the code was modified is included with the above copyright notice.

The code has been modified.

Copyright (c) 1994 Hewlett-Packard Company

Copyright (c) 1996-1999 Silicon Graphics Computer Systems, Inc.

Copyright (c) 1997 Moscow Center for SPARC Technology

Permission to use, copy, modify, distribute and sell this software and its documentation for any purpose is hereby granted without fee, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Hewlett-Packard Company makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

Permission to use, copy, modify, distribute and sell this software and its documentation for any purpose is hereby granted without fee, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Silicon Graphics makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

Permission to use, copy, modify, distribute and sell this software and its documentation for any purpose is hereby granted without fee, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation. Moscow Center for SPARC Technology makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

SmartFilter

Copyright (c) 2003 Secure Computing Corporation. All rights reserved.

SurfControl

Copyright (c) 2003 SurfControl, Inc. All rights reserved.

Symantec AntiVirus Scan Engine

Copyright (c) 2003 Symantec Corporation. All rights reserved.

TCPIP

Some of the files in this project were derived from the 4.X BSD (Berkeley Software Distribution) source.

Their copyright header follows:

Copyright (c) 1982, 1986, 1988, 1990, 1993, 1994, 1995

The Regents of the University of California. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

This product includes software developed by the University of California, Berkeley and its contributors.

4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Trend Micro

Copyright (c) 1989-2003 Trend Micro, Inc. All rights reserved.

zlib

Copyright (c) 2003 by the [Open Source Initiative](#)

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Contents

Contact Information

Third Party Copyright Notices

Chapter 1: Introducing the ProxySG

Web Security Solution	21
New Features in this Release	26
Security Services	30
Protocols Supported.....	30
Supported Browsers.....	31
Upgrading and Upgrade Enhancements	31
About the Document Organization	31
Related Blue Coat Documentation.....	33
Document Conventions.....	33

Chapter 2: Licensing

About Licensing.....	35
Licensable Components.....	35
About the Trial Period	37
About License Expiration.....	37
Obtaining a WebPower Account.....	38
Registering the Hardware	38
Installing a License Key File	40
Viewing License Information	43
Updating a License.....	44
Automatically Updating a License	45

Chapter 3: Accessing the ProxySG

Before You Begin: Understanding Modes	47
Accessing the ProxySG	48
Management Console Home Page.....	49
Changing the Logon Parameters.....	50
Configuring the SSH Console.....	55

Chapter 4: Configuring the System

Global Configurations	61
Archive Configuration.....	66
Adapters	71
Software and Hardware Bridges.....	75
Gateways	81
Defining Static Routes	83
Using RIP.....	87

DNS	90
Attack Detection.....	95
Using a Bypass List.....	100
Installing WCCP Settings.....	108
Virtual IP Addresses.....	112
Configuring Failover	113
TCP-IP Configuration.....	117

Chapter 5: Managing Port Services

Section A: Managing Multiple Management Consoles

HTTPS Console (Secure Console).....	122
HTTP Console.....	125
SSH Console.....	126
Telnet Console	127

Section B: Creating and Editing Services

About Service Attributes.....	130
DNS-Proxy	131
Endpoint Mapper Proxy	133
FTP	135
HTTP	136
HTTPS.....	138
Instant Messaging Protocols.....	140
Streaming Protocols.....	141
SOCKS	142
TCP Tunneling.....	143
Telnet Shell Proxy Service.....	145

Chapter 6: Configuring Proxies

About Explicit and Transparent Proxy	149
--	-----

Section A: Configuring Explicit Proxies

Creating an Explicit Proxy Server.....	151
Configuring the FTP Proxy.....	152
Configuring FTP Connection Welcome Banners.....	158
HTTP Proxy	159
HTTP Terms.....	161
Configuring Refresh Bandwidth for the HTTP Proxy	162
Setting Default HTTP Proxy Policy	164
Choosing the HTTP Proxy Profile	168
Additional Configuration Affecting Bandwidth Gain	175
Viewing HTTP Settings through the CLI	177
HTTP Compression	178
Troubleshooting HTTP Proxy Issues	185
Configuring a SOCKS Proxy	188
Shell Proxies.....	193

Section B: Transparent Proxies	
Configuring the Transparent Proxy Hardware	199
IP Forwarding.....	201
Creating a Transparent Proxy Service	201
Chapter 7: Using Secure Services	
HTTPS Termination Overview	203
Configuring HTTPS Termination	207
Managing the SSL Client.....	231
Configuring HTTP or HTTPS Origination to the Origin Content Server.....	236
Configuring DNS Resolution to the Origin Content Server	239
Chapter 8: Security and Authentication	
Section A: Controlling Access to the ProxySG	
Limiting Access to the ProxySG Appliance	243
About Password Security.....	244
Limiting User Access to the ProxySG—Overview.....	245
Moderate Security: Restricting Management Console Access Through the Console Access Control List (ACL).....	247
Maximum Security: Administrative Authentication and Authorization Policy	249
Section B: Controlling Access to the Internet and Intranet	
Using Authentication and Proxies.....	255
Using SSL with Authentication and Authorization Services	260
Creating a Proxy Layer to Manage Proxy Operations.....	261
Chapter 9: Using Authentication Services	
Understanding Realms.....	271
SSL Between the ProxySG and the Authentication Server	271
Section A: NTLM Realm Authentication and Authorization	
How Blue Coat Works with NTLM.....	273
Creating an NTLM Realm	273
NTLM Servers.....	274
Defining NTLM Realm General Properties.....	276
Creating the CPL.....	278
Tips and Boundary Conditions.....	278
Section B: LDAP Realm Authentication and Authorization	
Overview	279
Creating an LDAP Realm	280
LDAP Servers	281
Defining LDAP Base Distinguished Names	285
LDAP Search & Groups Tab (Authorization and Group Information)	288
Customizing LDAP Objectclass Attribute Values.....	291
Defining LDAP General Realm Properties.....	292
Creating the CPL.....	294

Section C: RADIUS Realm Authentication and Authorization

Creating a RADIUS Realm..... 295
 Defining RADIUS Realm Properties 296
 Defining RADIUS Realm General Properties 298
 Creating the CPL..... 301

Section D: Local Realm Authentication and Authorization

Creating a Local Realm 302
 Changing Local Realm Properties 303
 Defining the Local User List 305
 Creating the CPL..... 312

Section E: Certificate Realm Authentication

How Certificate Realm Works 313
 Creating a Certificate Realm..... 313
 Defining a Certificate Realm 315
 Defining Certificate Realm General Properties 316
 Revoking User Certificates 318
 Creating the Certificate Authorization Policy 318
 Tips..... 319

Section F: Netegrity SiteMinder

Understanding SiteMinder Interaction with Blue Coat 321
 Participating in a Single Sign-On (SSO) Scheme 323
 Creating a SiteMinder Realm 324
 SiteMinder Servers..... 327
 Defining SiteMinder Server General Properties..... 330
 Creating the CPL..... 335

Section G: Oblix COREid

Understanding COREid Interaction with Blue Coat 336
 Configuring the COREid Access System..... 336
 Additional COREid Configuration Notes 337
 Configuring the ProxySG Realm 337
 Participating in a Single Sign-On (SSO) Scheme 338
 Creating a COREid Realm 339
 Agents..... 340
 COREid Access Server..... 342
 COREid General..... 344
 Creating the CPL..... 346

Section H: Policy Substitution Realm

How Policy Substitution Realms Work 347
 Creating a Policy Substitution Realm 349
 Defining a Policy Substitution Realm 350
 Defining Policy Substitution Realm General Properties 352
 Tips and Boundary Conditions..... 354
 Creating the Policy Substitution Policy 355

Section I: Sequence Realm Authentication

Adding Realms to a Sequence Realm..... 356
 Creating a Sequence Realm 356
 Adding Realms to a Sequence Realm..... 357
 Defining Sequence Realm General Properties 359

Section J: Forms-Based Authentication

Understanding Authentication Forms..... 361
 Creating and Editing an Authentication Form..... 363
 Setting Storage Options..... 369
 Using CPL with Forms-Based Authentication..... 370
 Tips and Boundary Conditions..... 371

Section K: Managing the Credential Cache

Boundary Conditions 374

Chapter 10: Bandwidth Management

Licensing..... 375
 Bandwidth Management Terms 375
 Bandwidth Management Overview..... 376
 Configuring Bandwidth Allocation..... 381
 Using Policy to Manage Bandwidth..... 387

Chapter 11: External Services

Section A: ICAP

Supported ICAP Servers 400
 ICAP v1.0 Features..... 400
 About Content Scanning 401
 Installing the ICAP Server 403
 Creating an ICAP Service 403
 Deleting an ICAP Service..... 407
 Customizing ICAP Patience Text 408
 Creating ICAP Policy..... 413
 Managing Virus Scanning..... 418
 Access Logging..... 419
 References..... 420

Section B: Websense

Creating a Websense Service..... 421
 Deleting a Websense Service 423

Section C: Service Groups

Creating a Service Group..... 425
 Deleting a Service Group or Group Entry 427
 About Weighted Load Balancing..... 427

Section D: Displaying External Service and Group Information

Chapter 12: Health Checks

About General Health Checks..... 431
 Configuring Service-Specific Health Checks 432
 About Global Forwarding and SOCKS Gateway Health Checks 435
 Configuring Global Health Checks 435
 Pausing or Resuming Global Health Checking 437

Chapter 13: Managing Policy Files

About Policy Files 439
 Creating and Editing Policy Files 442
 Managing the Central Policy File 447
 Viewing Policy Files 449

Chapter 14: The Visual Policy Manager

Section A: About the Visual Policy Manager

System Requirements 455
 Launching the Visual Policy Manager 456
 About the Visual Policy Manager User Interface 457
 About VPM Components..... 460
 The Set Object Dialog 463
 The Add/Edit Object Dialog 465
 Online Help..... 465

Section B: Policy Layer and Rule Object Reference

About the Reference Tables 467
 Administration Authentication Policy Layer Reference 467
 Administration Access Policy Layer Reference 467
 DNS Access Policy Layer Reference 468
 SOCKS Authentication Policy Layer Reference 468
 Web Authentication Policy Layer Reference 469
 Web Access Policy Layer Reference 469
 Web Content Policy Layer Reference 472
 Forwarding Policy Layer Reference 473

Section C: Detailed Object Column Reference

Source Column Object Reference..... 475
 Destination Column Object Reference 487
 Service Column Object Reference..... 493
 Time Column Object Reference 498
 Action Column Object Reference..... 501
 Track Object Column Reference 527
 Comment Object Reference 530
 Using Combined Objects 530
 Centralized Object Viewing and Managing..... 534
 Creating Categories 536
 Restricting DNS Lookups 537
 Restricting Reverse DNS Lookups 538

Setting the Group Log Order.....	538
Section D: Managing Policy Layers and Files	
How Policy Layers, Rules, and Files Interact.....	541
Installing Policies	543
Managing Policy Files	544
Installing VPM-Created Policy Files	545
Viewing the Policy/Created CPL.....	547
Section E: Tutorials	
Tutorial—Creating a Web Authentication Policy	549
Tutorial—Creating a Web Access Policy	554
Chapter 15: Advanced Policy	
Section A: Blocking Pop Up Windows	
About Pop Up Blocking	562
Limitations	562
Recommendations.....	562
Section B: Stripping or Replacing Active Content	
About Active Content.....	564
About Active Content Types.....	564
Section C: Modifying Headers	
Section D: Defining Exceptions	
Built-in Exceptions	568
..... User-Defined Exceptions.....	572
About Exception Definitions	572
About the Exceptions Hierarchy.....	573
About the Exceptions Installable List.....	574
Creating or Editing Exceptions	576
Viewing Exceptions	579
Section E: Managing Peer-to-Peer Services	
About Peer-to-Peer Communications	580
The Blue Coat Solution.....	580
Policy Control	581
Proxy Authentication	582
Access Logging.....	582
Chapter 16: Streaming Media	
Section A: About Streaming Media	
Streaming Media Overview.....	584
Streaming Media Protocols.....	585
Streaming Media Player Support	587
Streaming Media Authentication	588
Streaming Media Caching Behavior.....	590

Section B: Configuring Streaming Media

Limiting Bandwidth 593

Configuring the Refresh Rate 597

Configuring HTTP Handoff 598

Forwarding Client Logs to the Media Server 599

Configuring Media Server Authentication Type (Windows Media) 600

About Multicast Streaming 600

Managing Multicast Streaming for Windows Media 602

Managing Multicast Streaming for Real Media 605

Managing Simulated Live Content (Windows Media) 606

ASX Rewriting (Windows Media) 608

About Fast Streaming (Windows Media) 612

Section C: Windows Media Player

Configuring Windows Media Player 613

Limitations 614

Windows Media Access Log Formats 615

Section D: RealPlayer

Configuring RealPlayer 616

Real Media Access Log Formats 618

Limitations and Known Issues 618

Section E: QuickTime Player

Configuring QuickTime Player 619

QuickTime Access Log Formats 619

Limitations 619

Access Log Format 620

Chapter 17: Instant Messaging

About Securing Instant Messaging 621

Recommended Deployments 621

About the Instant Messaging Protocol Services 621

About HTTP Proxy Support 622

About Instant Messaging Reflection 622

IM Reflection Diagrams 622

About Instant Messaging Proxy Authentication 626

Securing AOL Encryption Capability 626

Instant Message Proxies 627

Configuring Instant Messenger Clients 630

VPM Examples 633

Statistics 634

Related Material 634

Chapter 18: Content Filtering

Overview 635

Selecting Category Providers 636

Configuring a Local Database 640

Configuring Blue Coat Web Filter	645
Configuring InterSafe	653
Configuring Optenet	658
Configuring Proventia Web Filter	663
Configuring SmartFilter	667
Configuring SurfControl.....	677
Configuring Websense	681
Configuring Webwasher URL Filter	688
How to Apply Policy to Categorized URLs	692
Using Content-Filtering Vendors with ProxySG Policies	695
Tips.....	698
Chapter 19: Configuring the Upstream Networking Environment	
Forwarding Configuration	701
SOCKS Gateway Configuration.....	722
Internet Caching Protocol (ICP) Configuration.....	730
Using Policy to Manage Forwarding	736
Chapter 20: Access Logging	
Overview	743
Enabling or Disabling Access Logging	746
Customizing the Log Procedures	746
Editing Upload Clients.....	765
Testing Access Log Uploading.....	780
Viewing Access-Log Statistics.....	781
Using Access Logging with Policy Rules	781
Chapter 21: Maintaining the ProxySG	
Restarting the ProxySG	785
Restoring System Defaults.....	787
Purging the DNS Cache	790
Clearing the System Cache	790
Upgrading the ProxySG.....	791
Managing ProxySG Systems	794
Event Logging and Notification.....	797
Configuring SNMP	804
Disk Reinitialization	807
Deleting Objects from the ProxySG.....	808
Chapter 22: Statistics	
Selecting the Graph Scale	811
General Statistics	811
System Usage Statistics	816
HTTP/FTP History Statistics	821
IM History Statistics	826
P2P History Statistics.....	829

Streaming History Statistics	832
SOCKS History Statistics.....	837
Shell History Statistics	840
Resources Statistics	840
Efficiency Statistics.....	843
Contents Statistics	847
Event Logging.....	848
Bandwidth Management Statistics.....	849
Access-Log Statistics.....	853
Failover Statistics.....	857
Advanced Statistics.....	858

Appendix A: Using the Authentication/Authorization Agent

Installing the BCAA Service on a Windows or Windows NT System.....	862
Installing the BCAA Service on a Solaris System.....	869
Troubleshooting Authentication Agent Problems	869
Common BCAA Event Messages	870

Appendix B: Access Log Formats

Custom or W3C ELFF Format.....	877
SQUID-Compatible Format.....	880
NCSA Common Access Log Format	881
Fields Available for Creating Access Log Formats	882

Appendix C: Using WCCP

Overview	911
Quick Start.....	913
Configuring a WCCP Version 2 Service on the Router	914
Creating a ProxySG WCCP Configuration File	921
Examples	928
Troubleshooting: Home Router	932
Tips.....	935

Appendix D: RIP Commands

net	937
host.....	937
RIP Parameters	938
ProxySG-Specific RIP Parameters.....	939
Using Passwords with RIP	939

Appendix E: Diagnostics

Diagnostic Reporting (Service Information)	942
Packet Capturing (the PCAP Utility)	951
Core Image Restart Options	957

Diagnostic Reporting (Heartbeats) 958
 Diagnostic Reporting (CPU Monitoring)..... 959

Appendix F: Using Blue Coat Director to Manage Multiple Appliances

How Director Works with ProxySG 963
 Backing Up a ProxySG's SSL Settings..... 967
 Creating Profiles..... 967
 Creating Overlays 968
 Director Documentation 968

Index

Chapter 1: *Introducing the ProxySG*

Blue Coat® Systems ProxySG™ Appliance represents the latest in perimeter defense for securing and controlling Web-based content and applications. The Blue Coat ProxySG is designed to integrate protection and control functions for Internet and intranet traffic without sacrificing performance and employee productivity.

The ProxySG series of proxy appliances is designed specifically to manage and control user communication over the Internet. Acting on behalf of the user and the application, the ProxySG does not replace existing perimeter security devices; rather, it complements them by giving organizations the ability to control communications in a number of ways that firewalls and other externally focused devices cannot.

Web Security Solution

The Blue Coat ProxySG provides a point of integration, control, and acceleration for enterprise Web security applications, including:

- Layered security approach with content-level protection to combat Web-based threats using port 80.
- Abundant policy controls wrapped in performance-based hardware and a custom operating system to give organizations visibility and control over employee Web communications.
- A preventative spyware defense that combines multiple techniques in a high-performance solution acceptable for web-based business communications.
- Integrated reverse proxy caching and SSL support to offload content delivery and encryption tasks from Web servers, reducing server bottlenecks and enhancing Web site performance and scalability.
- Control over which users are allowed to use Instant Messaging, and which IM protocols are allowed, what features are to be enabled, to whom users may IM or chat with (inside the company or outside the company), what time of the day they can IM, and how logging is managed.
- Immediate and dynamic Peer-to-Peer (P2P) control, allowing an administrator to identify, log, and block P2P traffic.
- Integrated caching, content positioning, bandwidth savings, and bandwidth management to provide superior performance for controlling Web content.
- Control over Windows Media, RealTime, and QuickTime video and audio streams as the file is being downloaded over the Internet.
- Prevention of the spread of viruses and other malicious code by using the Blue Coat ProxyAV™ Appliance in conjunction with the Blue Coat ProxySG. The ProxySG with ProxyAV integration is a high-performance Web anti-virus (AV) solution.

- Control over the type of content retrieved by the ProxySG. You can also filter requests made by clients. If you use Blue Coat Web Filter (BCWF), a highly effective content filtering service that quickly learns and adapts to the working set of its users, you can also use a network service that dynamically examines and categorizes Web pages as they are requested.

Ease of Deployment

The ProxySG is specifically designed to increase security and reduce costs associated with central, regional, and branch office Web protection. For example, the SG200 and SG400 platforms easily *drop in* to remote environments where technical support staff is not always available, and features simple installation and remote management.

Other platforms also feature a simple-to-manage system that installs in minutes with little ongoing maintenance. In addition, they also provide configuration restoration that allows system configuration to be archived, including all system settings, filtering and policies; removable, hot-swappable disk drives for true fault tolerance, and are field serviceable and upgradeable.

Policy and Management Architecture

Networking environments have become increasingly complex, with a variety of security and access management issues. Enterprises face challenges in configuring products and ensuring the result supports enterprise policies. *Policies* enhance ProxySG features, such as authentication and virus scanning, allowing you to manage Web access specific to the enterprise's needs.

Blue Coat policies provide:

- Fine-grained control over various aspects of ProxySG behavior.
- Multiple policy decisions for each request.
- Multiple actions triggered by a particular condition.
- Bandwidth limits.
- Authentication awareness, including user and group configuration.
- Flexibility of user-defined conditions and actions.
- Convenience of predefined common actions and transformations.
- Support for multiple authentication realms.
- Configurable policy event logging.
- Built-in debugging.

The ProxySG uses policies and system configuration together to provide the best possible security for your network environment.

Blue Coat's unique architecture allows for scalable decision making. Effectively turning on multiple combinations of granular policy requires a unique level of performance.

Blue Coat's flexible logging features, coupled with integrated authentication and identification capabilities, give organizations the power to monitor Web access for every user in the network at any time, regardless of where they are. Internet access traffic flowing through the ProxySG gives administrators and managers the ability to audit Web traffic as needed.

Content Filtering

As the number of users and the total amount of traffic grows, policy enforcement demands higher performance to provide adequate end-user quality of experience. To satisfy the management level and scalability that enterprise traffic demands, ProxySG Appliances have emerged as a new layer of infrastructure that provide the performance and manageability required for enterprise-wide policy-based content filtering.

SGOS 4.1 offers a dynamic categorization service if you use the Blue Coat Web Filter (BCWF). The BCWF categorization service is an Internet service, available from designated service points with high-bandwidth connections and dedicated hardware. It analyzes data externally so that content (offensive, distasteful, or perhaps even potentially a legal liability) never enters the network.




Figure 1-1: Content Filtering

The ProxySG enforces Internet access policies based on:

- **Content categories (gambling, sex, etc.)**— Besides BCWF, which includes a database and a dynamic categorization service, databases from leading third-party filtering vendors are offered.
- **Content type and protocols (HTTP, FTP, streaming, MIME type, etc.)**— Adds the ability to block certain types of content transported on certain types of protocols.
- **Identity (user, group, network)**—Customize policy based on who the users are regardless of location.
- **Network conditions**—Customize based on real-time conditions.

Content and Virus Scanning

When integrated with a supported Internet Content Adaptation Protocol (ICAP) server such as the Blue Coat ProxyAV appliance, Blue Coat provides content scanning and filtering. ICAP is an evolving protocol that allows an enterprise to dynamically scan and change Web content. *Content scanning* includes actions like sending a given request for content to an ICAP server for virus scanning or malicious mobile code detection.

To eliminate threats to the network and to maintain caching performance, the ProxySG sends objects to the integrated ICAP server for evaluation and saves the scanned objects in its object store. With subsequent content requests, the ProxySG serves the scanned object rather than rescanning the same object for each request.




Figure 1-2: Content and Virus Scanning

The ProxySG blocks viruses from Web content behind and in front of the firewall. Blue Coat architecture is optimized to handle Web requests and responses that require scanning for potentially malicious mobile code and viruses. The ProxySG uses ICAP to vector responses to supported virus scanning servers to deliver unmatched flexibility and performance in scanning Web content.

Spyware

Spyware leverages multiple vectors, making silver bullet defenses using coarse-grained controls useless and unproductive and impeding critical Web-based business communications. No single technique can filter out spyware and adware to defend against the threat.

Blue Coat combines multiple techniques in a high-performance solution acceptable for Web-based business communications. Latency is minimal and the protection layers are comprehensive to stop, block, and scan spyware. With Blue Coat, you can:

- stop spyware installations;
- block spyware Web sites;
- scan for spyware signatures;
- detect desktop spyware and target for cleanup.



Figure 1-3: Preventing Spyware

For information on using the ProxySG and ProxyAV together, refer to the *Blue Coat ProxyAV Configuration and Management Guide*.

Instant Messaging

Instant Message (IM) usage in an enterprise environment creates security concerns because, regardless of how network security is configured, IM connections can be made from any established protocol, such as HTTP or SOCKS, on any open port. Because it is common for coworkers to use IM to communicate, especially in remote offices, classified company information can be exposed outside the network. Viruses and other malicious code can also be introduced to the network from file sharing through IM clients.

The ProxySG serves as an IM proxy, both in transparent and explicit modes. You can control IM 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. You can also log and archive all IM chats.

Using policy, administrators can quickly deploy sophisticated IM usage policies that integrate with existing authentication directories through LDAP, NTLM and RADIUS.

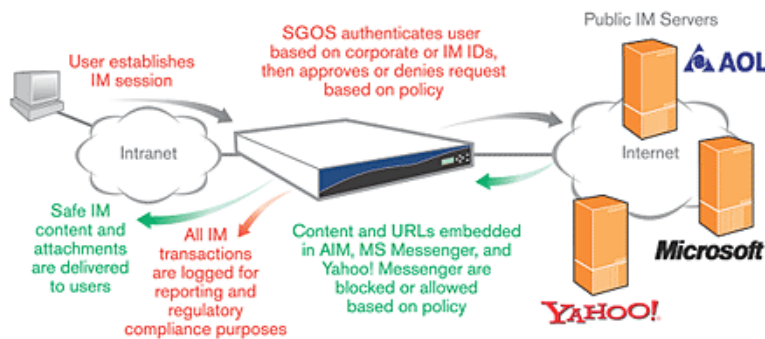


Figure 1-4: Controlling Instant Messaging

Peer-to-Peer

The very nature of the Peer-to-Peer (P2P) client architecture is to evade firewalls and general network security. Additionally, blocking a P2P client at the firewall has proved to be extremely difficult because:

- port blocking, as a means to controlling P2P, is very limited.
- P2P packets cannot be classified simply by looking at packet headers such as an IP address and port number.

Blue Coat ProxySG Appliances provide a powerful platform for immediate and dynamic P2P control.

Integrated Reverse Proxy

ProxySG Appliances are easily configured for reverse proxy mode, providing optimized Web server acceleration and featuring a high RAM-to-disk ratio and a built-in Secure Sockets Layer (SSL) encryption/decryption processor. This processor can manage 10 to 40 times more secure sessions than a standard Web server, allowing the appliances to accelerate the delivery of both public (HTTP) and private (HTTPS) content. The product is packaged in a compact 1U form factor (ProxySG 400 and ProxySG 800 models) a major advantage in space-constrained data centers, or a 4U form factor (ProxySG 8000) that allows for modular expansion of network interface cards, SSL cards, processors, and RAM.

The ProxySG system software is easily tuned for the workload of high traffic Web sites. This environment is characterized by a finite amount of site content accessed by many remote users, often resulting in flash crowds. The ProxySG Appliances allow efficient scaling of Web farms to address flash or peak periods of traffic, and includes advanced features such as protection against Denial-of-Service attacks and dynamic content acceleration.

Bandwidth Management

Bandwidth management allows you to classify, control, and, if required, limit the amount of bandwidth used by different classes of network traffic flowing into or out of the ProxySG. Network resource sharing (or link sharing) is accomplished using a bandwidth-management hierarchy where multiple traffic classes share available bandwidth in a controlled manner.

You can also create policies to constrain who can use certain media types, and how much of it. For example, you can allow your executives to view high-bandwidth streaming media, but only allow the accounting group to view streams up to 56k on corporate sites.

With Blue Coat, you can limit access based on user, group, network address, and the time of day. You can also prevent all access to the Internet except for a group of users who need access to do their jobs, effectively freeing bandwidth for mission-critical needs.

New Features in this Release

Blue Coat has long been the leader in proxy appliances. For SGOS 4.1.x, Blue Coat built upon this leadership by adding:

- New Authentication Realms
- Enhancements to Access Logging
- Bandwidth Management
- CPU Monitoring
- HTTP Object Compression
- SOCKS Compression and Endpoint Mapper proxy
- Content Filtering vendors
- Enhancements to the Blue Coat Patience Page
- New policy to support new SGOS 4.x features

For information on each of these features, continue with the following sections.

New Authentication Realms

In 4.x, two new authentication realms are available, bringing the total to eleven:

- **Oblix COREid:** With Oblix COREid (formerly NetPoint), the ProxySG acts as a custom AccessGate. The ProxySG supports authentication with Oblix COREid v6.5 and v7.0.
- **Policy Substitution:** A Policy Substitution realm provides a mechanism for identifying and authorizing users based on information in the request to the ProxySG. 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.

For more information on these realms, refer to "Section G: Oblix COREid" on 336 and "Section H: Policy Substitution Realm" on 347".

Access Logging

Access Logging has added several new features in SGOS 4.1.x:

- A switch to enable or disable access logging on a global basis, both through the Management Console (Access Logging>General>Global Settings) and the CLI.
- A P2P format and log to support the new P2P functionality.
- Signed access logs that certify that a specific ProxySG generated and uploaded a specific log file.
- New substitutions to support SGOS 4.x functionality. (For more information on new substitutions, refer to the *Blue Coat SGOS 4.x Upgrade Guide*.)

For information on access logging, see Chapter 20: "Access Logging" on page 743.

Bandwidth Management

Bandwidth Management is used to classify, control, and if required, limit the amount of bandwidth used by a class (a unit of bandwidth allocation) of network traffic flowing in or flowing out of the proxy. Network resource sharing (or link sharing) is accomplished in a hierarchical method where multiple traffic classes share the available bandwidth in a controlled manner. The hierarchy specifies how excess bandwidth is shared between the classes.

Note: Bandwidth Management is a licensable feature. It is enabled by default if you have a valid license.

For more information on Bandwidth Gain Management, see Chapter 10: "Bandwidth Management" on page 375.

CPU Monitoring

You can enable CPU monitoring to see the percentage of CPU being used by specific functional groups. CPU monitoring is disabled by default.

You can also view CPU monitoring statistics through Statistics>Advanced>Diagnostics.

For more information, see "Diagnostic Reporting (CPU Monitoring)" on page 959.

Content Filtering

Blue Coat Web Filter (BCWF), formerly Cerberian, is a highly effective content filtering service that quickly learns and adapts to the working set of its users.

You can evaluate BCWF free during the SG trial period (60 days). The auto-evaluation is available to new users of ProxySG as well as those upgrading from SGOS v3.x to SGOS v4.x.

SGOS 4.1 introduces integration with a dynamic categorization service. The BCWF categorization service is an Internet service, available from designated service points with high-bandwidth connections and dedicated hardware.

No username or password is required to use BCWF during the trial period (60 days). If you want to download the database on demand or on a schedule, or if you want to try out dynamic categorization, you'll need to configure the BCWF service.

For more information on BCWF, see "Configuring Blue Coat Web Filter" on page 645.

Also new in this release are three new third-party content filtering vendors—InterSafe, Optenet, and Webwasher. For more information, see "Configuring InterSafe" on page 653, "Configuring Optenet" on page 658, or "Configuring Webwasher URL Filter" on page 688.

HTTP Object Compression

With compression, 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 the ProxySG to send the response as is when the server sends compressed data, including non-cacheable responses. Any unsolicited encoded response is treated as an error.

Note: HTTP object compression is a licensable feature. It is disabled by default, even if you have a valid license.

Whether and where you 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 OCS. However, if CPU is comparatively expensive, configure the ProxySG to ask the OCS for the same compressions that the client asked for and to forward whatever the server returns.

For more information on compression, see "HTTP Compression" on page 178.

SOCKS Compression and Endpoint Mapper Proxy

Compression over SOCKS is supported for TCP/IP tunnels, which can compress the data transferred between the branch (downstream proxy) and main office (upstream proxy), reducing bandwidth consumption and improving latency. For information on enabling SOCKS compression, see "Understanding SOCKS Compression" on page 188.

For SOCKS compression to be successful, you must 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 "Endpoint Mapper Proxy" on page 133.

When configuration is complete, you can set policy to forward TCP traffic through a 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 192.

Patience Page

The ProxySG allows you to customize the patience pages that are displayed when HTTP clients experience delays as Web content is scanned.

In SGOS 4.1.x, patience page behavior has been modified to

- Refresh every five seconds, using Javascript.
- Update the status bar every second with patience page information.
- Manage a popup blocker. If a popup blocker is active, the patience page initiates the download of the scanned object when the root window gets the final patience page. The final patience page also updates the status bar to indicate that the scan is complete.

For information on using patience pages, see Chapter 11: "External Services" on page 399.

Policy

A number of properties (actions) and conditions (source) have been added to support the new features in SGOS 4.1.x. (For a complete list of new CPL and VPM objects, refer to the *Blue Coat SGOS 4.x Upgrade Guide*.)

VPM Object Naming

Objects that can be named by the user no longer start with "_" (underscore character). The underscore character is now used internally to prevent name collisions between objects that can be named by the user and internally generated names.

Exception Pages

A number of built-in exception pages have been added to SGOS 4.1.x to send information back to the user under operational contexts that are known to occur. New exception pages include:

- **HTML Notification**
 - `notify`
 - `notify_missing_cookie`
- **HTTP Compression**
 - `transformation_error`
 - `unensupported_encoding`

Documentation References

- Chapter 14: "The Visual Policy Manager" on page 453
- Chapter 13: "Managing Policy Files" on page 439
- *Blue Coat SGOS 4.x Upgrade Guide*

- *Blue Coat ProxySG Content Policy Language Guide*

Security Services

The Blue Coat ProxySG allows you to control content, instant messaging, and file sharing. In SGOS 4.x, Blue Coat has also added increased support for:

- P2P
- SSL Key Management

For information on each of these features, continue with the following sections.

P2P

The ProxySG recognizes P2P activity relating to P2P file sharing applications. By constructing policy, you can control, block, and log P2P activity and limit the bandwidth consumed by P2P traffic.

For more information, see Chapter 15: "Section E: Managing Peer-to-Peer Services" on 580.

SSL Key Management

In this release, SSL key management has been extended to interact more efficiently with Blue Coat Director.

Director allows you to configure a ProxySG and then push that configuration out to as many ProxySG Appliances as you need. Director also allows you to delegate network and content control to multiple administrators.

For information on using SSL key management with Director, refer to the *Blue Coat Director Configuration and Management Guide*.

Protocols Supported

Blue Coat ProxySGs are multi-protocol. For administrative purposes, you can connect to the Blue Coat ProxySG Appliances through the:

- **HTTPS-Console:** This is the default protocol used by the Management Console. It is configured and enabled by default.
- **SSH-Console:** This is the default protocol for connecting to the ProxySG through the CLI. It is configured and enabled by default.

If you prefer and are in a secure environment, you can use the HTTP-Console or Telnet-Console for administrative access to the system.

Note: HTTP-Console and Telnet-Console are security risks. They should not be used for administrative access in insecure situations.

Supported Browsers

The ProxySG Management Console supports Microsoft® Internet Explorer 6, Netscape® Communicator 7.2, and Firefox 1.0.

The Management Console uses the Java Runtime Environment. Because of security concerns, you should use JRE 1.5.0 (also called J2SE 5.0) if you plan to access external Internet sites.

Upgrading and Upgrade Enhancements

For information on doing upgrades or downgrades, or for restoring default system settings, refer to the *Blue Coat SGOS 4.x Upgrade Guide*.

About the Document Organization

This document is organized for easy reference, and is divided into the following sections and chapters:

Table 1.1: Document Organization

Chapter Title	Description
Chapter 1 – <i>Introducing the ProxySG</i>	This chapter discusses the ProxySG Security Solution and new features and enhancements in SGOS 3.x. It also covers document conventions.
Chapter 2 – <i>Licensing</i>	Several features must be licensed to be used beyond the evaluation trial date. This chapter describes which features require licenses and how to download licenses.
Chapter 3 – <i>Accessing the ProxySG</i>	This chapter explains how to log in to the ProxySG CLI and Web-based Management Console; how to change the administrator username, password, privileged-mode password; and how to make a secure connection using SSH and HTTPS.
Chapter 4 – <i>Configuring the System</i>	Instructions on setting the ProxySG name and system time, configuring the network adapter, load balancing, and FTP port services, and specifying DNS servers. This chapter also describes how to track client IP addresses using server-side transparency or virtual IP addresses.
Chapter 5 – <i>Managing Port Services</i>	This chapter describes port services configurable on the ProxySG, including several kinds of Management Consoles, such as HTTPS, HTTP, SSH, and Telnet Consoles, and application proxies such as Instant Messenger (IM), SOCKS, FTP, MMS, and RTSP, HTTP and HTTPS.
Chapter 6 – <i>Configuring Proxies</i>	Explicit and Transparent proxies are discussed in this chapter, as well as the recommended types of proxy.
Chapter 7 – <i>Using Secure Services</i>	HTTPS termination, including SSL, Certificates, keyrings, and keypairs are discussed in this chapter.

Table 1.1: Document Organization (Continued)

Chapter Title	Description
Chapter 8 – <i>Security and Authentication</i>	Enabling and maintaining security on the ProxySG is discussed in this chapter.
Chapter 9 – <i>Using Authentication Services</i>	Blue Coat supports six kinds of authentication, discussed here: LDAP, NTLM, RADIUS, Local (formerly UNIX), Certificate (which allows you to authenticate using certificates), and Sequence (which allows you to authenticate using multiple authentication servers).
Chapter 10 – <i>Bandwidth Management</i>	Managing the amount of bandwidth used by different classes of network traffic is discussed in this chapter.
Chapter 11 – <i>External Services</i>	ICAP and Websense off-box are described in this chapter.
Chapter 12 – <i>Health Checks</i>	The health of services, such as SOCKS, ICAP, and forwarding services, is discussed in this chapter.
Chapter 13 – <i>Managing Policy Files</i>	Four policy files are used to manage policy: Central, Local, Visual Policy Manager, and Forwarding. This chapter discusses how to manage them.
Chapter 14 – <i>The Visual Policy Manager</i>	This chapter contains a reference guide and several tutorials for using the Visual Policy Manager.
Chapter 15 – <i>Advanced Policy</i>	This chapter discusses using features such as pop-up ad blocking, managing active content, and creating exceptions.
Chapter 16 – <i>Streaming Media</i>	This chapter discusses streaming, including the new RTSP proxy.
Chapter 17 – <i>Instant Messaging</i>	How to configure and use the ProxySG's instant messaging capabilities is discussed in this chapter.
Chapter 18 – <i>Content Filtering</i>	This chapter discusses how to configure and use the ProxySG's content filtering capabilities, as well as configuring and using content filtering vendors to work with the ProxySG.
Chapter 19 – <i>Configuring the Upstream Networking Environment</i>	This chapter discusses how to control upstream interaction with the ProxySG.
Chapter 20 – <i>Access Logging</i>	Log formats, upload clients, upload schedules, and protocols are discussed in this chapter.
Chapter 21 – <i>Maintaining the ProxySG</i>	This chapter discusses upgrading the system and configuring event logs, SMNP, STMP, heartbeats, and core images.
Chapter 22 – <i>Statistics</i>	This chapter discusses viewing various kinds of statistics—system usage, efficiency, resources, and logs of all kinds.

Table 1.1: Document Organization (Continued)

Chapter Title	Description
Appendix A – <i>Using the Authentication/Authorization Agent</i>	The ProxySG BCAA agent is discussed in this appendix.
Appendix B – <i>Access Log Formats</i>	ELFF, SQUID, NCSA/Common, and custom logs are discussed in this appendix.
Appendix C – <i>Using WCCP</i>	Configuring and using a WCCP router with the ProxySG is discussed in this appendix.
Appendix D – <i>RIP Commands</i>	Commands supported for the Routing Information Protocol (RIP) configuration text file are discussed in the appendix.
Appendix E – <i>Diagnostics</i>	Determining and resolving ProxySG problems are discussed in this appendix.
Appendix F – <i>Using Blue Coat Director to Manage Multiple ProxySG Appliances</i>	Discusses how Blue Coat Director works with multiple ProxySG Appliances.

Related Blue Coat Documentation

- *Blue Coat 6000 and 7000 Installation Guide*
- *Blue Coat 200 Series Installation Guide*
- *Blue Coat 400 Series Installation Guide*
- *Blue Coat ProxySG 800 Series Installation Guide*
- *Blue Coat ProxySG 8000 Series Installation Guide*
- *Blue Coat ProxySG Content Policy Language Guide*
- *Blue Coat ProxySG Command Line Reference*

Document Conventions

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

Table 1.2: Typographic Conventions

Conventions	Definition
<i>Italics</i>	The first use of a new or Blue Coat-proprietary term.
Courier font	Command line text that appears on your administrator workstation.
<i>Courier Italics</i>	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 ProxySG literal to be entered as shown.

Table 1.2: Typographic Conventions

{ }	One of the parameters enclosed within the braces must be supplied
[]	An optional parameter or parameters.
	Either the parameter before or after the pipe character can or must be selected, but not both.

Chapter 2: *Licensing*

This chapter describes the ProxySG licensing behavior.

About Licensing

SGOS 4.x features a global licensing system for the ProxySG. License key files are issued on a per-appliance basis. One license key file includes all of the component licenses for whichever ProxySG features you have elected to use.

Note: When your ProxySG order was completed, you received an e-mail that contains serial numbers for licensable components. Those numbers are required for the procedures in this chapter.

Licensable Components

There are two types of licensable components: required and optional. The SGOS X base is required component of the license key file. Optional components license additional feature, and are added to the license key file. The following table lists the ProxySG licensable components.

Table 2.1: Licensable Components

Component	Type	Description
SGOS 4 Base	Required	<p>The ProxySG operating system, plus base features: HTTP, FTP, TCP-Tunnel, SOCKS, and DNS proxy. The following additional features are also included in the base license:</p> <ul style="list-style-type: none"> • 3rd Party Onbox Content Filtering: Allows use with third-party vendor databases, such as SmartFilter, Websense, and SurfControl. • Websense Offbox Content Filtering: For Websense off-box support only. • ICAP Services: External virus and content scanning with ICAP servers. • Bandwidth Management: Allows you to classify, control, and, if required, limit the amount of bandwidth used by different classes of network traffic flowing into or out of the ProxySG. • Windows Media Standard: MMS proxy; no caching or splitting; content pass-through. Full policy control over MMS. • Real Media Standard: RTSP proxy; no caching or splitting; content pass-through. Full policy control over RTSP. • Apple QuickTime Basic: RTSP proxy; no caching or splitting; content pass-through. Full policy control over RTSP. • Netegrity SiteMinder: Allows realm initialization and user authentication to SiteMinder servers. • Oblix COREid: Allows realm initialization and user authentication to COREid servers • Peer-to-Peer: Allows you to recognize and manage peer-to-peer P2P activity relating to P2P file sharing applications.
Compression	Optional	Allows reduction to file sizes without losing any data.
SSL	Optional	SSL Termination; includes an SSL termination card to be installed on the appliance.
IM	Optional	<ul style="list-style-type: none"> • AOL Instant Messaging: AIM proxy with policy support for AOL Instant Messenger. • MSN Instant Messaging: MSN proxy with policy support for MSN Instant Messenger. • Yahoo Instant Messaging: Yahoo proxy with policy support for Yahoo Instant Messenger.
Windows Media Premium	Optional	<ul style="list-style-type: none"> • MMS proxy; content caching and splitting. • Full policy control over MMS. • When the maximum concurrent streams is reached, all further streams are denied and the client receives a message.
Real Media Premium	Optional	<ul style="list-style-type: none"> • RTSP proxy; content caching and splitting. • Full policy control over RTSP. • When the maximum concurrent streams is reached, all further streams are denied and the client receives a message.

About the Trial Period

Blue Coat provides a trial period. The initial system boot-up triggers the 60-day trial period, during which you can evaluate the ProxySG functionality. For the first 60 days, all licensable components are active and available to use. Furthermore, when a license is installed during the trial period (or while using a demo license), components that are *not* part of that license remain available and active during the trial period.

Note: The ProxySG Licensing feature has slight changes from SGOS 3.x. The Blue Coat SGOS 4.x Upgrade Guide (in Chapter 2) describes licensing behavior concerning an upgrade to SGOS 4.x from SGOS 3.x.

Each time you navigate to the Management Console home page or click the Maintenance>Licensing tab, a pop-up dialog appears warning you that the trial period expires in so many days (a text message is displayed on a Telnet, SSH, or serial console). If you require more time to explore the ProxySG features, a demo license is available; refer to your reseller or contact Blue Coat Sales.

The trial period streaming and IM licenses are no-count licenses—unlimited streams and IM clients are accessible.

Upon installing licenses after or during the trial period, the Base SGOS, Instant Messaging (IM), Windows Media basic, and Real Media premium licenses are also unlimited, but Windows Media premium and IM licenses impose user limits established by each license type.

Note: If you invoke the `restore-defaults` command after you have installed licenses, and the serial number of your system is configurable (older boxes only), the licenses fail to install and you return to the trial period (if any time is left).

About License Expiration

At the end of the trial or demo period or, subsequently, when any normally licensed component expires, components that have not been licensed do not process requests. A license expiration notification message is logged in the Event Log (see "Viewing the Event Log" on page 848 for information).

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 ProxySG license has expired.
- Streaming media clients—If the Windows Media Player, RealPlayer, or QuickTime player version supports it, a message is displayed stating the ProxySG license has expired.
- Instant Messaging clients—Users do not receive a message that the ProxySG license has expired. Any IM activity is denied, and to the user it appears that the logon connection has failed.
- FTP clients—If the FTP clients supports it, a message is displayed stating the ProxySG license has expired.

You can still perform ProxySG configuration tasks through the Management Console, CLI, SSH console, serial console, or Telnet connection. Although the component becomes disabled, feature configurations are *not* altered. Also, policy restrictions remain independent of component availability.

Obtaining a WebPower Account

Before you can generate the license key file, you must have a Blue Coat WebPower user account and register the ProxySG.

If you do not have a WebPower account or forgot your account information, perform the following procedure.

To obtain a WebPower account:

1. Select Maintenance>Licensing>Install.
2. In the License Administration field, click Register/Manage. The License Configuration and Management Web page appears (ignore the dialog at this time).
3. Perform one of the following:
 - To obtain a new account, click the link for Need a WebPower User ID. Enter the information as prompted.
 - To obtain your current information for an existing information, click the link for Forgot your password.

Registering the Hardware

This section describes how to enter the appliance serial number and register the appliance with Blue Coat.

System Serial Number Prerequisite

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 ProxySG recognizes the serial number upon system boot-up.

The serial number is visible by navigating to Configuration>General>Identification.

The License Warning Dialog

When you first access the ProxySG Management Console, or when you select Maintenance>Licensing>Install, a License Warning dialog appears.

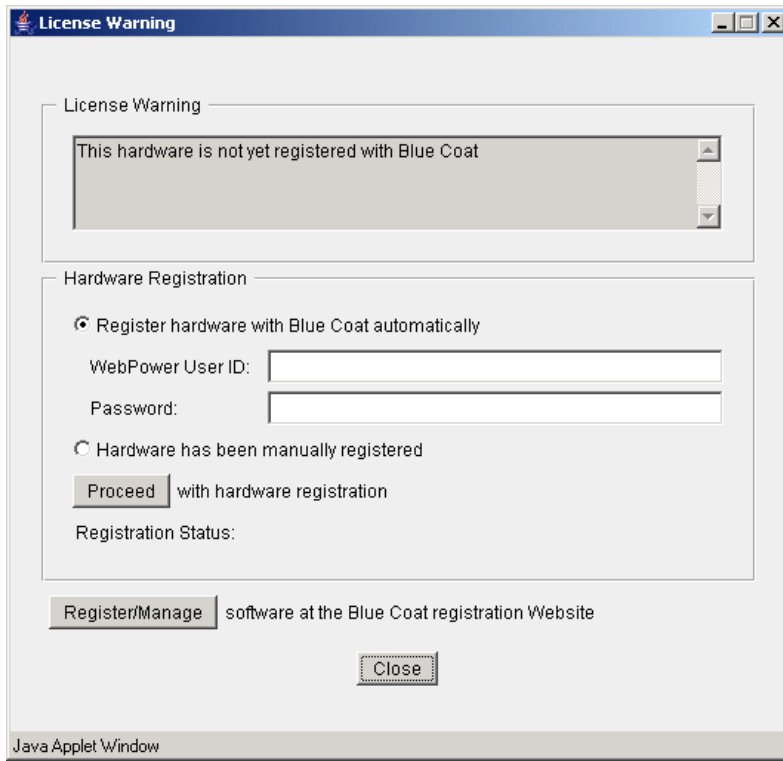


Figure 2-1: License Warning dialog: Hardware not registered.

You cannot install a license key until the hardware has been registered. The License Warning field indicates this status.

If you know the hardware has been manually registered, select Hardware has been manually registered and click Close. The system searches for the last instance and value of hardware registration. Proceed to "Installing a License Key File" on page 40.

Registering the ProxySG

This section describes how to register the ProxySG.

To register the hardware:

1. If the License Warning dialog is not displayed, select Maintenance>Licensing. The License Warning dialog appears.
2. Select Register hardware with Blue Coat automatically.
3. Enter your WebPower username and password.
4. Click Proceed. The Registration Status field displays relative information.

The ProxySG connects to the Blue Coat License Self-Service page. The next step is to obtain and install the license key file that allows access to the ProxySG features you require.

Installing a License Key File

This section describes how to register the ProxySG with Blue Coat and install the license key file.

Creating a License Key File

The License Self-Service Web page allows you to create a license key file that allows the use of the base and additional features for this ProxySG.



Figure 2-2: The License Self-Service Web page.

Upon purchasing the ProxySG from Blue Coat or a reseller, you received an e-mail that contains license serial numbers. These serial numbers are required to create the license key file.

To create a license key file:

1. In the first field under Add a new software solution to this appliance, enter the serial number for the SGOS 4.x base license.
2. In the subsequent fields, enter the serial numbers for any optional licenses you obtained (for example, Compression and IM).




Figure 2-3: Enter license serial numbers.

3. Click Apply.

A license key file, which contains either just the base license or the base combined with optional licenses, is generated and is ready to be downloaded and installed.

Downloading the License Key File

Downloading the license key file is accomplished by using the automatic installation feature or by receiving the key through e-mail and manually installing it from a Web server or a local file.

Automatic License Installation

If the ProxySG has Internet access, you can use the automatic license installation feature to retrieve and install the license from Blue Coat.

To automatically obtain and install the license from the Management Console:

1. Select Maintenance>Licensing>Install.
2. In the License Key Automatic Installation field, click Retrieve. The Request License Key dialog appears.

Figure 2-1: Requesting a License

3. Enter your Blue Coat WebPower user ID and password.
4. Click Send Request.
The ProxySG fetches the license associated with the serial number that is displayed.
5. The Installation Status field displays relevant information. When installation is complete, click Results; examine the results and click OK; click Close. The ProxySG is now licensed.

Manual License Installation

If the ProxySG does not have Internet access, Blue Coat can send you the license in an e-mail. The file can then be installed from a Web server or a local directory.

To Manually Obtain and Install the License:

1. Select Maintenance>Licensing>Install.
2. Click Register/Manage. A new window opens to the Blue Coat ProxySG Registration page. This Web page provides instructions for requesting that the license (associated to the ProxySG by the serial number) be sent through e-mail.
3. When the e-mail arrives, save the attached license file on a Web server or to a local file.
4. In the License Key Manual Installation field, select one of the following from the drop-down list and click Install:

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

- Remote URL—If the file resides on a Web server. The Install License Key dialog displays.

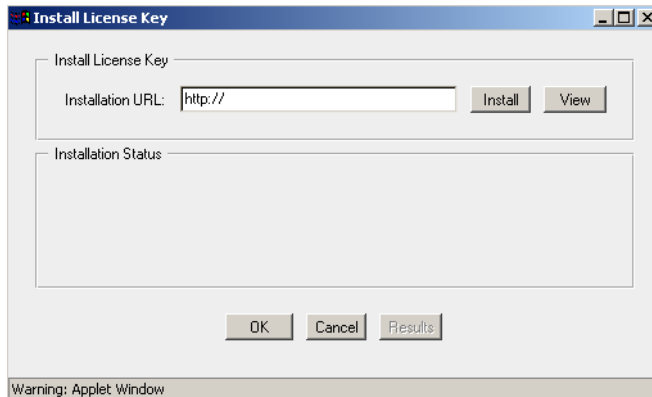


Figure 2-2: Installing a License from a Web Server

Enter the URL path and click Install. The Installation Status field displays relevant information. When installation is complete, click Results; examine the results, close the window, and click OK. Click Apply.

- Local File—If the file resides in a local directory. The Upload and Install File window opens.

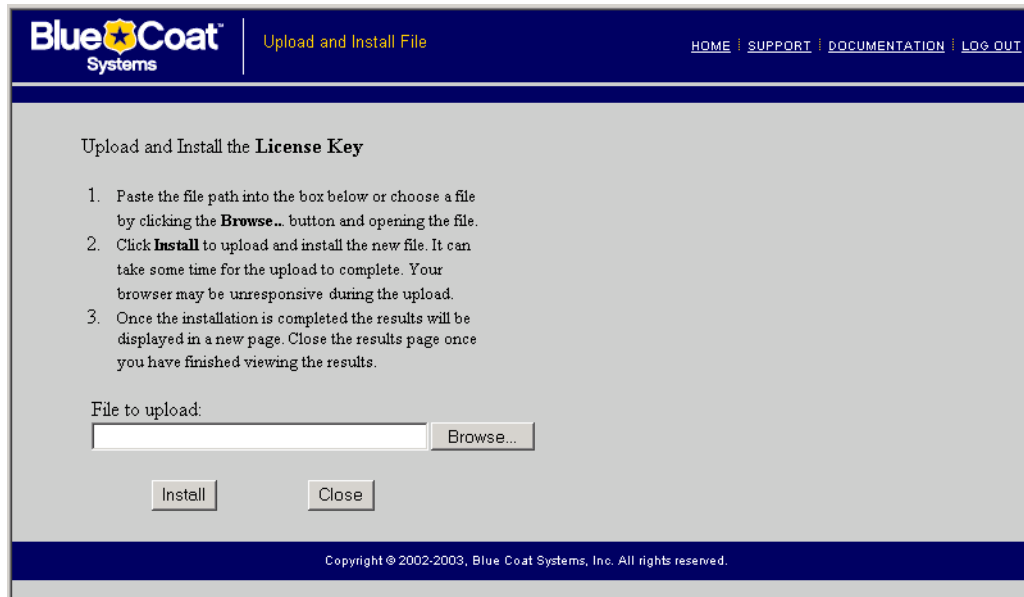


Figure 2-3: Uploading a License from a Local File

Enter a path to the license file or click Browse and navigate to the file. Click Install. A results window opens. Examine the license installation results; close the window. Click Close. Click Apply.

The ProxySG license is now installed. All features that you subscribed to are fully operational.

Viewing License Information

You can review the validity and expiration date of any licensed feature.

To View the License Information through the Management Console:

Select Maintenance>Licensing>View.

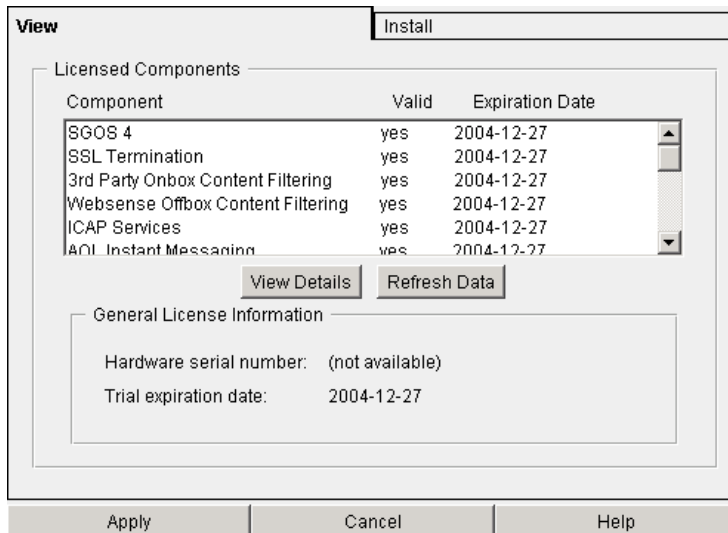


Figure 2-4: Viewing License Information

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

Note: To view the most current information, click Refresh Data.

You can also highlight a license component and click View Details. A dialog appears displaying more detailed information about that component. For example, a streaming component displays the maximum number of streams allowed.

Updating a License

After the initial license installation, you might decide to use another feature that requires a license. For example, you currently support Windows Media, but want to add Real Media support. The license must be updated to allow this support.

To Update a License through the Management Console:

1. Select Maintenance>Licensing>Install.
2. Click Register/Manage.
3. Follow the instructions on the Blue Coat License Self-Service Web page.
4. If using the automatic license installation feature, click Update; otherwise, manually install the license as described in "Manual License Installation" on page 41.

To Update a License through the CLI:

At the enable prompt, enter the following command:

```
SGOS# licensing update-key
```


Automatically Updating a License

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. The ProxySG continues to contact the Web site up to 30 days after the license is set to expire. 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 Feature through the Management Console:

1. Select Maintenance>Licensing>Install.
2. Select Use Auto-Update.
3. Click Apply.

To Configure the License Auto-Update Feature through the CLI:

At the (config) prompt, enter the following command:

```
SGOS# (config) license-key path url
SGOS# (config) license-key auto-update {enable | disable}
```

Note: If the automatic license update fails and you receive a Load from Blue Coat error, you must log on to your License Management account:
https://services.bluecoat.com/eservice_enu/licensing/mgr.cgi. Click Update License Key.

Chapter 3: *Accessing the ProxySG*

The Blue Coat Systems ProxySG uses the Secure Shell (SSH) and HTTPS protocols to securely access the ProxySG CLI and Management Console. Both SSHv1 and SSHv2 are enabled by default, and host keys have already been created on the ProxySG.

All data transmitted between the client and the ProxySG using SSH/HTTPS is encrypted.

During initial configuration, you assigned the ProxySG a username and password and a privileged-mode (enabled/configuration) password. These passwords are always stored and displayed hashed.

This chapter discusses:

- "Before You Begin: Understanding Modes"
- "Accessing the ProxySG"
- "Management Console Home Page"
- "Changing the Logon Parameters"
- "Configuring the SSH Console"

Important: This chapter assumes that you have completed the first-time setup of the ProxySG using either the front panel or serial console, and that the appliance is running on the network. These steps must be completed before accessing the appliance.

You can manage the ProxySG by logging on to and using one of the following:

- An SSH session to access the CLI.
- The Management Console graphical interface.

You can also use a serial console to access the CLI.

Note: To use a Telnet session, you must use a serial console connection until you have configured Telnet for use. (For security reasons Blue Coat does not recommend using Telnet).

Before You Begin: Understanding Modes

SGOS 4.x 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 box. This is the level you enter when you first access the Management Console.

- Configuration is a mode within the enabled mode. From this level, you can perform permanent changes to the ProxySG configuration.

If you use the Management Console, you are in configuration mode when you are completely logged on to the system.

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 *Blue Coat ProxySG Command Line Reference*.

Note: Although most administrator tasks can be performed using either the Management Console or the CLI, there is the occasional task that can only be done using one of the two: these are specified in the manual.

Accessing the ProxySG

You can access the ProxySG through either the CLI or the Management Console. By default, SSHv2 (CLI) and HTTPS (Management Console) are used to connect to the appliance.

The SSH and HTTPS ports are configured and enabled. For SSH, you can use either version 1 or version 2 (with password or RSA client key authentication).

Accessing the CLI

If you use the CLI, you can use SSHv2 to access the ProxySG, but you cannot use SSHv1 or Telnet without additional configuration.

Note: Enabling the Telnet-Console introduces a security risk, so it is not recommended.

To use SSHv1, you must first create an SSHv1 host key. For more information on creating SSH host keys, see "Configuring the SSH Console" on page 55.

To log on 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 (8082 is the default port number)

You must log on from your SSH client.

Accessing 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.

In the Web browser, enter HTTPS, the ProxySG IP address, and port 8082 (the default management port). For example, if the IP address configured during first-time installation is 10.25.36.47, enter the URL `https://10.25.36.47:8082` in the Web browser.

The Management Console consists of a set of Web pages and Java applets stored on the ProxySG. The appliance acts as a Web server on the management port to serve these pages and applets. From the ProxySG home page on the appliance, you can access the management applets, statistics applets, and documentation. The Management Console is supported with a complete online help facility to assist you in defining the various configuration options.

Note: If, when you access the Management Console home page, you get a “host mismatch” or an “invalid certificate” message, you need to recreate the security certificate used by the HTTPS-Console. For information on changing the security certificate, see "HTTPS Console (Secure Console)" on page 122.

Management Console Home Page

When you access the Management Console home page (see "Accessing the Management Console" on page 49), you are prompted to log on to the box.

Logging On

Each time you access the Management Console, you must log on.

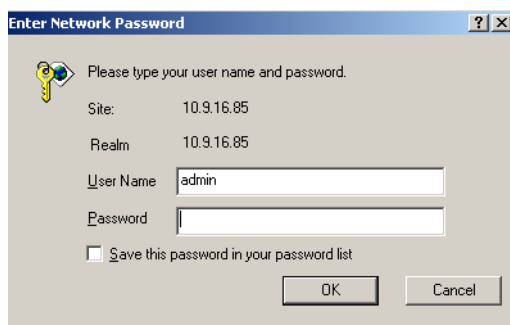


Figure 3-1: Logon Dialog

- The Site is the IP address of the ProxySG to which you are logging on.
- The Realm is a configurable name that can be anything you choose. The ProxySG IP address is the default. For more information on configuring the realm name, see "Changing the ProxySG Realm Name" on page 53.
- The User Name is the name of the account you are using on this ProxySG. The name must already exist. It cannot be created here.
- The Password is the password for the account you are using. It cannot be changed here.

You can change the username and password for the console through the Management Console or the CLI. See "Changing the Logon Parameters" on page 50.

Note: All successful and failed logon attempts are logged to the ProxySG event log.

Logging Out

Once you have logged on, you do not have to log on again unless you exit the current session or the session times out. The session timeout period, with a default of 900 seconds (15 minutes), is configurable.

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

Note: The Keep Working button saves your changes to the current applet. You cannot work in other applets without logging back on to the ProxySG.

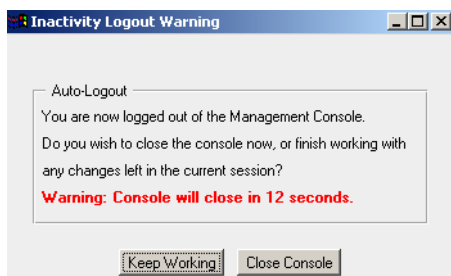


Figure 3-2: Automatic Logout Warning

If you do not click Keep Working or the X in the upper-right-hand corner within the thirty-second period, you are logged out. You must log back on to access the Management Console.

You have logged out. Please close the browser window.

[You need to log in again to use the console](#)

Figure 3-3: Logout Dialog

Click the hyperlink to log back on to the ProxySG.

Note: If no applet is running when the session times out (you are on the Management Console home page), you are logged out without seeing the logout warning dialog. You might not be aware that you are logged out until you try to access an applet. You must enter the logon information again.

Changing the Logon Parameters

You can change the console username and password, the console realm name (which displays when you log on to the ProxySG), and the auto-logout timeout (in seconds; the default 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 ProxySG.

Changing the Username and Password through the Management Console

You can change either the username or the password without changing both.

Changing the Username through the Management Console

The console account username was assigned during initial setup of the system. You can change the username at any time.

To Change the Username through the Management Console:

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

The Console Account tab displays.

Figure 3-4: Console Account Tab

Note: Changing the Console Account username or password causes the Management Console to refresh and log back on using the new information. Note that each parameter must be changed and individually refreshed. You cannot change both parameters at the same time.

2. Enter the username of the administrator or administrator group who 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. The username change was successfully applied, but the configuration could not be fetched from the ProxySG, as the username offered in the fetch request is still the old username.

4. Refresh the screen. You are then challenged for the new username.

To Change the Password through the Management Console:

The console password and privileged-mode password were defined during initial configuration of the system. The console password can be changed at any time through the Management Console. 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.

The Console Account tab displays.

2. Click Change Password.

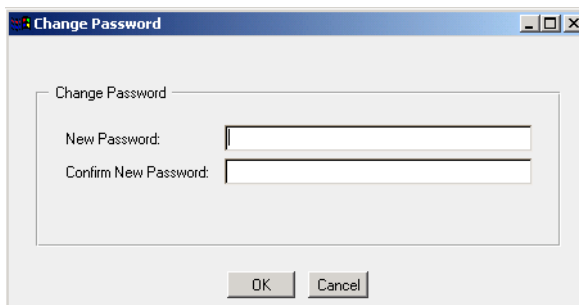


Figure 3-5: Setting or Changing a Password

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 ProxySG 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 "Moderate Security: Restricting Management Console Access Through the Console Access Control List (ACL)" on page 247.

Changing the Username and Password through the CLI

To Change the Console Account Username or Password, Privileged-Mode Password, and the Front-Panel PIN through the CLI:

1. Open a terminal session with the ProxySG and enter the current username and password as prompted.
2. At the command prompt, enter the following command:

```
SGOS> enable
```


3. Enter the privileged-mode password when prompted.
4. At the command prompt, enter the following commands (note that usernames and passwords can each be from 1 to 64 characters in length, but that passwords need to be in quotes):

```
SGOS#configure terminal
SGOS#(config) security username username
```

This command specifies the administrator username.

```
SGOS#(config) security password "password"
-or-
SGOS#(config) security hashed-password hashed_password
```

These commands specify the administrator console password.

```
SGOS#(config) security enable-password "password"
-or-
SGOS#(config) security hashed-enable-password hashed_password
```

These commands specify the administrator privileged-mode password. The ProxySG hashes the password if you enter it in clear text.

5. (Optional, for maximum security. Note that these commands are not available if the ProxySG does not have a front panel.) At the command prompt, change the ProxySG front panel PIN:


```
SGOS#(config) security front-panel-pin pin
-or-
SGOS#(config) security hashed-front-panel-pin hashed-pin
```
6. (Optional) Restrict access by creating an access control list or by creating a policy file containing <Admin> layer rules. For more information, see Section A: “Controlling Access to the ProxySG” on page 243.

Changing the ProxySG Realm Name

The realm name displays when you log on to the ProxySG Management Console. The default realm name is the connection used to access the ProxySG, usually the IP address of the system.

To Change the Realm Name through the Management Console:

1. Select Configuration>Authentication>Console Access>Console Account.
The Console Account tab displays.
2. Enter a new realm name.
The new realm name displays the next time you log on to the ProxySG Management Console.
3. Click Apply.

To Change the Realm Name through the CLI:

1. At the (config) prompt, enter the following command to change the name from the default.


```
SGOS#(config) security management display-realm name
```

 The new realm name displays the next time you log on to the ProxySG Management Console.
2. (Optional) View the results. As the show security command displays lengthy output, only the relevant section is displayed in the following example:

```
SGOS#(config) show security
Account:
  Username:          "admin"
  Hashed Password:  $1$aWmpN$/dsvVrZK6R68KH8r2SQxt/
  Hashed Enable Password: $1$P41pm$ZqFXg4J4A/T.ORGUbr0B/1
  Hashed Front Panel PIN: "$1$GGSf2$1EhLm9oITgny9PDF2kVfp."
  Management console display realm name: ""
  Management console auto-logout timeout: Never
```

You can negate the security management display-realm values by entering `no` before the command; for example, `security management no display-realm`.

Changing the ProxySG Timeout

The timeout is the length of time a session persists before you are logged out. The default timeout is 900 seconds (15 minutes).

To Change the Timeout through the Management Console:

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

The Console Account tab displays.

2. Either deselect Enforce auto-logout (which eliminates auto-logout entirely) or change the auto-logout timeout from its default of 900 seconds (15 minutes) to another time (in seconds). This is the allowable time on the ProxySG before the current session times out. Acceptable values are between 300 and 86400 seconds (5 minutes to 24 hours).

If you change the timeout value, the change takes effect on the next refresh of any applet on the Management Console.

3. Click Apply.

To Change the Timeout through the CLI:

1. To change the timeout from its default of 900 seconds (15 minutes), enter:

```
SGOS#(config) security management auto-logout-timeout seconds
```

The change takes effect on the next refresh of any applet in the Management Console. Acceptable values are between 300 and 86400 seconds (5 minutes to 24 hours).

2. (Optional) View the results. As the `show security` command displays lengthy output, only the relevant section is displayed in the following example:

```
SGOS#(config) show security"
Account:
  Username:          "admin"
  Hashed Password:  $1$a2zTlEE$1b88R3SXUTXS.z07lh8db0
  Hashed Enable Password: $1$xQnqGerX$LU65b20trsIAF6yJox26L.
  Hashed Front Panel PIN: "$1$ThSEiB1v$seyBhSxtTXEtUGDZ5NOB1/"
  Management console display realm name: "Aurora"
  Management console auto-logout timeout: Never
```

You can negate the security management auto-logout-timeout values by entering `no` before the command; for example, `security management no auto-logout-timeout`.

Configuring the SSH Console

By default, the ProxySG uses Secure Shell (SSH) and password authentication so administrators can access the ProxySG CLI or Management Console securely. SSH is a protocol for secure remote logon over an insecure network. No action is required unless you want to change the existing SSH host key, disable a version of SSH, or import RSA host keys. Only one SSH service is allowed on the ProxySG.

To disable the SSH port, see "Managing the SSH Host Connection" below.

Managing the SSH Host Connection

You can manage the SSH host connection through either the Management Console or the CLI.

To Manage the SSH Connection through the Management Console:

Note: Only one SSH Console can be enabled at a time. By default, both SSHv1 and SSHv2 are enabled and assigned to port 22. You do not need to create a new host key unless you want to change the existing configuration.

1. Select Configuration>Services>SSH Console>SSH Host.

The SSH Host tab displays.

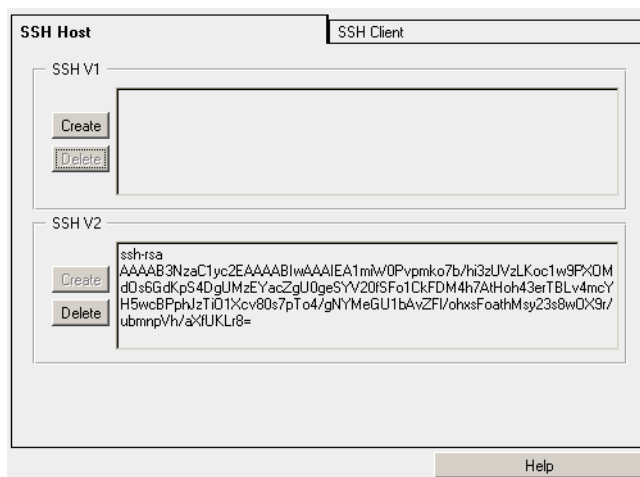


Figure 3-6: SSH Host Tab

2. To delete either SSHv1 or SSHv2 support on the ProxySG, click the appropriate Delete button.

The change is made on the ProxySG instantly.

Important: Do not delete both versions. This disables the SSH Console. Even if you add SSHv1 or SSHv2 client keys back, you will have to enable the service through Configuration>Services>Service Ports.

The SSH host tab redisplay with the appropriate host key deleted.

3. To add SSHv1 or v2 support, select the Create checkbox for the version you want. Remember that if both versions are deleted, you must re-enable the SSH service on port 22.
4. The SSH host key displays in the appropriate pane.

To Manage SSH Host Keys through the CLI:

Note: Only one SSH Console can be enabled at a time. By default, both SSHv1 and SSHv2 are enabled and set up on port 22. You do not need to create a new host key unless you want to change the existing configuration. In fact, you cannot create a new host key unless you delete one of the existing client keys.

You must set up RSA client keys to connect to the ProxySG using RSA. To set up RSA client keys, see "Managing the SSH Client" below.

1. From the `(config)` prompt of the ProxySG, enter the following commands to create a host key.

```
SGOS#(config) services
SGOS#(config services) ssh-console
SGOS#(config services ssh-console) create host-keypair [sshv1 | sshv2]
```

The client key, either SSHv1 or SSHv2 or both, is created, depending on which client key was previously deleted.

2. (Optional) View the results.

```
SGOS#(config services ssh-console) view host-public-key [sshv1 | sshv2]
1024 35
190118975106704546356706163851813093052627858203406609264841510464285480824
068799445880489701889675368436600545643174140823440610328520806007156774811
989754027101280816905716431491183274963949027032871437205903863441301419664
1366408168414061584835486361481236628643756053169543839452802141370496747163
3977037
ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEA2rSeDb3vhr78AFmd7TbdtziYfUQybaDxdMBbSLuyJVgwVbq+
tIvS4L6kDsTuFYGVR8Cg74Xqsj2kO6iwo71YGwdUnDXEzIFBwl0nvS4LkV2UINUwbuP0R0hD4Dt
jVTKsURrOHbTxcXkFipplDwFPDiCKOIqLm4ypcaC/Pj+Juq0=
```

3. To disable SSH, enter:

```
SGOS#(config services ssh-console) delete host-keypair [sshv1 | sshv2]
```

Deleting both of the client keys disables the SSH service on port 22, which then must be re-enabled before you can use SSH Console services again, even if you re-create the host keys.

Managing the SSH Client

You can have multiple RSA client keys on the ProxySG to allow for actions such as logging on to the ProxySG from different locations. You cannot create an RSA client key through the appliance, only through an SSH client. Many SSH clients are commercially available for UNIX and Windows.

Once you have created 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.

Understanding 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
AAAAB3NzaC1yc2EAAAABIwAAAIEAwFI78MKYvL8DrFgcVxpNRHMFkjrBMeBn2PKcv5oAJ2qz+uZ7
hiv7Zn43A6hXwY+DekhtNLOk3HCWmgsrDBE/NOOEnDpLQjBC6t/T3cSQKZjh3NmBbpE4U49rPdu
iiufvWkuoEiHUb5ylzRGdXRSNJHxxmg5LiGEiKaoELJfsDMc= user@machine
```

One of the public key format examples (this one created by the SSH client) is similar to the following:

```
---- BEGIN SSH2 PUBLIC KEY ----
Comment: "[1024-bit rsa, user.name@machine, Wed Feb 19 2003 19:2\8:09]"
AAAAB3NzaC1yc2EAAAADAQABAAQgQCw52JeWr6Fv4kLkzbpZePvapCpaTadPYQwqsGnCIYdf1W
e7/8336EmzV918G1jb/VT1SI1tM1KulBTal7uWai+aUBGKLlYuyhCTo03IZFMnsQC7QYzY1y3ju
fUP3H0be52fg7n7p7gNZR1lyzWhVeilvIKiyVKpjqi6hxCbMb2Q==
---- END SSH2 PUBLIC KEY ----
```

The OpenSSH.pub format appends a space and a user ID to the end of the client key.

The user ID for the key must be unique. Because the ProxySG manages the keys through the user, no two can be the same.

Other caveats:

- 1024 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.

To Import RSA Client Keys through the Management Console:

1. From your SSH client, create a client key and copy it to the clipboard.

Note: The above step must be done with your SSH client. The ProxySG cannot create client keys.

2. Select Configuration>Services>SSH Console>SSH Client.

The SSH Client tab displays.

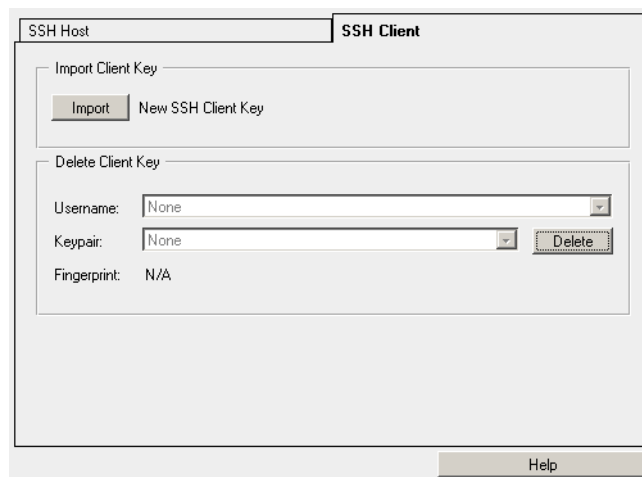


Figure 3-7: SSH Client Tab

3. Click Import to import a new host key.

The Import Client Key dialog displays.

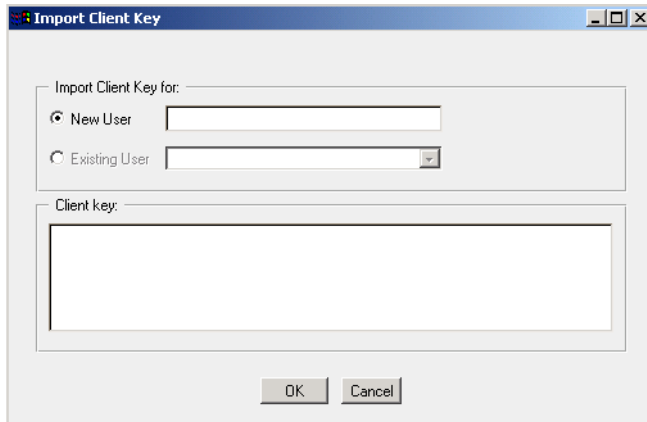


Figure 3-8: Import Client Key Dialog

4. Specify whether the client key is associated with an existing user or a new user, and enter the name.
5. 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.
6. Click OK.

The SSH Client tab reappears, with the fingerprint of the imported key displayed.

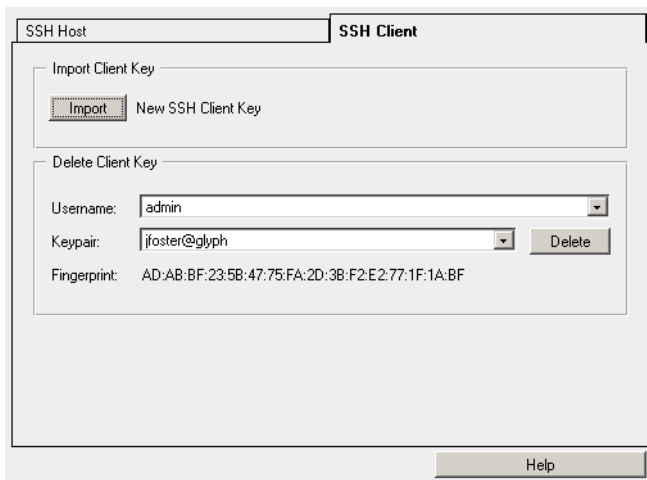


Figure 3-9: SSH Client with Imported Client Key

To Import a Client Key through the CLI:

1. From your SSH client, create a client key and copy it to the clipboard.
2. From the (config) prompt, enter the following commands to import a client key.

```
SGOS#(config) services
SGOS#(config services) ssh-console
SGOS#(config ssh-console) import client-key user_name
Paste client key here, end with "... " (three periods)
ssh-rsaAAAAB3NzaC1yc2EAAAABIwAAAIEAtAy+axsx0iwroFN7B9qSJYjfvbsxPfyC0aoZpSMBd
g97/oiFozDXPhrRmPI3c42EiVdJtVo65r0Aerpu4ybCYVeq6MjRwdsszaezY+VdqtftyYVptC6V1
7Pmj2erw4+A9AggKHTp56BECm3mEPQDdVW7J6QBrJ+U1C1FS/sMcbV8=laptop@GLYPH
...
ok
```

3. (Optional) View the results.

```
SGOS#(config services ssh-console) view client-key username
user_ID@PC 45:5C:3F:5F:EA:65:6E:CF:EE:4A:05:58:9A:C5:FB:4F
user_ID@LAPTOP 61:ED:79:23:F5:2A:1A:6D:84:81:A0:5B:25:36:C7:5F
```

Note: If you have upgraded from an older version ProxySG, and you want to view a previously imported client key, you might not need to enter a username.

Chapter 4: *Configuring the System*

This chapter describes how to configure various ProxySG system configurations, such as setting the time, configuring adapters, and creating software bridges.

This chapter contains the following topics:

- "Global Configurations"
- "Archive Configuration"
- "Adapters"
- "Software and Hardware Bridges"
- "Gateways"
- "Defining Static Routes"
- "Using RIP"
- "DNS"
- "Attack Detection"
- "Installing WCCP Settings"
- "Virtual IP Addresses"
- "Configuring Failover"
- "TCP-IP Configuration"

During initial configuration, Interface 0 was configured by default. The NTP server was defined to keep the system time correct. You also optionally configured a bridge, a gateway, and a DNS server.

These configurations require no further modification. These procedures are provided if you need to configure other adapters in the system or if the need to change a configuration occurs.

Global Configurations

The ProxySG global configurations include: defining the ProxySG name and serial number, setting the time, and configuring NTP for your environment.

Configuring the ProxySG Name

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

To Set the ProxySG Name through the Management Console:

1. Select Configuration>General>Identification.

The Identification tab displays.

The screenshot shows a configuration window titled "Identification". It has two main sections for input. The first section is labeled "Unique name for this ProxySG Appliance:" and contains a text input field with the value "10.9.16.85 - Blue Coat SG3000". The second section is labeled "Serial number of the ProxySG Appliance:" and contains an empty text input field. At the bottom of the window, there are three buttons: "Apply", "Cancel", and "Help".

Figure 4-1: General Identification Tab

2. In the Unique name for this ProxySG Appliance field, enter a ProxySG name.
3. Click Apply.

To Set the ProxySG Name through the CLI:

At the `(config)` command prompt, enter the following command:

```
SGOS#(config) hostname name
```

Configuring the Serial Number

The ProxySG serial number assists Blue Coat Systems Customer Support when analyzing configuration information, including heartbeat reports. This number is found on the ProxySG. Once the serial number is entered, the ProxySG does not verify the validity of the number, only that it is numeric.

Note: If the EPROM contains the ProxySG serial number, you cannot manually enter a serial number.

To Enter the Serial Number through the Management Console:

1. Select Configuration>General>Identification.
The Identification tab displays.
2. In the Serial Number field, enter the serial number.
3. Click Apply.

To Enter the Serial Number through the CLI:

At the `(config)` command prompt, enter the following command:

```
SGOS#(config) serial-number serial_number
```

Displayed Information

The serial number is visible on the Management Console home page, and is displayed using the `show serial-number` command. If the serial number was entered through the Management Console or the CLI, it is appended with `(configured)` to indicate a manual entry.

Configuring the System Time

To manage objects, the ProxySG must know the current Universal Time Coordinates (UTC) time, which is the international time standard and is based on a 24-hour clock.

By default, the ProxySG attempts to connect to an NTP server to acquire the UTC time. The appliance 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. If the appliance cannot access any of the listed NTP servers, you must manually set the UTC time.

To Set UTC Time through the Management Console:

1. Select Configuration>General>Clock>Clock.

The Clock tab displays.

Figure 4-2: General Clock Tab

2. Verify that Enable NTP is selected.
3. To set your local time, select a time zone from the Timezone drop-down list.

Once the local time zone is selected, event logs record the local time instead of GMT.

4. Click Acquire UTC time.
5. Click Apply.

To Set UTC Time through the CLI:

At the enable prompt, enter the following command:

```
SGOS# acquire-utc
```

If NTP is disabled, an error is displayed.

To Manually Set UTC Time through the Management Console:

1. Select Configuration>General>Clock>Clock.

The Clock tab displays.

2. De-select Enable NTP.

The UTC time and date fields become editable when NTP is disabled.

3. To set your local time, select a time zone from the Timezone drop-down list.

Once the local time zone is selected, event and access logs record the local time instead of GMT.

4. Click Pause in the upper-right-hand corner to stop the system clock.

5. Enter the current UTC time and date in the UTC time and date fields.

6. Click Resume to start the system clock.

7. Click Apply.

To Manually Set UTC Time through the CLI:

1. At the (config) command prompt, enter the following commands

```
SGOS# (config) clock day 1-31
SGOS# (config) clock hour 0-23
SGOS# (config) clock minute 0-59
SGOS# (config) clock month 1-12
SGOS# (config) clock second 0-59
SGOS# (config) clock year year
```

2. (Optional) View the results.

```
SGOS# (config) show clock
2003-08-28 22:50:56+00:00UTC
2003-08-28 22:50:56+00:00UTC
```

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 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 Universal Time Coordinates (UTC) to keep the system time accurate.

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

To Add an NTP Server through the Management Console:

1. Select Configuration>General>Clock>NTP.

The NTP tab displays.

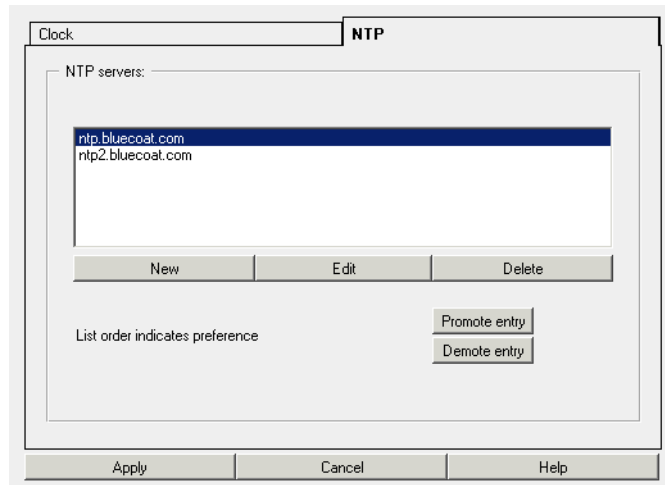


Figure 4-3: General Clock NTP Tab

2. Click New to add a new server to the list.
3. Enter either the domain name or IP address of the NTP server and click OK.
4. Click Apply.

To Add an NTP Server through the CLI:

1. At the (config) command prompt, enter:


```
SGOS#(config) ntp server domain_name
SGOS#(config) ntp interval minutes
SGOS#(config) ntp enable
```

2. (Optional) View the results.

```
SGOS#(config) show ntp
NTP is enabled
NTP servers:
  ntp.bluecoat.com
  ntp2.bluecoat.com
Query NTP server every 60 minutes
```

3. To remove a server from the NTP server list:

```
SGOS#(config) ntp no server domain_name
```

To Change the Access Order through the Management Console:

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.

The NTP tab displays.

2. Select an NTP server to promote or demote.
3. Click Promote entry or Demote entry as appropriate.
4. Click Apply.

Configuring HTTP Timeout

You can configure various network receive timeout settings for HTTP transactions. You can also configure the maximum time that the HTTP proxy waits before reusing a client-side or server-side persistent connection. You must use the CLI to configure these settings.

To Configure the HTTP Receive Timeout Setting through the CLI:

At the `(config)` command prompt, enter the following command:

```
SGOS#(config) http receive-timeout {client | refresh | server} #_seconds
```

where:

<code>client</code>	<code>#_seconds</code>	Sets the receive timeout for client to <code>#_seconds</code> . The default is 120 seconds.
<code>refresh</code>	<code>#_seconds</code>	Sets receive timeout for refresh to <code>#_seconds</code> . The default is 90 seconds.
<code>server</code>	<code>#_seconds</code>	Sets receive timeout for server to <code>#_seconds</code> . The default is 180 seconds.

To Configure the HTTP Persistent Timeout Setting through the CLI:

At the `(config)` command prompt, enter the following command:

```
SGOS#(config) http persistent-timeout {client | server} #_seconds
```

where:

<code>client</code>	<code>#_seconds</code>	The maximum amount of time the HTTP proxy waits before closing the persistent client connection if another request is not made. The default is 360 seconds.
<code>server</code>	<code>#_seconds</code>	The maximum amount of time the HTTP proxy waits before closing the persistent server connection if that connection is not re-used for any subsequent request from the proxy. The default is 900 seconds.

Archive Configuration

Blue Coat allows you to both use an existing configuration (modified to include only general parameters, not system-specific settings) to quickly set up a newly-manufactured ProxySG and to save the running configuration off-box for archival purposes.

To share configurations among systems, continue with the next section; to archive a configuration for later use, skip to "Archiving a Configuration" on page 69.

Sharing Configurations

You can share configuration between two ProxySG Appliances. 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 system.

Note: Blue Coat Director allows you to push configuration from one ProxySG to multiple ProxySG Appliances at the same time. For more information on using Director, see Appendix F: “Using Blue Coat Director to Manage Multiple Appliances” on page 963.

The new configuration 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*.

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 and Push a Configuration to a Newly Manufactured ProxySG through the Management Console:

From the already configured ProxySG:

1. Select Configuration>General>Archive.

The Archive Configuration tab displays.




Figure 4-4: Archive Configuration Tab

2. In the View Current Configuration panel, select the configuration from the drop-down list that you want to use for the newly-manufactured machine:

- ❑ 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.
 - ❑ Configuration - brief: This displays the configuration on the current system, but does not include the installable lists.
 - ❑ Configuration - expanded: This is the most complete snapshot of the system configuration, but it contains system-specific settings that should not be pushed to a new system.
 - ❑ Results of Configuration Load: This displays the results of the last configuration pushed to the system.
3. View the configuration you selected by clicking View. 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:
 - ❑ Save it 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.)

From the newly-manufactured ProxySG:

1. Launch the Management Console in a new browser window.
2. Select Configuration>General>Archive.
3. The Archive Configuration tab displays.
4. In the Install Configuration panel, select either Local File or Text Editor from the drop-down list (depending on whether you saved the file to your system or just copied it to the clipboard) and click Install.
 - ❑ If you saved the file to your system, browse to the location of the Local File, highlight the file, and click Install. The configuration is installed, and the results screen displays.
 - ❑ If you copied the contents of the file, paste it into the Text Editor and click Install. The configuration is installed, and the results screen displays.

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

5. Click Close.

To Create and Push a Configuration to a Newly Manufactured ProxySG through the CLI:

From the already configured ProxySG:

1. From the enable prompt (#), determine which configuration you want to use for the new system. The syntax is:

```
show configuration post-setup | brief | expanded
```

where:

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.
----------------------------	---

Configuration - brief:	This displays the configuration on the current system, but does not include the installable lists.
Configuration - expanded	This is the most complete snapshot of the system configuration, but it contains system-specific settings that should not be pushed to a new system.

```
SGOS# show configuration post-setup
```

The selected configuration displays on the screen.

2. Save the configuration. You can save the file two ways:

- Copy the contents of the configuration to the clipboard. (You will paste the file into the terminal on the newly-manufactured system.)
- Save it as a text file on a download FTP server accessible to the ProxySG. This is advised if you want to re-use the file.

From the newly-manufactured ProxySG, do 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 the FTP server:

At the enable command prompt, enter the following command:

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

where *url* must be in quotes and is fully-qualified (including the protocol, server name or IP address, path, and filename of the configuration file). The configuration file is downloaded from the server, and the ProxySG settings are updated.

Note: If you rename the archived configuration file so that it does not contain any spaces, the quotes surrounding the URL are unnecessary.

The username and password used to connect to the FTP server can be embedded into the URL. 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@
```

Archiving a Configuration

In the rare case of a complete system failure, restoring a ProxySG to its previous state is simplified by loading an archived system configuration from an FTP or TFTP server. The archive, taken from the running configuration, contains all system settings differing from system defaults, along with any installable lists configured on the ProxySG.

Archive and restore operations must be done through the CLI.

Note: You can archive a system configuration to an FTP or TFTP server that allows either anonymous logon or requires a specific username and password. Likewise, to restore a system configuration, the server storing the archive can be configured either to allow anonymous logon or to require a username and password.

Preparing to Archive a System Configuration

1. Obtain write permission to a directory on an FTP server. This is where the archive will be stored.

The system configuration must be stored using FTP.

2. At the `(config)` command prompt, enter the following commands:

```
SGOS#(config) archive-configuration protocol {ftp | tftp}
SGOS#(config) archive-configuration host host_name
```

where *host_name* is the IP address of the server.

Note: TFTP does not require a password, path, or username.

```
SGOS#(config) archive-configuration password password
```

-or-

```
SGOS#(config) archive-configuration encrypted-password encrypted-password
```

where *password* is the password (or encrypted password) used to access the server.

```
SGOS#(config) archive-configuration path path
```

where *path* is the directory on the server where the archive is to be stored, relative to the preset FTP directory.

```
SGOS#(config) archive-configuration filename-prefix filename
```

where *filename* can contain % strings that represent the information in the upload filename. If you do not use the filename command, the ProxySG creates a name with a timestamp and the filename *SG_last-ip-octet_timestamp*. For % string substitutions, see "Fields Available for Creating Access Log Formats" on page 882.

```
SGOS#(config) archive-configuration username username
```

where *user_name* is the username used to access the server.

Example Session

```
SGOS#(config) archive-configuration host 10.25.36.47
ok
SGOS#(config) archive-configuration password access
ok
SGOS#(config) archive-configuration username admin1
ok
SGOS#(config) archive-configuration path ftp://archive.server/stored
ok
SGOS#(config) archive-configuration protocol ftp
ok
```

Note: To clear the host, password, or path, type the above commands using empty double-quotes instead of the variable. For example, to clear the path, enter `archive-configuration path ""`.

To Archive a System Configuration through the CLI:

At the enable command prompt, enter the following command:

```
SGOS# upload configuration
```

To Restore a System Configuration through the CLI:

At the enable command prompt, enter the following command:

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

where *url* must be in quotes and is fully-qualified (including the protocol, server name or IP address, path, and filename of the configuration file). The configuration file is downloaded from the server, and the ProxySG settings are updated.

Note: If you rename the archived configuration file so that it does not contain any spaces, the quotes surrounding the URL are unnecessary.

The username and password used to connect to the FTP server can be embedded into the URL. 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@
```

Adapters

This section describes ProxySG network adapters and the adapter interfaces.

Note: In Blue Coat documentation, the convention for adapters and their interfaces (the connections on the adapter) is Adapter 0, Interface 0, or 0:0.

About Adapters

ProxySG Appliances ship with one or more network adapters installed on the system, each with one or more interfaces. You can change interface parameters or configure additional adapters in the appliance. You can also accept or reject inbound connections, change link settings in the event the system did not correctly determine them, and configure the browser for proxy settings.

Network Interface States

As you select adapters from the picklist, the Adapter panel (Configuration>Network>Adapters) displays the state of the configured adapter and its interfaces. When you initially set up the ProxySG, you optionally configured Adapter 0, Interface 0. If your system has only one adapter, you can skip this section. If your system shipped with other adapters, you can configure them through these procedures.

Configuring an Adapter

The following procedure describes how to configure an adapter. Repeat the process if the system has additional adapters.

To Configure a Network Adapter through the Management Console:

1. Select Configuration>Network>Adapters>Adapters.

The Adapters tab displays.

Note: Different ProxySG models have different adapter configurations, and the appearance of the Adapters tab varies accordingly.




Figure 4-5: Network Adapters Tab

2. Select an adapter from the Adapter drop-down list.

Notice that in the Interfaces field, a message displays stating whether the interface belongs to a bridge. For more information about network bridging, see "Software and Hardware Bridges" on page 75.

3. (Optional) If you have a dual interface adapter, select an interface from the drop-down list.
4. Enter the IP address and subnet mask for the interface into the IP address for interface x and Subnet mask for interface x fields (where interface x refers to the interface selected in the Interfaces drop-down list.)
5. (Optional) To configure link settings, restrict inbound connections, or set up browser proxy behavior for the adapter, select the adapter (under Interfaces) and click Settings. Enter any changes and click OK to close the Settings dialog.

Note: The default is to permit all inbound connections. Link settings are automatically determined and should not need to be modified. The browser default is to use the proxy's default PAC file. (See "About the Settings Button" below for more information on link settings and inbound connections.)

6. Click Apply.

REVIEWERS: Either the command is wrong (bug-worthy) or the documentation is wrong.

To Configure a Network Adapter through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) interface fast-ethernet interface_number
```

where *interface_number* is 0, 1, or *n*, up to one number less than the number of adapters in the system.

```
SGOS#(config interface interface_#) ip-address ip_address
```

```
SGOS#(config interface interface_#) subnet-mask subnet
```

```
SGOS#(config interface interface_#) exit
```

About the Settings Button

The Settings button in the Interfaces field allows you to restrict inbound connections on the selected adapter, and to choose manual or automatic configuration of the adapter link settings.

The default for Inbound connections is to permit all incoming connections. The link settings are automatically determined and should not normally require modification.

Note: Rejecting inbound connections improperly, or manually configuring link settings improperly, can cause the ProxySG to malfunction. Make sure that you know the correct settings before attempting either of these. If the ProxySG fails to operate properly after changing these settings, contact Blue Coat Support.

Rejecting Inbound Connections

The default setting allows inbound connections on all network adapters.

To Reject Inbound Connections through the Management Console:

1. Select Configuration>Network>Adapters>Adapters.

The Adapters tab displays.

2. Select an adapter from the Adapter drop-down list.

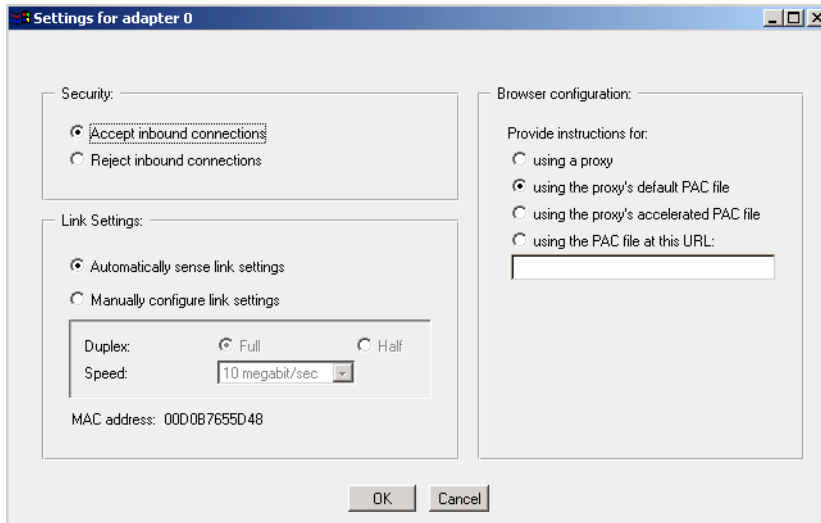


Figure 4-6: Settings for Individual Network Adapters

3. Click Settings.
4. To allow inbound connections, select the Accept inbound connections radio button. To reject inbound connections, select Reject inbound connections.
5. Click OK to close the Settings dialog.
6. Click Apply.

To Reject Inbound Connections through the CLI:

At the (config) command prompt, switch to the interface submode to enter the following commands:

```
SGOS#(config) interface interface_#
SGOS#(config interface interface_#) no accept inbound
SGOS#(config interface interface_#) exit
```

Manually Configuring Link Settings

By default, the ProxySG automatically determines the link settings for all network adapters. If the device incorrectly identifies the network adapter, you can manually configure the link settings.

To Manually Configure Link Settings through the Management Console:

1. Select Configuration>Network>Adapters>Adapters.
The Adapters tab displays.
2. Select an adapter from the Adapters drop-down list.
3. Click Settings.
4. Select Manually configure link settings.
5. Select Half or Full duplex.
6. Select the correct network speed.

7. Click OK to close the Advanced Settings dialog.
8. Click Apply.

To Manually Configure Link Settings through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) interface fast-ethernet interface_#
SGOS#(config interface interface_#) full-duplex | half-duplex
SGOS#(config interface interface_#) speed 10 | 100 | 1gb
SGOS#(config interface interface_#) exit
```

Setting Up Proxies

To set up proxies, see "Configuring Proxies" on page 149.

Detecting Network Adapter Faults

The ProxySG can detect whether the network adapters in an appliance are functioning properly. If the appliance finds that an adapter is faulty, it stops using it. When the fault is remedied, the ProxySG 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.

If an adapter fault is detected, and the adapter has an IP address assigned to it, the ProxySG logs a severe event. When an adapter does not have an IP address, the appliance does not log an entry.

Software and Hardware Bridges

This section describes the ProxySG hardware and software bridging capabilities.

About Bridging

Network bridging through the ProxySG provides transparent proxy pass-through and failover support. This functionality allows ProxySG Appliances 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.
- Hardware—A hardware, or *pass-through*, bridge uses a 10/100 dual interface Ethernet adapter. This type of bridge provides pass-through support.

About the Pass-Through Adapter

A pass-through adapter is a 10/100 dual interface Ethernet adapter designed by Blue Coat to provide an efficient fault-tolerant bridging solution. If this adapter is installed on a ProxySG, SGOS detects the adapter upon 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, Web traffic passes from one Ethernet interface to the other. Therefore, Web traffic is uninterrupted, but does not route through the appliance.

Important: This scenario creates a security vulnerability.

Once power is restored to the ProxySG, the bridge opens and Web traffic is routed to the appliance and thus is subject to that appliance's configured features, defined policies, and content scanning redirection instructions.

Note: Bridging supports only failover; it does not support load balancing.

The following figure provides an example of how the ProxySG indicates that an installed adapter is a pass-through adapter.




Figure 4-7: Pass-through Adapter

Note: The adapter state is displayed on Configuration>Network>Adapters>Adapters.

ProxySG Prerequisites

Before configuring a software bridge, the following conditions must be satisfied:

- **Adapters**—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, for example), the resultant behavior is unpredictable.
- **IP addresses**—If the bridge already has an IP address configured, IP addresses must be removed from any of adapter interfaces to be added. If the bridge does not have an IP address configured, the bridge can inherit the IP address from the first interface to be added.

Setting Bandwidth Management for Bridging

After you have created and configured a bandwidth management class for bridging (see Chapter 10: “Bandwidth Management” on page 375), you can manage the bandwidth used by all bridges.

Note: Before you can manage the bandwidth for bridging, you must first enable bandwidth management and create a bandwidth-management class configured for bridging. Bandwidth management is enabled by default if you have a valid license for this feature. See Chapter 10: “Bandwidth Management” on page 375 for information about enabling bandwidth management and creating and configuring the bandwidth class.

To Configure Bandwidth Management for Bridging through the Management Console:

1. Select Configuration>Network>Adapters>Bridges.

The Bridges tab displays.

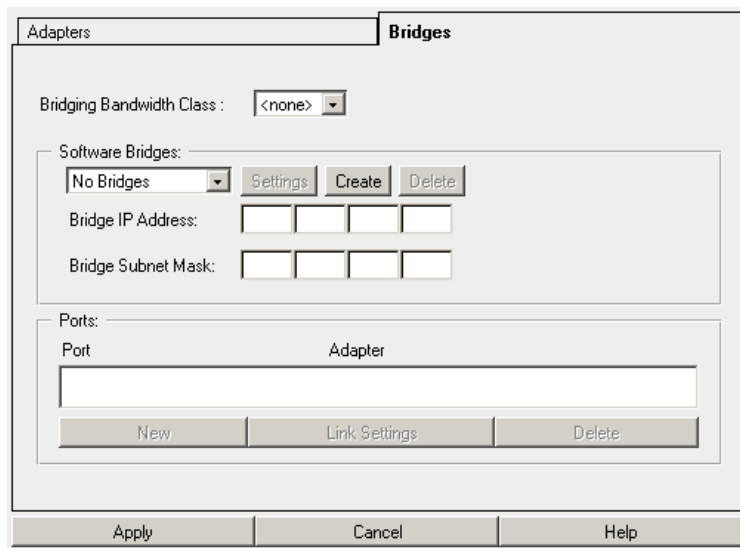


Figure 4-8: Bridges Tab

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.
3. Click Apply.

To Configure Bandwidth Management for Bridging through the CLI:

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) bridge
SGOS#(config bridge) bandwidth-class bw_class_name
```

where *bw_class_name* designates the name of the bandwidth class that you have created and configured to manage the bandwidth for software bridging.

2. (Optional) To disable bandwidth management for software bridging, enter the following command:

```
SGOS#(config bridge) no bandwidth-class
```

Configuring a Software Bridge

This section describes how to use the Management Console or the CLI to link adapters and interfaces to create a network bridge.

To Create and Configure a Software Bridge through the Management Console:

1. Select Configuration>Network>Adapters>Bridges.
The Bridges tab displays.
2. In the Software Bridges area, click Create.
3. In the New Bridge Name field of the dialog that appears, enter a name for the bridge, up to 16 characters; click OK.
4. In the Bridge IP Address field, enter the IP address of the interface you previously configured (see "Configuring an Adapter" on page 72).
5. In the Bridge Subnet Mask field, enter the subnet mask of the interface.
6. To add a port to the bridge:
 - a. In the Ports field, click New; the Create port for bridge dialog appears.
 - b. From the drop-down lists, select a port number and adapter interface number; click OK.
 - c. By default, link settings are automatically sensed. To change the Duplex and Speed options, click Link Settings, select Manually configure link settings, and change as required.
 - d. Click OK.
7. Further customize the bridge:
 - a. In the Software Bridges field, click Settings; the Settings for bridge dialog appears.
 - b. In the Security field, the default is to accept inbound connections on this interface. To disallow inbound connections, select Reject inbound connections.
 - c. The default browser instruction is to use the browser's default PAC file. To instruct the browser to use a proxy or other PAC file type, make a selection from the list in the Browser Configuration field.
 - d. Click OK.
8. Click Apply.

The Bridge Settings options allow you to clear bridge forwarding table and clear bridge statistics.

To Create or Edit a Software Bridge through the CLI:

1. To create a new software bridge, enter the following commands at the (config) command prompt:

```
SGOS#(config) bridge
SGOS#(config bridge) create bridge_name
      where bridge_name designates the name of the new bridge. The limit is 16 characters.
```
2. To edit the configuration of an existing software bridge, enter the following commands:

```
SGOS#(config bridge) edit bridge_name
      where bridge_name designates the name of the bridge that you want to configure. The prompt changes to a submode for that bridge.
```

```
SGOS#(config bridge bridge_name) ip-address ip_address
```

where *ip_address* designates the IP address of the adapter interface you previously configured (see "Configuring an Adapter" on page 72).

```
SGOS#(config bridge bridge_name) subnet-mask subnet_mask
```

where *subnet_mask* designates the subnet mask of the interface you previously configured.

3. To configure a port on a bridge, enter the following commands (repeat to add more ports):

```
SGOS#(config bridge bridge_name) port port_number
```

where *port_number* identifies a port on the interface. This changes the prompt to a submode for that port number on that bridge.

- To attach port to an interface or change the Duplex and Speed options, enter the following commands:

```
SGOS#(config bridge bridge_name port port_number) attach-interface interface_number
```

```
SGOS#(config bridge bridge_name port port_number) {full-duplex | half-duplex}
```

```
SGOS#(config bridge bridge_name port port_number) speed {10 | 100 | 1gb}
```

where:

<code>attach-interface</code>	<code><i>interface_number</i></code>	Attaches an interface for this port.
<code>full-duplex</code>		Configures this port for full duplex.
<code>half-duplex</code>		Configures this port for half duplex.
<code>speed</code>	<code>10 100 1gb</code>	Configures speed for this port.

```
SGOS#(config bridge bridge_name port port_number) exit
```

```
SGOS#(config bridge bridge_name)
```

- By default, link settings are automatically sensed. To perform an auto-sense, enter the following command:

```
SGOS#(config bridge bridge_name port port_number) link-autosense
```

- Return to the *bridge_name* submode:

```
SGOS#(config bridge bridge_name port port_number) exit
```

```
SGOS#(config bridge bridge_name)
```

4. To specify the maximum transmission unit (MTU), enter the following command:

```
SGOS#(config bridge bridge_name) mtu-size size
```

where *size* is a value from 72 to 1500.

5. The default is to accept inbound connections on this interface. To disallow inbound connections, enter the following command:

```
SGOS#(config bridge bridge_name) no accept-inbound
```

6. The default browser instruction is to use the browser's default PAC file. To instruct to use a proxy or other PAC file type, enter the following command:

```
SGOS#(config bridge bridge_name) instructions {proxy | default-pac | central-pac url | accelerated-pac}
```

where:

proxy		Use a proxy.
default-pac		Use the Blue Coat default PAC file.
central-pac	<i>url</i>	Use the PAC file specified at the given URL.
accelerated-pac		Use the proxy's accelerated PAC file.

Configuring Failover

You can configure failover for software bridges, but not for hardware bridges. Failover is accomplished by creating virtual IP addresses on each proxy, creating a failover group, and attaching the bridge configuration. One of the proxies *must* be designated with a higher priority (a master proxy).

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 port 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*.

```
ProxySG_A#(config) virtual-ip address 10.0.0.4
ProxySG_A#(config) failover
ProxySG_A#(config failover) create 10.0.0.4
ProxySG_A#(config failover) edit 10.0.0.4
ProxySG_A#(config failover 10.0.0.3) master
ProxySG_A#(config failover 10.0.0.3) priority 100
ProxySG_A#(config failover 10.0.0.3) interval 1
```

- **ProxySG B**—software bridge IP address: 10.0.0.3. Create a virtual IP address and a failover group.

```
ProxySG_B#(config) virtual-ip address 10.0.0.4
ProxySG_B#(config) failover
ProxySG_B#(config failover) create 10.0.0.4
ProxySG_B#(config failover) edit 10.0.0.4
ProxySG_B#(config failover 10.0.0.3) priority 100
ProxySG_B#(config failover 10.0.0.3) interval 1
```

- In the bridge configuration on *each* ProxySG, attach the bridge configuration to the failover group:

```
ProxySG_A#(config bridge bridge_name) failover 10.0.0.4
ProxySG_B#(config bridge bridge_name) failover 10.0.0.4
```

Static Forwarding Table Entries

Certain firewall configurations require the use of static forwarding table entries. 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 Entry through the CLI:

1. At the (config) prompt, enter the following commands:

```
SGOS# (config) bridge
SGOS# (config bridge) bridge_name
SGOS# (config bridge bridge_name) port port_number
SGOS# (config bridge_name port port_number) static-fwtable-entry mac_address
```

2. Add up to 256 entries per bridge.

To Clear a Static Forwarding Table Entry through the CLI:

At the (config) prompt, enter the following commands:

```
SGOS# (config) bridge
SGOS# (config bridge) bridge_name
SGOS# (config bridge bridge_name) clear-fwtable
```

Gateways

A key feature of the ProxySG is the ability to distribute traffic originating at the appliance through multiple gateways. You can also fine tune how the traffic is distributed to different gateways. This feature works with any routing protocol (such as static routes or RIP).

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.

About Gateways

During the initial setup of the ProxySG, you optionally defined a gateway (a device that serves as entrance and exit into a communications network) for the ProxySG.

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.

The gateway specified applies to all network adapters in the system.

ProxySG Specifics

Which gateway the ProxySG chooses to use at a given time is determined by how the administrator configures the assignment of preference groups to default gateways. You can define multiple gateways within the same preference group. A ProxySG can have from 1 to 10 preference groups. If you have only one gateway, it automatically has a weight of 100.

Initially, all gateways in the lowest preference group are considered to be the active gateways. If a gateway becomes unreachable, it is dropped from the active gateway list, but the remaining gateways within the group continue to be used until they all become unreachable, or until an unreachable gateway in a lower preference group becomes reachable again. If all gateways in the lowest preference group become unreachable, the gateways in the next lowest preference group become the active gateways.

In addition to a preference group, each gateway within a group can be assigned a relative weight value from 1 to 100. The weight value 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 value, whether 1 or 100, results in the same traffic distribution pattern. In a group with two gateways, 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.

Switching to a Secondary Gateway

When a gateway goes down, the ProxySG takes from 120 to 180 seconds to determine that the gateway is unreachable. At that point, the ProxySG switches to a secondary gateway if one is configured.

The ProxySG continues to check failed gateways once a minute using Address Resolution Protocol (ARP). The gateways are declared unreachable after three attempts. When a preferred gateway comes back on line, however, it might take up to 180 seconds for the ProxySG to confirm the preferred gateway is reachable and to switch back to that gateway.

These times are not user-configurable.

To Configure Multiple Gateway Load Balancing through the Management Console:

1. Select Configuration>Network>Routing>Gateways.

The Gateways tab displays.

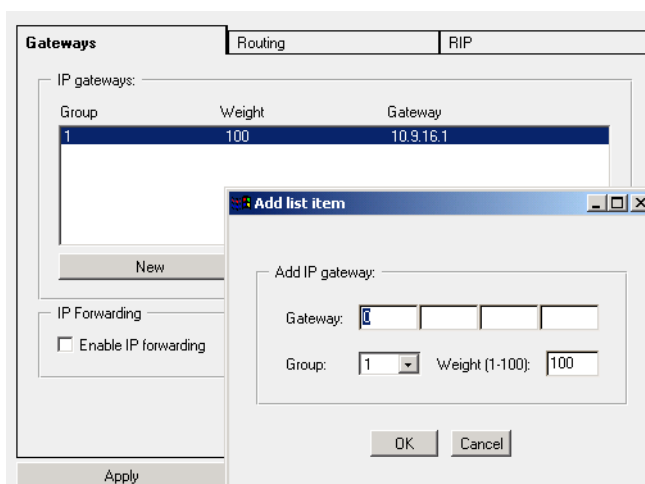


Figure 4-9: Network Routing Gateways Tab and Add List Item Dialog

2. Click New.

3. Enter the IP address, group, and weight for the gateway into the Add list item dialog that appears.
4. Click OK.
5. Repeat steps 2 to 4 until IP addresses, groups, and weights have been defined for all of your gateways.
6. Click Apply.

To Configure Multiple Gateway Load Balancing through the CLI:

1. At the `(config)` command prompt, enter the following command:

```
SGOS#(config) ip-default-gateway ip_address preference_group weight
```

The first value is the IP address of the gateway, the second value is the preference group, and the third value is the relative weighting for this gateway. For example, to use the gateway 10.25.36.1, the preference group 1, and the relative weighting 100, enter:

```
ip-default-gateway 10.25.36.1 1 100
```

2. Repeat until all IP addresses, groups, and weights of your IP gateways have been defined.
3. (Optional) View the results.

```
SGOS#(config) show ip-default-gateway
```

```
Default IP gateways
```

Gateway	Weight	Group
10.25.36.1	100	1

Defining Static Routes

The ProxySG can be configured to use *static routes*, a manually-configured route that specifies the transmission path a packet must follow, based on the packet's destination address. A static route specifies a transmission path to another network.

Note: You are limited to 10,000 entries in the static routes table.

You can install the routing table several ways.

- Using the ProxySG 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 static-route-table` command, which allows you to paste a static route table into the ProxySG.
- Using the CLI `static-routes` command, which requires that you place an already-created file on an FTP or HTTP server and enter the URL into the ProxySG.

The routing table is a text file containing a list of IP addresses, subnet masks, and gateways. The following is a sample router table:

10.25.36.0	255.255.255.0	10.25.46.57
10.25.37.0	255.255.255.0	10.25.46.58
10.25.38.0	255.255.255.0	10.25.46.59

When a routing table is loaded, all requested URLs are compared to the list and routed based on the best match.

To Install a Routing Table through the Management Console:

1. Select Configuration>Network>Routing>Routing.

The Routing tab displays.

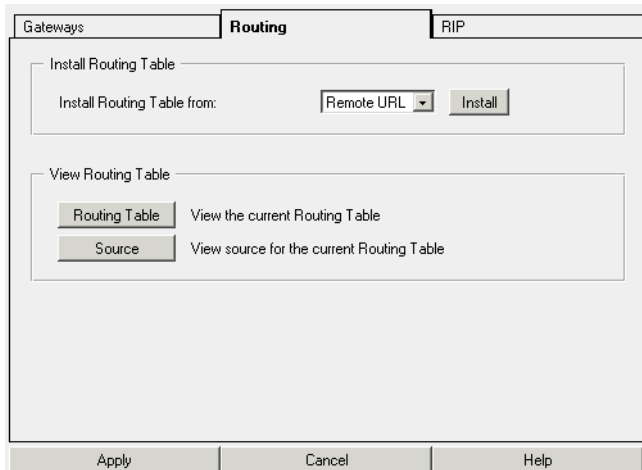


Figure 4-10: Network Routing Tab

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.

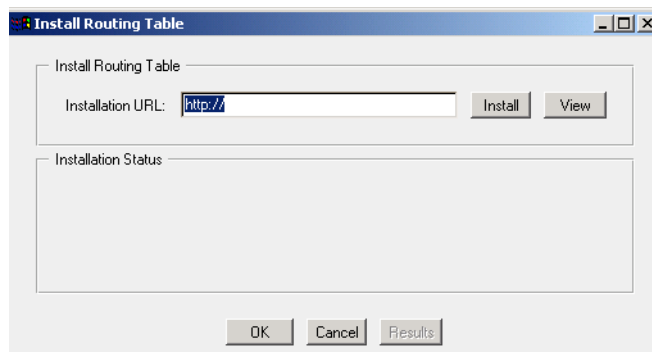


Figure 4-11: Specifying the Remote Location of a Routing Table

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.

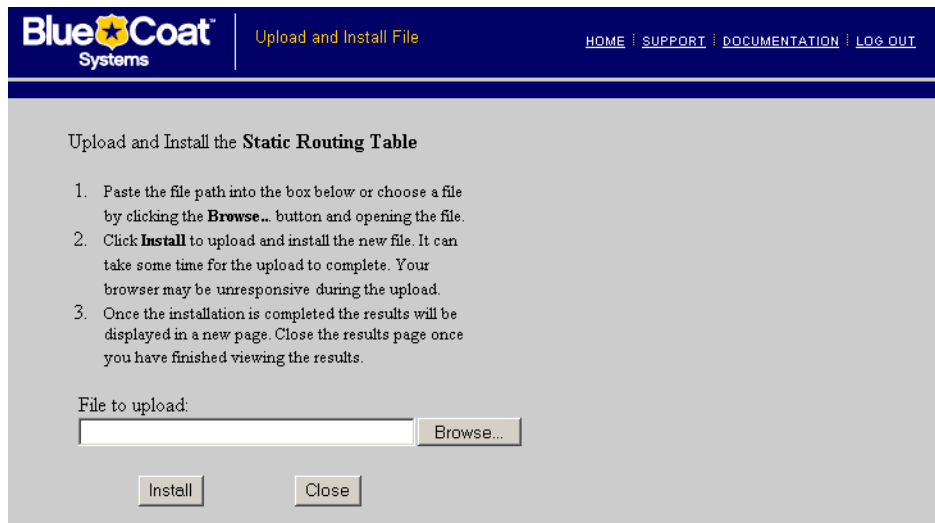


Figure 4-12: Specifying the Local Location of a Routing Table

□ **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**.



Figure 4-13: Creating a Static Routing Table on the ProxySG

3. Click Apply.

Installing a Routing Table Through the CLI

To install a routing table through the CLI, you can use the `inline` command to install the table directly, or enter a path to a remote URL that has an already-created text file ready to download.

When entering input for the `inline` command, you can correct mistakes on the current line using the `<Backspace>` key. If you detect a mistake in a line that has already been terminated using the `<Enter>` key, you can abort the `inline` command by typing `<Ctrl-c>`. If the mistake is detected after you terminate input to the `inline` command, type the same `inline` command again, but with the correct configuration information. The corrected information replaces the information from the last `inline` command.

The end-of-input marker is an arbitrary string chosen by the you to mark the end of input for the current `inline` command. The string can be composed of standard characters and numbers, but cannot contain any spaces, punctuation marks, or other symbols.

Take care to choose a unique end-of-input string that does not match any string of characters in the configuration information.

To Install a Routing Table through the CLI:

Do one of the following:

- To paste a static route table directly into the CLI, enter the following command at the (config) command prompt, then paste the table on the line after the first *end-of-file* marker:

```
SGOS#(config) inline static-route-table end-of-file_marker
paste static routing table
eof
ok
```

- To enter the static route table manually, enter the following command, then enter each IP address/subnet on the second line, following the first *end-of-file* marker:

```
SGOS#(config) inline static-route-table end-of-file_marker
10.25.36.0 255.255.255.0 10.25.46.57
10.25.37.0 255.255.255.0 10.25.46.58
10.25.38.0 255.255.255.0 10.25.46.59
eof
ok
```

- To enter a path to a remote URL where you have placed an already-created static route table, enter the following commands at the (config) command prompt:

```
SGOS#(config) static-routes path url
SGOS#(config) load static-route-table
```

Using RIP

The Routing Information Protocol (RIP) is designed to select the fastest route to a destination. RIP support is built into the ProxySG, and is configured by created and installing an RIP configuration text file onto the ProxySG. (No RIP configuration file is shipped with the appliance.) For commands that can be entered into the RIP configuration file, see Appendix D: "RIP Commands" on page 937.

Once you have created an RIP configuration file, you can install it several ways:

- Using the ProxySG 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 through the Management Console:

Note: When entering RIP settings that will change current settings (for instance, 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.

The RIP tab displays.

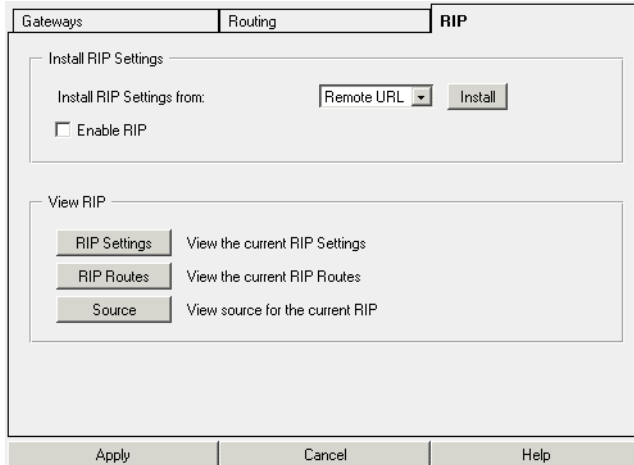


Figure 4-14: Network Routing RIP Tab

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 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.

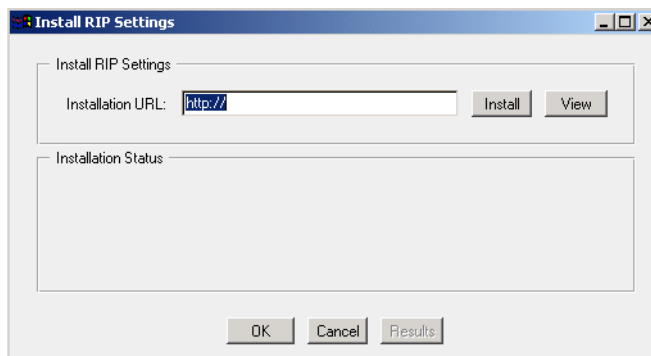


Figure 4-15: Specifying the Remote Location of a RIP Configuration File

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.



Figure 4-16: Specifying the Local Location of a RIP File

□ **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**.



Figure 4-17: Creating an RIP file on the ProxySG

4. Click **Apply**.
5. Select **Enable RIP**.
6. Click **Apply**.

Configuring RIP through the CLI

Note: When entering RIP settings that will change current settings (for instance, when switching from `ripv1` to `ripv2`), disable RIP before you change the settings; re-enable RIP when you have finished.

To Disable/Enable RIP through the CLI:

Enter the following command at the `(config)` command prompt:

```
SGOS#(config) rip {disable | enable}
```

To Install an RIP Configuration through the CLI:

Do one of the following:

- To enter a path to a remote URL where you have placed an already-created RIP configuration file, enter the following commands at the `(config)` command prompt:

```
SGOS#(config) rip path url
SGOS#(config) load rip-settings
```

- To paste an RIP configuration directly into the CLI, enter the following command at the `(config)` command prompt:

```
SGOS#(config) inline rip-settings end-of-file_marker
```

At this point you can paste RIP settings into the `inline` command, or you can enter values for specific settings. When you finish, enter your `end-of-file_marker` command.

Example

```
SGOS#(config) inline rip-settings eofm
ripv2
ripv1_out
no_rdisc eofm
ok
```

DNS

During first-time installation of the ProxySG, you configured the IP address of a single primary Domain Name Service (DNS) server. Using the Configuration>Network>DNS tab, you can change this primary DNS server at any time, and you can also define additional primary DNS servers and one or more alternate DNS servers.

ProxySG Specifics

If you have defined more than one DNS server, the ProxySG uses the following logic to determine which servers will be used to resolve a DNS host name and when to return an error to the client:

- The ProxySG first sends requests to DNS servers in the primary DNS server list.
- Servers are always contacted in the order in which they appear in a list.
- The next server in a list is only contacted if the ProxySG does not receive a response from the current server.

- If none of the servers in a list returns a response, the ProxySG returns an error to the client.
- The ProxySG only sends requests to servers in the alternate DNS server list if a server in the primary list indicates that a DNS host name cannot be resolved.

If a DNS server returns any other error (other than an indication that a DNS host name could not be resolved), the ProxySG returns the error to the client.

If a server in both the primary and alternate DNS server lists are unable to resolve a DNS host name, an error is returned to the client.

The ProxySG always attempts to contact the first server in the primary DNS server. If a response is received from this server, no attempts are made to contact any other DNS servers in the primary list.

If the response from the first primary DNS server indicates a name error, the ProxySG sends a DNS request to the first alternate DNS server, if one is defined. If no alternate DNS servers have been defined, an error is returned to the client indicating a name error. If the first alternate DNS server is unable to resolve the IP address, a name error is returned to the client, and no attempt is made to contact any other DNS servers in either the primary or alternate DNS server lists.

If a response is not received from any DNS server in a particular DNS server list, the ProxySG sends a DNS request to the next server in the list. The ProxySG returns a name error to the client if none of the servers in a DNS server list responds to the DNS request.

Note: The alternate DNS server is not used as a failover DNS server. It is only used when DNS resolution of primary DNS server returns name error. If a timeout occurs when looking up the primary DNS server, no alternate DNS server is contacted.

If the ProxySG receives a negative DNS response (a response with an error code set to Name Error), it caches that negative response. You can configure the ProxySG's negative response time-to-live value. (A value of zero disables negative caching.) If the ProxySG is not configured (the default), the ProxySG caches the negative response and uses the TTL value from the DNS response to determine how long it should be cached.

Configuring Split DNS Support

Customers with split DNS server configuration (for example, environments that maintain private internal DNS servers and external DNS servers) might choose to populate an Alternate DNS server list as well as the Primary DNS server list. In the ProxySG, the internal DNS servers are placed in the Primary list, while external DNS servers (with the Internet information) populate the Alternate list.

Complete the following procedures to configure Primary and Alternate DNS servers.

To Add a Primary DNS Server through the Management Console:

1. Select Configuration>Network>DNS>DNS.

The DNS tab displays.

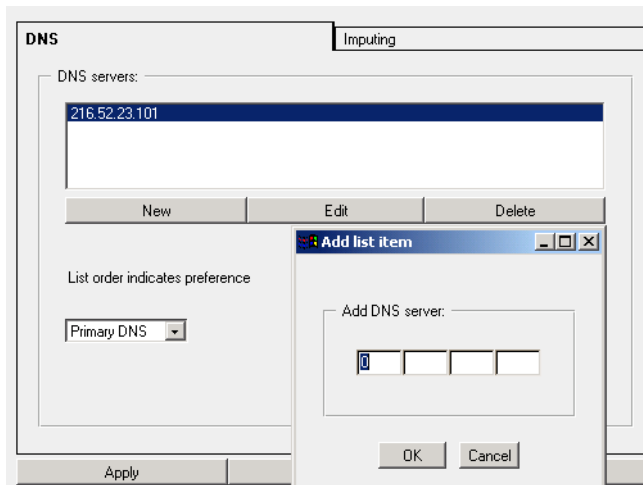


Figure 4-18: Network DNS Tab and Add List Item Dialog

2. Click New.
3. Enter the IP address of the DNS server in the dialog that appears and click OK.
4. Click Apply.

To Add a Primary DNS Server through the CLI:

At the (config) command prompt, enter the following command:

```
SGOS#(config) dns server ip_address
```

To Add an Alternate DNS Server through the Management Console:

1. Select Configuration>Network>DNS>DNS.
The DNS tab displays.
2. Select Alternate DNS in the drop-down list.
3. Click New.
4. Enter the IP address of the DNS server in the dialog that appears and click OK.
5. Click Apply.

To Add an Alternate DNS Server through the CLI:

1. At the (config) command prompt, enter the following command:
SGOS#(config) dns alternate ip_address
2. Repeat until alternate DNS servers have been defined.

Changing the Order of DNS Servers

The ProxySG uses DNS servers in the order displayed. You can organize the list of servers so that the preferred servers appear at the top of the list. This functionality is not available through the CLI.

To Change the Order of DNS Servers through the Management Console:

1. Select Configuration>Network>DNS>Imputing.

The Imputing tab displays.

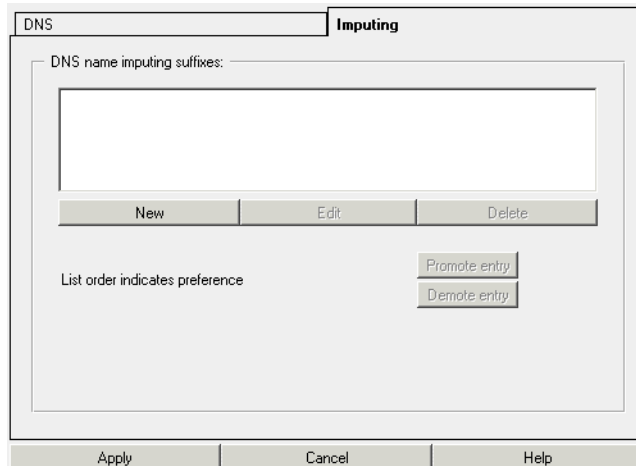


Figure 4-19: Network DNS Imputing Tab

2. Select the DNS server to promote or demote.
3. Click Promote entry or Demote entry as appropriate.
4. Click Apply.

Unresolved Host Names (Name Imputing)

Name imputing allows the ProxySG to resolve host names based on a partial name specification. When the ProxySG submits a host name to the DNS server, the DNS server resolves the name to an IP address. The ProxySG queries the original host name before checking imputing entries unless there is no period in the host name, in which case imputing is applied first. The ProxySG tries each entry in the name-imputing list until the name is resolved or it comes to 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 `company.com` and `com`, and a user submits the URL `http://www.eedept`, the ProxySG resolves the host names in the following order.

```
http://www.eedept
http://www.eedept.company.com
http://www.eedept.com
```

To Add Names to the Imputing List through the Management Console:

1. Select Configuration>Network>DNS>Imputing.
The Imputing tab displays.
2. Click New to add a new name to the imputing list.
3. Enter the name in the dialog that appears and click OK.
4. Click Apply.

To Add Names to the Imputing List through the CLI:

1. At the (config) command prompt, enter the following command:

```
SGOS#(config) dns imputing suffix
```

For example, to use `company.com` as the imputing suffix, enter `dns-imputing company.com`.

2. Repeat until all imputing suffixes have been entered.

Changing the Order of DNS Name Imputing Suffixes

The ProxySG uses imputing suffixes in the order displayed. You can organize the list of suffixes so the preferred suffix appears at the top of the list. This functionality is only available through the Management Console. You cannot configure it through the CLI.

To Change the Order DNS Name Imputing Suffixes through the Management Console:

1. Select Configuration>Network>DNS>Imputing.
The Imputing tab displays.
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.

Both type A and type PTR DNS responses are affected by negative caching.

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
```

Attack Detection

The ProxySG can reduce the effects of distributed denial of service (DDoS) attacks and port scanning, two of the most common virus infections.

A DDoS attack occurs when a pool of machines that have been infected with a DDoS-type of virus attack a specific Web site. As the attack progresses, the target host shows decreased responsiveness and often stops responding. Legitimate HTTP traffic is unable to proceed because the ProxySG is still waiting for a response from the target host.

Port scanning involves viruses attempting to self-propagate to other machines by arbitrarily trying to connect to other hosts on the Internet. If the randomly selected host is unavailable or behind a firewall or does not exist, the ProxySG continues to wait for a response, thus denying legitimate HTTP traffic.

The ProxySG prevents attacks by limiting the number of simultaneous TCP connections from each client IP address and either does not respond to connection attempts from a client already at this limit or resets the connection. It also limits connections to servers known to be overloaded.

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 virus-infected machines behind the ProxySG. The *server* attack-detection configuration is used when an administrator knows ahead of time that a virus is set to attack a specific host.

This feature is only available through the CLI. You cannot use the Management Console to enable attack detection.

For information on configuring a client, continue with the next section. To configure a server for attack detection, continue with "Configuring Attack-Detection Mode for a Server or Server Group" on page 99.

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)
```

To Change 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: true
- client interval: 20 minutes
- block-action: drop (for each client)
- connection-limit: 100 (for each client)
- failure-limit: 50 (for each client)
- unblock-time: unlimited

- warning-limit: 10 (for each client)

To Change the Global Defaults:

Remember that enable/disable limits and interval affect all clients. The values cannot be changed for individual clients. Other limits can be modified on a per-client basis.

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_65535
SGOS#(config client) default failure-limit integer_between_1_and_500
SGOS#(config client) default unblock-time minutes_between_10_and_1440
SGOS#(config client) default warning-limit integer_between_1_and_100
```

where:

<pre>enable-limits disable-limits</pre>		<p>Toggles between enabled and disabled. The default is disabled. This is a global setting and cannot be modified for individual clients.</p>
<pre>interval</pre>	<pre>integer</pre>	<p>Indicates the amount of time, in multiples of 10 minutes, that client activity is monitored. The default is 20. This is a global setting and cannot be modified for individual clients.</p>
<pre>block unblock</pre>	<pre><i>ip_address</i> [<i>minutes</i>]</pre>	<p>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.</p>
<pre>default block-action</pre>	<pre>drop send-tcp-rst</pre>	<p>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.</p>
<pre>default connection-limit</pre>	<pre><i>integer</i></pre>	<p>Indicates the number of simultaneous connections between 1 and 65535. The default is 100. This limit can be modified on a per-client basis.</p>
<pre>default failure-limit</pre>	<pre><i>integer</i></pre>	<p>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.</p>
<pre>default unblock-time</pre>	<pre><i>minutes</i></pre>	<p>Indicates the amount of time a client is blocked at the network level when the client-warning-limit is exceeded. Time must be a multiple of 10 minutes, up to a maximum of 1440. The default is unlimited. This limit can be modified on a per-client basis.</p>

default	<i>integer</i>	Indicates the number of warnings sent to the client before the client is blocked at the network level and the administrator is notified. The default is 10; the maximum is 100. This limit can be modified on a per-client basis.
warning-limit		

To Create and Edit a Client IP Address through the CLI:

1. Make sure you are in the attack-detection client submenu.

```
SGOS#(config) attack-detection
SGOS#(config attack-detection) client
SGOS#(config client)
```

2. Create a client.

```
SGOS#(config client) create client ip_address or ip_and_length
```

3. Move to edit client submenu.

```
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) connection-limit integer_between_1_and_65535
SGOS#(config client ip_address) failure-limit integer_between_1_and_65535
SGOS#(config client ip_address) unblock-time minutes
SGOS#(config client ip_address) warning-limit integer_between_1_and_65535
```

where:

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.
connection-limit	<i>integer</i>	Indicates the number of simultaneous connections between 1 and 65535. The default is 100.
failure-limit	<i>integer</i>	Indicates the behavior when the specified 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 50.
unblock-time	<i>minutes</i>	Indicates the amount of time a client is locked out at the network level when the client-warning-limit is exceeded. Time must be a multiple of 10 minutes, up to a maximum of 1440. The default is unlimited.
warning-limit	<i>integer</i>	Indicates the number of warnings sent to the client before the client is locked out at the network level 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 connection limit:      700
Client failure limit:         50
Client warning limit:         10
Blocked client action:        Drop
Client connection unblock time: unlimited
```

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 <cr> | blocked | connections | statistics
```

To View All Settings:

```
SGOS#(config client) view
Client limits enabled:          true
Client interval:                20 minutes
Default client limits:
Client connection limit:        100
Client failure limit:           50
Client warning limit:           10
Blocked client action:          Drop
Client connection unblock time: unlimited
Client limits for 10.25.36.47:
Client connection limit:        700
Client failure limit:           50
Client warning limit:           10
Blocked client action:          Drop
Client connection unblock time: unlimited
```

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
```

To View Client Statistics:

```
SGOS#(config client) view statistics
Client IP           Failure Count      Warning Count
10.9.44.72         1                  0
```

To Disable Attack-Detection Mode for all Clients:

```
SGOS#(config client) disable-limits
```

Configuring Attack-Detection Mode for a Server or Server Group

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 servers in the group are blocked.

To Create a Server or Server Group:

1. At the (config) prompt, enter the following commands:

```
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:

1. At the (config server) prompt, enter the following commands:

```
SGOS#(config server) edit hostname
```

The prompt changes to (config server hostname).

```
SGOS#(config server hostname) {add | remove} hostname
SGOS#(config server hostname) request-limit integer_from_1_to_65535
```

where:

<code>hostname</code>	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.
<code>add remove hostname</code>	Adds or removes a server from this server group.
<code>request-limit integer</code>	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:
Request limit: 1500
```

Using a Bypass List

A bypass list prevents the ProxySG from transparently accelerating requests to servers that perform IP authentication with clients. The bypass list contains IP addresses, subnet masks, and gateways. When a request matches an IP address and subnet mask specification in the bypass list, the request is sent to the designated gateway. A bypass list is only used for transparent caching.

There are three types of bypass lists: local list, central list, and policy-based list. Each of these bypass lists are discussed below.

The first two lists are not the same as the Local Policy file and the Central Policy file. The policy-based bypass list is a list maintained in the Forward Policy file or Local Policy file.

The local and central bypass lists can be managed two ways: either through the Management Console or through the CLI. For installation procedures for the two lists, see "Creating and Installing Local or Central Bypass Lists" on page 104.

Using the Local Bypass List

The local bypass list is one you create and maintain on your network. You can use a local bypass list alone, or in conjunction with a central list. You can also use a dynamic local bypass list to increase ProxySG efficiency. For more information on dynamic bypass lists, see "Using Dynamic Bypass" on page 101.

The gateways specified in the bypass list must be on the same subnet as the ProxySG. The local bypass list limit is 10,000 entries.

The local bypass list contains a list of IP addresses, subnet masks, and gateways. It can also define the default bypass gateway to be used by both the local bypass list and central bypass list. The gateways specified in the bypass list must be on the same subnet as the ProxySG. When you download a bypass list, the list is stored in the appliance until it is replaced by downloading a new list.

The following is a sample of a local bypass list:

```
;define the default gateway for the local and central bypass list
BYPASS_GATEWAY 10.25.46.57
;define addresses to bypass
;IP address    subnet                gateway (or use default gateway)
10.25.36.47    255.255.255.255
10.25.36.48    255.255.255.255
10.25.0.0      255.255.255.0      10.25.46.58
```

Note: The `BYPASS_GATEWAY` and default gateway must be on a different subnet from the IP addresses.

If you do not specify the `BYPASS_GATEWAY`, and you do not designate the gateway in the address specification, the ProxySG forwards the request to the default gateway defined in the network configuration.

For installation procedures for the local bypass list, see "Creating and Installing Local or Central Bypass Lists" on page 104.

Using Dynamic Bypass

Dynamic bypass provides a maintenance-free method for improving performance of the ProxySG by automatically compiling a list of requested URLs that return various kinds of errors.

With dynamic bypass, the ProxySG adds dynamic bypass entries containing the server IP address of sites that have returned an error to the appliance's local bypass list. For a configured period of time, further requests for the error-causing URLs are sent immediately to the origin content server (OCS), saving the ProxySG processing time. 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.

Once the dynamic bypass timeout for a URL has ended, the ProxySG removes the URL from the bypass list. On the next client request for the URL, the ProxySG attempts to contact the OCS. If the OCS still returns an error, the URL is once again added to the local bypass list for the configured dynamic bypass timeout. If the URL does not return an error, the request is handled in the normal manner.

Dynamic bypass increases ProxySG efficiency because redundant attempts to contact the OCS are minimized.

Limitations

- Dynamic bypass applies to transparent proxy connections only.
- Dynamic bypass entries are lost when the ProxySG is restarted or the static bypass file is reinstalled.
- No filtering checks are performed on client requests that match entries in the dynamic bypass list.
- Requests to sites that are put into the dynamic bypass list bypass future policy evaluation. If a site that requires forwarding policy to reach its destination is populated into the bypass list, the site might be inaccessible.

Sites requiring that client accesses always be subjected to ProxySG filtering considerations must either use the appliance in explicit proxy mode or leave dynamic bypass functionality disabled.

Configuring Dynamic Bypass

Dynamic bypass is disabled by default. Enabling and fine-tuning dynamic bypass is a two-step process:

- Edit or create a local bypass list, adding the desired dynamic bypass timeout and threshold parameters.
- Use 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, which is optional, in configuring dynamic bypass is to edit the local bypass list to set the `SERVER_BYPASS_THRESHOLD`, `MAX_DYNAMIC_BYPASS_ENTRY`, and/or `DYNAMIC_TIMEOUT` values. This step is optional because the ProxySG uses default configurations if you do not specify them in the local bypass list. 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_BYPASS_THRESHOLD` value defines the maximum number of entries in the dynamically generated portion of the local bypass list before the ProxySG consolidates client-server pair entries into a single server entry. 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 `DYNAMIC_TIMEOUT`.

The `MAX_DYNAMIC_BYPASS_ENTRY` defines the maximum number of total dynamic bypass entries. The range is 1 to 50,000. The default value is 16,000. When the number of entries exceeds the `MAX_DYNAMIC_BYPASS_ENTRY` value, the oldest entries are removed to make way for new entries.

The `DYNAMIC_TIMEOUT` value defines the number of minutes a dynamic bypass entry can remain unreferenced before it is deleted from the bypass list. The range is 1 to 6000. 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.

To Enable Dynamic Bypass and Trigger Events through the CLI:

At the `(config)` command prompt, enter the following commands:

```
SGOS#(config) dynamic-bypass enable
SGOS#(config) dynamic-bypass trigger trigger_event
```

where `trigger_event` can be any item in listed in Table 4.1 on page 102.

Enabling dynamic bypass causes the following warning to appear:

```
WARNING:
Requests to sites that are put into the dynamic bypass list will bypass
future policy evaluation. This could result in subversion of on-box policy.
The use of dynamic bypass is cautioned.
```

Table 4.1: Values for the Dynamic-Bypass Trigger Event

Event	Description
all	Enables all dynamic bypass triggers.
non-http	Enables dynamic bypass for non-HTTP responses.
connect-error	Enables dynamic bypass for any connection failure to the origin content server, including timeouts.
receive-error	Enables dynamic bypass for when a TCP connection to an origin content server succeeds, but the cache does not receive an HTTP response.
400	Enables dynamic bypass for HTTP 400 responses.
401	Enables dynamic bypass for HTTP 401 responses.
403	Enables dynamic bypass for HTTP 403 responses.
405	Enables dynamic bypass for HTTP 405 responses.
406	Enables dynamic bypass for HTTP 406 responses.
500	Enables dynamic bypass for HTTP 500 responses.

502	Enables dynamic bypass for HTTP 502 responses.
503	Enables dynamic bypass for HTTP 503 responses.
504	Enables dynamic bypass for HTTP 504 responses.

Example

For instance, the following command will enable connection error events:

```
SGOS#(config) dynamic-bypass trigger connect-error
```

Bypassing Connection and Receiving Errors

In addition to HTTP code triggers, you can configure the ProxySG to trigger dynamic bypass for connection and receiving errors.

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. In this instance, the ProxySG bypasses any connection attempts from the client to this IP address. By default, the timeout duration is 20 seconds, and the retry count is 3. These parameters are not configurable. Both the timeout duration and the retry attempt, whichever occurs first, triggers `connect-error`.

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. In this instance, the appliance bypasses any attempts from the client to this IP address. Server timeouts can also trigger `receive-error`. The default timeout value is 180 seconds, which can be changed (see "Configuring HTTP Timeout" on page 66).

Disabling Dynamic Bypass Triggers

Disabling one or more specific dynamic bypass triggers is an easy way to customize which errors cause a dynamic bypass entry to be created. For example, if you want all error events except 401 responses to create a dynamic bypass entry, you can enable all triggers and then disable only the 401-event trigger.

To Disable One or More Dynamic Bypass Triggers through the CLI:

At the `(config)` command prompt, enter the following command:

```
SGOS#(config) dynamic-bypass no trigger event
```

where `event` can be any item listed above in Table 4.1.

To Clear the Dynamic Bypass List through the CLI:

At the `(config)` command prompt, enter the following command:

```
SGOS#(config) dynamic-bypass clear
```

To Disable Dynamic Bypass through the CLI:

At the `(config)` command prompt, enter the following command:

```
SGOS#(config) dynamic-bypass disable
```

Viewing the Dynamic Bypass List

You can view the dynamic bypass list several ways:

To Display the Dynamic Bypass List through the CLI:

At the (config) command prompt, enter the following command:

```
SGOS#(config) show bypass-list
```

To Display the Dynamic Bypass List through the Management Console:

In a Web browser, enter the following URL:

```
https://ip_address_of_ProxySG:8082/TCP/IP-bypass
```

To View the Current Dynamic Bypass Configuration through the CLI:

At the (config) command prompt, enter the following command:

```
SGOS#(config) show dynamic-bypass
```

To Disable Dynamic Bypass through the CLI:

At the (config) command prompt, enter the following command:

```
SGOS#(config) dynamic-bypass disable
```

Using the Central Bypass List

The central bypass list is a shared list of addresses that is used by multiple ProxySG Appliances. The central list contains addresses to bypass, but does not specify gateways (because the ProxySG Appliances are located on different subnets, using different gateways). The gateway used for matches in the central bypass list is defined using the `BYPASS_GATEWAY` command in the local bypass list. If there is no `BYPASS_GATEWAY` option, the ProxySG uses the default gateway defined by the network configuration.

You can create your own central bypass list to manage multiple ProxySG Appliances, or you can use the central bypass list maintained by Blue Coat Technical Support at:

<https://download.bluecoat.com/release/SG4/files/CentralBypassList.txt>

Note: The central bypass list is limited to 10,000 entries.

The central bypass list maintained by Blue Coat contains addresses Blue Coat has identified as using client authentication. You can determine whether to download the list automatically when it changes or to just be sent an e-mail notifying you of the update. By default, neither is enabled.

For installation procedures for the central bypass list, continue with the next section.

Creating and Installing Local or Central Bypass Lists

You can install the local and central bypass lists several ways:

- Use the ProxySG Text Editor, which allows you to enter the lists (or copy and paste the contents of an already-created file) directly onto the ProxySG through the Management Console (see the instructions below).

- Create a local file on your local system; use the Management Console to browse to the file and install it (see the instructions below).
- Use a remote URL, where you place an already-created file on an FTP or HTTP server to be downloaded to the ProxySG. This can be done through either the Management Console or the CLI (see the instructions below).
- Use the CLI `inline bypass-list central | local` command, which allows you to paste the configurations onto the ProxySG (see the instructions below). For more information on using the CLI `inline` command, see "Using the Local Bypass List" on page 100 or "Using the Central Bypass List" on page 104.

To Install Bypass Lists through the Management Console:

1. Select Configuration>Network>Advanced>Bypass List.

The Bypass List tab displays.

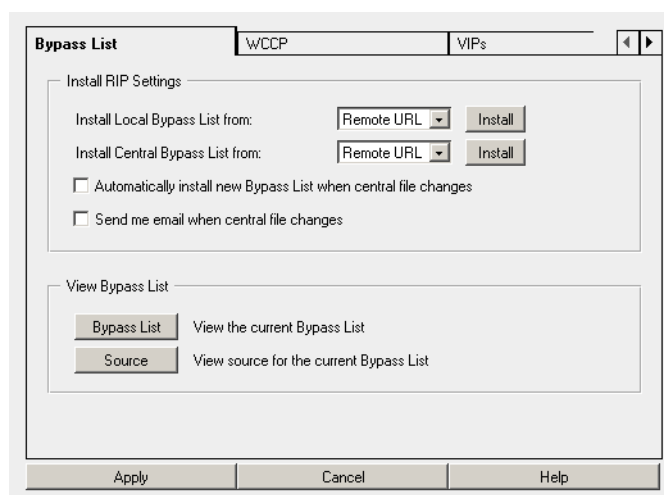


Figure 4-20: Bypass List Tab

2. To view the current bypass list or the source for the current bypass list before installing it, click Bypass List or Source.
3. (Optional) If installing the central bypass list, you can select whether to download the list automatically when it changes, or be sent an e-mail notifying you of the update. By default, neither is enabled.
4. Select a method to install the file for either the local or central bypass list; 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. View the installation status that displays; click OK.

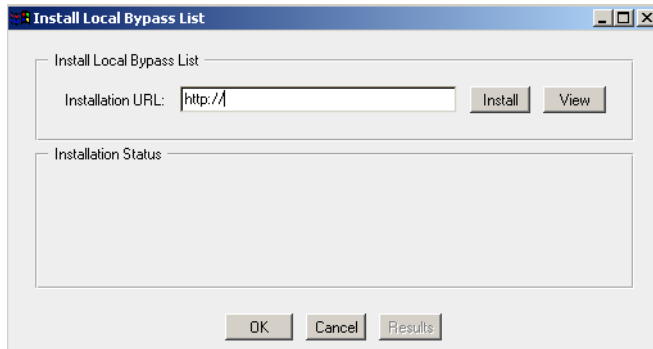


Figure 4-21: Specifying the Remote Location of a Local Bypass List Configuration File

□ Local File:

Click Browse to bring up the Local File Browse window. Browse for the file on your local system. Open it and click Install. When the installation is complete, a results window opens. View the results, close the window, and click Close.

Figure 4-22: Specifying the Local Location of a Local Bypass List

□ 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 Close.

Figure 4-23: Creating a Local Bypass List on the ProxySG

5. Click Apply.

To Install an Already-Existing Bypass List through the CLI:

At the `(config)` command prompt, enter the following commands:

```
SGOS#(config) bypass-list {local-path | central-path} url
SGOS#(config) load bypass-list {local | central}
```

To Install a Bypass List through the CLI inline Command:

At the `(config)` command prompt, enter the following commands:

```
SGOS#(config) inline bypass-list {local | central} end-of-file_marker
```

At this point you can paste in local or central configuration files, or you can enter values for specific settings, such as `server_bypass_threshold`, `max_dynamic_bypass_entry` or `dynamic_timeout`. When you finish, enter your *end-of-file* string.

Example

```
SGOS#(config) inline bypass-list local eof
max_dynamic_bypass_entry 20000
server_bypass_threshold 30
dynamic_timeout 100 eof
ok
```

Policy-Based Bypass Lists

ProxySG policies support the ability to define bypass lists. This section describes a property used to define a policy-based bypass list that can go into the Local Policy or Forward Policy file. For more information on defining policies, refer to the *Blue Coat ProxySG Content Policy Language Guide*.

While static and dynamic bypass lists allow traffic to bypass the ProxySG based on a destination IP address, the `bypass_cache` property is intended to allow a bypass based on the properties of the client. This property uses the following syntax:

```
bypass_cache(yes | no)
```

If set to `yes`, the ProxySG is not queried and the response is not stored. Set to `no` to specify the default behavior, which is to follow standard caching behavior. This property is available only in the `<proxy>` layer.

This property has no effect on streaming objects, but does affect the following types of transactions: HTTP, HTTPS, FTP over HTTP, and transparent FTP.

Example

```
; Bypass the cache for requests from this client IP.
client_address=10.25.198.0 bypass_cache(yes)
```

Installing WCCP Settings

The ProxySG can be configured to participate in a WCCP (Web Cache Control Protocol) scheme, where a WCCP-capable router collaborates with a set of WCCP-configured ProxySG Appliances to service requests.

Before you can install the WCCP configurations, you must create a WCCP configuration file for the ProxySG. The ProxySG does not ship with a default WCCP configuration file.

You can install the WCCP settings several ways:

- Using the ProxySG 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 wccp-settings` command, which allows you to paste the WCCP settings into the CLI.
- Using the CLI `wccp` command, which requires that you place an already-created file on an FTP or HTTP server and enter the URL into the CLI.

For more information about WCCP, see Appendix C: "Using WCCP" on page 911.

To Install WCCP Settings through the Management Console:

1. Select Configuration>Network>Advanced>WCCP.

The WCCP tab displays.

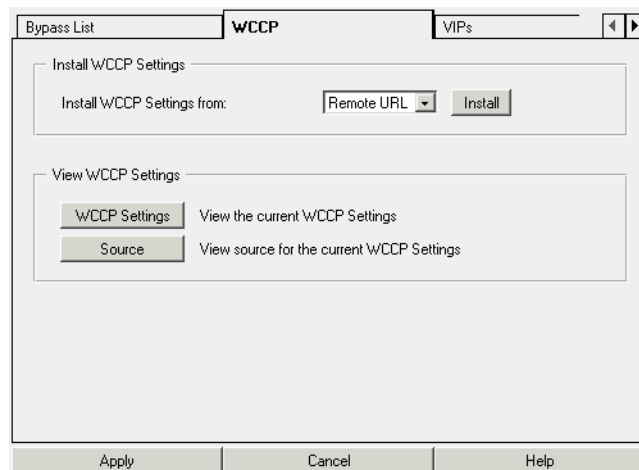


Figure 4-24: Network Advanced WCCP Tab

2. From the drop-down list, select the method used to install the WCCP settings; click Install.

❑ **Remote URL:**

Enter the fully-qualified URL, including the filename, where the WCCP file is located. To view the file before installing it, click View. Click Install. Viewing the installation status that displays; click OK.

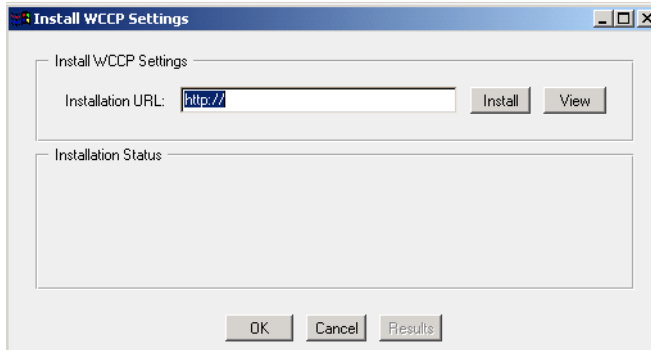


Figure 4-25: Specifying the Remote Location of a WCCP Settings File

□ **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, close the window, and click **Close**.

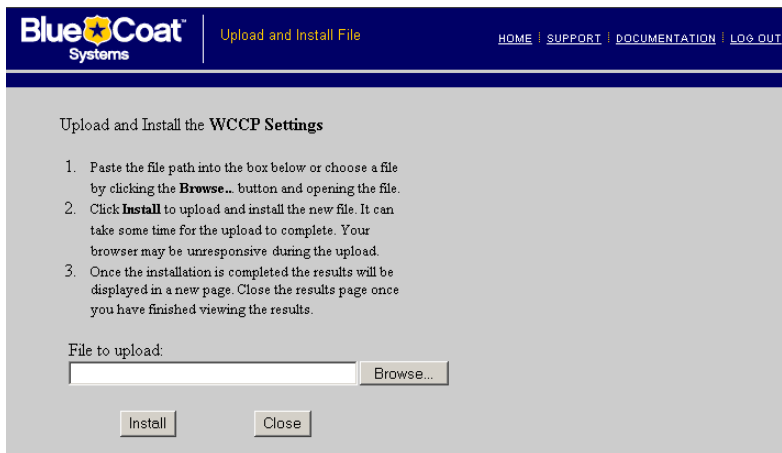


Figure 4-26: Specifying the Local Location of a WCCP Settings File

□ **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 **Close**.

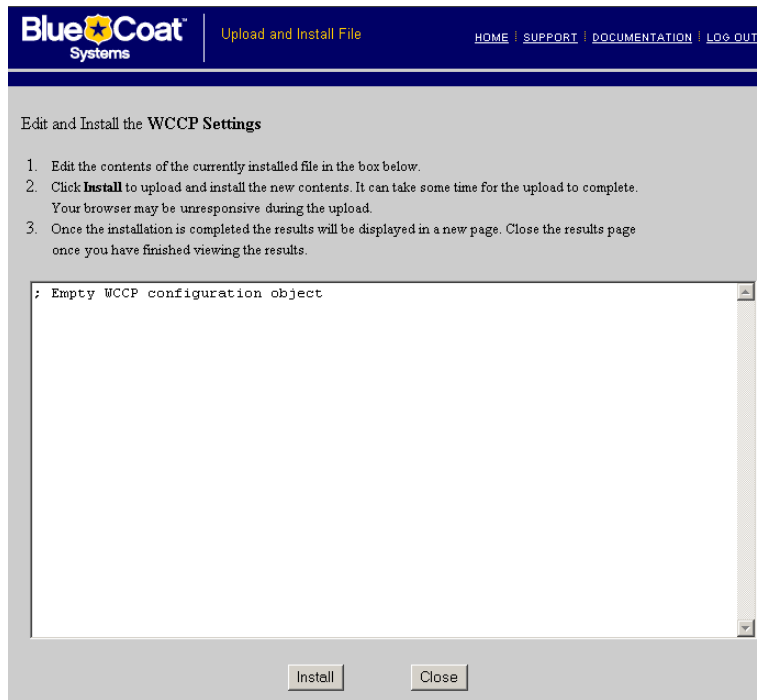


Figure 4-27: Creating a WCCP Settings File on the ProxySG

3. Click Apply.

To Install WCCP settings through the CLI:

Do one of the following:

- To enter WCCP settings directly onto the ProxySG, enter the following commands at the (config) command prompt:

```
SGOS#(config) inline wccp-settings end-of-file_marker
wccp enable
wccp version 2
service-group 9
priority 1
protocol 6
service-flags destination-ip-hash
service-flags ports-defined
ports 80 21 1755 554 80 80 80 80
interface 6
home-router 10.16.18.2
forwarding 12
eof
```

Note: For detailed instructions on configuring an WCCP file, see Appendix C: "Using WCCP" on page 911.

- To enter a path to a remote URL where you have placed an already-created static route table, enter the following commands at the (config) command prompt:

```
SGOS#(config) wccp path url
```

where *url* is a fully qualified URL, including the filename, where the configuration file is located.

```
SGOS#(config) load wccp-settings
```

```
SGOS#(config) wccp enable
```

Virtual IP Addresses

Virtual IP (VIP) addresses are addresses assigned to a system that are recognized by other systems on the network. Up to 255 VIPs can be configured on each ProxySG. They 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 s on the same subnet.

For information on creating an HTTPS Console, see "Creating and Editing Services" on page 129; for information on using VIPs with authentication realms, see Chapter 9: "Using Authentication Services" on page 271; to use VIPs with failover, see "Configuring Failover" on page 113.

To Create a VIP through the Management Console:

1. Select Configuration>Network>Advanced>VIPs.

The VIPs tab displays.




Figure 4-28: Network Advanced VIPs Tab

2. Click New.

The Add VIP dialog displays.

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 and forwarding to point to the origin content server's real IP address.

4. Click OK; click Apply.

The VIP address can now be used.

To Create a VIP through the CLI:

At the (config) command prompt, run the virtual IP address command:

```
SGOS#(config) virtual address ip_address
ok
```

To Delete a VIP through the CLI:

Note that VIP addresses are deleted silently. If you are using a VIP for a service, the service will no longer work once the VIP is deleted.

```
SGOS#(config) virtual no address ip_address
ok
```

To Clear All VIP Addresses in the System:

```
SGOS#(config) virtual clear
ok
```

To View All the VIPs in the System:

```
SGOS#(config) show virtual
Virtual IP addresses:
SGOS#(config) accelerated-pac path 10.25.36.47
10.9.36.47
10.25.36.48
10.25.36.47
```

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 "About Bridging" on page 75.

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 76.

Using a pool of IP addresses to provide redundancy and load balancing, Blue Coat migrates these IP addresses among a group of machines.

About Failover

Failover allows a second machine to take over if a first machine 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 keepalive 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 advertisements' 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

Before you begin, be aware that software bridges must already exist before you can use them to configure failover. For information on configuring bridges, see "Adapters" on page 71.

You also need to decide which machine is the master and which machines is 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 ProxySG Appliances have equal priority, the one with the highest physical address ranks higher.

To Configure Failover through the Management Console:

1. Go to Configuration>Network>Advanced>Failover.

The Failover tab displays.

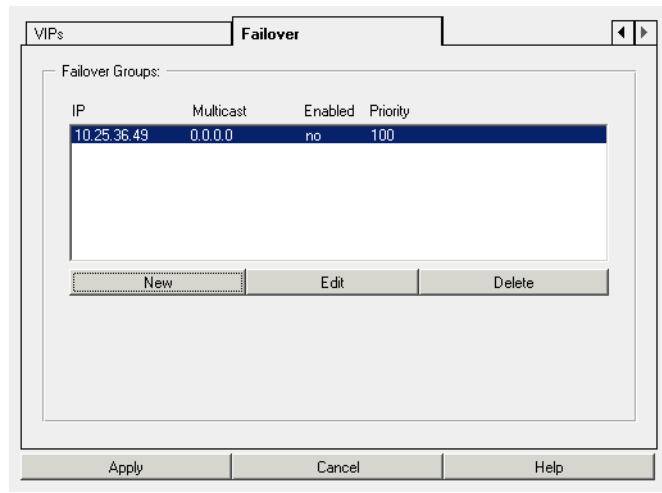


Figure 4-29: Network Advanced Failover Tab

2. Click New.

The Add Failover Group dialog displays.

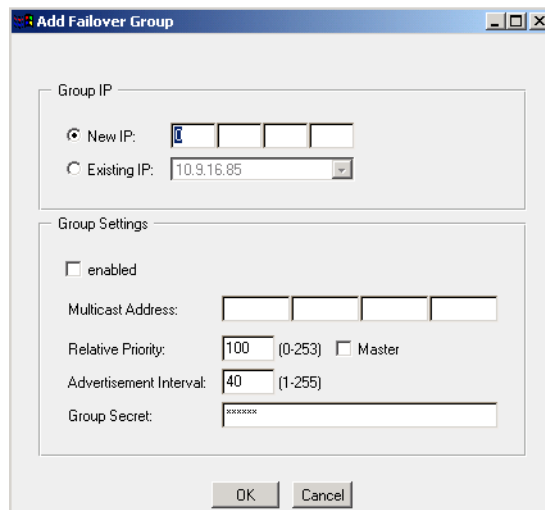


Figure 4-30: Add Failover Group Dialog

3. In the Add Failover Group dialog that appears, fill in the fields as appropriate:
 - 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.
 - The enabled option specifies whether this group is active or inactive. Select enabled to enable the failover group.
 - 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 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.

- ❑ 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.
- ❑ (Optional) Advertisement Interval refers to the length of time between advertisements sent by the group master. The default is 40 seconds. Once the group master has failed, the slave with the highest priority takes over (after approximately three times the interval value). The failover time of the group can be controlled by setting this value.
- ❑ (Optional, but recommended) Group Secret refers to a password shared only with the group.

4. Click OK; click Apply.

To Configure Failover through the CLI:

1. At the `(config)` command prompt, enter the following commands:

```
SGOS#(config) failover
SGOS#(config failover) create group_address

    The IP address does not have to exist.

SGOS#(config failover) edit group_address
SGOS#(config failover group_address) multicast-address multicast_address
SGOS#(config failover group_address) master
SGOS#(config failover group_address) priority number
SGOS#(config failover group_address) interval seconds
SGOS#(config failover group_address) secret secret
-or-
SGOS#(config failover group_address) encrypted-secret encrypted_secret
SGOS#(config failover group_address) enable
```

where:

<i>group_address</i>	Refers to the IP address or VIP address that is monitored by this group. Once the group has been named, you cannot change the name. To change the name, you must delete the group and start over.
<i>multicast-address</i> <i>multicast_address</i>	Refers to a multicast address where the master sends the keepalives (advertisements) to the slave systems.
<i>master</i>	(Optional) Identifies the system to be used as the master.
<i>no</i>	Negates these settings: <i>multicast-address</i> , <i>priority</i> , <i>interval</i> , <i>secret</i> , and <i>master</i> .
<i>priority number</i>	(Optional) Refers to the rank of slave systems. The range is from 1 to 254. (The master system, the one whose IP address matches the group address, gets 255.) Output of <code>show config</code> and <code>show failover</code> might differ when the master system is also the holder of the physical IP address.
<i>interval seconds</i>	(Optional) Refers to the time between advertisements from the master to the multicast address. The default is 40 seconds. Entering <code>no interval</code> resets the interval to the default time of 40 seconds.

<code>secret secret</code>	(Optional but recommended) Refers to a password shared only with the group. You can create a secret, which then is hashed, or you can provide an encrypted secret.
<code>-or-</code>	
<code>encrypted-secret</code>	
<code>encrypted_secret</code>	
<code>enable disable</code>	Enables or disables failover on the ProxySG.

2. (Optional) View the results.

```
SGOS#(config) show failover configuration group_address
Failover Config
Group Address: 10.25.36.47
Multicast Address      : 224.1.2.3
Local Address          : 10.9.17.159
Secret                 : none
Advertisement Interval: 40
Priority                : 100
Current State          : DISABLED
Flags                  : V M
```

Three flags exist, set as you configure the group.

V—Specifies the group name is a virtual IP address.

R—Specifies the group name is a physical IP address

M—Specifies this machine can be configured to be the master if it is available

3. (Optional) You can view Failover Group Statistics

These are all integers/counters that count various events.

```
SGOS#(config) show failover statistics
Failover Statistics
Advertisements Received      : 0
Advertisements Sent         : 194
States Changes               : 2
Bad Version                  : 0
Bad Packet                   : 0
Bad Checksum                 : 0
Packet Too Short             : 0
Bad Packet Header            : 0
Invalid Group                : 0
```

Viewing Statistics

To view statistics on failover, see "Failover Statistics" on page 857

TCP-IP Configuration

Use the TCP-IP configuration options to enhance the performance and security of the ProxySG. Except for IP Forwarding (see "IP Forwarding" on page 201), these commands are only available through the CLI.

- **RFC-1323:** Enabling RFC-1323 support enhances the high-bandwidth and long-delay operation of the ProxySG over very high-speed paths, ideal for satellite environments.
- **TCP NewReno:** Enabling TCP NewReno support improves the fast recovery of the ProxySG.

- **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 "Viewing the TCP-IP Configuration" on page 120.

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 only be configured through the CLI, and is enabled by default.

To Enable or Disable RFC-1323 Support through the CLI:

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 disabled by default.

To Enable or Disable TCP NewReno Support through the CLI:

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 through the CLI:

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 "Attack Detection" on page 95.

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 through the CLI:

At the `(config)` command prompt, enter the following command:

```
SGOS#(config) tcp-ip icmp-timestamp-echo {enable | disable}
```

TCP Window Size

Adjusting the TCP window-size regulates the amount of unacknowledged data that the ProxySG receives before sending an acknowledgement.

To Configure the TCP Window Size through the CLI:

At the `(config)` command prompt, enter the following command:

```
SGOS#(config) tcp-ip window-size window_size
```

where `window_size` indicates the number of bytes allowed before acknowledgement (the value must be between 8192 and 4194304).

PMTU Discovery

PMTU (Path Maximum Transmission Unit) is a mechanism designed to discover the largest packet size sent that is not fragmented anywhere along the path between two communicating ProxySG Appliances that are not directly attached to the same link. A ProxySG doing 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. The first ProxySG then reduces the PMTU size and re-transmits the transmissions.

The discovery period temporarily ends when the ProxySG's estimate of the PMTU is low enough that its packets can be delivered without fragmentation or when the ProxySG stops setting the `Do-Not-Fragment` bit. Five minutes later (this value is configurable), rediscovery is used to see if the transmittable packet size has changed.

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.

PMTU is disabled by default.

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.

To Configure PMTU Discovery through the CLI:

Note: PMTU discovery can only be configured through the CLI. It is not available through the Management Console.

At the `(config)` command prompt, enter the following commands:

```
SGOS#(config) tcp-ip pmtu-discovery enable | disable
SGOS#(config) tcp-ip pmtu-discovery expire-period seconds
SGOS#(config) tcp-ip pmtu-discovery probe-interval seconds
```

where

<code>tcp-ip</code>	<code>enable disable</code>	Allows you to enable PMTU discovery. The default is disabled.
<code>pmtu-discovery</code>		
	<code>expire-period</code> <code>seconds</code>	Determines the time, in seconds, when PMTU rediscovery takes place after receiving the ICMP Host Unreachable - Needs Fragmentation error message. The default is 600 seconds.
	<code>probe-interval</code> <code>seconds</code>	Determines the time, in seconds, when the next PMTU rediscovery takes place following a previous consecutive successful expansion of the PMTU value. The default is 120 seconds.

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:        enabled
PMTU expiration period:    600 seconds
PMTU probe interval:       120 seconds
TCP window size:           65535 bytes
```

Chapter 5: *Managing Port Services*

This chapter describes port services that are configurable on the ProxySG. These services run on the ProxySG, and include Management Consoles such as HTTPS, HTTP, SSH, and Telnet Consoles, and application proxies such as Instant Messenger (IM), SOCKS, FTP, MMS, and RTSP, HTTP and HTTPS.

Other proxy services, like ICAP and Websense, are remote to the ProxySG and are discussed in Chapter 11: “External Services” on page 399.

This chapter discusses

- "Managing Multiple Management Consoles"
- "Creating and Editing Services"

This chapter does not discuss configuration of some of the port services that are enabled here. The following are discussed in Chapter 6: “Configuring Proxies” on page 149:

- FTP Proxy
- HTTP Proxy
- SOCKS Proxy
- Shell Proxies (Telnet)

Section A: Managing Multiple Management Consoles

Section A: Managing Multiple Management Consoles

The ProxySG ships with a number of already existing consoles designed to manage the system and communication with the system:

- **HTTP and HTTPS Consoles:** These consoles are designed to allow you access to the ProxySG. The HTTPS Console is created and enabled; 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 to gain access to the ProxySG through the CLI with your SSH service.
- **Telnet Console:** This console is created but is disabled by default because of security concerns. You must enable the service before you can access the ProxySG through a Telnet client (not recommended).

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 box as well as any of the ProxySG's virtual IP (VIP) addresses. The default is HTTPS over port 8082.

The ProxySG ships with an HTTPS Console already created and enabled. You do not need to create other HTTPS Consoles unless you need them for other purposes.

An HTTPS Console and an HTTPS service are not the same. The HTTPS Console is meant only for accessing the ProxySG. An HTTPS service is meant to allow secure access to other systems.

Creating a new HTTPS Console port requires three steps, discussed more fully in the following sections:

- Selecting a keyring (a keypair and a certificate that is stored together)
- Selecting an IP address and port on the system that the service will use, including virtual IP addresses
- Putting the keyring and service together into an HTTPS Console

Selecting a Keyring

The ProxySG ships with a default keyring that can be reused with each HTTPS service that you create. You can also create your own keyrings for other purposes.

To use the default keyring, simply accept the default keyring through the Management Console, or, if you're using the CLI, enter `default` for the keyring ID when using the `services https-console create` command.

Section A: Managing Multiple Management Consoles

Note: When using certificates for the HTTPS Console or for HTTPS termination services that are issued by Certificate Signing Authorities that are not well-known, see "Troubleshooting Certificate Problems" on page 230.

If you get "host mismatch" errors or if the security certificate is called out as invalid, you need to create a different certificate and use it for the HTTPS Console.

For information on creating a keypair and a certificate to make a keyring, see "Configuring HTTPS Termination" on page 207.

Selecting an IP Address

You can use any IP address on the ProxySG for the HTTPS Console service, including virtual IP addresses. To create a virtual IP address, see "Virtual IP Addresses" on page 112.

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 through the Management Console:

1. Select Configuration>Services>Service Ports.

The Service Ports tab displays.

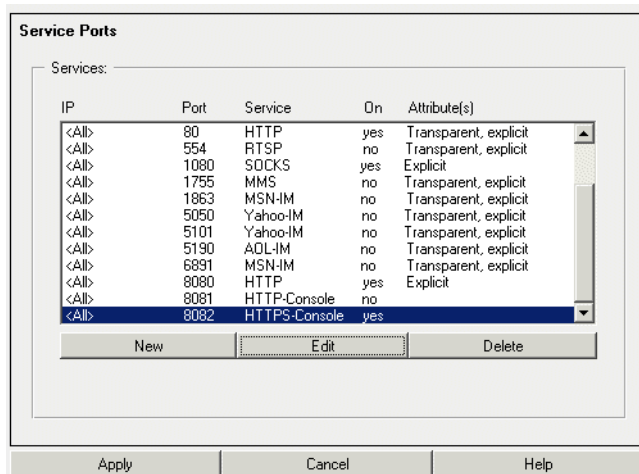


Figure 5-1: Service Ports Tab

2. Do one of the following:
 - To create a new HTTPS Console port service, click New; the Add Service dialog appears. Select HTTPS-Console from the Protocol drop-down list.
 - To edit an existing HTTPS Console port service, highlight the HTTPS Console and click Edit; the Edit Service dialog appears.

Section A: Managing Multiple Management Consoles

Continue with the next step.

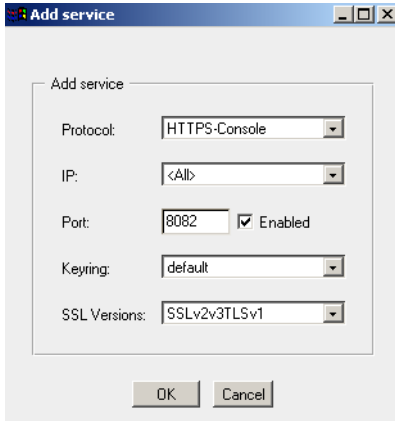


Figure 5-2: HTTPS-Console Add Service Dialog

3. The default IP address value is <All>. To limit the service to a specific IP address, select the IP address from the drop-down list. It must already exist.
4. Identify the port you want to use for this service.
5. In the Keyring drop-down list, select any already created keyring that is on the system. The system ships with a default keyring that can be reused for each HTTPS service.

Note: The configuration-passwords-key keyring that shipped with the ProxySG does not contain a certificate and cannot be used for HTTPS Consoles.

6. (Optional) In the SSL Versions drop-down list, select the version that you want to use for this service. The default is SSL v2/v3 and TLS v1.
7. Click OK; click Apply.

Note: For information on creating keyrings and client certification lists, see "Configuring HTTPS Termination" on page 207.

To Create Another HTTPS Console Port Service through the CLI:

1. At the (config) command prompt, enter the following commands:


```
SGOS#(config) services
SGOS#(config services) https-console
SGOS#(config services https-console) create [ip_address:]port [keyring id]
```

If you do not specify a keyring, the default is used.

```
SGOS#(config services https-console) attribute cipher-suite ip_address:port
```
2. (Optional) View the results:

Section A: Managing Multiple Management Consoles

```

SGOS#(config services https-console) view
Port:      8082      IP: 0.0.0.0      Type: https-console
Keyring: default
Properties: explicit, enabled
Cipher suite:
RC4-MD5:RC4-SHA:DES-CBC3-SHA:DES-CBC3-MD5:RC2-CBC-MD5:RC4-64-MD5:DES-CBC-SHA
:DES-CBC-MD5:EXP1024-RC4-MD5:EXP1024-RC4-SHA:EXP1024-RC2-CBC-MD5:EXP1024-DES
-CBC-SHA:EXP-RC4-MD5:EXP-RC2-CBC-MD5:EXP-DES-CBC-SHA:
+SSLv2:+SSLv3+LOW:+SSLv2+LOW:+EXPOHTTP

```

Note: To create client-certification lists and keyrings, see "Configuring HTTPS Termination" on page 207. To set the cipher-suite to the ciphers you want to use, see "Changing the Cipher Suites of the SSL Client" on page 233.

HTTP Console

The HTTP Console is meant to allow you to access the ProxySG if you require a less secure environment. The default HTTP Console is already configured; you must enable it before it can be used.

You can create and use more than one HTTP Console as long the IP address and the port do not match the existing HTTP Console settings.

To Create or Edit an HTTP Console Port Service through the Management Console:

1. Select Configuration>Services>Service Ports.

The Service Ports tab displays.

2. Do one of the following:

- To create a new HTTP-Console port service, click New; the Add Service dialog appears. Select HTTP-Console from the Protocol drop-down list.
- To edit an existing HTTP-Console port service, highlight the HTTP-Console and click Edit; the Edit Service dialog appears.




Figure 5-3: HTTP-Console Add Service Dialog

In either case, continue with the next step.

Section A: Managing Multiple Management Consoles

3. The default IP address value is <All>. To limit the service to a specific IP address, select the IP address from the drop-down list. It must already exist.
4. Identify the port you want to use for this service.
5. Click OK; click Apply.

To Create or Edit an HTTP Console Port Service and Enable It through the CLI:

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) services
SGOS#(config services) http-console
SGOS#(config services http-console) create [ip_address:]port
```
2. (Optional) View the results:

```
SGOS#(config services http-console) view
Port:      8085 IP: 0.0.0.0           Type: http-console
Properties: enabled
```

SSH Console

The SSH Console is created and enabled by default. Only one SSH Console can exist on the ProxySG. If you inadvertently deleted the SSHv1 and SSHv2 host keys from the system at the same time, you automatically disabled the SSH Console and will have to enable the SSH Console after you create a host key.

For information on managing SSH, see "Configuring the SSH Console" on page 55.

To Edit an SSH Console Service through the Management Console:

1. Select Configuration>Services>Service Ports.
The Service Ports tab displays.
2. To edit the existing SSH-Console port service, highlight the SSH-Console and click Edit.
The Edit Service dialog appears.

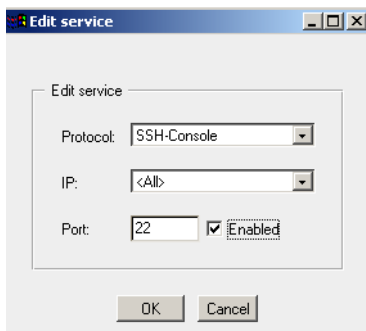


Figure 5-4: SSH-Console Add Service Dialog

3. The default IP address value is all. To limit the service to a specific IP address, select the IP address from the drop-down list.

Section A: Managing Multiple Management Consoles

4. In the Port field, specify a port number; select Enable.
5. Click OK; click Apply.

To Create an SSH Port Service through the CLI:

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) services
SGOS#(config services) ssh-console
SGOS#(config services ssh-console) create [ip_address:]port
SGOS#(config services ssh-console) enable [ip_address:]port
```

2. (Optional) View the results:

```
SGOS#(config services ssh-console) view
Port:      22      IP: 0.0.0.0 Type: ssh-console
Properties: enabled
```

Telnet Console

The Telnet Console allows you to connect to and manage the ProxySG using the Telnet protocol. Remember that Telnet is an insecure protocol that should not be used in insecure conditions. By default, only SSH is created and enabled.

Blue Coat Systems recommends against using Telnet because of the security hole it creates.

Note: If you do enable the Telnet Console, be aware that you cannot use Telnet everywhere in the CLI. Some modules, such as SSL, respond with the error message:

```
Telnet sessions are not allowed access to ssl commands.
```

To Create or Edit a Telnet Console Port Service through the Management Console:

Before you begin, make sure that no Telnet service exists on the default telnet port (23). If it does exist, delete it and apply the changes before continuing. If you also want a Telnet service, you can re-create it later, being sure to use a different port. For information on the Telnet service, see "Telnet Shell Proxy Service" on page 145.

1. Select Configuration>Services>Service Ports.

The Service Ports tab displays.

2. Do one of the following:

- To create a new Telnet-Console port service, click New; the Add Service dialog appears. Select Telnet-Console from the Protocol drop-down list.
- To edit an existing Telnet-Console port service, highlight the Telnet-Console and click Edit; the Edit Service dialog appears.

In either case, continue with the next step.

Section A: Managing Multiple Management Consoles

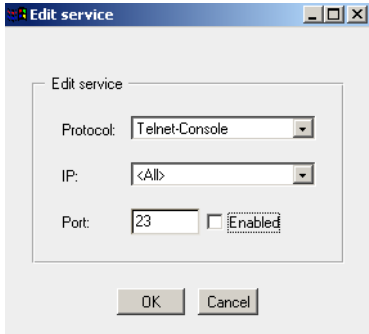


Figure 5-5: Telnet Console Edit Service Dialog

3. Select Telnet protocol from the drop-down list.
4. The default IP address value is all. To limit the service to a specific IP address, select the IP address from the drop-down list.
5. In the Port field, specify a port number; 23 is the default.

Note: If you want to use the Telnet shell proxy and retain the Telnet Console as well, you must change the port number on one of them. Only one service is permitted on a port. For more information on the Telnet shell proxy, see "Telnet Shell Proxies" on page 195.

6. Select Enabled.
7. Click OK; click Apply.

To Create or Edit a Telnet Port Service through the CLI:

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) services
SGOS#(config services) telnet-console
SGOS#(config services telnet-console) create [ip_address:]port
```
2. (Optional) View the results.

```
SGOS#(config services telnet-console) view
Port:      23      IP: 0.0.0.0      Type: telnet-console
Properties: enabled
```

Section B: Creating and Editing Services

Section B: Creating and Editing Services

Port services define attributes for ports on which the ProxySG listens for Web requests. Each service applies to all IP addresses, or can be limited to a specific address.

You can create as many services as you require, keeping in mind that every newly created service uses up resources.

Note: When multiple non-wildcard services are created on a port, all of them must be of the same service type. So you can have HTTP listening on a given port on some subset of adapter interfaces (or VIPs), but you can't have HTTP on one adapter interface and HTTPS on a different adapter interface with both using the same port.

Also note that wildcard services and non-wildcard services cannot both exist at the same time on a given port.

The following table lists the available ProxySG services, including their attributes and default status. The defaults are for a new ProxySG. If you have an upgraded appliance, the settings do not change.

Table 5.1: Proxy Port Services

Proxy Service	Default Port	Status	Configuration Discussed
DNS-Proxy	53 (both transparent and explicit)	Disabled	"DNS-Proxy"
EPMapper	135 (both transparent and explicit)	Disabled	"Endpoint Mapper Proxy"
FTP	21 (transparent and explicit)	Disabled	"FTP"
HTTP	80 (transparent and explicit) 8080 (explicit only)	Enabled	"HTTP"
HTTP-Console	8081	Disabled	"HTTP Console"
HTTPS		Disabled	"HTTPS"
HTTPS-Console	8082	Enabled	"HTTPS Console (Secure Console)"
MSN-IM	1863 (transparent and explicit) and 6891 (transparent and explicit)	Disabled	"Instant Messaging Protocols"
Yahoo-IM	5050 (transparent and explicit) and 5101 (transparent and explicit)	Disabled	"Instant Messaging Protocols"
AOL-IM	5190 (transparent and explicit)	Disabled	"Instant Messaging Protocols"
MMS	1755 (transparent and explicit)	Disabled	"Streaming Protocols"
RTSP	554 (transparent and explicit)	Disabled	"Streaming Protocols"
SOCKS	1080	Disabled	"SOCKS"
SSH-Console	22	Enabled	"SSH Console"
TCP-Tunnel		Not Created	"TCP Tunneling"

Section B: Creating and Editing Services

Table 5.1: Proxy Port Services

Proxy Service	Default Port	Status	Configuration Discussed
Telnet-Console	23	Not Created	"Telnet Console"
Telnet shell proxy	23	Disabled	"Telnet Shell Proxy Service"

Note: If HTTP is configured to be explicit, Internet Explorer version 6.0 users accessing FTP sites over HTTP must disable the browser setting Enable folder view for FTP sites. To access this attribute in Internet Explorer, select Tools>Internet Options, click the Advanced tab, deselect Enable folder view for FTP sites, and click OK.

About Service Attributes

The service attributes define the parameters the ProxySG uses for a particular service. In addition to configuring the default port services, you can create additional ports and define their attributes.

Note: For all service types except HTTPS, a specific listener cannot be posted on a port if the same port has a wildcard listener of any service type already present.

The following table describes the attributes; however, depending on the protocol, not all attributes are available.

Table 5.2: Attributes

Attribute	Description
Explicit	Enables or disables explicit attribute for the port. (Explicit allows connections to a ProxySG IP address.) Note: If DNS redirection is used to direct traffic to the ProxySG, the explicit flag on its services must be enabled, as these connections will be routed through DNS to the ProxySG's IP address.
Transparent	Enables or disables transparent-proxy attribute for port. (This allows connections to any IP address other than those belonging to the ProxySG.)
Authenticate-401	All transparent and explicit requests received on the port always use transparent authentication (cookie or IP, depending on the configuration). This is especially useful to force transparent proxy authentication in some proxy-chaining scenarios.
Send client IP	Enables or disables sending of client's IP address instead of the ProxySG's IP address. For more information, see the section on tracking client IP addresses using server-side transparency.

Note: If you use the CLI to create a service, specify 0.0.0.0 to define that the service listens on all IP addresses; specify the individual IP address to limit the service to one IP address.

Section B: Creating and Editing Services

DNS-Proxy

When a DNS-Proxy service is enabled, 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 does a lookup of the DNS cache to determine if requests can be answered. If yes, the ProxySG responds. If not, the DNS-Proxy forwards the request to the DNS server list configured on the ProxySG. (To configure the DNS server list, see Configuration>Network>DNS.)

Note: The ProxySG is not a DNS server. It does not perform zone transfers, and recursive queries are forwarded to other name servers.

Through policy, you can configure the list of resolved domain names (the *resolving name list*) the DNS-Proxy 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.

To Create or Edit a DNS-Proxy Service through the Management Console:

1. Select Configuration>Services>Service Ports.

The Service Ports tab displays.

2. Click New or Edit; the Add (or Edit) Service dialog appears.
3. Select DNS-Proxy from the Protocol drop-down list.

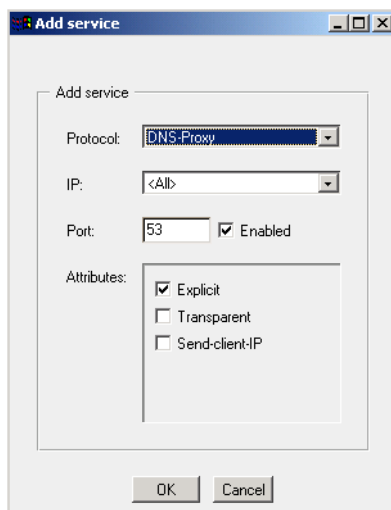


Figure 5-6: DNS-Proxy Add Service Dialog

4. The default IP address value is All. To limit the service to a specific IP address, select the IP address from the drop-down list.
5. In the Port field, 53 displays; you can change it to any unused port.

Section B: Creating and Editing Services

6. Select Enabled.
7. In the Attributes field, select Transparent, Explicit, Send-client-IP (spoofing), or all three. Explicit is the default.

Note: The send-client-IP attribute allows the ProxySG to pretend to be a client, allowing the origin content server to see the client's IP address. If an alternate path exists for traffic returning from the Internet to the client, the Send-client-IP attribute does not work.

8. Click OK; click Apply.

To Create or Edit a DNS-Proxy Service through the CLI:

1. At the (config) command prompt, enter the following commands to set the value returned to the client before configuring the DNS service:

```
SGOS#(config) services
SGOS#(config services) dns
SGOS#(config services dns) create ip_address:port
```

2. If you do not need to change the defaults, you have completed the procedure. To change the attributes, enter the following command:

```
SGOS#(config services dns) attribute {explicit | transparent | send-client-ip}
{enable | disable} [ip_address:] port
```

where:

```
attribute  explicit |
           transparent
           |send-client-ip
           enable
           [ip_address:]port
```

Give the DNS proxy explicit and transparent attributes, and create IP spoofing (where the ProxySG pretends to be a client so the OCS can see the client's IP address).

Note: The send-client-IP attribute allows the ProxySG to pretend to be a client, allowing the origin content server to see the client's IP address. If an alternate path exists for traffic returning from the Internet to the client, the Send-client-IP attribute does not work.

```
enable    [ip_address:]port
```

Enable the new DNS proxy.

3. (Optional) View the results:

```
SGOS#(config services dns) view
Port:      53      IP: 0.0.0.0      Type: dns
Properties: transparent, explicit, enabled
Port:      54      IP: 0.0.0.0      Type: dns
Properties: transparent, enabled
```

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 *Blue Coat ProxySG Content Policy Language Guide*.)

Each name resolving list entry contains a domain-name matching pattern. The matching rules are:

Section B: Creating and Editing Services

- `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*).

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)
```

Note: You can also create a resolving name list using VPM. For more information on using the DNS-Proxy layer in VPM, see "Web Content Policy Layer Reference" on page 472.

Endpoint Mapper Proxy

The Endpoint Mapper proxy accelerates Microsoft RPC traffic (applications that use dynamic port numbers) between branch and main offices, automatically creating TCP tunnels to ports where RPC services are running. The Endpoint Mapper proxy can be used in both explicit and transparent mode.

Endpoint Mapper works by intercepting and tunnelling RPC traffic in the branch office (downstream proxy). The tunneled data is compressed and forwarded to the main office (upstream proxy). The upstream proxy, using SOCKS gateways, decompresses the traffic and forwards it to RPC server. (For information on SOCKS compression, see "Understanding SOCKS Compression" on page 188.)

Note: Only Microsoft RPC version 5.0 is supported. Traffic for unsupported Microsoft RPC versions is passed through the ProxySG without processing.

For information on using SOCKS proxy and EPMapper together, see the *Blue Coat Edge Deployment Guide*.

By default, the service is created but not enabled.

To Create or Edit Endpoint Mapper Service through the Management Console:

1. Select Configuration>Services>Service Ports.
The Service Ports tab displays.
2. Click New or highlight the existing Endpoint Mapper proxy service and click Edit; the Add (or Edit) Service dialog appears.

Section B: Creating and Editing Services

3. Select EndpointMapper from the Protocol drop-down list.

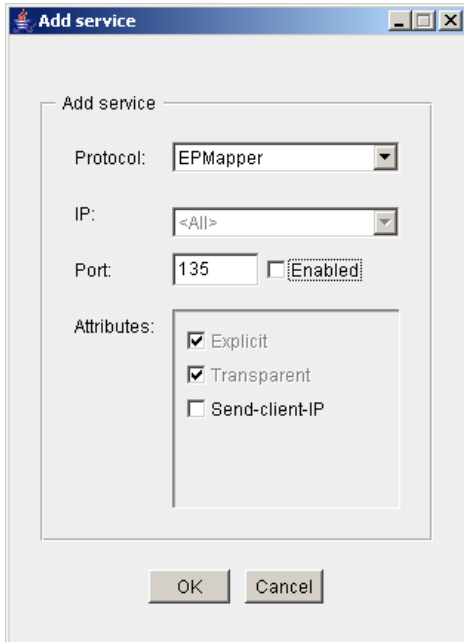


Figure 5-7: Endpoint Mapper Edit Service Dialog

4. The default IP address value is All. It cannot be changed.
5. In the Port field, 135 displays. Port 135 is the standard port for Microsoft RPC traffic.
6. Select Enabled.
7. In the Attributes field, select Send-client-IP, if necessary. Explicit and Transparent attributes are not user configurable.

Note: The send-client-IP attribute allows the ProxySG to pretend to be a client, allowing the origin content server to see the client's IP address. If an alternate path exists for traffic returning from the Internet to the client, the Send-client-IP attribute does not work.

8. Click OK; click Apply.

To Create or Edit an Endpoint Mapper Proxy Service through the CLI:

1. At the (config) command prompt, enter the following commands to create a new Endpoint Mapper proxy service. If you want to edit the existing Endpoint Mapper proxy, skip to step 2.:

```
SGOS#(config) services  
SGOS#(config services) epmapper  
SGOS#(config services epmapper) create port
```

2. To enable the Endpoint Mapper proxy service or enable the send-client-ip attribute, enter the following commands:

Section B: Creating and Editing Services

```
SGOS#(config services epmapper) enable port
SGOS#(config services epmapper) attribute send-client-ip {enable | disable}
port
```

where:

```
attribute send-client-ip enable port
```

Enable sending the client's IP address instead of the ProxySG's IP address.

Note: If an alternate path exists for traffic returning from the Internet to the client, the Send-client-IP attribute does not work.

```
enable port
```

Enable the new Endpoint Mapper proxy. Port 135 is the standard port for Microsoft RPC traffic.

3. (Optional) View the results:

```
SGOS#(config services epmapper) view
Port:      135      IP: 0.0.0.0      Type: epmapper
Properties: transparent, explicit, disabled
```

FTP

To configure the native FTP proxy, see "Configuring the FTP Proxy" on page 152.

To Create or Edit an FTP Port Service through the Management Console:

1. Select Configuration>Services>Service Ports.
The Service Ports tab displays.
2. Click New or Edit; the Add (or Edit) Service dialog appears.
3. Select FTP from the Protocol drop-down list.

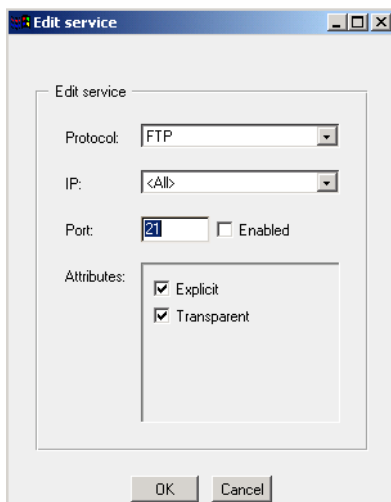


Figure 5-8: FTP Edit Service Dialog

Section B: Creating and Editing Services

4. The default IP address value is all. To limit the service to a specific IP address, select the IP address from the drop-down list.
5. In the Port field, specify a port number; select the Enabled checkbox.
6. In the Attributes field, both Explicit and Transparent are selected. You can de-select one of them if necessary
7. Click OK; click Apply.

To Create an FTP service through the CLI:

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) services
SGOS#(config services) ftp
SGOS#(config services ftp) create [ip_address:]port
SGOS#(config services ftp) attribute passive-mode {enable | disable}
-or-
SGOS#(config services ftp) attribute {explicit | transparent} {enable |
disable} [ip_address:]port
```

2. (Optional) View the results.

```
10.9.17.159 - Blue Coat SG3000#(config services ftp) view
Port:      21          IP: 0.0.0.0          Type: ftp
Properties: transparent, enabled, passive-allowed
```

HTTP

Two HTTP services exist by default and are enabled, one with explicit and transparent attributes on port 80 and one with explicit attributes on port 8080. You can change the attributes or create other HTTP ports if needed.

To Create or Edit an HTTP Port Service through the Management Console:

1. Select Configuration>Services>Service Ports.
The Service Ports tab displays.
2. Click New or highlight the service and click Edit; the Add (or Edit) Service dialog appears.
3. Make sure HTTP is selected from the Protocol drop-down list.

Section B: Creating and Editing Services

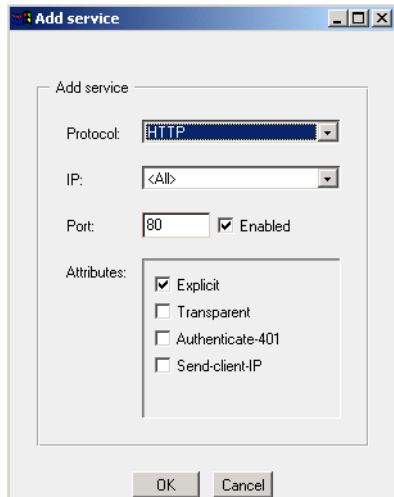


Figure 5-9: HTTP Edit Service Dialog

4. The default IP address value is all. To limit the service to a specific IP address, select the IP address from the drop-down list.
5. In the Port field, specify a port number; be sure Enabled is selected.
6. In the Attributes field, select all that apply: Explicit, Transparent, Authenticate-401, or Send-client-IP.

Note: The send-client-IP attribute allows the ProxySG to pretend to be a client, allowing the origin content server to see the client's IP address. If an alternate path exists for traffic returning from the Internet to the client, the Send-client-IP attribute does not work.

7. Click OK; click Apply.

To Create an HTTP Service through the CLI:

Two HTTP services exist and are enabled on the ProxySG. If you need to create another at a different port in addition to the services already existing on the system, complete the following steps:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) services
SGOS#(config services) http
SGOS#(config services http) create [ip_address:]port
SGOS#(config services http) attribute {authenticate-401 | explicit |
send-client-ip | transparent} {enable | disable} [ip_address:]port
-or-
SGOS#(config services http) attribute {connect | head} {enable | disable
{drop | error}} [ip_address:]port
```

Note: The send-client-IP attribute allows the ProxySG to pretend to be a client, allowing the origin content server to see the client's IP address. If an alternate path exists for traffic returning from the Internet to the client, the Send-client-IP attribute does not work.

Section B: Creating and Editing Services

To view the results:

```
SGOS#(config services http) view
Port:      8080      IP: 0.0.0.0          Type: http
Properties: explicit, enabled
Port:      80       IP: 0.0.0.0          Type: http
Properties: transparent, explicit, enabled
```

HTTPS

The HTTPS service is not configured or enabled by default when the ProxySG ships. You can configure and use multiple HTTPS services.

To Create an HTTPS Port Service through the Management Console:

1. Select Configuration>Services>Service Ports.
The Service Ports tab displays.
2. Click New; the Add Service dialog appears.
3. Select HTTPS from the Protocol drop-down list.

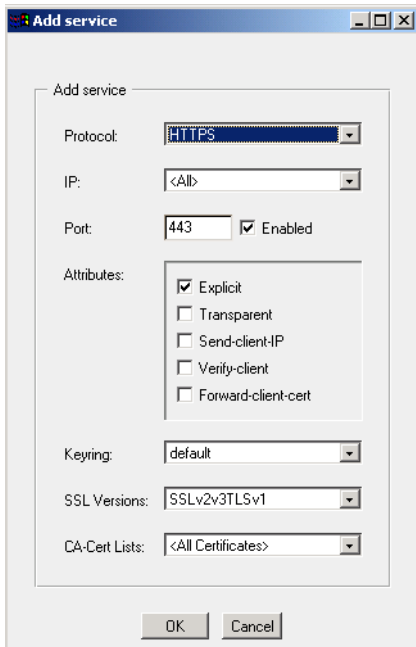


Figure 5-10: HTTPS Add Service Dialog

4. To select or add an IP address, do one of the following:
 - ❑ To select a local address, specify a real IP address from the IP drop-down list. All is not a selection option.
 - ❑ To add a non-local IP address, first select the Transparent attribute, then enter a non-local IP address that is not bound to the ProxySG.

Section B: Creating and Editing Services

5. In the Port field, specify a port number; select Enable.
6. In the Attributes field, select all that apply: Explicit, Transparent, Send-client-IP, Verify-client, or Forward-client-cert. The send-client-IP attribute allows the ProxySG to pretend to be a client, allowing the origin content server to see the client's IP address.

Note: If the ProxySG HTTPS service is configured to require a client certificate (if the Verify-client checkbox is selected), information from the client certificate is extracted and put into a header that is included in the request when it is forwarded to the OCS.

The name of the header is `Client-Cert`. The header contains the certificate serial number, subject, validity dates and issuer (all as name=value) pairs. The actual certificate itself is not forwarded.

7. In the Keyring drop-down list, select any already-created keyring that is on the system. The system ships with a default keyring that can be reused for each HTTPS service. Keep in mind that the default certificate associated with the default keyring is self-signed and might not be trusted by all clients.

Note: The configuration-passwords-key keyring that shipped with the ProxySG does not contain a certificate and cannot be used for HTTPS services.

8. In the SSL Versions drop-down list, select the version that you want to use for this service. The default is SSL v2/v3 and TLS v1.
9. In the CA-Cert Lists drop-down list, select the list (already created) for the HTTPS service to use.
10. Click OK; click Apply.

Note: To create client-certification lists and keyrings, see "Configuring HTTPS Termination" on page 207.

To Create an HTTPS Service through the CLI:

1. At the `(config)` command prompt, enter the following commands:

```
SGOS#(config) services
SGOS#(config services) https
SGOS#(config services https) create ip_address:port keyring
SGOS#(config services https) attribute ccl list_name ip_address:port
-or-
SGOS#(config services https) attribute cipher-suite ip_address:port
-or-
```

Section B: Creating and Editing Services

```
SGOS#(config services https) attribute {forward-client-cert | send-client-ip
| verify-client} {enable | disable} ip_address:port
-or-
SGOS#(config services https) attribute ssl-protocol-version {ssl2 | ssl3 |
tlsv1 | ssl2v3| ssl2tlsv1 | ssl3tlsv1 | ssl2v3tlsv1} ip_address:port
```

Note: If the ProxySG HTTPS service is configured to require a client certificate (if you enter the command `SGOS#(config services https) attribute verify-client enable ip_address:port`), information from the client certificate is extracted and put into a header that is included in the request when it is forwarded to the OCS.

The name of the header is `Client-Cert`. The header contains the certificate serial number, subject, validity dates and issuer (all as name=value) pairs. The actual certificate itself is not forwarded.

2. (Optional) View the results:

```
SGOS#(config services https) view
Port:      1000      IP: 10.9.17.159      Type: https
Keyring: default
Properties: explicit, enabled
SSL Protocol version: SSLv2v3TLsv1
CA Certificate List: not configured
Cipher suite:
RC4-MD5:RC4-SHA:DES-CBC3-SHA:DES-CBC3-MD5:RC2-CBC-MD5:RC4-64-MD5:DES-CBC-SHA
:DES-CBC-MD5:EXP1024-RC4-MD5:EXP1024-RC4-SHA:EXP1024-RC2-CBC-MD5:EXP1024-DES
-CBC-SHA:EXP-RC4-MD5:EXP-RC2-CBC-MD5:EXP-DES-CBC-SHA:+SSLv2:+SSLv3+LOW:+SSLv
2+LOW:+EXPO
```

Instant Messaging Protocols

Supported instant messaging (IM) services are present by default with the transparent and explicit attributes selected and listening on all IP addresses; none of them are enabled. Note that the explicit attribute is not user-configurable.

To Create or Enable an AOL, Yahoo, or MSN Port Service through the Management Console:

1. Select Configuration>Services>Service Ports.
The Service Ports tab displays.
2. Click New or highlight the service you want and select Edit; the Add (or Edit) Service dialog appears.
3. Select the IM service you want to create or edit from the Protocol drop-down list.
4. The default port is determined by the protocol:
 - AOL—Port 5190
 - Yahoo—Ports 5050 and 5101
 - MSN—1863 and 6891

Section B: Creating and Editing Services

5. Click OK; click Apply.

To Manage an Instant Messaging Service through the CLI:

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) services
SGOS#(config services) aol-im | msn-im | yahoo-im
SGOS#(config services protocol) create port
SGOS#(config services protocol) attribute send-client-ip {enable | disable}
port
```

Note: The send-client-IP attribute allows the ProxySG to pretend to be a client, allowing the origin content server to see the client's IP address. If an alternate path exists for traffic returning from the Internet to the client, the Send-client-IP attribute does not work.

2. (Optional) View the results:

```
SGOS#(config services aol-im) view
Port:      5190 IP: 0.0.0.0           Type: aol-im
Properties: transparent, explicit, enabled
SGOS#(config services aol-im) exit
SGOS#(config services) yahoo-im
SGOS#(config services yahoo-im) view
Port:      5050 IP: 0.0.0.0         Type: yahoo-im
Properties: transparent, explicit, enabled
```

Streaming Protocols

MMS and RTSP services are configured on the system, but are disabled by default. To enable the default MMS and RTSP service, follow the steps below.

To Enable an MMS or RTSP Port Service through the Management Console:

1. Select Configuration>Services>Service Ports.
The Service Ports tab displays.
2. Click New to create a new MMS or RTSP port service or highlight the existing service and click Edit.
The Add (or Edit) Service dialog appears.
3. Select MMS or RTSP from the Protocol drop-down list.
4. The default IP address value is All. To limit the service to a specific IP address, select the IP address from the drop-down list.
5. In the Port field, specify a port number; select Enabled.
6. In the Attributes field, select the attributes you want the service to have.
7. Click OK; click Apply.

Section B: Creating and Editing Services

To Enable an MMS or RTSP Service through the CLI:

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) services
SGOS#(config services) {mms | rtsp}
SGOS#(config services protocol) create [ip_address:]port
SGOS#(config services protocol) attribute {explicit | send-client-ip |
transparent} {enable | disable} [ip_address:]port
```

Note: The send-client-IP attribute allows the ProxySG to pretend to be a client, allowing the origin content server to see the client's IP address. If an alternate path exists for traffic returning from the Internet to the client, the Send-client-IP attribute does not work.

2. (Optional) View the results:

```
SGOS#(config services mms) view
Port:      1755      IP: 0.0.0.0      Type: mms
Properties: transparent, explicit, enabled
SGOS#(config services mms) exit
SGOS#(config services) rtsp
SGOS#(config services rtsp) view
Port:      554      IP: 0.0.0.0      Type: rtsp
Properties: transparent, explicit, enabled
```

SOCKS

By default, a SOCKS service is configured with explicit attribute on port 1080, but not enabled. You can create additional SOCKS services.

To enable a SOCKS port service, complete the steps below. To configure SOCKS gateway forwarding, see "SOCKS Gateway Configuration" on page 722.

Note: The version of SOCKS used is controlled through policy. For example, to use only SOCKSv5:

```
<proxy> client.protocol=socks
        ALLOW socks.version=4 deny
        DENY
```

To Create or Edit a SOCKS Port Service through the Management Console:

1. Select Configuration>Services>Service Ports.
The Service Ports tab displays.
2. Click New to create a new SOCKS service or select Edit to enable the existing service; the Add (or Edit) Service dialog appears.
3. Select SOCKS from the Protocol drop-down list.

Section B: Creating and Editing Services

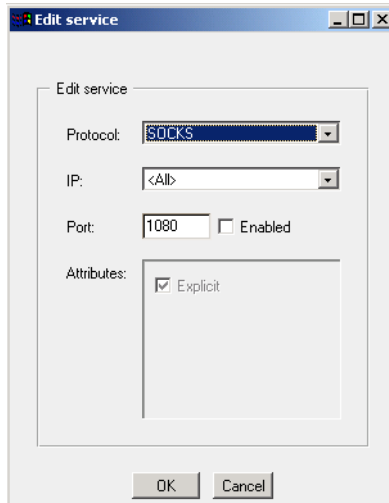


Figure 5-11: SOCKS Edit Service Dialog

4. The default IP address value is All. To limit the service to a specific IP address, select the IP address from the drop-down list.
5. In the Port field, specify a port number; select Enable.
6. Click OK; click Apply.

To Create a SOCKS Port Service through the CLI:

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) services
SGOS#(config services) socks
SGOS#(config services socks) create [ip_address:]port
SGOS#(config services socks) enable [ip_address:]port
```

2. (Optional) View the results:

```
SGOS#(config services socks) view
Port:      1080      IP: 10.25.36.48 Type: socks
Properties: explicit, enabled
```

TCP Tunneling

Tunneling, or port forwarding, is a way to forward TCP traffic. Any application protocol running over TCP can be tunneled using this service. Client-server applications carry out any authentication procedures just as they do when TCP tunneling is not involved.

SGOS uses a tcp:// scheme for tcp-tunnel transactions instead of HTTPS because SGOS does not actually know that it is HTTPS that is being tunneled.

You can use the SOCKS proxy in conjunction with TCP tunnels to compress and accelerate the tunneled traffic. For information on using the SOCKS proxy, see "Configuring a SOCKS Proxy" on page 188.

Section B: Creating and Editing Services

Both explicit and transparent TCP tunneling are supported. Which one you use depends on your needs.

Explicit TCP tunneling allows connections to one of the ProxySG's IP addresses.

Transparent TCP tunneling allows connections to any IP address other than those belonging to the ProxySG. TCP tunneling in transparent mode supports categorization as well as blocking of destination IP address, port, host, and domain.

Note: The TCP-Tunnel service does not support content filtering with Websense offbox or ICAP.

If you want to create a transparent TCP tunneling protocol, you can do so from either the CLI or the Management Console. When a TCP-Tunnel service is created, it is by default created as an explicit service and is also enabled automatically.

To Create a Transparent or Explicit TCP-Tunnel Port Service through the Management Console:

1. Select Configuration>Services>Service Ports.
The Service Ports tab displays.
2. Click New; the Add Service dialog appears.
3. Select TCP-Tunnel from the Protocol drop-down list.

The Add Service dialog displays.

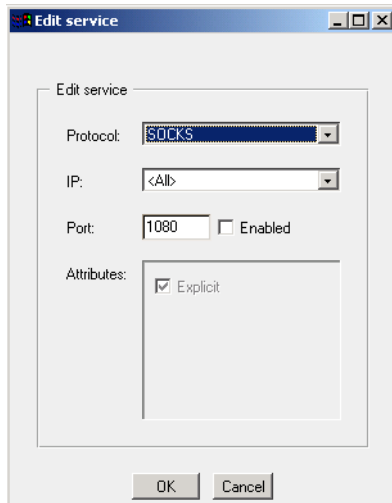


Figure 5-12: TCP-Tunnel Add Service Dialog

4. The default IP address value is All. To limit the service to a specific IP address, select the IP address from the drop-down list.
5. In the Port field, specify a port number; select Enabled.
6. If you are configuring a transparent TCP-Tunnel service, make sure Transparent is selected in the Attributes field; if you are configuring an explicit TCP-Tunnel service, make sure Explicit is selected.
7. Click OK; click Apply.

To Create a TCP-Tunnel Transparent or Explicit Port Service through the CLI:

1. At the (config) prompt, enter the following commands to create a transparent or explicit service:

```
SGOS#(config) services
SGOS#(config services) tcp-tunnel
SGOS#(config services tcp-tunnel) create [ip_address:]port
```

where *ip_address* is the IP address of the ProxySG (use 0.0.0.0 to indicate all available IP addresses), and *port* is the number of the port where you want the ProxySG to listen. You must choose a port that is not configured for any other service.

2. Enable the service to be transparent or explicit. By default, the port service is explicit.

```
SGOS#(config services tcp-tunnel) attribute {explicit | transparent} {enable
| disable} [ip_address:]port
```

3. (Optional) View the results.

```
SGOS#(config services tcp-tunnel) view
Port:      7080      IP: 0.0.0.0      Type: tcp-tunnel
Properties: transparent, explicit, enabled
```

If you created a transparent TCP-Tunnel service, you have completed the procedure. If you created an explicit TCP-Tunnel service, you must configure a forwarding destination port.

Configuring a Forwarding Destination Port through the CLI:

1. Create a forwarding destination port, where the ProxySG directs traffic.

```
SGOS#(config services tcp-tunnel) exit
SGOS#(config services) exit
SGOS#(config) forwarding
SGOS#(config forwarding) create host_alias ip_address tcp=port
```

2. (Optional) View the results:

```
SGOS#(config forwarding) view
Forwarding Groups: (* = host unresolved)
No forwarding groups defined.
Individual Hosts: (* = host unresolved)
Host_Alias 10.25.36.47 tcp=port_number
```

Telnet Shell Proxy Service

On a new system, Telnet proxy service is configured and disabled on port 23. On an upgrade, Telnet proxy service is not created.

To Enable or Create a Telnet Proxy Service through the Management Console

Important: If you want to use Telnet to manage the ProxySG, create a Telnet-Console rather than a Telnet service. The Telnet service allows you to use Telnet for outbound connections, and the ProxySG functions as Shell proxy in that situation. For more information on the Telnet-Console, see "Telnet Console" on page 127.

1. Select Configuration>Services>Service Ports.
2. Click New if you are creating a new Telnet service; highlight the Telnet service and click Edit if you are enabling an existing Telnet service;

The Add or Edit Service dialog appears.

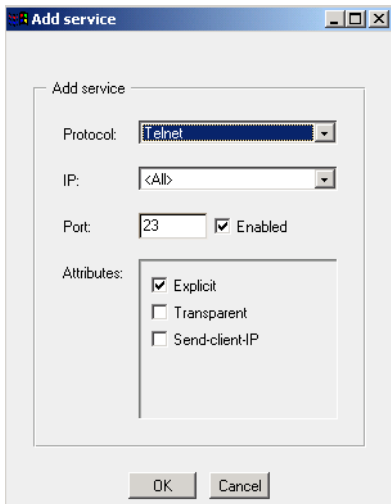


Figure 5-13: Creating a Telnet Service

3. In the Protocol drop-down list, select Telnet.
4. The default IP address value is all. To limit the service to a specific IP address, select the IP address from the drop-down list.
5. In the Port field, specify a port number; select Enable. Port 23 is the default.

Important: You can have only one service on a port, so you must choose a port number for the Telnet service that is different from the port chosen for the Telnet Console.

6. In the Attributes field, select Transparent, Explicit, Send-client-IP (spoofing), or all three. Explicit is the default.

Note: The `send-client-ip` attribute allows the ProxySG to pretend to be a client, allowing the origin content server to see the client's IP address. If an alternate path exists for traffic returning from the Internet to the client, the `Send-client-IP` attribute does not work.

7. Click OK; Click Apply.

To Enable or Create a Telnet Proxy Service through the CLI

Note: The `explicit` attribute is enabled by default and the `transparent` and `send-client-ip` attributes are disabled by default. Note also that only one service can use a port, so if you have Telnet-Console enabled on Port 23, you must choose a different port number for the Telnet shell proxy.

From the (config) prompt, enter the following commands:

```

SGOS#(config) services
SGOS#(config services) telnet
SGOS#(config services telnet) create [ip_address:]port
SGOS#(config services telnet) attribute {explicit | transparent |
send-client-ip} enable [ip_address:]port
SGOS#(config services telnet) enable [ip_address:]port

```

where:

<code>create</code>	<code>[<i>ip_address</i>:]<i>port</i></code>	Create a Telnet shell proxy service at the (optional) address and port number.
<code>attribute</code>	<code>explicit transparent send-client-ip enable [<i>ip_address</i>:]<i>port</i></code>	Give the Telnet shell proxy explicit and transparent attributes, and create IP spoofing (where the ProxySG pretends to be a client so the OCS can see the client's IP address). Note: The send-client-IP attribute allows the ProxySG to pretend to be a client, allowing the origin content server to see the client's IP address. If an alternate path exists for traffic returning from the Internet to the client, the Send-client-IP attribute does not work.
<code>enable</code>	<code>[<i>ip_address</i>:]<i>port</i></code>	Enable the new Telnet shell proxy.

To view the results:

```

SGOS#(config services telnet) view
Port:      23          IP: 0.0.0.0          Type: telnet
Properties: transparent, explicit, disabled
Port:      24          IP: 10.25.36.47    Type: telnet
Properties: explicit, enabled

```


Chapter 6: *Configuring Proxies*

A *proxy* filters traffic, monitors Internet and intranet resource usage, blocks specific Internet and intranet resources for individuals or groups, and enhances the quality of Internet or intranet user experiences.

A proxy also serves as an intermediary between a Web client and a Web server and can require authentication to allow identity-based policy and logging for the client. The rules used to authenticate a client are based on the policies created and implemented through your existing security framework, such as LDAP, RADIUS, and NTLM, and are further discussed in "Using Authentication Services" on page 271.

Explicit/Transparent proxy specifies the mode the client requests get to the proxy.

- **Explicit**—The default, requiring software configuration for both browser and service.
- **Transparent**—Requires 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 ProxySG IP address. You can also use the ProxySG software bridge to transparently proxy requests.

Some software configuration on the ProxySG is also required to allow the appliance to know what traffic to intercept.

You might also configure both proxy types, depending on the services you require.

This chapter contains the following topics:

- "About Explicit and Transparent Proxy"
- "Configuring Explicit Proxies"
- "Configuring the Transparent Proxy Hardware"

About 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.

Explicit Proxy

In an explicit proxy configuration, the client (browser) is explicitly configured to use a proxy server. The browser is given the IP address and port number of the proxy service (the ProxySG). It is also possible to configure the browser to download the proxy settings from a Web server. This is called a Proxy Auto-Configuration (PAC) file. When a user makes a request, the browser connects to the proxy service and sends the request. Since the browser knows it is talking to a proxy, the browser provides the proxy server with the destination server.

The proxy service accepts the explicit connection to it, and fetches the request from the browser. The request identifies the desired origin content server (OCS) and the resource on that server. The proxy service uses this information to contact the OCS if necessary.

The disadvantage to explicit proxy is that each desktop must be properly configured to use the proxy, which might not be feasible in a large organization.

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 make sure that the appropriate traffic is directed to the ProxySG, deploy hardware such as a Layer-4 switch or a WCCP router, or the ProxySG appliance's software bridge that can redirect selected traffic to the appliance. Traffic redirection is managed through polices 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 199.

Section A: Configuring Explicit Proxies

Section A: Configuring Explicit Proxies

You can configure several different explicit proxy servers and services:

- Native FTP—See "Configuring the FTP Proxy" on page 152.
- HTTP Proxy—See "HTTP Proxy" on page 159.
- SOCKS—See "Configuring a SOCKS Proxy" on page 188.
- Shell Proxies—See "Customizing Policy Settings for Shell Proxies" on page 193

For information on creating an explicit proxy server, regardless of type, continue with "Creating an Explicit Proxy Server".

Creating an Explicit Proxy Server

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. The ProxySG generates client instructions that describe how to configure Microsoft Internet Explorer, Netscape Communicator, and other browsers based on instructions selected by the ProxySG administrator. You can configure client instructions for each network adapter in the ProxySG with the Configuration>Network>Adapters>Interface>Settings button.

After selecting client instructions, the ProxySG administrator directs clients to go to the ProxySG home page and follow the instructions in the Browser Configuration section. The ProxySG detects the browser installed on the client and displays the appropriate instructions.

Using the ProxySG as an Explicit Proxy

To use the ProxySG as an explicit proxy and use services such as SOCKS or FTP, you must provide custom instructions to clients instructing them how to configure their browsers to use the ProxySG as a proxy server.

This is a two-step process, requiring that you add the proxy IP address to the browser and also instruct the ProxySG which adapter interface uses the proxy IP address.

Before the proxy can be used, you must:

- Configure the proxy server.
- Enable the explicit proxy (whether a service or a server).

The browsers described here are Internet Explorer 6.0 and Netscape 6.2. If you have different browsers or different versions of Internet Explorer or Netscape, refer to the vendor documentation for information on configuring proxies.

From Internet Explorer

1. Select Tools>Internet Options>Connections>LAN Settings.
2. Select Use a proxy server.

Section A: Configuring Explicit Proxies

3. Enter the IP address and port number for the proxy, or click Advanced to set proxy server IP addresses and port numbers for services such as HTTP, FTP, and SOCKS. (Configure HTTPS through the Secure field.)
4. Click OK to exit the Advanced Settings tab, then continue to click OK until you exit the Tools menu.

From Netscape 6.2

1. Select Edit>Preferences>Advanced>Proxies.
2. Select Manual proxy configuration.
3. Enter proxy server IP addresses and port numbers for services such as HTTP, FTP, SOCKS and SSL.
4. Click OK.

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 "Configuring Failover" on page 113.

Adapter Proxy Settings

Once the explicit proxy is configured on the browser, decide which adapter interfaces listen for which service. Each adapter interface can listen for only one IP address; you can configure multiple proxies on one ProxySG using the same IP address.

To Provide Configuration Instructions through the Management Console

1. Select Configuration>Network>Adapters.
2. Select an adapter and the correct interface and click Settings.
3. Select Using a proxy.
4. Click OK to close the Settings dialog.
5. Click Apply.

To Provide Configuration Instructions through the CLI

At the (config) command prompt, enter the following commands:

```
SGOS#(config) interface fast-ethernet interface_#  
SGOS#(config interface interface_#) instructions proxy
```

Configuring the FTP Proxy

In previous SGOS releases, connections to FTP origin content servers were only accomplished over HTTP. SGOS 4.x supports Native FTP proxy.

Note: As in previous releases, FTP requests sent through the HTTP proxy are still valid.

Section A: Configuring Explicit Proxies

Configuring an FTP proxy requires ProxySG configuration and specific configuration of the FTP client. The service must be enabled on the ProxySG before it can be used.

Data connections initiated by an FTP client to an FTP server are known as passive mode data connections. This type of connection is useful in situations where an FTP server is unable to make a 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.

This functionality allows administrators to select how the ProxySG responds to a request from an FTP client for a passive mode data connection (PASV command). This functionality does not affect HTTP requests for FTP objects (for example, those originating from browsers that are explicitly proxied to a ProxySG).

If the FTP server responds that it supports PASV, but the ProxySG is unable to connect because of a firewall blocking the port, the ProxySG only attempts a PORT command. 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.

Disabling passive mode data connections on the ProxySG servicing requests from this type of FTP client might provide a more acceptable response to the end user.

When passive mode data connections are disabled, the ProxySG returns a response to the FTP client indicating that the server does not support passive mode. The FTP client software controls any messages displayed to the end user as a result of this response from the ProxySG.

Limitations

- Internet Explorer does not support proxy authentication for Native FTP.
- The ProxySG FTP proxy does not support exceptions.

FTP Spoofing

Using policy, you can spoof the IP addresses for FTP data connections in both transparent and explicit deployments, for both active and passive modes; certain deployments are subject to limitations. The client and server-side policies are:

- `ftp.match_client_data_ip(yes)`—Matches the source IP address of the ACTIVE data connection with the destination IP address of the control connection (client side).
- `ftp.match_server_data_ip(yes)`—Matches the source IP address of the PASV data connection with the source IP address of the ProxySG control connection (server side).

Note: To always use the ProxySG physical IP address (no spoofing), define policy as `ftp.match_[client | server]_data_ip(no)`.

The following points describe the various data flow scenarios:

- Outbound client data connection (ProxySG to client)—When the client issues a PORT command, the ProxySG opens a data connection to the FTP client with the source IP address of whatever destination IP address the client used when opening the control connection.

Section A: Configuring Explicit Proxies

- Inbound client data connection (client to ProxySG)—When the client issues a PASV command, the ProxySG returns the IP address and port to which client makes a data connection.
 - Explicit—The ProxySG returns the destination IP address of the control connection; this can be a physical or virtual ProxySG IP address.
 - Transparent—The ProxySG returns the IP address of the physical adapter on which the control connection arrived.
- Outbound server data connection (ProxySG to FTP server)—When the ProxySG issues a PASV command upstream, the server returns an IP address and port to connect to. The ProxySG then opens a data connection to the server with the same source IP address it used to open the control connection. This address is defined by the `reflect_ip` property.
- Inbound server data connection (FTP server to ProxySG)—When the ProxySG issues a PORT command, the ProxySG provides the IP address and port number to which the server makes a data connection.
 - The ProxySG sends the control connection's source IP address if that IP is a local ProxySG (virtual or physical) IP address; or
 - The ProxySG sends the IP address of the physical adapter that was used to make the outgoing control connection.

FTP Server Limitations

Consider the following limitations when defining FTP spoofing policy:

- IIS and WS_FTP servers do not support PASV data connections with a source IP address that is different from the source IP address of the control connection.
- IIS and WS_FTP servers do not support ACTIVE data connections with a destination IP address that differs from the source IP address of the control connection.

Configuring the ProxySG for Native FTP Proxy

This section describes how to configure the ProxySG through the Management Console and the CLI.

To Configure Native FTP Proxy through the Management Console

1. Select Configuration>Services>FTP Proxy.

The FTP Proxy tab displays.

Section A: Configuring Explicit Proxies

FTP Proxy

FTP Options

Allow caching of FTP objects

Cache FTP objects for % of the time since last modified date

Cache FTP objects without last modified date for hours

Allow use of PASV mode to clients

Welcome Banner

Blue Coat FTP Service

Figure 6-1: FTP Proxy Tab

2. Select the checkbox to Allow caching of FTP objects. The default is enabled.
3. Determine the amount of time in percentage of how long since the object was last modified. The default is 10%.
4. Enter an amount, in hours, that the object remains in the cache before becoming eligible for deletion. The default is 24 hours.
5. Select the checkbox to allow use of PASV mode to clients. The default is enabled.

To Configure Native FTP Proxy through the CLI

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config)  caching
SGOS#(config caching)  max-cache-size 18
SGOS#(config caching)  ftp
SGOS#(config caching ftp)  enable
SGOS#(config caching ftp)  type-m-percent 20
SGOS#(config caching ftp)  type-n-initial 12
```

where:

<code>max-cache-size</code>	<i>megabytes</i>	The maximum size, in megabytes, of the largest object that can stored on the ProxySG. Note that <code>max-cache-size</code> sets the maximum object size for both HTTP <i>and</i> FTP.
<code>enable disable</code>		Enables or disables the caching of FTP objects.
<code>type-m-percent</code>	<i>percent</i>	Time to live for objects with a last-modified time.
<code>type-n-initial</code>	<i>hours</i>	Time to live for objects without a last-modified time.

2. (Optional) View the result.

Section A: Configuring Explicit Proxies

```
SGOS#(config caching ftp) view
Caching FTP objects is enabled
FTP objects with last modified date, cached for 20% of last modified time
FTP objects without last modified date, initially cached for 12 hours
```

3. (Optional) Change the default login syntax. The default syntax is Raptor. The ProxySG also supports the Checkpoint authentication syntax. The supported Checkpoint formats are:

- remoteuser@proxyuser@host (in USER command) for explicit FTP.
- remotepass@proxypass (in PASS command) for explicit FTP.
- remoteuser@proxyuser (in USER command) for transparent FTP.
- remotepass@proxypass (in PASS command) for transparent FTP.

Enter the following command to change the login syntax:

```
SGOS# (config) ftp login-syntax {raptor | checkpoint}
```

Note: Neither proxy authentication for transparent FTP nor proxy chaining are supported with the Checkpoint syntax.

Enabling the FTP Service

By default, an FTP service is already created with explicit and transparent attributes, but it is disabled. You must enable the FTP port before it can be used.

To Create and Enable a Native FTP Port Service through the Management Console

1. Select Configuration>Services>Service Ports.
The Service Ports tab displays.
2. Click New; the Add Service dialog appears.

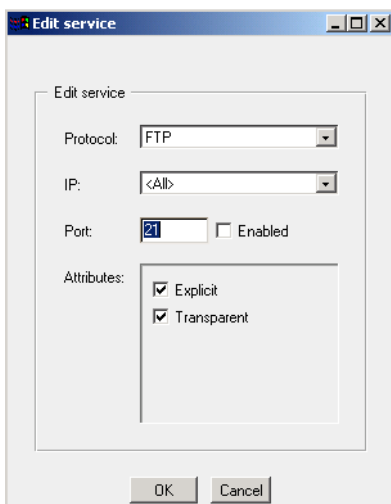


Figure 6-2: FTP Add Service Dialog

Section A: Configuring Explicit Proxies

3. In the Protocol drop-down list, select FTP.
4. The default IP address value is All. To limit the service to a specific IP, select the IP from the drop-down list.
5. In the Port field, specify a port number; select Enabled.
6. Choose the attributes you want the FTP proxy to have: Explicit, Transparent, or both.
7. Click OK; Click Apply.

To Create a Native FTP Port Service through the CLI

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) services
SGOS#(config services) ftp
SGOS#(config services ftp) create [ip_address:]port
SGOS#(config services ftp) attribute passive-mode {enable | disable}
SGOS#(config services ftp) attribute explicit enable [ip_address:]port
SGOS#(config services ftp) attribute transparent enable [ip_address:]port
```

2. (Optional) View the results.

```
SGOS#(config services ftp) view
Port:      25      IP: 0.0.0.0      Type: ftp
Properties: transparent, explicit, enabled, passive-allowed
```

Configuring FTP Clients

FTP clients must be configured as follows:

- Enable firewall.
- Select USER with no logon.
- For proxy authentication, select USER remoteID@remoteHost fireID and specify a proxy username and password.

Example

The following graphic demonstrates configuring a WSFtp client.

Section A: Configuring Explicit Proxies

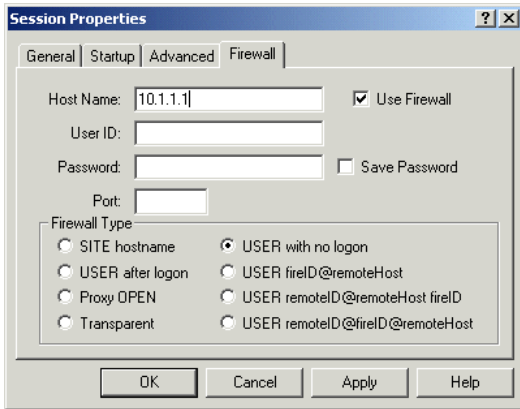


Figure 6-3: Configuring the WSFtp Client for Native FTP

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 Blue Coat 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 displayable when FTP is explicit through the ProxySG. In transparent deployments, the banner is sent to the client when proxy authentication is required; otherwise, the banner is forwarded from the FTP server.

To Define the Default FTP Banner through the Management Console

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.
3. Click Apply.

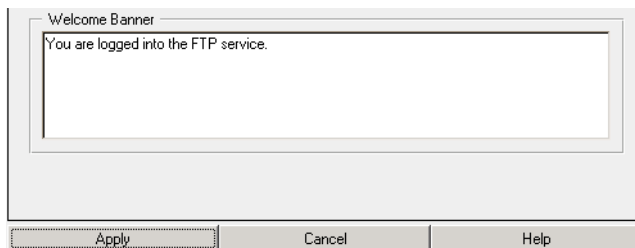


Figure 6-4: Configuring an FTP Connection Welcome Banner

Section A: Configuring Explicit Proxies

To Define the Default FTP Banner through the CLI

At the `(config)` prompt, enter the following command:

```
#SGOS#(config) ftp
#SGOS#(config) ftp welcome-banner "message"
```

To Create Policy that Overrides the Default Banner

Add the following property to a policy:

```
<Proxy>
ftp.welcome_banner "message"
```

If entering text that spans more than one line, use `$(crlf)` for line breaks.

HTTP Proxy

By default, an HTTP proxy service, with both explicit and transparent attributes set, is enabled on port 80. To change the attributes of the proxy service or create new HTTP proxy services, see "HTTP" on page 136.

The HTTP proxy is the first line of defense for the ProxySG, controlling all traffic that arrives on port 80 to the ProxySG. To control that traffic and improve performance, you can:

- Set default caching policies to configure the length of time an object or negative response is cached, whether objects are always revalidated before being served, whether server certificates are verified by default, and how headers are parsed. For more information, see "Setting Default HTTP Proxy Policy" on page 164.
- Configure the HTTP proxy as a server accelerator. For more information, see "Choosing the HTTP Proxy Profile" on page 168.
- Set a limit to the maximum bandwidth the ProxySG is allowed to use for refreshing objects in the background. For more information, see "Configuring Refresh Bandwidth for the HTTP Proxy" on page 162.
- Compress and decompress content. For more information, see "HTTP Compression" on page 178.

Note: Use of the compression feature is a trade-off among three resources: server-side bandwidth, client side-bandwidth, and CPU. If server-side bandwidth is expensive compared to CPU, then you should always request compressed content from the OCS. If CPU is comparatively expensive, then the ProxySG should ask the server for the compression formats that the client asked for and forward whatever the server returns.

The HTTP proxy is designed to control Web traffic, providing:

- Security
- Authentication
- Virus Scanning and Patience Pages
- Performance
 - Default HTTP Proxy Policy

Section A: Configuring Explicit Proxies

- ❑ HTTP Proxy Caching Profiles
- ❑ Byte-Range Support
- ❑ Refresh Bandwidth
- ❑ Compression

This chapter deals with HTTP proxy performance. See also:

- Chapter 8: “Security and Authentication” on page 241 to learn about controlling access to the ProxySG, Internet, intranet, and making decisions based on user identity.
- "Forms-Based Authentication" on page 361 for information about using Web forms for authentication.
- See "About Content Scanning" on page 401 for information about virus scanning and sending patience pages to explain the delays that can occur when scanning for viruses before serving data.

HTTP Proxy Performance

Default HTTP Proxy Policy

Using the ProxySG Management Console or CLI, you can configure global defaults that determine HTTP proxy caching policy, such as the maximum size of cacheable objects, the length of time that negative responses remain in the cache, whether the ProxySG revalidates each object before serving it, whether the server certificate is verified by default, and how headers are parsed.

For information about setting default policy for HTTP proxy caching, see "Setting Default HTTP Proxy Policy" on page 164.

HTTP Proxy Acceleration Profiles

You have a choice of three profiles to use for the ProxySG:

- Normal (the default setting) acts as a client-accelerator, and is used for enterprise deployments
- Portal acts as a server accelerator, and is used for web-hosting
- Bandwidth Gain is used for ISP deployments

For information on HTTP profiles, see "Choosing the HTTP Proxy Profile" on page 168.

Byte-Range Support

If a client makes a request using the `Range: HTTP` header, the ProxySG can serve the requested portions of the file from the cache if byte-range support is enabled (the default). If byte range support is disabled, all such requests will be forwarded to the origin content server and the response will not be cached. For information on using byte-range support to determine how the ProxySG handles byte-range requests, see "Additional Configuration Affecting Bandwidth Gain" on page 175.

Section A: Configuring Explicit Proxies

Refresh Bandwidth

Refresh bandwidth refers to server-side bandwidth used for all forms of asynchronous refresh activity. The default configuration is to allow the ProxySG to manage refresh bandwidth. The ProxySG manages the bandwidth in order to preserve the maximum freshness of accessed objects. It can sometimes be the case, however, that the automatic refresh bandwidth amount is too high. It is not unusual for refresh bandwidth to spike up occasionally, depending on access patterns at the time. If necessary, you can choose to impose a limit on refresh bandwidth. To limit the refresh bandwidth to a specified amount, you must disable automatic management of the bandwidth and explicitly set a bandwidth limit. Note that setting the refresh bandwidth amount too low can lower the estimated freshness of objects in the cache. If you set the refresh bandwidth amount to zero, the ProxySG does not do active refresh at all.

For information about configuring refresh bandwidth, see "Configuring Refresh Bandwidth for the HTTP Proxy" on page 162.

Compression

Compression is disabled by default (even if you have a valid license for this feature). 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 the ProxySG to send the response as is when server sends compressed data, including non-cacheable responses. Any unsolicited encoded response is forwarded to the client as is.

For more information on compression, see "HTTP Compression" on page 178.

HTTP Terms

- **Asynchronous Adaptive Refresh (AAR)**—This allows the ProxySG to keep cached objects as fresh as possible, thus reducing response times. The AAR algorithm allows HTTP proxy to manage cached objects based on their rate of change and popularity: an object that changes frequently and/or is requested frequently will be more eligible for asynchronous refresh compared to an object with a lower rate of change and/or popularity.
- **Asynchronous Refresh Activity**—Refresh activity that does not wait for a request to occur, but that occurs *asynchronously* from the request.
- **Bandwidth Gain**—A measure of the difference in client-side and server-side internet traffic expressed in relation to server-side internet traffic. It is managed in two ways: you can enable or disable bandwidth gain mode or you can select the Bandwidth Gain profile (this also enables bandwidth gain mode). See "Configuring the HTTP Proxy Profile" on page 173 for information about configuring bandwidth gain.
- **Byte-Range Support**—The ability of the ProxySG to respond to byte-range requests (requests with a `Range`: HTTP header).
- **Cache-hit**—An object that is in the ProxySG and can be retrieved when an end user requests the information.

Section A: Configuring Explicit Proxies

- **Cache-miss**—An object that can be stored but has never been requested before; it was not in the ProxySG to start, so it must be brought in and stored there as a side effect of processing the end-user's request. If the object is cacheable, it is stored and served the next time it is requested.
- **Compression**—An algorithm that reduces a file's size but does not lose any data. The ability to compress or decompress objects in the cache is based on policies you create. Compression can have a huge performance benefit, and it can be customized based on the needs of your environment: Whether CPU is more expensive (the default assumption), server-side bandwidth is more expensive, or whether client-side bandwidth is more expensive.
- **Freshness**—A percentage that reflects the objects in the ProxySG cache that are expected to be fresh; that is, the content of those objects is expected to be identical to that on the OCS (origin content server).
- **Maximum Object Size**—The maximum object size stored in the ProxySG. All objects retrieved that are greater than the maximum size are delivered to the client but are not stored in the ProxySG.
- **Negative Responses**—An error response received from the OCS when a page or image is requested. If the ProxySG is configured to cache such negative responses, it returns that response in subsequent requests for that page or image for the specified number of minutes. If it is not configured, which is the default, the ProxySG attempts to retrieve the page or image every time it is requested.
- **Refresh Bandwidth**—The amount of bandwidth used to keep stored objects fresh. By default, the ProxySG is set to manage refresh bandwidth automatically. You can configure refresh bandwidth yourself, although Blue Coat does not recommend this.
- **Variants**—Objects that are stored in the cache in various forms: the original form, fetched from the OCS; the transformed (compressed or uncompressed) form (if compression is used). If a required compression variant is not available, then one might be created upon a cache-hit. (Note: policy-based content transformations are not stored in the ProxySG.)

Configuring Refresh Bandwidth for the HTTP Proxy

The ProxySG uses as much bandwidth as necessary for refreshing to achieve the desired access freshness.

The amount of bandwidth used varies depending on client demands. If you determine that the ProxySG is using too much bandwidth (by reviewing the logged statistics and examining current bandwidth used shown in the Refresh bandwidth field), you can specify a limit to the amount of bandwidth the ProxySG uses to try to achieve the desired freshness. Be aware, however, that if you limit the amount of bandwidth the ProxySG can use, you might prohibit the ProxySG from achieving the desired freshness. If the refresh bandwidth configuration remains at the recommended default—Let the ProxySG Appliance manage refresh bandwidth (recommended) in the Management Console or `SGOS#(config caching) refresh automatic` in the CLI—then the ProxySG uses whatever bandwidth is available in its efforts to maintain 99.9% estimated freshness of the next access.

To Set Refresh Bandwidth through the Management Console

1. Select Configuration>Services>HTTP Proxy>Freshness.

The Freshness tab displays.

Section A: Configuring Explicit Proxies

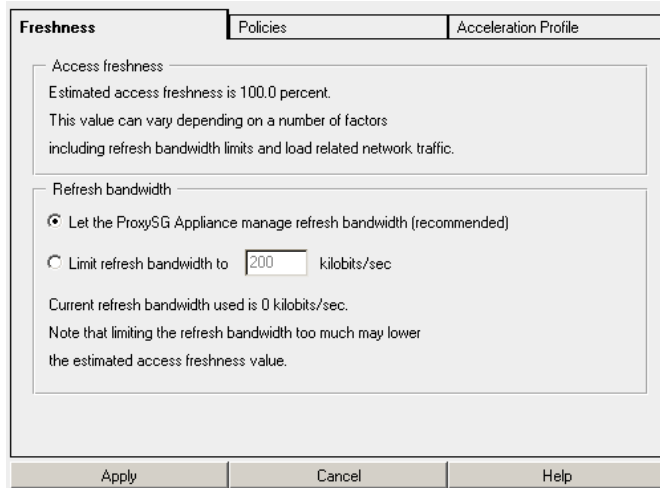


Figure 6-5: Freshness Tab

The Refresh bandwidth field displays the refresh bandwidth options. The default setting is to allow the ProxySG to manage refresh bandwidth automatically.

Important: Blue Coat strongly recommends that you not change the setting from the default.

2. Do one of the following:
 - To turn off automatic bandwidth refresh, select Limit refresh bandwidth to (not recommended). Enter a new value into the kilobits/sec field, if necessary.
 - To return the ProxySG to automatic bandwidth refresh, select Let the ProxySG Appliance manage refresh bandwidth (recommended).
3. Click Apply.

To Set Refresh Bandwidth through the CLI

1. To disable automatic bandwidth refresh (not recommended), enter the following commands at the (config) command prompt:


```
SGOS#(config) caching
SGOS#(config caching) refresh no automatic
```
2. (Optional) To adjust the kilobit/sec refresh bandwidth value, enter the following commands:

Note: Adjusting the refresh bandwidth value has no effect if you do not also turn off the automatic refresh bandwidth option (you must perform Step 1). You can skip Step 2 if the refresh bandwidth value is acceptable (200 Kbps is the default).

```
SGOS#(config) caching
SGOS#(config caching) refresh bandwidth kbps
```

3. To return the ProxySG to automatic bandwidth refresh (recommended), enter the following commands:

Section A: Configuring Explicit Proxies

```
SGOS#(config)  caching  
SGOS#(config caching)  refresh automatic
```

4. (Optional) View the (truncated) results:

```
SGOS#(config caching)  view  
Refresh:  
  Estimated access freshness is 100.0%  
  Let the ProxySG Appliance manage refresh bandwidth  
  Current bandwidth used is 0 kilobits/sec
```

To view all HTTP settings, see "Viewing HTTP Settings through the CLI" on page 177.

Setting Default HTTP Proxy Policy

Using the ProxySG Management Console or CLI, you can configure global defaults that determine HTTP proxy policy, such as the maximum size of cacheable objects, the length of time that negative responses remain in the cache, whether the ProxySG revalidates each object before serving it, whether the server certificate is verified by default, and how headers are parsed.

Other policy can be controlled only by using Blue Coat Content Policy Language (CPL). This section is about using the Management Console or CLI to set default HTTP proxy policy; see "Creating a Proxy Layer to Manage Proxy Operations" on page 261 for information about using CPL to configure HTTP proxy caching.

Note: Tolerant HTTP request parsing can only be done through the CLI; it is not available through the Management Console.

To Set HTTP Default Proxy Policy through the Management Console:

1. Select Configuration>Services>HTTP Proxy>Policies.

The Policies tab displays.

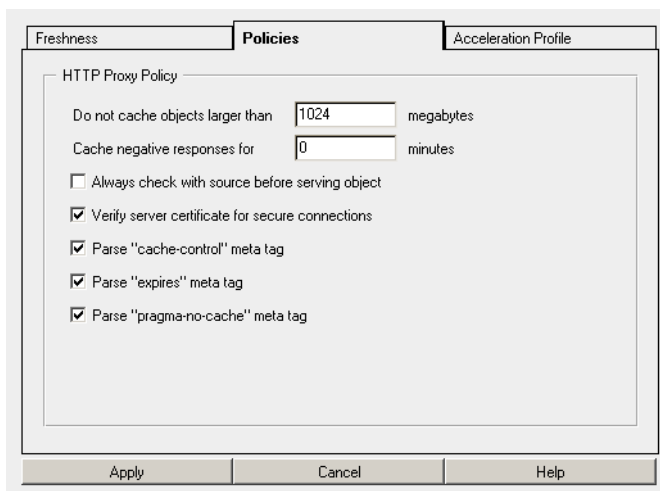


Figure 6-6: Policies Tab

Section A: Configuring Explicit Proxies

2. Fill in the fields as appropriate:

- ❑ In the Do not cache objects larger than field, enter the maximum object size to cache. The default is 1024 MB. This configuration determines the maximum object size to store in the ProxySG. All objects retrieved that are greater than the maximum size are delivered to the client but are not stored in the ProxySG.
- ❑ In the Cache negative responses for field, enter the number of minutes the ProxySG stores negative responses. The default is 0, meaning that the ProxySG will not cache negative responses and will always attempt to retrieve the object.

The OCS might send a client error code (4xx HTTP response) or a server error code (5xx HTTP response) as a response to some requests. If the ProxySG is configured to cache such negative responses, it returns that response in subsequent requests for that page or image for the specified number of minutes. If it is not configured, which is the default, the ProxySG attempts to retrieve the page or image every time it is requested.

If you enter a number of minutes into this field, then your response times will improve, but you could receive negative responses to requests that might otherwise have been served for that period of time.

- ❑ To always verify that each object is fresh upon access, select the Always check with source before serving object checkbox. Enabling this setting has a significant impact on performance, because HTTP proxy will revalidate requested cached objects with the OCS before serving them to the client, resulting in a negative impact on response times and bandwidth gain. Therefore, this configuration item should not be enabled unless absolutely required.
- ❑ If you communicate with an OCS through HTTPS and want the OCS's certificate to be verified by the ProxySG, make sure that the Verify server certificate for secure connections checkbox is checked.
- ❑ The default is to parse HTTP meta tag headers in HTML documents if the MIME type of the object is text/HTML. The function of all meta tags is same as the corresponding HTTP headers.

To disable meta-tag parsing, remove the check from the checkbox for:

- Parse "cache-control" meta tag
- Parse "expires" meta tag
- Parse "pragma-no-cache" meta tag

The following sub-headers are parsed when this checkbox is selected: private, no-store, no-cache, max-age, s-maxage, must-revalidate, proxy-revalidate.

3. Click Apply.

To Set HTTP Proxy Default Policy through the CLI

1. At the (config) command prompt, enter the following commands:

Section A: Configuring Explicit Proxies

```

SGOS#(config)  caching 
SGOS#(config caching) max-cache-size megabytes
SGOS#(config caching) negative-response minutes
SGOS#(config caching) always-verify-source
-or-
SGOS#(config caching) no always-verify-source

```

where:

<code>max-cache-size</code>	<i>megabytes</i>	The maximum size, in megabytes, of the largest object that can stored on the ProxySG. Note that <code>max-cache-size</code> sets the maximum object size for both HTTP and FTP.
<code>negative-response</code>	<i>minutes</i>	The amount of time, in minutes, that the ProxySG remembers that an object is not stored.
<code>always-verify-source</code>		Ensures that every object is always fresh upon access. This has a significant impact on performance, because HTTP proxy will revalidate requested cached objects with the OCS before serving them to the client, resulting in a negative impact on response times and bandwidth gain. Therefore, this configuration item should not be enabled unless absolutely required.
<code>no</code>	<code>always-verify source</code>	The default setting. This tells the ProxySG never to check objects on the source before serving them to the client.

Note: If you use HTTPS, you might want to change the verify-server certificate from the default of enabled to suppress verification of the OCS certificate (step 2).

2. (Optional) To enable or disable the verify-server certificate setting, enter one of the following commands:

```

SGOS#(config caching) exit
SGOS#(config) http ssl-verify-server
-or-
SGOS#(config) http no ssl-verify-server

```

3. (Optional) To enable or disable meta-tag parsing (parsing is enabled by default), enter one of the following commands:

```

SGOS#(config services) exit
SGOS#(config) http parse meta-tag {cache-control | expires | pragma-no-cache}
-or-
SGOS#(config) http no parse meta-tag {cache-control | expires | pragma-no-cache}

```

To view all HTTP settings, see "Viewing HTTP Settings through the CLI" on page 177.

Section A: Configuring Explicit Proxies

Tips on Parsing Meta Tags

- If ICAP response modification is occurring, the response body modified by the ICAP server is not parsed.
- Relevant HTTP meta tags must appear within the first 1000 bytes of HTTP object body. If the meta tag does not appear within the first 1000 bytes, it is ignored.

Tips on Using Meta Tags With Policy

- The following CPL properties can be used in the <Cache> layer to control meta tag processing for HTTP proxy, HTTP refresh, and HTTP pipeline transactions:

```
http.response.parse_meta_tag.Pragma.no-cache (yes|no)
http.response.parse_meta_tag.Cache-Control (yes|no)
http.response.parse_meta_tag.Expires (yes|no)
```

- VPM support for this feature is not available.

Tolerant HTTP Request Parsing

By default, the ProxySG blocks malformed HTTP requests, returning a *400 Invalid Request* error. The tolerant HTTP request parsing flag causes certain types of malformed requests to be processed instead of being rejected.

By default, a header line not beginning with a <Tab> or space character must consist of a header name (which contains no <Tab> or space characters), followed by a colon, followed by an optional value, or an error is reported. 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 one or more <Tab> or space characters, and nothing else, is considered ambiguous. Blue Coat doesn't know if this is a blank continuation line or if it is the 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 signals the end of the header section.

To Enable the HTTP Tolerant Request Parsing Flag through the CLI

Note: This feature is only available through the CLI. It cannot be set through the Management Console.

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, enter the following command:

```
SGOS#(config) http no tolerant-request-parsing
```

To view all HTTP settings, including `http tolerant-request-parsing` if it is enabled, see "Viewing HTTP Settings through the CLI" on page 177.

Section A: Configuring Explicit Proxies

Choosing the HTTP Proxy Profile

You can select from among three profiles for the HTTP proxy, depending on your needs, and you can also create a customized profile from the three available.

The three profiles are:

- Normal, which acts as a client-accelerator and is used for enterprise deployments
- Portal, which acts as a server accelerator and is used for web-hosting
- Bandwidth, which is used for ISP deployments

Table 6.1 shows the configuration for each profile. For a description of each configuration setting, see Table 6.2 on page 170.

Table 6.1: 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	Does not change
Substitute GET for HTTP 1.1 conditionals	Disabled	Enabled	Enabled
Substitute GET for IE (Internet Explorer) reload	Disabled	Enabled	Does not change
Never refresh before expiration	Disabled	Enabled	Enabled
Never serve after expiration	Disabled	Enabled	Does not change

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 needed elements of all three profiles.

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, prefetching (pipelining) of embedded objects and redirects is enabled, which reduces response time for clients.

Section A: Configuring Explicit Proxies

Portal Profile

When configured as a server accelerator, the ProxySG improves object response time to client requests, scalability of the 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.

Because an OCS can actually consist of many servers—a single Web server or an entire server farm—OCSs are identified by domain name or IP address. To the ProxySG, the domain name or IP address is treated as the OCS, regardless of how many back-end Web servers might be installed.

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, the freshness of the object is not as important as 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 that reduces the amount of server-side bandwidth required.

HTTP Object Types

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's asynchronous adaptive refresh (AAR) algorithm is normally applied to all three types of HTTP objects. To maximize the freshness of the next access to objects in the ProxySG's cache, asynchronous revalidations are performed on those objects in accordance with their relative popularity and the amount of time remaining before their estimated time of expiration. Estimated expiration times will vary as object content changes are observed during such asynchronous revalidations. This will happen even for type-T objects, because the expiration times of type-T objects are not always perfectly managed by webmasters of content servers. However, for situations where such management can be trusted, certain configuration items exist to reduce speculative revalidation of type-T objects. In the following section, the terms revalidation and refresh mean the same thing—to assess the freshness of an object by sending a conditional GET request to the object's OCS.

Section A: Configuring Explicit Proxies

HTTP Proxy Profile Configuration Components

Table 6.2 gives a definition of the customizable HTTP proxy profile settings. Both the Management Console field and CLI (`config`) command text is included.

Table 6.2: Description of Profile Configuration Components in the Management Console and CLI

Management Console Checkbox Field	CLI (<code>config</code>) 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 will acquire the object's content before the client requests the object. This improves response time dramatically. If this setting is disabled, HTTP proxy will 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.
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 pipelines those embedded objects as well. This nested pipelining behavior can occur three levels deep at most. If this setting is disabled, HTTP proxy does not engage in nested pipelining behavior.
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.

Section A: Configuring Explicit Proxies

Table 6.2: Description of Profile Configuration Components in the Management Console and CLI (Continued)

Management Console Checkbox Field	CLI (<code>config</code>) Command	Definition
Substitute Get for IMS	<code>http [no] substitute if-modified-since</code>	<p>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 is a strong indication that the copy in the cache is stale. When that is the case, HTTP proxy will do a conditional GET to the OCS, based on the last modified time of the cached object.</p> <p>To control this aspect of the ProxySG's treatment of the <code>If-Modified-Since:</code> header, disable the Substitute Get for IMS setting. When this setting is disabled, a client time condition greater than the last modified time of the object in the cache does not trigger revalidation of the object.</p> <p>Note, however, that not all objects necessarily have a last-modified time specified by the OCS.</p>
Substitute Get for HTTP 1.1 conditionals	<code>http [no] substitute conditional</code>	<p>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.</p> <p>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:</p> <ul style="list-style-type: none"> • <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>	<p>Typically, if a client sends an HTTP GET request with a <code>Pragma: no-cache</code> or <code>Cache-Control: no-cache</code> header (both are hereby referred to as PNC for convenience), a cache must consult the OCS before serving the content. This means that HTTP proxy will always re-fetch 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).</p>

Section A: Configuring Explicit Proxies

Table 6.2: Description of Profile Configuration Components in the Management Console and CLI (Continued)

Management Console Checkbox Field	CLI (config) Command	Definition
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 will ignore the PNC interpretation of the <code>Accept: */*</code> header.
Never refresh before expiration	<code>http [no] strict-expiration refresh</code>	Applies only to cached type-T objects. When this setting is enabled, the ProxySG will 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 will be 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 will be 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 will be 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.

Section A: Configuring Explicit Proxies

Table 6.2: Description of Profile Configuration Components in the Management Console and CLI (Continued)

Management Console Checkbox Field	CLI (config) Command	Definition
Enable Bandwidth Gain Mode	<code>bandwidth-gain {disable enable}</code>	<p>This setting controls both HTTP-object acquisition after client-side abandonment and AAR (asynchronous adaptive refresh) revalidation frequency.</p> <ul style="list-style-type: none"> • 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. Note that not all asynchronously refreshable objects are guaranteed to be revalidated.

Configuring the HTTP Proxy Profile

You can configure the profile you want from either the Management Console or the CLI.

To Configure the HTTP Proxy Profile through the Management Console

1. Select Configuration>Services>HTTP Proxy>Acceleration Profile.

The Acceleration Profile tab displays (Normal is the default profile). Text appears at the bottom of this tab indicating which profile is selected. If you have a customized profile, this text does not appear.

Section A: Configuring Explicit Proxies

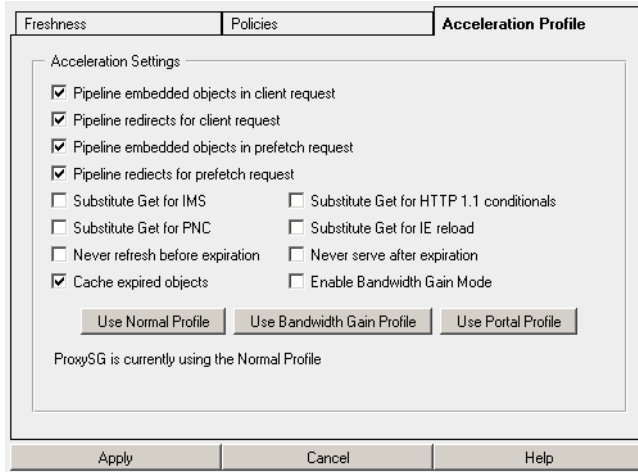


Figure 6-7: Acceleration Profile Tab

Important: If you have a customized profile and you click one of the Use Profile buttons, no record of your customized settings remains. However, once the ProxySG is set to a specific profile, the profile is maintained in the event the ProxySG is upgraded.

2. 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.

3. (Optional) To customize the profile settings, select or deselect any of the checkboxes (see Table 6.2 on page 170 for information about each setting).
4. Click Apply.

To Configure the HTTP Proxy Profile through the CLI

1. At the `(config)` command prompt, enter the profile you want:
`SGOS#(config) profile {normal | portal | bwgain}`
2. (Optional) Use the following commands to customize the profile settings (see Table 6.2 on page 170 for information about these settings):

```
SGOS#(config) http [no] pipeline client requests
SGOS#(config) http [no] pipeline client redirects
SGOS#(config) http [no] pipeline prefetch requests
SGOS#(config) http [no] pipeline prefetch redirects
SGOS#(config) http [no] substitute if-modified-since
SGOS#(config) http [no] substitute conditional
SGOS#(config) http [no] substitute pragma-no-cache
SGOS#(config) http [no] substitute ie-reload
```

Section A: Configuring Explicit Proxies

```
SGOS#(config) http [no] strict-expiration refresh
SGOS#(config) http [no] strict-expiration serve
SGOS#(config) http [no] cache expired
SGOS#(config) bandwidth-gain {disable | enable}
```

3. (Optional) View the settings. (This example assumes you have selected the Portal profile.)

```
SGOS#(config) show profile
SG is currently using the Normal Profile
Pipeline client requests:           Enabled
Pipeline client redirects:          Enabled
Pipeline prefetch requests:         Enabled
Pipeline prefetch redirects:        Enabled
Substitute Get "if-modified-since": Disabled
Substitute Get "pragma: no-cache":  Disabled
Substitute HTTP 1.1 Conditional Get: Disabled
Substitute Internet Explorer reload: Disabled
Never refresh before expiration:     Disabled
Never serve after expiration:        Disabled
Cache expired objects:              Enabled
Bandwidth gain mode:                Disabled
```

You can view all HTTP settings. See "Viewing HTTP Settings through the CLI" on page 177 for more information.

Additional Configuration Affecting Bandwidth Gain

In addition to the configuration items related to top-level profiles, other configurable items also affect bandwidth gain. You can set the top-level profile and adjust various related configuration items to fine tune ProxySG for your needs (see "Configuring the HTTP Proxy Profile" on page 173), and you can provide additional fine-tuning with the following configuration items:

- Byte-range support
- Revalidate pragma-no-cache

Byte-range requests can be made with a pragma-no-cache header. In order to serve these requests from the cache, you will need to enable the revalidate pragma-no-cache setting (see "Revalidate Pragma-No-Cache" below).

Byte-Range Support

If a client requests a byte range using the `Range:` HTTP header, the ProxySG can serve the requested portions of the file from the cache if byte-range support is enabled (the default). If byte range support is disabled, all such requests will be forwarded in a non-cacheable way to the origin content server.

Byte-range configuration can significantly affect bandwidth gain where heavy use of range requests is expected. Download managers (such as NetAnts®) typically use byte-range requests heavily.

Section A: Configuring Explicit Proxies

With byte-range support enabled, if the object is already cached and does not need to be reloaded from the OCS, the ProxySG serves the byte-range request from the cache only. But if the object is not in the cache, or if a reload of the object is required, the ProxySG might treat the byte-range request as if byte-range support is disabled and serve the object from the cache. It is important to note that HTTP proxy never caches partial objects, even if byte-range support is enabled.

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.

HTTP proxy categorizes the range requests in following three categories when byte-range support is enabled:

- Type-1: 0-N: Range request for first N bytes of the object
- Type-2: N-M: Range request from N bytes to M bytes of the object
- Type-3: -N: Range request for last N bytes of the object

If the object does not exist in the cache, and a byte-range request is received with the first range being type-1 or type-2, and the start byte of the first requested range is within 14336 bytes (hard coded threshold), then the entire object is retrieved from the OCS and cached in the ProxySG. Even though HTTP proxy retrieves the entire object from the OCS, it sends an appropriate byte-range response to the client. If the object does not exist in the cache, and if the first range in the request is not of type-1 or type-2, or if the start byte of the first requested range is beyond 14336 bytes, then the client's request is sent unaltered to the OCS and the response is not cached.

If the object exists in the cache, and if a range request with an effective PNC (the PNC header is not substituted or revalidated—see "Revalidate Pragma-No-Cache" below) is made, and the first range in the request is either type-3 or type-1 or 2 with a start byte offset greater than 14336 bytes, then, even if the object exists in the cache, the transaction is made non-cacheable (the request is sent to the OCS without any modification and the response is not cached). If an object exists in the cache and a byte-range request is made without the PNC header, then the byte-range response is satisfied from the cache.

Most download managers make byte-range requests with a PNC (pragma-no-cache) header. To serve such requests from the cache, the revalidate pragma-no-cache option should be configured along with byte-range support (see "Revalidate Pragma-No-Cache" below).

To Configure Byte-Range Support through the CLI

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
```

To view all HTTP settings, see "Viewing HTTP Settings through the CLI" on page 177.

 Section A: Configuring Explicit Proxies

Revalidate Pragma-No-Cache

The pragma-no-cache (PNC) header in a client's request can affect the efficiency of the proxy from a bandwidth gain perspective (this behavior is described in Table 6.2 in the Substitute Get for PNC configuration description). If you do not want to completely ignore PNC in client requests (which you can do by using the Substitute Get for PNC configuration), you can lower the impact of the PNC by enabling the `revalidate-pragma-no-cache` setting. When the `revalidate-pragma-no-cache` setting is enabled, a client's non-conditional PNC-GET request will result in a conditional GET request sent to the OCS if the object is already in the cache. This gives the OCS a chance to return the 304 Not Modified response, thus consuming less server-side bandwidth, because it has not been forced to return full content even though the contents have not actually changed. By default, the `revalidate-pragma-no-cache` configuration is disabled and is not affected by changes in the top-level profile. Note that when the Substitute Get for PNC configuration is enabled (see "Configuring the HTTP Proxy Profile" for configuration information), the `revalidate-pragma-no-cache` configuration has no effect.

To Configure the Revalidate PNC Setting through the CLI

Note: The `revalidate-pragma-no-cache` setting can only be configured through the CLI.

To enable or disable the `revalidate-pragma-no-cache` 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
```

To view all HTTP settings, see "Viewing HTTP Settings through the CLI" below.

Viewing HTTP Settings through the CLI

You can view the existing HTTP settings by entering the following command:

```
SGOS#(config) show http
Supported protocol version: HTTP 1.1
Caching options:
  Cache authenticated data: enabled
  Cache expired objects:    enabled
  Cache personal pages:    disabled
  Reverse DNS lookup on IP: disabled
  Strip From Headers:      disabled
  Byte range support:      enabled
  Force NTLM on proxy IE:  disabled
  Rewrite redirects for XP: disabled
  Revalidate "pragma: no-cache": disabled
  WWW redirect if host not found: enabled
Force explicit expirations:
  Never refresh before:    disabled
  Never serve after:      disabled
Add headers:
  "Front-end-https":      disabled
  "Via":                  disabled
  "X-forwarded-for":     disabled
```

Section A: Configuring Explicit Proxies

```
"Client-ip":                disabled
Parsing options:
  HTML meta tag "Cache-Control":  enabled
  HTML meta tag "Expires":        enabled
  HTML meta tag "Pragma: no-cache": enabled
Persistent connections:
  Client connections:            enabled
  Server connections:            enabled
Pipeline:
  Client requests:               enabled
  Client redirects:              enabled
  Prefetch requests:             enabled
  Prefetch redirects:            enabled
Substitute simple Get for:
  Get "if-modified-since":       disabled
  Get "pragma: no-cache":         disabled
HTTP 1.1 Conditional get: disabled
Internet Explorer reload: disabled
Proprietary header extensions:
  Blue Coat extensions:          disabled
FTP proxy:
  Url path is:                   absolute from root
  Configuration/access log uploads: will use PASV
Persistent connection timeouts:
  Server:                         900
  Client:                          360
Receive timeouts:
  Server:                          180
  Client:                           120
  Refresh:                          90
Https:
  ssl-verify-server:              enabled
  tolerant-request-parsing:       enabled
```

HTTP Compression

Compression is an algorithm that 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, then you should always request compressed content from the origin content server (OCS). However, if CPU is comparatively expensive, the ProxySG 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 behavior.

Section A: Configuring Explicit Proxies

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 from the origin content server, bandwidth gain statistics reflect such requests/responses as negative bandwidth gain.

Compression is disabled by default (even if you have a valid license for this feature). 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 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: Compression is a licensed feature of SGOS. If the ProxySG has no valid license for compression, it does not compress the content if the server sends uncompressed content. However, the ProxySG continues to uncompress content if necessary to apply transformations.

Any unsolicited encoded response is forwarded to the client as is.

For information on licensing, see "Licensing" on page 35.

Compression is controlled by policy only.

You can view compression statistics by going to Statistics>System Usage>Client Comp. Gain and Server Comp. Gain and Statistics>HTTP/FTP History>Client Comp. Gain and Server Comp. Gain. For information on these statistics, see "System Usage Statistics" on page 816 and "HTTP/FTP History Statistics" on page 821.

Compression Behavior

The ProxySG compression behavior is detailed in the tables below.

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.

Compression increases the overall percentage of cacheable content, increasing the hit rate in terms of number of objects served from the cache.

Section A: Configuring Explicit Proxies

For cache-hit compression behavior, see Table 6.3 below. For cache-miss compression behavior, see Table 6.4 on page 180.

Table 6.3: 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 does not convert from gzip to deflate.)

Table 6.4: 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 issues a `transformation_error` exception (HTTP response code 403), when the server sends an unknown encoding and the ProxySG is configured to do content transformation.

Section A: Configuring Explicit Proxies

- The ProxySG issues an `unsupported_encoding` exception (HTTP response code 415 - Unsupported Media Type) when the ProxySG is unable to deliver content due to configured policy.

The messages in the exception pages can be customized. For information on using exception pages, see "Section D: Defining Exceptions".

Configuring Compression

Compression behavior can only be configured through policy—VPM or CPL.

Using VPM to Configure Compression Behavior

Two objects can be used to configure compression through VPM:

- HTTP client compression object: Allows you to determine the behavior when the client wants the content in a different form than is in the cache.
- HTTP server compression object: Allows you to enable or disable compression and to set options.

Complete the following steps to manage HTTP server and client compression.

To Add or Edit Client Compression

1. Create a Web Access Layer:
 - Select Configuration>Policy>Visual Policy Manager; click the Launch button.
 - 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 Client Compression. The Add Client HTTP Compression Object dialog displays.

Section A: Configuring Explicit Proxies

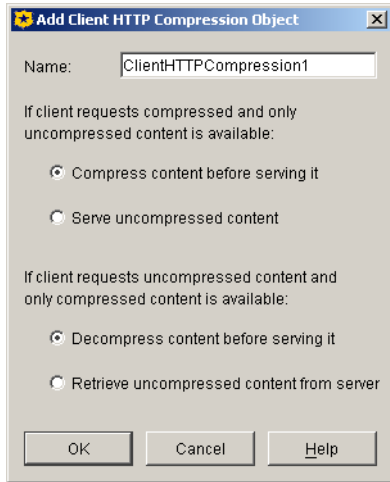


Figure 6-8: Add Client HTTP Compression Object Dialog

- ❑ Select the compression options you want to use; click OK.
- ❑ 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:
 - ❑ Select Configuration>Policy>Visual Policy Manager; click the Launch button.
 - ❑ 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 Server HTTP Compression.

The Add Server HTTP Compression Object dialog displays.




Figure 6-9: Add Server HTTP Compression Object Dialog

Section A: Configuring Explicit Proxies

- ❑ Select the compression options you want to use; 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 you have a valid compression license and compression is enabled. If you do not have a valid compression license, or if compression is not enabled, neither compression nor decompression are allowed.

Policy controls the compression or decompression of content on the ProxySG. If compression is turned off, uncompressed content is served to the client if a compressed variant is not available. If decompression is turned off, an uncompressed version is fetched from the OCS if the variant doesn't exist and the client requested uncompressed content.

Note: The ProxySG decompresses the content if transformation is to be applied, even if the compression license is expired or not present.

You can use server-side or client-side controls to manage compression through policy, as described in the following table.

Table 6.5: Compression Properties

Compression Properties	Meaning
<code>http.allow_compression(yes no)</code>	Allow the ProxySG to decompress content on the fly if needed.
<code>http.allow_decompression(yes no)</code>	Allow the ProxySG to decompress content on the fly if needed.
<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).

Section A: Configuring Explicit Proxies

Table 6.5: Compression Properties (Continued)

Compression Properties	Meaning
<code>http.client.allow_encoding(client);</code>	Allow all client encodings. This is the default regardless of the state of the compression license.
<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 the compression license is expired or not present, 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 6.5. The `Accept-Encoding header modification` can continue to be used if no compression policies are applied, or if the compression license is not present or expired. 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:

Section A: Configuring Explicit Proxies

```

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

```

Boundary Conditions

- Policy-based content transformations are not stored as variant objects. If content transformation is configured, it will be 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. If you want to change this, you can write policy to disallow compression or decompression if `Cache-Control: no-transform` response header is present.
- The ProxySG 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 receives unknown content encoding and if content transformation is configured (such as popup blocking), an error results.
- 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 doesn't 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.

Troubleshooting HTTP Proxy Issues

This section covers problems you might encounter using the HTTP proxy.

Section A: Configuring Explicit Proxies

Using Explicit HTTP Proxy with Internet Explorer

Internet Explorer does not allow OCS NTLM authentication through a ProxySG when explicitly proxied. To correct this, Blue Coat has added a `Proxy-Support: Session-based-authentication` header that is sent by default when the ProxySG receives a 401 authentication challenge from upstream when the client connection is an explicit proxy connection.

For older browsers or if both the ProxySG and the OCS do NTLM authentication, the Proxy-Support header might not work. In this case, you can disable the header and instead enable NTLM-force, which converts the 401-type server authentication challenge to a 407-type proxy authentication challenge, supported by Internet Explorer. The ProxySG also converts the resulting Proxy-Authentication headers in client requests to standard server authorization headers, which allows an OCS NTLM authentication challenge to pass through when Internet Explorer is explicitly proxied through the ProxySG.

Disabling the Proxy-Support Header

You can control the header using header modification policy. Suppression or modification of the Proxy-Support custom header keeps 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, see Chapter 14: “The Visual Policy Manager” on page 453.

Note: If you want to suppress the Proxy-Support header globally, you can use the `http force-ntlm` command to change the option. To suppress the header only in certain situations, continue with the procedures below.

Suppress Proxy-Support Header through VPM

To suppress the header using VPM, create a new Web Access Layer. Then:

1. Right click in the Action field to see the drop-down list; select Set.
The Existing Action Object dialog displays.
2. Click New to see the drop-down list; select Control Response Header.
The Add Control Response Header Object dialog displays.

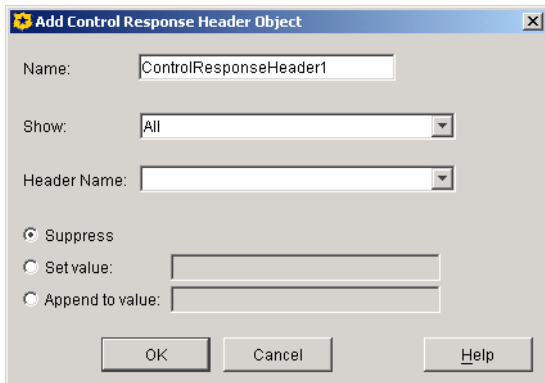


Figure 6-10: Add Control Response Header Object

Section A: Configuring Explicit Proxies

3. Fill in the fields as follows:
 - Name: Enter a meaningful name. This name will display in the Existing Action Objects dialog.
 - Show: Select Custom from the drop-down list.
 - Header Name: Enter Proxy-Support.
 - Make sure the Suppress radio button is selected.
4. Click OK.
5. Scroll to the bottom of the Add Control Response Header Object dialog to see the Proxy-Support header.
6. Click OK.

Suppress Proxy-Support Header through CPL

Use CPL to define the Proxy-Support custom header object and to specify what action you want 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
```

Enabling or Disabling NTLM Authentication for Internet Explorer Clients

The following procedure forces Internet Explorer clients explicitly-proxied through a ProxySG to participate in NTLM authentication. Note that this CLI setting is global, affecting all clients. You can also use VPM or CPL to provide granular control for NTLM authentication. (See "Force NTLM Authentication through VPM" on page 188 and "Force NTLM Authentication through CPL" on page 188.) These commands should only be used if the Proxy-support header is not suitable for the situation.

Note: These procedures can only be done through the CLI. The Management Console is not available.

Do one of the following (note that the default is `http no force-ntlm`):

- To force NTLM authentication for Internet Explorer clients, enter the following command at the (config) command prompt:


```
SGOS#(config) http force-ntlm
```
- To disable NTLM authentication for Internet Explorer clients, enter the following command at the (config) command prompt:


```
SGOS#(config) http no force-ntlm
```

To view all HTTP settings, see "Viewing HTTP Settings through the CLI" on page 177.

Section A: Configuring Explicit Proxies

Force NTLM Authentication through VPM

To use VPM to force NTLM authentication, create a new Web Access Layer. Then:

1. Right click in the Action field to see the drop-down list; select Set.
The Existing Action Object dialog displays.
2. Scroll to the Force NTLM for Server Auth static object; select it.
3. Click OK.

Force NTLM Authentication through CPL

Global configuration of NTLM authentication behavior is set through the CLI command `http force-ntlm` (the default is `http no force-ntlm`). The `http.force_ntlm_for_server_auth()` CPL property can be used to override the global settings for a particular subnet.

To create a rule to force NTLM authentication for explicitly proxied Internet Explorer clients, first define the action, then define the rule.

This example implements the following policies:

- All clients from the “ForceNTLM_subnet” have `force-ntlm` turned on. These clients do not use the Proxy-Support header.
- Requests for all other hosts have `force-ntlm` turned off. These hosts use the Proxy-Support header.

```
define subnet ForceNTLM_subnet
  10.10.0.0/16
end
<Proxy>
  client.address=ForceNTLM_subnet http.force_ntlm_for_server_auth(yes)
  http.force_ntlm_for_server_auth(no)
end
```

Configuring a SOCKS Proxy

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 or UDP protocols. The ProxySG supports both SOCKSv4/4a and SOCKSv5; however, because of increased username and password authentication capabilities and compression support, Blue Coat recommends that you use SOCKS v5.

Note: For Blue Coat compatibility with SOCKS clients, check with customer support.

Understanding SOCKS Compression

Compression over SOCKS is supported for TCP/IP tunnels, which can compress the data transferred between the branch (downstream proxy) and main office (upstream proxy), reducing bandwidth consumption and improving latency.

Section A: Configuring Explicit Proxies

TCP tunnels are created by posting a listener on a static port for protocols that have a well-known port; applications that use dynamic port numbers are handled through the Endpoint Mapper proxy that automatically creates TCP tunnels to ports where Microsoft RPC services are running. (For information on using the Endpoint Mapper proxy, see "Endpoint Mapper Proxy" on page 133.)

Except for enabling the SOCKS proxy, no configuration is required on the main office ProxySG to support SOCKS compression. However, configuration is required on the branch ProxySG to forward data through the SOCKS gateway. You can use policy or the `socks-gateway` CLI options to enable SOCKS compression globally. Using policy, you can enable or disable compression on a per-connection basis on either the client side or the server side.

If SOCKS compression is enabled and the upstream SOCKS gateway does not support it, the connection fails.

Note: Enabling compression on TCP tunnels impacts performance and should be done only when the ProxySG is sized correctly to handle the incremental CPU load.

In a typical 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 "Endpoint Mapper Proxy" on page 133.
- Create any other TCP tunnel proxies you need at the remote office: SMTP, DNS, and the like. For information on configuring TCP tunnels, see "TCP Tunneling" on page 143.
- 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 "SOCKS Gateway Configuration" on page 722.
- 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 192.

Note: In cases where more than two proxies exist in the chain and intermediate proxy nodes are configured to do compression the traffic is forwarded as is. If the intermediate proxy is not configured to do compression, traffic is decompressed before being forwarded to the next proxy.

For more information on deploying SOCKS compression and Endpoint Mapper proxy, refer to the *Blue Coat Edge Deployment Guide*.

Creating and Configuring the Service

Complete the following steps to create a SOCKS proxy and to configure SOCKS-proxy connection and timeout values.

Section A: Configuring Explicit Proxies

To Create a SOCKS Proxy Server through the Management Console

1. Select Configuration>Services>SOCKS Proxy.

The SOCKS Proxy tab displays.

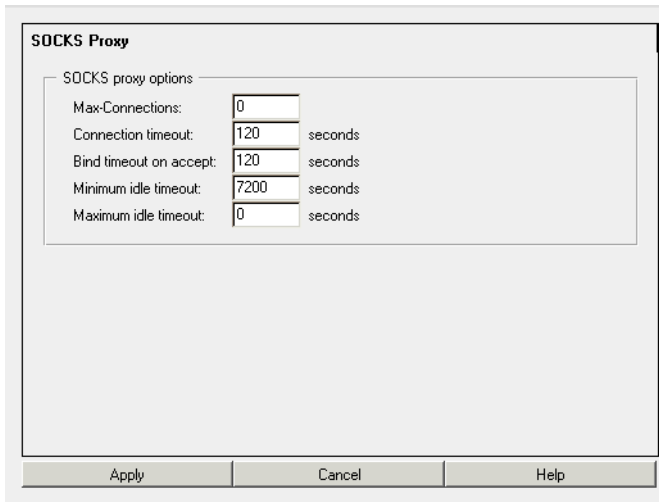


Figure 6-11: SOCKS Proxy Tab

2. Fill in the option fields (described below) as needed. The defaults are displayed and should be sufficient for most purposes.

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>	Set minimum SOCKS client idle time threshold.
Maximum idle timeout	<i>seconds</i>	Set maximum SOCKS client idle time threshold.

To Configure the SOCKS Proxy through the CLI

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) socks-proxy accept-timeout seconds | connect-timeout seconds |
max-connections number | max-idle-timeout seconds | min-idle-timeout seconds
```

2. (Optional) View the results.

```
SGOS#(config) show socks-proxy
max-connections: 0
accept-timeout: 120
connect-timeout: 120
min-idle-timeout: 7200
max-idle-timeout: 0
```

Section A: Configuring Explicit Proxies

Enabling the SOCKS Proxy

Note that a SOCKS port is already configured on port 1080 and enabled.

To Edit an Existing SOCKS Port Service through the Management Console

1. Select Configuration>Services>Service Ports.
2. Highlight the SOCKS server.
3. Click Edit; the Edit Service dialog appears.

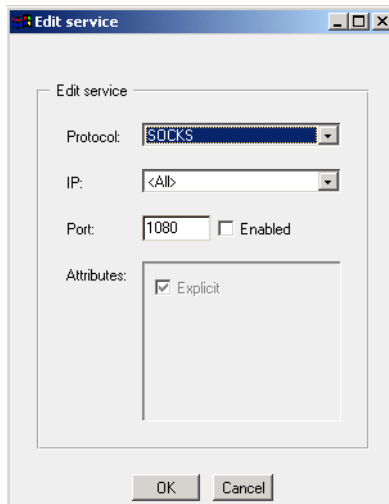


Figure 6-12: SOCKS Add Service Dialog

4. In the Protocol drop-down list, select SOCKS.
5. The default IP address value is all. To limit the service to a specific IP, select the IP from the drop-down list.
6. In the Port field, specify a port number; select Enable.
7. Click OK; Click Apply.

To Edit an Existing SOCKS Port Service through the CLI

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) services
SGOS#(config services) socks
SGOS#(config services socks) enable [ip_address:]port
```

2. (Optional) View the results:

```
SGOS#(config services socks) view
Port:      1080      IP: 10.9.87.85      Type: socks
Properties: explicit, enabled
```

Section A: Configuring Explicit Proxies

Using Policy to Control the SOCKS Proxy

Once the basic configuration for the SOCKS proxy has been set through the Management Console or the CLI, you can use policy to control the SOCKS proxy.

Note: SOCKS compression requires that a SOCKS gateway be set up with SOCKS compression enabled. You can either use policy to configure a gateway for SOCKS compression, or you can configure SOCKS compression while you are configuring the SOCKS gateway. To configure the SOCKS gateway, see "SOCKS Gateway Configuration" on page 722

- To use SOCKS version 5, which allows you to use a SOCKS username/password and SOCKS compression, you must set the version through policy. Note that SOCKS version 4 does not support compression.
 - If using VPM, go to the 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
```
- To use SOCKS compression, you must request SOCKS compression through policy.
 - If using VPM:
 - For global outbound connections (the downstream proxy or branch office location): go to the Forwarding layer, select Source>Set Source Object>New>SOCKS Gateway Compression Object. (Request compression is enabled by default.)
 - For global inbound connections (the upstream proxy or the main office location): go to the Web Access Layer, select Action>New>SOCKS Compression Object. (Allow compression is enabled by default.)
 - If using CPL:
 - For global outbound connections (the downstream proxy or branch office location):

```
<forward>
        client.protocol=tcp socks_gateway(socks_gateway_alias)
        socks_gateway.request_compression(yes|no|default)
```

where `default` refers to the current configuration.
To enable SOCKS compression on a per-connection basis, use a policy similar to the following:

```
<forward>
        client_address=ip_address
        socks_gateway.request_compression(yes|no|default)
```
 - For global inbound connections (the upstream proxy or the main office location):

```
<proxy>
        socks.method=CONNECT socks.allow_compression(yes|no)
```

Allow compression is enabled by default.

 Section A: Configuring Explicit Proxies

Shell Proxies

Shell proxies are those that provide a shell allowing a client to connect to the ProxySG. In this version, only a Telnet shell proxy is supported.

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 6.6: Substitutions Available at New Connection Time

proxy.name or appliance.name	Configured name of the ProxySG.
proxy.address	IP address of the appliance on which this connection is accepted.
proxy.card	Adapter number of the appliance on which this connection is accepted.
client.protocol	This is "telnet".
client.address	IP address of the client.
proxy.primary_address or appliance.primary_address	Primary address of the proxy, not where the user is connected.
release.id	SGOS version.

Customizing Policy Settings for Shell Proxies

To manage a shell proxy through policy, you can use the conditions, properties, and actions list below. For information on using CPL to manage shell proxies, refer to the *Blue Coat ProxySG Content Policy Language Guide*.

Conditions:

All time and date related triggers	proxy.address=
All exception related triggers	proxy.card=
All server_url triggers	proxy.port=
All url triggers	client.protocol=

Section A: Configuring Explicit Proxies

```
All authentication related triggers      user-defined conditions
category=                                client.protocol=telnet
client.address=                           url.scheme=telnet
```

Properties:

```
allow, deny, force_deny                  force_exception(exception_id[, details])
action.action_name{yes|no}               forward(alias_list | no)
All trace() properties                   forward.fail_open(yes | no)
All access_log() properties              reflect_ip(auto|no|client|vip|ip-address)
All log.xxx() properties                 socks_gateway(alias_list | no)
access_server(yes|no)                    socks_gateway.fail_open(yes | no)
authenticate.force(yes|no)               telnet.prompt(no | string)
authenticate(realm)                       telnet.realm_banner(no | string)
exception(exception_id[,                 telnet.welcome_banner(no | string)
details])
```

The banner strings support \$-sign substitutions.

Actions:

```
rewrite(url.host, host_regex_pattern,    log_message()
replacement_pattern)

rewrite(url, url_regex_pattern,          notify_email(subject, body)
replacement_pattern)

set(url_port, port_number)               notify_snmp(message)
```

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 will not be 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.

 Section A: Configuring Explicit Proxies

Telnet Shell Proxies

The Telnet shell proxy allows you to manage a Telnet protocol connection to the ProxySG. Using the Telnet shell proxy, you can do:

- Explicit termination without proxy authentication, where you explicitly connect, through Telnet, to the ProxySG hostname 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 user, the ProxySG challenges for proxy credentials. Once the correct proxy credentials are provided and authenticated, the ProxySG makes an upstream connection and goes into tunnel mode. In this case, the ProxySG 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 "Transparent Proxies" on page 199.
- Transparent termination with proxy authentication, where, after intercepting the transparent connection, the ProxySG challenges for proxy credentials. Once the correct proxy credentials are provided and authenticated, the ProxySG makes an upstream connection and goes into tunnel mode. For more information on configuring a transparent proxy, see "Transparent Proxies" on page 199.

Once in the shell, the following commands are available:

- `help`: Displays available commands and their effects.
- `telnet <server[:port]>`: Makes an outgoing telnet connection to specified server. The colon (:) between server and port can be replaced with a space, if preferred.
- `exit`: Terminates the shell session.

Creating a Telnet Shell Proxy Service

On a new system, Telnet proxy service is configured but disabled on port 23. On an upgrade, a Telnet proxy service is not created.

To enable or create a Telnet proxy service, use `Services>Service Ports` on the Management Console, or `config>services>telnet` from the CLI. For more information, see "Telnet Shell Proxy Service" on page 145.

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 through the Management Console

1. Select `Configuration>Services>Shell Proxies>Telnet Proxy Settings`.

The Telnet Proxy Settings Tab displays.

Section A: Configuring Explicit Proxies

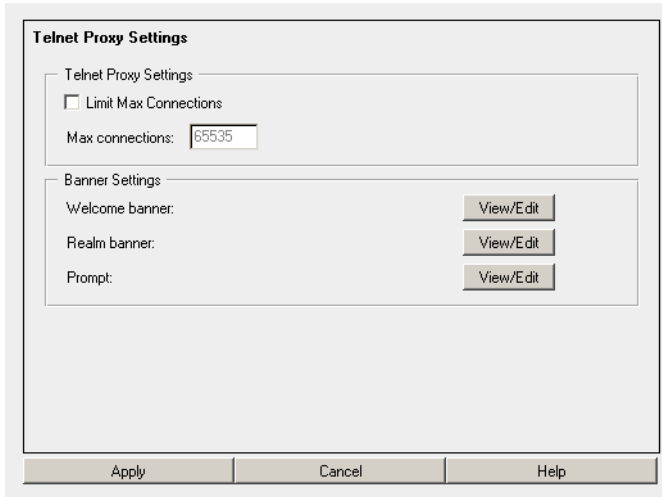


Figure 6-13: Telnet Proxy Settings

2. If you want to set the maximum concurrent connections, check the Limit Max Connections checkbox. Then enter the number of maximum concurrent connections allowed for this service. Allowed values are between 1 and 65535.
3. Set the banner settings:
 - a. To set the Welcome banner message (users see this when they enter the shell), click View/Edit next to the Welcome Banner. The Edit Welcome Banner dialog displays. (If you do not want this banner displayed, remove the text.)

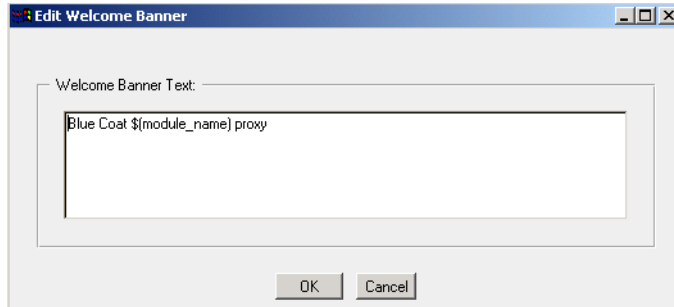


Figure 6-14: Editing Welcome Banner Properties.

Change the banner as necessary. The `$(client.protocol)` text is a CPL variable indicating that Telnet is the protocol. You do not have to use a variable. (For a list of available substitutions, see "Substitutions Available at New Connection Time" on page 193.) When finished, click OK. Click Apply.

Section A: Configuring Explicit Proxies

- b. To set the realms banner message (users see this help message just before they see the Username prompt for proxy authentication), click View/Edit next to the Realms Banner. The Edit Realms Banner dialog displays. (If you do not want this banner displayed, remove the text.)

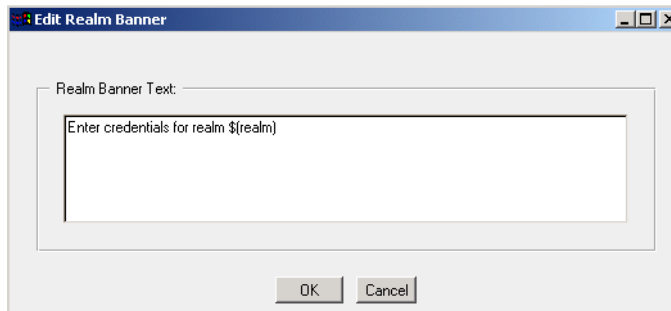


Figure 6-15: Editing Realm Banner Properties

Change the banner as necessary. The `$(realm)` text is a CPL variable indicating the name of the realm. You do not have to use a variable. (For a list of available substitutions, see "Substitutions Available at New Connection Time" on page 193.) When finished, click OK. Click Apply.

- c. To set the prompt, click View/Edit next to the Prompt line. The Prompt dialog displays.

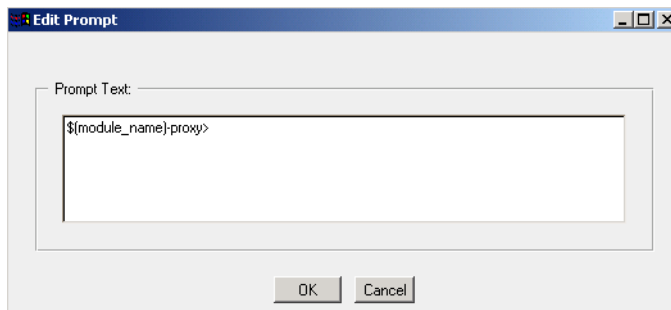


Figure 6-16: Editing the Prompt

Change the banner as necessary. The default is `$(client-protocol)>`, where `$(client-protocol)` is Telnet. You do not have to use a variable. (For a list of available substitutions, see "Substitutions Available at New Connection Time" on page 193.) When finished, click OK. Click Apply.

To Customize Telnet Shell Proxy Settings through the CLI

You can use CPL substitutions when creating welcome and realm banners and Telnet prompts. For a list of available CPL substitutions, see "Substitutions Available at New Connection Time" on page 193.

1. From the (config) prompt, enter the following commands:

Section A: Configuring Explicit Proxies

```
SGOS#(config) shell max-connections number
SGOS#(config) shell welcome-banner welcome-banner-string (Enclose string in
quotes if string includes spaces)
SGOS#(config) shell realm-banner realm-banner-string (Enclose string in
quotes if string includes spaces)
SGOS#(config) shell prompt prompt-string (Enclose string in quotes if string
includes spaces)
```

where:

<code>max-connections</code>	<i>number</i>	Allowed values are between 1 and 65535.
<code>welcome-banner</code>	<i>string</i>	The text a user sees when the shell is entered. You can hide this banner by using <code>shell no welcome-banner</code> .
<code>realm-banner</code>	<i>string</i>	The text a user sees when the realm is entered. You can hide this banner by using <code>shell no realm-banner</code> .
<code>prompt</code>	<i>string</i>	The prompt a user sees when the shell is entered. You can hide the prompt by using <code>shell no prompt</code> .

2. (Optional) To view the shell's settings:

```
SGOS#(config) show shell
max-connections:    Unlimited
prompt:            Telnet #
realm-banner:      Enter credentials for realm Test
welcome-banner:    Welcome to Blue Coat Telnet shell proxy
```

To hide the shell's settings:

```
SGOS#(config) shell no welcome-banner
SGOS#(config) shell no realm-banner
SGOS#(config) shell no prompt
SGOS#(config) shell no max-connections
```

Boundary Conditions for Telnet Shell Proxies

- Telnet credential exchange is in clear text.
- A Telnet proxy cannot be used to communicate with non-Telnet servers (such as webservers on port 80) because Telnet proxies negotiate Telnet options with the client before a server connection can be established.

Section B: Transparent Proxies

Section B: Transparent Proxies

To use transparent proxy, you must:

- Configure the network to redirect client requests
- Create a transparent proxy service

Configuring the Transparent Proxy Hardware

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

- ProxySG Pass-Through card
- ProxySG software bridge
- Layer-4 switch
- WCCP

Configuring the Pass-Through Card for Hardware Bridging

The Blue Coat Pass-Through card is a 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.

Configure a transparent service on the bridge's IP address just like for any other IP address, and it intercepts traffic as usual.

The differences are:

- Forwards traffic: it does not intercept without enabling global IP packet forwarding.
- Proxies for requests on either adapters, so if you have connected one side of the bridge to your Internet connection, you must be careful.

Configuring the ProxySG for Software Bridging

Blue Coat supports a software or *dynamic* bridge that is constructed using a set of installed adapters. Keep in mind the following about software bridges:

- The adapters must 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.
- IP addresses—If any of the adapter interfaces to be added to the bridge already have IP addresses assigned to them, those IP addresses must be removed.

To set up a software bridge, see "Configuring a Software Bridge" on page 78.

Section B: Transparent Proxies

Configuring a Layer-4 Switch for Transparent Proxy

In Transparent Proxy Acceleration, as traffic is sent to the OCS, any traffic sent on TCP port 80 is redirected to the ProxySG Appliances 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 ProxySG 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 Blue Coat ProxySG machine even for requests to a specific IP address.
- ProxySG bypass enabling. You can configure a Layer 4 device to never go through the ProxySG.

The following are very generic directions for configuring transparent proxy using a Layer 4 switch and ProxySG Appliances. The steps to perform depend on the brand of Layer 4 switch. Refer to the Layer 4 switch manufacturer's documentation for details.

To Set up Transparent Proxy Using a Layer-4 Switch and the ProxySG

From the Layer 4 switch:

1. Configure the Layer 4 switch according to the manufacturer's instructions.
2. Configure for global transparent cache switching (TCS). With global TCS, incoming traffic from all devices attached to all ports of the Layer-4 switch is redirected to the ProxySG. Assign an IP address, default gateway, and subnet mask to the Layer-4 switch.
3. Configure TCS using a global policy, enabling redirection for all ports.
4. Identify one or more ProxySG Appliances.
5. Create a device server group.
6. Apply the ProxySG name to the device group.
7. Configure Ethernet interface 2.
8. Disable the redirection policy for the port to which the ProxySG is connected.
9. Configure Ethernet interface 4.
10. Disable the redirection policy for the port to which the router is connected.
11. (Optional) Configure the Layer-4 switch for server load balancing.
12. Save the Layer-4 switch configuration.

From the ProxySG, all you need to do is:

- Define the appropriate IP configurations per the instructions in the *Installation Guide* that accompanied the ProxySG.
- Test the new network configuration.

Configuring WCCP for Transparent Proxy

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 ProxySG Appliances.
- Redirection safeguards—If no ProxySG Appliances are available, redirection stops and the router forwards traffic to the original destination address.

For information on using WCCP with a Blue Coat ProxySG, see Appendix C: “Using WCCP” on page 911.

IP Forwarding

IP Forwarding is a special type of transparent proxy. The ProxySG is configured to act as a gateway. The gateway is configured so that if a packet is addressed to the gateway’s adapter, but not to its IP address, the packet is forwarded toward the final destination. (If IP forwarding is turned off, the packet is rejected as being mis-addressed).

By default, IP forwarding is set to off (disabled) to maintain a secure network.

To Enable IP Forwarding through the Management Console

1. Select Configuration>Network>Routing>Gateways.
2. Select the Enable IP forwarding checkbox.
3. Click Apply.

To Enable IP Forwarding through the CLI

At the (config) command prompt, enter the following command:

```
SGOS#(config) tcp-ip ip-forwarding enable
```

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 explicitly defined (Configuration> Services>Service Ports).

Creating a Transparent Proxy Service

As noted earlier, Blue Coat recommends that you ignore authentication until the proxy service is configured and running.

The below example uses HTTP. Note that two HTTP services are already configured and enabled on SGOS 4.x.

To Create a Transparent HTTP Port Service through the Management Console

1. Select Configuration>Services>Service Ports.

2. Click New; the Add Service dialog appears.

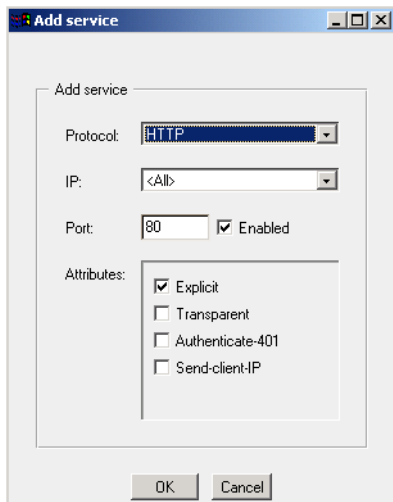


Figure 6-17: HTTP Add Service Dialog

3. In the Protocol drop-down list, select HTTP.
4. The default IP address value is all. To limit the service to a specific IP, select the IP from the drop-down list.
5. In the Port field, specify a port number; select Enable.
6. In the Attributes field, select Transparent.
7. Click OK; Click Apply.

To Create a Transparent HTTP Port Service through the CLI

At the (config) command prompt, enter the following commands:

```
SGOS#(config) services
SGOS#(config services) http
SGOS#(config services http) create [ip_address:]port
SGOS#(config services http) attribute transparent enable [ip_address:]port
SGOS#(config services http) enable [ip_address:]port
```

Example

```
SGOS#(config services http)attribute transparent enable 80
```

To View the Results

```
SGOS#(config services http) view
Port:      8080      IP: 0.0.0.0          Type: http
Properties: explicit, enabled
Port:      80       IP: 0.0.0.0          Type: http
Properties: transparent, enabled
```

Chapter 7: *Using Secure Services*

Secure services allow you to provide the maximum security level for your enterprise. Maximum security is provided by using:

- SSH (with RSA authentication) instead of Telnet for basic communication between machines.
- HTTPS instead of HTTP for secure communication over insecure channels.
- A method of authenticating (identifying your users) and authorizing (limiting what a user can do).

Configuring secure services requires creating and using keypairs and certificates to verify trusted hosts.

This chapter discusses:

- "HTTPS Termination Overview"
- "Configuring HTTPS Termination"
- "Managing the SSL Client"
- "Enabling an HTTPS Service"
- "Configuring HTTP or HTTPS Origination to the Origin Content Server"
- "Configuring DNS Resolution to the Origin Content Server"

HTTPS Termination Overview

Offloading SSL processing from the origin server (referred to as *HTTPS termination*), allows a large number of requests to be processed very quickly from the ProxySG.

The HTTPS termination implementation:

- Combines hardware-based SSL acceleration with full caching functionality.
- Establishes and services incoming SSL sessions.
- Provides SSL v2.0, v3.0, and TLSv1 protocol support.

A common scenario in using HTTPS termination is in conjunction with HTTPS origination. HTTPS termination is used to connect the client to the ProxySG; HTTPS origination is used to connect from the ProxySG to the Origin Content Server (OCS).

Before discussing the specifics of how a ProxySG accelerates HTTPS requests, it is important to understand securing data using HTTPS. There are several RFCs and books on the public key cryptographic system (PKCS). This discussion of the elements of PKCS is relevant to their implementation in SGOS.

The key concepts to understand are:

- Public keys and private keys

- Certificates
- Keyrings
- Cipher Suites
- SSL client

There are many network infrastructure variables that must be considered in your key and certificate management plan. A good publication that addresses such issues is *Understanding Public-Key Infrastructure; Concepts, Standards, and Deployment Considerations* by Carlisle Adams and Steve Lloyd - ISBN 1-57870-166-X.

Public Keys and Private Keys

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.

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

Certificates

Certificates are used to authenticate the identity of a server by associating a public key with a particular server. A certificate is electronic confirmation of the owner of a public key, and contains other information, such as its expiration date. Several kinds of certificates are used.

Self-Signed Certificates

A self-signed certificate is a certificate that you create and authorize yourself that has no official guarantees or authority in the real world. It is mainly used for intranet security.

CA Certificates

The association between a public key and a particular server is performed by a certificate signing authority (CA), who verifies the identity of a server and then signs the server's public key. The resulting certificate can then be offered by the server to clients who can recognize the CA's signature and trust that the server is who it claims to be. Such use of certificates issued by CAs has become the primary infrastructure for authentication of communications over the Internet.

ProxySG appliances come with many popular CA certificates already installed. You can review these certificates using the Management Console or the CLI.

CA certificates installed on the ProxySG are used to verify client certificates (when browsers are configured to offer them during negotiation) and are also required to verify secure servers in communication with the ProxySG.

External Certificates

An external certificate is an X.509 certificate created outside the ProxySG for the purpose of encrypting data, such as access logs, with a public key on the ProxySG so that it can only be decrypted by someone off-box who has the corresponding private key. When you import an external certificate to the ProxySG, you can use it to encrypt your access logs so that only those with the appropriate security credential can decrypt them. See "Customizing the Log Facility: Configuring the Upload Client" on page 759 for information about encrypting access logs.

Wildcard Certificates

Wildcard certificates are certificates that contain wildcard characters in the common name field of an X.509 certificate. Wildcards certificates are typically used in order to share a single certificate among multiple hosts belonging to the same DNS domain.

Wildcard certificates during HTTPS termination are supported. Keep in mind that 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.

Cipher Suites Supported by SGOS

A cipher suite is an object that 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.

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 is configured to allow SSLv2 and v3 as well as TLSv1 traffic. The cipher suites available to use differ depending on whether you configure SSL for version 2, version 3, TLS, or a combination of these.

Table 7.1: SGOS Cipher Suites Shipped with the ProxySG

SGOS Cipher #	Cipher Name	Strength	Exportable	Description
1	RC4-MD5	Medium	No	128-bit key size.
2	RC4-SHA	Medium	No	128-bit key size.
3	DES-CBC3-SHA	High	No	168-bit key size.
4	DES-CBC3-MD5	High	No	168-bit key size.
5	RC2-CBC-MD5	Medium	No	128-bit key size.
6	RC4-64-MD5	Low	No	64-bit key size.
7	DES-CBC-SHA	Low	No	56-bit key size.

Table 7.1: SGOS Cipher Suites Shipped with the ProxySG (Continued)

8	DES-CBC-MD5	Low	No	56-bit key size.
9	EXP1024-RC4-MD5	Export	Yes	56-bit key size.
10	EXP1024-RC4-SHA	Export	Yes	56-bit key size.
11	EXP1024-RC2-CBC-MD5	Export	Yes	56-bit key size.
12	EXP1024-DES-CBC-SHA	Export	Yes	56-bit key size.
13	EXP-RC4-MD5	Export	Yes	40-bit key size.
14	EXP-RC2-CBC-MD5	Export	Yes	40-bit key size.
15	EXP-DES-CBC-SHA	Export	Yes	40-bit key size.

Server Gated Cryptography and International Step-Up

Due to US export restrictions, international access to a secure site requires the site negotiate 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.

The ProxySG 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, the ProxySG supports VeriSign's Global ID Certificate product. The Global ID certificate contains the extra information necessary to implement SGC and International Step-up.

Note: When requesting a Global ID certificate, be sure to specify bluecoat as the server.

SSL Client

The SSL client is used to determine the protocol of outgoing HTTPS connections. The protocol must be specified when a ProxySG obtains content from the origin server using an encrypted connection.

The ProxySG uses one SSL client. The role of the SSL client is to:

- Determine which certificate can be presented to origin servers by associating a keyring with the SSL client.
- Identify the protocol version the ProxySG uses in negotiations with origin servers.
- Identify the cipher suites used.

Configuring HTTPS Termination

To configure HTTPS termination, you must complete the following tasks:

- (Optional) Create or import a keyring. A default keyring is shipped with the system. You can create others.
- (Optional) Create Certificate Signing Requests (CSRs) that can be sent to Certificate Signing Authorities (CAs).
- (Optional, if creating a new keyring) Create or import certificates and associate them with the keyring.
- (Optional) If connections will be forwarded upstream using HTTPS, configure the SSL client appropriately. See "Managing the SSL Client" on page 231.
- (Optional) Set the SSL configuration timeout period.
- Create the HTTPS Service. The keyring should contain the server certificate to present to clients connecting to this service.

Do these steps in order.

Note: These steps must be done with a serial console or SSH connection; you cannot use Telnet.

Before you begin, you should be familiar with the following terms:

CA Certificates	This is a certificate that has been given out by a CA identifying the authority and what public key to use to verify certificates signed by them. CA certificates are used to verify certificates presented by clients during HTTPS termination or to verify certificates presented by servers during HTTPS origination. You only need this certificate if the ProxySG will be obtaining data through an encrypted source.
CA-Certificate Lists	CA-Certificate lists allow you to associate a specific CA certificate (or a list of CA certificates) with the HTTPS service you create.
Certificates	Regular certificates are presented by the ProxySG as server certificates when doing HTTPS Termination or as client certificates when doing HTTPS origination. A certificate can be created (self-signed) or imported from another machine. Certificates and CA Certificates are imported differently on the ProxySG and have different purposes.
Certificate Signing Authority (CA)	CAs receive Certificate Signing Requests and create certificates from the information and the keypair provided. The certificate is then returned to the originator, who can import it into the ProxySG.
Certificate Signing Request (CSR)	CSRs are used to send a keypair and critical information to a Certificate Signing Authority. You can use Blue Coat to create a CSR or you can create a CA Certificate off-line. Once the certificate is sent from the CA, you can import into the ProxySG. (For information on importing CA certificates, see "Importing a CA Certificate" on page 225.)
SSL Client	Only one SSL client can be used on the ProxySG, and only one keyring can be associated with it. If a keyring is associated with the SSL client and you change the association, the old association is overwritten by the new.

SSL Server	When the ProxySG is acting as an SSL server (HTTPS termination), SSL sessions are cached for one hour.
HTTPS Service	A service on which the ProxySG listens for Web requests sent through the HTTPS protocol.
Keyring	A keyring holds a keypair and a certificate, and can be used when configuring secure connections on the ProxySG. When a keyring is created, it only contains a keypair. You can associate a certificate with this keyring. If you have multiple certificates, you can configure multiple keyrings and associate the certificates and the keyrings.

Creating a Keyring

The ProxySG ships with two keyrings already created:

- `default`
- `configuration-passwords-key`

The default keyring contains a certificate and an automatically-generated keypair. Because the default keyring is self-signed, you might want to create other keyrings signed by a well-known Certificate Signing Authority (CA).

You must associate a keyring with the SSL client if the ProxySG will be obtaining content through HTTPS from an origin content server (OCS) that requires a client certificate to be presented. If the OCS requires a client certificate and no keyring is associated with the SSL client, the connections will fail. For information on associating a keyring with the SSL client, "Managing the SSL Client" on page 231.

The `configuration-passwords-key` keyring contains a keypair but does not contain a certificate. It is a keyring created for encrypting passwords in the `show config` command and should not be used for other purposes.

To Create a Keyring through the Management Console

1. Select Configuration>SSL>Keyrings>SSL Keyrings.

The SSL Keyrings tab displays.

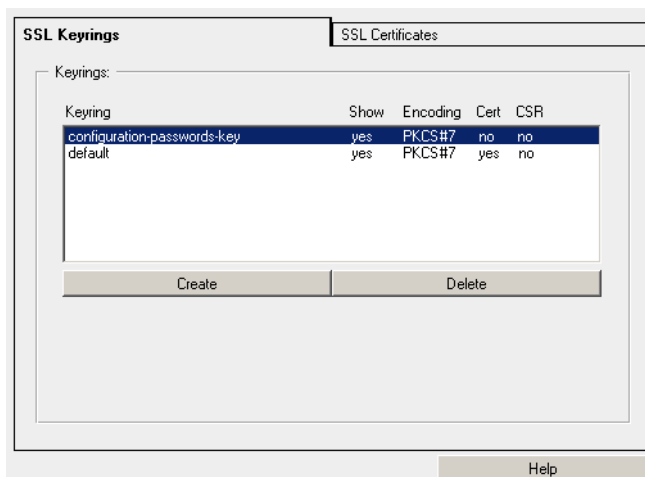


Figure 7-1: SSL Keyring Tab

2. Click Create; the Create Keyring dialog appears.

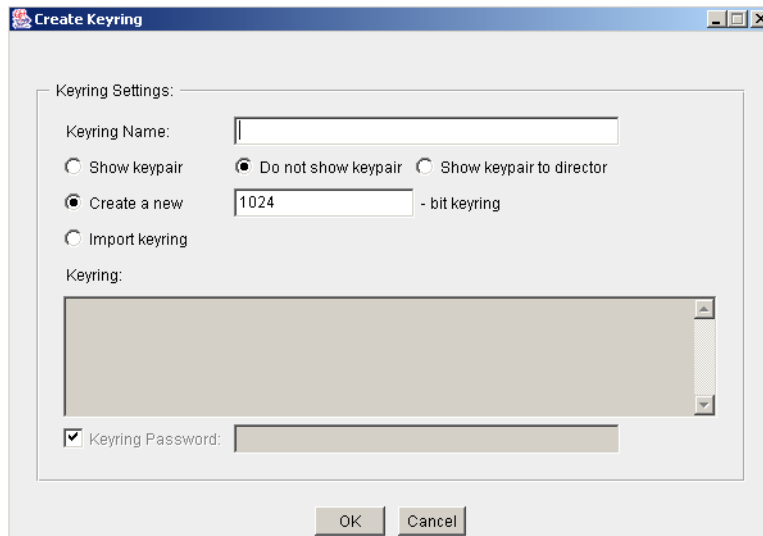


Figure 7-2: Create Keyring Dialog

3. Fill in the dialog window as follows:

- Keyring Name: Give the keyring a meaningful name to you.
- Select the show option you need:
 - Show keypair allows the keys, and everything in 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 of show/show-director/no-show has implications for whether keyrings are included in profiles and backups created by Director. For more information, refer to *Blue Coat Director Configuration and Management Guide*.

- Select the keyring length in the Create a new _____-bit keyring field. A length of 1024 bits is the maximum (and default). Longer keypairs provide better security, but with a slight performance expense on the ProxySG. Be aware that the maximum key length allowed for international export might be different than the default. For deployments reaching outside of the U.S., determine the maximum key length allowed for export.

Click OK. The keyring, containing a keypair, is created with the name you chose. It does not have a certificate associated with it yet. To associate a certificate, see "Managing SSL Certificates" on page 212

-or-

- Select the Import keyring radio button.

The grayed-out Keyring field becomes enabled, allowing you to paste in an already existing keypair. The certificate associated with this keypair must be imported separately. For information on importing a certificate, see "Importing an Existing Keypair and Certificate" on page 216.

If the keypair that is being imported has been encrypted with a password, select Keyring Password and enter the password into the field.

Click OK.

To Create an SSL Keyring through the CLI

At the (config) command prompt, enter the following commands to create an SSL keyring:

```
SGOS#(config) ssl
SGOS#(config ssl) create keyring {show | show-director | no-show} keyring_id
[key_length]
```

where:

- | | |
|---------------|---|
| show | • show: Allows the keys, and everything in the keys, to be exported. |
| show-director | • show-director: Prevents the keypair from being exported. |
| no-show | • no-show: A keyring viewable only if Director is issuing the command using a SSH-RSA connection. |

Note: The choice of show/show-director/no-show has implications for whether keyrings are included in profiles and backups created by Director. For more information, refer to *Blue Coat Director Configuration and Management Guide*.

keyring_id	The name, meaningful to you, of the keyring.
key_length	Longer keypairs provide better security, but with a slight performance expense on the ProxySG Appliance. The default key length used in SGOS and most U.S.-based servers is 1024, which is the maximum key length. Be aware that the maximum key length allowed for international export might be different than the default. For deployments reaching outside of the U.S., determine the maximum key length allowed for export.

To Import a Keyring through the CLI

1. Copy the keyring to the clipboard.
2. At the (config) command prompt, enter the following commands:

```
SGOS#(config) ssl
SGOS#(config ssl) inline keyring show | show-director | no-show keyring_id
[<password>|<"">] eof
Paste keypair here
eof
```

where:

- show: Allows the keys, and everything in the keys, to be exported

- `show-director`: Prevents the keypair from being exported.
- `no-show`: A keyring viewable only if Director is issuing the command using a SSH-RSA connection.
- `password`: If the optional password is provided on the command line, the CLI does not prompt for a password when importing an encrypted keyring. If the optional password is not provided on the command line and if you are trying to import an encrypted keyring, the CLI asks for the password (interactive). (You can also use "" to specify an empty password to make the command non-interactive.)

Note: Director uses non-interactive commands in profiles and overlays to create keyrings.

- `eof`: End-of-file marker. This can be anything, as long as it doesn't also appear in the inline text. (If it appears in the inline text, the `inline` command completes at that point.)

To View the Results of a New or Imported Keyring through the CLI

Note: This example shows the default keyring.

```
SGOS#(config ssl) view keyring
KeyringID: default
  Is private key showable? yes
  Have CSR? no
  Have certificate? yes
  Is certificate valid? yes
  CA: Blue Coat SG110
  Expiration Date: Dec 16 22:37:30 2013 GMT
  Fingerprint: AA:E2:34:DB:5D:06:A7:FF:D8:69:BE:0D:12:FC:34:D5
KeyringID: configuration-passwords-key
  Is private key showable? yes
  Have CSR? no
  Have certificate? no
```

To View a Keypair

Note: This example shows the default keypair, unencrypted.

```
SGOS#(config ssl) view keypair [des | des3 | unencrypted] [keyring_id]
[password]
-----BEGIN RSA PRIVATE KEY-----
MIICWwIBAAKBgQC6t/IhFTYuyczvEN/wT4dcJl3Ar/aEKs/CjL9DPG+ND79sscFe
tfzmLrjBvxJmZYnim6VEMtKb0qH37YQjXwtQFqYAdWe+yKS6kqJ+Rky/mgHX8awL
RvijF1BkLYMG2SOa1YphOTg/v/dPm28TyJ5ZcavM5Atdpa+RRGPPDR1YQwIDAQAB
AoGAE4TVL/Yqsttvq/Ikptd5e/2awWDjsU9UZq8V825m7uUdirxOTZtSs7FgqQhT
YRbuQh0pOqbhc16ihetza8sswGXJe7YYF7d2zQAFwDsvStcsVu1mXQmdhddItGuv
+nZWVMqP/tQIk/NtRhp6IJ2qg4Mu3yEVfDEeHP1Um2nGPbECQQDltYIaoiLW27sa
+O7Rz12geVoVvdROjKg0g0gyT65tRCgqyGv6AXI1+gW11TcP5rh01B9XX3i0wiUp
HejKsompAkEA0BbQNCRXUXZTPyK6R6JaHE0Ji8SSXtLCUN9RZrChdjGc263D6/IV
```

```
/jqppqkLLR2qSibmKDX1ADmYAP9U18ta+CwJAecPBd8TCmwpXIHEch3LRBqPNMQEz
bX/6GfwNzT3/xEQA1szvD9N8a0hhfgqL6Y3v3Rd/lZ0yKv9PG4CTSf9iIQJAL7Jq
+uixkxyaLEibhjvyh7Yoz/64xj9tBviJQg6Ok/b/S2NjGzwcSm/L4Bj7W11URXlf
6YOiISrEN915RjZuzQJAYUlytdCj7pM2nziyO7jrWnY8MmIod3+kHlQajov/OI6Q
Z5gaJ2nLOWicS1SY4MFewHavvRS18yI9JP2q1+6Y/g==
-----END RSA PRIVATE KEY-----
```

Notes

- If you want 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. You can also use "" to specify an empty password to make the command non-interactive.
- 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 will be 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.

Managing SSL Certificates

The ProxySG ships with a certificate associated with a default keyring. The certificate, self-signed and associated with the default keyring, can be reused in other keyrings meant for internal use.

You can add three kinds of SSL certificates:

- A self-signed certificate
- A certificate signed by a CA
- An external certificate

Note: You can create a Certificate Signing Request either on the ProxySG or off-box to send to a Certificate Signing Authority.

To create a self-signed certificate for internal use, continue with the next section. To import an existing certificate, continue with "Importing an Existing Keypair and Certificate" on page 216; to import an external certificate, see "Importing an External Certificate" on page 219; to import a CA certificate, see "Importing a CA Certificate" on page 225.

Adding a Self-Signed Certificate

Self-signed certificates are generally meant for intranet use, not internet.

To Create a Self-Signed Certificate through the Management Console

1. Select Configuration>SSL>Keyrings>SSL Certificates.

The SSL Certificates tab displays.

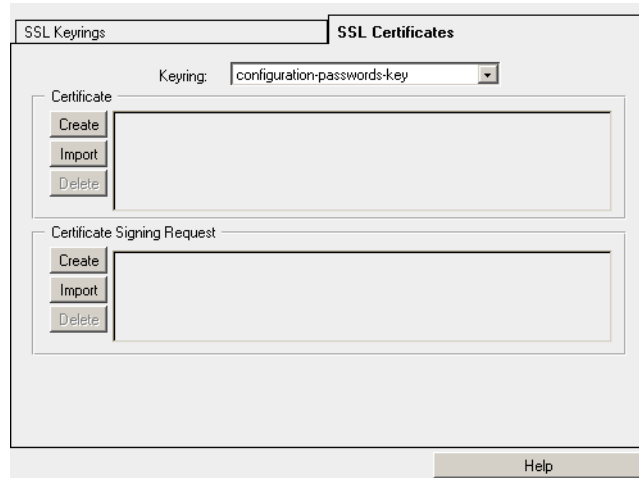


Figure 7-3: SSL Certificates Tab

2. Select the keyring for which you want to add a certificate in the keyring drop-down list.
3. Click Create in the Certificate tab.

The Create Certificate dialog displays.

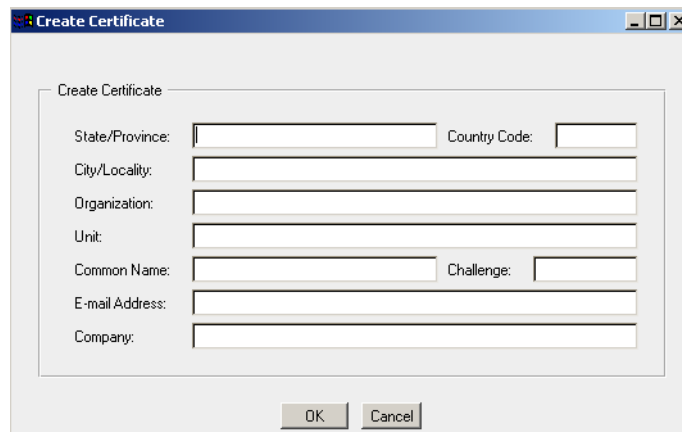


Figure 7-4: Create Certificate Dialog

4. Fill in the fields as appropriate:
 - 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 will be 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 4-16 character alphanumeric challenge.
- ❑ E-mail Address—The email address you enter must be 40 characters or less. A longer email address will generate an error.
- ❑ Company—Enter the name of the company.

5. The Create tab displays the message: *Creating.....*

To Create a Self-Signed Certificate through the CLI

You can create a self-signed certificate two ways: interactively or non-interactively.

Note: Director uses non-interactive commands in profiles and overlays to create self-signed certificates.

To create a certificate using the:

- **interactive version of the `create certificate` command:** continue with the next section.
- **non-interactive version of the `create certificate` command:** skip to "To Create a Self-Signed Certificate Non-interactively Using Create Commands" on page 215.

Note: If you want the certificate to be part of a profile or overlay, the keyring must be configured as showable.

To Create a Self-Signed Certificate Interactively Using Create Commands

1. At the `(config)` command prompt, enter the following commands to interactively create a self-signed certificate.

```
SGOS#(config ssl) create certificate keyring_id
Country code []: US
State or province []: CA
Locality or city []: SV
Organization name []: Blue Coat
Organization unit []: Docs
Common name []: www.bluecoat.com
Email address []: test@bluecoat.com
Challenge []: test
Company name []: Blue Coat
ok
```

where:

Country code	At the Country code prompt, enter the two-character ISO code of the country.
State or province	Name of the state or province where the machine is located.

Locality or city	Name of the town where the machine is located.
Organization name	Name of the company.
Organization unit	Name of the group within the company.
Common name	Verify the <code>Common name</code> is the same as the domain name of the Web site being terminated. If the <code>Common name</code> and site domain name do not match, a client browser generates a warning whenever the ProxySG terminates an HTTPS request for that site. The use of wildcards is supported in the <code>Common name</code> .
Email address	The email address you enter must be 40 characters or less. A longer email address will generate an error
Challenge	At the <code>Challenge</code> prompt, enter a 4-16 character alphanumeric secret.
Company name	Name of the company.

2. View the certificate.

```
SGOS#(config ssl) view certificate keyring_id
-----BEGIN CERTIFICATE-----
MIIB3zCCAzmAwIBAgIBADANBgkqhkiG9w0BAQQFADCBhzELMAkGA1UEBhMCVVMxCzAJBgNVBAGT
AkNBMQswCQYDVQQHEwJTVjESMBAGA1UEChMJQmx1ZSBDb2F0MQ0wCwYDVQQLEwREb2NzMRkwFwY
DVQQDExB3d3cuYmx1ZWNvYXQuY29tMSAwHgYJKoZIhvcNAQkBFhF0ZXN0QGUJsdWVjb2F0LmNvbT
AeFw0wMzAzMDQyMTA2NThaFw0wMzA0MDMyMTA2NThaMIGHMQswCQYDVQQGEwJVUzELMAkGA1UEC
BMCQ0ExCzAJBgNVBACtAlNWMRIwEAYDVQQKEw1CbHV1IENvYXQxDTALBgNVBAsTBBERvY3MxGTAX
BgNVBAMTEHd3dy5ibHV1Y29hdC5jb20xIDAeBgkqhkiG9w0BCQEWEXRlc3RAYmx1ZWNvYXQuY29
tMEwwDQYJKoZIhvcNAQEBBQADAwAwOAIxAK+AGYRMBnjyGr7U0oZUYds106y8uQnxq2PV6qCr4Q
BpN1Vqyr1Fi7ZEaw01yMs5FwIDAQABMA0GCSqGSIb3DQEBAUAUAzEAe8zoYW0igTcGRGG7sBpca
U95J907ZVm8qSU/PQfx1IrDzKdRSQP09Gs1I8MqXi0D
-----END CERTIFICATE-----
```

To Create a Self-Signed Certificate Non-interactively Using Create Commands

Note: If you want the keyring to part of an overlay or profile, the keyring must be configured as showable.

At the `(config)` command prompt, use the following syntax to create a self-signed certificate

```
SGOS#(config ssl)create certificate keyring-id [attribute value] [attribute value]
```

where any or all of the following attribute and value pairs are accepted:

Mandatory:

- `cn` <common name>
- `challenge` <at least four characters>

Optional:

- `c` <2 character country code>
- `o` <organization name>
- `ou` <organizational unit>
- `email` <email-id>
- `state` <state or province>
- `city` <locality or city>
- `company` <company name>

Notes:

- If you do not specify any attributes, the interactive mode is assumed, meaning that the self-signed certificate cannot be created by Director in profiles or overlays.
- The name of the attribute is predefined and the value of the attribute is a *string*. The value can be quoted if it contains white space or other special characters.
- You must specify the name and value together; the order of appearance of multiple name value pairs does not matter. If you omit an attribute, an empty *string* is assumed for the value of the attribute.

Example:

```
SGOS#(config ssl)create certificate keyring-id cn bluecoat challenge test  
c US state CA company bluecoat
```

Importing an Existing Keypair and Certificate

If you have a keypair and certificate from another system, you can import it for use on a different system. You can also import a certificate chain containing multiple certificates in a single operation. Use the `inline certificate` command to import multiple certificates through the CLI.

If you are importing a keyring and one or more certificates onto a ProxySG, first import the keyring, followed by the related certificates. Note that the certificates contain the public key from the keyring, and the keyring and certificates are related.

To Import a Keyring through the Management Console

1. Copy the already-created certificate onto the clipboard.
2. Select Configuration>SSL>Keyrings>SSL Keyrings.
3. Click Create.

The Create Keyring dialog appears.

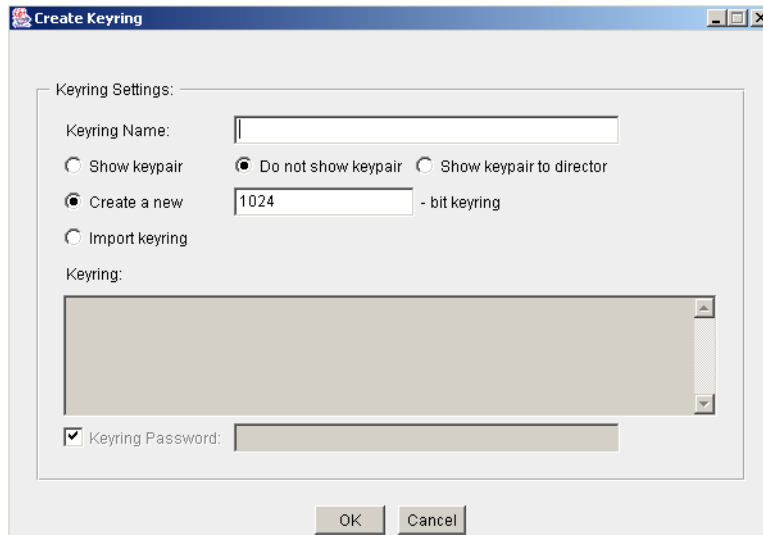


Figure 7-5: Import a Keyring

4. Fill in the dialog window as follows:

- Keyring Name: Give the keyring a meaningful name to you.
- Select the show option you need:
 - `show`: Keyrings created with this attribute can be included as part of a profile or overlay pushed by Director.
 - `show-director`: Keyrings created with this attribute can be included as part of a profile or overlay pushed by Director.
 - `no-show`: Keyrings created with this attribute cannot be part of a profile. The `no-show` option is provided as additional security for environments where the keys will never be used outside of the particular ProxySG.
- Select the keyring length in the Create a new _____ -bit keyring field. A length of 1024 bits is the maximum (and default). Longer keypairs provide better security, but with a slight performance expense on the ProxySG. Be aware that the maximum key length allowed for international export might be different than the default. For deployments reaching outside of the U.S., determine the maximum key length allowed for export.

Click OK. The keyring, containing a keypair, is created with the name you chose. It does not have a certificate associated with it yet. To associate a certificate, see "Managing SSL Certificates" on page 212.

-or-

- Select the Import keyring radio button.

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.

If the keypair that is being imported has been encrypted with a password, select Keyring Password and enter the password into the field.

Click OK.

To Import a Certificate and Associate it with a Keyring through the Management Console

1. Copy the certificate onto the clipboard.
2. Select Configuration>SSL>Keyrings>SSL Certificates and select the keyring that you just imported from the Keyring drop-down list.
3. Click Import in the Certificate field.
4. Paste the certificate into the Import Certificate dialog that appears. Be sure to include the ---BEGIN CERTIFICATE--- and -----END CERTIFICATE---- statements.
5. Click OK.

To Import a Keyring through the CLI Using Inline Commands

1. Copy the keyring to the clipboard.
2. At the (config) command prompt, enter the following commands:

```
SGOS#(config) ssl
SGOS#(config ssl) inline {keyring show | show-director | no-show} keyring_id
eof
Paste keyring here
eof
```

where:

- `show` allows the keys, and everything in the keys, to be exported.
- `no-show` prevents the keypair from being exported.
- `show-director` is a keyring viewable only if Director is issuing the command using a SSH-RSA connection.

Note: The choice of `show/show-director/no-show` 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*.

- `eof`: End-of-file marker. This can be anything, as long as it doesn't also appear in the inline text. (If it appears in the inline text, the inline command completes at that point.)

To Import a Certificate and Associate it with a Keyring through the CLI

Note: The keyring you want to associate with the certificate must already be on this ProxySG. The key and certificate must be imported onto the ProxySG in PEM (base64 encoded text) format.

1. Copy the certificate or certificate chain to the clipboard. Be sure to include the ---BEGIN CERTIFICATE--- and -----END CERTIFICATE---- statements.
2. At the (config) command prompt, enter the following commands:

```
SGOS#(config) ssl
SGOS#(config ssl) inline certificate keyring_id eof
Paste certificate here
```

eof

Deleting an Existing Keyring and Certificate

To Delete a Keyring and the Associated Certificate through the Management Console

1. Select Configuration>SSL>Keyrings>SSL Keyrings.
2. Highlight the name of the keyring that you want to delete.
3. Click Delete.

The Confirm delete dialog appears.

4. Click OK in the Confirm delete dialog that appears.

To Delete a Keyring and the Associated Certificate through the CLI

From the (config) prompt, enter the following commands:

```
SGOS#(config) ssl
SGOS#(config ssl) delete keyring keyring_id
```

Importing an External Certificate

You can import an X.509 certificate into ProxySG to use for encrypting data (see "Customizing the Log Facility: Configuring the Upload Client" on page 759).

To Import an External Certificate through the Management Console

1. Copy the certificate onto the clipboard.
2. Select Configuration>SSL>External Certificates.

The External Certificates tab displays.

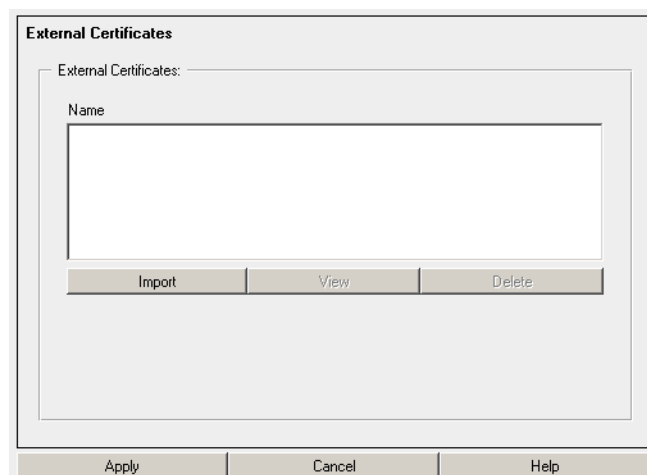


Figure 7-6: External Certificates Tab

3. Click Import.

The Import External Certificate dialog displays.

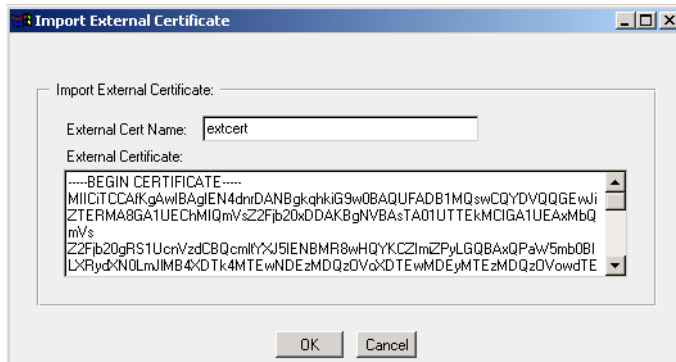


Figure 7-7: Import External Certificate Dialog

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 Import an External Certificate through the CLI Using Inline Commands

1. Copy the certificate or certificate chain to the clipboard. Be sure to include the `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` statements.
2. From the `(config)` prompt, enter the following commands to paste the certificate and enter the eof marker:

```
SGOS#(config) ssl
SGOS#(config ssl) inline external-certificate keyring_id eof
Paste certificate here
eof
```

Viewing an External Certificate

To View an External Certificate through the CLI

```
SGOS#(config) ssl
SGOS#(config ssl) view external-certificate certificate_name
-----BEGIN CERTIFICATE-----
MIICiTCcAfKgAwIBAgIEN4dnrDANBgkqhkiG9w0BAQUFADB1MQswCQYDVQQGEwJi
ZTERMA8GA1UEChMIQmVsZ2Fjb20xDDAKBgNVBAsTA01UTTEkMCIGA1UEAxMmVz
Z2Fjb20gRS1UcnVzdCBQcm1tYXJ5J5IENBMR8wHQYKCCZImZPyLGQBAxQPaw5mb0B1
LXRydXN0LmJlMB4XDtk4MTEwNDEzMDQzOVoXDTEwMDEyMTEzMDQzOVowdTELMakG
A1UEBhMCYmUxETAPBgNVBAoTCEJlbGdhY29tMQwwCgYDVQQLEwNnVE0xJDAiBgNV
BAMTG0JlbGdhY29tIEUHVHJlc3QgUHJpbWVyeSBDOUeFMB0GCgmSjomT8ixkAQMUD2
luZm9AZS10cnVzdC5iZTCBnzANBgkqhkiG9w0BAQEFAAOBjQAwgYkCgYEAqtm5
s9VPak3FQd7BGFqi3Gbb9pk41huJ1XCrc4XsPz6ko0I8Bxy/7LDMf7gaoeXTMxD
V6coeTq1g12kHwrxasU+FCIdWQZv8KYxd9ywSTjmywwP/qpyNIjaKDohWu50Kxuk
21sTFRvzX80ujNLAPj2wy/Dsi4YLwsFEGFPjqNUCAwEAAAMmMCQwDwYDVR0TBAGw
```



```

EgEB/wIBATARBg1ghkgBhvCAQEEBAMCAAcwDQYJKoZIhvcNAQEFBQADgYEAerKx
pbF9M+nC4RvO05OMfwh9Gx1amq6rB1Ev7Ymr3VBCux//SrWknLFhKQpM6oNZSY2v
hmnXgaxHqqRxbInvyxqblSK2qiSyfVms3lf1IsBniFjRjWTpcJfImIDcB1jI+hr
SB0jECfY9t9HorrsgFBKbMRwprkdCJ/9oRiMn7=
-----END CERTIFICATE-----

```

To View the External Certificate Summary through the CLI

```

SGOS#(config) ssl
SGOS#(config ssl)view summary external-certificate
Certificate ID: test1
Is certificate valid? yes
CA: Blue Coat SG3000
Expiration Date: Sep 24 19:33:30 2014 GMT
Fingerprint: 72:D5:7F:9F:B0:CA:D2:54:24:47:A4:7A:37:48:63:D9

```

Deleting an External Certificate

To Delete an External Certificate through the Management Console

1. Select Configuration>SSL>External Certificates.
The External Certificates tab displays.
2. Highlight the name of the external certificate that you want to delete.
3. Click Delete.
The Confirm delete dialog appears.
4. Click OK in the Confirm delete dialog that appears; click Apply.

To Delete an External Certificate through the CLI

From the (config) prompt, enter the following commands:

```

SGOS#(config) ssl
SGOS#(config ssl) delete external-certificate certificate_name

```

About Certificate Chains

A certificate chain is one that requires that the certificates form a chain where the next certificate in the chain validates the previous certificate, going up the chain to the root, which is signed by a well-known root certificate provider. However, expiration is done at the single certificate level and is checked independently of the chain verification. Each certificate in the chain must not have expired for the entire chain to be valid. You can import a certificate chain containing multiple certificates in a single operation.

The valid certificate chain can be presented to a browser. To get the ProxySG to present a valid certificate chain, the keyring for the HTTPS service must be updated.

The ProxySG Appliance's CA-certificate list must also be updated if the ProxySG uses HTTPS to communicate with the origin server and if the ProxySG is configured, through the `ssl-verify-server` option, to verify the certificate (chain) presented by HTTPS server. If the ProxySG uses HTTP to communicate with the origin server, updating the CA-certificate list has no effect.

Creating Certificate Signing Requests

While you must create certificate signing requests (CSR) to get a certificate signed by a Certificate Authority, CSRs are also used for the configuration of regular certificates, certificates that are sent out to clients or servers for external validation.

Creating a CSR through the Management Console

1. Select Configuration>SSL>SSL Keyrings>SSL Certificates.

The SSL Certificates tab displays.

2. Select, from the drop-down list, the keyring for which you need a signed certificate.
3. From the Certificate Signing Request tab, click the Create button.

The Create Certificate-signing-request dialog displays.

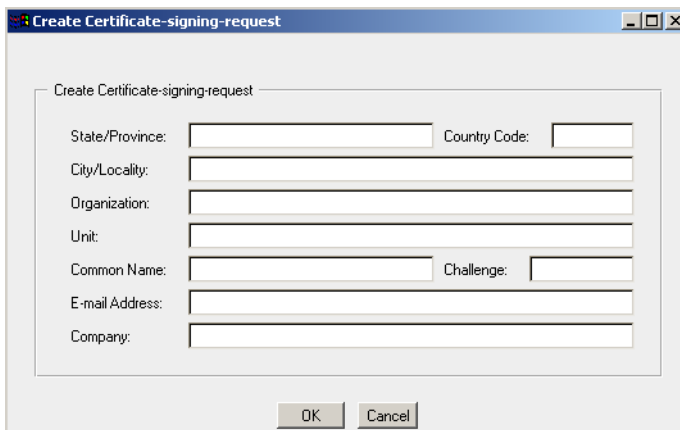


Figure 7-8: Create Certificate-Signing-Request Dialog

4. Fill in the fields as appropriate:
 - 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 will be managing the machine.
 - Common Name—Enter the URL of the company.
 - Challenge—Enter a 4-16 character alphanumeric challenge.
 - E-mail Address—The email address you enter must be 40 characters or less. A longer email address will generate an error.
 - Company—Enter the name of the company.
5. The Create tab displays the message: *Creating...*
6. Click OK.

Creating a CSR through the CLI

You have a choice of using the interactive or non-interactive `create` command.

Note: Director uses non-interactive commands in profiles and overlays to create certificate signing requests.

For more information on Director, refer to the *Blue Coat Director User Guide*.)

To create a CSR using the:

- `interactive create signing-request` command: continue with the next section.
- `non-interactive create signing-request` command: skip to "To Create a Signing Request Non-interactively Using Create Commands" on page 224.

To Create a CSR Interactively using Create Commands

1. At the `(config)` command prompt, enter the following commands to create an SSL CSR:

```
SGOS#(config) ssl
SGOS#(config ssl) create signing-request keyring_id
Country code []: US
State or province []: CA
Locality or city []: SV
Organization name []: Blue Coat
Organization unit []: Docs
Common name []: www.bluecoat.com
Email address []: test@bluecoat.com
Challenge []: test
Company name []: Blue Coat
ok
```

where:

Country code	At the <code>country code</code> prompt, enter the two-character ISO code of the country.
State or province	Name of the state or province where the machine is located.
Locality or city	Name of the town where the machine is located.
Organization name	Name of the company.
Organization unit	Name of the group within the company.
Common name	Verify the <code>Common name</code> is the same as the domain name of the Web site being terminated. If the <code>Common name</code> and site domain name do not match, a client browser generates a warning whenever the ProxySG terminates an HTTPS request for that site. The use of wildcards is supported in the <code>Common name</code> .
Email address	The email address you enter must be 40 characters or less. A longer email address will generate an error
Challenge	At the <code>challenge</code> prompt, enter a 4-16 character alphanumeric secret.
Company name	Name of the company.

2. View the results.

```
SGOS#(config ssl) view signing-request keyring_id
-----BEGIN CERTIFICATE REQUEST-----
MIIBVDCCAQ4CAQAwYcxCzAJBgNVBAYTA1VTMQswCQYDVQQIEwJDQTELMakGA1UEBxMCU1YxEjAQ
BgNVBAoTCUJsdWUgQ29hdDENMAsgAlUECxmERG9jczEZMBcGAlUEAxMQd3d3LmJsdWVjb2F0LmN
vbTEgMB4GCSqGSIb3DQEJARYRdGVzdEBibHV1Y29hdC5jb20wTDANBgkqhkiG9w0BAQEFAAM7AD
A4AjEAobHjK0AsnKV0TcsntWCdfTaNyCgwNDXffxT5FwM0xkzQi0pCSku27CJXn7TahrKRAGMBA
AGgMTAUBGkqhkiG9w0BCQcxBxMFdGVzdAAwGQYJKoZIhvcNAQkCMQwWCkJsdlWUgQ29hdAAwDQYJ
KoZIhvcNAQEEBQADMQBooZfEnzZT2WMMiu3oT9EP3CdtddOTtdBImWUXPdHJGfmlvEJ7HI0cE0W
71JP6pUY=
-----END CERTIFICATE REQUEST-----
```

To Create a Signing Request Non-interactively Using Create Commands

At the (config) command prompt, enter the following commands to create a signing request:

```
SGOS#(config) ssl
SGOS#(config ssl) create signing-request keyring_id [attribute value] [attribute
value]
```

where the following attribute and value pairs are accepted:

Mandatory:

- cn <common name>
- challenge <at least four characters>

Optional:

- c <2 character country code>
- o <organization name>
- ou <organizational unit>
- email <email-id>
- state <state or province>
- city <locality or city>
- company <company name>

Notes:

- If you do not specify any attributes, the interactive mode is assumed, meaning that the CSR cannot be created by Director in profiles or overlays.
- The name of the attribute is predefined and the value of the attribute is a *string*. The value can be quoted if it contains white space or other special characters.
- You must specify the name and value together; the order of appearance of multiple name value pairs does not matter. If you omit an attribute, an empty *string* is assumed for the value of the attribute.

Example:

```
#(config ssl) create signing-request keyring_id cn bluecoat challenge test
c US state CA company bluecoat
```

Managing CA Certificates

If you plan to use certificates issued by well-known Certificate Authorities, you can use the ProxySG to create certificate signing requests (CSRs). These can be sent to the Certificate Authority for signing.

Obtain the keypair and CSR to send to the CA in one of two ways:

- Use the Blue Coat Certificate Signing Request (CSR)
- Obtain the keypair and CSR off-box

Once the signed request is returned to you from the CA, you can import the certificate into the ProxySG. To create a Blue Coat CSR, see "Creating Certificate Signing Requests" on page 222.

Note: If you have a CA certificate that is not on the ProxySG default CA certificate list, you might receive the following message when you attempt to connect to a web site:

```
Network Error (ssl_failed)
A secure SSL session could not be established with the Web Site:.
```

You must import the CA Certificate before the ProxySG can trust the site.

To import a CA Certificate, continue with "Importing a CA Certificate" below.

Importing a CA Certificate

A CA Certificate is a certificate that verifies the identity of a Certificate Authority. The certificate is used by the ProxySG to verify server certificates and client certificates.

To Import an Approved CA Certificate through the Management Console

1. Copy the certificate to the clipboard.
2. Select Configuration>SSL>CA Certificates>CA Certificates.

The CA Certificates tab displays, with its list of existing CA certificates.

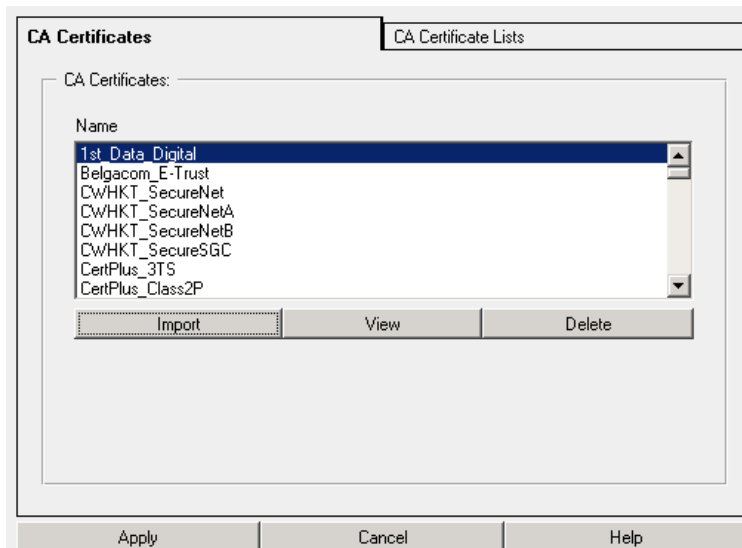


Figure 7-9: CA Certificates

3. Click Import.

The Import CA Certificate dialog displays.




Figure 7-10: Import CA Certificate Dialog

4. Paste the signed CA Certificate into the Import CA Certificate field.
5. Click OK.
6. When the certificate displays in the Certificate tab, click Apply.

To View a CA Certificate through the Management Console

1. Select Configuration>SSL>CA Certificates>CA Certificates.

The CA Certificates tab displays.

2. Select the certificate you want to view.
3. Click View.

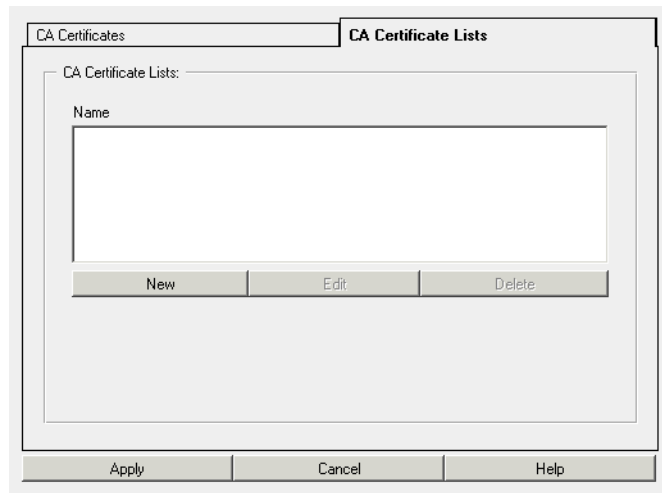


Figure 7-12: SSL CA-Certificates Lists Dialog

The current CA-Certificate lists display in the pane.

2. Click New to create a new list.

The Create CA Certificate List dialog displays.

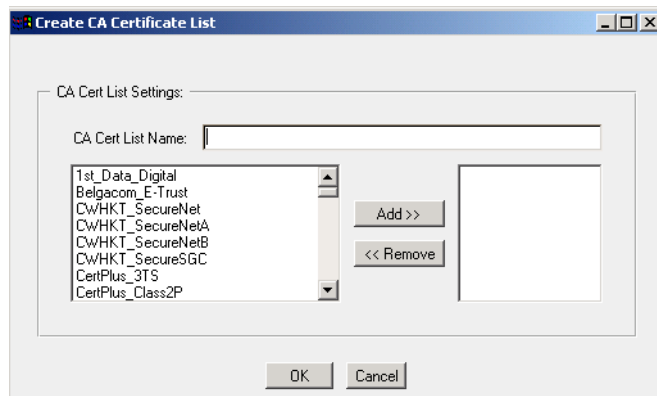


Figure 7-13: Create CA Certificate List Dialog

3. Enter a name meaningful to you for the list in the CA-Certificate List Name.
4. To add CA Certificates to the list, highlight the certificate and click the Add button. You cannot add a certificate to a certificate list if it is not already present.
5. To remove CA Certificates from the list, highlight the certificate in the Add list and click the Remove button.
6. Click OK when you finish; click Apply.

To Create CA-Certificate Lists through the CLI

1. At the `(config)` command prompt, view the CA certificates already existing on the system. You cannot add a certificate to a certificate list if it is not already present.

```
SGOS#(config) ssl
SGOS#(config ssl) view summary ca-certificate
```

All the CA Certificates on the system display.

2. Enter the followings commands to create a list and add existing certificates to the list you just generated.

```
SGOS#(config ssl) create ccl list_name
SGOS#(config ssl) edit ccl list_name
```

The prompt changes, putting you in `ccl` submenu.

```
SGOS#(config ssl ccl list_name) add ca_cert_name
```

3. Repeat the above command until you have entered all the needed certificates. You can have more than one CA-Certificate list. Each list can have an unlimited number of certificates.
4. (Optional) View the list.

```
SGOS#(config ssl ccl list_name) view
CA Certificate ID: VRSN_Secure_Server_CA
Is certificate valid? yes
CA: RSA Data Security, Inc.
Expiration Date: Jan 07 23:59:59 2010 GMT
Fingerprint: 74:7B:82:03:43:F0:00:9E:6B:B3:EC:47:BF:85:A5:93
```

```
CA Certificate ID: DeutscheTelekom
Is certificate valid? yes
CA: Deutsche Telekom AG
Expiration Date: Jul 09 23:59:00 2019 GMT
Fingerprint: 9B:34:0D:1A:31:5B:97:46:26:98:BC:A6:13:6A:71:96
```

```
CA Certificate ID: CWHKT_SecureNetA
Is certificate valid? yes
CA: C&W HKT SecureNet CA Class A
Expiration Date: Oct 15 23:59:00 2009
Fingerprint: E2:D5:20:23:EC:EE:B8:72:E1:2B:5D:29:6F:FA:43:DA
```

Troubleshooting Certificate Problems

- If the client does not trust the Certificate Signing Authority that has signed the ProxySG Appliance's certificate, you will see an error message in the event log similar to the following:

```
2004-02-13 07:29:28-05:00EST "CFSSL:SSL_accept error:14094416:SSL
routines:SSL3_READ_BYTES:sslv3 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 shows 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 from a trusted Certificate Signing Authority in Internet Explorer.

- Import a certificate on the ProxySG that is signed by a well-known Certificate Signing Authority and use that for HTTPS Console access and HTTPS termination.
- If the ProxySG'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 "HTTPS Console (Secure Console)" on page 122.

Managing the SSL Client

Although only one SSL client exists on a ProxySG, the SSL client:

- Determines which certificates can be presented to origin servers if the secure server requires the ProxySG to present a certificate.
- Identifies the protocol the ProxySG uses in negotiations with origin servers.
- Identifies the cipher suites to be used with the certificate.

You can change the protocol and the cipher suites used.

Creating an SSL Client

Only one SSL client can be created on a ProxySG. Creation of the SSL client means that for every HTTPS connection to the destination server, the ProxySG picks the parameters needed for negotiating the SSL connection from the SSL-client configuration. Thus, multiple SSL connections to different HTTPS destination servers can be supported with a single SSL-client configuration. This is similar to a browser where one configuration is used to negotiate multiple connections with different hosts.

When the ProxySG is acting as an SSL client (SSL origination), SSL sessions are re-used until the server forces a fresh handshake or until the same session ID has been used 255 times.

If you just need to change the protocol, the cipher suites, or the keyring associated with the SSL client, you do not need to recreate the client. Continue with "Associating a Keyring and Protocol with the SSL Client" on page 231 or "Changing the Cipher Suites of the SSL Client" on page 233.

To Create the SSL Client through the CLI

```
SGOS#(config ssl) create ssl-client default
defaulting protocol to SSLv2v3TLSv1
defaulting associated keyring-id to default
ok
```

To Delete the SSL Client through the CLI

```
SGOS#(config ssl) delete ssl-client default
ok
```

Associating a Keyring and Protocol with the SSL Client

The SSL client, called default, already exists on the ProxySG. Keyrings that are not used to authenticate encrypted connections do not need to be associated with the SSL client.

Important: Only one keyring can be associated with the SSL client at a time.

To Associate a Keyring with the SSL Client and Change the Protocol Version through the Management Console

1. Select Configuration>SSL>SSL Client.

The SSL Client tab displays.




Figure 7-14: SSL Client

2. To use the SSL client, verify Use SSL Client is selected.
3. Only keyrings with certificates can be associated with the SSL client, displayed in the Keyring drop-down list. Select the keyring you want to use to negotiate with origin content servers through an encrypted connection.
4. You can change the SSL Versions default from SSLv2v3TLSv1 to any other protocol listed in the drop-down list.
5. Click Apply.

To Associate a Keyring and Protocol with the SSL Client through the CLI

1. To associate a keyring with the SSL client, enter the following commands at the (config) command prompt:

```
SGOS#(config) ssl
SGOS#(config ssl) edit ssl-client default
SGOS#(config ssl ssl-client default) keyring-id keyring_id
SGOS#(config ssl ssl-client default) protocol {sslv2 | sslv3 | tlsv1 |
sslv2v3 | sslv2tlsv1 | sslv3tlsv1 | sslv2v3tlsv1}
```

Note: To configure the ProxySG to accept only SSL version 3 traffic, for example, use the `sslv3` parameter. To configure the ProxySG to accept SSL version 2 and version 3 traffic, use the `sslv2v3` parameter.

2. View the results. The results also show the current value of the cipher suites, which is discussed in "Changing the Cipher Suites of the SSL Client" on page 233.

```
SGOS#(config ssl ssl-client default) view
SSL-Client Name      Keyring Name        Protocol
default              default              SSLv2v3TLsv1
```

Changing the Cipher Suites of the SSL Client

The cipher suite sets the encryption method used by the ProxySG. As the encryption key strength is determined by the signed certificate, configuring a higher cipher suite than defined by the certificate will have no affect. Conversely, the cipher suite configuration must be high enough to accommodate certification encryption values.

This can only be done through the CLI.

To Change the Cipher Suite of the SSL Client through the CLI

The default is to use all ciphers.

You have a choice of using the interactive or non-interactive `create` command.

Note: Director uses non-interactive commands in profiles and overlays to create cipher suites.

For more information on Director, refer to the Blue Coat Director *User Guide*.)

To change the cipher suites used through the:

- interactive command: continue with the next procedure.
- non-interactive command: skip to "To Change the Cipher Suites Non-interactively" on page 234.

To Change the Cipher Suites using the Interactive Cipher-Suites Command:

Note that the `Use` Column in the `set cipher-suite` output below indicates that the default is to use all ciphers.

1. Choose the cipher suites 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  RC4-MD5              Medium
2        no   RC4-SHA              Medium
3        no   DES-CBC3-SHA        High
4        no   DES-CBC3-MD5        High
5        no   RC2-CBC-MD5         Medium
6        no   RC4-64-MD5          Low
7        no   DES-CBC-SHA         Low
8        no   DES-CBC-MD5         Low
```

```

9          no          EXP1024-RC4-MD5          Export
10         no          EXP1024-RC4-SHA          Export
11         no          EXP1024-RC2-CBC-MD5       Export
12         no          EXP1024-DES-CBC-SHA       Export
13         no          EXP-RC4-MD5              Export
14         no          EXP-RC2-CBC-MD5          Export
15         no          EXP-DES-CBC-SHA          Export

```

```

Select cipher numbers to use, separated by commas: 1,3,4
ok

```

2. (Optional) View the results. Note the change in the Use column.

```
SGOS#(config ssl ssl-client default) view
```

```

SSL-Client Name      Keyring Name        Protocol
-----
default              default             SSLv2v3TLSv1

```

```

Cipher#  Use  Description          Strength
-----  ---  -----
1         yes  RC4-MD5              Medium
2         no   RC4-SHA              Medium
3         yes  DES-CBC3-SHA        High
4         yes  DES-CBC3-MD5        High
5         no   RC2-CBC-MD5         Medium
6         no   RC4-64-MD5          Low
7         no   DES-CBC-SHA         Low
8         no   DES-CBC-MD5         Low
9         no   EXP1024-RC4-MD5     Export
10        no   EXP1024-RC4-SHA     Export
11        no   EXP1024-RC2-CBC-MD5 Export
12        no   EXP1024-DES-CBC-SHA Export
13        no   EXP-RC4-MD5         Export
14        no   EXP-RC2-CBC-MD5     Export
15        no   EXP-DES-CBC-SHA     Export

```

To Change the Cipher Suites 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-suite cipher-suite

```

where [cipher-suite] can be any combination of the following:

1. rc4-md5
2. rc4-sha
3. des-cbc3-sha
4. des-cbc3-md5
5. rc2-cbc-md5
6. rc4-64-md5
7. des-cbc-sha
8. des-cbc-md5
9. exp1024-rc4-md5
10. exp1024-rc4-sha

```

11.exp1024-rc2-cbc-md5
12.exp1024-des-cbc-sha
13.exp-rc4-md5
14.exp-rc2-cbc-md5
15.exp-des-cbc-sha

```

Notes:

- If you do not specify any attributes, the interactive mode is assumed, meaning that the cipher suites cannot be used by Director in profiles or overlays.
- Multiple cipher suites can be specified on the command line.

Example

```

SGOS#(config ssl ssl-client default) cipher-suite rc4-md5 des-cbc3-md5
exp1024-rc4-md5 exp-des-cbc-sha
ok

```

```

SGOS#(config ssl ssl-client default)view

```

SSL-Client Name	Keyring Name	Protocol
default	default	SSLv2v3TLSv1

Cipher#	Use	Description	Strength
1	no	RC4-MD5	Medium
2	no	RC4-SHA	Medium
3	no	DES-CBC3-SHA	High
4	no	DES-CBC3-MD5	High
5	no	RC2-CBC-MD5	Medium
6	no	RC4-64-MD5	Low
7	no	DES-CBC-SHA	Low
8	no	DES-CBC-MD5	Low
9	no	EXP1024-RC4-MD5	Export
10	no	EXP1024-RC4-SHA	Export
11	no	EXP1024-RC2-CBC-MD5	Export
12	no	EXP1024-DES-CBC-SHA	Export
13	no	EXP-RC4-MD5	Export
14	no	EXP-RC2-CBC-MD5	Export
15	yes	EXP-DES-CBC-SHA	Export

Note: You can use two policy conditions to test the negotiated cipher strength of a securely connected client:

```

client.connection.negotiated_cipher=cipher_suite_list
client.connection.negotiated_cipher.strength=cipher_strength_list

```

For more information on using these two conditions, refer to Chapter 3 of the *Blue Coat ProxySG Content Policy Language Guide*.

Troubleshooting Server Certificate Verification

Server certificate verification can be disabled for all upstream hosts or specific upstream hosts. The ProxySG, by default, verifies the SSL certificate presented by the upstream HTTPS server. However, it fails to negotiate the SSL connection if SSL certificate verification fails. The most common cause of server certificate verification failure is the absence of a suitable CA certificate on the ProxySG. Ensure that the SG is configured with the relevant CA certificates to avoid unwanted verification failures. The default behavior can be changed by using the `http ssl-verify-server` option.

If a forwarding host of type HTTPS server is being used, you can override the default behavior by changing the `ssl-verify-server` option on a per-host basis.

Setting the SSL Negotiation Timeout

The SSL negotiation timeout value dictates the time a ProxySG waits for a new SSL handshake to complete. This value applies to both HTTPS termination and SSL origination.

You can change the default SSL negotiation timeout value if the default, 300 seconds, is not sufficient for your environment. Note that this value can only be changed through the CLI; it cannot be set from the Management Console.

To change the HTTPS termination timeout period, enter the follow commands from the command prompt:

```
SGOS#(config) ssl
SGOS#(config ssl) view ssl-nego-timeout
300
SGOS#(config ssl) ssl-nego-timeout seconds
```

Enabling an HTTPS Service

The final step in creating HTTPS termination is to select a port and enable the HTTPS service on that port. For general information on enabling services, see Chapter 5: “Managing Port Services” on page 121. For more information on enabling the HTTPS Service, see “HTTPS” on page 138.

Configuring HTTP or HTTPS Origination to the Origin Content Server

In previous procedures, you configured HTTPS termination 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 7.2: Scenario 1: HTTPS Termination with HTTPS Origination

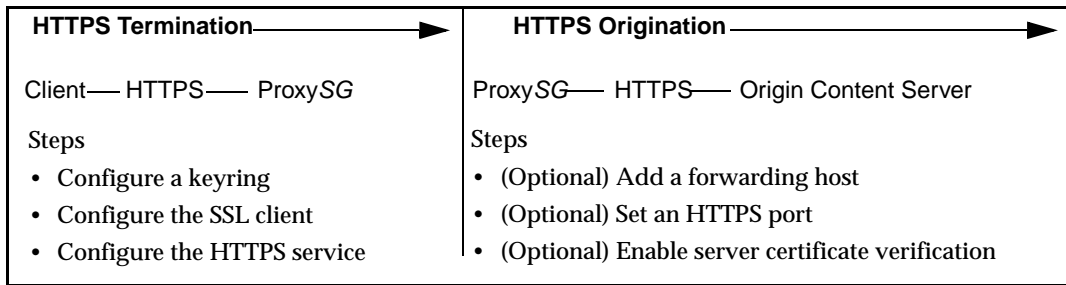
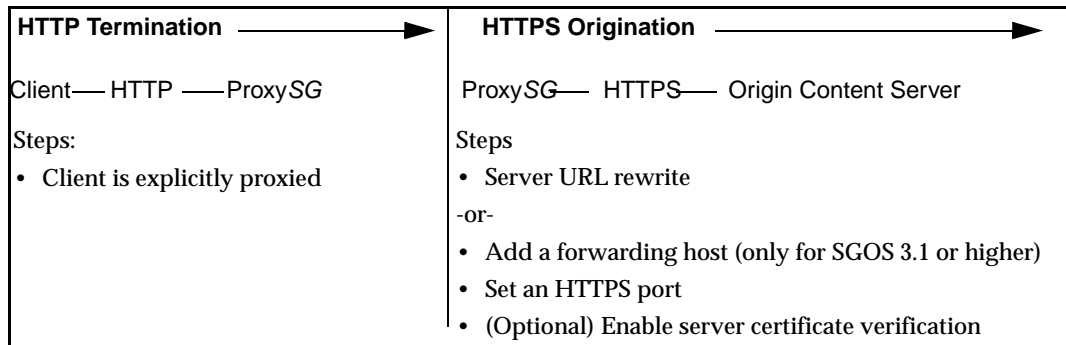


Table 7.3: Scenario 2: HTTP Termination with HTTPS Origination



Using server URL rewrite is the preferred method. For information on rewriting the server URL, refer to the *Blue Coat ProxySG Content Policy Language Guide*.

You can only configure HTTPS origination through the CLI. You cannot use the Management Console.

To Configure HTTPS Origination:

At the (config) command prompt, enter the following commands:

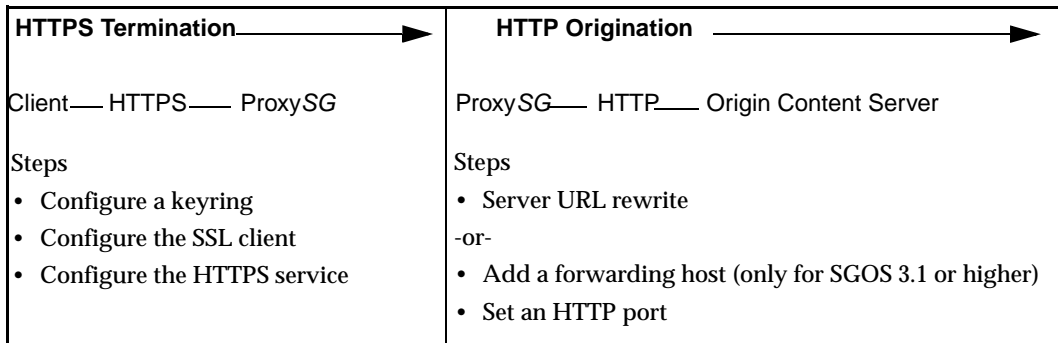
```
SGOS#(config forwarding) create host_alias host_name https[=port_number]
server ssl-verify-server=yes
```

where:

<code>host_alias</code>	<code>ip_address</code>	Specifies the IP address of the OCS.
<code>host_name</code>	<code>url</code>	Specifies the URL of the OCS, such as www.bluecoat.com .
<code>https</code>	<code>[=port_number]</code>	Specifies the port number on the OCS in which HTTPS is listening.
<code>server</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. Proxy is the default.
<code>ssl-verify-server=</code>	<code>yes no</code>	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's not necessary for a secure connection between the proxy and server. For information on using the ProxySG as a reverse proxy, see "Choosing the HTTP Proxy Profile" on page 168.

Table 7.4: Scenario 3: HTTPS Termination with HTTP Origination



Using server URL rewrite is the preferred method. For information on rewriting the server URL, refer to the *Blue Coat ProxySG Content Policy Language Guide*.

You can only configure HTTP origination through the CLI. You cannot use the Management Console.

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:

<code>host_alias</code>	<code>ip_address</code>	Specifies the IP address of the OCS.
<code>host_name</code>	<code>url</code>	Specifies the URL of the OCS, such as www.bluecoat.com .
<code>http</code>	<code>[=port_number]</code>	Specifies the port number on the OCS in which HTTP is listening.
<code>server</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. 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 module, 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.

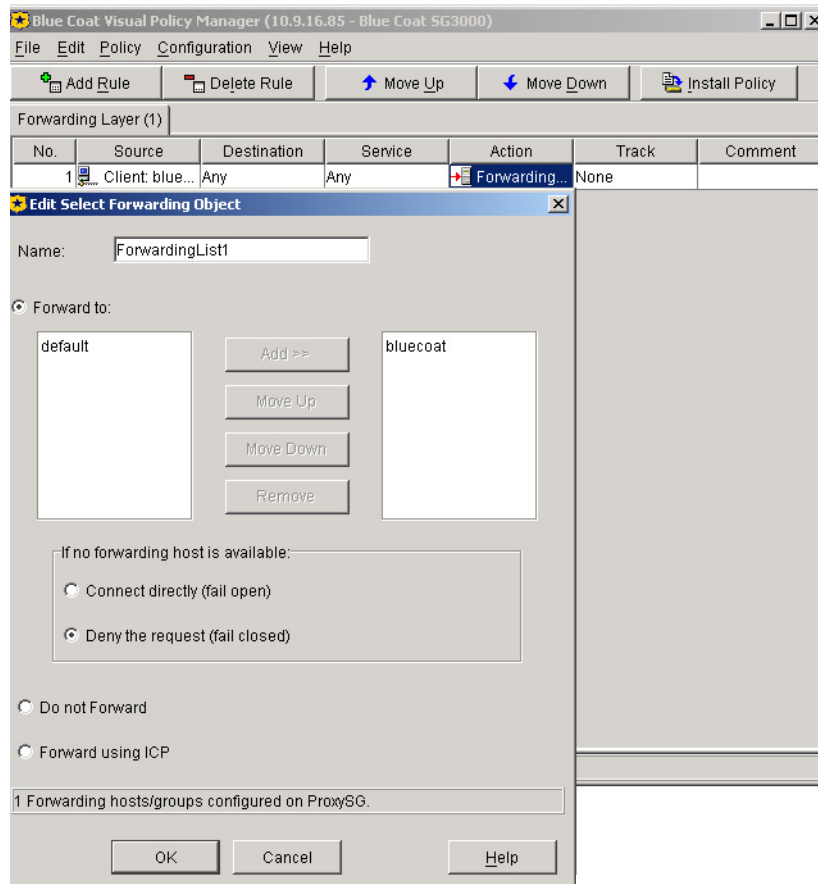


Figure 7-15: Creating a Forwarding Object

Configuring DNS Resolution to the Origin Content Server

In different server accelerator scenarios, you might be required to use DNS resolution to the OCS instead of HTTPS origination.

As long as the DNS that the ProxySG points to correctly resolves the domain name that the client seeks to access, no addition configuration is required. Verify that the ProxySG has the certificate of the Certificate Authority that signs the certificate on the OCS.

Chapter 8: *Security and Authentication*

Enterprise-wide security begins with security on the ProxySG itself, and continues with controlling user access to the Intranet and Internet.

Terms:

proxy	<p>Caches content, filters traffic, monitors Internet and intranet resource usage, blocks specific Internet and intranet resources for individuals or groups, and enhances the quality of Internet or intranet user experiences.</p> <p>A proxy can also serve as an intermediary between a Web client and a Web server and can require authentication to allow identity based policy and logging for the client.</p> <p>The rules used to authenticate a client are based on the policies you create on the ProxySG, which can reference an existing security infrastructure—LDAP, RADIUS, NTLM, and the like, discussed in more detail in Chapter 9: “Using Authentication Services” on page 271.</p>
explicit proxy	<p>A configuration in which the browser is explicitly configured to communicate with the proxy server for access to content.</p> <p>This is the default for the ProxySG, and requires configuration for both browser and the interface card.</p>
transparent proxy	<p>A configuration in which traffic is redirected to the ProxySG without the knowledge of the client browser. No configuration is required on the browser, but network configuration, such as an L4 switch or a WCCP-compliant router, is required.</p>
forward proxy	<p>A proxy server deployed close to the clients and used to access many servers. A forward proxy can be explicit or transparent.</p>
reverse proxy	<p>A proxy that acts as a front-end to a small number of pre-defined servers, typically to improve performance. Many clients can use it to access the small number of predefined servers.</p>
SSL	<p>A standard protocol for secure communication over the network. Blue Coat recommends using this protocol to protect sensitive information.</p>
authentication	<p>The process of identifying a specific user.</p>
authorization	<p>The permissions given to a specific user.</p>
realms	<p>A realm is a named collection of information about users and groups. The name is referenced in policy to control authentication and authorization of users for access to Blue Coat Systems ProxySG services. Note that multiple authentication realms can be used on a single ProxySG. Realm services include NTLM, LDAP, Local, and RADIUS. For detailed information on realms, see Chapter 9: “Using Authentication Services” on page 271.</p>

serial console

A device that allows you to connect to the ProxySG when it is otherwise unreachable, without using the network. It can be used to administer the ProxySG through the CLI. You must use the CLI to use a serial console.

Anyone with access to the serial console can change the administrative access controls, so physical security of the serial console is critical.

SSH and HTTPS are the recommended (and default) methods for managing access to the ProxySG. SSL is the recommended protocol for communication between the ProxySG and a realm's off-box authentication server.

This chapter contains the following sections:

- "Controlling Access to the ProxySG"
- "Controlling Access to the Internet and Intranet"

Section A: Controlling Access to the ProxySG

Section A: Controlling Access to the ProxySG

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.

This section contains:

- "Limiting Access to the ProxySG Appliance"
- "About Password Security"
- "Limiting User Access to the ProxySG—Overview"
- "Moderate Security: Restricting Management Console Access Through the Console Access Control List (ACL)"
- "Maximum Security: Administrative Authentication and Authorization Policy"

Limiting Access to the ProxySG Appliance

You can limit access to the ProxySG Appliance 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 "Changing the Username and Password through the Management Console" on page 43.

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"
- "Securing the Serial Port"

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.

Section A: Controlling Access to the ProxySG

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 CLI. If this option is not enabled, all workstations are allowed to access the CLI. 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 247.)

Securing the Serial Port

If you choose to secure the serial port, you must provide a Setup Console password that will be 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 787.
- Use the reset button (if the appliance has a reset button) to delete all system settings.

To enable the secure serial port, refer to the *Installation Guide* for your platform.

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 plaintext 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, be sure it supports RSA encryption, OAEP padding, and Base64 encoded with no new lines.

Section A: Controlling Access to the ProxySG

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 766
- Archive configuration FTP password—For configuration information, see "Archive Configuration" on page 66
- RADIUS primary and alternate secret—For configuration information, see "Defining RADIUS Realm Properties" on page 296
- LDAP search password—For configuration information, see "LDAP Search & Groups Tab (Authorization and Group Information)" on page 288
- Content filter download passwords—For configuration information, see Chapter 18: "Content Filtering" on page 635

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 applies to Telnet as well. Use of Telnet is not recommended because it is not a secure protocol.

Section A: Controlling Access to the ProxySG

- 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 247.

- 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 can't 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 NTLM with BASIC credentials enabled. For more information on realms, see Chapter 9: "Using Authentication Services" on page 271.
- 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, see "Defining Policies Using the Visual Policy Manager" on page 250.

Section A: Controlling Access to the ProxySG

- 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 247. 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 8.1: 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)		✓		✓
Console Access List evaluated		✓ (if console credentials are offered)	✓	✓ (if console credentials are offered)
CPL <Admin> Layer evaluated		✓ (see Note 1 below)		✓ (see Note 2 below)
Enable password required to enter privileged mode (see Note 2 below)	✓	✓	✓	
CLI <code>line-vty timeout</code> command applies.	✓	✓	✓	
Management Console Login/Logout				✓

Note 1: 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.

Note 2: 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.

Section A: Controlling Access to the ProxySG

To Create an ACL through the Management Console:

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

The Console Access tab displays.

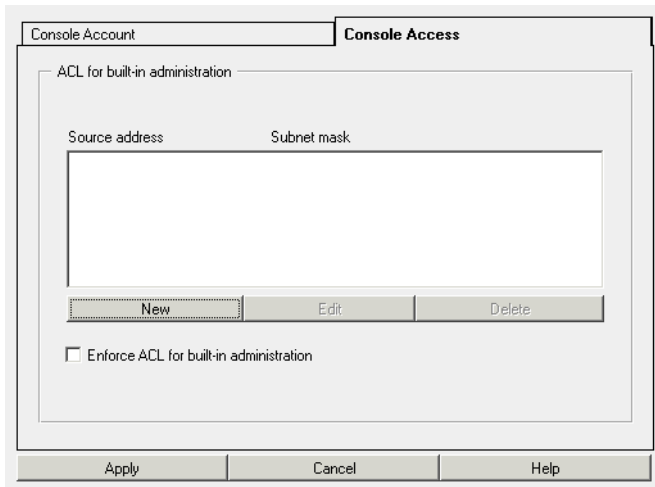


Figure 8-1: Console Access Tab

2. (Optional) To add a new address to the ACL, click New.

The Add List Item dialog is displayed.

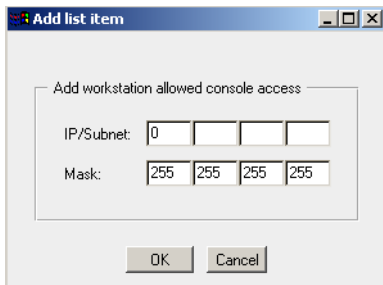


Figure 8-2: Add List Item Dialog

- a. In the IP/Subnet fields, enter a static IP address.
 - b. 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 page.
 - d. Repeat step 2 to add other IP addresses.
3. (Optional) To remove a source address from the ACL, select the address to remove from the Console Access page and click Delete.
 4. (Optional) To change a source IP address, select the IP address to revise and click Edit. See step 2, above, for details.

Section A: Controlling Access to the ProxySG

- 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, deselect the checkbox. The ACL is ignored.

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

- Click Apply.

To Create an ACL through the CLI:

- At the `(config)` command prompt, enter the following command to add workstation IP addresses to the ACL:

```
SGOS#(config) security allowed-access add ip_address [subnet_mask]
```

Note: If you omit the subnet mask, the default subnet mask of 255.255.255.255 is assumed.

- Repeat step 1 for each workstation that you need to add to the console access list.
- At the `(config)` command prompt, enter the following command to enforce the ACL created in step 1

```
SGOS#(config) security enforce-acl enable
```

Only those workstation IP addresses added to the ACL will be able to use the Management console account to administer the ProxySG. Make sure the IP address for the workstation you are using is included in the list.

- To disable the ACL and open through the access to the console account user, enter the following command:

```
security enforce-acl disable
```

- To remove an IP address and subnet mask from the ACL, enter the following command:

```
SGOS#(config) security allowed-access remove ip_address [subnet_mask]
```

Note: If you omit the subnet mask, the default subnet mask of 255.255.255.255 is assumed.

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

Section A: Controlling Access to the ProxySG

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 `config>services>ssh-console`.)

Defining Administrator Authentication and Authorization Policies

The ProxySG uses CPL to define policies, including administrator, authentication, and authorization policies. CPL also allows you to give administrator privileges to users in any external authentication service.

The following summarizes the steps required to define Administrator Authentication and Authorization policies on the ProxySG:

- (Optional) If you need to give administrative access to existing users or groups, create and configure the authentication realm. See Chapter 9: “Using Authentication Services” on page 271 for details on configuring authentication realms.
- Define the policies in the appropriate policy file where you keep the <Admin> Layer layers and rules.
- Load the policy file on the ProxySG.

When you define such policies, make sure you define them in the appropriate policy file(s). For more information on policy files and how they are used, see Chapter 14: “The Visual Policy Manager” on page 453.

Defining Policies Using the Visual Policy Manager

To define policies through the Management Console, use the Visual Policy Manager. When you use the VPM, policies are configured in CPL and saved in the VPM policy file. For examples of Administrator authentication or authorization policy CPL, continue with the next section. The VPM is described in detail in Chapter 14: “The Visual Policy Manager” on page 453.

Defining Policies Directly in Policy Files

To define policies manually, type CPL *rules* directly in one of the two policy files, Central or Local.

Important: Do not manually enter CPL rules directly into the VPM file. The file becomes corrupted.

Section A: Controlling Access to the ProxySG

For specific information on creating policies within the policy files, refer to the *Blue Coat ProxySG Content Policy Language Guide*.

Following are the CPL elements that can be used to define administrator policies for the ProxySG.

To Define Administrator Policies by Editing a Policy File:

1. Open the policy file in a text editor.
2. Define the policies, using the correct CPL syntax.
3. Save the file.
4. Load the policy file (see "Creating and Editing Policy Files" on page 442).

Admin Transactions and <Admin> Layers

Admin transactions execute <Admin> layers. Only a restricted set of conditions, properties, and actions are permitted in <Admin> layers. Table 8.2 lists the conditions permitted in the <Admin> layer:

Table 8.2: <Admin> Layer Conditions

<Admin> Network Connection Conditions	
<code>client_address=<i>ip_address</i> [.subnetmask]</code>	Tests for a match between <i>ip_address</i> and the IP address of the client transaction source.
<code>proxy.port=<i>number</i></code>	Tests for a match between <i>number</i> and the port number for which the request is destined.
<code>proxy.address=<i>ip_address</i></code>	Tests for a match between <i>ip_address</i> and the IP address of the network interface card for which the request is destined.
<code>proxy.card=<i>number</i></code>	Tests for a match between <i>number</i> and the ordinal number associated with the network interface card for which the request is destined.
<Admin> General Conditions	
<code>condition=<i>condition.label</i></code>	Tests if the specified defined condition is true.
<code>release.id=</code>	Tests the ProxySG release id.
<Admin> Date/Time Conditions	
<code>date[.utc]=[<i>date</i> <i>date...date</i>]</code>	Tests for a match between <i>date</i> and the date timestamp associated with the source of the transaction. <i>date</i> specifies a single date of the form YYYY-MM-DD or an inclusive range, as in YYYY-MM-DD...YYYY-MM-DD. By default, date is calculated based on local time. To calculate year based on the Coordinated Universal Time, include the .utc qualifier

Section A: Controlling Access to the ProxySG

Table 8.2: <Admin> Layer Conditions (Continued)

<code>year[.utc]=[year year..year]</code>	Tests for a match between <i>year</i> and the year timestamp associated with the source of the transaction. <i>year</i> specifies a single Gregorian calendar year of the form YYYY or an inclusive range of years, as in YYYY...YYYY. By default, year is calculated based on local time. To calculate year based on the Coordinated Universal Time, include the .utc qualifier.
<code>month[.utc]=[month month...month]</code>	Tests for a match between <i>month</i> and the month timestamp associated with the source of the transaction. <i>month</i> specifies a single Gregorian calendar month of the form MM or an inclusive range of months, as in MM...MM. By default, month is calculated based on local time. To calculate month based on the Coordinated Universal Time, include the .utc qualifier.
<code>weekday[.utc]=[number number...number]</code>	Tests for a match between <i>weekday</i> and the weekday timestamp associated with the source of the transaction. <i>weekday</i> specifies a single day of the week (where Monday=1, Tuesday=2, and Sunday=7) or an inclusive range of weekdays, as in <i>number...number</i> . By default, weekday is calculated based on local time. To calculate weekday based on the Coordinated Universal Time, include the .utc qualifier.
<code>day[.utc]=[day day...day]</code>	Tests for a match between <i>day</i> and the day timestamp associated with the source of the transaction. <i>day</i> specifies a single Gregorian calendar day of the month of the form DD or an inclusive range of days, as in DD...DD. By default, day is calculated based on local time. To calculate day based on the Coordinated Universal Time, include the .utc qualifier.
<code>hour[.utc]=[hour hour...hour]</code>	Tests for a match between <i>hour</i> and the hour timestamp associated with the source of the transaction. <i>hour</i> specifies a single Gregorian hour of the form HH (00, 01, and so forth, through 23) or an inclusive range of hours, as in HH...HH. By default, hour is calculated based on local time. To calculate hour based on the Coordinated Universal Time, include the .utc qualifier.
<code>minute[.utc]=[minute minute...minute]</code>	Tests for a match between <i>minute</i> and the minute timestamp associated with the source of the transaction. <i>minute</i> specifies a single Gregorian minute of the form MM (00, 01, and so forth, through 59) or an inclusive range of minutes, as in MM...MM. By default, minute is calculated based on local time. To calculate minute based on the Coordinated Universal Time, include the .utc qualifier.
<code>time[.utc]=[time time...time]</code>	Tests for a match between <i>time</i> and the time timestamp associated with the source of the transaction. <i>time</i> specifies military time of the form TTTT (0000 through 2359) or an inclusive range of times, as in TTTT...TTTT. By default, time is calculated based on local time. To calculate time based on the Coordinated Universal Time, include the .utc qualifier.

<Admin> Authorization Conditions

Section A: Controlling Access to the ProxySG

Table 8.2: <Admin> Layer Conditions (Continued)

<code>attribute.name =value</code>	Tests if the current transaction is authorized in a RADIUS or LDAP realm, and if the authenticated user has the specified attribute with the specified value. This trigger is unavailable if the current transaction is not authenticated
<code>authenticated={yes no}</code>	Tests if authentication was requested and the credentials could be verified.
<code>group=group_name</code>	If <code>authenticate=yes</code> , the group condition tests the source of the transaction for membership in the specified groupname.
<code>has_attribute.name=boolean</code>	Tests if the current transaction is authorized in an LDAP realm and if the authenticated user has the specified LDAP attribute.
<code>realm=realm_name</code>	If <code>authenticate=yes</code> , the realm condition tests the source of the transaction for membership in the specified realm name.
<code>user=username</code>	If <code>authenticate=yes</code> , the user condition tests the source of the transaction for the expected username.
<code>user.domain= windows_domain_name</code>	(This condition is NTLM-realm specific.) If <code>authenticate=yes</code> , the <code>user_domain</code> condition tests whether the realm type is NTLM and whether the domain component of the username is the expected domain name.

<Admin> Read-only or Read-write Conditions

<code>admin_access=read write</code>	<p><code>read</code> tests whether the source of the transaction has read-only permission for the ProxySG console. <code>write</code> tests whether the source has read-write permission.</p> <p>When an Administrator logs into the CLI, the ProxySG executes an <Admin> transaction that includes the condition <code>admin_access=read</code>. If the transaction is ultimately allowed (all conditions have been met), the user will have read-only access to configuration information through the CLI. Further, when that user executes the CLI <code>enable</code> command, or logs into the Management Console, the ProxySG executes an <Admin> transaction with <code>admin_access=write</code>. If the transaction is allowed, the user will have read-write access within the CLI or the Management Console.</p>
--	---

Table 8.3 lists the properties permitted in the <Admin> layer:

Table 8.3: <Admin> Layer Properties

<Admin> Properties

<code>deny</code>	Refuse service to the source of the transaction.
<code>authenticate(realm_name)</code>	Requests authentication of the transaction source for the specified realm.

Section A: Controlling Access to the ProxySG

Table 8.3: <Admin> Layer Properties (Continued)

<code>authenticate.force()</code>	If 'yes' is specified then forces authentication even if the transaction will be denied. This results in the user information being available for logging. If <code>no</code> , then early denial without authentication will be possible.
<code>allow</code>	Permit further service to the source of the transaction.
<code>log.suppress.field-id()</code>	Controls suppression of the specified field-id in all facilities
<code>log.suppress.field-id[log_list]()</code>	Controls suppression of the specified field-id in the specified facilities.
<code>log.rewrite.field-id()</code>	Controls rewrites of a specific log field in all facilities.
<code>log.rewrite.field-id[log_list]()</code>	Controls rewrites of a specific log field in a specified list of log facilities.

Table 8.4 lists the actions permitted in the <Admin> layer:

Table 8.4: <Admin> Layer Actions

<Admin> Actions

<code>notify_email()</code>	Sends an e-mail notification to the list of recipients specified in the Event Log mail configuration when the transaction terminates.
<code>notify_snmp()</code>	The SNMP trap is sent when the transaction terminates.

Example Policy Using CPL Syntax

To authenticate users against an LDAP realm, use the following syntax in the Local Policy file:

```
<admin>
  authenticate(LDAP_Realm)

<admin>
  group="cn=Administrators,cn=Groups,dc=bluecoat,dc=com" allow
```

This authenticates users against the specified LDAP realm. If the users are successfully authenticated and belong to group *Administrators*, they are allowed to administer the ProxySG.

Section B: Controlling Access to the Internet and Intranet

Section B: Controlling Access to the Internet and Intranet

Once 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.

This section contains:

- "Using Authentication and Proxies"
- "Using SSL with Authentication and Authorization Services"
- "Creating a Proxy Layer to Manage Proxy Operations"

Using Authentication and Proxies

Authentication means that the ProxySG requires proof of user identity in order to make decisions based on that identity. This proof is obtained by sending the client (a browser, for example) a *challenge*—a request to provide credentials. 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 may not be expecting it.

Once the browser supplies the credentials, the ProxySG authenticates them.

Authentication Modes

You can control the way the ProxySG interacts with the client for authentication by controlling the authentication mode. The mode specifies the challenge type and the accepted surrogate credential.

Note: *Challenge type* is the kind of challenge (for example, proxy or origin-ip-redirect) issued.

Surrogate credentials are credentials accepted in place of the user's real credentials.

Section B: Controlling Access to the Internet and Intranet

- **Auto:** The default; the mode is automatically selected, based on the request. Chooses among proxy, origin-IP, and origin-IP-redirect, depending on the kind of connection (explicit or transparent) and the transparent authentication cookie configuration. For streaming transactions, `authenticate.mode(auto)` uses origin mode.
- **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 will not work; origin challenges are then issued.
- **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 will 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 NTLM 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 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. Note that the ProxySG does not support origin-redirects with the CONNECT method.

Note: During cookie-based authentication, the redirect 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. Note that the ProxySG does not support origin-redirects with the CONNECT method.
- **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.

Section B: Controlling Access to the Internet and Intranet

- **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.

Important: Modes that use an IP 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 on using authentication modes, see the *Blue Coat ProxySG 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 NTLM uses connection surrogate credentials. In `sg2` mode, explicit NTLM uses IP surrogate credentials.

To Configure the NTLM Default authenticate.mode Settings:

At the `(config)` command prompt, enter the following commands:

```
SGOS#(config) security default-authenticate-mode {auto | sg2}
```

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.

Section B: Controlling Access to the Internet and Intranet

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 258 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. You will receive an error message similar to the following:

```
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 surrogate credentials are less secure than cookie surrogate credentials and should be avoided if possible. Note that 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 mechanism cannot distinguish between those users

Configuring Transparent Proxy Authentication

The following sections provide general instructions on configuring for transparent proxy authentication.

In addition to configuring transparent proxy authentication, you must also enable a transparent proxy port before the transparent proxy is functional. To enable a transparent proxy port, see "Creating and Editing Services" on page 129.

To Set Transparent Proxy Options through the Management Console:

1. Select Configuration>Authentication>Transparent Proxy.

The Transparent Proxy tab displays.

Section B: Controlling Access to the Internet and Intranet

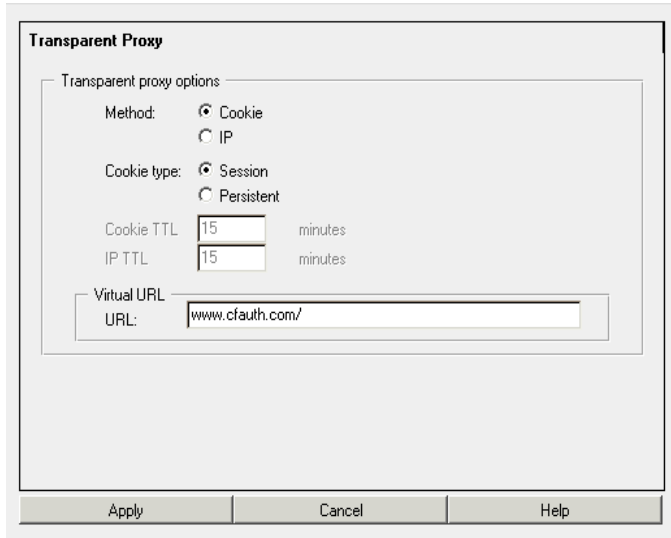


Figure 8-3: Transparent Proxy Tab

2. Select the transparent proxy method—cookie-based or IP address-based. The default is Cookie.
 - If you select Cookie, the Cookie Type radio buttons are available. Click either Session, for cookies that will be deleted at the end of a session, or Persistent, for cookies that will remain on a client machine until the cookie TTL (Time To Live) is reached or the credentials cache is flushed. The default is Session.
 - If you select Persistent Cookies, enter the Cookie TTL. If you choose IP address-based, enter the IP address TTL. The default for each is 15 minutes.

Note: A value of 0 (zero) for the IP address TTL re-prompts the user for credentials once the specified cache duration for the particular realm has expired.

3. Select the Virtual URL. The default is `www.cfauth.com`. Blue Coat recommends you change the virtual hostname to something meaningful to you, preferably the IP address of the ProxySG, unless you are doing secure credentials over SSL. Using the IP address of the ProxySG enables you to be sure that the correct ProxySG is addressed in a cluster configuration.
4. Click Apply.

To Set Transparent Proxy Options through the CLI:

1. At the `(config)` command prompt, enter the following command:


```
SGOS#(config) security transparent-proxy-auth method {cookie | ip}
```

 - a. If you select cookie-based transparent proxy authentication, enter the following command to specify persistent cookies or cookies that persist for the current session only:


```
SGOS#(config) security transparent-proxy-auth cookie {persistent | session}
```

Section B: Controlling Access to the Internet and Intranet

- b. If you select persistent cookies, enter the following command to specify the minutes that the cookie persists:

```
SGOS#(config) security transparent-proxy-auth time-to-live  
persistent-cookie minutes
```

- c. If you choose IP-based transparent proxy authentication, enter the following command to specify that the user be re-prompted for credentials after the number of TTL minutes specified:

```
SGOS#(config) security transparent-proxy-auth time-to-live ip minute
```

A value of 0 (zero) for the IP address TTL re-prompts the user for credentials once the specified cache duration for the particular realm has expired.

2. (Optional step for single ProxySG scenarios, only needed if specifying a different virtual URL than supplied by Blue Coat—www.cfauth.com) To specify the virtual URL for cookie-based authentication, enter the following command:

```
SGOS#(config) security transparent-proxy-auth cookie virtual-url url
```

3. (Optional, if you choose cookie-based) Add the virtual host domain to the DNS service for your organization so that browsers, when redirected to the virtual URL, can resolve the hostname in the URL. (If you use the virtual hostname provided by Blue Coat—www.cfauth.com—you do not need to add the hostname to the DNS service.)

Using SSL with Authentication and Authorization Services

Blue Coat recommends that you use SSL during authentication to secure your user credentials. Blue Coat now supports SSL between the client and the ProxySG and between the ProxySG to LDAP and NTLM authentication servers.

SSL Between the Client and the ProxySG

To configure SSL for to use 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 only use SSL between the client and the ProxySG for origin-style challenges on transparent 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 will be encrypted using the certificate.

Section B: Controlling Access to the Internet and Intranet

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 may not important).

For information on creating a keyring and a certificate, see "Configuring HTTPS Termination" on page 207.

You can use SSL between the ProxySG and NTLM and LDAP authentication servers. For more information, see Chapter 9: "Using Authentication Services" on page 271.

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 on creating effective CPL, refer to the *Blue Coat ProxySG Content Policy Language Guide*.

Table 8.5: <Proxy> Layer Conditions

<Proxy> Layer Conditions	Meaning
<code>admin.access=</code>	Tests the administrative access requested by the current transaction. Can also be used in <Admin> layers.
<code>attribute.name=</code>	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.
<code>authenticated=</code>	Tests if authentication was requested and the credentials could be verified; otherwise, false. Can also be used in <Admin> layers.
<code>bitrate=</code>	Tests if a streaming transaction requests bandwidth within the specified range or an exact match. Can also be used in <Cache> layers.
<code>category=</code>	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.
<code>client_address=</code>	Tests the IP address of the client. Can also be used in <Admin> layers.
<code>client.connection. negotiated_cipher=</code>	Test the cipher suite negotiated with a securely connected client. Can also be used in <Exception> layers.
<code>client.connection. negotiated_cipher. strength=</code>	Test the cipher strength negotiated with a securely connected client. Can also be used in <Exception> layers.

Section B: Controlling Access to the Internet and Intranet

Table 8.5: <Proxy> Layer Conditions (Continued)

<code>client.host=</code>	Test the hostname of the client (obtained through RDNS). Can also be used in <Admin>, <Forward>, and <Exception > layers.
<code>client.host.has_name=</code>	Test the status of the RDNS performed to determine 'client.host'. Can also be used in <Admin>, <Forward>, and <Exception> layers.
<code>client_protocol=</code>	Tests true if the client transport protocol matches the specification. Can also be used in <Exception> layers.
<code>condition=</code>	Tests if the specified defined condition is true. Can be used in all layers.
<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.

Section B: Controlling Access to the Internet and Intranet

Table 8.5: <Proxy> Layer Conditions (Continued)

<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.
<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.

Section B: Controlling Access to the Internet and Intranet

Table 8.5: <Proxy> Layer Conditions (Continued)

<code>im.message.text.substring=</code>	Performs a signed 64-bit range test. 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>send_unknown</code> or <code>receive_unknown</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.
<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.

Section B: Controlling Access to the Internet and Intranet

Table 8.5: <Proxy> Layer Conditions (Continued)

<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 (<i>header_name</i>) 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 <i>header_name</i> . 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 <i>header_name</i> . Can also be used in <Exception> layers.
<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.

Section B: Controlling Access to the Internet and Intranet

Table 8.5: <Proxy> Layer Conditions (Continued)

tunneled=	
url.domain=	Tests if the requested URL, including the domain-suffix portion, matches the specified pattern. Can also be used in <Forward> layers.
url.extension=	Tests if the filename extension at the end of the path matches the specified string. Can also be used in <Forward> layers.
url.host=	Tests if the host component of the requested URL matches the IP address or domain name. Can also be used in <Forward> layers.
url.host.has_name	Test whether the request URL has a resolved DNS hostname. Can also be used in <Exception> layers
url.is_absolute	Test whether the request URL is expressed in absolute form. Can also be used in <Exception> layers
url.host.is_numeric=	This is true if the URL host was specified as an IP address. Can also be used in <Forward> layers.
url.host.no_name=	This is true if no domain name can be found for the URL host. Can also be used in <Forward> layers.
url.host.regex=	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.
url.host.suffix=	Can also be used in <Forward> layers.
url.path=	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.
url.path.regex=	Tests if the regex matches a substring of the path component of the request URL. Can also be used in <Forward> layers.
url.port=	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.
url.query.regex=	Tests if the regex matches a substring of the query string component of the request URL. Can also be used in <Forward> layers.
url.regex=	Tests if the requested URL matches the specified pattern. Can also be used in <Forward> layers.
url.scheme=	Tests if the scheme of the requested URL matches the specified string. Can also be used in <Forward> layers.
user=	Tests the authenticated user name of the transaction. Can also be used in <Admin> layers.
user.domain=	Tests if the authenticated condition is set to yes, the client is authenticated, the logged-into realm is an NTLM realm, and the domain component of the user name is the specified domain. Can also be used in <Admin> layers.

Table 8.5: <Proxy> Layer Conditions (Continued)

<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 8.6: <Proxy> Layer Properties

<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 NTLM uses connection surrogate credentials. In <code>sg2.mode</code> , explicit NTLM uses IP surrogate credentials.
<code>authenticate.redirect_sto</code> <code>red_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 will be 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 <code>yes</code> , then if all clients requesting an object close their connections prior to the object being delivered, the object fetch from the origin server will be abandoned. Can also be used in <Cache> layers.
<code>deny</code>	Denies service. Can be used in all layers except <Exception> and <Forward> layers.

Table 8.6: <Proxy> Layer Properties (Continued)

<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.
<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 'Cache-Control' 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 'Expires' 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 'Pragma: no-cache' 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.

Table 8.6: <Proxy> Layer Properties (Continued)

<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.filter_service ()</code>	Websense is the built in service name for the off-box content filtering service. Can also be used in <Cache> 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.
<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 8.7: <Proxy> Layer Actions

<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 email 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.

Chapter 9: *Using Authentication Services*

Determining and configuring the type of security (such as LDAP, local list, and NTLM) to implement on your network (authorization) is a critical part of managing enterprise security.

Understanding Realms

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 "Configuring Proxies" on page 149.

Multiple authentication realms can be used on a single ProxySG. Multiple realms are essential if the enterprise is a managed service 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—(NTLM, LDAP, RADIUS, Local, Certificate, Sequences, Netegrity SiteMinder[®], Oblix COREid[™], Policy Substitution)
- 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.

SSL Between the ProxySG and the Authentication Server

SSL communication between the ProxySG and LDAP and NTLM authentication servers is supported. In addition, you can also use SSL between the client and the ProxySG. For more information on using SSL between the client and the ProxySG, see "SSL Between the Client and the ProxySG" on page 260.

Configuring a realm to use SSL between the ProxySG and the authentication server is performed on a per-realm basis. Part of the SSL configuration is specifying whether to verify the server's certificate. If the server certificate is to be verified, then the server's certificate must be signed by a Certificate Authority that the ProxySG trusts, and the common name in the server certificate must match the server host as specified in the realm configuration.

The realms use the default SSL client defined on the ProxySG for SSL communications to the authentication servers.

Note: If the browser is configured for on-line checking of certificate revocation, the status check must be configured to bypass authentication.

The chapter contains the following sections:

- "NTLM Realm Authentication and Authorization"
- "LDAP Realm Authentication and Authorization"
- "RADIUS Realm Authentication and Authorization"
- "Local Realm Authentication and Authorization"
- "Certificate Realm Authentication"
- "Netegrity SiteMinder"
- "Obliv COREid"
- "Policy Substitution Realm"
- "Sequence Realm Authentication"
- "Forms-Based Authentication"
- "Managing the Credential Cache"

Section A: NTLM Realm Authentication and Authorization

Section A: NTLM Realm Authentication and Authorization

Windows NT LAN Manager (NTLM) is the authentication protocol used on Windows NT networks.

NTLM is a Microsoft-proprietary protocol that authenticates users and computers based on an authentication challenge and response. When an NTLM realm is used and a resource is requested by the client from the ProxySG, the appliance contacts the user's or computer's account domain to verify identity and then requests an access token. The access token is generated by the domain controller and passed to (and if valid, accepted by) the ProxySG.

Refer to the Microsoft Web site for detailed information about the NTLM protocol and a list of which versions of the Microsoft operating systems use NTLM.

This section discusses the following topics:

- "How Blue Coat Works with NTLM"
- "Creating an NTLM Realm"
- "NTLM Servers"
- "Defining NTLM Realm General Properties"
- "Creating the CPL"

How Blue Coat Works with NTLM

Blue Coat uses a proprietary NTLM agent to better manage NTLM connections.

For NTLM, a single BCAA (Blue Coat Authentication and Authorization Agent) can support multiple ProxySG Appliances; however, only one agent is permitted per realm.

Important: You cannot use CAASNT with SGOS 3.2 and higher.

BCAAA must be installed on a domain controller or member server. If the server where the BCAA is installed and its domain have a trust relationship with other domains, the user is authenticated automatically by the other domains.

Creating an NTLM Realm

To create an NTLM realm, you must provide at least the primary host of the NTLM server for that realm.

To Create an NTLM Realm through the Management Console:

1. Select Configuration>Authentication>NTLM>NTLM Realms.

The NTLM Realms tab displays.

Section A: NTLM Realm Authentication and Authorization

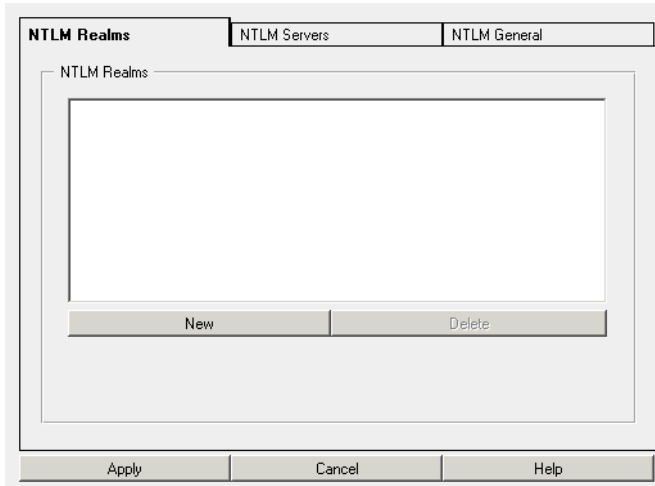


Figure 9-1: NTLM Realms Tab

2. Click New; the Add NTLM Realm dialog displays.

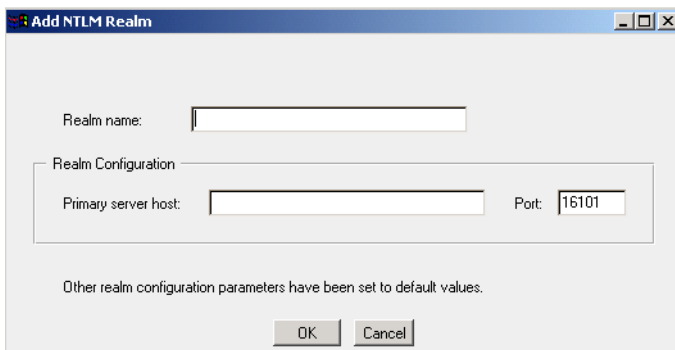


Figure 9-2: Add NTLM Realm

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. Identify the primary server host. You must enter a valid host or an error message is generated.
5. (Optional) The default port is 16101. You can change the port number if the primary server is listening on a different port.
6. Click OK; click Apply.

NTLM Servers

Once you have created an NTLM realm, you can use the NTLM Servers page to change the current default settings.

1. Select Configuration>Authentication>NTLM>NTLM Servers.

The NTLM Servers page displays.

Section A: NTLM Realm Authentication and Authorization

Figure 9-3: NTLM Servers Tab

2. From the Realm Name drop-down list, select the NTLM realm for which you want to change server properties.
3. You must have defined at least one NTLM realm (using the NTLM Realms page) before attempting to set NTLM server properties. If the message `Realms must be added in the NTLM Realms tab before editing this tab is displayed in red at the bottom of this page, you do not currently have any NTLM realms defined`
4. Specify the host and port for the primary NTLM server. The default port is 16101.
5. (Optional) Specify the host and port for the alternate NTLM server. The default port is 16101.
6. (Optional) Under SSL Options, click the SSL enable checkbox to enable SSL.
7. (Optional) By default, if SSL is enabled, the BCAA certificate is verified. If you do not want to verify the BCAA certificate, deselect this checkbox.
8. 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.)
9. Click Apply. Repeat the above steps for additional NTLM realms, up to a total of 40.

To Create and Define an NTLM Realm through the CLI:

1. At the `(config)` prompt, enter the following command to create an NTLM realm:

```
SGOS#(config) security ntlm create-realm realm_name primary_host
[primary_port]
```

where:

<code>realm_name</code>	The name of the NTLM realm.
<code>primary_host</code>	The host for the primary NTLM server.
<code>primary_port</code>	The port for the primary NTLM server. The default port is 16101.

2. To redefine the NTLM realm configuration for the realm you just created, enter the following commands:

Section A: NTLM Realm Authentication and Authorization

```
SGOS#(config) security ntlm edit-realm realm_name
SGOS#(config ntlm realm_name) primary-server primary_host [primary_port]
```

and optionally,

```
SGOS#(config ntlm realm_name) alternate-server alternate_host
[alternate_port]
```

where:

<i>primary_host</i>	The host for the primary NTLM server.
<i>primary_port</i>	The port for the primary NTLM server. The default port is 16101.
<i>alternate_host</i>	The host for the alternate NTLM server.
<i>alternate_port</i>	The port for the alternate NTLM server. The default port is 16101.

3. To enable SSL for this realm and to have the BCAA certificate verified, enter:

```
SGOS#(config ntlm realm_name) ssl enable
SGOS#(config ntlm realm_name) ssl-verify-server enable
```

Defining NTLM Realm General Properties

The NTLM General tab allows you to specify the display name, whether to support Basic and NTLM credentials, the credential cache duration and a virtual URL.

To Configure General Settings through the Management Console:

1. Select Configuration>Authentication>NTLM>NTLM General.

The NTLM General tab displays.

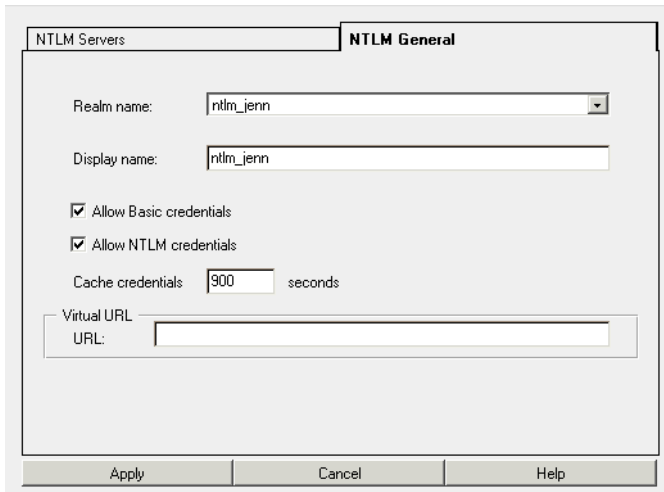


Figure 9-4: NTLM General Tab

2. From the Realm Name drop-down list, select the NTLM realm for which you want to change properties.

Section A: NTLM Realm Authentication and Authorization

Note: You must have defined at least one NTLM realm (using the NTLM Realms tab) before attempting to set NTLM general properties. If the message `Realms must be added in the NTLM Realms tab before editing this tab` is displayed in red at the bottom of this page, you do not currently have any NTLM realms defined.

3. If needed, change the NTLM realm display name. 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 enable or disable support for Basic credentials in the realm by selecting or deselecting the Allow Basic credentials checkbox

Note: At least one Basic or NTLM credential must be supported. Also, if the NTLM realm is part of a sequence realm and is not the first realm in the sequence with try NTLM authentication only once enabled that Basic credentials cannot be disabled in the NTLM realm.

5. You can enable or disable support for NTLM credentials in the realm by selecting or deselecting the Allow NTLM credentials checkbox. Note that at least one of Basic or NTLM credentials must be supported.
6. Specify the length of time, in seconds, that user and administrator credentials received from the NTLM server are cached. Credentials can be cached for up to 3932100 seconds. The default cache duration is 900 seconds (15 minutes).

Note: If you specify 0, traffic is increased to the NTLM server because each authentication request generates an authentication and authorization request to the server. You can specify a virtual URL based on the individual realm. For more information on the virtual URL, see Chapter 8: “Security and Authentication” on page 241.

7. Click Apply.

To Configure General Settings through the CLI:

At the `(config)` command prompt, enter the following commands to configure general settings:

```
SGOS#(config ntlm realm_name) cache-duration seconds
SGOS#(config ntlm realm_name) credentials-basic enable | disable
SGOS#(config ntlm realm_name) credentials-ntlm enable | disable
SGOS#(config ntlm realm_name) display-name name
SGOS#(config ntlm realm_name) virtual-url URL
```

where:

<code>cache-duration</code>	<code>seconds</code>	Specifies the length of time in seconds that user and administrator credentials received from the NTLM server are cached. Credentials can be cached for up to 3932100 seconds. The default value is 900 seconds (15 minutes).
<code>credentials-basic</code>	<code>enable disable</code>	Enables or disables Basic credential support.
<code>credentials-ntlm</code>	<code>enable disable</code>	Enables or disables NTLM credential support.
<code>display-name</code>	<code>name</code>	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.

Section A: NTLM Realm Authentication and Authorization

`virtual-url`

URL

The URL to redirect to when the user needs to be challenged for credentials. see Chapter 8: “Security and Authentication” on page 241 for more details.

Creating the CPL

You can create CPL policies now that you have completed NTLM 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 4.x systems, the default policy condition is *deny*.

Note: Refer to the *Blue Coat ProxySG Content Policy Language Guide* for details about CPL and how transactions trigger the evaluation of policy file layers.

- Every NTLM-authenticated user is allowed access the ProxySG.

```
<Proxy>
  authenticate (NTLMRealm)
```

- Group membership is the determining factor in granting access to the ProxySG.

```
<Proxy>
  authenticate (NTLMRealm)
<Proxy>
  group="Domain\internetusers"
  deny
```

Tips and Boundary Conditions

- Forms authentication modes cannot be used with an NTLM realm that allows only NTLM credentials, a Policy Substitution realm, or a Certificate realm. If a form mode is in use and the authentication realm is any of them, you will receive a configuration error.
- For Windows Internet Explorer NTLM users who want true single-sign-on (allowing Internet Explorer to provide your credentials automatically when challenged), you must set the virtual URL to a hostname that is resolvable to the IP address of the ProxySG by the client machines. Dots (for example, 10.1.1.1) are not allowed.

To define the information in Internet Explorer, go to Internet Options>Security>Local intranet>Sites>Advanced...>Web sites. (If you are an XP user, go to Internet Options>Security>Internet>Custom Level, then check Automatic logon with current username and password.)

For Windows Internet Explorer 6.x users, add the virtual host address to Internet Options>Privacy>Web Sites>Managed Web Sites>Always Allow

Section B: LDAP Realm Authentication and Authorization

Section B: LDAP Realm Authentication and Authorization

Many companies and organizations use the Lightweight Directory Access Protocol (LDAP) as the directory protocol of choice, enabling software to find an individual user without knowing where that user is located in the network topography.

This section discusses the following topics:

- "Overview"
- "Creating an LDAP Realm"
- "LDAP Servers"
- "Defining LDAP Base Distinguished Names"
- "LDAP Search & Groups Tab (Authorization and Group Information)"
- "Customizing LDAP Objectclass Attribute Values"
- "Defining Sequence Realm General Properties"
- "Creating the CPL"

Overview

Blue Coat supports both LDAP v2 and LDAP v3, but recommends LDAP v3 because it uses Transport Layer ProxySG (TLS) and SSL to provide a secure connection between the ProxySG and the LDAP server.

An LDAP directory, either version 2 or version 3, consists of a simple tree hierarchy. An LDAP directory might span multiple LDAP servers. In LDAP v3, servers can return referrals to other servers back to the client, allowing the client to follow those referrals if desired.

Directory services simplify administration; any additions or changes made once to the information in the directory are immediately available to all users and directory-enabled applications, devices, and ProxySG Appliances.

The ProxySG supports the use of external LDAP database servers to authenticate and authorize users on a per-group or per-attribute basis.

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
- Other

The ProxySG also provides the ability to search for a single user in a single root of an LDAP directory information tree (DIT), and to search in multiple Base Distinguished Names (DNs).

You can configure a LDAP realm to use SSL when communicating to the LDAP server.

Section B: LDAP Realm Authentication and Authorization

Configuring LDAP involves the following steps:

- Creating a realm (up to 40) and configuring basic settings.
- Configuring an LDAP server
- Defining LDAP Base Distinguished Names
- Defining Authorization and Group information
- Configuring general LDAP realm settings
- Creating policy

Creating an LDAP Realm

To Create an LDAP Realm through the Management Console:

1. Select Configuration>Authentication>LDAP>LDAP Realms.

The LDAP Realms tab displays.

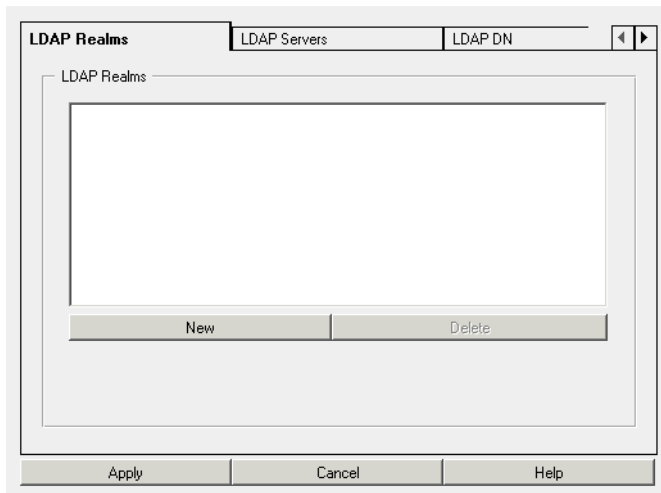


Figure 9-5: LDAP Realms Tab

2. Click New; the Add LDAP Realm dialog displays.

Section B: LDAP Realm Authentication and Authorization

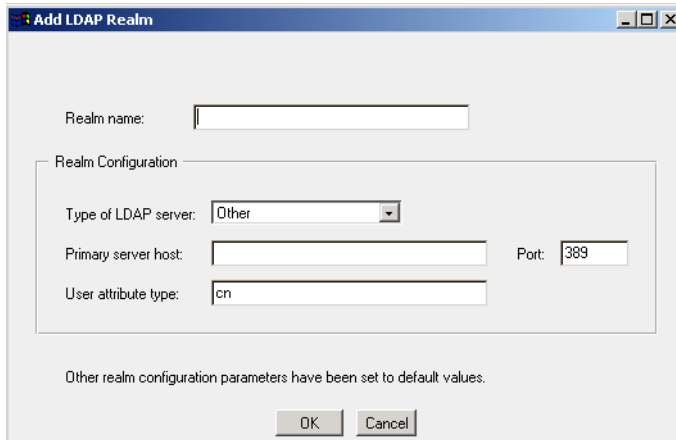


Figure 9-6: Add LDAP Realm

3. In the Real 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. From the Type of LDAP server drop-down list, select the specific LDAP server.
5. Specify the host and port for the primary LDAP server. The host must be entered. The default port number is 389.
6. In the User attribute type field, specify the default user attribute type for the type of LDAP server.

Microsoft Active Directory Server	sAMAccountName=
Novell NDS/eDirectory Server/Other	cn=
Netscape/iPlanet Directory Server	uid=

7. Click OK; click Apply.

LDAP Servers

Once you have created an LDAP realm, you can use the LDAP Servers page to change the current default settings.

To Edit LDAP Server Properties through the Management Console:

Note that the default values exist. You do not need to change these values if the default settings are acceptable.

1. Select Configuration>Authentication>>LDAP>LDAP Servers.

The LDAP Servers page displays.

Section B: LDAP Realm Authentication and Authorization

Figure 9-7: LDAP Servers Tab

- From the Realm Name drop-down list, select the LDAP realm for which you want to change server properties.

Note: You must have defined at least one LDAP realm (using the LDAP Realms tab) before attempting to set LDAP server properties. If the message `Realms must be added in the LDAP Realms tab before editing this tab` is displayed in red at the bottom of this page, you do not currently have any LDAP realms defined.

- From the Type of LDAP server drop-down list, select the specific LDAP server.
- 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 click the Follow referrals checkbox to allow the client to follow referrals to other servers. (This feature is not available with LDAP v2.) The default is Disabled.
- Specify the host and port for the primary LDAP server. The host must be entered. The default port number is 389.
- (Optional) Specify the host and port for the alternate LDAP server. The default port is 389.
- (Optional) Under SSL Options, click Enable SSL to enable SSL. You can only select this option if you are using LDAP v3.
- (Optional) By default, if SSL is enabled, the LDAP server certificate is verified. If you do not want to verify the server certificate, disable this setting.
- (Optional) Change the timeout request for the server from its default of 60 seconds.
- Click Apply. Repeat the above steps for additional LDAP realms, up to a total of 40.

To Define a Realm and Edit LDAP Server Properties through the CLI:

- At the `(config)` command prompt, enter the following command to create an LDAP realm:

Section B: LDAP Realm Authentication and Authorization

```
SGOS#(config) security ldap create-realm {ad | iplanet | nds | other}
realm_name [base_dn] primary_host [primary_port]
```

where:

<code>{ad iplanet nds other}</code>	The type of LDAP realm to create. <code>ad</code> specifies a Microsoft Active Directory realm; <code>iplanet</code> specifies a Netscape/Sun iPlanet realm; <code>nds</code> specifies a Novell NDS/eDirectory realm; <code>other</code> specifies a realm of any other type.
<code>realm_name</code>	The name of the new LDAP realm.
<code>base_dn</code>	The distinguished name (DN) that will be used as the unique key for the LDAP group database; the distinguished name of the key entry and all entries below it in the directory tree. You can specify additional Base DNs after the realm has been created. For example: <code>ou=insidesales, o=toolsdivision</code> . A Base DN can be up to 128 characters long. (In Netscape/iPlanet Directory Server, Base DN is also known as the Root DN.) See Table 9.1 for sample DN entries. Note that at least one base DN is required for authentication to succeed, although you can create a realm without a base DN.
<code>primary_host</code>	The host for the primary LDAP server.
<code>primary_port</code>	The port for the primary LDAP server. The default port is 389.

- To redefine the newly-created LDAP realm authentication properties, enter the following commands:

```
SGOS#(config) security ldap edit-realm realm_name
SGOS#(config ldap realm_name) primary-server host [port]
```

and, optionally:

```
SGOS#(config ldap realm_name) alternate-server host [port]
SGOS#(config ldap realm_name) distinguished-name base-dn clear
SGOS#(config ldap realm_name) distinguished-name base-dn add base_DN
SGOS#(config ldap realm_name) protocol-version {2 | 3}
SGOS#(config ldap realm_name) referrals-follow {enable | disable}
SGOS#(config ldap realm_name) spooof-authentication none | origin | proxy
SGOS#(config ldap realm_name) ssl enable | disable
SGOS#(config ldap realm_name) ssl-verify-server enable | disable
SGOS#(config ldap realm_name) exit
SGOS#(config ldap realm_name) timeout seconds
```

where

<code>alternate-server</code>	<code>host [port]</code>	The host for the secondary LDAP server. The port can also be added, if you need it to be other than the default (389).
-------------------------------	--------------------------	--

Section B: LDAP Realm Authentication and Authorization

distinguished name base-dn	clear add <i>base_DN</i>	<p>Clears the existing <i>base_DN</i> or adds the specified <i>base_DN</i>. The distinguished name (DN) that will be used as the unique key for the LDAP group database; the distinguished name of the key entry and all entries below it in the directory tree. You can specify additional Base DNs after the realm has been created. For example: <i>ou=insidesales, o=toolsdivision</i>. A Base DN can be up to 128 characters long. (In Netscape/iPlanet Directory Server, Base DN is also known as the Root DN.) See Table 9.1 for sample DN entries.</p> <p>Note that at least one base DN is required for authentication to succeed, although you can create a realm without a base DN.</p>
protocol-version	2 3	<p>The LDAP version you want to use. LDAP v3 is the default, allowing you to use the <i>referrals-follow</i> argument and to use SSL.</p>
referrals-follow	enable disable	<p>Allows the client to follow referrals to other servers. This argument is not available if you use LDAP v2.</p>
spooof-authentication	none origin proxy	<p>Enables/disables the forwarding of authenticated credentials to the origin content server or for proxy authentication. You can only choose one.</p> <ul style="list-style-type: none"> • If set to <i>origin</i>, the spoofed header will be an Authorization: header. • If set to <i>proxy</i>, the spoofed header will be a Proxy-Authorization: header. • If set to <i>none</i>, no spoofing will be done. <p>Flush the entries for a realm if the <i>spooof-authentication</i> value is changed to ensure that the <i>spooof-authentication</i> value is immediately applied.</p>
ssl	enable disable	<p>Enables or disables SSL. This argument is not available if you use LDAP v2.</p>
ssl-verify-server	enable disable	<p>By default, if SSL is enabled, the LDAP server certificate is verified. If you do not want to verify the server certificate, disable this setting.</p>

Section B: LDAP Realm Authentication and Authorization

```
SGOS#(config ldap          seconds
realm_name) timeout
```

Note that this command is not in the `edit-realm` submenu. Changes the timeout request for the server from its default of 60 seconds.

- (Optional) Once in the `edit-realm` submenu, use the `?` command to view all of the `edit-realm` commands available.

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. Table 9.1 lists some examples of distinguished name attributes.

Table 9.1: 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>mail=email address</code>	User or group email 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.
<code>streetAddress=street address</code>	Street number and address of user or group defined by the entry. Example: <code>streetAddress= 650 Almanor Avenue Sunnyvale, California 94085-3515</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.

Section B: LDAP Realm Authentication and Authorization

Table 9.1: Distinguished Name Attributes (Continued)

<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 through the Management Console:

1. Select Configuration>Authentication>LDAP>LDAP DN.

The LDAP DN tab displays.

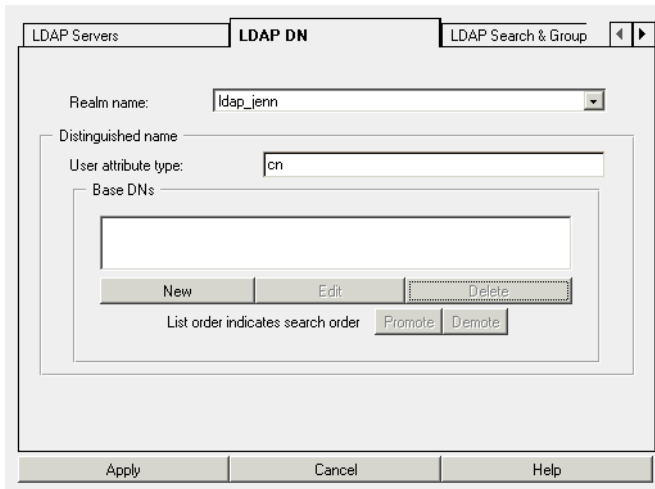


Figure 9-8: LDAP DN Tab

2. From the Realm Name drop-down list, select the LDAP realm for which you want to change DN properties.

Note: You must have defined at least one LDAP realm (using the LDAP Realms tab) before attempting to set LDAP server properties. If the message `Realms must be added in the LDAP Realms tab before editing this tab` is displayed in red at the bottom of this page, you do not currently have any LDAP realms defined.

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 Server	<code>sAMAccountName=</code>
Novell NDS/eDirectory Server/Other	<code>cn=</code>
Netscape/iPlanet Directory Server	<code>uid=</code>

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.

4. Enter as many Base DN's as you need for the realm. Assume, for example, that Sample_Company has offices in New York and Lisbon, each with its own Base DN.

Section B: LDAP Realm Authentication and Authorization




Figure 9-9: Simplified Directory Information Trees

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:




Figure 9-10: Searching SampleCompany

To search the manufacturing organizations, rather than starting at the top, enter *ou* and *o* values:




Figure 9-11: Searching Part of SampleCompany

You can add, edit, and delete Base DN's for a ProxySG to search. You can also select an individual DN and move it up or down in the list with the **Promote** and **Demote** buttons. The ProxySG searches multiple DN's in the order listed, starting at the top and working down.

5. Click **Apply** to save the changes.

To Define One or More Searchable LDAP Base DN's through the CLI:

1. To define a Base DN, enter the following command:

```
SGOS#(config ldap realm_name) distinguished-name base-dn add base-dn
```

where *base-dn* is a string up to 128 characters long in the format appropriate to the type of LDAP server represented by the realm name. The *base-dn* should be the Fully-Qualified Domain Name (FQDN) of the base of the search.

Repeat this step for each additional Base DN you want added to the list. Entries in the list start with the first Base DN created; subsequent additions are appended to the list. The list is searched from the top down.

2. (Optional) To remove a Base DN:

```
SGOS#(config ldap realm_name) distinguished-name base-dn remove base_dn
```

3. (Optional) To remove all Base DN's and clear the list:

```
SGOS#(config ldap realm_name) distinguished-name base-dn clear
```

4. (Optional) To move a Base DN up or down in the list of Base DN's:

Section B: LDAP Realm Authentication and Authorization

```
SGOS#(config ldap realm_name) distinguished-name base-dn {promote | demote}
base_dn
```

where `promote` moves the specified Base DN up one level in the list and `demote` moves it down one level. You need to issue the command for each level you want to move the Base DN.

LDAP Search & Groups Tab (Authorization and Group Information)

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 ProxySG looks up and queries are derived from the groups specified in policy in `group=` conditions, `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.

To Define LDAP Realm Authorization Properties through the Management Console:

1. Select Configuration>Authentication>LDAP>LDAP Search & Groups.

The LDAP Search & Groups tab displays.

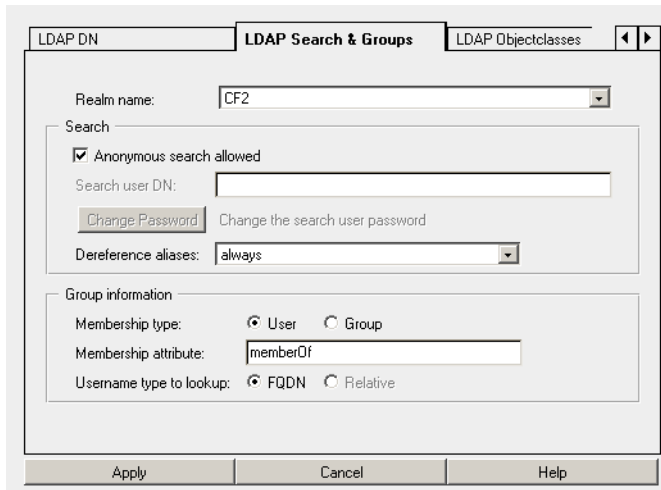


Figure 9-12: LDAP Search & Groups Tab

2. From the Realm Name drop-down list, select the LDAP realm for which you want to specify authorization information.

Note: You must have defined at least one LDAP realm (using the LDAP Realms tab) before attempting to set LDAP Search & Group properties. If the message `Realms must be added in the LDAP Realms tab before editing this tab` is displayed in red at the bottom of this page, you do not currently have any LDAP realms defined.

Section B: LDAP Realm Authentication and Authorization

- Specify whether to allow anonymous search or to enforce user authentication before allowing a search.

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.)

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.

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.

To enforce user authentication before binding to the LDAP service, deselect Anonymous Search Allowed, and set the Search User DN and Search User Password. 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.

You can set or change the user password by clicking Change Password. This password can be up to 64 alphanumeric characters long.

You might want to create a separate user (such as Blue Coat, for example) instead of using an Administrator distinguished name and password.

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.

- Group Information

Membership type and Membership attribute: The ProxySG enters the appropriate default:

- Microsoft Active Directory:
Membership type: `user`
Membership attribute type: `memberOf`
- Netscape/Sun iPlanet:
Membership type: `group`
Membership attribute type: `uniqueMember`
- Novell NDS eDirectory/Other
Membership type: `user`
Membership attribute type: `member`

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.

Section B: LDAP Realm Authentication and Authorization

5. Click Apply.

To Define LDAP Realm Authorization Properties through the CLI:

1. Define the search criteria for the LDAP realm:

```
SGOS#(config ldap realm_name) search {anonymous {disable | enable} |
dereferece {always | finding | never | searching} | password password |
encrypted-password encrypted_password | user-dn user_dn}
```

where:

anonymous	disable enable	<p>If disabled, users will not be permitted to anonymously bind to the LDAP service.</p> <p>If enabled, users will be permitted to anonymously bind to the LDAP service. When anonymous access is allowed, no password is required by the LDAP client to retrieve information, however, one can be specified, if extra security is desirable.</p> <p>The LDAP directory attributes available for an anonymous client are typically a subset of those available to clients that have been authenticated through a user distinguished name and password.</p>
dereferece	always finding never searching	<p>Sets dereference options.</p> <p><i>always</i> dereference aliases is the default.</p> <p><i>finding</i> dereferences aliases only during name resolution.</p> <p><i>searching</i> dereferences aliases only after name resolution.</p> <p><i>never</i> means that aliases will never be dereferenced.</p>
password encrypted- password	password encrypted_ password	<p>Specifies the user password (or encrypted password) associated with the user distinguished name. The non-encrypted (or plain-text) password can be up to 64 alphanumeric characters long.</p> <p>The primary use of the encrypted-password command is to allow the ProxySG to reload a password that it encrypted. If you choose to use a third-party encryption application, be sure it supports RSA encryption, OAEP padding and Base64 encoded with no newlines.</p>
user_dn	user_dn	<p>Specifies a user distinguished name. 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). Search user distinguished name can be up to 512 characters long.</p>

2. To define LDAP realm membership properties:

```
SGOS#(config ldap realm_name) membership-attribute membership_attribute
```

where *membership_attribute* is the name of the attribute that has the group information. (For Active Directory, the attribute name is *memberOf*. For iPlanet, the attribute name is *uniquemember*. For Novell Directory service, the attribute name is *member*.)

```
SGOS#(config ldap realm_name) membership-type {group | user}
```

Section B: LDAP Realm Authentication and Authorization

where `group` specifies that this realm is composed of individual members belonging to a group defined elsewhere, and `user` specifies that this realm is composed of individual disparate members whose only link to each other is membership in this group.

```
SGOS#(config ldap realm_name) membership-username (full | relative)
```

where `full` specifies that the user's FQDN will be used during membership lookups, and `relative` specifies that the user's relative username will be used during membership lookups. Only one can be selected at a time.

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 DN's 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 through the Management Console:

1. Select Configuration>Authentication>LDAP>LDAP Objectclasses.

The LDAP Objectclasses tab displays.

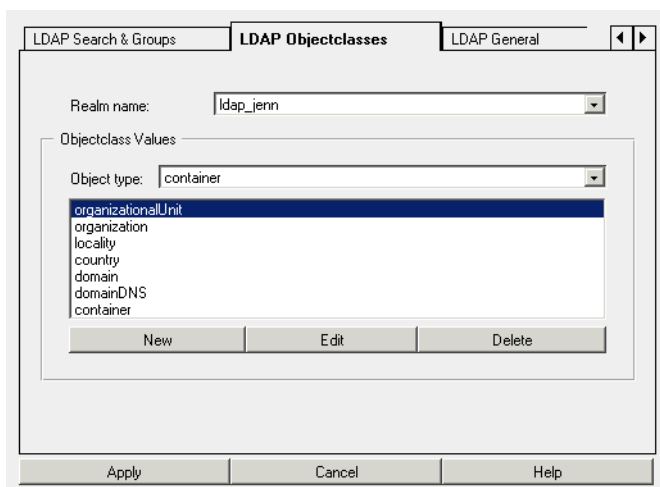


Figure 9-13: LDAP Objectclasses Tab

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.

Section B: LDAP Realm Authentication and Authorization

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.)

The Add/Edit Objectclass Value dialog displays.




Figure 9-14: Add Objectclass Value

5. Enter or edit the objectclass, and click OK; click Apply. For example, objectclass=organization.

Defining LDAP General Realm Properties

The LDAP General page allows you to indicate whether an LDAP server is configured to expect case-sensitive usernames and passwords, the length of time that credentials are cached, the display name, and if you want to use a special virtual host for this realm.

To Configure General LDAP Settings through the Management Console:

1. Select Configuration>Authentication>LDAP>LDAP General.

The LDAP General tab displays.

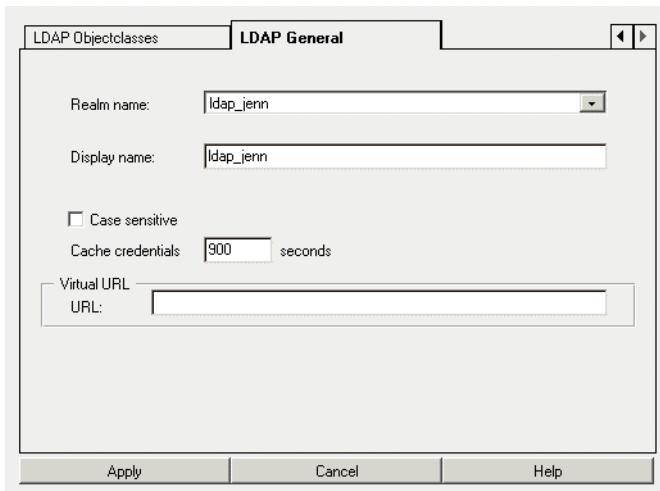


Figure 9-15: LDAP General Tab

Section B: LDAP Realm Authentication and Authorization

- From the Realm Name drop-down list, select the LDAP realm for which you want to change properties.

Note: You must have defined at least one LDAP realm (using the LDAP Realms tab) before attempting to set LDAP general properties. If the message `Realms must be added in the LDAP Realms tab before editing this tab` is displayed in red at the bottom of this page, you do not currently have any LDAP realms defined.

- 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 longer than 128 characters and it cannot be null.
- If the LDAP server is configured to expect case-sensitive usernames and passwords, select Case sensitive.
- Specify the length of time in seconds that user and administrator credentials received from the LDAP server are cached. Credentials can be cached for up to 3932100 seconds. The default value is 900 seconds (15 minutes).

Note: If you specify 0, this increases traffic to the LDAP server because each authentication request generates an authentication and authorization request to the server.

- You can specify a virtual URL based on the individual realm. For information on the virtual URL, see Chapter 8: “Security and Authentication” on page 241.

To Configure General Settings through the CLI:

At the `(config)` prompt, enter the following command to configure general settings:

```
SGOS#(config ldap realm_name) cache-duration seconds
SGOS#(config ldap realm_name) case-sensitive {enable | disable}
SGOS#(config ldap realm_name) virtual-url URL
SGOS#(config ldap realm_name) display-name name
```

where:

<code>cache-duration</code>	<code>seconds</code>	Specifies the length of time in seconds that user and administrator credentials received from the LDAP server are cached. Credentials can be cached for up to 3932100 seconds. The default value is 900 seconds (15 minutes). If you specify 0, cached user and administrator credentials are not re-used.
<code>case-sensitive</code>	<code>enable disable</code>	Enable this setting if the LDAP server is configured to expect case-sensitive usernames and passwords.
<code>virtual-url</code>	<code>URL</code>	The URL to redirect to when the user needs to be challenged for credentials. See Chapter 8: “Security and Authentication” on page 241.
<code>display-name</code>	<code>name</code>	The default value for the display name is the realm name. The display name cannot be longer than 128 characters and cannot be null.

Section B: LDAP Realm Authentication and Authorization

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 *Blue Coat ProxySG 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 4.x systems is *deny*.

- Every LDAP-authenticated user is allowed access 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
```

Section C: RADIUS Realm Authentication and Authorization

Section C: RADIUS Realm Authentication and Authorization

RADIUS is often the protocol of choice for ISPs or enterprises with very large numbers of users. RADIUS was 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. Even if someone were able to retrieve another person's password off the wire, they would not be able to reuse it. The ProxySG's one-time password support works with products such as Secure Computing's SafeWord. It is important to note that the ProxySG does not currently support SafeWord's two-part challenge mechanism.

This section discusses the following topics:

- "Creating a RADIUS Realm"
- "Defining RADIUS Realm Properties"
- "Defining RADIUS Realm General Properties"
- "Creating the CPL"

Creating a RADIUS Realm

To Create a RADIUS Realm through the Management Console:

1. Select Configuration>Authentication>RADIUS>RADIUS Realms.

The RADIUS Realms tab displays.

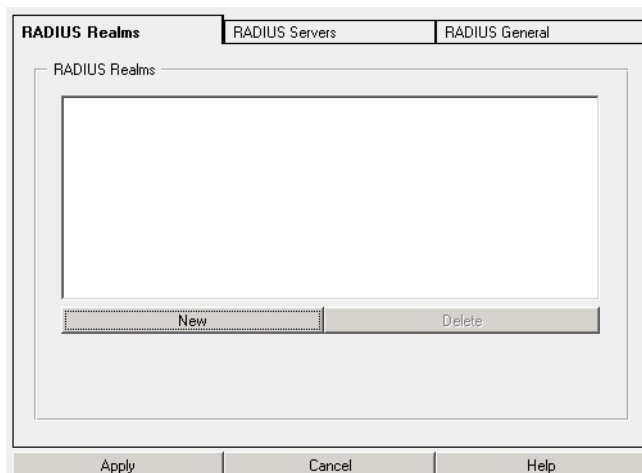


Figure 9-16: RADIUS Realms Tab

2. Click New; the Add RADIUS Realm dialog displays.

Section C: RADIUS Realm Authentication and Authorization

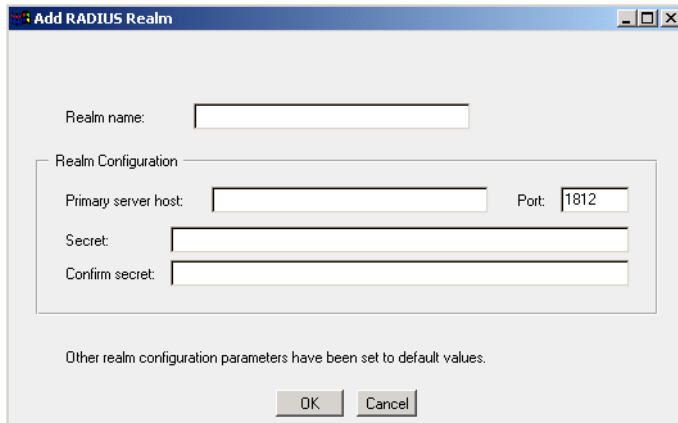


Figure 9-17: Add RADIUS Realm

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. Confirm the secret.
7. Click OK; 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 through the Management Console:

Note: To make these settings through the CLI, see "To Create and Define a RADIUS Realm through the CLI:" on page 299.

1. Select Configuration>Authentication>RADIUS>RADIUS Servers.

The RADIUS Servers tab displays.

Section C: RADIUS Realm Authentication and Authorization

Figure 9-18: RADIUS Servers Tab

Note: You must have defined a RADIUS realm (using the RADIUS Realms tab) before attempting to set RADIUS server properties. If the message `Realms must be added in the RADIUS Realms tab before editing this tab is displayed in red at the bottom of this page`, you do not currently have a RADIUS realm defined.

2. 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.)
3. Specify the Service type, which can be one of the following:
 - Login
 - Framed
 - Callback Login
 - Callback Framed
 - Outbound
 - Administrative
 - NAS Prompt
 - Authenticate Only
 - Callback NAS Prompt
 - Call Check
 - Callback Administrative

Framed is the default. If the user record contains `Check-list ServiceType` attributes, then at least one of the `ServiceType` values must match the service-type of the RADIUS server as configured on the `ProxySG`.

Section C: RADIUS Realm Authentication and Authorization

4. (Optional) Specify the host and port for the alternate 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.)
5. Specify the service type. (See step 3, above, for information on the allowed services types.)
Framed is the default. If the user record contains Check-list ServiceType attributes, then at least one of the ServiceType values must match the service-type of the RADIUS server as configured on the ProxySG.
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 5 seconds. In the Retry field, enter the number of attempts permitted. The default number of retries is 5.
7. If you are using one-time passwords, select the One-time passwords checkbox. (For more information on using one-time passwords, see page 295.)
8. Click Apply.

Defining RADIUS Realm General Properties

The RADIUS General tab allows you to specify the display name and a virtual URL.

To Configure General Settings through the Management Console:

1. Select Configuration>Authentication>RADIUS>RADIUS General.

The RADIUS General tab displays.

The screenshot shows a configuration window with three tabs: 'RADIUS Realms', 'RADIUS Servers', and 'RADIUS General'. The 'RADIUS General' tab is active. It contains the following fields and controls:

- Realm name: radius_1
- Display name: radius_1 (text input field)
- Case sensitive
- Cache credentials: 900 seconds (text input field)
- Virtual URL: URL: (text input field)
- Buttons: Apply, Cancel, Help

Figure 9-19: RADIUS General Tab

Note: You must have defined a RADIUS realm (using the RADIUS Realms tab) before attempting to set RADIUS server properties. If the message `Realms must be added in the RADIUS Realms tab before editing this tab` is displayed in red at the bottom of this page, you do not currently have a RADIUS realm defined.

Section C: RADIUS Realm Authentication and Authorization

2. If needed, change the RADIUS realm display name. 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.
3. If the RADIUS server is configured to expect case-sensitive usernames and passwords, make sure the Case sensitive checkbox is selected.
4. Specify the length of time, in seconds, that user credentials received from the RADIUS server are cached. Credentials can be cached for up to 3932100 seconds. The default is 900 seconds (15 minutes).

Note: If you specify 0, traffic is increased to the RADIUS server because each authentication request generates an authentication and authorization request.

5. (Optional) You can specify a virtual URL based on the individual realm. For more information on the virtual URL, see Chapter 8: “Security and Authentication” on page 241.
6. Click Apply.

To Create and Define a RADIUS Realm through the CLI:

1. At the (config) prompt, enter the following command to create a RADIUS realm:

```
SGOS#(config) security radius create-realm realm_name secret
primary-server_host [primary-server_port]
-or-
SGOS#(config) security radius create-realm-encrypted realm_name
encrypted_secret primary_host [primary_port]
```

where:

<code>realm_name</code>	The name of the RADIUS realm.
<code>secret</code> <code>encrypted_secret</code>	The shared secret (or encrypted secret) associated with the primary RADIUS server. (RADIUS secrets can be up to 64 characters long and are always case sensitive.)
	The primary use of the encrypted-password command is to allow the ProxySG to reload a password that it encrypted. If you choose to use a third-party encryption application, be sure it supports RSA encryption, OAEP padding and Base64 encoded with no new lines.
<code>primary_host</code>	The host for the primary RADIUS server.
<code>primary_port</code>	The port for the primary RADIUS server. The default port is 1812.

2. To set the newly-created RADIUS realm primary and alternate hosts and passwords, enter the following commands:

```
SGOS#(config) security radius edit-realm realm_name
SGOS#(config radius realm_name) primary-server primary_host [primary_port]
SGOS#(config radius realm_name) primary-server service-type type
SGOS#(config radius realm_name) primary-server secret secret
-or-
SGOS#(config radius realm_name) primary-server encrypted-secret
encrypted_secret
```

Section C: RADIUS Realm Authentication and Authorization

and optionally:

```
SGOS#(config radius realm_name) alternate-server alternate_host
[alternate_port]
SGOS#(config radius realm_name) alternate-server secret secret
-or-
SGOS#(config radius realm_name) alternate-server encrypted-secret
encrypted_secret
SGOS#(config radius realm_name) alternate-server service-type type
```

where:

secret | *encrypted_secret* The shared secret (or encrypted secret) associated with the primary or alternate RADIUS server. (RADIUS secrets can be up to 64 characters long and are always case sensitive.)

The primary use of the encrypted-password command is to allow the ProxySG to reload a password that it encrypted. If you choose to use a third-party encryption application, be sure it supports RSA encryption, OAEP padding and Base64 encoded with no newlines.

type *type* stands for the service type, which can be one of the following:

1. Login
2. Framed
3. Callback Login
4. Callback Framed
5. Outbound
6. Administrative
7. NAS Prompt
8. Authenticate Only
9. Callback NAS Prompt
10. Call Check
11. Callback Administrative

If the user record contains Check-list ServiceType attributes, then at least one of the ServiceType values must match the service-type of the RADIUS server as configured on the ProxySG.

primary_server The host for the primary RADIUS server.

primary_port The port for the primary RADIUS server. The default port is 1812.

alternate_host The host for the alternate RADIUS server.

alternate_port The port for the alternate RADIUS server. The default port is 1812.

3. To complete configuration of the RADIUS realm, enter the following commands:

```
SGOS#(config radius realm_name) timeout seconds
SGOS#(config radius realm_name) server-retry count
SGOS#(config radius realm_name) cache-duration seconds
SGOS#(config radius realm_name) case-sensitive enable | disable
SGOS#(config radius realm_name) display-name name
SGOS#(config radius realm_name) spoof-authentication none | origin | proxy
SGOS#(config radius realm_name) one-time-passwords enable | disable
```

Section C: RADIUS Realm Authentication and Authorization

where:

<code>timeout</code>	<code>seconds</code>	The length of time permitted for RADIUS requests to be received before timing out. The default is 5 seconds
<code>server-retry</code>	<code>count</code>	The maximum number of attempts to access the server.
<code>cache-duration</code>	<code>seconds</code>	The length of time that credentials should be cached for this RADIUS realm. The default is 900 seconds (15 minutes)
<code>display-name</code>	<code>name</code>	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.
<code>spooof-authentication</code>	<code>none origin proxy</code>	<p>Enables/disables the forwarding of authenticated credentials to the origin content server or for proxy authentication. You can only choose one.</p> <ul style="list-style-type: none"> • If set to <i>origin</i>, the spoofed header will be an Authorization: header. • If set to <i>proxy</i>, the spoofed header will be a Proxy-Authorization: header. • If set to <i>none</i>, no spoofing will be done. <p>Flush the entries for a realm if the <code>spooof-authentication</code> value is changed to ensure that the <code>spooof-authentication</code> value is immediately applied.</p>
<code>one-time-passwords</code>	<code>enable disable</code>	Allows you to use one-time passwords for authentication. The default is disabled. For more information on one-time passwords, see page 295.

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 *Blue Coat ProxySG 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.

```
<Proxy>
  authenticate (RADIUSRealm)
<Proxy>
  allow hasAttribute.servicetype=yes
  deny
```

Section D: Local Realm Authentication and Authorization

Section D: Local Realm Authentication and Authorization

Using a Local realm is appropriate when the network topography 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 on the ProxySG. 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.

This section discusses the following topics:

- "Creating a Local Realm"
- "Changing Local Realm Properties"
- "Defining the Local User List"
- "Creating the CPL"

Creating a Local Realm

To Create a Local Realm through the Management Console:

1. Select Configuration>Authentication>Local >Local Realms.

The Local Realms tab displays.




Figure 9-20: Local Realms Tab

2. Click New; the Add Local Realm dialog displays.

Section D: Local Realm Authentication and Authorization




Figure 9-21: Add Local Realm

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; click Apply.

To Create a Local Realm through the CLI:

Up to 40 Local realms can be configured per ProxySG.

At the (config) command prompt, enter the following command to create a Local realm:

```
SGOS#(config) security local create-realm realm_name
```

where *realm_name* is the name of the new Local realm.

Changing Local Realm Properties

Once you have created a Local realm, you can modify the properties through the Management Console or the CLI.

To Define or Change Local Realm Properties through the Management Console:

1. Select Configuration>Authentication>Local >Local Main.

The Local Main tab displays.

Section D: Local Realm Authentication and Authorization

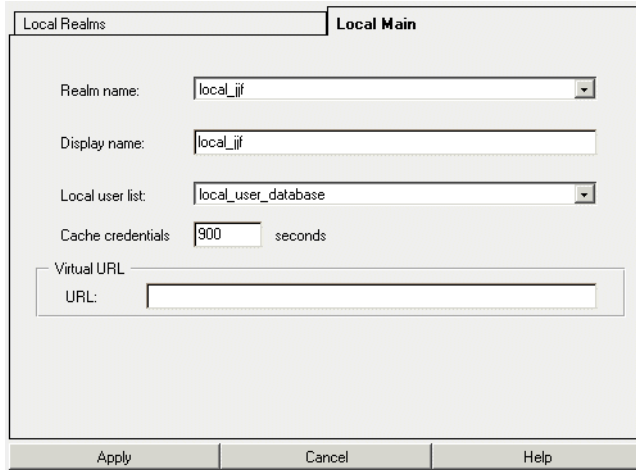


Figure 9-22: Local Main Tab

Note: You must define a Local realm (using the Local Realms tab) before attempting to set realm properties. If the message `Realms must be added in the Local Realms tab` before editing this tab is displayed in red at the bottom of this page, you do not have a Local realm defined.

2. Display name: 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.
3. Local User List: Specify the local user list you want to use from the drop-down list.
4. Specify the length of time, in seconds, that user and administrator credentials received from the Local password file should be cached. Credentials can be cached for up to 3932100 seconds. The default is 900 seconds (15 minutes).
5. You can specify a virtual URL based on the individual realm. For information on using virtual URLs, see Chapter 8: “Security and Authentication” on page 241.
6. Click Apply.

To Define or Change Local Realm Properties through the CLI:

1. From the `(config)` prompt, enter the following commands to modify realm properties:

```
SGOS#(config) security local edit-realm realm_name
SGOS#(config local realm_name) cache-duration 600
SGOS#(config local realm_name) display-name name
SGOS#(config local realm_name) local-user-list list_name
SGOS#(config local realm_name) rename new_name
SGOS#(config local realm_name) spoof-authentication disable | enable
SGOS#(config local realm_name) virtual-url url
```

Section D: Local Realm Authentication and Authorization

where:

<code>cache-duration</code>	<code>seconds</code>	The number of seconds that user and administrator credentials received from the Local password file should be cached. The default is 900 seconds (15 minutes).
<code>display-name</code>	<code>name</code>	The display name for a realm, presented to the user as part of the authentication challenge, is equivalent to the <code>display-name</code> option in the CPL authenticate action. 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.
<code>local-user-list</code>	<code>list_name</code>	The list you want to associate with this realm. The list must exist before it is added. The local user list is set to the default list when the realm is created. For more information on creating a local list, see "Defining the Local User List" on page 305.
<code>rename</code>	<code>new_name</code>	Allows you to change the display name of an existing realm.
<code>spooof-authentication</code>	<code>none origin proxy</code>	Enables/disables the forwarding of authenticated credentials to the origin content server or for proxy authentication. You can only choose one. <ul style="list-style-type: none"> • If set to <i>origin</i>, the spoofed header will be an Authorization: header. • If set to <i>proxy</i>, the spoofed header will be a Proxy-Authorization: header. • If set to <i>none</i>, no spoofing will be done. Flush the entries for a realm if the spoof-authentication value is changed to ensure that the spoof-authentication value is immediately applied.
<code>virtual-url</code>	<code>URL</code>	The URL to redirect to when the user needs to be challenged for credentials. See Chapter 8: "Security and Authentication" on page 241 for more details.

2. (Optional) View the configuration:

```
SGOS#(config local realm_name) view
  Realm name:      local1
  Display name:    local1
  Local user list: list20
  Cache duration:  600
  Virtual URL:     10.9.87.85
```

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.

Section D: Local Realm Authentication and Authorization

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 listname
SGOS#(config local-user-list listname) 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 listname
```

Section D: Local Realm Authentication and Authorization

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 305.

How to Populate 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.

Once you have added 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 ProxySG, see "Populating a Local User List through the ProxySG" on page 308.

Section D: Local Realm Authentication and Authorization

How to Upload 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 listname`. 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`.

Uploading the .htpasswd File:

The `.htpasswd` file is loaded onto the ProxySG with a Perl script found at:

http://download.bluecoat.com/release/tools/set_auth.zip

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 listname
```

To Modify an Existing Local User List (Can be Empty or Contain Users):

1. From the `(config)` prompt, enter:

```
SGOS#(config) security local-user-list edit listname
SGOS#(config local-user-list listname)
```

2. 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.

Section D: Local Realm Authentication and Authorization

```
SGOS#(config local-user-list listname) group create group1
ok
SGOS#(config local-user-list listname) group create group2
ok
SGOS#(config local-user-list listname) group create group3
ok
SGOS#(config local-user-list listname) user create username
```

3. Add the user information to the user account.

```
SGOS#(config local-user-list listname) user edit username
SGOS#(config local-user-list listname username) group add groupname1
SGOS#(config local-user-list listname username) group add groupname2
SGOS#(config local-user-list listname username) password password
-or-
SGOS#(config local-user-list listname 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 ProxySG does not hash it again.

4. (Optional) The user account is enabled by default. To disable a user account:

```
SGOS#(config local-user-list listname username) disable
ok
```

5. Repeat the above steps 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 listname 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 listname username) exit
SGOS#(config local-user-list listname) view
list20
Lockout parameters:
  Max failed attempts: 60
  Lockout duration: 3600
  Reset interval: 7200
Users:
```

Section D: Local Realm Authentication and Authorization

```
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 listname username) group remove group_name
```

To Delete Users from a List:

```
SGOS#(config local-user-list listname) 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 listname) user clear
ok
```

The groups remain but have no users.

Section D: Local Realm Authentication and Authorization

To Delete all Groups from a List:

```
SGOS#(config local-user-list listname) 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 will be disabled (locked). If this is zero, there is no limit. The default is 60 attempts.
- **Lockout duration:** The time after which a locked account will be re-enabled. If this is zero, the account will not automatically re-enable, but will instead stay locked until manually enabled. The default is 3600 seconds (one hour).
- **Reset interval:** The time after which a failed password count will be reset after the last failed password attempt. If this is zero, the failed password count will be reset 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 listname
SGOS#(config local-user-list listname) lockout-duration seconds
SGOS#(config local-user-list listname) max-failed-attempts attempts
SGOS#(config local-user-list listname) reset-interval seconds
```

2. (Optional) View the settings:

```
SGOS#(config local-user-list listname) 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 listname) no [lockout-duration |
max-failed-attempts | reset-interval]
```

Section D: Local Realm Authentication and Authorization

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 *Blue Coat ProxySG Content Policy Language Guide* 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
```

Section E: Certificate Realm Authentication

Section E: Certificate Realm Authentication

Certificate realms are used to authenticate users. If the users are members of an LDAP or Local group, the Certificate Realm can also forward the user credentials to the specified authorization realm, which determines the user's authorization (permissions).

This section discusses the following topics:

- "How Certificate Realm Works"
- "Creating a Certificate Realm"
- "Defining a Certificate Realm"
- "Defining Certificate Realm General Properties"
- "Revoking User Certificates"

How Certificate Realm Works

Once an SSL session has been established, the user is asked to select the certificate to send to the ProxySG. If the certificate was signed by a Certificate Signing Authority that the ProxySG trusts, including itself, then the user is considered authenticated. The username for the user is the one extracted from the certificate during authentication.

At this point the user is authenticated. If an authorization realm has been specified, such as LDAP 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 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, then you do not need to specify an authorization realm. Also, if your policy is such that it works as desired when all certificate realm-authenticated users are not in any group, you do not have to specify an authorization realm.

To use a Certificate Realm, you must:

- Configure SSL between the client and ProxySG (for more information, see "SSL Between the Client and the ProxySG" on page 260)
- Enable verify-client on the HTTPS service that will be used (for more information, see "HTTPS" on page 138).
- Verify that the certificate authority that signed the client's certificates is in the ProxySG *trusted* list.

Creating a Certificate Realm

To Create a Certificate Realm through the Management Console:

1. Select Configuration>Authentication>Certificate>Certificate Realms.

Section E: Certificate Realm Authentication

The Certificate Realms tab displays.

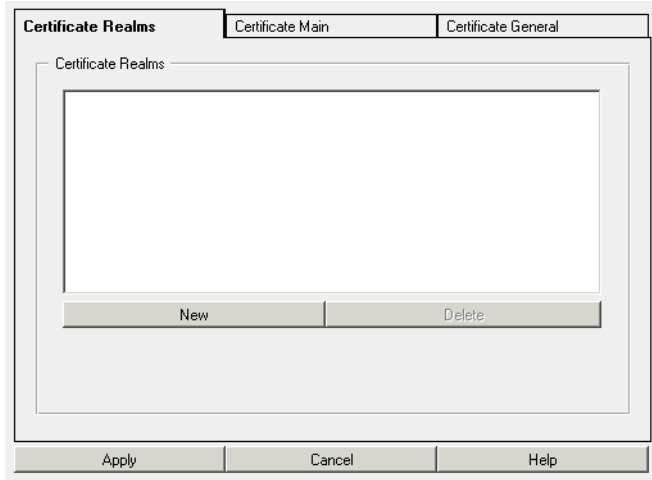


Figure 9-23: Certificate Realms Tab

2. Click New; the Add Certificate Realm dialog displays.




Figure 9-24: Add Certificate Realm

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; click Apply.

To Create a Certificate Realm through the CLI:

Up to 40 Certificate realms can be configured per ProxySG.

At the (config) command prompt, enter the following command to create a Certificate realm:

```
SGOS#(config) security certificate create-realm realm_name
```

where *realm_name* is the name of the new Certificate realm.

Section E: Certificate Realm Authentication

Defining a Certificate Realm

To Define Certificate Authentication Properties through the Management Console:

Note: You can also define certificate authentication properties through the CLI. For information, see "To Create and Define a Certificate Realm through the CLI" on page 317.

1. Select Configuration>Authentication>Certificate>Certificate Main.

The Certificate Main tab displays.

Figure 9-25: Certificate Main Tab

2. From the Realm Name drop-down list, select the Certificate realm for which you want to change realm properties.

Note: You must have defined at least one Certificate realm (using the Certificate Realms tab) before attempting to set Certificate realm properties. If the message `Realms must be added in the Certificate Realms tab before editing this tab` is displayed in red at the bottom of this page, you do not currently have any Certificate realms defined.

3. (Optional) From the Authorization Realm Name drop-down list, select the LDAP or Local realm you want to use to authorize users.
4. From the username attribute field, enter the attribute that specifies the common name in the subject of the certificate. CN is the default.
5. (Optional, if you are configuring a Certificate realm with LDAP authorization) Enter the list of attributes (the container attribute field) that should be used to construct the user's distinguished name.
For example, `$(OU) $(O)` substitutes the OU and O fields from the certificate.
6. (Optional, if you are configuring a Certificate realm with LDAP authorization) Select or deselect Append Base DN.

Section E: Certificate Realm Authentication

7. (Optional, if you are configuring a Certificate realm with LDAP authorization) Enter the Base DN where the search starts. If no BASE DN is specified and Append Base DN is enabled, the first Base DN defined in the LDAP realm used for authorization is appended.
8. Cache credentials: Specify the length of time, in seconds, that user and administrator credentials received from the Local password file should be cached. Credentials can be cached for up to 3932100 seconds. The default is 900 seconds (15 minutes).

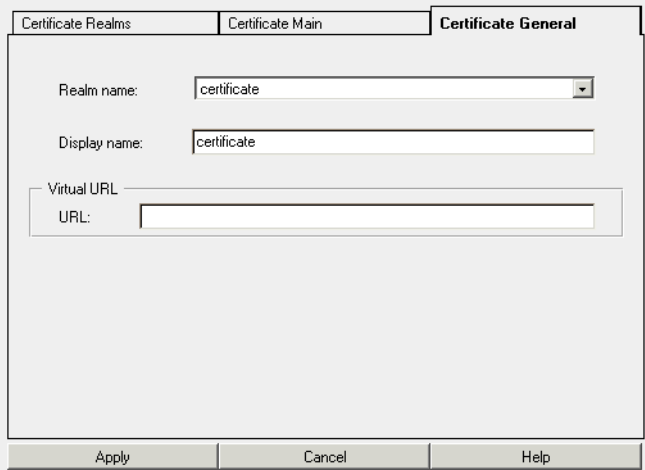
Defining Certificate Realm General Properties

The Certificate General tab allows you to specify the display name and a virtual URL.

To Configure Certificate Realm General Settings through the Management Console:

1. Select Configuration>Authentication>Certificate>Certificate General.

The Certificate General tab displays.



The screenshot shows a web-based configuration interface for a Certificate Realm. It has three tabs: 'Certificate Realms', 'Certificate Main', and 'Certificate General'. The 'Certificate General' tab is active. It contains three input fields: 'Realm name' (a drop-down menu with 'certificate' selected), 'Display name' (a text box with 'certificate'), and 'Virtual URL' (a text box with 'URL:'). At the bottom are 'Apply', 'Cancel', and 'Help' buttons.

Figure 9-26: Certificate General Tab

2. From the Realm name drop-down list, select the Certificate realm for which to change properties.

Note: You must have defined at least one Certificate realm (using the Certificate Realms tab) before attempting to set Certificate general properties. If the message `Realms must be added in the Certificate Realms tab before editing this tab` is displayed in red at the bottom of this page, you do not currently have any Certificate realms defined.

3. If needed, change the Certificate realm display name. 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. For more information on the virtual URL, see Chapter 8: “Security and Authentication” on page 241.
5. Click Apply.

Section E: Certificate Realm Authentication

To Create and Define a Certificate Realm through the CLI

1. At the (config) prompt:

```
SGOS#(config) security certificate create-realm realm_name
```

2. To define an authorization realm for the Certificate realm configuration for the realm you just created, enter the following commands:

```
SGOS#(config) security certificate edit-realm realm_name
SGOS#(config certificate realm_name) authorization {append-base-dn {enable |
disable | dn dn_to_append} | container-attr-list list | realm-name realm |
username-attribute attribute}
```

where:

append-base-dn	enable disable dn dn_to_append	Used only if an LDAP authorization realm is present.
container-attr-list	list	Used only if an LDAP authorization realm is present. If the CLI contains spaces, quotes must be used, as in "ou=Research and Development,ou=Sales,o=Blue Coat".
realm-name	realm_name	The name of the LDAP or Local realm that will be used for authorization. The realm name must already exist.
username-attribute	attribute	The attribute that specifies the common name in the subject of the certificate. CN is the default.

3. Enter the following commands to modify Certificate realm properties:

```
SGOS#(config certificate realm_name) cache-duration 600
SGOS#(config certificate new_realm_name) virtual-url cfauth.com
SGOS#(config certificate new_realm_name) display-name display_name
```

where:

cache-duration	seconds	The number of seconds that user and administrator credentials received from the Credential realm should be cached. The default is 900 seconds (15 minutes).
virtual-url	URL	The URL to redirect to when the user needs to be challenged for credentials. See Chapter 8: "Security and Authentication" on page 241 for more details.
display-name	display_name	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. (Optional) View the results:

```
SGOS#(config certificate certificate-name) view
Realm name:          certificate-name
Display name:        certificate-name
Cache duration:      900
Virtual URL:         cfauth.com
```

Section E: Certificate Realm Authentication

```
Authorization realm:  ldap-realm
Username attribute:  cn
Container attr. list: ou=Sales,ou=Manufacturing
Append DN:          enabled
Base DN:
```

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.

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=2CB06E9F00000000000B
```

Creating the 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 *Blue Coat ProxySG Content Policy Language Guide* for details about CPL and how transactions trigger the evaluation of policy file `<Proxy>` and other layers.

Section E: Certificate Realm Authentication

Be aware that the default policy condition for these examples is *allow*. On new SGOS4.x 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, since 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.

Section E: Certificate Realm Authentication

- Forms authentication modes cannot be used with a Certificate realm. If a form mode is in use and the authentication realm is a Certificate realm, a Policy Substitution realm, or an NTLM realm, you will receive a configuration error.

Section F: Netegrity SiteMinder

Section F: Netegrity SiteMinder

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.

Important: Use of this feature is subject to obtaining the appropriate license. The license check is on the ProxySG.

Access to the SiteMinder policy server is done through the Blue Coat Authentication and Authorization Agent (BCAAA), which must be installed on a Windows 2000 system or higher with access to the SiteMinder policy servers.

Understanding 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 ProxySG.

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 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: The BCAA service is not supported on Solaris in this release. However, Blue Coat can communicate with SiteMinder, regardless of the system it runs on.

Configuration of the ProxySG SiteMinder 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.

Section F: Netegrity SiteMinder

A suitable web agent must be created and configured on the SiteMinder server. This must be configured to support 4.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).

Important: 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.

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`.

Note that if the policy server returns an LDAP FQDN as part of the authentication response, the ProxySG will use 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-signon (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`. Note that this precludes that agent from doing SSO with 4.x agents.

Section F: Netegrity SiteMinder

- 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 Appliance'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 Blue Coat agent(s) that will 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 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 324.

Note: All ProxySG and agent configuration is done on the ProxySG. 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 Proxy SG can use the SiteMinder realm as an ordinary realm.

Section F: Netegrity SiteMinder

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`). Note that, 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 and 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. Once authentication is complete, the request is redirected 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.

Note that 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 through the Management Console:

1. Select Configuration>Authentication>Netegrity SiteMinder>SiteMinder Realms.

The SiteMinder Realms tab displays.

Section F: Netegrity SiteMinder

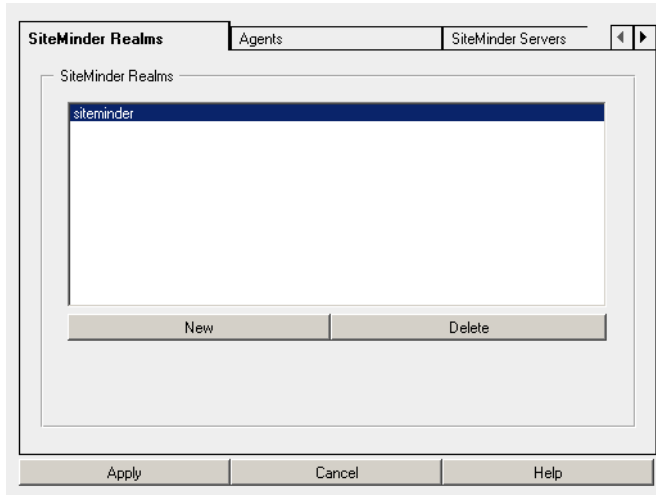


Figure 9-27: SiteMinder Realms Tab

2. Click New; the Add SiteMinder Realm dialog displays.




Figure 9-28: Add SiteMinder Realm

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.

To Create a SiteMinder Realm through the CLI:

At the (config) prompt, enter the following command to create a SiteMinder realm:

```
SGOS#(config) security siteminder create-realm realm_name
      where realm_name is the name of the SiteMinder realm.
```

Agents

You must configure the SiteMinder realm so that it can find the Blue Coat Authentication and Authorization Agent (BCAAA).

1. Select Configuration>Authentication>Netegrity SiteMinder>Agents.

Section F: Netegrity SiteMinder

The Agents page displays.

Figure 9-29: SiteMinder Agents Page

2. Select the realm name to edit from the drop-down list.

Note: You must have defined at least one SiteMinder realm (using the SiteMinder Realms tab) before attempting to configure SiteMinder agents. If the message *Realms must be added in the SiteMinder Realms tab before editing this tab is displayed in red at the bottom of this page*, you do not currently have any SiteMinder realms defined.

3. In the Primary agent section, enter the hostname or IP address where the agent resides.
4. Change the port from the default of 16101 if necessary.
5. Enter the agent name in the Agent name field. The agent name is the name of the agent as configured on the SiteMinder policy server.
6. 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.
7. (Optional) Enter an alternate agent host and agent name in the Alternate agent section.
8. (Optional) Click Enable SSL to enable SSL between the ProxySG and the BC AAA.
9. (Optional) By default, if SSL is enabled, the SiteMinder BC AAA certificate is verified. If you do not want to verify the agent certificate, disable this setting.

To Edit a SiteMinder Agent through the CLI:

1. To define the primary and alternate agent configuration for the realm you just created, enter the following commands at the (config) prompt:

Section F: Netegrity SiteMinder

```

SGOS#(config) security siteminder edit-realm realm_name
SGOS#(config siteminder realm_name) primary-agent agent-name agent_name
SGOS#(config siteminder realm_name) primary-agent host host_name_or_IP
SGOS#(config siteminder realm_name) primary-agent port port_number
SGOS#(config siteminder realm_name) primary-agent encrypted-shared-secret
encrypted_shared_secret
-or-
SGOS#(config siteminder realm_name) primary-agent shared-secret
shared_secret

SGOS#(config siteminder realm_name) alternate-agent agent-name agent_name
SGOS#(config siteminder realm_name) alternate-agent host host_name_or_IP
SGOS#(config siteminder realm_name) alternate-agent port port_number
SGOS#(config siteminder realm_name) alternate-agent encrypted-shared-secret
encrypted_shared_secret
-or-
SGOS#(config siteminder realm_name) alternate-agent shared-secret
shared_secret

```

where:

primary-agent/ alternate agent		These commands allow you to configure either the primary or alternate agent for the SiteMinder realm.
agent-name	<i>agent_name</i>	The name of the agent.
host	<i>host_name</i> or <i>IP</i> <i>address</i>	The host ID or the IP address of the system that contains the agent.
port	<i>port_number</i>	The port where the agent listens.
encrypted-shared-secret /shared-secret	<i>secret</i>	The shared secret (or encrypted secret) associated with the primary or alternate agent. (Secrets can be up to 64 characters long and are always case sensitive.) The primary use of the encrypted-password command is to allow the ProxySG to reload a password that it encrypted. If you choose to use a third-party encryption application, be sure it supports RSA encryption, OAEP padding, and Base64 encoded with no newlines.

2. To enable SSL for this realm and to have the BCAA certificate verified, enter:

```

SGOS#(config siteminder realm_name) ssl enable
SGOS#(config siteminder realm_name) ssl-verify-agent enable

```

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 Configuration>Authentication>Netegrity SiteMinder>SiteMinder Servers.

The SiteMinder Servers page displays.

Section F: Netegrity SiteMinder

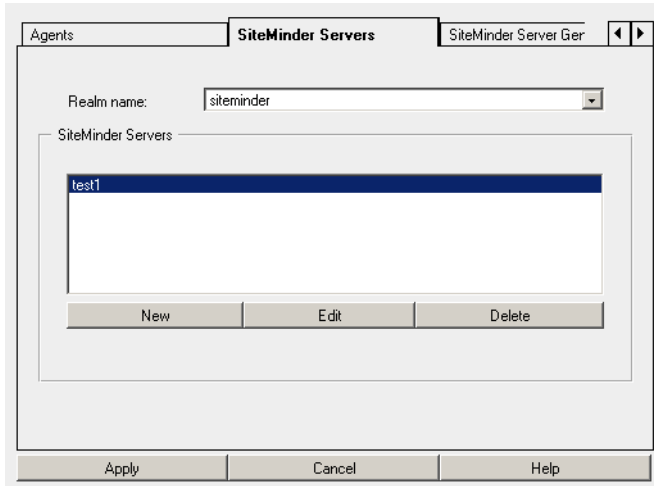


Figure 9-30: SiteMinder Servers Tab

2. From the Realm Name drop-down list, select the SiteMinder realm for which you want to add servers or change server properties.

Note: You must have defined at least one SiteMinder realm (using the SiteMinder Realms page) before attempting to set SiteMinder policy server properties. If the message `Realms must be added in the SiteMinder Realms tab before editing this tab` is displayed in red Click Apply. Repeat the above steps for additional SiteMinder realms, up to a total of 40.

3. To create a new SiteMinder policy server, click New.

The Add List dialog displays.




Figure 9-31: SiteMinder Add List Item Dialog

- a. Enter the name of the server in the dialog. This name is used only to identify the server in the ProxySG Appliance's configuration; it usually is the real hostname of the SiteMinder policy server.
 - b. Click OK.
4. To edit an existing SiteMinder policy server, click Edit.
The Edit dialog displays.

Section F: Netegrity SiteMinder

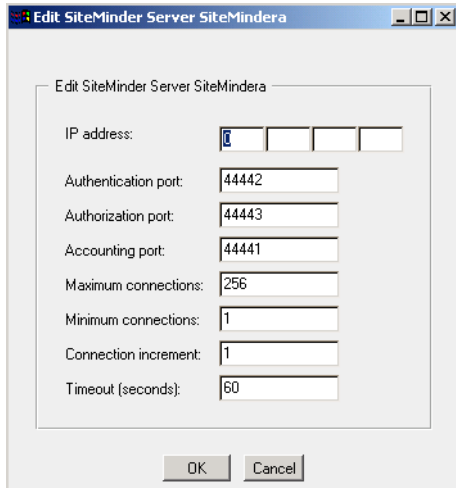


Figure 9-32: SiteMinder Edit Server Dialog

- a. Enter the IP address of the SiteMinder policy server in the IP address field.
 - b. Enter the correct port number for the Authentication, Authorization, and Accounting ports. The ports should be the same as the ports configured on their SiteMinder policy server. The valid port range is 1-65535.
 - c. The maximum number of connections 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 is not exceeded. One is the default.
 - e. The timeout value has a default of 60 seconds, which can be changed.
5. Click OK.
 6. Click Apply.

Editing SiteMinder Policy Servers through the CLI:

To create and edit the SiteMinder policy server for the realm you just created, enter the following commands:

Note: The only required option is the IP address. The other options need only be used if you want to change the defaults.

```
SGOS#(config) security siteminder edit-realm realm_name
SGOS#(config siteminder realm_name) siteminder-server create server_name
SGOS#(config siteminder realm_name) siteminder-server edit server_name
SGOS#(config siteminder realm_name server_name) ip-address ip_address
SGOS#(config siteminder realm_name server_name) authentication-port
port_number
SGOS#(config siteminder realm_name server_name) authorization-port
port_number
```

Section F: Netegrity SiteMinder

```
SGOS#(config siteminder realm_name server_name) accounting-port port_number
SGOS#(config siteminder realm_name server_name) connection-increment number
SGOS#(config siteminder realm_name server_name) max-connections number
SGOS#(config siteminder realm_name server_name) min-connections number
SGOS#(config siteminder realm_name server_name) timeout seconds
```

where:

siteminder-server	create server_name edit server_name delete	You can create a SiteMinder policy server, edit it, or delete it.
edit server_name	ip-address ip_address	The IP address of the SiteMinder policy server.
edit server_name	authentication-port port_number	The default is 44442. The ports should be the same as the ports configured on the SiteMinder policy server. The valid port range is 1-65535.
edit server_name	authorization-port port_number	The default is 44443. The ports should be the same as the ports configured on the SiteMinder policy server. The valid port range is 1-65535.
edit server_name	accounting-port port_number	The default is 44441. The ports should be the same as the ports configured on the SiteMinder policy server. The valid port range is 1-65535.
edit server_name	connection-increment number	The default is 1. The connection increment specifies how many connections to open at a time if more are needed and the maximum is not exceeded.
edit server_name	max-connections number	The default is 256. The maximum number of connections is 32768.
edit server_name	min-connections number	The default is 1.
edit server_name	timeout seconds	The default is 60.

To View the SiteMinder Policy Server Configuration:

```
SGOS#(config siteminder realm_name server_name) view
Server name:      test
IP address:       10.25.36.47
Min connections:  1
Max connections:  256
Connection inc:   1
Timeout:          60
Authentication Port: 44442
Authorization Port: 44443
Accounting Port:  44441
```

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.

Section F: Netegrity SiteMinder

To Configure General Settings through the Management Console:

1. Select Configuration>Authentication>Netegrity SiteMinder>SiteMinder Server General.

The SiteMinder Server General tab displays.

Figure 9-33: SiteMinder Server General Tab

2. From the Realm Name drop-down list, select the SiteMinder realm for which you want to change properties.

Note: You must have defined at least one SiteMinder realm (using the SiteMinder Realms tab) before attempting to set SiteMinder general properties. If the message `Realms must be added in the SiteMinder Realms tab before editing this tab` is displayed in red at the bottom of this page, you do not currently have any SiteMinder realms defined.

3. 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.
4. 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 (BCAAA) interacts with the SiteMinder policy server, not the way that ProxySG interacts with BCAA.

5. To force authentication challenges to always be redirected to an off-box URL, check the Always redirect off-box checkbox.

Note: All SiteMinder Web agents involved must have the setting `EncryptAgentName=no` in their configurations to go off-box for any reason.

Section F: Netegrity SiteMinder

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.

6. If your web applications need information from the SiteMinder policy server responses, you can check the `Add Header Responses` checkbox. When this is checked, responses from the policy server obtained during authentication are added to each request forwarded by the ProxySG. Note that header responses will replace any existing header of the same name; if no such header exists, the header will be added. Cookie responses will replace a cookie header with the same cookie name; if no such cookie header exists, one will be added.
7. To enable validation of the client IP address, select the `Validate client IP address` checkbox. 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, deselect the `Validate client IP address` checkbox for the realm. SiteMinder agents participating in SSO with the ProxySG should also be modified; the `TransientIPCheck` variable should be set to `yes` to enable IP address validation and `no` to disable it.
8. Click `Apply`.

To Configure General Settings through the CLI:

At the `(config)` command prompt, enter the following commands to configure general server settings:

```
SGOS#(config siteminder realm_name) protected-resource-name
protected_resource_name
SGOS#(config siteminder realm_name) server-mode failover | round-robin
(optional)SGOS#(config siteminder realm_name) always-redirect-offbox enable
| disable
(optional)SGOS#(config siteminder realm_name) add-header-responses enable |
disable
(optional)SGOS#(config siteminder realm_name) validate-client-IP disable |
enable
```

where:

<code>protected-resource-name</code>	<code>protected_resource-</code> <code>name</code>	The resource name on the SiteMinder policy server that has rules and policy defined for it.
<code>server-mode</code>	<code>failover </code> <code>round-robin</code>	Behavior of the server. 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.

Section F: Netegrity SiteMinder

`always-redirect-offbox` `enable` | `disable`

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.

All agents involved must have the setting `EncryptAgentName=no` in their configurations to go off-box for any reason.

`add-header-responses` `enable` | `disable`

Enable if your web applications need information from the SiteMinder policy server responses. Note that header responses will replace any existing header of the same name; if no such header exists, the header will be added. Cookie responses will replace a cookie header with the same cookie name; if no such cookie header exists, one will be added.

`validate-client-IP` `enable` | `disable`

Enables validation of the client IP address. If the client IP address in the SSO cookie may be valid yet different from the current request client IP address, due to downstream proxies or other devices, disable client IP validation. The SiteMinder agents participating in SSO with the ProxySG should also be modified. The `TransientIPCheck` variable should be set to `yes` to enable IP validation and `no` to disable it.

SiteMinder General

The SiteMinder General tab allows you to set a display name, cache credentials, timeout value, and create a virtual URL.

To Manage General Settings for the SiteMinder realm

1. Select Authentication>Netegrity SiteMinder>SiteMinder General.

The SiteMinder General tab displays.

Section F: Netegrity SiteMinder

Figure 9-34: SiteMinder General Page

- From the Realm Name drop-down list, select the SiteMinder realm for which you want to change properties.

Note: You must have defined at least one SiteMinder realm (using the SiteMinder Realms tab) before attempting to set SiteMinder general properties. If the message `Realms must be added in the SiteMinder Realms tab before editing this tab` is displayed in red at the bottom of this page, you do not currently have any SiteMinder realms defined.

- If needed, change the SiteMinder realm display name. 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.
- Specify the length of time, in seconds, that user and administrator credentials received from the SiteMinder policy server are cached. Credentials can be cached for up to 3932100 seconds. The default cache-duration is 900 seconds (15 minutes).
- If you want group comparisons for SiteMinder groups to be case sensitive, select Case sensitive.
- 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`.
- Click Apply.

To Set SiteMinder General Settings through the CLI:

At the `(config)` command prompt, enter the following commands to configure general server settings:

```
SGOS#(config siteminder realm_name) cache-duration seconds
SGOS#(config siteminder realm_name) case-sensitive enable | disable
SGOS#(config siteminder realm_name) display-name name
SGOS#(config siteminder realm_name) virtual-url URL
```

Section F: Netegrity SiteMinder

where:

<code>cache-duration</code>	<code>seconds</code>	Specifies the length of time in seconds that user and administrator credentials received from the SiteMinder policy server are cached. Credentials can be cached for up to 3932100 seconds. The default value is 900 seconds (15 minutes).
<code>case-sensitive</code>	<code>enable</code> <code>disable</code>	Specifies whether the SiteMinder policy server is configured to expect case-sensitive usernames and passwords.
<code>display-name</code>	<code>name</code>	Equivalent to the <code>display-name</code> option in the CPL <code>authenticate</code> action. 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.
<code>virtual-url</code>	<code>URL</code>	The URL to redirect to when the user needs to be challenged for credentials. If the ProxySG 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, <code>www.cfauth.com</code> .

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 4.x systems, the default policy condition is *deny*.

Note: Refer to the *Blue Coat ProxySG 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(LDAPRealm)
```

```
<Proxy>
```

```
    group="cn=proxyusers, ou=groups, o=myco"
```

```
    deny
```

Section G: Oblix COREid

Section G: Oblix COREid

The ProxySG can be configured to consult an Oblix COREid (formerly known as Oblix NetPoint) Access Server for authentication and session management decisions. This requires that a COREid realm be configured on the ProxySG and policy written to use that realm for authentication.

The ProxySG supports authentication with Oblix COREid v6.5 and v7.0.

Important: Use of this feature is subject to obtaining the appropriate license. The license check is on the ProxySG.

Access to the COREid Access System is done through the Blue Coat Authentication and Authorization Agent (BCAAA), which must be installed on a Windows 2000 system or higher with access to the COREid Access Servers.

Understanding COREid Interaction with Blue Coat

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 ProxySG.

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 Oblix COREid files needed for configuration is created for that process. 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.

Configuration of the ProxySG COREid realm must be coordinated with configuration of the Access System. Each must be aware of the AccessGate. In addition, certain authorization actions must be configured in the Access System so that BCAA gets the information the ProxySG needs.

Configuring the COREid Access System

Note: Blue Coat assumes you are familiar with the configuration of the COREid Access System and WebGates.

Since 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.

Section G: Oblix COREid

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 thememberOf attribute.

Once the COREid AccessGate, authentication scheme, policy domain, rules, and actions have been defined, the ProxySG can be configured.

Additional COREid Configuration Notes

The ProxySG Appliance's 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 will act on its behalf (hostname or IP address, port, SSL options, and the like).
- Provide BCAA with the information necessary to allow it to identify itself as an AccessGate (AccessGate id, shared secret).
- Provide BCAA with the information that allows it to contact the primary COREid Access Server (IP address, port, connection information).

Section G: Oblix COREid

- Provide BCAA 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 339.

Note: All ProxySG and agent configuration is done on the ProxySG. The ProxySG sends the necessary information to BCAA 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`). Note that, 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.

Section G: Oblix COREid

If the COREid authentication scheme is configured to use a forms-based authentication, the ProxySG will redirect 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 will only redirect 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 through the Management Console:

1. Select Configuration>Authentication>Oblix COREid>COREid Realms.

The COREid Realms tab displays.

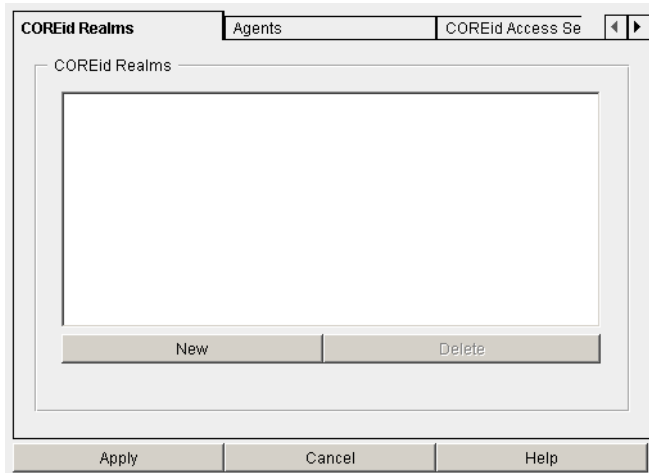


Figure 9-35: Creating a COREid Realm

2. Click New; the Add COREid Realm dialog displays.

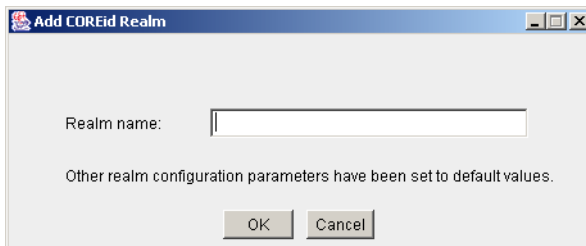


Figure 9-36: Adding the COREid Realm Name

Section G: Oblix COREid

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.
5. Click Apply.

To Create a COREid Realm through the CLI:

At the (config) prompt, enter the following command to create a COREid realm:

```
SGOS#(config) security coreid create-realm realm_name  
where realm_name is the name of the COREid realm.
```

Agents

You must configure the COREid realm so that it can find the Blue Coat Authentication and Authorization Agent (BCAAA).

1. Select Configuration>Authentication>Oblix COREid>Agents.

The Agents page displays.

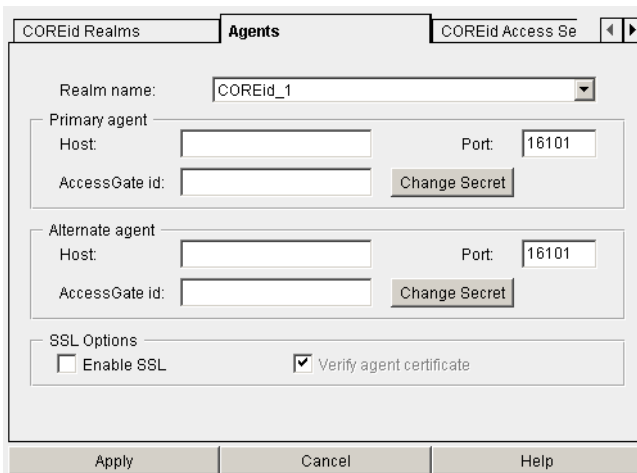


Figure 9-37: Configuring COREid Agents

2. Select the realm name to edit from the drop-down list.

Note: You must have defined at least one COREid realm (using the COREid Realms tab) before attempting to configure COREid agents. If the message Realms must be added in the COREid Realms tab before editing this tab is displayed in red at the bottom of this page, you do not currently have any COREid realms defined.

3. In the Primary agent section, enter the hostname or IP address where the agent resides.
4. Change the port from the default of 16101 if necessary.

Section G: Oblix COREid

5. Enter the AccessGate id in the AccessGate id field. The AccessGate id is the id of the AccessGate as configured in the Access System.
6. 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.
7. (Optional) Enter an alternate agent host and AccessGate id in the Alternate agent section.
8. (Optional) Click Enable SSL to enable SSL between the ProxySG and the BCAA agent.
9. (Optional) By default, if SSL is enabled, the COREid BCAA certificate is verified. If you do not want to verify the agent certificate, disable this setting.

To Edit a COREid Agent through the CLI:

1. To define the primary and alternate agent configuration for the realm you just created, enter the following commands at the (config) prompt:

```
SGOS#(config) security coreid edit-realm realm_name
SGOS#(config coreid realm_name) primary-agent accessgate-id id
SGOS#(config coreid realm_name) primary-agent host host
SGOS#(config coreid realm_name) primary-agent port port
SGOS#(config coreid realm_name) primary-agent encrypted-secret
encrypted_shared_secret
-or-
SGOS#(config coreid realm_name) primary-agent secret shared_secret
SGOS#(config coreid realm_name) alternate-agent accessgate-id id
SGOS#(config coreid realm_name) alternate-agent host host
SGOS#(config coreid realm_name) alternate-agent port port
SGOS#(config coreid realm_name) alternate-agent encrypted-secret
encrypted_shared_secret
-or-
SGOS#(config coreid realm_name) alternate-agent secret shared_secret
```

where

primary-agent/ alternate agent		These commands allow you to configure either the primary or alternate agent for the COREid realm.
accessgate-id	<i>id</i>	The id of the AccessGate.
host	<i>host</i>	The hostname or the IP address of the system that contains the agent.
port	<i>port</i>	The port where the agent listens.
encrypted-secret/ secret	<i>shared_secret</i>	The password (or encrypted password) associated with the primary or alternate AccessGate. (Passwords can be up to 64 characters long and are always case sensitive.) The primary use of the encrypted-secret command is to allow the ProxySG to reload a password that it encrypted. If you choose to use a third-party encryption application, be sure it supports RSA encryption, OAEP padding, and is Base64 encoded with no newlines.

Section G: Oblix COREid

- To enable SSL between the ProxySG and the BCAA agent and to have the BCAA certificate verified, enter:

```
SGOS#(config coreid realm_name) ssl enable
SGOS#(config coreid realm_name) ssl-verify-agent enable
```

COREid Access Server

Once you create a COREid realm, use the COREid Access Server page to specify the primary Access Server information.

- Select Configuration>Authentication>Oblix COREid>COREid Access Server.

The COREid Access Server page displays.

Figure 9-38: Configuring the COREid Access Server

- Select the realm name to edit from the drop-down list.

Note: You must have defined at least one COREid realm (using the COREid Realms tab) before attempting to configure COREid agents. If the message Realms must be added in the COREid Realms tab before editing this tab is displayed in red at the bottom of this page, you do not currently have any COREid realms defined.

- Enter the protected resource name. The protected resource name is the same as the resource name defined in the Access System policy domain.
- Select the Security Transport Mode for the AccessGate to use when communicating with the Access System.
- 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.
- If Cert mode is used, specify the location on the BCAA 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.

Section G: Oblix COREid

7. To force authentication challenges to always be redirected to an off-box URL, check the Always redirect off-box checkbox.
8. To enable validation of the client IP address in SSO cookies, select the Validate client IP address checkbox. 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, then deselect the Validate client IP address checkbox in the realm. The WebGates participating in SSO with the ProxySG should also be modified. The WebGateStatic.lst file should be modified 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, you can check the Add Header Responses checkbox. When this is checked, authorization actions from the policy domain obtained during authentication are added to each request forwarded by the ProxySG. Note that header responses will replace any existing header of the same name; if no such header exists, the header will be added. Cookie responses will replace a cookie header with the same cookie name, if no such cookie header exists, one will be 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

Editing COREid Access Server through the CLI:

To create and edit the COREid Access Server configuration for the realm you just created, enter the following commands:

```
SGOS#(config) security coreid edit-realm realm_name
SGOS#(config coreid realm_name) protected-resource-name resource_name
SGOS#(config coreid realm_name) security-mode cert | open | simple
SGOS#(config coreid realm_name) transport-pass-phrase pass_phrase
-or-
SGOS#(config coreid realm_name) encrypted-transport-pass-phrase
encrypted_pass_phrase
SGOS#(config coreid realm_name) certificate-path certificate_path
SGOS#(config coreid realm_name) always-redirect-offbox disable | enable
SGOS#(config coreid realm_name) validate-client-IP disable | enable
SGOS#(config coreid realm_name) add-header-responses disable | enable
SGOS#(config coreid realm_name) access-server-id id
SGOS#(config coreid realm_name) access-server-hostname hostname
SGOS#(config coreid realm_name) access-server-port port
```

where:

protected-resource-name	protected_resource_name	The resource name defined in the Access System policy domain.
security-mode	cert open simple	The Security Transport Mode for the AccessGate to use when communicating with the Access System

Section G: Oblix COREid

<code>transport-pass-phrase</code>	<code>pass_phrase</code>	If Simple or Cert mode is used, the Transport passphrase (or encrypted passphrase) configured in the Access System.
<code>-or-</code>	<code>-or-</code>	
<code>encrypted-transport-pass-phrase</code>	<code>encrypted_pass_phrase</code>	
<code>certificate-path</code>	<code>certificate_path</code>	If Cert mode is used, the location on the BCAA host machine where the key, server and CA chain certificates reside. The certificate files must be named <code>aaa_key.pem</code> , <code>aaa_cert.pem</code> and <code>aaa_chain.pem</code> respectively.
<code>always-redirect-offbox</code>	<code>disable enable</code>	Forces authentication challenges to always be redirected to an off-box URL.
<code>validate-client-IP</code>	<code>disable enable</code>	Enables validation of the client IP address in SSO cookies. 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, then disable client IP address validation. The WebGates participating in SSO with the ProxySG should also be modified. The <code>WebGateStatic.lst</code> file should be modified to either set the <code>ipvalidation</code> parameter to <code>false</code> or to add the downstream proxy/device to the <code>IPValidationExceptions</code> lists.
<code>add-header-responses</code>	<code>disable enable</code>	When enabled, authorization actions from the policy domain obtained during authentication are added to each request forwarded by the ProxySG. Note that header responses will replace any existing header of the same name; if no such header exists, the header will be added. Cookie responses will replace a cookie header with the same cookie name; if no such cookie header exists, one will be added.
<code>access-server-id</code>	<code>id</code>	The ID of the primary Access Server.
<code>access-server-hostname</code>	<code>hostname</code>	The hostname of the primary Access Server.
<code>access-server-port</code>	<code>port</code>	The port of the primary Access Server.

COREid General

The COREid General tab allows you to set a display name, cache credentials timeout, request timeout value, and case-sensitivity and create a virtual URL.

Section G: Oblix COREid

To Manage General Settings for the COREid Realm through the Management Console:

1. Select Authentication>Oblix COREid>COREid General.

The COREid General tab displays.

Figure 9-39: Configuring COREid General Properties

2. From the Realm Name drop-down list, select the COREid realm for which you want to change properties.

Note: You must have defined at least one COREid realm (using the COREid Realms tab) before attempting to configure COREid agents. If the message Realms must be added in the COREid Realms tab before editing this tab is displayed in red at the bottom of this page, you do not currently have any COREid realms defined.

3. If needed, change the COREid realm display name. 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. Specify the length of time, in seconds, to elapse before timeout if a response from BCAA is not received.
5. Specify the length of time, in seconds, that user and administrator credentials are cached. Credentials can be cached for up to 3932100 seconds. The default cache-duration is 900 seconds (15 minutes).
6. If you want username and group comparisons on the ProxySG to be case sensitive, select Case sensitive.
7. Specify the virtual URL to redirect the user to when they need to be challenged by the ProxySG. If the ProxySG 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.
8. Click Apply.

Section G: Oblix COREid

To Set COREid General Settings through the CLI:

At the (config) command prompt, enter the following commands to configure general settings:

```
SGOS#(config coreid realm_name) display-name name
SGOS#(config coreid realm_name) timeout seconds
SGOS#(config coreid realm_name) cache-duration seconds
SGOS#(config coreid realm_name) case-sensitive disable | enable
SGOS#(config coreid realm_name) virtual-url URL
```

where:

display-name	name	Equivalent to the display-name option in the CPL authenticate action. 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.
timeout	seconds	Specifies the length of time, in seconds, to elapse before timeout if a response from BCAAA is not received.
cache-duration	seconds	Specifies the length of time in seconds that user and administrator credentials received are cached. Credentials can be cached for up to 3932100 seconds. The default value is 900 seconds (15 minutes).
case-sensitive	disable enable	Specifies whether the username and group comparisons on the ProxySG should be case-sensitive.
virtual-url	URL	The URL to redirect to when the user needs to be challenged for credentials. If the ProxySG 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.

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 4.x systems, the default policy condition is *deny*.

Note: Refer to the *Blue Coat ProxySG 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
```

Section H: Policy Substitution Realm

Section H: Policy Substitution Realm

A Policy Substitution realm provides a mechanism for identifying and authorizing users based on information in the request to the ProxySG. 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.

If authorization data (such as group membership) is needed, the realm can be configured 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.

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.

This section discusses the following topics:

- "How Policy Substitution Realms Work"
- "Creating a Policy Substitution Realm"
- "Defining a Policy Substitution Realm"
- "Defining Policy Substitution Realm General Properties"

How Policy Substitution Realms Work

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 two other attributes:

- A user field: A string containing policy substitutions that describes how to construct the simple username.

Section H: Policy Substitution Realm

- 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: 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.

Once 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`

Section H: Policy Substitution Realm

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”.

Creating a Policy Substitution Realm

To Create a Policy Substitution Realm through the Management Console:

1. Select Configuration>Authentication>Policy Substitution>Policy Substitution Realms.

The Policy Substitution Realms tab displays.

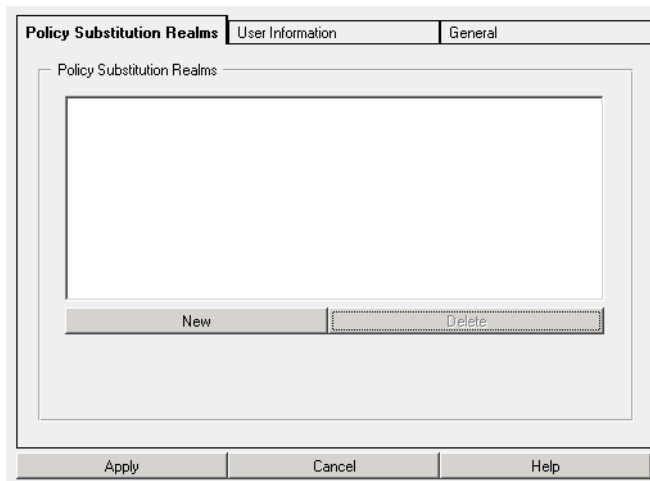


Figure 9-40: Policy Substitution Realms Tab

2. Click New; the Add Policy Substitution Realm dialog displays.




Figure 9-41: Add Policy Substitution Realm Dialog

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; click Apply.

Section H: Policy Substitution Realm

To Create a Policy Substitution Realm through the CLI:

Up to 40 Policy Substitution realms can be configured per ProxySG.

At the `(config)` command prompt, enter the following command to create a Policy Substitution realm:

```
SGOS#(config) security policy-substitution create-realm realm_name
```

where `realm_name` is the name of the new Policy Substitution realm.

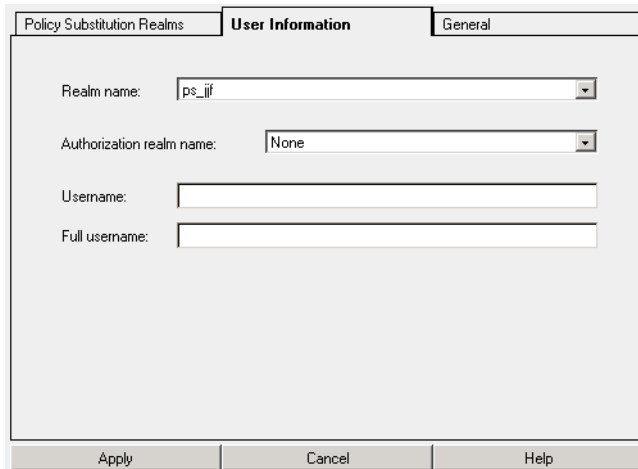
Defining a Policy Substitution Realm

You can define a Policy Substitution realm through either the Management Console or the CLI.

To Define Policy Substitution User Information through the Management Console:

1. Select Configuration>Authentication>Policy Substitution>User Information.

The Policy Substitution User Information tab displays.



Policy Substitution Realms	User Information	General
Realm name: <input type="text" value="ps_lijf"/>		
Authorization realm name: <input type="text" value="None"/>		
Username: <input type="text"/>		
Full username: <input type="text"/>		
<input type="button" value="Apply"/>	<input type="button" value="Cancel"/>	<input type="button" value="Help"/>

Figure 9-42: Policy Substitution User Information Tab

2. From the Realm Name drop-down list, select the Policy Substitution realm for which you want to change realm properties.

Note: 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.

Section H: Policy Substitution Realm

- (Optional) From the Authorization Realm Name drop-down list, select the realm you want to use to authorize users.

Note: Remember that Policy Substitution 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.

- To construct usernames and full usernames, 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 "How Policy Substitution Realms Work" on page 347.
- Click Apply.

To Define Policy Substitution User Information through the CLI:

- From the (config) command prompt, enter the following commands:

```
SGOS#(config) security policy-substitution create-realm realm_name
SGOS#(config) security policy-substitution edit-realm realm_name
(Optional)SGOS#(config policy-substitution realm_name)
authorization-realm-name realm_name
SGOS#(config policy-substitution realm_name) username policy
SGOS#(config policy-substitution realm_name) full-username policy
```

where

<code>edit-realm</code>	<code>realm_name</code>	The name of the realm you want to edit. This command puts you in the edit submenu.
<code>authorization-realm-name</code>	<code>realm_name</code>	This option is only required if you are associating an authorization realm with the Policy Substitution realm.

Section H: Policy Substitution Realm

username	<i>construction</i> <i>_rule</i>	<p>The username as created through policy substitutions. The construction rule is made up any of the substitutions whose values are available at client logon, listed in Appendix D, “CPL Substitutions,” in the <i>Blue Coat ProxySG Content Policy Language Guide</i>.</p> <p>Note: 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 is stored in the user object in the transaction, and becomes the user’s identity.</p> <p>To create usernames for various uses in Policy Substitution realms, see the <i>Blue Coat ProxySG Content Policy Language Guide</i>.</p>
full-username	<i>construction</i> <i>_rule</i>	<p>The username as created through policy substitutions. Note that the full username is only required if you are using an authorization realm. The construction rule is made up any of the policy substitutions whose values are available at client logon listed in Appendix D, “CPL Substitutions,” in the <i>Blue Coat ProxySG Content Policy Language Guide</i>.</p> <p>Note: 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 is stored in the user object in the transaction, and becomes the user’s identity.</p> <p>To create full usernames for the various uses of Policy Substitution realms, see the <i>Blue Coat ProxySG Content Policy Language Guide</i>.</p>

Defining Policy Substitution Realm General Properties

The Policy Substitution General tab allows you to specify the display name and a virtual URL.

Section H: Policy Substitution Realm

To Configure Policy Substitution Realm General Settings through the Management Console:

1. Select Configuration>Authentication>Policy Substitution>General.

The Policy Substitution General tab displays.

Figure 9-43: Policy Substitution General Tab

2. From the Realm name drop-down list, select the Policy Substitution realm for which to change properties.

Note: 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.`

3. Specify the length of time, in seconds, that user and administrator credentials are cached. Credentials can be cached for up to 3932100 seconds. The default cache-duration is 900 seconds (15 minutes).
4. You can specify a virtual URL. For more information on the virtual URL, see Chapter 8: “Security and Authentication” on page 241.
5. Click Apply.

To Configure Policy Substitution Realm General Settings through the CLI:

1. Enter the following commands to modify Policy Substitution realm properties:


```
SGOS#(config) security policy-substitution edit-realm realm_name
SGOS#(config policy-substitution realm_name) cache-duration seconds
SGOS#(config policy-substitution realm_name) virtual-url URL
```

Section H: Policy Substitution Realm

where:

cache-duration	<i>seconds</i>	The number of seconds that user and administrator credentials received from the Policy Substitution realm should be cached. The default is 900 seconds (15 minutes).
virtual-url	<i>URL</i>	The authentication virtual URL for this Policy Substitution realm.

2. (Optional) View the results:

```
SGOS#(config Policy Substitution realm_name) view
Realm name: PS_1
Username:          $(netbios.messenger-user-name)
Full username:     cn=$(netbios.messenger-user-name),cn=users,
dc=$(netbios.computer-domain)
Authorization realm: LDAP_1
Cache duration:    600
Virtual URL:
```

Tips and Boundary Conditions

- Following is an example of how to configure three different types of Policy Substitution realms. For a list of available substitutions, see "Fields Available for Creating Access Log Formats" on page 882.

- 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
```

Section H: Policy Substitution Realm

- 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

```

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.

Note: Refer to the *Blue Coat ProxySG 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 4.x systems, the default policy condition is *deny*.

Note: The Policy Substitution realm cannot be used in an <Admin> layer.

- Every Policy Substitution realm authenticated user is allowed to access the ProxySG.

```

<Proxy>
  authenticate (PolicySubstitutionRealm)

```

Section I: Sequence Realm Authentication

Section I: Sequence Realm Authentication

Once a realm is configured, you can associate it with other realms to allow Blue Coat 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*.

This section discusses the following topics:

- "Adding Realms to a Sequence Realm"
- "Creating a Sequence Realm"

Adding Realms to a Sequence Realm

Keep in mind 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 NTLM, verify that the Allow Basic Credentials checkbox 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 NTLM realm is allowed in a realm sequence.
- If an NTLM realm is in a realm sequence, it must be either the first or last realm in the list.
- If an NTLM realm is in a realm sequence and the NTLM realm does not support Basic credentials, the realm must be the first realm in the sequence and try NTLM authentication once must be enabled.
- Multiple BASIC 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 through the Management Console:

1. Select Configuration>Authentication>Sequences>Sequence Realms.
The Sequence Realms tab displays.

Section I: Sequence Realm Authentication

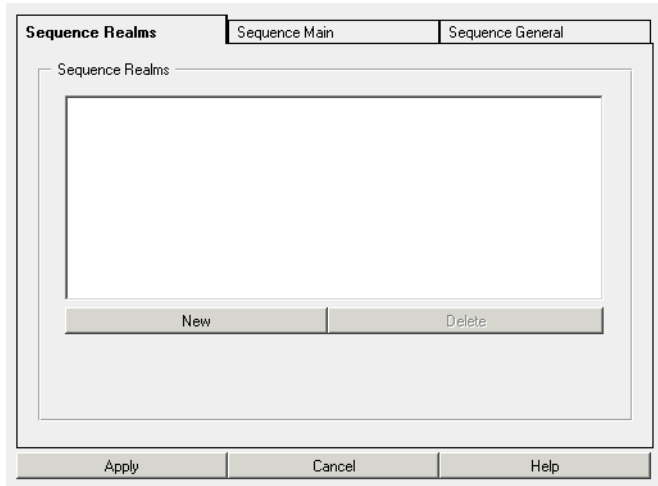


Figure 9-44: Sequence Realms Tab

2. Click New; the Add Sequence Realm dialog displays.




Figure 9-45: Add Sequence Realm

3. In the Realm name, 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.

To Create a Sequence Realm through the CLI:

Up to 40 Sequence realms can be configured per ProxySG.

At the (config) command prompt, enter the following command to create a Sequence realm:

```
SGOS#(config) security sequence create-realm realm_name
      where realm_name is the name of the new Sequence realm.
```

Adding Realms to a Sequence Realm

To Add Realms to a Sequence Realm through the Management Console:

1. Select Configuration>Authentication>Sequences>Sequence Main.

Section I: Sequence Realm Authentication

The Sequences tab displays with the Sequence realm that you want to add realms to.

Note: You must have defined at least one sequence realm (using the Sequence Realms tab) before attempting to set Sequence general properties. If the message Realms must be added in the Sequence Realms tab before editing this tab is displayed in red at the bottom of this page, you do not currently have any Sequence realms defined.

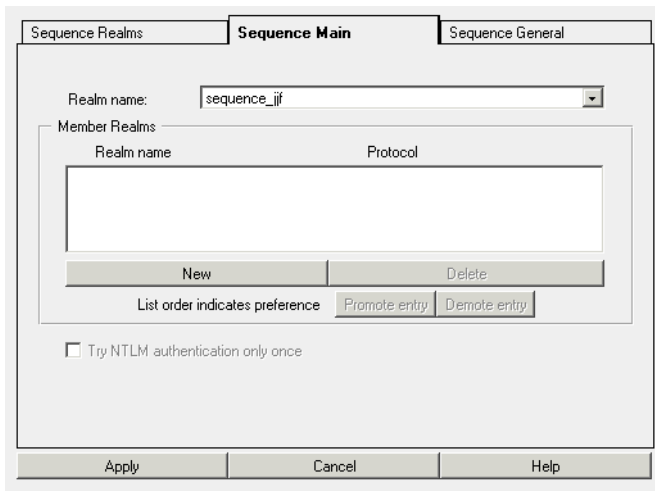


Figure 9-46: Sequence Main Tab

2. Click New to add an existing realm to the realm sequence from the drop-down list. Remember that each realm can be used only once in a realm sequence.



Figure 9-47: Add Member Realm

3. From the drop-down list, select the Sequence realm you wanted added to the realm sequence.
4. Click OK.
You are returned to the main Sequences menu.
5. Click Apply.
6. Repeat from step 2 until you have added all necessary realms.
7. To change the order that the realms are checked, use the promote/demote buttons. Note that when you add an NTLM realm, it is placed first in the list and you can allow the realm sequence to try NTLM authentication only once. If you demote the NTLM entry, it becomes last in the sequence and the default of checking NTLM multiple times is enabled.
8. Click Apply.

Section I: Sequence Realm Authentication

To Add Authentication Realms to a Sequence Realm through the CLI:

1. From the (config) prompt, add existing realms to the new specified sequence realm name:

```
SGOS#(config) security sequence edit-realm realm_sequence_name
SGOS#(config sequence realm_sequence_name) realm add realm_name
```

2. Repeat the `realm add realm_name` command until all necessary realms have been added.

3. (Optional) Give the new sequence realm a display name. 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.

```
SGOS#(config sequence realm_sequence_name) display-name display_name
```

Defining Sequence Realm General Properties

The Sequence General tab allows you to specify the display name and a virtual URL.

1. Select Configuration>Authentication>Sequences>Sequence General.

The Sequence General tab displays.

Figure 9-48: Sequence General Tab

2. From the Realm name drop-down list, select the Sequence realm for which you want to change properties.

Note: You must have defined at least one sequence realm (using the Sequence Realms tab) before attempting to set Sequence general properties. If the message `Realms must be added in the Sequence Realms tab before editing this tab` is displayed in red at the bottom of this page, you do not currently have any Sequence realms defined

3. If needed, change the Sequence realm display name. 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 Chapter 8: “Security and Authentication” on page 241.

Section I: Sequence Realm Authentication

5. Click Apply.

To Manage Authentication Realms in a Sequence Realm through the CLI:

1. When you add an NTLM realm it is placed first in the list, and you have the option of allowing the realm sequence to try NTLM authentication only once. If you demote the NTLM entry, it becomes last in the sequence and the default of checking NTLM multiple times is enabled.

```
SGOS#(config sequence realm_sequence_name) ntlm-only-once-enable
% An NTLM realm must be the first member of the realm sequence before
specifying to try NTLM authentication only once
SGOS#(config sequence realm_sequence_name) realm promote ntlm1
SGOS#(config sequence realm_sequence_name) ntlm-only-once-enable
```

2. (Optional) You can specify a virtual URL based on the individual realm sequence. For information on the virtual URL, see "Security and Authentication" on page 241.

```
SGOS#(config sequence realm_sequence_name) virtual-url 10.25.36.47
ok
```

3. View the configuration.

- a. To view the configuration of the current realm sequence, remain in the submode and enter:

```
SGOS#(config sequence realm_sequence_name) view
Realm name:          seq1
Display name:seq1
Virtual URL:         10.25.36.47
Try NTLM only once:  yes
Member realms:      ntlm1
                   radius1
                   test
                   ldap1
```

- b. To view the configurations of all realm-sequence-names, exit the `edit-realm` submode, and enter:

```
SGOS#(config sequence realm_sequence_name) exit
SGOS#(config) security sequence view
Realm name:          seq1
Display name:seq1
Virtual URL:         10.25.36.47
Try NTLM only once:  yes
Member realms:      ntlm1
                   radius1
                   test
                   ldap1
Realm name:          seq2
Virtual URL:
Try NTLM only once:  no
Member realms:      ldap1
                   ldap2
```

 Section J: Forms-Based Authentication

Section J: Forms-Based Authentication

You can use forms-based authentication exceptions to control what your users see during authentication. 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 ProxySG 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 will be presented with the authentication form again.

Note: You can configure and install the authentication form and several properties through the Management Console and the CLI, but you must use policy to dictate the authentication form's use.

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.

To create and put into use forms-based authentication, you must complete the following steps:

- Create a new form or edit the existing authentication form exception
- Set storage options
- Set CPL policies

Understanding Authentication Forms

You can customize the default authentication form exception or you can use it as a template to 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.)

The default authentication form contains 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 NTLM realm. If it is used, the value is prepended to the username value with a backslash.

Section J: Forms-Based Authentication

- **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.

- **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 will 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_DOMAIN. (optional) If specified, its value will be prepended to the username and separated with a backslash.

The default 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: $(exception.last_error)
<P>
<FORM METHOD="POST" ACTION=$(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=$(x-cs-auth-request-id)>
<P><INPUT TYPE="SUBMIT" VALUE="Submit"> <INPUT TYPE="RESET"></P>
</FORM>
<P>$(exception.contact)
</BODY>
</HTML>
```

Section J: Forms-Based Authentication

If the realm is an NTLM 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.

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	

Note: Any substitutions that are valid in CPL and in other exceptions are valid in authentication form exceptions.

For a discussion of using CPL and a complete list of CPL substitutions, as well as a description of each substitution, refer to the *Blue Coat ProxySG Content Policy Language Guide*.

Tip

There is no realm restriction on the number of authentication form exceptions you can create. You can have as many forms as you want, although you might want to make them as generic as possible to cut down on maintenance.

Creating and Editing an Authentication Form

You can create a new authentication form or you can edit the existing one. If you create a new form, the new form uses the default authentication form as a template. If you edit the default form, all forms created after that contain the modification.

Creating or Editing an Authentication Form through the Management Console

1. Select Configuration>Authentication>Forms.

The Authentication Forms page displays.

Section J: Forms-Based Authentication

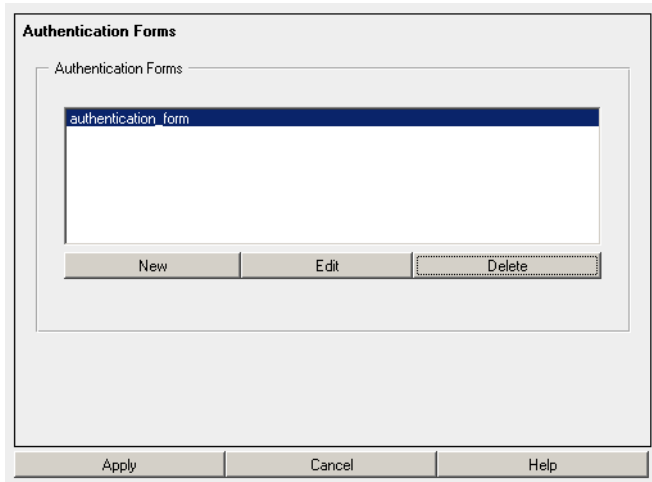


Figure 9-49: Authentication Forms

Click Edit to edit the default authentication form; click New to create a new authentication form based on the existing default settings.

2. If you click New, the Add Authentication Form dialog displays.

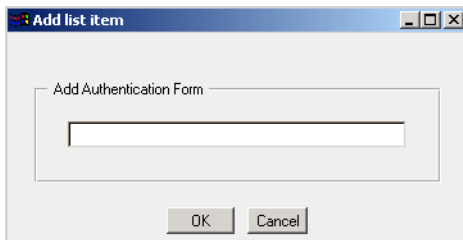


Figure 9-50: Authentication Form Dialog

3. Enter the form name. Click OK.
4. If you highlight the form you want to edit and click Edit, the Edit Authentication *authentication form name* dialog displays.




Figure 9-51: Edit Authentication Form Dialog

5. From the drop-down list, select the method to use to install the authentication form; click Install.

Section J: Forms-Based Authentication

□ **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.

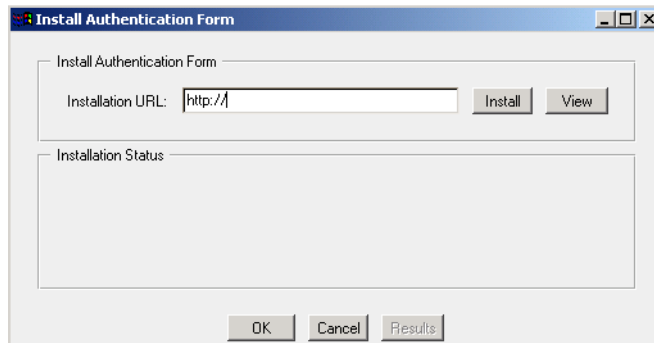
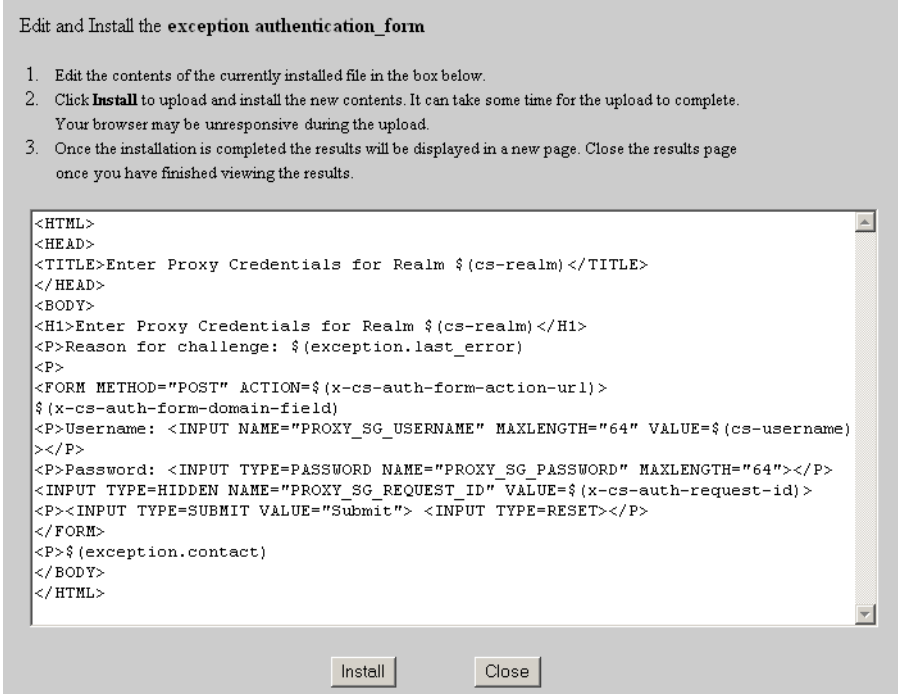


Figure 9-52: Specifying the Remote Location of an Authentication Form

Section J: Forms-Based Authentication

□ 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**.



Edit and Install the **exception authentication_form**

1. Edit the contents of the currently installed file in the box below.
2. Click **Install** to upload and install the new contents. It can take some time for the upload to complete. Your browser may be unresponsive during the upload.
3. Once the installation is completed the results will be displayed in a new page. Close the results page once you have finished viewing the results.

```
<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: ${exception.last_error}
<P>
<FORM METHOD="POST" ACTION=${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=${x-cs-auth-request-id} >
<P><INPUT TYPE="SUBMIT" VALUE="Submit"> <INPUT TYPE="RESET"></P>
</FORM>
<P>${exception.contact}
</BODY>
</HTML>
```

Figure 9-53: Specifying the Local Location of an Authentication Form

Section J: Forms-Based Authentication

□ Text Editor:

The current authentication form is displayed in the text editor. You can edit the form in place. Click **Install** to install the form. When the installation is complete, a results window opens. View the results; to close the window, click **Close**.

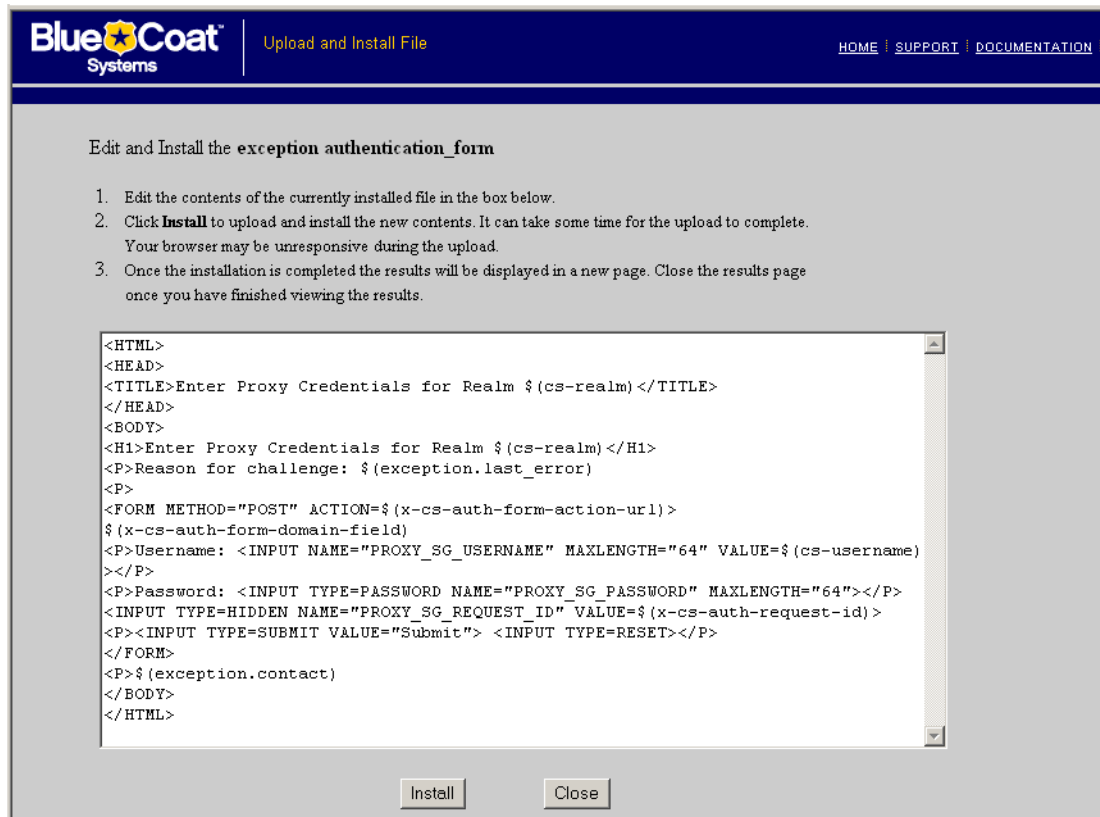


Figure 9-54: Using the ProxySG Text Editor

To Create an Authentication Form through the CLI

Remember that if you create a new form, the new form uses the default authentication form as a template. If you edit the default form, all forms created after that will also have the modification.

To create a new form, enter the following command from the (config) prompt:

```
SGOS#(config) security authentication-form create form_name
ok
```

To view the authentication forms on the system, enter the following command:

```
SGOS#(config) show security authentication-form
Authentication forms:
authentication_form
authentication_test
```

Section J: Forms-Based Authentication

To Edit an Authentication Form through the CLI

You cannot edit an authentication form in place through the CLI. You can replace a form though using either remote download or through the ProxySG Appliance's inline commands.

To Edit an Authentication Form using Inline Commands:

```
SGOS#(config) security authentication-form inline form_name
end-of-file_marker
-or-
SGOS# inline authentication-form form_name end-of-file_marker
```

Remember that any form you modify must contain the username, password and request ID. A form that is missing these values will result in the user receiving an error page when the authentication form is submitted. For more information on required fields in a new authentication form, see "Understanding Authentication Forms" on page 361.

Note: You can also import the entire set of forms through the `inline authentication-forms` command.

Notes on using inline commands:

- If you make a mistake on the current line of the script you are typing, you can backspace to correct the problem.
- If you notice a mistake on a previous line, you must quit the script (by using `Ctrl<c>`) and start over.
- The inline script overwrites the existing template.

To Create and Download an Authentication Form using a Text Editor:

1. Create the authentication form as a text file.
2. Place the form on a server that is accessible to the ProxySG.
3. Enter the following commands to give the ProxySG the file's location and to download the file:

```
SGOS#(config) security authentication-form path [form_name] path
SGOS#(config) security authentication-form load form_name
-or-
SGOS#load authentication-form form_name
```

Note: You can also download the entire set of forms through the `security authentication-form path` and `load authentication-forms` commands.

To Delete an Authentication Form

From the `(config)` prompt, enter the following commands:

```
SGOS#(config) security authentication-form delete form_name
```

Section J: Forms-Based Authentication

Setting Storage Options

When a request requiring the user to be challenged with a form contains a body, the request is stored on the ProxySG 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.

To Set Storage Options through the Management Console

1. Select Configuration>Authentication>Request Storage.

The Request Storage tab displays.

Figure 9-55: Request Storage Tab

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 don't want the ProxySG to Verify the IP address against the original request, doubleclick the checkbox to deselect the option. The default is to verify the IP address.

Section J: Forms-Based Authentication

5. If you want to Allow redirects from the origin servers, click the checkbox. The default is to not allow redirects from origin servers.

Note: During authentication, the user's POST is redirected to a GET request. The client will therefore automatically follow redirects from the origin server. Since the ProxySG is converting the GET to a POST and adding the post data to the request before contacting the origin server, the administrator needs to explicitly specify that redirects to these POSTs requests can be automatically followed.

6. Click Apply.

To Set Storage Options through the CLI

From the (config) prompt, enter the following commands to select storage options:

```
SGOS# (config) security request-storage max-size megabytes
SGOS# (config) security request-storage expiry-time seconds
SGOS# (config) security request-storage verify-ip enable | disable
SGOS# (config) security request-storage allow-redirects enable | disable
```

where

max-size	<i>megabytes</i>	Sets the maximum POST request size during authentication. The default is 50 megabytes.
expiry-time	<i>seconds</i>	Sets the amount of time before the stored request expires. The default is 300 seconds (five minutes)
verify-ip	enable disable	Enables or disables the verify-ip option. The default is to enable the ProxySG to verify the IP address against the original request.
allow-redirect	enable disable	Specifies whether to allow redirects. The default is disable.

Using CPL with Forms-Based Authentication

To use forms-based authentication, you must create policies that enable it and also control which form will be 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 the CPL condition `authenticate.form(form_name)`.

Note: The `authenticate.form(form.name)` condition can be used with the form authentication modes only. If no form is specified the form defaults to `authentication_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:

Section J: 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 `authenticate.mode()` property will be defined as a form mode (above) in policy, you don't 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 the *Blue Coat ProxySG Content Policy Language Guide*.

Tips and Boundary Conditions

- 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 should result 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 will be 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, it is recommended 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 HTTP browsers only.
- Since forms only support BASIC authentication, authentication-form exceptions cannot be used with a Certificate realm, a Policy Substitution realm, or with an NTLM realm that allows only NTLM credentials. If a form is in use and the authentication realm is a NTLM credentials, a Policy Substitution realm, or a Certificate realm, the user receives a configuration error.
- User credentials are sent in plain text. However, they can be sent securely using SSL if the virtual URL is HTTPS.

Section J: Forms-Based Authentication

- Since not all user requests support forms (such as WebDAV and streaming), you should create policy to bypass authentication or use a different authentication mode with the same realm for those requests.

Section K: Managing the Credential Cache

When you have configured all your realms, you can view your realms and manage the credentials cache for a specific realm.

To Manage the Credential Cache through the Management Console:

1. Select Configuration>Authentication>Realms.

The Realms page displays, with all realms that you have created.

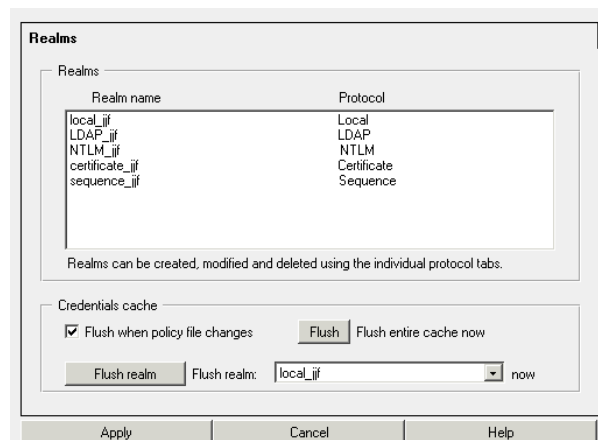


Figure 9-56: Viewing All Realms on the ProxySG

2. To manage the credential cache:
 - ❑ To purge the credentials cache when you make policy changes, select Flush When Policy File Changes (this option is selected by default).
 - ❑ To flush the entire credentials cache immediately, click Flush and confirm.
 - ❑ To flush only the entries for a particular realm in the credentials cache, select the realm from the drop-down list, click Flush Realm confirm.

All of these actions force users to be re-authenticated.

3. Click Apply.

To Manage the Credential Cache through the CLI:

From the (config) command prompt, enter the following command:

```
SGOS#(config) security flush-credentials <cr> [on-policy-change {enable | disable} | realm realm]
```

where:

<cr>

Press Enter to flush the credential cache now.

on-policy-change enable | disable

Flush the cache only if the policy changes.

realm realm

Flush the credential cache for the specified realm.

Section K: Managing the Credential Cache

Boundary Conditions

- For all realms except NTLM, SiteMinder, and COREid, the maximum number of entries stored in the credential cache is 80,000.

For NTLM, SiteMinder, and COREid authentication, the maximum number of entries stored in the credential cache is dependent on the system. You can have at least 2500 entries but potentially more depending on the system resources.

- XFTP users are not prompted for proxy authentication if the credentials are in the cache and the credentials have not expired.

Chapter 10: *Bandwidth Management*

Bandwidth management (BWM) allows you to classify, control, and, if required, limit the amount of bandwidth used by different classes of network traffic flowing into or out of the ProxySG. Network resource sharing (or link sharing) is done using a bandwidth-management hierarchy where multiple traffic classes share available bandwidth in a controlled manner.

Note: The ProxySG does not try 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 do 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.

Licensing

Bandwidth-management is enabled by default if you have a valid license for this feature. For information about obtaining a bandwidth-management license, refer to Chapter 2: “Licensing” on page 35.

Bandwidth Management Terms

Some of the terms used in this document are described below.

- **Bandwidth Class:** A defined unit of bandwidth allocation. An administrator uses bandwidth classes to allocate bandwidth to a particular type of traffic flowing through the ProxySG.
- **Bandwidth Class Hierarchy:** Bandwidth classes can be grouped together in a class hierarchy, which is a tree structure that specifies the relationship among different classes. You create a hierarchy by creating at least one parent class and assigning other classes to be its children.
- **Bandwidth Policy:** The set of rules that you define in the policy layer to identify and classify the traffic in the ProxySG, using the bandwidth classes that you create. You must use policy (through either VPM or CPL) in order to manage bandwidth.
- **Child Class:** The child of a parent class is dependent upon that parent class for available bandwidth (they share the bandwidth in proportion to their minimum/maximum bandwidth values and priority levels). A child class with siblings (classes with the same parent class) shares bandwidth with those siblings in the same manner.
- **Inbound Traffic:** Network packets flowing into the ProxySG. Inbound traffic mainly consists of the following:

- ❑ Server inbound: packets originating at the origin content server (OCS) and sent to the ProxySG to load a web object.
- ❑ Client inbound: packets originating at the client and sent to the ProxySG for Web requests.
- OCS: Origin content server.
- Outbound Traffic: Network packets flowing out of the ProxySG. Outbound traffic mainly consists of the following:
 - ❑ Client outbound: packets sent to the client in response to a Web request.
 - ❑ Server outbound: packets sent to an OCS or upstream proxy to request a service.
- Parent Class: A class with at least one child. The parent class must share its bandwidth with its child classes in proportion to the minimum/maximum bandwidth values or priority levels.
- Sibling Class: A bandwidth class with the same parent class as another class.
- Traffic Flow: Also referred to as *flow*. A set of packets belonging to the same TCP/UDP connection that terminate at, originate at, or flow through the ProxySG. A single request from a client involves two separate connections. One of them is from the client to the ProxySG, and the other is from the ProxySG to the OCS. Within each of these connections, traffic flows in two directions—in one direction, packets flow out of the ProxySG (outbound traffic), and in the other direction, packets flow into the ProxySG (inbound traffic). Connections can come from the client or the server. Thus, traffic can be classified into one of four types:
 - ❑ Server inbound
 - ❑ Server outbound
 - ❑ Client inbound
 - ❑ Client outbound

These four traffic flows represent each of the four combinations described above. Each flow represents a single direction from a single connection.

Bandwidth Management Overview

To manage the bandwidth of different types of traffic that flow into, out of, or through the ProxySG, you must do 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.

- 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). For more information about using VPM to create policy rules, see Chapter 14: “The Visual Policy Manager” on page 453. For information about composing CPL, refer to the *Blue Coat ProxySG Content Policy Language Guide*.

Bandwidth Allocation

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 388.

Bandwidth Classes

To define a bandwidth class, you will 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 configuration settings available 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 379).

Minimum Bandwidth

Setting a minimum for a bandwidth class guarantees that class will receive 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 will not be able to guarantee the minimums defined for each class.

Note: The ProxySG does not try 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 will be available to that class. It does not matter how much bandwidth is available; a class can never receive more bandwidth than its maximum.

To keep 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 gets 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 slow down traffic on purpose. Unused bandwidth can go to waste with the maximum-bandwidth setting, while the minimum-bandwidth settings and priority levels will always distribute 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.

In order 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 will flow 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 will always be counted as belonging to a single class. If multiple policy rules match a single flow and try to classify it into multiple bandwidth classes, the last classification done by policy will apply.
- When a flow is classified as belonging to a bandwidth class, all packets belonging to that flow will be 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 in order 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.

Keep in mind that, 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 will have to 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 will be denied—no class with a specified maximum is ever allowed more than that amount.
- If a class does not have a minimum specified, it will have to 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 ProxySG).
2. Server outbound—traffic flowing out of the ProxySG to a server.
3. Server inbound—traffic flowing back into the ProxySG from a server (the entity responding to the request).
4. Client outbound—traffic flowing back out of the ProxySG to a client.

Figure 10-1 shows the traffic flows between a client and server via the ProxySG.




Figure 10-1: Network Configuration Showing Traffic Flow Directions

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 ProxySG then forwards this GET to a server (server outbound). The server responds to the ProxySG with the appropriate content (server inbound), and then the ProxySG 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 387 for information about classifying traffic flows with policy.

Configuring Bandwidth Allocation

You can use either the Management Console or the CLI to do the following tasks:

- Enable or disable bandwidth management.
- Create and configure bandwidth classes.
- Delete bandwidth classes.
- View bandwidth management class configurations.

Note: If you are planning to manage the bandwidth of streaming media protocols (Windows Media, Real Media, or QuickTime), you might want to use the streaming features instead of the bandwidth management features described in this section. For most circumstances, Blue Coat recommends that you use the streaming features to control streaming bandwidth rather than the bandwidth management features. For information about the differences between these two methods, see "Choosing a Method to Limit Streaming Bandwidth" on page 594.

Enabling or Disabling Bandwidth Management

The following procedures explain how to enable or disable bandwidth management through the Management Console or the CLI.

Note: Bandwidth management is enabled by default if you have a valid license for this feature. For information about getting a license for bandwidth management, see Chapter 2: "Licensing" on page 35.

To Enable or Disable Bandwidth Management through the Management Console

1. Select Configuration>Bandwidth Management>BWM Classes>Bandwidth Classes.

The Bandwidth Classes tab displays.



Figure 10-2: Bandwidth Classes Tab

2. To enable or disable bandwidth management, select or deselect the Enable Bandwidth Management checkbox.
3. Click Apply.

To Enable or Disable Bandwidth Management through the CLI

At the `(config)` command prompt, enter the following commands to enable or disable bandwidth management:

```
SGOS#(config) bandwidth-management  
SGOS#(config bandwidth-management) enable | disable
```

Creating and Editing Bandwidth Classes

The following procedures detail how to create and edit a bandwidth management class.

To Create a BWM Class through the Management Console

1. Select Configuration>Bandwidth Management>BWM Classes>Bandwidth Classes.

The Bandwidth Classes tab displays.

2. To create a new BWM class, click New.

The Create Bandwidth Class dialog displays.

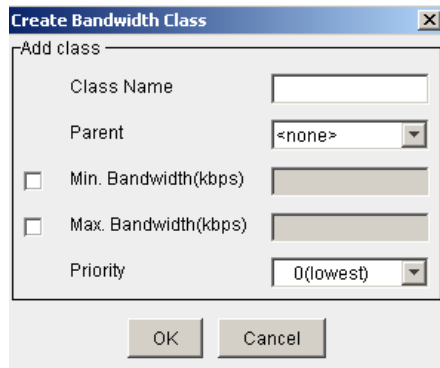


Figure 10-3: Create Bandwidth Class Dialog

3. Fill in the fields as appropriate:
 - Class name: Assign a meaningful name for this class. The name can be up to 64 characters long; spaces are not allowed.
 - Parent: If you want the class you are creating to be the child of another class in the bandwidth class hierarchy, choose a class from the Parent drop-down list. This class must already exist.
 - Min. Bandwidth: To set a minimum bandwidth for this class in kilobits per second (kbps), select the Min. Bandwidth checkbox and enter a minimum bandwidth value in the field. The default minimum bandwidth setting is *Unspecified*, meaning the class is not guaranteed a minimum amount of bandwidth.
 - Max. Bandwidth: To set a maximum bandwidth for this class in kbps, select the Max. Bandwidth checkbox 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.
 - 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.
4. Click OK.
5. Click Apply.

To Create a BWM Class through the CLI

1. At the (config) command prompt, enter the following commands to create a new BWM class:


```
SGOS#(config) bandwidth-management
SGOS#(config bandwidth-management) create bwm_class
```

 where *bwm_class* will be the name of the new BWM class.
2. Configure the newly created bandwidth class (see "To Edit a BWM Class through the CLI" on page 384 for instructions).

To Edit a BWM Class through the Management Console

1. Select Configuration>Bandwidth Management>BWM Classes>Bandwidth Classes.
The Bandwidth Classes tab displays.

2. Highlight the class that you want to edit and click Edit.

The Edit Class dialog displays.

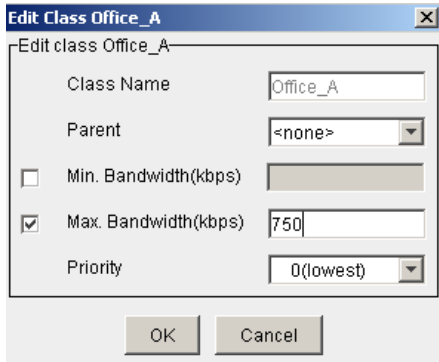


Figure 10-4: Edit Class Dialog

3. Fill in the fields as appropriate:

- Class Name: this field cannot be edited. To change the name of a class, you must delete the class and create a new one with the new name.
- Parent: If you want the class you are editing to be the child of another class in the bandwidth class hierarchy, choose a class from the Parent drop-down list.
- Min. Bandwidth: To set a minimum bandwidth for this class in kilobits per second (kbps), select the Min. Bandwidth checkbox and enter a minimum bandwidth value in the field. The default minimum bandwidth setting is *Unspecified*, meaning the class is not guaranteed a minimum amount of bandwidth.
- Max. Bandwidth: To set a maximum bandwidth for this class in kbps, select the Max. Bandwidth checkbox 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.
- 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.

4. Click OK.
5. Click Apply.

To Edit a BWM Class through the CLI

1. To set the priority level and minimum/maximum bandwidth values for an existing BWM class, enter the following commands:

```
SGOS#(config) bandwidth-management
SGOS#(config bandwidth-management) edit bwm_class
```

This changes the prompt and puts you into the Bandwidth-Class submode.

```
SGOS#(config bw-class bwm_class) min-bandwidth minimum_in_kbps
SGOS#(config bw-class bwm_class) max-bandwidth maximum_in_kbps
SGOS#(config bw-class bwm_class) priority value_from_0_to_7
```

where:

<code>min-bandwidth</code>	<code>minimum_in_kbps</code>	Sets the minimum bandwidth for this class in kilobits per second. The default for this setting is unspecified, meaning that the class is not guaranteed a minimum amount of bandwidth.
<code>max-bandwidth</code>	<code>maximum_in_kbps</code>	Sets the maximum bandwidth for this class in kilobits per second. The default for this setting is unlimited (no maximum).
<code>priority</code>	<code>value_from_0_to_7</code>	Sets the priority level for this class—0 is the lowest priority level and 7 is the highest. The default priority is 0.

2. (Optional) To reset the values to the defaults, enter the following commands:

```
SGOS#(config bandwidth-management bwm_class) no {min-bandwidth |
max-bandwidth}
```

where:

<code>no min-bandwidth</code>	Sets the default minimum to the default, unspecified (no minimum bandwidth guarantee).
<code>no max-bandwidth</code>	Sets the maximum-bandwidth setting to the default, unlimited (no maximum).

3. To make this class a child of another class or to clear the parent class from this class, enter one of the following commands:

```
SGOS#(config bandwidth-management bwm_class) parent parent_class_name
-or-
SGOS#(config bandwidth-management bwm_class) no parent
```

4. To view the configuration for this class, enter the following command:

```
SGOS#(config bandwidth-management bwm_class) view
```

For example:

```
SGOS#(config bandwidth-management Office_A) view
Class Name:      Office_A
Parent:          <none>
Minimum Bandwidth:  unspecified
Maximum Bandwidth: 750 kbps
Priority:         0
```

5. To view the configuration of any child classes of this class, enter the following command:

```
SGOS#(config bandwidth-management bwm_class) view children
```

Deleting a Bandwidth Management Class

The following procedures explain how to delete a bandwidth management class through the Management Console or the CLI.

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 will display explaining why this class cannot be deleted.

To Delete a BWM Class through the Management Console

1. Select Configuration>Bandwidth Management>BWM Classes>Bandwidth Classes.
The Bandwidth Classes tab displays.
2. Highlight the class you want to delete and click the Delete button.
The Remove Object dialog displays.
3. Click Yes to delete the class.
4. Click Apply.

To Delete a BWM Class through the CLI

At the `(config)` command prompt, enter the following command to delete the specified BWM class:

```
SGOS#(config) bandwidth-management  
SGOS#(config bandwidth-management) delete bwm_class
```

Viewing Bandwidth Management Configurations and Statistics

You can view bandwidth management configurations to see what the settings are for each class, and you can view bandwidth management statistics to see the current and total bandwidth, packet rate, and number of drops (the total number of packets dropped).

Bandwidth management configurations (minimum/maximum bandwidth, priority level, and hierarchy relationships) are visible in the Management Console. The `view` commands allow you to view the same information in the CLI. See "Bandwidth Management Statistics" on page 849 for information about viewing bandwidth management statistics.

Viewing Bandwidth Management Configurations

You can view the following bandwidth class configurations through the Management Console or CLI:

- Level in the hierarchy (parent/child relationships)
- Priority level
- Maximum bandwidth value
- Minimum bandwidth value

To View BWM Configuration through the Management Console

1. Select Configuration>Bandwidth Management>BWM Classes>Bandwidth Classes.

The Bandwidth Classes tab displays. On this tab, you can view a class's minimum, maximum and priority value. Top level classes are visible—classes with children have a folder icon on the left.

2. To view the configurations of the child class(es) of a class, double-click the folder icon.

The child classes become visible. A second double-click closes the folder.

To View BWM Configuration through the CLI

1. To view all BWM configuration information, enter the following commands at the (config) command prompt:

```
SGOS#(config) bandwidth-management
SGOS#(config bandwidth-management) view configuration
```

2. To view the BWM configuration for a specific class, enter the following command:

```
SGOS#(config bandwidth-management) view configuration bwm_class
```

For example:

```
SGOS#(config bandwidth-management) view configuration Office_A
Class Name:           Office_A
Parent:               <none>
Minimum Bandwidth:   unspecified
Maximum Bandwidth:   750 kbps
Priority:              0
```

3. To view the BWM configuration for the children of a specific class, enter the following commands:

```
SGOS#(config bandwidth-management) edit bwm_class
SGOS#(config bw-class bwm_class) view children
```

Viewing Bandwidth Management Statistics

See "Bandwidth Management Statistics" on page 849 for information about viewing BWM statistics.

Using Policy to Manage Bandwidth

Once you have created and configured bandwidth management classes, you need to 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, so that one class can manage bandwidth for several types of flows based on different criteria. However, any given flow will 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 will apply.

To manage the bandwidth classes you have created, you can either compose CPL (see "CPL Support for Bandwidth Management" below) or you can use VPM (see "VPM Support for Bandwidth Management" on page 388). To see examples of policy using these methods, see "Bandwidth Allocation and VPM Examples" on page 388 or "Policy Examples: CPL" on page 396.

CPL Support for Bandwidth Management

You must use policy to classify traffic flows to different bandwidth classes. Refer to the *Blue Coat ProxySG Content Policy Language Guide* for more information about writing and managing policy.

CPL Triggers

You can use all of the CPL triggers for BWM classification (refer to the *Blue Coat ProxySG Content Policy Language Guide* for information about using CPL triggers). Note that basing a bandwidth decision on a trigger means that the decision will not take effect until after the information needed to make that decision becomes available. For example, if you set the CPL to trigger on the MIME type of the HTTP response, then the HTTP headers must be retrieved from the OCS before a classification can be made. The decision to retrieve those headers is made too late to count any of the request bytes from the client or the bytes in the HTTP response headers. However, the decision would affect the bytes in the body of the HTTP response and any bytes sent back to the client.

Supported CPL

Bandwidth class can be set with policy on each of these four traffic flows:

- `limit_bandwidth.client.inbound(none | bwm_class)`
- `limit_bandwidth.client.outbound(none | bwm_class)`
- `limit_bandwidth.server.inbound(none | bwm_class)`
- `limit_bandwidth.server.outbound(none | bwm_class)`

If you set policy to `none`, the traffic is unclassified and it will not be bandwidth-managed.

VPM Support for Bandwidth Management

You can manage bandwidth using VPM in the Action column of four policy layers: Web Access, DNS Access, Web Content, and Forwarding Layers. For more information about using VPM to manage bandwidth, see "Manage Bandwidth" on page 518. For examples of bandwidth management scenarios using VPM, see "Bandwidth Allocation and VPM Examples" below.

Bandwidth Allocation and VPM Examples

This section illustrates how to allocate bandwidth, arrange hierarchies, and create policy using the Visual Policy Manager. 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 through either the Management Console or the CLI, see "Configuring Bandwidth Allocation" on page 381. For examples of bandwidth management tasks done by composing CPL, see "Policy Examples: CPL" on page 396.

Task One: Bandwidth Allocation

The administrator is responsible for managing the bandwidth of three branch offices. He has been told to make sure 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 in Figure 10-5. He also creates policy rules for each class, as described below in "Task One: VPM".




Figure 10-5: Bandwidth Hierarchy Diagram One

Each of the classes in Figure 10-5 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 VPM and creates a new Web Access Layer, calling 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, as shown in Figure 10-6.

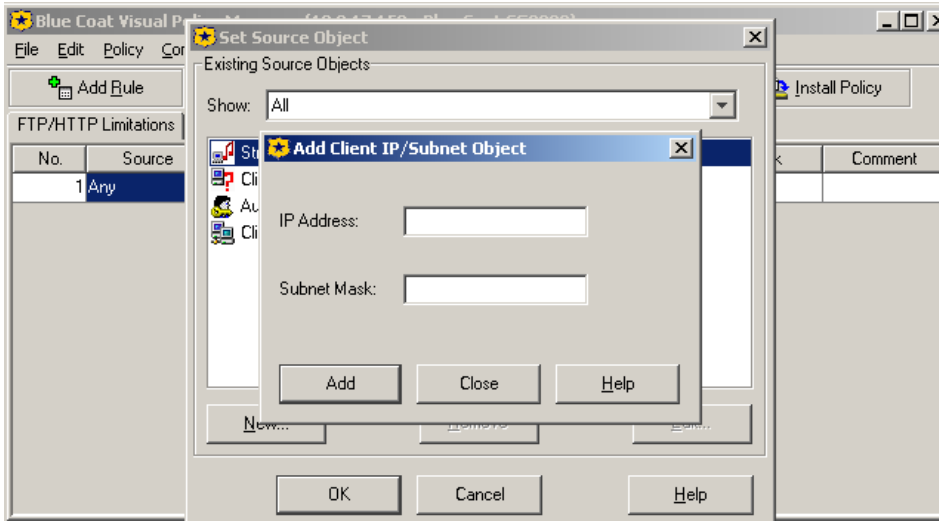


Figure 10-6: Adding the Client IP Address and Subnet Mask to the Source Column

He selects a Combined Service Object in the Service column, calling it FTP/HTTP and adding a Client Protocol for FTP and for HTTP. In the Add Combined Service Object dialog, he adds both protocols to the top box, as shown in Figure 10-7.

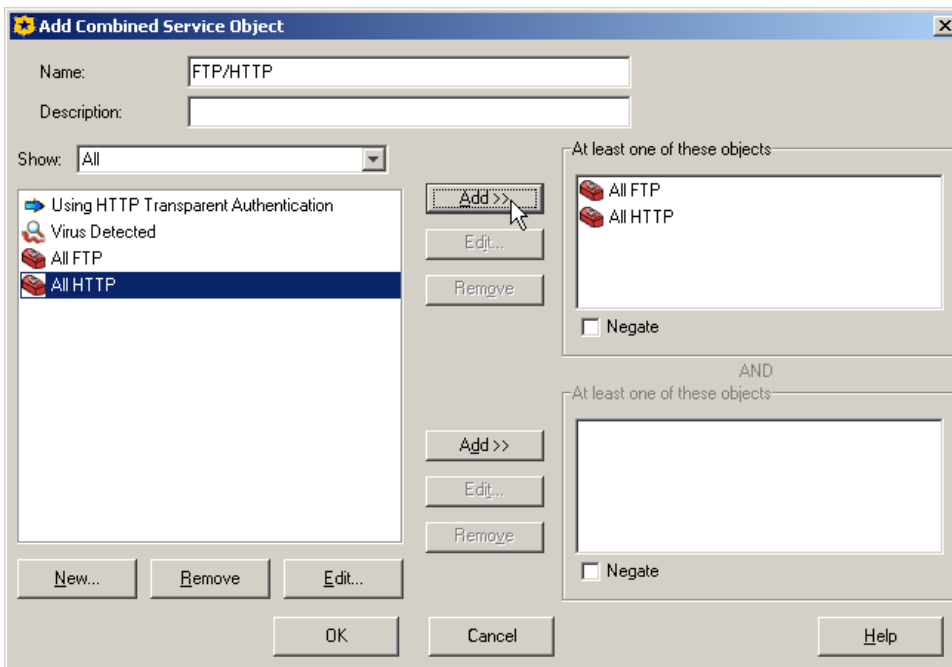


Figure 10-7: Adding Protocols to a Combined Service Object

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, as shown in Figure 10-8.

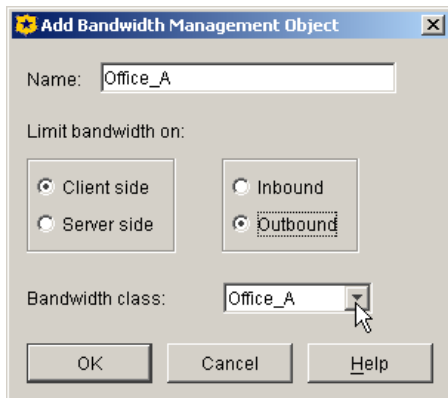


Figure 10-8: Manage Bandwidth Action Object

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 instance, 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" below). Figure 10-9 shows the hierarchical relationship among all of the classes.




Figure 10-9: Bandwidth Hierarchy Diagram Two

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 (see Figure 10-10).

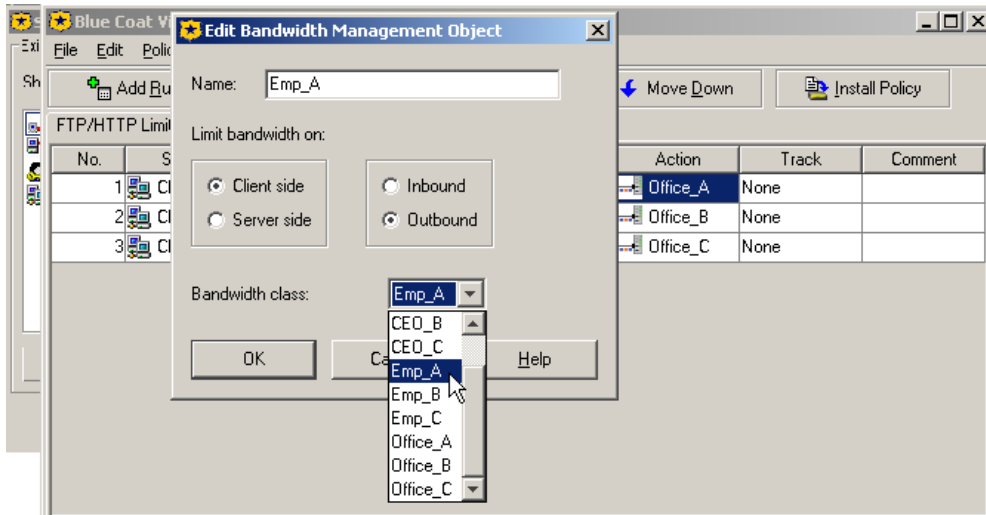


Figure 10-10: Editing the Bandwidth Management Object

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 the Client IP/subnet object already created for each office with a User object that he creates for the CEO, as shown in Figure 10-11.

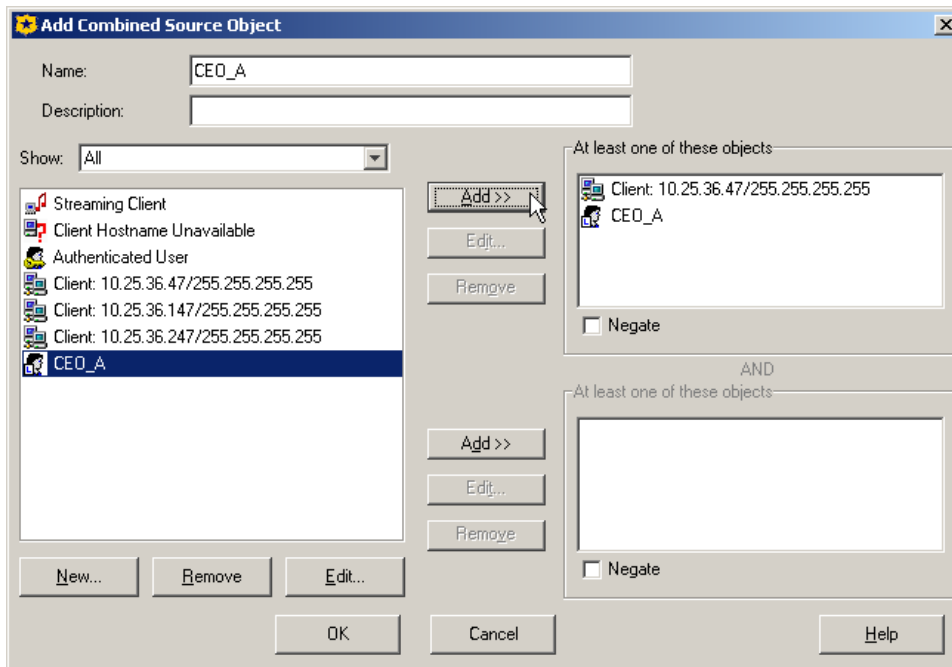


Figure 10-11: Adding a Combined Source Object

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 ProxySG 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 ProxySG will apply the CEO rules to the CEO's Web surfing, and an employee's Web surfing will not trigger the CEO rules and will instead skip 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 will 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 in Figure 10-12. To see what the administrator did to the policy rules, see "Task Three: VPM" below.

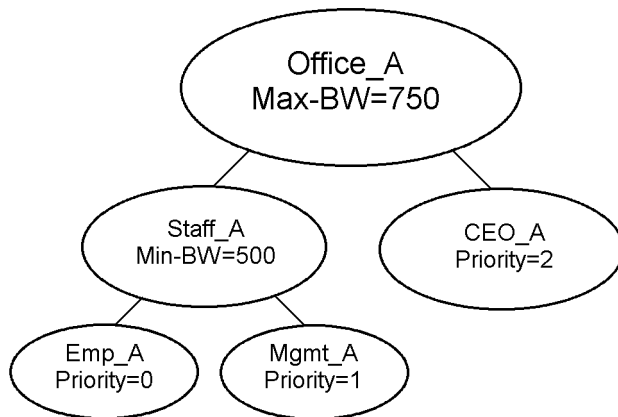


Figure 10-12: Bandwidth Hierarchy Diagram Three

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". One of the new classes he added for each office is a parent class that will 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 in Figure 10-13.

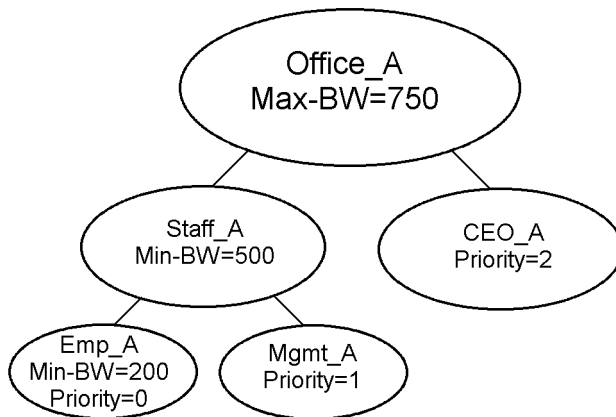


Figure 10-13: Bandwidth Hierarchy Diagram Four

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 would probably need to create a minimum for management as well.

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 making sure that bandwidth management is enabled, he's 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 388.

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 commands:

```
SGOS#(config) bandwidth-management
SGOS#(config 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 commands:

```
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 11: External Services

This chapter describes how to configure the ProxySG to interact with external ICAP and Websense servers to provide content scanning, content transformation, and content filtering services.

This chapter contains the following sections:

- "Section A: ICAP"—Describes the ICAP protocol and describes how to create and manage ICAP services and patience pages on the ProxySG.
- "Section B: Websense"—Describes how to create a Websense service
- "Section C: Service Groups"—Describes how to create service groups of ICAP or Websense entries and configure load balancing.
- "Section D: Displaying External Service and Group Information"—Describes how to display external service configurations through the CLI.

Related Topics:

- Appendix 12: "Health Checks"
- Appendix 18: "Content Filtering"

Section A: ICAP

Section A: ICAP

This section describes the Internet Content Adaptation Protocol (ICAP) solution of content scanning and modification.

When integrated with a supported ICAP server, the ProxySG provides content scanning, filtering, and repair service for internet-based malicious code. ICAP is an evolving architecture that allows an enterprise to dynamically scan and change Web content. To eliminate threats to the network and to maintain caching performance, the ProxySG sends objects to the integrated ICAP server for checking and saves the scanned objects in its object store. With subsequent content requests, the appliance serves the scanned object rather than rescan the same object for each request.

Configuring ICAP on the ProxySG involves the following steps:

1. Install the ICAP server.
2. Configure the ProxySG to use ICAP and configure basic features.
3. Define scanning policies, then load the policy file on the ProxySG.

Supported ICAP Servers

The ProxySG supports the following ICAP servers:

- Symantec AntiVirus Scan Engine (SAVSE) 4.0, version 4.04.46; ICAP 1.0
- Trend Micro InterScan WebProtect (ISWP) 1.5, build_SOL_1266; ICAP 1.0
- WebWasher 4.4, build 552; ICAP 1.0
- Finjan Vital Security for Web v7.0; Service Pack 2; build 4.552; ICAP 1.0

Note: While SGOS 2.x supported ICAP v0.95 servers and services, SGOS 3.2.x does not. Upon upgrading to SGOS 3.2.x, any configured v0.95 services become inactive.

ICAP v1.0 Features

This section describes features of the ICAP v1.0 protocol.

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 403.

ISTags

An ICAP v1.0 server is required to return in each response an ICAP header IStag indicating the current state of the ICAP server. This eliminates the need to designate artificial pattern version numbers, as is required in v0.95.

Section A: ICAP

Note: 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 value, use the Sense Settings option as described in "Creating an ICAP Service" on page 403.

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 further requests. If a connection error occurs, the connection closes to prevent further errors.

About Content Scanning

The ProxySG ICAP scanning solution prevents the spread of viruses and other malicious code by serving content that has been scanned by a supported ICAP server.

Determining Which Files to Scan

In determining which files to scan, this integrated solution uses the content scanning server's filtering in addition to ProxySG capabilities. Table 11.1 describes the supported content types and protocols.

Table 11.1: Content Types Scanned By ICAP Server and the ProxySG

ICAP Server supported content types	ProxySG supported protocols	Unsupported content protocols
<p>All or specified file types, based on file extension, as configured on the server.</p> <p>Examples: <code>.exe</code> (executable programs), <code>.bat</code> (batch files), <code>.doc</code> and <code>.rtf</code> (document files), and <code>.zip</code> (archive files), or with specific MIME types.</p>	<ul style="list-style-type: none"> • HTTP objects • FTP objects (uploads and downloads) • Transparent FTP responses 	<ul style="list-style-type: none"> • Streaming content (for example, RTSP and MMS) • Live HTTP streams (for example, HTTP radio streams)
	HTTPS connections terminated at a ProxySG.	HTTPS connections tunneled through a Blue Coat client-side ProxySG.

After the ProxySG retrieves a requested Web object from the origin server, it uses content scanning policies to determine whether the object should be sent to the ICAP server for scanning. If cached objects are not scanned or are scanned before a new pattern file was updated, the ProxySG rescans those objects upon:

- the next request for that object, or
- the object is refreshed.

Section A: ICAP

With the ProxySG, you can define flexible, enterprise-specific content scanning policies. Consider the following example. A business wants to scan software downloaded by employees from popular shareware Web sites. To do this, the business defines an appliance policy that includes a custom *scanshareware* action for the purpose. This rule includes URL domains related to the relevant shareware Web sites.

Before continuing, plan the types of policies you want to use. For more information, see "Creating ICAP Policy" on page 413.

Performing 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 ICAP server features and configuration determine how scanning works, including the following:

- Handling of certain objects, including those that are infected and cannot be repaired.
- Whether to attempt to repair infected files.
- Whether to delete infected files that cannot be repaired from the ICAP server's archive.

Performing Request Modification

The ProxySG sends the client request to a ICAP server that supports request modification. The server might then 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.

Note: Some ICAP servers do not support virus scanning for request modification, only content filtering.

Returning the Object to the ProxySG

This object may 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.

Section A: ICAP

Caching and Serving the Object

Once an object has been scanned and is determined cacheable, the ProxySG saves it and serves it for the subsequent content requests. 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, distributes, and refreshes, as well as pipelining fetches.

For more information on policies, see "Creating ICAP Policy" on page 413. For more information on the <Cache> layer, refer to the *Blue Coat ProxySG Content Policy Language Guide*.

Installing the ICAP Server

Follow the manufacturer instructions for installing the ICAP server, including any configuration necessary to work with the Blue Coat 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 error message displayed to end users in the event the requested object was modified or blocked.

Creating an ICAP Service

An ICAP service on the ProxySG is specific to the ICAP server and includes the server IP address or hostname, as well as the supported number of connections. If you are using the ProxySG with multiple ICAP servers or multiple scanning services on the same server, add an ICAP service for each server or scanning service.

To Create and Configure an ICAP Service through the Management Console:

1. Select Configuration>External Services>ICAP Services.
2. Click New; the Add List Item dialog appears.
3. In the ICAP service name field, enter an alphanumeric name; click OK.
4. Highlight the new ICAP service name and click Edit; the Edit ICAP Service dialog appears.

Section A: ICAP

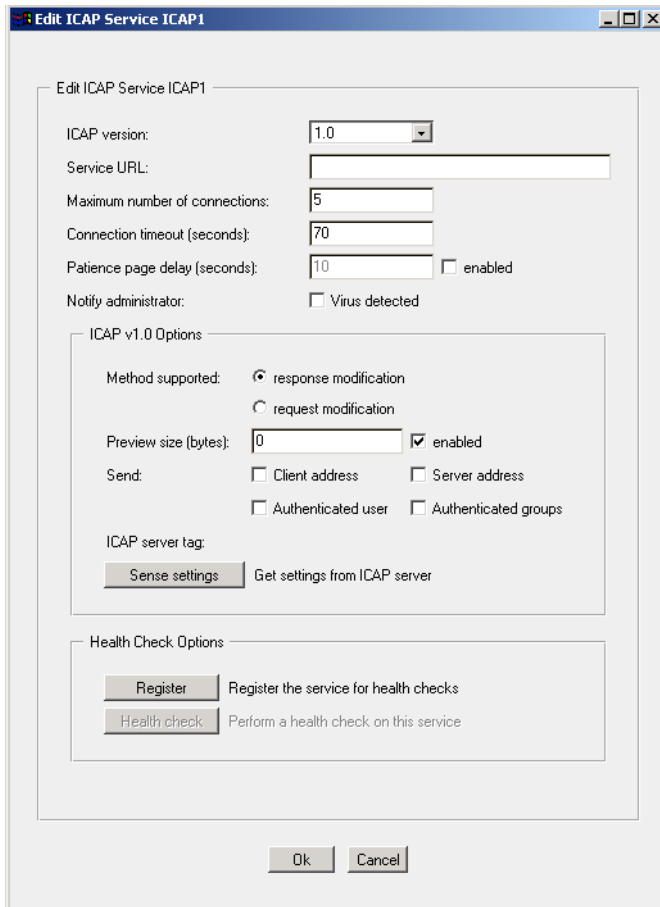


Figure 11-14: ICAP Service Dialog

The default ICAP version is 1.0 and cannot be changed.

5. Enter or select the following information:
 - a. The service URL, which includes the URL schema, ICAP server hostname or IP address, and the ICAP port number. For example:

`icap://10.x.x.x/`

The default port number is 1344, which can be changed; for example:

`icap://10.x.x.x:99`. You can also give an HTTP URL, but you must define a port number.

Note: An ICAP service pointing to a WebWasher server must use `icap` as the protocol in the URL. Blue Coat also recommends that you review your specific ICAP server documentation, as each vendor might require additional URL information.

Section A: ICAP

- b. The maximum number of connections possible at any given time between the ProxySG and the ICAP server. The range is a number from 1 to 65535. The default is 5. The number of recommended connections is dependent on the capabilities of the ICAP server. Refer to the vendor's product information.
- c. The number of seconds the ProxySG waits for replies from the ICAP server. The range is 60 to 65536. The default timeout is 70 seconds.
- d. Optional: You can enable the ProxySG to display a default patience page when an ICAP server is processing a relatively large object. The page informs users that a content scan is in process. If enabled, the patience page is displayed after the designated time value is reached for scanned objects. Patience pages might not be displayed for truncated objects; Blue Coat recommends increasing the maximum cacheable object size (up to 1 GB) to reduce the amount of truncated objects.

Note: Patience pages display regardless of any pop-up blocking policy that is in effect.

To enable the patience page, in the Patience page delay field, enter the number of seconds the ProxySG waits before displaying the page. The range is 5 to 65535. Select Enable.

- e. Select Notify administrator: Virus detected to send an email to the administrator if the ICAP scan detects a virus. The notification is also sent to the Event Log and the Event Log email list.
6. The following steps configure ICAP v1.0 features:
- a. Select the ICAP method: response modification or request modification.

Note: An ICAP server might have separate URLs for response modification and request modification services.

- b. Enter the preview size (in bytes) and select preview size enable. 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. The default is 0. Only response headers are sent to the ICAP server; more object data is only sent if requested by the ICAP server.

Note: Trend Micro does not support previews for request modification mode.

- c. (Optional) Click Send: Client address or Server address to specify what is sent to the ICAP server: Send: Client address, Server address, Authenticated user, or Authenticated groups.
- d. (Optional) Clicking Sense Settings automatically configures the ICAP service using the ICAP server parameters. If you use the sense settings feature, no further steps are required; the ICAP service is configured. Otherwise, proceed with the manual configuration.

7. Click OK; click Apply.

To Register a Newly Created ICAP Service for Health Checking:

For convenience, the Edit ICAP Service dialog allows you to register a newly-created ICAP service for health checking (this duplicates the functionality on the Configuration>Health Checks>General tab). Registering for health checking requires that a valid ICAP server URL was entered.

Section A: ICAP

- Click Register; a dialog prompts confirmation; click OK.
- You can also click Health check to perform an immediate health check on this service.

To Monitor ICAP Health Checks:

In a browser, enter one of the following URLs to list health check information.

- To list all health check entries and their configuration parameters, enter:

```
http://ProxySG_IP_address:8081/health_check/view
```

- To list the statistics for all currently active entries, enter:

```
http://ProxySG_IP_address:8081/health_check/statistics
```

For more information about health checks, see Chapter 12: “Health Checks” on page 431.

To Create and Configure an ICAP Service through the CLI:

1. At the (config) command prompt, enter the following commands:

```
SGOS# (config) external-services  
SGOS# (config external-services) create icap service_name
```

Specify a unique alphanumeric name for each service.

2. To configure the service, enter the following commands:

```
SGOS# (config external-services) edit service_name  
SGOS# (config icap service_name) url icap://url
```

where *url* specifies the URL schema, ICAP server hostname or IP address, and the ICAP port number. The default port number is 1344.

Note: While `http://url:1344` is valid, an ICAP service pointing to a WebWasher server *must* use `icap` as the protocol in the URL.

```
SGOS# (config icap service_name) max-conn number
```

where *number* is the maximum number, from 1 to 65535, of connections the ICAP service uses to connect to the ICAP server. The default is 5. Blue Coat recommends that the setting not exceed 200.

```
SGOS# (config icap service_name) timeout timeout_seconds
```

where *timeout_seconds* is the number of seconds, from 60 to 65535, the ProxySG waits for replies from the ICAP server. The default timeout is 70 seconds.

```
SGOS# (config icap service_name) notify virus-detected
```

Sends an email to the administrator if the ICAP scan detects a virus. The notification is also sent to the Event Log and the Event Log email list.

3. The following commands configure ICAP v1.0 features:

```
SGOS# (config icap service_name) methods {REQMOD | RESPMOD}
```

Specifies the ICAP service type: request modification or response modification.

Note: On most ICAP servers, one URL is designated for response modification and one for request modification.

```
SGOS# (config icap service_name) preview-size bytes
```

Section A: ICAP

where *number* is the preview size in bytes. If specified, 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.

The default is 0. Only response headers are sent to the ICAP server; more object data is only sent if requested by the ICAP server.

Optional:

```
SGOS# (config icap service_name) send {client-address | server-address}
```

Specifies to send the client IP address or the server IP address to the ICAP server.

```
SGOS# (config icap service_name) send {authenticated-user |
authenticated-groups}
```

Specifies to send authenticated user and group information to the ICAP server.

4. Optional: If the ICAP server is a version 1.0 server, you can use the `sense-settings` command to automatically configure the ICAP service using ICAP server parameters. Otherwise, proceed with the manual configuration in Step 3. To use the ICAP server parameters, enter the following command:

```
SGOS# (config icap services service_name) sense-settings
```

The ICAP service is now configured. No further steps are required.

5. Optional: You can enable the ProxySG to display a default patience page when an ICAP server is processing a relatively large object. The page informs users that a content scan is in process. If enabled, the patience page is displayed after the designated time value is reached for scanned objects. Patience pages might not be displayed for truncated objects; Blue Coat recommends increasing the maximum cacheable object size (up to 1 GB) to reduce the amount of truncated objects. To customize patience pages, see "Customizing ICAP Patience Text" on page 408.

```
SGOS# (config icap services service_name) patience-page seconds
```

where *seconds* is the duration before the patience page is displayed. The range is 5 to 65535. The default is disabled.

Note: Patience pages display regardless of any pop-up blocking policy that is in effect.

Deleting an ICAP Service

The following steps describe how to delete an ICAP service.

Note: You cannot delete an ICAP service used in a ProxySG policy (that is, if a policy rule uses the ICAP service name) or that belongs to a service group.

To Delete an ICAP Service through the Management Console:

1. Select Configuration>External Services>ICAP.
2. Select the service to be deleted.
3. Click Delete; click OK to confirm.
4. Click Apply.

Section A: ICAP

To Delete an ICAP Service through the CLI:

At the (config) prompt, enter the following commands:

```
SGOS# (config) external-services  
SGOS# (config external-service) delete service_name
```

Customizing ICAP Patience Text

This section describes how to customize text displayed during ICAP scanning.

HTTP Patience Text

The ProxySG allows you to customize the patience pages that are displayed when HTTP clients experience delays as Web content is scanned. You can customize the following patience page components:

- Header—Contains HTML tags that define what appears in the dialog title bar. This component also contains the `<meta http-equiv>` tag, which is used to specify a non-English character set.




Figure 11-15: Customizing the Header Component

- Summary—HTML and text that informs users that a content scan is occurring.

Section A: ICAP

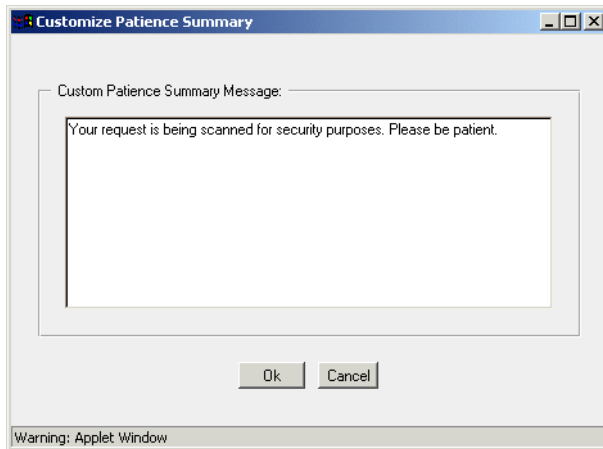


Figure 11-16: Customizing the Summary Component

- **Details**—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

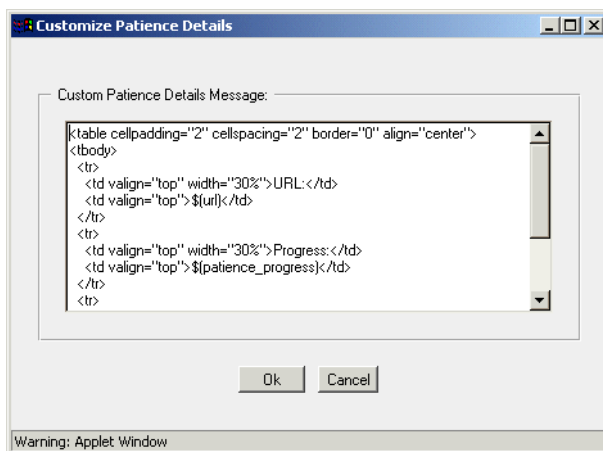


Figure 11-17: Customizing the Details Component

- **Help**—Displays instructions for users should they experience a problem with the patience page.

Section A: ICAP

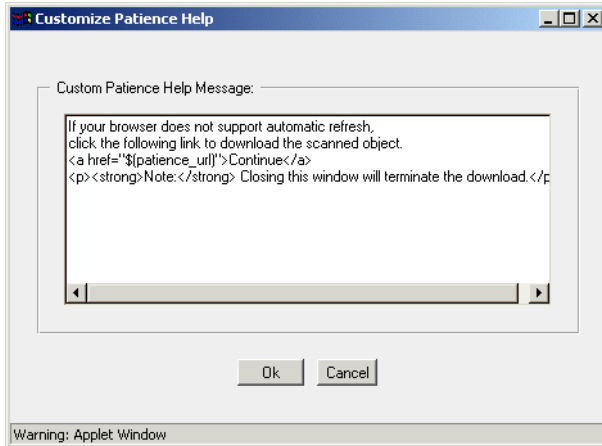


Figure 11-18: Customizing the Help Component

All of these components are displayed on the patience page.

To Customize ICAP Patience Text through the Management Console:

1. Select Configuration>External Services>ICAP>ICAP Patience Page.
2. In the HTTP Patience Page Customization field, click Header, Summary, Details, or Help; the appropriate customize dialog appears. Customize the information as appropriate.
3. Click OK; click Apply.

Example

The following example demonstrates customizing the message summary.

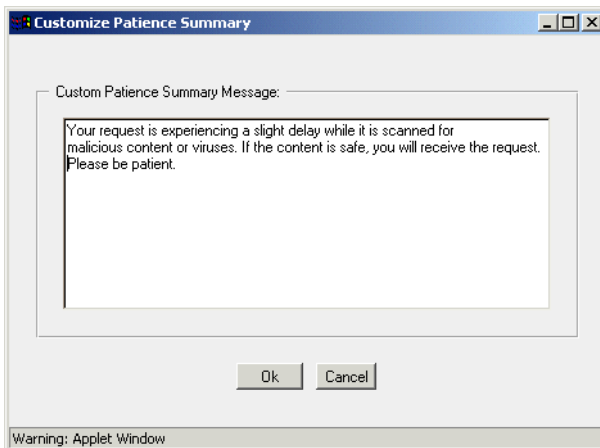


Figure 11-19: Entering a Custom Summary Message

Section A: ICAP

To Customize ICAP Patience Text through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS# (config) external-services
SGOS# (config external-services) icap-patience {details-inline |
header-inline | help-inline | summary-inline} eof
```

where:

<code>eof</code>	Specifies the end-of-file marker. After entering customized text, enter the end-of-file marker to end the customizing process.
<code>details-inline</code>	The string that displays the progress of the content scanning.
<code>header-inline</code>	The title of the page. Appears in the dialog title bar. The default is: Please be patient
<code>help-inline</code>	Clients with browsers that do not support automatic refresh must click a link to load the content after scanning is complete. The default is: If your browser does not support automatic refresh, click the following link to download the scanned object. Continue. Note: Closing this window terminates the download.
<code>summary-inline</code>	The text message informing users that a content scan is occurring. The default is: Your request is being scanned for security purposes. Please be patient.

Example:

```
SGOS# (config) external-services
SGOS# (config external-services) icap-patience summary-inline eof
Your request is experiencing a slight delay while it is scanned for malicious
content or viruses. If the content is safe, you will receive the request.
Please be patient. eof
SGOS# (config external-services)
```

Windows XP, Service Pack 2 Behavior

With Windows XP, Microsoft is continually updating the 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.

Section A: ICAP

- Users experience two patience page responses for non-cacheable objects.

Interactivity and Limitations

- Patience pages and exceptions can only be triggered by left-clicking a link. If a user right-clicks a link and attempt 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*); 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. Once 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 retain the ability to provide a patience page, configure HTTP:

```
#ProxySG (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 is displayed in another 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

The patience text displayed to FTP clients during an ICAP scan can be modified.

To Customize ICAP Patience Text through the Management Console:

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 appears. Customize the FTP client patience text. Customize the information as appropriate.
3. Click OK; click Apply.

To Customize ICAP Patience Text through the CLI:

At the `(config)` command prompt, enter the following commands:

```
SGOS# (config) external-services
SGOS# (config external-services) inline ftp icap-patience-text eof
```


Section A: ICAP

Creating ICAP Policy

Defined ICAP policy dictates the anti-virus behavior for your enterprise. You can either use the Visual Policy Manager (VPM) or you can manually edit policy files. For more information on the VPM and defining policies, see Chapter 14: “The Visual Policy Manager” on page 453.

Use the `request.icap_service()` (request modification) or `response.icap_service()` (response modification) properties to manage the ProxySG ICAP services.

VPM Objects

The VPM contains the following objects specific to AV scanning (linked to their descriptions in the VPM chapter).

Table 11.1: VPM ICAP Objects

Object	Layer>Column
"Virus Detected"	Web Access>Service
"ICAP Error Code"	Web Access>Service
"Return ICAP Patience Page"	Web Access>Action
"Set ICAP Request Service"	Web Access>Action
"Set ICAP Request Service"	Web Content>Action
"Set ICAP Response Service"	Web Content>Action

Note: For CPL policy, refer to the *Blue Coat Systems ProxySG Content Policy Language Guide*.

Example ICAP 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).

For this example:

- The ProxySG has configured ICAP services. The response service is `corporateav1` and the request service is `corporateav2`.
- The ProxyAV is the virus scanner and is configured to serve password-protected files.
- A group named IT is configured on the ProxySG.
- The IT group wants to be allowed to download password protected files, but deny everyone else.

Procedure—To perform virus scanning, protecting both the server side and client side:

1. In the VPM, select Policy>Web Content Layer. Name the layer RequestAV.
2. Right-click the Action column; select Set. The Set Action Object dialog appears.
 - a. Select Set ICAP Request Service; the Add ICAP Request Service Object dialog appears.
 - b. From the Use ICAP request service drop-down list, select `corporateav2`.

Section A: ICAP

- c. Select 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.

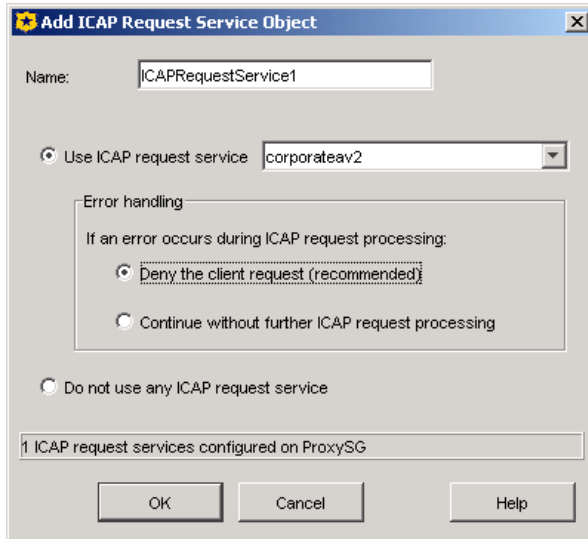


Figure 11-20: Specifying an ICAP Request Service Object.

- d. Click OK; click OK again to add the object to the rule.

No.	Source	Destination	Service	Time	Action	Track	Comment
1	Any	Any	Any	Any	ICAPRequestService1	None	

Figure 11-21: The Web Content Layer policy.

3. In the VPM, select Policy>Web Access Rule. Name the rule ResponseAV.
4. Right-click the Action column; select Set. The Set Action Object dialog appears.
 - a. Select Set ICAP Response Service; the Add ICAP Response Service Object dialog appears.
 - b. From the Use ICAP response service drop-down list, select corporateav1.
5. 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.

Procedure—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 appears.
 - a. Select Virus Detected (static object).
 - b. Click OK to add the object to the rule.

Section A: ICAP

3. Right-click the Action column. Select Delete.
4. Right-click the Track column. Select Set; the Set Track Object dialog appears.
 - a. Click New; select Event Log. The Event Log dialog appears.
 - 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	None	 VirusLog1	

Figure 11-22: The AVErrors rule.

Procedure—Create an exception for IT group:

1. In the 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 appears.
 - a. Click New; select ICAP Error Code. The Add ICAP Error Code Object appears.
 - b. Name the object `password_protected`.
 - c. Select Selected Errors.

Section A: ICAP

- d. From the list of errors, select Password Protected Archive; click Add.

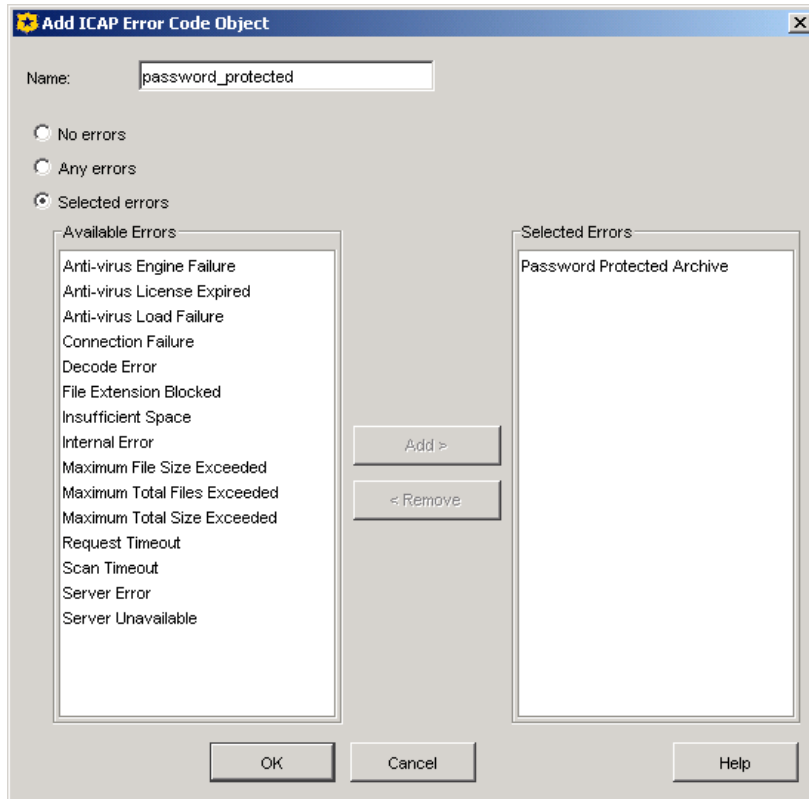


Figure 11-23: Specifying an ICAP Error Code object.

- e. Click OK; click OK again to add the object to the rule.
- 4. Right-click the Action column and select Allow.
- 5. Click Add Rule.
- 6. In the Service column, add the password_protected object.
- 7. Right-click the Action column; select Deny.

No.	Source	Destination	Service	Time	Action	Track	Comment
1	IT	Any	password_protected	Any	Allow	None	
2	Any	Any	password_protected	Any	Deny	None	

Figure 11-24: The AVException layer.

Once this policy is installed:

- Virus scanning is performed for client attempts to upload content and content responses to client requests.

Section A: ICAP

- If a virus is detected and there were no scanning process errors, a log entry occurs.
- As the 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.

```
<proxy>
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 Limitation

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 the following workaround policy to bypass the ICAP request modification service for HTTP progressive downloads:

```
<proxy>
url.scheme=http request_header.User-Agent="user_agent"
  request.icap_service(no)
url.scheme=http request_header.User-Agent="user_agent"
  request.icap_service(no)
```

where *user_agent* specifies a media player attribute that is disrupting service. 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)
```

CPL Notes

- 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 objects to be served without ICAP processing if the server is down.

```
request.icap_service(service_name, fail_open)
response.icap_service(service_name, fail_open)
```

Section A: ICAP

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 the ICAP service is restored, these objects are scanned and served from the cache if they are requested again.

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.

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 "Creating a Service Group" on page 425.

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 "Creating a Service Group" on page 425. For more information on defining policies, see Chapter 13: "Managing Policy Files" on page 439, as well as the *Blue Coat ProxySG Content Policy Language Guide*.

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 may need to configure the preview size the appliance uses. For information, see "Creating an ICAP Service" on page 403.

Section A: ICAP

Replacing the ICAP Server

If you replace an ICAP server with another supported ICAP server, reconfigure the ICAP service on the ProxySG:

```
SGOS# (config) external-services
SGOS# (config external-service) edit service_name
SGOS# (config service_name) url url
```

For information about these commands, see "Creating an ICAP Service" on page 403.

Access Logging

The ProxySG provides access log support for Symantec, Trend Micro, 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= <i>number</i>	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.

Trend Micro Interscan WebProtect v 1.5

When of these Trend Micro ICAP servers performs a scan, identifies a problem (for example, a virus), and performs a content transformation, the action is logged. For example:

```
"virus-id: name"
```

where *name* specifies the name of the virus.

Important: The `ivscan.ini` ISWP configuration file on the Trend Micro server must contain the following entry:

```
'yes': security_gateway_virus_log=yes.
```

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:

Section A: ICAP

```
"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 Access Logging, see Chapter 20: "Access Logging" on page 743.

References

The following are selected references for this feature.

Note: As with any Web site, addresses are subject to change or deletion at any time.

- **Symantec**—A provider of internet security technology, including content and network security software and appliance solutions.
<http://www.symantec.com/>
<http://enterprisesecurity.symantec.com/products/>
- **Trend Micro**—A provider of network anti-virus and internet content security software and services.
<http://www.trendmicro.com/>
- **Finjan**—A provider of proactive active content defense, virus protection, and Web and email content filtering solutions.
<http://www.finjan.com/>
- **ICAP Forum**—A resource on Internet Content Adaptation Protocol (ICAP), an evolving Web architecture. ICAP effectively adapts content for user needs.
<http://www.i-cap.org/>

Section B: Websense

Section B: Websense

This section describes how to create and manage Websense off-box services on the ProxySG. The ProxySG supports Websense off-box server versions 4.3 and higher.

For more information about Websense and content filtering, see Chapter 18: “Content Filtering” on page 635.

Creating a Websense Service

To Configure a Websense Off-box Service through the Management Console:

1. Select Configuration>External Services>Websense.
2. Click New; the Add List Item dialog appears.
3. In the Add Websense Service field, enter an alphanumeric name; click OK.
4. Highlight the new Websense service name and click Edit; the Edit Websense Service *Name* dialog appears.

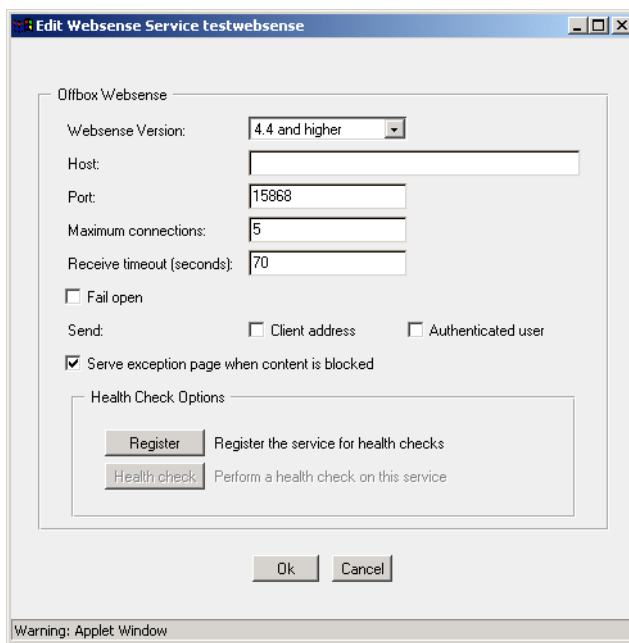


Figure 11-25: The Edit Websense Service Dialog

5. Enter following information:

Section B: Websense

- a. Select the Websense server version: 4.3 or 4.4 and higher.
 - b. In the Host field, enter the hostname or IP address of the remote Websense server.
 - c. In the Port field, enter the port number of the Websense server; or leave as is to accept the default (15868).
 - d. In the Maximum connections field, enter the maximum number of connections. The range is a number from 1 to 65535. The default is 5. Blue Coat recommends that the setting not exceed 200.
 - e. In the Receive Timeout field, enter the number of seconds the ProxySG waits for replies from the Websense server. The range is 60 to 65535. The default timeout is 70 seconds.
6. Select the following options, as required:
 - a. Fail open—If a default Websense service is selected (from the External Services>Websense tab), a connection error with the Websense server results in requests and responses proceeding, as the default Websense service is subjected to policy.
 - b. Send client address—Sends the client IP address to the Websense server.
 - c. Send Authenticated user—Sends user information to the Websense server.
 - d. Serve exception page when content is blocked—If the requested content is defined by Websense as inappropriate, the client receives a page with information stating the content is blocked. When this option is selected, the exception page originates from the ProxySG; if not selected, the Websense server provides the exception page.
 7. Click OK.
 8. Optional: You can designate a default Websense service. On the Configuration>External Services>Websense tab, select a service from the Default service to use drop-down list.

To Register a Newly Created Websense Service for Health Checking:

For convenience, the Edit Websense Service dialog allows you to register a newly-created Websense service for health checking (this duplicates the functionality on the Configuration>Health Checks>General tab). Registering for health checking requires that a valid ICAP server URL was entered.

- Click Register; a dialog prompts confirmation; click OK.
- You can also click Health check to perform an immediate health check on this service.

For more information about health checks, see Chapter 12: “Health Checks” on page 431.

To Configure an ICAP Service through the CLI:

1. At the `(config)` command prompt, enter the following commands:

```
SGOS# (config) external-services
SGOS# (config external-services) create websense service_name
```

Specify a unique alphanumeric name for each service.
2. To configure the service, enter the following commands:

```
SGOS# (config external-services) edit service_name
SGOS# (config websense service_name) version {4.3 | 4.4}
```

where *version* specifies 4.3 or 4.4 and higher.

```
SGOS# (config websense service_name) host {hostname | IP_address}
```

Section B: Websense

where *hostname* or *IP_address* specifies the Websense server.

```
SGOS# (config websense service_name) port port_number
```

where *port_number* specifies the port number of the Websense server. The default port number is 15868.

```
SGOS# (config websense service_name) max-conn number
```

where *number* is the maximum number, from 1 to 65535, of connections the Websense service uses to connect to the Websense server. The default number is 5. Blue Coat recommends that the setting not exceed 200.

```
SGOS# (config websense service_name) timeout timeout_seconds
```

where *timeout_seconds* is the number of seconds, from 60 to 65535, the ProxySG waits for replies from the Websense server. The default timeout is 70 seconds.

```
SGOS# (config websense service_name) send {client-address |
authenticated-user}
```

Specifies to send the client IP address or authenticated user information to the Websense server.

- Optional: You can automatically detect the categories defined on the Websense server.

```
SGOS# (config websense service_name) sense-categories
```

- Optional: You can designate a default Websense service.

```
SGOS# (config websense service_name) apply-by-default
```

This Websense service is now the default and is used if failover is enabled.

- Optional: You can enable failover. If a default Websense service is selected (from the External Services>Websense tab), a connection error with the Websense server results in requests and responses proceeding, as the default Websense service is subjected to policy.

```
SGOS# (config websense service_name) fail-open
```

- Optional: You can send a test URL to the Websense server to verify content filtering is active.

```
SGOS# (config websense service_name) test-url url
```

where *url* is a valid URL that points to a site determined categorized by Websense as inappropriate.

Deleting a Websense Service

The following steps describe how to delete an Websense service.

Note: You cannot delete a Websense service used in a ProxySG policy (that is, if a policy rule uses the Websense service name) or if the service belongs to a service group.

To Delete a Websense Service through the Management Console:

- Select Configuration>External Service>Websense.
- Select the service to be deleted.
- Click Delete; click OK to confirm.

Section B: Websense

4. Click Apply.

To Delete an Websense Service through the CLI:

At the (config) prompt, enter the following commands:

```
SGOS# (config) external-services  
SGOS# (config external-services) delete service_name
```

Section C: Service Groups

Section C: Service Groups

This section describes how to create and manage ICAP or Websense service groups. In high-traffic network environments, a service group accelerates response time by performing a higher volume of scanning.

Creating a Service Group

Create the service group and add the relevant ICAP or Websense services to the group. Services within group must be the same type (ICAP or Websense).

To Configure a Service Group through the Management Console:

1. Select Configuration>External Services>Service-Groups.
2. Click New; the Add List Item dialog appears.
3. In the Add Service Group field, enter an alphanumeric name; click OK.
4. Highlight the new service group name and click Edit; the Edit Service Group dialog appears.



Figure 11-26: The Edit Service Group Dialog

5. Click New to add a service to the service group; the Add Service Group Entry dialog appears.

Section C: Service Groups

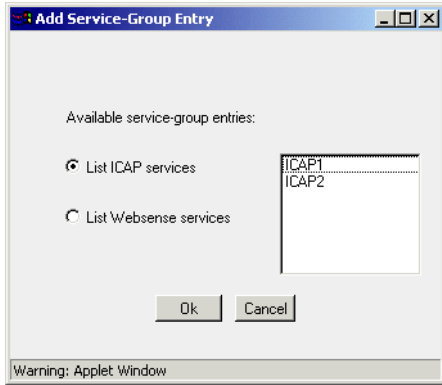


Figure 11-27: The Add Service Group Entry Dialog

6. Select List ICAP services or List Websense services. The picklist displays the available configured services that are eligible for this service group.
7. Select a service; to select multiple services, use Ctrl-click. Click OK.
8. To assign a weight value to a service, select a service and click Edit; the Edit Service Group Entry weight dialog appears. In the Entry Weight field, assign a weight value. The valid range is 0-255. For detailed information about service weighting, see the next topic, "About Weighted Load Balancing" on page 427.
9. Click OK; click OK again to close the Edit Service Group Entry dialog
10. Click Apply.

To Configure a Service Group through the CLI:

1. At the (config) command prompt, enter the following commands:


```
SGOS# (config) external-services
SGOS# (config external-services) create service-group name
SGOS# (config service-group name) add service_name
```

Enter a unique alphanumeric name for each service; the ICAP or Websense service must already exist on the ProxySG.
2. Repeat the `add service_name` command for each service to be added.

The type of service group (ICAP or Websense) is determined by the first service added. For example, if the first added service is an ICAP service, the service group is automatically defined as an ICAP service group. If you attempt to add a Websense service, an error is displayed.
3. To assign weights to each service, enter the following commands:


```
SGOS# (config service-group name) edit service_name
SGOS# (config service-group name) weight value
```

where *value* is from 0 to 255. For information about weight values, see "About Weighted Load Balancing" on page 427.

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 through the Management Console:

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 through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS# (config) external-services
SGOS# (config external-services) delete service_group_name
```

To Delete a Service Group Entry through the Management Console:

1. Select Configuration>External Services>Service-Groups.
2. Select the service group to be modified.
3. Click Edit.
4. Select the service entry; click Delete.
5. Click OK; click Apply.

To Delete a Service Group Configuration through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS# (config) external-services
SGOS# (config external-services) edit service_group_name
SGOS# (config type name) remove entry_name
```

About Weighted Load Balancing

The ProxySG supports weighted load balancing in forwarding requests to service groups. By default, the ProxySG performs typical round-robin load balancing and evenly forwards requests sequentially to servers as defined within the service group. Manually assigning weights takes advantage of round-robin load balancing in service groups that are not homogeneous, or where the servers have different capacities.

Weighting determines what proportion of the load one server bears relative to the others. If all servers have either the default weight (1) or the same weight, each share an equal proportion of the load. 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 relative weights to assign to each server. Factors that could affect assigned weight of a ICAP server include the following:

- 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)
- The maximum number of connections configured for the service. Note that the maximum connections setting pertains to how many simultaneous scans can be performed on the server, while weighting applies to throughput in the integration. While these settings are not directly related, consider both when configuring weighted load balancing. For more information on maximum connections, see "Creating an ICAP Service" on page 403 and "Creating a Websense Service" on page 421.

The table below provides an example of how weighting works with a service group of three ICAP servers, Server1, Server2, and Server3. Because Server3 is a higher-capacity server (including dual CPUs and multiple network interface cards (NICs)) compared to Server1 and Server2, it is assigned a heavier weight. Using the weights below, for every 100 requests forwarded to the service group, Server3 receives 60 requests, while Server1 and Server2 each receive 20 requests.

Table 11.1: Example of Weighted Load Balancing for an ICAP Service Group

ICAP server	Capacity	ICAP service / Maximum connections	Weight
Server1	Standard	Service1 / 10	1
Server2	Standard	Service2 / 10	1
Server3	High	Service3 / 25	3

Note: 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.

While you cannot specifically designate an ICAP server in a group as a backup, you can specify weight values that create a large differential between a server that is used continuously and one that is rarely used, thus simulating a backup server.

Section D: Displaying External Service and Group Information

After configuring a service or service group, you can display information either for all or individual service groups.

To Display Information about all External Services and Groups through the CLI:

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
icap4
ICAP-Version:          1.0
URL:                   icap://10.1.1.1
Max-conn:              5
Timeout(secs):        70
Health-checks:        no
Patience-page(secs): disabled
Notification:          never
Methods:               RESPMOD
Preview-size:          0
Send:                  nothing
ISTag:

websense4
Version:               4.4
Host:                  www.websense.com/list
Port:                  15868
Max-conn:              5
Timeout(secs):        70
Send:                  nothing
Fail-by-default:      closed
Apply-by-default:     no
Serve-exception-page: yes

; External Service-Groups
CorpICAP
total weight 5
entries:
  ICAP1
    weight 4
  ICAP2
    weight 1
BranchWebsense
total weight 2
entries:
  Websense1
    weight 1
  Websense2
    weight 1
```

To Display Information about an Individual Service or Service Group through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS (config) external-services  
SGOS# (config external-services) edit {service_name | service_group_name}  
SGOS# (config type name) view
```

Chapter 12: Health Checks

This chapter discusses health checks for services and hosts and describes how to configure the ProxySG.

About General Health Checks

The ProxySG can perform health checks on a forwarding host or external server that is providing a service. The supported server types are HTTP, HTTPS, ICAP, Websense (off-box), and SOCKS gateways, Layer-3, and Layer 4 forwarding hosts.

Based on the health check type, the ProxySG periodically verifies the health status, and thus the availability, of the host. The time interval between checks is configurable. If the health check is successful, the ProxySG considers the host available. If the initial health check is not successful for a host, the ProxySG retries, using the number of attempts in the health check failure count. If the health check is not successful for every server in a domain, the ProxySG might not serve stale content from its object store, depending on the ProxySG configuration.

The following table describes the types of health checks.

Table 12.1: Types of Health Checks

Health check type	Description
HTTP	Use this type to confirm that the host can fulfill a content request over HTTP by the ProxySG. The ProxySG accepts only a 200 OK as a healthy response.
Criterion for success	The ProxySG fetches the object.
Criterion for failure	The ProxySG cannot fetch the object.
HTTPS	Use this type to confirm that the host can fulfill a content request over HTTPS by the ProxySG. The ProxySG accepts only a 200 OK as a healthy response.
Criterion for success	The ProxySG fetches the object.
Criterion for failure	The ProxySG cannot fetch the object.
Layer-3 health check	Use this type to confirm the basic connection between the ProxySG and the origin server. The server must recognize ICMP echoing. The ProxySG sends a ping (three Internet Control Message Protocol [ICMP] echo requests) to the host.
Criterion for success	The ProxySG receives at least one ICMP echo reply.
Criterion for failure	The ProxySG does not receive a single ICMP echo reply.
Layer-4 health check	Use this type to confirm that the ProxySG can connect to the host HTTP and FTP ports. The ProxySG attempts to establish a TCP connection to an HTTP port or FTP port on the host.

Table 12.1: Types of Health Checks (Continued)

Health check type	Description
Criterion for success	The ProxySG establishes the connection to the defined port (of any type), then closes it. For global forwarding checks, the first defined port in the forwarding host port list is used for the attempt (except for SOCKS gateways, in which the SOCKS port is used).
Criterion for failure	The ProxySG cannot establish the connection.
ICAP health check and Websense 4 off-box	Requests are not sent to <i>sick</i> services. If a health check determines the service is healthy, requests resume.

Configuring Service-Specific Health Checks

This section describes how to create a health check service for a specific host (for example, an ICAP server). A failed health check results in administrator notification; however, unlike global forwarding health checks, the ProxySG does not recognize the healthy or sick status of the host and thus alters where it sends transactions.

To Configure Health Checks through the Management Console:

Part 1: General Tasks

This part of the procedures is the same for all health check types.

1. Select Configuration>Health Checks>General.
2. Click New.
3. In the Add Health Check dialog, specify a name for the health check service; click OK.
4. In the Health Check list, select the newly created service and click Edit; the Edit Health Check dialog displays.

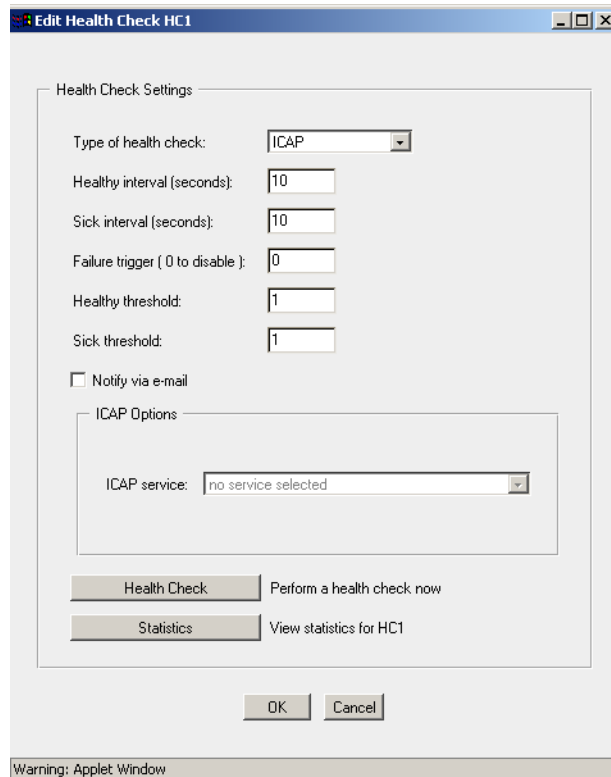


Figure 12-1: Edit Health Check Dialog

5. Select the health check type (HTTP, HTTPS, ICAP, Layer-3, Layer-4, or Websense off-box).
6. Specify the healthy interval, in seconds, between health checks to the server. The default is 10.
7. Specify the sick interval, in seconds, between health checks to the server that has been determined to be sick, or out of service. The default is 10.
8. Specify the failure trigger, or the number of failed connections to the server before a health check is triggered. Valid values are 0-65535, where 0 disables the trigger. The default is 0.
9. Specify the healthy threshold, or the number of successful health checks before an entry is considered healthy. Valid values are 1-65535. The default is 1.
10. Specify the sick threshold, or the number of failed health checks before an entry is considered sick. Valid values are 1-65535. The default is 1.
11. Optional: Select the Notify via email checkbox to send notification mail when the health of a service changes. Recipients are specified in Management>Event Logging>Mail.

Part 2: Health Check Type Specific Tasks

This part of the procedure configures the health check based upon the type selected.

1. Upon selecting the health check type, the Options section of the dialog changes to display the appropriate configuration fields. Enter the required information:
 - HTTP and HTTPS: Enter the URL of the server to be checked.

- ❑ ICAP: Select the ICAP service. The ICAP service must already be configured on the ProxySG (see Appendix 11: "External Services").
 - ❑ Layer-3 and Layer-4: Enter the host name; for Layer-4, also enter the port number.
 - ❑ Websense off-box: Select the Websense service. The Websense service must already be configured on the ProxySG (see Appendix 11: "External Services"). Enter the URL to be test-categorized, or click Use default.
2. Click OK to close the Edit Health Check dialog; Click Apply to apply the configuration to the ProxySG.

To Specify a Health Check through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) health-check
SGOS#(config health-check) create name
SGOS#(config health-check) edit name
SGOS#(config health-check name) type {layer-3 | layer-4 | http | https | icap
| websense-offbox}
```

where *type* specifies the type of health check.

```
SGOS#(config health-check name) type parameter
```

where *type* is the type of health check and *parameter* is the required attribute:

```
layer-3 hostname host_name
layer-4 hostname host_name
layer-4 port port
{http | https} url url
```

icap servicename *service_name*—The ICAP service must already be configured on the ProxySG. See Appendix 11: "External Services".

websense-offbox servicename *service_name*—The Websense service must already be configured on the ProxySG. For more information, see Appendix 11: "External Services".

```
websense-offbox {url | default-url}
```

```
SGOS#(config health-check name) interval healthy seconds
```

where *seconds* specifies the interval between health checks to the server. The default is 10.

```
SGOS#(config health-check name) interval sick seconds
```

where *seconds* specifies the interval between health checks to the server that has been determined to be sick. The default is 10.

```
SGOS#(config health-check name) threshold healthy number
```

where *number* is the number of successful health checks before an entry is considered healthy. Valid values are 1-65535. The default is 1.

```
SGOS#(config health-check name) threshold sick number
```

where *number* is the number of failed health checks before an entry is considered sick. Valid values are 1-65535. The default is 1.

```
SGOS#(config health-check name) failure-trigger number
```

where *number* is the number of failed connections to the server before a health check is triggered. Valid values are 0-65535, where 0 disables the trigger. The default is 0.

Optional:

```
SGOS#(config) health-check name) notify
```

Sends email notification when the health of a service changes. The recipients are specified

```
in (config event-log) mail add option.
```

Perform an Instant Health Check

You can manually issue a health check request.

To Do a Health Check through the Management Console:

1. Select Health Checks>General.
2. Select a health check name.
3. Click Edit.
4. Click Health Check.

To Do a Health Check through the CLI:

At the (config) prompt, enter the following commands:

```
SGOS#(config) health-check
SGOS#(config) health-check) edit health_check_name
SGOS#(config) health-check name) perform-health-check
```

About Global Forwarding and SOCKS Gateway Health Checks

This section describes health checks that can be configured on the ProxySG that apply to all forwarding hosts and SOCKS gateway hosts.

When the ProxySG performs a health check on one or more hosts, it determines whether the host returns a response and is available to fill a content request. A positive health check indicates that there is an end-to-end connection and that the host is healthy and is able to return a response.

With multiple forwarding hosts, health checks are vital to ProxySG efficiency. When hosts respond positively to health checks, the ProxySG forwards requests to those hosts and not to unavailable hosts, which provides quicker content fill requests. With a single forwarding host, health checking is also important determine whether the host is available.

Note: When a forwarding host or SOCKS gateway is created, it is automatically registered for health checks. Similarly, when a forwarding host or SOCKS gateway is deleted, it is removed from the health check registry.

Configuring Global Health Checks

This section describes how to configure the ProxySG to perform global health checks.

To Configure Global Forwarding Host or SOCKS Gateway Health Checks through the Management Console:

1. Select Configuration>Health Checks>Forwarding or SOCKS Gateway.

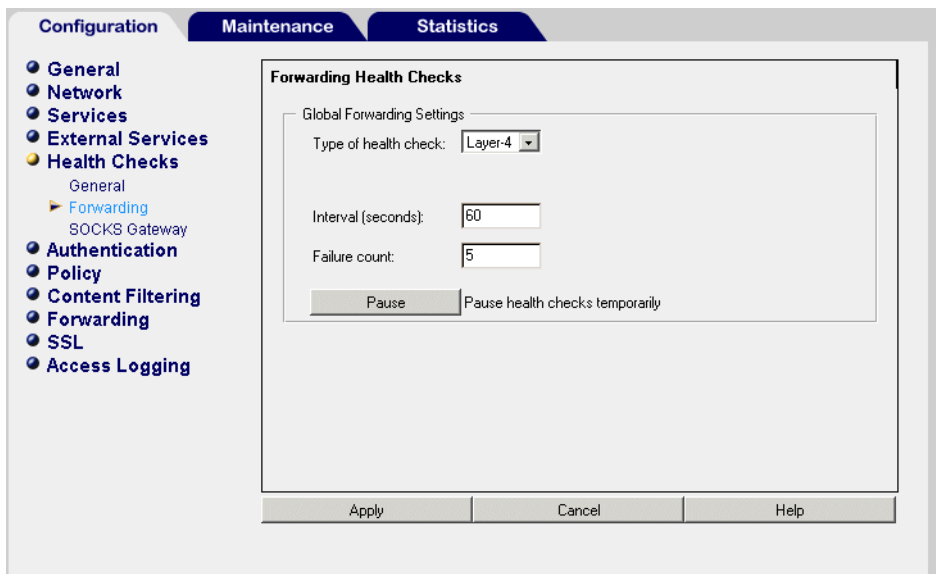


Figure 12-2: Global Forward Health Check Tab

2. Select the health check type:
 - Forwarding—HTTP, HTTPS, Layer-3, or Layer-4.
 - SOCKS Gateway—Layer-3 or Layer-4.
3. Specify the interval, in seconds, between health checks. The default is 60.
4. Specify the failure count, which specifies the number of sequential failures before the host is considered down. The default is 5.
5. Click Apply.

To Configure Global Forwarding Host Health Checks through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) health-check
SGOS#(config health-check) forwarding type {http | https | layer-3 | layer-4}
SGOS#(config health-check) forwarding interval seconds
    where seconds specifies the time between health checks.
SGOS#(config health-check) forwarding failcount count
    where count specifies the number of sequential failures before the host is considered
    down. The default is 5.
```

To Configure Global SOCKS Gateways Health Checks through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) health-check
SGOS#(config health-check) socks-gateways type {layer-3 | layer-4}
SGOS#(config health-check) socks-gateways interval seconds
    where seconds specifies the time between health checks.
```



```
SGOS#(config) health-check) socks-gateways failcount count
```

where count specifies the number of sequential failures before the host is considered down. The default is 5.

Pausing or Resuming Global Health Checking

You can temporarily halt global health checks and resume when ready. This is helpful if the ProxySG needs to be temporarily taken out of service.

Note: If the health check is paused, the state remains paused until the resume option is invoked. The paused state remains even after a reboot.

To Pause or Resume Health Checking through the Management Console:

1. Select Configuration>Health Checks>Forwarding or SOCKS Gateway.
2. Click Pause.
3. To resume health checks, click Resume.

To Pause or Resume Health Checking through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) health-check
SGOS#(config) health-check) {forwarding | socks-gateway} {pause | resume}
```


Chapter 13: *Managing Policy Files*

Policy files contain the policies that manage every aspect of the ProxySG, from controlling user authentication and privileges to disabling access logging or determining the version of SOCKS.

The policy for a given system can contain several files with many layers and rules in each. Policies can be defined through the Visual Policy Manager (VPM) or composed in Content Policy Language (CPL). (Some advanced policy features are not available in VPM and can only be configured through CPL.)

Policies are managed through four files:

- Central policy file—Contains global settings to improve performance and behavior and filters for important and emerging viruses (such as Code Red and Nimda). This file is usually managed by Blue Coat, although you can point the ProxySG to a custom Central policy file instead.
- Forward policy file—Usually used to supplement any policy created in the other three policy files. The Forward policy file contains Forwarding rules when the system is upgraded from a previous version of SGOS (2.x) or CacheOS (4.x).
- Local policy file—A file you create yourself. When VPM is not the primary tool used to define policy, the Local file contains the majority of the policy rules for a system. If VPM is the primary tool, this file is either empty or includes rules for advanced policy features that are not available in VPM.
- Visual Policy Manager—The policy created by VPM can either supplement or override the policies created in the other policy files.

This chapter contains the following sections:

- "About Policy Files"
- "Creating and Editing Policy Files"
- "Managing the Central Policy File"
- "Viewing Policy Files"

To learn about writing policies, refer to the *Blue Coat ProxySG Content Policy Language Guide*.

About Policy Files

When creating the files, keep in mind:

- The order in which the files are evaluated.
- The transaction default settings, which control whether you allow everything or deny everything by default.
- Whether to use VPM.

Policy File Evaluation

The order in which the ProxySG evaluates policy rules is important. Changes to the evaluation order can result in different effective policy, as the order of policy evaluation defines general rules and exceptions. While this order is configurable, the default and recommended order is:

```
VPM File-Local Policy File-Central Policy File-Forward File
```

This prevents policies in the Central file that block virus signatures from being inadvertently overridden by allow (access-granting) policy rules in the VPM and Local files.

When changing the policy file evaluation order, remember that final decisions can differ because decisions from files later in the order can override decisions from earlier files (the Forward policy file order cannot be changed).

For a new ProxySG, the default evaluation order is: VPM, Local, Central, and Forward.

For an upgraded ProxySG, the policy evaluation order is the order already existing on the appliance before the upgrade.

To Change Policy Order through the Management Console:

1. Select Configuration>Policy>Policy Options.

The Policy Options tab displays.

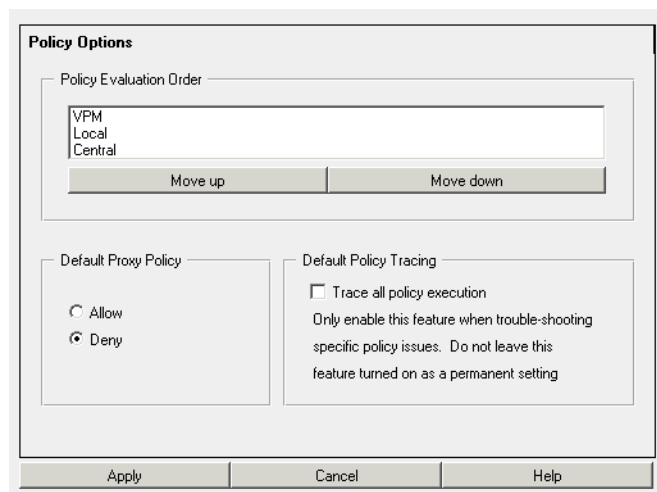


Figure 13-1: Policy Options Tab

2. To change the order, select the file to move and click the Move Up or Move Down button. Remember that the last file in the list overwrites decisions in files evaluated earlier.

To Change Policy Order through the CLI:

At the (config) command prompt, enter the command:

```
SGOS# (config) policy order v l c
```

where v (VPM), l (local), and c (central) specify the order of evaluation. These are case-insensitive, but you must enter all three in any order, including a space between each letter.

Note: Use the `show policy order` command to check the current settings.

Transaction Settings: Deny and Allow

The default `<Proxy>` transaction policy is *deny proxied transactions* or *allow proxied transactions*, depending on whether this is a new installation or an upgrade. You can change the default policy.

- A default `<Proxy>` transaction policy of `Deny` prohibits proxy-type access to the ProxySG; you must then create policies to explicitly grant access on a case-by-case basis. This is the default for those who are installing a new release of SGOS without an upgrade and for administrator transactions.
- A default `<Proxy>` transaction policy of `Allow` permits any and all proxy-types access to the ProxySG; you must then create policies to explicitly deny access on a case-by-case basis. This is the default for those upgrading from a previous version of SGOS and for `<Cache>` transactions.

Changing the default `<Proxy>` transaction policy affects the basic environment in which the overall policy is evaluated. It is likely that you must revise policies to retain expected behavior after such a change.

Also consider:

- Changes to the evaluation order might result in different effective policy, because the order of policy evaluation defines general rules and exceptions.
- Changes made to `<Proxy>` transactions do not affect `<Cache>` transactions and `<Admin>` transactions.

To Configure Deny or Allow Default Policy through the Management Console:

1. Select Configuration>Policy>Policy Options.
2. Under Default Proxy Policy, select either Deny or Allow.
3. Click Apply.

To Configure the Deny or Allow <Proxy> Transaction Policy through the CLI:

At the `(config)` command prompt, enter the following command

```
SGOS#(config) policy proxy-default {allow | deny}
```

Policy Tracing

Tracing enabled with the Management Console or CLI is global; that is, it records every policy-related event in every layer. It should be used only while troubleshooting. For information on troubleshooting policy, refer to the *Blue Coat ProxySG Content Policy Language Guide*. Turning on policy tracing of any kind is expensive in terms of system resource usage, and it will slow down the ProxySG's ability to handle traffic.

To Enable Policy Tracing through the Management Console:

1. Select Configuration>Policy>Policy Options.
2. Select Trace all policy execution.

3. Click Apply.

To Enable Policy Tracing through the CLI:

From the command prompt, enter the following command:

```
SGOS# policy trace {all | none}
```

Creating and Editing Policy Files

You can create and edit policy files two ways:

- Through the Management console (recommended).
- Through the CLI inline policy command (not recommended because the policies can grow large and using `inline policy` overwrites any existing policy on the ProxySG).

You can use VPM to create policy layers and rules in the VPM file. For information on managing the VPM file, see Chapter 14: “The Visual Policy Manager” on page 453.

To create or edit policy files, use CPL to define policy rules (refer to the *Blue Coat ProxySG Content Policy Language Guide*). You can use the Management Console or CLI to create or edit policy files directly, or create a file that can be uploaded to the ProxySG through the Management Console or CLI.

Create and Edit Policy Files

You can install the policy files in the following ways.

- Using the ProxySG Text Editor, which allows you to enter directives (or copy and paste the contents of an already-created file) directly onto the ProxySG.
- 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.
- Through the CLI `inline` command.

The ProxySG compiles the new policy from all source files and installs the policy, if the compilation is successful.

Important: If errors or warnings are produced when you load the policy file, a summary of the errors and/or warnings is displayed automatically. If errors are present, the policy file is not installed. If warnings are present, the policy file is installed, but the warnings should be examined.

To Define and Install Policy Files Directly through the Management Console:

1. Select Configuration>Policy>Policy Files>Policy Files.

The Policy Files tab displays.




Figure 13-2: Policy Files Tab

2. From the appropriate Install Local/Forward/Central File from drop-down list, select the method you want to use to install the local, forward, or central policy configuration; click Install and complete one of the three procedures below:

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

- Installing a policy file using a Remote URL:

In the Install Local/Forward/Central File dialog that appears, enter the fully-qualified URL, including the filename, where the policy configuration is located. To view the file before installing it, click View. Click Install. The Installation Status field summarizes the results; click the Results button to open the policy installation results window. Close the window when you are finished viewing the results; click OK in the Install Local/Forward/Central File dialog.

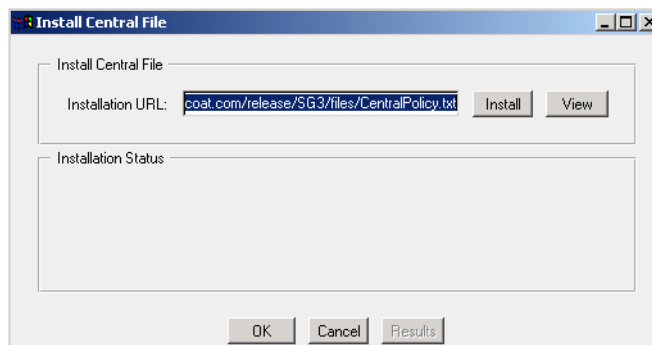


Figure 13-3: Policy Files Remote Installation Dialog

Note: If you use the default Blue Coat Central policy file, load it from:
`https://download.bluecoat.com/release/SG3/files/CentralPolicy.txt`

If you install a Central policy file, the default is already entered; change this field only if you want to create a custom Central policy file.

To load a Forward, Local, or a custom Central policy file, move it to an HTTP or FTP server, and then use that URL to download the file to the ProxySG.

❑ **Installing a policy file using a Local File:**

In the Upload and Install File window that opens, either enter the path to the file into the File to upload field, or click Browse to display the Choose file dialog, locate the file on the local system, and open it. Click Install. When the installation is complete, the installation results display. View the results and close the window.



Figure 13-4: Specifying the Local Location of a Policy File

- Installing a policy file using a Text Editor:

The current configuration is displayed in installable list format. Define the policy rules using CPL in the Edit and Install File window that opens (refer to the *Blue Coat ProxySG Content Policy Language Guide*); click Install. When the installation is complete, a results window opens. View the results, close the results window and click Close in the Edit and Install File window.



Figure 13-5: Edit and Install File Window

Using the CLI Inline Command

To create policies using the CLI, you can use the ProxySG `inline policy` command. This command either creates a new policy file or, if the specified file already exists, overwrites an existing policy file. You cannot edit an existing policy file using this command.

Note: If you are not sure whether a policy file is already defined, check before using the `inline policy` command. For more information, see "Viewing Policy Source Files" on page 450.

To Create Policy Files through the CLI:

1. At the `(config)` command prompt, enter the following command:

```
SGOS#(config) inline policy file end-of-input-marker
```

where *file* specifies the type of policy you want to define: `Central` (Central policy file), `Forward` (Forward policy file), or `local` (local policy file).

Note: Do not use the `inline policy` command with files created using the VPM module.

end-of-file-marker—Specifies the string that marks the end of the current `inline` command input; `eof` usually works as a string. The CLI buffers all input until you enter the marker string.

2. Define the policy rules using CPL (refer to the *Blue Coat ProxySG Content Policy Language Guide*). Enter each line and press <Enter>. To correct mistakes on the current line, use <Backspace>. If a mistake has been made in a line that has already been terminated by <Enter>, exit the `inline policy` command by typing `Ctrlc` to prevent the file from being saved.
3. Enter the `eof` marker to save the policies and exit the `inline` mode.

For more information on the `inline` command, refer to the *Blue Coat ProxySG Command Line Reference*.

To Load Policy Files through the CLI:

At the `(config)` command prompt, enter the following commands:

```
SGOS#(config) policy {forward-path | local-path | central-path} url
SGOS#(config) load policy {forward | local | central}
```

The ProxySG compiles and installs the new policy. The ProxySG might display a warning if the new policy causes conflicts. If a syntax error is found, the appliance displays an error message. For information about these messages, refer to the *Blue Coat ProxySG Content Policy Language Guide*. Correct the error, then reload the file.

Unloading Policy Files

To disable policies, do the following procedure to unload the compiled policy file from the ProxySG memory. These steps describe how to replace a current policy file with an empty policy file.

To keep a current policy file, either make a backup copy or rename the file before unloading it. By renaming the file, you can later reload the original policy file. If you use multiple policy files, back up or rename files as necessary. Alternatively, rather than use an empty policy file, you can delete the entire contents of the file, then reload it.

To unload policies defined using the VPM, you can either:

- Do the procedure below for unloading policies through the CLI.
- Use the VPM and individually delete all layers.

To Unload Policies through the Management Console:

1. Select Configuration>Policy>Policy Files>Policy Files.
2. Select Text Editor in the Install Local/Forward/Central File from drop-down list and click the appropriate Install button.

The Edit and Install the Local/Forward/Central Policy File window opens.

3. Delete the text and click Install.
4. View the results in the results page that opens; close the page.
5. Click Close.

To Unload Policies through the CLI:

1. At the `(config)` command prompt, enter the following command:

```
SGOS#(config) inline policy file end-of-input-marker
```

where:

<code>file</code>	Specifies the type of policy you want to define: <code>central</code> (central policy file), <code>local</code> (local policy file), <code>vpm-cpl</code> , or <code>vpm-xml</code> (VPM policy files, usually defined using the VPM).
<code>end-of-input-marker</code>	Specifies the string that marks the end of the current <code>inline</code> command input. The CLI buffers all input until you enter the marker string. <code>eof</code> is commonly used as the marker.

Note: If you use the CLI to unload VPM-generated policies, you must run the `inline` command twice; once for the CPL file and once for the XML file.

2. Enter an `end-of-input-marker` to save the policies and exit inline mode. Enter nothing else.
3. If you use multiple policy files, repeat step 1 and step 2 for each policy file used.

For more information on the `inline policy` command, refer to the *Blue Coat ProxySG Command Line Reference*.

Managing the Central Policy File

The Central policy file is updated when needed by Blue Coat. The file can be updated automatically or you can request e-mail notification. You can also configure the path to point to your own custom Central policy file.

Configuring Automatic Installation

You can specify whether the ProxySG checks for a new version of the Central policy file. If a new version exists, the appliance can install it automatically.

Configuring the ProxySG for Automatic Installation

Do the following procedure to configure the ProxySG to check for and install a new version of the Central policy file.

To Configure Automatic Installation through the Management Console:

1. Select Configuration>Policy>Policy Files>Policy Files.
2. Select Automatically install new Policy when central file changes.
3. Click Apply.

To Configure Automatic Installation through the CLI:

At the `(config)` command prompt, enter the following command:

```
SGOS#(config) policy subscribe
```

Configuring a Custom Central Policy File for Automatic Installation

If you define your own Central policy file, you can configure the ProxySG to automatically install any subsequent updated version of the file. To use this capability, you must change the Central policy file's first line with each version update. With automatic installation, the ProxySG checks for a change to the first line of the file. In defining a custom Central policy file, add an item, such as a comment, to the first line of the Central policy file that changes with each update. The following is a sample first line, containing date information that is routinely updated with each version:

```
; Central policy file MonthDate, Year version
```

When you update and save the file in the original location, the ProxySG automatically loads the updated version.

Configuring E-mail Notification

You can specify whether the ProxySG sends e-mail when the Central policy file changes. The e-mail address used is the same as that used in diagnostic reporting; the event recipient for the custom heartbeat e-mail. For information about diagnostic reporting, see "Diagnostic Reporting (Heartbeats)" on page 958.

To Configure E-mail Notification through the Management Console:

1. Select Configuration>Policy>Policy Files>Policy Files.
2. Select Send me email when central file changes.
3. Click Apply.

To Configure E-mail Notification through the CLI:

At the (config) command prompt, enter the following command:

```
SGOS#(config) policy notify
```

Configuring the Update Interval

You can specify how frequently the ProxySG checks for a new version of the Central policy file. By default, the appliance checks for an updated Central policy file once every 24 hours (1440 minutes). You must use the CLI to configure the update interval. You cannot configure the update interval through the Management Console.

To Configure the Update Interval through the CLI:

At the (config) command prompt, enter the following command:

```
SGOS#(config) policy poll-interval minutes
```

Checking for an Updated Central Policy File

You can manually check whether the Central policy file has changed. You must use the CLI. You cannot check for updates through the Management Console.

To Check for an Updated Central File through the CLI:

At the (config) command prompt, enter the following command:

```
SGOS#(config) policy poll-now
```

The ProxySG displays a message indicating whether the Central file has changed.

Resetting the Policy Files

You can clear all the policy files automatically through the CLI.

To Clear all Policy Files through the CLI:

1. At the (config) command prompt, enter the following command:

```
SGOS#(config) policy reset
```

```
WARNING: This will clear local, central, forward and VPM policy. Are you sure
you want to reset ALL policy files? (y or n)
```

The ProxySG displays a warning that you will be resetting all of your policy files.

2. Enter **y** to continue or **n** to cancel.

Note: This command does not change the default proxy policy settings.

Moving VPM Policy Files from One ProxySG to Another

VPM policy files are specific to the ProxySG where they were created. But just as you can use the same Central, Local, and Forward policy files on multiple ProxySG Appliances, you can use VPM policies created on one appliance on other appliances.

For detailed information on moving VPM policy files, see "Installing VPM-Created Policy Files" on page 545 in Appendix 14: "The Visual Policy Manager".

Viewing Policy Files

You can view either the compiled policy or the source policy files. Use these procedures to view policies defined in a single policy file (for example, using VPM) or in multiple policy files (for example, using the Blue Coat Central policy file and VPM).

Viewing the Installed Policy

Use the Management Console or a browser to display installed Central, Local, or Forward policy files.

Note: You can view VPM policy files through the Visual Policy Files tab.

To View Installed Policy through the Management Console:

1. Select Configuration>Policy>Policy Files>Policy Files.
2. In the View File drop-down list, select Current Policy to view the installed and running policy, as assembled from all policy source files. You can also select Results of Policy Load to view any warnings or errors resulting from the last attempt (successful or not) to install policy.
3. Click View.

The ProxySG opens a separate browser window and displays the installed policy file.

To View the Currently Installed Policy through a Browser:

1. Enter a URL in one of the following formats:

- ❑ If an HTTPS-Console is configured, use `https://ip_address_of_ProxySG:HTTPS-Console_port/Policy/current` (the default port is 8082).
- ❑ If an HTTP-Console is configured, use `http://ip_address_of_ProxySG:HTTP-Console_port/Policy/current` (the default port is 8081).

The ProxySG opens a separate browser window and displays the policy.

2. Review the policy, then close the browser.

To View the Currently Installed Policy through the CLI:

At the (config) command prompt, enter the following command:

```
SGOS#(config) show policy
```

Viewing Policy Source Files

You can display source (uncompiled) policy files on the ProxySG.

To View Policy Source Files through the Management Console:

1. Select Configuration>Policy>Policy Files>Policy Files.
2. To view a policy source file, select the file you want to view (Local, Forward, or Central) from the View File drop-down list and click View.

The ProxySG opens a separate browser window and displays the appropriate source policy file.

To View Policy Source Files through the CLI:

At the (config) command prompt, enter one of the following commands:

```
SGOS#(config) show configuration
-or-
SGOS#(config) show sources policy {central | local | forward | vpm-cpl |
vpm-xml }
```

The `show configuration` command displays general configuration information, followed by the policy source file contents. If the ProxySG is using multiple policy files, file source displays in this sequence: Central file, local file, VPM. The `show sources policy` command allows you to specify the policy files you want to view.

Note: You can use the `show configuration` command to save the output to a file for reference, in addition to displaying the current configuration. For more information, refer to the *Blue Coat ProxySG Command Line Reference*.

Viewing Policy Statistics

You can view policy statistics on all requests processed by the ProxySG. Use the Management Console or a browser. You cannot view policy statistics through the CLI.

To Review Policy Statistics through the Management Console:

1. Select Statistics>Advanced.
2. Click the Policy link.
3. Click the Show policy statistics link.

A separate browser window opens and displays the statistics.

4. Examine the statistics, then close the browser.

To Review Policy Statistics through a Browser:

1. Enter a URL in one of the following formats:
 - If an HTTPS-Console is configured, use `https://ip_address_of_ProxySG:HTTPS-Console_port/Policy/statistics` (the default port is 8082).
 - If an HTTP-Console is configured, use `http://ip_address_of_ProxySG:HTTP-Console_port/Policy/statistics` (the default port is 8081).

The ProxySG opens a separate browser window and displays the statistics.

2. Examine the statistics, then close the browser.

Chapter 14: *The Visual Policy Manager*

The Visual Policy Manager (VPM) is a graphical policy editor included with the ProxySG. The VPM allows you to define Web access and resource control policies without having an in-depth knowledge of Blue Coat Systems Content Policy Language (CPL) and without the need to manually edit policy files.

This chapter assumes that you are familiar with basic concepts of ProxySG policy functionality as described in Appendix 13: "Managing Policy Files".

While VPM creates only a subset of everything you can achieve by writing policies directly in CPL, it is sufficient for most purposes. If your needs require more advanced policies, consult the *Blue Coat ProxySG Content Policy Language Guide*.

This chapter contains the following sections:

- "Section A: About the Visual Policy Manager"
- "Section B: Policy Layer and Rule Object Reference"
- "Section C: Detailed Object Column Reference"
- "Section D: Managing Policy Layers and Files"
- "Section E: Tutorials"

Related topics:

- *Blue Coat ProxySG Content Policy Language Guide*
- Appendix 13: "Managing Policy Files"
- Appendix 18: "Content Filtering"

Section A: About the Visual Policy Manager

This section contains the following topics:

- "System Requirements" on page 455—Discusses the Java Runtime Environment component requirement.
- "Launching the Visual Policy Manager" on page 456—Describes how to start VPM from the Management Console.
- "About the Visual Policy Manager User Interface" on page 457—Describes VPM menu items, tool bars, and work areas.
- "About VPM Components" on page 460—Provides definitions of the policy layers and describes how rule objects comprise the layers.
- "The Set Object Dialog" on page 463—Describes the dialog used to select objects to be added or edited.
- "The Add/Edit Object Dialog" on page 465—Describes the dialog used to add and edit rule objects.

Section A: About the Visual Policy Manager

System Requirements

Before launching the VPM, verify client computers that are to access the VPM meet the basic requirements.

Supported Operating Systems

This VPM version supports the following operating systems:

- Windows 2000 Professional; SP4 or later
- Windows XP; SP 2 or later

Supported Browsers

This VPM version supports the following browsers on the supported operating systems:

- Internet Explorer 6.0; SP or later
- Netscape 7.2
- Firefox 1.0

The VPM *might* operate on other browsers; however, Blue Coat has not tested other browsers and support is not available.

JRE Requirement

The VPM requires the Java Runtime Environment Standard Edition (JRE). This VPM version supports JRE versions 1.4.1_07 and 5.0 (also listed as 1.5).

If a client attempting to start the VPM does not have a valid JRE version, the ProxySG automatically connects to the Sun Microsystems download center to begin the download and installation. Follow the on-screen instructions to download v5.0 (the default version for this release).

The VPM is completely independent from the Management Console. If the browser is configured properly with its default JRE, the VPM uses the later of the valid versions.

Notes and Limitations

- If you are updating to JRE v5.0, Blue Coat recommends the option to remove any version previous to 1.4.1_07 before the update. Removing the obsolete version after updating might cause the browser to not recognize v5.0, though it remains installed, and instigate a download prompt.
- JRE v5.0: When viewing objects in a drop-down list, you can press a letter key to skip to the first object name that starts with that letter. Pressing the same letter cycles to the next object. However, immediately pressing another letter key does not take you to the next object if you have not waited a few seconds. For example, in searching protocols you can repeatedly press H to cycle through the protocols HTTP, HTTPS, and so on, but if you do not wait a short interval to press F to go to FTP, no action occurs.

Section A: About the Visual Policy Manager

Launching the Visual Policy Manager

To Launch VPM:

1. Select Configuration>Policy>Visual Policy Manager.

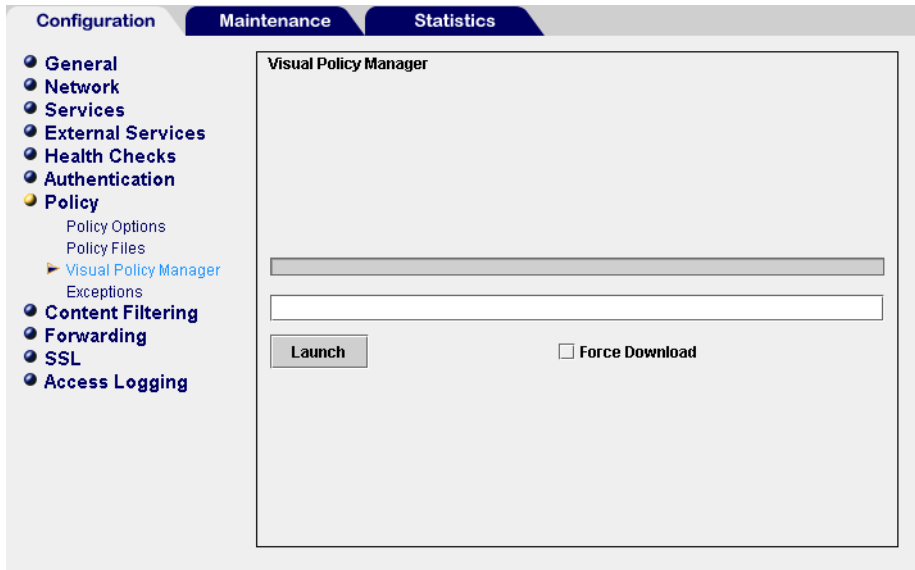


Figure 14-1: Launching VPM from the Management Console

2. Launch VPM. If this is the first time launching VPM following an OS upgrade, or to ensure you are launching the most current VPM version, click Force Download before clicking Launch.

If a valid JRE is already installed on your workstation, the ProxySG opens a separate browser window and starts VPM. The first time you start the policy editor, it displays an empty policy.

If a valid JRE is *not* installed on your workstation, a security warning dialog box appears. Click Yes to continue. Follow the instructions to install the JRE. After installation completes, a Launch VPM tab briefly displays before VPM starts.

Note: If using Internet Explorer: Depending on the browser version and settings, launching VPM from an HTTPS URL might display security warning messages (page contains secure and non-secure items). This is caused by the HTML link in the VPM launch page that links to the JRE download site at the Sun Web site. This message can be ignored. You can adjust the browser settings to display or not to display this message concerning secure and non-secure items.

Section A: About the Visual Policy Manager

About the Visual Policy Manager User Interface

The following figure labels VPM components.

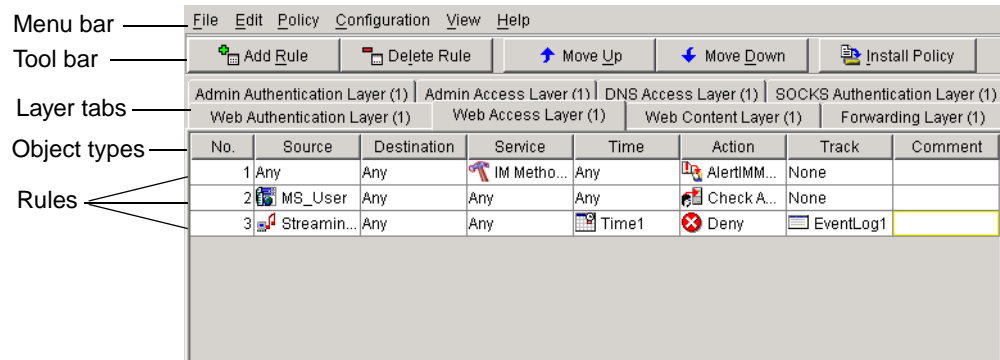


Figure 14-2: Visual Policy Manager

Section A: About the Visual Policy Manager

Menu Bar

The following table describes VPM Menu Bar items.

Table 14.1: VPM Menu Bar Items

File	Install Policy On ProxySG	Saves all new policy rules.
	Revert to existing Policy on ProxySG	Ignores changes and reloads installed policy rules.
	Exit	Exits the application.
Edit	Add Rule Delete Rule	Adds a new blank rule to the visible policy layer or removes a rule from the visible policy layer.
	Cut Rule Copy Rule Paste Rule	Standard cut, copy, and paste operations.
	Move Rule Up Move Rule Down	Moves rules up or down one position in a policy layer.
	Reorder Layers Delete Layer	Reorders the policy layers. Deletes a specific policy layer.
	Add Admin Authentication Layer Add Admin Access Layer Add DNS Access Layer Add SOCKS Authentication Layer Add Web Authentication Layer Add Web Access Layer Add Web Content Layer Add Forwarding Layer	The Policy menu items add policy layers to be populated with policy rules.
Configuration	Set DNS Lookup Restrictions	Restricts DNS lookups during policy evaluation.
	Set Reverse DNS Lookup Restrictions	Restricts reverse DNS lookups during policy evaluation.
	Set Group Log Order	Configures the order in which the group information is logged.
	Edit Categories	Edits content filtering categories.

Section A: About the Visual Policy Manager

Table 14.1: VPM Menu Bar Items

View	Generated CPL	Displays the CPL generated by VPM.
	Current ProxySG VPM Policy Files	Displays the currently stored VPM policy files.
	Object Occurrences	Lists the user-created object(s) in the selected rule; lists use in other rules as well.
	All Objects	Displays a dialog that displays current static and user-defined VPM objects. You can also create, edit, and delete objects. See "Centralized Object Viewing and Managing" on page 534.
	Tool Tips	Toggles the tool-tip display on and off.
Help	Help Topics	Displays the online help.
	About	Displays copyright and version information.

Tool Bar

The VPM Tool Bar contains the following functions:

- Add Rule—Adds a blank rule to visible policy layer; all values for the rule are the defaults.
- Delete Rule—Deletes the selected rule from the visible policy layer.
- Move Up—Moves a rule up one position in the visible policy layer.
- Move Down—Moves a rule down one position in the visible policy layer.
- Install Policy—Converts the policies created in VPM into Blue Coat Content Policy Language (CPL) and installs them on the ProxySG.

Policy Layer Tabs

Every policy layer you create from the Policy>Add Layer menu is displayed as a tab. Click a tab and the rules included in that policy layer display below in the main body of the pane. Right-clicking a tab displays the options of renaming and deleting the policy layer.

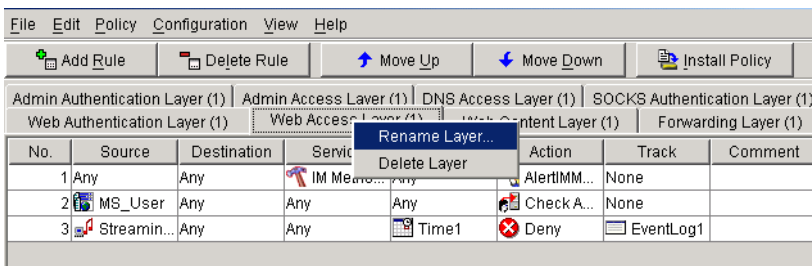


Figure 14-3: Right-click a Policy Tab to Rename or Delete a Policy Layer
Each VPM policy layer is described in later sections in this chapter.

Section A: About the Visual Policy Manager

Rules and Objects

A policy layer can contain multiple rules. Every rule is numbered and listed in a separate row. To create a new rule, click the Add Rule button; a new rule is added to the bottom of the list. If multiple rules exist within a policy layer, the ProxySG finds the first one that matches a given situation and ignores the remaining rules. Therefore, rule order is important. Use the Move buttons on the rule bar to reorder the rules in a policy.

Each rule is comprised of objects. The objects are the individual elements of a rule you specify. With the exception of No. (number), which indicates the order of the rule in the layer and is filled in automatically, all objects are configurable.

To specify or edit an object setting, position the mouse in the appropriate object cell within a rule and right-click to display the drop-down menu.

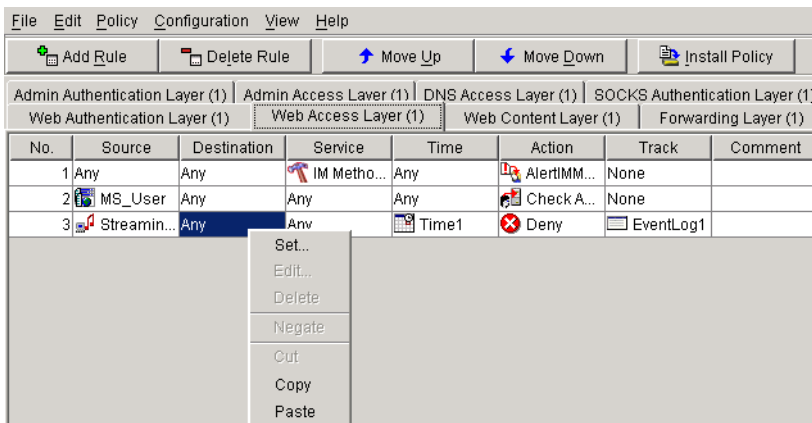


Figure 14-4: Right-click a Rule Cell to Set or Edit the Object Properties

Each object type is described in "Policy Layer and Rule Object Reference" on page 466.

About VPM Components

This section describes the specific policy layer types and rule objects.

Policy Layers

The layers are:

- Administration Authentication—Determines how administrators accessing ProxySG must authenticate.
- Administration Access—Determines who can access the ProxySG to perform administration tasks.
- DNS Access—Determines how the ProxySG processes DNS requests.
- SOCKS Authentication—Determines the method of authentication for accessing the proxy through SOCKS.
- Web Authentication—Determines whether user clients that access the proxy or the Web must authenticate.

Section A: About the Visual Policy Manager

- **Web Access**—Determines what user clients accessing the proxy or the Web can access and any restrictions that apply.
- **Web Content**—Determines caching behavior, such as verification and ICAP redirection.
- **Forwarding**—Determines forwarding hosts and methods.

As you create policy layers, you will create many different layers of the same type. Often, an overall policy requires layers of different types designed to work together to perform a task. For example, Authentication and Access layers usually accompany each other; an Authentication layer determines if a user or client must authenticate, and an Access layer subsequently determines where that user or client can go (what ProxySG or Web sites they can access) once they are authenticated.

Each object type is described in "Policy Layer and Rule Object Reference" on page 466.

Rule Objects

Policy layers contain rule objects. Only the objects available for that policy layer type are displayed. There are two types of objects:

- **Static Objects**—A self-contained object that cannot be edited or removed. For example, if you write a rule that prohibits users from accessing a specific Web site, the Action object you select is Deny.

Static objects are part of the system and are always displayed.

- **Configurable Objects**—A configurable object requires parameters. For example, consider the rule mentioned in the previous item that prohibits users from accessing a specific Web site. In this case, the user is a Source object. That object can be a specific IP Address, user, group, user agent (such as a specific browser), and so on. Select one and then enter the required information (such as a verifiable user name or group name).

Configurable objects do not exist until you create them. A created object is listed along with all static objects in the list dialog, and you can reuse it in other applicable policy layers. For example, an IP address can be a Source or Destination object in many different policy-layer types.

Important: The orders of policy layers, and the order of rules within a layer are important. For more information, see "How Policy Layers, Rules, and Files Interact" on page 541.

While individual object-type menus occasionally contain entries specific to the object type, the basic menu options are:

- **Allow**—(Web Access Layer Action column only) Quick menu access; sets the policy to allow.
- **Deny**—(Web Access Layer Action column only) Quick menu access; sets the policy to deny.
- **Set**—Displays the Set Object dialog where you select an object or create a new one.
- **Edit**—Opens the Edit Object dialog where you edit an object or change to another.
- **Delete**—Removes the selected object from the current rule and restores the default.
- **Negate**—Defined as *not*. Negate provides flexibility in writing rules and designing the structure of policies. The following is a simple Web Access rule that states: "When any client tries to access a URL contained in an object of JobSearch, allow access."

Section A: About the Visual Policy Manager




Figure 14-5: A Simple Web Access Policy Rule

The rule now specifies: “Allow access to everywhere but these JobSearch sites.”

Dragging the pointer to the Destination list, right-clicking to display the drop-down list, and clicking Negate invokes a red circle with a horizontal white line in the icon in the cell.




Figure 14-6: The Red Icon in the Cell Indicates Negation, or “Not”

- Cut, Copy, and Paste are the standard paste operations with the following restrictions: you can only paste anything cut or copied from the same column in the same table and the copy and paste functions do not work across multiple layers.

The following table describes the general function of each object type:

Object	Description
Source	Specifies the source attribute, such as an IP address, user, or group.
Destination	Specifies the destination attribute, such as a URL, IP address, and file extension.
Service	Specifies the service attribute, such as protocols, protocol methods, and IM file transfer limitations.
Time	Specifies day and time restrictions.
Action	Specifies what to do when the rule matches.
Track	Specifies tracking attributes, such as event log and E-mail triggers.
Comment	Optional. You can provide a comment regarding the rule.

Section A: About the Visual Policy Manager

Policy Layer/Object Matrix

The following table displays which object types are available in each policy layer

Policy Layer	Source	Destination	Service	Time	Action	Track	Comment
Admin Authentication	x				x	x	x
Admin Access	x				x	x	x
DNS Access	x	x		x	x	x	x
SOCKS Authentication	x				x	x	x
Web Authentication	x	x			x	x	x
Web Access	x	x	x	x	x	x	x
Web Content		x	x		x	x	x
Forwarding	x	x	x		x	x	x

The Set Object Dialog

This section discusses the Set Object dialog used to select objects for configuration.

The object rules in all policy layer types determine the conditions for a particular policy rule. Depending on the type of policy layer, an object can be anything from a user or group to an IP address or a URL and so forth.

To create a rule, right-click a cell in an object cell. The relevant Set Object dialog displays. In this dialog, select the objects for the rule or create new objects as necessary.

Objects have type-specific icons to provide a visual aid in distinguishing among different types in the list.

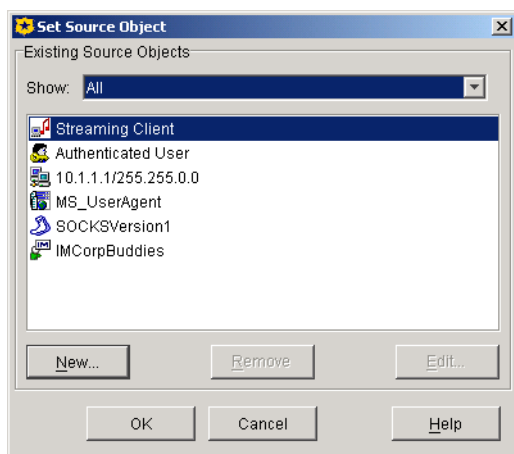


Figure 14-7: Set Source Object Dialog with Selectable Objects

Section A: About the Visual Policy Manager

The Set Object dialog only displays or allows you to create the objects allowable in the specific option of the rule type you are creating. But if more than one policy-layer type uses the same object type (for example, IP address can be a source in rules for four of the five types of policies), then those existing objects display in all Set Object dialogs, regardless of policy-layer type.

Controlling the List of Objects in the Set Object Window

As you create more policies, it is likely that the lists of existing objects in the various Set Object dialogs expand. You can restrict the display of objects in the list to a specific type by selecting an object type from the Show drop-down list above the list. The following figure demonstrates the window displayed above with the list restricted to Client IP addresses.

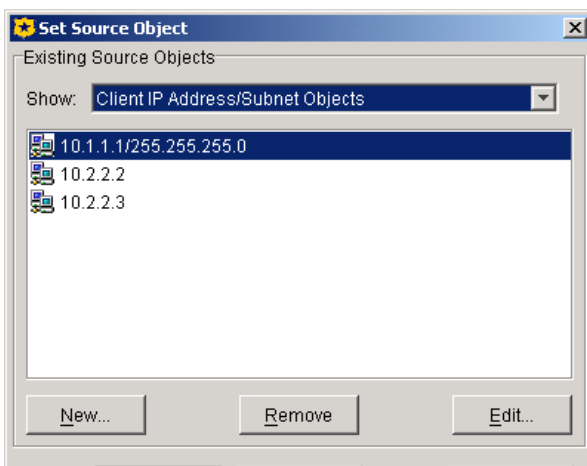


Figure 14-8: Limiting the Set Object Dialog View

The Add/Edit Object Dialog

From the Set Object dialog, the Add Object dialog is used to define configurable objects. Existing configurable options can be altered using the Edit Object dialog. In terms of functionality, the two dialogs are identical.

For the initial configuration of an object, click New on the Set Object dialog to display the Add Object dialog. Perform the tasks required to configure the object and click OK. The newly named and configured object appears in the list of selectable objects in the Set Object dialog and is ready to be selected for the rule.

To edit an existing object, select an object from the list and click Edit. The Edit Object dialog appears with the existing parameters on display. Edit as necessary and click OK.

To remove an existing object, select an object from the list and click Remove. A secondary prompt verifies your attempt to remove the object; click OK. The object is deleted.

Section A: About the Visual Policy Manager

Online Help

The VPM contains its own Help module (a porting of this chapter). Each object in the VPM contains a Help button that links to the corresponding object reference in the Help file. This reference describes the purpose of the object. Interaction with other policy and references to feature-related sections in the *Configuration and Management Guide* are provided, if relevant. Also, this Help module contains navigation buttons and its own Table of Contents.

Note: The online Help file is displayed in a separate window and requires a few seconds to load and scroll to the correct object. The speed of your system might impact this slight lag time. Furthermore, this lag time increases on slower machines running JRE v1.5.

Section B: Policy Layer and Rule Object Reference

Section B: Policy Layer and Rule Object Reference

This section contains the following topics:

- "About the Reference Tables" on page 467—Describes the table conventions used in this section.
- "Administration Authentication Policy Layer Reference" on page 467—Describes the objects available in this policy layer.
- "Administration Access Policy Layer Reference" on page 467—Describes the objects available in this policy layer.
- "DNS Access Policy Layer Reference" on page 468—Describes the objects available in this policy layer.
- "SOCKS Authentication Policy Layer Reference" on page 468—Describes the objects available in this policy layer.
- "Web Authentication Policy Layer Reference" on page 469—Describes the objects available in this policy layer.
- "Web Access Policy Layer Reference" on page 469—Describes the objects available in this policy layer.
- "Web Content Policy Layer Reference" on page 472—Describes the objects available in this policy layer.
- "Forwarding Policy Layer Reference" on page 473—Describes the objects available in this policy layer.

Section B: Policy Layer and Rule Object Reference

About the Reference Tables

The tables in this section list the static and configurable objects available for each policy layer.

Note: If viewing this document as a PDF, you can click an object name to jump to a description of that object (all objects are described in Section C). To jump back to a specific policy layer reference, click policy layer name in any object reference table that appears in the next section.

Administration Authentication Policy Layer Reference

The following table provides the objects available in the Administration Authentication policy layer.

Source Objects	Action Objects	Track Objects
Client IP Address/Subnet	Do Not Authenticate	Trace
Client Hostname	Deny	
Proxy IP Address/Port	Authenticate	
Combined Objects	Force Authenticate	

Administration Access Policy Layer Reference

The following table provides the objects available in the Administration Access policy layer.

Source Objects	Action Objects	Track Objects
Client IP Address/Subnet	Allow Read-Only Access	Event Log
Client Hostname	Allow Read-Write Access	Email
Proxy IP Address/Port	Deny	SNMP
User	Force Deny	Trace
Group		Combined Objects
Attribute		
Combined Objects		

Section B: Policy Layer and Rule Object Reference

DNS Access Policy Layer Reference

The following table provides the objects available in the DNS Access policy layer.

Source Objects	Destination Objects	Time Objects	Action Objects	Track Objects
Client IP Address/Subnet	DNS Response Contains No Data	Time	Bypass DNS Cache	Event Log
Proxy IP Address/Port	DNS Response IP Address/Subnet	Combined Objects	Do Not Bypass DNS Cache	Email
DNS Request Name	RDNS Response Host		Allow DNS From Upstream Server	SNMP
RDNS Request IP Address/Subnet	DNS Response CNAME		Serve DNS Only From Cache	Trace
DNS Request Opcode	DNS Response Code		Enable/Disable DNS Imputing	Combined Objects
DNS Request Class	Category		Send DNS/RDNS Response Code	
DNS Request Type	Combined Objects		Send DNS Response	
DNS Client Transport			Send Reverse DNS Response	
Combined Objects			Reflect IP	
			Manage Bandwidth	
			Combined Objects	

SOCKS Authentication Policy Layer Reference

The following table provides the objects available in the SOCKS Authentication policy layer.

Source Objects	Action Objects	Track Objects
Client IP Address/Subnet	Do Not Authenticate	Trace
Client Hostname	Authenticate	
Proxy IP Address/Port	Force Authenticate	
SOCKS Version		
Combined Objects		

Section B: Policy Layer and Rule Object Reference

Web Authentication Policy Layer Reference

The following table provides the objects available in the Web Authentication policy layer.

Source Objects	Destination Objects	Action Objects	Track Objects
Client Hostname Unavailable	Destination IP Address/Subnet	Do Not Authenticate	Trace
Client IP Address/Subnet	Destination Host/Port	Deny	
Client Hostname	URL	Authenticate	
Proxy IP Address/Port	Category	Force Authenticate	
User Agent	Combined Objects		
Request Header			
Combined Objects			

Web Access Policy Layer Reference

The following table provides the objects available in the Web Access policy layer.

Web Access policy layers regulate, from a general to a granular level, who or what can access specific Web locations or content.

- Users, groups, individual IP addresses, and subnets, as well as object lists comprised of any combination of these, can be subject to rules.
- Rules can include access control for specific Web sites, specific content from any Web site, individual IP addresses, and subnets.
- Actions taken can range from allowing and denying access to more finely tuned changes or limitations.
- Rules can also be subject to day and time specifications and protocol, file type, and agent delimiters.

Source Objects	Destination Objects	Service Objects	Time Objects	Action Objects	Track Objects
Streaming Client	Destination IP Address/Subnet	Using HTTP Transparent Authentication	Time	Allow	Event Log
Client Hostname Unavailable	Destination Host/Port	Virus Detected	Combined Objects	Deny	Email
Authenticated User	URL	Client Protocol		Force Deny	SNMP
Client IP Address/Subnet	Category	Protocol Methods		Bypass Cache	Trace

Section B: Policy Layer and Rule Object Reference

Source Objects	Destination Objects	Service Objects	Time Objects	Action Objects	Track Objects
Client Hostname	File Extensions	IM File Transfer		Do Not Bypass Cache	Combined Objects
Proxy IP Address/Port	HTTP MIME Types	IM Message Text		Check/Do Not Check Authorization	
User	Response Code	IM Message Reflection		Always Verify	
Group	Response Header	Streaming Content Type		Use Default Verification	
Attribute	IM Buddy	ICAP Error Code		Block/Do Not Block PopUp Ads	
User Agent	IM Chat Room	Combined Objects		Force/Do Not Force NTLM for Server Auth	
IM User Agent	Combined Objects			Reflect/Do Not Reflect IM Messages	
Request Header				Block/Do Not Block IM Encryption	
SOCKS Version				Return Exception	
IM User				Return Redirect	
P2P Client				Send IM Alert	
Client Negotiated Cipher				Modify Access Logging	
Client Negotiated Cipher Strength				Override Access Log Field	
Combined Objects				Rewrite Host	
				Reflect IP	
				Suppress Header	
				Control Request Header/Control Response Header	

Section B: Policy Layer and Rule Object Reference

Source Objects	Destination Objects	Service Objects	Time Objects	Action Objects	Track Objects
				Notify User Strip Active Content Set Client HTTP Compression Set Server HTTP Compression Set SOCKS Compression Manage Bandwidth Modify IM Message Return ICAP Patience Page Set External Filter Service Set ICAP Request Service Set FTP Connection Set SOCKS Acceleration Set Streaming Max Bitrate Combined Objects	

Section B: Policy Layer and Rule Object Reference

Web Content Policy Layer Reference

The following table provides the objects available in the Web Content policy layer.

The Web Content policy layer applies to requests independent of user identity.

Content scanning policy layers scan requested URLs and file types for viruses and other malicious code. You must have an ICAP service installed on the ProxySG to use this policy type.

Destination Objects	Action Objects	Track Objects
Destination IP Address/Subnet	Check/Do Not Check Authorization	Event Log
Destination Host/Port	Always Verify	
URL	Use Default Verification	Email
Category	Use Default Caching	SNMP
File Extensions	Do Not Cache	Trace
HTTP MIME Types	Force Cache	Combined Objects
Combined Objects	Mark/Do Not Mark As Advertisement	
	Enable/Disable Pipelining	
	Set Dynamic Categorization	
	Set External Filter Service	
	Set Client HTTP Compression	
	Set Server HTTP Compression	
	Manage Bandwidth	
	Set ICAP Request Service	
	Set ICAP Response Service	
	Set TTL	
	Modify Access Logging	
	Override Access Log Field	
	Combined Objects	

Section B: Policy Layer and Rule Object Reference

Forwarding Policy Layer Reference

The following table provides the objects available in the Forwarding policy layer.

Source Objects	Destination Objects	Service Objects	Action Objects	Track Objects
Streaming Client	Destination IP Address/Subnet	Client Protocol	Send Direct	Trace
Authenticated User	Destination Host/Port	Combined Objects	Integrate/Do Not Integrate New Hosts	
Client IP Address/Subnet	URL		Allow Content From Origin Server	
Client Hostname	Combined Objects		Serve Content Only From Cache	
Proxy IP Address/Port			Select SOCKS Gateway	
User			Select Forwarding	
Group			Reflect IP	
Attribute			Manage Bandwidth	
SOCKS Version			Set IM Transport	
P2P Client			Set SOCKS Gateway Compression	
Combined Objects			Set Streaming Transport	
			Combined Objects	

Section C: Detailed Object Column Reference

This section contains the following topics:

- "Source Column Object Reference"
- "Destination Column Object Reference"
- "Service Column Object Reference"
- "Time Column Object Reference"
- "Action Column Object Reference"
- "Track Object Column Reference"
- "Comment Object Reference"
- "Using Combined Objects"
- "Creating Categories"

Section C: Detailed Object Column Reference

Source Column Object Reference

A *source* object specifies the communication or Web transaction origin that is evaluated by the policy. Not all policy layers contain the same source objects.

Any

Applies to any source.

Streaming Client

This is a static object. This rule applies to any request from a streaming client.

Client Hostname Unavailable

This is a static object. This rule applies if the client IP address could not be looked up with a reverse DNS query.

Authenticated User

This is a static object. This rule applies to any authenticated user.

Client IP Address/Subnet

Specifies the IP address and, optionally, a subnet mask of a client. The policy defined in this rule applies only to this address or addresses on this subnet. This object is automatically named using the prefix *Client*; for example, *Client: 1.2.0.0/255.255.0.0*.

Client Hostname

Specifies a reverse DNS hostname resolved in the reverse lookup of a client IP address. Enter the host name and select matching criteria. This object is automatically named using the prefix *Client*; for example, *Client: host.com*. If you select a matching qualifier, that attribute is appended to the object in parentheses. For example, *Client: host.com (RegEx)*.

Proxy IP Address/Port

Specifies the IP address and, optionally, a port on the *ProxySG*. The policy defined in this rule applies only to this address or addresses with this subnet.

User

Specifies an individual user in the form of a verifiable username or login name. Enter a user name and an authentication realm. The dialog then displays different information depending on the type of authentication realm specified. Select the appropriate realm from the drop-down list. Items in the list are taken from the realms configured by the administrator in the *ProxySG*.

Section C: Detailed Object Column Reference

LDAP

You can optionally select a User Base DN from a drop-down list. Entries in the User Base DN list come from those specified by the administrator in the ProxySG. You can also edit an entry selected in the list, or type a new one. Edited names and new names are retained in the list. Notice in the Full Name field that VPM takes the User Attribute type specified by the administrator in the ProxySG (*cn=* in the following illustration), and associates it with the user name and Base DN entered here.

Important: When you configure a realm, the ProxySG assumes a default primary user attribute (*sAMAccountName* for Active Directory; *uid* for Netscape/iPlanet Directory Server/SunOne; *cn* for Novel NDS). You can accept the default or change it. Whatever is entered there is what VPM uses here, entering it in the Full Name display field once a Base DN is selected.

If the primary user attribute specified in the ProxySG differs from the primary user attribute specified in the directory server, enter the latter in the User field with the appropriate value (in the format *attribute=value*). This replaces the entry in the Full Name field. Examine the following screenshot. Assume that the organization uses *phone* as the primary attribute in its LDAP directory:

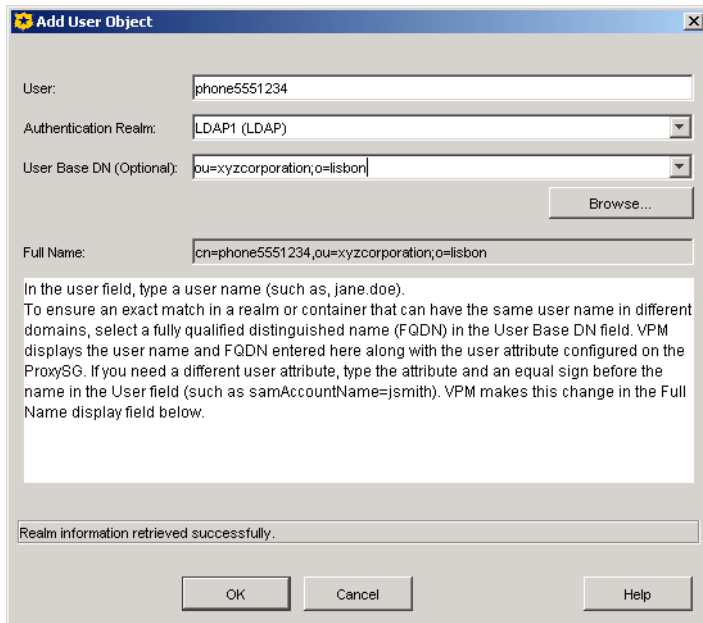


Figure 14-9: Specifying an LDAP Primary User Attribute

You can only enter a user attribute and equal sign in the User field if a User Base DN is selected.

NTLM

Entries in this list are not prepopulated. You must enter a name in the Domain Name field. An entered name is retained and can subsequently be selected and edited. Notice in the Full Name field that VPM displays domain name and user name entered above.

Section C: Detailed Object Column Reference

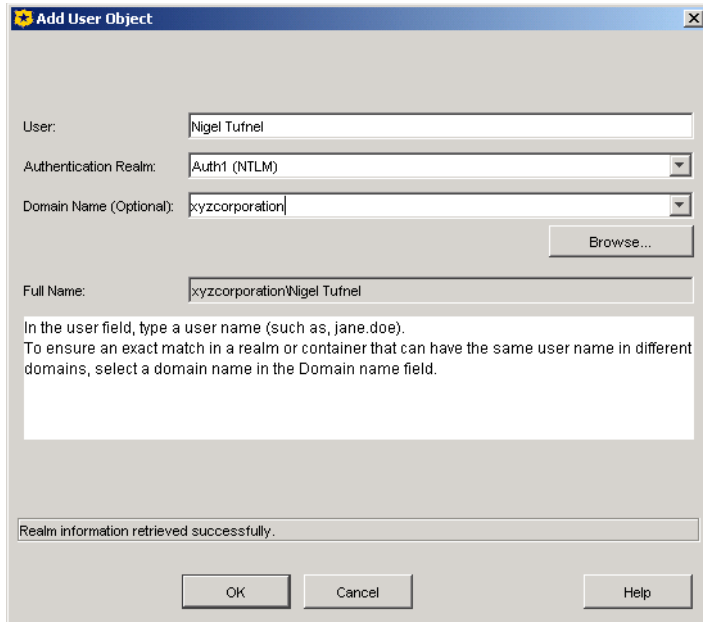


Figure 14-10: Adding an NTLM User with a FQDN or DN

RADIUS

Entries in this list are not prepopulated. You must enter a name in the User field. An entered name is retained and can subsequently be selected and edited. Notice in the Full Name field that VPM displays domain name and user name entered above.

Local

Entries in this list are not prepopulated. You must enter a name in the User field. An entered name is retained and can subsequently be selected and edited. Notice in the Full Name field that VPM displays domain name and user name entered above.

Certificate

If a Certificate realm is selected and that realm uses an LDAP realm as authentication realm, the Browse button is clickable. This option allows you to browse the LDAP database and select users, thus preventing typing errors possible from manually entering names in the fields. If the Certificate realm does not use an LDAP authentication realm, Browse is not displayed.

Netegrity SiteMinder

Entries in this list are not prepopulated. You must enter a name in the User field. An entered name is retained and can subsequently be selected and edited. Notice in the Full Name field that VPM displays domain name and user name entered above.

Section C: Detailed Object Column Reference

Obliv COREid

Entries in this list are not prepopulated. You must enter a name in the User field. An entered name is retained and can subsequently be selected and edited. Notice in the Full Name field that VPM displays domain name and user name entered above.

Policy Substitution

Entries in this list are not prepopulated. You must enter a name in the User field. An entered name is retained and can subsequently be selected and edited. Notice in the Full Name field that VPM displays domain name and user name entered above.

Sequences

Entries in this list are not prepopulated. You must enter a name in the User field. An entered name is retained and can subsequently be selected and edited. Notice in the Full Name field that VPM displays domain name and user name entered above. From the Member Realm drop-down list, select an authentication realm (already configured on the ProxySG). Depending on the realm type, new fields appear.

Group

Specifies a verifiable group name. Enter a user group and an authentication realm. The dialog then displays different information depending on the type of authentication realm specified.

- Group field—Replace the default with a verifiable group name.
- Authentication Realm field—Select the appropriate realm from the drop-down list. Items in the list are taken from the realms configured by the administrator in the ProxySG.
 - LDAP—Entries in the Group Base DN list come from those specified by the administrator in the ProxySG. You can also edit an entry selected in the list, or type a new one. Edited names and new names are retained in the list. Notice in the Full Name field that VPM takes the User Attribute type specified by the administrator in the ProxySG (cn= in the following illustration), and conjoins it with the group name and Base DN entered here.

Important: When you create a group, the default attribute is cn= in the Full Name display field.

Section C: Detailed Object Column Reference

Add Group Object

Group:

Authentication Realm:

Group Base DN (Required):

Full Name:

In the Group Base DN field, select a DN from the pick list. You can also edit a DN from the list or type a DN in the field. In the group field, replace the default entry with a name for the group (such as , sales). VPM displays the groupname and FQDN entered here along with the probable user attribute. If you need a different user attribute, type the attribute and an equal sign before the name in the Group field (such as samAccountName=jsmith). VPM makes this change in the Full Name display field below.

Realm information retrieved successfully.

Figure 14-11: Creating an LDAP Group Object

If the primary user attribute specified in the ProxySG differs from the primary user attribute specified in the directory server, you need to enter the latter here. Do that by typing it in the Group field with the appropriate value (in the format attribute=value). Doing so replaces the entry in the Full Name field. Unlike the comparable situation when creating a user (described immediately above), when creating a group, the Group Base DN does not need to be selected in order to type the attribute=value pair in the Group field.

- NTLM—Entries in this list are not prepopulated. You must enter a name in the Domain Name field. A name typed in is retained and can subsequently be selected and edited. Notice in the Full Name field that VPM displays the domain name and group name entered above.

Section C: Detailed Object Column Reference

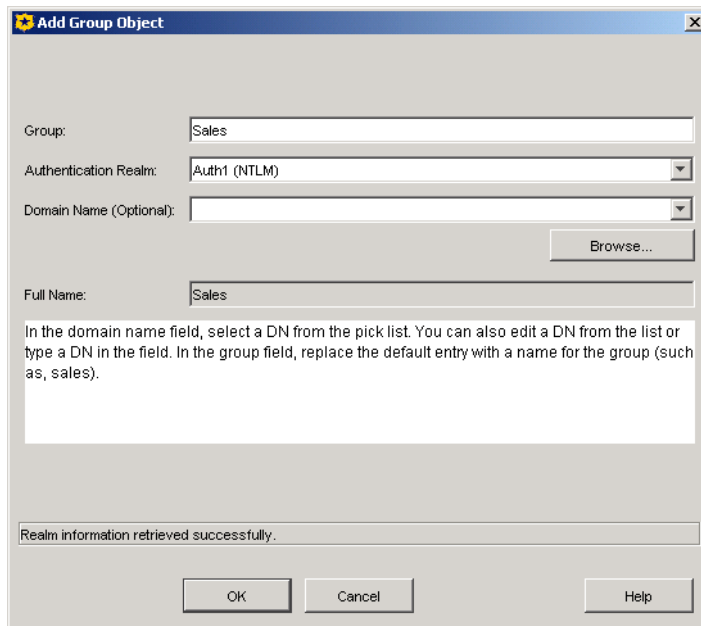


Figure 14-12: Creating an NTLM Group Object

- ❑ Local—Entries in this list are not prepopulated. You must enter a name in the Group field. A name typed in is retained and can subsequently be selected and edited. Notice in the Full Name field that VPM displays the group name entered above.
- ❑ Certificate—If a Certificate realm is selected and that realm uses an LDAP realm as authentication realm, the Browse button is clickable. This option allows you to browse the LDAP database and select users, thus preventing typing errors possible from manually entering names in the fields. If the Certificate realm does not use an LDAP authentication realm, Browse is not displayed.
- ❑ Netegrity SiteMinder—Entries in this list are not prepopulated. You must enter a name in the Group field. A name typed in is retained and can subsequently be selected and edited. Notice in the Full Name field that VPM displays the group name entered above.
- ❑ Oblix COREid—Entries in this list are not prepopulated. You must enter a name in the Group field. A name typed in is retained and can subsequently be selected and edited. Notice in the Full Name field that VPM displays the group name entered above.
- ❑ Policy Substitution—Entries in this list are not prepopulated. You must enter a name in the Group field. A name typed in is retained and can subsequently be selected and edited. Notice in the Full Name field that VPM displays the group name entered above.
- ❑ Sequences—Entries in this list are not prepopulated. You must enter a name in the Group field. An entered name is retained and can subsequently be selected and edited. Notice in the Full Name field that VPM displays domain name and user name entered above. From the Member Realm drop-down list, select an authentication realm (already configured on the ProxySG). Depending on the realm type, new fields appear.

Section C: Detailed Object Column Reference

Attribute

Specifies an LDAP or Radius realm attribute or service.

LDAP

Specifies a specific LDAP attribute (and optional value).

To Specify an LDAP Attribute:

1. In the Name field, enter a name for the object or leave as is to accept the default.
2. Select All LDAP or a specific realm.
3. Enter an Attribute Type.
4. Enter an Attribute Value, or leave blank to accept any value.

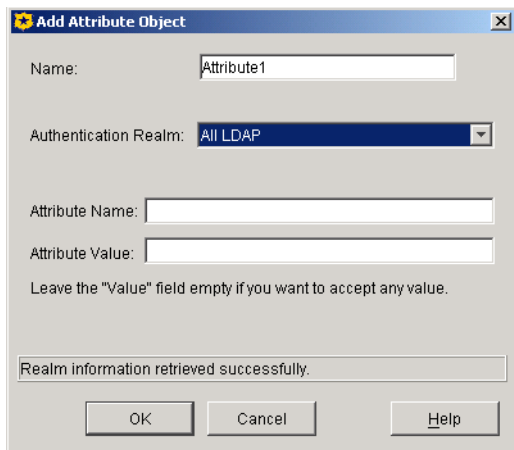


Figure 14-13: Specifying an LDAP Attribute

RADIUS

Specifies a RADIUS attribute.

To Specify a RADIUS Attribute:

Select from a drop-down list of available services.

1. In the Name field, enter a name for the object or leave as is to accept the default.
2. Select All RADIUS or a specific realm.
3. Select a RADIUS Attribute.

Section C: Detailed Object Column Reference

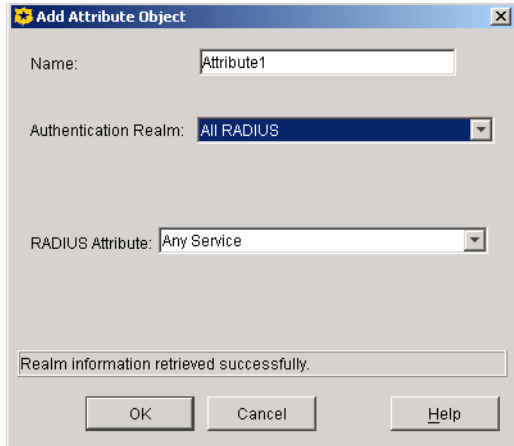


Figure 14-14: Specifying a RADIUS Attribute

DNS Request Name

Specifies a DNS request. Enter the host name and select matching criteria. This object is automatically named using the prefix DNS; for example, DNS: host.com. If you select a matching qualifier, that attribute is appended to the object in parentheses. For example, DNS: host.com (RegEx).

RDNS Request IP Address/Subnet

Specifies the reverse DNS IP address and, optionally, a subnet mask. The policy defined in this rule applies only to this address or addresses on this subnet. This object is automatically named using the prefix RDNS; for example, RDNS: 5.6.0.0/255.255.0.0.

DNS Request Opcode

Specifies OPCODEs to represent in the DNS header.

To Specify a DNS Request OPCODE Object:

1. In the Name field, enter a custom name or leave as is to accept the default.
2. Select one or more of the OPCODEs.
3. Click OK.

DNS Request Class

Specifies the DNS request class (QCLASS) properties.

To Specify a DNS Request Class Object:

1. In the Name field, enter a custom name or leave as is to accept the default.
2. Select one or more of the request classes.

Section C: Detailed Object Column Reference

3. Click OK.

DNS Request Type

Specifies the DNS request types (QTYPE) attributes.

To Specify a DNS Request Type Object:

1. In the Name field, enter a custom name or leave as is to accept the default.
2. Select one or more of the request types.
3. Click OK.

DNS Client Transport

Specifies the DNS client transport method, UDP or TCP.

To Specify a DNS Client Transport Object:

1. Select UDP Transport or TCP Transport. This object is automatically named using the prefix DNS; for example, DNS: Client Transport UDP.
2. Click OK.

SOCKS Version

Specifies the SOCKS version, 4 or 5. This object is automatically named as SOCKSVersion4 or SOCKSVersion5.

User Agent

Specifies one or more agents a client might use to request content. The choices include specific versions of: Microsoft Internet Explorer, Netscape Communicator, Microsoft Windows Media Player and NetShow, Real Media RealPlayer and RealDownload, Apple QuickTime, Opera, and Wget.

The policy defined in this rule applies to these selected agents. You can name this list and create other custom lists to use with other policy layer rules.

Section C: Detailed Object Column Reference

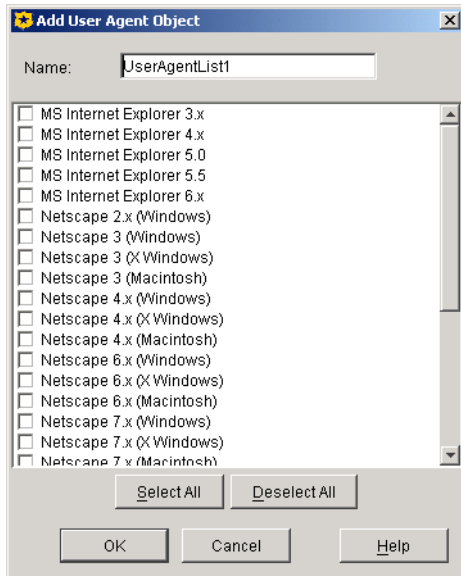


Figure 14-15: Selecting User Agents

Note: If you require a user agent not contained in this list, use the Request Header object, which can contain user agent specified as a header.

IM User Agent

Checks the specified string for a match in the user agent provided by the IM client. For example, specify the string Lotus to distinguish between the Lotus AOL client and the standard AOL client.

To Specify a Header:

1. In the IM User Agent field, enter a string.
2. From the drop-down list, select a matching criteria.
3. Click Add.

Request Header

Specifies the rule applies to requests containing a specific header. Blue Coat supplies a list of standard headers, but you can also select a custom header.

To Specify a Header:

1. In the Name field, enter a custom name or leave as is to accept the default.
2. From the Show drop-list select the viewing field from All to Standard or Custom, as desired. Standard displays only the default standard headers. Custom displays any admin-defined headers that exist.

Section C: Detailed Object Column Reference

3. From the Header Name drop-list, select a standard or custom header or enter a new custom header name.
4. In the Header Regex field, enter the header values to which this rule applies.

Example

An object named CorporateHeader with client IP address 10.1.1.1.

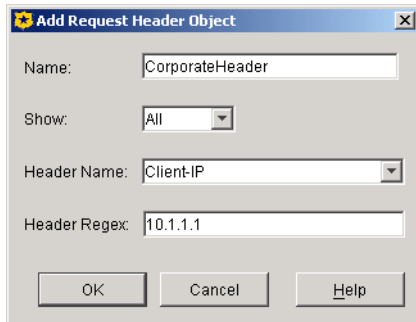


Figure 14-16: Specifying a Header

IM User

Specifies an IM user by their handle. IM traffic sent to or from this user is subject to this rule. You can enter a complete User ID, a string that is part of a User ID, or a string with a regular expression. Select the match type from the drop-down list to the right (Exact, Contains, or RegEx).




Figure 14-17: Specifying an IM User

P2P Client

Specifies peer-to-peer (P2P) clients.

To Specify P2P Clients:

1. In the Name field, enter a name for the object or accept the default.
2. Select All P2P Clients (all protocols become selected), or one or more P2P protocols.
3. Click OK.

Section C: Detailed Object Column Reference

Client Negotiated Cipher

Specifies a cipher code. Select a code from the drop-down list. This object is automatically named using the prefix Cipher; for example, Cipher: RC4-MD5.

Client Negotiated Cipher Strength

Specifies the cipher strength. Select Export, Low, Medium, or High from the drop-down list. This object is automatically named using the prefix Cipher Strength; for example, Cipher Strength: Medium.

Combined Source Object

Allows you to create an object that combines different source types. Refer to "Using Combined Objects" on page 530.

Source Column/Policy Layer Matrix

The following matrix lists all of the Source column objects and indicates which policy layer they apply to.

Object	Admin Auth	Admin Access	DNS Access	SOCKS Auth	Web Auth	Web Access	Web Content	Forwarding
Streaming Client						x		
Client Hostname Unavailable					x	x		
Authenticated User						x		x
Client IP Address/Subnet	x	x	x	x	x	x		x
Client Hostname	x			x	x	x		x
Proxy IP Address/Port	x	x	x	x	x	x		x
User		x				x		x
Group		x				x		x
Attribute		x				x		x
DNS Request Name			x					
RDNS Request IP Address/Subnet			x					
DNS Request Opcode			x					
DNS Request Class			x					
DNS Request Type			x					
DNS Client Transport			x					

Section C: Detailed Object Column Reference

Object	Admin Auth	Admin Access	DNS Access	SOCKS Auth	Web Auth	Web Access	Web Content	Forwarding
SOCKS Version				x		x		x
User Agent					x	x		
IM User Agent						x		
Request Header					x	x		
IM User						x		
P2P Client						x		x
Client Negotiated Cipher						x		
Client Negotiated Cipher Strength						x		
Combined Objects	x	x	x	x	x	x		x

Destination Column Object Reference

A *destination* object specifies the communication or Web traffic destination that is evaluated by the policy. Not all policy layers contain the same destination objects.

Any

Applies to any destination.

DNS Response Contains No Data

This is a static object.

Destination IP Address/Subnet

Specifies the IP address and, optionally, a subnet mask of a destination server. The policy defined in this rule only applies to this address only or addresses within this subnet. This object is automatically named using the prefix *Destination*; for example, *Destination: 1.2.0.0/255.255.0.0*.

Destination Host/Port

Specifies the hostname or port of a destination server. The policy defined in this rule applies to this host on this port only. Enter the host name and port number, and select matching criteria. This object is automatically named using the prefix *Destination*; for example, *Destination: company.com:80*.

URL

Specifies a URL entered by a user.

Section C: Detailed Object Column Reference

To Specify a URL:

Select a radio button and enter the required information in the fields:

- **Simple Match**—Matches a partial URL. If a host name is specified, all hosts in that domain or subdomain match; if a path is specified, all paths with that path prefix match; if a scheme or port number is specified, only URLs with that scheme or port match. This object is automatically named using the prefix URL; therefore, the object is displayed as URL: host.com.
- **Regular Expression Match**—Specifies a regular expression. This object is automatically named using the prefix URL; therefore, the object is displayed as URL: host.com (RegEx).
- **Advanced Match**—Specifies a scheme (protocol), host, port range, and/or path. Enter a name at the top of the dialog to name the object. With host and path, you can select from the drop-down list to match exactly as entered or parts thereof: Exact Match, Contains, At Beginning, At End, or RegEx. If you select a matching qualifier, that attribute is appended to the object in parentheses. For example, URL: host.com (Contains).

Figure 14-18: Specifying Destination URLs

Category

Allows you to create and customize categories of URLs.

- **Policy**—Displays all current pre-defined and user created URL categories. This includes all category-related configurations created in the VPM, as well as in the Local and Central policy files (once installed). Select and deselect categories as required.

Section C: Detailed Object Column Reference

You can also create new categories from this dialog, which is similar to the dialog accessed through the VPM Menu Bar as described in "Creating Categories" on page 536.

If you enable a service, such a content filter, those relevant categories appear in this object.

- System—Displays hard-coded ProxySG configurations. These are not editable, but you can select or deselect them.

To create a policy category:


1. Select Policy; click Add. The Object Name dialog appears.
2. Name the category and click OK.
3. Drop the Policy list and select the created category; click Edit URLs. The Edit Locally Defined Category Object dialog appears.
4. Enter URLs appropriate for the content filter category you are creating; click OK.
5. Click OK.

Note: If one or more other administrators have access to the ProxySG through other workstations and are creating categories either through VPM or with `inline` commands, consider that newly-created or edited categories are not synchronized until the policy is installed. When the policy is installed with VPM, the categories are refreshed. If confusion occurs, select the File>Revert to Existing Policy on ProxySG Appliance option to restore the policy to the previous state and reconfigure categories.

Category Hierarchy Behavior


Once categories have been created, they can be selected and deselected as required. If you create sub-categories (a parent and child category hierarchy), the category selection behavior is the following:

- Selecting a parent category automatically selects all child categories if no child categories are already selected.



- Deselecting a parent category automatically deselects all child categories if all child categories are already selected.
- If one or more of the child categories are already selected or deselected, selecting or deselecting the parent category does *not* affect child categories—the status of selected or deselected remains the same.

Section C: Detailed Object Column Reference



This behavior applies to as many levels as you create.

File Extensions

Creates a list of file extensions. The rule is triggered for content with an extension matching any on the list. You can create multiple lists that contain various extensions to use in different rules. For example, create a list called Pictures, and select file extension types such as GIF, JPEG, BMP, XPM, and so on.

HTTP MIME Types

Creates a list of HTTP MIME content types. The rule is triggered for content matching any on the list. You can create multiple lists that contain various MIME types to use in different rules. For example, create a list called MicrosoftApps, and select MIME types application/vnd.ms-excel, application/vnd.ms-powerpoint, application/vnd.ms-project, and application/vnd.works.

Note: If you require a MIME type not contained in this list, use a URL object that uses the At End matching criteria.

Response Code

Specifies the rule applies to content responses containing a specific HTTP code. Select a code from the drop-down list. You can name the rule object or accept the default name.




Figure 14-19: Specifying a Response Code

Response Header

Specifies the rule applies to content responses containing a specific header. Blue Coat supplies a list of standard headers, but you can also enter a custom header.

Section C: Detailed Object Column Reference

To Specify a Header:

1. In the Name field, enter a custom name or leave as is to accept the default.
2. From the Show drop-down list select the viewing field from All to Standard or Custom, as desired. Standard displays only the default standard headers. Custom displays any admin-defined headers that exist.
3. From the Header Name drop-down list, select a standard or custom header.
4. In the Header Regex field, enter the header string this rule applies to.

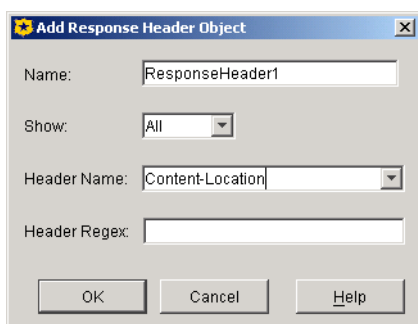


Figure 14-20: Specifying a Response Header

IM Buddy

Specifies an IM buddy by their handle. IM traffic sent to or from this buddy is subject to this rule. You can enter a complete buddy ID, a string that is part of a buddy ID, or a string with a regular expression. Select the match type from the drop-down list to the right (Exact, Contains, or RegEx).

IM Chat Room

Specifies an IM chat room by name or other triggers. IM traffic sent to this chat room is subject to this rule.

To Create a Chat Room Trigger:

1. In the Name field, enter a name for the object or leave as is to accept the default.
2. Select one or more of the following triggers:
 - Room ID—Specifies a specific IM chat room by its name. Enter a name and from the drop-down list select an option: Exact Match, Contains, or RegEx.
 - Type—Specifies the type of room. Select Private or Public.
 - Invite Only—Specifies to trigger if user must be invited or not.
 - Voice Enabled—Specifies whether room supports voice chat or not.
 - Conference—Specifies whether room is a conference or not.
3. Click OK.

Section C: Detailed Object Column Reference

Example

An object named ChatRoom1 that triggers the rule if the room is private.

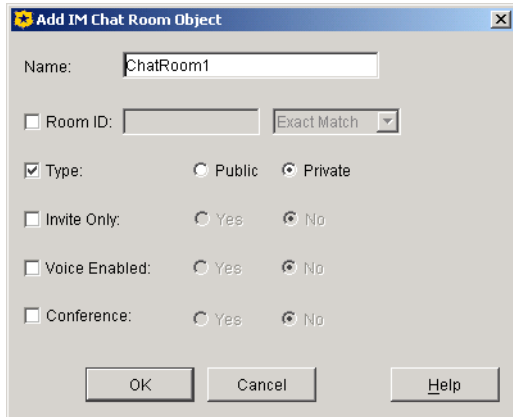


Figure 14-21: Chat Room Object

DNS Response IP Address/Subnet

Specifies the destination DNS IP address and, optionally, a subnet mask. The policy defined in this rule only applies to DNS responses containing this address or addresses within this subnet. This object is automatically named using the prefix DNS; for example, DNS: 1.2.3.4/255.255.0.0.

RDNS Response Host

Specifies a reverse DNS response hostname resolved in the reverse lookup of a client IP address. Enter the host name and select matching criteria. This object is automatically named using the prefix RDNS; for example, RDNS: host.com. If you select a matching qualifier, that attribute is appended to the object in parentheses. For example, RDNS: host.com (RegEx).

DNS Response CNAME

Specifies the rule applies to DNS CNAME responses matching a given hostname. Enter the host name and select matching criteria. This object is automatically named using the prefix DNS CNAME; therefore, the object is displayed as DNS CNAME: host.com.

DNS Response Code

Specifies the rule applies to DNS responses containing a specific DNS Response code. Select one or more codes from the list. You can name the rule object or accept the default name.

Combined Destination Objects

Allows you to create an object that combines different destination types. Refer to "Using Combined Objects" on page 530.

Section C: Detailed Object Column Reference

Destination Column/Policy Layer Matrix

The following matrix lists all of the Destination column objects and indicates which policy layer they apply to.

Object	Admin Auth	Admin Access	DNS Access	SOCKS Auth	Web Auth	Web Access	Web Content	Forwarding
Destination IP Address/Subnet					x	x	x	x
Destination Port					x	x	x	x
URL					x	x	x	x
Category			x		x	x	x	
File Extensions						x	x	
HTTP MIME Types						x	x	
Response Header						x		
Response Code						x		
IM Buddy						x		
IM Chat Room						x		
DNS Response IP Address/Subnet			x					
RDNS Response Host			x					
DNS Response CNAME			x					
DNS Response Code			x					
Combined Objects			x		x	x	x	x

Service Column Object Reference

A *service* object specifies a service type, such as a protocol, that is evaluated by the policy. Not all policy layers contain the same service objects.

Any

Applies to any service.

Using HTTP Transparent Authentication

This is a static object. The rule applies if the service is using HTTP transparent authentication.

Section C: Detailed Object Column Reference

Virus Detected

This is a static object. The rule applies if ICAP scanning detects a virus.

Client Protocol

Specifies the client protocol types and subsets. From the first drop-down list, select a type from the drop-down list: Endpoint Mapper, FTP, HTTP, HTTPS, Instant Messaging, P2P, Shell, SOCKS, Streaming, or TCP Tunneling.

The second drop-down list allows you to select a protocol subset:

- All—Applies to all communication using this type of protocol.
- Pure—Applies if the protocol is using a direct connection.
- Over—Applies if a protocol is communicating through a specific transport method.

For example, you can select the Instant Messaging protocol and select Pure Yahoo Instant Messaging.

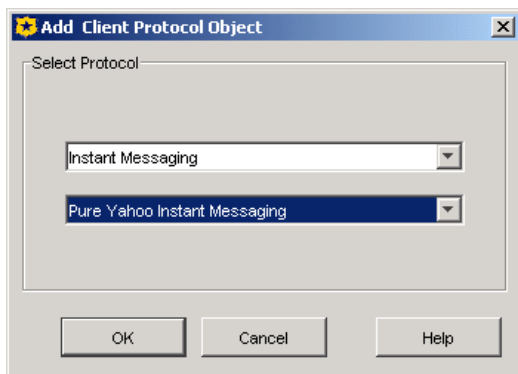


Figure 14-22: Specifying Client Protocol Type

Protocol Methods

Specifies the protocol methods. Select a protocol from the drop-down list: FTP, HTTP, HTTPS, Instant Messaging, SOCKS. Next, select each specific method.

Section C: Detailed Object Column Reference

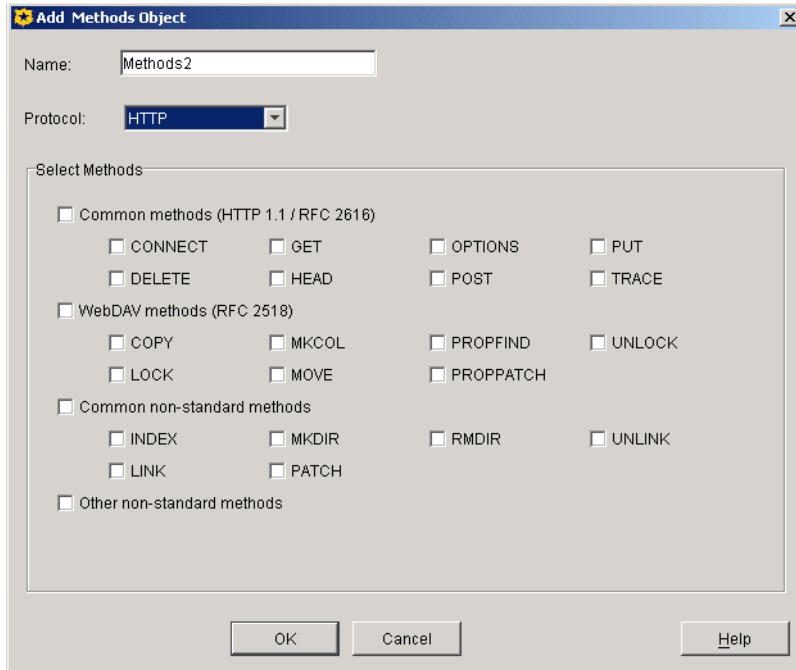


Figure 14-23: Specifying HTTP Protocol Methods

IM File Transfer

Specifies the rule is applied to IM file transfers, which can be triggered by matching the file name, file size, or both.

To Specify IM File Transfer Parameters:

1. In the Name field, enter a name for the object or accept the default.
2. To trigger by file name, select File. In the File field, specify a file name; from the drop-down list, select if file is matched exactly (Exact Match), if the file contains the name (Contains), or matched by regular expression (RegEx).
3. To trigger by message size, select Size. Enter a range; from the drop-down list, select the size attribute: bytes, kilobytes, megabytes, or gigabytes.

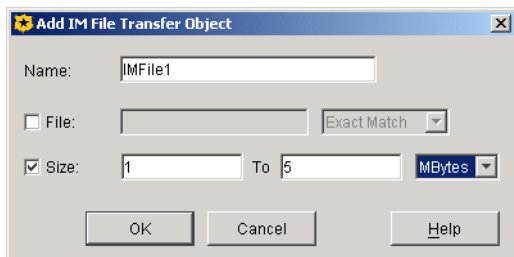


Figure 14-24: Limiting IM File Transfers

Section C: Detailed Object Column Reference

IM Message Text

Specifies the rule is applied to IM message text, which can be triggered by any or all of the following: matching content keywords, message size, service type, and whether the content type is text or application.

To Specify IM Message Text Parameters:

1. In the Name field, enter a name for the object or accept the default.
2. To trigger by content keywords, select Text. In the Text field, specify a keyword; from the drop-down list, select if the file contains the text (Contains), or if it is to be matched by regular expression (RegEx).
3. To trigger by message size, select Size. Enter a range; from the drop-down list, select the size attribute: bytes, kilobytes, megabytes, or gigabytes.
4. To specify the message route, select Route. From the drop-down list, select Service, Direct, or Chat.
5. To specify message type, select Text or Application.

Text specifies messages entered by a user; Application specifies messages sent by the client application, such as typing notifications.

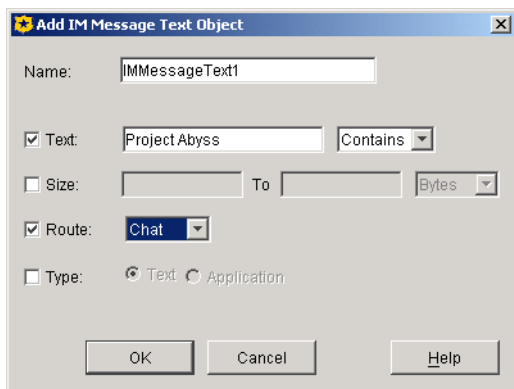


Figure 14-25: Triggering IM Rule with Text and Route Type

IM Message Reflection

Allows policy to test whether or not reflection is enabled for the current message and, if enabled, whether it is possible.

- Succeeded—IM reflection is enabled and is performed for this message.
- Failed—IM reflection is enabled, but not possible for this message because the recipient is not connected through this ProxySG.
- Disabled—IM reflection is not enabled for this message.

The objects are automatically named based on the selection and can be used in any rule.

Section C: Detailed Object Column Reference

Streaming Content Type

Specifies streaming protocols.

To Specify Streaming Protocols:

1. In the Name field, enter a name for the object or accept the default.
2. Select All Streaming Content (all protocols become selected), or one or more streaming protocols.
3. Click OK.

ICAP Error Code

Defines an object that recognizes one or more ICAP error codes returned during an antivirus scan. The rule applies if the scan returns the specified errors.

To Specify ICAP Error Codes:

1. In the Name field, enter a name for the object or accept the default.
2. Select an option:
 - No errors—An ICAP scan was performed without scanning errors.
 - Any errors—An ICAP error code was returned during a scan.
 - Selected errors—An ICAP error code of a specific type or types.

In the Available Errors field, select one or more ICAP error codes (press and hold the Control key to select more than one type or the Shift key to select a block of types). Click Add.

3. Click OK.

Combined Service Objects

Allows you to create an object that combines different service types. Refer to "Using Combined Objects" on page 530.

Service Column/Policy Layer Matrix

The following matrix lists all of the Service column objects and indicates which policy layer they apply to.

Object	Admin Auth	Admin Access	DNS Access	SOCKS Auth	Web Auth	Web Access	Web Content	Forwarding
Using HTTP Transparent Authentication						x		
Client Protocol						x	x	x
Protocol Methods						x	x	x

Section C: Detailed Object Column Reference

Object	Admin Auth	Admin Access	DNS Access	SOCKS Auth	Web Auth	Web Access	Web Content	Forwarding
IM File Transfer						x		
IM Message Text						x		
IM Message Reflection						x		
Streaming Content Type						x		
ICAP Error Code						x		
Combined Objects						x	x	x

Time Column Object Reference

A *time* object specifies a block of time or time trigger that determines client access regarding other parameters in the rule (such Web sites and content types). Currently, the Time object is only applicable to the Web Access Layer.

Any

Applies anytime.

Time

Specifies the time restrictions.

To Configure Time Restrictions:

1. In the Name field, enter a name for the object or leave to accept the default.
2. Select Use Local Time Zone or Use UTC Time Zone.
Local time sets the rule to follow the ProxySG internal clock. UTC sets the rule to use the Universal Coordinated Time (also known as Greenwich Mean Time or GMT).
3. To specify a range for any given day, select Enable; in the Specify Time of Day Restriction (hh:mm) field, configure the times. The time style is military.
The range can be contained within one 24-hour calendar day, or overlap days. For example, configuring the time to range from 22:00 to 06:00 sets a limit from 10 at night to 6 the following morning.
4. To specify a day of the week restriction, select Enable; in the Specific Weekday Restriction field, select one or more days.
5. To specify a day of the month range restriction, select Enable; in the Specify Day of Month Restriction field, select the days, which are numbered from 01 to 31. To limit the range to specific day, configure the numbers to be the same. For example, selecting 22 and 22 specifies the rule to apply only the 22nd day of every month.

Section C: Detailed Object Column Reference

- To specify a restriction that spans one or more months, select Enable; in the Specify Annually-Recurring Date Restriction field, select the month and day ranges. This calendar restriction applies every year unless the restriction is altered.

Overlapping months is allowed, similar to the behavior of day ranges in Step 3.

- To specify a one-time only restriction, select Enable; in the Specify Non-Recurring Date Restriction field, select the year, month, and day ranges. This calendar restriction applies only during the time specified and will not repeat.

Figure 14-26: Configuring Time Limitations

Combined Time Object

Allows you to combine a time object that adheres to multiple time restrictions. See "Using Combined Objects" on page 530.

Time Column/Policy Layer Matrix

The following matrix lists all of the Time column objects and indicates which policy layer they apply to.

Section C: Detailed Object Column Reference

Object	Admin Auth	Admin Access	DNS Access	SOCKS Auth	Web Auth	Web Access	Web Content	Forwarding
Time			x			x		
Combined Objects			x			x		

Section C: Detailed Object Column Reference

Action Column Object Reference

An *action* object determines which action to take if other parameters, such as source, destination, service, and time requirements validate the rule.

Important: Because of character limitations required by the generated CPL, only alphanumeric, underscore, and dash characters can be used to define an action object name.

Allow

This is a static object. Selecting this overrides other related configurations and the specified user requests are allowed.

Deny

This is a static object. Selecting this overrides other related configurations and denies the specified user requests.

Force Deny

This is a static object. Forces a request to be denied, regardless if rules in subsequent layers would have allowed the request.

Allow Read-Only Access

This is a static object. Grants full access to view data on the appliance.

Allow Read-Write Access

This is a static object. Grants full access to view and manipulate data on the appliance.

Do Not Authenticate

This is a static object. Selecting this overrides other configurations and the specified users are not authenticated when requesting content.

Authenticate

Creates an authentication object to verify users. An authentication realm must exist on the ProxySG to be selected through VPM.

Note: In the SOCKS Authentication policy layer, the object is SOCKS Authenticate.

To Create an Authentication Object:

1. In the Name field, enter a name for the object or leave as is to accept the default.

Section C: Detailed Object Column Reference

2. From the Realm drop-down list, select an authentication realm, which must already exist on the ProxySG.
3. Optional (in the Web Authentication policy layer only): from the Mode drop-down list, select a mode. The mode determines the way the ProxySG interacts with the client for authentication specifying the challenge type and the accepted surrogate credential:
 - ❑ Auto—The default; the mode is automatically selected, based on the request. Selects among proxy, origin-IP, and origin-IP-redirect, depending on the type of connection (explicit or transparent) and the transparent authentication cookie settings.
 - ❑ Form Cookie—For forms-based authentication: 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—The user's IP address is used as a surrogate credential. The form is presented whenever the user's 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.
 - ❑ Proxy—For explicit forward proxies: the ProxySG uses an explicit proxy challenge. No surrogate credentials are used. This is the typical mode for an authenticating explicit proxy.
 - ❑ Proxy IP—The ProxySG uses an explicit proxy challenge and the client's IP address as a surrogate credential.
 - ❑ 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 Cookie—For transparent proxies: for clients that understand cookies but do not understand redirects; the ProxySG acts like an origin server and issues origin server challenges. The surrogate credential is used.
 - ❑ Origin Cookie Redirect—For transparent forward proxies: 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.
 - ❑ Origin IP Redirect—Significantly reduces security; only useful for reverse proxy and when clients have unique IP addresses and do not understand cookies (or you cannot set a cookie). Provides partial control of transparently intercepted HTTPS requests. 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.
 - ❑ SG2—The mode is selected automatically, based on the request using the SGOS 2.x-defined rules.

Section C: Detailed Object Column Reference

- (Optional) If you selected a Form mode in Step 3, the Form drop-down list becomes active. Select
- Click OK.

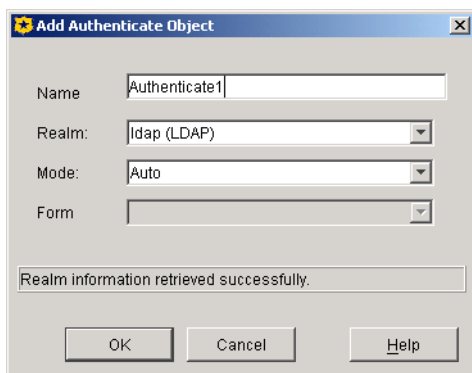


Figure 14-27: Specifying an Authentication Realm

Users are prompted to provide a valid user name and password.

Force Authenticate

Forces the user to authenticate even though the request is going to be denied for reasons that do not depend on authentication. This action is useful to identify a user before the denial so that the username is logged along with the denial.

Note: In the SOCKS Authentication policy layer, the object is Force SOCKS Authenticate.

Bypass Cache

This is a static object. Prevents the cache from being queried when serving a proxy request, and prevents the response from the origin server from being cached.

Do Not Bypass Cache

This is a static object. The ProxySG always checks if the destination is cached before going to the origin server; also, the content is cached if cacheable.

Bypass DNS Cache

This is a static object. Prevents the request from querying the DNS cache list of resolved lookup names or addresses.

Do Not Bypass DNS Cache

This is a static object. The ProxySG always queries the DNS cache list of resolved lookup names or addresses.

Section C: Detailed Object Column Reference

Allow DNS From Upstream Server

This is a static object. Allows the ProxySG to send requests for data not currently cached to DNS servers.

Serve DNS Only From Cache

This is a static object. Instructs the ProxySG to only serve a DNS request from content that is already cached.

Enable/Disable DNS Imputing

These are static objects. If DNS imputing is enabled, the ProxySG appends the suffixes in the DNS imputing list to hostnames looked up when the original name is not found.

Check/Do Not Check Authorization

These are static objects. These actions control whether or not the ProxySG forces a request to be sent to an upstream server every time to check authorization, even if the content is already cached. The check action is not usually required for upstream origin content servers performing authentication, as the ProxySG automatically tracks whether content required authentication in each case. However, it can be required when an upstream proxy is performing proxy authentication because of the way some proxies cache credential information, causing them not to reliably challenge every request. When requests are directed to an upstream proxy which operates in this manner, enabling Check Authorization ensures that all such requests are properly authorized by the upstream proxy before the content is served from the local cache.

Always Verify

This is a static object. Cached content is always verified for freshness for the sources, destinations, or service specified in the rule. For example, the CEO and Executive Staff always require content to be the most recent, but everyone else can be served from the cache.

Use Default Verification

This is a static object. Overrides the Always Verify action and instructs the ProxySG to use its default freshness verification.

Block/Do Not Block PopUp Ads

These are a static objects. Blocks or allows pop up windows. Blue Coat recommends creating separate Web Access policy layers that only contain pop up blocking actions. Furthermore, many Web applications require pop up windows. As it is unlikely that your Intranet contains pages that pop up unwanted advertising windows, Blue Coat recommends disabling pop up blocking for your Intranet. For example:

- Web Access rule 1: Specify the Intranet IP address and subnet mask in the Destination column and select Do Not Block Popup Ads in the Action column.

Section C: Detailed Object Column Reference

- Web Access rule 2: Select Block Popup Ads in the Action column.

As you continue to modify policy, you can add more policy layers to block or allow specific IP addresses, but the policy layer as defined in the Web Access rule 2 above *must* always be positioned last. Blocking pop up ads is the default if a previous policy rule does not trigger.

For more concept information about pop up windows, refer to "Blocking Pop Up Windows" on page 562.

Force/Do Not Force NTLM for Server Auth

These are static objects. When configured for explicit proxy, Internet Explorer (IE) does not support an NTLM challenge from an origin server. If Force NTLM for Server Auth is applied, the ProxySG converts the 401-type server authentication challenge to a 407-type proxy authentication challenge, which IE supports. The ProxySG also converts the resulting Proxy-Authentication headers in client requests to standard server authorization headers, which allows an origin server NTLM authentication challenge to pass through when IE is explicitly proxied through the ProxySG.

Reflect/Do Not Reflect IM Messages

These are static objects. IM traffic can be contained and restricted to the network so that it never reaches the IM server. A hierarchy of ProxySG appliances manage the traffic and routes it depending on each ProxySG fail open and fail closed configurations. For detailed information about this deployment, refer to Chapter 16, Instant Messaging, of the *Blue Coat ProxySG Configuration and Management Guide*.

Block/Do Not Block IM Encryption

These are static objects. AOL IM provides the option for clients to send encrypted messages through both standard messaging (through the IM service) and direct connection messaging. These objects allow you to block or not block the ability to send encrypted messages through AOL IM. For detailed information about security benefits of this feature, see Chapter 16, Instant Messaging, of the *Blue Coat ProxySG Configuration and Management Guide*.

Return Exception

Allows you to select exception types and associate a custom message, if desired. Blue Coat provides a list of standard exceptions, but VPM also accepts user-defined values.

To Create a Return Exception Object:

1. In the Name field, enter a name for the object or leave as is to accept the default.
2. Perform one of the following:
 - Standard exception type: select one from the Built-in exception drop-down list.
 - Custom exception (which already must be defined on the ProxySG) type: select one from the User-defined exception drop-down list.
3. Optional: Select Force exception even if later policy would allow request to

Section C: Detailed Object Column Reference

4. Optional: Select Allow re-authentication to
5. Optional: in the Details field, enter a message that is displayed along with the summary and exception ID on the exception page displayed to the user when the exception is returned.

Example

An object named DNSException2 that upon a DNS server failure returns a message to the user instructing them to contact their support person.

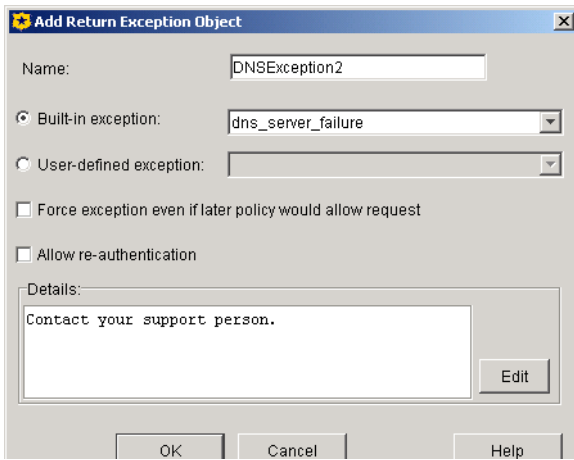


Figure 14-28: Return Exception Object

To create custom exceptions, see "Defining Exceptions" on page 568.

Return Redirect

Aborts the current transaction and forces a client request to redirect to a specified URL. For example, used to redirect clients to a changed URL; or redirecting a request to a generic page stating the Internet access policy. Applies only to HTTP transactions.

Note: Internet Explorer (IE) ignores redirect responses from FTP over HTTP requests, although Netscape Navigator obeys the redirect. To avoid problems with IE, do not use redirect when `url.scheme=ftp`.

If the URL that you are redirecting the browser to also triggers a redirect response from your policy, then you can put the browser into an infinite loop.

In the Name field, enter a name for the object (or leave as is to accept the default); in the URL field, enter the redirect destination URL.

Example

An object that redirects clients to an HTML policy statement page.

Section C: Detailed Object Column Reference




Figure 14-29: Return Redirect Object

Send IM Alert

Defines a message that is sent to an IM user by the ProxySG. The message is triggered by the IM parameters defined in the policy (for example, client login, sent or received messages, and buddy notification). Chapter 17: “Instant Messaging” on page 621 provides more information about regulating IM through the ProxySG, as well as VPM examples.

Example

A message that informs IM users their messaging is logged.




Figure 14-30: Send IM Alert Object

Modify Access Logging

Defines access logging behavior.

- Disable all access logging—No activity is logged for the requests matched by the rule.
- Reset to default logging—Resets to logging the request to the default log specified by the ProxySG configuration, if one exists.
- Enable logging to—Enables logging of requests matched by this rule to the specified log.
- Disable logging to—Disables logging of requests matched by this rule to the specified log.

Section C: Detailed Object Column Reference

Override Access Log Field

Allows you to manipulate access log entries. For any specific log value, you can suppress the value, encode the value in Base64, or rewrite the value.

To Override Access Log Fields:

1. In the Name field, enter a name for the object or leave as is to accept the default.
2. From the Log Name drop-down list, select a log (must already be configured on the ProxySG).
3. From the Field Name drop-down list, select an access log field.
4. Select one of the following:

- Log original value—Records unmodified value in the access log.
- Suppress value—Prevents value from appearing in the access log.
- Base64 encode value—Records an encoded version of the value in the access log.
- Rewrite value—In the field, enter a string that replaces the value.
- Edit—Clicking Edit calls the Select The Rewrite String dialog. The substitution variables instruct the ProxySG to append specific information to the object. The variables are categorized alphabetically, according to prefix.

Note: Some variables do not have prefixes, which allows you to substitute the value with information defined by other field types.

5. Click OK.

Example

An object called CEOLogRewrite that suppresses log entries so persons, such as support personal, cannot view economically sensitive information to gain improper knowledge.

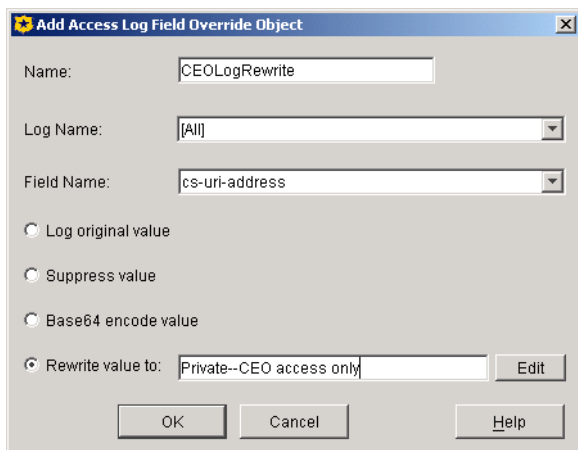


Figure 14-31: Overriding Access Log Fields

Section C: Detailed Object Column Reference

Rewrite Host

Rewrites host component of a URL, specifying either Windows Media, Real Media, or all protocols. Use this to redirect the request to a different host. For example, rewrite `www.foo.com` to `www.bar.com`. You can create and name multiple rewrites, but you can only specify one per rule.

To Specify a Rewrite:

1. In the Name field, enter a name or leave as is to accept to the default.
2. From the Scheme drop-down list, Windows Media, Real Media, or All to rewrite all URLs, regardless of protocol.
3. In the Pattern field, enter a host name (for example, `foo`).
4. In the Replacement field, enter the name the patter is rewritten as (for example, `bar`).
5. Click OK.

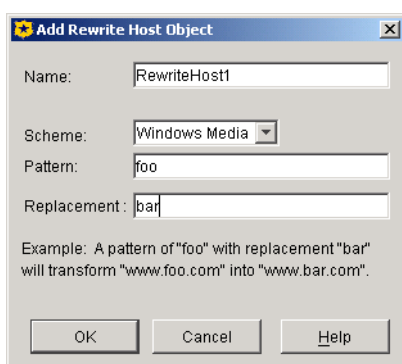


Figure 14-32: Specifying a Host Rewrite

Reflect IP

Specifies which IP address is used when making connections to upstream hosts.

To Create a Reflect IP Object:

1. In the Name field, enter name for the object or leave as is to accept the default.
2. In the In outgoing client IP, reflect field, select one of the following:
 - Do not reflect IP—Disables reflecting IPs; the ProxySG uses the IP address of the interface that request is sent out on.
 - Incoming client IP [IP spoofing]—Reflects the client IP address.
 - Incoming proxy IP—Reflects the IP address of where the request arrived to.
 - Proxy IP—Specifies to reflect a specific IP of the ProxySG; enter the IP address in the field.
 - Use services configuration—Specifies whether to reflect IP in the configuration of the service which is used to process the request.

Section C: Detailed Object Column Reference

3. Click OK.

Example

This object reflects another IP address configured on the ProxySG.

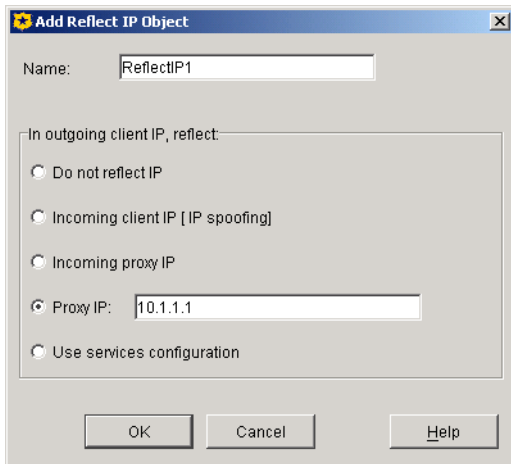


Figure 14-33: Reflect IP Object

Suppress Header

Specifies one or more standard headers that are suppressed (not transmitted) on the outbound request, the outbound response, or both.

To Create a Suppress Header Object:

1. In the Name field, enter name for the object or leave as is to accept the default.
2. Select Request, Response, or Both.
3. Select one or more header types from the list.
4. Click OK.

Example

An object called IntranetHeaders that suppresses headers so specified users can access economically sensitive information without people, such as support personal, being able to gain knowledge of sources.

Section C: Detailed Object Column Reference

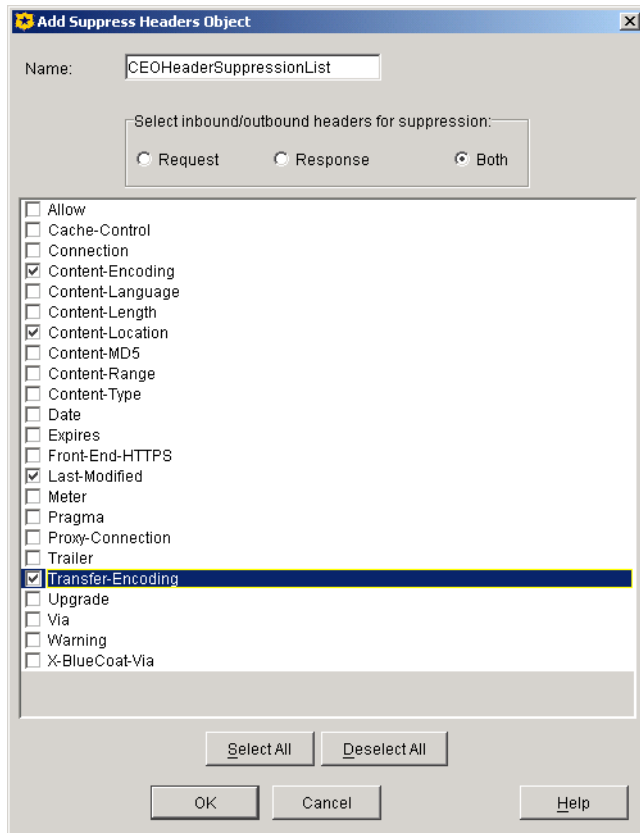


Figure 14-34: Creating a Suppressed Header List Object

"Section C: Modifying Headers" provides more conceptual information about header modification.

Control Request Header/Control Response Header

For any request or response, specifies to:

- Insert a header with a specific value.
- Rewrite the value of a specific header.
- Suppress a specific header.

Section C: Detailed Object Column Reference

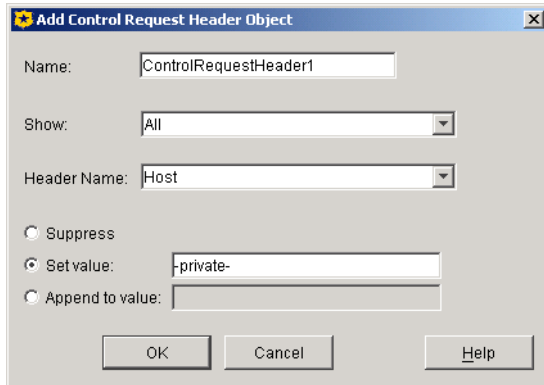


Figure 14-35: Rewriting the Request Header

Notify User

This action displays a notification page in the user's Web browser. A user must read the notification and click an Accept button before being allowed to access the Web content. You can customize the following:

- The page title, notification message, and the Accept button.
- The conditions that cause a notification to be displayed again. By default, the notification is displayed each time a user begins a new Web browsing session (reboots, logs out, or closes all Web browser windows). You can configure re-notification to occur for each new visited host or Web site, or after a time interval.

Note: The Accept button click action is logged if HTTP access logging is enabled. A URL is logged that contains the string: `accepted-NotifyName`, where `NotifyName` is the name of the Notify User object.

This feature is designed to provide the following functionality:

- **Web-use compliance:** A compliance page is a customized notification page displayed on a user's Web browser when attempting to access the Internet. This page ensures employees read and understand the company's Acceptable Use Policy before Internet use is granted. Typically, a compliance notification is displayed each time a browser is opened, but you can configure a time condition to display the page at specific intervals or times of the day, week, or month.
- **Coach users:** A coaching page displays when a user visits a Web site that is blocked by content filtering policy. This page explains why the site is blocked, the consequences of un-authorized access, and a link to the site if business purposes warrants access. A coaching page is configured to display each time a user visits a new Web page that is barred by content filtering policy; however, you can also configure this page to appear at different time intervals.

To configure HTML notification:

1. In the Name field, enter a name for the object or leave as is to accept the default.
2. In the Title field, enter a name that is the title of the page (text only; no HTML is allowed).

Section C: Detailed Object Column Reference

3. In the Body field, compose a block of HTML that displays the message to the user. You can also customize the Accept link or button text. The HTML body must contain an Accept button or link. The default is:

```
<p><a href="$ (exception.details) " onclick="Accept ();">Accept</a> </body>
```

You can also use a button image (the image resides on an external Web server, as in the following example):

```
<p><a href="$ (exception.details) " onclick="Accept ();">
 </a> </body>
```

If you use an HTML editor to compose code, you can paste it into the VPM; however, only copy the HTML from the `<body>` tag to the `</body>` tag.

4. Under Notify mode, select an option that determines notification when visiting a new Web site:
- Notify once for all hosts—The notification page is displayed only once; this is used for configuring compliance pages. This option uses a Virtual Notify URL. If you must change the URL from the default value, please read the limitation section following this procedure.

Note: This option might cause users to experience some noticeable Web browsing slowness.
 - Notify only once for related domains—The notify page reappears each time the user visits a new Web site; this is used for configuring coaching pages.

Note: This option interferes with some Web advertising banners. In some cases, the notification page appears inside the banner. In other cases, banner ads are disabled by javascript errors. To fix these problems, do not serve notification pages for URLs that belong to the *Web Advertising*, *Advertising*, or *Web Ads* category. The actual name of this category varies with the content filtering vendor, and some vendors do not have an equivalent.
 - Notify on every host—The notify page reappears each time the user visits a new Web host. Blue Coat recommends that only highly experienced administrators employ this option. In addition to breaking banner ads, as described above in the previous option, this option, on some internet Web sites, might cause Javascript errors that impair the functionality of the site.
5. Under Notify users again, select an option that specifies when the notification expires and re-notification is required:
- At next browser session— The notification page does not reappear until the next browser session. When a user reboots, logs out, or closes all Web browser windows, this ends the browser session.
 - After (time interval)—Notification reoccurs after the defined elapsed time (minutes or hours); this is useful for coaching.
 - After (specific time)—Notification reoccurs at a specific time of day. You can specify an interval of days; this is useful for compliance.

Note: The time is referenced from the local workstation. If a compliance page is configured, verify the workstations and ProxySG clocks are synchronized.

Section C: Detailed Object Column Reference

Add Notify User Object

Name:

Notification page

Title:

Body:

```
<body>
Notice: The company expects all employees to read the corporate Web
usage policy. By clicking Accept, you acknowledge your understanding
of this policy.

<!-- The following is the Accept button, which you can customize. -->
<p><a href="#{exception.details}" onclick="Accept();">Accept</a>
```

Notify mode

Notify once for all hosts
Virtual notify URL:

Notify only once for related domains

Notify on every host

Notify users again:

At next browser session

After minute(s)

After 7:00 AM in 1 day(s)

OK Cancel Help

Figure 14-36: A User Notify object with a custom message and set to display once a day after 7AM.

Limitations and Workarounds

If you must change the default Virtual Notify URL, consider the following:

- The Virtual Notify URL consists of an HTTP domain name or IP address (`http://`); a port number is optional.
- Do *not* use a host name that is explicitly defined as a *trusted site* on Internet Explorer 6 for Windows XP, Service Pack 2. Furthermore, only use domain names that contain dots. If you use domain names that do not contain dots, the HTTP redirects generated by the notification action causes Internet Explorer to display false warning messages each time the user is redirected from an untrusted site to a trusted site, or the other way around.
- For transparent proxy deployments, the domain name *must* be DNS-resolvable to an IP address that is in the range of destination IP addresses that are routed to the ProxySG.

Policy Interactions

This action generates CPL that might interfere with other policy or cause undesired behavior. Enhancements will occur in future SGOS releases. For this release, consider the following guidelines:

Section C: Detailed Object Column Reference

- Do not create VPM policy that modifies the `Cookie` request header.
- Do not create VPM policy that modifies the `Set-Cookie` and `P3P` response headers.
- Notification pages exist in the browser history. Therefore, if you click `Accept` and are taken to the requested page, then click the back button, you get the notification page again.
- If you have a chain of `ProxySGs`, with different notification pages configured on each appliance in the chain, then each notification page *must* have a different object name.

Strip Active Content

Strips HTTP tags from specified active content HTML pages. For each item you select for removal, you can also create a customized message that is displayed to the user.

Note: Pages served over an HTTPS tunneled connection are encrypted, so the content cannot be modified.

"Section B: Stripping or Replacing Active Content" provides detailed information about the different types of active content.

To Create a Strip Active Content Object:

1. In the Name field, enter name for the object or leave as is to accept the default.
2. Select the active content to be stripped.
3. The default message in the Replacement Text column is Active Content Removed. To replace the default message, double click the field and enter a message. For example, Java applets have been removed. To not display a message, delete the default and leave blank.

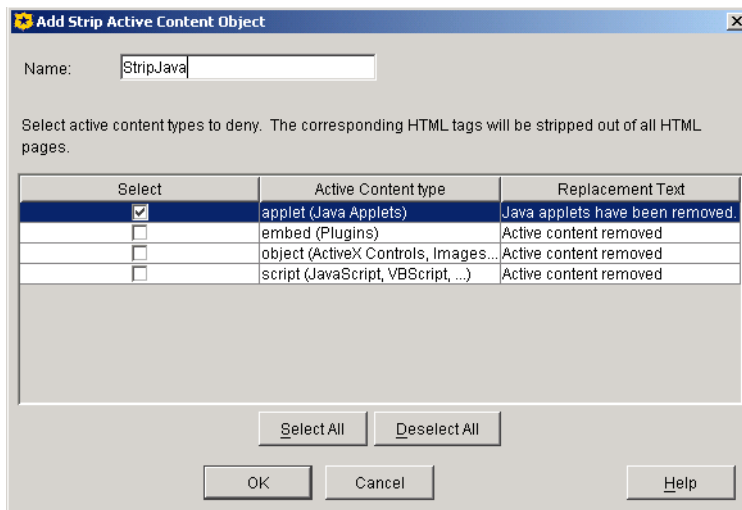


Figure 14-37: Creating a Strip Active Content Action Object

Section C: Detailed Object Column Reference

Exempting the ProxySG

Stripping active content might interfere with Web applications deployed on your intranet. For example, if you create a policy rule that removes Java applets, and the destination defined in the rule contains an IP address of a ProxySG functioning as a proxy, the policy rule actually disables the Management Console because the Console itself is comprised of Java applets.

To prevent this, for each ProxySG functioning as a proxy, create a rule that exempts the IP address of the ProxySG from the stripping action.

1. Click Add Rule.
2. Click Move Up; the rule to exempt the ProxySG must precede the rule that strips active content.
3. In the Destination field, enter the ProxySG IP address.
4. With the IP address entered, right-click it in the Destination field and select Negate from the drop-down list.
5. In the Action field, enter the Remove Active Contents, Java Apps action.

Web Authentication Layer (1)		Admin Authentication Layer (1)		Web Content Layer (1)		Web Access Layer (1)	
No.	Source	Destination	Service	Time	Action	Track	Comment
1	Any	10.1.1.1	Any	Any	StripJava	None	

Figure 14-38: Exempting a ProxySG IP Address

Set Client HTTP Compression

Specifies the behavior when the client wants the content in a different compression form than is in the cache.

To specify compression actions:

1. In the Name field, enter name for the object or leave as is to accept the default.
2. This object has two instructions:
 - A client requests compressed content, but only uncompressed content is available. Select to either compress the content before serving it, or serve uncompressed content.
 - A client requests uncompressed content, but only compressed content is available. Select to either uncompress the content before serving it, or serve compressed content.

The default is to compress or decompress content, respectively, before serving it.

3. Click OK.

For recommended compression configurations, see "HTTP Compression" on page 178.

Set Server HTTP Compression

Enables or disables HTTP compression.

Section C: Detailed Object Column Reference

To specify compression options:

1. In the Name field, enter name for the object or leave as is to accept the default.
2. Select a compression option:
 - Disable HTTP compression—The default. Objects are not compressed.
 - Use client HTTP compression options—Default to the type of content requested by the client.
 - Always request HTTP compression—Force clients to always request compressed content.
3. Click OK.

For recommended compression configurations, see "HTTP Compression" on page 178.

Set SOCKS Compression

SOCKS compression reduces bandwidth and improves the latency between the main office, or *core*, and the *edge* proxies. This also applies to non-Web protocols (for example, Microsoft Exchange) that comprise large percentages of the enterprise traffic.

For incoming SOCKS connections, this object determines if compression is allowed or not. This is typically used by a core (upstream) SOCKS proxy to allow SOCKS connections from an edge (downstream) SOCKS proxy.

- Allow compression—The gateway request to allow compression is granted.
- Do not allow compression—The gateway request to allow compression is not granted and the connection fails.

For detailed information about SOCKS compression and policy, refer to "Understanding SOCKS Compression" on page 188.

Set SOCKS Gateway Compression

SOCKS compression reduces bandwidth and improves the latency between the main office, or *core*, and the *edge* proxies. This also applies to non-Web protocols (for example, Microsoft Exchange) that comprise large percentages of the enterprise traffic.

This object determines whether a forwarded SOCKS connection requests compression or not. This is typically an edge (downstream) SOCKS proxy request to a core (upstream) SOCKS proxy.

Note: The success of the compression request is determined by the upstream proxy, which can allow or deny a compression request.

- Request compression—Asks the upstream proxy to allow a compressed SOCKS connection.
- Do not request compression—The request is forwarded, but compression is not required.
- Use gateway configuration setting—The request is forwarded, using the

For detailed information about SOCKS compression and policy, refer to "Understanding SOCKS Compression" on page 188.

Section C: Detailed Object Column Reference

Manage Bandwidth

Allows you to manage bandwidth for all protocols or specific protocols, on both inbound and outbound traffic.

To manage bandwidth:

1. In the Name field, enter name for the object or leave as is to accept the default.
2. Select to limit bandwidth on the: Client side or Server side.
 - Client side—Traffic flowing between a client and the ProxySG.
 - Server side—Traffic flowing between a server and the ProxySG.
3. Select to limit bandwidth for: Inbound or Outbound traffic.
 - Inbound—Network packets flowing into the ProxySG. Inbound traffic mainly consists of packets originating at the origin content server (OCS) and sent to the ProxySG to load a Web object and packets originating at the client and sent to the ProxySG for Web requests.
 - Outbound—Network packets flowing out of the ProxySG. Outbound traffic mainly consists of packets sent to the client in response to a Web request and packets sent to an OCS or other service (such as a virus scanner) to request a service.
4. Select a Bandwidth Class from the drop-down list.
5. Click OK; click Save Changes.

For complete information about Bandwidth Management, see Chapter 10: “Bandwidth Management” on page 375.

Modify IM Message

In IM clients, replaces or appends the given text that is displayed to IM messages in clients that are logged in through the ProxySG. For example, inform users that their IM messaging activity is monitored.

1. In the Name field, enter a name for the object, or accept the default.
2. In the large, blank field, enter text to be displayed.
3. Select Set message text to replace the text displayed to the user. For example, Instant Messaging is not allowed during these hours. Alternatively, select Append to message text to add the text to the displayed message.

Chapter 17: “Instant Messaging” on page 621 provides more information about regulating IM through the ProxySG, as well as VPM examples.

Return ICAP Patience Page

Specifies to display an ICAP patience page after a predetermined amount of time. Enter a time value (in seconds) that the ProxySG waits for content to be serviced from the origin content server before displaying the page that instructs users an ICAP scan is in progress.

Section C: Detailed Object Column Reference

Note: Patience pages display regardless of any pop up blocking policy that is in effect.

Patience page management and limitations are described in "Customizing ICAP Patience Text" on page 408.

Set Dynamic Categorization

Dynamic categorization extends the process of categorizing a URL. Traditional content filtering involves searching of massive URL pattern databases, which are published by vendors and downloaded to the ProxySG at specified intervals. As new content constantly reaches the Web, the limitation is that it cannot be filtered until its existence is discovered, added, and uploaded. Dynamic categorization enhances content filtering by scanning a new Web page, attempting to determine its contents, and categorizing accordingly in real time.

When an un-categorized page is first encountered, the ProxySG calls an external service with a categorization request. Once the content is scanned, a category is assigned (a majority of the time).

For related information, see "Configuring Dynamic Categorization for Blue Coat Web Filter" on page 651.

To configure Dynamic Categorization:

1. Select a mode:
 - Do not categorize dynamically—The loaded database is consulted for category information. URLs not in the database show up as category none.
 - Categorize dynamically in the background—Objects not categorized by the database are dynamically categorized as time permits. Proxy requests are not blocked while DRTR is consulted. Objects not found in the database appear as category pending, indicating that DRTR was requested, but the object was served before the DRTR response was available.
 - Categorize dynamically in realtime—The default. Objects not categorized by the database are dynamically categorized on first access. If this entails consulting the DRTR service, the proxy request is blocked until DRTR responds.
 - Use dynamic categorizing setting from configuration—Default to the ProxySG configuration (Content Filtering>Blue Coat>Dynamic Categorization).
2. Click OK.

Set External Filter Service

Specifies which installed content filtering service or service group a content request is subjected to or bypasses, and specifies what occurs if a communication error occurs between the ProxySG and the external service.

To Determine External Filter Service Request Behavior:

1. In the Name field, enter a name for the object or leave as is to accept the default.

Section C: Detailed Object Column Reference

2. To instruct all request types defined in the rule to not route through an external filter service for content filtering, select Bypass external filter request.
3. To instruct all requests defined in the rule to route to a specific external filter service, select Use External Filter Service; from the drop-down list, select the external filter service or service group (which must already exist on the ProxySG).
4. Select an On communication error with external filter service option. To deny all requests if a communication error occurs, select Deny the request; to allow requests to go through without content filtering, select Bypass external filter service and process the request.

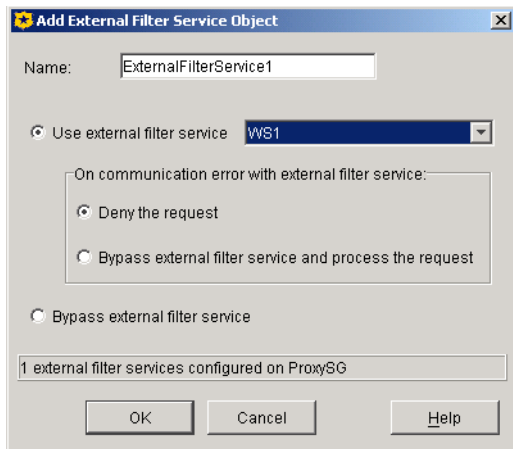


Figure 14-39: External Service Object

Set ICAP Request Service

Specifies which installed ICAP service or service group a content request routes to or bypasses, and specifies what occurs if a communication error occurs between the ProxySG and the ICAP server.

To Determine ICAP Request Behavior:

1. In the Name field, enter a name for the object or leave as is to accept the default.
2. To instruct all request types defined in the rule to not route through an ICAP service for content scanning, select Bypass ICAP request service.
3. To instruct all request or response types defined in the rule to route to a specific ICAP service, select Use ICAP Request Service; from the drop-down list, select the ICAP service or service group (which must already exist on the ProxySG).
4. Select an On communication error with ICAP service option. To deny all requests or responses if a communication error occurs, select Deny the request; to allow requests or responses to go through without ICAP scanning, select Bypass ICAP service and process the request.

When the ICAP service is restored, these objects are scanned and served from the cache if they are requested again.

Section C: Detailed Object Column Reference

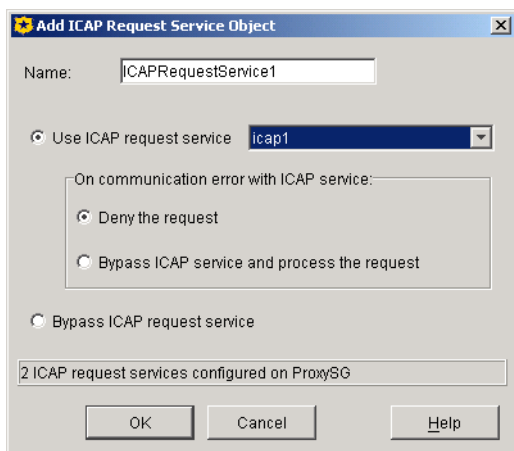


Figure 14-40: ICAP Service Object

Set ICAP Response Service

Identical to "Set ICAP Request Service", but applies to other protocol responses, such as HTTP and FTP.

Set FTP Connection

For an outgoing request over FTP, specifies whether the FTP connection should be made immediately or deferred, if possible. The benefit of deferring connections is that requests for previously cached content can be served without contacting the origin server, which reduces the FTP load on that server.

Set SOCKS Acceleration

Specifies whether or not accelerate SOCKS requests, and defines the transport method.

To Set SOCKS Acceleration:

1. In the Name field, enter a name for the object or leave as is to accept the default.
2. Select one of the following:
 - Automatically—Accelerates SOCKS requests automatically, based on the destination port receiving the connection.
 - Do Not Accelerate—Never accelerate SOCKS requests matched by this rule.
 - Accelerate via [HTTP | AOL IM | MSN IM | Yahoo IM]—Specifies the type of acceleration applied to requests matched by this rule.
3. Click OK.

Set Streaming Max Bitrate

Specifies the maximum bitrate, in kilobits per second, of requested streaming media. If a request exceeds this rule, the request is denied.

Section C: Detailed Object Column Reference

Send DNS/RDNS Response Code

Specifies to send out the default response code or a selectable error response code. Perform one of the following:

- Select Send Default DNS Response; optionally, enter a TTL value.
- Select Send Error Response Code and select a code from the drop-down list.

Send DNS Response

Specifies which IP address to return for a specified host.

To Set a DNS Response:

1. In the Host field, enter a host name that is returned.
2. To respond with the incoming IP address, select Respond with proxy IP.
3. To respond with one or more IP addresses:
 - a. Select Respond with listed IPs.
 - b. Click Add. The Add DNS Response IP dialog appears.
 - c. Enter an IP address and click Add.
 - d. Repeat as desired; click Close when finished.
4. (Optional) In the TTL field, enter a time-to-live value.
5. Click OK.

Send Reverse DNS Response

Specifies which host to return for a reverse DNS response. Optional: define a time-to-live value.

Do Not Cache

This is a static object. Specifies that objects are never cached.

Force Cache

This is a static object. Specifies that (cacheable) objects are always cached. Objects that are not cacheable (for example, RealMedia file types) and supported in pass-through mode only are not cached.

Use Default Caching

This is a static object. Overrides the Do Not Cache and Force Cache actions and instructs the ProxySG to use its default determination of whether or not to cache the content.

Section C: Detailed Object Column Reference

Mark/Do Not Mark As Advertisement

These are static objects. Specifies content to be identified as an advertisement. The ProxySG still fetches content from the cache (if present); however, just after serving to the client, the content is re-fetched from the ad server so that hit counters are updated.

Enable/Disable Pipelining

These are static objects. Enables or disables the ProxySG pipelining feature, which, when enabled, examines Web pages for embedded objects and requests them from the origin server in anticipation of a client request.

Set TTL

Specifies the time-to-live (TTL) an object is stored in the ProxySG. In the Name field, enter a name for the object (or leave as is to accept the default); in the TTL field, enter the amount of time in seconds.

Send Direct

This is a static object. Overrides forwarding host, SOCKS gateway, or ICP configurations and instructs the ProxySG to request the content directly from the origin server.

Integrate/Do Not Integrate New Hosts

This is a static object. Used in server accelerator deployments. When enabled, the corresponding host that is accessed is added to the list of hosts for which the ProxySG performs health checks. If that host name resolves to multiple IP addresses that correspond to different servers, the ProxySG fetches content from the available servers and ignores the servers that fail the health check.

Allow Content From Origin Server

This is a static object. Allows request to access content from an origin server if the content is not cached.

Serve Content Only From Cache

This is a static object. Requests to access content that is not cached are denied. If the content is cached, the content is served.

Select SOCKS Gateway

Specifies which SOCKS gateway, if any, to use; defines behavior if communication between the SOCKS gateway and the ProxySG is down.

- To instruct the rule to connect directly without routing through a SOCKS service, select Do not use SOCKS gateway.

Section C: Detailed Object Column Reference

- To instruct the rule to connect through a SOCKS gateway, select Use SOCKS Gateway and select an installed SOCKS service from the drop-down list.

In the If no SOCKS gateway is available field, select Deny the request or Connect directly, which allows requests to bypass the SOCKS service.

Select Forwarding

Specifies which forwarding host or group, if any, to use; defines behavior if communication between the forwarding and the ProxySG is down.

- To instruct the rule to connect directly without redirecting to a forwarding host or group, select Do not forward.
- To instruct the rule to redirect to a forwarding host, select Use Forwarding and select an installed forwarding host from the drop-down list.

In the If no forwarding is available field, select Deny the request (fail closed) or Connect directly (fail open), which allows requests to bypass the forwarding host.

- To instruct the rule to forward using the ICP configuration, select Forward using ICP.

Set IM Transport

Specifies the transport method used for IM traffic.

- Auto—Connects using the transport method used by the client.
- HTTP—Tunnels the IM requests over HTTP.
- Native—Connects using the native transport used by the service.

Set Streaming Transport

Specifies which streaming transport method the rule uses.

- Auto—Connects using the transport method used by the client.
- HTTP—Streaming over HTTP.
- TCP—Streaming over TCP.

Combined Action Objects

Allows you to combine an action object that invokes multiple actions. See "Using Combined Objects" on page 530.

Action Column/Policy Layer Matrix

The following matrix lists all of the Action column objects and indicates which policy layer they apply to.

Section C: Detailed Object Column Reference

Object	Admin Auth	Admin Access	DNS Access	SOCKS Auth	Web Auth	Web Access	Web Content	Forwarding
Allow						x		
Deny	x	x			x	x		
Allow Read-Only Access		x						
Allow Read-Write Access		x						
Do Not Authenticate	x			x	x			
Authenticate	x			x	x			
Force Authenticate	x			x	x			
Bypass Cache						x		
Do Not Bypass Cache						x		
Check Authorization						x	x	
Do Not Check Authorization						x	x	
Always Verify						x	x	
Use Default Verification						x	x	
Block Up Ads						x		
Do Not Block PopUp Ads						x		
Force NTLM For Server Auth						x		
Do Not Force NTLM For Server Auth						x		
Reflect IM Messages						x		
Do Not Reflect IM Messages						x		
Block IM Encryption						x		
Do Not Block IM Encryption						x		
Return Exception						x		
Return Redirect						x		
Send IM Alert						x		
Modify Access Logging						x	x	
Override Access Log Field						x	x	
Rewrite Host						x		

Section C: Detailed Object Column Reference

Object	Admin Auth	Admin Access	DNS Access	SOCKS Auth	Web Auth	Web Access	Web Content	Forwarding
Reflect IP			x			x		
Suppress Header						x		
Control Request Header						x		
Control Response Header						x		
Notify User						x		
Strip Active Content						x		
Set Client HTTP Compression						x		
Set Server HTTP Compression						x		
Set SOCKS Compression						x		
Set SOCKS Gateway Compression								x
Modify IM Message						x		
Return ICAP Patience Page						x		
Set Dynamic Categorization							x	
Set External Filter Service						x		
Set ICAP Request Service						x	x	
Set ICAP Response Service							x	
Use Default Caching							x	
Set FTP Connection						x		
Set SOCKS Acceleration						x		
Set Streaming Max Bitrate						x		
Send DNS/RDNS Response Code			x					
Send DNS Response			x					
Send Reverse DNS Response			x					
Do Not Cache							x	
Force Cache							x	
Mark As Advertisement							x	

Section C: Detailed Object Column Reference

Object	Admin Auth	Admin Access	DNS Access	SOCKS Auth	Web Auth	Web Access	Web Content	Forwarding
Do Not Mark as Advertisement							x	
Enable Pipelining							x	
Disable Pipelining							x	
Set TTL							x	
Send Direct								x
Integrate New Hosts								x
Do Not Integrate New Hosts								x
Allow Content From Origin Server								x
Serve Content Only From Cache								x
Select SOCKS Gateway								x
Select Forwarding								x
Reflect IP								x
Set IM Transport								x
Set Streaming Transport								x
Combined Objects			x			x	x	x

Track Object Column Reference

A *track* object defines the parameters for tracking and tracing traffic. All policy layers contain the same trace objects, but tracking parameters are layer-specific.

Note: Because of character limitations required by the generated CPL, only alphanumeric, underscore, and dash characters can be used to define an action object name.

Event Log, E-mail, and SNMP

You can customize the event log, E-mail notification, and SNMP with triggers. These triggers are the same for all three object types.

Customize an event log, E-mail, or SNMP object:

1. Right-click the Tracking cell in a policy layer and select Set; the Set Track Object dialog appears.

Section C: Detailed Object Column Reference

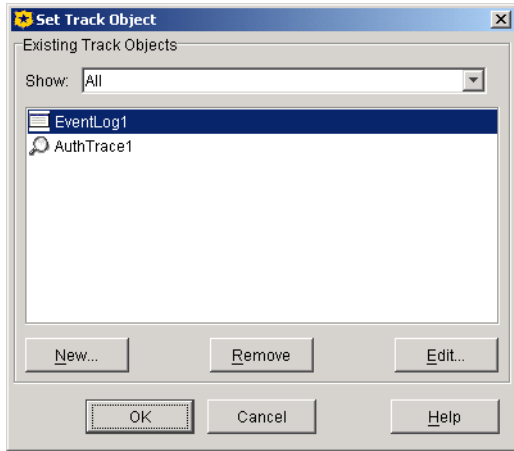


Figure 14-41: Add Track Object Dialog

2. Click New and select Event Log, Email, or SNMP; the appropriate add object dialog appears.
3. In the Name field, enter a name for this object or leave as is to accept the default.
Note: The E-mail object also contains a Subject field.
4. In the Message Text field, enter a customized message that appears with each entry.
5. Optional: In the Substitution Variables field, select a variable and click Insert. Repeat as required.

The substitution variables instruct the ProxySG to append specific information to the tracking object. The variables are categorized alphabetically, according to prefix.

Note: Some variables do not have prefixes.

Tracing Objects

This object specifies rule and Web traffic tracing.

Click Trace Level and select one of the following trace options:

- No Tracing—The default.
- Request Tracing—Generates trace output for the current request. The trace output contains request parameters (such as URL and client address), the final values of property settings, and descriptions of all actions taken.
- Rule and Request—Generates trace output that displays each rule that was executed
- Verbose Tracing—Generates the same output as Rule and Request, but also lists which rules were skipped because one or more of their conditions were false, and displays the specific condition in the rule that was false.

Furthermore, a trace destination can be entered that specifies the destination for any trace produced by the current transaction. To specify a destination path, select Trace File and enter a path in the field. For example, `abc.html`.

Section C: Detailed Object Column Reference

If a trace destination is configured in multiple layers, the actual trace destination value displayed is the one specified in the last layer that had a rule evaluated (which has a destination property configured). Consider the following multiple Web Access Layer example, demonstrated by the generated CPL:

```
<PROXY>
url.domain=aol.com trace.request(yes) trace.rules(all)
trace.destination("aol_tracing.html")
url.domain=msn.com trace.request(yes) trace.rules(all)
trace.destination("msn_tracing.html")

<PROXY>
client.address=10.10.10.1 trace.request(yes) trace.rules(all)
```

The resulting actions are:

- Requests to the `aol.com` domain are logged to `aol_tracing.html`.
- Requests to the `msn.com` domain are logged to `msn_tracing.html`.
- Requests from the client with the IP of `10.10.10.1` are logged to the default location of `default.html`.

Note: After using a trace to troubleshoot, remove the trace to save log space.

The Trace File option can be used in conjunction or separately from the Trace Level option.

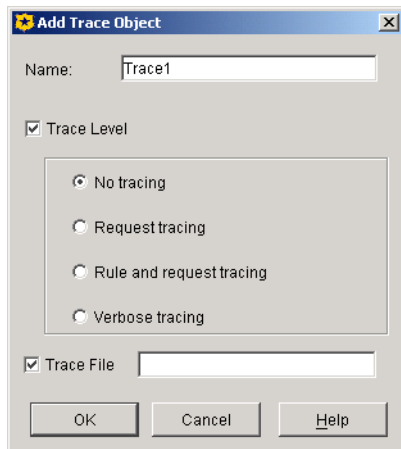


Figure 14-42: Creating a Trace Object

The default path of the trace file is accessible through one of the following URLs.

If the Management Console secure mode is enabled (the default on a new or upgraded system):

```
https://ProxySG_IP_address:8082/Policy/Trace/default_trace.html
```

If the Management Console is deployed in non-secure mode:

```
http://ProxySG_IP_address:8081/Policy/Trace/default_trace.html
```

Section C: Detailed Object Column Reference

Combined Track Object

Allows you to combine track objects into one. See "Using Combined Objects".

Track Objects/Policy Layer Matrix

The following matrix lists all of the Track and column objects and indicates which policy layer they apply to.

Object	Admin Auth	Admin Access	DNS Access	SOCKS Auth	Web Auth	Web Access	Web Content	Forwarding
Event Log		x	x			x	x	
Email Log		x	x			x	x	
SNMP Objects		x	x			x	x	
Trace	x	x	x	x	x	x	x	x
Combined Objects		x	x			x	x	

Comment Object Reference

The Comment object allows you to write any text to aid in labeling the policy layer. The text in this field does not impact the performance of the rule.

Using Combined Objects

As previously discussed, you select one object for as many object types as required for a given rule. Most object types also have the option of using a combined object. This feature allows you to select multiple objects for a given type, thus creating more complex tools. There are two uses for combined conditions: lists and multiple object types. Also consider the Negate option, which exempts the objects in the list.

Example One

Consider the following example. You want a Web Content policy layer that as an action forces authorization *and* sends the response to an ICAP service for content scanning.

Section C: Detailed Object Column Reference

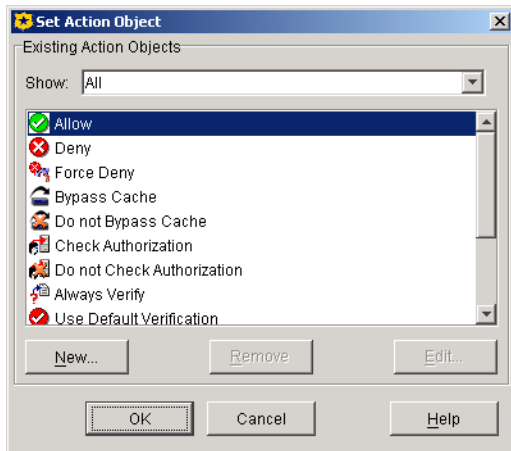


Figure 14-43: Set Action Object dialog

1. In the Set Action Object dialog, select New>Combined Action Object.

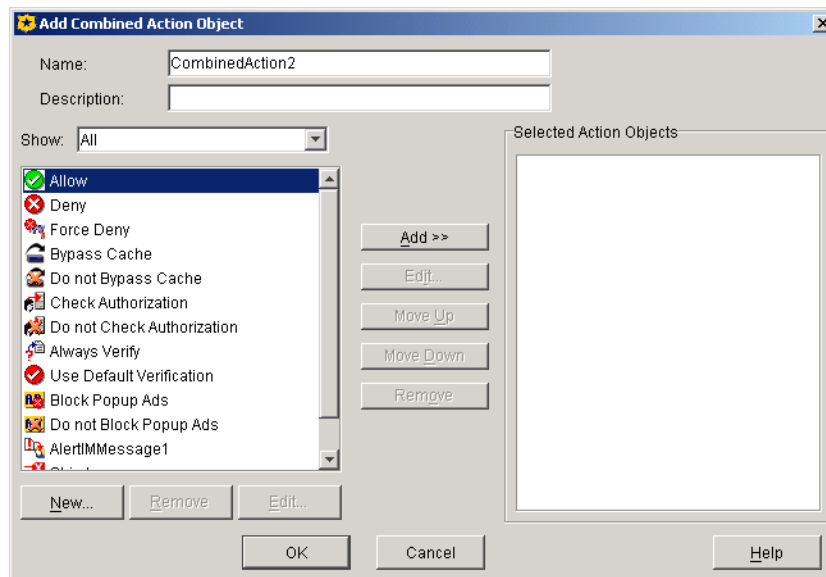


Figure 14-44: Add Combined Object Dialog

2. Select Check Authorization; click Add.
3. Select an existing ICAP service object. For this example, the object is named ICAPReqServ1. Click Add.

Section C: Detailed Object Column Reference

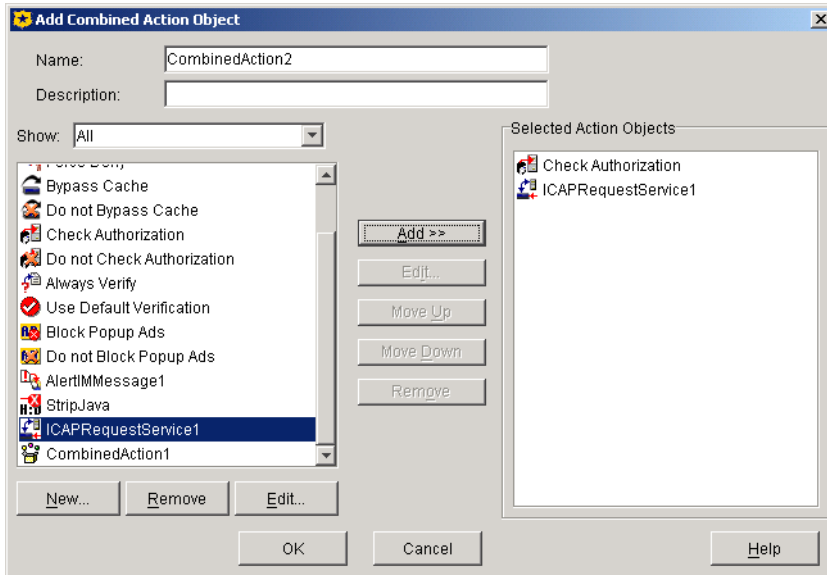


Figure 14-45: Creating a Combined Object

4. Click OK. The CombinedAction1 object appears as a separate, selectable object.

5. Select CombinedAction1; click OK. The object is now part of the rule.

Based on the other parameters specified in the rule, all requests are forced to an upstream server for authorization and the Web responses are subject to content scanning through the ICAP service.

Example Two

In the following example, the rule searches for one of the IP Address/Subnet objects and one of the streaming client user agents.

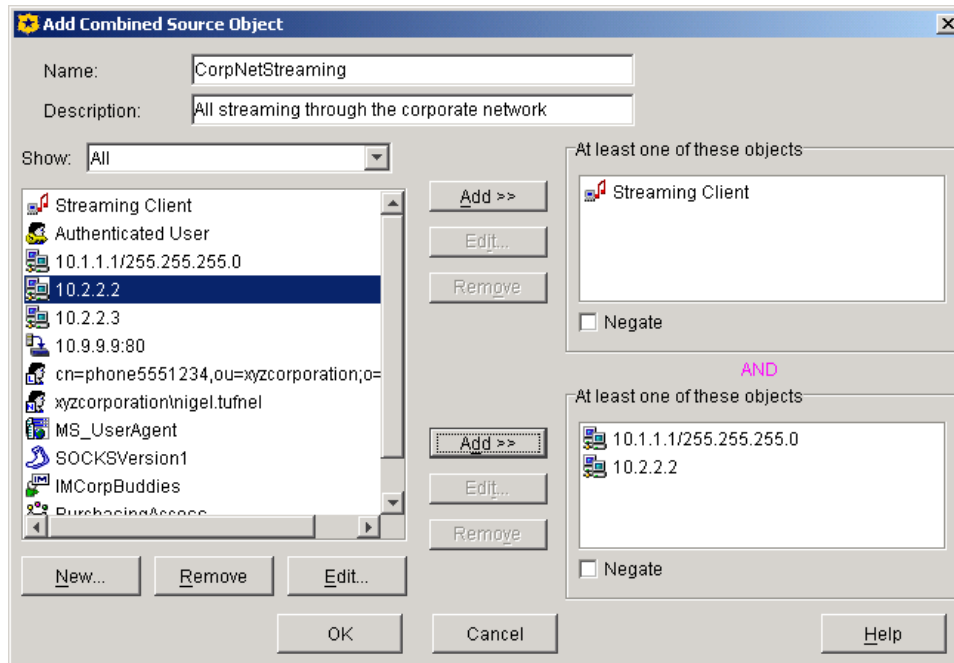


Figure 14-46: Combined Object with Multiple Object Types

Note

The VPM displays various warning messages if you attempt to add objects that creates an invalid combined object. However, it is possible to add a combined object to another combined object, even if doing so presents duplication of simple object definitions without receiving validation warnings. For example, the contents of a child combined object may have already been included either within the parent combined object directly, or indirectly within other child combined objects. This is allowable because of the complexity some combined objects and policies can achieve.

Centralized Object Viewing and Managing

The All Objects feature allows you view a list of all objects—both static and user-defined—that currently exist across all layers and columns. To view all configured VPM objects, in the Menu Bar select View>All Objects.

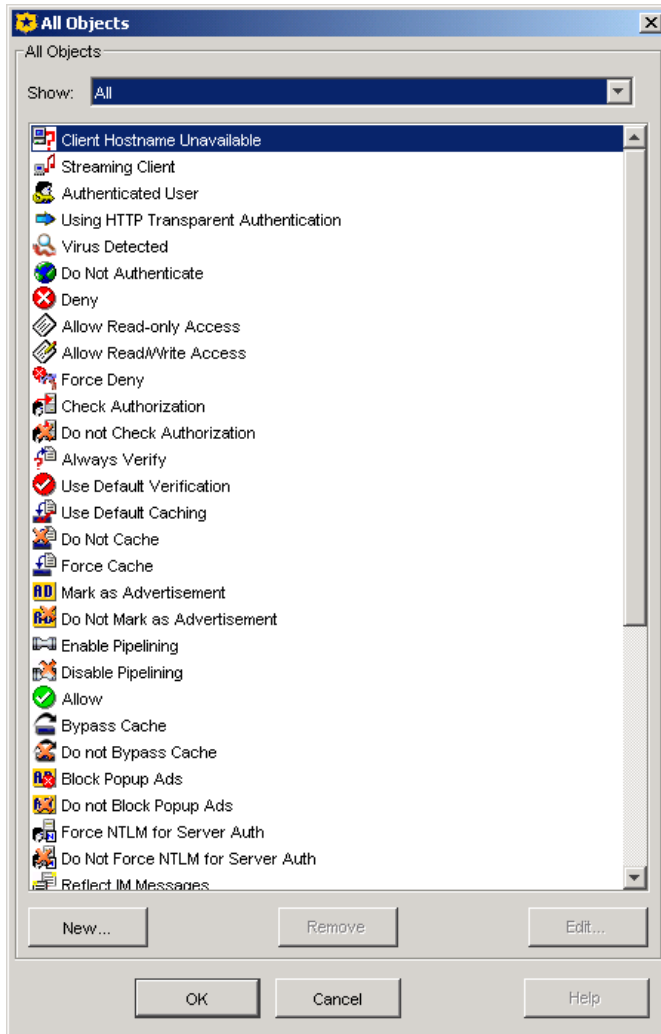


Figure 14-47: The All Objects dialog.

The objects are displayed according to the policy layer order (click Policy in the menu bar) and the column order (as presented in "Policy Layer and Rule Object Reference" on page 466). To narrow the scope of the displayed objects, use the Show drop-down list at the top. You can select to display only the static objects for the Source, Destination, Service, and Action columns, or any object type. For example, you want to only view the user-defined P2P Client objects. Scroll down and select P2P Client Objects.




Figure 14-48: Narrowing the list to user-defined P2P objects.

Creating Objects

The All Objects dialog also allows you to create objects. Once an object is created, it appears in the list. When creating or editing policy layers, the objects are available to add to rules.

To create an object:

1. Select New. The available columns and relevant objects are displayed in a cascade style.
2. Select *Column>Object*. The Add dialog for that object appears.
3. Define the object as required.
4. Click OK.

Note: When creating Combined Objects, not all objects that appear in the left column are valid for more than one policy layer type. For example, the IM User object is only valid in the Web Access Layer>Source column. If you attempt to add an object that is not valid, a dialog appears with that information.

Editing Objects

Any user-defined object can be modified. Highlight the object and click Edit. After editing the object, re-install the policy to apply the modified object in every policy layer it exists in.

Deleting Objects

You cannot delete an object that is currently part of an installed policy or combined object. Before removing an object, you can use the View>Object Occurrences feature to identify which policy layers contain the object.

Creating Categories

This feature allows you create the content filter URL categories that can be used in the Category object. The Destination column in the DNS Access, Web Access, Web Authentication, and Web Content policy layers contain the Category object. Similarly, categories created in the Category object (see "Category" on page 488) appear in this dialog and can be edited.

To Create a Category:

1. In VPM, select Configuration>Edit Categories. The Edit Categories dialog appears.

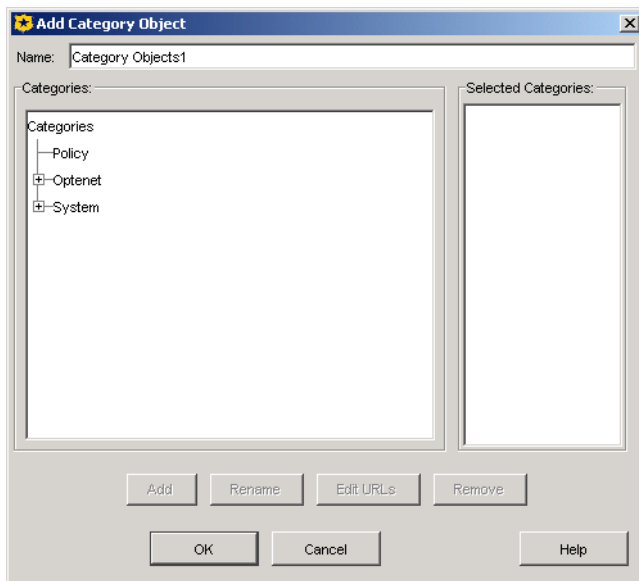


Figure 14-49: The Edit Categories Dialog

2. Select Policy; click Add. The Object Name dialog appears.
3. Name the category and click OK.
4. Drop the Policy list and select the created category; click Edit URLs. The Edit Locally Defined Category Object dialog appears.
5. Enter URLs appropriate for the content filter category you are creating; click OK.
6. Click OK.

Note: If other administrators have access to the ProxySG through other workstations and are creating categories either through VPM or with `inline` commands, consider that newly-created or edited categories are not synchronized until the policy is installed. When the policy is installed with VPM, the categories are refreshed. If too many categories are created at the same time and confusion occurs, select the File>Revert to Existing Policy on ProxySG Appliance option to restore the policy to the previous state and reconfigure categories.

Refreshing Policy

In between occurrences when either VPM is closed and reopened or Install Policies is invoked, VPM does not recognize changes to VPM-managed policy that were made on the ProxySG through another method. For example:

- Another administrator opens a separate VPM to make changes.
- Another administrator edits the local or central policy file through the serial console.
- Another administrator makes edits the local or central policy file through the Management Console.
- A new content filter database is downloaded automatically and the new update contains category changes.
- A new content filter database is downloaded manually by an administrator through the CLI or the Management Console.

Restricting DNS Lookups

This section discusses DNS lookup restrictions and describes how to create a list.

About DNS Lookup Restriction

The DNS lookup restriction list is a list of domain names that apply globally, regardless of policy layer definitions. Once a domain name is added to the list, DNS lookup requests do not occur for that domain name while policy is evaluated. For more detailed information about using DNS lookups, refer to the *Blue Coat ProxySG Content Policy Language Guide*.

Creating the DNS Lookup Restriction List

The list is created from the VPM Menu bar.

To Create the DNS Lookup Restriction List:

1. Select Configuration>Set DNS Lookup Restrictions; the Set DNS lookup restrictions dialog appears.
The default is None; no domain names are restricted.
2. To restrict every domain name, select All.
3. To add specific domain names, perform the following steps.
 - a. Select Listed Host Patterns.
This enables the Host Patterns field.
 - b. Click Add; the Add Host Pattern dialog appears.
 - c. Enter a domain name; click OK.
 - d. Repeat to add other domain names.
 - e. Click OK.

Restricting Reverse DNS Lookups

This section discusses reverse DNS lookup restrictions and describes how to create a list.

About Reverse DNS Lookup Restriction

The Reverse DNS lookup restriction list is a list of subnets that apply globally, regardless of policy layer definitions. Once a subnet is added to the list, the ProxySG will not perform a reverse lookup of addresses on that subnet during policy evaluation. For more detailed information about using reverse DNS lookups, refer to the *Blue Coat ProxySG Content Policy Language Guide*.

Creating the Reverse DNS Lookup Restriction List

The list is created from the VPM Menu bar. This prevents the ProxySG from performing reverse DNS lookups of addresses in the list while evaluating policy.

To Create the Reverse DNS Lookup Restriction List:

1. Select Configuration>Set Reverse DNS Lookup Restrictions; the Set Reverse DNS lookup restrictions dialog appears.

The default is None; no subnets are restricted.

2. To restrict every subnet, select All.
3. To add specific subnets, perform the following steps.
 - a. Select Listed Subnets.
This enables the Subnets field.
 - b. Click Add; the Add Subnet dialog appears.
 - c. Enter a subnet; click OK.
 - d. Repeat to add other subnets.
 - e. Click OK.

Setting the Group Log Order

This section discusses the group log order and describes how to create a list.

About the Group Log Order

The Group Log Order object allows you to establish the order group data appears in the access logs. For more detailed information about using group log ordering, refer to the *Blue Coat ProxySG Content Policy Language Guide*.

Creating the Group Log Order List

The list is created from the VPM Menu bar.

To Create the Group Log Order List:

1. Select Configuration>Set Group Log Order; the Set Group Log Order dialog appears.

2. Click Add; the Add Group Object dialog appears.
3. In the Group Name field, enter the name of a group.
The group must be already configured on the ProxySG.
4. From the Authentication Realm drop-down list, select a realm.
5. Click OK.
6. Repeat as required to add more groups.
7. To order the list, select a group and click Move Up or Move Down until you achieve the desired order.
8. Click OK.

Section D: Managing Policy Layers and Files

This section contains the following topics:

- "How Policy Layers, Rules, and Files Interact" on page 541—Describes the importance of rule order policy layer order.
- "Managing Policy Files" on page 544—Describes how to save and install policies on the ProxySG.
- "Installing VPM-Created Policy Files" on page 545—Describes how to propagate a policy file created on one ProxySG to another.
- "Viewing the Policy/Created CPL" on page 547—Describes how to view the underlying CPL that is created with VPM.

How Policy Layers, Rules, and Files Interact

The following critical points discuss the behaviors and priorities of policy rules, layers, and files:

- Rules in different policy layers of the same type work together, and the order of policy layers is important.
- The order of policy layers of different types is important.
- The order of rules in a policy layer is important.
- Policy created in VPM is saved in a file on the ProxySG; the state of the VPM user interface is also stored as an XML file on the ProxySG.

Note: These files are stored *only* if the policy is installed without any errors.

- How the appliance evaluates those rules in relation to policy layers that exist in the central and local policy files is important. For more information, see Appendix 13: "Managing Policy Files".

How VPM Layers Relate to CPL Layers

VPM generates CPL in various layers, but the concept of layers presented in VPM is slightly different. VPM provides policy layers for special purposes. For example, Web Authentication and Web Authorization, which both generate CPL `<Proxy>` layers. This minimizes timing conflicts by restricting the choices of triggers and properties to those compatible timing requirements. The following table summarizes how to use VPM layers and which CPL layers result.

Table 14.2: VPM-Generated CPL Layers

Policy Purpose	VPM Layer	CPL Layer
Establish Administrator identities.	Admin Authentication	<code><Admin></code>
Control Administrator access.	Admin Authorization	<code><Admin></code>
Control DNS access.	DNS Access	<code><Proxy></code>
Establish SOCKS user identities.	SOCKS Authentication	<code><Proxy></code>
Establish user identities.	Web Authentication	<code><Proxy></code>
Control user access.	Web Access	<code><Proxy></code>
Control content independent of users.	Web Content	<code><Cache></code>
Control forwarding.	Forwarding	<code><Forward></code>

Note: VPM currently does not support the `<Exception>` layer.

Ordering Rules in a Policy Layer

The ProxySG evaluates the rules in the order in which they are listed in a policy layer. When it finds a rule that applies to the situation, it skips the remaining rules in the policy layer and goes on to the next policy layer.

Consider the following simple example. Assume that a company has a policy that prohibits everyone from accessing the Web. This is a policy that is easy to create with a Web Access layer rule.

There are, however, likely to be exceptions to such a broad policy. For example, you require the manager of the purchasing department to be able to access the Web sites of suppliers. Members of the sales department need to access their customer Web sites. Creating Web Access rules for both these situations is also simple. But if you put all these rules in a single policy layer, then the rule prohibiting access to everyone must be ordered last, or the other two rules will not be applied.

The principle design of rules within policy layers is:
Always go from the specific to the general.

Using Policy Layers of the Same Type

Because the ProxySG skips the remaining rules in a policy layer as soon as it finds one that meets the condition, multiple policy layers and a combination of rules might be required to accomplish a task.

Consider the following example. A company does not want to prohibit its employees from accessing the Web, but it does not want them to abuse the privilege. To this end, the company wants employees who access the Web to authenticate when they do so; that is, enter a username and password. So the company creates a Web Authentication policy layer with a rule that says: “If anyone from anywhere in the company sends a request to a URL on the Web, authenticate the client before granting access.”

The company also allows members of the group Sales to access various sports Web sites only during non-work hours. Given the Web Authentication rule above, these people must authenticate when they do this. But the company feels that it is not important for people going to these sites after hours to authenticate. So the company creates the following Web Access policy-layer rule:

- Grant Sales personnel access to sports Web sites from 5:00 PM to midnight.

But there are additional issues. Some members of the sales department spend a lot of time watching game highlights on video clips, and this takes up a lot of bandwidth. At the same time, a lot of customers access the company Web site in the evening (during non-work hours), so internal bandwidth should remain manageable. The company, therefore, limits the bandwidth available to the people in the Sales department with a Web Access layer rule that is identical to the one above in all respects except for the action:

- Grant Sales personnel access to sports Web sites from 5:00 PM to midnight, but limit the maximum streaming bitrate to 300 kilobits per second.

For both these rules to work, they need to be in separate policy layers. If they were in the same policy layer, the rule listed second would never be applied.

Ordering Policy Layers

The order of policy layers is also important. The ProxySG evaluates policy layers in the order in which they are listed in VPM. When the ProxySG is going through policy layers, it does not execute a given rule as soon as it finds that it meets the specific situation. Rather, it compiles a list of all the rules that meet the condition; when it has gone through all the policy layers, it evaluates the list, resolves any apparent conflicts, and then executes the required actions. If there is a conflict between rules in different policy layers, the matching rule in the policy layer evaluated last takes precedence.

The principle design of policy layers is: Always go from the general to the specific; that is, establish a general rule in an early policy layer, and then write exception rules in later policy layers.

In the above example, there are two Web Access policy layers: one contains a rule stating that Sales personnel can access certain Web sites without authenticating, and the other states that when they do access these Web sites, limit the available bandwidth. The order of these policy layers is irrelevant. The order is irrelevant because there is no conflict between the rules in the layers.

The following is an example in which the order of policy layers does matter. Assume all URL requests from members of the purchasing department are directed to a single proxy server. To discourage employees from surfing the Web excessively during business hours, a company creates a Web Authentication Policy rule that says: “Whenever a client request comes in to the proxy server, prompt the client to authenticate.”

Members of the purchasing department, however, need to access specific Web sites for business reasons, and the company does not want to require authentication every time they do this. So they create a Web Access policy rule that says: “If any member of the purchasing department sends a request to a specific URL contained in a combined-object list, allow access.”

The policy layer with the first rule needs to come first in evaluation order; it is then overridden by the second rule in a subsequent policy layer.

Installing Policies

As you add policy layers and rules, your work is saved in a file on the ProxySG. However, policies only take effect after you install the policies and the generated XML has been validated. The ProxySG then compiles the policies into CPL format and saves the resulting policies in the `vpm.cpl` file. This overwrites any policies previously created using VPM. The appliance saves VPM-generated policies in a single file and loads it all at once. You do not need to load policies separately, as is the case with the local or central policy files.

To Install Policies:

- Select File>Install Policies, or
- Click Install Policies on the Rule bar.

The VPM validates the generated XML for any issues, such as missing layers. If the validation passes, the CPL is generated and the policies are loaded.

If the XML fails the validation, a dialog appears allowing you to:

- Revert to the policy currently installed on the ProxySG, or
- Continue to edit the policy and attempt another installation.

Furthermore, the failed XML file is written to your hard disk; view this file to troubleshoot the failed XML. The default location for this file is:

```
C:\Documents and Settings\user.name\bluecoat\vpm_err.xml
```

Notes:

The Category and Notify User objects and the DNS Lookup Restrictions, Reverse DNS Lookup Restrictions, and Group Log Order configuration objects generate CPL, regardless if they are or are not included in rules. These specific objects and features allow users to edit categories and lists that might or might not be used in current policies.

Managing Policy Files

This section describes how to manage the policy files generated by the VPM.

Refreshing Policy

In between occurrences when either VPM is closed and reopened or Install Policies is invoked, VPM does not recognize changes to VPM-managed policy that were made on the ProxySG through another method. For example:

- Another administrator opens a separate VPM to make changes.
- Another administrator edits the local or central policy file through the serial console.
- Another administrator makes edits the local or central policy file through the Management Console.
- A new content filter database is downloaded automatically and the new update contains category changes.
- A new content filter database is downloaded manually by an administrator through the CLI or the Management Console.

Reverting to a Previous Policy

If after creating new policies or editing an existing policy you decide to abandon the process and continue with the existing policy installed on the ProxySG, you can revert to that version. All current changes are deleted (VPM provides a verification prompt).

To Revert to an Existing Installed Policy:

Select File>Revert to Existing Policy on ProxySG Appliance.

Changing Policies

You can change, edit, delete, add to, and otherwise manage policies created in VPM at any time by returning to VPM and working with policy layers and rules just as you did when creating them.

Disabling and Enabling Policy Rule

Occasionally, you may need to temporarily disable rules in a policy layer; for example, when troubleshooting compiles errors and warnings. This might help confirm that the ProxySG can successfully compile the remaining policy. After disabling a rule, you can edit the objects and re-enable the rule.

To Disable or Enable A Rule:

1. Click the appropriate policy layer tab.
2. Right-click in the No. column.
3. Click Disable Rule on the shortcut menu. The policy editor changes the rule text color to red.

4. To enable the rule, repeat step 3. After you enable a disabled rule, the policy editor changes the rule text color to black.

Deleting a Policy Layer

To Delete a Policy Layer:

1. Right-click the tab of the policy layer to be deleted.
2. Select Delete Policy from the drop-down list.

You can also right-click the policy layer tab and select Edit>Delete Layer.

Installing VPM-Created Policy Files

Policies created with VPM are saved on the specific ProxySG on which they are created. SGOS automatically creates the following files when saving VPM-created policies:

```
config_policy_source.xml
config_policy_source.txt
```

You can install VPM policies that were created on another ProxySG. This requires the following steps:

1. Copy the two VPM files, to be shared, to a Web server from the ProxySG on which they reside.
2. Use the Management Console or CLI to load VPM files on another ProxySG.

To Copy VPM Files from a ProxySG to a Web server:

1. Select Statistics>Advanced.
2. Scroll down and click Policy.

The page jumps down to the Policy files links.

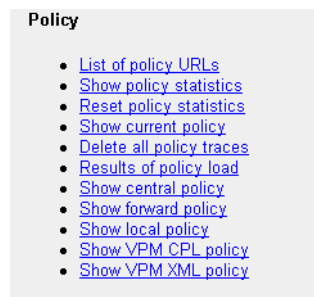


Figure 14-50: Policy Files in Custom URLs

3. Right-click the Show VPM CPL policy link.
4. In the Save As dialog, enter the full path to a directory on the Web server before the file name and click OK.

Important: The Save As dialog offers the appropriate default file name (config_policy_source.xml or config_policy_source.txt). You can change the names, including the extension. This can be helpful if an enterprise is using various sets of shared VPM files. You could rename files to indicate the ProxySG on which they were created, for example, or for a department that has a set of VPM-specific policies, used perhaps in multiple locations (sales_vpm.cpl and sales_vpm.xml).

5. Repeat the previous step for the second VPM file.

To Load VPM Files to a ProxySG through the Management Console:

1. Select Configuration>Policy>Policy Files>Visual Policy Files.

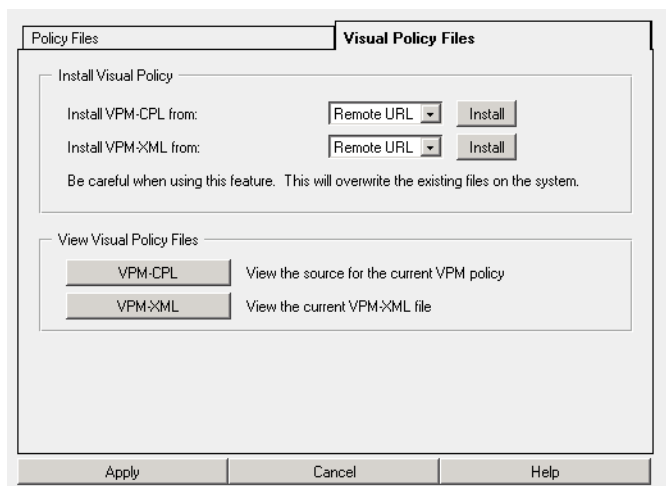


Figure 14-51: Visual Policy Files Tab

2. In the Install Visual Policy field:

- a. Select Remote URL from the Install VPM-CPL from drop-down list.
- b. Click Install. The Install VPM-CPL dialog appears.
- c. In the Installation URL field, enter the URL to the VPM CPL file copied to the Web server (this is the file with the default .txt extension) and click Install.
- d. Repeat Steps a through c to enter the URL to the second VPM XML file copied to the Web server (this is the file with the default .xml extension) and click Install.

3. Click Apply.

Notes

- If VPM files already exist on the ProxySG, the URLs to those files display in the two file fields. To replace them, delete the URLs and type new ones. Installing new files overwrites any that are already present.
- To review VPM-generated policies before installing them, enter the URL to the CPL file on the Web server and click View.

- Regardless of whether you are installing new VPM files, you can review the CPL or XML files of the policies currently on the ProxySG. Click VPM-CPL and VPM-XML in the View Visual Policy Files box at the bottom of the dialog.
- Never edit either of the VPM files directly. Change the files only by working with the policies in VPM and saving changes there.

To Load VPM Files to a ProxySG through the CLI:

The two commands in the first step load one of the VPM policy files; the commands in the second step load the other policy file. In each case, *url* is the complete path, including file name, to the appropriate file on the Web server.

1. At the `config` command prompt, enter the following commands:

```
SGOS#(config) policy vpm-cpl-path url
SGOS#(config) load policy vpm-cpl
```

2. At the `config` command prompt, enter the following commands:

```
SGOS#(config) policy vpm-xml-path url
SGOS#(config) load policy vpm-xml
```

Viewing the Policy/Generated CPL

View the CPL generated by installing VPM-created policy from VPM or the Management Console.

To View the Generated CPL through VPM:

Select View>Generated CPL.

To View the VPM Policy File:

Select View>Current ProxySG Appliance VPM Policy Files.

Important: Do *not* edit or alter VPM-generated files by opening the VPM policy file and working in the generated CPL. To edit, change, or add to VPM policies, edit the policy layers and re-install the policy.

Section E: Tutorials

This section contains the following topics:

- "Tutorial—Creating a Web Authentication Policy" on page 549
- "Tutorial—Creating a Web Access Policy" on page 554

Tutorial—Creating a Web Authentication Policy

This section is a tutorial that demonstrates how to create policies and rules for Web authentication.

Use Web Authentication policies to specify whether the individual making a request is prompted to authenticate by entering a username and password. In this example, a company uses a PAC file to configure most employee browsers to connect to a specific IP address on the ProxySG. It wants these users to authenticate when their browsers send a request to the proxy.

Create a Policy Layer

1. Start VPM from the Management Console, Configuration>Policy>Visual Policy Manager.
2. Select Policy>Add Web Authentication Layer.

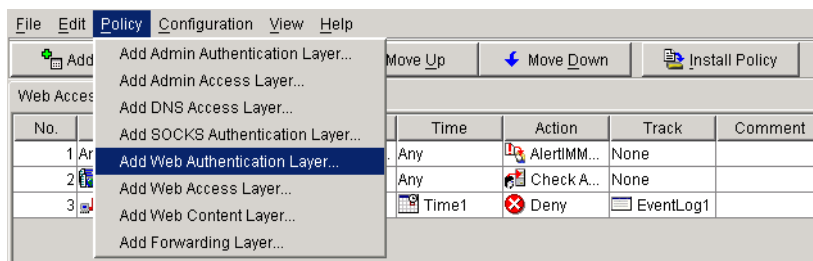


Figure 14-52: Creating a Policy Layer

3. A dialog displays offering a default name for the layer, consisting of the layer type and a number. Rename the layer or accept the default and click OK.




Figure 14-53: Add New Policy Table Dialog

4. VPM creates the new layer tab and adds a blank rule.

Example 1: Create an Authentication Rule

1. By default, the unmodified rule applies to everyone whose browsers connect to a specific IP address.

No.	Source	Destination	Action	Track	Comment
1	Any	Any	None	None	

Figure 14-54: Creating a Web Authentication Rule

2. Right-click the Source cell to drop the menu.
3. Select Set to open the Set Source Object dialog.

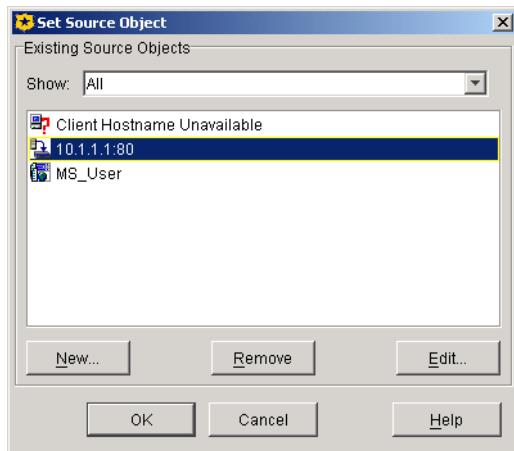


Figure 14-55: Selecting an IP Address Source Object

4. Select a proxy IP address or port; if necessary, click **New** to create a new one. This example selects IP address on the ProxySG where the PAC file sends most employee browsers.
5. Click **OK** to enter the IP address in the Source cell.

Web Access Layer (1)		Web Authentication Layer (1)			
No.	Source	Destination	Action	Track	Comment
1	10.1.1.1:80	Any	None	None	

Figure 14-56: Completed Source Object

6. Create an authentication Action object. Right-click the Action cell to drop the menu and select **Set**; the Set Action Object dialog displays.
7. The only objects available are the pre-existing static objects, so you must create a new **Authenticate** object. Click **New** and select **Authenticate**.

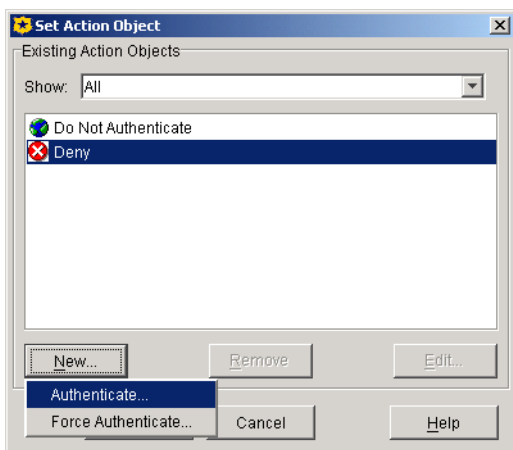


Figure 14-57: Selecting Authenticate

8. The Add Authenticate Object window displays. Note the following:
 - Name—Every configurable object has a name. The default name Authenticate1; change to Authenticate_XYZ_Corp, which is how it will be listed in the Add Object window.
 - Realm—Specifies the realm to be authenticated against from the drop-down list.
 - Mode—Specifies the authentication realm mode.




Figure 14-58: Specifying the Authentication Realm

9. Click OK to close the Add Action Object window, with the new Authenticate object in the list.




Figure 14-59: New Authentication Action Object

10. Click OK.

Web Authentication Layer (1)		Web Content Layer (1)		Forwarding Layer (1)		Web Access Layer (2)	
No.	Source	Destination	Action	Track	Comment		
1	Any	Any	Authenticate_X...	None			

Figure 14-60: Completed Action Object

11. Create a Trace object to log all authentication activity. Right-click the Track cell to drop the menu and select Set; the Set Track Object dialog appears.
12. You must create a new Trace object. Click New and select Trace; the Add Trace Object appears.

13. In the Name field, enter AuthTrace.
14. Name the object AuthTrace1. Click Trace Level and Verbose to enable verbose tracing, which lists the rules that were skipped because one or more of their conditions were false and displays the specific condition in the rule that was false.

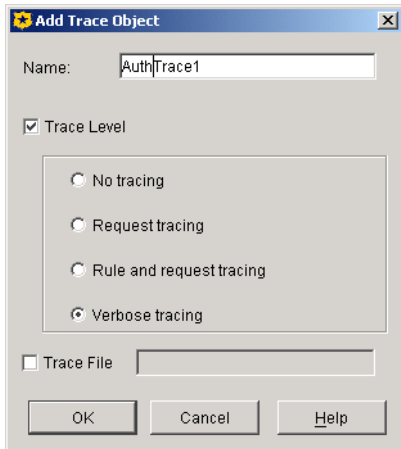


Figure 14-61: Creating a Trace Object

15. Click OK; click OK again to add the object. The rule is complete.

Web Access Layer (1)		Web Authentication Layer (1)				
No.	Source	Destination	Action	Track	Comment	
1	10.1.1.1:80	Any	Authenticate_X...	AuthTrace		

Figure 14-62: Completed Rule

Example 2: Exempt Specific Users from Authentication

Certain individuals and groups are exempt from the above restriction. Individuals in the purchasing department are allowed to the Web so they can order online from supplier Web sites. And the company does not want them to authenticate.

1. To create a new rule, click Add Rule.

Web Access Layer (1)		Web Authentication Layer (1)				
No.	Source	Destination	Action	Track	Comment	
1	10.1.1.1:80	Any	Authenticate_X...	AuthTrace		
2	Any	Any	None	None		

Figure 14-63: Adding a Second Web Authentication Policy Layer Rule

2. People in the purchasing group use the same PAC file and thus their browsers are directed to the same IP address. A Combined Source Object is needed that includes the purchasing group, who are their own subnet.

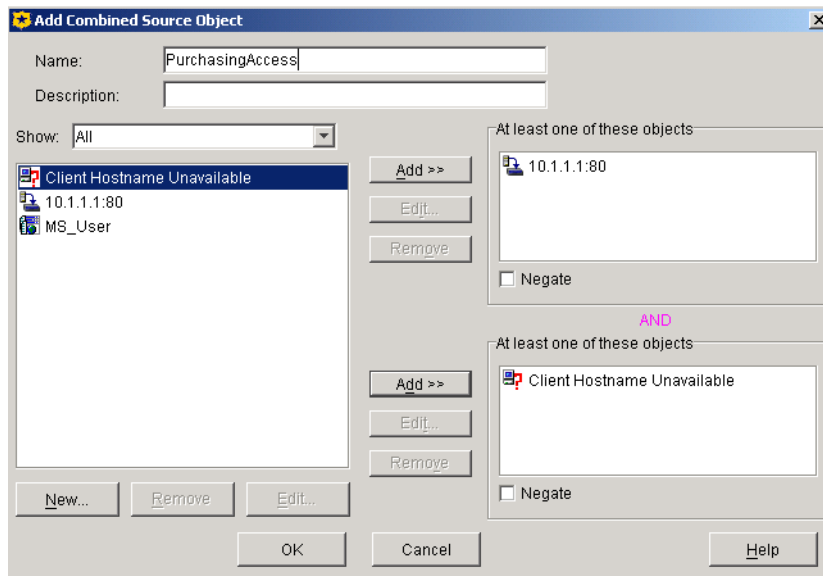


Figure 14-64: A Combined Object

The new rule in the policy layer accepts the default Action Object to not authenticate and does not require a Trace Object.

Web Access Layer (1)		Web Authentication Layer (1)			
No.	Source	Destination	Action	Track	Comment
1	10.1.1.1:80	Any	Authenticate_X...	AuthTrace	
2	PurchasingAcc...	Any	None	None	

Figure 14-65: Updated Second Rule

However, a problem exists. The second rule cannot be evaluated because the first rule affects everyone who goes through the proxy. The rules need to be reversed.

3. Select the second rule and click Move Up to reorder the rules.

Web Access Layer (1)		Web Authentication Layer (1)			
No.	Source	Destination	Action	Track	Comment
1	PurchasingAcc...	Any	None	None	
2	10.1.1.1:80	Any	Authenticate_X...	AuthTrace	

Figure 14-66: Reordered Rules

Tutorial—Creating a Web Access Policy

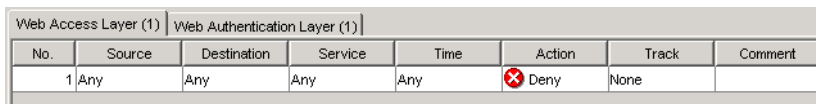
This section is a tutorial that demonstrates how to create policies and rules for Web access.

Use ProxySG policies to define end-user access to Web resources. For more information about Web access policies, see Appendix 18: "Content Filtering". This section provides examples.

Example 1: Restrict Access to Specific Websites

This example demonstrates a simple rule that denies everyone access to specific job searching Web sites. This rule requires you to configure only one rule option; it uses the defaults for all other options.

1. Start the policy editor and select Policy>Add Web Access Layer. VPM displays a tab with the name of the new policy; beneath that is a new rule-specific row. Notice that the default Action is Deny.



No.	Source	Destination	Service	Time	Action	Track	Comment
1	Any	Any	Any	Any	Deny	None	

Figure 14-67: Creating a New Policy Layer

2. Right-click Destination and select Set; the Set Destination Object dialog appears.

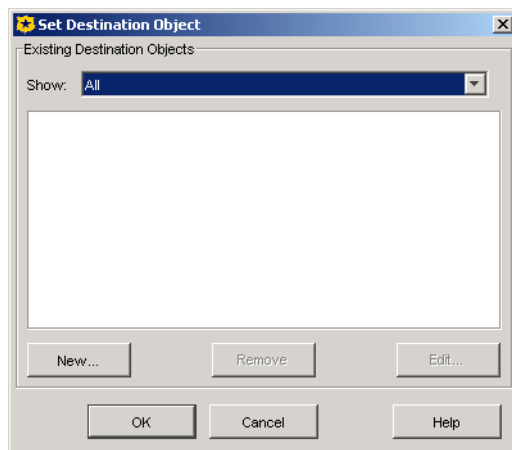


Figure 14-68: Set Destination Dialog

3. Click New; select Combined Destination Object. The Add Combined Destination Object dialog appears.

- a. Name the object JobSearchURLs.
- b. Select New>URL. The Add URL Object dialog appears.

Add URL Object

Name :

Simple Match

URL:

If the host specified is a domain name, all hosts in that domain (or any subdomain) will match. If a path is specified, all paths with that prefix will match. If a scheme or port number is specified, only URLs with that scheme or port will match.

Regular Expression Match

RegEx:

Advanced Match

Scheme:

Host:

Port: e.g. 80 or 1800-2000

Path:

Figure 14-69: Add URL Object Dialog

- c. Click Simple Match; in the URL field, enter hotjobs.com and click Add.
- d. Repeat Step c twice, entering bajobs.com and monster.com.
- e. Click Close; the entered URLs appear as selectable items.

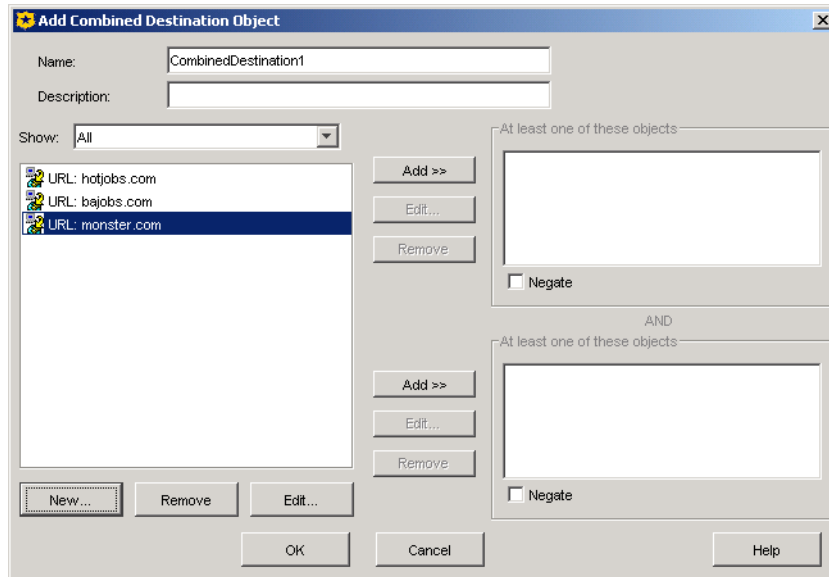


Figure 14-70: Jobsite URLs

- 4. Select each newly added URL and click the first Add button.

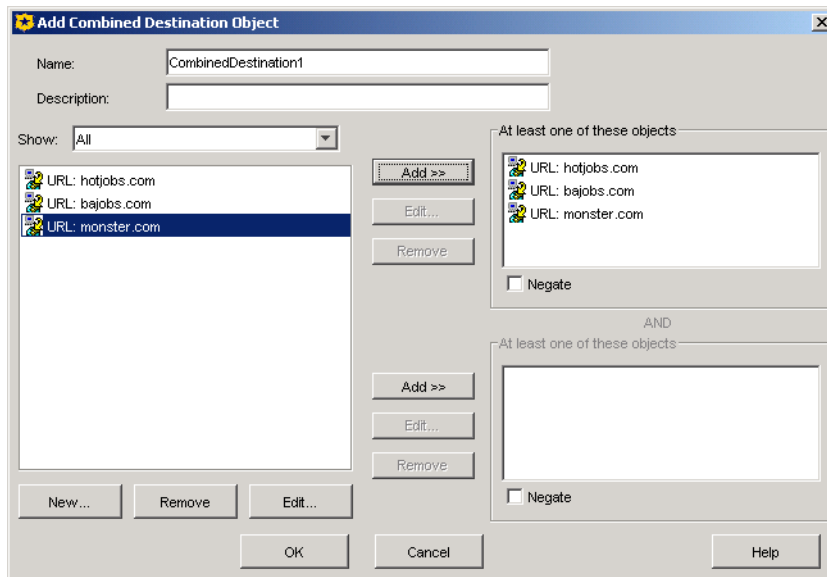


Figure 14-71: Creating a Combined Destination Object with URLs

- 5. Click OK. The Set Destination Object now contains the individual URL objects and the combined object.

6. Select the JobSearchURLs combined object and click OK. The object is now part of the rule.


Web Access Layer (1)		Web Authentication Layer (1)					
No.	Source	Destination	Service	Time	Action	Track	Comment
1	Any	 JobSearchURLs	Any	Any	 Deny	None	

Figure 14-72: Completed Rule

As the default action is deny, the rule is complete. No one can access these Web sites.

7. To activate the rule, click Install Policies.

Example 2: Allow Specific Users to Access Specific Websites

The after-hours IT shift is comprised of part-time college interns who are on call to handle small problems, but are not involved in major projects. Therefore, you allow them to browse certain sports and entertainment Web sites when all is quiet; access is allowed from two workstations and you still want to track their browsing activity.

Configuring the Source Object

1. Add a new rule to the policy and position the pointer in the Source cell.


Web Access Layer (1)		Web Authentication Layer (1)					
No.	Source	Destination	Service	Time	Action	Track	Comment
1	Any	 JobSearchURLs	Any	Any	 Deny	None	
2	Any	Any	Any	Any	 Deny	None	

Figure 14-73: Setting a Source Rule Option

2. Right-click the Source cell and select Set to display the Add Source Object dialog.
3. Click New and select Combined Source Object; the Add Combined Source Object appears.
4. Name the object IT_PM_Shift.
5. Under the selectable list of objects, click New and select Client IP Address/Subnet; the Add Client IP Address/Subnet Object dialog appears.
6. Enter the IP address of the first workstation and click Add; repeat for the second; click Close.
7. Select each IP address and click the first Add.

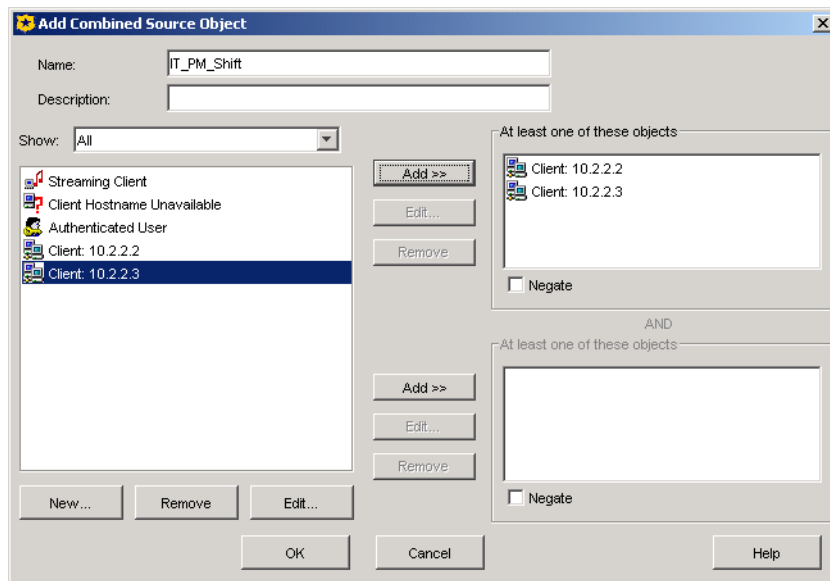


Figure 14-74: Combined IP Address Object

8. Click OK; click OK again to add the Source object to the rule.

Configuring the Destination Object

1. Right-click the Destination field and select Set; the Set Destination Object dialog appears.
2. Click New and select Category; the Add Category Object dialog appears.

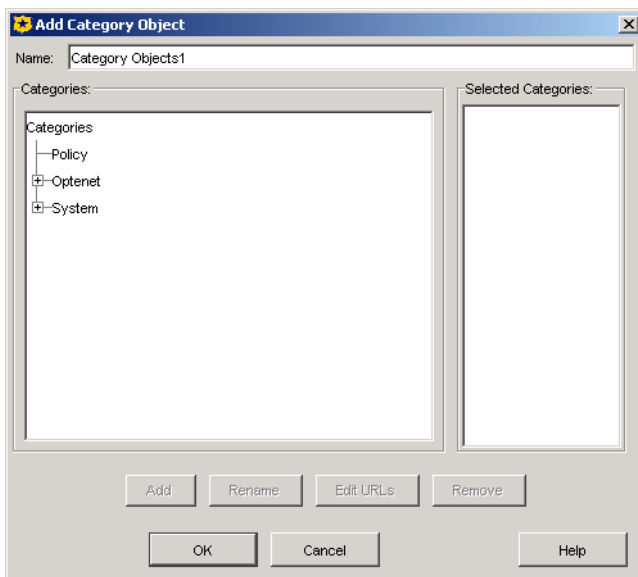


Figure 14-75: Category Dialog

3. Select Policy and click Add; the Enter Name for New Category dialog appears.
4. Name the object Allowable_Sports and click OK.

5. Select Sports URLs and click Edit URLs. The Edit Locally Defined Category Object dialog appears.
6. Enter the URLs for the allowable sports Web sites.

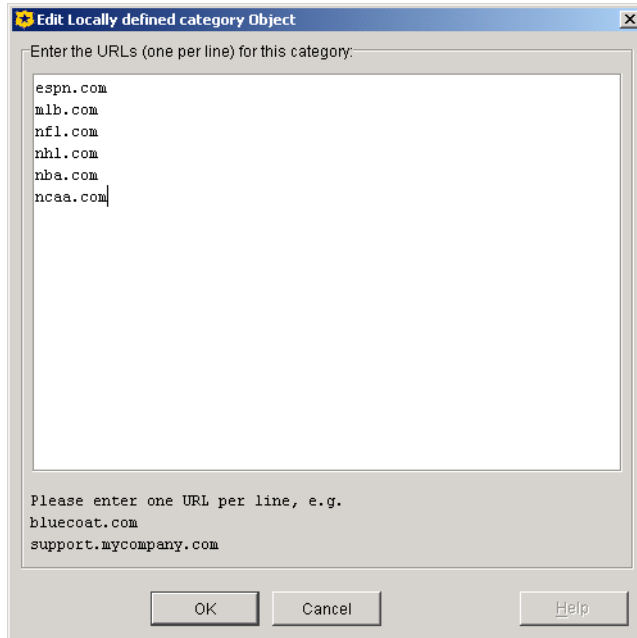


Figure 14-76: Sports URLs Category

7. Click OK. Under Policy, select Allowable_Sports; click OK.
8. Repeat Steps 3 through 7, creating a category called Allowable_Entertainment with the URLs ew.com, rollingstone.com, and variety.com.
9. Create a Combined Object named Allowable PM IT Websites with the two categories: Allowable_Sports and Allowable_Entertainment. Click OK twice to add the object to the rule.

Configuring the Time Object

This example allows the specified users to access the sports and entertainment Websites after business hours.

1. Right-click the Time field and select Set; the Set Time Object dialog appears.
2. Click New and select Time Object; the Add Time Object dialog appears.
3. Name the object After Hours.
4. Select Local Time. In the Specific Time of Day Restriction field, select Enable and set the time from 18:00 to 05:59.

This defines after hours as 6:00 PM to 6:00 AM.

5. In the Specific Weekday Restriction field, select Enable and select Monday, Tuesday, Wednesday, Thursday, and Friday.

This defines the days of the week to which this rule applies.

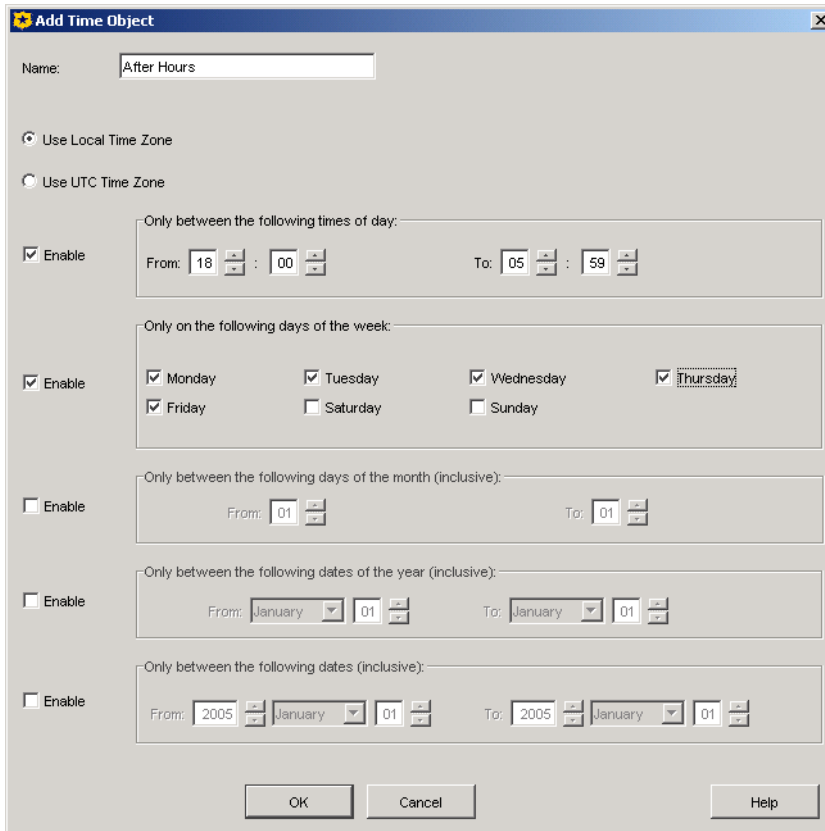


Figure 14-77: After Hours Time Object

6. Click OK to add the Time Object to the rule.

Configuring the Action Option

1. Right-click Action select Set; the Set Action Object dialog appears.
2. Allow is a static object. Select Allow and click OK.

The rule is now complete.

Web Access Layer (1)		Web Authentication Layer (1)					
No.	Source	Destination	Service	Time	Action	Track	Comment
1	Any	JobSearchURLs	Any	Any	Deny	None	
2	IT_PM_Shift	Allowable PM IT ...	Any	After Ho...	Allow	None	

Figure 14-78: Completed Rule

Chapter 15: *Advanced Policy*

This chapter provides conceptual information about ProxySG advanced policy features. While many Blue Coat Systems features have a policy component, some features have no configuration component outside policy. Configuring advanced policy is done by defining rules in the Visual Policy Manager (VPM) or by composing Content Policy Language (CPL). While some examples are provided in this chapter, references to the relevant VPM chapter component are included in each section.

This chapter contains the following topics:

- "Blocking Pop Up Windows"
- "Stripping or Replacing Active Content"
- "Modifying Headers"
- "Defining Exceptions"
- "Managing Peer-to-Peer Services"

Excluding exceptions, you *must* use policy to implement these capabilities. (For exceptions, you can create a list outside of policy to install on the system.)

Section A: Blocking Pop Up Windows

Section A: Blocking Pop Up Windows

This section describes the Blue Coat solution for blocking unwanted pop up windows.

About Pop Up Blocking

The ProxySG allows you to block pop up windows, which are usually in the form of unsolicited advertisements. Pop up windows are blocked by inserting Javascript code into each HTML Web page. Every time the Web page tries to open a new window, the code attempts to determine if the window is a result of user click. The window is allowed to open if the ProxySG determines a user clicked a button or link; otherwise, the window does not open.

Limitations

Because of the dynamic nature of the Web, blocking pop up windows is not a perfect solution. Keep in mind the following limitations before configuring this feature:

- Windows that contain desired or useful information cannot be distinguished from undesired content, such as advertisements.
- If the Web browser caches a page that spawns pop up windows before the blocking policy was installed, pop up ads continue to be served from that page regardless of current policy.
- Animated ads contained within Web pages are not blocked. Commonly seen in scrolling or drop-down form, these are not true pop up windows but are contained within the page. Users who see these ads might believe that pop up window blocking is not implemented.
- Pop up windows that are delivered through HTTPS are not blocked.
- Although the Blue Coat request headers tell a Web server not to use compression, it is possible (though not likely) for a Web server to be configured to send compressed responses anyway. The pop up blocking feature does not work on compressed HTML pages.

Recommendations

- To compensate for limiting factors, administrators and users can override pop up blocking:
 - ❑ Administrators—Use VPM to create policy rules that exempt pop up blocking for specific Web sites and IP address ranges. For example, Blue Coat recommends disabling pop up blocking for your Intranet, which commonly resides on a IP address range.
 - ❑ Users—When a pop up window is blocked, a message is displayed in the status bar:
`blocked popup window -- use CTRL Refresh to see all popups.`
While pressing the Control key, click the Web browser Refresh button; the page is reloaded with pop up blocking disabled for that action.
- Create a separate Web Access policy layer for pop up blocking actions. This alleviates interference with Web applications deployed on your Intranet that require pop up windows.
- To prevent a cached Web page from spawning pop up windows, clear the browser cache, then reload the page without holding down the `CTRL` key.

Section A: Blocking Pop Up Windows

Blocking pop up windows is accomplished through the Visual Policy Manager. See "Block/Do Not Block PopUp Ads" in Chapter 14: "The Visual Policy Manager" on page 453 for information about how to create blocking actions in a policy layers.

Section B: Stripping or Replacing Active Content

Section B: Stripping or Replacing Active Content

This section describes the Blue Coat solution for stripping or replacing unwanted active content.

About Active Content

Scripts activated within Web pages can pose a security concern. The ProxySG policy can be configured to supplement standard virus scanning of Web content by detecting and removing the HTML tags that launch active content such as Java applets or scripts. In addition, the removed content can be replaced with predefined material, a process referred to as *active content transformation*.

When the ProxySG is configured to perform active content transformation, Web pages requested by a client are scanned before they are served and any specified tags and the content they define are either removed or replaced. Since the transformed content is not cached, the transformation process is based on a variety of conditions, including time of day, client identity, or URL.

Note: Pages served over an HTTPS tunneled connection are encrypted, so the content cannot be modified.

The following tags and related content can be removed or replaced:

- `<APPLET>`—Java applets, as defined by HTML `<applet>` elements.
- `<EMBED>`—Embedded multimedia objects displayed using Netscape Navigator plug-ins as defined by HTML `<embed>` elements.
- `<OBJECT>`—Embedded multimedia objects displayed using Internet Explorer Active-X controls and other multimedia elements, as defined by HTML `<object>` elements
- `<SCRIPT>`—Embedded Javascript and VBScript programs, whether these are represented as HTML `<script>` elements, Javascript entities, Javascript URLs, or event handler attributes. The `<noscript>` tag is *not* affected by this features.

Stripping active content is accomplished through the Visual Policy Manager or by composing CPL.

- See "Strip Active Content" in Chapter 14: "The Visual Policy Manager" on page 453 for information about how to create a strip active content object in a Web Access policy layer.
- Refer to the *Blue Coat ProxySG Content Policy Language Guide*.

About Active Content Types

The following sections provide more detail about the types of active content that can be removed or replaced.

Script Tags

Scripts are generally placed between the start and end tags `<SCRIPT>` and `</SCRIPT>`. The type of script used is defined by the `LANGUAGE` attribute; for example, `<SCRIPT LANGUAGE="JavaScript 1.0">`). When the `LANGUAGE` attribute is undefined, the browser assumes JavaScript.

Section B: Stripping or Replacing Active Content

When `transform active_content` is configured to remove scripts, the basic operation is to remove all content between and including `<SCRIPT>` and `</SCRIPT>`, regardless of the language type, and substitute any defined replacement text. A notable exception when a script is defined in the header portion of the HTML document (defined by the `<HEAD>` tag). In this case, the script is simply removed. This is because images, objects, and text are not allowed in the header of an HTML document. If the end script tag `</SCRIPT>` is missing from the document (the end of the document is defined as either up to the `</BODY>` or `</HTML>` tag, or the last character of the document), then all content from the start `<SCRIPT>` tag to the end of the document is removed.

JavaScript Entities

JavaScript entities have the following format: `&{javascript code}` and are found anywhere in the value part of an attribute (that is, ``). You can define more than one entity in the value portion of the attribute. When `transform active_content` is configured to remove scripts, all JavaScript entities attribute/value pairs are removed. No replacement text is put in its place.

JavaScript Strings

JavaScript strings have the following format: `javascript: javascript code` and are found anywhere in the value part of an attribute, though usually only one of them can be defined in an attribute. Most modern browsers support JavaScript strings. When `transform active_content` is configured to remove scripts, all JavaScript string attribute/value pairs are removed. No replacement text is put in its place.

JavaScript Events

JavaScript events are attributes that start with the keyword `on`. For example, ``. The HTML 4.01 specification defines 21 different JavaScript events:

```
onBlur, onChange, onClick, onDblClick, onDragDrop, onFocus, onKeyDown,
onKeyPress, onKeyUp, onLoad, onMouseDown, onMouseMove, onMouseOut,
onMouseOver, onMouseUp, onMove, onReset, OnResize, onSelect, onsubmit,
onUnload
```

Both Microsoft Internet Explorer and Netscape have defined variations on these events as well as many new events. To catch all JavaScript events, the active content transformer identifies any attribute beginning with the keyword `on`, not including `on` itself. For example, the attribute `onDonner` in the tag `` is removed even though `onDonner` does not exist as a valid JavaScript event in the browser. In this case, the transformed file would show ``.

Section B: Stripping or Replacing Active Content

Embed Tags

HTML `<EMBED>` tags are not required to have an `</EMBED>` end tag. Many Web browsers do, however, support the `<EMBED>` `</EMBED>` tag pair. The text between the tags is supposed to be rendered by the browsers when there is no support for the embed tag, or if the MIME-type of the embed object is not supported. Thus, when `transform active_content` is configured to transform embed tags, only the `<EMBED>` tag is removed and replaced with any replacement text. Any occurrence of the end tag `</EMBED>` is simply removed, leaving the text between the beginning and end tags intact.

Object Tags

Objects tags have a start `<OBJECT>` and end `</OBJECT>` tag pair, and the attributes `CODETYPE` and `TYPE` determine the type of object. The text between the tags is supposed to be rendered by the browsers when the object tag is not supported, so when `transform active_content` is configured to transform object tags, only the `<OBJECT>` and `</OBJECT>` tags are removed and replaced with any replacement text. The text between the tags remains. The `CODETYPE` or `TYPE` attributes do not affect the transformation. Also, if the end `</OBJECT>` tag is missing, the transformation will not be affected.

Section C: Modifying Headers

Section C: Modifying Headers

The request headers are sent when users access Web objects that contain a lot of information. This can raise a concern that such details compromise the privacy or security of the enterprise or user.

When a user clicks on a link, the Web browser sets the request's Referer header to the URL of the Web page that contained the link. (This header is not set if the URL was entered or selected from a favorites or bookmarks list.) If an internal Web page provides links to external Web sites, users clicking those links sends the URL of the internal pages, and are logged in the Web logs of those external sites. This is not usually an issue; however, if the external Web site is a competitor Web site or another site with interest in the internal details of your enterprise, this might be a concern.

For example, how you structure your intranet might suggest something about your company's current or future direction. Certain project names or codewords might show up in directory or file names. Exposing the structure of the intranet makes it easier for hackers to attack the network.

The broad solution of deleting Referer headers from all requests presents a problem because some Web sites will not serve images or other linked objects unless the Referer header is set to a referring page on that same Web site. The solution implemented by Blue Coat is to strip the Referer header only when the target Web page resides on the Internet and the referring page is on an internal host.

Suppressing headers is accomplished through the Visual Policy Manager or by composing CPL.

- See "Suppress Header" in Chapter 14: "The Visual Policy Manager" on page 453 for information about how to create a strip active content object in a Web Access policy layer.
- Refer to the *Blue Coat ProxySG Content Policy Language Guide*.

Section D: Defining Exceptions

Section D: Defining Exceptions

Exceptions (formerly called message or RMG pages) are sent in response to certain ProxySG client requests, such as denial by policy, failure to handle the request, and authentication failure. Exceptions are returned to users based on policy rules defined by the ProxySG administrator. For example, if a client sends a request for content that is not allowed, an exception HTML page (for HTTP connections) or an exceptions string (for non-HTTP connections) is returned, informing the client that access is denied.

Two types of exceptions are used: built-in and user-defined.

Built-in Exceptions

Built-in exceptions are a set of pre-defined exceptions contained on the ProxySG. Built-in exceptions send information back to the user under operational contexts that are known to occur, such as *policy_denied* or *invalid_request*.

Built-in exceptions are always available and can also have their contents customized; however, built-in exceptions cannot be deleted, and you cannot create new built-in exceptions.

The table below lists the built-in exceptions and the context under which they are issued.

Table 15.1: Built-in Exceptions

Exception Type	Issued When
<code>authentication_failed</code>	The transaction cannot be authenticated, usually because the credentials were incorrect. <code>authentication_failed</code> is a synonym for <code>deny.unauthorized</code> .
<code>authentication_failed_password_expired</code>	The authentication server reports that the credentials provided have expired, and a new password must be obtained.
<code>authentication_mode_not_supported</code>	The configured authentication mode is not supported for the current request.
<code>authentication_redirect_from_virtual_host</code>	Transparent redirect authentication is being used. This exception redirects the transaction from the virtual authentication host to the original request.
<code>authentication_redirect_off_box</code>	The request is being redirected to an authentication service on another device.
<code>authentication_redirect_to_virtual_host</code>	Transparent redirect authentication is being used. This exception redirects the transaction to the virtual authentication host.
<code>authentication_success</code>	Transparent redirect authentication with cookies is being used. This exception redirects the transaction to the original request, but removes the authentication cookie from the request URL.

Section D: Defining Exceptions

Table 15.1: Built-in Exceptions (Continued)

Exception Type	Issued When
<code>authorization_failed</code>	The <code>deny.unauthorized</code> policy action is matched. This exception notifies the user that their currently authenticated identity is not permitted to perform the requested operation, but they may have some other credentials that would allow their request through (for example, they get an opportunity to enter new credentials).
<code>client_failure_limit_exceeded</code>	Too many requests from your ip address (<code>\$(client.address)</code>) have failed.
<code>configuration_error</code>	A configuration error on the ProxySG was detected, and the requested operation could not be handled because of the configuration error. This exception is a likely indicator that the administrator of the ProxySG will have to intervene to resolve the problem.
<code>connect_method_denied</code>	A user attempted an CONNECT method to a non-standard port when explicitly proxied. Blue Coat does not allow CONNECT methods to non-standard ports by default because it is considered a security risk to do so.
<code>content_filter_denied</code>	A particular request is not permitted because of its content categorization.
<code>content_filter_unavailable</code>	An external content-filtering service could not be contacted, and the ProxySG is failing closed in such a situation.
<code>dns_server_failure</code>	The request could not be processed because the ProxySG was unable to communicate with the DNS server in order to resolve the destination address of the request.
<code>dns_unresolved_hostname</code>	The request could not be processed because the ProxySG was unable to resolve the hostname in the request with DNS.
<code>dynamic_bypass_reload</code>	The <code>dynamic_bypass</code> policy action is matched.
<code>gateway_error</code>	There was a network error while attempting to communicate with the upstream gateway.
<code>icap_communication_error</code>	A network error occurred while the ProxySG was attempting to communicate with an external ICAP server.
<code>internal_error</code>	The ProxySG encountered an unexpected error that resulted in the inability to handle the current transaction.
<code>invalid_auth_form</code>	The submitted authentication form is invalid. The form data must contain the username, password, and valid original request information.
<code>invalid_request</code>	The request received by the ProxySG was unable to handle the request because it detected that there was something fundamentally wrong with the syntax of the request.

Section D: Defining Exceptions

Table 15.1: Built-in Exceptions (Continued)

Exception Type	Issued When
license_expired	The requested operation cannot proceed because it would require the usage of an unlicensed feature.
method_denied	The requested operation utilizes a method that has been explicitly denied because of the service properties associated with the request.
not_implemented	The protocol cannot handle the requested operation because it utilizes a feature that is not currently implemented.
notify	Used internally by VPM. You do not need to customize the text of this exception, since in this case the entire HTML response is generated by VPM and is not taken from the exception definition.
notify_missing_cookie	This exception is returned when a VPM Notify User action is being used to notify the user, and the user has disabled cookies in the web browser.
policy_denied	<code>policy_denied</code> is a synonym for <code>deny</code> .
policy_redirect	A <code>redirect</code> action is matched in policy.
radius_splash_page	The RADIUS/TACACS splash generator feature is in use, and the user must be authorized by RADIUS/TACACS. (The RADIUS/TACACS secrets must be configured through the <code>(config) splash-generator</code> commands.)
redirected_stored_requests_not_supported	This applies to forms authentication with POST requests only): The origin server returned a redirect for the request. The ProxySG is configured to not allow stored requests to be redirected.
refresh	A refresh (using the <code>HTTP Refresh:</code> header) is required. The refresh exception (by default) refreshes the originally requested URL (or in some cases, its post-imputed form).
server_request_limit_exceeded	Too many simultaneous requests are in progress to <code>\$(url.host)</code> .
silent_denied	An exception (<code>silent_denied</code>) is matched in policy. This exception is pre-defined to have no body text, and is “silent” in that it results in only the status code being sent to the client.
ssl_domain_invalid	There was a failure contacting an upstream host through HTTPS because the certificate presented by the upstream host was either the incorrect one or invalid.
ssl_failed	A secure connection could not be established to an upstream host. This is typically because the upstream host is not configured to accept SSL connections.

Section D: Defining Exceptions

Table 15.1: Built-in Exceptions (Continued)

Exception Type	Issued When
tcp_error	A network error occurred attempting to communicate with an upstream host.
transformation_error	The server sends an unknown encoding and the ProxySG is configured to do content transformation.
unsupported_encoding	The client makes a request with an “Accept-Encoding: Identity;q=0, ...” header. Only uncompressed content is available in cache, the ProxySG is not configured to compress the content, or the compression license is expired, or the client request results in to “Accept-Encoding: Identity;q=0” due to the combination of request and configured policy.
unsupported_protocol	The protocol used in the request is not understood.

Most of the above exceptions can be initiated directly through the policy exception property. However, some require additional state that makes initiating them either problematic or out of context. The following are exceptions that cannot be initiated through the exception property:

- authentication_failed
- authentication_failed_password_expired
- authentication_redirect_from_virtual_host
- authentication_redirect_to_virtual_host
- authentication_success
- dynamic_bypass_reload
- license_expired
- radius_splash_page
- ssl_domain_invalid
- ssl_failed

To view the content of a built-in exception, enter the following commands at the (config) prompt:

```
SGOS#(config) exceptions
SGOS#(config exceptions) show exceptions configuration_error
configuration_error exception:
all protocols:
summary text:
    ProxySG configuration error
details text:
    Your request could not be processed because of a configuration error:
$(exception.last_error)
help text:
```

Section D: Defining Exceptions

```
The problem is most likely because of a configuration error,  
$(exception.contact) and provide them with any pertinent information from  
this message.  
http protocol:  
code: 403
```

User-Defined Exceptions

User-defined exceptions are created and deleted by the administrator. If a user-defined exception is referenced by policy, it cannot be deleted. The default HTTP response code for user-defined exceptions is 403.

Note: For users who are explicitly proxied and use Internet Explorer to request an HTTPS URL, an exception body longer than 900 characters might be truncated. The workaround is to shorten the exception body.

Note also that an exception body less than 512 characters might cause a *page does not exist* 404 error. If this occurs, use the `exception.autopad(yes|no)` property to pad the body to more than 513 characters. For more information on the `exception.autopad` property, refer to the *Blue Coat ProxySG Content Policy Language Guide*.

About Exception Definitions

Each exception definition (whether built-in or user-defined) contains the following elements:

- **Identifier**—Identifies the type of exception. Table 15.1 on page 568 lists the built-in exception types. For user-defined exceptions, the identifier is the name specified upon creation.
- **Format**—Defines the appearance of the exception. For an HTTP exception response, the format is an HTML file. For other protocols, where the user agents are not able to render HTML, the format is commonly a single line.
- **Summary**—A short description of the exception that labels the exception cause. For example, the default `policy_denied` exception summary is "Access Denied".
- **Details**—The default text that describes reason for displaying the exception. For example, the default `policy_denied` exception (for the HTTP protocol) detail is: Your request has been denied by system policy.
- **Help**—An informative description of common possible causes and potential solutions for users to take. For example, if you want the categorization of a URL reviewed, you can append the `$(exception.category_review_url)` and `$(exception.category_review_message)` substitutions to the `$(exception.help)` definition. Note that you must first enable this capability through content filtering configuration. For information on enabling review categorization, see "Selecting Category Providers" on page 636.
- **Contact**—Used to configure site-specific contact information that can be substituted in all exceptions. Although it is possible to customize contact information on a per-exception basis, customizing the top-level contact information, which is used for all exceptions, is sufficient in most environments.

Section D: Defining Exceptions

- **HTTP-Code**—The HTTP response code to use when the exception is issued. For example, the `policy_denied` exception by default returns the 403 `Forbidden` HTTP response code.

Important: Fields other than `Format` must be less than 8000 characters. If they are greater than this, they will not be displayed.

When defining the above fields, you can use substitution variables that are particular to the given request. Some of the above fields are also available as substitutions:

- `$(exception.id)`
- `$(exception.summary)`
- `$(exception.details)`
- `$(exception.help)`
- `$(exception.contact)`

Additionally, the `Format`, `Summary`, `Details`, `Help` and `Contact` fields can be configured specifically for HTTP, or configured commonly for all protocols.

The `Format` field, the body of the exception, is not available as a substitution. However, the `Format` field usually includes other substitutions. For example, the following is a simple HTML format:

```
<html>
<title>$(exception.id): $(exception.summary)</title>
<body><pre>
Request: $(method) $(url)
Details: $(exception.details)
Help: $(exception.help)
Contact: $(exception.contact)
</pre></body></html>
```

Some additionally useful substitutions related to exceptions are:

- `$(exception.last_error)`—For certain requests, the ProxySG determines additional details on why the exception was issued. This substitution includes that extra information.
- `$(exception.reason)`—This substitution is determined internally by the ProxySG when it terminates a transaction and indicates the reason that the transaction was terminated. For example, a transaction that matches a DENY rule in policy has its `$(exception.reason)` set to "Either 'deny' or 'exception' was matched in policy".

About the Exceptions Hierarchy

Unlike the error pages in previous SGOS releases, exceptions are not required to have its entire contents defined. Exceptions are stored in a hierarchical model, and *parent* exceptions can provide default values for *child* exceptions. There are two parent exceptions from which other exceptions are derived: `exception.all` and `exception.user-defined.all`.

Each built-in and user-defined exception derives its default values from the `all` exception. For example, by default the built-in exceptions do not define the `format` field. Instead, they depend on the `all` exception's `format` field definition. To change the format text for all built-in and user-defined exceptions, customize the `format` field for the `all` exception.

Section D: Defining Exceptions

The `user-defined.all` exception is the parent of all user-defined exceptions, but it is also a child of the `all` exception. Configuring `exception.user-defined.all` is only necessary if you want certain fields to be common for all user-defined exceptions, but not common for built-in exceptions.

The following example demonstrates using the `exception inline` command to configure the `$(exception.contact)` substitution for every HTTP exception:

```
!(config exceptions) inline http contact EOF
For assistance, contact <a
href="mailto:sysadmin@example.com">sysadmin</a>EOF
```

The following example configures a different `$(exception.contact)` substitution for every HTTP exception:

```
!(config exceptions) user-defined inline http contact EOF
For assistance, contact <a
href="mailto:policyadmin@example.com">policyadmin</a>EOF
```

About the Exceptions Installable List

The Exceptions Installable List uses the Structured Data Language (SDL) format. This format provides an effective method to express a hierarchy of key/value pairs. For example, the following is SDL file before customization:

```
(exception.all
  (format "This is an exception: $(exception.details)")
  (details "")
  (exception.policy_denied
    (format "")
    (details "your request has been denied by system policy")
  )
)
```

This SDL file defines an exception called `policy_denied` that defines the `$(exception.details)` substitution as "Your request has been denied by system policy". Because the exception does not define the `format` field, it inherits the `format` field from its parent exception (`exception.all`). When the `policy_denied` exception is issued, the resulting text is: This is an exception: your request has been denied by system policy.

Suppose you want to customize the `$(exception.contact)` substitution for every HTTP exception. Edit the `exception.all` component.

Note: The default HTTP format and built-in exception definitions have been removed for example purposes.

```
(exception.all
  (contact "For assistance, contact your network support team.")
  (details "")
  (format "$(exception.id): $(exception.details)")
  (help "")
  (summary ""))
```

Section D: Defining Exceptions

```

    (http
      (code "200")
      (contact "")
      (details "")
      (format <<EOF
<format removed>
      EOF
    )
    (help "")
    (summary "")
  )
<built-in exceptions removed>
)

```

To add the `$(exception.contact)` information, modify the `contact` substitution under the `http` node:

```

(exception.all
  (contact "For assistance, contact your network support team.")
  (details "")
  (format "$(exception.id): $(exception.details)")
  (help "")
  (summary "")
  (http
    (code "200")
    (contact "For assistance, contact <a
      href='mailto:sysadmin@example.com'>sysadmin</a>") EOF
    (details "")
    (format <<EOF
<format removed>
    EOF
  )
  (help "")
  (summary "")
)
<built-in exceptions removed>
)

```

Keep in mind the following conditions when modifying exception installable lists:

- Every exception installable list must begin with a definition for `exception.all`.
- In the exceptions' installable list, all definitions must be enclosed by `exception.all` and its accompanying closing parenthesis; that is,


```

(exception.all
(exception.policy_denied)
)

```
- Keep the definition strings under the enclosed parentheses short, no longer than one line if possible.
- Blue Coat strongly recommends downloading the existing exceptions installable list, then modifying it.

Section D: Defining Exceptions

Creating or Editing Exceptions

You can create or edit an exception with the CLI or with installable lists on the Management Console.

Note: You cannot create user-defined exceptions for Patience Pages.

To Create or Edit an Exception through the CLI:

1. At the (config) prompt, enter the following commands:

```
SGOS#(config) exceptions
SGOS#(config exceptions) create definition_name
SGOS#(config exceptions) edit definition_name
SGOS#(config exceptions user-defined.definition_name) http-code numeric HTTP
response code
SGOS#(config exceptions user-defined.definition_name) inline ?
  contact    Set the $(exceptions.contact) substitution
  details    Set the $(exceptions.details) substitution
  format     Set the format for this exception
  help       Set the $(exceptions.help) substitution
  http       Configure substitution fields for just HTTP exceptions
  summary    Set the $(exception.summary) substitution
SGOS#(config exceptions user-defined.definition_name) inline contact eof
string eof
SGOS#(config exceptions user-defined.definition_name) inline details eof
string eof
SGOS#(config exceptions user-defined.definition_name) inline format eof
string eof
SGOS#(config exceptions user-defined.definition_name) inline help eof
string eof
SGOS#(config exceptions user-defined definition_name) inline summary eof
string eof
```

2. (Optional) View the results.

```
SGOS#(config exceptions user-defined.test) show exceptions user-defined.test
$(exception.id):
  test
$(exception.summary):
  Connection failed
$(exception.details):
  Connection failed with stack error
$(exception.contact):
  Tech Support
```

To Delete a User-Defined Exception:

From the (config) prompt, enter the following commands:

Section D: Defining Exceptions

```
SGOS#(config) exceptions
SGOS#(config exceptions) delete exception_name
ok
```

Note: You cannot delete a user-defined exception that is referenced by policy. You must remove the reference to the exception from the policy before deleting the exception.

Using the Management Console to Create and Install an Exceptions List

The Management Console allows you to create and install exceptions with the following methods:

- Using the ProxySG Text Editor, which allows you to customize the existing exceptions file.
- Creating a local file on your local system; the ProxySG can browse to the already-created file and install it.
- Using a remote URL, where you place an already-created exceptions list on an FTP or HTTP server to be downloaded to the ProxySG.

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

When the Exceptions file is customized, it updates the existing exceptions already on the ProxySG. The configuration remains in effect until it is overwritten by another update; it can be modified or overwritten using CLI commands.

To Install an Exceptions Definition through the Management Console:

1. Select Configuration>Policy>Exceptions.

The Exceptions tab displays.

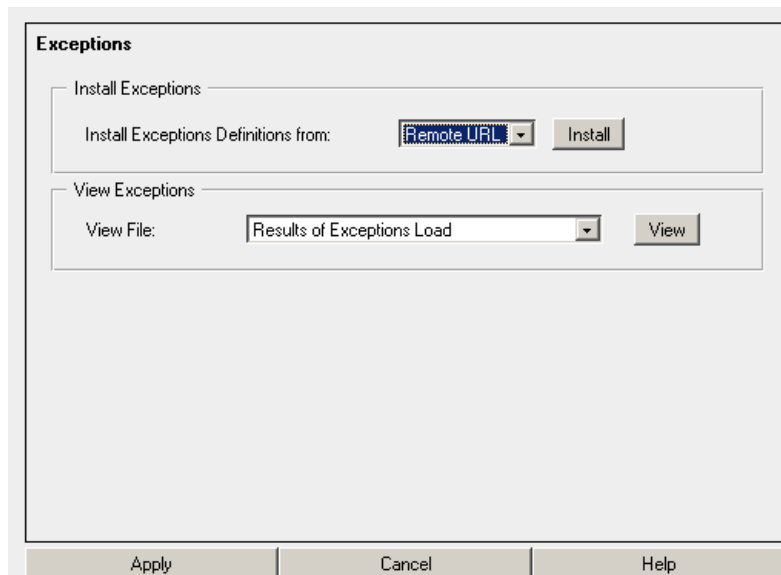


Figure 15-1: Selecting the Exceptions Definitions Download Method

Section D: Defining Exceptions

Note: Click View to examine the existing definitions: Current Exceptions, Default Exceptions Source, Exceptions Configuration, and Results of Exception Load.

2. From the Install Exceptions Definitions From drop-down list, select the method used to install the exceptions 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. View the installation status; click OK.

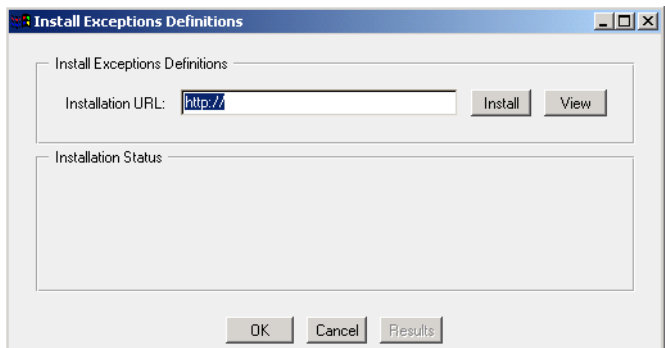


Figure 15-2: Specifying the Remote Location of an Exceptions Configuration

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, close the window, and click Close.



Figure 15-3: Specifying the Local Location of a Exception Definition

Section D: Defining Exceptions

Viewing Exceptions

You can view the exceptions defined on the ProxySG, including how the defined HTML appears to users. The following are the viewable defined exception components:

- Current Exceptions—Displays all of the exceptions as they are currently defined.
- Default Exceptions Source—Displays the default ProxySG exceptions.
- Exceptions Configuration—Displays a page from which you can click links to view how exceptions appear in HTML to users.
- Results of Exception Load—Displays the results of the last installable list load, including any errors and warning to be fixed.

To View Exceptions through the Management Console:

1. Select Configuration>Policy>Exceptions.

The Exceptions tab displays.

2. From the View Exceptions Definitions From drop-down list, select the page to view; click View.

- 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 Close.

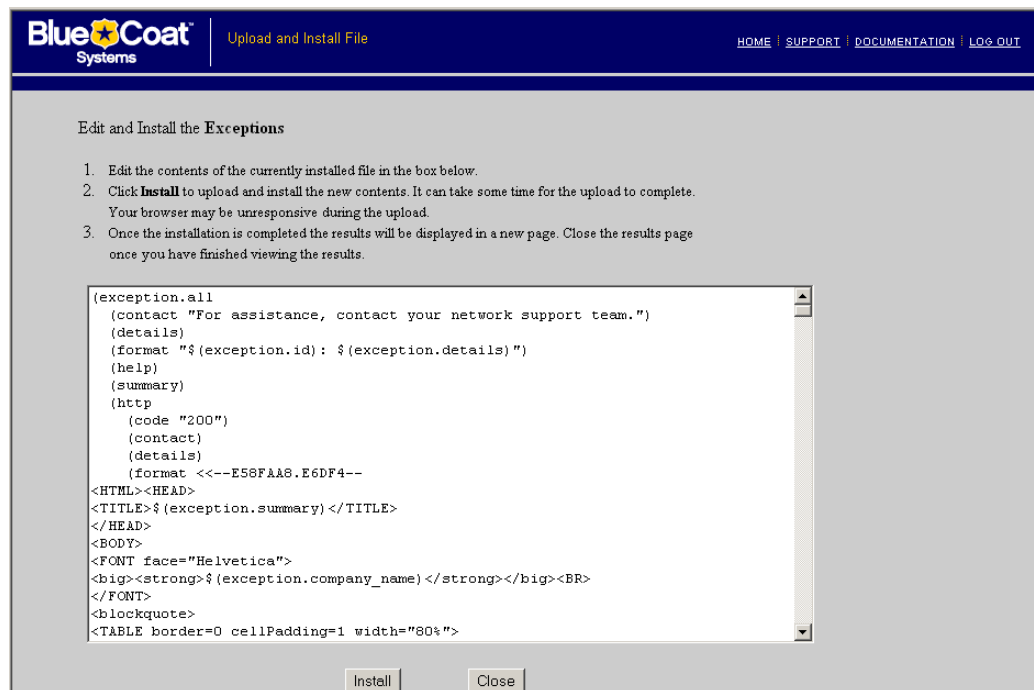


Figure 15-4: Using the ProxySG Text Editor

3. Click Apply.

Section E: Managing Peer-to-Peer Services

This section describes the Blue Coat solution for managing and blocking peer-to-peer traffic.

About Peer-to-Peer Communications

The use of peer-to-peer (P2P) technologies and services consumes an estimated 60% of broadband ISP bandwidth. By design, most P2P services are port-agnostic, which makes attempting to block them at the firewall extremely difficult. One peer finds another IP address and port that is willing to share the file, but different peers can use different ports. Furthermore, P2P is not based on any standards, which makes it nearly impossible for network administrations to control or even detect.

Although P2P provides some practical business uses in enterprises, unmanaged P2P activity creates risks:

- Excessive bandwidth consumptions affects mission-critical applications.
- Exponential security risk of exposure to viruses, spyware, and other malicious content.
- The threat of legal action concerning the unlawful downloading of copyrighted music and movies.

Managing P2P is a dynamic challenge, as the administrator must be able to evaluate both P2P use and enterprise requirements.

The Blue Coat Solution

The ProxySG recognizes P2P activity relating to P2P file sharing applications. By constructing policy, you can control, block, and log P2P activity and limit the bandwidth consumed by P2P traffic.

Note: Neither caching nor acceleration are provided with this feature.

Supported Services

This version of SGOS supports the following P2P services:

- FastTrack (Kazaa)
- EDonkey
- BitTorrent
- Gnutella

Note: Refer to the Release Notes for the most current list of P2P services and versions the ProxySG supports.

Section E: Managing Peer-to-Peer Services

Deployment

To effectively manage P2P activity, the ProxySG must be deployed to intercept outbound network traffic and the firewall configured to block outbound connections that are *not* initiated by the ProxySG.

Notes:

- The ProxySG intercepts outbound TCP network connections, as routed through an L4 switch or a ProxySG in bridging mode.
- Configure ProxySG HTTP, SOCKS, and TCP tunnel services for destination ports to be monitored.
- Create firewall rules that allow only outbound connections that are initiated by the ProxySG.
- You can block all known P2P ports and define policy to stop P2P traffic attempting to come through over HTTP.

Note: This features does not include additional configurations for intercepting or controlling UDP traffic.

Licensing

This feature requires a license. To obtain a license, see Chapter 2: “Licensing” on page 35.

Policy Control

This section lists the policy used to manage P2P.

VPM Support

The following VPM components relate to P2P control:

- Web Access Layer; Source column; P2P Client object. See "P2P Client" on page 485.
- Web Access Layer, Service column; Client Protocols. See "Client Protocol" on page 494.

CPL Support

CPL Triggers

- `http.connect={yes | no}`
- `p2p.client={yes | no | bittorrent | edonkey | fasttrack | gnutella}`

CPL Properties

- `force_protocol()`
- `detect_protocol.protocol`(yes | no)
- `detect_protocol.[protocol1, protocol2, ...]`(yes | no)

- `detect_protocol(all | none)`
- `detect_protocol(protocol1, protocol2, ...)`

Where **protocol** is: `http`, `bittorrent`, `edonkey`, `fasttrack`, or `gnutella`.

The default is `detect_protocol(all)`.

Support CPL

The following properties can be used in conjunction with the P2P-specific CPL:

- `allow`, `deny`, `force_deny`
- `access_server(yes | no)`—If the value is determined as `no`, the client is disconnected.
- `authenticate(realm)`—Unauthenticated clients are disconnected.
- `socks_gateway(alias_list | no)`
- `socks_gateway.fail_open(yes | no)`
- `forward(alias_list) | no`—Only forwarding hosts currently supported by TCP tunnels are supported.
- `forward.fail_open(yes | no)`
- `reflect_ip(auto | no | client | vip | ip_address)`

For complete CPL references, refer to the *Blue Coat Systems Content Policy Language Guide*.

Policy Example

The following policy example demonstrates how to deny network traffic that the ProxySG recognizes as P2P:

```
<proxy>
  p2p.client=yes deny
```

Proxy Authentication

While P2P protocols do not support native proxy authentication, most P2P clients support SOCKS v5 and HTTP 1.1 proxies. P2P proxy authentication is supported only for clients using these protocols (that are configured for proxy authentication).

For information about proxy authentication, see Section B: “Controlling Access to the Internet and Intranet” on page 255. For a list of P2P clients suspected of not supporting SOCKS v5 with authentication, see the Release Notes for this release.

Access Logging

P2P activity is logged and reviewable. See Chapter 20: “Access Logging” on page 743.

Chapter 16: Streaming Media

This chapter contains the following sections:

- "Section A: About Streaming Media"—Provides streaming media terminology, general concepts, and information, such as player limitations and supported formats.
- "Section B: Configuring Streaming Media"—Provides feature-related concepts and procedures for configuring the ProxySG to manage streaming media applications and bandwidth.
- "Section C: Windows Media Player"—Describes how to configure the Windows Media client and describes associated limitations and access log conventions.
- "Section D: RealPlayer"—Describes how to configure the Real Media client and describes associated limitations and access log conventions.
- "Section E: QuickTime Player"—Describes how to configure the QuickTime client and describes associated limitations and access log conventions.

Related Topics:

- Chapter 5: "Managing Port Services" on page 121
- Chapter 6: "Configuring Proxies" on page 149
- Chapter 22: "Statistics" on page 811

Section A: About Streaming Media

Section A: About Streaming Media

This section contains the following topics:

- "Streaming Media Overview"
- "Streaming Media Protocols"
- "Streaming Media Player Support"
- "Streaming Media Authentication"
- "Streaming Media Caching Behavior"

Streaming Media Overview

Streaming is a method of content delivery. With media streaming, video and audio are delivered over the Internet rather than the user having to wait for an entire file to be downloaded before it can be played.

Streaming media support on the ProxySG provides the following features:

- Streaming media files can be live as well as prerecorded.
- Employs flexible delivery methods: unicast, multicast, HTTP, TCP, and UDP.
- Ability to seek, fast-forward, reverse, and pause.
- Ability to play entire file and control media playback, even before it is downloaded.
- Adjust media delivery to available bandwidth, including multi-bit-rate and thinning support.

Important: The ProxySG streaming media components require valid licenses. For more information, see Chapter 2: "Licensing" on page 35.

Supported Streaming Media Clients

The ProxySG supports Microsoft Windows Media, RealNetworks RealPlayer, and Apple QuickTime clients. The specific protocols are discussed in "Streaming Media Protocols" on page 585.

Delivery Method

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 ProxySG provides caching, splitting, and multicast functionality.

Section A: About Streaming Media

Serving Content

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 if the content is live or video-on-demand.

Live Unicast Content

A ProxySG can serve many clients through one unicast connection by receiving the content from the origin server and then splitting that stream to the clients that request it. This method saves server-side bandwidth and reduces the server load. You cannot pause or rewind live broadcasts. A live broadcast can be of prerecorded content. A common example is a company president making a speech to all employees.

Video-on-Demand Unicast Content

A ProxySG can store frequently requested data and distribute it upon client requests. Since the ProxySG is closer to the client than the origin server, the data is served locally, which saves firewall bandwidth and increases quality of service by reducing pauses or buffering during playback. The ProxySG provides higher quality streams (also dependent on the client connection rate) than the origin server because of its closer proximity to the end user. VOD content can be paused, rewound, and played back. Common examples include training videos or news broadcasts.

Multicast Content

The ProxySG can take a unicast stream from the origin media server and deliver it as a multicast broadcast. This enables 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.

For Windows Media multicast, an NSC file is downloaded through HTTP to acquire the control information required to set up content delivery.

For Real Media and QuickTime (through RTSP), 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.

Streaming Media Protocols

This section describes the vendor-specific streaming protocols supported by the ProxySG.

Windows Media Protocols

The ProxySG supports the following protocols:

- MMS-UDP (Microsoft Media Streaming—User Data Protocol)
- MMS-TCP (Microsoft Media Streaming—Transmission Control Protocol)
- HTTP streaming.
- All protocols between the client and the ProxySG for video-on-demand and live unicast content.

Section A: About Streaming Media

- MMS-TCP and HTTP streaming between the ProxySG and origin server for video-on-demand and live unicast content.
- Multicast-UDP is the only delivery protocol supported for multicast. No TCP control connection exists for multicast delivery.

The following briefly describes each of the supported delivery protocols:

- MMS-UDP—UDP provides the most efficient network throughput from server to client. The disadvantage to UDP is that many network administrators close their firewalls to UDP traffic, limiting the potential audience for Multicast-UDP-based streams.

The Windows Media Player attempts to connect in the following order:

- Multicast session. Multicast-UDP uses a TCP connection for control messages and UDP for streaming data. TCP provides packet receipt acknowledgement back to the sender. This insures control message delivery.
- MMS-TCP session. If an MMS-UDP session cannot be established, the client falls back to MMS-TCP automatically.

The ProxySG then establishes a connection to the origin server running the Microsoft Windows Media service.

- MMS-TCP—TCP provides a reliable protocol for delivering streaming media content from a server to a client. At the expense of less efficiency compared to MMS-UDP data transfer, MMS-TCP provides a reliable method for streaming content from the origin server to the ProxySG.

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.

- HTTP Streaming—The Windows Media server also supports HTTP-based media control commands along with TCP-based streaming data delivery. This combination has the benefit of working with all firewalls that let only Web traffic through (port 80).

Depending on the configuration, if MMS-UDP is used between the ProxySG and the client, the appliance can use MMS-TCP, HTTP, or multicast-UDP as the connection to the media server. No protocol relationship exists between the ProxySG and the media server, or between the ProxySG and the client.

Real Media Protocols

The ProxySG supports the following Real Media protocols:

Client-side

- 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)

Section A: About Streaming Media

- 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.

Server-side

- Interleaved RTSP
- HTTP streaming

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 protocols:

- RTP over unicast UDP (RTSP over TCP, RDT over unicast UDP)
- Interleaved RTSP (RTSP over TCP, RDT over TCP on the same connection)
- 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.

Server-side

- Interleaved RTSP
- HTTP streaming

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.

Streaming Media Player Support

This section describes which media player and server versions the ProxySG supports. It also provides the supported streaming content formats.

Consider that the various players might experience limitations dependent upon certain SGOS configurations and features. Feature sections list such limitations, as necessary.

Supported Windows Media Players and Servers

The ProxySG supports the following versions and formats:

Section A: About Streaming Media

- Windows Media Player 6.4, 7, and 9
- Windows Media Server 4.1
- Windows Media Server 9

Supported Real Media Players and Servers

The ProxySG supports the following versions:

- RealOne Player, version 2
- RealPlayer 8
- RealServer 8
- Helix Universal Server

Note: Blue Coat recommends that you not deploy a Helix proxy in 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 Players 6.x and 5.x
- Darwin Streaming Server 4.1.x and 3.x.
- Helix Universal Server

Streaming Media Authentication

The following sections discuss authentication between streaming media clients and ProxySG appliances and between ProxySG appliances and origin content servers (OCS).

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 NT LAN Manager (NTLM) Account Database
- NTLM Authentication and NTLM Account Database

The ProxySG supports the caching and live-splitting of server-authenticated data. The functionality is also integrated with partial caching functionality so that multiple security challenges are not issued to the Windows Media Player when it accesses different portions of the same media file.

Section A: About Streaming Media

When Windows Media content on the server is accessed for the first time, the ProxySG caches the content along with the authentication type enabled on the server. 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).

Authentication type on the server refers to the authentication type enabled on the origin server at the time when the client sends a request for the content.

Windows Media Proxy Authentication

If proxy authentication is configured, Windows Media clients are authenticated based on the policy settings. The proxy (the ProxySG) evaluates the request from the client and verifies the accessibility against the set policies. The Windows Media player then prompts the client for the proper password. If the client 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 Limitations

Consider the following proxy authentication limitations with the Windows Media player (except when specified, these do not apply to HTTP streaming):

- For Windows Media Player 6.4 only: if the media server is configured for NTLM authentication, Windows Media Player 6.4 uses the credentials of the logged-on user to satisfy the challenged request. If the media server or proxy authentication type is NTLM, configure the Windows Media server to accept logged-on user credentials.
- For Windows Media Player 6.4 only: if proxy authentication is not configured and the media server is configured as BASIC and the user fails to provide a valid username and password, the user fails to receive another dialog box. Instead, the request fails to open the stream.
- If the proxy authentication type is configured as BASIC and the server authentication type is configured as NTLM, the default is denial of service.
- If proxy authentication is configured as NTLM 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.

Section A: About Streaming Media

- 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 the 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 transparent request from player or hard-coded service specified in the ProxySG—and request of cache-local (ASX-rewritten or unicast alias) URLs.

Real Media Proxy Authentication

If proxy authentication is configured, Real Media clients are authenticated based on the policy settings. The proxy (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 is accepted, the Real Media server may 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 NTLM challenge is issued, the mode automatically downgrades to BASIC.

Streaming Media Caching Behavior

The following sections describe how the ProxySG and SGOS process and store streaming media requests. Discussed are caching, video on demand (VOD), live splitting, bit rate support, and pre-populating content.

Streaming Media Caching Behavior

Windows Media

The ProxySG caches Windows Media-encoded video and audio files. The standard extensions for these file types are: `.wm`, `.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: `.rm`, `.ra`, and `.rv`.

Section A: About Streaming Media

QuickTime

The ProxySG does not cache QuickTime content (.mov files). All QuickTime content is served in *pass-through* mode only.

Miscellaneous

The ProxySG supports all other files *pass-through* mode only; such files are not cacheable. Examples of these types of files are: MPEG (including .mp3), .rt, .rp, .swf, .gif, .smil, .jpg, and .jpeg.

Video On Demand (VOD)

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 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.

Live Splitting

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 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, MMSU, MMST, and HTTP are considered as identical.

Note: If the origin server is made up of multiple servers, stream splitting sometimes does not occur because Windows Media player 6.4 does not send domain information to the ProxySG; the appliance can only split streams based on the host IP address. In addition, if the URL is composed of hostnames instead of IP addresses, splitting does not occur across WMP 6.4 and WMP 7.0 clients.

Splitting of live unicast streams provides bandwidth savings, since subsequent requests do not increase network traffic.

Section A: About Streaming Media

Multiple Bit Rate Support

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 traffic). This 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.

Only the requested bitrate is cached. Therefore, a media client that requests a 50Kbps stream receives that stream, and the Blue Coat ProxySG caches only the 50Kbps bitrate content.

Bitrate Thinning

Thinning support is closely related to MBR, but different in that 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

The ProxySG supports pre-population of streaming files (QuickTime content is *not* supported) from HTTP servers and origin content servers. Downloading streaming files from HTTP servers reduces the time required to pre-populate the file. With previous SGOS versions, pre-population was accomplished through streaming from the media server. The required download time was equivalent to the file length; for example, a two-hour movie required two hours to download. With the pre-population content management feature, if the media file is hosted on a 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 a HTTP server in addition to the media server.

Using the `content pull` 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.

Section B: Configuring Streaming Media

Section B: Configuring Streaming Media

This section contains the following topics:

- "Limiting Bandwidth"
- "Configuring the Refresh Rate"
- "Configuring HTTP Handoff"
- "Forwarding Client Logs to the Media Server"
- "Configuring Media Server Authentication Type (Windows Media)"
- "About Multicast Streaming"
- "Managing Multicast Streaming for Windows Media"
- "Managing Multicast Streaming for Real Media"
- "Managing Simulated Live Content (Windows Media)"
- "ASX Rewriting (Windows Media)"
- "About Fast Streaming (Windows Media)"

Related Topics

You must also configure the network service (Configuration>Network>Services) to assign port numbers and modes (transparent or proxy). For more information, see Chapter 6: "Configuring Proxies" on page 149.

Limiting Bandwidth

The following sections describe bandwidth limitation and how to configure the ProxySG to limit global and protocol-specific media bandwidth.

About Bandwidth Limitation

Streaming media bandwidth management is achieved by configuring 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 limitations for the ProxySG, as well as for each streaming protocol (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.

Once 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.

Section B: Configuring Streaming Media

Note: If a maximum bandwidth limitation has been specified for the ProxySG, the following condition may 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.
- 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.
- ProxySG to server—The total kilobits per second allowed between the Appliance and the media server. Setting this option to 0 effectively prevents the ProxySG from accepting media content. Limiting ProxySG bandwidth restricts the following streaming media-related functions:

- Live and video-on-demand media, the sum of all bitrates
- Limits the ability to fetch new data for an object that is partially cached
- Reception of multicast streams

- Client to ProxySG—The total kilobits per second allowed between Windows Media streaming media clients and the ProxySG. Setting this option to 0 effectively prevents streaming clients from making connections to the ProxySG.

Limiting server bandwidth restricts the following streaming media-related functions:

- MBR is supported; the ProxySG assumes the client is using the maximum bit rate
- Limits the transmission of multicast streams

- Client connections—The total number of clients that can connect concurrently. Once 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.

Choosing a Method to Limit Streaming Bandwidth

You can control streaming bandwidth using two different methods: you can use the streaming features described below, or you can use the bandwidth management features described in Chapter 10: “Bandwidth Management” on page 375. You should not, however, use both methods to control streaming bandwidth. The way that each method controls bandwidth differs—read the information below to decide which method works best for you.

Limiting streaming bandwidth using the streaming features (described in this section) works this way: if a new stream comes in that pushes above the specified bandwidth limit, that new stream is denied. This allows existing streams to continue to get the same level of quality they currently receive.

Section B: Configuring Streaming Media

Limiting streaming bandwidth using the bandwidth management features (described in Chapter 10: “Bandwidth Management” on page 375) works this way: 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 will be allowed, and the amount of bandwidth available for existing streams will thus be reduced. This can cause 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 will see a reduction in stream quality.

For most circumstances, Blue Coat recommends that you use the streaming features to control streaming bandwidth rather than the bandwidth management features.

Configuring Bandwidth Limitation—Global

This section describes how to limit all bandwidth use through the ProxySG.

To Specify the Bandwidth Limit for all Streaming Protocols through the Management Console:

1. Select Configuration>Services>Streaming Proxies>General.

The General tab displays.

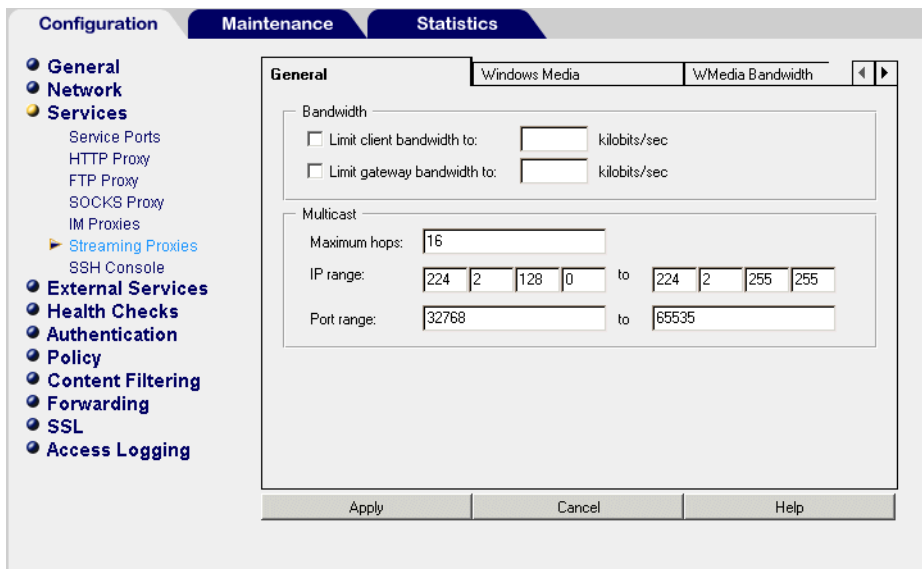


Figure 16-1: Streaming Media General Tab

2. To limit the client connection bandwidth:
 - a. In the Bandwidth field, select Limit client bandwidth to.
 - b. In the Kbits/sec field, enter the number of kilobits to specify the maximum number of kilobits per second that the ProxySG allows for all streaming client connections.
3. To limit the ProxySG (origin server/upstream connection) bandwidth:

Section B: Configuring Streaming Media

- a. In the Bandwidth field, select Limit gateway bandwidth to.
- b. In the Kbits/sec field, enter the number of kilobits to specify the maximum number of kilobits per second that the ProxySG allows for all streaming connections to origin media servers.

To Specify Bandwidth Limit for all Streaming Protocols through the CLI:

To limit the client connection bandwidth, at the `(config)` command prompt, enter the following command:

```
SGOS#(config) streaming max-client-bandwidth kbits_second
```

To limit the ProxySG (origin server/upstream connection) bandwidth, at the `(config)` command prompt, enter the following command:

```
SGOS#(config) streaming max-gateway-bandwidth kbits_second
```

Note: To allow maximum client bandwidth, use the `streaming windows-media no max-client-bandwidth` or the `streaming windows-media no max-gateway-bandwidth` command.

Configuring Bandwidth Limitation—Protocol-Specific

This section describes how to limit bandwidth use per-protocol (Windows Media and Real Media) through the ProxySG.

To Specify the Bandwidth Limit for Windows Media, Real Media, or QuickTime through the Management Console:

1. Select Configuration>Services>Streaming Proxies>WMedia Bandwidth or RMedia Bandwidth or QuickTime Bandwidth.
2. To limit the bandwidth for client connections to the ProxySG:
 - a. Select Limit client bandwidth to.
 - b. In the Kbits/sec field, enter the number of kilobits to specify the maximum number of kilobits per second that the ProxySG allows for all streaming client connections.
3. To limit the bandwidth for connections from the ProxySG to origin content servers:
 - a. Select Limit gateway bandwidth to.
 - b. In the Kbits/sec field, enter the number of kilobits to specify the maximum number of kilobits per second that the ProxySG allows for all streaming connections to origin media servers.

To Specify the Bandwidth Limit for Windows Media, Real Media, or QuickTime through the CLI:

To limit the client connection bandwidth, at the `(config)` prompt, enter the following command:

```
SGOS#(config) streaming {windows-media | real-media | quicktime}
max-client-bandwidth kbits_second
```

To limit the ProxySG (origin server/upstream connection) bandwidth, at the `(config)` command prompt, enter the following command:

Section B: Configuring Streaming Media

```
SGOS#(config) streaming {windows-media | real-media | quicktime}
max-gateway-bandwidth kbits_second
```

Note: To allow maximum client bandwidth, use the `streaming windows-media no max-client-bandwidth` or the `streaming windows-media no max-gateway-bandwidth` command.

Configuring Bandwidth Limitation—Fast Start (WM)

Note: This section applies to Windows Media only.

This section describes how to configure the maximum bandwidth (in kilobytes per second) each Windows Media Player can start with. Upon connection to the ProxySG, streaming media clients will not consume more bandwidth (in kilobits per second) than the defined value.

To Specify the Maximum Starting Bandwidth through the CLI:

At the `(config)` prompt, enter the following command:

```
SGOS#(config) streaming windows-media max-fast-bandwidth kbps
```

Maximum Connections

This section describes how to configure the maximum number of streaming clients, on a per-protocol basis, that can connect to the ProxySG.

To Specify the Maximum Number of Client Connections through the Management Console:

1. Select Configuration>Services>Streaming Proxies>WMedia Bandwidth or Real Media Bandwidth or QuickTime Bandwidth.
2. To limit the bandwidth for connections from the ProxySG to Windows Media origin servers:
 - a. Select Limit maximum connections.
 - b. In the clients field, enter the total number of clients that can connect concurrently.

To Specify the Maximum Number of Client Connections through the CLI:

At the `(config)` prompt, enter the following command:

```
SGOS#(config) streaming {windows-media | real-media | quicktime}
max-connections number
```

Note: To allow maximum number of connections, invoke the `streaming {windows-media | real-media | quicktime} no max-connection` command.

Configuring the Refresh Rate

The refresh feature specifies the length of time before cached streaming content is checked for freshness.

Section B: Configuring Streaming Media

The default is never refresh. Blue Coat recommends that you change this setting.

To Set the Refresh Rate through the Management Console:

1. Select Configuration>Services>Streaming Proxies>Windows Media or Real Media.
2. Perform one of the following:
 - a. In the Check freshness every n.nn hours field, enter the length of time before the cached streaming content is checked for freshness.
 - b. To force the ProxySG to check every time for freshness, select Check freshness every access.
3. Click Apply.

To Set the Refresh Rate through the CLI:

At the (config) prompt, enter the following commands:

```
SGOS#(config) streaming {windows-media | real-media} refresh-interval  
number.number
```

where *number.number* is the length of time before the cached streaming content should be checked for freshness.

Note: A value of 0 requires the streaming content to always be checked for freshness.

To disable freshness checking, enter the following command:

```
SGOS#(config) streaming {windows-media | real-media} no refresh-interval
```

Configuring HTTP Handoff

HTTP handoff is enabled by default. This section describes HTTP handoff and how to disable the feature.

About HTTP Handoff

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 allowing traffic through a firewall, the HTTP module passes control to the streaming module so HTTP streaming can be supported through the HTTP proxy port.

The ProxySG supports HTTP streaming. It does not support HTTP downloading of media files from HTTP servers and their subsequent caching and serving as streaming files. An HTTP connection is established through port 80 that allows you to send streaming data from the origin server to the clients through the ProxySG.

Note: The default setting for HTTP Handoff is enabled. If you do not want HTTP streams to be cached or split, change this setting to disabled.

 Section B: Configuring Streaming Media

Disabling HTTP Handoff

To Disable HTTP Port Handoff through the Management Console:

1. Select Configuration>Services>Streaming Proxies>Windows Media or Real Media or QuickTime.
2. Deselect Enable HTTP handoff.
3. Click Apply.

To Disable the HTTP Port Handoff through the CLI:

At the (config) prompt, enter the following command:

```
SGOS#(config) streaming {windows-media | real-media | quicktime} http-handoff
disable
```

Forwarding Client Logs to the Media Server

This section describes media server compatibility and how to forward client logs.

About Forwarding Client Logs

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.

Note: For Real Media, the log is only forwarded before a streaming session is halted; QuickTime log forwarding is not supported.

Both HTTP streaming and MMS-TCP support logging to the origin server. Logging information is generated for both the server-side connection and the client-side connection that the ProxySG makes to the server. (For more information on what is logged to the origin server, see Chapter 20: “Access Logging” on page 743).

Logging messages are embedded in a log message sent to the content server when:

- The ProxySG receives an end-of-file notification.
- The ProxySG-server connection is closed.
- A user stops the stream. The *connection* is not stopped; the same connection to the OCS remains and is used to send client information. This prevents starting another connection to the OCS.
- A user opens a new file without closing or stopping the current one.

Windows Media only:

- A user seeks to a new position or uses fast forward or reverse.
- A file is looped in a live scenario. Logging occurs whenever playing of a file ends before going on to play another. (Windows Media Player 6.4 does not log this instance.)

Section B: Configuring Streaming Media

When the ProxySG receives a log record from the client, the ProxySG records the log in the access log and then it forwards the log to the origin content server. As necessary, the ProxySG sends log records to the origin media server for whatever content the appliance fetched to populate itself.

Configuring the ProxySG to Forward Client Logs

To Enable Forwarding Client-generated Logging to the Origin Media Server through the Management Console:

1. Select Configuration>Services>Streaming Proxies>Windows Media or Real Media.
2. Select Forward client-generated logs to origin media server.
3. Click Apply.

To Enable or Disable Forwarding Client-generated Logging through the CLI:

At the (config) command prompt, enter the following command:

```
SGOS#(config) streaming {windows-media | real-media} log-forwarding {enable | disable}
```

To Enable or Disable Forwarding Client-generated Logging through the CLI:

```
SGOS#(config) streaming {windows-media} log-compatibility {enable | disable}
```

Configuring Media Server Authentication Type (Windows Media)

Note: This section applies to Windows Media streaming only.

Configure the ProxySG to recognize the type of authentication the origin content server is using: BASIC or NTLM.

To Configure the Media Server Authentication Type through the CLI:

At the (config) prompt, enter the following command:

```
SGOS#(config) streaming windows-media server-auth-type {basic | ntlm}
```

About Multicast Streaming

This section describes multicast streaming and how to configure the ProxySG to manage multicast broadcasts.

Section B: Configuring Streaming Media

About Serving Multicast Content

- How multicast content is handled through the ProxySG depends on whether the ProxySG is delivering Windows Media or Real Media multicast broadcasts (QuickTime is not supported). For Windows Media, 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 Real Media, the ProxySG takes a unicast stream from the origin RealServer and delivers it as a multicast stream. This enables 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.

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 and Real 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.

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 through the Management Console:

1. Select Configuration>Services>Streaming Proxies>General.
2. In the Maximum Hops field, enter a time-to-live (TTL) value.
3. In the IP Range fields, enter the IP address range.
4. In the Port Range fields, enter the port range.
5. Enable Windows and Real Media multicast; see the next section, "Managing Multicast Streaming for Windows Media" and "Managing Multicast Streaming for Real Media" on page 605.

Section B: Configuring Streaming Media

Managing Multicast Streaming for Windows Media

This section describes multicast station and `.nsc` files, and describes how to configure the ProxySG to send multicast broadcasts to Windows Media clients.

About Multicast Stations

A multicast station is a defined location from where the Windows Media player retrieves live streams. This defined location allows `.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 the 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 that is created from the multicast station configuration. 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 will not be able to send packets to the network. For video-on-demand files, use the `broadcast-alias` command, discussed below.

Multicast to Multicast

Use the `multicast-alias` command to get the source stream for the multicast station.

About Broadcast Aliases

A broadcast alias defines a playlist, specify a starting time, date, and the number of times the content will be repeated.

Creating a Multicast Station

To create a multicast station, you must perform the following:

- Define a name for the multicast station.

Section B: Configuring Streaming Media

- Define the source of the multicast stream.
- The port range to be used.
- Define the address range of the multicast stream.
- Define the TTL value.
- Create the multicast alias, unicast alias, and broadcast alias commands to enable the functionality.

Note: You must configure multicast stations through the CLI.

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 `mms://10.25.36.47/tenchi`.
- Accepts the address, port, and TTL default values.

```
SGOS#(config) streaming windows-media multicast-station station1
mms://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

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
```

Section B: Configuring Streaming Media

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 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 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 get an `.nsc` file from 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 the Windows Media Player to start the stream.

The newly created file is not editable; the settings come from streaming configuration file. In that file, you have already defined the following pertinent information for the file:

- ❑ 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 the Windows Media Player cannot find the streaming packets.
- ❑ The description, which references the MMS URL that you defined.
- ❑ The format, which contains important ASF header information. All streams delivered by the multicast station definition have their ASF headers defined here.

 Section B: Configuring Streaming Media

Monitoring the Multicast Station

You can determine the multicast station definitions by viewing the streaming windows configuration. To determine the current client connections and current ProxySG connections, use the `show streaming windows-media statistics` command.

To View the Multicast Station Setup through the CLI:

```
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 mms://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 among the default ranges allowed.

To View the Multicast Station Statistics through the CLI:

```
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
```

Managing Multicast Streaming for Real Media

This section describes how to configure Real Media multicast streaming.

Section B: Configuring Streaming Media

About Real Media Multicast Broadcasts

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, who 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.

Enabling Real Media Multicast

To Enable Multicast through the Management Console:

1. Select Configuration>Services>Streaming Proxies>RMedia Bandwidth.
2. Select Enable multicast.
3. Click Apply.

To Set the Refresh Rate through the CLI:

At the `(config)` prompt, enter the following commands:

```
SGOS#(config) streaming real-media multicast enable
```

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.

Note: This section applies only to Windows Media.

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.

Once a starting date and time for the simulated live content have been set, the broadcast of the content starts when there is at least one client requesting the file. 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 prepopulating the content on the Appliance provides better streaming quality.

Before configuring simulated live, consider the following:

Section B: Configuring Streaming Media

- 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.
- Once 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.

Three scenarios can occur when a client requests the simulated live content:

- Clients connect before the scheduled start time of the simulated live content: clients are put into *wait* mode.
- Clients connect during the scheduled playback time of the simulated live content: clients receive cached content for playback.
- Clients connect after the scheduled playback time of the simulated live: the client receives an error message.

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.

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, ...

Section B: Configuring Streaming Media

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. Up to 128 URLs can be added.

```
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 `mms://ocs.bca.com/bca1.asf` and `mms://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  
mms://ocs.bca.com/bca1.asf 2 today 4pm,6pm,8pm  
SGOS#(config) streaming windows-media broadcast-alias bca  
mms://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.

About ASX Rewrite

The ProxySG provides proxy support for Windows Media Player 6.4, although the player itself does not support the specification of explicit proxies using the MMS protocol.

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 6.4 clients by rewriting the Windows Media metafile (which contains entries with URL links to the actual location of the streaming content) to point to the appliance 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 refers to 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 will not be processed by the Windows Media module. Also note that the `.asx` file with the appropriate syntax must be located on an HTTP (not 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`).

Section B: Configuring Streaming Media

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 6.4 requires the following:

- The client is explicitly proxied for HTTP content to the ProxySG that will rewrite 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 should be 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 Limitations

Before creating rules, consider the following.

- Each rule you create must be checked for a match; therefore, performance might be affected if you create large amounts of rules.
- Lower numbers have a higher priority than high numbers.

Note: Rules can only be created through the CLI.

- 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 that the client proxies to). The combination of proxy-chained ProxySGs in the HTTP path coupled with ASX rewrite configured for multiple ProxySGs in the chain can create a rewritten URL requested by the client in the example form of:

Section B: Configuring Streaming Media

```
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 may 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 Streaming UDP
 - `mmst` specifies Microsoft Media Streaming 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. This specifies that the protocol specified in the original URL will 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.

Section B: Configuring Streaming Media

To Set Up the .asx Rewrite Rules through the CLI:

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 has been modified immediately, 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 NTLM 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 NTLM authentication, the following procedure allows you to designate any virtual IP addresses to the NTLM authentication type. If you know that all of the activity through the ProxySG requires NTLM authentication, you can use the IP address of the appliance.

To Designate an IP Address to an Authentication Type through the CLI:

1. If necessary, create a virtual IP address that will be 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 B: Configuring Streaming Media

About Fast Streaming (Windows Media)

Note: This feature applies to Windows Media only.

Windows Media Server version 9 contains a feature called Fast Streaming that allows clients to provide streams with extremely low buffering time.

SGOS 4.x supports the following functionality for both cached and uncached content:

- Fast Start
- Fast Cache

Fast Recovery and Fast Reconnect are currently not supported.

Section C: Windows Media Player

Section C: Windows Media Player

This section describes how to configure the Windows Media Player client and describes associated limitations and access log conventions.

Configuring Windows Media Player

To apply the ProxySG Windows Media streaming services, Windows Media Player 6.4 or higher must be installed and configured to use explicit proxy.

Note: If using Windows Media Player 6.4, you must define the HTTP explicit proxy.

MMS explicit proxy is defined with the `asx-rewrite` command (discussed earlier in this chapter) or with CPL (`url_host_rewrite`).

Note: The example below uses Windows Media Player 9.0. 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>Network.

The Network tab displays.

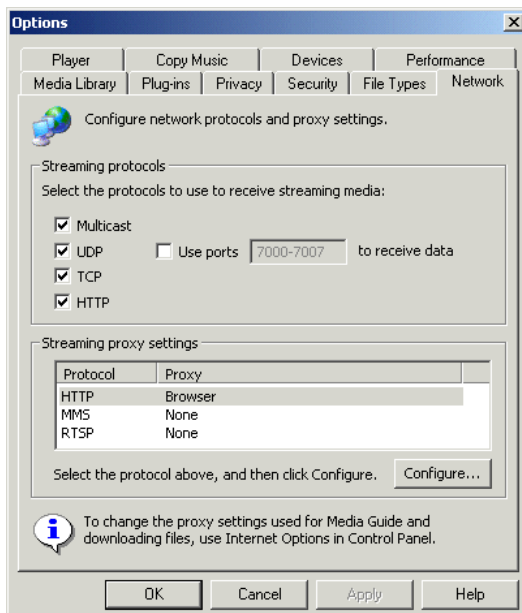


Figure 16-2: Configuring Windows Media Player Proxy

3. In the Streaming proxy settings section, select MMS and click Configure.

Section C: Windows Media Player

The Configure Protocol window displays for the selected protocol.

4. Select Use the following proxy server and enter the ProxySG IP address and the port number used for the explicit proxy (the default MMS port is 1755).
5. Click OK; click OK again to close the Options dialog.

Limitations

This section describes Windows Media Player limitations that might affect performance.

Striding Limitations

When you use the Windows Media Player, consider the following limitations 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 Limitations

- Applies to Version 6.4 only: for ASX rewriting to occur, the player must be configured to use the ProxySG as the HTTP proxy. Configuring the browser only as the HTTP proxy is not sufficient.
- Applies to Versions 6.4 and 9: 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 server the content.

Section C: Windows Media Player

- If you request an un-cached file using the HTTP protocol, the file is likely to stop playing if the authentication type is set to BASIC or NTLM and you initiate rapid seeks before the buffering begins for a previous seek. The 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, the Windows Media Player enters a waiting stage. This is normal. However, if HTTP is used as the protocol, after a minute or two the Windows Media Player closes the HTTP connection, but remains in the waiting stage, even when the stream is broadcasting.

Note: For authentication-specific limitations, see "Windows Media Player Authentication Limitations".

Windows Media Access Log Formats

See Appendix B: "Access Log Formats" on page 877.

Section D: RealPlayer

Section D: RealPlayer

This section describes how to configure Real Player and describes associated limitations and access log formats.

Configuring RealPlayer

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 the RealPlayer.

To Configure RealPlayer:

Note: This procedure features RealOne Basic, version 2.0. Installation and setup menus vary with different versions of RealPlayer. Refer to the RealPlayer documentation to configure earlier versions of RealPlayer.

1. Start RealPlayer.
2. Select Tools>Preferences.

The Preferences dialog appears.

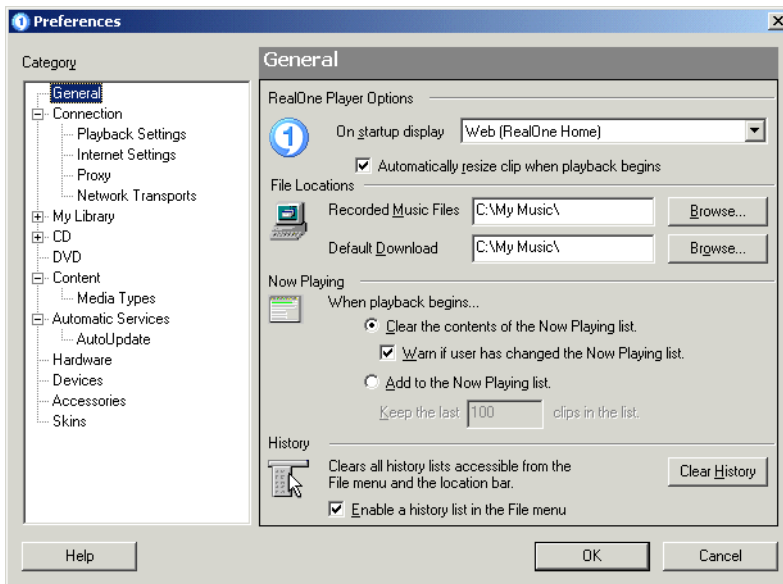


Figure 16-3: RealOne Preferences Dialog

3. Click Proxy. In the Streaming Setting section, click Change Settings; the Streaming Proxy Settings dialog appears.
4. In the PNA and RTSP proxies: field, click Use Proxies and in the RTSP field enter the IP address of the proxy ProxySG. Also enter the RTSP port number (the default is 554).

Section D: RealPlayer

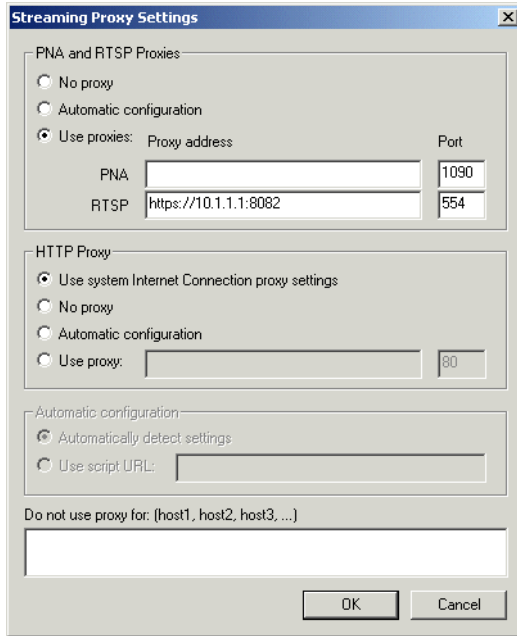


Figure 16-4: Configuring the RealPlayer to Proxy through the ProxySG

These settings must match the settings configured in the ProxySG. If you change the ProxySG explicit proxy configuration, you must also reconfigure the RealPlayer.

5. For HTTP Proxy, if you have an HTTP proxy already configured in your browser, select Use system Internet Connection proxy settings.

Note: If using transparent proxy, RTSP port 554 is set by default and cannot be changed.

6. 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, which is the recommended method.

7. Click OK to close the Streaming Proxy Settings dialog.
8. To configure RealPlayer transport settings, select Network Transports.
9. Click RTSP Settings.

The RTSP Transport Settings dialog appears.

Section D: RealPlayer

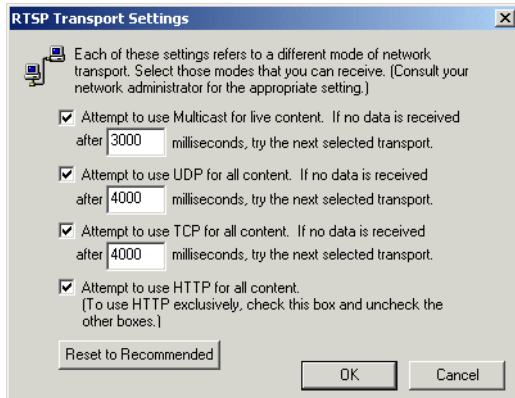


Figure 16-5: Configuring RealPlayer RTSP Transport Settings

10. Click the appropriate boxes based on your network configuration. For example, if your firewall does not accept UDP, select Attempt to use TCP for all content. Blue Coat recommends using the default settings.
11. Click OK.
12. To allow the creation of access log entries, RealPlayer must be instructed to communicate with the RealServer. Perform one of the following or both as necessary:
 - RealPlayer 8—Select View>Preferences>Support; click Send connection-quality data to RealServers; click OK.
 - RealOne Player—Select Tools>Preferences>Internet Settings; in the Internet Settings field, click Send connection-quality data to RealServers; click OK.

Real Media Access Log Formats

See Appendix B: “Access Log Formats” on page 877.

Limitations and Known Issues

For authentication-specific limitations, see "Real Media Player Authentication Limitation" on page 590.

Section E: QuickTime Player

Section E: QuickTime Player

This section describes how to configure the QuickTime client and describes associated limitations and access log formats.

Configuring QuickTime Player

This section describes how to configure the QuickTime player for explicit proxy to the ProxySG.

To Configure QuickTime:

1. Select Edit>Preferences>QuickTime Preferences.

The QuickTime Settings dialog appears.

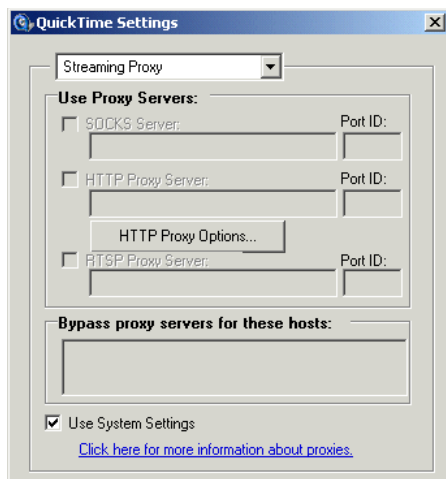


Figure 16-6: Configuring the QuickTime Client Proxy

2. Deselect Use System Settings.
3. Select RTSP proxy server; enter the IP address of the ProxySG to connect to and 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 RealPlayer.

4. Close the dialog.

QuickTime Access Log Formats

See Appendix B: "Access Log Formats" on page 877.

Limitations

For authentication-specific limitations, see "QuickTime Proxy Authentication" on page 590.

Access Log Format

See Appendix B: “Access Log Formats” on page 877.

Chapter 17: *Instant Messaging*

This chapter discusses how to control Instant Messaging (IM) activity through the ProxySG.

About Securing Instant Messaging

Instant Messaging usage in an enterprise environment creates security concerns because regardless of how network security is configured, IM connections can be made from any established protocol, such as HTTP or SOCKS, on any open port. Because it is common for coworkers to use IM to communicate, especially in remote offices, classified company information can be exposed outside the network. Viruses and other malicious code can also be introduced into the network from file sharing through IM clients.

The ProxySG serves as an IM proxy. You can control IM 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.

The ProxySG supports the AOL, MSN, and Yahoo IM protocols.

Recommended Deployments

For large networks with unimpeded Internet access, Blue Coat Systems recommends transparently redirecting the IM protocols to the ProxySG, which requires the ProxySG bridging feature or an L4 switch or WCCP.

For networks that do not allow outbound access, Blue Coat recommends using the SOCKS proxy and configuring policy and content filtering denials for HTTP requests to IM servers.

About the Instant Messaging Protocol Services

The ProxySG accepts connections for the supported IM protocols on ports specified in services. The following are the default service ports (transparent, but disabled):

- AOL-IM: 5190
- MSN-IM: 1863 and 6891
- Yahoo-IM: 5050 and 5101

These ports are disabled by default.

MSN port 1863 and Yahoo port 5050 are the default client login ports. MSN port 6891 and Yahoo port 5101 are the default for client-to-client direct connections and file transfers. If these ports are not enabled:

- Client-to-client direct connections do not occur.
- After a file transfer request is allowed by the ProxySG, the resulting data is sent directly from one client to another without passing through the ProxySG:

- ❑ For MSN: The above bullet point only applies to MSN version previous to and including 6.0. Post-6.0 versions use a dynamic port for file transfers; therefore, port 6891 is not required for the ProxySG to intercept file transfers.
- ❑ For Yahoo: The above bullet only applies to standard file transfer requests. Port 5101 must be enabled to allow file list requests.

Note: All file transfers for AOL clients are handled through the default (5190) or specified client login port.

To enable a default IM port or configure additional IM services, see Chapter 5: “Managing Port Services” on page 121.

About HTTP Proxy Support

SGOS 4.x supports instant messaging through HTTP proxy. IM clients can be configured to connect to IM services through HTTP, which allows IM activity from behind restrictive firewalls.

The ProxySG supports HTTP proxy for Yahoo, MSN, and AOL IM clients, including application of policies and IM activity logging. This is accomplished by the HTTP proxy handing off IM communications to the IM proxy.

Limitations

AOL and Yahoo clients are lose certain features when connected through HTTP proxy rather than through SOCKS or transparent connections:

- AOL—Direct connections, file transfers, and files sharing are not available.
- Yahoo—Client cannot create a chat room.

About Instant Messaging Reflection

IM reflection allows you to contain IM traffic within the enterprise network, which further reduces the risk of exposing company-confidential information through public IM networks. Normally, an IM sent from one buddy to another is sent to and from an IM service. With IM reflection, IM traffic between buddies, including chat messaging, on the same network never has to travel beyond the ProxySG. This includes IM users who login to two different ProxySG appliances configured in a hierarchy (proxy chaining).

IM Reflection Diagrams

The following diagrams depict how the ProxySG manages IM reflection.

IM Reflection with Fail Open

The following diagram demonstrates IM reflection deployment with fail open on a ProxySG that is configured to attempt to reflect all IM activity. IM clients 1 and 2 logged into the same ProxySG, while client 3 is outside the network. IM activity between clients 1 and 2 are reflected by the ProxySG; IM activity between clients 1 and 3 are forwarded to the IM service provider for normal delivery.

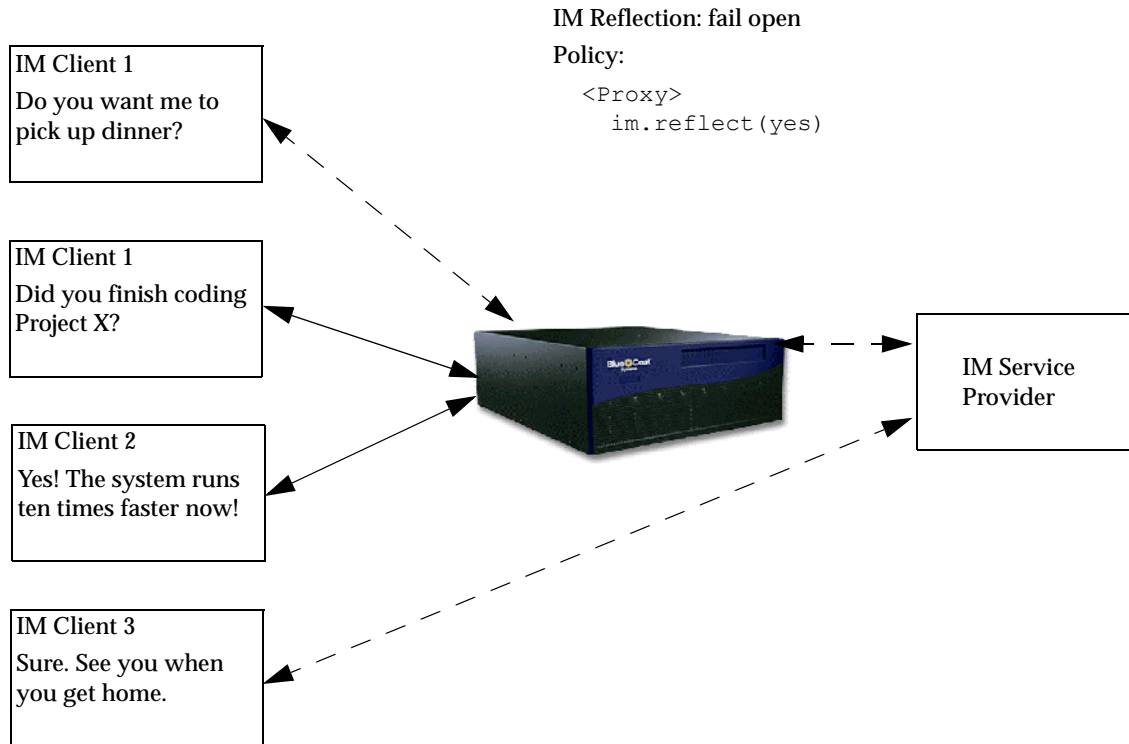


Figure 17-1: IM Reflection with Fail Open

IM Reflection With Fail Closed

By adding a policy rule to deny IM service to clients not logged into the ProxySG, client 1 receives a denial of service message when trying to message client 3.

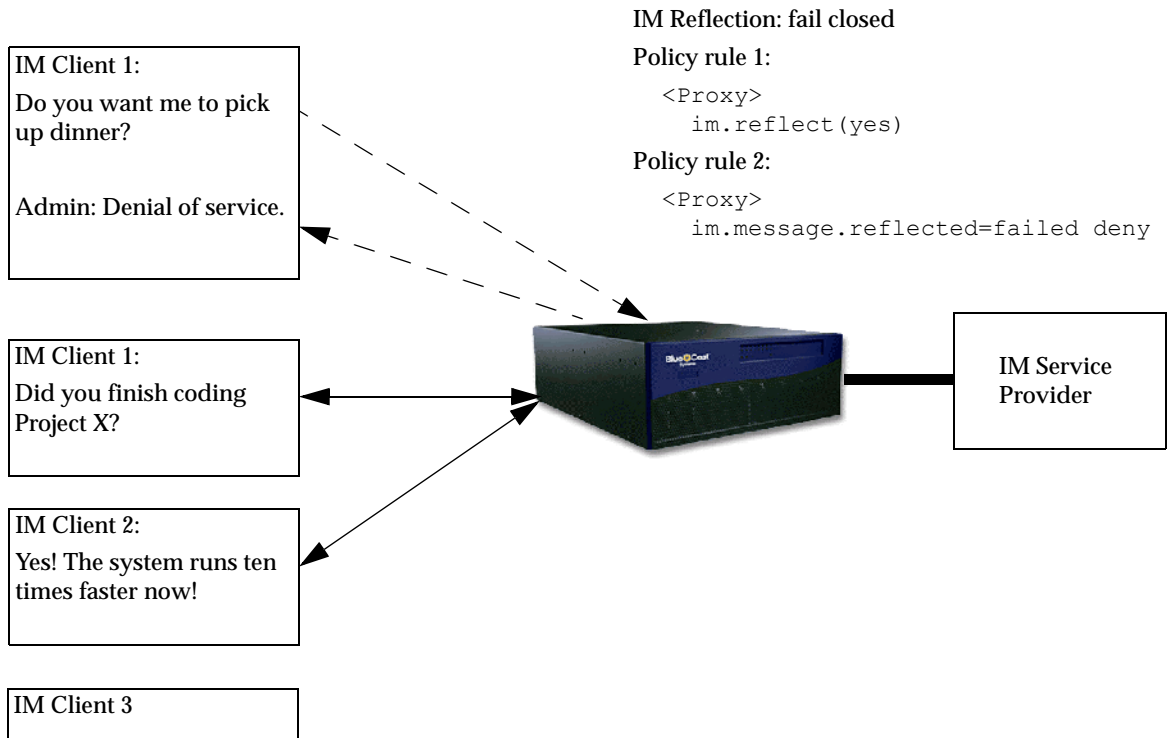


Figure 17-2: IM Reflection with Fail Closed

IM Reflection With A Hierarchy Of Proxies

Larger enterprise networks have users logging in through different primary ProxySG appliances. IM reflection is still possible by using SOCKS and HTTP forwarding, policy, and a ProxySG hierarchy.

Consider the following deployment. IM Clients 1 and 2 are located on the same main campus, but log into different primary ProxySG appliances, PSG 1 and PSG 2, which proxy to the intermediate ProxySG, PSG 3. IM Client 3 is an employee in a remote location and logs into PSG 4. PSG 5 is the corporate root appliance. IM Client 4 is a buddy of IM Client 2, but is not on the employee network.

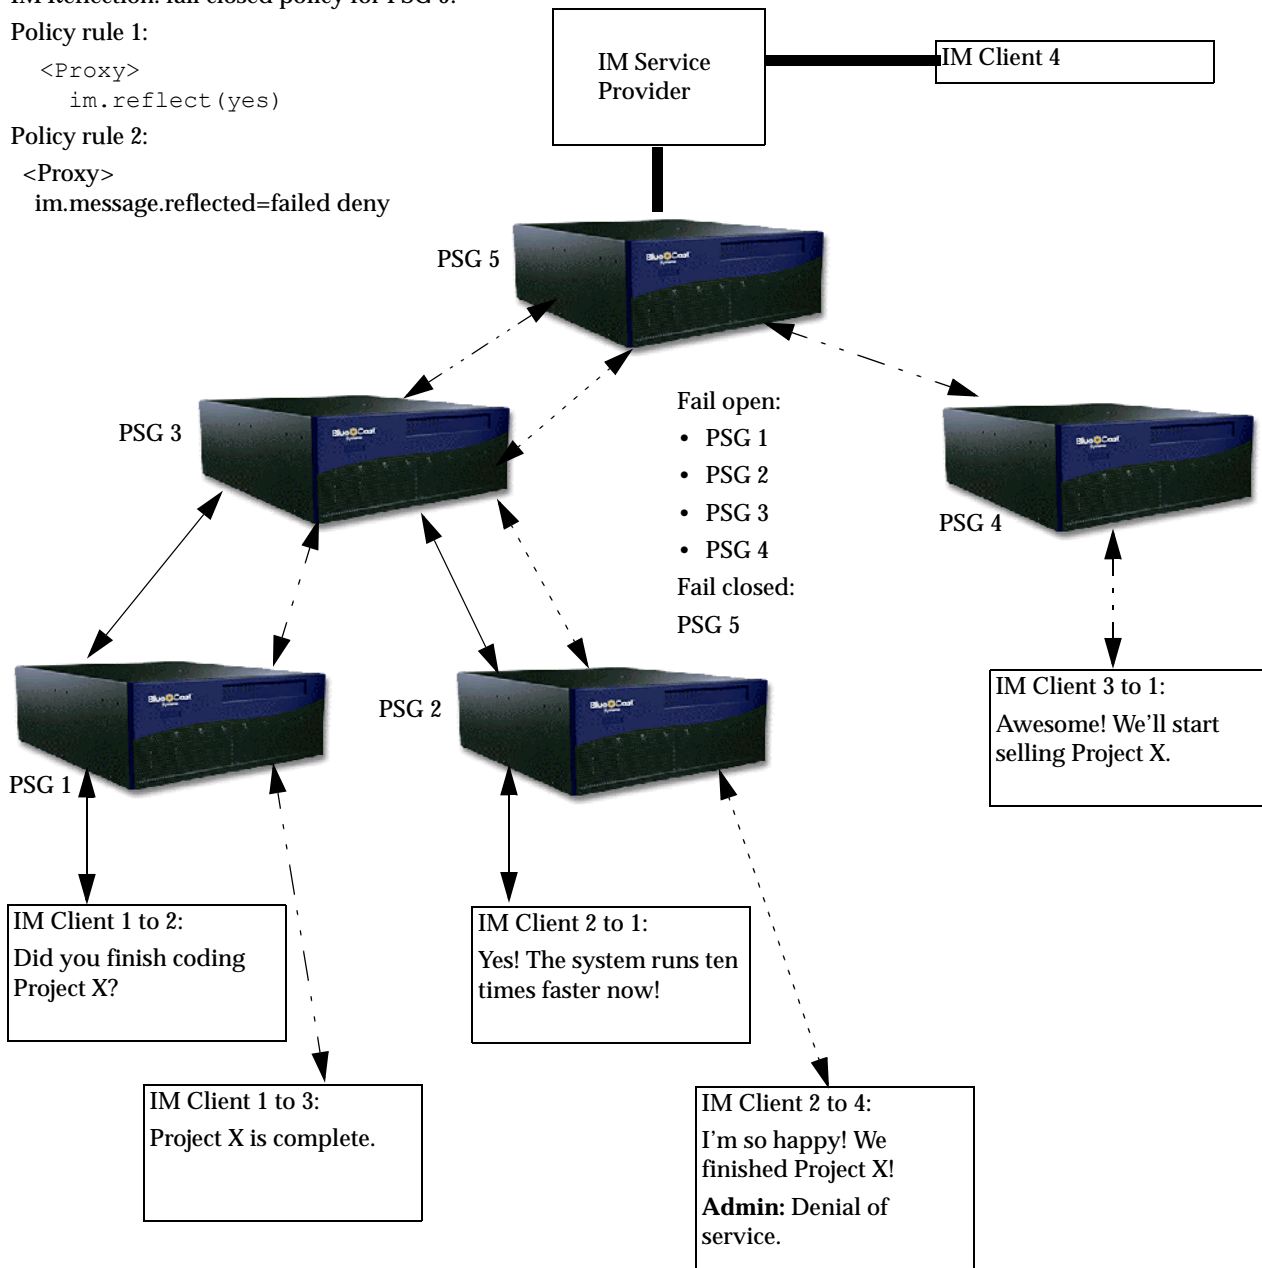
IM Reflection: fail closed policy for PSG 5:

Policy rule 1:

```
<Proxy>
  im.reflect (yes)
```

Policy rule 2:

```
<Proxy>
  im.message.reflected=failed deny
```



IM Reflection: proxy hierarchy, fail closed policy for PSG 1-4:

```
define condition "IM protocols" client.protocol=(aol-im,msn-im,yahoo-im)end
condition "IM protocols"
```

```
<Forward>
```

```
condition="IM protocols" socks_gateway(gateway_1) socks_gateway.fail_open(no)
```

Figure 17-3: IM Reflection with SOCKS Forwarding in a Proxy Hierarchy

Each primary and intermediate ProxySG (PSGs 1, 2, 3, and 4) forward IM traffic that is not reflectable (policy to fail open) to the next ProxySG in the chain. If the next ProxySG services the appropriate IM client, the message is reflected and delivered. The root ProxySG, PSG 5, has a policy to fail closed. Therefore, all IM traffic forwarded to it that cannot be reflected, such as IM Client 2's attempt to contact IM Client 4, is denied access to the public IM service.

Further policy fine-tuning can allow or disallow IM forwarding based on other triggers. For example, the group Corp-Market can send messages to anyone inside or outside the network, but all other groups are prohibited from sending messages to the outside.

About Instant Messaging Proxy Authentication

The ProxySG supports explicit proxy authentication if explicit SOCKS V5 proxy is specified in the IM client configuration.

Because the IM protocols do not support proxy authentication natively, authentication for transparently redirected clients is not supported because policies requiring authentication would deny transparently redirected clients.

HTTP Proxy Limitations

The following proxy authentication limitations apply to IM clients using HTTP proxy:

- AOL IM—Proxy authentication is supported.
- MSN IM (5.0 and above)—While the MSN IM client does support user credentials, it cannot respond to HTTP proxy authentication requests from the ProxySG and the MSN passport service login fails. You can, however, add policy to pass-through the traffic to the MSN passport.com site without requiring authentication.
- Yahoo IM—Yahoo IM clients do not have proxy authentication configuration abilities.

Securing AOL Encryption Capability

This section describes AOL encryption capabilities and how to manage them with ProxySG policy.

About AOL Encryption

AOL IM provides the option for clients to send encrypted messages through both standard messaging (through a service) and direct connection messaging. While this encryption benefits IM users, it provides a security risk for corporate network administrators implementing a communication policy through a proxy. Encryption-capable AOL IM buddies can enable encryption and communicate. Because the ProxySG cannot decrypt these communications, policy cannot be applied and sensitive material can be transferred between buddies without a denial of service or access logging for key-word matching. The ProxySG also cannot replace or append encrypted text, rendering that IM proxy feature useless.

To allow unabated proxy control of IM traffic, SGOS 4.x can strip the encryption capabilities from AOL and Trillian IM clients. While this might appear counter-intuitive to securing communications, greater security and control are gained from the ability to apply policy to message content and to log communications. Determine the need to strip encryption based on your enterprise proxy requirements.

Note: If encryption is blocked, the service does not recognize the logged-in IM client as capable of encryption. If a proxied client attempts to create a chatroom with encryption on, the client receives a create error. This behavior is expected.

Policy for Stripping AOL Encryption

The policy property is only applicable to the `im.method=login` trigger; other properties are not affected. Once encryption stripping is enabled, any existing encryption capabilities and certificates are stripped when a client logs in, and the IM service recognizes the clients as not able to send encrypted messages.

VPM

In a Web Access Layer, select Block IM Encryption in the Action column.

CPL

Add the following property to the policy file:

```
<Proxy>
  im.block_encryption(yes)
```

Instant Message Proxies

This section discusses the IM proxy behavior and configurations on the ProxySG.

Configuring Instant Message DNS Redirection

The ProxySG can be configured as an IM proxy that performs a DNS redirection for client requests. This provides greater control because it prevents IM clients from making outside connections.

The IM clients provide the DNS lookup to the IM server, which the ProxySG DNS module uses to connect to the IM server. To the client, the ProxySG appears to be the IM server. A virtual IP address used only for IM must be configured, as it is used to represent the IM server address for all IM protocols.

To Configure Instant Message DNS Redirection through the Management Console:

1. Create a virtual IP address to be used for IM DNS redirection through the ProxySG. Navigate to Configuration>Network>Advanced>VIPs.
2. Select Configuration>Services>IM Proxies>IM Proxy Settings.
3. In the General Settings field, select the configured VIP from the Explicit Proxy Virtual IP drop-down list.
4. In the Protocol Settings field, select an IM protocol to define: AOL, MSN, or Yahoo. Once selected, the appropriate host fields display below. Each field contains the default hosts used by clients to connect to the IM service.
 - AOL: Native IM Host, HTTP IM Host, and Direct IM Proxy Host.
 - MSN: Native IM Host and HTTP IM Host.

- Yahoo: Native IM Host, HTTP IM Host, HTTP Chat Host, Upload Host, and Download Host.

Important: Only edit these hosts if the client experiences a change in its hardcoded value.

To Configure Instant Message DNS Redirection through the CLI:

At the (config) prompt, enter the following commands:

```
SGOS#(config) virtual-ip address
SGOS#(config) im explicit-proxy-vip address
    where address is the same VIP defined with the previous command.
SGOS#(config) im host
    where host is:
    aol-direct-proxy-host host
    aol-http-host host
    aol-native-host host
    msn-http-host host
    msn-native-host host
    yahoo-download-host host
    yahoo-http-host host
    yahoo-http-chat-host host
    yahoo-native-host host
    yahoo-upload-host host
```

To view the current default or configured hosts, enter the `show im` command.

Configuring Instant Message Alert Settings

This section describes how to configure the IM proxy settings on the ProxySG. You can assign an administrator buddy name for each client type, and determine how exception messages are sent.

An administrator buddy name can be a registered name user handle or a fictitious handle. The benefit of using a registered name is that users can send IM messages to the administrator directly to report any issues, and that communication can be logged for tracking and record-keeping.

To Configure the IM Proxy Setting through the Management Console:

1. Select Configuration>Services>IM Proxies>IM Alert Settings.

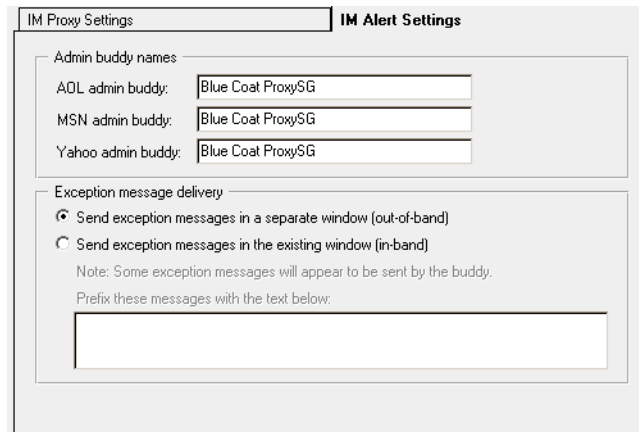


Figure 17-4: The IM Proxy Screen

2. In the Admin buddy names field, enter the handle or handles for the administrator.
3. In the Exception message delivery field, select the method that exception messages are delivered to IM users.
 - Send exception messages in a separate window (out-of-band)—if an exception occurs, the user receives the message in a separate IM window.
 - Send exception messages in the existing window (in-band)—If an exception occurs, the message appears in the same IM window.

If in-band is selected, the message appears to be sent by the buddy on the other end, with the exception that when in a chat room, the message always appears to be sent by the configured Admin buddy name. You can enter a prefix message that appears in the client window before the message. For example: “From the Company Administrator: Inappropriate IM use. Refer to Employee Conduct Handbook concerning Internet usage.”

Note: Regardless of the IM exception delivery configuration, IM alert messages triggered by policy based on certain protocol methods are always sent out-of-band because a specific buddy is not associated.

4. Click Apply.

To Configure the IM Proxy Setting through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) im [aol | msn | yahoo]-admin-buddy admin_handle
```

Specifies the handle or handles for the administrator; configure for each IM client type.

```
SGOS#(config) im exceptions [in-band | out-of-band]
```

Specifies the method the exception messages are delivered to IM users. If in-band is selected, enter the following command to specify a prefix message:

```
SGOS#(config) im buddy-spoof-message text
```

Configuring Instant Messaging HTTP Handoff

IM Handoff allows the Blue Coat HTTP proxy to handle requests from supported IM protocols. If IM HTTP handoff is disabled, requests are passed through, and IM-specific policies are not applied.

Handoff should be enabled (the default) if you write IM policy.

If you want to allow a specific IM client to connect through HTTP through the ProxySG and that IM protocol has not been licensed, disable IM HTTP handoff to allow the traffic to be treated as plain HTTP traffic and to avoid an error in the licensing check done by the IM module. This might be also be necessary to temporarily pass through traffic from new versions of IM clients that are not yet supported by Blue Coat.

To Disable Instant Messaging HTTP Handoff through the Management Console:

1. Select Configuration>Services>IM Proxies>IM Proxy Settings.
2. Deselect Enable HTTP Handoff.

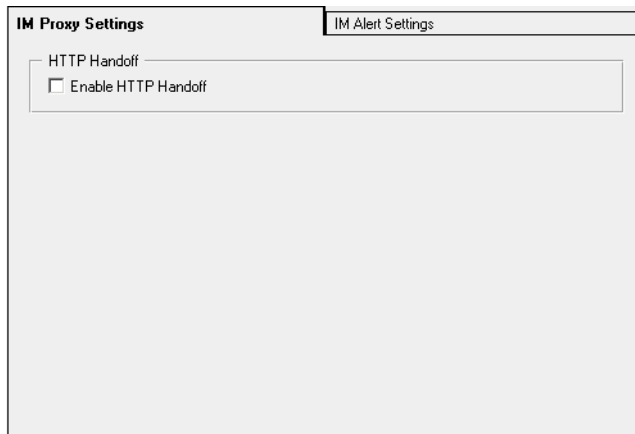


Figure 17-5: Disabling IM HTTP Handoff

To Disable Instant Messaging HTTP Handoff through the CLI:

At the (config) command prompt, enter the following command:

```
SGOS#(config) im http-handoff disable
```

Configuring Instant Messenger Clients

This section describes how to configure the IM clients to send traffic through the ProxySG.

General Configuration

As each IM client has different menu structures, the procedures to configure them differ. This section provides the generic tasks that need to be completed.

Explicit Proxy

Perform the following tasks on the IM client:

1. Navigate to the Connection Preferences dialog.
2. Select Use Proxies.
3. Select proxy type as SOCKS V5.
4. Enter the ProxySG IP address.
5. Enter the SOCKS port number; the default is 1080.
6. Enter authentication information, if required.

Transparent Proxy

IM clients do not require any configuration changes for transparent proxy. An L4 switch or inline ProxySG routes the traffic.

Yahoo Messenger Client Explicit Proxy Configuration Screen

The following example configures a Yahoo Messenger client for explicit proxy.

1. Select Login>Preferences>Connection.
2. Click Connection.
3. Select Use proxies.
4. Select Enable SOCKS proxy; select Ver 5.
5. Enter the server name.
6. Enter the port number (the default is 1080).
7. If authentication is required on the ProxySG, enter the authentication user name and password.

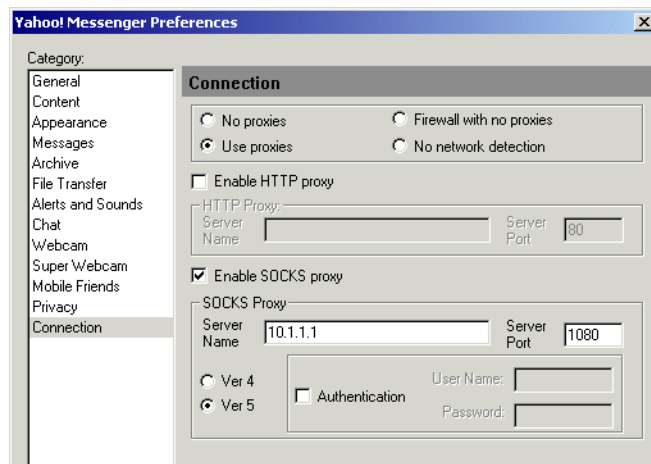


Figure 17-6: Yahoo IM Client Explicit Proxy Configuration

Notes

If Yahoo Messenger is configured for explicit proxy (SOCKS) through the ProxySG, the IM voice chat feature is disabled. Any client attempting a voice chat with a client behind the ProxySG firewall receives an error message. The voice data stream is carried by default on port 5001; therefore, you can create and open this port and configure Yahoo IM to use transparent proxy. However, the ProxySG only supports the voice data in pass-through mode.

AOL Messenger Client Explicit Proxy Configuration Screen

The following example configures an AOL Messenger client for explicit proxy.

1. Select My AIM>Edit Options>Edit Preferences>Sign On/Off.
2. Click Connection.
3. Select Connect using proxy.
4. Select SOCKS 5.
5. Enter the server name.
6. Enter the port number (the default is 1080).
7. If authentication is required on the ProxySG, enter the authentication user name and password.

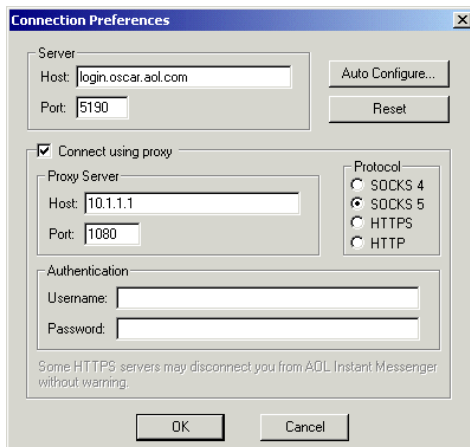


Figure 17-7: AOL IM Client Explicit Proxy Configuration

MSN Messenger Client Explicit Proxy Configuration Screen

The following example configures an MSN Messenger client for explicit proxy.

1. Select Tools>Options.
2. Click Connection.
3. Select I use a proxy server.
4. Select SOCKS Version 5.
5. Enter the server name.

6. Enter the port number (the default is 1080).
7. If authentication is required on the ProxySG, enter the authentication user name and password.

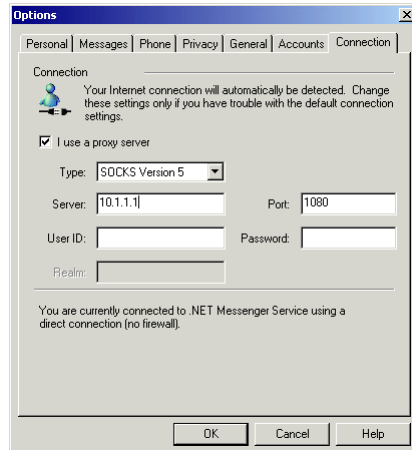


Figure 17-8: MSN IM Client Explicit Proxy Configuration

VPM Examples

Once the IM clients are configured to send traffic through the ProxySG, you can control and limit IM activity. The Visual Policy Manager (VPM) allows you to create rules that control and track IM communications, including IM activities based on users and groups, IM handle, chat room handle, file name, and other triggers.

To learn about the VPM, see Appendix 14: "The Visual Policy Manager".

Example 1: File Transfer

The following example demonstrates an IM rule created with the VPM that IM handle Nigel1 can perform a file transfer at any time, but the file must be between 1 and 5 MB in size, and the handle, the file path, and file size are logged:

1. In the VPM, select Policy>Add Web Access Layer; name it IM_FileTransfer.
2. Right-click the Source field; select Set. The Set Source Object dialog appears.
3. Click New; select IM User. The Add IM User Object dialog appears.
4. In the IM User field, enter Nigel1; click OK in each dialog.
5. Right-click the Service field; select Set. The Set Service Object dialog appears.
6. Click New; select IM File Transfer. The Add IM File Transfer dialog appears.
7. Select Size and enter a range 1 and 5; select MBytes from the drop-down list; click OK in each dialog.
8. Right-click the Track field; select Set. The Add Track Object dialog appears.
9. Click New; select Event Log. The Add Event Log Object dialog appears.

10. From the Substitution Variables list, select x-im-buddy-name and click insert. Repeat for x-im-file-path and x-im-file-size. Click OK in each dialog.
11. Click Install Policy.

Example 2: Send an IM Alert Message

The following example demonstrates a rule created with the VPM that informs all IM users when they login that their IM activity is tracked and logged.

1. In the VPM, select Policy>Add Web Access Layer; name it IM_NotifyMessage.
2. Right-click the Service field; select Set. The Set Service Object dialog appears.
3. Click New; select Protocol Methods. The Add Methods Object dialog appears.
4. From the Protocol drop-down list, select Instant Messaging.
5. Click Login/Logout; LOGIN; click OK to close the dialog; click OK to insert the object in the rule.
6. Right-click the Service field; select Set. The Set Service Object dialog appears.
7. Click New; select Send IM Alert. The Add Send IM Alert Object dialog appears.
8. In the Alert Text field, enter a message that appears to users. For example, Notice: Your Instant Messaging message activity is tracked and logged.
9. Click OK to close the dialog; click OK to insert the object in the rule.

Statistics

The IM statistics allow you to track IM connections, file transfers, and messages that are currently in use and in total, or have been allowed and denied. The information can be displayed for each IM client type or combined.

For information about viewing IM statistics, see "IM History Statistics" on page 826.

Related Material

Refer to the following Blue Coat documentation for related IM information:

- Chapter 14: "The Visual Policy Manager" on page 453
- *Blue Coat ProxySG Content Policy Language Guide*

Chapter 18: *Content Filtering*

The ProxySG allows the use of *content filtering* to control the type of content retrieved by the ProxySG and to filter requests made by clients. You can use a local content-filtering database and/or content-filtering policy to reduce the infinite number of URLs to a small number of categories and then manage those categories. Categories can be used anywhere you would use a URL-based trigger.

You can also combine the ProxySG local database and policies with a content-filtering vendor to provide a cohesive approach to managing access to the Web.

This chapter contains the following topics:

- "Overview"
- "Selecting Category Providers"
- "Configuring a Local Database"
- "Configuring Blue Coat Web Filter"
- "Configuring InterSafe"
- "Configuring Optenet"
- "Configuring SmartFilter"
- "Configuring SurfControl"
- "Configuring Websense"
- "Configuring Webwasher URL Filter"
- "How to Apply Policy to Categorized URLs"
- "Using Content-Filtering Vendors with ProxySG Policies"

Overview

Content filtering allows you to categorize Web sites (such as sports and gambling). Once the Web sites and content are categorized, access to that content can be controlled through policy.

The ProxySG content filtering feature (which requires a license—see Chapter 2: “Licensing” on page 35) allows you to integrate subscription-based filtering lists that are automatically updated and categorized as the Web changes.

Content filtering allows you to block sites based on what you believe to be in them. You can either filter URLs yourself, allowing or denying permission to them using your own local content-filtering database, you can use the Blue Coat Web Filter, or you can use a third-party content-filtering vendor to provide the categories and assign the categories to URLs. Consider that, based on the ever-changing nature of the Web, manually maintaining a local content filter database is an overwhelming task.

Categories and their meanings are defined by the specific category providers. For third-party databases, the most up-to-date information on how categories are assigned to URLs can be obtained from the provider's Web site. You can request that specific URLs be reviewed for correct categorization, if your content-filtering provider supports this.

Examples in this document are believed to be correct at the time of publication, but could be affected by subsequent changes in third-party databases.

Once the content is categorized, you can control access to the categories (using policy) by username, department, time of day, and other criteria.

To use a third-party vendor for content filtering, contact the vendor for license and authorization information. Continue with the appropriate section to configure the properties.

- "Selecting Category Providers" on page 636
- "Configuring a Local Database" on page 640
- "Configuring Blue Coat Web Filter" on page 645
- "Configuring InterSafe" on page 653
- "Configuring Optenet" on page 658
- "Configuring SmartFilter" on page 667
- "Configuring SurfControl" on page 677
- "Configuring Websense" on page 681
- "Configuring Webwasher URL Filter" on page 688

To use policy to create and manage categories, see "How to Apply Policy to Categorized URLs" on page 692. To use policy to refine third-party vendor content filtering, see "Using Content-Filtering Vendors with ProxySG Policies" on page 695.

Selecting Category Providers

You can select a local database, Blue Coat Web Filter (BCWF), or a third-party vendor database for content filtering, view the filtering categories available, and test a URL through either the Management Console or the CLI.

To Select a Local Database, Blue Coat Web Filter, or Third-Party Vendor Database through the Management Console:

1. Select Configuration>Content Filtering>General.

The General tab displays.

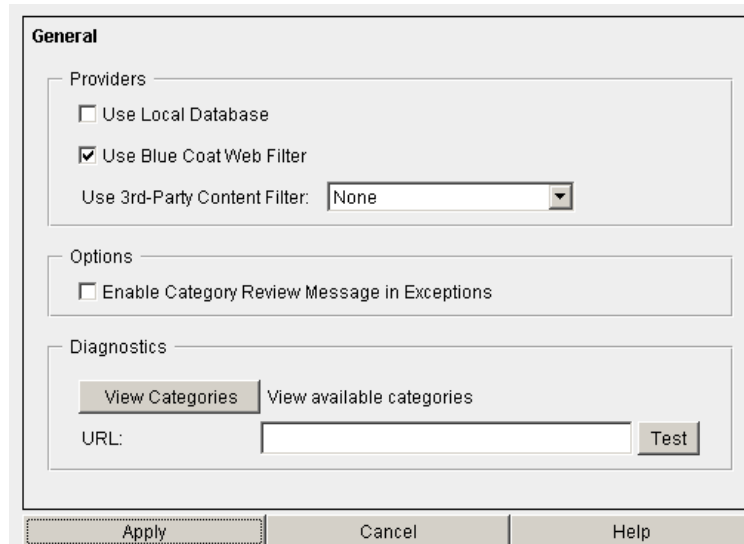


Figure 18-1: Content Filtering, General Tab

2. (Optional) To use a local database for content filtering, select the Use Local Database checkbox.
3. (Optional) To use Blue Coat Web Filter for content filtering, select the Use Blue Coat Web Filter checkbox.

Note: If you select Blue Coat Web Filter, a small database that contains the category list is downloaded immediately, allowing immediate policy creation.

No username or password is required during the trial period (60 days). If you want to download the database on demand or on a schedule, or if you want to try out dynamic categorization, you'll need to configure the BCWF service.

4. (Optional) If you have a licensed third-party vendor that you want to use (either instead of or in addition to a local database or the Blue Coat Web Filter service), select the vendor from the Use 3rd party database drop-down list. Select None to stop using a vendor.
5. (Optional) If you are using a provider that supports it, you can select the Enable Category Review Message in Exceptions. In conjunction with two substitutions—`$(exception.category_review_url)` and `$(exception.category_review_message)`—you can request that specific URLs be reviewed for correct categorization.

If you enable the Category Review Message, the two substitutions are automatically appended to the “help” element of all exception definitions. For information on using the `$(exception.help)` element, see "User-Defined Exceptions" on page 572.

Note: The substitution values are empty if the selected content filter provider does not support review messages, or if the provider was not consulted for categorization, or if the categorization process failed due to an error.

6. Click Apply.

7. (Optional) To see all categories available for use in policy, click View Categories. Note that categories will not be displayed for a vendor or local database if no database has been downloaded.
8. To see what categories a Web site is assigned by your current configuration, enter the URL into the URL field and click Test.

To Select a Local or Third-Party Vendor Database through the CLI:

1. At the (config) command prompt, enter the following command to enter content-filter mode:
SGOS#(config) **content-filter**
2. (Optional) To use a local database for content filtering, enter the following command:
SGOS#(config content-filter) **use-local-database**
3. (Optional) To stop using a local database for content filtering, enter the following command:
SGOS#(config content-filter) **no use-local-database**
4. (Optional) If you have a licensed third-party vendor that you want to use (either instead of or in addition to a local database), enter the following command to select a vendor:
SGOS#(config content-filter) **select-provider {bluecoat | intersafe | proventia | smartfilter | surfcontrol | websense | webwasher}**
5. (Optional) You can request that specific URLs be reviewed for correct categorization.
SGOS#(config content-filter) **review-message | no review-message**
If you enable `review-message`, two substitutions—`$(exception.category_review_url)` and `$(exception.category_review_message)`—are automatically appended to the `help` element of all exception definitions. For information on using the `$(exception.help)` element, see "User-Defined Exceptions" on page 572.
Note: The substitution values are empty if the selected content filter provider does not support review messages, or if the provider was not consulted for categorization, or if the categorization process failed due to an error.
6. (Optional) To stop using any third-party vendor, enter the following command:
SGOS#(config content-filter) **select-provider none**
7. (Optional) To identify the categories assigned by the current configuration to a particular URL, enter the following command:
SGOS#(config content-filter) **test-url url**
where `url` specifies the URL for which you want to identify categories.
8. (Optional) To view all available categories, which might include those created by policy, a local database if enabled, a selected vendor, and the system, enter the following command:
SGOS#(config content-filter) **categories**
Categories defined by Policy:
Sports URLs
Entertainment
Categories defined by Local:
cat1
cat2
cat3
cat4

Categories defined by SurfControl:

```
Web-based Email
Motor Vehicles
.
.
.
Chat
```

(Long list truncated)

Categories defined by System:

```
none
unavailable
unlicensed
```

9. (Optional) View the content-filtering configuration.

```
SGOS#(config content-filter) view
Provider                               Local
Status:                                Ready
Download URL:                           ftp://10.25.36.47/list-1000000-cat.t
Download Username:                       anonymous
Automatic download:                     Enabled
Download time of day (UTC):              0
Download on:                             sun, mon, tue, wed, thu, fri, sat
Download log:
  Local database download at: Wed, 28 Jan 2004 00:19:48 UTC
  Downloading from ftp://10.25.36.47/list-1000000-cat.txt
  Download size:                          16274465
  Database date:                          Wed, 28 Jan 2004 00:22:04 UTC
  Total URL patterns: 1000000
  Total categories: 10
  Provider:                               Websense
Status:                                  Ready
Download License key:                    TUVW67XYZ89ABC0
Download Server:                          download.websense.com
Email contact:
Automatic download:                      Enabled
Download time of day (UTC):              0
Download on:                             sun, mon, tue, wed, thu, fri, sat
Use regular expression filters: No
Config Server:                            Disabled
Config Server listening port:            15870
Download log:
  Websense download at: Wed, 28 Jan 2004 22:11:37 UTC
  Downloading from download.websense.com
  Download size:                          63642227
  Database version:                       71617
  Database date:                          2004-01-28
  License expires:                        Sun, 28 Nov 2004 08:00:00 UTC
  License max users:                      25
  Licenses in use:                        1
```

Configuring a Local Database

You can create your own local database file and download it to the ProxySG. This file is created in the same way that policy files are created, except that only *define category* statements are allowed in the local database. Refer to the *Blue Coat ProxySG Content Policy Language Guide* for information on define category statements, or see "Defining Custom Categories in Policy" on page 696.

Note: You might find it convenient to put your local database on the same server as any policy files you are using.

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 administrating 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.

You can choose to use any combination of the local database, policy files, or VPM to manage your category definitions. See "How to Apply Policy to Categorized URLs" on page 692 for more information. You can also use both a local database and a third-party vendor for your content filtering needs.

Use the ProxySG Management Console or the CLI to configure local database content filtering and to schedule automatic downloads. For information about scheduling automatic downloads, see "Scheduling Automatic Downloads for a Local Database" on page 644.

To Configure Local Database Content Filtering through the Management Console:

1. Select Configuration>Content Filtering>Local Database.

The Local Database tab displays.

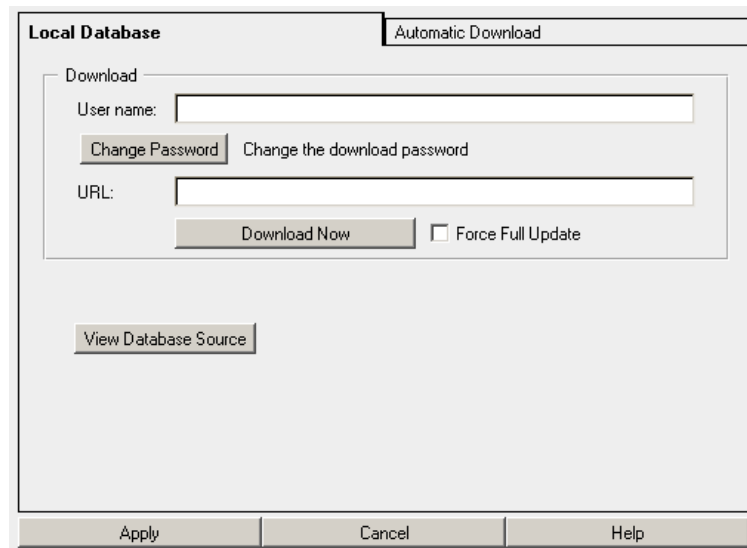


Figure 18-2: Local Database Configuration Tab

2. (Optional) If you need a password to access the download site, click Change Password, enter the password in the Change Password dialog, and click OK.
3. Enter the database download URL in the URL field.
4. (Optional) To see a display of the currently installed text file, click View Database Source; close the display when you are finished.
5. Click Apply.
6. (Optional) To download the local database immediately, complete the following steps.

Note: You can return here at any time to download a database on demand (remember that the automatic download feature, if configured, keeps you up-to-date—see "Scheduling Automatic Downloads for a Local Database" on page 644).

Ordinarily, the ProxySG will check to see if the database has changed before initiating a download. If the database is up to date, then no download is necessary and none is performed. You can override this check and force a download by selecting the Force Full Update checkbox; this option is not needed under normal circumstances.

- a. Click Download Now (if you want to download a full database, first select the Force Full Update checkbox—this option is unnecessary under normal circumstances).

The Local Installation status dialog box displays with the message Local download in progress.

When the operation is complete, the dialog changes to indicate installation status.

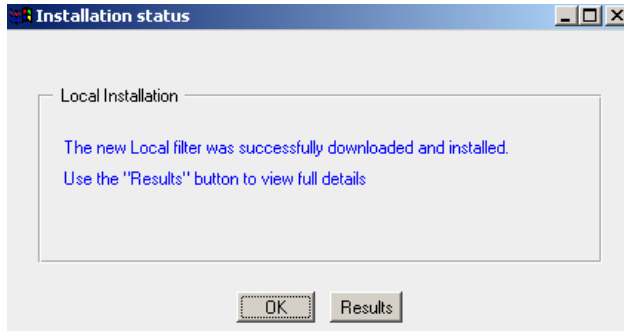


Figure 18-3: Local Database Successfully Downloaded

- b. Click Results to see the completion message:

Download log:

```
Local database download at: Wed, 20 Oct 2004 19:29:39 UTC
Downloading from ftp://10.25.36.47/list-1000000-cat.txt
Download size:      16274465
Database date: Wed, 20 Oct 2004 19:31:58 UTC
Total URL patterns: 1000000
Total categories:  10
```

To Configure Local Database Content Filtering through the CLI:

The following commands allow you to enter the username, specify the URL from which the database is to be downloaded, and do an immediate download of the local database. You can even clear the database and all associated files, if necessary

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) content-filter
SGOS#(config content-filter) local
SGOS#(config local) download username username
SGOS#(config local) download password password
-or-
SGOS#(config local) download encrypted-password encrypted_password
SGOS#(config local) download url url
SGOS#(config local) clear
```

where:

clear		Clears the database from the system.
download username	<i>username</i>	Identifies the username needed to access the download site, if any.
download encrypted- password	<i>encrypted_password</i>	Allows you to take a password previously encrypted by the ProxySG and cut and paste the encrypted password on the same appliance (or another appliance if it shares the same password-display keyring). The primary use of the encrypted-password command is to allow the ProxySG to reload a password that it encrypted.
download password	<i>password</i>	Identifies the password needed to access the download site, if any.
download url	<i>url</i>	The local URL.

2. (Optional) To download the database now, enter one of the following commands:

```
SGOS#(config local) download get-now
-or-
SGOS#(config local) download full-get-now
```

where:

download get-now	Initiates an immediate database download. If the database is already up-to-date, no download will be initiated.
download full-get-now	Initiates an immediate database download, forcing a download whether or not an update is necessary. This option is unnecessary under most circumstances.

3. (Optional) To view the local database source file, enter the following command:

```
SGOS#(config local) source
```

4. (Optional) To view the configuration, enter the following command:

```
SGOS#(config local) view
Status: Ready
Download URL: ftp://10.25.36.47/list-1000000-cat.txt
Download Username: user1
Automatic download: Enabled
Download time of day (UTC): 0
Download on: sun, mon, tue, wed, thu, fri, sat
Download log:
Local database download at: Fri, 02 Jul 2004 18:40:11 UTC
Downloading from ftp://10.25.36.47/list-1000000-cat.txt
Download size: 612
Database date: Fri, 02 Jul 2004 18:38:57 UTC
Total URL patterns: 8
Total categories: 5
```

Scheduling Automatic Downloads for a Local Database

Note: By default, the automatic download setting is enabled (for every day at midnight, UTC) and does not need to be configured unless you want to change the schedule or disable auto-download.

To download the local database without creating a schedule, see "Configuring a Local Database" on page 640.

The Automatic Download tab allows you to set the times the local database will be downloaded. You can specify an automatic download on the day and time you prefer. Because sites become stale quickly, Blue Coat recommends downloading on an automatic schedule frequently.

When the database is downloaded, a log is available that includes the information about how the database was updated, but in a more detailed form. You can view the download log through the Management Console (Statistics>Advanced>Content Filter Service) or the CLI (SGOS#(config) `show content-filter status`).

Setting Local Database Automatic Download Times through the Management Console:

1. Select Configuration>Content Filtering>Local Database>Automatic Download.

The Automatic Download tab displays.

Figure 18-4: Local Database Automatic Download Tab

2. To set up a schedule for local database downloads, select the Download new filter at checkbox and select the time of day from the drop-down list. The default is Midnight.
3. All days are selected by default. Deselect days as needed.
4. Click Apply when finished.

Setting Local Database Automatic Download Times through the CLI:

1. At the `(config)` command prompt, enter the following commands to enable or disable automatic downloading of the local database.

```
SGOS#(config) content-filter
SGOS#(config content-filter) local
SGOS#(config local) download auto
-or-
SGOS#(config local) no download auto
```

2. At the `(config local)` command prompt, enter the following command to select the day(s) to automatically download the local database.

```
SGOS#(config local) download day-of-week {all | none | sun | mon | tue | wed
| thu | fri | sat}
-or-
SGOS#(config local) no download day-of-week {sun | mon | tue | wed | thu | fri
| sat}
```

where `all` selects all days of the week, and `none` clears all days of the week from the schedule.

All days are selected by default; to deselect days, enter `none` and enter specific days. You can only select one day each time, but it is appended to the list. You can also use the `no download day-of-week` command to clear specific days from the schedule.

3. Enter the following command to specify the hour (UTC) of the selected days during which the download should be performed.

```
SGOS#(config local) download time-of-day 0-23
```

4. (Optional) To download the local database now, enter the following command:

```
SGOS#(config local) download get-now
```

Downloading the database now does not affect the automatic database download schedule.

Configuring Blue Coat Web Filter

Blue Coat Web Filter (BCWF) is a highly effective content filtering service that can quickly learn and adapt to the working set of its users. Also, BCWF provides a network service that can dynamically examine and categorize web pages as they are requested. This dynamic real-time categorization enhances both the accuracy and freshness of the BCWF filtering solution.

Note: If you selected Blue Coat Web Filter on the General tab, a small database that contains the category list was downloaded immediately, while the full BCWF database loaded in the background. All filtering was done by the BCWF dynamic categorization service while the full database downloaded.

No username or password is required during the trial period (60 days). If you want to download the database on demand or on a schedule or use the dynamic categorization service, you'll need to configure the BCWF service as described in the following sections.

SGOS 4.1 introduces integration with this dynamic categorization service. The BCWF categorization service is an Internet service, available from designated service points with high-bandwidth connections and dedicated hardware.

Note: Dynamic analysis of content is done on a remote network service, and not locally on the ProxySG. If the category returned by this service is blocked by policy, the offending material never enters the network in any form.

For information on configuring dynamic categorization, see "Configuring Dynamic Categorization for Blue Coat Web Filter" on page 651.

Use the ProxySG Management Console or the CLI to configure BCWF.

To Configure Blue Coat Web Filter through the Management Console:

1. Select Configuration>Content Filtering>Blue Coat Web Filter.

The Blue Coat Web Filter tab displays.

The screenshot shows a web-based configuration interface for the Blue Coat Web Filter. The main window has a title bar with 'Blue Coat Web Filter' and two sub-tabs: 'Automatic Download' and 'Dynamic Categorization'. The 'Dynamic Categorization' sub-tab is selected. Below the sub-tabs is a 'Download' section. It contains a 'User name:' label followed by an empty text input field. Below this is a 'Change Password' button with the text 'Change the download password' next to it. Below that is a 'URL:' label followed by a text input field containing the URL 'https://list.bluecoat.com/bcwf/activity/download/bcwf.db'. Below the URL field are two buttons: 'Set to default' and 'Download Now'. To the right of the 'Download Now' button is a checkbox labeled 'Force Full Update'. At the bottom of the window are three buttons: 'Apply', 'Cancel', and 'Help'.

Figure 18-5: Blue Coat Web Filter Configuration Tab

2. When you subscribed to the BCWF Service, you received a username and password for access to download updates. Enter your username into the Username field and click the Change Password button to enter or change your password. (If you are in the trial period, no username or password is required.)
3. The default database download location is displayed in the URL field. If you have been instructed to use a different URL, enter it here. You can restore the default at any time by clicking Set to default.
4. (Optional) To download the Blue Coat Web Filter database immediately, complete the following steps.

Note: You can return here at any time to download a database on demand (remember that the automatic download feature, if configured, keeps you up to date—see "Scheduling Automatic Downloads for Blue Coat Web Filter" on page 649).

Ordinarily, the ProxySG checks to see if the database has changed before initiating a download. If the database is up to date, no download is necessary and none is performed. If an incremental update is available on the server, then it is downloaded (an incremental update contains only the changes between the current installed version and the latest published version of the database, and is much smaller than a full copy of the database). You can override this process and force a download of the full database by selecting the Force Full Update checkbox; this option is not needed under normal circumstances.

- a. Click Download Now (if you want to download a full database, first select the Force Full Update checkbox—this option is unnecessary under normal circumstances).

The Blue Coat Web Filter Installation status dialog box displays with the message Blue Coat Web Filter download in progress.

When the operation is complete, the dialog changes to indicate installation status.

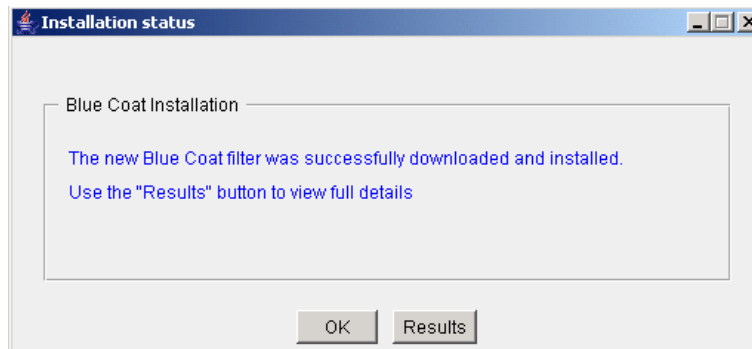


Figure 18-6: Blue Coat Web Filter Database Successfully Downloaded

- b. Click Results to see the Blue Coat Web Filter download log:

```
Download log:
Blue Coat download at: Thu, 10 Feb 2005 00:04:06 UTC
Downloading from https://list.bluecoat.com/bcwf/activity/download/bcwf.db
Requesting differential update
Differential update applied successfully
Download size:      84103448
Database date:     Wed, 09 Feb 2005 08:11:51 UTC
Database expires: Fri, 11 Mar 2005 08:11:51 UTC
Database version:  2005040
```

To Configure Blue Coat Web Filter Content Filtering through the CLI:

The following commands allow you to enter the Blue Coat Web Filter username and password and define the default URL and the default URL location.

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) content-filter
SGOS#(config content-filter) bluecoat
```

```
SGOS#(config bluecoat) download username username
SGOS#(config bluecoat) download password password
-or-
SGOS#(config bluecoat) download encrypted-password encrypted-password
SGOS#(config bluecoat) download url {default | url}
```

where:

download username	<i>username</i>	Specifies the username assigned to you for database download. If you are in the trial period, no username is required.
download encrypted-password	<i>encrypted_password</i>	Allows you to take a password previously encrypted by the ProxySG and cut and paste the encrypted password on the same appliance (or another appliance if it shares the same password-display keyring). The primary use of the encrypted-password command is to allow the ProxySG to reload a password that it encrypted. If you are in the trial period, no password is required.
download password	<i>password</i>	Specifies the password assigned to you for database download. If you are in the trial period, no password is required.
download url	default	Specifies the use of the default download URL.
	<i>url</i>	The URL is the Blue Coat Web Filter URL. You can change it if directed to do so.

2. (Optional) To download the database now, enter one of the following commands:

```
SGOS#(config bluecoat) download get-now
-or-
SGOS#(config bluecoat) download full-get-now
```

where:

download get-now	Initiates an immediate database download. An incremental update is requested.
download full-get-now	Initiates an immediate full-size database download. This option is unnecessary under most circumstances.

3. (Optional) View the configuration.

```
SGOS#(config bluecoat) view
Status: Ready
Download URL: https://list.bluecoat.com/bcwf/activity/download/bcwf.db
Download Username:
Automatic download: Enabled
Download time of day (UTC): 0
Download on: sun, mon, tue, wed, thu, fri, sat
Download log:
```

```
Download log:
Blue Coat download at: Thu, 10 Feb 2005 00:04:06 UTC
Downloading from
  https://list.bluecoat.com/bcwf/activity/download/bcwf.db
Requesting differential update
Differential update applied successfully
Download size:      84103448
Database date:     Wed, 09 Feb 2005 08:11:51 UTC
Database expires:  Fri, 11 Mar 2005 08:11:51 UTC
Database version:  2005040
```

Scheduling Automatic Downloads for Blue Coat Web Filter

Note: By default, the automatic download setting is enabled (for every day at midnight, UTC) and does not need to be configured unless you want to change the schedule or disable auto-download.

To download the Blue Coat Web Filter database without creating a schedule, see "Configuring Blue Coat Web Filter" on page 645.

The Automatic Download tab allows you to set the times at which the Blue Coat Web Filter database is downloaded. You can specify an automatic download on the day and time you prefer. Because sites become stale quickly, Blue Coat recommends downloading on an automatic schedule frequently.

When the database is downloaded, a log is available that includes the information about how the database was updated, but in a more detailed form. You can view the download log through the Management Console (Statistics>Advanced>Content Filter Service) or the CLI (SGOS#(config) `show content-filter status`).

Setting Blue Coat Web Filter Automatic Download Times through the Management Console:

1. Select Configuration>Content Filtering>Blue Coat Web Filter>Automatic Download.

The Automatic Download tab displays.

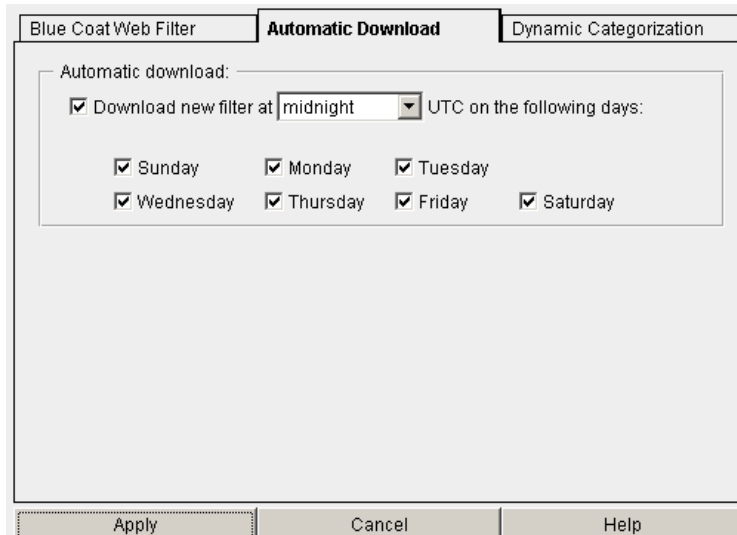


Figure 18-7: Blue Coat Web Filter Automatic Download Tab

2. To set up a schedule for Blue Coat Web Filter database downloads, select the Download new filter at checkbox and select the time of day from the drop-down list. The default is Midnight.
3. All days are selected by default. Clear checkboxes as needed.
4. Click Apply when finished.

Setting Blue Coat Web Filter Automatic Download Times through the CLI:

1. At the `(config)` command prompt, enter the following commands to enable or disable automatic downloading of the Blue Coat Web Filter database.

```
SGOS#(config) content-filter
SGOS#(config content-filter) bluecoat
SGOS#(config bluecoat) download auto
-or-
SGOS#(config bluecoat) no download auto
```

2. At the `(config bluecoat)` command prompt, enter the following commands to select or deselect the day(s) to automatically download the local database.

```
SGOS#(config bluecoat) download day-of-week {all | none | sun | mon | tue |
wed | thu | fri | sat}
-or-
SGOS#(config bluecoat) no download day-of-week {sun | mon | tue | wed | thu |
fri | sat}
```

where `all` selects all days of the week, and `none` clears all days of the week from the schedule.

All days are selected by default; to deselect days, enter `none` and enter specific days. You can only select one day each time, but it is appended to the list. You can also use the `no download day-of-week` command to clear specific days from the schedule.

3. Enter the following command to specify the hour (UTC) of the selected days during which the download should be performed.

```
SGOS#(config bluecoat) download time-of-day 0-23
```

4. (Optional) To download the Blue Coat Web Filter database now, enter the following command:

```
SGOS#(config bluecoat) download get-now
```

Downloading the database now does not affect the automatic database download schedule.

Configuring Dynamic Categorization for Blue Coat Web Filter

The Dynamic Categorization tab allows you to enable or disable dynamic categorization and its various forms of behavior on the ProxySG. By default, dynamic categorization is enabled.

Note: The dynamic service is consulted only when the installed BCWF database does not contain complete categorization for an object. This dispatch mechanism is independent of results from other categorization services.

Administrators might need to process certain URL requests in real time while other requests can be done in the background. Some requests might avoid dynamic categorization entirely as circumstances dictate. Therefore, the choice of real-time, background mode, or no dynamic categorization for each URL categorization can be made on a per-transaction basis using Blue Coat policy. The configuration establishes a default mode; Blue Coat policy can override that default. For more information, see Chapter 14: "The Visual Policy Manager" or refer to the *Blue Coat ProxySG Content Policy Language Guide*.

Dynamic categorization has two types of cost:

- **Bandwidth:** Represents the round trip request/response from the ProxySG to the service. Because the dynamic categorization protocol is compact, this cost is minimal.
- **Latency:** Represents the time spent waiting for the dynamic categorization service to provide a result.

Remember that these costs are only incurred when a URL cannot be categorized by a database lookup on the ProxySG. SGOS 4.x offers three modes of operation to compensate for some of this cost:

- **Real-time mode (default).** Real-time mode incurs both bandwidth and latency costs. The advantage of real-time mode dynamic categorization is that Blue Coat policy has access to the results of dynamic categorization, which means that policy decisions are made immediately upon receiving all available information.
- **Background mode.** Background mode incurs only the bandwidth cost. In background mode once a call is made to the dynamic categorization service, the URL request immediately proceeds without waiting for the external service to respond. The system category *pending* is assigned to the request, indicating that the policy was evaluated with potentially incomplete category information.

Once received, the results of dynamic categorization are entered into a categorization cache (as are the results of real-time requests). This cache ensures that any subsequent requests for the same or similar URLs can be categorized quickly, without needing to query the external service again.

- Do not categorize. Dynamic categorization is not done (unless explicitly requested by policy). This mode is distinct from disabling the service. When *Do not categorize* is set as the default, dynamic categorization (in either real time or background mode) can be explicitly invoked by policy. When the service is disabled, no dynamic categorization is done, regardless of policy, and the ProxySG does not make any contact with the dynamic categorization service.

To Configure Dynamic Categorization through the Management Console

1. Select Configuration>Content Filtering>Blue Coat>Dynamic Categorization.

The Dynamic Categorization tab displays.

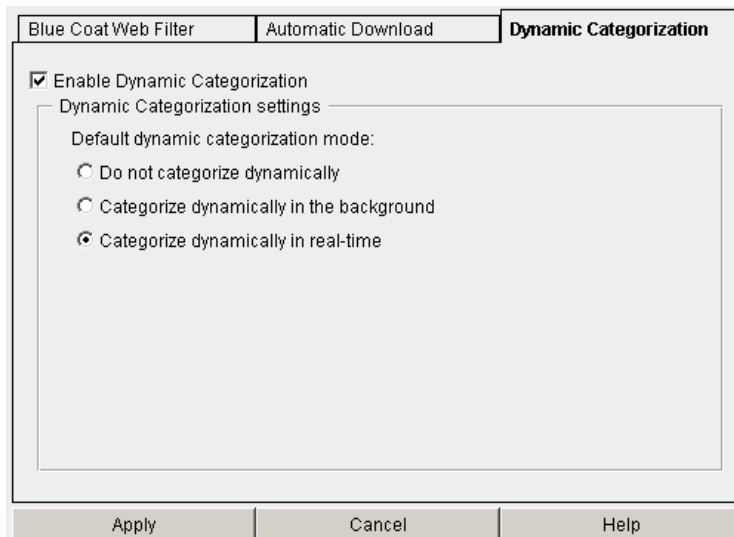


Figure 18-8: Blue Coat Web Filter Dynamic Categorization

Dynamic Categorization is enabled by default. To disable it, clear the checkbox. If dynamic categorization is disabled, then the ProxySG does not contact the dynamic categorization service, even when no category is found for a URL in the database, and any dynamic categorization properties specified in policy are ignored. Note that even if dynamic categorization is enabled, it is only invoked while BCWF is in use

2. To change the Dynamic Categorization Settings, choose the radio button for the choice you need.
 - Do not categorize dynamically. The loaded database is consulted for category information. URLs not found in the database show up as category *none*.
 - Categorize dynamically in the background. Objects not categorized by the database are dynamically categorized when time permits.
 - Categorize dynamically in real-time, the default. Objects not categorized by the database are dynamically categorized.

To Configure Dynamic Categorization through the CLI

The following commands allow you to analyze and manage requested Web pages in real time. You can also configure dynamic categorization settings and specify behavior when doing dynamic categorization.

1. At the `(config)` command prompt, enter the following commands:


```

SGOS#(config) content-filter
SGOS#(config content-filter) bluecoat
SGOS#(config bluecoat) service enable | disable
SGOS#(config bluecoat) service mode {background | realtime | none}

```

where:

service	enable disable	Enable or disable dynamic categorization. Dynamic categorization is enabled by default.
service mode	background realtime none	Do dynamic categorization one of three ways: <ul style="list-style-type: none"> • background: Objects not categorized by the database are dynamically categorized when time permits. • realtime: The default. Objects not categorized by the database are dynamically categorized. • none: The loaded database is consulted for category information. URLs not in the database show up as category <i>none</i>.

Configuring InterSafe

Use the ProxySG Management Console or the CLI to configure InterSafe content filtering.

To Configure InterSafe Content Filtering through the Management Console:

1. Select Configuration>Content Filtering>InterSafe.

The InterSafe tab displays.

Figure 18-9: InterSafe Configuration Tab

2. Enter the username and password assigned to you for downloading the InterSafe database: enter your username into the Username field and click the Change Password button to enter or change your password.

3. The default database download location is displayed in the URL field. If you have been instructed to use a different URL, enter it here. You can restore the default at any time by clicking Set to default.
4. (Optional) To download the InterSafe database immediately, complete the following steps.

Note: You can return here at any time to download a database on demand (remember that the automatic download feature, if configured, keeps you up-to-date—see "Scheduling Automatic Downloads for InterSafe" on page 656).

Ordinarily, the ProxySG checks to see if the database has changed before initiating a download. If the database is up to date, then no download is necessary and none is performed. If an incremental update is available on the server, then it is downloaded (an incremental update contains only the changes between the current installed version and the latest published version of the database, and is much smaller than a full copy of the database). You can override this process and force a download of the full database by selecting the Force Full Update checkbox; this option is not needed under normal circumstances.

- a. Click Download Now (if you want to download a full database, first select the Force Full Update checkbox—this option is unnecessary under normal circumstances).

The InterSafe Installation status dialog box displays with the message InterSafe download in progress.

When the operation is complete, the dialog changes to indicate installation status.

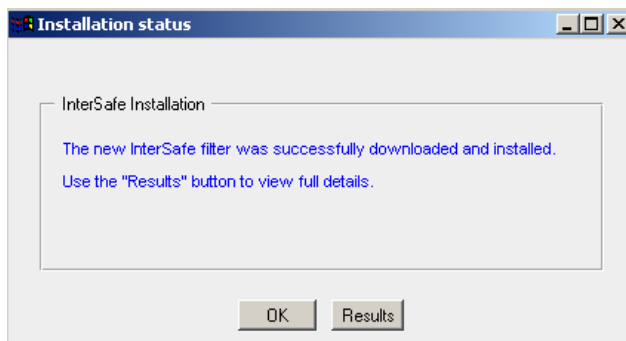


Figure 18-10: InterSafe Database Successfully Downloaded

- b. Click Results to see the InterSafe download log:

Download log:

```
InterSafe download at: Tue, 28 Sep 2004 20:16:16 UTC
Downloading from https://list.bluecoat.com/.../download/intersafe.db
Warning: Unable to determine current database version; requesting full
update
Download size:      8106572
Database date:     Fri, 10 Sep 2004 07:02:08 UTC
Database expires:  Sun, 10 Oct 2004 07:02:08 UTC
Database version:  3
```

To Configure InterSafe Content Filtering through the CLI:

The following commands allow you to enter the InterSafe username and password and define the default URL and the default URL location.

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) content-filter
SGOS#(config content-filter) intersafe
SGOS#(config intersafe) download username username
SGOS#(config intersafe) download password password
-or-
SGOS#(config intersafe) download encrypted-password encrypted-password
SGOS#(config intersafe) download url {default | url}
```

where:

download username	<i>username</i>	Specifies the username assigned to you for database download.
download encrypted-password	<i>encrypted_password</i>	Allows you to take a password previously encrypted by the ProxySG and cut and paste the encrypted password on the same appliance (or another appliance if it shares the same password-display keyring). The primary use of the encrypted-password command is to allow the ProxySG to reload a password that it encrypted.
download password	<i>password</i>	Specifies the password assigned to you for database download.
download url	default	Specifies the use of the default download URL.
	<i>url</i>	The URL is the InterSafe URL. You can change it if directed to do so.

2. (Optional) To download the database now, enter one of the following commands:

```
SGOS#(config intersafe) download get-now
-or-
SGOS#(config intersafe) download full-get-now
```

where:

download get-now	Initiates an immediate database download. An incremental update is requested.
download full-get-now	Initiates an immediate full-size database download. This option is unnecessary under most circumstances.

3. (Optional) View the configuration.

```
SGOS#(config intersafe) view
Status:                               Ready
Download URL:                          https://list.bluecoat.com/.../download/intersafe.db
Download Username:                      admin
Automatic download:                    Enabled
Download time of day (UTC):             0
Download on:                            sun, mon, tue, wed, thu, fri, sat
Download log:
  InterSafe download at: Tue, 28 Sep 2004 20:23:16 UTC
  Downloading from https://list.bluecoat.com/.../download/intersafe.db
  Requesting differential update
  Warning: Unable to determine current database version; requesting full update
  Download size:                        8106572
  Database date:                        Fri, 10 Sep 2004 07:02:08 UTC
  Database expires: Sun, 10 Oct 2004 07:02:08 UTC
  Database version: 3
```

Scheduling Automatic Downloads for InterSafe

Note: By default, the automatic download setting is enabled (for every day at midnight, UTC) and does not need to be configured unless you want to change the schedule or disable auto-download.

To download the InterSafe database without creating a schedule, see "Configuring InterSafe" on page 653.

The Automatic Download tab allows you to set the times at which the InterSafe database is downloaded. You can specify an automatic download on the day and time you prefer. Because sites become stale quickly, Blue Coat recommends downloading on an automatic schedule frequently.

When the database is downloaded, a log is available that includes the information about how the database was updated, but in a more detailed form. You can view the download log through the Management Console (Statistics>Advanced>Content Filter Service) or the CLI (SGOS#(config) show content-filter status).

Setting InterSafe Automatic Download Times through the Management Console:

1. Select Configuration>Content Filtering>InterSafe>Automatic Download.

The Automatic Download tab displays.

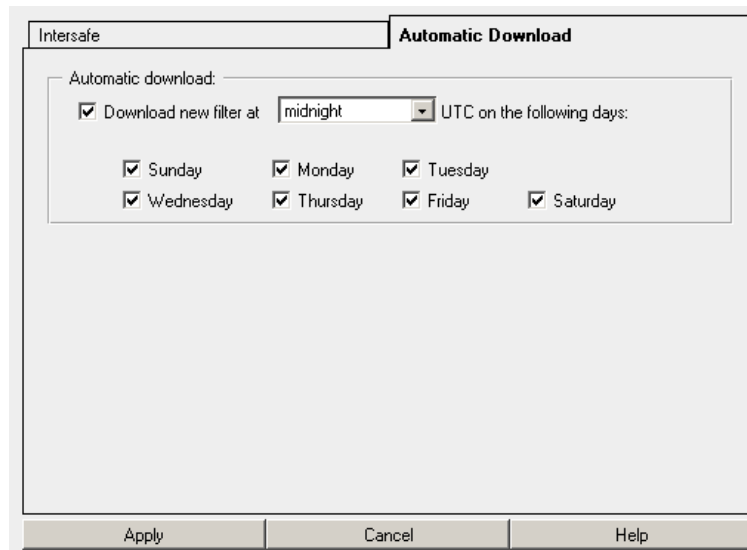


Figure 18-11: InterSafe Automatic Download Tab

2. To set up a schedule for InterSafe database downloads, select the Download new filter at checkbox and select the time of day from the drop-down list. The default is Midnight.
3. All days are selected by default. Deselect days as desired.
4. Click Apply when finished.

Setting InterSafe Automatic Download Times through the CLI:

1. At the `(config)` command prompt, enter the following commands to enable or disable automatic downloading of the InterSafe database.

```
SGOS#(config) content-filter
SGOS#(config content-filter) intersafe
SGOS#(config intersafe) download auto
-or-
SGOS#(config intersafe) no download auto
```

2. At the `(config intersafe)` command prompt, enter the following commands to select or deselect the day(s) to automatically download the local database.

```
SGOS#(config intersafe) download day-of-week {all | none | sun | mon | tue | wed | thu | fri | sat}
-or-
SGOS#(config intersafe) no download day-of-week {sun | mon | tue | wed | thu | fri | sat}
```

where `all` selects all days of the week, and `none` clears all days of the week from the schedule.

All days are selected by default; to deselect days, enter `none` and enter specific days. You can only select one day each time, but it is appended to the list. You can also use the `no download day-of-week` command to clear specific days from the schedule.

3. Enter the following command to specify the hour (UTC) of the selected days during which the download should be performed.

```
SGOS#(config intersafe) download time-of-day 0-23
```

4. (Optional) To download the InterSafe database now, enter the following command:

```
SGOS#(config intersafe) download get-now
```

Downloading the database now does not affect the automatic database download schedule.

Configuring Optenet

Use the ProxySG Management Console or the CLI to configure Optenet content filtering.

To Configure Optenet Content Filtering through the Management Console:

1. Select Configuration>Content Filtering>Optenet.

The Optenet tab displays.

Figure 18-12: Optenet Configuration Tab

2. Enter the Optenet username and password: enter your username into the Username field and click the Change Password button to enter or change your password.
3. The default database download location is displayed in the URL field. If you have been instructed to use a different URL, enter it here. You can restore the default at any time by clicking Set to default.
4. (Optional) To download the Optenet database immediately, complete the following steps.

Note: You can return here at any time to download a database on demand (remember that the automatic download feature, if configured, keeps you up-to-date—see "Scheduling Automatic Downloads for Optenet" on page 661).

Ordinarily, the ProxySG checks to see if the database has changed before

initiating a download. If the database is up to date, then no download is necessary and none is performed. If an incremental update is available on the server, then it is downloaded (an incremental update contains only the changes between the current installed version and the latest published version of the database, and is much smaller than a full copy of the database). You can override this process and force a download of the full database by selecting the Force Full Update checkbox; this option is not needed under normal circumstances.

- a. Click Download Now (if you want to download a full database, first select the Force Full Update checkbox—this option is unnecessary under normal circumstances).

The Optenet Installation status dialog box displays with the message Optenet download in progress.

When the operation is complete, the dialog changes to indicate installation status.

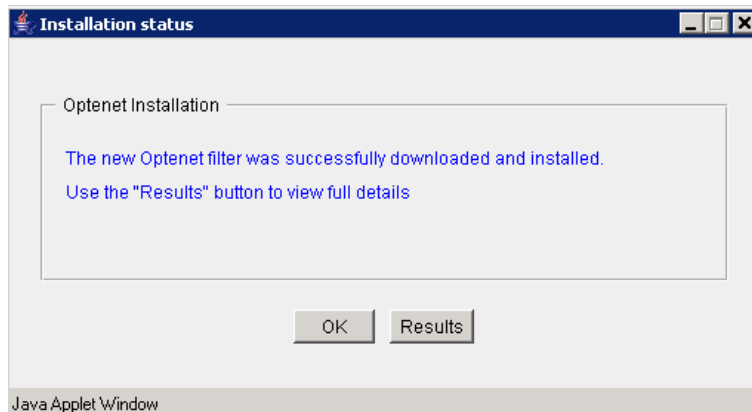


Figure 18-13: Optenet Database Successfully Downloaded

- b. Click Results to see the Optenet download log:

Download log:

```
Optenet download at: Fri, 04 Mar 2005 21:21:06 UTC
Downloading from
https://list.bluecoat.com/optenet/activity/download/optenet.db
Warning: Unable to determine current database version; requesting full update
Download size:      8681732
Database date:     Tue, 01 Mar 2005 17:27:03 UTC
Database expires:  Thu, 31 Mar 2005 17:27:03 UTC
Database version:  2
```

To Configure Optenet Content Filtering through the CLI:

The following commands allow you to enter the Optenet username and password and define the default URL and the default URL location.

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) content-filter
SGOS#(config content-filter) optenet
```

```

SGOS#(config optenet) download username username
SGOS#(config optenet) download password password
-or-
SGOS#(config optenet) download encrypted-password encrypted-password
SGOS#(config optenet) download url {default | url}
    
```

where:

download username	<i>username</i>	Specifies the username assigned to you for database download.
download encrypted-password	<i>encrypted_password</i>	Allows you to take a password previously encrypted by the ProxySG and cut and paste the encrypted password on the same appliance (or another appliance if it shares the same password-display keyring). The primary use of the encrypted-password command is to allow the ProxySG to reload a password that it encrypted.
download password	<i>password</i>	Specifies the password assigned to you for database download.
download url	default	Specifies the use of the default download URL.
	<i>url</i>	The URL is the Optenet URL. You can change it if directed to do so.

2. (Optional) To download the database now, enter one of the following commands:

```

SGOS#(config optenet) download get-now
-or-
SGOS#(config optenet) download full-get-now
    
```

where:

download get-now	Initiates an immediate database download. An incremental update is requested.
download full-get-now	Initiates an immediate full-size database download. This option is unnecessary under most circumstances.

3. (Optional) View the configuration.

```

SGOS#(config optenet) view
Status:                               Ready
Download URL:
https://list.bluecoat.com/optenet/activity/download/optenet.db
Download Username:                     OPTENET-USER
Automatic download:                   Enabled
Download time of day (UTC):           0
Download on:                           sun, mon, tue, wed, thu, fri, sat
Download log:
  Optenet download at: Fri, 04 Mar 2005 21:29:41 UTC
  Downloading from
  https://list.bluecoat.com/optenet/activity/download/optenet.db
  Requesting differential update
    
```



```
File has not changed since last download attempt; no download required
Previous download:
Optenet download at: Fri, 04 Mar 2005 21:21:06 UTC
Downloading from
https://list.bluecoat.com/optenet/activity/download/optenet.db
Warning: Unable to determine current database version; requesting full update
Download size:      8681732
Database date:     Tue, 01 Mar 2005 17:27:03 UTC
Database expires:  Thu, 31 Mar 2005 17:27:03 UTC
Database version:  2
```

Scheduling Automatic Downloads for Optenet

Note: By default, the automatic download setting is enabled (for every day at midnight, UTC) and does not need to be configured unless you want to change the schedule or disable auto-download.

To download the Optenet database without creating a schedule, see "Configuring Optenet" on page 658.

The Automatic Download tab allows you to set the times at which the Optenet database is downloaded. You can specify an automatic download on the day and time you prefer. Because sites become stale quickly, Blue Coat recommends downloading on an automatic schedule frequently.

When the database is downloaded, a log is available that includes the information about how the database was updated, but in a more detailed form. You can view the download log through the Management Console (Statistics>Advanced>Content Filter Service) or the CLI (SGOS#(config) show content-filter status).

Setting Optenet Automatic Download Times through the Management Console:

1. Select Configuration>Content Filtering>Optenet>Automatic Download.

The Automatic Download tab displays.

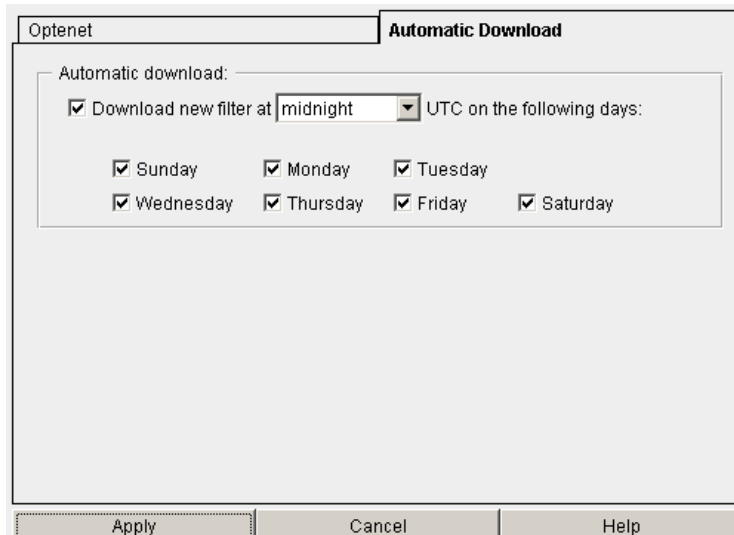


Figure 18-14: Optenet Automatic Download Tab

2. To set up a schedule for Optenet database downloads, select the Download new filter at checkbox and select the time of day from the drop-down list. The default is Midnight.
3. All days are selected by default. Deselect days as desired.
4. Click Apply when finished.

Setting Optenet Automatic Download Times through the CLI:

1. At the `(config)` command prompt, enter the following commands to enable or disable automatic downloading of the Optenet database.

```
SGOS#(config) content-filter
SGOS#(config content-filter) optenet
SGOS#(config optenet) download auto
-or-
SGOS#(config optenet) no download auto
```

2. At the `(config optenet)` command prompt, enter the following commands to select or deselect the day(s) to automatically download the local database.

```
SGOS#(config optenet) download day-of-week {all | none | sun | mon | tue |
wed | thu | fri | sat}
-or-
SGOS#(config optenet) no download day-of-week {sun | mon | tue | wed | thu |
fri | sat}
```

where `all` selects all days of the week, and `none` clears all days of the week from the schedule.

All days are selected by default; to deselect days, enter `none` and enter specific days. You can only select one day each time, but it is appended to the list. You can also use the `no download day-of-week` command to clear specific days from the schedule.

3. Enter the following command to specify the hour (UTC) of the selected days during which the download should be performed.

```
SGOS#(config optenet) download time-of-day 0-23
```

4. (Optional) To download the Optenet database now, enter the following command:

```
SGOS#(config optenet) download get-now
```

Downloading the database now does not affect the automatic database download schedule.

Configuring Proventia Web Filter

Use the ProxySG Management Console or the CLI to configure Proventia Web Filter content filtering.

To Configure Proventia Web Filter Content Filtering through the Management Console:

1. Select Configuration>Content Filtering>Proventia.

The Proventia Web Filter tab displays.

Figure 18-15: Proventia Web Filter Configuration Tab

2. Enter the Proventia Web Filter username and password: enter your username into the Username field and click the Change Password button to enter or change your password.
3. The default database download location is displayed in the URL field. If you have been instructed to use a different URL, enter it here. You can restore the default at any time by clicking Set to default.
4. (Optional) To download the Proventia Web Filter database immediately, complete the following steps.

Note: You can return here at any time to download a database on demand (remember that the automatic download feature, if configured, keeps you up-to-date—see "Scheduling Automatic Downloads for Proventia Web Filter" on page 666).

Ordinarily, the ProxySG will check to see if the database has changed before initiating a download. If the database is up to date, then no download is necessary and none is performed. If an incremental update is available on the server, then it

is downloaded (an incremental update contains only the changes between the current installed version and the latest published version of the database, and is much smaller than a full copy of the database). You can override this process and force a download of the full database by selecting the Force Full Update checkbox; this option is not needed under normal circumstances.

- a. Click Download Now (if you want to download a full database, first select the Force Full Update checkbox—this option is unnecessary under normal circumstances).

The Proventia Installation status dialog box displays with the message Proventia download in progress.

When the operation is complete, the dialog changes to indicate installation status.

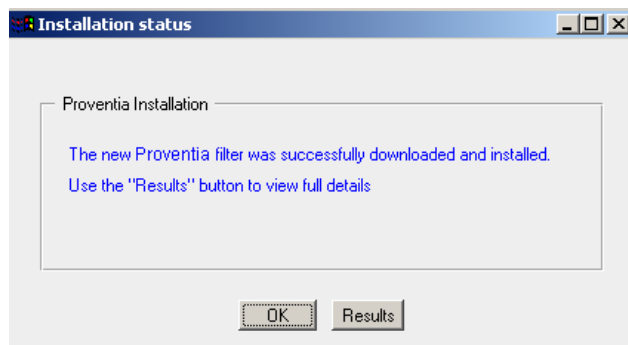


Figure 18-16: Proventia Web Filter Database Successfully Downloaded

- b. Click Results to see the Proventia Web Filter download log:

Download log:

```
Proventia download at: Sat, 10 Jul 2004 18:54:43 UTC
Downloading from
http://list.bluecoat.com/proventia/activity/download/proventia.db
Requesting differential update
Download size:      144913364
Database date:     Wed, 16 Jun 2004 09:40:34 UTC
Database expires: Sat, 06 Feb 2106 06:28:16 UTC
Database version:  16777216
```

To Configure Proventia Web Filter Content Filtering through the CLI:

The following commands allow you to enter the Proventia Web Filter username and password and define the default URL and the default URL location.

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) content-filter
SGOS#(config content-filter) proventia
SGOS#(config proventia) download username username
SGOS#(config proventia) download password password
-or-
SGOS#(config proventia) download encrypted-password encrypted-password
SGOS#(config proventia) download url {default | url}
```

where:

download username	<i>username</i>	Specifies the username assigned to you for database download.
download encrypted-password	<i>encrypted_password</i>	Allows you to take a password previously encrypted by the ProxySG and cut and paste the encrypted password on the same appliance (or another appliance if it shares the same password-display keyring). The primary use of the encrypted-password command is to allow the ProxySG to reload a password that it encrypted.
download password	<i>password</i>	Specifies the password assigned to you for database download.
download url	default	Specifies the use of the default download URL.
	<i>url</i>	The URL is the Proventia Web Filter URL. You can change it if directed to do so.

2. (Optional) To download the database now, enter one of the following commands:

```
SGOS#(config proventia) download get-now
```

-or-

```
SGOS#(config proventia) download full-get-now
```

where:

download get-now	Initiates an immediate database download. An incremental update will be requested.
download full-get-now	Initiates an immediate full-size database download. This option is unnecessary under most circumstances.

3. (Optional) View the configuration.

```
SGOS#(config proventia) view
```

```
Status: Ready
```

```
Download URL:
```

```
http://list.bluecoat.com/proventia/activity/download/proventia.db
```

```
Download Username:
```

```
Automatic download: Enabled
```

```
Download time of day (UTC): 0
```

```
Download on: sun, mon, tue, wed, thu, fri, sat
```

```
Download log:
```

```
Proventia download at: Sat, 10 Jul 2004 19:21:51 UTC
```

```
Downloading from
```

```
http://list.bluecoat.com/proventia/activity/download/proventia.db
```

```
Requesting differential update
```

```
Download size: 144913364
```

```
Database date: Wed, 16 Jun 2004 09:40:34 UTC
```

```
Database expires: Sat, 06 Feb 2106 06:28:16 UTC
```

```
Database version: 16777216
```

Scheduling Automatic Downloads for Proventia Web Filter

Note: By default, the automatic download setting is enabled (for every day at midnight, UTC) and does not need to be configured unless you want to change the schedule or disable auto-download.

To download the Proventia Web Filter database without creating a schedule, see "Configuring Optenet" on page 658.

The Automatic Download tab allows you to set the times at which the Proventia Web Filter database is downloaded. You can specify an automatic download on the day and time you prefer. Because sites become stale quickly, Blue Coat recommends downloading on an automatic schedule frequently.

When the database is downloaded, a log is available that includes the information about how the database was updated, but in a more detailed form. You can view the download log through the Management Console (Statistics>Advanced>Content Filter Service) or the CLI (SGOS#(config) `show content-filter status`).

Setting Proventia Web Filter Automatic Download Times through the Management Console:

1. Select Configuration>Content Filtering>Proventia Web Filter>Automatic Download.

The Automatic Download tab displays.

Figure 18-17: Proventia Web Filter Automatic Download Tab

2. To set up a schedule for Proventia Web Filter database downloads, select the Download new filter at checkbox and select the time of day from the drop-down list. The default is Midnight.
3. All days are selected by default. Deselect days as desired.
4. Click Apply when finished.

Setting Proventia Web Filter Automatic Download Times through the CLI:

1. At the `(config)` command prompt, enter the following commands to enable or disable automatic downloading of the Proventia Web Filter database.

```
SGOS#(config) content-filter
SGOS#(config content-filter) proventia
SGOS#(config proventia) download auto
-or-
SGOS#(config proventia) no download auto
```

2. At the `(config proventia)` command prompt, enter the following commands to select or deselect the day(s) to automatically download the local database.

```
SGOS#(config proventia) download day-of-week {all | none | sun | mon | tue |
wed | thu | fri | sat}
-or-
SGOS#(config proventia) no download day-of-week {sun | mon | tue | wed | thu
| fri | sat}
```

where `all` selects all days of the week, and `none` clears all days of the week from the schedule.

All days are selected by default; to deselect days, enter `none` and enter specific days. You can only select one day each time, but it is appended to the list. You can also use the `no download day-of-week` command to clear specific days from the schedule.

3. Enter the following command to specify the hour (UTC) of the selected days during which the download should be performed.

```
SGOS#(config proventia) download time-of-day 0-23
```

4. (Optional) To download the Proventia Web Filter database now, enter the following command:

```
SGOS#(config proventia) download get-now
```

Downloading the database now does not affect the automatic database download schedule.

Configuring SmartFilter

Use the ProxySG Management Console or the CLI to configure SmartFilter content filtering.

Differences Between SmartFilter Version 3 and Version 4

SmartFilter version 3 and version 4 have different category lists. Even when categories have similar names, they might have changed in meaning or cover different URLs. Before you switch from version 3 to version 4, you should consider how your policy will be affected, and take steps to avoid unwanted transitional behavior. Table 18.1 is a table of version equivalencies between categories. Where a version 3 category has an obvious version 4 equivalent, CPL will recognize the version 3 name as a deprecated synonym for the version 4 name. A CPL compilation warning will be generated for references to these categories. In two cases, there is no direct equivalent for the version 3 category (“Lifestyle” has been dropped as a category by SmartFilter, and “Mature” has been split into several sub-categories). These names will not be recognized, and you will need to update policy accordingly.

For example, the following policy was compiled using version 3:

```
<proxy>
  category=Sex exception(content_filter_denied)
  category=Nudity exception(content_filter_denied)
  category=Mature exception(content_filter_denied)
```

After configuring and downloading a version 4 list, the policy will still compile, but will generate the following warnings:

```
Warning: Obsolete category name: 'Sex' is now 'Pornography'
Warning: Unknown category: 'Mature'
```

While the first warning is just advising you of a name change (the content continues to be blocked), none of the content formerly categorized as `Mature` will be blocked. You can modify policy to include some or all of the equivalent new subcategories:

```
<proxy>
  category=Pornography exception(content_filter_denied)
  category=Nudity exception(content_filter_denied)
  category=("Sexual Material", Alcohol, Tobacco, Weapons, Profanity, \
  "Provocative Attire", Tasteless/Gross) exception(content_filter_denied)
```

The version 4 category names shown above are subject to change, because they are defined by the control list and can be updated by SmartFilter. See "Selecting Category Providers" on page 636 for instructions on how to view a list of all current category names.

Table 18.1: Category Names in Version 3 and Version 4

Version 3 Category Names	Version 4 Category Names
Art/Culture	Art/Culture/Heritage
Anonymizer/Translator	Anonymizers, Anonymizing Utilities
Chat	Chat, Instant Messaging, Forum/Bulletin Boards
Criminal Skills	Criminal Skills, Malicious Sites
Cults/Occult	Religion and Ideology
Dating	Dating/Personals
Drugs	Drugs
Entertainment	Entertainment/Recreation/Hobbies
Extreme	Extreme, Violence, Profanity, Tasteless/Gross
Gambling	Gambling
Games	Games
Gen. News	General News
Hate Speech	Hate Speech
High Bandwidth	Media Downloads, Internet Radio/TV, Streaming Media, Shareware/Freeware
Humor	Humor
Investing	Stock Trading
Job Search	Job Search
<i>Lifestyle</i>	<i>(deprecated)</i>

Table 18.1: Category Names in Version 3 and Version 4 (Continued)

<i>Mature</i>	<i>Alcohol, Provocative Attire, Profanity, Sexual Materials, Tobacco, Weapons</i>
Nudity	Nudity
Online Sales	Shopping/Merchandizing, Auction
Politics/Religion	Politics/Opinion, Religion and Ideology
Personal	Personal Pages
Portal Sites	Portal Sites
Self-Help/Health	Health
Sex	Pornography
Sports	Sports
Travel	Travel
Usenet News	Usenet News
Webmail	Web Mail

Note: The categories in italics are NOT translated from version 3 to version 4. When a version 3 category that IS translated has more than one version 4 category (for instance, “Politics/Religion”), the version 3 category is automatically translated to the first version 4 category that appears in the table (the version 3 example, “Politics/Religion,” would translate to the version 4 “Politics/Opinion”).

To Configure SmartFilter Content Filtering through the Management Console:

1. Select Configuration>Content Filtering>SmartFilter.
2. The SmartFilter tab displays. Which screen you see depends on which Control List Version you select.

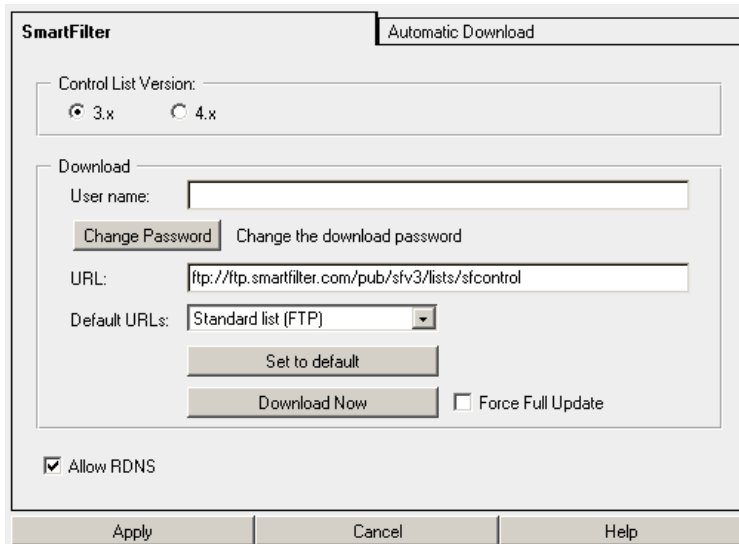


Figure 18-18: SmartFilter Configuration Tab, Version 3.x

3. To change the version of the SmartFilter control list, select a version in the Control List Version field (3.x or 4.x) and click OK in the SmartFilter version changed dialog that appears. If you configure 4.x, skip to step 5.
4. If you configure 3.x:
 - a. Enter the username and password assigned to you for downloading the SmartFilter database: enter your username into the Username field and click the Change Password button to enter or change your password.
 - b. To set the download URL, complete one of the following two steps:
 - The default database download URL is displayed in the URL field. If you have been instructed to use a different URL, enter it in the URL field.
 - You can restore or change the default database download URL at any time. For version 4.x, click Set to default. For version 3.x, select the protocol (FTP or HTTP) and the type of control list (Standard or Premier) that you want to use to download the SmartFilter database by choosing one of the four options in the Default URLs drop-down list; click the Set to default button.

Note: If you are using SmartFilter 3.x and have a low-RAM platform, such as the 400-0 model platform, use the Standard list rather than the Premier list to avoid performance degradation.

5. If you configure 4.x:
The 4.x screen displays.

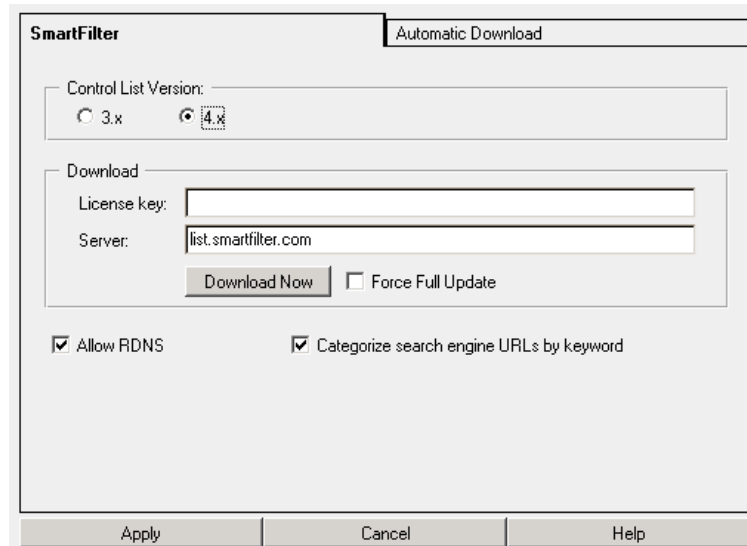


Figure 18-19: SmartFilter Configuration Tab, Version 4.x

To complete the change, set the download URL as appropriate for the new version (step b, below) and download the new version of the SmartFilter database (step 9, below). See "Differences Between SmartFilter Version 3 and Version 4" on page 667 for instructions about category differences that could affect your content-filtering policy.

- c. In the license key field, enter the customer serial number assigned you by SmartFilter.
 - d. The default server is displayed. If you have been instructed to use a different server, enter the hostname or IP address here.
6. (Optional) SmartFilter lookups can require use of reverse DNS to properly categorize a Web site. To disable the use of reverse DNS by Smartfilter, uncheck the Allow RDNS checkbox.

Important: Disabling reverse DNS prevents SmartFilter from correctly classifying some sites and can increase the likelihood of the ProxySG serving inappropriate content.

7. (Optional) By default, SmartFilter categorizes search engines based on keywords in the URL query. To disable this setting, remove the check from the Categorize search engine URLs based on keywords checkbox.

Note: Leaving keywords enabled can cause unexpected results. For example, the keyword *electoral college* falls into the educational category.

8. Click Apply.
9. (Optional) To download the SmartFilter database immediately, complete the following steps.

Note: You can return here at any time to download a database on demand (remember that the automatic download feature, if configured, keeps you up-to-date—see "Scheduling Automatic Downloads for SmartFilter" on page 675).

Ordinarily, the ProxySG will check to see if the database has changed before initiating

a download. If the database is up to date, then no download is necessary and none is performed. If an incremental update is available on the server, then it is downloaded (an incremental update contains only the changes between the current installed version and the latest published version of the database, and is much smaller than a full copy of the database). You can override this process and force a download of the full database by selecting the Force Full Update checkbox; this option is not needed under normal circumstances.

- a. Click Download Now (if you want to download a full database, first select the Force Full Update checkbox—this option is unnecessary under normal circumstances).

The SmartFilter Installation status dialog box displays with the message SmartFilter download in progress.

When the operation is complete, the dialog changes to indicate installation status.

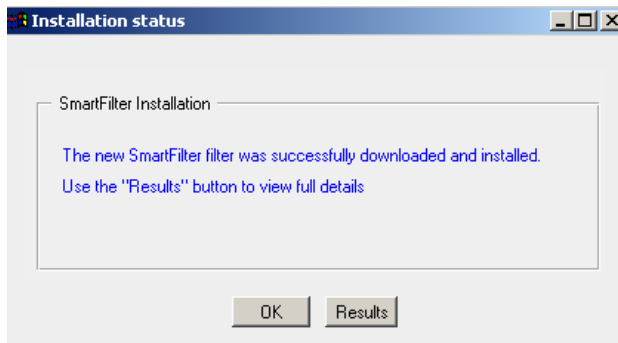


Figure 18-20: SmartFilter Database Successfully Downloaded

- b. Click Results to see the completion message:

Download log:

```
SmartFilter download at: Tue, 06 Apr 2004 20:27:14 UTC
Checking incremental update
Warning: Unable to open input control list
Warning: Unable to open installed control list
Downloading full control file
SmartFilter download at: Tue, 06 Apr 2004 20:27:14 UTC
Downloading from http://example.com/...version=4.0
Download size:      45854194
Database version:   95
Database date:     Tue, 06 Apr 2004 07:05:01 UTC
Database expires:  Tue, 11 May 2004 07:05:01 UTC
```

Note: The first time you download a SmartFilter database, warnings appear in the results message under Checking incremental update. These are expected, and represent the normal process of checking to see if an incremental update is possible. The next time you download a SmartFilter database, the ProxySG will check the previously downloaded database, and download only what is necessary to keep the database current.

You can expect similar warnings the first time you attempt a download after changing SmartFilter versions from 3.x to 4.x (or vice versa). A full download will be necessary the first time the new version is downloaded; subsequent downloads should only need an incremental update.

To Configure SmartFilter Content Filtering through the CLI:

The following commands allow you to select a SmartFilter version, enter a username, specify a URL from which the database is to be downloaded, and do an immediate download of the SmartFilter database.

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) content-filter
SGOS#(config content-filter) smartfilter
SGOS#(config smartfilter) list-version {3 | 4}
```

2. If you chose list-version 3, enter the following commands:

```
SGOS#(config smartfilter) download username username
SGOS#(config smartfilter) download password password
-or-
SGOS#(config smartfilter) download encrypted-password encrypted_password
SGOS#(config smartfilter) download url url
-or-
SGOS#(config smartfilter) download url {standard-list | premier-list {http | ftp}}
SGOS#(config smartfilter) allow-rdns | no allow-rdns
```

3. If you chose list-version 4, enter the following commands:

```
SGOS#(config smartfilter) download license license_key
SGOS#(config smartfilter) download server ip_address_or_hostname
SGOS#(config smartfilter) allow-rdns | no allow-rdns
SGOS#(config smartfilter) use-search-keywords
```

where:

list-version	3 4	Specifies the version (3.x or 4.x) of the SmartFilter control list. See "Differences Between SmartFilter Version 3 and Version 4" on page 667 for instructions about category differences that could affect your content-filtering policy.
download username	<i>username</i>	Version 3.x only. Specifies the username assigned to you for database download.

download encrypted-password	<i>encrypted_password</i>	Version 3.x only. Allows you to take a password previously encrypted by the ProxySG and cut and paste the encrypted password on the same appliance (or another appliance if it shares the same password-display keyring). The primary use of the encrypted-password command is to allow the ProxySG to reload a password that it encrypted.
download license	<i>license_key</i>	Version 4.x only. The customer serial number assigned you by SmartFilter.
download password	<i>password</i>	Version 3.x only. Specifies the password assigned to you for database download.
download url	<i>url</i>	Version 3.x only. Specifies the URL of the downloaded database. The URL is the SmartFilter URL. You can change it if directed to do so.
	{standard-list {http ftp} premier-list {http ftp}}	Version 3.x only. Select the type of control list (Standard or Premier) and the protocol (FTP or HTTP) that you want to use to download the SmartFilter database. If you have a low-RAM platform, such as the 400-0, select <i>standard-list</i> to avoid performance degradation.
download server	<i>ip_address_or_hostname</i>	Version 4.x only. Enter the IP address or hostname of the server you should use for downloads.
allow-rdns	no	A toggle that enables or disables reverse DNS lookup.
use-search-keywords	no	Version 4.x only. Allows you to categorize search engines based on keywords in the URL query.

4. (Optional) To download the database now, enter one of the following commands:

```
SGOS#(config smartfilter) download get-now
-or-
SGOS#(config smartfilter) download full-get-now
```

where:

download get-now	Initiates an immediate database download. An incremental update will be requested.
download full-get-now	Initiates an immediate full-size database download. This option is unnecessary under most circumstances.

5. (Optional) View the configuration.

```

SGOS#(config smartfilter) view
Status: Ready
Use control list version: 3
Download URL: ftp://ftp.smartfilter.com/pub/sfv3/lists/...
Download Username: cf00100002
Automatic download: Enabled
Download time of day (UTC): 0
Download on: sun, mon, tue, wed, thu, fri, sat
Allow RDNS for lookups: No
Download log:
  SmartFilter download at: Fri, 02 Jul 2004 00:13:08 UTC
  Checking incremental update
    Installed database version: 152
    Current published version: 153
  Incremental download complete
Download size: 42821832
Database version: 153
Database date: Thu, 01 Jul 2004 07:05:00 UTC
Database expires: Thu, 05 Aug 2004 07:05:00 UTC

```

Scheduling Automatic Downloads for SmartFilter

Note: By default, the automatic download setting is enabled (for every day at midnight, UTC) and does not need to be configured unless you want to change the schedule or disable auto-download.

To download the SmartFilter database without creating a schedule, see "Configuring SmartFilter" on page 667.

The Automatic Download tab allows you to set the times at which the SmartFilter database is downloaded. You can specify an automatic download on the day and time you prefer. Because sites become stale quickly, Blue Coat recommends downloading on an automatic schedule frequently.

When the database is downloaded, a log is available that includes the information about how the database was updated, but in a more detailed form. You can view the download log through the Management Console (Statistics>Advanced>Content Filter Service) or the CLI (`enable>show content-filter status`).

Setting SmartFilter Automatic Download Times through the Management Console:

1. Select Configuration>Content Filtering>SmartFilter>Automatic Download.

The Automatic Download tab displays.

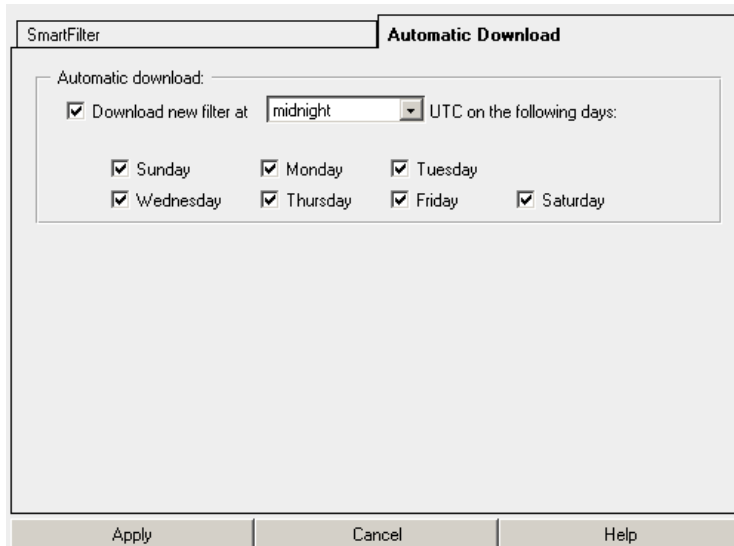


Figure 18-21: SmartFilter Automatic Download Tab

2. To set up a schedule for SmartFilter database downloads, select the Download new filter at checkbox and select the time of day from the drop-down list. The default is Midnight.
3. All days are selected by default. Deselect days as desired.
4. Click Apply when finished.

Setting SmartFilter Automatic Download Times through the CLI:

1. At the `(config)` command prompt, enter the following commands to enable or disable automatic downloading of the SmartFilter database.

```
SGOS#(config) content-filter
SGOS#(config content-filter) smartfilter
SGOS#(config smartfilter) download auto
-or-
SGOS#(config smartfilter) no download auto
```

2. At the `(config smartfilter)` command prompt, enter the following command to select the day(s) to automatically download the local database.

```
SGOS#(config smartfilter) download day-of-week {all | none | sun | mon | tue | wed | thu | fri | sat}
-or-
SGOS#(config smartfilter) no download day-of-week {sun | mon | tue | wed | thu | fri | sat}
```

where `all` selects all days of the week, and `none` clears all days of the week from the schedule.

All days are selected by default; to deselect days, enter `none` and enter specific days. You can only select one day each time, but it is appended to the list. You can also use the `no download day-of-week` command to clear specific days from the schedule.

3. Enter the following command to specify the hour (UTC) of the selected days during which the download should be performed.


```
SGOS#(config smartfilter) download time-of-day 0-23
```

4. (Optional) To download the SmartFilter database now, enter the following command:

```
SGOS#(config smartfilter) download get-now
```

Note: Downloading the database now does not affect the automatic database download schedule.

Configuring SurfControl

Use the ProxySG Management Console or the CLI to configure SurfControl content filtering.

To Configure SurfControl Content Filtering through the Management Console:

1. Select Configuration>Content Filtering>SurfControl.

The SurfControl tab displays.

Figure 18-22: SurfControl Configuration Tab

2. Enter the license key assigned to you for downloading the SurfControl database.
3. The default database download location is displayed in the URL field. If you have been instructed to use a different URL, enter it here. You can restore the default at any time by clicking Set to default.
4. (Optional) To download the SurfControl database immediately, complete the following steps.

Note: You can return here at any time to download a database on demand (remember that the automatic download feature, if configured, keeps you up-to-date—see "Scheduling Automatic Downloads for SurfControl" on page 680).

Ordinarily, the ProxySG will check to see if the database has changed before initiating a download. If the database is up to date, then no download is necessary and none is performed. You can override this check and force a download by selecting the Force Full Update checkbox; this option is not needed under normal circumstances.

- a. Click Download Now (if you want to download a full database, first select the Force Full Update checkbox—this option is unnecessary under normal circumstances).

The SurfControl Installation status dialog box displays with the message SurfControl download in progress.

When the operation is complete, the dialog changes to indicate installation status.

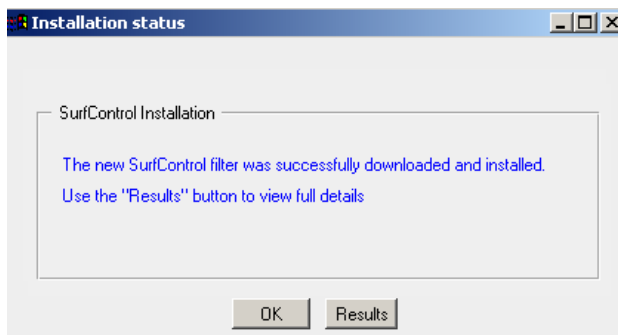


Figure 18-23: SurfControl Database Successfully Downloaded

- b. Click Results to see the SurfControl download log:

Download log:

```
SurfControl download at: Thu, 08 Jul 2004 22:43:42 UTC
Downloading from http://listsrv.surfcontrol.com/bluecoat/v4/scdb.md5
Download size:      99279946
Database version:   96
Database date:      Thu, 24 Jun 2004 18:27:25 UTC
Database expires:   Mon, 23 Aug 2004 18:27:25 UTC
```

5. Click Apply.

To Configure SurfControl Content Filtering through the CLI:

The following commands allow you to enter the username and define the default URL and the default URL location.

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) content-filter
SGOS#(config content-filter) surfcontrol
SGOS#(config surfcontrol) download license license_key
SGOS#(config surfcontrol) download url {default | url}
```

where:

download license	<i>license_key</i>	Specifies the license key assigned to you for database download.
download url	default	Specifies the use of the default download URL.
	<i>url</i>	The URL is the SurfControl URL. You can change it if directed to do so.

2. (Optional) To download the database now, enter one of the following commands:

```
SGOS#(config surfcontrol) download get-now
-or-
SGOS#(config surfcontrol) download full-get-now
```

where:

download get-now	Initiates an immediate database download. If the database is already up-to-date, no download will be initiated.
download full-get-now	Initiates immediate database download, forcing a download whether or not an update is necessary. This option is unnecessary under most circumstances.

3. (Optional) View the configuration.

```
SGOS#(config surfcontrol) view
Status: Ready
Download License key:
Download URL:
http://listsrv.surfcontrol.com/bluecoat/v4/scdb.md5
Automatic download: Enabled
Download time of day (UTC): 0
Download on: sun, mon, tue, wed, thu, fri, sat
Download log:
SurfControl download at: Thu, 08 Jul 2004 22:43:42 UTC
Downloading from http://listsrv.surfcontrol.com/bluecoat/v4/scdb.md5
Download size: 99279946
Database version: 96
Database date: Thu, 24 Jun 2004 18:27:25 UTC
Database expires: Mon, 23 Aug 2004 18:27:25 UTC
```

Using SurfControl Reporter with SGOS 4.x

You can use the SurfControl Reporter with SGOS 4.x access logging to periodically upload information to the SurfControl Reporter reports database.

After you create a SurfControl access logging client, Reporter periodically uploads flat files from the ProxySG. The files are then edited and deleted before loaded into the reports database.

Working with SurfControl Reporter and SGOS 4.x requires several configuration steps. You must:

- Create and configure an access log with SurfControl as the client. For information on configuring a SurfControl access logging client, see "Editing the Custom SurfControl Client" on page 774.

- Download a SurfControl database and configure SurfControl as the content-filtering vendor. For information on downloading a SurfControl database and configuring SurfControl, see "Configuring SurfControl" on page 677.
- Configure the SurfControl server by installing Reporter and configuring the SurfControl Schedule. (Note that the schedule should not be the same as the ProxySG appliance's upload time.) For information on configuring the SurfControl server, refer to the SurfControl server documentation.

Scheduling Automatic Downloads for SurfControl

Note: By default, the automatic download setting is enabled (for every day at midnight, UTC) and does not need to be configured unless you want to change the schedule or disable auto-download.

To download the SurfControl database without creating a schedule, see "Configuring SurfControl" on page 677.

The Automatic Download tab allows you to set the times at which the SurfControl database is downloaded. You can specify an automatic download on the day and time you prefer. Because sites become stale quickly, Blue Coat recommends downloading on an automatic schedule frequently.

When the database is downloaded, a log is available that includes the information about how the database was updated, but in a more detailed form. You can view the download log through the Management Console (Statistics>Advanced>Content Filter Service) or the CLI (`enable>show content-filter status`).

Setting SurfControl Automatic Download Times through the Management Console:

1. Select Configuration>Content Filtering>SurfControl>Automatic Download.

The Automatic Download tab displays.

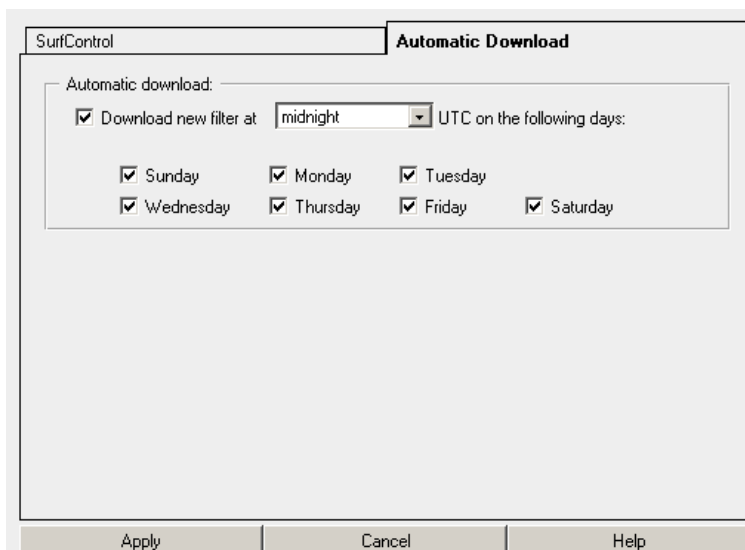


Figure 18-24: SurfControl Automatic Download Tab

2. To set up a schedule for SurfControl database downloads, select the Download new filter at checkbox and select the time of day from the drop-down list. The default is Midnight.
3. All days are selected by default. Deselect days as desired.
4. Click Apply when finished.

Setting SurfControl Automatic Download Times through the CLI:

1. At the `(config)` command prompt, enter the following commands to enable or disable automatic downloading of the SurfControl database.

```
SGOS#(config) content-filter
SGOS#(config content-filter) surfcontrol
SGOS#(config surfcontrol) download auto
-or-
SGOS#(config surfcontrol) no download auto
```

2. At the `(config surfcontrol)` command prompt, enter the following command to select the day(s) to automatically download the local database.

```
SGOS#(config surfcontrol) download day-of-week {all | none | sun | mon | tue
| wed | thu | fri | sat}
-or-
SGOS#(config surfcontrol) no download day-of-week {sun | mon | tue | wed | thu
| fri | sat}
```

where `all` selects all days of the week, and `none` clears all days of the week from the schedule.

All days are selected by default; to deselect days, enter `none` and enter specific days. You can only select one day each time, but it is appended to the list. You can also use the `no download day-of-week` command to clear specific days from the schedule.

3. Enter the following command to specify the hour (UTC) of the selected days during which the download should be performed.

```
SGOS#(config surfcontrol) download time-of-day 0-23
```

4. (Optional) To download the SurfControl database now, enter the following command:

```
SGOS#(config surfcontrol) download get-now
```

Downloading the database now does not affect the automatic database download schedule.

Configuring Websense

Use the ProxySG Management Console or the CLI to configure Websense content filtering.

Note: Websense databases contain a category called *User-Defined* to support locally-specified categorizations on other platforms. This category should not be used on the ProxySG. Instead, you can define your own categories through the ProxySG and assign URLs to them using Policy (see page "Defining Custom Categories in Policy" on page 696), or using a local category database (see "Create and Edit Policy Files" on page 442).

To Configure the Database through the Management Console:

1. Select Configuration>Content Filtering>Websense>Websense.

The Websense download tab displays.

Figure 18-25: Websense Configuration Tab

2. Fill in the fields as appropriate:

- License Key—Enter the license key assigned to you for downloading the Websense database.
- Server—Enter the Websense server from which you wish to download. Your licensing information might suggest an alternate value; otherwise, use the default (`download.websense.com`).
- Contact e-mail—(Optional) Enter an e-mail address through which Websense can contact you.
- Always apply regular expressions to urls—(Optional)

Select this option if you want to force an additional regular expression lookup for each URL to be categorized. Normally, regular expression lookups are done only when no category is found in the Websense database. If this option is selected, regular expression lookups will always be done, even for categorized URLs. Selecting this option can cause a significant reduction in lookup performance, but can allow certain sites (such as translation, search engine, and link-cache sites) to be categorized more accurately.

3. To use the Websense Reporter, you must enable the Websense Integration Service.

- a. In the Integrator Service Host field, enter the Integrator Service Host IP (which has the same IP address as the Websense Log Server).
- b. In the Port field, specify the port of the Websense Integration Service. It must be between 0 and 65535 and match the port selected on the Integration Service host.
- c. Click the enabled checkbox to enable the service.

Note: The Policy Server, the Log Server, and Reporter must be installed and enabled on your PC before Reporter can be used. For information on Websense products, refer to: <http://www.websense.com/support/documentation/integrationservice>.

You must also set up access logging on the ProxySG with Websense as the client. For more information on configuring a Websense access logging client, see "Editing the Websense Client" on page 775.

4. (Optional) To download the Websense database immediately, complete the following steps.

Note: You can return here at any time to download a database on demand (remember that the automatic download feature, if configured, keeps you up-to-date—see "Scheduling Automatic Downloads for Websense" on page 686).

Ordinarily, the ProxySG will check to see if the database has changed before initiating a download. If the database is up to date, then no download is necessary and none is performed. If an incremental update is available on the server, then it is downloaded (an incremental update contains only the changes between the current installed version and the latest published version of the database, and is much smaller than a full copy of the database). You can override this process and force a download of the full database by selecting the Force Full Update checkbox; this option is not needed under normal circumstances.

- a. Click Download Now (if you want to download a full database, first select the Force Full Update checkbox—this option is unnecessary under normal circumstances).

The Websense Installation status dialog box displays with the message Websense download in progress.

When the operation is complete, the dialog changes to indicate installation status.




Figure 18-26: Websense Database Successfully Downloaded

b. Click Results to view the Websense download log:

```

Download log:
Websense download at: Fri, 04 Mar 2005 19:07:56 UTC
No database is currently installed
Attempting full download
Downloading from download.websense.com
Processing download file
  Retrieved full update
Download size:      116596744
Database version:  82220
Database date:     2005-03-03
License expires:   Tue, 29 Mar 2005 08:00:00 UTC
License max users: 25
Licenses in use:   1
    
```

5. Click Apply.

To Configure Websense through the CLI:

1. At the (config) command prompt, enter the following commands to configure the Websense download:

```

SGOS#(config) content-filter
SGOS#(config content-filter) websense
SGOS#(config websense) download email-contact email_address
SGOS#(config websense) download server ip_address_or_hostname
SGOS#(config websense) download license license_key
    
```

2. (Optional) Enter the following command to configure regular expression lookups for each URL to be categorized.

```

SGOS#(config websense) always-apply-regexes
-or-
SGOS#(config websense) no always-apply-regexes
    
```

where:

download email-contact	<i>email_address</i>	(Optional) Specifies an email address through which Websense can contact you.
download server	<i>ip_address_or_hostname</i>	Specifies the Websense server from which you wish to download. Your licensing information might suggest an alternate value; otherwise, use the default (download.websense.com).
download license	<i>license_key</i>	Specifies the license key assigned to you for downloading the Websense database.

always-apply-regexes		<p>Forces an additional regular expression lookup for each URL to be categorized.</p> <p>Normally, regular expression lookups are done only when no category is found in the Websense database. If this optional command is selected, regular expression lookups will always be done, even for categorized URLs. Selecting this option can cause a significant reduction in lookup performance, but can allow certain sites (such as translation, search engine, and link-cache sites) to be categorized more accurately.</p> <p>The default setting is <code>no</code> <code>always-apply-regexes</code>; you should change the default only if you are certain that you need the advanced setting.</p>
no always-apply-regexes		<p>Causes regular expression lookups to be done only when no category is found in the Websense database. This is the default setting.</p>

3. (Optional) To download the database now, enter one of the following commands:

```
SGOS#(config websense) download get-now
-or-
SGOS#(config websense) download full-get-now
```

where:

download get-now	Initiates an immediate database download. An incremental update will be requested.
download full-get-now	Initiates an immediate full-size database download. This option is unnecessary under most circumstances.

4. (Optional) View the configuration.

```
SGOS#(config websense) view
Status: Ready
Download License key: EBC123DEF456GHI789
Download Server: download.websense.com
Email contact:
Automatic download: Enabled
Download time of day (UTC): 0
Download on: sun, mon, tue, wed, thu, fri, sat
Use regular expression filters: No
Always apply regex filters: Yes
Integration Server: Disabled
Integration Server host:
Integration Server port: 0
Download log:
Websense download at: Fri, 04 Mar 2005 19:07:56 UTC
No database is currently installed
Attempting full download
Downloading from download.websense.com
Processing download file
```

```
Retrieved full update
Download size:      116596744
Database version:   82220
Database date:      2005-03-03
License expires:    Tue, 29 Mar 2005 08:00:00 UTC
License max users:  25
```

To Configure the Websense Integration Service through the CLI:

Enter the following commands to enable (or disable) and configure the Websense Integration Service through the CLI:

```
SGOS#(config) content-filter
SGOS#(config content-filter) websense
SGOS#(config websense) integration-service {enable | disable}
SGOS#(config websense) integration-service host ip_address_or_hostname
SGOS#(config websense) integration-service port integer
```

where:

- *host* specifies the hostname or IP address of the Websense Integration Service, which is the name or IP address of the Websense Log Server.
- *port* specifies the port of the Websense Integration Service, must be between 0 and 65535, and match the port selected on the Integration Service host.

Note: The Policy Server, the Log Server, and Reporter must be installed and enabled on your PC before Reporter can be used. For information on Websense products, refer to: <http://www.websense.com/support/documentation/integrationservice>

You must also set up access logging on the ProxySG with Websense as the client. For more information on configuring a Websense access logging client, see "Editing the Websense Client" on page 775.

Scheduling Automatic Downloads for Websense

Note: By default, the automatic download setting is enabled (for every day at midnight, UTC) and does not need to be configured unless you want to change the schedule or disable auto-download.

To download the Websense database without creating a schedule, see "Configuring Websense" on page 681.

The Automatic Download tab allows you to set the times the Websense database will download. You can specify an automatic download on the day and time you prefer. Because sites become stale quickly, Blue Coat recommends downloading on an automatic schedule frequently.

When you download the database, a log is available that includes the information about how the database was updated, but in a more detailed form. You can view the download log through the Management Console (Statistics>Advanced>Content Filter Service) or the CLI (`enable>show content-filter status`).

Setting Websense Automatic Download Times through the Management Console:

1. Select Configuration>Content Filtering>Websense>Automatic Download.

The Automatic Download tab displays.

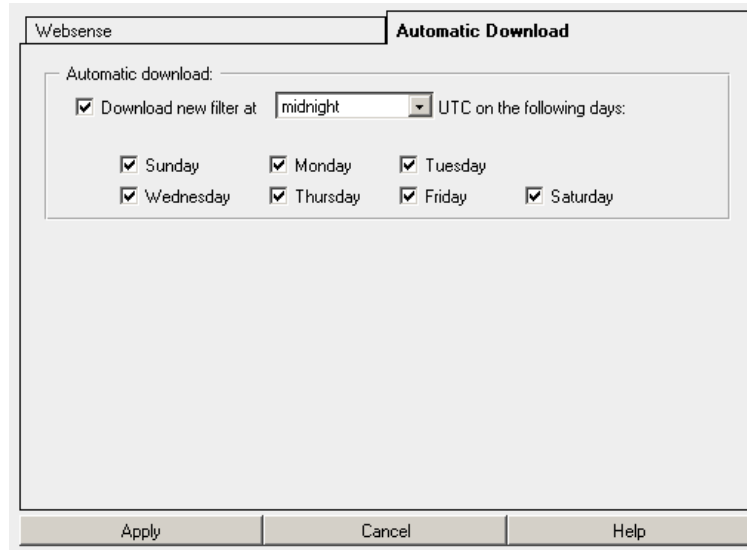


Figure 18-27: Websense Automatic Download Tab

2. To set up a schedule for Websense database downloads, select the Download new filter at checkbox and select the time of day from the drop-down list. The default is Midnight.
3. All days are selected by default. Deselect days as desired.
4. Click Apply when finished.

Setting Websense Automatic Download Times through the CLI:

1. At the (config) command prompt, enter the following commands to enable or disable automatic downloading of the Websense database.

```
SGOS#(config) content-filter
SGOS#(config content-filter) websense
SGOS#(config websense) download auto
-or-
SGOS#(config websense) no download auto
```

2. At the (config websense) command prompt, enter the following command to select the day(s) to automatically download the local database.

```
SGOS#(config websense) download day-of-week {all | none | sun | mon | tue | wed | thu | fri | sat}
-or-
SGOS#(config websense) no download day-of-week {sun | mon | tue | wed | thu | fri | sat}
```

where `all` selects all days of the week, and `none` clears all days of the week from the schedule. .

All days are selected by default; to deselect days, enter `none` and enter specific days. You can only select one day each time, but it is appended to the list. You can also use the `no download day-of-week` command to clear specific days from the schedule.

3. Enter the following command to specify the hour (UTC) of the selected days during which the download should be done.

```
SGOS#(config websense) download time-of-day 0-23
```

4. (Optional) To download the Websense database now, enter the following command:

```
SGOS#(config websense) download get-now
```

Downloading the database now does not affect the automatic database download schedule.

Configuring Webwasher URL Filter

Use the ProxySG Management Console or the CLI to configure Webwasher URL Filter content filtering.

To Configure Webwasher URL Filter Content Filtering through the Management Console:

1. Select Configuration>Content Filtering>Webwasher.

The Webwasher URL Filter tab displays.

Figure 18-28: Webwasher URL Filter Configuration Tab

2. Enter the Webwasher URL Filter username and password: enter your username into the Username field and click the Change Password button to enter or change your password.
3. The default database download location is displayed in the URL field. If you have been instructed to use a different URL, enter it here. You can restore the default at any time by clicking Set to default.

4. (Optional) To download the Webwasher URL Filter database immediately, complete the following steps.

Note: You can return here at any time to download a database on demand (remember that the automatic download feature, if configured, keeps you up-to-date—see "Scheduling Automatic Downloads for Webwasher URL Filter" on page 691).

Ordinarily, the ProxySG will check to see if the database has changed before initiating a download. If the database is up to date, then no download is necessary and none is performed. If an incremental update is available on the server, then it is downloaded (an incremental update contains only the changes between the current installed version and the latest published version of the database, and is much smaller than a full copy of the database). You can override this process and force a download of the full database by selecting the Force Full Update checkbox; this option is not needed under normal circumstances.

- a. Click Download Now (if you want to download a full database, first select the Force Full Update checkbox—this option is unnecessary under normal circumstances).

The Webwasher Installation status dialog box displays with the message Webwasher download in progress.

When the operation is complete, the dialog changes to indicate installation status.

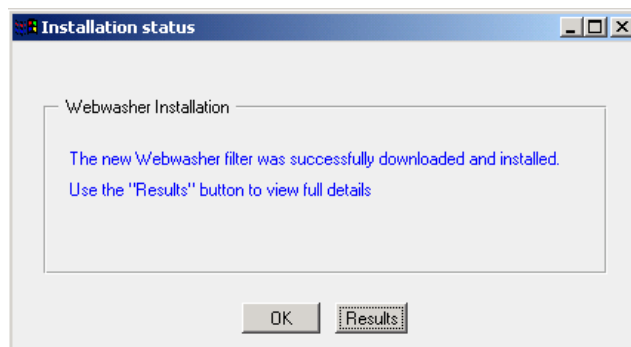


Figure 18-29: Webwasher URL Filter Database Successfully Downloaded

- b. Click Results to see the Webwasher URL Filter download log:

Download log:

```
Webwasherdownload at: Sat, 10 Jul 2004 18:54:43 UTC
Downloading from
http://list.bluecoat.com/webwasher/activity/download/webwasher.db
Requesting full update
Download size:          93484280
Database date:         Tue, 14 Dec 2004 22:38:14 UTC
Database expires:     Mon, 11 Jan 2016 06:31:29 UTC
Database version:     900
```

To Configure Webwasher URL Filter Content Filtering through the CLI:

The following commands allow you to enter the Webwasher URL Filter username and password and define the default URL and the default URL location.

1. At the `(config)` command prompt, enter the following commands:

```

SGOS#(config) content-filter
SGOS#(config content-filter) webwasher
SGOS#(config webwasher) download username username
SGOS#(config webwasher) download password password
-or-
SGOS#(config webwasher) download encrypted-password encrypted-password
SGOS#(config webwasher) download url {default | url}
    
```

where:

download username	<i>username</i>	Specifies the username assigned to you for database download.
download encrypted-password	<i>encrypted_password</i>	Allows you to take a password previously encrypted by the ProxySG and cut and paste the encrypted password on the same appliance (or another appliance if it shares the same password-display keyring). The primary use of the encrypted-password command is to allow the ProxySG to reload a password that it encrypted.
download password	<i>password</i>	Specifies the password assigned to you for database download.
download url	<i>default</i>	Specifies the use of the default download URL.
	<i>url</i>	The URL is the Webwasher URL Filter URL. You can change it if directed to do so.

2. (Optional) To download the database now, enter one of the following commands:

```

SGOS#(config webwasher) download get-now
-or-
SGOS#(config webwasher) download full-get-now
    
```

where:

download get-now	Initiates an immediate database download. An incremental update will be requested.
download full-get-now	Initiates an immediate full-size database download. This option is unnecessary under most circumstances.

3. (Optional) View the configuration.

```

SGOS#(config webwasher) view
Status: Ready
Download URL:
https://list.bluecoat.com/webwasher/activity/download/webwasher.db
Download Username:
Automatic download: Enabled
Download time of day (UTC): 0
Download on: sun, mon, tue, wed, thu, fri, sat
Download log:
Webwasher download at: Thu, 16 Dec 2004 20:52:58 UTC
    
```

```

Downloading from
https://list.bluecoat.com/webwasher/activity/download/webwasher.db
Requesting full update
Download size:      93484280
Database date:     Tue, 14 Dec 2004 22:38:14 UTC
Database expires:  Mon, 11 Jan 2016 06:31:29 UTC
Database version:  900

```

Scheduling Automatic Downloads for Webwasher URL Filter

Note: By default, the automatic download setting is enabled (for every day at midnight, UTC) and does not need to be configured unless you want to change the schedule or disable auto-download.

To download the Webwasher URL Filter database without creating a schedule, see "Configuring Webwasher URL Filter" on page 688.

The Automatic Download tab allows you to set the times at which the Webwasher URL Filter database is downloaded. You can specify an automatic download on the day and time you prefer. Because sites become stale quickly, Blue Coat recommends downloading on an automatic schedule frequently.

When the database is downloaded, a log is available that includes the information about how the database was updated, but in a more detailed form. You can view the download log through the Management Console (Statistics>Advanced>Content Filter Service) or the CLI (SGOS#(config) `show content-filter status`).

Setting Webwasher URL Filter Automatic Download Times through the Management Console:

1. Select Configuration>Content Filtering>Webwasher>Automatic Download.

The Automatic Download tab displays.

Figure 18-30: Webwasher URL Filter Automatic Download Tab

2. To set up a schedule for Webwasher URL Filter database downloads, select the Download new filter at checkbox and select the time of day from the drop-down list. The default is Midnight.
3. All days are selected by default. Deselect days as desired.
4. Click Apply when finished.

Setting Webwasher URL Filter Automatic Download Times through the CLI:

1. At the `(config)` command prompt, enter the following commands to enable or disable automatic downloading of the Webwasher URL Filter database.

```
SGOS#(config) content-filter
SGOS#(config content-filter) webwasher
SGOS#(config webwasher) download auto
-or-
SGOS#(config webwasher) no download auto
```

2. At the `(config webwasher)` command prompt, enter the following commands to select or deselect the day(s) to automatically download the local database.

```
SGOS#(config webwasher) download day-of-week {all | none | sun | mon | tue |
wed | thu | fri | sat}
-or-
SGOS#(config webwasher) no download day-of-week {sun | mon | tue | wed | thu
| fri | sat}
```

where `all` selects all days of the week, and `none` clears all days of the week from the schedule.

All days are selected by default; to deselect days, enter `none` and enter specific days. You can only select one day each time, but it is appended to the list. You can also use the `no download day-of-week` command to clear specific days from the schedule.

3. Enter the following command to specify the hour (UTC) of the selected days during which the download should be performed.

```
SGOS#(config webwasher) download time-of-day 0-23
```

4. (Optional) To download the Webwasher URL Filter database now, enter the following command:

```
SGOS#(config webwasher) download get-now
```

Downloading the database now does not affect the automatic database download schedule.

How to Apply Policy to Categorized URLs

You apply policy to categories in the same way you apply policy to individual URLs: using Content Policy Language (CPL). To define policies on the ProxySG, you can either use the Visual Policy Manager (VPM) or you can manually edit policy files. For information about the VPM, see Chapter 14: “The Visual Policy Manager” on page 453; for information about managing policy files, see Chapter 13: “Managing Policy Files” on page 439.

Note: 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. A local database does, however, have some restrictions that policy does not: no more than 200 separate categories are allowed, category names must be 32 characters or less, and a given URL pattern can appear in no more than four category definitions. You can choose to use any combination of the local database, policy files, and VPM to manage your category definitions. See "Configuring a Local Database" on page 640 for more information.

The CPL trigger `category=` is used to test the category or categories assigned to the request URL, and thus make a policy decision. For example, to block all requests for URLs that are categorized as Sports:

```
DENY category=Sports
```

The following example demonstrates a condition that is true when a request contains the Websense content categories Sexuality and Drugs:

```
<proxy>
  category=(sexuality, drugs)
```

You can block multiple categories with a single rule:

```
category=(Sports, Gambling, Shopping) exception(content_filter_denied)
```

In this example, three categories are blocked and instead the predefined exception page `content_filter_denied` is served; by default this indicates that the request was denied due to its content and specifies the categories found.

The following example shows a condition that includes an extensive number of categories:

```
category=(Abortion, Activist, Adult, Gambling, Illegal, Hacking, Militancy,
  Racism, Shopping, Tasteless, Violence, Weapons)
```

Note that URLs which 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)
```

Such category tests can also be combined with other types of triggers to produce more complex policy, such as:

- Restrict access by category and time: block sports from 6 am to 6 pm:

```
category=Sports time=0600..1800 DENY
```

- Restrict by category and user identity: only members of the group Sales are permitted to visit Shopping sites:

```
category=Shopping group=!Sales DENY
```

- Require special authentication for access to certain categories:

```
category=Hacking authenticate(restricted_realm)
```

where *restricted_realm* is an authentication realm you have configured.

- Log certain types of access:

```
category=Adult action.Log_adult_site_access(yes)
```

where *Log_adult_site_access* is a policy action defined elsewhere that records extra information about this request in the event log.

In general, `category=` can be used in policy anywhere that a basic URL test can be used. Refer to the *Blue Coat ProxySG Content Policy Language Guide* for more details.

Depending on which provider you have selected and whether you have defined any of your own categories in policy (see "Defining Custom Categories in Policy" on page 696), you have a number of possible category names that can be used with `category=`. To review the valid category names, use the `categories` CLI command or click View Categories in the Management Console (as described in "Selecting Category Providers" on page 636).

The `category=` expressions are normally put in <Proxy> Layers (Web Access Layers in the VPM), because the goal of content-filtering policy is usually to control requests from users. They can also be used in <Cache> (Web Content in the VPM) Layers. Either way, policy is enforced on all user requests.

It is possible for an attempt to categorize a URL to fail—for example, if no database is loaded, your license is expired, or if a system error occurs. In such a case, the category is considered *unavailable* and triggers such as:

```
category=Sports
```

are false, even if the URL is actually a Sports site, 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. You can control how the ProxySG treats these situations with the condition:

```
category=unavailable
```

which is true in these cases. In continuing with the example, to make sure that Sports is always blocked, even when errors occur (this is a mode of operation called *fail-closed*), use a rule such as:

```
category=(sports, unavailable) exception(name_of_exception_page)
```

This rule is true if the category is sports or if the category could not be determined, and in either case the proper exception page is served instead of the restricted content.

The category *unlicensed* is assigned in addition to *unavailable* when the failure to categorize occurred because of license expiry. That 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. You can use

```
category=unlicensed
```

to detect this situation as a distinct case from other causes of unavailability.

You can also use this feature with custom exception pages (see Chapter 15: "Advanced Policy" on page 561):

```
<proxy>
  category=sports time=0800..1800 exception(sports_during_bus_hrs)
  category=unlicensed exception(contact_admin_re_license)
  category=unavailable exception(content_filter_unavailable)
  where sports_during_bus_hrs is a custom exception page you have created to respond
  to requests for Sports pages between 8 am and 6 pm local time.
  contact_admin_re_license is another page that instructs the user to inform the
  administrator about license expiry, and is served if a license check fails. When the
  category is unavailable for some other reason, the pre-defined exception
  (content_filter_unavailable) is served.
```

Note that the most common reason (other than license expiry) why categories are unavailable is that a provider is selected but no database is installed. Barring hardware or network problems that might cause a downloaded database to become corrupted and unreadable, it is unlikely that the database will suddenly become unavailable.

To define policies on the ProxySG, use either the Visual Policy Manager or manually edit Policy files.

Content filtering policies are usually found in <Proxy> and <Cache> layers.

If you are using content filtering to manage a type of content globally, create these rules in the <Cache> layer.

However, if your content filtering policy is dependent on user identity or request characteristics, create these rules in the <Proxy> layer.

Using Content-Filtering Vendors with ProxySG Policies

The ProxySG provides the ability to define flexible Web access and control policies. With content filtering, you can set up policies to provide a customized level of Web-site access control. With vendor-based content filtering, these policies use and can supplement vendor categories. By supplementing content-filtering vendor categories, you can further refine the type of content filtering the ProxySG performs. For example, if *Travel* is a vendor-defined content category, you can define a policy that allows only Human Resources staff to access travel sites. You can define policies that filter by a variety of conditions, including category, protocol (including MMS and RTSP streaming protocols), time of day, and user or user groups.

Example

Policy: 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 Chapter 9: “Using Authentication Services” on page 271.

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
```

Example

Policy: 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

You can use CPL to create your own categories and assign URLs to them. This is done with the `define category` construct (for more complete information on the `define category` construct, refer to *Blue Coat ProxySG Content Policy Language Guide*). To add URLs to a category, simply list them in the definition. You only need to specify a partial URL:

- hosts and subdomains within the domain you specify will automatically be included
- if you specify a path, all paths with that prefix will be included (if you specify no path, the whole site is included)

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 instance, 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) using the Test button on the Management Console or the `test-url` command in the CLI, as described in "Selecting Category Providers" on page 636.

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
```

Remember that you can use 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 *Blue Coat ProxySG 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 640 for information about local databases.

Tips

- When you use an expired database, the category *unlicensed* will be assigned to all URLs and no lookups will be done on the database. This can occur even if your download license with the database vendor is still valid, but you haven't downloaded a database for a long time (databases expire after a certain number of days). You can view the date that your database expires (or expired) in the download log or by using the `view` command in the CLI.

When you download a database through the CLI, you can see the download log as soon as the download is complete. To see the download log when you download a database through the Management Console, click the Results button in the Installation Status dialog when the download is complete.

To see the last download log without doing another download, enter the following CLI (`config`) commands:

```
SGOS#(config) content-filter
SGOS#(config content-filter) view
```

- When your license with the database vendor expires, you can no longer download. This does not have an immediate effect—you can still use the database you have for a period of time. But eventually, the database will expire and you will get the category *unlicensed*, as described above.
- If a requested HTTPS host is categorized in a content filtering database, then filtering will apply. However, if the request contains a path and the categorization relies on the host/relative path, content filtering only filters on the host name because the path is not accessible. This might result in a different categorization than if the host plus path were used.

- If you receive an error message when downloading a content-filtering database, check the error message (in the Management Console, click the Results button on the Installation status dialog; in the CLI, the results message appears in the event of an error). If you see an error message such as “ERROR: HTTP 401 - Unauthorized,” check that you entered your username and password correctly. For example, the following error message was generated by entering an incorrect username and attempting to download a SmartFilter database:

Download log:

```
SmartFilter download at: Thu, 08 Apr 2004 18:03:08 UTC
Checking incremental update
  Checking download parameters
  Fetching:http://example.com/
  Warning: HTTP 401 - Unauthorized
Downloading full control file
SmartFilter download at: Thu, 08 Apr 2004 18:03:17 UTC
Downloading from http://example.com/
  Fetching:http://example.com/
  ERROR: HTTP 401 - Unauthorized
  Download failed
Download failed
Previous download:
...
```


Chapter 19: *Configuring the Upstream Networking Environment*

The ProxySG must interact not only with the local network, but with the upstream network environment to fill requests. To control upstream interaction, various options are supported, such as forwarding, SOCKS gateways, ICP (Internet Caching Protocol), and WCCP (Web Cache Control Protocol).

- Forwarding—Allows you to define the hosts and groups of hosts to which client requests can be redirected. Those hosts can be servers or proxies, including additional ProxySG Appliances. Rules to redirect requests are set up in policy.

Forwarding is available for HTTP, HTTPS, FTP, Windows Media, RTSP, Telnet, and TCP tunnels.

- SOCKS gateways—SOCKS servers provide application level firewall protection for an enterprise. The SOCKS protocol provides a generic way to proxy HTTP.
- ICP—ICP is a service to handle ICP queries from other caching devices looking for cached data. The devices that can access this service can be controlled. ICP can also be used by the ProxySG to locate cached data in other systems.
- WCCP—WCCP is a Cisco[®]-developed protocol that allows you to establish redirection of the traffic that flows through routers. (For more information on WCCP, see Appendix C: "Using WCCP" on page 911.)

This chapter contains the following topics:

- "Forwarding Configuration"
- "SOCKS Gateway Configuration"
- "Internet Caching Protocol (ICP) Configuration"

Forwarding, SOCKS gateways, and ICP all work together to identify the network topology.

Forwarding and SOCKS Gateway each have a single configuration created either through the CLI or through installable lists. ICP configuration must be created through installable lists.

Forwarding Configuration

Forwarding allows you to define the hosts and groups of hosts to which client requests can be redirected. Those hosts can be servers or proxies, including additional ProxySG Appliances. Rules to redirect requests are configured in policy.

Forwarding stores forwarding hosts' configuration information in the registry; policy is used to create rules for forwarding hosts.

Configuration can be created through the CLI or through installable lists that you can create.

To set the default load-balancing and host-affinity values, you must use the high level load-balance and host-affinity commands (see "Configuring Load Balancing" on page 708 or "Configuring Host Affinity" on page 710). However, three methods are available to you for setting per host or per group settings. You can:

- Use the `(config forwarding) create` command (see "Creating Forwarding Hosts or Host Groups" on page 702).
- Use the `(config forwarding host_alias)` or `(config forwarding group_alias)` commands (see "Editing a Forwarding Host" on page 705 or "Editing a Forwarding Host Group" on page 707).
- Use the `(config forwarding) load-balance` or `(config forwarding) host-affinity` commands (see "Configuring Load Balancing" on page 708 or "Configuring Host Affinity" on page 710).

Note: Forwarding is available for HTTP, HTTPS, FTP, Windows Media, RTSP, Telnet, and TCP tunnels.

The host/group aliases cannot be CPL keywords, such as `no`, `default`, or `forward`.

Configuring Forwarding through the CLI

You can use the CLI to configure your forwarding hosts and describe the upstream network. If you prefer to create an installable list instead of using the CLI, see "Using Forwarding Directives to Create an Installable List" on page 712.

Creating Forwarding Hosts or Host Groups

You can create a maximum of 32 groups, and each group can contain a maximum of 512 hosts. You can create 512 individual hosts that do not belong to any group.

To Create a Host or Host Group through the CLI:

The only required entries under the `create` option (for a host) are the `host_alias`, `host_name`, a protocol, and a port number. The port number can be defined explicitly (such as `http=8080`), or it can take on the default port value of the protocol, if one exists (such as enter `http`, and the default port value of 80 is entered automatically).

To create a host group, you must also include the `group=group_name` command. If this is the first mention of the group, `group_name`, then that group is automatically created with this host as its first member. Do not use this command when creating an independent host.

1. At the `(config)` command prompt, create a forwarding host:

```
SGOS#(config) forwarding
SGOS#(config forwarding) create host_alias host_name [default-schemes]
[http[=port | =no]] [https[=port | =no]] [ftp[=port | =no]] [mms[=port |
=no]] [rtsp[=port | =no]] [tcp=port] [telnet[=port | =no]]
[ssl-verify-server[=yes | =no]] [group=group_name] [server | proxy]
[load-balance={no | round-robin | least-connections}] [host-affinity={no |
client-ip-address | accelerator-cookie}] [host-affinity-ssl={no |
client-ip-address | accelerator-cookie | ssl-session-id}]
```

where:

`host_alias`

This is the alias for use in policy. Define a meaningful name.

<code>host_name</code>		The name of the host domain, such as <code>www.bluecoat.com</code> , or its IP address.
<code>default-schemes</code>		If you use <code>default-schemes</code> in the directive, all protocols, along with their default ports are selected. This directive is only available for proxy hosts.
<code>http</code>	<code>=port =no</code>	No protocol is selected by default if the forwarding host is a server. You must choose at least one protocol where <code>port=0</code> to <code>65535</code> . If only one protocol is configured, the ProxySG configures the default port for that protocol. You can use <code>default-schemes</code> and then eliminate protocols by selecting the protocol you do not want; for example, <code>http=no</code> . If you do not want to use the default ports for the protocols, you must also specify them here. HTTPS protocols are not allowed if the host is a proxy.
<code>https</code>		
<code>ftp</code>		
<code>mms</code>		
<code>rtsp</code>		
<code>telnet</code>		
<code>tcp</code>	<code>=port</code>	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.
<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> . To disable this feature, you must specify <code>ssl-verify-server=no</code> in the installable list or CLI.
<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. Do not use the <code>group=</code> option when creating independent hosts.
<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. Proxy is the default.

<code>load-balance</code>	<code>=no =round-robin =least-connections</code>	<p>Specifies either the least-connections or round-robin method of load balancing. Select <code>no</code> to disable load balancing for this forwarding host or host group.</p> <p>If these settings are not specified for a particular host or host group, then the global default settings are used. To configure the settings for a specific host or host group, use the <code>edit host_alias</code> or <code>edit group_alias</code> commands (see "Editing a Forwarding Host" on page 705 or "Editing a Forwarding Host Group" on page 707).</p>
<code>host-affinity</code>	<code>=no =client-ip-address =accelerator-cookie</code>	<p>Specifies non-SSL host affinity via either a client IP address or an accelerator cookie. Select <code>no</code> to disable non-SSL host affinity for this forwarding host or host group.</p> <p>If these settings are not specified for a particular host or host group, then the global default settings are used. To configure the settings for a specific host or host group, use the <code>edit host_alias</code> or <code>edit group_alias</code> commands (see "Editing a Forwarding Host" on page 705 or "Editing a Forwarding Host Group" on page 707).</p>
<code>host-affinity-ssl</code>	<code>=no =client-ip-address =accelerator-cookie =ssl-session-id</code>	<p>Specifies SSL host affinity via a client IP address, an accelerator cookie, or an SSL session ID. Select <code>no</code> to disable SSL host affinity for this forwarding host or host group.</p> <p>If these settings are not specified for a particular host or host group, then the global default settings are used. To configure the settings for a specific host or host group, use the <code>edit host_alias</code> or <code>edit group_alias</code> commands (see "Editing a Forwarding Host" or "Editing a Forwarding Host Group").</p>

2. Repeat step 1 to create additional forwarding hosts or host groups.
3. Complete the configuration by entering the following commands as necessary:

```
SGOS#(config forwarding) download-via-forwarding disable | enable
SGOS#(config forwarding) failure-mode closed | open
SGOS#(config forwarding) integrated-host-timeout minutes
SGOS#(config forwarding) delete {all | group group_name | host host_alias}
SGOS#(config forwarding) path url
SGOS#(config forwarding) no path
```

where:

download-via-forwarding	enable disable	Specifies whether to allow configuration file downloads using forwarding.
failure-mode	closed open	Specifies the default failure mode for forwarding hosts if an operation is unsuccessful.
integrated-host-timeout	<i>minutes</i>	An integrated host is an Origin Content Server (OCS) that has been added to the health check list. The host, added through the <code>integrate_new_hosts</code> property, ages out after being idle for the specified time. The default is 60 minutes.
delete	all group <i>group_name</i> host <i>host_alias</i>	Deletes all forwarding hosts and groups (delete all) or a specific forwarding group (delete group <i>group_name</i>) or host (delete host <i>host_alias</i>).
path	<i>url</i>	(Optional) Specifies the download path to use if you download forwarding settings through directives.
no	path	Clears the network path URL to download forwarding settings.

Editing a Forwarding Host

Once you have created a forwarding host, you can edit its configuration.

To Edit the Settings of a Forwarding Host through the CLI:

1. At the `(config)` command prompt, enter the following commands to configure the settings of a forwarding host:

```
SGOS#(config) forwarding
SGOS#(config forwarding) edit host_alias
SGOS#(config forwarding host_alias) {ftp | http | https | mms | rtsp |
telnet} [port]
SGOS#(config forwarding host_alias) group group_name
SGOS#(config forwarding host_alias) host host_name
SGOS#(config forwarding host_alias) host-affinity method {accelerator-cookie
| client-ip-address | default}
-or-
SGOS#(config forwarding host_alias) host-affinity ssl-method
{accelerator-cookie | client-ip-address | default | ssl-session-id}
SGOS#(config forwarding host_alias) load-balance method {default |
least-connections | round-robin}
SGOS#(config forwarding host_alias) proxy | server
SGOS#(config forwarding host_alias) ssl-verify-server
SGOS#(config forwarding host_alias) tcp port
```

where:

<pre>ftp http [port] https mms rtsp telnet</pre>	<pre>[port]</pre>	<p>Adds the protocol and optional port for this host if it was not set previously or changes the port number for the specified protocol if it was. If you do not enter a port number, the default port number is used.</p> <p>HTTPS protocols are not allowed if the host is a proxy.</p>
<pre>tcp</pre>	<pre>port</pre>	<p>Changes the port number for the TCP protocol for this host. You must enter a port number if you use the TCP protocol.</p> <p>TCP protocols are not allowed if the host is a proxy.</p>
<pre>group</pre>	<pre>group_name</pre>	<p>Changes the group membership for this host.</p>
<pre>host</pre>	<pre>host_name</pre>	<p>Changes this host's name.</p>
<pre>host-affinity</pre>	<pre>method {accelerator-cookie client-ip-address default}</pre> <pre>ssl-method {accelerator-cookie client-ip-address default ssl-session-id}</pre>	<p>Changes the non-SSL host affinity method for this host.</p> <p>Changes the SSL host affinity method for this host.</p>
<pre>load-balance</pre>	<pre>method {default least-connections round-robin}</pre>	<p>Changes the load balancing method for this host.</p>
<pre>proxy</pre>		<p>Defines this host as a proxy instead of a server; any HTTPS or TCP port is deleted.</p>
<pre>server</pre>		<p>Defines this host as a server instead of a proxy.</p>
<pre>ssl-verify-server</pre>		<p>Sets SSL to specify that the ProxySG checks the CA certificate of the upstream server for this host.</p>

2. (Optional) Enter the following commands to negate or disable settings for this host:

```

SGOS#(config forwarding host_alias) no {ftp | http | https | mms | rtsp | tcp
| telnet}
-or-
SGOS#(config forwarding host_alias) no group
SGOS#(config forwarding host_alias) no host-affinity {method | ssl-method}
-or-
SGOS#(config forwarding host_alias) no load-balance method
-or-
SGOS#(config forwarding host_alias) no ssl-verify-server

```

where:

no {ftp http https mms rtsp tcp telnet}		Clears the specified protocol and port from this host.
no group		Removes this host from any and all groups.
no host-affinity	method ssl-method	Disables the host affinity method (non-SSL or SSL) for this host.
no load-balance	method	Disables the load balancing method for this host.
no ssl-verify-server		Disables SSL verification for this host.

Example

```

SGOS#(config) forwarding
SGOS#(config forwarding) edit testhost
SGOS#(config forwarding testhost) server
ok
SGOS#(config forwarding testhost) no ftp
ok
SGOS#(config forwarding testhost) exit
SGOS#(config forwarding) exit
SGOS#(config)

```

Editing a Forwarding Host Group

Once you have created a forwarding host group, you can edit its configuration.

To Edit the Settings of a Forwarding Host Group through the CLI:

1. At the (config) command prompt, enter the following commands to configure the settings of a forwarding host group:

```

SGOS#(config) forwarding
SGOS#(config forwarding) edit group_alias
SGOS#(config forwarding group_alias) host-affinity method {accelerator-cookie
| client-ip-address | default}
-or-
SGOS#(config forwarding group_alias) host-affinity ssl-method
{accelerator-cookie | client-ip-address | default | ssl-session-id}

```

```
SGOS#(config forwarding group_alias) load-balance hash {default | domain | url}
```

-or-

```
SGOS#(config forwarding group_alias) load-balance method {default | least-connections | round-robin}
```

where:

<code>host-affinity</code>	<code>method {accelerator-cookie client-ip-address default}</code>	Changes the non-SSL host affinity method for this group.
	<code>ssl-method {accelerator-cookie client-ip-address default ssl-session-id}</code>	Changes the SSL host affinity method for this group.
<code>load-balance</code>	<code>hash {default domain url}</code>	Changes if and how load balancing hashes between group members for this group.
	<code>method {default least-connections round-robin}</code>	Changes the load balancing settings of the method specified for this group.

2. (Optional) Enter the following commands to disable settings for a forwarding host group:

```
SGOS#(config forwarding group_alias) no host-affinity {method | ssl-method}
```

-or-

```
SGOS#(config forwarding group_alias) no load-balance {hash | method}
```

where:

<code>no host-affinity</code>	<code>method ssl-method</code>	Disables a host affinity method (non-SSL or SSL) for this group.
<code>no load-balance</code>	<code>hash method</code>	Disables the specified load balancing setting for this group.

Example

```
SGOS#(config) forwarding  
SGOS#(config forwarding) edit testgroup  
SGOS#(config forwarding testgroup) host-affinity method client-ip-address  
ok  
SGOS#(config forwarding testgroup) no load-balance hash  
ok  
SGOS#(config forwarding testgroup) exit  
SGOS#(config forwarding) exit  
SGOS#(config)
```

Configuring Load Balancing

Load-balancing of groups has two stages:

- In the first stage, you optionally apply a hash (a domain name or full URL) to select one host from the group.

- In the second stage, you apply a method (`round-robin`, `least-connections`, or `none`) to the IP addresses of the selected host (the host must have more than one IP address).

If the hash in the first stage is disabled, all of the IP addresses in the group are gathered together and the selected method is applied to the whole set of IP addresses. If the hash in the first stage is enabled and has more than one IP address, it will select one of the member hosts of the group, and that selected host will then have the load balancing method applied to its set of IP addresses.

Load balancing defaults can be configured globally. Each individual host or group will use those defaults unless it is optionally configured to its own private values which override those defaults. See "Creating Forwarding Hosts or Host Groups" on page 702, "Editing a Forwarding Host" on page 705, or "Editing a Forwarding Host Group" on page 707 for ways to set a host or group's individual load-balance configuration (including ways to return it to using the global defaults). The top level load-balance commands shown below can be used to set the global defaults. Those top level commands also provide an alternative way to set the override values for a particular host or group or to return a host or group to using the global defaults.

To Set Load Balancing through the CLI:

```
SGOS#(config) forwarding
SGOS#(config forwarding) load-balance hash {domain | no | url}
SGOS#(config forwarding) load-balance method {least-connections | no |
round-robin}
```

where:

hash	{domain no url}	If you use the hash for load balancing, you can choose to hash the domain or the full URL or you can choose no to disable hashing, and the load balancing method applies across a group.
method	{least-connections no round-robin}	If you use method for load balancing, you can select the round-robin method or the least-connections method, or you can specify no to disable load balancing.

To Configure Group Load Balancing through the CLI:

```
SGOS#(config) forwarding
SGOS#(config forwarding) load-balance hash {default | domain | no | url}
group_alias
SGOS#(config forwarding) load-balance method {default | least-connections |
no | round-robin} host_or_group_alias
```

where:

hash	{default domain no url} group_alias	You can specify a group to apply the load-balancing hash setting to only that group.
method	{default least-connections no round-robin} host_or_group_alias	You can specify a host or group to apply the load-balancing method to only that host or group.

Example

```
SGOS#(config forwarding) load-balance method least-connections
test-host-name
ok
```

Configuring Host Affinity

Host affinity is the attempt to direct multiple connections by a single user to the same group member. Host affinity affects load balancing behavior.

Host affinity defaults can be configured globally. Each individual host or group will use those defaults unless it is optionally configured to its own private values which override those defaults. See "Creating Forwarding Hosts or Host Groups" on page 702, "Editing a Forwarding Host" on page 705, or "Editing a Forwarding Host Group" on page 707 for ways to set a host or group's individual host-affinity configuration (including ways to return it to using the global defaults). The top level host-affinity commands shown below can be used to set the global defaults. Those top level commands also provide an alternative way to set the override values for a particular host or group or to return a host or group to using the global defaults.

Note: The non-SSL host affinity methods are implemented for HTTP only and the SSL host affinity methods are implemented for HTTPS only.

To Configure Host Affinity through the CLI:

```
SGOS#(config) forwarding
SGOS#(config forwarding) host-affinity method {accelerator-cookie |
client-ip-address | no}
-or-
SGOS#(config forwarding) host-affinity ssl-method {accelerator-cookie |
client-ip-address | ssl-session-id | no}
SGOS#(config forwarding) host-affinity timeout minutes
```

where:

method	{accelerator-cookie client-ip-address no}	Sets which non-SSL host-affinity method to use (accelerator cookie or client-ip-address) or you can use no to disable non-SSL host affinity.
ssl-method	{accelerator-cookie client-ip-address ssl-session-id no}	Sets which SSL host-affinity method to use (accelerator cookie, client-ip-address, or ssl-session-id) or you can use no to disable SSL host affinity.
timeout	<i>minutes</i>	Determines how long a user's IP address, SSL ID, or cookie remains valid.

To Configure Group-Specific Host Affinity Settings through the CLI:

```

SGOS#(config) forwarding
SGOS#(config forwarding) host-affinity method {accelerator-cookie |
client-ip-address | default | no} host_or_group_alias
-or-
SGOS#(config forwarding) host-affinity ssl-method {accelerator-cookie |
client-ip-address | default | no | ssl-session-id} host_or_group_alias

```

where:

<pre> method {accelerator-cookie client-ip-address default no ssl-session-id} host_or_group_alias </pre>	<p>You can choose which non-SSL host-affinity method to use (accelerator cookie or client-ip-address) for a specific host or group, or you can use no to disable non-SSL host affinity for a specific host or group. You can also apply the global non-SSL host-affinity method to a specific host or group.</p>
<pre> ssl_method {accelerator-cookie client-ip-address default no ssl-session-id} host_or_group_alias </pre>	<p>You can choose which SSL host-affinity method to use (accelerator cookie, client-ip-address, or ssl-session-id) for a specific host or group, or you can use no to disable SSL host affinity for a specific host or group. You can also apply the global SSL host-affinity method to a specific host or group (use the default command).</p>

Example

```

SGOS#(config forwarding) host-affinity method client-ip-address
ok
SGOS#(config forwarding) host-affinity ssl-method no test-group-name
ok
SGOS#(config forwarding) host-affinity timeout 45
ok

```

Creating a Default Sequence

The default sequence defines the order in which forwarding hosts are used in case of failover and which host to use first (only one default sequence is allowed). All members must be pre-existing hosts and groups, and no member can be in the group more than once.

Note: The default sequence replaces the deprecated `default` and `backup` settings. The default sequence (if present) is applied only if no applicable forwarding gesture is in policy.

A default failover sequence (and any sequence specified in policy) 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 through the CLI. 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 a Default Sequence through the CLI:

From the (config) prompt, enter the following commands:

```
SGOS#(config forwarding) sequence add alias_name
SGOS#(config forwarding) sequence clear
SGOS#(config forwarding) sequence demote alias_name
SGOS#(config forwarding) sequence promote alias_name
SGOS#(config forwarding) sequence remove alias_name
```

where:

add	<i>alias_name</i>	Adds an alias to the end of the default failover sequence.
clear		Clears the default failover sequence.
demote	<i>alias_name</i>	Moves an alias one place towards the end of the default failover sequence.
promote	<i>alias_name</i>	Moves an alias one place towards the start of the default failover sequence.
remove	<i>alias_name</i>	Removes an alias from the default failover sequence.

Example

```
SGOS#(config forwarding) sequence clear
ok
```

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 will get an error message. You must remove the host/group from the sequence first, then delete.

Using Forwarding Directives to Create an Installable List

You can use either directives or the CLI to create and configure forwarding hosts. To use the CLI, see "Configuring Forwarding through the CLI" on page 702.

The forwarding configuration includes directives that:

- Create the forwarding hosts and groups
- Provide load balancing and host affinity

Table 19.1: Forwarding Directives

Directive	Meaning	See
fwd_fail	Determines whether the forwarding host should fail open or fail closed if an operation does not succeed. Fail open is a security risk.	"Setting Fail Open/Closed and Host Timeout Values" on page 715.

Table 19.1: Forwarding Directives

Directive	Meaning	See
<code>fwd_host</code>	Create a forwarding host and set configuration parameters for it, including protocols and ports.	"Creating Forwarding Host and Group Directives" on page 713.
<code>host_affinity</code>	The attempt to direct multiple connections by a single user to the same group member.	"Configuring Host Affinity Directives" on page 716.
<code>integrated_host_timeout</code>	An origin content server 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.	"Setting Fail Open/Closed and Host Timeout Values" on page 715.
<code>load_balance</code>	The attempt to manage the load among forwarding hosts in a group, or among multiple IP addresses of a host.	"Configuring Load Balancing Directives" on page 716.
<code>sequence alias_list</code>	where <i>alias_list</i> is a space separated list of one or more forwarding host and group aliases.	"Creating a Default Sequence" on page 717.

Creating Forwarding Host and Group Directives

You can add directives into the forwarding installable list that allows you to create and delete the forwarding host and associate protocols and ports with the host.

You can create a maximum of 32 groups, and each group can contain a maximum of 512 hosts. You can create 512 individual hosts that do not belong to any group.

To create a forwarding host, choose the protocols you want to use, or optionally add the forwarding host to a group, enter the following into your installable list. You should create a `fwd_host` directive for each forwarding host you want to create.

```
fwd_host host_alias host_name [default-schemes] [http[=port | =no]]
[https[=port | =no]] [ftp[=port | =no]] [mms[=port | =no]] [rtsp[=port |
=no]] [tcp=port] [telnet[=port | =no]] [ssl-verify-server[=yes | =no]]
[group=group_name] [server | proxy] [load-balance={no | round-robin |
least-connections}] [host-affinity={no | client-ip-address |
accelerator-cookie}] [host-affinity-ssl={no | client-ip-address |
accelerator-cookie | ssl-session-id}]
```

where:

`host_alias`

This is the alias for use in policy. Define a name meaningful to you.

`host_name`

The name of the host domain, such `www.bluecoat.com`, or its IP address.

`default-schemes`

If you use `default-schemes` in the directive, all protocols, along with their default ports are selected. This directive is only available for proxy hosts.

<code>http</code>	<code>=port =no</code>	No protocol is selected by default if the forwarding host is a server. You must choose at least one protocol where <code>port=0</code> to 65535. If only one protocol is configured, the ProxySG configures the default port for that protocol.
<code>https</code>		
<code>ftp</code>		
<code>mms</code>		
<code>rtsp</code>		You can use <code>default-schemes</code> and then eliminate protocols by selecting the protocol you do not want; for example, <code>http=no</code> . If you do not want to use the default ports for the protocols, you must also specify them here.
<code>telnet</code>		HTTPS protocols are not allowed if the host is a proxy.
<code>tcp</code>	<code>=port</code>	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.
<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> . 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</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. Do not use the <code>group=</code> option when creating independent hosts.
<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> .
<code>load-balance</code>	<code>=no =round-robin =least-connections</code>	Specifies either the least-connections or round-robin method of load balancing. Select <code>no</code> to disable load balancing for this forwarding host or host group. If these settings are not specified for a particular host or host group, then the global default settings are used. To configure the settings for a specific host or host group, use the <code>edit host_alias</code> or <code>edit group_alias</code> commands (see "Editing a Forwarding Host" on page 705 or "Editing a Forwarding Host Group" on page 707).

<pre>host-affinity =no =client-ip-address =accelerator-cookie</pre>	<p>Specifies non-SSL host affinity via either a client IP address or an accelerator cookie. Select <code>no</code> to disable non-SSL host affinity for this forwarding host or host group.</p> <p>If these settings are not specified for a particular host or host group, then the global default settings are used. To configure the settings for a specific host or host group, use the <code>edit host_alias</code> or <code>edit group_alias</code> commands (see "Editing a Forwarding Host" on page 705 or "Editing a Forwarding Host Group" on page 707).</p>
<pre>host-affinity =no -ssl =client-ip-address =accelerator-cookie =ssl-session-id</pre>	<p>Specifies SSL host affinity via a client IP address, an accelerator cookie, or an SSL session ID. Select <code>no</code> to disable SSL host affinity for this forwarding host or host group.</p> <p>If these settings are not specified for a particular host or host group, then the global default settings are used. To configure the settings for a specific host or host group, use the <code>edit host_alias</code> or <code>edit group_alias</code> commands (see "Editing a Forwarding Host" on page 705 or "Editing a Forwarding Host Group" on page 707).</p>

Example

```
fwd_host www.bluecoat1.com 10.25.36.48 default-schemes ssl-verify-server=no
group=bluecoat
```

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, Blue Coat 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.

The syntax is:

```
fwd_fail {open | closed}
integrated_host_timeout minutes
```

where:

<pre>fwd_fail</pre>	<pre>{open closed}</pre>	<p>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.</p> <p>This setting can be overridden by policy, (using the <code>forward.fail_open</code> (yes no) property).</p>
---------------------	----------------------------	---

```
integrated_host_timeout minutes
```

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
```

Configuring Load Balancing Directives

Load balancing shares the load among a set of IP addresses, whether a group or a host with multiple IPs.

The syntax is:

```
load_balance hash {domain | no | url} [group_alias]
load_balance method {least-connections | round-robin | no}
[host_or_group_alias]
```

where:

```
hash      {domain | no | url}
          [group_alias]
```

If you use the hash for load balancing, you can choose to hash the domain or the full URL, or you can choose no to disable hashing and the load-balancing method will apply across a group. If you do not specify a group, the settings apply as the default for all groups.

```
method    {least-connections |
          no | round-robin}
          [host_or_group_alias]
```

If you use method for load balancing, you can select the least-connections method or the round-robin method, or you can specify no to disable load balancing (hashing will still occur if it is set). If you do not specify a host or group, the settings apply as the default for all hosts or groups.

Example

```
load_balance method 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 method {accelerator-cookie | client-ip-address | no}
[host_or_group_alias]
host_affinity ssl_method {accelerator-cookie | client-ip-address | no |
ssl-session-id} [host_or_group_alias]
host_affinity timeout seconds
```


where:

<pre>method {accelerator-cookie client-ip-address no} [host_or_group_alias]</pre>	<p>You can choose which non-SSL host-affinity method to use (accelerator cookie or client-ip-address), or you can use no to disable non-SSL host affinity. If you do not specify a host or group, the settings apply as the default for all hosts or groups.</p>
<pre>ssl_method {accelerator-cookie client-ip-address no ssl-session-id} [host_or_group_alias]</pre>	<p>You can choose which SSL host-affinity method to use (accelerator cookie, client-ip-address, or ssl-session-id), or you can use no to disable SSL host affinity. If you do not specify a host or group, the settings apply as the default for all hosts or groups.</p>
<pre>timeout minutes</pre>	<p>Determines how long a user's IP address, SSL ID, or cookie remains valid.</p>

Example

```
host_affinity ssl_method 10.25.36.48
host_affinity timeout 5
```

Creating a Default Sequence

A default sequence defines the order in which forwarding hosts are used. Only one default sequence is allowed. All members must be pre-existing hosts and groups, and no member can be in the group more than once.

Note: The default sequence, completely overridden by policy, replaces the deprecated `default` and `backup` settings.

A default failover 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 through the CLI. 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.

The syntax is

```
sequence alias_list alias_list
  where alias_list is a space-separated list of one or more forwarding host and group aliases.
```

Example

```
sequence bluecoat
```

Creating a Forwarding Installable List

You can create and install the forwarding installable list with the following methods:

- Use the ProxySG 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 ProxySG.
- Create a local file on your local system; the ProxySG can browse to the file and install it.
- 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.

When the Forwarding Installable List is installed, it updates the forwarding directives on the ProxySG. The directives remain in effect until they are overwritten by another installable list; the list can be modified or overwritten using CLI commands.

Note: During the time that a forwarding installable list is being compiled and installed, forwarding is not available. Any transactions that come into the ProxySG during this time will not be forwarded properly and will be denied.

Installation of forwarding installable lists should be done outside peak traffic times.

To Create a Forwarding Installable List through the Management Console:

1. Select Configuration>Forwarding>Forwarding Hosts.

The Forwarding Hosts tab displays.

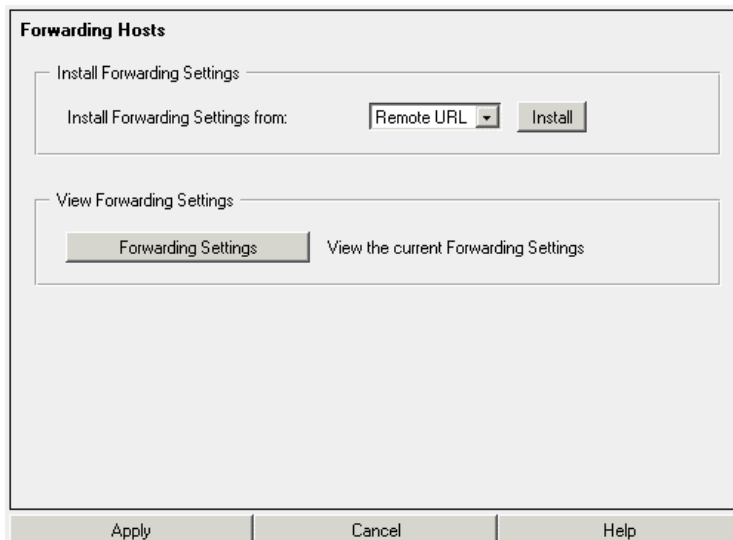


Figure 19-1: Selecting the Forwarding Hosts Download Method

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 ProxySG.

- 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.

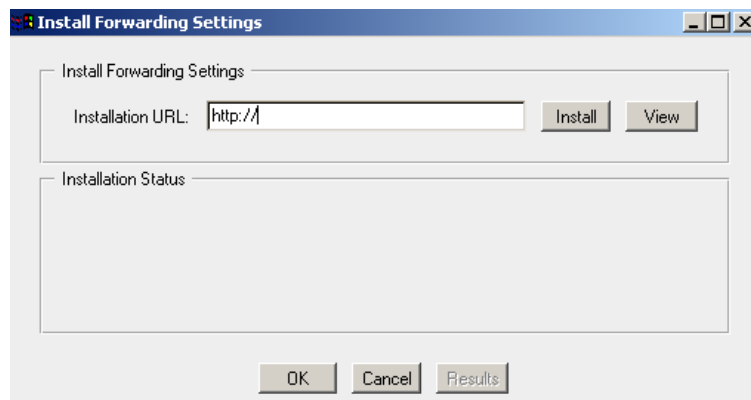


Figure 19-2: Specifying the Remote Location of a Forwarding Configuration

□ Local File:

Click Browse to bring up 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.

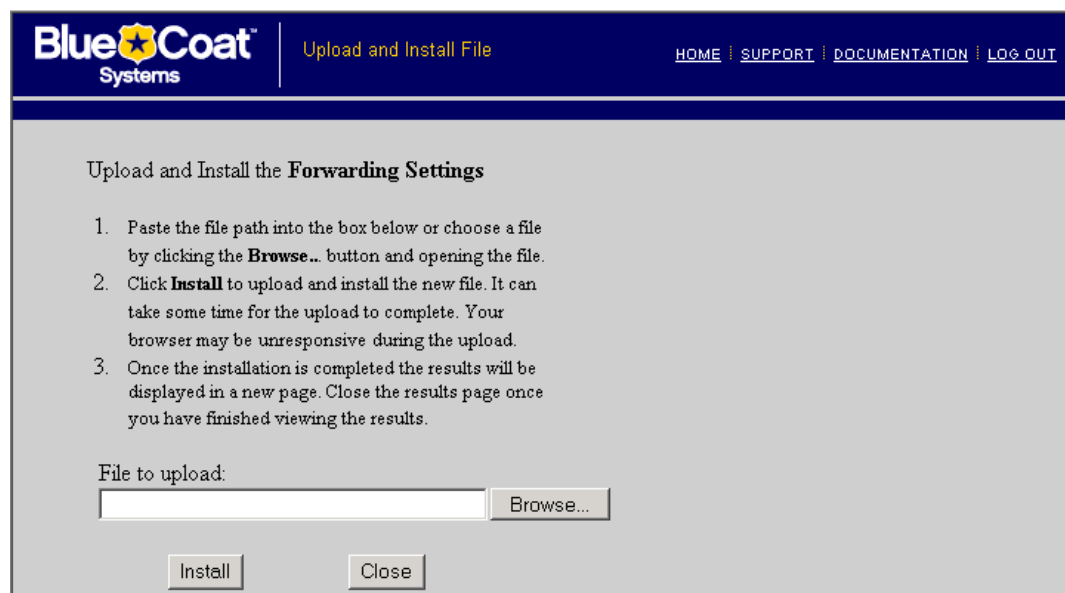


Figure 19-3: Specifying the Local Location of a Forwarding Configuration

□ 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.

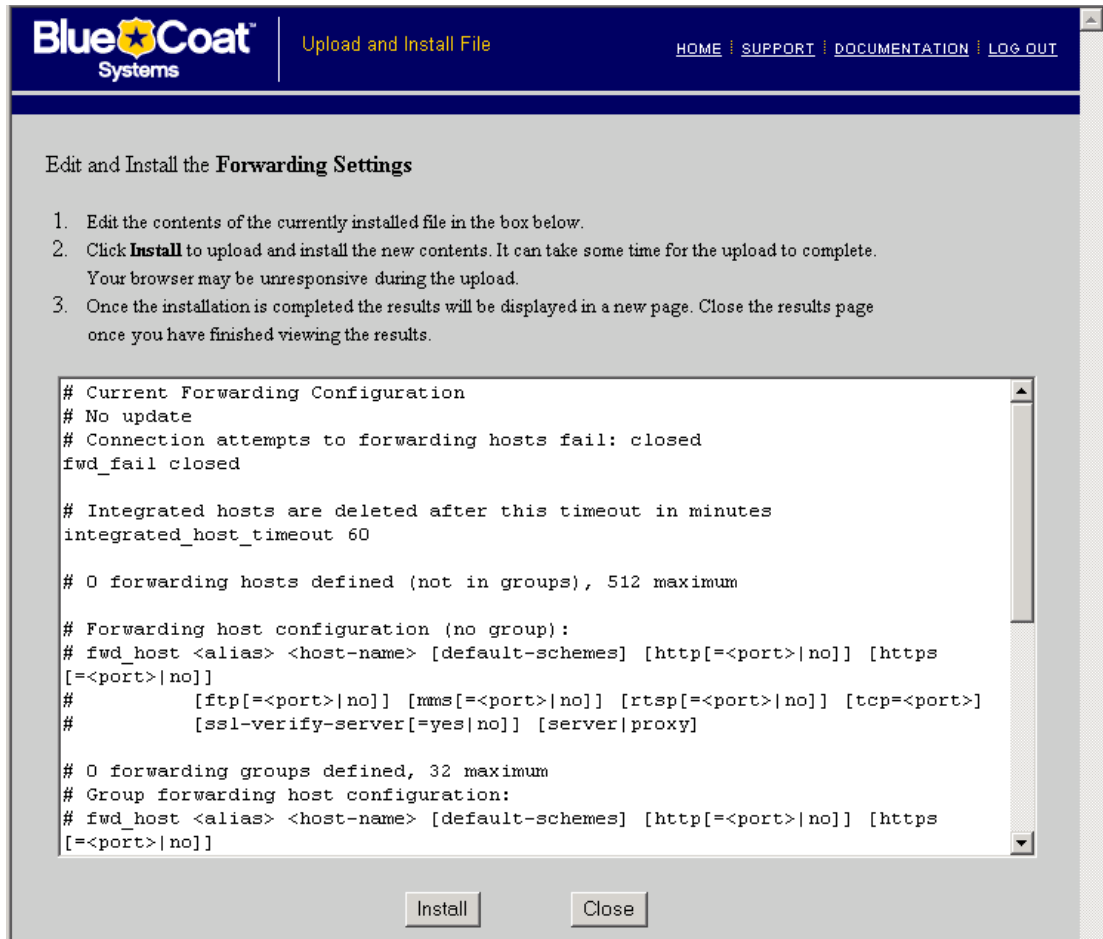


Figure 19-4: Using the ProxySG Text Editor

3. Click Apply.

To Create a Remote Forwarding Installable List through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) forwarding
SGOS#(config forwarding) path url
```

where *url* is a fully-qualified URL, including the filename, where the installable list is located.

```
SGOS#(config forwarding) exit
SGOS#(config) load forwarding
```

To Create Forwarding Settings on the ProxySG through the CLI:

1. At the `(config)` command prompt, enter the following commands to create an inline set of commands. You can use any of the forwarding directives, but host affinity and load balancing are mutually exclusive. The procedure below demonstrate the creation of an inline forwarding configuration, using the non-SSL host affinity method:

```
SGOS#(config) inline forwarding eof
fwd_host test 10.25.36.47 default-schemes
host_affinity method client-ip-address
host_affinity timeout 45
eof
ok
```

where:

<code>forwarding</code>	Identifies the kind of inline settings you are creating.
<code>eof</code>	Indicates the marker you use to tell the CLI that you are beginning or ending the set of commands. You can use any characters as the end-of-file marker.

The limitation to using the `inline` command to create a configuration is that you cannot create mistakes except on the current line. If you find an error farther back than that, you must start over after exiting the current file.

2. View the results.

```
SGOS#(config) show forwarding
download-via-forwarding: enabled
Connection attempts to forwarding hosts fail: closed.
Forwarding Groups: (* = host unresolved)
  Group: techpubs
    test3          10.25.36.47 http=80 ftp=21 rtsp=554
Individual Hosts: (* = host unresolved)
  No individual hosts defined.
Load balancing hash: domain
Load balancing method: no
Host affinity method (non-SSL): client-ip-address
Host affinity method (SSL): client-ip-address
Host affinity timeout: 45 minutes
```

To Delete Forwarding Settings on the ProxySG through the CLI:

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 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.

SOCKS Gateway Configuration

The ProxySG implementation of SOCKS includes the following:

- A SOCKS proxy server that supports both SOCKSv4/4a and SOCKSv5, running on the ProxySG.
- Support for forwarding through SOCKS gateways.

To configure a SOCKS proxy server on the ProxySG, see "Configuring a SOCKS Proxy" on page 188. To use SOCKS gateways when forwarding, continue with the next section.

Note: SOCKS gateway aliases cannot be CPL keywords, such as `no`, `default`, `forward`, or `socks_gateways`.

Using SOCKS Gateways

SOCKS servers provide application level firewall protection for an enterprise. The SOCKS protocol provides generic way to proxy HTTP.

SOCKS gateways, like ICP and forwarding, can use installable lists for configuration. You can configure the installable list using directives. You can also use the CLI to create a SOCKS gateways configuration.

Using the CLI to Create SOCKS Gateways Settings

If you prefer, you can use SOCKS gateways CLI commands, instead of an installable list, to create SOCKS gateways settings. For information about using an installable list, see "Using SOCKS Gateways Configuration Directives to Create an Installable List" on page 726.

To Create a SOCKS Gateways Host through the CLI:

1. At the `(config)` command prompt, enter the following commands:

```
SGOS#(config) socks-gateways
SGOS#(config socks-gateways) create gateway_alias gateway_host SOCKS_port
[version {=4 | =5 [user=username password=password] [request-compression
{=yes | =no}]]]
```

where:

<code>gateway_alias</code>	A name, meaningful to you
<code>gateway_host</code>	The IP address or the host name of the gateway where traffic will be directed. The host name must DNS resolve.
<code>SOCKS_port</code>	The port number of the SOCKS gateway.
<code>version</code> <code>=4 =5</code>	The version that SOCKS gateways can support. (SOCKS v5 is recommended, if you have a choice). If no version is configured, the default is version 4.

<code>user</code>	<code>=username</code>	(Optional, and only if you use v5) The username of the user on the SOCKS gateway. The username already must exist on the gateway. If you use <code>user=</code> , you must also use <code>password=</code> .
<code>password</code>	<code>=password</code>	(Optional, and only if you use v5) The password of the user on the SOCKS gateway. The password must match the gateway's information. If you use <code>user=</code> , you must also use <code>password=</code> .
<code>request-compression</code>	<code>=yes =no</code>	(Optional, and only if you use v5) Enable or disable SOCKS compression. The default is <code>no</code> . To use SOCKS compression, you must enable compression on a SOCKS gateway, enable an Endpoint Mapper proxy, and create policy to forward TCP traffic through the SOCKS gateway. For more information, see "Understanding SOCKS Compression" on page 188

- Repeat for step 1 for each gateway you want to create. The `failure-mode` command applies to all SOCKS gateways configured on the system. The default failure mode can be overridden using policy.
- Complete the configuration by entering the following commands as necessary:

```
SGOS#(config socks-gateways) failure-mode {open | closed}
SGOS#(config socks-gateways) delete {all | gateway gateway_alias}
SGOS#(config socks-gateways) path url
SGOS#(config socks-gateways) no path
```

where

<code>failure-mode</code>	<code>open closed</code>	If the health checks fail, <code>open</code> specifies that the connection be attempted without use of any SOCKS gateway (whether to an origin content server or a forwarding target); <code>closed</code> specifies that the connection be aborted.
<code>delete</code>	<code>all gateway <i>gateway_alias</i></code>	Deletes all SOCKS gateways (<code>delete all</code>) or a specific SOCKS gateway (<code>delete gateway <i>gateway_alias</i></code>).
<code>path</code>	<code><i>url</i></code>	(Optional) Specifies the download path to use if you download SOCKS-gateways settings through directives.
<code>no</code>	<code>path</code>	Clears the network path URL to download SOCKS gateway settings.

- View the results.

```
SGOS#(config socks-gateways) view
SOCKS Gateways: (* = gateway unresolved)
Sec_Appl          10.25.36.47 1080 v5
```

Editing a SOCKS Gateways Host

Once you have created a SOCKS gateways host, you can edit the settings.

To Edit the Settings of a SOCKS Gateways Host through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) socks-gateways
SGOS#(config socks-gateways) edit gateway_alias
SGOS#(config socks-gateways gateway_alias) host gateway_host
SGOS#(config socks-gateways gateway_alias) no password | user
SGOS#(config socks-gateways gateway_alias) password password
SGOS#(config socks-gateways gateway_alias) port socks_port
SGOS#(config socks-gateways gateway_alias) user username
SGOS#(config socks-gateways gateway_alias) version 4 | 5
SGOS#(config socks-gateways gateway_alias) request-compression enable |
disable
```

where:

host	gateway_host	Changes the host name.
no	password user	Optional, and only if you use version 5. Deletes the version 5 password or username.
password	password	Optional, and only if you use version 5. Changes the version 5 password.
port	socks_port	Changes the SOCKS port.
user	username	Optional, and only if you use version 5. Changes the version 5 username.
version	4 5	Changes the SOCKS version.
request-compression	enable disable	(Optional, and only if you use v5) Enable or disable SOCKS compression. The default is disable. To use SOCKS compression, you must enable compression on a SOCKS gateway, enable an Endpoint Mapper proxy, and create policy to forward TCP traffic through the SOCKS gateway. For more information, see "Understanding SOCKS Compression" on page 188

Example

```

SGOS#(config) socks-gateways
SGOS#(config socks-gateways) edit testsocks
SGOS#(config socks-gateways testsocks) port 23
    ok
SGOS#(config socks-gateways testsocks) version 5
    ok
SGOS#(config socks-gateways testsocks) exit
SGOS#(config socks-gateways) exit
SGOS#(config)

```

Creating a Default Sequence

A default sequence defines the order in which SOCKS gateway hosts are used. Only one default sequence is allowed. All members must be pre-existing hosts, and no member can be in the group more than once.

Note: The default sequence replaces the deprecated `default` and `backup` settings. The default sequence (if present) is applied only if no applicable forwarding gesture is in policy.

A default failover 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 SOCKS-gateways policy applies, you can create a default sequence through the CLI. 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.

The syntax is

```
sequence alias_name alias_name
```

where *alias_name* is a space-separated list of one or more SOCKS gateways.

To create a default failover sequence, enter the following commands from the `(config)` prompt:

```

SGOS#(config) socks-gateways
SGOS#(config socks-gateways) sequence add gateway-alias
SGOS#(config socks-gateways) sequence promote | demote gateway-alias
SGOS#(config socks-gateways) sequence clear | remove gateway-alias

```

where:

sequence	add	Adds an alias to the end of the default fail-over sequence
	clear	Clears the default fail-over sequence
	demote	Demotes an alias one place towards the end of the default fail-over sequence
	promote	Promotes an alias one place towards the start of the default fail-over sequence
	remove	Removes an alias from the default fail-over sequence.

Using SOCKS Gateways Configuration Directives to Create an Installable List

To configure a SOCKS gateway you must create an installable list and load it on the ProxySG. Alternately, you can use the CLI to configure SOCKS gateways. To use the CLI, see "Using the CLI to Create SOCKS Gateways Settings" on page 722.

For information on installing the file itself, see "Creating a SOCKS Gateway Installable List" on page 728.

The SOCKS gateways configuration includes SOCKS directives that:

- Names the SOCKS gateway hosts
- Specifies the SOCKS version
- (Optional, if using Version 5) Specifies user name and password

Available directives are described in the table below.

Table 19.2: SOCKS Gateway Directives

Directive	Meaning
gateway	Specifies the gateway alias and name, SOCKS port, version supported, usernames and password.
socks_fail	In case connections cannot be made, specifies whether to abort the connection attempt or to connect to the origin content server
sequence	Specifies the order in which hosts should be used for failover.

Syntax for the SOCKS directives are:

```
gateway gateway_alias gateway_host SOCKS_port [version={4 | 5 [user=username
password=password] [request-compression={yes | no}]]]
socks_fail {open | closed}
sequence gateway_name
```

where:

gateway	Configures the SOCKS gateway host.
gateway_alias	A meaningful name that is used for policy rules.
gateway_name	The IP address or host name of the gateway where traffic will be directed. The host name must DNS resolve.
SOCKS-port	The port number of the SOCKS gateway.
version={4 5}	The version that SOCKS gateways can support.
user=username	(Optional, if you use v5) The username of the user on the SOCKS gateway. 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.

<code>request-compression</code>	<code>request-compression =yes =no</code>	(Optional, if you use v5) Enables or disables SOCKS compression. The default is no. To use SOCKS compression, you must enable compression on a SOCKS gateway, enable an Endpoint Mapper proxy, and create policy to forward TCP traffic through the SOCKS gateway. For more information, see "Understanding SOCKS Compression" on page 188
<code>socks_fail</code>	<code>{open closed}</code>	If health checks fail, <code>socks_gateway.fail_open</code> specifies that the connection be attempted without using a SOCKS gateway (for example, go to the original server or forwarding target); <code>socks_gateway.fail_closed</code> specifies that the connection be aborted. The default is <code>closed</code> . 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>sequence</code>	<code>gateway_name</code>	Specifies the order in which hosts should be used for failover.

Example

```
gateway Sec_App1 10.25.36.47 1022 version=5 user=username password=password
socks_gateway.fail_open no
```

Important: The username and password display in clear text if you run the `show config` command.

A default sequence defines the order in which forwarding hosts are used. Only one default sequence is allowed. All members must be pre-existing hosts and groups, and no member can be in the sequence more than once.

Note: The default sequence replaces the deprecated `default` and `backup` settings. The default sequence (if present) is applied only if no applicable forwarding gesture is in policy.

A default failover 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 SOCKS-gateways policy applies, you can create a default sequence through the CLI. 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.

The syntax is

```
sequence gateway_name gateway_name
```

where `gateway_name` is a space-separated list of one or more SOCKS gateway aliases.

Example

```
sequence gateway_alias
```

Creating a SOCKS Gateway Installable List

You can create and install the SOCKS gateway installable list with the following methods:

- Use the ProxySG 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.

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 CLI commands.

Note: During the time that a SOCKS gateway installable list is being compiled and installed, forwarding is not available. Any transactions that come into the ProxySG during this time will not be forwarded properly and will be denied.

Installation of SOCKS gateways installable-list configuration should be done outside peak traffic times.

To Create a SOCKS Gateways Installable List through the Management Console:

1. Select Configuration>Forwarding>SOCKS Gateways.

The SOCKS Gateways tab displays.

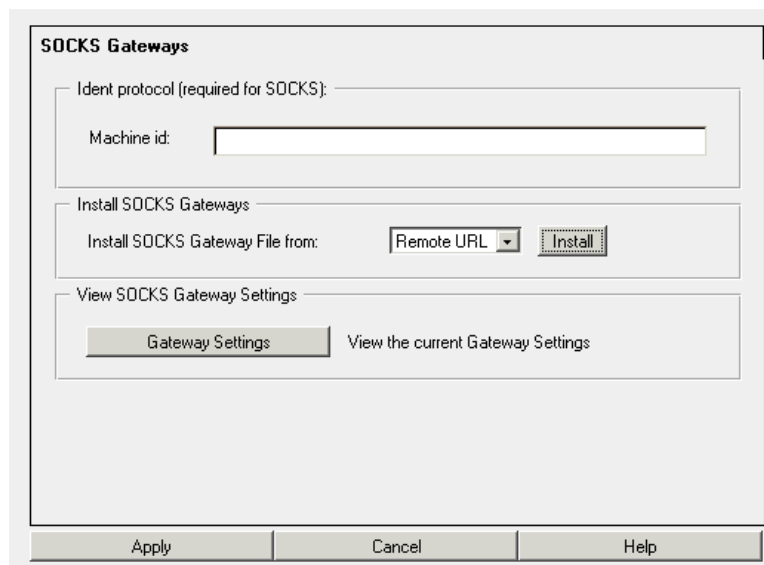


Figure 19-5: Selecting the SOCKS Gateways Tab

2. If you use a SOCKS gateway server for the primary or alternate forwarding gateway, you must specify the ProxySG 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.

To Specify the SOCKS Gateway Machine ID through the CLI:

Note: This is an optional command. The default is BLUE COAT SYSTEMS.

At the `config` command prompt, enter the following command:

```
SGOS#(config) socks-machine-id machine_ID
```

To Create a Remote SOCKS Gateways Installable List through the CLI:

At the `(config)` prompt, enter the following commands:

```
SGOS#(config) socks-gateways
```

```
SGOS#(config socks-gateways) path url
```

where `url` is a fully-qualified URL, including the filename, where the configuration is located.

```
SGOS#(config) socks-gateways) exit
```

```
SGOS#(config) load socks-gateways
```

Tip for SOCKS Configuration

By default, SOCKS treats all incoming requests destined to port 80 as HTTP, allowing the usual HTTP policy to be done on them, including ICAP scanning. If the SOCKS connection is being made to a server on another port, you should write policy on the ProxySG to match on the server host and port and specify that it is HTTP using SOCKS.

Internet Caching Protocol (ICP) Configuration

ICP is a communication protocol for caches. It allows a cache (not necessarily a ProxySG) to query other caches for an object, without actually requesting the object. By using ICP, the cache can determine if the object is available from a neighboring cache, and which cache will provide the fastest response.

Note: The ProxySG (assuming ICP is configured) does ICP queries only if no forwarding host or SOCKS gateway is identified as an upstream target. If ICP is used by the ProxySG, it prompts other cache devices for the item, and upon a positive response re-directs the upstream request to that cache device instead of the content origin server.

Only use ICP if you have ICP hosts available or if you want the ProxySG to support requests from other ICP hosts.

By default, the ICP protocol requires the requesting host to wait up to two seconds for all ICP hosts to respond to the request for an object (the time is configurable).

If the ICP service is configured and running, the service is used if no forwarding or SOCKS gateway target was specified. In other words, the policy rule `icp(yes)` is the default, assuming that the ICP service is available. You can disable ICP with the policy rule `icp(no)` to control ICP queries for requests.

Configuring ICP

An ICP *hierarchy* is comprised of a group of caches, with defined parent and sibling relationships. A cache parent is one that can return the object if it is in the cache, or request the object from the source on behalf of the requester if the object is not in the cache. A cache sibling is a device that can only return the object if it is in the cache. One cache acting as a parent can also act as a sibling to other cache devices.

- When an object is not cached, the cache device sends an ICP query to its neighbors (parents and siblings) to see if any of its peers holds the object.
- Each neighbor that holds the requested object returns an `ICP_HIT` reply.
- Each neighbor that does not hold the object returns an `ICP_MISS` reply.

Based on the responses, the cache can determine where to request the object: from one of its neighbors or from the source. If an `ICP_HIT` reply is received, the request is sent to the host that returned the first reply. If no `ICP_HIT` reply is received, the request is forwarded to the first parent that replied. If no parents respond or are configured, the request is made directly to the source.

Using ICP Configuration Directives to Create an Installable List

To configure ICP you must create an installable list and load it on the ProxySG. The ICP protocol contains a number of *directives*, commands used to create a list that can be installed on the ProxySG.

For information on installing the file itself, see "Creating an ICP Installable List" on page 734.

The ICP configuration includes directives that:

- Name the ICP hosts

- Restrict ICP access to only these hosts

Available directives are listed in Table 19.3.

Table 19.3: ICP Directives

Directive	Meaning	Where used
icp_host	The <code>icp_host</code> directive describes cache peers in the hierarchy. There should be one entry for each ProxySG you want to use.	Names the ICP hosts. See "Naming the IP Hosts" on page 732.
icp_access_domain	The <code>icp_access_domain</code> directive is used to control which ICP queries are accepted. The <code>icp_access_domain</code> directive requires a reverse DNS lookup of each ICP query to validate the IP address.	Restricts access. See "Restricting Access" on page 732.
icp_access_ip	The <code>icp_access_ip</code> directive works like the <code>icp_access_domain</code> command, except you can specify an IP address and subnet mask rather than a domain.	Restricts access. See "Restricting Access" on page 732.
icp_port	The <code>icp_port</code> directive sets the port the ProxySG uses to listen for ICP requests. The default port is 3130. If you set the port to 0, ICP is disabled.	Connects to other ICP hosts. See "Connecting to other ICP Hosts" on page 733.
neighbor_timeout	The <code>neighbor_timeout</code> directive sets the number of seconds the ProxySG waits for ICP replies. When the cache device sends an ICP request, it waits for all hosts to reply or for the <code>neighbor_timeout</code> to expire. The default timeout is 2 seconds.	Connects to other ICP hosts. See "Connecting to other ICP Hosts" on page 733.
icp_failcount	The <code>icp_failcount</code> directive sets the number of consecutive failures the cache device can receive before considering the ICP host as failed. By default, the ICP failure count is set to 20. Each time a request fails, the failure count is incremented. When a request succeeds, the failure count is reset to zero.	Connects to other ICP hosts. See "Connecting to other ICP Hosts" on page 733.
http_failcount	The <code>http_failcount</code> directive sets the number of consecutive failures the cache device can receive before considering the HTTP host as failed. By default, the HTTP failure count is set to 5. The failure count increments each time a request fails. When a request succeeds, the failure count is reset to zero. When an HTTP host fails, the cache device waits five minutes before attempting to use it again as a forwarding target. If the next request fails, the cache device continues to wait five minutes between attempts until the cache becomes available.	Connects to other ICP hosts. See "Connecting to other ICP Hosts" on page 733.

Table 19.3: ICP Directives (Continued)

Directive	Meaning	Where used
<code>host_fail_notify</code>	The <code>host_fail_notify</code> directive tells the cache device to send event notification email when a connect fails persistently.	Connects to other ICP hosts. See "Connecting to other ICP Hosts" on page 733.
<code>host_recover_notify</code>	The <code>host_recover_notify</code> directive tells the cache device to send event notification email when a failed host recovers.	Connects to other ICP hosts. See "Connecting to other ICP Hosts" on page 733.

Naming the IP Hosts

The `icp_host` directive describes peers in the hierarchy. One entry is required for each ProxySG you want to use.

```
icp_host hostname peertype HTTPport ICPport [default | backup | feeder]
```

where:

<code>hostname</code>		The host name of the ProxySG.
<code>peertype</code>	{parent sibling}	Relationship of the ProxySG to the cache device you are configuring.
<code>HTTPport</code>		TCP port where the ProxySG accepts HTTP requests. The common HTTP port is 80 or 8080.
<code>ICPport</code>		UDP port where the ProxySG accepts ICP requests. The common ICP port is 3130.
<code>default</code>		If specified, designates a ProxySG host parent to be the default ICP parent. If no ICP reply is received, all requests will be forwarded to the default parent.
<code>backup</code>		If specified, designates the cache device host parent to be the backup default ICP parent. If the default parent is not available, the cache device uses the backup default parent.
<code>feeder</code>		If specified, designates the ProxySG host sibling as a feeder-type host, using ICP request loops to populate the ProxySG.

The following are sample `icp_host` directives that can be entered into the ICP configuration:

```
; Define ICP parent and sibling hosts.
icp_host cm1.bluecoat.com parent 8080 3130 default
icp_host cm2.bluecoat.com sibling 8080 3130
icp_host cm3.bluecoat.com sibling 8080 3130
icp_host cm4.bluecoat.com sibling 8080 3130
icp_host cm5.bluecoat.com parent 8080 3130
```

Restricting Access

You can restrict access to ProxySG acting as caches by other ICP hosts using the `icp_access_domain` and `icp_access_ip` directives. By default, when ICP is configured, all ICP hosts are allowed access. You should deny access to all domains other than the ICP hosts you want to use.

icp_access_domain Directive

The `icp_access_domain` directive defines which hosts can request objects from the Web cache using ICP. The default action is to allow all requests. When you use `icp_access_domain`, each ICP query requires a reverse DNS lookup to validate the IP address. Depending on the number of ICP requests, these lookups can consume ProxySG resources.

```
icp_access_domain {allow | deny} domain
```

where:

allow deny	Allows or denies ICP queries from neighbors that match the domain specification.
domain	The domain to match. All ICP queries from neighbors that match the specified domain are handled by the host. The special domain of <i>all</i> defines the default action when there is no domain match.

The following are sample `icp_access_domain` directives to be entered into the ICP configuration:

```
; allow ICP access to this Blue Coat Systems ProxySG Appliance from the
; bluecoat.com domain
icp_access_domain allow bluecoat.com
icp_access_domain deny all
; the deny all option should always be specified to deny all other
; domains
```

icp_access_ip Directive

The `icp_access_ip` directive works like the `icp_access_domain` command, except you can specify an IP address and subnet mask rather than a domain. The following describes the parameters for the `icp_access_ip` command:

```
icp_access_ip {allow | deny} subnet mask
```

where:

allow deny	Allow or deny ICP queries from neighbors that match the address specification.
address/subnet mask	The address and subnet mask to match. All ICP queries that match the specified address will be handled by the ICP host. The special address of 0.0.0.0 defines the default action when there is no address match.

The following are sample `icp_access_ip` directives to be entered into the ICP configuration:

```
; allow ICP access to this Blue Coat Systems ProxySG Appliance from the local
subnet
icp_access_ip allow 192.168.10.0/255.255.255.0
icp_access_ip deny 10.25.36.47
; the deny all option should always be specified to deny all other domains
```

Connecting to other ICP Hosts

In addition to the ICP directives described in the sections above, you can specify the following directives in the ICP configuration:

```
icp_port 0
neighbor_timeout 2
icp_failcount 20
http_failcount 5
host_fail_notify on
host_recover_notify on
```

where:

<code>icp_port</code>	The default port is 3130. If you set the port to 0, ICP is disabled.
<code>neighbor_timeout</code>	When the cache device sends an ICP request, it waits for all hosts to reply or for the <code>neighbor_timeout</code> to expire. The default timeout is 2 seconds.
<code>http_failcount</code>	By default, the HTTP failure count is set to 5. The failure count increments each time a request fails. When a request succeeds, the failure count is reset to zero. When an HTTP host fails, the cache device waits five minutes before attempting to use it again as a forwarding target.
<code>icp_failcount</code>	By default, the ICP failure count is set to 20. Each time a request fails, the failure count is incremented. When a request succeeds, the failure count is reset to zero.
<code>host_fail_notify</code>	<code>on</code> tells the cache to send event notification email when a connect fails persistently; <code>off</code> disables this setting.
<code>host_recover_notify</code>	<code>on</code> tells the cache to send event notification email when a failed host recovers; <code>off</code> disables this setting.

Creating an ICP Installable List

You can create the ICP installable list with the following methods:

- Use the ProxySG 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.
- Use the CLI `inline` command.

When the ICP installable list is created and installed, it overwrites any ICP settings on the ProxySG.

To Create an ICP Installable List through the Management Console:

1. Select Configuration>Forwarding>ICP.

The ICP tab displays.

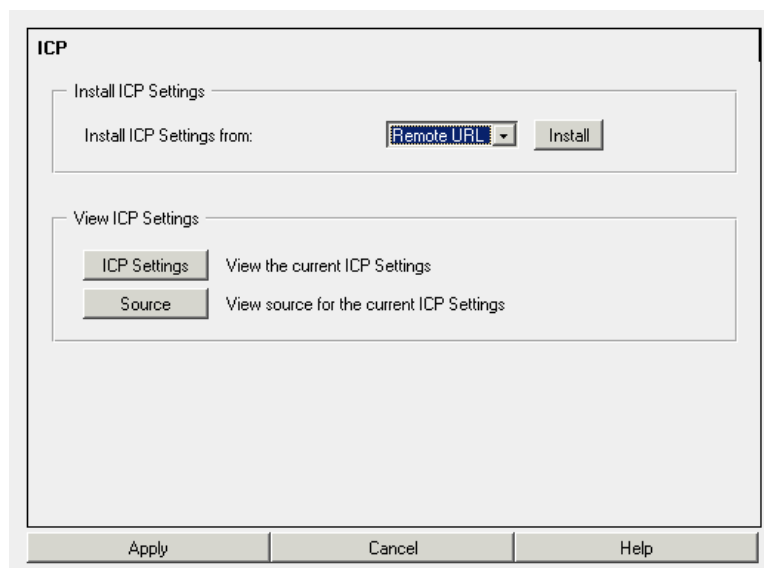


Figure 19-6: Selecting the ICP Download Method

2. From the drop-down list, select the method you want to use to install the ICP configuration; then 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.
3. Click Apply.

To Create a Remote ICP Installable List through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) icp path url
```

where *url* is a fully -qualified URL, including the filename, where the configuration is located.

```
SGOS#(config) load icp-settings
```

To Create ICP Settings on the ProxySG through the CLI:

From the (config) prompt, enter the following commands to create an inline set of commands. You can use any of the ICP directives, not just the ones displayed here.

```
SGOS#(config) inline icp-settings eof
icp_port 0
neighbor_timeout 2
icp_failcount 20
http_failcount 5
eof
ok
```

where:

icp-settings	Identifies the type of inline settings you are creating.
eof	Specifies the marker that tells the CLI that you are using to begin and end the set of commands. You can use any characters as the end-of-file marker.

The drawback to using the inline command to create a configuration is that you cannot correct mistakes except on the current line. If you find an error farther back than that, you must start over after exiting the current file.

Enabling ICP

ICP must be running and at least one forwarding host configured before ICP can be used in the ProxySG environment. ICP can be enabled or disabled through the policy rule `icp`. The default is `icp(yes)`. You can disable ICP with the policy rule `icp(no)` to control ICP queries for requests.

Using Policy to Manage Forwarding

Once ICP, forwarding, and the SOCKS gateways are configured, you can use policy to create and manage forwarding rules. Forwarding, ICP, and SOCKS gateway rules should go in the `<Forward>` layer of your Forwarding Policy file or your VPM Policy file (if you use the VPM).

Note: Because the contents of the Forward policy file are overwritten by the CLI `restore-sgos2-config` or `restore-cacheos4-config` commands, you should back up the file before using them.

The separate `<Forward>` layer (and `server_url` triggers in place of `url` triggers) is provided because the `url` can undergo URL rewrites before the request is fetched. This rewritten URL is accessed as `server_url` and decisions about upstream connections are based on that, requiring a separate layer. All policy commands allowed in the `<Forward>` layer are described in Table 19.4.

Table 19.4: Forwarding Conditions, Properties, Actions, and Definitions

Forwarding	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.

Table 19.4: Forwarding Conditions, Properties, Actions, and Definitions (Continued)

Forwarding	Description
<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.
<code>minute[.utc]=month[.utc]=</code>	Tests if the minute of the hour 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 <code><Admin></code> and <code><Proxy></code> layers.
<code>proxy.card=</code>	Tests the ordinal number of the network interface card (NIC) used by a request. Can also be used in <code><Admin></code> and <code><Proxy></code> 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 <code><Admin></code> and <code><Proxy></code> layers.
<code>server_url[.case_sensitive .no_lookup]=</code>	Tests if a portion of the requested URL exactly matches the specified pattern.

Table 19.4: Forwarding Conditions, Properties, Actions, and Definitions (Continued)

Forwarding	Description
<code>server_url.address=</code>	Tests if the host IP address of the requested URL matches the specified IP address, IP subnet, or subnet definition.
<code>server_url.domain[.case_sensitive][.no_lookup]=</code>	Tests if the requested URL, including the domain-suffix portion, matches the specified pattern.
<code>server_url.extension[.case_sensitive]=</code>	Tests if the filename extension at the end of the path matches the specified string.
<code>server_url.host.has_name=</code>	Tests whether the server URL has a resolved DNS hostname.
<code>server_url.host[.exact .substring .prefix .suffix .regex][.no_lookup]=</code>	Tests if the host component of the requested URL matches the IP address or domain name.
<code>server_url.host.is_numeric=</code>	This is true if the URL host was specified as an IP address.
<code>server_url.host.no_name=</code>	This is true if no domain name can be found for the URL host.
<code>server_url.host.regex=</code>	Tests if the specified regular expression matches a substring of the domain name component of the requested URL.
<code>server_url.is_absolute=</code>	Tests whether the server URL is expressed in absolute form.
<code>server_url.path[.exact .substring .prefix .suffix .regex][.case_sensitive]=</code>	Tests if a prefix of the complete path component of the requested URL, as well as any query component, matches the specified string.
<code>server_url.path.regex=</code>	Tests if the regex matches a substring of the path component of the request URL.
<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 <code><Exception></code> and <code><Proxy></code> layers.

Table 19.4: Forwarding Conditions, Properties, Actions, and Definitions (Continued)

Forwarding	Description
<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. Note that this setting is not definitive. It depends on the capabilities of the selected forwarding host.
<code>forward()</code>	Determines forwarding behavior. Note that 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 terminates or continues to process the request if the specified forwarding host or any designated backup or default cannot be contacted.
<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 <code><Cache></code> 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 <code><Proxy></code> layers.
<code>icp()</code>	Determines when to consult ICP. The default is <code>yes</code> if ICP hosts are configured and if no forwarding host or SOCKS gateway is identified as an upstream target.

Table 19.4: Forwarding Conditions, Properties, Actions, and Definitions (Continued)

Forwarding	Description
<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. Note that 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. Note that 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:8081/Policy/Trace/path</code>
Actions	
<code>notify_email()</code>	Sends an email notification to the list of recipients specified in the Event Log mail configuration. Can be used in all layers.

Table 19.4: Forwarding Conditions, Properties, Actions, and Definitions (Continued)

Forwarding	Description
notify_snmp()	The SNMP trap is sent when the transaction terminates. Can be used in all layers.
log_message	Writes the specified string to the ProxySG event log.
Definitions	
define server_url.domain condition name	Binds a user-defined label to a set of domain suffix patterns for use in a condition= expression.

Chapter 20: 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.

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.

Overview

The ProxySG 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).

These log records may 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 prior to upload.

Data stored in log facilities can be automatically uploaded to a remote location for analysis and archive purposes. The uploads may take place using HTTP, FTP, or one of several proprietary protocols. Once 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 User Guide*.

Understanding 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 in SGOS 4.x, 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 via policy.

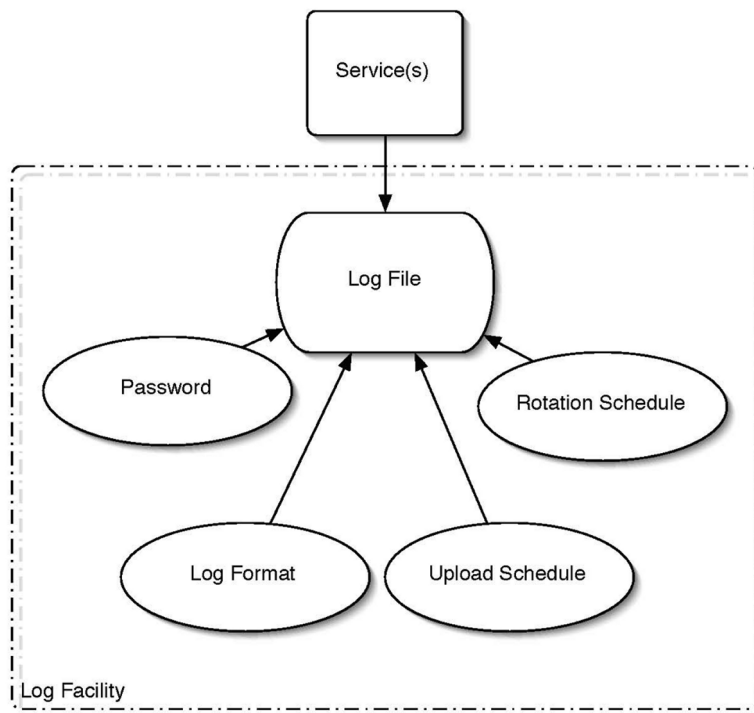


Figure 20-1: Log Facility

Understanding Protocols and Formats

The following protocols support configurable access logging:

- Endpoint Mapper
- FTP
- HTTP/HTTPS
- ICP
- Instant Messaging
- Peer-to-peer (P2P)
- RealMedia/QuickTime
- SOCKS
- TCP Tunnel
- Telnet Proxy
- Windows Media

The ProxySG 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
- SmartReporter, an ELFF log format compatible with SmartFilter's SmartReporter
- SurfControl, a log format compatible with the SurfControl Reporter tool
- Websense, a log format compatible with the Websense Reporter tool

The log facilities, each containing a single logical file and supporting a single log format, are managed by policy (created through VPM or CPL), which specifies the destination log format and log file.

Terms

- *Log Facility*: A separate log that contains a single logical file and supports a single log format. It also 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.
- *Encrypted Log*: A log is encrypted using an external certificate associated with a private key. Encrypted logs can only be decrypted by someone with access to the private key. The private key is not accessible to the ProxySG.
- *Log Format*: The type of log that is used: NCSA/Common, SQUID, ELFF, SurfControl, or Websense.

The proprietary log types each have a corresponding pre-defined log format that has been set up to produce exactly that type of log (these logs cannot be edited). In addition, a number of other ELFF type log formats are also pre-defined (im, main, p2p, streaming). These can be edited, but they start out with a useful set of log fields for logging particular protocols understood by the ProxySG. It is also possible to create new log formats of type ELFF or Custom which can contain any desired combination of log fields.

- *Log Tail*: The access log tail shows the log entries as they get logged. With high traffic on the ProxySG, not all access log entries are necessarily displayed. However, you can view all access log information after uploading the log.
- *NCSA common log format*: A log type that contains only basic HTTP access information.
- *SQUID-compatible format*: A log type that was designed for cache statistics.
- *ELFF-compatible format*: A log type defined by the W3C that is general enough to be used with any protocol.
- *SmartReporter*: A proprietary ELFF log type that is compatible with the SmartFilter SmartReporter tool.
- *SurfControl*: A proprietary log type that is compatible with the SurfControl reporter tool.
- *Websense*: A proprietary log type that is compatible with the Websense reporter tool.

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.

Once 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 through the Management Console:

1. Select Configuration>Access Logging>General>Default Logging.

The Default Logging tab displays.

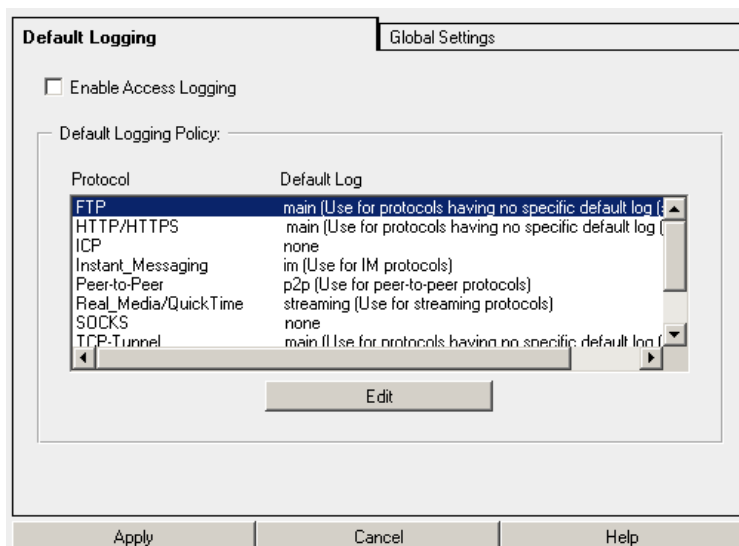


Figure 20-2: Enabling Access Logging

2. Select the Enable checkbox to enable access logging or deselect it to disable access logging.
3. Click Apply.

To Enable or Disable Access Logging through the CLI:

From the (config) command prompt, enter the following commands:

```
SGOS#(config) access-log
SGOS#(config access-log) enable | disable
```

Customizing the Log Procedures

You should first decide what protocols and log formats you want to use, the logging policy, and the upload schedule. Then you can do 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.

Customizing the Log: Creating and Editing Log Formats

The Format tab allows you to create a format to use for your log facilities. Several log formats ship with the ProxySG, and they might be sufficient for your needs. If so, you do not need to use the Format tab and can skip to "Customizing the Log: Creating an Access Log Facility" on page 750. If the formats that exist do not meet your needs, you can use the Format tab to create a custom or ELFF format and specify the string and other qualifiers used.

Several log formats already exist. For a description of each value, see Appendix B: "Access Log Formats" on page 877:

- **im (Instant Messaging):** This is an ELFF format with the custom strings of:


```
date time c-ip cs-username 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-path cs-uri-query cs-username s-hierarchy
s-supplier-name cs(Content-Type) cs(User-Agent) sc-filter-result
sc-filter-category x-virus-id s-ip s-sitename
```
- **nca:** 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-protocol x-p2p-client-type
x-p2p-client-info x-p2p-client-bytes x-p2p-peer-bytes duration s-action
```
- **smartreporter:** This is a reserved format that cannot be edited. It contains the following string:


```
localtime s-computername c-ip c-uri sc-filter-result cs-categories cs-user
sc-bytes
```
- **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.

- streaming:** This is an ELFF format with custom strings of:

```
c-ip date time c-dns cs-uri-stem 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 x-cache-info x-client-address
```
- surfcontrol, surfcontrolv5, and smartfilter:** These are reserved formats that cannot be edited.
- websense:** This is a reserved format that cannot be edited.

Note: If you had previously created formats with the name *smartreporter* or *surfcontrolv5* and you upgrade your ProxySG, those formats are changed to *smartreporter_user* or *surfcontrolv5_user*. If you already have a log format named *smartreporter_user* or "*surfcontrolv5_user*", then the names will be *smartreporter_user1* or *surfcontrolv5_user1*. This naming protocol continues (*_user2*, *_user3*...) as long as necessary. The logs associated with these formats are automatically associated with the new format name.

Creating a Custom or ELFF Log Format

If you are using one of the already-existing formats, skip to "Customizing the Log: Creating an Access Log Facility" on page 750. Complete the following steps to create a custom or ELFF log format.

To Create and Edit the Log Format through the Management Console:

1. Select Configuration>Access Logging>Formats.

The Formats tab displays, with the current log formats.

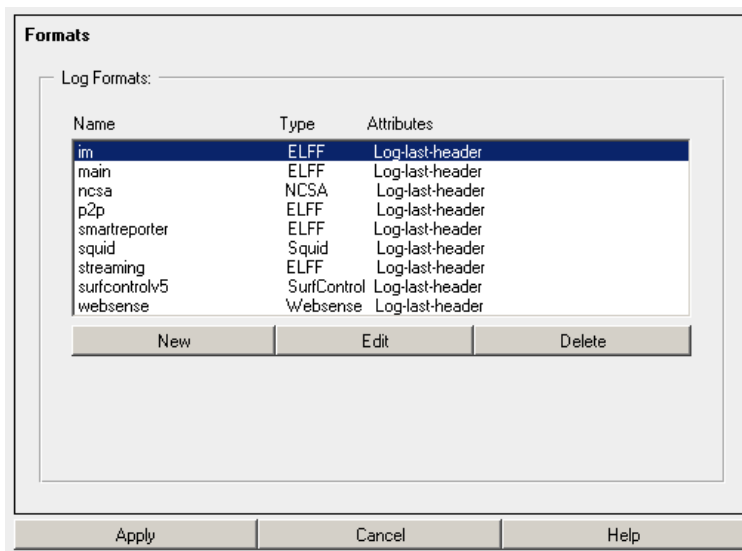


Figure 20-3: Formats Tab

- To create or edit a new custom or ELFF log format, click the New button; to edit an existing ELFF log format (such as im, main, p2p, or streaming), highlight the format to be changed and click the Edit button. If you choose an unconfigurable format, you will see an error message.

The Create/Edit Format dialog displays.

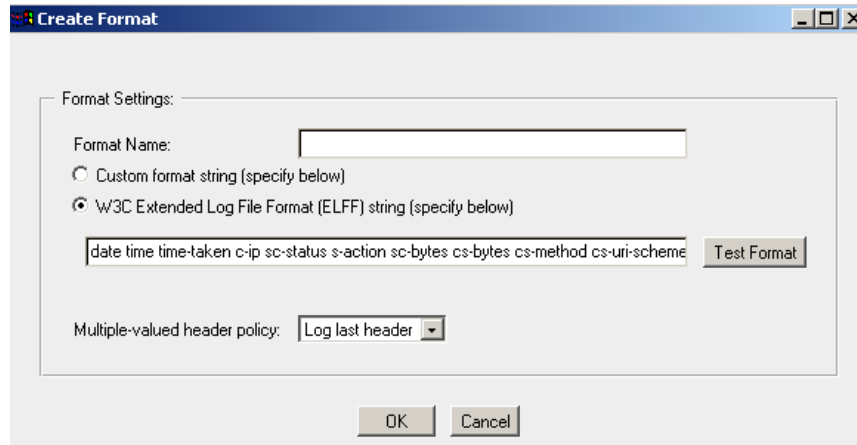


Figure 20-4: Create Format Dialog

- If you are creating a new format, provide a name meaningful to you.
- Use the Format Settings radio buttons to select the log format you need; specify the string in the field below.

Note: ELFF strings cannot start with spaces.

- Click the Test Format button to test whether the format-string syntax is correct.

When you click the Test Format button, a line displays below the field that indicates that testing is in progress and then gives a result, such as Format is valid.

Note: To doublecheck the format-string syntax, see "Creating a Custom or ELFF Log Format" on page 748 or Appendix B: "Access Log Formats" on page 877.

- From the Multiple-valued header policy drop-down list, select the header you want to log: Log last header, log first header, log all headers.

The Multiple valued header policy allows you to determine what will be done with HTTP-headers that have multiple headers.

- Click OK; click Apply.

To Create and Edit a Custom or ELFF Log Format through the CLI:

- To create a custom or ELFF log format name, enter the following commands from the (config) command prompt (skip to step 2 to edit an existing ELFF format log):

```
SGOS#(config) access-log
SGOS#(config access-log) create format format_name
```

- To edit a newly created or existing log format:

```
SGOS#(config access-log) edit format format_name
```

The prompt changes to:

```
SGOS#(config format format_name)
```

3. To customize the log format:

```
SGOS#(config format format_name) type {custom | elff} format_string  
SGOS#(config format format_name) multi-valued-header-policy {log-all-headers  
| log-first-header | log-last-header}
```

where

<code>type</code>	<code>{custom elff}</code> <code><i>format_string</i></code>	Specifies the log format.
<code>multi-valued-header-policy</code>	<code>log-all-headers </code> <code>log-first-header</code> <code> log-last-header</code>	(Optional) Specifies which headers should be logged. The default is log-last-header.

4. (Optional) View the results.

```
SGOS#(config format format_name) view  
Settings:  
Format name: format_name  
Type elff "date time time-taken c-ip sc-status s-action sc-bytes cs-bytes  
cs-method cs-uri-scheme cs-host cs-uri-path cs-uri-query cs-username  
s-hierarchy s-supplier-name rs(Content-Type) cs(User-Agent) sc-filter-result  
sc-filter-category x-virus-id s-ip s-sitename"  
Multiple-header-policy log-last-header
```

5. (Optional) To delete a log format:

```
SGOS#(config) access-log  
SGOS#(config access-log) delete format format_name
```

Boundary Condition: Creating a Custom or ELFF Log Format

The access log ignores any ELFF or custom format fields it doesn't 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.

Customizing the Log: Creating an Access Log Facility

You can use existing log facilities and modify them for your needs. You can also create new log facilities for special circumstances, such as associating the SurfControl log format with a log facility. To create new log facilities, continue with the next section. If you need to edit an existing log facility, skip to "Customizing the Log: Editing an Existing Log Facility" on page 752.

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 "Customizing the Log: Creating and Editing Log Formats" on page 747).

To Create a Log Facility through the Management Console:

1. Select Configuration>Access Logging>Logs>Logs.

The Logs tab displays.

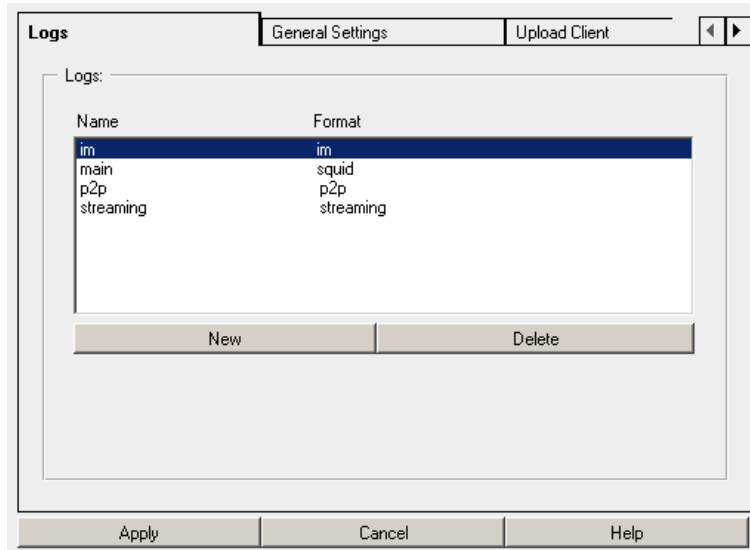


Figure 20-5: Logs Tab

2. The log facilities already created are displayed in the Logs tab. To create a new log, click New.

The Create Log dialog displays.

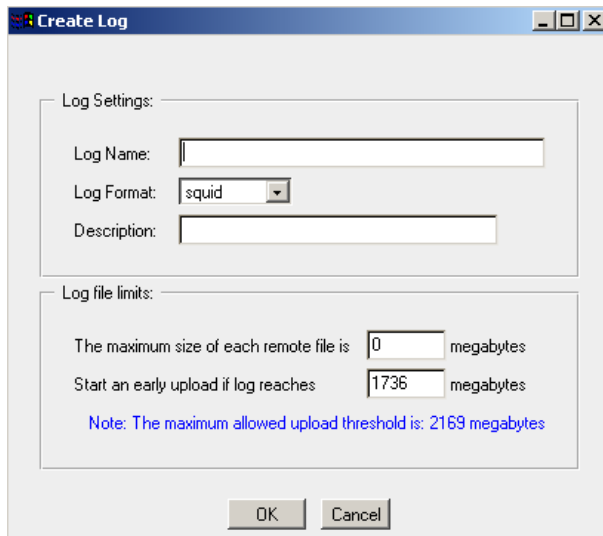


Figure 20-6: Create Log Dialog

3. Fill in the fields as appropriate:
 - Log Name: Enter a log facility name that is meaningful to you.
 - Log Format: Select a log format from the drop-down list.

- ❑ 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 "Customizing the Log: Editing an Existing Log Facility" below.)
 - ❑ 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.
 - ❑ Specify a size that will trigger an early upload—the maximum upload size varies depending on the size of the ProxySG disks (the maximum allowed upload threshold appears below this field).
 5. Click OK; click Apply.

To Create a Log Facility through the CLI:

From the (config) command prompt, enter the following commands:

```
SGOS#(config) access-log
SGOS#(config access-log) create log log_name
```

See "Customizing the Log: Editing an Existing Log Facility" below for information on configuring the newly created log.

Customizing the Log: Editing an Existing Log Facility

Four log facilities exist, each associated with a log format. For a description of the format, see "Customizing the Log: Creating and Editing Log Formats" on page 747.

- im (Instant Messaging): Associated with the im format.
- main: Associated with the main format.
- p2p (Peer-to-Peer): Associated with the p2p 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, keep in mind 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 through the Management Console:

1. Select Configuration>Access Logging>Logs>General Settings.
The General Settings tab displays.

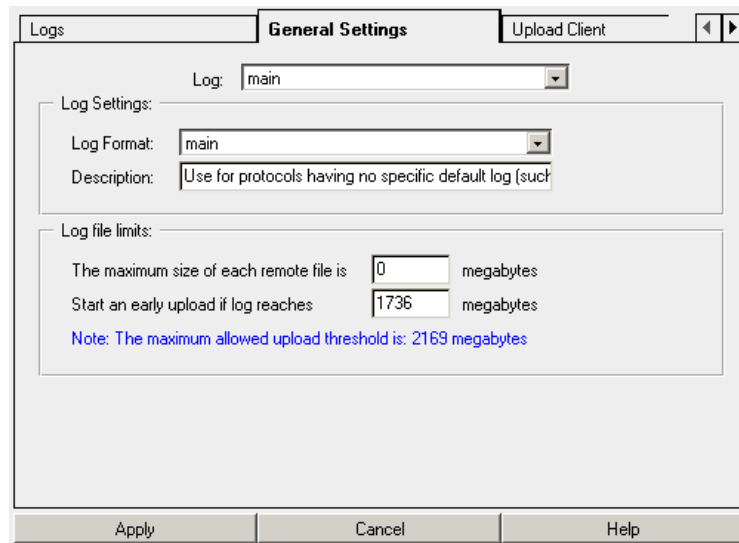


Figure 20-7: General Settings Tab

2. Fill in the fields as appropriate:

- Log: Select an already-existing log facility from the Log drop-down list.
- Log Format: Select the log format from the drop-down list.
- 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:

- 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.
- Specify a size that will trigger an early upload—the maximum upload size varies depending on the size of the ProxySG disks (the maximum allowed upload threshold appears below this field).

4. Click OK; click Apply.

To View an Existing Log Facility through the CLI:

A log facility must exist before you can edit it. You can view all the created log facilities with their configured settings through the CLI. The example below shows the settings for only one log facility. To view settings for a particular log facility only, include the optional `log_name` argument.

To view the existing log formats on the system, enter the following command:

```
SGOS#(config) show access-log log [log_name]
Settings:
Log name: main
Format name: main
Description: Use for protocols having no specific default log (such as im)
Logs uploaded using FTP client
```

```
Logs upload as gzip file
Wait 60 seconds between server connection attempts
Log encryption disabled
FTP client:
Filename format: SG_%f_%c_%l%m%d%H%M%S.log
Filename uses utc time
Use PASV: yes
Use secure connections: no
Primary host site:
Host:
Port: 21
Path:
Username:
Password: *****
Alternate host site:
Host:
Port: 21
Path:
Username:
Password: *****
HTTP client:
Filename format: SG_%f_%c_%l%m%d%H%M%S.log
Filename uses utc time
Use secure connections: no
Primary host site:
Host:
Port: 80
Path:
Username:
Password: *****
Alternate host site:
Host:
Port: 80
Path:
Username:
Password: *****
Custom client:
Primary server: :69
Alternate server: :69
Use secure connections: no
Websense client:
Primary server: :55805
Alternate server: :55805
Log uploading:
Log is uploaded daily at 02:00
No bandwidth class has been set for uploads
A keep-alive log packet is sent every 300 seconds
Start an early upload when log reaches 1736 megabytes
Remote log file rotation by size is disabled
```

To Edit an Existing Log Facility through the CLI:

Once you know which log facility you want to edit, complete the following procedure.

1. From the (config) command prompt, enter the following commands:

```

SGOS#(config) access-log
SGOS#(config access-log) edit log log_name
SGOS#(config log log_name) format-name format_name
SGOS#(config log log_name) early-upload megabytes
SGOS#(config log log_name) remote-size megabytes

```

where

<code>format-name</code>	<code><i>format_name</i></code>	Specifies a log format for this log facility. The format name can be any format that already exists on the ProxySG.
<code>early-upload</code>	<code><i>megabytes</i></code>	Specifies the size that will trigger an early upload—the maximum upload size varies depending on the size of the ProxySG disks.
<code>remote-size</code>	<code><i>megabytes</i></code>	Specifies the maximum size for each remote log file (the file on the upload server). The default is 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

2. (Optional) View the results.

```
SGOS#(config log log_name) view
```

Note: The output includes all the defaults for the log facility, whether or not you configured them.

3. (Optional) To delete a log facility :

```

SGOS#(config) access-log
SGOS#(config access-log) delete log log_name

```

Note: Deleting the log deletes any existing log entries on the ProxySG. To avoid this, upload the access log entries before deleting the logs.

Customizing the Log: Associating a Log Facility with a Protocol

You can associate a log facility with a protocol at any point in the process. New systems have certain protocols associated with certain logs by default. This allows you to begin access logging as soon as it is enabled (see "Enabling or Disabling Access Logging" on page 746).

Note: If you have a policy that defines protocol and log association, that policy will override any settings you make here.

The following list shows the protocols supported and the default log facilities assigned to them, if any:

- FTP: main
- HTTP/HTTPS: main

- ICP: none
- Instant Messaging: im
- Peer to Peer: p2p
- Real Media/QuickTime: streaming
- SOCKS: none
- TCP Tunneling: main
- Telnet: main
- Windows Media: streaming

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 757) or modify the access logging policy in VPM (see "Modify Access Logging" on page 507).

To Associate a Log Facility with a Protocol through the Management Console:

1. Select Configuration>Access Logging>General>Default Logging.

The Default Logging tab displays.

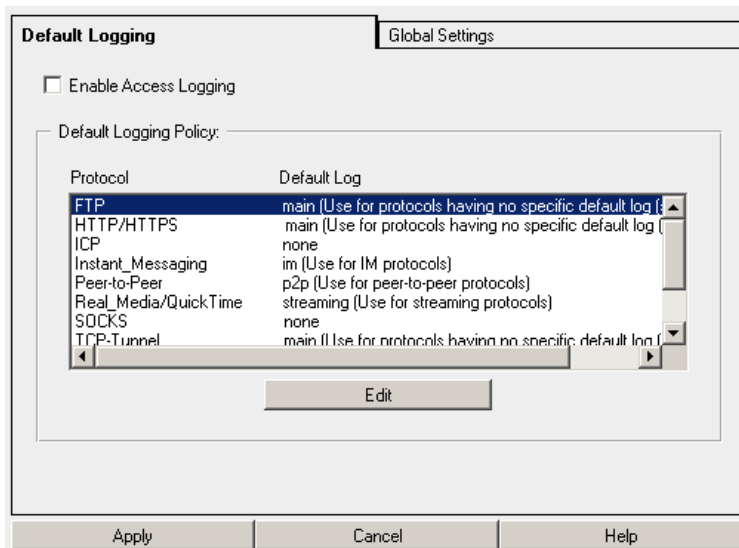


Figure 20-8: Default Logging Tab

2. Highlight the protocol you want to associate with a log facility and click Edit.

The appropriate Edit Logging dialog appears.

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.

5. Click Apply.

To Associate a Log Facility with a Protocol through the CLI:

1. From the (config) command prompt, enter the following commands:

```
SGOS#(config) access-log
SGOS#(config access-log) default-logging {icp | ftp | http | im | mms | p2p |
rtsp | socks | tcp-tunnel | telnet}} log_name
```

where:

<code>icp</code>	<code>log_name</code>	Sets the default log facility for ICP.
<code>ftp</code>	<code>log_name</code>	Sets the default log facility for FTP.
<code>http</code>	<code>log_name</code>	Sets the default log facility for HTTP/HTTPS.
<code>im</code>	<code>log_name</code>	Sets the default log facility for IM.
<code>mms</code>	<code>log_name</code>	Sets the default log facility for MMS.
<code>p2p</code>	<code>log_name</code>	Sets the default log facility for Peer-to-Peer.
<code>rtsp</code>	<code>log_name</code>	Sets the default log facility for Real Media/QuickTime.
<code>socks</code>	<code>log_name</code>	Sets the default log facility for SOCKS.
<code>tcp-tunnel</code>	<code>log_name</code>	Sets the default log facility for TCP tunneling.
<code>telnet</code>	<code>log_name</code>	Sets the default log facility for Telnet Proxy.

2. (Optional) View the results.

```
SGOS#(config access-log) view default-logging
Default Logging:
Protocol      Log
-----
http          main
ftp           main
socks         none
tcp tunnel    main
telnet        main
im            im
icp           none
mms           streaming
p2p           p2p
rtsp          streaming
```

Disabling Access Logging for a Particular Protocol

To Disable Access Logging for a Particular Protocol through the Management Console:

1. Select Configuration>Access Logging>General>Default Logging.

The Default Logging tab displays.

2. Highlight the protocol for which you want to disable access logging and click Edit.

The appropriate Edit Logging dialog appears.

3. Select none from the drop-down menu.
4. Click OK.
5. Click Apply.

To Disable Access Logging for a Particular Protocol through the CLI:

From the (config) command prompt, enter the following commands:

```
SGOS#(config) access-log
SGOS#(config access-log) no default-logging {icp | ftp | http | im | mms |
p2p | rtsp | socks | tcp-tunnel | telnet}
```

where access logging will be disabled for the protocol command you enter.

Customizing the Log Facility: Configuring Global Settings

You can set global limits for log size and early upload times. These settings can be overridden by individual log facilities.

To Set Global Log Facility Limits through the Management Console:

1. Select Configuration>Access Logging>General>Global Settings.

The Global Settings tab displays.

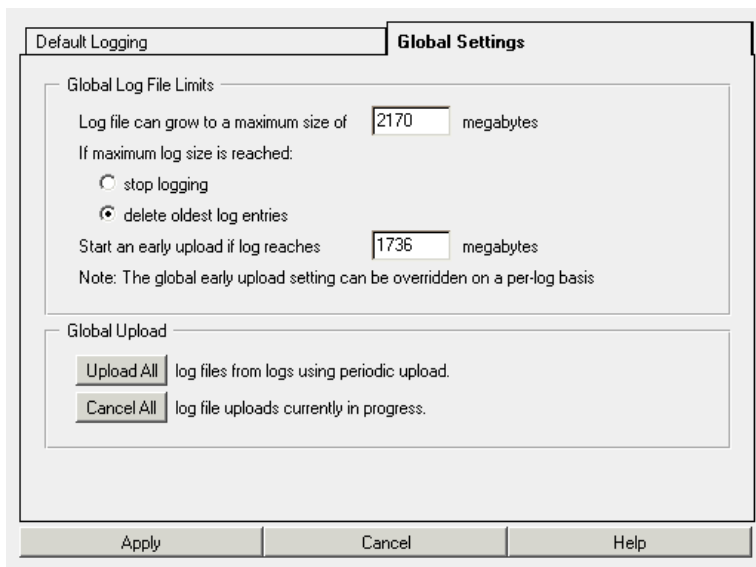


Figure 20-9: Global Settings Tab

2. Fill in the Global Log File Limits panel as appropriate:
 - Configure the maximum size occupied by all of the log files (in megabytes).
 - Determine the behavior of the log when the maximum size is reached. You can have the log stop logging (and do an immediate upload) or have it delete the oldest log entries.
 - Specify the size of the log that triggers an early upload.

3. The Global Upload options affect all log facilities currently available. They do not affect scheduled upload times. You can upload logs now, using the periodic upload method, or you can cancel all the uploads that are currently in progress.

To Set Global Log Facility Limits through the CLI:

From the `(config)` command prompt, enter the following commands:

```
SGOS#(config) access-log
SGOS#(config access-log) overflow-policy {delete | stop}
SGOS#(config access-log) early-upload megabytes
SGOS#(config access-log) upload {all | log log_name}
SGOS#(config access-log) cancel-upload {all | log log_name}
```

where

<code>overflow-policy</code>	<code>delete stop</code>	When the log reaches its maximum size, you can delete the oldest log entries, or you can stop logging (and do an immediate upload).
<code>early-upload</code>	<code>megabytes</code>	Specifies the size of the log before an upload can take place.
<code>upload</code>	<code>{all log log_name}</code>	An immediate upload for all logs or a specified log.
<code>cancel-upload</code>	<code>{all log log_name}</code>	Cancels the current upload for all logs or a specified log.

Customizing the Log Facility: Configuring the Upload Client

Blue Coat supports four types of upload client:

- FTP client, the default
- HTTP client
- Custom client
- Websense client

Blue Coat also supports secure FTP, HTTP, and Custom client.

The Custom client can be used for special circumstances, such as working with SurfControl Reporter. Custom client is based on plain sockets.

Note: You must have a socket server to use the Custom client.

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, Custom client, or Websense client panes or the `access-log ftp-client`, `https-client`, `custom-client`, or `websense-client` CLI commands.

Keep in mind that only one client can be used at any one time. All four can be configured, but only the selected client is used.

The ProxySG provides access logging with two types of uploads to a remote server:

- continuous uploading, where the ProxySG continuously streams new access log entries from the ProxySG memory to a remote server
- scheduled (periodic) uploading, where the ProxySG transmits log entries on a scheduled basis. See "Customizing the Log: Configuring the Upload Schedule" on page 776 for more information.

The ProxySG allows you to upload either compressed access logs or plain-text access logs. The ProxySG uses the gzip format to compress access logs. Gzip-compressed files allow more log entries to be stored in the ProxySG. 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 ProxySG 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`.

Note: You cannot upload gzip access-log files for the Websense client.

For greater security, you can configure the ProxySG 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 219). The ProxySG 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 765.

Note: The encryption feature is not available for custom or Websense clients.

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.

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 a private key and a corresponding x.509 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`.

Note: The signing feature is not available for custom or Websense clients.

For information about verifying a log, see "Verifying a Digital Signature" on page 765.

To Configure the Upload Client through the Management Console:

1. Select Configuration>Access Logging>Logs>Upload Client.

The Upload Client tab displays.

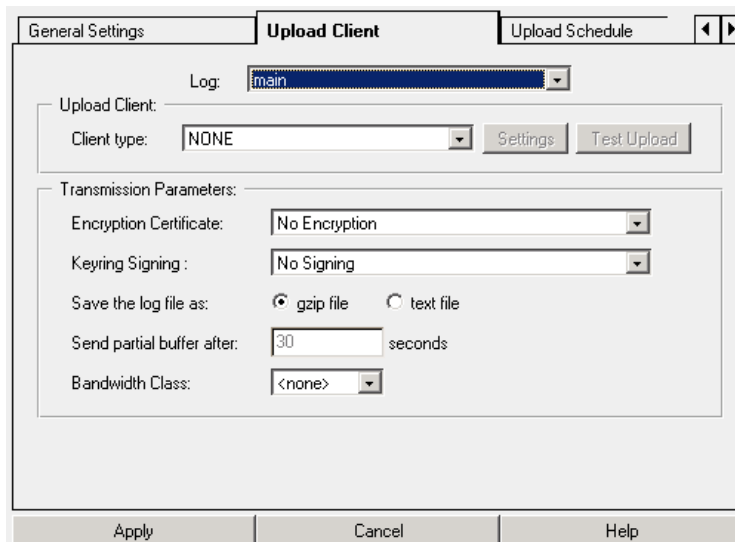


Figure 20-10: Upload Client Tab

2. From the Log drop-down list, select the log facility you want to configure. The facility must exist before it displays in this list.
3. From the Client type drop-down list, select the upload client to use. Only one client can be configured for each log facility.
4. Click the Settings button to customize the upload client.

For information on customizing the clients, skip to "Editing the FTP Client" on page 766, "Editing the HTTP Client" on page 769, "Editing the Custom Client" on page 773, "Editing the Custom SurfControl Client" on page 774, or "Editing the Websense Client" on page 775.

For information about testing the upload client, see "Testing Access Log Uploading" on page 780.

5. (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 (see "Importing an External Certificate" on page 219).

The encryption option is not available for Websense or Custom clients.

6. (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 Websense or Custom clients.

7. Select one of the Save the log file as radio buttons to determine whether the access log that is uploaded will be compressed (gzip file, the default) or not (text file).

Note: If you are configuring a SurfControl Custom client, select the text file radio button.

8. If you chose 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 will force 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 chose gzip file, the Send partial buffer after *n* seconds field is not configurable. Also, this setting is only valid for continuous uploading (see "Customizing the Log: Configuring the Upload Schedule" on page 776 for information about continuous uploading).

9. (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 enable bandwidth management and create a bandwidth-management class. Bandwidth management is enabled by default if you have a valid license for this feature. Note that it is the log facility that is bandwidth-managed—the upload client type does not affect this setting. See Chapter 10: "Bandwidth Management" on page 375 for information about enabling bandwidth management and creating and configuring the bandwidth class.

Remember that less bandwidth slows down the upload, while more could flood the network.

10. Click Apply.

To Configure the Upload Client through the CLI:

From the (config) command prompt, enter the following commands to make general settings for the upload client.

```
SGOS#(config) access-log
SGOS#(config access-log) edit log log_name
SGOS#(config log log_name) client-type {custom | ftp | http | websense}
SGOS#(config log log_name) upload-type {gzip | text}
SGOS#(config log log_name) bandwidth-class class_name
SGOS#(config log log_name) encryption certificate certificate_name
SGOS#(config log log_name) signing keyring_id
SGOS#(config log log_name) ftp-client | http-client | custom-client |
websense-client
```

where

<code>client-type</code>	<code>custom</code> <code>ftp</code> <code>http</code> <code>websense</code>	Specifies which upload client to use. Only one client can be configured for each log.
<code>upload-type</code>	<code>gzip</code> <code>text</code>	Specifies upload as a gzip or a text file. Websense client always uploads a text file.

<code>bandwidth-class</code>	<code>class_name</code>	Specifies a bandwidth-management class for managing the bandwidth of this log. IMPORTANT: Before you can manage the bandwidth for this log, you must first enable bandwidth management and create a bandwidth-management class. Bandwidth management is enabled by default if you have a valid license for this feature. See Chapter 10: "Bandwidth Management" on page 375 for information about enabling bandwidth management and creating and configuring bandwidth classes.
<code>no</code>	<code>bandwidth-class</code>	Disables bandwidth management for this log.
<code>encryption</code>	<code>certificate</code> <code>certificate_name</code>	Specifies the access log encryption certificate. Cannot be used for Websense or Custom clients.
<code>no</code>	<code>encryption</code>	Disables access log encryption.
<code>signing</code>	<code>keyring_id</code>	Specifies the keyring to be used for digital signatures.
<code>no</code>	<code>signing</code>	(Default) Disables access log digital signature.
<code>ftp-client</code>		Edits the FTP client configuration. Skip to "Editing the FTP Client" on page 766 for more information.
<code>http-client</code>		Edits the HTTP client configuration. Skip to "Editing the HTTP Client" on page 769 for more information.
<code>custom-client</code>		Edits the Custom client configuration. Skip to "Editing the Custom Client" on page 773 for more information.
<code>websense-client</code>		Edits the Websense client configuration. Skip to "Editing the Websense Client" on page 775 for more information.

Disabling Log Uploads

To disable log uploads, set the upload client-type to none.

To Disable an Upload through the Management Console:

1. Select Configuration>Access Logging>Logs>Upload Client.

The Upload Client tab displays.

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.

To Disable an Upload through the CLI:

From the (config) command prompt, enter the following commands:

```
SGOS#(config) access-log
SGOS#(config access-log) edit log log_name
SGOS#(config log log_name) client-type none
```

where *log_name* is the name of the log for which you want to disable an upload.

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 -recip
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" 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, Custom, and Websense. Each of these clients are described below. You can also create a SurfControl or SmartFilter 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 through the Management Console:

1. Select Configuration>Access Logging>Logs>Upload Client.

The Upload Client tab displays. See "Customizing the Log Facility: Configuring the Upload Client" on page 759 for configuration information.

2. Select FTP Client from the Client type drop-down list. Click the Settings button.

The FTP Client Settings dialog displays.

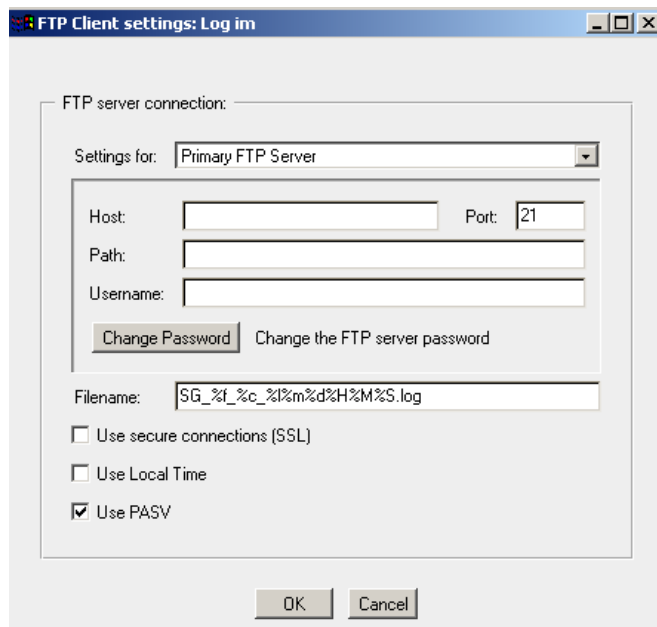


Figure 20-11: Edit FTP Client Dialog

3. Select the primary or alternate FTP server you want to configure from the Settings for drop-down list.
4. Fill in the fields as appropriate:
 - Host: The name of the upload client host. If the Use secure connections (SSL) checkbox is selected, the hostname must match the hostname in the certificate presented by the server.
 - Port: The default is 21; it can be changed.
 - Path: The directory path where the access log will be uploaded on the server.
 - Username: This is the username that is known on the host you are configuring.
 - Change Password: Change the password on the FTP host by clicking this button: the Change Password dialog displays; enter and confirm the new password; click OK.

- ❑ **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 will overwrite the old file. Also, if you use more than one external certificate to encrypt logs, you should include the %c specifier in the Filename field to keep track of which external certificate was used to encrypt the uploaded log file.

If you are creating a SurfControl client, change the `.log` file extension to `.tmp`.

- ❑ **Secure Connections:** If you use FTPS, select the Use secure connections (SSL) checkbox. The remote FTP server must support FTPS.
- ❑ **Local Time:** If you want the upload to reflect the local time it was uploaded instead of Universal Time Coordinates (UTC), select the Local Time checkbox.
- ❑ **Use PASV:** With the Use PASV checkbox (the default) selected, the ProxySG connects to the FTP server. With the Use PASV checkbox de-selected, the FTP server uses the PORT command to connect to the ProxySG.

5. Click OK; click Apply.

To Edit the FTP Client through the CLI:

1. At the `(config)` command prompt, configure the FTP client's primary or secondary server information:

```
SGOS#(config) access-log
SGOS#(config access-log) edit log log_name
SGOS#(config log log_name) ftp-client primary host hostname [port]
SGOS#(config log log_name) ftp-client no primary
SGOS#(config log log_name) ftp-client primary path path
SGOS#(config log log_name) ftp-client primary username user_name
SGOS#(config log log_name) ftp-client primary password password
-or-
SGOS#(config log log_name) ftp-client primary encrypted-password
encrypted_password
```

where

<code>primary host</code>	<code><i>hostname</i></code> <code>[<i>port</i>]</code>	Specifies the primary FTP server to which logs should be uploaded. By default, the ProxySG uses port 21.
<code>no primary</code>		Deletes the primary FTP host site.
<code>primary path</code>	<code><i>path</i></code>	The path is the directory on the primary FTP server to which logs should be uploaded.
<code>primary username</code>	<code><i>user_name</i></code>	Specifies the username on the primary FTP server to which logs should be uploaded. The <code><i>user_name</i></code> must have write privileges in the access log file upload directory.

primary	<i>password</i>	Specifies the password for the <i>user_name</i> in the previous command. Note that the primary use of the encrypted-password command is to allow the ProxySG to load a password that it encrypted.
password	<i>encrypted_</i>	
-or-	<i>password</i>	
primary		
encrypted-		
password		

2. (Optional) Repeat these steps for the secondary server, replacing primary with alternate.

```
SGOS#(config log log_name) ftp-client alternate host hostname [port]
SGOS#(config log log_name) ftp-client no alternate
SGOS#(config log log_name) ftp-client alternate path path
SGOS#(config log log_name) ftp-client alternate username user_name
SGOS#(config log log_name) ftp-client alternate password password
-or-
SGOS#(config log log_name) ftp-client alternate encrypted-password
encrypted_password
```

3. Enter the following commands to complete configuration of the FTP client.

```
SGOS#(config log log_name) ftp-client filename format
-or-
SGOS#(config log log_name) ftp-client no filename
SGOS#(config log log_name) ftp-client pasv {yes | no}
SGOS#(config log log_name) ftp-client secure {yes | no}
SGOS#(config log log_name) ftp-client time-format {local | utc}
```

where

<i>filename</i>	<i>format</i>	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. Be cautious if you change the Filename field. If an ongoing series of access log files are produced and you do not have a time-specifier in this field, each access log file produced will overwrite the old file. Also, if you use more than one external certificate to encrypt logs, you should include the %c specifier in the Filename field to keep track of which external certificate can decrypt the uploaded file. If you are creating a SurfControl client, you must change the .log file extension to .tmp.
no filename		Deletes the FTP client configuration parameters.
pasv	yes no	Specifies whether the ProxySG connects to the FTP server or if the FTP server connects to the ProxySG. The default is yes, using the PORT command only on failure.
secure	yes no	Specifies whether FTPS is used. The default is no. If yes, the <i>hostname</i> in Step 2 must match the hostname in the certificate presented by the server.

```
time-format local | utc
```

Specifies whether Universal Time Coordinates (UTC) or the local time is used. UTC is the default. UTC was formerly known as Greenwich Mean Time (GMT).

4. (Optional) View the results.

```
SGOS#(config log log_name) view
```

Tip: Doing a Manual Upload for FTP Upload Client through the CLI

Sometimes, an FTP connection is established with the FTP server and is left open. If you try to use the `upload-now` command while the connection is still open, the command fails with the error message:

```
User upload request failed. There is an open-connection. Try closing the connection.
```

To Close the Connection:

```
SGOS#(config access-log) edit log log_name
SGOS#(config log log_name) commands close-connection
ok
```

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. (You create an HTTPS client by checking the Use secure connections (SSL) checkbox.)

Note: To create an HTTPS client, you must also import the appropriate CA Certificate. For information, see "Importing a CA Certificate" on page 225.

To Edit the HTTP Client through the Management Console:

1. Select Configuration>Access Logging>Logs>Upload Client.

The Upload Client tab displays. See "Customizing the Log Facility: Configuring the Upload Client" on page 759 for configuration information.

2. Select HTTP Client from the Client type drop-down list. Click the Settings button.

The HTTP Client Settings dialog displays.

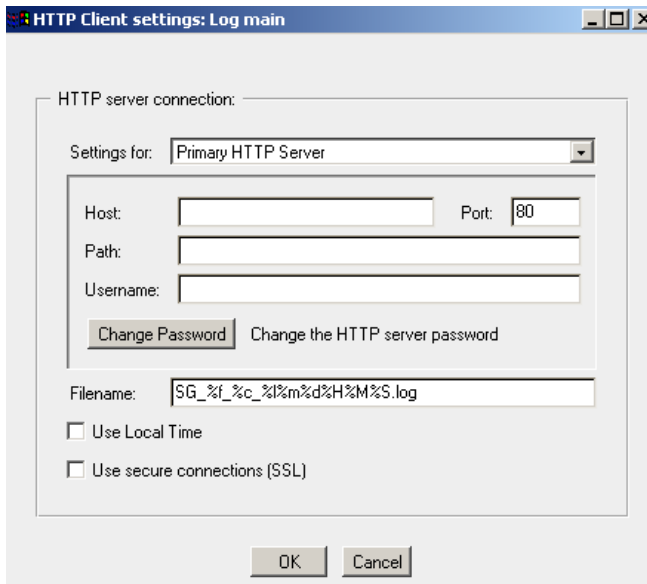


Figure 20-12: Edit HTTP Client Dialog

3. From the Settings for drop-down list, select the primary or alternate HTTP server you want to configure.
4. Fill in the fields as appropriate:
 - Host: The name of the upload host. If the Use secure connections (SSL) checkbox is selected, the hostname must match the hostname in the certificate presented by the server.
 - Port: The default is 80, but you can change it.

Note: For HTTPS, change the port to 443.

- Path: The directory path where the access log facility will be uploaded on the server.
- Username: This is the username that is known on the host you are configuring.
- Change Password: Change the password on the HTTP host by clicking this button: the Change Password dialog displays; enter and confirm the new password and click OK.
- 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 will overwrite the old file. Also, if you use more than one external certificate to encrypt logs you should include the %c specifier in the Filename field to keep track of which external certificate can decrypt the uploaded log file.

If you are creating a SurfControl client, change the .log file extension to .tmp.

- ❑ Local Time: If you want the upload to reflect the local time it was uploaded instead of Universal Time Coordinate (UTC), select the Local Time checkbox.
- ❑ Use secure connections (SSL): Select this checkbox to create an HTTPS client. To create an HTTPS client, you must also create a keypair, import or create a certificate, and, if necessary, associate the keypair and certificate (called a keyring), with the SSL-client.

5. Click OK; click Apply.

To Edit the HTTP Client through the CLI:

1. At the `(config)` command prompt, configure the HTTP client's primary or secondary server information:

```
SGOS#(config) access-log
SGOS#(config access-log) edit log log_name
SGOS#(config log log_name) http-client primary host hostname [port]
SGOS#(config log log_name) http-client no primary
SGOS#(config log log_name) http-client primary path path
SGOS#(config log log_name) http-client primary username user_name
SGOS#(config log log_name) http-client primary password password
-or-
SGOS#(config log log_name) http-client primary encrypted-password
encrypted_password
```

where

primary host	<i>hostname</i> [<i>port</i>]	Specifies the primary HTTP server to which logs should be uploaded. By default, the ProxySG uses port 80. For HTTPS, change the port to 443.
no primary		Deletes the primary HTTP host site.
primary path	<i>path</i>	The path is the directory on the primary HTTP server to which logs should be uploaded.
primary username	<i>user_name</i>	Specifies the username on the primary HTTP server to which logs should be uploaded. The <i>user_name</i> must have write privileges in the access log file upload directory.
primary password	<i>password</i>	Specifies the password (or encrypted password) for the <i>user_name</i> in the previous command. Note that the primary use of the <code>encrypted-password</code> command is to allow the ProxySG to load a password that it encrypted.
-or-	<i>encrypted_</i>	
primary encrypted-password	<i>password</i>	

2. Repeat these steps for the secondary server, replacing primary with alternate

```
SGOS#(config access-log) edit log log_name
SGOS#(config log log_name) http-client alternate host hostname [port]
SGOS#(config log log_name) http-client no alternate
SGOS#(config log log_name) http-client alternate path path
SGOS#(config log log_name) http-client alternate username user_name
SGOS#(config log log_name) http-client alternate password password
-or-
SGOS#(config log log_name) http-client alternate encrypted-password
encrypted_password
```

3. (Optional) To stop the log from being uploaded to a primary or secondary server in the future, clear the hostname by entering an empty string (that is., double-quotes) in the following command:

```
SGOS#(config log log_name) http-client primary ""
-or-
SGOS#(config log log_name) http-client alternate ""
```

4. Enter the following commands to complete configuration of the HTTP client.

```
SGOS#(config log log_name) http-client secure {no | yes}
SGOS#(config log log_name) http-client filename log_name
SGOS#(config log log_name) http-client no filename
SGOS#(config log log_name) http-client time-format {utc | local}
```

where

<code>secure</code>	<code>no yes</code>	Specifies if you want to use SSL connections. The default is <code>no</code> . If <code>yes</code> , the <code>hostname</code> in Step 2 must match the hostname in the certificate presented by the server.
<code>filename</code>	<code><i>log_name</i></code>	The Filename field is comprised of text and/or specifiers. The default filename includes specifiers and text that indicate the log name (<code>%f</code>), name of the external certificate used for encryption, if any (<code>%c</code>), the fourth parameter of the ProxySG IP address (<code>%l</code>), the date and time (Month: <code>%m</code> , Day: <code>%d</code> , Hour: <code>%H</code> , Minute: <code>%M</code> , Second: <code>%S</code>), and the <code>.log</code> or <code>.gzip.log</code> file extension. Be cautious if you change the Filename field. If an ongoing series of access log files are produced and you do not have a time-specifier in this field, each access log file produced will overwrite the old file. Also, if you use more than one external certificate to encrypt logs, you should include the <code>%c</code> specifier in the Filename field to keep track of which external certificate can decrypt the uploaded log file. If you are creating a SurfControl client, you must change the <code>.log</code> file extension to <code>.tmp</code> .
<code>no filename</code>		Deletes the HTTP client configuration parameters.
<code>time-format</code>	<code>utc local</code>	Specifies whether Universal Time Coordinates (UTC) or the local time is used. UTC is the default.

5. (Optional) View the results.

```
SGOS#(config log log_name) view
```


Editing the Custom Client

To Edit the Custom Client through the Management Console:

1. Select Configuration>Access Logging>Logs>Upload Client.

The Upload Client tab displays. See "Customizing the Log Facility: Configuring the Upload Client" on page 759 for configuration information.

2. Select Custom Client from the Client type drop-down list. Click the Settings button.

The Custom Client Settings dialog displays.

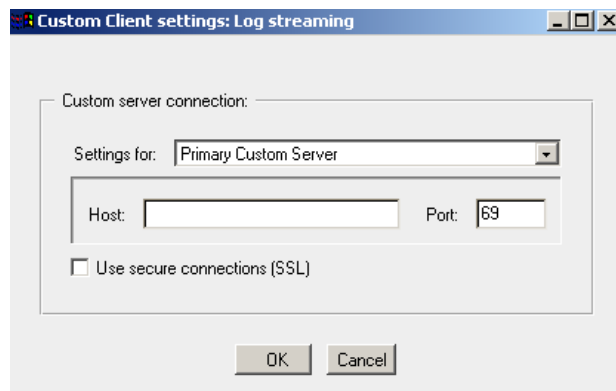


Figure 20-13: Edit Custom Client Dialog

3. From the Settings for drop-down list, select the primary or alternate custom server you want to configure.
4. Fill in the fields as appropriate:
 - Host: Enter the hostname of the upload destination. If the Use secure connections (SSL) checkbox is selected, the hostname must match the hostname in the certificate presented by the server.
 - Port: The default is 69; it can be changed.
 - Use secure connections (SSL): Select this checkbox if you are using secure connections.
5. Click OK; click Apply.

To Edit the Custom Client through the CLI:

1. At the `(config)` command prompt, configure the Custom client's primary or secondary server information:

```
SGOS#(config) access-log
SGOS#(config access-log) edit log log_name
SGOS#(config log log_name) custom-client primary hostname [port]
-or-
SGOS#(config log log_name) custom-client alternate hostname [port]
```

where *hostname* specifies the primary or alternate server to which logs should be uploaded. By default the ProxySG uses port 69.

2. Enter the following command to complete configuration of the Custom client:

```
SGOS#(config log log_name) custom-client secure {no | yes}
```

which specifies whether SSL connections are used. The default is `no`. If `yes`, the hostname in Step 1 must match the hostname in the certificate presented by the server.

3. (Optional) To stop the log from being uploaded to a primary or secondary server in the future, clear the hostname by entering an empty string (that is, double-quotes) in the following command:

```
SGOS#(config log log_name) custom-client primary ""
```

-or-

```
SGOS#(config log log_name) custom-client alternate ""
```

4. (Optional) View the results.

```
SGOS#(config log log_name) view
```

Editing the Custom SurfControl Client

You can use the Custom Client to create an upload client that will upload information to SurfControl Reporter. Before you begin, be sure that:

- You have created a log (see "Customizing the Log: Creating an Access Log Facility" on page 750).
- You have associated the SurfControl log format with the log you created (see "Customizing the Log: Editing an Existing Log Facility" on page 752).

To Edit the SurfControl Client through the Management Console:

1. Select Configuration>Access Logging>Logs>Upload Client.
The Upload Client tab displays. See "Customizing the Log Facility: Configuring the Upload Client" on page 759 for configuration information.
2. From the Log drop-down list, select the SurfControl log that you associated with the SurfControl log format.
3. Make sure 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 773.

5. Click the Settings button for that client.
6. Customize the upload client for SurfControl Reporter.
 - Enter the hostname, path, and username, if necessary, for the SurfControl Reporter server.
 - Make sure the filename extension is `.tmp` and not `.gzip` or `.log`. SurfControl only recognizes files with a `.tmp` extension.
 - If your SurfControl server supports SSL, select the Use secure connections (SSL) checkbox.
7. Click OK; click Apply.

To Edit the Custom SurfControl Client through the CLI:

1. At the `(config)` command prompt, configure the customized settings for the SurfControl upload client:

```
SGOS#(config) access-log
SGOS#(config access-log) edit log log_name
SGOS#(config access-log log_name) upload-type text
SGOS#(config access-log log_name) periodic-upload upload-interval {daily 0-23
| hourly hours [minutes]}
SGOS#(config access-log log_name) periodic-upload enable
SGOS#(config access-log log_name) custom-client
```

For specific information on managing upload clients, see "Editing the Custom Client" on page 773.

2. (Optional) View the results.

```
SGOS#(config log log_name) view
```

Editing the Websense Client

Before you begin, make sure you have created a Websense log using the Websense log format and configured the log to your environment. See "Customizing the Log: Creating an Access Log Facility" on page 750.

Note: You cannot upload gzip access log files with the Websense client.

To Edit the Websense Client through the Management Console:

1. Select Configuration>Access Logging>Logs>Upload Client.

The Upload Client tab displays. See "Customizing the Log Facility: Configuring the Upload Client" on page 759 for configuration information.

2. Select the Websense Client from the Client type drop-down list. Click the Settings button.

The Websense Client Settings dialog displays.




Figure 20-14: Edit Websense Client Dialog

3. From the Settings for drop-down list, select the primary or alternate server you want to configure.
4. Fill in the fields as appropriate:
 - ❑ Host: Enter the hostname of the primary Websense Server.

- Port: The default is 55805, but you can change it if the Websense Server is using a different port.
5. Repeat for the Alternate Websense Server.
 6. Click OK; click Apply.

To Edit the Websense Client through the CLI:

1. At the `(config)` command prompt, configure the Websense client's primary or secondary server information:

```
SGOS#(config) access-log
SGOS#(config access-log) edit log log_name
SGOS#(config log log_name) websense-client primary hostname [:port]
-or-
SGOS#(config log log_name) websense-client alternate hostname [:port]
```

where *hostname* specifies the primary or alternate server to which logs should be uploaded. By default the ProxySG uses port, which is 55805 by default.

2. (Optional) To stop the log from being uploaded to a primary or secondary server in the future, clear the hostname by entering an empty string (that is, double-quotes) in the following command:

```
SGOS#(config log log_name) websense-client primary ""
-or-
SGOS#(config log log_name) websense-client alternate ""
```

3. (Optional) View the results.

```
SGOS#(config log log_name) view
```

Customizing the Log: Configuring the Upload Schedule

The Upload Schedule allows you to configure 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.

You can specify either *periodic uploading* or *continuous uploading*. Both periodic and continuous uploading can send log information from a ProxySG farm to a single log analysis tool. This allows you to treat multiple ProxySG 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 ProxySG 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 ProxySG memory to a remote server. Note that *streaming* here refers to the real-time transmission of access log information. The ProxySG 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 ProxySG performance by sending log information to a remote server (avoiding disk writes), and save ProxySG disk space by saving log information on the remote server.

If the remote server is unavailable to receive continuous upload log entries, the ProxySG saves the log information on the ProxySG 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 through the Management Console:

1. Select Configuration>Access Logging>Logs>Upload Schedule.

The Upload Schedule page displays.

The screenshot shows the 'Upload Schedule' configuration window. At the top, there are two tabs: 'Upload Client' and 'Upload Schedule'. The 'Upload Schedule' tab is active. Below the tabs, there is a 'Log Facility' dropdown menu set to 'main'. Under the 'Upload type:' section, there are two radio buttons: 'continuously' (unselected) and 'periodically' (selected). Below these are two input fields: 'Wait between connect attempts:' with a value of '60' and 'seconds', and 'Time between keep-alive log packets:' with a value of '300' and 'seconds'. Under the 'Upload the access log:' section, there are two radio buttons: 'Daily at' (selected) and 'Every'. The 'Daily at' option has a dropdown menu showing '2:00 a.m.' and a button labeled 'Upload Now'. The 'Every' option has input fields for '8' hours and '0' minutes, and a button labeled 'Cancel Upload'. At the bottom of the window, there are three buttons: 'Apply', 'Cancel', and 'Help'.

Figure 20-15: Upload Schedule Tab

2. From the Log drop-down list, choose the log whose schedule you are configuring.
3. Select an upload method by choosing the continuously or the periodically radio button; click Apply.
4. To change the time between connection attempts, enter the new time (in seconds) in the Wait between connect attempts field.
5. (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.

Keep-alives maintain the connection during low periods of system usage. When no logging information is being uploaded, the ProxySG 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.

6. (Optional) From the Daily at drop-down list, specify the time of day you want the access log updated or rotated (if you are doing continuous uploads).
7. (Optional) If you don't want the log uploaded or rotated on a daily basis, select the Every radio button and enter the time between uploads.

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 ProxySG periodically creates a new log file, and archives the older one without disturbing the current log file.

8. (Optional) 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 will not affect the next scheduled upload time.

The Cancel upload button (for periodic uploads) allows you to stop repeated upload attempts if the Web server becomes unreachable while an upload is in progress. Clicking this button 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 will take time for it to take effect.

9. Click OK; click Apply.

To Configure an Upload Schedule through the CLI:

1. From the (config) command prompt, enter the following commands.

```
SGOS#(config access-log) edit log log_name
SGOS#(config log log_name) upload-type {gzip | text}
```

2. Configure either a continuous upload schedule or a periodic upload schedule by using the options in the continuous-upload or periodic-upload commands.

Note: If you are configuring a SurfControl upload client you must use periodic-upload, not continuous-upload. If you configure a Websense upload client, you should set it to continuous-upload.

```
SGOS#(config log log_name) continuous-upload {enable | keep-alive seconds |
lag-time seconds | rotate-remote {daily 0-23 | hourly hours [minutes]}}
-or-
SGOS#(config log name) periodic-upload {enable | upload-interval {daily 0-23
| hourly hours [minutes]}}
```

where

upload-type	gzip text	Specifies using a compressed file (gzip) or a text file for uploading.
continuous-upload	enable	Specifies continuous upload (automatically disables periodic upload).

<code>keep-alive <i>seconds</i></code>	Specifies the interval between keep-alive log packets. Acceptable values are between 0 and 65535 seconds.
<code>lag-time <i>seconds</i></code>	Specifies the maximum time between log packets (text upload only). Acceptable values are between 0 and 65535 seconds. This setting configures the maximum time between text log packets, meaning that it will force 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.
<code>rotate-remote {daily 0-23 hourly <i>hours</i> [<i>minutes</i>]}</code>	Specifies when to rotate to new remote log file: enter the time of day for a <code>daily</code> rotation or enter how often to rotate (every <code>n</code> hours/minutes) for an <code>hourly</code> rotation. To rotate more than once an hour, enter 0 hours and specify <code>n</code> minutes.
<code>periodic-upload enable</code>	Specifies periodic upload (automatically disables continuous upload).
<code>upload-interval {daily 0-23 hourly <i>hours</i> [<i>minutes</i>]}</code>	Specifies when to upload a log file: enter the time of day for a <code>daily</code> upload or enter how often to upload (every <code>n</code> hours/minutes) for an <code>hourly</code> upload. To upload more than once an hour, enter 0 hours and specify <code>n</code> minutes.

3. Specify the time between connection attempts between the ProxySG and the remote server:

```
SGOS#(config log log_name) connect-wait-time seconds
```

4. (Optional) Use the following options to upload an access log immediately, to cancel an access log upload, to switch immediately to a new remote log file, or to permanently delete all access logs on the ProxySG:

```
SGOS#(config log log_name) commands upload-now
SGOS#(config log log_name) commands cancel-upload
SGOS#(config log log_name) commands rotate-remote-log
SGOS#(config log log_name) commands delete-logs
```

Ordinarily, the ProxySG automatically deletes the local copies of access logs from the ProxySG after the logs have been uploaded. You can manually delete access logs from the ProxySG, but it is not recommended.

5. (Optional) View the results.

```
SGOS#(config log log_name) view
```

6. (Optional) To delete an individual log:

```
SGOS#(config) access-log
SGOS#(config access-log) delete log log_name
```

Testing Access Log Uploading

For the duration of the test, configure the event log to use the verbose event level (see "Configuring Which Events to Log" on page 798). This logs more complete log information. After you test uploading, you can check the event log through the Management Console for the test upload event and determine whether any errors occurred (go to Statistics>Event Logging). You cannot check the event log through the CLI.

To Test Access Log Uploading through the Management Console:

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 the Test Upload button.
The Test upload dialog appears.
3. Click OK in the Test upload dialog that appears.
4. Check the event log for upload results: go to Statistics>Event Logging.

To Test Access Log Uploading through the CLI:

For the duration of the test, configure the event log to use the verbose event level. This logs more complete log information. After you test uploading, you can check the event log under the Statistics tab for the test upload event and determine whether any errors occurred.

1. At the `(config)` command prompt, enter the following commands:

```
SGOS#(config) access-log  
SGOS#(config access-log) edit log log_name  
SGOS#(config log log_name) commands test-upload
```
2. In the Management Console, select Statistics>Event Logging.
3. Click the forward arrow and back arrow buttons to move through the event list.
4. Locate the following event log message entry: `Access Log: Transfer complete`.
If you do not see the message, check for other `Access Log` messages from the appropriate time frame. Use these to help in troubleshooting.
5. To check for a successful test upload at the remote server, locate the file `logname_upload_result` or, for an encrypted test upload, locate the file `logname_upload_result.enc`. If you locate the file, the test was successful.

If you cannot locate the file, use the steps above to open the ProxySG event log, locate the test upload event, and determine whether any errors occurred.

The following is a sample of the test file contents.

```
***** START OF TEST FILE *****  
  
NOTE: This is a verification file sent to test access log uploading.  
Please check the file for correctness and the event log for errors.  
For security purposes, please delete this file after perusal.
```



```

ProxySG Appliance Date:      2003-04-05
ProxySG Appliance Time:     02:17:23 UTC
ProxySG Appliance Name:     10.25.36.47 - Blue Coat SG800
ProxySG Appliance IP:       10.25.36.47
ProxySG Appliance Type:     Blue Coat SG800

```

Sent to FTP server using the following configuration:

```

Host:          10.25.45.35
Port:          21
Path:
User:          joe
Password:      *****
This file should contain approx. 819 bytes
***** END OF TEST FILE *****

```

Viewing Access-Log Statistics

Access-log statistics can be viewed from the Management Console Statistics>Access Logging tab or the CLI `show access-log statistics [log_name]` command. See "Access-Log Statistics" on page 853 for information.

Using Access Logging with Policy Rules

After configuration is complete, you must create rules to manage the access logs you set up. You can create rules through the Visual Policy Manager module of the Management Console, or you can use Content Policy Language (CPL) directly (refer to the *Blue Coat ProxySG Content Policy Language Guide*).

Actions you can do to manage access logging:

- Reset logging to its default
- Disable all logging
- Add logging to a log file
- Disable logging to a log file
- Override specific access-log fields

You can also set the list of logs to be used, but you must use CPL to create this action. It is not available through VPM.

The first two actions—reset logging to its default and disable all logging—are referred to as constant actions, just like the allow/deny actions. You should select only one per rule.

All of the actions are allowed in all layers. If you use VPM, the access-logging actions display in the VPM policy; if you use CPL, you can put the actions into any file, but Blue Coat recommends you use the Local file.

Example: Using VPM to Prevent Entries Matching a Source IP Address from Being Logged

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:

- ❑ Select Configuration>Policy>Visual Policy Manager; click the Launch button.
- ❑ Select Policy>Add Web Access Layer from the menu of the Blue Coat VPM window that appears.

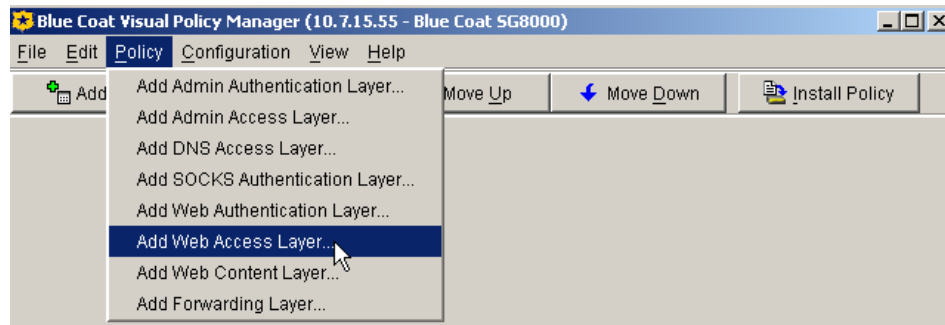


Figure 20-16: Select Add Web Access Layer

- ❑ Type a layer name into the dialog that appears and click OK.
2. Add a Source object:
- ❑ Right click on the item in the Source column; select Set.
 - ❑ Click New in the Set Source Object dialog that appears; select Client IP Address/Subnet.

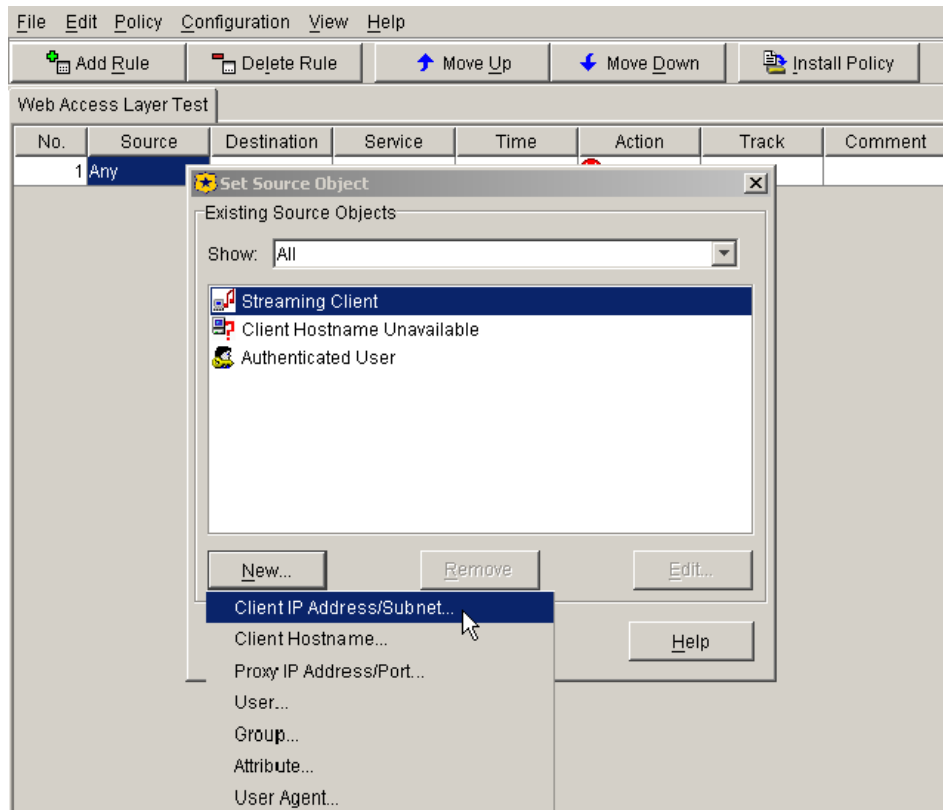


Figure 20-17: Select Client IP Address/Subnet

- 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.
3. Add an Action object to this rule:
- Right-click on the item in the Action column; select Set.

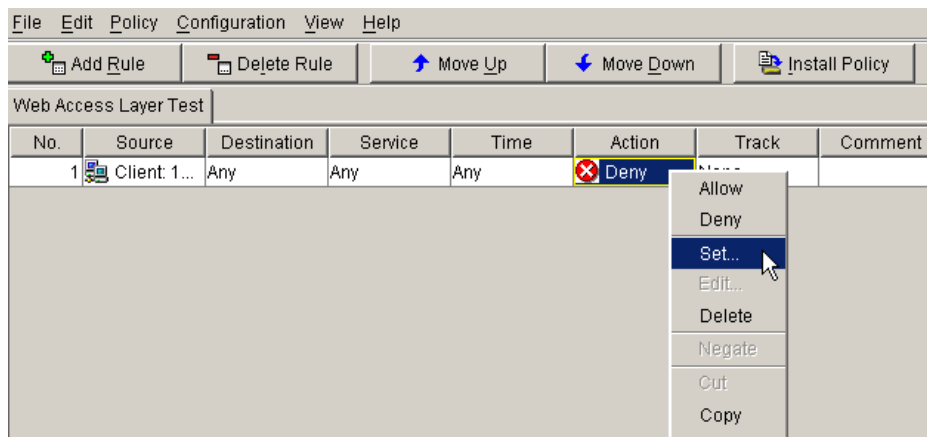


Figure 20-18: Right-Click Action and Select Set

- Click the New button in the Set Action Object dialog that appears; select Modify Access Logging.

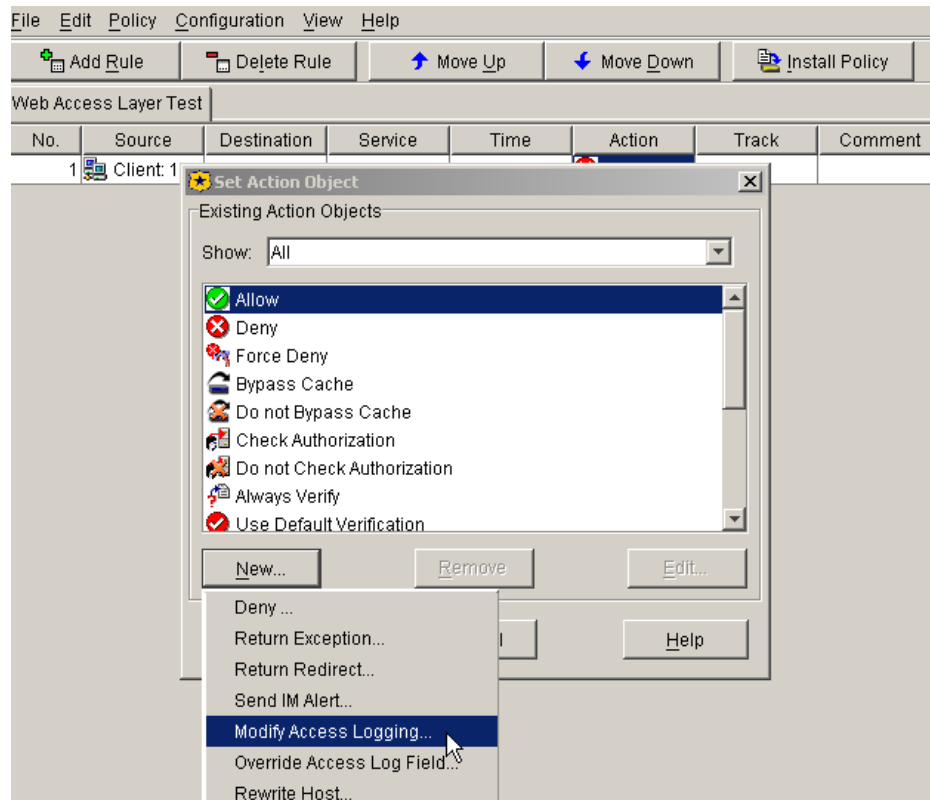


Figure 20-19: Disable Access Logging

- ❑ To disable a particular log, click the Disable logging to radio button and select that log from the drop-down list; to disable all access logging, click the Disable all access logging radio button.
- ❑ Click OK; click OK again; close the VPM window and click Yes in the dialog to save your changes.

Chapter 21: *Maintaining the ProxySG*

The Maintenance tabs provide a set of tools for managing and configuring an array of system-wide parameters, such as restarting the ProxySG, restoring system defaults, configuring SNMP, and managing the ProxySG.

This chapter contains the following sections:

- Restarting the ProxySG
- Restoring System Defaults
- Purging the DNS Cache
- Clearing the System Cache
- Upgrading the ProxySG
- Managing ProxySG Systems
- Event Logging and Notification
- Configuring SNMP
- Disk Reinitialization
- Deleting Objects from the ProxySG

Restarting the ProxySG

The restart options control the restart attributes of the ProxySG in case a restart is needed due to a system fault.

Important: The default settings of the Restart option suits most systems. Changing them without assistance from Blue Coat Systems Technical Support is not recommended.

Hardware and Software Restart Options

The Restart settings determine if the ProxySG performs a faster software-only restart, or a more comprehensive hardware and software restart. The latter can take several minutes longer, depending upon the amount of memory and number of disk drives in the ProxySG.

The default setting of Software only will suit most situations. Restarting both the hardware and software is recommended in situations where a hardware fault is suspected.

For information about the Core Image settings, see "Core Image Restart Options" on page 957.

Note: If you change restart option settings and you want them to apply to the next ProxySG restart, be sure to click Apply.

To Restart the ProxySG through the Management Console:

1. Select Maintenance>General.

The General tab displays.

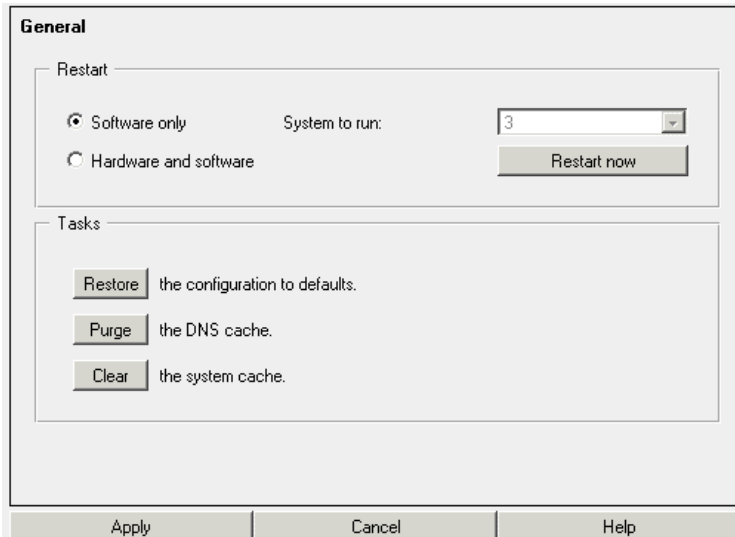


Figure 21-1: Restarting the ProxySG

2. In the Restart field, select either Software only or Hardware and software.
3. If you select the Hardware and software option, select a system from the System to run drop-down list.

The default system is pre-selected.

4. Click Apply.
5. Click the Restart now button.
6. Click OK to confirm and restart the ProxySG.

To Configure the Hardware/Software Restart Settings through the CLI:

At the (config) command prompt, enter the following command:

```
SGOS#(config) restart mode {hardware | software}
```

where:

- | | |
|----------|---|
| hardware | Configures the ProxySG for hardware (and software) restart. |
| software | Configures the ProxySG for software only restart. |

To Restart the ProxySG through the CLI:

Do one of the following:

- To restart the system according to the restart settings, enter the following command:
SGOS# **restart regular**
- To restart hardware and software with a full core image, enter the following command:

```
SGOS# restart abrupt
```

- If you have added a new software image and want to restart the system using that image, enter the following command:

```
SGOS# restart upgrade
```

Restoring System Defaults

The ProxySG 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.
- The `restore-defaults` command with the `keep-console` option allows you to restore default settings without losing all IP addresses on the system.

Restore-Defaults

Settings that are deleted when you use the `restore-defaults` command include:

- All IP addresses (these will have to be restored through the CLI before you can access the Management Console again)
- DNS server addresses (these will have to be restored through the CLI before you can access the Management Console again)
- Installable lists
- All customized configurations
- Third-party vendor licenses, such as SmartFilter or Websense. If you use the `restore-defaults` command after you have installed licenses, and the serial number of your system is configurable (older boxes only), the licenses will fail to install and you will return to the trial period (if any time is left). To correct the problem, you must configure your serial number and install your license-key again.
- 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.

Factory-Defaults

All system settings are deleted when you use the `restore-defaults` command with the `factory-defaults` option.

The only settings that are kept when you use the `restore-defaults` command with the `factory-defaults` option are:

- Trial period information
- The last five installed appliance systems, from which you can pick one for rebooting

The Setup Console password is also deleted if you use `restore-defaults factory-defaults`. For information on the Setup Console password, see "Securing the Serial Port" on page 244.

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 addresses, including default gateway and bridging, except for virtual IP addresses)
- Ethernet maximum transmission unit (MTU) size
- TCP round trip time
- Static routes table information

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 Restore System Defaults through the Management Console:

Note: The `keep-console` and `factory-defaults` options are not available through the Management Console.

1. Select Maintenance>General.
The General tab displays.

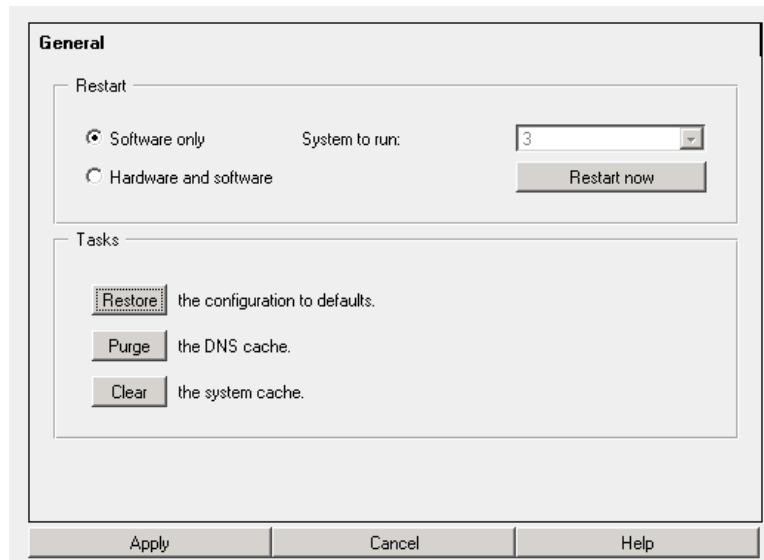


Figure 21-2: Restoring System Defaults

- From the Tasks field of the General Tab, click Restore the configuration to defaults. Keep in mind that, if you restore the configuration from the Management Console, most settings are lost because you cannot use the keep-console option.

The Restore Configuration dialog appears.

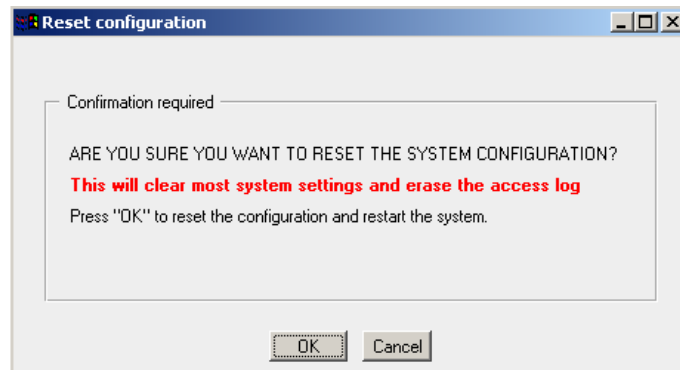


Figure 21-3: Reset Configuration Confirmation

- Click OK.

To Restore System Defaults through the CLI:

At the command prompt, enter the following command:

```
SGOS# restore-defaults [keep-console]
```

To Restore System Defaults without Confirmation:

At the command prompt, enter the following command:

```
SGOS# restore-defaults [keep-console] force
```

To Restore Factory Defaults through the CLI:

At the command prompt, enter the following command:

```
SGOS# restore-defaults factory-defaults
```

Purging the DNS Cache

You can purge 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 Purge the DNS Cache through the Management Console:

1. Select Maintenance>General.
2. In the Tasks field, click the Purge button.
3. Click OK to confirm in the Purge system DNS cache dialog that appears.

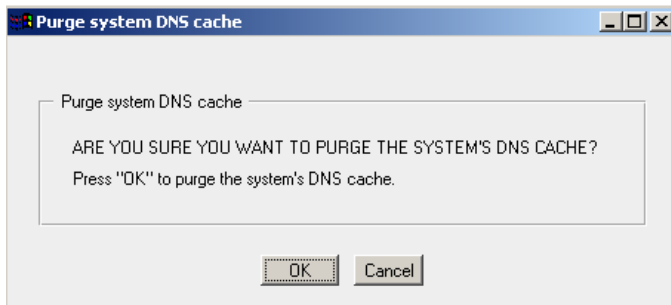


Figure 21-4: Purging the DNS Cache

To Purge the DNS Cache through the CLI:

At the enable command prompt, enter the following command:

```
SGOS# purge-dns-cache
```

Clearing the System Cache

You can clear the system 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 System Cache through the Management Console:

1. Select Maintenance>General.
2. In the Tasks field, click the Clear button.
3. Click OK to confirm in the Clear cache dialog that appears.




Figure 21-5: Clearing the System Cache

To Clear the System Cache through the CLI:

At the enable command prompt, enter the following command:

```
SGOS#clear-cache
```

Troubleshooting Tip

Occasionally, the Management Console might behave incorrectly because of browser caching, particularly if the browser was used to run different versions of the ProxySG Management Console. This problem might be resolved by clearing the browser cache.

Upgrading the ProxySG

When an upgrade to the ProxySG becomes available, you can download it through the Internet and install it. You can also download it to your PC and install it from there.

Important: The auto-detect encoding feature on your browser should be enabled so that it uses the encoding specified in the console URLs. The browser does not use the auto-detect encoding feature by default. If auto-detect encoding is not enabled, the browser ignores the charset header and uses the native OS language encoding for its display.

The ProxySG 4.x Version Upgrade

The appliance must be running version SGOS 3.2.3 or later in order to upgrade to SGOS 4.x. You cannot directly upgrade from any previous version.

When upgrading from the SGOS 3.2.3 or higher release, a copy of the settings is saved prior to any transformations by SGOS 4.x so that the original settings are available if the ProxySG is downgraded to SGOS 3.2.x. Any changes made to the system when it is running SGOS 4.x are not reflected in the saved copy of the original settings. Any changes made when running the previous operating systems after upgrading once to SGOS 4.x are not reflected in the SGOS 4.x settings when the system is re-upgraded to SGOS 4.x. In other words, the upgrade process only happens one time between SGOS 3.2.3 and SGOS 4.x. You can override this behavior by using the `load upgrade [ignore-warnings]` command. For information on using the `load upgrade [ignore-warnings]` command, see "To Upgrade the ProxySG through the CLI:" on page 793.

Note: At least one other system must be unlocked to do the upgrade. If all systems are locked, or all systems except the running system are locked, the Download button in the Management Console is disabled. Similarly, the `load upgrade` command in the CLI generates an error.

To Upgrade the ProxySG through the Management Console:

1. Select Maintenance>Upgrade>Upgrade.

The Upgrade tab displays.

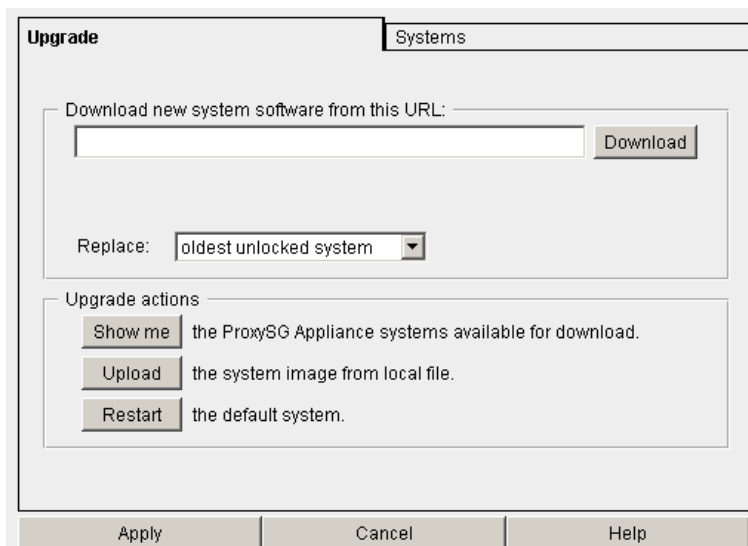


Figure 21-6: Upgrading the ProxySG

2. Click Show me to connect to the Blue Coat download page, follow the instructions, and note the URL of the ProxySG upgrade for your system model. Then enter the URL in the Download new system software from this URL field and click Download.

-or-

(Only if you previously downloaded a system image to your PC) Click Upload and Browse to the file location, then click Install. Note that the upload might take several minutes.

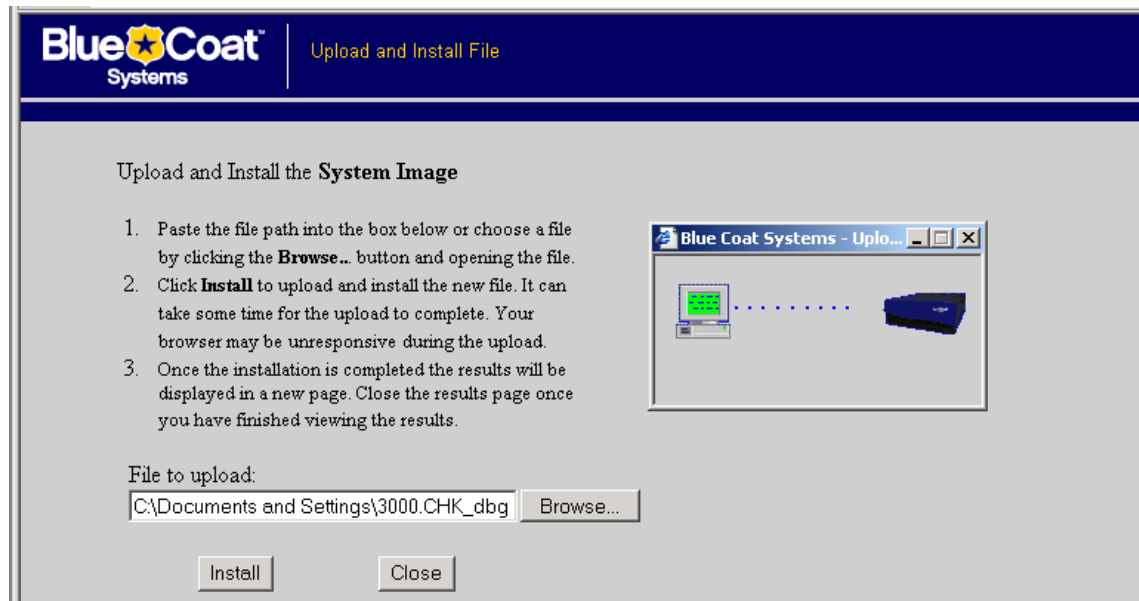


Figure 21-7: Uploading a System Image from a PC

3. (Optional) Select the system to replace in the Replace drop-down list. If you uploaded an image from your PC, refresh the Systems pane to see the new system image.
4. Click Restart.

The Restart system dialog displays.

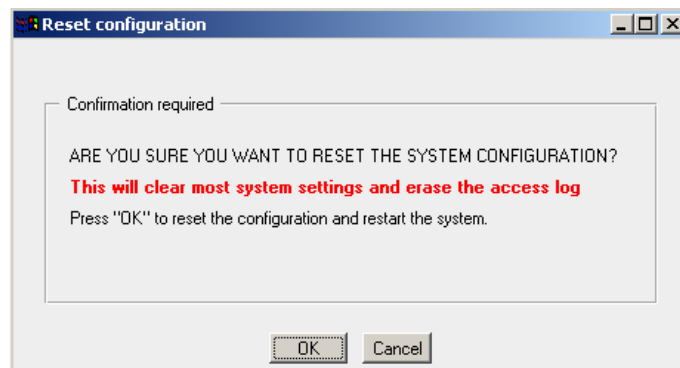


Figure 21-8: Restart System Dialog

5. Click OK to reboot the ProxySG to the default system.

To Upgrade the ProxySG through the CLI:

From the serial console, enter the following commands:

```
SGOS# (config) upgrade-path url
```

where *url* is the location of the SGOS upgrade image.

```
SGOS# (config) exit
SGOS# load upgrade [ignore-warnings]
```

where `ignore-warnings` allows you to force an upgrade even if you receive policy deprecation warnings. Note that using the `load upgrade ignore-warnings` command to force an upgrade while the system emits deprecation warnings results in a policy load failure; all traffic will be allowed or denied according to default policy.

```
SGOS# restart upgrade
```

Managing ProxySG Systems

The ProxySG Systems tab displays the five available ProxySG systems. Empty systems are indicated by the word Empty.

The ProxySG system currently running is highlighted in blue and cannot be replaced or deleted.

From this screen, you can:

- Select the SGOS system version to boot.
- Lock one or more of the available SGOS system versions.
- Select the SGOS system version to be replaced.
- Delete one or more of the available SGOS system versions (CLI only).
- View details of the available SGOS system versions.

To View ProxySG System Replacement Options through the Management Console:

Select Maintenance>Upgrade>Systems.

The Systems tab displays.

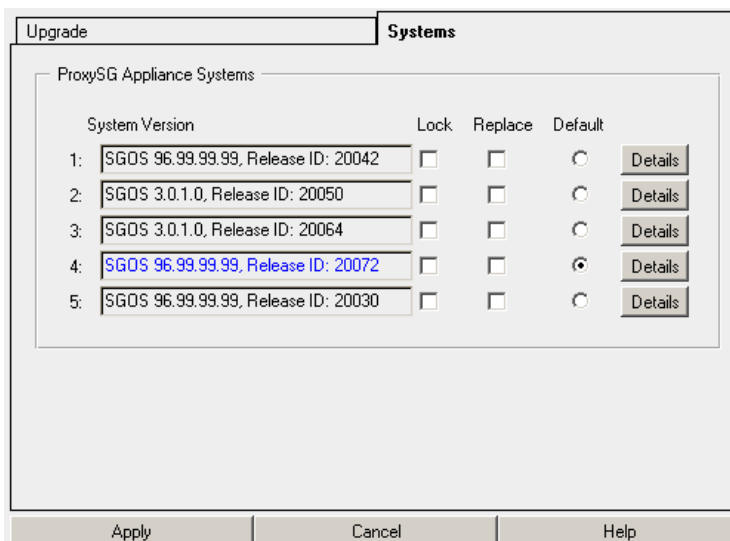


Figure 21-9: Setting SGOS System Version Replacement Properties

To View Details for an SGOS System Version through the Management Console:

1. Select Maintenance>Upgrade>Systems.
2. Click the Details button next to the system for which you want to view detailed information; click OK when you are finished.

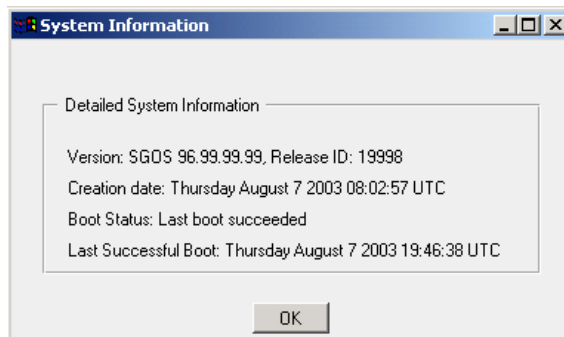


Figure 21-10: SGOS System Version Details

To View Details for an SGOS System Version through the CLI:

At the command prompt:

```
SGOS> show installed-systems
```

Example Session

```
SGOS> show installed-systems
ProxySG Appliance Systems
1. Version: SGOS 96.99.99.99, Release ID: 20042
   Thursday August 21 2003 08:08:58 UTC, Lock Status: Unlocked
   Boot Status: Last boot succeeded, Last Successful Boot: Thursday August 21
   2003 17:51:50 UTC
2. Version: SGOS 3.0.1.0, Release ID: 20050
   Friday August 22 2003 04:43:34 UTC, Lock Status: Unlocked
   Boot Status: Last boot succeeded, Last Successful Boot: Friday August 22
   2003
   16:47:53 UTC
3. Version: SGOS 96.99.99.99, Release ID: 20021
   Tuesday August 12 2003 20:02:42 UTC, Lock Status: Unlocked
   Boot Status: Last boot succeeded, Last Successful Boot: Thursday August 14
   2003 17:57:06 UTC
4. Version: SGOS 96.99.99.99, Release ID: 20029
   Thursday August 14 2003 20:01:55 UTC, Lock Status: Unlocked
   Boot Status: Last boot succeeded, Last Successful Boot: Thursday August 14
   2003 20:49:02 UTC
5. Version: SGOS 96.99.99.99, Release ID: 20030
   Friday August 15 2003 08:01:47 UTC, Lock Status: Unlocked
   Boot Status: Last boot succeeded, Last Successful Boot: Friday August 15
   2003
   19:20:32 UTC
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 to Run on the Next Hardware Restart through the Management Console:

1. Select Maintenance>Upgrade>Systems.
2. Select the preferred ProxySG 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.

To Set the ProxySG to Run on the Next Hardware Restart through the CLI:

At the (config) command prompt:

```
SGOS#(config) installed-systems  
SGOS#(config installed-systems) default system_number  
where system_number is the default system version.
```

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 through the Management Console:

1. Select Maintenance>Upgrade>Systems.
2. Select the system(s) to lock in the Lock column.
3. Click Apply.

To Lock a System through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) installed-systems  
SGOS#(config installed-systems) lock system_number  
where system_number is the system you want to lock.
```

To Unlock a System through the Management Console:

1. Select Maintenance>Upgrade>Systems.

2. Deselect the system(s) to unlock in the Lock column.
3. Click Apply.

To Unlock a System through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) installed-systems
SGOS#(config installed-systems) no lock system_number
      where system_number is the system you want to unlock.
```

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 will be replaced by default. You cannot specify a locked system for replacement.

To Specify the System to Replace through the Management Console:

1. Select Maintenance>Upgrade>Systems.
2. Select the system to replace in the Replace column.
3. Click Apply.

To Specify the System to Replace through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) installed-systems
SGOS#(config installed-systems) replace system_number
      where system_number is the system to be replaced.
```

Deleting a ProxySG System

You can delete any of the ProxySG 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 through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) installed-systems
SGOS#(config installed-systems) delete system_number
      where system_number is the system you want to delete.
```

Event Logging and Notification

You can configure the ProxySG to log system events as they occur. *Event logging* allows you to specify the types of system events logged, the size of the event log, and to configure Syslog monitoring. The ProxySG can also notify you by e-mail if an event is logged.

Configuring Which Events to Log

The event level options are listed from the most to least important events. Because each event requires some disk space, setting the event logging to log all events fills the event log more quickly.

To Set the Event Logging Level through the Management Console:

1. Select Maintenance>Event Logging>Level.

The Level tab displays.

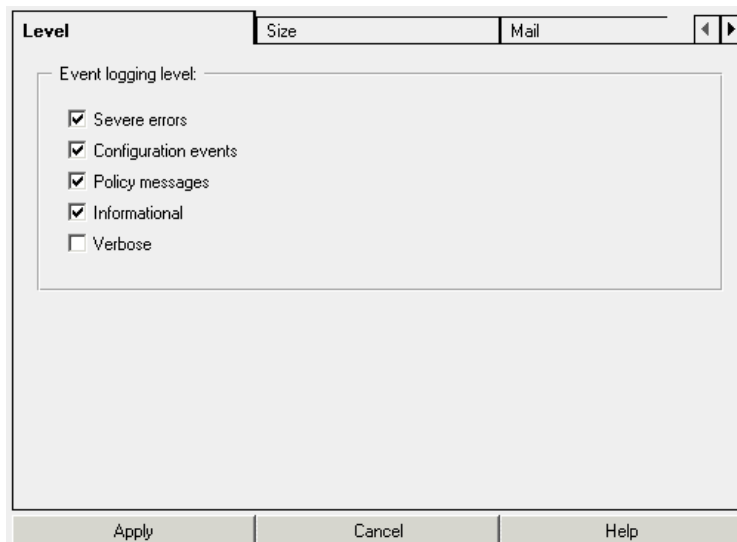


Figure 21-11: Selecting Which Events are Logged

2. Select the events you want to log.

When you select an event level, all levels above the selection are included. For example, if you select Verbose, all event levels are included.

3. Click Apply.

To Set the Event Logging Level through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) event-log
SGOS#(config event-log) level {severe | configuration | policy |
informational | verbose}
```

where:

severe	Writes only severe error messages to the event log.
configuration	Writes severe and configuration change error messages to the event log.
policy	Writes severe, configuration change, and policy event error messages to the event log.
informational	Writes severe, configuration change, policy event, and information error messages to the event log.
verbose	Writes all error messages to the event log.

Setting Event Log Size

You can limit the size of the ProxySG's event log and specify what the appliance should do if the log size limit is reached.

To Set Event Log Size through the Management Console:

1. Select Maintenance>Event Logging>Size.

The Size tab displays.

Figure 21-12: Configuring Event Log Size

2. In the Event log size field, enter the maximum size of the event log in megabytes.
3. Select either Overwrite earlier events or Stop logging new events to specify the desired behavior when the event log reaches maximum size.
4. Click Apply.

To Set Event Log Size through the CLI:

At the (config) command prompt, enter the following command:

```
SGOS#(config) event-log
SGOS#(config event-log) log-size megabytes
SGOS#(config event-log) when-full {overwrite | stop}
```

Specifies event logging behavior should the event log become full.

Enabling Event Notification

The ProxySG can send event notifications to Internet e-mail addresses using SMTP. You can also send event notifications directly to Blue Coat for support purposes. For information on configuring diagnostic reporting, see Appendix E: "Diagnostics" on page 941.

Note: The ProxySG must know the host name or IP address of your SMTP mail gateway to mail event messages to the e-mail address(es) you have entered. If you do not have access to an SMTP gateway, you can use the Blue Coat default gateway to send event messages directly to Blue Coat.

The Blue Coat SMTP gateway will only send mail to Blue Coat. It will not forward mail to other domains.

To Enable Event Notifications through the Management Console:

1. Select Maintenance>Event Logging>Mail.

The Mail tab displays.

Figure 21-13: Enabling Event Notification

2. Click New to add a new email address; click OK in the Add list item dialog that appears.
3. In the SMTP gateway name field, enter the host name of your mail server; or in the SMTP gateway IP field, enter the IP address of your mail server.
4. Click Apply.

To Enable Event Notifications through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) event-log
SGOS#(config event-log) mail smtp-gateway gateway
      where gateway is a domain name or an IP address.
```

```
SGOS#(config event-log) mail add recipient@url
SGOS#(config event-log) exit
SGOS#(config) policy notify
```

Sends event notifications directly to Blue Coat for support purposes.

Syslog Event Monitoring

Syslog is an event-monitoring scheme that is especially popular in UNIX environments. Sites that use syslog typically have a log host node, which acts as a sink (respository) 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.

To Enable Syslog Monitoring through the Management Console:

1. Select Maintenance>Event Logging>Syslog.

The Syslog tab displays.

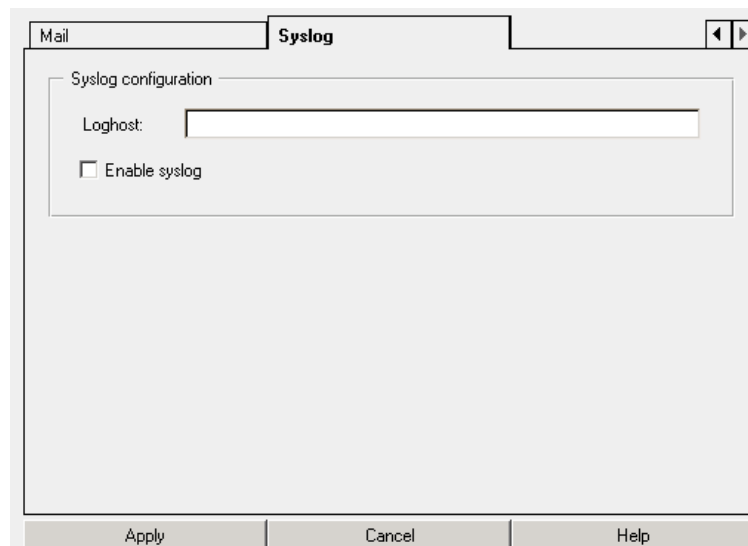


Figure 21-14: Setting Up Syslog Monitoring

2. In the Loghost field, enter the domain name or IP address of your loghost server.
3. Select Enable Syslog.
4. Click Apply.

To Enable Syslog Monitoring through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) event-log
SGOS#(config event-log) syslog loghost loghost
```

where *loghost* is the log host domain name or IP address.

```
SGOS#(config event-log) syslog enable
```

Viewing Event Log Configuration and Content through the CLI

You can view the system event log through the CLI, either in its entirety or selected portions of it.

Viewing the Event Log Configuration through the CLI

You can view the event log configuration, from `show` or from `view` in the event-log configuration mode.

To View the Event Log Configuration through the CLI:

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
```

Viewing the Event Log Contents through the CLI

Again, you can view the event log contents from the `show` command or from the event-log configuration mode.

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 timezone of the ProxySG.

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 will be 00:00:00
- If the time is omitted for the ending time, it will be 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 will be 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 "2004-10-22 9:00:00" end "2004-10-22 9:15:00"
2004-10-22 09:00:02+00:00UTC "Snapshot sysinfo_stats has fetched /sysinfo-stats
" 0 2D0006:96 ../Snapshot_worker.cpp:183
2004-10-22 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
```

Configuring SNMP

You can view a ProxySG using a Simple Network Management Protocol (SNMP) management station. The ProxySG supports MIB-2 (RFC 1213), Proxy MIB, and the RFC2594 MIB, and can be downloaded at the following URL: <http://download.bluecoat.com/release/SGOS3/index.html>. (The SNMP link is in the lower right-hand corner).

Enabling SNMP

To view a ProxySG from an SNMP management station, you must enable and configure SNMP support on the ProxySG.

To Enable and Configure SNMP through the Management Console:

1. Select Maintenance>SNMP>SNMP General.

The SNMP General tab displays.

SNMP General		Community Strings	Traps
General settings:			
<input type="checkbox"/> Enable SNMP	Reset SNMP settings		
Values for MIB variables:			
sysLocation:	3C 6785		
sysContact:	admin@company.com		
Apply	Cancel	Help	

Figure 21-15: Enabling SNMP

2. Select Enable SNMP.
3. In the sysLocation field, enter a string that describes the ProxySG's physical location.
4. In the sysContact field, enter a string that identifies the person responsible for administering the ProxySG.

5. Click Apply.

To Enable and Configure SNMP through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) snmp
SGOS#(config snmp) enable
SGOS#(config snmp) sys-location location
```

where *location* specifies the ProxySG's physical location.

```
SGOS#(config snmp) sys-contact contact
```

where *contact* identifies the person responsible for administering the ProxySG.

Configuring SNMP Community Strings

Use *community strings* to restrict access to SNMP data. To read SNMP data on the ProxySG, specify a *read community string*. To write SNMP data to the ProxySG, specify a *write community string*. To receive traps, specify a *trap community string*. By default, all community string passwords are set to public.

Note: If you enable SNMP, make sure to change all three community-string passwords to values that are difficult to guess. Use a combination of uppercase, lowercase, and numeric characters. An easily-guessed community-string password makes it easier to gain unauthorized access to the ProxySG and network.

To Set or Change Community Strings through the Management Console:

1. Select Maintenance>SNMP>Community Strings.

The Community Strings tab displays.

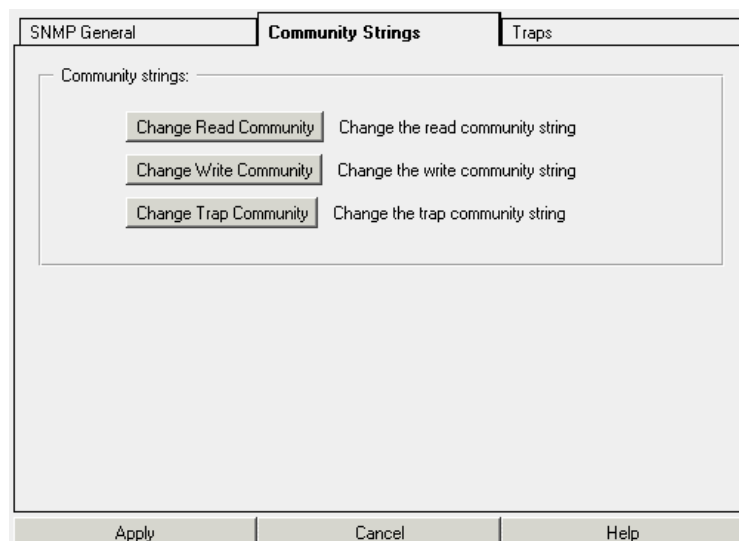


Figure 21-16: Configuring SNMP Community Strings

2. Click the community string button you want to change.

The Change Read/Write/Trap Community dialog displays.

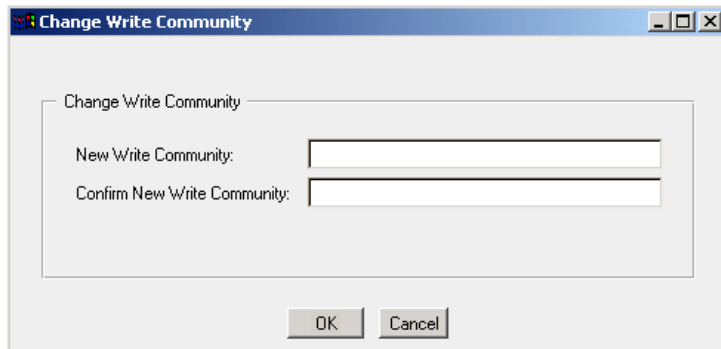


Figure 21-17: SNMP Change Community String Dialog

3. Enter and confirm the community string; click OK.
4. Click Apply.

To Set or Change Community Strings through the CLI:

You can set the community strings in either cleartext or encrypted form.

To set them in cleartext:

```
SGOS#(config) snmp
SGOS#(config snmp) enable
SGOS#(config snmp) read-community password
SGOS#(config snmp) write-community password
SGOS#(config snmp) trap-community password
```

To set them as encrypted:

```
SGOS#(config) snmp
SGOS#(config snmp) enable
SGOS#(config snmp) encrypted-read-community encrypted-password
SGOS#(config snmp) encrypted-write-community encrypted-password
SGOS#(config snmp) encrypted-trap-community encrypted-password
```

Configuring SNMP Traps

The ProxySG can send SNMP traps to a management station as they occur. By default, all system-level traps are sent to the address specified. You can also enable authorization traps to send notification of attempts to access the ProxySG Management Console.

Note: The SNMP trap for CPU utilization is sent only if the CPU continues to stay up for 32 or more seconds.

To Enable SNMP Traps through the Management Console:

Note: You cannot configure SNMP traps to go out through a particular interface.

1. Select Maintenance>SNMP>Traps.

The Traps tab displays.

Figure 21-18: Configuring SNMP Traps

2. In the Send traps to fields, enter the IP address(es) of the workstation(s) where traps are to be sent.
3. To receive authorization traps, select Enable authorization traps.
4. Click Apply.

To Enable SNMP Traps through the CLI:

Note: You cannot configure SNMP traps to go out through a particular interface.

1. At the (config) command prompt, enter the following commands:

```
SGOS#(config) snmp
SGOS#(config snmp) enable
SGOS#(config snmp) trap-address 1 ip_address
```

To add additional trap addresses, repeat using trap-address 2 or trap-address 3 to specify the IP address for traps 2 and 3.

2. (Optional) To enable authorization traps, enter the following command:

```
SGOS#(config snmp) authorize-traps
```

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 ProxySG, contact Blue Coat. Technical Support for assistance.

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.

Multi-Disk ProxySG

On a multi-disk ProxySG, the master disk is the leftmost valid disk. “Valid” means that the disk is online, has been properly initialized, and is not marked as invalid or unusable.

If the current master disk is taken offline, reinitialized, or declared invalid or unusable, the leftmost valid disk that has not been reinitialized since restart becomes the master disk. Thus, as disks are reinitialized in sequence, a point is reached where no disk can be chosen as the master. At this point, the current master disk is the last disk. If this disk is taken offline, reinitialized, or declared invalid or unusable, the ProxySG is restarted.

On a multi-disk ProxySG, a disk is reinitialized by setting it to empty and copying pre-boot programs, boot programs, and starter programs, and system images from the master disk to the reinitialized disk.

Reinitialization is done online without rebooting the ProxySG. (For more information, refer to the `#disk` command in the *Blue Coat ProxySG Command Line Reference*.) ProxySG operations, in turn, are not affected, although during the time the disk is being reinitialized, that disk is not available for caching. Note that only the master disk reinitialization will restart 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 will contain 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 ProxySG 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.

Single-Disk ProxySG

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

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 roughly a million. The number is based on the 4GB RAM on the motherboard and is not user-configurable.

This feature is not available in the Management Console. Use the CLI instead.

To Delete a Single Object from the ProxySG through the CLI:

At the `(config)` prompt, enter the following command:

```
SGOS#(config) content delete url url
```

To Delete Multiple Objects from the ProxySG through the CLI:

At the (config) prompt, enter the following command:

```
SGOS#(config) content delete regex regex
```


Chapter 22: *Statistics*

The Statistics tabs of the Management Console allows you to graphically view the status of many system operations, as well as to take disks offline and put them online. Many statistics are available through the CLI, but without the benefit of graphical display.

You can also view detailed system information through the CLI using the `show` command. You can access this command through either the enable command prompt (`SGOS#`) or the config command prompt (`SGOS#(config)`). For convenience, the procedures in this chapter show only the enable command prompt.

Selecting the Graph Scale

Some statistics are reported in the form of bar graphs. Most bar graphs offer the option to show all values in the graph or to clip a percentage of the peak values, which means that a percentage is allowed to fall off the scale. For example, if you select clip 25% of peaks, the top 25% of the values will be allowed to exceed the scale for the graph, showing greater detail for the remaining 75% of the values. To set the graph scale, select a value from the Graph scale should drop-down list. Some of the graphs offer the option of viewing statistics in bytes or objects. On these pages, you can switch among viewing modes by selecting bytes served or objects served mode from the Graph shows or Percentages reflect drop-down list.

You can also move your cursor over the bar graphs to dynamically display color-coded statistical information. See Figure 22-7 for an example of this.

General Statistics

The General statistics tabs in the Management Console (Summary, Environment, and Disks) provide information about system configuration and the status of hardware sensors and allow you to take disks offline and put them online. These statistics are also available in the CLI.

Note: The ProxySG 400 Series Appliances do not have an Environment tab.

System Summary

The device provides a variety of information on its status. The fields on the Summary tab are described below:

- Disks Installed—the number of disk drives installed in the device. The Disks tab displays the status of each drive.
- Memory installed—the amount of RAM installed in the device.
- CPUs installed—the number of CPUs installed in the device.
- Software image—the version and release number of the device image.
- Serial number—the serial number of the machine, if available.

- System started—the time and date the device was started.
- CPU utilization—the current percent utilization of the device CPU.

Viewing the System Summary

To View the System Summary Statistics through the Management Console

Select **Statistics>General>Summary**.

The **General Summary** tab displays.

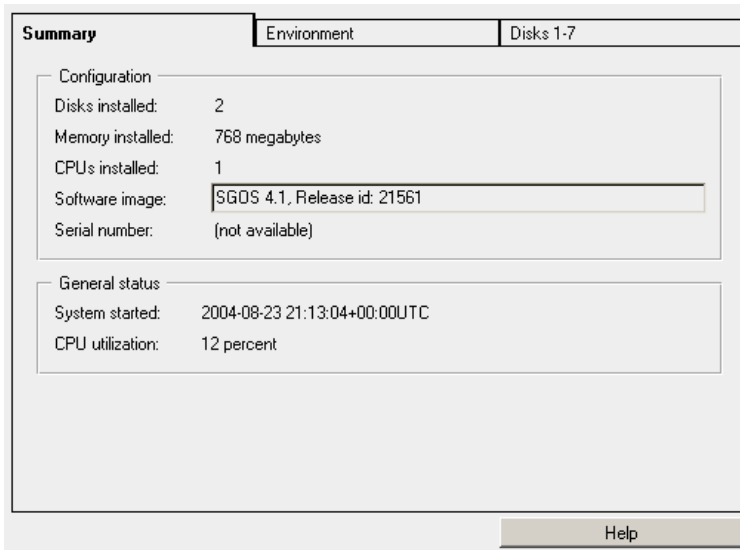


Figure 22-1: General Summary Tab

To View the System Summary through the CLI

Enter the following command at the prompt:

```
SGOS# show status
Configuration:
  Disks installed:          2
  Memory installed:       768 megabytes
  CPUs installed:         1
  Software version:       SG 4.1
  Release id:             21574
  Machine id:             00D0B7655D48
  Serial number:          (not available)
  NIC 0 MAC:              00D0B7655D01
General status:
  System started:         2004-09-10 18:05:14+00:00UTC
  CPU utilization:        17%
```


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: You cannot view environment statistics on a ProxySG 400 Series Appliance.

To View the System Environment Statistics through the Management Console

1. Select Statistics>General>Environment.

The Environment tab displays.

Note: This tab can vary depending on the type of ProxySG that you are using.

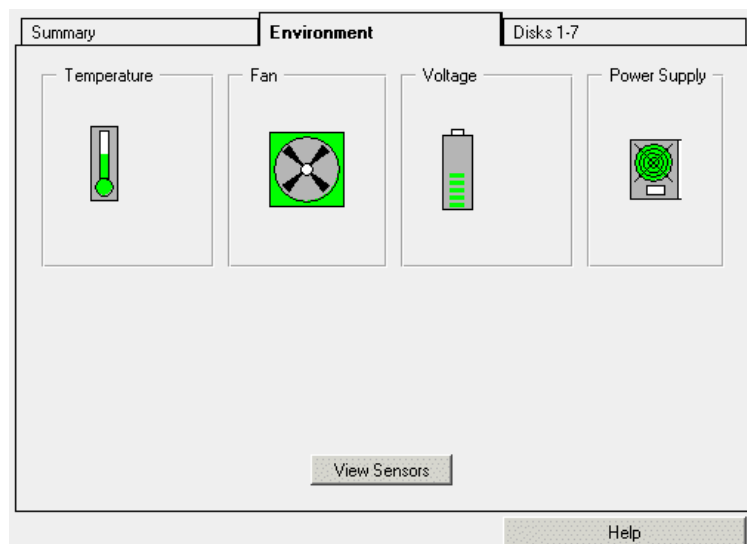


Figure 22-2: Environment Tab

2. Click View Sensors to see detailed sensor values; close the window when you are finished.

Sensor statistics

Sensor Name	Reading	Status
Bus temperature #1	26.0 C	OK
Bus temperature #2	23.0 C	OK
CPU temperature #1	49.0 C	OK
Fan #1		OK
Fan #2		OK
Bus voltage #1	1.4 volts	OK
CPU voltage #1	2.0 volts	OK
Power supply #1		OK
Power supply #2		OK

Figure 22-3: Sensor Statistics Window

To View the System Environmental Statistics through the CLI

Note: You cannot view environmental statistics on a ProxySG 400 Series Appliance.

Enter the following command at the prompt (the results that display vary among ProxySG platforms):

```
SGOS# show environmental
Environmental Sensor Information
Baseboard Temperature # 1 :
Temperature Reading: 27.0 C
Current Threshold Status : NOMINAL -- OK
% UPPER CRITICAL      : 60.0
% UPPER NON CRITICAL  : 55.0
% LOWER NON CRITICAL  : 0.0
% LOWER CRITICAL      : -10.0
Baseboard Temperature # 2 :
Temperature Reading: 25.0 C
Current Threshold Status : NOMINAL -- OK
% UPPER CRITICAL      : 60.0
% UPPER NON CRITICAL  : 55.0
% LOWER NON CRITICAL  : 0.0
% LOWER CRITICAL      : -10.0
Baseboard Voltage # : 1
Voltage Reading: 1.4
Current Threshold Status : NOMINAL -- OK
% UPPER CRITICAL      : 1.7
% LOWER CRITICAL      : 1.2
Fans
Fan #1 : Running OK
Fan #2 : Running OK
```

Power Supplies
 Power Supply #1 : OK
 Power Supply #2 : OK

where the Upper (non) Critical and Lower (non) Critical Temperature and Voltage values are for reference and indicate values that are (critically or non-critically) too high or too low. The Temperature/Voltage Reading indicates the current status of the ProxySG. The Current Threshold Status indicates whether a problem exists.

Viewing Disk Status

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 through the Management Console

1. Select Statistics>General>Disks.

The Disks tab displays, providing information about the disk in slot 1.

Note: The name and appearance of this tab differs, depending on the range of disks available to the ProxySG model you use.

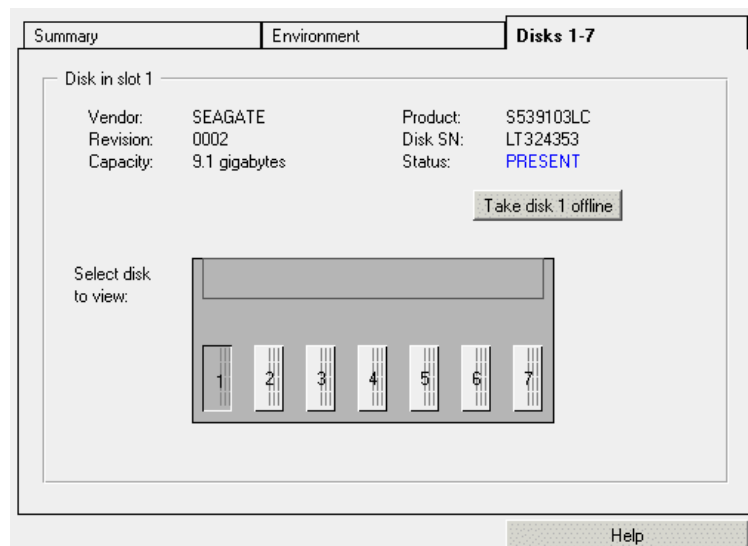


Figure 22-4: Disks Tab

2. Select the disk to view or to take offline by clicking the appropriate disk icon.
3. (Optional) To take the selected disk offline, click the Take disk x offline button (where x is the number of the disk you have selected); click OK in the Take disk offline dialog that displays.

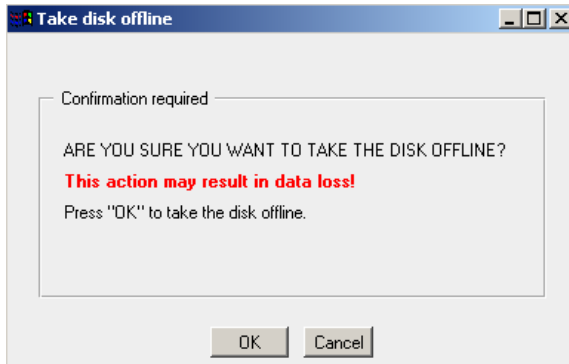


Figure 22-5: Take Disk Offline Dialog

To View Disk Statistics through the CLI

Enter the following command at the prompt:

```
SGOS# show disk {all | disk_number}
```

where *all* displays information about all disks and *disk_number* displays information about the disk specified.

To Take a Disk Offline through the CLI

Enter the following command at the prompt:

```
SGOS# disk offline disk_number
```

where *disk_number* is the number of the disk that you want to take offline.

System Usage Statistics

The System Usage tabs (CPU, Bandwidth Gain, Client Comp. Gain, Server Comp. Gain, Freshness, and Refresh Bandwidth) display bar graphs that illustrate the last 60 minutes, 24 hours, and 30 days for CPU utilization, bandwidth gain, client and server compression gain, freshness of objects in the cache, and the average network bandwidth used to maintain freshness.

Viewing CPU Utilization

Through the Management Console, you can view the average CPU utilization percentages for the ProxySG over the last 60 minutes, 24 hours, and 30 days. You can see the current CPU utilization statistic through the CLI.

To View CPU Utilization through the Management Console

1. Select Statistics>System Usage>CPU.

The CPU tab displays.

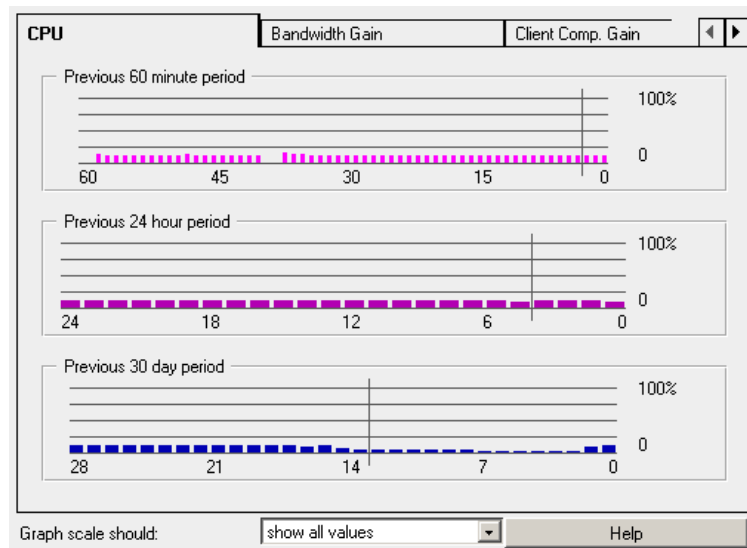


Figure 22-6: CPU Tab

- (Optional) To set the graph scale to clip a percentage of the peaks in value, select a percentage from the Graph scale should drop-down list.

To View CPU Utilization through the CLI

Enter the following command at the prompt:

```
SGOS# show cpu
```

Viewing Bandwidth Gain

Through the Management Console, you can view bandwidth-gain statistics for the ProxySG over the last 60 minutes, 24 hours, and 30 days. These statistics are not available through the CLI.

The green display on the bar graph represents client data; the blue display represents server data. Hover your cursor over the graph to see the bandwidth gain data.

It is normal to see 100% markers in places where there has been no client-use for the activity. This means that, of the server-side traffic being expended, 100% of it is being expended for ProxySG internal usage, such as asynchronous adaptive refresh.

To View Bandwidth Gain Statistics through the Management Console

- Select Statistics>System Usage>Bandwidth Gain.

The Bandwidth Gain tab displays.

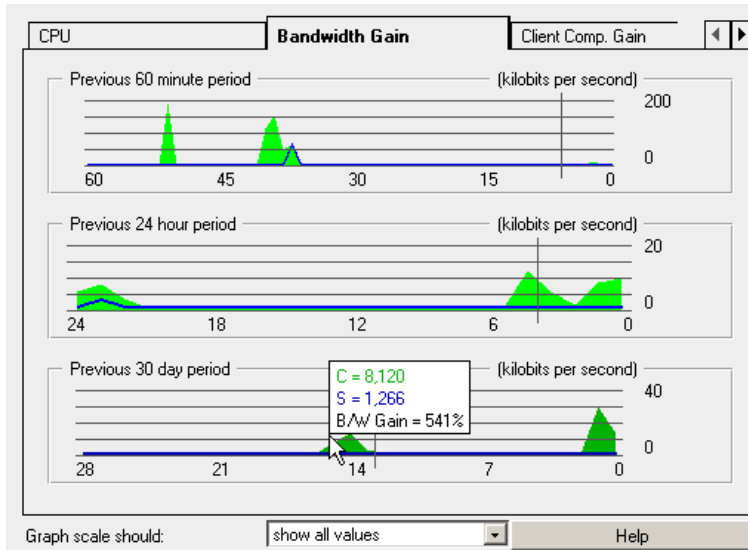


Figure 22-7: Bandwidth Gain Tab

2. (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

Viewing Overall Client and Server Compression Gain Statistics

Under System Usage, you can view overall 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. HTTP/FTP client and server compression-gain statistics are displayed under HTTP/FTP History. 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.

To View Overall Client Compressed Gain Statistics through the Management Console

1. Select Statistics>System Usage>Client Comp. Gain.

The Client Comp. Gain tab displays.

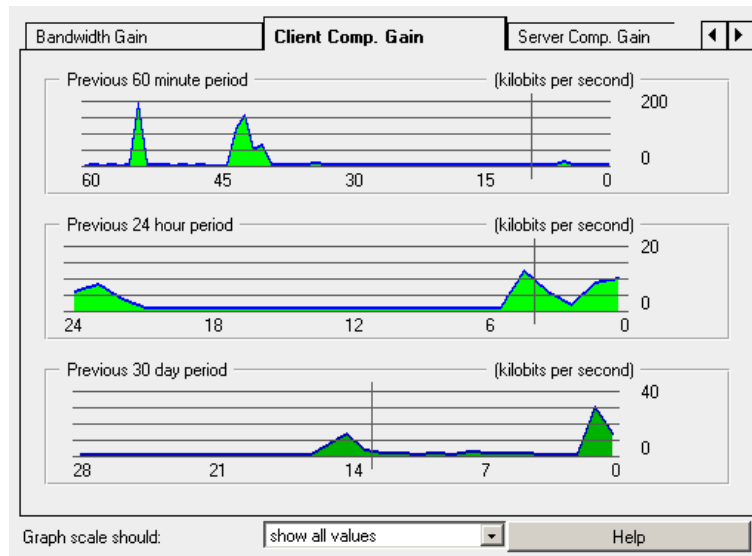


Figure 22-8: System Usage Client Compression Gain Tab

- (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

To View Overall Server Compressed Gain Statistics through the Management Console

- Select Statistics>System Usage>Server Comp. Gain.

The Server Comp. Gain tab displays.

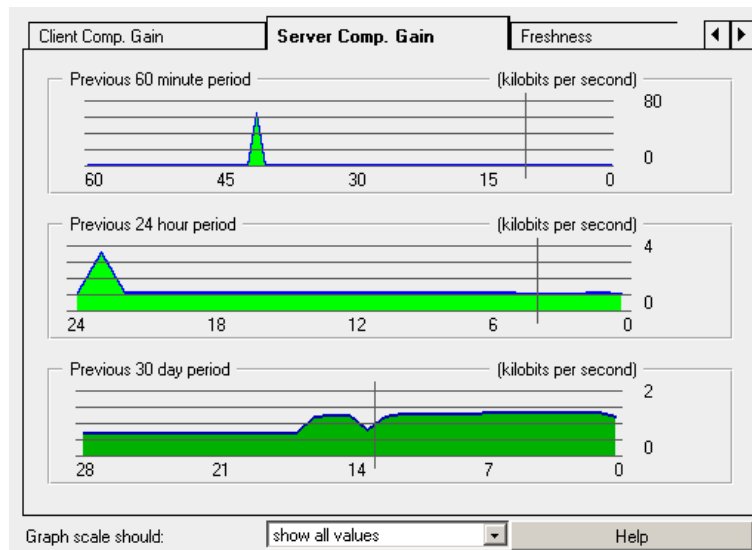


Figure 22-9: System Usage Server Compression Gain Tab

- (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

Viewing Cache Freshness

The Freshness tab illustrates the estimated freshness of objects in the cache over the last 60 minutes, 24 hours, and 30 days. These statistics are not available through the CLI.

Freshness applies only to objects that are cached (all objects that are not cached are always 100% fresh). For example, if the estimated freshness is 99%, that means when you request an object there is a 99% chance that object is fresh in the cache.

To View Cache Freshness through the Management Console

1. Select Statistics>System Usage>Freshness.

The Freshness tab displays.

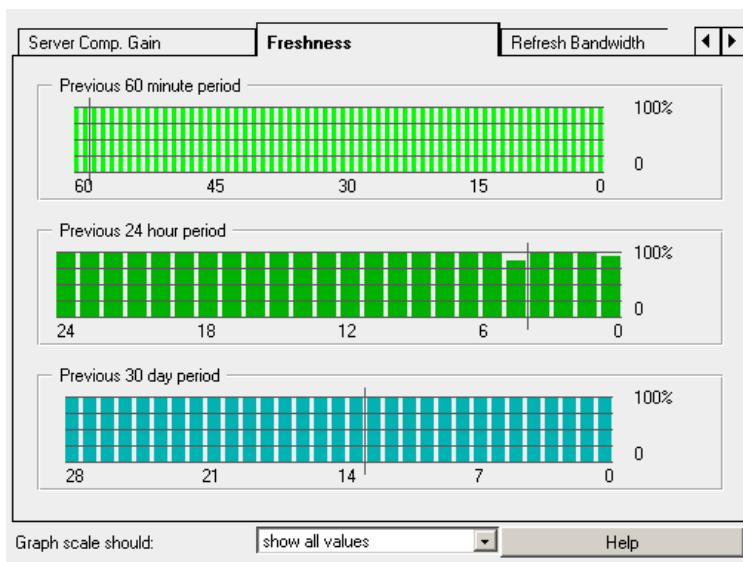


Figure 22-10: Freshness Tab

2. (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

Viewing Refresh Bandwidth Statistics

The Refresh Bandwidth tab illustrates the average network bandwidth used to maintain freshness in the cache over the last 60 minutes, 24 hours, and 30 days. These statistics are not available through the CLI.

To View Refresh Bandwidth through the Management Console

1. Select Statistics>System Usage>Refresh Bandwidth.

The Refresh Bandwidth tab displays.

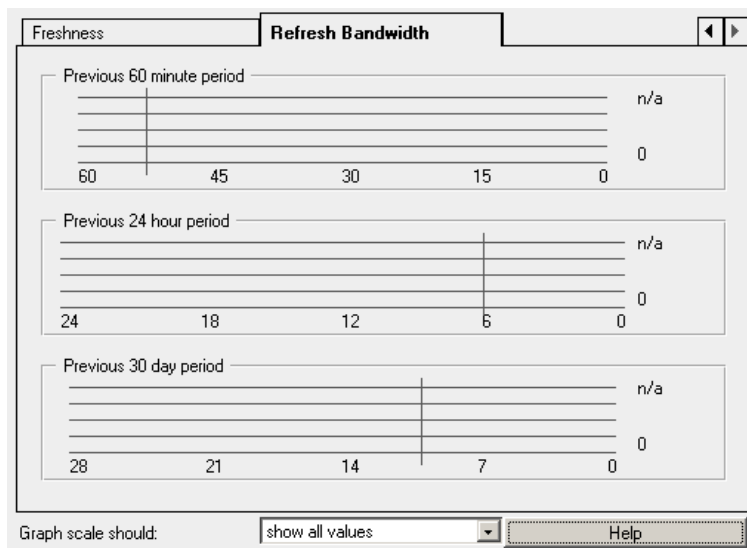


Figure 22-11: Refresh Bandwidth Tab

- (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

HTTP/FTP History Statistics

The HTTP/FTP History tabs (HTTP/FTP Objects, HTTP/FTP Bytes, HTTP/FTP Clients, Client Comp. Gain, and Server Comp. Gain) display bar graphs that illustrate the last 60 minutes, 24 hours, and 30 days for the number of HTTP/FTP objects served, the number of HTTP/FTP bytes served, the maximum number of active HTTP/FTP clients processed, and the HTTP/FTP client and server compression-gain statistics. Overall client and server compression-gain statistics are displayed under System Usage.

Note: You can view current HTTP configurations and statistics through the CLI using the `show http` and `show http-stats` commands.

Viewing the Number of HTTP/FTP Objects Served

The HTTP/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. To review the number of cached objects versus non-cached objects, view the Efficiency tabs.

Note: The maximum number of objects that can be stored in a ProxySG is roughly a million. The number is based on the 4GB RAM on the motherboard and is not user configurable.

To View the Number of HTTP/FTP Objects Served through the Management Console

- Select Statistics>HTTP/FTP History>HTTP/FTP Objects.

The HTTP/FTP Objects tab displays.




Figure 22-12: HTTP/FTP Objects Tab

2. (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/FTP Bytes Served

The 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/FTP Bytes Served through the Management Console

1. Select Statistics>HTTP/FTP History>HTTP/FTP Bytes.

The HTTP/FTP Bytes tab displays.




Figure 22-13: HTTP/FTP Bytes Tab

2. (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/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/FTP Objects and HTTP/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 through the Management Console

1. Select Statistics>HTTP/FTP History>HTTP/FTP Clients.

The HTTP/FTP Clients tab displays.

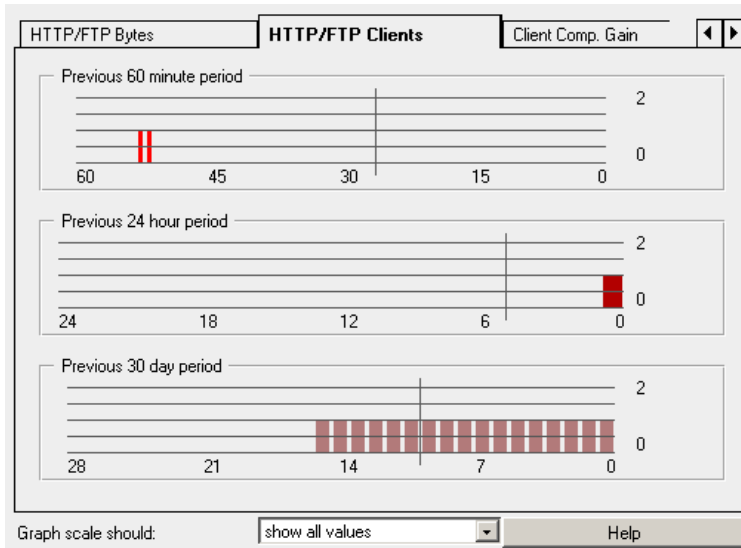


Figure 22-14: HTTP/FTP Clients Tab

2. (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

Viewing HTTP/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.

To View HTTP/FTP Client Compressed Gain Statistics through the Management Console

1. Select Statistics>HTTP/FTP History>Client Comp. Gain.

The Client Comp. Gain tab displays.

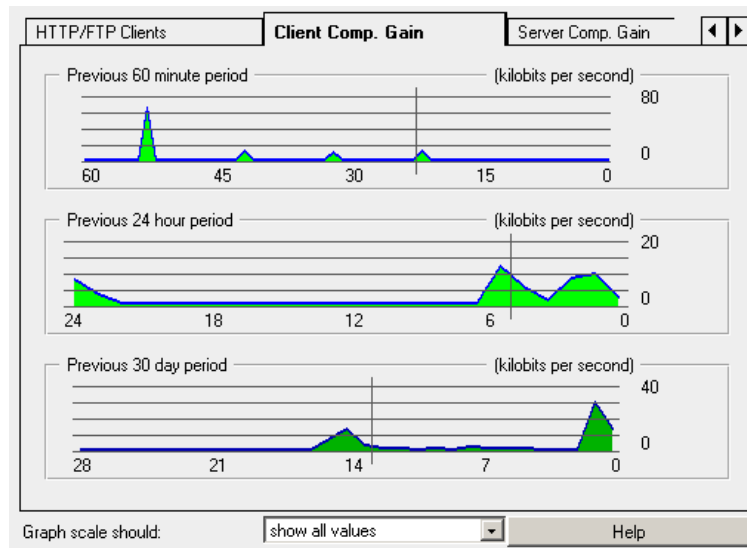


Figure 22-15: HTTP/FTP Client Comp. Gain Tab

- (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

To View HTTP/FTP Server Compressed Gain Statistics through the Management Console

- Select Statistics>HTTP/FTP History>Server Comp. Gain.

The Server Comp. Gain tab displays.

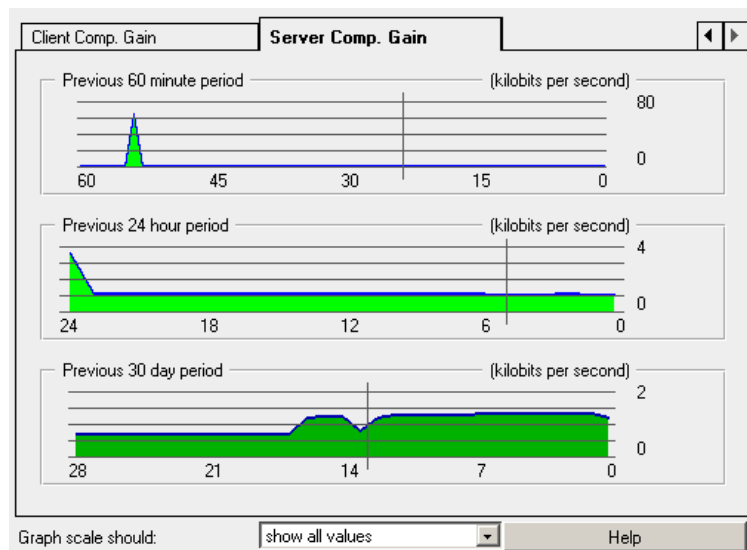


Figure 22-16: HTTP/FTP Server Comp. Gain Tab

- (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

IM History Statistics

The IM statistics allow you to track IM connections, file transfers, and messages that are currently in use and in total, or have been allowed and denied. The information can be displayed for each IM client type or combined.

IM Connection Data Tab

The following IM Connection Data statistics indicate current and overall connection data since the last statistics clear:

- Native Clients—The number of native IM clients connected.
- HTTP Clients—The number of HTTP IM clients connected.
- Chat Sessions—The number of IM chats occurring.
- Direct IM Sessions—The number of chats using direct connections.
- File Transfers—The number of file transfers sent through IM clients.

To view the Connection Data Statistics through the Management Console

1. Select Statistics>IM History>IM Connection Data.

The IM Connection Data tab displays.

	Current	Total
Native Clients:	0	0
HTTP Clients:	0	0
Chat Sessions:	0	0
Direct IM Sessions:	0	0
File Transfers:	0	0

Figure 22-17: IM Connection Statistics Data Tab

2. The default protocol is All. To select a specific protocol, select AOL, MSN, or Yahoo from the drop-down list.

IM Activity Data Tab

The following IM Activity Data statistics indicate allowed and denied connections since the last statistics clear:

- Logins—The number of times IM clients have logged in.
- Messages—The number of IM messages.
- File Transfers—The number of file transfers sent through IM clients.
- Voice Chats—The number of voice conversations through IM clients.
- Messages—The number of IM messages reflected or not reflected (if IM Reflection policy is enabled).

Note: The IM activity data statistics are available only through the Management Console.

To View the Activity Data Statistics through the Management Console

1. Select Statistics>IM History>IM Activity Data.

The IM Activity Data tab displays.

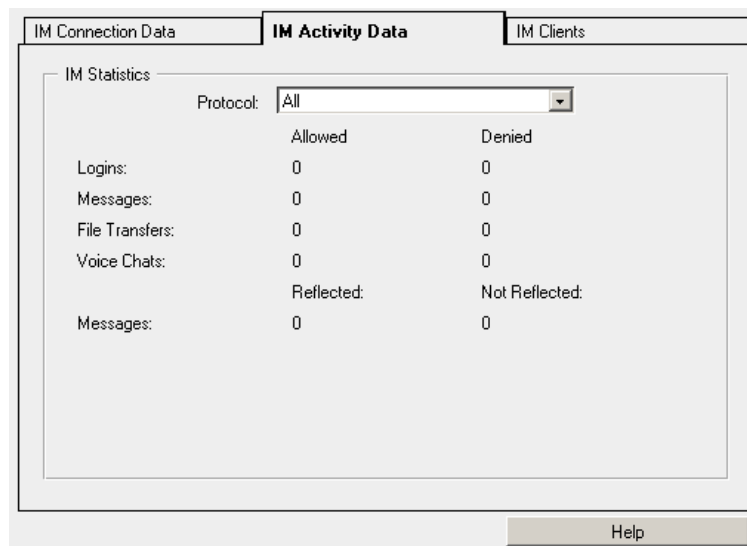


Figure 22-18: IM Activity Statistics Data Tab

2. The default protocol is All. To select a specific protocol, select AOL, MSN, or Yahoo from the drop-down list.

IM Clients Tab

The IM Clients tab displays dynamic graphical statistics for connections over 60 minutes, 24 hours and 30 days. The page can display all values in the graph or clip a percentage of peak values. When peak values are clipped by a percentage, that percentage is allowed to fall off the top of the scale.

For example, if you clip 25% of the peaks, the top 25% of the values are allowed to exceed the scale for the graph, showing greater detail for the remaining 75% of the values.

Move the cursor over the graphs to dynamically display the color-coded AOL, MSN, Yahoo, and total statistics.

Note: The IM clients statistics are available only through the Management Console.

To View the Client Connection Statistics through the Management Console

1. Select Statistics>IM History>IM Clients.

The IM Clients tab displays.

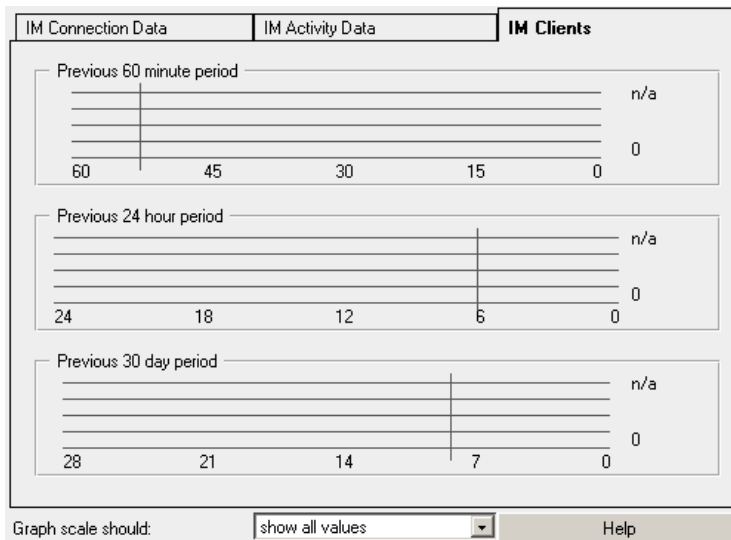


Figure 22-19: IM Client Data Statistics

2. (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

IM Statistics and Configuration in the CLI

To View IM Statistics through the CLI

Enter the following command at the prompt:

```
SGOS# show im {aol-statistics | msn-statistics | yahoo-statistics}
```

To View the IM Configuration through the CLI

Enter the following command at the prompt:

```
SGOS# show im configuration
```


P2P History Statistics

You can construct policy that controls, blocks, and logs peer-to-peer (P2P) activity and limits the bandwidth consumed by P2P traffic (see Chapter 15: “Advanced Policy” on page 561 for information about constructing P2P policy). The following section explains how to view P2P statistics, using either the Management Console or the CLI.

Note: Some P2P statistics (P2P client connections and total bytes sent and received over a period of time) can only be viewed through the Management Console (see “P2P Clients” and “P2P Bytes”, below).

P2P Data

The P2P Data tab on the Management Console displays P2P statistics, either all P2P services at once or one service at a time.

Table 22.1 details the statistics provided through the Management Console P2P Data tab or through the CLI.

Table 22.1: P2P Data Statistics

Status	Description
Current Tunneled Sessions	The current number of P2P client connections using native transport.
Current HTTP Requests	The current number of HTTP requests from P2P clients.
Total Tunneled Sessions	The cumulative number of P2P client connections using native transport since the ProxySG was last rebooted.
Total HTTP Requests	The cumulative number of HTTP requests from P2P clients since the ProxySG was last rebooted.
Total Bytes Received	The total number of bytes received from all P2P clients.
Total Bytes Sent	The total number of bytes sent to all P2P clients.

To View P2P Data Statistics through the Management Console

1. Select Statistics>P2P History>P2P Data.

The P2P Data tab displays. The default view shows all P2P protocols.

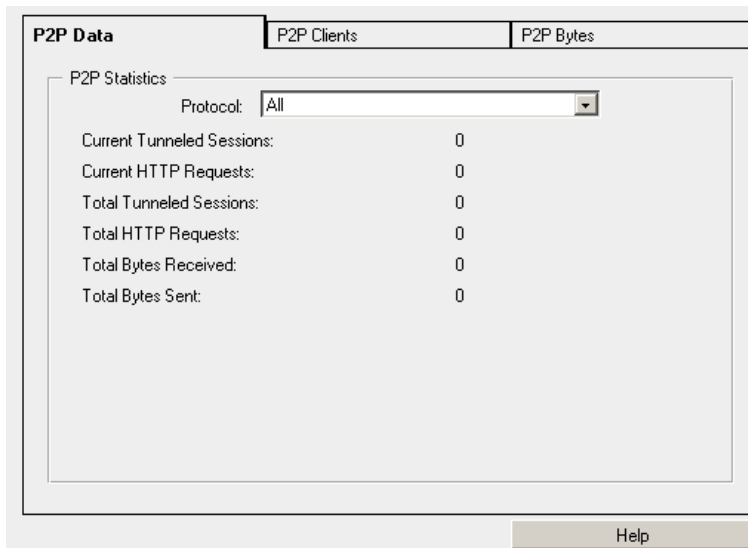


Figure 22-20: P2P Data Tab

2. (Optional) To view the statistics for a specific P2P protocol, make a selection from the Protocol drop-down list.

To View P2P Data Statistics through the CLI

Enter the following command at the prompt:

```
SGOS# show p2p statistics
```

P2P Clients

You can view the total number of P2P client connections received in the last 60 minute, 24 hour, or 30 day period.

Note: The P2P client statistics are available only through the Management Console.

To View P2P Client Statistics through the Management Console

1. Select Statistics>P2P History>P2P Clients.

The P2P Clients tab displays.

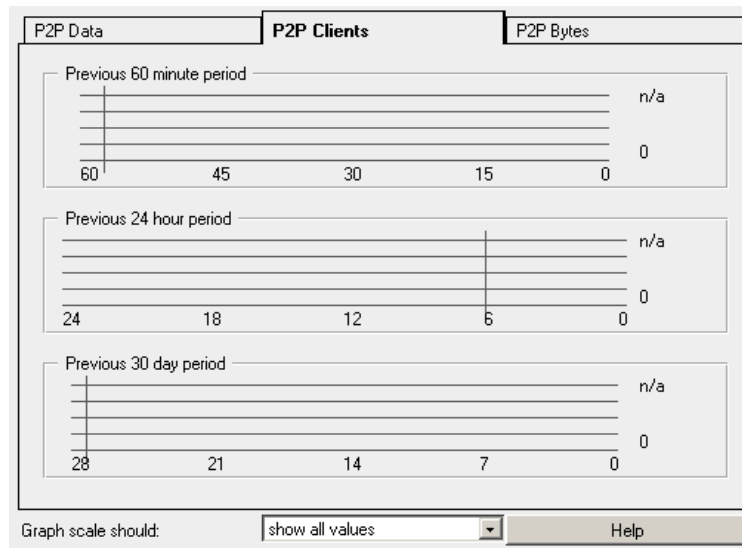


Figure 22-21: P2P Clients Tab

- (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

P2P Bytes

You can view the total number of bytes sent to and received from P2P clients in the last 60 minute, 24 hour, or 30 day period.

Note: The P2P bytes statistics are available only through the Management Console.

To View P2P Byte Statistics through the Management Console

- Select Statistics>P2P History>P2P Bytes.

The P2P Clients tab displays.

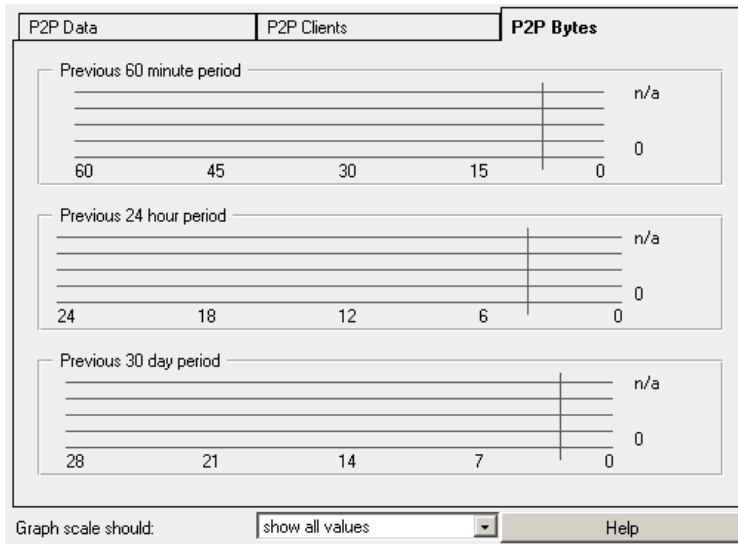


Figure 22-22: P2P Bytes Tab

2. (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

Streaming History Statistics

The Streaming History tabs (Windows Media, Real Media, and QuickTime) display bar graphs that illustrate the number of active client connections over the last 60 minutes, 24 hours, and 30 days. These statistics are not available through the CLI. The Current Streaming Data and Total Streaming Data tabs display real-time values for current connection and live traffic activity on the ProxySG. Current and total streaming data statistics are available through the CLI.

Viewing Windows Media Statistics

The Windows Media tab shows the number of active Windows Media client connections over the last 60 minutes, 24 hours, and 30 days.

To View Windows Media Client Statistics through the Management Console

1. Select Statistics>Streaming History>Windows Media.

The Windows Media tab displays.

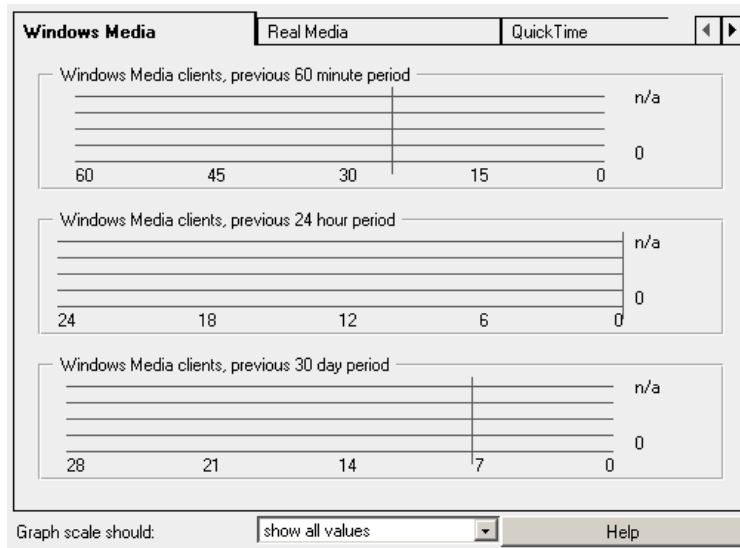


Figure 22-23: Windows Media Tab

2. (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

Viewing Real Media Statistics

The Real Media tab shows the number of active Real Media client connections over the last 60 minutes, 24 hours, and 30 days.

To View Real Media Data Statistics through the Management Console

1. Select Statistics>Streaming History>Real Media.

The Real Media tab displays.

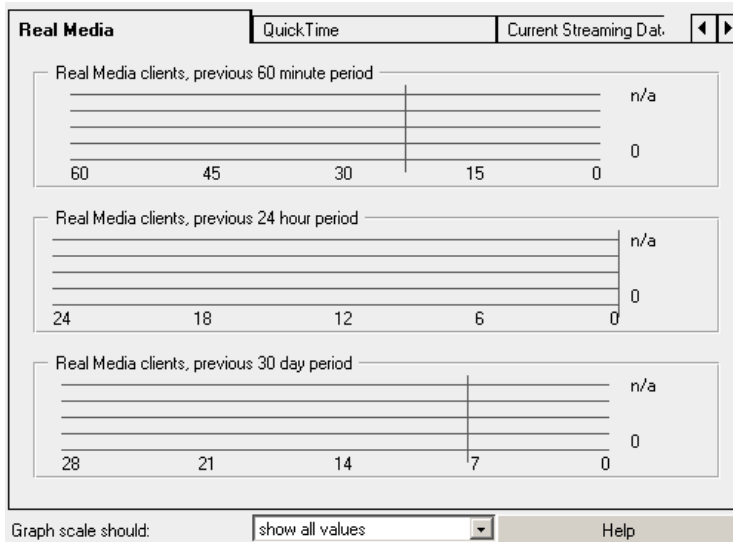


Figure 22-24: Real Media Tab

2. (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

Viewing QuickTime Statistics

The QuickTime tab shows the number of active QuickTime client connections over the last 60 minutes, 24 hours and 30 days.

To View QuickTime Data Statistics through the Management Console

1. Select Statistics>Streaming History>QuickTime.

The QuickTime tab displays.

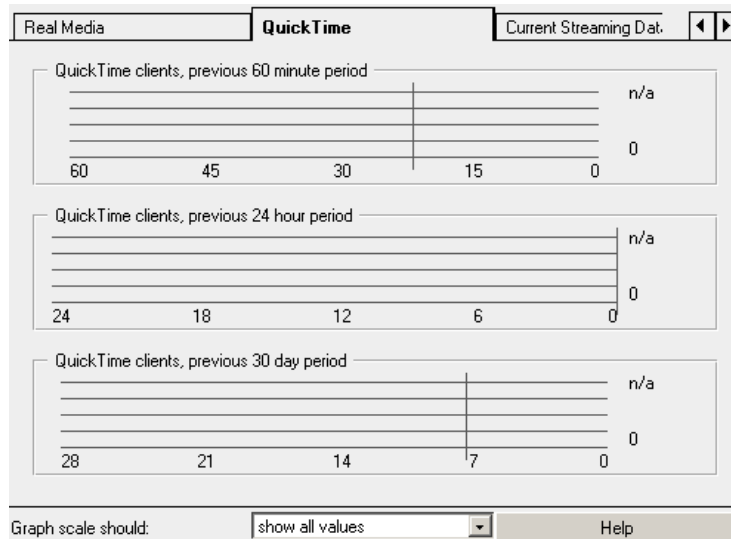


Figure 22-25: QuickTime Tab

- (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, and QuickTime activity on the ProxySG. These statistics can also be viewed through the CLI.

To View Current Streaming Data Statistics through the Management Console

- Select Statistics>Streaming History>Current Streaming Data.

The Current Streaming Data tab displays.

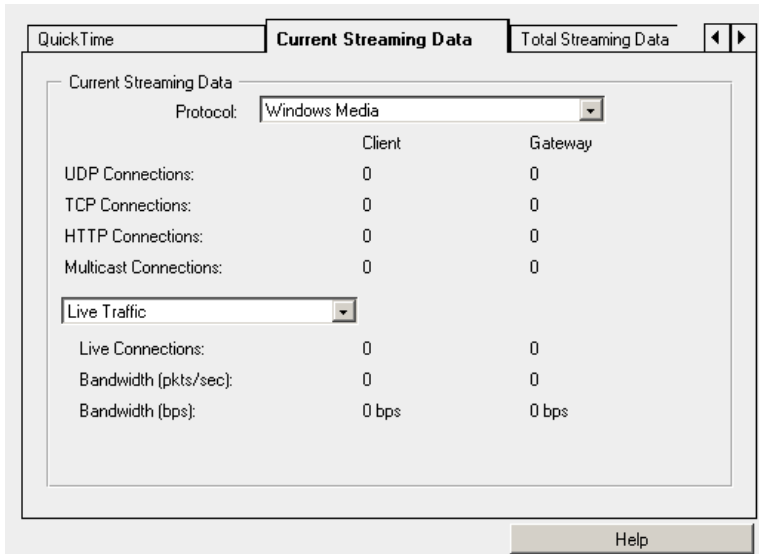


Figure 22-26: Current Streaming Data Tab

2. Select a streaming protocol from the Protocol drop-down list.
3. Select a traffic connection type (Live, On-Demand, or Pass-thru) from the drop-down list.

To View Total Streaming Data Statistics through the Management Console

1. Select Statistics>Streaming History>Total Streaming Data.

The Total Streaming Data tab displays.

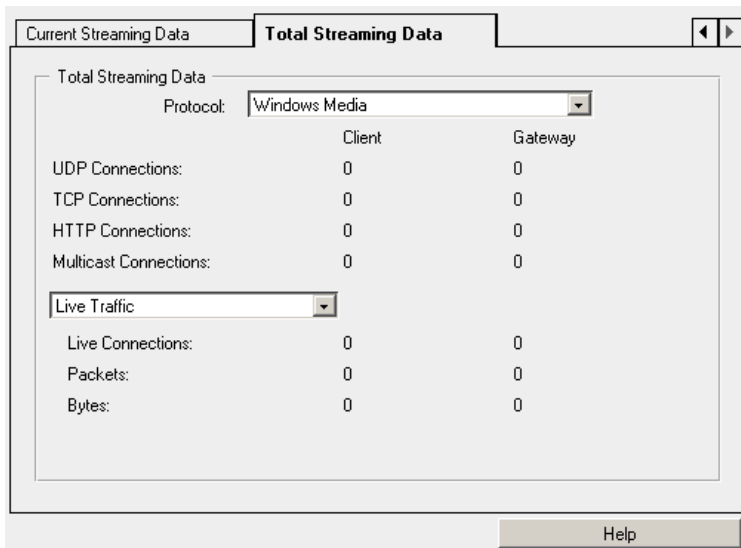


Figure 22-27: Total Streaming Data Tab

2. Select a streaming protocol from the Protocol drop-down list.
3. Select a traffic connection type (Live, On-Demand, or Passthru) from the drop-down list.

To View Current and Total Streaming Data Statistics through the CLI

Enter the following command at the prompt:

```
SGOS# show streaming {quicktime | real-media | windows-media} statistics
```

To Clear Streaming Statistics through the CLI

Enter the following command at the prompt:

```
SGOS# clear-statistics {quicktime | real-media | windows-media}
```

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

Select Statistics>SOCKS History>SOCKS Clients.

The SOCKS Clients tab displays.

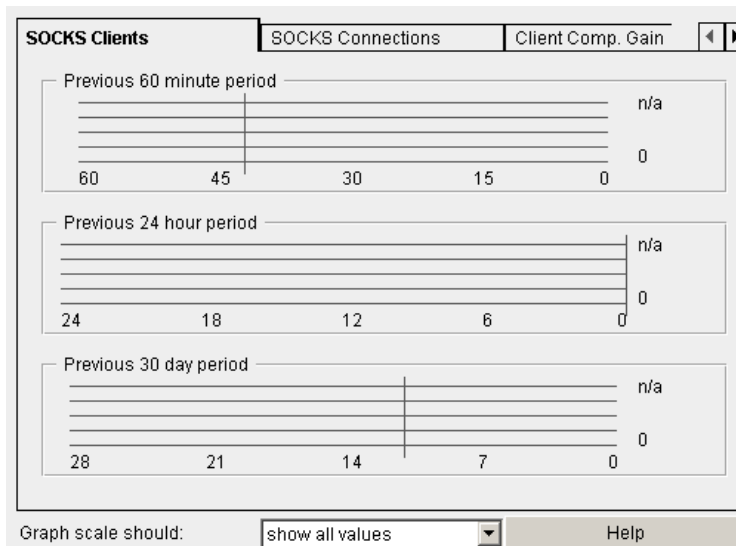


Figure 22-28: SOCKS Client Tab

Viewing SOCKS Connections

The SOCKS Connections tab displays SOCKS Connection data.

To View SOCKS Connection Data through the Management Console

Select Statistics>SOCKS History>SOCKS Connections.

The SOCKS Connections tab displays.

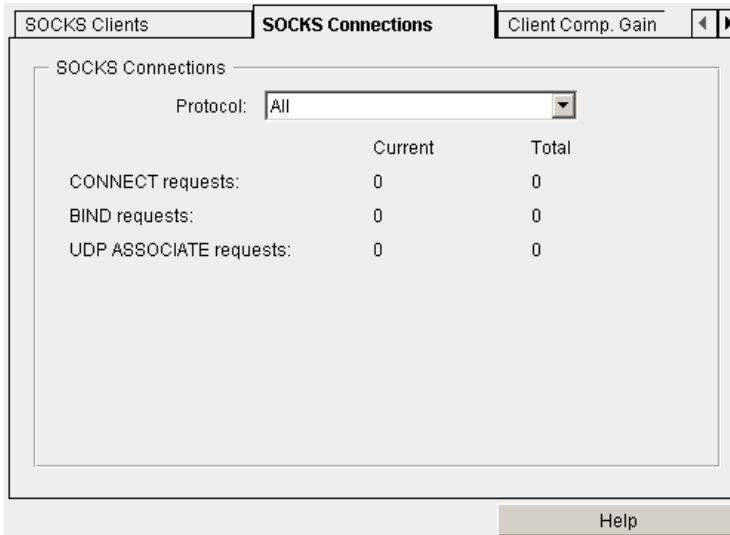


Figure 22-29: SOCKS Connections Tab

Viewing SOCKS Client and Server Compression Gain Statistics

Under SOCKS History, 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.

- SOCKS client/server uncompressed bytes are the bytes that would have been received or sent if compression were disabled.
- SOCKS client/server compressed bytes are the bytes that were actually received/sent when compression was enabled.

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.

To View SOCKS Client Compressed Gain Statistics through the Management Console

1. Select Statistics>SOCKS History>Client Comp. Gain.

The Client Comp. Gain tab displays.

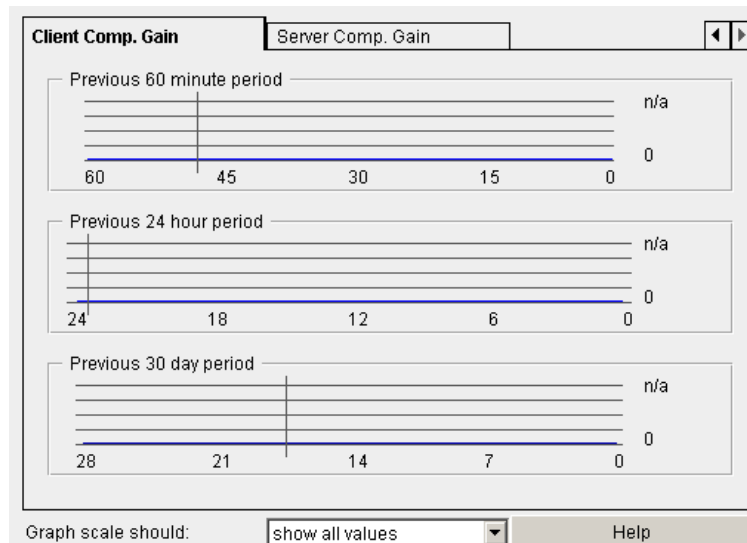


Figure 22-30: SOCKS Client Comp. Gain Tab

- (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

To View SOCKS Server Compressed Gain Statistics through the Management Console

- Select Statistics>SOCKS History>Server Comp. Gain.

The Server Comp. Gain tab displays.

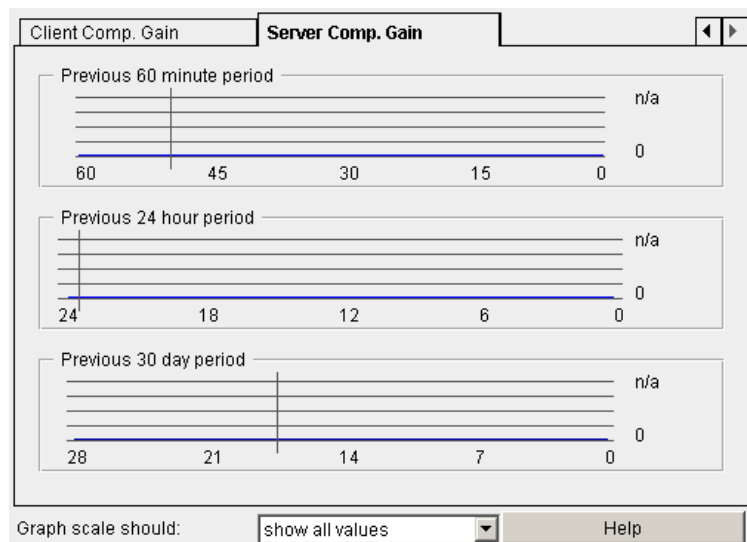


Figure 22-31: SOCKS Server Comp. Gain Tab

- (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

Shell History Statistics

The Shell History tab displays client connections on a per hour, per day, and per month basis.

Note: The Shell history statistics are available only through the Management Console.

To View Shell History Statistics through the Management Console

1. Select Statistics>Shell History.

The Shell Clients tab displays.




Figure 22-32: Shell Clients History Tab

2. (Optional) To set the graph scale to a different value, select a value from the Graph scale should drop-down list.

Resources Statistics

The Resources tabs (Disk Use, Memory Use, and Data) allow you to view information about how 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 Disk Use Statistics

The Disk Use tab shows the ProxySG disk usage. The fields on the tab are:

- System Objects—the percentage of storage resources currently used for non-access-log system objects.
- Access log—the percentage of storage resources currently used for the access log.

- Cache in Use—the percentage of non-system, non-access-log resources currently in use for cached objects.
- Cache available—the percentage of non-system, non-access-log resources still available for caching objects.

To View Disk Use Statistics through the Management Console

Select Statistics>Resources>Disk Use.

The Disk Use tab displays.

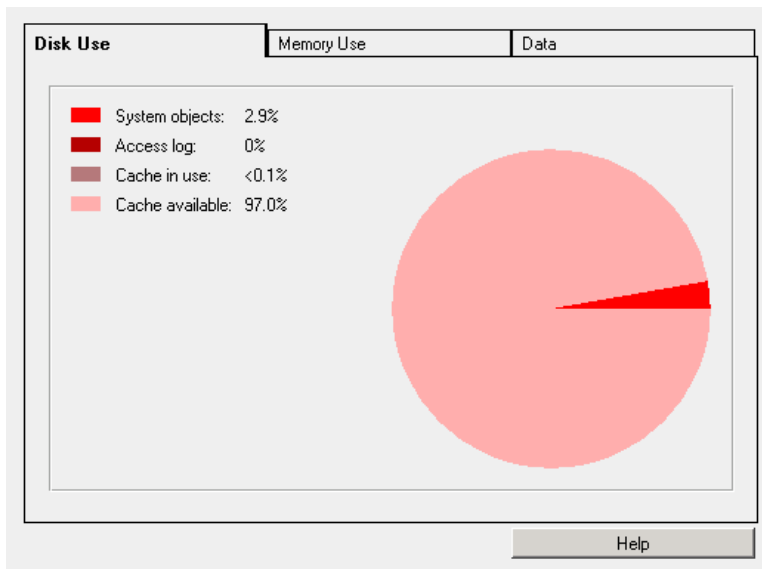


Figure 22-33: Disk Use Tab

Viewing Memory Use Statistics

The Memory Use tab shows the amount of memory used for RAM, the ProxySG itself, and for network buffers. The fields on the Memory Use tab are:

- RAM Cache—the amount of RAM that is used for caching.
- System allocation—the amount of RAM allocated for the device system.
- Network buffers—the amount of RAM currently allocated for network buffers.

To View Memory Use Statistics through the Management Console

Select Statistics>Resources>Memory Use.

The Memory Use tab displays.

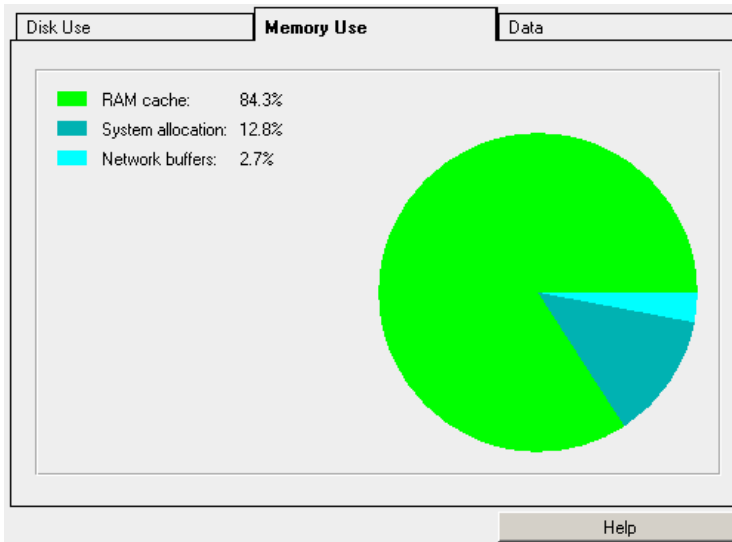


Figure 22-34: Memory Use Tab

Viewing Data Allocation Statistics in RAM and on Disk

The Data tab shows the total and available disk space and RAM, and how they are currently allocated. The fields on the Data tab are described below. This information can also be viewed through the CLI.

- Maximum objects supported—the maximum number of objects that can be supported.
- Cached objects—the number of objects that are currently cached.
- Disk used by system objects—the amount of disk space used by the system objects.
- Disk used by access log—the amount of disk space used for access logs.
- Total disk installed—the total amount of disk space installed on the device.
- RAM used by cache—the amount of RAM allocated for caching.
- RAM used by system—the amount of RAM allocated for system use.
- RAM used by network—the amount of RAM allocated for network use.
- Total RAM installed—the total amount of RAM installed.

To View Data Allocation Statistics through the Management Console

Select Statistics>Resources>Data.

The Data tab displays.

Disk Use	Memory Use	Data																				
<table> <tr> <td>Maximum objects supported:</td> <td>1,119,930 objects</td> </tr> <tr> <td>Cached Objects:</td> <td>0 objects</td> </tr> <tr> <td>Disk used by system objects:</td> <td>512.63 megabytes</td> </tr> <tr> <td>Disk used by access log:</td> <td>0 bytes</td> </tr> <tr> <td>Total disk installed:</td> <td>16.95 gigabytes</td> </tr> <tr> <td colspan="2"> </td> </tr> <tr> <td>RAM used by cache:</td> <td>653.83 megabytes</td> </tr> <tr> <td>RAM used by system:</td> <td>93.28 megabytes</td> </tr> <tr> <td>RAM used by network:</td> <td>20.88 megabytes</td> </tr> <tr> <td>Total RAM installed:</td> <td>768 megabytes</td> </tr> </table>			Maximum objects supported:	1,119,930 objects	Cached Objects:	0 objects	Disk used by system objects:	512.63 megabytes	Disk used by access log:	0 bytes	Total disk installed:	16.95 gigabytes			RAM used by cache:	653.83 megabytes	RAM used by system:	93.28 megabytes	RAM used by network:	20.88 megabytes	Total RAM installed:	768 megabytes
Maximum objects supported:	1,119,930 objects																					
Cached Objects:	0 objects																					
Disk used by system objects:	512.63 megabytes																					
Disk used by access log:	0 bytes																					
Total disk installed:	16.95 gigabytes																					
RAM used by cache:	653.83 megabytes																					
RAM used by system:	93.28 megabytes																					
RAM used by network:	20.88 megabytes																					
Total RAM installed:	768 megabytes																					
Help																						

Figure 22-35: Resources Data Tab

To View Data Allocation Statistics through the CLI

Enter the following command at the prompt:

```
SGOS# show resources
```

Efficiency Statistics

The Efficiency tabs (Summary, Non-cacheable, Access Pattern, and Data) allow you to see information about the flow of both cacheable and non-cacheable data through the ProxySG. You can also see information about how data is being served (such as, RAM, disk, origin).

Viewing the Cache Efficiency Summary

The Summary tab shows the percent of objects served from cache, the percent loaded from the network, and the percent that were non-cacheable. The data dates from the last device reset. The values shown are either objects served or bytes served, based on the Values reflect field at the bottom of the tab. The fields on the Summary tab are:

- Served from cache—the percentage of requests the device was able to serve from the cache.
- Loaded from source—the percentage of requests the device had to retrieve from the Web and was able to store in the cache.
- Non-cacheable—the percentage of requests for non-cacheable objects.

To View the Cache Efficiency Summary

1. Select Statistics>Efficiency>Summary.

The Summary tab displays.




Figure 22-36: Efficiency Summary Tab

- (Optional) To switch the view between objects served and bytes served, select either bytes served or objects served from the Graph shows drop-down list.

Viewing a Breakdown of Non-Cacheable Data

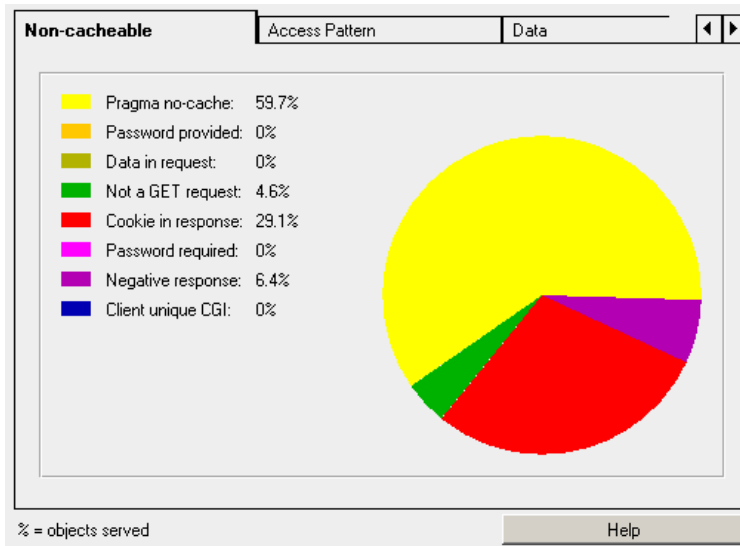
The Non-cacheable tab shows a breakdown of non-cacheable objects. It shows how many of the various types of non-cacheable requests have been handled. The non-cacheable request types are:

- Pragma no-cache—requests that specify non-cached objects, such as when a user clicks the refresh button in the Web browser.
- Password provided—requests that include a client password.
- Data in request—requests that include additional client data.
- Not a GET request—only the HTTP method GET request can be cached. These are all other methods (PUT, HEAD, POST, DELETE, LINK, and UNLINK).
- Cookie in response—responses that include an HTTP cookie.
- Password required—responses that require a client password.
- Negative response—failed responses, such as when a server or object is not available. This value is zero if the Cache Negative Responses option is enabled.
- Client unique CGI responses—unique responses generated by a CGI application for a specific client.

To View a Breakdown of Non-Cacheable Data

Select Statistics>Efficiency>Non-cacheable.

The Non-cacheable tab displays.



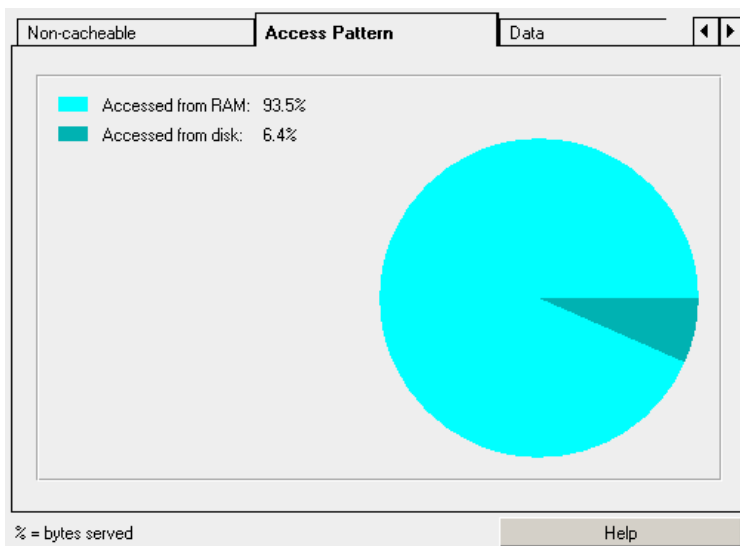
Viewing the Cache Data Access Pattern

The Access Pattern tab shows the number of cached requests served from RAM and disk. Cached objects are stored first in RAM. As time passes without additional requests for an object, the object is migrated to disk.

To View the Cache Data Access Pattern

Select Statistics>Efficiency>Access Pattern.

The Access Pattern tab displays.



Viewing Totals for Bytes Served

The Data tab lists a breakdown of all requests served. The fields are:

- Served from cache—the number of objects served from the cache.
- Loaded from source—the number of objects that could not be served from the cache and were retrieved from the Web.
- Non-cacheable—the number of objects served that could not be cached.
- Pragma no-cache—requests that specify non-cached objects, such as when a user clicks the refresh button in a Web browser.
- Password provided—requests that include a client password.
- Data in Request—requests that include additional client data.
- Not a GET request—requests that include an invalid HTTP method.
- Cookie in response—responses that include an HTTP cookie.
- Password required—responses that require a client password.
- Negative response—failed responses, such as when a server or object is not available. This information is only displayed if the Cache Negative Responses option is disabled.
- Client unique CGI—responses that contain unique CGI data.
- Accessed from RAM—the total number of bytes served from the RAM cache.
- Accessed from disk—the total number of bytes served from the disk cache.

To View Totals For Bytes Served

1. Select Statistics>Efficiency>Data.

The Data tab displays.

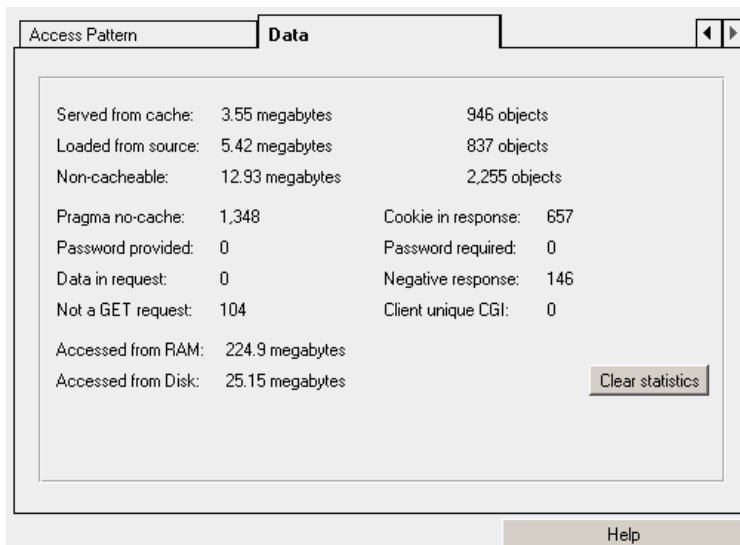


Figure 22-39: Efficiency Data Tab

- (Optional) To clear all statistics, click the Clear statistics button.

Contents Statistics

The Contents tabs (Distribution and Data) allow you to see information about objects currently stored or served organized by size. The cache contents include all objects currently stored by the ProxySG. The cache contents are not cleared when the ProxySG is powered off.

Viewing Cached Objects by Size

The Distribution tab shows the objects currently stored by the ProxySG, ordered by size.

To View the Distribution of Cache Contents

Select Statistics>Contents>Distribution.

The Distribution tab displays.

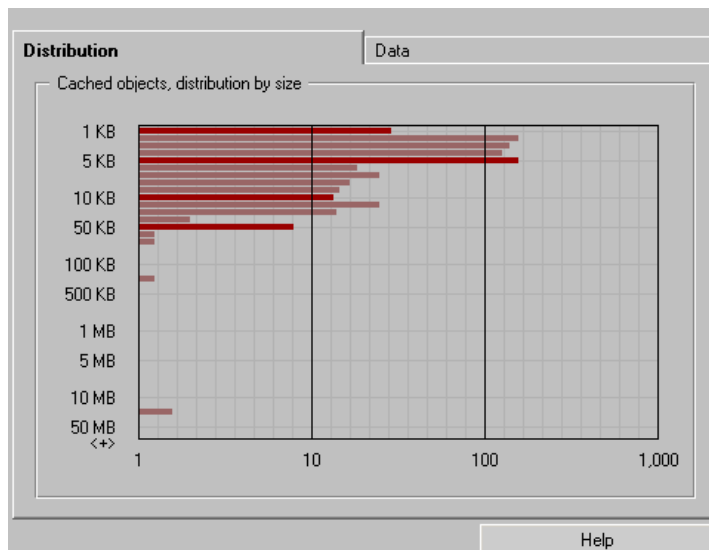


Figure 22-40: Contents Distribution Tab

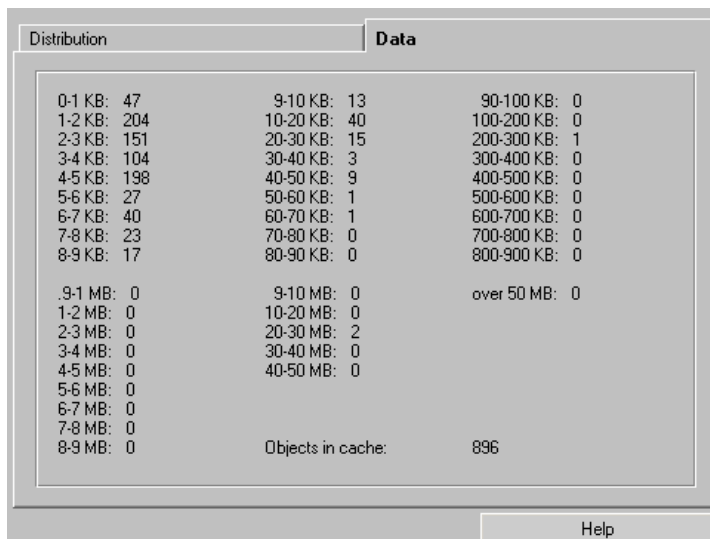
Viewing the Number of Objects Served by Size

The Data tab displays the number of objects served by the ProxySG, organized by size. This chart shows you how many objects of various sizes have been served.

To View the Number of Objects Served

Select Statistics>Contents>Data.

The Data tab displays.



Distribution		Data	
0-1 KB: 47	9-10 KB: 13	90-100 KB: 0	
1-2 KB: 204	10-20 KB: 40	100-200 KB: 0	
2-3 KB: 151	20-30 KB: 15	200-300 KB: 1	
3-4 KB: 104	30-40 KB: 3	300-400 KB: 0	
4-5 KB: 198	40-50 KB: 9	400-500 KB: 0	
5-6 KB: 27	50-60 KB: 1	500-600 KB: 0	
6-7 KB: 40	60-70 KB: 1	600-700 KB: 0	
7-8 KB: 23	70-80 KB: 0	700-800 KB: 0	
8-9 KB: 17	80-90 KB: 0	800-900 KB: 0	
.9-1 MB: 0	9-10 MB: 0	over 50 MB: 0	
1-2 MB: 0	10-20 MB: 0		
2-3 MB: 0	20-30 MB: 2		
3-4 MB: 0	30-40 MB: 0		
4-5 MB: 0	40-50 MB: 0		
5-6 MB: 0			
6-7 MB: 0			
7-8 MB: 0			
8-9 MB: 0			
	Objects in cache:	896	

Figure 22-41: Contents Data Tab

Event Logging

Viewing the Event Log

The event log contains all events that have occurred on the ProxySG. Configure the level of detail available by selecting Maintenance>Event Logging>Level (see "Configuring Which Events to Log" on page 798 for details).

To View the Event Log

1. Select Statistics>Event Logging.

The Event Viewer tab displays.

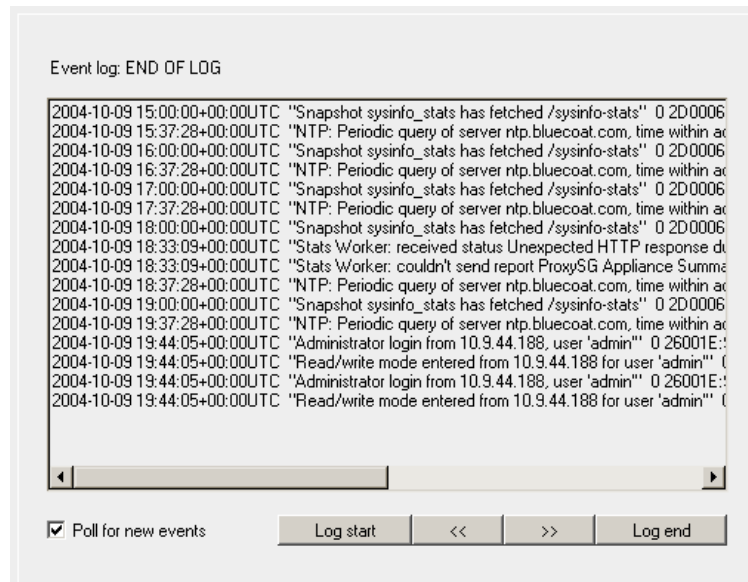


Figure 22-42: Event Viewer

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 checkbox to poll for new events that occurred while the log was being displayed.

Note: The Event Log cannot be cleared.

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.

Bandwidth management statistics are also available through the CLI.

Current Class Statistics Tab

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 through the Management Console

1. Go to Statistics>Bandwidth Management>Current Class Statistics.

The Current Class Statistics tab displays. The high level bandwidth classes and their statistics are visible.

Bandwidth Class	Current Packet Rate(pps)	Current Bandwidth(kbps)
● bridging	0	0
● ftp	0	0
● http	0	0
● im	0	0
● streaming	0	0
● p2p	0	0
● service-info	0	0
● socks	0	0
● tcp-tunnel	0	0

Figure 22-43: Current Class Statistics Tab

- 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 will close the folder.

Total Class Statistics Tab

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 through the Management Console

- Go to Statistics>Bandwidth Management>Total Class Statistics.

The Total Class Statistics tab displays. The high level bandwidth classes and their statistics are visible.

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

Figure 22-44: Total Class Statistics Tab

- To view the statistics of child bandwidth classes, double-click the folder icon of the parent class. A second double-click will close the folder.

Bandwidth Management Statistics in the CLI

To View Bandwidth Management Statistics through the CLI

- To view all bandwidth management statistics, enter the following commands at the prompt:

```
SGOS#(config) bandwidth-management
SGOS#(config bandwidth-management) view statistics
```

- To view the BWM statistics for a specific class, enter the following command at the (config) command prompt:

```
SGOS#(config bandwidth-management) view statistics bwm_class
```

Example

```
SGOS#(config bandwidth-management) view statistics http
Class Name:          http
Parent:              <none>
Minimum Bandwidth:  unspecified
Maximum Bandwidth:  unlimited
Priority:             0
Total Bytes:         0 bytes
Total Packets:       0 pkts
Dropped Packets:    0 pkts
Current Bandwidth:   0 kbps
Current Packet Rate: 0 pps
Queue Length:        0 bytes
```

where:

Parent	The class name of the parent of this class.
Minimum Bandwidth	The maximum bandwidth setting for this class.
Maximum Bandwidth	The minimum bandwidth setting for this class.
Priority	The priority level for this class.
Total Bytes	The total number of bytes served.
Total Packets	The total number of packets served.
Dropped Packets	Total number of packets dropped (packets in the queue that are dropped because the queue length is reached).
Current Bandwidth	Current bandwidth value (in kilobits per second).
Current Packet Rate	Current packets-per-second value.
Queue Length	Maximum length allowed for the queue of packets that lack available bandwidth but are waiting for bandwidth to become available.

To Clear Bandwidth Management Statistics through the CLI

1. To clear bandwidth management statistics for all bandwidth management classes, enter the following command at the prompt:
`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`

Access-Log Statistics

Access-log statistics can be viewed from the Management Console or the CLI, although not all statistics you can view in the Management Console are available in the CLI.

You can also view some access log statistics by going to Statistics>Advanced, and clicking on 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 logobject 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

This option is not available through the CLI.

To Display the Access Log Tail through the Management Console:

1. Select Statistics>Access Logging>Log Tail.

The Log Tail tab displays.

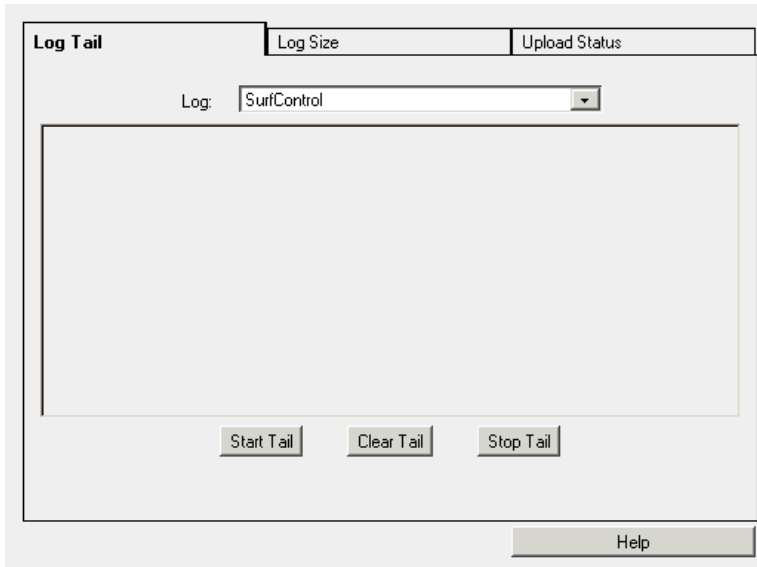


Figure 22-45: Viewing the Access Log Tail

2. From the Log drop-down list, select the log you want 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 22.2 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 22.2: 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.

Table 22.2: Log Writing Status Description

shutdown	The system is shutting down.
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 through the Management Console

1. Select Statistics>Access Logging>Log Size.

The Log Size tab displays.




Figure 22-46: Checking the Log Size Statistics

2. From the Log drop-down list, select the log you want to view.

Viewing Access Logging Status

The ProxySG 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 through the Management Console:

1. Select Statistics>Access Logging>Upload Status.

The Upload Status tab displays.

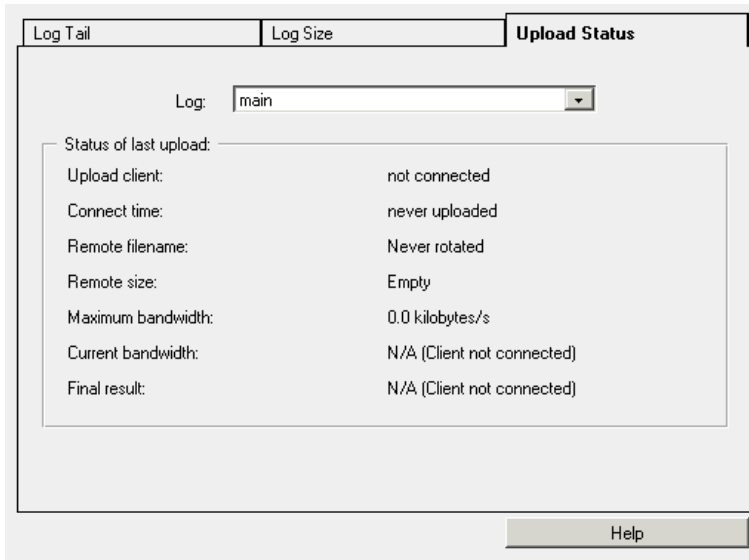


Figure 22-47: Viewing Upload Status Statistics

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 22.3: 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.

Viewing Access-Log Statistics through the CLI

In the CLI, you can view all access log statistics at once, or you can view the statistics of a specific access log. For details of the meaning of these statistics, see "Viewing the Log File Size" on page 854 and "Viewing Access Logging Status" on page 855.

To View Access Logging Statistics through the CLI

1. To view the statistics for all access logs at once, enter the following command:

```
SGOS# show access-log statistics
```

2. To view the statistics for a specific access log, enter the following command:

```
SGOS# show access-log statistics log_name
```

The statistics for the access log Main are displayed below as an example:

```
SGOS#(config) show access-log statistics main
Statistics:
Access Log (main) Statistics:
Log Manager Version 3
Log entry lifetime counter:      0
System Status:
  Log manager:                   enabled and running
  Upload client:                 disabled
  Log writer:                   idle
  Log reader:                   idle
Log Information:
  Current log size:              0 bytes
  Early upload threshold:        1736 MB
  Maximum log size:              2170 MB
  Max size policy:               stop logging
  Bytes in write buffer :        0
  Tail sockets in use :          0
  Modified time:                 2004-08-26 22:10:49+00:00UTC
Next Upload:
  Client type:                   ftp
  Next attempt:                  uploading disabled
  Connect type:                  daily upload
  Connect reason:                regular upload
  Estimated upload size:
  compressed:                    nothing to upload
  uncompressed:                  nothing to upload
  Upload format:                 gzip
Last Upload Attempt:
  Time:                           never uploaded
  Maximum bandwidth:             0.00 KB/sec
  Result:                         failure
Current/Last Upload File:
  Remote filename:               Never rotated
  Remote size:                   0 bytes
```

Failover Statistics

At any time, you can view statistics for any failover group you have configured on your system.

Viewing Failover Status

To View Failover Status

1. Go to Statistics>Failover.

The Status tab displays.

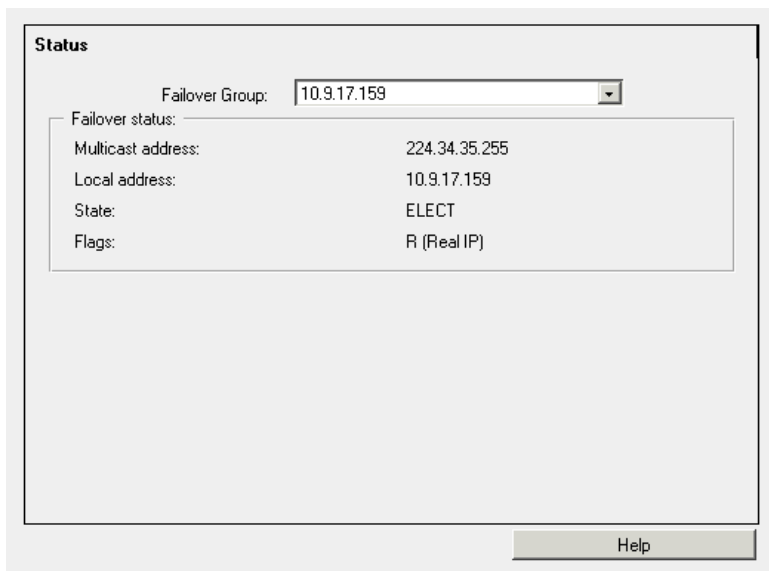


Figure 22-48: Failover Status Tab

2. From the drop-down list, select the group whose statistics you want 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.

Advanced Statistics

A variety of system statistics are conveniently located in one place and accessible by clicking the links listed in the Advanced tab of the Management Console.

To View System-Wide Advanced Statistics

1. Select Statistics>Advanced.

The Advanced tab displays.

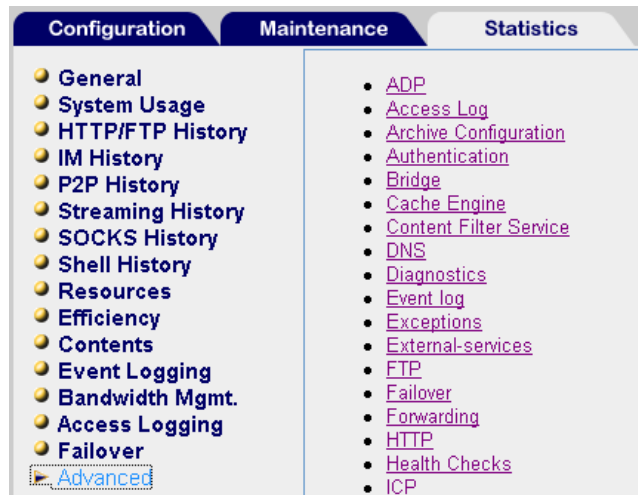


Figure 22-49: Advanced Tab

2. Click the appropriate link for the service you want to view.

A list of categories for that service will appear.

Note: If you upgraded from SGOS 2.x or CacheOS 4.x and have log files generated by those versions, you can view or retrieve them through the `Statistics>Advanced>Access Log>Show Old Logs URL`.

3. To view the statistics for a particular category, click that category's link.
A window will open detailing the relevant statistics.
4. Close the window when you have finished viewing the statistics.
5. To return to the list of links, either reselect `Statistics>Advanced` or click your browser's Back button.

Appendix A: Using the Authentication/Authorization Agent

The Blue Coat Systems Authentication and Authorization Agent (BCAAA) allows SGOS 4.x to manage authentication and authorization for NTLM, Netegrity SiteMinder realms, and Oblix COREid realms. Note that the agent is installed and configured separately from SGOS 4.x. The agent is available at the Blue Coat Web site.

- **NTLM:** The BCAA service does not talk directly to an NTLM server. The BCAA service must be installed on a domain controller or member server, allowing the ProxySG to access Windows NT Lan Manager (NTLM) domain controllers. The BCAA service authenticates users in all domains trusted by the computer on which it is running. A single installation of the BCAA service can support multiple ProxySG appliances.

Note: SSL is not supported on Windows NT 4.

- **SiteMinder and COREid:** When a SiteMinder or COREid realm is referenced in policy, a BCAA process is created. The ProxySG then sends a configuration request that describes the servers to use. The BCAA service logs in to the appropriate servers and determines configuration information to be passed back to the ProxySG (such as the kind of credentials required). Responses from the SiteMinder and COREid policy servers are translated into appropriate BCAA protocol responses and returned to the ProxySG.

Before you can use the BCAA service with SiteMinder or COREid, you must configure the appropriate ProxySG realm to work with the SiteMinder or COREid servers. The realm can be configured from the SiteMinder or COREid configuration tabs in the Management Console or from the CLI.

For specific information about configuring the SiteMinder realm to work with the Netegrity policy servers, see "Section F: Netegrity SiteMinder" on page 321 in Chapter 9: "Using Authentication Services". For specific information about configuring the COREid realm to work with Oblix Access Servers, see "Section G: Oblix COREid" on page 336.

Important: The BCAA service release version should match the released version of the operating system.

You can also use the BCAA service in place of the deprecated CAASNT application for SGOS 2.x and SGOS 3.1.x. You cannot use CAASNT with SGOS 3.2 and higher.

Operating system requirements are:

- **NTLM:** Windows® NT 4 (SP6 or greater and Windows installer 2.0 or greater), Windows 2000 or later. SSL is not supported on Windows NT 4.
- **SiteMinder and COREid:** Windows 2000 or later or Solaris™ 5.8 or 5.9.

The appendix discusses:

- "Installing the BCAA Service on a Windows or Windows NT System"
- "Installing the BCAA Service on a Solaris System"

- "Troubleshooting Authentication Agent Problems"
- "Common BCAA Event Messages"

Installing the BCAA Service on a Windows or Windows NT System

All images in this section are from a Windows 2000 system.

Note: If you have an existing CAASNT service on your system, it will be stopped and deleted as part of the BCAA installation procedure.

To Install the Authentication Agent:

1. Download the file from the Blue Coat download site at <https://download.bluecoat.com/>
2. Launch the install wizard.



Figure A-1: BCAA Installation Wizard Launch

3. Click Next to select the destination folder.

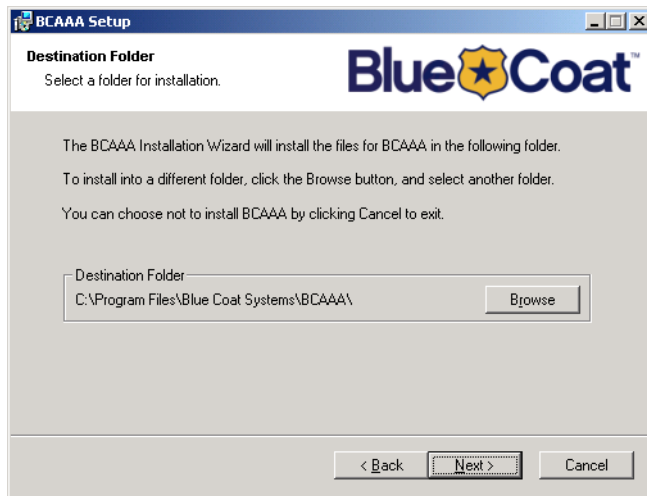


Figure A-2: Destination Folder for BCAA Application

Note: When doing an upgrade from one version of BCAA to another version of BCAA, you must install into the previous BCAA folder to retain your settings. If you install to a different folder, a new .ini file will be created with default settings.

When upgrading from CAASNT to BCAA, the settings from CAASNT are copied to the new installation directory.

4. Click Browse if you want to choose a different destination folder for the BCAA service.
5. Click Next to accept the default and select the port number.

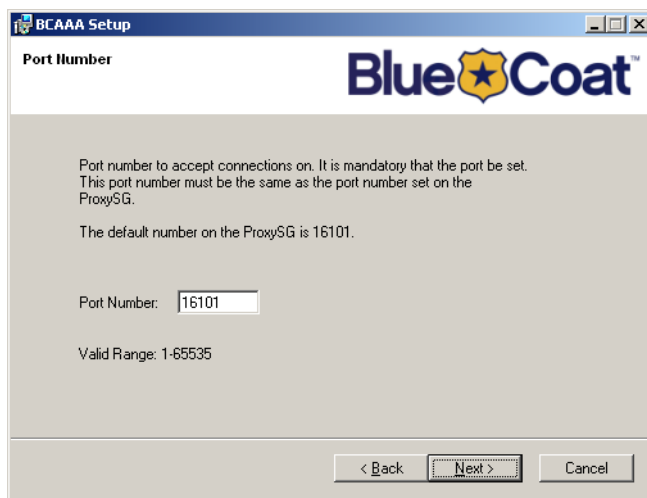


Figure A-3: Selecting BCAA Port Number

6. The port number must match the port number you specify on the ProxySG for the BCAA service. The default is 16101.

7. Click Next to select the number of threads.

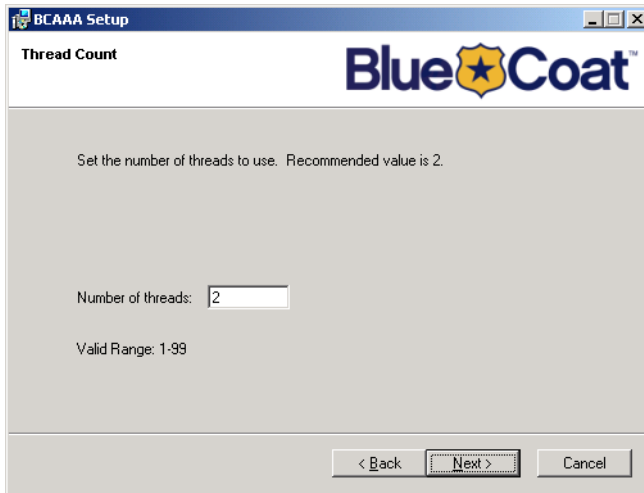


Figure A-4: Thread Count

8. The recommended (and default) value is 2. The maximum number of threads allowed is 99 per ProxySG. After selecting the number, click Next to specify the SSL requirements.

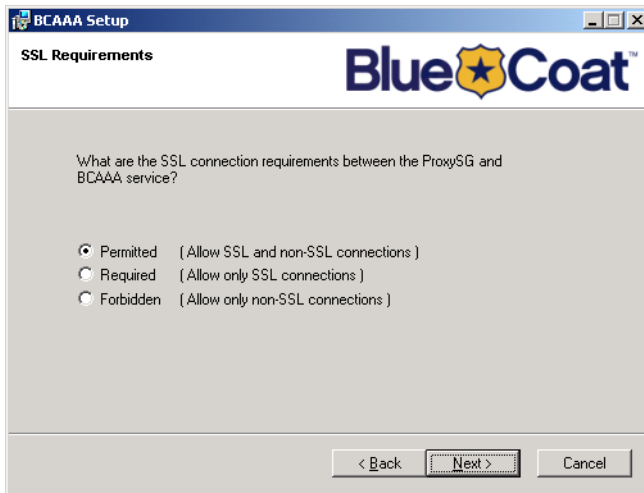


Figure A-5: SSL Requirements

9. The default is that SSL is Permitted, allowing both SSL and non-SSL connections. This setting must be compatible with the setting on the ProxySG.

Note: If you are installing the BC AAA service on an NT4 system, this screen is not available.

10. Click Next to specify the subject of the SSL certificate.

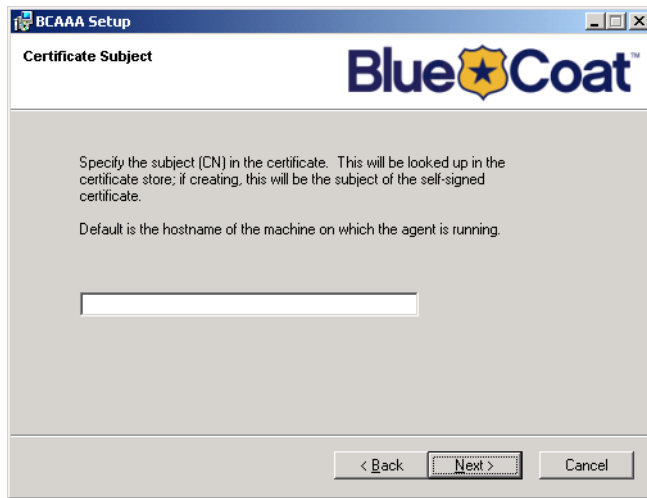


Figure A-6: Specifying the Subject of the Certificate

11. Specify the subject of the certificate.

Note: If you are installing the BCAAA service on an NT4 system, this screen is not available.

The BCAAA service looks up the specified subject in the service's certificate store. If it finds the subject, it uses it instead of generating a new certificate. If not, it generates a self-signed certificate with that subject. This generated certificate can be saved (as specified on the next screen).

12. Click Next to specify save options for the certificate.

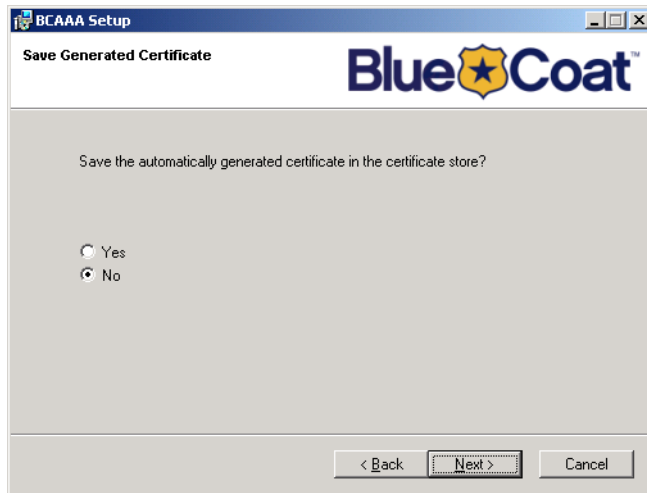


Figure A-7: Saving the Generated Certificate

Note: If you are installing the BCAAA service on an NT4 system, this screen is not available.

13. Click Next to specify whether the ProxySG must provide a valid certificate when connecting to the BCAA service.

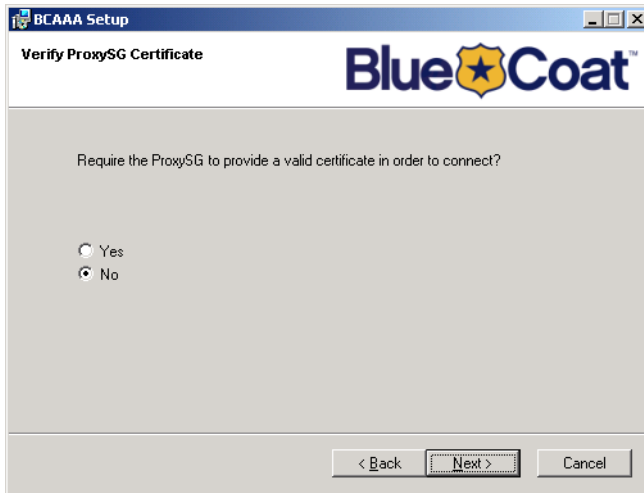


Figure A-8: Verify ProxySG Certificate

14. To force the ProxySG to provide a valid certificate to connect to the BCAA service, select the Yes radio button. The default is No.

Note: If you are installing the BCAA service on an NT4 system, this screen is not available.

15. Click Next to view the summary of the changes you made.

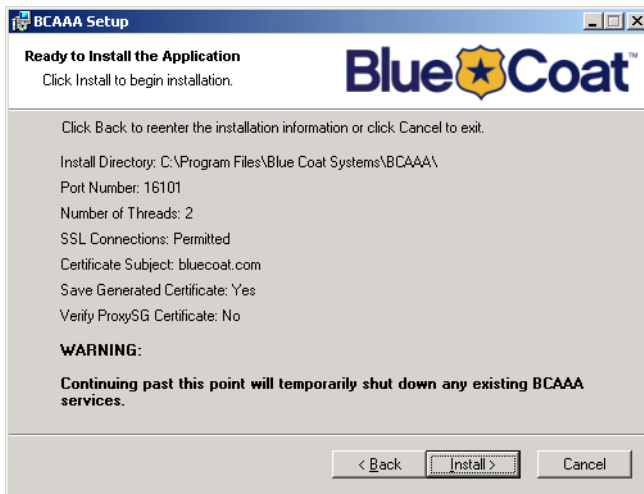


Figure A-9: BCAA Summary

16. Click Install to install the BCAA service using the settings you configured. When installation completes, the final BCAA screen displays.



Figure A-10: Completing BCAA Installation

To Modify Settings or Uninstall the Authentication Agent:

1. Launch the install wizard.

The Application Maintenance page displays.

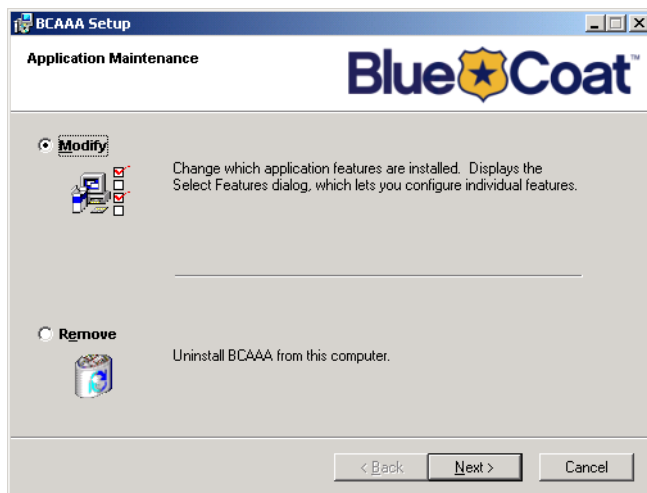


Figure A-11: Applications Maintenance Page

2. Click Modify to re-enter the installation wizard; click Remove to uninstall the BCAA service from the system

Note: For instructions on using the installation wizard, see "Installing the BCAA Service on a Windows or Windows NT System" on page 862.

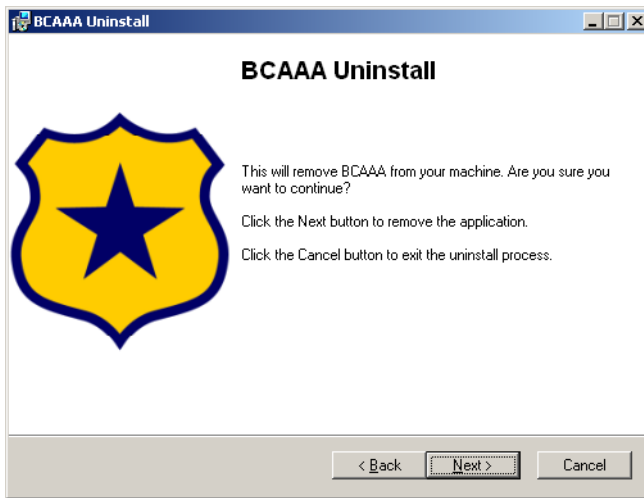


Figure A-12: Uninstalling the BCAA Service

3. Click Next to start the procedure.
4. Click Finish to exit the uninstall application.

To View the Application Event Log:

The BCAA service logs all errors to the Windows 2000 Application Event Log under the name BCAA.

1. Launch the Event Log.
2. Doubleclick the information message BCAA service to see that the BCAA service has been automatically started.




Figure A-13: BCAA Information Message

To View the BCAA Service:

The BCAA service logs all errors to the Windows 2000 Application Event Log under the name BCAA.

1. Launch the Event Viewer.
2. Right-click on BCAA and select Properties to manage the service. For example, to make the BCAA service start only manually, set the Startup Type to Manual. (Automatic is the default setting.)

Installing the BCAA Service on a Solaris System

To install the BCAA service on Solaris, complete the following instructions. You must be `root` to complete installation.

1. Download the shell script to your system.
2. Execute the shell script:


```
# sh bcaaa-version_number-SOLARIS-install.sh
```
3. Answer the questions to install the service on your Solaris system. A sample session is shown below:

```
Enter a path to a scratch directory [/tmp]:
Install Blue Coat Systems Authentication and Authorization Agent (BCAAA)? (y/n)y
Enter user that should own the installed files [root]
Enter group for the installed files [root]
/usr/local/bin/bcaaa installed
/usr/local/bin/bcaaa-99 installed
Libraries installed in /usr/local/lib/BlueCoatSystems/
/usr/local/etc/bcaaa.ini installed

If you use inetd, append the following line to /etc/services

bcaaa          16101/tcp          # Blue Coat Systems Authentication
Agent

If you use inetd, append the following line to /etc/inetd.conf, then signal inetd
to re-read the configuration file

If you use something else, make the equivalent changes

bcaaa stream tcp nowait root /usr/local/bin/bcaaa bcaaa -c
/usr/local/etc/bcaaa.ini

Installation complete
```

Troubleshooting Authentication Agent Problems

This section describes some common problems you might encounter when setting up or using the BCAA service on a Windows platform.

To troubleshoot the BCAA service, launch the event viewer.

The Properties pane displays, providing information about the status of the BCAA service at that time. What's of most interest to you is the Type and the Event ID. The description below the Type/Event ID lists the problem. You can often find more information about the problem and suggestions for its solution in "Common BCAA Event Messages" on page 870.

Common problems:

- If an attempt to start the BCAA service is issued when BCAA is already started, the following error message displays:

The requested service has already been started.

- If another application is using the same port number as the BCAA service, the following messages are displayed:

The BCAA service could not be started.

A system error has occurred.

System error 10048 has occurred.

Only one usage of each socket address (protocol/network address/port) is normally permitted.

Common BCAA Event Messages

Following are the most common event messages that can be logged to the Windows 2000 Application Event Log. Most of the event messages not listed here are error status messages returned by Win32 function calls. When a Win32 call fails, the error code and error text containing the reason for the error displays in the event log under the name BCAA.

To View the BCAA Event Log:

1. Right click on My Computer and select Manage.

The Computer Management window displays.

2. Select System Tools>Event Viewer>Application.

Note: When running under Windows NT4, the error text is unavailable for some error codes. In these cases, only the error code displays. This limitation does not apply to Windows 2000.

For each BCAA event message, the event message is displayed along with the event number.

Table A.1: BCAA Event Messages

Message ID	Message	Description
200	Various messages	The associated message provides information about a condition that is not an error.
300	Various messages	The associated message warns about an unexpected condition that does not prevent operation.
400	Various messages	The associated message describes an error condition that prevents normal operation.
1001	Authentication Agent service started: port=# threads=# socket=0x# process id=# agent version=# ProxySG Appliance version=#	This indicates successful startup and provides information about the agent.

Table A.1: BCAA Event Messages (Continued)

Message ID	Message	Description
1002	Authentication Agent stopped	This indicates normal shutdown of the service.
1003	ProxySG Appliance (a.b.c.d) connected; Process # spawned as #	This indicates a ProxySG has connected to the agent (Windows only).
1004	ProxySG Appliance agent process exited (normal logout)	This indicates normal logout by a ProxySG.
1005	Process %d has terminated, ExitCode=0x#, link=0x#	This indicates an unexpected termination of an agent process (Windows only).
1006	Service dispatcher exited.	This indicates an unexpected termination of the service dispatcher.
1007	CreateNamedPipe failed, pipe='%s'	The agent dispatcher could not create the named pipe for the reason given.
1008	ConnectNamedPipe failed, pipe='%s'	The agent process could not obtain the information from the dispatcher on the named pipe for the reason given.
1009	WriteFile failed, pipe='%s'	The dispatcher could not write information to the named pipe for the reason given.
1011	CreateThread (ProcessTimerThread) failed	The dispatcher could not create its timer thread.
1012	Failed to create ProxySG Appliance process '%s'	The dispatcher could not create an agent process.
1015	Various	Too many groups were configured in policy on the ProxySG, or the total length of the group names was too long.
1019	Various	The dispatcher was unable to determine the exit status of an agent process.
1020	Terminating ProxySG Appliance process #, ProcNum=# Handle=0x#	An agent process was active when the Windows service was shut down.
1022	Various	The associated message reports the status of a ProxySG login attempt.
1101	BasicAuth: CloseHandle failed; user xx\ \xx'	The agent was unable to close the login handle for the specified user.
1102	Username: '%s\ \s' too long	The ProxySG offered the specified username, which is too long.
1106	Various	An attempted authentication using BASIC credentials failed for the reason given.
1107	User Right 'Act as part of the operating system' required for Basic Authentication	The agent does not have the necessary privileges to do BASIC authentication
1108	Various	The agent was unable to determine information about the user for the reason given.
1202	Unable to create GroupsOfInterest mutex xx' - already exists	The agent could not create the Windows mutex needed for group authorization checks because it already exists.
1203	Unable to create GroupsOfInterest mutex xx	The agent could not create the Windows mutex needed for group authorization checks.

Table A.1: BCAA Event Messages (Continued)

Message ID	Message	Description
1204	OpenMutex failed for AuthGroups mutex '%s', group='%s'	The agent was unable to open the Windows mutex needed for group authorization checks.
1205	Various	The agent was unable to close the Windows mutex named for the reason given.
1207	GetAclInformation failed	The agent was unable to obtain ACL information needed to do group authorization checks.
1209	GetKernelObjectSecurity failed for AuthGroup='%s'	The agent was unable to obtain security information about the specified group.
1210	SetKernelObjectSecurity failed	The agent was unable to set up security information for the reason specified.
1211	InitializeSecurityDescriptor failed	The agent was unable to initialize the security descriptor for the reason specified.
1212	GetSecurityDescriptorDacl failed	The agent was unable to get the discretionary access control list (DACL) for the reason specified.
1213	SetSecurityDescriptorDacl failed	The agent was unable to set the discretionary access control list (DACL) for the reason specified.
1214	InitializeAcl failed	The agent was unable to initialize the access control list (ACL) for the reason specified.
1215	GetUserName failed for AuthGroup='%s'	The agent was unable to determine the username while processing the specified group.
1217	GetAce failed for AuthGroup='%s'	The agent was unable to get the access control entry (ACE) for the specified group.
1218	AddAce failed	The agent was unable to add the necessary access control entry (ACE) for the reason specified.
1219	AddAccessAllowedAce failed	The agent was unable to add the necessary "access allowed" access control entry (ACE).
1220	Could not establish groups-of-interest: result=0x##	The agent was unable to initialize groups-of-interest checking.
1221	AuthGroup '%s' does not exist	The specified group does not exist.
1222	NtLm RevertSecurityContext failed, user='%s'	The agent could not revert the security context for the specified user.
1223	BASIC: RevertToSelf failed, user='%s'	The agent could not revert the security context for the specified user.
1224	Error calling OpenProcessToken	The agent's call to OpenProcessToken failed for the specified reason.
1225	Error calling LookupPrivilegeValue	The agent could not get information about a needed privilege.
1226	Error calling AdjustTokenPrivileges	The agent could not adjust its privileges as required.
1227	ImpersonateLoggedOnUser failed; Group access denied for user '%s'	The agent could not impersonate the specified user.
1228	NtLm: ImpersonateSecurityContext failed; Group access denied for user '%s'	The agent could not impersonate the specified user.
1301	NOTE: Pending ContextLink=### timed out; deleting SecurityContext h=## TS=## now=##	The ProxySG did not provide a response to a challenge quickly enough.

Table A.1: BCAA Event Messages (Continued)

Message ID	Message	Description
1302	Various	An authentication request from a ProxySG referenced an in-progress request that has timed out or does not exist.
1304	Various	The agent was unable to delete a security context for the reason given.
1305	AcceptSecurityContext failure, SEC_E_INVALID_HANDLE, ContextLink=### count=#	The agent was provided with an invalid context handle.
1306	Various	The client provided an invalid token to the authentication system.
1308	AcceptSecurityContext failure, ContextLink=# count=#, detail=#(xxx)	Windows rejected the authentication attempt for the reason given.
1310	Various	This records the failure of NTLM authentication or group authorization.
1311	3:Failed NTLM Authentication for user: '%s'	This records the failure of NTLM authentication; the user name was supplied by the client.
1312	Various	The agent could not determine the username from the NTLM type 3 message supplied by the client.
1313	Invalid Type3 message	The client provided an NTLM type 3 message that was invalid.
1314	BASE64_Decode: Length of token exceeds max (%d)	The client provided an NTLM token that was too long.
1316	Unsupported version in request: %d(0x%x)	The ProxySG sent a request with an unsupported version number.
1401	Various	The agent lost communication with the ProxySG.
1403	Various	The agent is aborting for the reason given.
1402	Unexpected thread 0 exit	The agent exited unexpectedly.
1404	Unable to get ProcessInfo from parent process.	The agent could not obtain its information from the dispatcher.
1405	CreateFile failed, pipe='xx'	The agent could not create a handle for the dispatcher's named pipe.
1406	WaitNamedPipe failed, pipe='%s'	The agent could not wait for the dispatcher's named pipe.
1407	ReadFile failed, pipe='%s'	The agent could not read information from the dispatcher's named pipe.
1409	Various	The agent could not create the specified thread for the reason given.
1412	Various	The agent could not create a required Windows event object.
1413	AuthMethod 'xxs' not supported: returning _AuthResult=0x##	The ProxySG requested an unsupported authentication mechanism.
1414	Various	The specified request is unsupported.
1500	Various	The agent has a problem with memory allocation; typically this means there is not enough memory.

Table A.1: BCAA Event Messages (Continued)

Message ID	Message	Description
1501	Unable to allocate memory for ProcLink buffer.	The agent could not allocate some needed memory.
1502	Unable to allocate memory for ContextLink buffer.	The agent could not allocate some needed memory.
1503	Various	The agent was unable to allocate needed memory.
1604	Service dispatch failed	The Windows service dispatcher failed to start.
1605	RegisterServiceCtrlHandler failed	The agent dispatcher was unable to register the service control handler.
1608	SetServiceStatus failed, g_StatusHandle=%d	The agent was unable to set the service's status.
1610	Unsupported service control code: #	Windows sent a service control code that the agent does not support.
1701	WSASocket failed	The agent could not create a Windows socket for the reason given.
1702	WSAStartup failed.	The agent could not start the Windows socket for the reason given.
1703	Various	The agent could not send data to the ProxySG for the reason given.
1704	Various	The agent could not receive data from the ProxySG for the reason given.
1705	accept failed	The agent dispatcher could not initialize to accept new connections.
1706	bind failed, PortNumber=#	The agent dispatcher could not bind to the specified port.
1707	listen failed.	The agent dispatcher could not listen for new connections.
1708	Various	Windows reported an event wait failure to the agent while doing I/O on the socket.
1709	The agent is already running or the agent's port # is in use by another process	Some other process is already using the port needed by the agent.
1710	WSARecv failed reading bytes from socket	Windows reported an error when the agent tried to receive bytes from the ProxySG.
1711	WSASend failed sending bytes to socket.	Windows reported an error when the agent tried to send bytes to the ProxySG.
1712	Various	A socket I/O operation did not complete successfully.
1801	Error calling AcquireCredentialsHandle	The agent could not acquire its credentials from Windows.
1803	Various	The agent could not load a needed library (DLL).
1804	Various	The agent could not locate the needed services in a library (DLL).
1805	Unsupported SSPI Windows platform; PlatformId=#	The reported Windows platform is not supported for NTLM authentication.
1806	Error calling QueryContextAttributes	The agent could not determine the authenticated user's security attributes.

Table A.1: BCAA Event Messages (Continued)

Message ID	Message	Description
1807	QuerySecurityPackageInfo failed	The agent could not get needed security information from Windows.
1808	Max Token size too long (#); max size is #	The client supplied an NTLM token that is too long.
1809	FreeContextBuffer failed	An attempt to free the NTLM context buffer failed.
1811	Username 'x\\y' too long	The reported user name is too long.
1901	Admin Services Error: Access denied to domain/user/group information	The agent was unable to access necessary information.
1902	Admin Services Error: Invalid computer from which to fetch information	The computer to be used to get security information is invalid.
1903	Admin Services Error: Group not found	The requested group could not be found.
1904	Various	The reported error was encountered while browsing.
1905	Admin services error: could not translate context to Unicode	The requested object for browsing could not be translated to Unicode
1906	Admin service out of memory	The browsing service ran out of memory.
1907	Search request object too long: # > #	The requested object for browsing is too long.
2000	AcquireCredentialsHandle failed: 0x#	The agent could not acquire the credentials needed for an SSL session.
2001	Various	The agent was unable to negotiate an SSL session for the reason given.
2002	Various	An I/O error occurred during an SSL session .
2003	Various	The specified cryptographic error occurred during an SSL session.
2004	Various	The specified problem occurred with a certificate during SSL negotiation.

Appendix B: Access Log Formats

The ProxySG can create access logs in one of the following formats:

- "Custom or W3C ELFF Format"
- "SQUID-Compatible Format"
- "NCSA Common Access Log Format"

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 Table B.1.

Table B.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. Note that ELFF does not support character strings and require a space between fields, unlike the Blue Coat custom format.

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 Table B.2.

Table B.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.
%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.
%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 appliance take to process this request.
%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.

Table B.2: Blue Coat Custom Format and Extended Log File Format (Continued)

Blue Coat Custom Format	Extended Log File Format	Description
%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.]
%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	Service 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 may change a little.

Action Field Values

Table B.3 describes the possible values for the action field.

Table B.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.
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_AUTH_HIT	The requested object requires upstream authentication, and was served from the cache.
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_REDIRECT	The client was redirected to another URL for authentication.
TCP_CLIENT_REFRESH	The client forces a revalidation with the origin server with a <code>Pragma: no-cache</code> . If the server returns 304 Not Modified, this appears in the <code>Statistics:Efficiency</code> file as <code>In Cache, verified Fresh</code> .
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_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_NC_MISS	The object returned from the origin server was non-cacheable.
TCP_PARTIAL_MISS	The object is in the cache, but retrieval from the origin server is in progress.
TCP_POLICY_REDIRECT	The client was redirected to another URL due to policy.

Table B.3: Action Field Values (Continued)

Value	Description
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_MISS	A GIMS request to the server was forced and new content was returned.
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_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).
UDP_	Refers to requests on the ICP port (3130).
UDP_DENIED	Access was denied for this request.
UDP_HIT	A valid copy of the requested object was in the cache. This value is also used with ICP queries.
UDP_INVALID	The ICP request was corrupt, short, or otherwise unintelligible.
UDP_MISS	The requested object was not in the cache. This value is also used with ICP queries.
UDP_MISS_NOFETCH	An ICP request was made to this cache for an object not in the cache. The requestor was informed that it could not use this cache as a parent to retrieve the object. (This is not supported at this time.)
UDP_OBJ	An ICP request was made to this cache for an object that was in cache, and the object was returned through UDP. (This is not supported at this time. This functionality is deprecated in the current ICP specification.)

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 Table B.4.

Table B.4: Common Log Format Entries

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

Table B.5 details the specifiers for the access log upload filenames.

Table B.5: Specifiers for the 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).
%l	The fourth part of the ProxySG's IP address, using three digits (001.002.003.004)
%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).
%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

The following table lists all fields available for creating access log formats. When creating an ELFF format, you must use the values from the ELFF column. When creating a custom format, you can use values from the ELFF, CPL, or custom column.

Table B.6: Access Log Format Fields

Category: bytes			
ELFF	CPL	Custom	Description
cs-bodylength			Number of bytes in the body (excludes header) sent from client to appliance

Table B.6: Access Log Format Fields (Continued)

cs-bytes		%B	Number of bytes sent from client to appliance
cs-headerlength			Number of bytes in the header sent from client to appliance
rs-bodylength			Number of bytes in the body (excludes header) sent from upstream host to appliance
rs-bytes			Number of bytes sent from upstream host to appliance
rs-headerlength			Number of bytes in the header sent from upstream host to appliance
sc-bodylength			Number of bytes in the body (excludes header) sent from appliance to client
sc-bytes		%b	Number of bytes sent from appliance to client
sc-headerlength			Number of bytes in the header sent from appliance to client
sr-bodylength			Number of bytes in the body (excludes header) sent from appliance to upstream host
sr-bytes			Number of bytes sent from appliance to upstream host
sr-headerlength			Number of bytes in the header sent from appliance to upstream host
Category: connection			
ELFF	CPL	Custom	Description
cs-ip	proxy.address		IP address of the destination of the client's connection
c-connect-type			The type of connection made by the client to the appliance -- 'Transparent' or 'Explicit'
c-dns		%h	Hostname of the client (uses the client's IP address to avoid reverse DNS)
x-cs-dns	client.host		The hostname of the client obtained through reverse DNS.
c-ip	client.address	%a	IP address of the client
x-cs-connection-negotiated-cipher	client.connection.negotiated_cipher		OpenSSL cipher suite negotiated for the client connection
x-cs-connection-negotiated-cipher-strength	client.connection.negotiated_cipher.strength		Strength of the OpenSSL cipher suite negotiated for the client connection
r-dns			Hostname from the outbound server URL

Table B.6: Access Log Format Fields (Continued)

r-ip			IP address from the outbound server URL
r-port		%p	Port from the outbound server URL
r-supplier-dns			Hostname of the upstream host (not available for a cache hit)
r-supplier-ip			IP address used to contact the upstream host (not available for a cache hit)
r-supplier-port			Port used to contact the upstream host (not available for a cache hit)
sc-adapter	proxy.card		Adapter number of the client's connection to the Appliance
sc-connection			Unique identifier of the client's connection (i.e. SOCKET)
x-bluecoat-server-connection-socket-errno	server_connection. socket_errno		Error message associated with a failed attempt to connect to an upstream host
s-computername	proxy.name	%N	Configured name of the appliance
s-connect-type			Upstream connection type (Direct, SOCKS gateway, etc.)
s-dns			Hostname of the appliance (uses the primary IP address to avoid reverse DNS)
s-ip		%I	IP address of the appliance on which the client established its connection
s-port	proxy.port	%P	Port of the appliance on which the client established its connection
s-sitename		%S	Service used to process the transaction
x-module-name	module_name		The SGOS module that is handling the transaction
s-supplier-ip		%D	IP address used to contact the upstream host (not available for a cache hit)
s-supplier-name		%d	Hostname of the upstream host (not available for a cache hit)
x-bluecoat-transaction-id	transaction.id		Unique per-request identifier generated by the appliance (note: this value is not unique across multiple appliances)
x-bluecoat-appliance-name	appliance.name		Configured name of the appliance
x-bluecoat-appliance-primary-address	appliance.primary_address		Primary IP address of the appliance
x-bluecoat-proxy-primary-address	proxy.primary_address		Primary IP address of the appliance

Table B.6: Access Log Format Fields (Continued)

x-client-address			IP address of the client
x-client-ip			IP address of the client
Category: dns			
ELFF	CPL	Custom	Description
x-dns-cs-transport	dns.client_transport		The transport protocol used by the client connection in a DNS query
x-dns-cs-address	dns.request.address		The address queried in a reverse DNS lookup
x-dns-cs-dns	dns.request.name		The hostname queried in a forward DNS lookup
x-dns-cs-opcode	dns.request.opcode		The DNS OPCODE used in the DNS query
x-dns-cs-qtype	dns.request.type		The DNS QTYPE used in the DNS query
x-dns-cs-qclass	dns.request.class		The DNS QCLASS used in the DNS query
x-dns-rs-rcode	dns.response.code		The DNS RCODE in the response from upstream
x-dns-rs-a-records	dns.response.a		The DNS A RRs in the response from upstream
x-dns-rs-cname-records	dns.response.cname		The DNS CNAME RRs in the response from upstream
x-dns-rs-ptr-records	dns.response.ptr		The DNS PTR RRs in the response from upstream
Category: im			
ELFF	CPL	Custom	Description
x-im-buddy-id			Instant messaging buddy ID
x-im-buddy-name			Instant messaging buddy display name
x-im-buddy-state			Instant messaging buddy state
x-im-chat-room-id			Instant messaging identifier of the chat room in use
x-im-chat-room-members			The list of chat room member Ids
x-im-chat-room-type			The chat room type, one of 'public' or 'public', and possibly 'invite_only', 'voice' and/or 'conference'
x-im-client-info			The instant messaging client information
x-im-user-agent	im.user_agent		The instant messaging user agent string
x-im-file-path			Path of the file associated with an instant message

Table B.6: Access Log Format Fields (Continued)

x-im-file-size			Size of the file associated with an instant message
x-im-http-gateway			The upstream HTTP gateway used for IM (if any)
x-im-message-opcode	im.message.opcode		The opcode utilized in the instant message
x-im-message-reflected	im.message.reflected		Indicates whether or not the IM message was reflected.
x-im-message-route			The route of the instance message
x-im-message-size			Length of the instant message
x-im-message-text			Text of the instant message
x-im-message-type			The type of the instant message
x-im-method			The method associated with the instant message
x-im-user-id			Instant messaging user identifier
x-im-user-name			Display name of the client
x-im-user-state			Instant messaging user state
Category: packets			
ELFF	CPL	Custom	Description
c-pkts-lost-client			Number of packets lost during transmission from server to client and not recovered at the client layer via error correction or at the network layer via UDP resends.
c-pkts-lost-cont-net			Maximum number of continuously lost packets on the network layer during transmission from server to client
c-pkts-lost-net			Number of packets lost on the network layer
c-pkts-received			Number of packets from the server (s-pkts-sent) that are received correctly by the client on the first try
c-pkts-recovered-ECC			Number of packets repaired and recovered on the client layer
c-pkts-recovered-resent			Number of packets recovered because they were resent via UDP.
c-quality			The percentage of packets that were received by the client, indicating the quality of the stream
c-resendreqs			Number of client requests to receive new packets
s-pkts-sent			Number of packets from the server

Table B.6: Access Log Format Fields (Continued)

Category: req_rsp_line			
ELFF	CPL	Custom	Description
cs-method	method	%m	Request method used from client to appliance
x-cs-http-method	http.method		HTTP request method used from client to appliance. Empty for non-HTTP transactions
cs-protocol	client.protocol		Protocol used in the client's request
cs-request-line	http.request_line	%r	First line of the client's request
x-cs-raw-headers-count	request.raw_headers.count		Total number of 'raw' headers in the request
x-cs-raw-headers-length	request.raw_headers.length		Total length of 'raw' headers in the request
cs-version	request.version	%V	Protocol and version from the client's request, e.g. HTTP/1.1
x-bluecoat-proxy-via-http-version	proxy.via_http_version		Default HTTP protocol version of the appliance without protocol decoration (e.g. 1.1 for HTTP/1.1)
x-bluecoat-redirect-location	redirect.location		Redirect location URL specified by a redirect CPL action
rs-response-line			First line (a.k.a. status line) of the response from an upstream host to the appliance
rs-status	response.code		Protocol status code of the response from an upstream host to the appliance
rs-version	response.version		Protocol and version of the response from an upstream host to the appliance, e.g. HTTP/1.1
sc-status		%s	Protocol status code from appliance to client
x-bluecoat-ssl-failure-reason	ssl_failure_reason		Upstream SSL negotiation failure reason
x-cs-http-version	http.request.version		HTTP protocol version of request from the client. Does not include protocol qualifier (e.g. 1.1 for HTTP/1.1)
x-cs-socks-ip	socks.destination_address		Destination IP address of a proxied SOCKS request
x-cs-socks-port	socks.destination_port		Destination port of a proxied SOCKS request
x-cs-socks-method	socks.method		Method of a proxied SOCKS request
x-cs-socks-version	socks.version		Version of a proxied SOCKS request.

Table B.6: Access Log Format Fields (Continued)

x-sc-http-status	http.response.code		HTTP response code sent from appliance to client
x-rs-http-version	http.response.version		HTTP protocol version of response from the upstream host. Does not include protocol qualifier (e.g. 1.1 for HTTP/1.1)
x-sc-http-version			HTTP protocol version of response to client. Does not include protocol qualifier (e.g. 1.1 for HTTP/1.1)
x-sr-http-version			HTTP protocol version of request to the upstream host. Does not include protocol qualifier (e.g. 1.1 for HTTP/1.1)
sc (Content-Encoding)			Client Response header: Content-Encoding
sr (Accept-Encoding)			Server Request header: Accept-Encoding
Category: special_token			
ELFF	CPL	Custom	Description
x-bluecoat-special-amp	amp		The ampersand character
x-bluecoat-special-apos	apos		The apostrophe character (a.k.a. single quote)
x-bluecoat-special-cr	cr		Resolves to the carriage return character
x-bluecoat-special-crlf	crlf		Resolves to a carriage return/line feed sequence
x-bluecoat-special-empty	empty	%1	Resolves to an empty string
x-bluecoat-special-esc	esc		Resolves to the escape character (ASCII HEX 1B)
x-bluecoat-special-gt	gt		The greater-than character
x-bluecoat-special-lf	lf		The line feed character
x-bluecoat-special-lt	lt		The less-than character
x-bluecoat-special-quot	quot		The double quote character
x-bluecoat-special-slash	slash		The forward slash character
Category: status			
ELFF	CPL	Custom	Description
x-bluecoat-release-id	release.id		The release ID of the ProxySG operating system
x-bluecoat-release-version	release.version		The release version of the ProxySG operating system

Table B.6: Access Log Format Fields (Continued)

cs-categories			All content categories of the request URL
cs-categories-external			All content categories of the request URL that are defined by an external service.
cs-categories-policy			All content categories of the request URL that are defined by CPL.
cs-categories-local			All content categories of the request URL that are defined by a Local database.
cs-categories-provider			All content categories of the request URL that are defined by the current 3rd-party provider.
cs-categories-qualified			All content categories of the request URL, qualified by the provider of the category.
cs-category			Single content category of the request URL (a.k.a. sc-filter-category)
r-hierarchy			How and where the object was retrieved in the cache hierarchy.
sc-filter-category	category	%f	Content filtering category of the request URL
sc-filter-result		%W	Content filtering result: Denied, Proxied or Observed
s-action		%w	What type of action did the Appliance take to process this request.
s-cpu-util			Average load on the proxy's processor (0%-100%)
s-hierarchy		%H	How and where the object was retrieved in the cache hierarchy.
s-icap-info		%Z	ICAP response information
s-icap-status		%z	ICAP response status
x-bluecoat-surfcontrol-category-id			The SurfControl specific content category ID.
x-bluecoat-surfcontrol-is-denied			'1' if the transaction was denied, else '0'
x-bluecoat-surfcontrol-is-proxied			'0' if transaction is explicitly proxied, '1' if transaction is transparently proxied
x-bluecoat-surfcontrol-reporter-id			Specialized value for SurfControl reporter
x-bluecoat-surfcontrol-reporter-v4			The SurfControl Reporter v4 format

Table B.6: Access Log Format Fields (Continued)

x-bluecoat-surfcontrol-reporter-v5			The SurfControl Reporter v5 format
x-bluecoat-websense-category-id			The Websense specific content category ID
x-bluecoat-websense-keyword			The Websense specific keyword
x-bluecoat-websense-reporter-id			The Websense specific reporter category ID
x-bluecoat-websense-status			The Websense specific numeric status
x-bluecoat-websense-user			The Websense form of the username
x-bluecoat-websense-reporter-protocol-3			The Websense reporter format protocol version 3
x-exception-company-name	exception.company_name		The company name configured under exceptions
x-exception-contact	exception.contact		Describes who to contact when certain classes of exceptions occur, configured under exceptions (empty if the transaction has not been terminated)
x-exception-details	exception.details		The configurable details of a selected policy-aware response page (empty if the transaction has not been terminated)
x-exception-header	exception.header		The header to be associated with an exception response (empty if the transaction has not been terminated)
x-exception-help	exception.help		Help text that accompanies the exception resolved (empty if the transaction has not been terminated)
x-exception-id	exception.id		Identifier of the exception resolved (empty if the transaction has not been terminated)
x-exception-last-error	exception.last_error		The last error recorded for the current transaction. This can provide insight when unexpected problems are occurring (empty if the transaction has not been terminated)
x-exception-reason	exception.reason		Indicates the reason why a particular request was terminated (empty if the transaction has not been terminated)

Table B.6: Access Log Format Fields (Continued)

x-exception-sourcefile	exception.sourcefile		Source filename from which the exception was generated (empty if the transaction has not been terminated)
x-exception-sourceline	exception.sourceline		Source file line number from which the exception was generated (empty if the transaction has not been terminated)
x-exception-summary	exception.summary		Summary of the exception resolved (empty if the transaction has not been terminated)
x-patience-javascript	patience_javascript		Javascript required to allow patience responses
x-patience-progress	patience_progress		The progress of the patience request
x-patience-time	patience_time		The elapsed time of the patience request
x-patience-url	patience_url		The url to be requested for more patience information
x-virus-id			Identifier of a virus if one was detected
Category: streaming			
ELFF	CPL	Custom	Description
audiocodec			Audio codec used in stream.
avgbandwidth			Average bandwidth (in bits per second) at which the client was connected to the server.
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-cpu			Client computer CPU type.
c-hostexe			Host application
c-hostexever			Host application version number
c-os			Client computer operating system
c-osversion			Client computer operating system version number
c-playerid			Globally unique identifier (GUID) of the player
c-playerlanguage			Client language-country code
c-playerversion			Version number of the player

Table B.6: Access Log Format Fields (Continued)

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-status			Codes that describe client status
c-totalbuffertime			Time (in seconds) the client used to buffer the stream
filelength			Length of the file (in seconds).
filesize			Size of the file (in bytes).
protocol			Protocol used to access the stream: mms, http, or asfm.
s-totalclients			Clients connected to the server (but not necessarily receiving streams).
transport			Transport protocol used (UDP, TCP, multicast, etc.)
videocodec			Video codec used to encode the stream.
x-cache-info			Values: UNKNOWN, DEMAND_MISS, DEMAND_PARTIAL_HIT, DEMAND_HIT, LIVE_FROM_ORIGIN, LIVE_PARTIAL_SPLIT, 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	streaming.client		Type of streaming client in use (windows_media, real_media, or quicktime).
x-rs-streaming-content	streaming.content		Type of streaming content served. (e.g. windows_media, quicktime)
x-streaming-bitrate	bitrate		The reported client-side bitrate for the stream
Category: time			
ELFF	CPL	Custom	Description
connect-time			Total ms required to connect to the origin server

Table B.6: Access Log Format Fields (Continued)

date	date.utc	%x	GMT Date in YYYY-MM-DD format
dnslookup-time			Total ms cache required to perform the DNS lookup
duration		%T	Time taken (in seconds) to process the request
gmttime		%t	GMT date and time of the user request in format: [DD/MM/YYYY:hh:mm:ss GMT]
x-bluecoat-day-utc	day.utc		GMT/UTC day (as a number) formatted to take up two spaces (e.g. 07 for the 7th of the month)
x-bluecoat-hour-utc	hour.utc		GMT/UTC hour formatted to always take up two spaces (e.g. 01 for 1AM)
x-bluecoat-minute-utc	minute.utc		GMT/UTC minute formatted to always take up two spaces (e.g. 01 for 1 minute past)
x-bluecoat-month-utc	month.utc		GMT/UTC month (as a number) formatted to take up two spaces (e.g. 01 for January)
x-bluecoat-monthname-utc	monthname.utc		GMT/UTC month in the short-form string representation (e.g. Jan for January)
x-bluecoat-second-utc	second.utc		GMT/UTC second formatted to always take up two spaces (e.g. 01 for 1 second past)
x-bluecoat-weekday-utc	weekday.utc		GMT/UTC weekday in the short-form string representation (e.g. Mon for Monday)
x-bluecoat-year-utc	year.utc		GMT/UTC year formatted to always take up four spaces
localtime		%L	Local date and time of the user request in format: [DD/MMM/YYYY:hh:mm:ss +nnnn]
x-bluecoat-day	day		Localtime day (as a number) formatted to take up two spaces (e.g. 07 for the 7th of the month)
x-bluecoat-hour	hour		Localtime hour formatted to always take up two spaces (e.g. 01 for 1AM)
x-bluecoat-minute	minute		Localtime minute formatted to always take up two spaces (e.g. 01 for 1 minute past)
x-bluecoat-month	month		Localtime month (as a number) formatted to take up two spaces (e.g. 01 for January)

Table B.6: Access Log Format Fields (Continued)

x-bluecoat-monthname	monthname		Localtime month in the short-form string representation (e.g. Jan for January)
x-bluecoat-second	second		Localtime second formatted to always take up two spaces (e.g. 01 for 1 second past)
x-bluecoat-weekday	weekday		Localtime weekday in the short-form string representation (e.g. Mon for Monday)
x-bluecoat-year	year		Localtime year formatted to always take up four spaces
time	time.utc	%y	GMT time in HH:MM:SS format
timestamp		%g	Unix type timestamp
time-taken		%e	Time taken (in milliseconds) to process the request
rs-time-taken			Total time taken (in milliseconds) to send the request and receive the response from the origin server
x-bluecoat-end-time-wft			End local time of the transaction represented as a windows file time
x-bluecoat-start-time-wft			Start local time of the transaction represented as a windows file time
x-bluecoat-end-time-mssql			End local time of the transaction represented as a serial date time
x-bluecoat-start-time-mssql			Start local time of the transaction represented as a serial date time
x-cookie-date	cookie_date		Current date in Cookie time format
x-http-date	http_date		Current date in HTTP time format
x-timestamp-unix			Seconds since UNIX epoch (Jan 1, 1970) (local time)
x-timestamp-unix-utc			Seconds since UNIX epoch (Jan 1, 1970) (GMT/UTC)
Category: url			
ELFF	CPL	Custom	Description
cs-host		%v	Hostname from the client's request URL. If URL rewrite policies are used, this field's value is derived from the 'log' URL
cs-uri	log_url	%i	The 'log' URL.
cs-uri-address	log_url.address		IP address from the 'log' URL. DNS is used if URL uses a hostname.
cs-uri-extension	log_url.extension		Document extension from the 'log' URL.
cs-uri-host	log_url.host		Hostname from the 'log' URL.

Table B.6: Access Log Format Fields (Continued)

cs-uri-hostname	log_url.hostname		Hostname from the 'log' URL. RDNS is used if the URL uses an IP address.
cs-uri-path	log_url.path	%U	Path from the 'log' URL. Does not include query.
cs-uri-pathquery	log_url.pathquery		Path and query from the 'log' URL.
cs-uri-port	log_url.port		Port from the 'log' URL.
cs-uri-query	log_url.query	%Q	Query from the 'log' URL.
cs-uri-scheme	log_url.scheme		Scheme from the 'log' URL.
cs-uri-stem			Stem from the 'log' URL. The stem includes everything up to the end of path, but does not include the query.
c-uri	url		The original URL requested.
c-uri-address	url.address		IP address from the original URL requested. DNS is used if the URL is expressed as a hostname.
c-uri-cookie-domain	url.cookie_domain		The cookie domain of the original URL requested
c-uri-extension	url.extension		Document extension from the original URL requested
c-uri-host	url.host		Hostname from the original URL requested
c-uri-hostname	url.hostname		Hostname from the original URL requested. RDNS is used if the URL is expressed as an IP address
c-uri-path	url.path		Path of the original URL requested without query.
c-uri-pathquery	url.pathquery		Path and query of the original URL requested
c-uri-port	url.port		Port from the original URL requested
c-uri-query	url.query		Query from the original URL requested
c-uri-scheme	url.scheme		Scheme of the original URL requested
c-uri-stem			Stem of the original URL requested
sr-uri	server_url		URL of the upstream request
sr-uri-address	server_url.address		IP address from the URL used in the upstream request. DNS is used if the URL is expressed as a hostname.
sr-uri-extension	server_url.extension		Document extension from the URL used in the upstream request
sr-uri-host	server_url.host		Hostname from the URL used in the upstream request

Table B.6: Access Log Format Fields (Continued)

sr-uri-hostname	server_url.hostname		Hostname from the URL used in the upstream request. RDNS is used if the URL is expressed as an IP address.
sr-uri-path	server_url.path		Path from the upstream request URL
sr-uri-pathquery	server_url.pathquery		Path and query from the upstream request URL
sr-uri-port	server_url.port		Port from the URL used in the upstream request.
sr-uri-query	server_url.query		Query from the upstream request URL
sr-uri-scheme	server_url.scheme		Scheme from the URL used in the upstream request
sr-uri-stem			Path from the upstream request URL
s-uri	cache_url		The URL used for cache access
s-uri-address	cache_url.address		IP address from the URL used for cache access. DNS is used if the URL is expressed as a hostname
s-uri-extension	cache_url.extension		Document extension from the URL used for cache access
s-uri-host	cache_url.host		Hostname from the URL used for cache access
s-uri-hostname	cache_url.hostname		Hostname from the URL used for cache access. RDNS is used if the URL uses an IP address
s-uri-path	cache_url.path		Path of the URL used for cache access
s-uri-pathquery	cache_url.pathquery		Path and query of the URL used for cache access
s-uri-port	cache_url.port		Port from the URL used for cache access
s-uri-query	cache_url.query		Query string of the URL used for cache access
s-uri-scheme	cache_url.scheme		Scheme from the URL used for cache access
s-uri-stem			Stem of the URL used for cache access
x-cs(Referer)-uri	request.header.Referer.url		The URL from the Referer header.
x-cs(Referer)-uri-address	request.header.Referer.url.address		IP address from the 'Referer' URL. DNS is used if URL uses a hostname.
x-cs(Referer)-uri-extension	request.header.Referer.url.extension		Document extension from the 'Referer' URL.

Table B.6: Access Log Format Fields (Continued)

x-cs (Referer) -uri-host	request.header.Referer.url.host		Hostname from the 'Referer' URL.
x-cs (Referer) -uri-hostname	request.header.Referer.url.hostname		Hostname from the 'Referer' URL. RDNS is used if the URL uses an IP address.
x-cs (Referer) -uri-path	request.header.Referer.url.path		Path from the 'Referer' URL. Does not include query.
x-cs (Referer) -uri-pathquery	request.header.Referer.url.pathquery		Path and query from the 'Referer' URL.
x-cs (Referer) -uri-port	request.header.Referer.url.port		Port from the 'Referer' URL.
x-cs (Referer) -uri-query	request.header.Referer.url.query		Query from the 'Referer' URL.
x-cs (Referer) -uri-scheme	request.header.Referer.url.scheme		Scheme from the 'Referer' URL.
x-cs (Referer) -uri-stem			Stem from the 'Referer' URL. The stem includes everything up to the end of path, but does not include the query.
x-cs-raw-uri	raw_url		The 'raw' request URL.
x-cs-raw-uri-host	raw_url.host		Hostname from the 'raw' URL.
x-cs-raw-uri-port	raw_url.port		Port string from the 'raw' URL.
x-cs-raw-uri-scheme	raw_url.scheme		Scheme string from the 'raw' URL.
x-cs-raw-uri-path	raw_url.path		Path from the 'raw' request URL. Does not include query.
x-cs-raw-uri-pathquery	raw_url.pathquery		Path and query from the 'raw' request URL.
x-cs-raw-uri-query	raw_url.query		Query from the 'raw' request URL.
x-cs-raw-uri-stem			Stem from the 'raw' request URL. The stem includes everything up to the end of path, but does not include the query.
Category: user			
ELFF	CPL	Custom	Description
cs-auth-group	group		One group that an authenticated user belongs to. If a user belongs to multiple groups, the group logged is determined by the Group Log Order configuration specified in VPM. If Group Log Order is not specified, an arbitrary group is logged.
cs-auth-groups	groups		List of groups that an authenticated user belongs to.

Table B.6: Access Log Format Fields (Continued)

cs-auth-type			Client-side: authentication type (basic, ntlm, etc.)
cs-realm	realm		Authentication realm that the user was challenged in.
cs-user		%u	Qualified username for NTLM. Relative username for other protocols
cs-userdn	user		Full username of a client authenticated to the proxy (fully distinguished)
cs-username	user.name		Relative username of a client authenticated to the proxy (i.e. not fully distinguished)
sc-auth-status			Client-side: Authorization status
x-agent-sso-cookie			The authentication agent single signon cookie
x-cache-user			Relative username of a client authenticated to the proxy (i.e. not fully distinguished) (same as cs-username)
x-cs-auth-domain			The domain of the authenticated user.
x-cs-auth-form-action-url			The URL to submit the authentication form to.
x-cs-auth-form-domain-field			The authentication form input field for the user's domain.
x-cs-auth-request-id			The bas64 encoded string containing the original request information during forms based authentication
x-cs-username-or-ip			Used to identify the user using either their authenticated proxy username or, if that is unavailable, their IP address.
x-radius-splash-session-id			Session ID made available through RADIUS when configured for session management
x-radius-splash-username			Username made available through RADIUS when configured for session management
x-user-x509-issuer	user.x509.issuer		If the user was authenticated via an X.509 certificate, this is the issuer of the certificate as an RFC2253 DN
x-user-x509-serial-number	user.x509.serialNumber		If the user was authenticated via an X.509 certificate, this is the serial number from the certificate as a hexadecimal number.

Table B.6: Access Log Format Fields (Continued)

x-user-x509-subject	user.x509.subject		If the user was authenticated via an X.509 certificate, this is the subject of the certificate as an RFC2253 DN
Category: ci_request.header			
ELFF	CPL	Custom	Description
cs (Accept)	request.header.Accept		Request header: Accept
cs (Accept) -length	request.header.Accept.length		Length of HTTP request header: Accept
cs (Accept) -count	request.header.Accept.count		Number of HTTP request header: Accept
cs (Accept-Charset)	request.header.Accept-Charset		Request header: Accept-Charset
cs (Accept-Charset) -length	request.header.Accept-Charset.length		Length of HTTP request header: Accept-Charset
cs (Accept-Charset) -count	request.header.Accept-Charset.count		Number of HTTP request header: Accept-Charset
cs (Accept-Encoding)	request.header.Accept-Encoding		Request header: Accept-Encoding
cs (Accept-Encoding) -length	request.header.Accept-Encoding.length		Length of HTTP request header: Accept-Encoding
cs (Accept-Encoding) -count	request.header.Accept-Encoding.count		Number of HTTP request header: Accept-Encoding
cs (Accept-Language)	request.header.Accept-Language		Request header: Accept-Language
cs (Accept-Language) -length	request.header.Accept-Language.length		Length of HTTP request header: Accept-Language
cs (Accept-Language) -count	request.header.Accept-Language.count		Number of HTTP request header: Accept-Language
cs (Accept-Ranges)	request.header.Accept-Ranges		Request header: Accept-Ranges
cs (Accept-Ranges) -length	request.header.Accept-Ranges.length		Length of HTTP request header: Accept-Ranges
cs (Accept-Ranges) -count	request.header.Accept-Ranges.count		Number of HTTP request header: Accept-Ranges
cs (Age)	request.header.Age		Request header: Age
cs (Age) -length	request.header.Age.length		Length of HTTP request header: Age
cs (Age) -count	request.header.Age.count		Number of HTTP request header: Age
cs (Allow)	request.header.Allow		Request header: Allow
cs (Allow) -length	request.header.Allow.length		Length of HTTP request header: Allow
cs (Allow) -count	request.header.Allow.count		Number of HTTP request header: Allow

Table B.6: Access Log Format Fields (Continued)

cs (Authentication-Info)	request.header. Authentication-Info		Request header: Authentication-Info
cs (Authentication-Info) - length	request.header. Authentication-Info. length		Length of HTTP request header: Authentication-Info
cs (Authentication-Info) - count	request.header. Authentication-Info. count		Number of HTTP request header: Authentication-Info
cs (Authorization)	request.header. Authorization		Request header: Authorization
cs (Authorization) - length	request.header. Authorization.length		Length of HTTP request header: Authorization
cs (Authorization) - count	request.header. Authorization.count		Number of HTTP request header: Authorization
cs (Cache-Control)	request.header.Cache- Control		Request header: Cache-Control
cs (Cache-Control) - length	request.header.Cache- Control.length		Length of HTTP request header: Cache-Control
cs (Cache-Control) - count	request.header.Cache- Control.count		Number of HTTP request header: Cache-Control
cs (Client-IP)	request.header.Client- IP		Request header: Client-IP
cs (Client-IP) - length	request.header.Client- IP.length		Length of HTTP request header: Client-IP
cs (Client-IP) - count	request.header.Client- IP.count		Number of HTTP request header: Client-IP
cs (Connection)	request.header. Connection		Request header: Connection
cs (Connection) - length	request.header. Connection.length		Length of HTTP request header: Connection
cs (Connection) - count	request.header. Connection.count		Number of HTTP request header: Connection
cs (Content-Encoding)	request.header.Content- Encoding		Request header: Content-Encoding
cs (Content-Encoding) - length	request.header.Content- Encoding.length		Length of HTTP request header: Content-Encoding
cs (Content-Encoding) - count	request.header.Content- Encoding.count		Number of HTTP request header: Content-Encoding
cs (Content-Language)	request.header.Content- Language		Request header: Content-Language
cs (Content-Language) - length	request.header.Content- Language.length		Length of HTTP request header: Content-Language
cs (Content-Language) - count	request.header.Content- Language.count		Number of HTTP request header: Content-Language
cs (Content-Length)	request.header.Content- Length		Request header: Content-Length

Table B.6: Access Log Format Fields (Continued)

cs(Content-Length)-length	request.header.Content- Length.length		Length of HTTP request header: Content-Length
cs(Content-Length)-count	request.header.Content- Length.count		Number of HTTP request header: Content-Length
cs(Content-Location)	request.header.Content- Location		Request header: Content-Location
cs(Content-Location)-length	request.header.Content- Location.length		Length of HTTP request header: Content-Location
cs(Content-Location)-count	request.header.Content- Location.count		Number of HTTP request header: Content-Location
cs(Content-MD5)	request.header.Content- MD5		Request header: Content-MD5
cs(Content-MD5)-length	request.header.Content- MD5.length		Length of HTTP request header: Content-MD5
cs(Content-MD5)-count	request.header.Content- MD5.count		Number of HTTP request header: Content-MD5
cs(Content-Range)	request.header.Content- Range		Request header: Content-Range
cs(Content-Range)-length	request.header.Content- Range.length		Length of HTTP request header: Content-Range
cs(Content-Range)-count	request.header.Content- Range.count		Number of HTTP request header: Content-Range
cs(Content-Type)	request.header.Content- Type		Request header: Content-Type
cs(Content-Type)-length	request.header.Content- Type.length		Length of HTTP request header: Content-Type
cs(Content-Type)-count	request.header.Content- Type.count		Number of HTTP request header: Content-Type
cs(Cookie)	request.header.Cookie	%C	Request header: Cookie
cs(Cookie)-length	request.header.Cookie. length		Length of HTTP request header: Cookie
cs(Cookie)-count	request.header.Cookie. count		Number of HTTP request header: Cookie
cs(Cookie2)	request.header.Cookie2		Request header: Cookie2
cs(Cookie2)-length	request.header.Cookie2. length		Length of HTTP request header: Cookie2
cs(Cookie2)-count	request.header.Cookie2. count		Number of HTTP request header: Cookie2
cs(Date)	request.header.Date		Request header: Date
cs(Date)-length	request.header.Date. length		Length of HTTP request header: Date
cs(Date)-count	request.header.Date. count		Number of HTTP request header: Date
cs(Etag)	request.header.Etag		Request header: Etag

Table B.6: Access Log Format Fields (Continued)

cs (Etag) -length	request.header.Etag.length		Length of HTTP request header: Etag
cs (Etag) -count	request.header.Etag.count		Number of HTTP request header: Etag
cs (Expect)	request.header.Expect		Request header: Expect
cs (Expect) -length	request.header.Expect.length		Length of HTTP request header: Expect
cs (Expect) -count	request.header.Expect.count		Number of HTTP request header: Expect
cs (Expires)	request.header.Expires		Request header: Expires
cs (Expires) -length	request.header.Expires.length		Length of HTTP request header: Expires
cs (Expires) -count	request.header.Expires.count		Number of HTTP request header: Expires
cs (From)	request.header.From		Request header: From
cs (From) -length	request.header.From.length		Length of HTTP request header: From
cs (From) -count	request.header.From.count		Number of HTTP request header: From
cs (Front-End-HTTPS)	request.header.Front-End-HTTPS		Request header: Front-End-HTTPS
cs (Front-End-HTTPS) -length	request.header.Front-End-HTTPS.length		Length of HTTP request header: Front-End-HTTPS
cs (Front-End-HTTPS) -count	request.header.Front-End-HTTPS.count		Number of HTTP request header: Front-End-HTTPS
cs (Host)	request.header.Host		Request header: Host
cs (Host) -length	request.header.Host.length		Length of HTTP request header: Host
cs (Host) -count	request.header.Host.count		Number of HTTP request header: Host
cs (If-Match)	request.header.If-Match		Request header: If-Match
cs (If-Match) -length	request.header.If-Match.length		Length of HTTP request header: If-Match
cs (If-Match) -count	request.header.If-Match.count		Number of HTTP request header: If-Match
cs (If-Modified-Since)	request.header.If-Modified-Since		Request header: If-Modified-Since
cs (If-Modified-Since) -length	request.header.If-Modified-Since.length		Length of HTTP request header: If-Modified-Since
cs (If-Modified-Since) -count	request.header.If-Modified-Since.count		Number of HTTP request header: If-Modified-Since
cs (If-None-Match)	request.header.If-None-Match		Request header: If-None-Match

Table B.6: Access Log Format Fields (Continued)

cs (If-None-Match) -length	request.header.If-None-Match.length		Length of HTTP request header: If-None-Match
cs (If-None-Match) -count	request.header.If-None-Match.count		Number of HTTP request header: If-None-Match
cs (If-Range)	request.header.If-Range		Request header: If-Range
cs (If-Range) -length	request.header.If-Range.length		Length of HTTP request header: If-Range
cs (If-Range) -count	request.header.If-Range.count		Number of HTTP request header: If-Range
cs (If-Unmodified-Since)	request.header.If-Unmodified-Since		Request header: If-Unmodified-Since
cs (If-Unmodified-Since) -length	request.header.If-Unmodified-Since.length		Length of HTTP request header: If-Unmodified-Since
cs (If-Unmodified-Since) -count	request.header.If-Unmodified-Since.count		Number of HTTP request header: If-Unmodified-Since
cs (Last-Modified)	request.header.Last-Modified		Request header: Last-Modified
cs (Last-Modified) -length	request.header.Last-Modified.length		Length of HTTP request header: Last-Modified
cs (Last-Modified) -count	request.header.Last-Modified.count		Number of HTTP request header: Last-Modified
cs (Location)	request.header.Location		Request header: Location
cs (Location) -length	request.header.Location.length		Length of HTTP request header: Location
cs (Location) -count	request.header.Location.count		Number of HTTP request header: Location
cs (Max-Forwards)	request.header.Max-Forwards		Request header: Max-Forwards
cs (Max-Forwards) -length	request.header.Max-Forwards.length		Length of HTTP request header: Max-Forwards
cs (Max-Forwards) -count	request.header.Max-Forwards.count		Number of HTTP request header: Max-Forwards
cs (Meter)	request.header.Meter		Request header: Meter
cs (Meter) -length	request.header.Meter.length		Length of HTTP request header: Meter
cs (Meter) -count	request.header.Meter.count		Number of HTTP request header: Meter
cs (P3P)	request.header.P3P		Request header: P3P
cs (P3P) -length	request.header.P3P.length		Length of HTTP request header: P3P
cs (P3P) -count	request.header.P3P.count		Number of HTTP request header: P3P
cs (Pragma)	request.header.Pragma		Request header: Pragma

Table B.6: Access Log Format Fields (Continued)

cs (Pragma) -length	request.header.Pragma.length		Length of HTTP request header: Pragma
cs (Pragma) -count	request.header.Pragma.count		Number of HTTP request header: Pragma
cs (Proxy-Authenticate)	request.header.Proxy-Authenticate		Request header: Proxy-Authenticate
cs (Proxy-Authenticate) -length	request.header.Proxy-Authenticate.length		Length of HTTP request header: Proxy-Authenticate
cs (Proxy-Authenticate) -count	request.header.Proxy-Authenticate.count		Number of HTTP request header: Proxy-Authenticate
cs (Proxy-Authorization)	request.header.Proxy-Authorization		Request header: Proxy-Authorization
cs (Proxy-Authorization) -length	request.header.Proxy-Authorization.length		Length of HTTP request header: Proxy-Authorization
cs (Proxy-Authorization) -count	request.header.Proxy-Authorization.count		Number of HTTP request header: Proxy-Authorization
cs (Proxy-Connection)	request.header.Proxy-Connection		Request header: Proxy-Connection
cs (Proxy-Connection) -length	request.header.Proxy-Connection.length		Length of HTTP request header: Proxy-Connection
cs (Proxy-Connection) -count	request.header.Proxy-Connection.count		Number of HTTP request header: Proxy-Connection
cs (Range)	request.header.Range		Request header: Range
cs (Range) -length	request.header.Range.length		Length of HTTP request header: Range
cs (Range) -count	request.header.Range.count		Number of HTTP request header: Range
cs (Referer)	request.header.Referer	%R	Request header: Referer
cs (Referer) -length	request.header.Referer.length		Length of HTTP request header: Referer
cs (Referer) -count	request.header.Referer.count		Number of HTTP request header: Referer
cs (Refresh)	request.header.Refresh		Request header: Refresh
cs (Refresh) -length	request.header.Refresh.length		Length of HTTP request header: Refresh
cs (Refresh) -count	request.header.Refresh.count		Number of HTTP request header: Refresh
cs (Retry-After)	request.header.Retry-After		Request header: Retry-After
cs (Retry-After) -length	request.header.Retry-After.length		Length of HTTP request header: Retry-After
cs (Retry-After) -count	request.header.Retry-After.count		Number of HTTP request header: Retry-After
cs (Server)	request.header.Server		Request header: Server

Table B.6: Access Log Format Fields (Continued)

cs(Server)-length	request.header.Server.length		Length of HTTP request header: Server
cs(Server)-count	request.header.Server.count		Number of HTTP request header: Server
cs(Set-Cookie)	request.header.Set-Cookie		Request header: Set-Cookie
cs(Set-Cookie)-length	request.header.Set-Cookie.length		Length of HTTP request header: Set-Cookie
cs(Set-Cookie)-count	request.header.Set-Cookie.count		Number of HTTP request header: Set-Cookie
cs(Set-Cookie2)	request.header.Set-Cookie2		Request header: Set-Cookie2
cs(Set-Cookie2)-length	request.header.Set-Cookie2.length		Length of HTTP request header: Set-Cookie2
cs(Set-Cookie2)-count	request.header.Set-Cookie2.count		Number of HTTP request header: Set-Cookie2
cs(TE)	request.header.TE		Request header: TE
cs(TE)-length	request.header.TE.length		Length of HTTP request header: TE
cs(TE)-count	request.header.TE.count		Number of HTTP request header: TE
cs(Trailer)	request.header.Trailer		Request header: Trailer
cs(Trailer)-length	request.header.Trailer.length		Length of HTTP request header: Trailer
cs(Trailer)-count	request.header.Trailer.count		Number of HTTP request header: Trailer
cs(Transfer-Encoding)	request.header.Transfer-Encoding		Request header: Transfer-Encoding
cs(Transfer-Encoding)-length	request.header.Transfer-Encoding.length		Length of HTTP request header: Transfer-Encoding
cs(Transfer-Encoding)-count	request.header.Transfer-Encoding.count		Number of HTTP request header: Transfer-Encoding
cs(Upgrade)	request.header.Upgrade		Request header: Upgrade
cs(Upgrade)-length	request.header.Upgrade.length		Length of HTTP request header: Upgrade
cs(Upgrade)-count	request.header.Upgrade.count		Number of HTTP request header: Upgrade
cs(User-Agent)	request.header.User-Agent	%A	Request header: User-Agent
cs(User-Agent)-length	request.header.User-Agent.length		Length of HTTP request header: User-Agent
cs(User-Agent)-count	request.header.User-Agent.count		Number of HTTP request header: User-Agent
cs(Vary)	request.header.Vary		Request header: Vary

Table B.6: Access Log Format Fields (Continued)

cs (Vary) -length	request.header.Vary.length		Length of HTTP request header: Vary
cs (Vary) -count	request.header.Vary.count		Number of HTTP request header: Vary
cs (Via)	request.header.Via		Request header: Via
cs (Via) -length	request.header.Via.length		Length of HTTP request header: Via
cs (Via) -count	request.header.Via.count		Number of HTTP request header: Via
cs (WWW-Authenticate)	request.header.WWW-Authenticate		Request header: WWW-Authenticate
cs (WWW-Authenticate) -length	request.header.WWW-Authenticate.length		Length of HTTP request header: WWW-Authenticate
cs (WWW-Authenticate) -count	request.header.WWW-Authenticate.count		Number of HTTP request header: WWW-Authenticate
cs (Warning)	request.header.Warning		Request header: Warning
cs (Warning) -length	request.header.Warning.length		Length of HTTP request header: Warning
cs (Warning) -count	request.header.Warning.count		Number of HTTP request header: Warning
cs (X-BlueCoat-Error)	request.header.X-BlueCoat-Error		Request header: X-BlueCoat-Error
cs (X-BlueCoat-Error) -length	request.header.X-BlueCoat-Error.length		Length of HTTP request header: X-BlueCoat-Error
cs (X-BlueCoat-Error) -count	request.header.X-BlueCoat-Error.count		Number of HTTP request header: X-BlueCoat-Error
cs (X-BlueCoat-MC-Client- - Ip)	request.header.X-BlueCoat-MC-Client- Ip		Request header: X-BlueCoat-MC-Client- Ip
cs (X-BlueCoat-MC-Client- - Ip) -length	request.header.X-BlueCoat-MC-Client- Ip.length		Length of HTTP request header: X-BlueCoat-MC-Client- Ip
cs (X-BlueCoat-MC-Client- - Ip) -count	request.header.X-BlueCoat-MC-Client- Ip.count		Number of HTTP request header: X-BlueCoat-MC-Client- Ip
cs (X-BlueCoat-Via)	request.header.X-BlueCoat-Via		Request header: X-BlueCoat-Via
cs (X-BlueCoat-Via) -length	request.header.X-BlueCoat-Via.length		Length of HTTP request header: X-BlueCoat-Via
cs (X-BlueCoat-Via) -count	request.header.X-BlueCoat-Via.count		Number of HTTP request header: X-BlueCoat-Via
cs (X-Forwarded-For)	request.header.X-Forwarded-For	%X	Request header: X-Forwarded-For
cs (X-Forwarded-For) -length	request.header.X-Forwarded-For.length		Length of HTTP request header: X-Forwarded-For
cs (X-Forwarded-For) -count	request.header.X-Forwarded-For.count		Number of HTTP request header: X-Forwarded-For

Table B.6: Access Log Format Fields (Continued)

Category: <code>si_response.header</code>			
ELFF	CPL	Custom	Description
<code>rs (Accept)</code>	<code>response.header.Accept</code>		Response header: Accept
<code>rs (Accept-Charset)</code>	<code>response.header.Accept-Charset</code>		Response header: Accept-Charset
<code>rs (Accept-Encoding)</code>	<code>response.header.Accept-Encoding</code>		Response header: Accept-Encoding
<code>rs (Accept-Language)</code>	<code>response.header.Accept-Language</code>		Response header: Accept-Language
<code>rs (Accept-Ranges)</code>	<code>response.header.Accept-Ranges</code>		Response header: Accept-Ranges
<code>rs (Age)</code>	<code>response.header.Age</code>		Response header: Age
<code>rs (Allow)</code>	<code>response.header.Allow</code>		Response header: Allow
<code>rs (Authentication-Info)</code>	<code>response.header.Authentication-Info</code>		Response header: Authentication-Info
<code>rs (Authorization)</code>	<code>response.header.Authorization</code>		Response header: Authorization
<code>rs (Cache-Control)</code>	<code>response.header.Cache-Control</code>		Response header: Cache-Control
<code>rs (Client-IP)</code>	<code>response.header.Client-IP</code>		Response header: Client-IP
<code>rs (Connection)</code>	<code>response.header.Connection</code>		Response header: Connection
<code>rs (Content-Encoding)</code>	<code>response.header.Content-Encoding</code>		Response header: Content-Encoding
<code>rs (Content-Language)</code>	<code>response.header.Content-Language</code>		Response header: Content-Language
<code>rs (Content-Length)</code>	<code>response.header.Content-Length</code>		Response header: Content-Length
<code>rs (Content-Location)</code>	<code>response.header.Content-Location</code>		Response header: Content-Location
<code>rs (Content-MD5)</code>	<code>response.header.Content-MD5</code>		Response header: Content-MD5
<code>rs (Content-Range)</code>	<code>response.header.Content-Range</code>		Response header: Content-Range
<code>rs (Content-Type)</code>	<code>response.header.Content-Type</code>	<code>%c</code>	Response header: Content-Type
<code>rs (Cookie)</code>	<code>response.header.Cookie</code>		Response header: Cookie
<code>rs (Cookie2)</code>	<code>response.header.Cookie2</code>		Response header: Cookie2
<code>rs (Date)</code>	<code>response.header.Date</code>		Response header: Date
<code>rs (Etag)</code>	<code>response.header.Etag</code>		Response header: Etag
<code>rs (Expect)</code>	<code>response.header.Expect</code>		Response header: Expect

Table B.6: Access Log Format Fields (Continued)

rs (Expires)	response.header.Expires		Response header: Expires
rs (From)	response.header.From		Response header: From
rs (Front-End-HTTPS)	response.header.Front-End-HTTPS		Response header: Front-End-HTTPS
rs (Host)	response.header.Host		Response header: Host
rs (If-Match)	response.header.If-Match		Response header: If-Match
rs (If-Modified-Since)	response.header.If-Modified-Since		Response header: If-Modified-Since
rs (If-None-Match)	response.header.If-None-Match		Response header: If-None-Match
rs (If-Range)	response.header.If-Range		Response header: If-Range
rs (If-Unmodified-Since)	response.header.If-Unmodified-Since		Response header: If-Unmodified-Since
rs (Last-Modified)	response.header.Last-Modified		Response header: Last-Modified
rs (Location)	response.header.Location		Response header: Location
rs (Max-Forwards)	response.header.Max-Forwards		Response header: Max-Forwards
rs (Meter)	response.header.Meter		Response header: Meter
rs (P3P)	response.header.P3P		Response header: P3P
rs (Pragma)	response.header.Pragma		Response header: Pragma
rs (Proxy-Authenticate)	response.header.Proxy-Authenticate		Response header: Proxy-Authenticate
rs (Proxy-Authorization)	response.header.Proxy-Authorization		Response header: Proxy-Authorization
rs (Proxy-Connection)	response.header.Proxy-Connection		Response header: Proxy-Connection
rs (Range)	response.header.Range		Response header: Range
rs (Referer)	response.header.Referer		Response header: Referer
rs (Refresh)	response.header.Refresh		Response header: Refresh
rs (Retry-After)	response.header.Retry-After		Response header: Retry-After
rs (Server)	response.header.Server		Response header: Server
rs (Set-Cookie)	response.header.Set-Cookie		Response header: Set-Cookie
rs (Set-Cookie2)	response.header.Set-Cookie2		Response header: Set-Cookie2
rs (TE)	response.header.TE		Response header: TE

Table B.6: Access Log Format Fields (Continued)

rs(Trailer)	response.header.Trailer		Response header: Trailer
rs(Transfer-Encoding)	response.header.Transfer-Encoding		Response header: Transfer-Encoding
rs(Upgrade)	response.header.Upgrade		Response header: Upgrade
rs(User-Agent)	response.header.User-Agent		Response header: User-Agent
rs(Vary)	response.header.Vary		Response header: Vary
rs(Via)	response.header.Via		Response header: Via
rs(WWW-Authenticate)	response.header.WWW-Authenticate		Response header: WWW-Authenticate
rs(Warning)	response.header.Warning		Response header: Warning
rs(X-BlueCoat-Error)	response.header.X-BlueCoat-Error		Response header: X-BlueCoat-Error
rs(X-BlueCoat-MC-Client-Ip)	response.header.X-BlueCoat-MC-Client-Ip		Response header: X-BlueCoat-MC-Client-Ip
rs(X-BlueCoat-Via)	response.header.X-BlueCoat-Via		Response header: X-BlueCoat-Via
rs(X-Forwarded-For)	response.header.X-Forwarded-For		Response header: X-Forwarded-For

Appendix C: Using WCCP

This appendix discusses how to configure a Blue Coat Systems ProxySG to participate in a Web Cache Communication Protocol (WCCP) scheme, when a WCCP-capable router collaborates with a set of WCCP-configured ProxySG Appliances to service requests. If you are already familiar with WCCP version 2 and want to get your router and ProxySG up and running right away, see the "Quick Start" on page 913.

Overview

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 ProxySG Appliances.
- **Redirection safeguards.** If no ProxySG Appliances are available, redirection stops and the router forwards traffic to the original destination address.

WCCP has two versions, version 1 and version 2, both of which are supported by Blue Coat. However, only one protocol version can be active on the ProxySG at a time. The active WCCP protocol set up in the ProxySG configuration must match the version running on the WCCP router.

Using WCCP and Transparent Redirection

A WCCP-capable router operates in conjunction with the ProxySG Appliances to transparently redirect traffic to a set of caches that participate in the specified WCCP protocol. IP packets are redirected based on fields within each packet. For instance, WCCP version 1 only redirects destination TCP port 80 (default HTTP traffic) IP packets. WCCP version 2 allows you to redirect traffic from other ports and protocols.

Load balancing is achieved through a redirection hash table to determine which ProxySG will receive the redirected packet.

WCCP Version 1

In WCCP version 1, the WCCP-configured home router transparently redirects TCP port 80 packets to a maximum of 32 ProxySG Appliances. (A ProxySG is seen as a cache in WCCP protocol.)

One of the caches participating in the WCCP service group is automatically elected to configure the home router's redirection tables. This way, caches can be transparently added and removed from the WCCP service group without requiring operator intervention. WCCP version 1 supports only a single service group.

Figure C-1: "A Typical WCCP Version 1 Configuration" on page 912 illustrates a typical WCCP version 1 implementation.

Each applicable client IP packet received by the home router is transparently redirected to a cache. A ProxySG from the group is selected to define the home router's redirection hash table for all caches. All caches periodically communicate with the home router to verify WCCP protocol synchronization and ProxySG availability within the service group. In return, the home router responds to each cache with information as to which ProxySG Appliances are available in the service group.




Figure C-1: A Typical WCCP Version 1 Configuration

Note the following WCCP version 1 caveats:

- The home router IP must be configured on all participating interfaces and must match the home router address configured on the ProxySG.
- The adapter connected to the ProxySG must be Ethernet or Fast Ethernet.
- For Cisco routers using WCCP version 1, minimum IOS releases are 11.1(18)CA and 11.2(13)P. Note that releases prior to IOS 12.0(3)T only support WCCP version 1. Ensure that you are using the correct IOS software for the router and that the ProxySG configuration protocol version number and router protocol version number match.

For more information on WCCP Version 1, refer to the Cisco web site. The rest of this appendix discusses WCCP version 2 only.

WCCP Version 2

For Cisco routers using WCCP version 2, minimum IOS releases are 12.0(3)T and 12.0(4). Note that release 12.0(5) and later releases support WCCP versions 1 and 2. Ensure that you use the correct IOS software for the router and that you have a match between the ProxySG configuration WCCP version number and router protocol version number.

WCCP version 2 protocol offers the same capabilities as version 1, along with increased protocol security and multicast protocol broadcasts. Version 2 multicasting allows caches and routers to discover each other via a common multicast service group and matching passwords. In addition, up to 32 WCCP-capable routers can transparently redirect traffic to a set of up to 32 ProxySG Appliances. Version 2 WCCP-capable routers are capable of redirecting IP traffic to a set of ProxySG Appliances based on various fields within those packets.

Version 2 allows routers and caches to participate in multiple, simultaneous service groups. Routers can transparently redirect IP packets based on their formats. For example, one service group could redirect HTTP traffic and another could redirect FTP traffic.

Note: Blue Coat recommends that WCCP-compliant caches from different vendors be kept separate and that only one vendor's routers be used in a service group.

One of the caches participating in the WCCP service group is automatically elected to configure the home router's redirection tables. This way, caches can be transparently added and removed from the WCCP service group without requiring operator intervention. WCCP version 2 supports multiple service groups.

Figure C-2, below, illustrates a WCCP version 2 implementation using multiple routers and ProxySG Appliances. In this scenario, routers 1 through n and caches 1 through m participate in the same service group. As in version 1, an appliance from the group is selected to define the redirection hash table in all routers for all caches. All caches periodically communicate with all routers to verify WCCP protocol synchronization and ProxySG and router availability within the service group. In return, each router responds to caches with information as to what caches and discovered routers are available in the service group.




Figure C-2: A Version 2 Configuration Using Packet Redirection to Multiple Routers and Caches

Quick Start

Two tasks must be completed to get WCCP running: configuring the router and configuring the ProxySG. If you have a standard router and ProxySG configuration, use the Quick Start below. Otherwise, begin with the instructions in the procedure "To Do Initial Router Configuration:", below, and "To Create a ProxySG WCCP Configuration File and Enable WCCP:" on page 914.

If you require a more complicated configuration, start with "Configuring a WCCP Version 2 Service on the Router".

To Do Initial Router Configuration:

1. From the router `(config)` mode, tell WCCP which service group you want use. The web-cache service group redirects port 80 (HTTP) traffic only.

```
Router(config)#ip wccp web-cache
```

2. Enter the `(config-if)` submode by telling WCCP which IP address to use.

```
Router(config)#int interface
```

where *interface* is the adapter interface with an IP address. The prompt changes to configuration interface submode.

3. Enable packet redirection on an outbound (Internet facing) interface.

```
Router(config-if)# ip wccp web-cache redirect out
```

4. Prevent packets received on an adapter interface from being checked for redirection and allow the use of Blue Coat bypass lists.

```
Router(config-if)# ip wccp redirect exclude in
```

For more information on WCCP router configuration, see "Configuring a WCCP Version 2 Service on the Router" on page 914.

To Create a ProxySG WCCP Configuration File and Enable WCCP:

1. Create a WCCP configuration file through either the ProxySG's CLI inline commands or through a text editor. Make sure that the home router you enter here is the home router that was named in the router's configuration. If you do have a mismatch, you must correct it before continuing. See "Identifying a Home Router/Router ID Mismatch" on page 933.

For more information on creating a configuration file, see "Creating a ProxySG WCCP Configuration File" on page 921.

If you used the `inline` commands, you have completed WCCP configuration for both the router and the ProxySG and you have enabled WCCP on the ProxySG. No further steps are needed.

2. If you used a text editor, copy the file to an HTTP server accessible to the ProxySG.
3. Enable WCCP and download the configuration file to the ProxySG.

```
SGOS#(config) wccp enable
SGOS#(config) wccp path http://205.66.255.10/files/wccp.txt
SGOS#(config) load wccp-settings
```

Configuring a WCCP Version 2 Service on the Router

Configuring a router requires that you work with two different types of configuration commands:

- Creating a service group (which uses global settings).
- Configuring the Internet-Connected Interface (which uses interface settings).

Define service group settings before defining adapter interface settings.

Setting up a Service Group

Services are of two types:

- Well known services (web-cache for port 80—HTTP— redirection)

The `web-cache` service group is supported by both Cisco and Blue Coat.

- Dynamic services (which can be used for other services, such as FTP, RTSP redirection, and reverse proxy).

Dynamic service uses identifiers ranging from 0-99 to name the service group.

WCCP global settings allow you to name the service group and then define the characteristics for that service group. Even if you use the pre-defined `web-cache` service group, you should:

- configure a multicast group address
- create and identify a redirection access list and associate it with a service group
- create and identify a cache bypass list and associate it with a service group
- create password authentication for messages sent by the service group to the router

Syntax for configuring a service group (global settings):

```
ip wccp {web-cache | service-number} [group-address groupaddress] [redirect-list
access-list] [group-list access-list] [password password]
```

where:

<code>web-cache</code>	Enables port 80 (HTTP) service.
<code>service-number</code>	The identification number of the cache service group being controlled by the router. Services are identified using a value from 0 to 99. The reverse-proxy service is indicated using the value 99, although any value can be used for reverse proxy.
<code>group-address</code> <code>groupaddress</code>	(Optional) If no redirect list is defined (the default), all traffic will be redirected. The group address option directs the router to use a specified multicast IP address to coalesce the “I See You” responses to the “Here I Am” messages that it has received on this address. The <code>groupaddress</code> argument requires a multicast address used by the router to determine which cache engine should receive redirected messages. The response is sent to the group address, as well. If no group address is defined (the default), all “Here I Am” messages are responded to with a unicast reply.
<code>redirect-list</code> <code>access-list</code>	(Optional) Directs the router to use an access list to control traffic redirected to the defined service group. The <code>access-list</code> parameter specifies either a number from 1 to 99 identifying a predefined standard or extended access list number, or a name (up to 64 characters long) identifying an existing standard or extended access list. The access list itself specifies which traffic may be redirected.
<code>group-list</code> <code>access-list</code>	(Optional) If no group list is defined (the default), all caches may participate in the service group. The <code>group-list</code> option directs the router to use an access list to determine which caches are allowed to participate in the service group. The <code>access-list</code> parameter specifies either a number from 1 to 99 identifying a predefined standard or extended access list number or a name (up to 64 characters long) identifying an existing standard or extended access list. The access list itself specifies which caches are permitted to participate in the service group.

<code>password</code> <code>password</code>	(Optional) By default, password authentication is not configured and authentication is disabled. The password option increases authentication security to messages received from the service group specified by the service-number. Messages that do not pass authentication are discarded. The password can be up to eight characters long. Note that if you specify a password in the router configuration, you must also configure the same password separately on each cache.
--	---

Naming a Service Group and Enabling WCCP

WCCP version 2 is enabled when you name a WCCP service group. (Version 1 requires a specific `enable` command.) The service group can already exist, such as `web-cache`, or it could be a new group, such as `36`.

To Name a Service Group and Enable WCCP:

From the router (`config`) mode, enter the following command:

```
Router#(config) ip wccp web-cache
-or-
Router#(config) ip wccp 36
```

Configuring a Global Multicast Group Address

Benefits of using a multicast address include reduced WCCP protocol traffic and the ability to easily add and remove caches and routers from a service group without having to reconfigure all service group members. Multicast addresses fall within the range of 224.0.0.0 to 239.255.255.255.

Use the following syntax to configure a global multicast group address for multicast cache discovery.

```
ip wccp {web-cache | service-number} [group-address group_address]
```

To Configure a Multicast Address:

From the router (`config`) mode, name the group that will use the multicast address, provide the address, then tell the router which adapter interface will be used:

```
Router(config)# ip wccp 36 group-address 225.1.1.1
Router(config)# interface ethernet 0
Router(config-if)# end
```

Creating a Redirection Access List and Associating it with a Service Group

Redirection access lists can contain commands redirecting packets from one network or cache to another. The lists also can be used to determine which caches participate in which service groups.

The two lists, although similar, have different purposes, and are applied to the router differently. The redirection lists are applied with the `redirect-list` option. The cache bypass lists are applied with the `group-list` argument. Both lists can be identified with either a name or a number.

Use the following syntax to create a redirection access list. Note that this is partial syntax for this command. Access lists are very complicated; refer to the Cisco web site for complete syntax.

```
access-list acl_ID [deny | permit] protocol {[source_addr source_mask] |
[local_addr local_mask]}
```


where:

<code>acl_ID</code>	Names the access list you are creating. You can use either a name or number.
<code>deny</code>	Indicates that you do not want to allow a packet to traverse the Cisco router. By default, the router firewall denies all inbound or outbound packets unless you specifically permit access.
<code>permit</code>	Selects a packet to traverse the PIX firewall. By default, the router firewall denies all inbound or outbound packets unless you specifically permit access.
<code>protocol</code>	Identifies, by name or number, an IP protocol. This parameter can be one of the keywords <code>icmp</code> , <code>ip</code> , <code>tcp</code> , or <code>udp</code> , or an integer in the range 1 to 254 representing an IP protocol number. To match any Internet protocol, including ICMP, TCP, and UDP, use the keyword <code>ip</code> .
<code>source_addr</code>	Indicates the address of the network or host from which the packet is being sent. Use the keyword <code>any</code> as an abbreviation for an address of 0.0.0.0.
<code>source_mask</code>	Specifies the netmask bits (mask) to be applied to <code>source_addr</code> , if the source address is for a network mask. Use the keyword <code>any</code> as an abbreviation for a mask of 0.0.0.0.
<code>local_addr</code>	Indicates the address of the network or host local to the PIX firewall. The <code>local_addr</code> is the address after NAT has been performed. Use the keyword <code>host</code> , followed by <code>address</code> , as an abbreviation for a mask of 255.255.255.255.
<code>local_mask</code>	Specifies the netmask bits (mask) to be applied to <code>local_addr</code> , if the local address is a network mask. Use the keyword <code>host</code> followed by <code>address</code> as an abbreviation for a mask of 255.255.255.255.

To Create a Redirection Access List or a Cache Bypass List:

From the router (config) prompt, name an access list and assign rules to it.

```
Router(config)# access-list 100 deny ip any host 126.10.10.10
Router(config)# access-list 100 permit ip any any
Router#
```

- The commands above gave the access list a name of 100.
- Denied packets from any protocol to be sent from any host on the 126.10.10.10 network.
- Permitted packets from any protocol to be sent from any other network.

To Associate a Redirection Access List with a Specific Service Group:

1. Create a redirection access list.
2. Associate the access list with a specified service group.

```
ip wccp {web-cache | service-number} [redirect-list access-list]
Router(config)# interface ethernet 0/0
Router(config-if)# ip wccp web-cache redirect-list 100
Router(config-if)# end
Router#
```

To Associate a Cache Bypass Access List with a Specific Service Group:

1. Create a redirection access list, using the syntax discussed above.
2. Associate the access list with a specified service group.

```
ip wccp {web-cache | service-number} [group-list access-list]
Router(config)# interface ethernet 0/0
Router(config-if)# ip wccp web-cache group-list 120
Router(config-if)# end
Router#
```

Configuring the Internet-Connected Interface

WCCP interface settings allow you to configure the Internet-connected adapter interface that will redirect Web traffic to the content engine.

Using the interface commands allows you to:

- Enable and prevent packet redirection
- Enable reception of multicast packets for service group member routers

Syntax for configuring an internet-connected adapter interface (interface settings):

```
ip wccp [{web-cache | service-number} redirect out | group-listen] | redirect
exclude in
```

where:

<code>web-cache</code>	Enables the web cache service group.
<code>service-number</code>	The identification number of the cache service group being controlled by the router. Services are identified using a value from 0 to 99. The reverse-proxy service is indicated using the value 99.
<code>redirect out</code>	Enables packet redirection on an outbound (Internet facing) adapter interface.
<code>group-listen</code>	On a router that is a member of a service group, enables the reception of pre-defined IP multicast packets.
<code>redirect</code> <code>exclude in</code>	Prevents packets received on an adapter interface from being checked for redirection. Note that if the cache <i>service-group</i> is located on a separate router interface, the possibility exists that bypass filters could be enabled on the cache.

Using Packet Redirection

WCCP communication among the routers and the ProxySG Appliances can be done by either directly addressing protocol packets to each router's and cache's IP address (as illustrated in Figure C-1 on page 912) or by sending these packets to a common multicast address as illustrated in Figure C-3, below:




Figure C-3: A Version 2 Configuration Using Multicast Packet Redirection

You can configure redirection on inbound or outbound interfaces.

To Configure Redirection on the Outbound Interfaces:

Use the following syntax to configure redirection on the outbound adapter interface.

```
ip wccp {web-cache | service-number} redirect out
```

From the router (config) prompt, enter the following:

```
Router(config)# interface ethernet 0
Router(config-if)# ip wccp web-cache redirect out
Router(config-if)# end
```

To Exclude Packet Redirection on an Inbound Adapter Interface:

Use the following command to prevent packets received on an adapter interface from being checked for redirection.

```
ip wccp redirect exclude in
```

The following example shows how to exclude Blue Coat adapter interface (xx, in this case) and allow use of Blue Coat bypass lists:

From the router (config) prompt, enter the following:

```
Router(config)# int xx
Router(config-if)# ip wccp redirect exclude in
Router(config-if)# end
```

Enabling Reception of Multicast Packets

Benefits of using a multicast address include reduced WCCP protocol traffic and the ability to easily add and remove caches and routers from a service group without having to reconfigure all service group members. You (optionally) set up a multicast group address in "Configuring a Global Multicast Group Address". In the following procedure, you enable the reception of the pre-defined IP multicast packets to routers that are members of the group.

Multicast addresses fall within the range 224.0.0.0 to 239.255.255.255.

Use the following syntax to configure for multicast discovery of the cache(s).

```
ip wccp {web-cache | service-number} group-listen
```

The following example configures the router to use the WCCP 36 service group to redirect port 80 destination traffic. WCCP protocol traffic will be using multicast address 225.1.1.1. Adapter interface “Ethernet 0” is used to receive the multicast WCCP traffic.

```
Router(config)# ip wccp 36 group-address 225.1.1.1
Router(config)# interface ethernet 0
Router(config-if)# ip wccp web-cache group-listen
Router(config-if)# end
```

Saving and Viewing Changes

Once you have made all the changes, you must permanently save them to disk. If not, the changes will be lost at the next reboot of the router.

To Save Router Configuration:

```
Router# write memory
```

To Display all Current WCCP Configuration Settings:

Use the following syntax to verify the settings in the new router configuration and to ensure that the appropriate cache engines are visible to the router.

```
show ip wccp {web-cache | service-number} [view | detail]
```

where

- view (Optional) Lists all members of the identified service group and whether they have been detected.
- detail (Optional) Displays IP and protocol version information about the router. Displays IP, protocol version, state, initial and assigned hash, hash allotment, redirected packet, and connection time information about the associated cache engine (ProxySG).

For example:

```
Router# show ip wccp web-cache view
```

```
Global WCCP Information:
Service Name: web-cache:
Number of Cache Engines:1
Number of Routers:1
Total Packets Redirected:186
Redirect Access-list:120
Total Packets Denied Redirect:57
Total Packets Unassigned:-none-
Group Access-list:0
Total Messaged Denied to Group:0
Total Authentication Failures:0
```

```
WCCP Router Informed of:
```

```
86.135.77.10
186.135.77.20

WCCP Cache Engines Visible:
186.135.77.11
186.135.77.12

WCCP Cache Engines Not Visible:
-none-
```

Creating a ProxySG WCCP Configuration File

Once you have the router global and adapter interface settings complete, you must create a WCCP configuration file for the ProxySG. These configurations should include the following:

- Identify the service group
- Identify the queuing priorities for all defined service groups
- Identify the protocol
- Load balancing caches in a service group
- Identify ports
- Identify the home router as defined in the router configuration
- Identify the packet forwarding method

Understanding Packet Forwarding

By default, Cisco's GRE encapsulation (Generic Routing Encapsulation) is used to forward packets from the WCCP router to the caches. If you have a version 2 WCCP router, you can alternatively use Layer 2 (L2) rewrites to forward packets, which is faster than GRE and saves network bandwidth.

Using GRE, redirected packets are encapsulated in a new IP packet with a GRE header.

Using L2, redirected packets are not encapsulated; the packet's destination MAC address is replaced with the MAC address of the target cache. This different way of directing packets saves you the overhead of creating the GRE packet at the router and decoding it at the cache. Also, you save network bandwidth that would otherwise be consumed by the GRE header.

If you want to continue using GRE, you need not change any settings. To use L2 packet redirection, you must add the forwarding option to the ProxySG configuration file.

If WCCP version 2 is supported, the router sends out a list of forwarding mechanisms supported by the router in the first WCCP2_I_SEE_YOU message. The cache responds with a WCCP2_HERE_I_AM message. If the router does not send the list, the cache aborts its attempt to join the WCCP service group. If the method of forwarding mechanism is not supported by the router, the WCCP2 messages from the cache are ignored.

Caveats for using L2 redirection:

- You must use WCCP version 2.
- If a cache is not connected directly to a router, the router will not allow the cache to negotiate the rewrite method.

- The same rewrite method must be used for both packet forwarding and packet return.

Understanding Cache Load Balancing

If you use WCCP version 2, you can balance the load on the caches in a service group. When a router receives an IP packet for redirection, it hashes fields within the packet to yield an index within the hash table. The packet then is forwarded to the “owner” ProxySG for servicing. The proportion of redirection hash table assigned to each ProxySG can be altered to provide a form of load balancing between caches in a service group.

A hash table is configured by a dynamically elected ProxySG participating in a service group, enabling the simultaneous interception of multiple protocols on multiple ports. You can configure up to 100 dynamic or standard service groups plus standard service groups. A single service can intercept up to eight port numbers.

Each element in this 256-entry hash table refers to an active ProxySG within the service group. By default, each ProxySG is assigned roughly an even percentage of the 256-element redirection hash table. Multiple network cards within a ProxySG can participate in the same service group. To the routers and other caches, each adapter interface appears as a unique cache. Using this strategy, redirected traffic can be better distributed among network interfaces in a cache.

Using Figure C-4, below, all caches would be assigned $1/m$ of the redirection hash table, but since Cache 2 and Cache 3 are physically located within the same ProxySG Appliance, that appliance would actually be assigned $2/m$ of the redirection hash table.




Figure C-4: A Version 2 Configuration Using Multicast Packet Redirection to Multiple Routers, Multiple Caches, and a Service Group

Assigning Percentages

You can override the default of each ProxySG being assigned roughly an even percentage; the relative distribution of the redirection hash table can be specified for each cache. Multiple hash-distributions are supported. Also, all, none, or part of a source and/or destination IP address or port number can be used in the hash. Each ProxySG can be assigned a primary-hash-weight value to determine the proportion of the 256-element hash table to be assigned.

If all caches are configured with a 0 primary-hash-weight value (the default) then each ProxySG is assigned an equal proportion of the redirection hash table. However, if any ProxySG is configured with a non-zero primary-hash-weight, each ProxySG is assigned a relative proportion of the table.

For instance, consider a configuration with five caches that use a primary-hash-weight defined as {25, 200, 0, 50, 25}. The total requested weight value is $25+200+0+50+25=300$ and, therefore, the proportion of the hash table assigned to each ProxySG will be $25/300$, $200/300$, $0/300$, $50/300$, and $25/300$.

Note that since one cache did not specify a non-zero primary-hash-weight, that cache will not be assigned any elements within the redirection hash table and, therefore, will not receive any redirected traffic. Also note that the hash weight can be specified for each caching member within a ProxySG. In Figure C-4, Cache 2 and Cache 3 can be assigned different weight values.

Alternate Hash Table

In some cases, a web site becomes an Internet *hot spot*, receiving a disproportional number of client traffic relative to other sites. This situation can cause a larger request load on a specific ProxySG since the hash element associated with the popular site receives more activity than other hash elements.

To balance the redirection traffic load among the caches, a service group can be configured to use an alternate hash function when the number of GRE packets forwarded to the cache exceeds a certain number. (If you use L2 forwarding, the ProxySG counts MAC addresses.) Therefore, when a router receives an IP packet that hashes to an element flagged as a hot spot, the alternate hash function is computed. The ProxySG specified by the new index in the redirection hash table receives the redirected packet.

Each ProxySG can dynamically determine a hot spot within its assigned portion of the redirection hash table.

Alternate hash tables are only used for dynamic service groups that specify alternate-hash flags within their service-flags. Note that the default web-cache service group cannot use an alternate hash table. Instead, a comparable dynamic service group must be created.

To use hot spot detection, the ProxySG's WCCP configuration file must specify:

```
service-flags source-ip-hash
service-flags destination-port-alternate-hash
```

Creating a Configuration File

An example of a file using a dynamic service, as opposed to the default web-cache service, is shown below:

Note that if you are using the default web-cache service, the service group settings *priority*, *protocol*, *service flags*, and *ports* are not used.

```
wccp enable
wccp version 2
service-group 9
forwarding-type L2
priority 1
protocol 6
service-flags destination-ip-hash
service-flags ports-defined
ports 80 21 1755 554 80 80 80 80
```

```
interface 6
home-router 10.16.18.2
end
```

You can create a configuration file customized for the environment two ways: CLI inline commands or through a text file. In either case, the configuration file must include the information required by the commands below.

Syntax to create a customized configuration file:

```
service-group {web-cache | service-number}
[priority priority-number]
[protocol protocol-number]
[service-flags hash-bit-identifier]
[ports port1 ... port8]
home-router [ip-address | domain-name]
interface [interface-number]
[password string]
[primary-hash-weight interface-number value]
forwarding-type [GRE | L2]
```

Using Optional Negation Syntax, you can create an alternative WCCP configuration file using these negative commands; this is especially helpful when testing and debugging. This functionality enables you to change some of the configuration settings without altering or reloading the main configuration file.

```
[no] service-group {web-cache | service-number}
[priority priority-number]
[protocol protocol-number]
[no] service-flags hash-bit-identifier
[ports port1 ...port8]
home-router [ip-address | domain-name]
[no] interface [interface-number]
[password string | no password]
[primary-hash-weight interface-number value]
```

where:

<i>web-cache</i>	Enables the web cache service group. Note that if you use the web-cache service group for WCCP, the dynamic service group settings (priority, protocol, service flags, and ports) are not applicable.
<i>service-number</i>	The identification number of the dynamic service group being controlled by the router. Services are identified using a value from 0 to 99. The reverse-proxy service is indicated using the value 99.
<i>priority-number</i>	(Applies to a dynamic service group only. A dynamic service group is one identified by a <i>service number</i> .) Establishes queuing priorities for all defined service groups, based on a priority number from 0 through 255, inclusive.
<i>protocol-number</i>	(Applies to a dynamic service group only. A dynamic service group is one identified by a <i>service number</i> .) Number of an Internet protocol. <i>Protocol-number</i> must be an integer in the range 0 through 255, inclusive, representing an IP protocol number.

<i>hash-bit-identifier</i>	<p>(Applies to a dynamic service group only. A dynamic service group is one identified by a <i>service number</i>.) Sets the hash index, for load balancing purposes.</p> <p>The key associated with the <i>hash-bit-identifier</i> you specify will be hashed to produce the primary redirection hash table index. For instance, if only the <i>destination-ip-hash</i> flag is set, then the packet destination IP address will be used to determine the index. The index is constructed by starting with an initial value of zero and then computing an exclusive OR (XOR) of the fields specified in the hash-bit identifier.</p> <p>If alternative hashing has been enabled, any alternate hash flags are processed in the same way and produce a secondary redirection hash table index. Alternate hash flags end with the suffix “-alternate-hash.”</p> <p>For more information using the hashing table, see "Understanding Cache Load Balancing" on page 922.</p>
<i>source-ip-hash</i> (<i>hash-bit-identifier</i>)	Sets the source IP bit definition within the redirection hash table index.
<i>destination-ip-hash</i> (<i>hash-bit-identifier</i>)	Sets the source IP bit definition within the redirection hash table index.
<i>source-port-hash</i> (<i>hash-bit-identifier</i>)	Sets the source port bit definition within the redirection hash table index.
<i>destination-port-hash</i> (<i>hash-bit-identifier</i>)	Sets the destination port bit definition within the redirection hash table index.
<i>ports-defined</i> (<i>hash-bit-identifier</i>)	Sets the port bit definition within the redirection hash table index.
<i>ports-source</i> (<i>hash-bit-identifier</i>)	Sets the source port bit definition within the redirection hash table index.
<i>source-ip-alternate-hash</i> (<i>hash-bit-identifier</i>)	Sets the alternate source IP bit definition within the redirection hash table index.
<i>destination-ip-alternate-hash</i> (<i>hash-bit-identifier</i>)	Sets the alternate destination IP bit definition within the redirection hash table index.
<i>source-port-alternate-hash</i> (<i>hash-bit-identifier</i>)	The alternate source port bit definition within the redirection hash table index.
<i>destination-port-alternate-hash</i> (<i>hash-bit-identifier</i>)	Sets the alternate destination port bit definition within the redirection hash table index.
<i>port1...port8</i>	<p>(Applies to a dynamic service group only. A dynamic service group is one identified by a <i>service number</i>.) A zero-terminated list of TCP port identifiers. If the <i>service-flags ports-defined</i> flag is set, packets will be matched against the set of ports supplied. If the <i>service-flags ports-source</i> flag is set, the ports are assumed to be source ports. Otherwise, the ports are assumed to be destination ports.</p>

<i>ip-address</i>	<p>Indicates the IP address of your network's home router. For version 2, <i>ip-address</i> can be a multicast address. (Multicast addresses are in the range 224.0.0.0 to 239.255.255.255, inclusive.)</p> <p>In version 2, multiple <i>ip-addresses</i> can be specified for unicast addressing. For multicast addresses, only one <i>ip-address</i> can be specified per service group.</p> <p>If you choose to specify the home router IP address, it is very important that the actual home router IP address and the home router IP address specified in this ProxySG configuration file match. If you do not already know the IP address of the home router, you can easily determine it from the router CLI by using the <code>show ip wccp</code> command.</p>
<i>domain-name</i>	<p>Specifies the domain name of your network's home router. Domain-name must be a valid domain name string that will successfully resolve on DNS lookup.</p>
<i>interface-number</i>	<p>Specifies the adapter interface number for the service group. Note that you cannot use a colon (0:0 or 0:1, for example).</p>
<i>string</i>	<p>(Applies to a dynamic service group only. A dynamic service group is one identified by a service number.) String can be at least one, and not more than eight, alphanumeric characters long.</p> <p>The password string specified here must match the password string declared for the router.</p>
<i>interface-number</i>	<p>(When used with the hash identifiers) Indicates the adapter interface to which the weight factor should be applied to alter the distribution of the primary hash table.</p>
<i>value</i>	<p>Specifies the weight factor value (0 through 255) that should be applied to the adapter interface indicated to alter the distribution of the primary hash table.</p>
<i>forwarding-type</i> [GRE L2]	<p>Switches between GRE encapsulation (the default) and L2 MAC address rewrite for forwarding packets. If this command is not present, GRE encapsulation is used.</p>

You can create a configuration file customized for the environment through the CLI inline commands or through a text file. The CLI inline commands enable WCCP on the ProxySG immediately; the drawback is that if any information changes, you must re-create the whole file using the inline command. With a text file, if any information changes, you can change the individual line; the drawback is that you must download the file again from an HTTP server to the ProxySG.

To use CLI commands to create a configuration file, continue with the next procedure. To use a text editor to create a configuration file, continue with "Creating a Configuration File using a Text File" on page 927.

Creating a Configuration File using CLI Inline Commands

For examples of various types of WCCP configurations, see "Examples" on page 928.

If you choose to configure through the CLI and the `inline` command, refer to the example below:

```
SGOS# configure terminal
SGOS#(config) inline wccp eof
```

where `eof` marks the beginning and end of the inline commands.

For example:

```
SGOS#(config) inline wccp eof
wccp enable
wccp version 2
service-group 9
forwarding-type L2
priority 1
protocol 6
service-flags destination-ip-hash
service-flags ports-defined
ports 80 21 1755 554 80 80 80 80
interface 6
home-router 10.16.18.2
end
eof
```

You created a WCCP configuration file and enabled WCCP on the ProxySG. WCCP setup is complete.

Creating a Configuration File using a Text File

If you create a configuration file using a text editor, assign the file the extension `.txt`. Note the following Blue Coat ProxySG configuration file rules:

- Only one command (and any associated parameters) is permitted, per line.
- Comments must begin with a semicolon (;) or a pound sign (#).
- Comments can begin in any column; however, all characters from the beginning of the comment to the end of the line are considered part of the comment and, therefore, are ignored.

For examples of various types of WCCP configurations, see "Examples" on page 928.

To Create a Configuration File using a Text Editor and Load the File on a ProxySG:

1. Open a text editor.
2. Using the commands described in "Syntax to create a customized configuration file:" on page 924, enter the arguments you need.
3. Copy the configuration file to an HTTP server so that it can be downloaded to the ProxySG.
4. Enable WCCP and download the WCCP configuration file using the following syntax:

```
wccp {enable | disable | no} [path config-file-url] | [version version-number]
```

where:

<code>enable</code>	Enables WCCP on the ProxySG.
<code>disable</code>	Disables WCCP on the ProxySG.
<code>no</code>	Indicates that you want to clear the current WCCP configuration settings.
<code><i>config-file-url</i></code>	Specifies the ProxySG WCCP configuration file or alternate configuration file.
<code><i>version-number</i></code>	Indicates the version of WCCP that your router is configured to use. If <code><i>version-number</i></code> is omitted, it is assumed to be 2.

For example:

```
SGOS# (config) wccp enable
SGOS# (config) wccp path http://205.66.255.10/files/wccp.txt
SGOS# (config) load wccp-settings
```

Examples

This section provides detailed examples of both the router and ProxySG configurations for:

- Standard HTTP redirection
- Standard HTTP redirection and a multicast address
- Standard HTTP redirection and a security password
- Standard transparent FTP
- A service group and alternate hashing

For information and examples about using WCCP, refer to

http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/fun_r/frprt3/frd3005.htm.

Displaying the Router's Known Caches

Use the router `show` command to display information about the ProxySG Appliances that are known to the router.

```
Router# show ip wccp web-caches
WCCP Web-Cache information:
IP Address:192.168.51.102
Protocol Version:0.3
State:Usable
Initial Hash
Info:FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
Assigned Hash:
Info:FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF
Hash Allotment:256 (100.00%)
Packets Redirected:0
Connect Time:00:00:31
Router# exit
```

Standard HTTP Redirection

The web-cache service group enables HTTP traffic redirection on port 80.

Router Configuration

The following example enables standard HTTP traffic redirection on a WCCP version 2-capable Cisco router.

```
Router (config) # ip wccp web-cache
Router (config) # interface ethernet 0/0
```

```
Router(config-if)# ip wccp web-cache redirect out
Router(config-if)# end
```

ProxySG Configuration

To enable the web-cache service group within the ProxySG, the following configuration file could be loaded.

```
# Enable WCCP to allow WCCP protocol communication between
# the ProxySG Appliance and the home router.
wccp enable
# By default, the WCCP version 2 protocol is assumed. An
# explicit "wccp version 2" command could be specified here.
service-group web-cache
# Specify the address for the router.
home-router 90.0.0.90
# Network interface 0 will participate.
interface 0
end
```

Standard HTTP Redirection and a Multicast Address

Configuring a multicast address on a WCCP-capable router provides reduced WCCP protocol traffic and the ability to easily add and remove caches and routers from a service group without having to reconfigure all service group members.

Router Configuration

The following example enables the standard HTTP traffic redirection on a WCCP version 2-capable Cisco router. In this case, WCCP protocol traffic is directed to the multicast address 226.1.1.1.

```
Router(config)# ip wccp web-cache group-address 226.1.1.1
Router(config)# interface ethernet 0/0
Router(config-if)# ip wccp web-cache group-listen
Router(config-if)# ip wccp web-cache redirect out
Router(config-if)# end
```

ProxySG Configuration

To enable the standard web-cache service group within the ProxySG, the following configuration file should be loaded. Note that in this example, both network interfaces 0 and 1 will participate within the service group. Both interfaces send and receive WCCP protocol packets by way of the multicast address.

```
# Enable WCCP to allow WCCP protocol communication between
# the ProxySG Appliance and the home router.
wccp enable
# By default, the WCCP version 2 protocol is assumed. An
# explicit "wccp version 2" command could be specified here.
service-group web-cache
# Specify the multicast address.
home-router 224.1.1.1
# Network interface 0 will participate.
interface 0
```

```
# Network interface 1 will also participate.  
interface 1  
end
```

Standard HTTP Redirection Using a Security Password

A simple eight-character password is configured within the router. This password must match the password configured within the ProxySG.

Router Configuration

The following example enables standard HTTP traffic redirection on a WCCP version 2-capable Cisco router.

```
Router(config)# ip wccp web-cache password 29gy8c2  
Router(config)# interface ethernet 0/0  
Router(config-if)# ip wccp web-cache redirect out  
Router(config-if)# end
```

ProxySG Configuration

To enable the standard WCCP version 2 service group within the ProxySG, the following configuration file could be loaded.

```
# Enable WCCP to allow WCCP protocol communication between  
# the ProxySG Appliance and the home router.  
wccp enable  
# By default, the WCCP version 2 protocol is assumed. An  
# explicit "wccp version 2" command could be specified  
# here.  
service-group web-cache  
# Specify the address for the router.  
home-router 90.0.0.90  
# Network interface 0 will participate.  
interface 0  
password 29gy8c2  
end
```

Standard Transparent FTP

In WCCP version 1, only HTTP traffic on port 80 could be redirected. In WCCP version 2, you can create a numbered service group that redirects other protocols on other ports.

You set the service group on the router, and tell the ProxySG which ports should be redirected.

Router Configuration

In this configuration, you create a new service group that you are dedicating to FTP redirects.

```
# Enables the service group that redirects ports besides 80.  
Router(config)# ip wccp 10  
# Enables a service group that allows user-defined  
# ports to be redirected.
```

```
Router(config)# int e0
Router(config-if)# ip wccp 10 redirect out
```

ProxySG Configuration

In this configuration, you take the service group created by the router and assign the characteristics to the group.

```
SGOS#(config) inline wccp eof
wccp enable
service-group 10
interface 0
home-router 10.1.1.1
protocol 6
priority 1
service-flags ports-defined
service-flags destination-port-hash
ports 20 21 80 80 80 80 80 80
eof
```

Reverse Proxy Service Group

This service group redirects IP packets for TCP destination port 80 traffic by hashing the source IP address.

Router Configuration

The following example enables the special ProxySG service group on a WCCP-capable router.

```
Router(config)# ip wccp 99
Router(config)#interface ethernet 0/0
Router(config-if)# ip wccp 99 redirect out
Router(config-if)# end
```

ProxySG Configuration

To configure the special ProxySG service group on the appliance, a dynamic service group must be created as illustrated by the following example.

```
# Enable WCCP to allow WCCP protocol communication between
# the ProxySG Appliance and the home router.
wccp enable
# By default, the WCCP version 2 protocol is assumed. An
# explicit "wccp version 2" command could be specified here.
# Service Group 99 is specially identified within the router
# as representing the ProxySG Appliance service.
service-group 99
# Specify the address for the router.
home-router 90.0.0.90
# Network interface 0 will participate.
interface 0
# Specify the TCP protocol.
protocol 6
# The hash should be based on the source IP address.
```

```
service-flags source-ip-hash
end
```

Service Group with Alternate Hashing

You can create a special service group on a WCCP-capable router that uses alternate hashing when hot spots are detected. This service group redirects IP packets by hashing the source IP address.

Router Configuration

In this configuration, you create a new service group that you are dedicating to website hot spots.

```
Router(config)# ip wccp 5
Router(config)# interface ethernet 0/0
Router(config-if)# ip wccp 5 redirect out
Router(config-if)# end
```

ProxySG Configuration

To configure this special service group on the ProxySG, a dynamic service group must be created.

```
# Enable WCCP to allow WCCP protocol communication between
# the ProxySG Appliance and the home router.
wccp enable
# By default, the WCCP version 2 protocol is assumed. An
# explicit "wccp version 2" command could be specified here.
# Service Group 5 will be created to redirect standard HTTP
# traffic and use an alternate hash function based on the
# source IP address, if necessary.
service-group 5
# Specify the address for router 1.
home-router 90.0.0.90
# Specify the address for router 2.
home-router 90.0.1.5
# Network interface 0 will participate.
interface 0
# Specify the TCP protocol.
protocol 6
# The following two flags specify that a hash function based
# on the destination IP address should be applied first. If
# a hot-spot is detected, then an alternate hash
# function using the source IP address should be used.
service-flags destination-ip-hash
service-flags source-ip-alternate-hash
end
```

Troubleshooting: Home Router

If you install WCCP settings and then later upgrade the Cisco IOS software or change network configuration by adding a device with a higher IP address, the change might result in a different home router IP assignment. WCCP might or might not work under these conditions, and performance might decrease. If you upgrade the router software or change the network configuration, verify that the actual home router IP address and home router IP address in the WCCP configuration match.

To Verify the Home Router IP Address Matches the Home Router IP Address Listed in the WCCP Configuration:

1. From the router CLI, view the WCCP configuration:

```
Router#(config) show ip wccp
```

The home router information appears, similar to the example below:

```
Global WCCP information:
Router information:
Home router Identifier:195.200.10.230
Protocol Version:2.0
```

2. From the Blue Coat ProxySG, verify that the home router IP address specified in the ProxySG WCCP configuration file is the same as the actual home router IP address discovered through the router CLI command. The following is a ProxySG WCCP configuration file showing the same home router IP as in the example above:

```
SGOS# show wccp config
;WCCP Settings
;Version 1.3
wccp enable
wccp version 2
service-group web-cache
interface 1
home-router 195.200.10.230
end
```

In this case, the two home router identifiers match.

Identifying a Home Router/Router ID Mismatch

The following is some helpful information for resolving a home-router/Router ID mis-match that results in the router crashing the ProxySG. This situation can occur when the router interface is set to a higher IP address than the home-router and WCCP messages show w/bad rcv_id.

Note that WCCP version 1 does not care what home router the cache had configured. So if you upgrade from WCCP version 1 to WCCP version 2, the router might pick a different IP address than was configured as a home router in the cache.

This means that a mismatch can occur after an upgrade.

ProxySG Configuration

Use the `show wccp statistics` command to identify the configured home router and the highest router IP.

```
SGOS#(config) show wccp statistics
Service Group ident.      :512,1,9, 1,6,18, 1755,554,20,21,80,80,80,80
Home Routers          :10.2.3.224 <<=====Configured Home Router IP
Hotspots announced      :0
Assignment state        :idle
Designated Cache      :10.2.3.228 <<=====Blue Coat IP
Announcement key #      :2
Cache view change #     :13 <<==== # times cache view changed
```

```
Router View Changed          :0
Recent hit count             :0
Primary hit count            :0
Alternate hit count          :0
Instance IP address :10.2.3.228 <<=====Blue Coat IP
Sequence info                :10.2.3.231,636
Query response info:
Active                       :1
Primary hash weight          :0
Hotspot information          :0,0,0,0
Total assign weight          :0
Router IP address :10.2.3.231 <<=====Router ID/Highest IP on Router
Receive #                    :636
Change #                     :4
Activation time               :Wed, Jan 30 2002 00:17:58 UTC
Last I-See-You time          :Wed, Jan 30 2002 01:08:58 UTC
Active caches                 :10.2.3.228
Assignment key                :10.2.3.228,2
Router state                  :active
Cache                         :10.2.3.228,L,D
Active                        :1
```

Notice that .231 is highest IP on router and is automatically selected as the home router, even though .224 is the configured home router IP.

You can also use the `show wccp configuration` command if you already know the highest IP and just want to know what the Security Gateway identifies as the home-router.

```
SGOS#(config) show wccp configuration
;WCCP Settings
;Version 1.3
wccp enable
wccp version 2
service-group 9
interface 0
home-router 10.2.3.224
protocol 6
priority 1
service-flags ports-defined
service-flags destination-ip-hash
ports 1755 554 20 21 80 80 80 80
```

Router Configuration

The configuration below reveals that two interfaces are active on the router, and that one of the IP addresses is higher than the home router configured in the ProxySG configuration file. The higher IP address takes over duties as the home router, causing a mismatch between the router and the ProxySG.

```
Router# show conf
Using 689 out of 129016 bytes
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
```

```

hostname NachoL3
enable secret 5 $1$r6nJ$dr58AZ.ZDg6RKA6MYeGRb.
enable password nacho

ip subnet-zero
no ip routing
ip wccp 9

interface FastEthernet0/0
  ip address 10.2.3.224 255.255.255.0
  ip wccp 9 redirect out
  no ip route-cache
  no ip mroute-cache
  speed 100
  half-duplex
!
interface FastEthernet0/1
  ip address 10.2.3.231 255.255.255.0
  no ip route-cache
  no ip mroute-cache
  speed 100
  half-duplex

```

Correcting a Home Router Mismatch

The home router must have the same IP address on both the router and the ProxySG. Every time a higher IP address is introduced to the router, the higher address becomes the home router.

On a WCCP router, the `Router Identifier` parameter is dynamically assigned. It cannot be manually configured.

To Set the Correct Home Router IP Address on the ProxySG:

You cannot edit a WCCP configuration file created by the SGOS inline commands. You must recreate the configuration file. For more information on creating a WCCP configuration file using CLI commands on a ProxySG, see "Creating a Configuration File using CLI Inline Commands" on page 926.

If you created a text file and downloaded it, you can edit the file and then download it again to the ProxySG. For more information for editing the WCCP text file and downloading it, see "Creating a Configuration File using a Text File" on page 927.

Tips

- If you use IP spoofing with WCCP, do the following for best results:
The `ip wccp redirect exclude in` command should be applied to the adapter to which the ProxySG is attached.
- For L2 forwarding, the ProxySG should be directly connected to the router interface.

Appendix D: RIP Commands

You can place any of the commands below into a Routing Information Protocol (RIP) configuration text file. Note that you cannot edit a RIP file through the command line. You can overwrite a RIP file using the `inline rip-settings` command.

Once the file is complete, place it on an HTTP or FTP server accessible to the ProxySG and use the following commands to install the file on the ProxySG:

At the `(config)` command prompt:

```
SGOS#(config) rip path url
SGOS#(config) load rip-settings
```

For more information on installing the RIP configuration file, see "Using RIP" on page 87.

net

```
net Nname[/mask] gateway Gname metric Value {passive | active | external}
```

Syntax

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 <code>net Nname/32</code> specification is equivalent to the <code>host Hname</code> command.
<code>passive active external</code>	Indicates whether the gateway should be 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}
```

Syntax

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 <code>net Nname/32</code> specification is equivalent to the <code>host Hname</code> command.
<code>passive active external</code>	Indicates whether the gateway should be 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:

Parameters:	Description
if=[0 1 2 3]	Indicates 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 that will be 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 via 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.
no_solicit	Disables the transmission of Router Discovery Solicitations.
send_solicit	Specifies that Router Discovery solicitations should be sent, even on point-to-point links, which by default only listen to Router Discovery messages.
no_rdisc_adv	Disables the transmission of Router Discovery Advertisements.
rdisc_adv	Specifies that Router Discovery Advertisements should be sent, even on point-to-point links, which by default only listen to Router Discovery messages.
bcast_rdisc	Specifies that Router Discovery packets should be broadcast instead of multicast.
rdisc_pref=N	Sets the preference in Router Discovery Advertisements to the integer N.
rdisc_interval=N	Sets the nominal interval with which Router Discovery Advertisements are transmitted to N seconds and their lifetime to 3*N.
trust_gateway=rname	Causes RIP packets from that router and other routers named in other trust_gateway keywords to be accept, and packets from other routers to be ignored.
redirect_ok	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 configuration:

Parameters:	Description
supply_routing_info -or- advertise_routes	<p>-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.</p> <p>-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.</p> <p>-h option: Suppress_extra_host_routes advertise_host_route</p> <p>-m option: Advertise_host_route on multi-homed hosts</p> <p>-A option: Ignore_authentication //</p>
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`.

Appendix E: Diagnostics

Blue Coat Systems has a number of resources to provide diagnostic information:

- **Heartbeats:** Enabled by default, Heartbeats (statistics) are a primary diagnostic tool used by Blue Coat, allowing them to proactively monitor the health of ProxySG 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 a 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 Appendix B: “Troubleshooting” in the *Blue Coat Content Policy Language Guide*.
- **Event Logging:** The event log files contain messages generated by software or hardware events encountered by the ProxySG. For information on configuring event logging, see “Event Logging and Notification” on page 797.
- **Access Logging:** Access logs allow for analysis of Quality of Service, content retrieved, and other troubleshooting. For information on Access Logging, see Chapter 20: “Access Logging” on page 743.
- **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.
- `tracert`: 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 device).
- `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 *ProxySG Command Line Interface Reference Guide*.

Note: If you cannot access the Management Console at all, be sure that you are using HTTPS (`https://ProxySG_IP_address:8082`). This more secure option was added in SGOS 4.x. If you want to use HTTP, you must explicitly enable it before you can access the Management Console.

This appendix discusses the following topics:

- "Diagnostic Reporting (Service Information)" on page 942. This includes taking snapshots of the system.
- "Packet Capturing (the PCAP Utility)" on page 951.
- "Core Image Restart Options" on page 957.
- "Diagnostic Reporting (Heartbeats)" on page 958.
- "Diagnostic Reporting (CPU Monitoring)" on page 959

If the ProxySG does not appear to work correctly and you are unable to diagnose the problem, contact Blue Coat Technical Support.

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 will also be sent.

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 through the Management Console:

1. Select Maintenance>Service Information>Send Information>General.
The General tab displays.

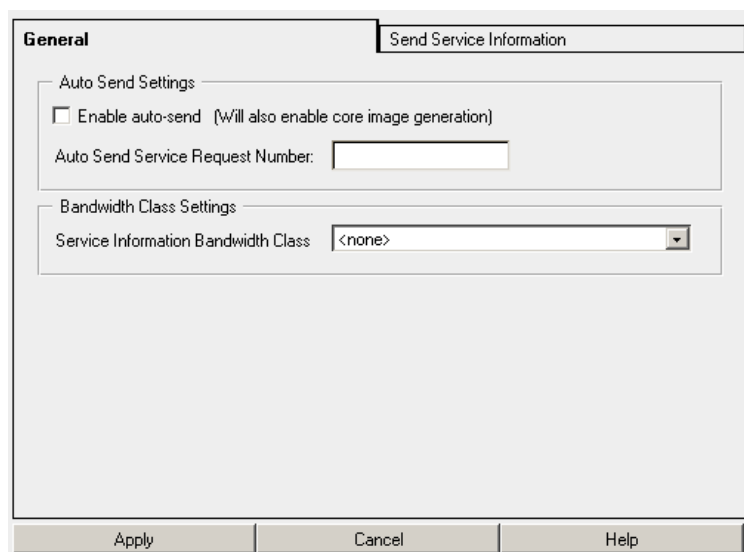


Figure E-1: Service Information General Tab

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.

To Send Service Information Automatically through the CLI:

1. To enable (or disable) the automatic service information feature, enter the following commands at the (config) command prompt:

```
SGOS#(config) diagnostics
SGOS#(config diagnostics) service-info
SGOS#(diagnostics service-info) auto {enable | disable}
SGOS#(diagnostics service-info) auto sr-number sr_number
```

where:

enable		Enables the automatic service information feature.
disable		Disables the automatic service information feature.
sr-number	<i>sr_number</i>	Sets the service-request number for the automatic service information feature.

2. (Optional) To clear the service-request number, enter the following command:

```
SGOS#(diagnostics service-info) auto no sr-number
```

Managing the Bandwidth for Service Information

If you have a license for the bandwidth management feature, 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. See Chapter 10: “Bandwidth Management” on page 375 for information about using bandwidth management. Note that changing to a new bandwidth management class will 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, will affect transfers already in progress if that class was selected prior to initiating the transfer.

To Manage Bandwidth for Service Information through the Management Console:

1. Select Maintenance>Service Information>Send Information>General.

The General tab displays.

2. To manage the bandwidth of automatic service information, select a bandwidth class from the Service Information Bandwidth Class drop-down menu.

Note: Before you can manage the bandwidth for the automatic service information feature, you must first enable bandwidth management and create an appropriate bandwidth-management class. Bandwidth management is enabled by default if you have a valid license for this feature. See Chapter 10: “Bandwidth Management” on page 375 for information about enabling bandwidth management and creating and configuring bandwidth classes.

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.

To Manage Bandwidth for Service Information through the CLI:

1. To manage the bandwidth of automatic-service information, enter the following command:

```
SGOS#(diagnostics service-info) bandwidth-class bandwidth_class_name
```

where *bandwidth_class_name* is the name of the bandwidth class that you have created and configured to manage the bandwidth of service information.

Note: Before you can manage the bandwidth for the automatic service information feature, you must first enable bandwidth management and create an appropriate bandwidth-management class. Bandwidth management is enabled by default if you have a valid license for this feature. See Chapter 10: “Bandwidth Management” on page 357 for information about enabling bandwidth management and creating and configuring bandwidth classes.

2. (Optional) To disable the bandwidth-management of service information, enter the following command:

```
SGOS#(diagnostics service-info) no bandwidth-class
```

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 942.

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/index.html> for details on opening a service request ticket.

The following list details information that you can send:

- Packet Capture
- Event Log
- Memory Core
- SYSInfo
- Access Logs (can specify multiple)
- Snapshots (can specify multiple)
- Contexts (can specify multiple)

To Send Information through the Management Console:

1. Select Maintenance>Service Information>Send Information>Send Service Information.

The Send Service Information tab displays.

The screenshot shows a window titled "Send Service Information" with a "General" tab selected. Inside the window, there is a "Service Request Number:" label followed by an empty text input field. Below this, the section "Information to send:" contains several checkboxes: "Packet Capture (Unknown)", "Event Log (40,960)", "Memory Core (Unknown)", "SYSInfo (Unknown)", "Access Logs", "Snapshots", and "Contexts". To the right of the "Access Logs", "Snapshots", and "Contexts" checkboxes are buttons labeled "Select access logs to send", "Select snapshots to send", and "Select contexts to send" respectively. Below the "Contexts" checkbox is a "Select Newest" button. At the bottom of the main content area are two buttons: "Send" and "View Progress". At the bottom right of the window is a "Help" button.

Figure E-2: Send Service Information Tab

2. Enter the service-request number that you received from a Technical Support representative (the service-request number is in the form xx-xxxxxxx or x-xxxxxxx).
3. Select the appropriate checkboxes (as indicated by a Technical Support representative) in the Information to send field.

Note: Options for items that you do not have on your system will be grayed out and you will not be able to select that checkbox.

4. (Optional) If you select Access Logs, Snapshots, or Contexts checkbox, 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 appears:

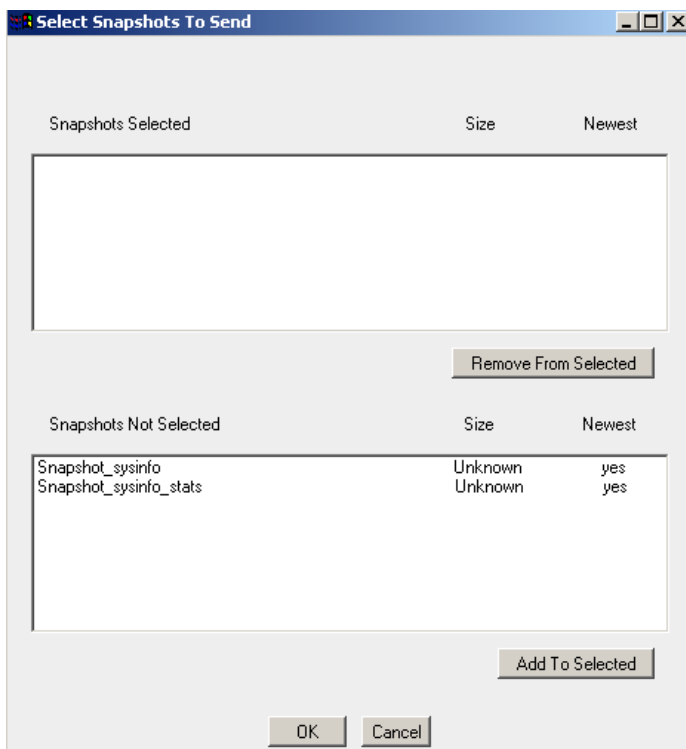


Figure E-3: Select Snapshots to Send Dialog

- To select information to send, highlight the appropriate selection in the Access Logs/Snapshots/Contexts Not Selected field and click Add to Selected.
 - To remove information from the Access Logs/Snapshots/Contexts Selected field, highlight the appropriate selection and click Remove from Selected.
 - Click Ok.
5. Click Send.
 6. Click Ok in the Information upload started dialog that appears.

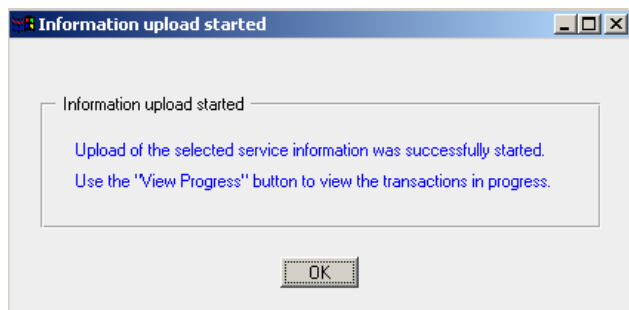


Figure E-4: Information Upload Started Dialog

7. (Optional) Click View Progress to open a window displaying the current transactions in progress; click Ok to close the window.

To Send Information through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) diagnostics
SGOS#(config diagnostics) service-info
SGOS#(config service-info) view available
SGOS#(config service-info) send sr_number one_or_more_commands_from_view_available
SGOS#(config service-info) view status
SGOS#(config service-info) cancel {all | one_or_more_from_view_status}
SGOS#(config service-info) exit
```

where:

cancel	all	Cancels all service information being sent to Blue Coat.
	<i>one_or_more_from_view_status</i>	Cancels certain service information items being sent to Blue Coat. These items can be chosen from the list provided by the <code>view status</code> command.
send	<i>sr_number</i>	Specifies the service-request number to send to Blue Coat (you must also select one or more of the <code>view available</code> commands).
	<i>one_or_more_commands_from_view_available</i>	Specifies the command or commands to send to Blue Coat (you must also specify a service-request number). Choose commands from those listed under the <code>view available</code> command.
view	available	Shows the list of service information than can be sent to Blue Coat.
	status	Shows the transfer status of service information to Blue Coat.
exit		Exits configure diagnostics service-info mode and returns to configure diagnostics mode.

Example:

```

SGOS#(config) diagnostics
SGOS#(config diagnostics) service-info
SGOS#(diagnostics service-info) view available
Service information that can be sent to Blue Coat

Name                                     Approx Size (bytes)
Event_log                               188,416
System_information                       Unknown
Snapshot_sysinfo                         Unknown
Snapshot_sysinfo_stats                   Unknown

SGOS#(diagnostics service-info) send 1-4974446 event_log system_information
snapshot_sysinfo
Sending the following reports
Event_log
System_information
Snapshot_sysinfo

SGOS#(diagnostics service-info) view status
Name                                     Transferred      Total Size      % Done
Event_log                               Transferred successfully
Snapshot_sysinfo                         Transferred successfully
Event_log                               Transferred successfully
System_information                       Transferred successfully
SGOS#(diagnostics service-info) exit
SGOS#(config diagnostics) exit
SGOS#(config)

```

Creating and Editing Snapshot Jobs

To Create a New Snapshot Job through the Management Console:

1. Select Maintenance>Service Information>Snapshots.

The Snapshots tab displays.

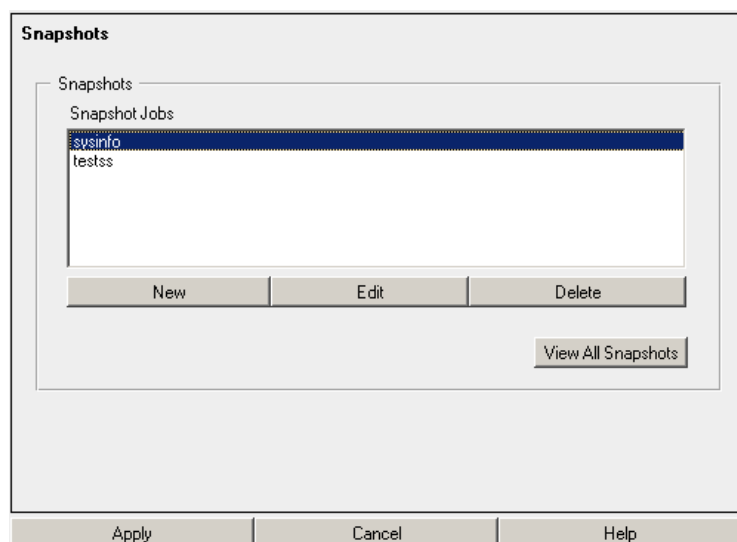


Figure E-5: Snapshots Tab

2. Click New.
3. Enter a snapshot job into the Add list item dialog that displays; click Ok.
4. Click Apply.
5. (Optional) To view snapshot job information, click View All Snapshots. Close the window that opens when you are finished viewing.

To Create a New Snapshot Job through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) diagnostics
SGOS#(config diagnostics) snapshot create snapshot_name
```

To Edit an Existing Snapshot Job through the Management Console:

1. Select Maintenance>Service Information>Snapshots.
The Snapshots tab displays.
2. Select the snapshot job you want to edit (highlight it).
3. Click Edit.
The Edit Snapshot dialog displays.

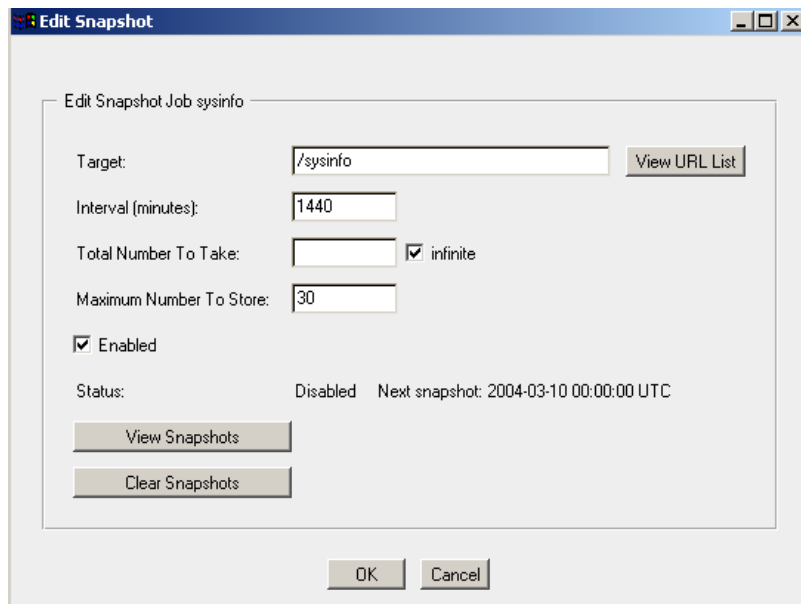


Figure E-6: Edit Snapshot Dialog

4. Enter the following information into the Edit Snapshot fields:
 - Target: Enter the object to snapshot.
 - Interval (minutes): Enter the interval between snapshot reports.
 - Total Number To Take: Enter the total number of snapshots to take or select the Infinite checkbox to take an infinite number of snapshots.
 - Maximum Number To Store: Enter the maximum number of snapshots to store.
 - Enabled: Check this box to enable this snapshot job or uncheck 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.

To Edit an Existing Snapshot Job through the CLI:

At the (config) command prompt, enter the following commands:

```
SGOS#(config) diagnostics
SGOS#(config diagnostics) snapshot edit snapshot_name
SGOS#(config snapshot snapshot_name) clear-reports
SGOS#(config snapshot snapshot_name) disable
SGOS#(config snapshot snapshot_name) enable
SGOS#(config snapshot snapshot_name) exit
SGOS#(config snapshot snapshot_name) interval minutes
SGOS#(config snapshot snapshot_name) keep number_to_keep (from 1 - 100)
SGOS#(config snapshot snapshot_name) take infinite | number_to_take
```

```
SGOS#(config snapshot snapshot_name) target object_to_fetch
SGOS#(config snapshot snapshot_name) view
```

where:

clear-reports		Clears all stored snapshots reports.
disable		Disables this snapshot job.
enable		Enables this snapshot job.
exit		Exits configure diagnostics snapshot name mode and returns to configure diagnostics service-info mode.
interval	<i>minutes</i>	Specifies the interval between snapshots reports in minutes.
keep	<i>number_to_keep</i> (from 1 - 100)	Specifies the number of snapshot reports to keep.
take	<i>infinite</i> <i>number_to_take</i>	Specifies the number of snapshot reports to take.
target	<i>object_to_fetch</i>	Specifies the object to snapshot.
view		Displays snapshot status and configuration.

Packet Capturing (the PCAP 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. The maximum PCAP size allowed is 100MB. 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 via an interface's MAC address or via an interface's broadcast address. The utility also captures transmitted packets that are sent from the ProxySG. 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, Ethereal or Packet Sniffer Pro 3.0).

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 as well as 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 953.

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 E.1: Common Filter Expressions

Filter Expression	Packets Captured
<code>ip host 10.25.36.47</code>	Captures packets from a specific host with IP address <code>10.25.36.47</code> .
<code>not ip host 10.25.36.47</code>	Captures packets from all IP addresses except <code>10.25.36.47</code> .
<code>ip host 10.25.36.47 and ip host 10.25.36.48</code>	Captures packets from two IP addresses: <code>10.25.36.47</code> and <code>10.25.36.48</code> .
<code>ether host 00:e0:81:01:f8:fc</code>	Captures packets from MAC address <code>00:e0:81:01:f8:fc</code> .
<code>port 80</code>	Captures packets to port <code>80</code> .
<code>Ip src bluecoat.com</code>	Captures all packets that came from the host <code>bluecoat.com</code> to the ProxySG.
<code>Host example.com and tcp</code>	Captures all TCP packets sent between the host <code>example.com</code> and the ProxySG.

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" above.

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 through the Management Console:

1. Select Maintenance>Service Information>Packet Captures.

The Packet Captures tab displays.

Figure E-7: Packet Captures Tab

2. To configure packet capturing, complete the following steps:
 - To define or change the PCAP filter, enter the filter information into the Capture filter field. (See "Common PCAP Filter Expressions" on page 952 for information about PCAP filter expressions for this field.) To remove the filter, clear this field.

- To specify the number of kilobytes to capture, check the Include packets in core checkbox and enter a number. 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.
 - To capture all packets, even those that are bridged, check the Capture bridging packets checkbox. Normally, the packets that are bridged from one interface to another (see "Software and Hardware Bridges" on page 75) are not included in the packet capture.
3. Choose one of the following three radio buttons:
 - 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.
 - 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.
 4. Click Apply.
 5. To start the capture, click the Start capture button. This button will be grayed out if a packet capture is already started.
 6. To stop the capture, click the Stop capture button. This button will be grayed out if a packet capture is already stopped.
 7. To download the capture, click the Download capture button. This button will be grayed out if no file is available for downloading.

To Define Packet Capturing Settings through the CLI:

1. To define PCAP filter parameters, enter the following command at the enable command prompt:

```
SGOS# pcap filter parameters
```

This captures packets according to the parameters set. If no parameters are set, all packets are captured until the `pcap stop` command is issued.

See "Using Filter Expressions in the CLI" on page 952 for information about CLI filter parameters.
2. To begin capturing packets, enter the following command at the enable command prompt:

```
SGOS# pcap start {first number | last number | capsizesize number (kilobytes) | trunc number}
```

where:

first *number* allows you to enter the number of matching packets (*number*) 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.

last *number* allows you to enter the number of matching packets (*number*) 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 `last` and `first` options supersede each other.

capsize *number (kilobytes)* allows you to stop the collection after *number* kilobytes (up to 100 MB) of packets have been captured. This command prevents packet capturing from taking up too much memory and degrading performance. If no parameter is specified, the default is to capture packets until the stop directive is issued.

trunc *number* allows collecting, at most, *number* bytes of packets from each frame.

To Enable, Stop, and Download Packet Captures through a Browser:

1. Start your Web browser.
2. Enter the URL: `https://ProxySG_IP_address:8082/PCAP/Statistics` and log on to the ProxySG as needed.

The Packet Capture Web page opens.

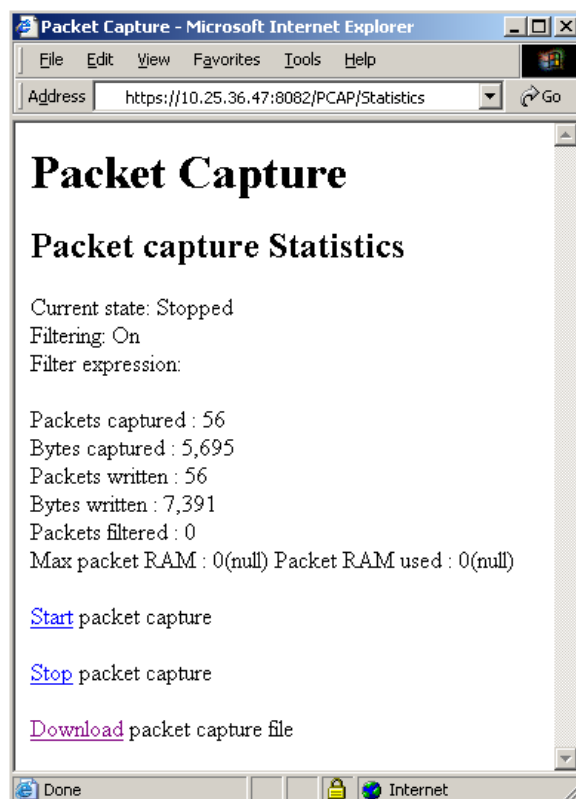


Figure E-8: Packet Capture Web Page

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 enable 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 Data through the Management Console:

1. Select Maintenance>Service Information>Packet Captures.

The Packet Captures tab displays.

2. To view the packet capture statistics, click the Show statistics button.

A window opens displaying the statistics on the current packet capture settings. Close the window when you are finished viewing the statistics.

To View Current Packet Capture Data through the CLI:

At the enable command prompt, enter the following command:

```
SGOS# pcap info
packet capture information:
Packets captured:          12
Bytes captured:           1879
Packets written:          12
Bytes written:            2343
Max packet ram:           16384
Packet ram used:          2167
Packets filtered:         405
Bridge capture all:       Disabled
Current state:             Stopped
Filtering:                 On
Filter expression:         iface out expr ""
```

Uploading Packet Capture Data

Use the following steps to transfer packet capture data from the ProxySG to an FTP site through the CLI. You cannot use the Management Console. After uploading is complete, you can analyze the packet capture data.

To Upload Packet Captures to a Server through the CLI:

At the enable command prompt, enter the following command:

```
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 ProxySG 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 will be retained. The default value is 2; the range is between 1 and 10.

To Configure Core Image Restart Options through the Management Console:

1. Select Maintenance>Core Images.

The Core Images tab displays.

Figure E-9: Configuring Core Image Restart Options

2. Select a core image restart option.
3. (Optional) Select the number of core images that will be retained from the Number of stored images drop-down list.
4. Click Apply.

To Configure Core Image Restart Options through the CLI:

1. At the (config) command prompt, enter the following command:

```
SGOS#(config) restart core-image {context | full | none}
```

- (Optional) To select the number of core images that will be retained, enter the following command:

```
SGOS# (config) restart core-image keep number
```

Diagnostic Reporting (Heartbeats)

The ProxySG diagnostic reporting configurations are located in the Management Console (under the Maintenance>Heartbeats tab), and in the CLI (under the configuration diagnostics submenu).

The daily heartbeat is a periodic message that is sent every 24 hours and contains ProxySG statistical data. Besides telling the recipient that the device is alive, heartbeats also are an indicator of the ProxySG Appliance's health. Heartbeats do not contain any private information; they contain only aggregate statistics that can be used to preemptively diagnose support issues. The daily heartbeat is encrypted and transferred to Blue Coat using HTTPS. Administrators can have the daily heartbeat messages e-mailed to them by configuring event log notification. The content that is e-mailed to the administrator is the same content sent to Blue Coat. For more information about e-mail notification, see "Enabling Event Notification" on page 800.

If Blue Coat monitoring is enabled, Blue Coat receives encrypted information over HTTPS whenever the ProxySG is rebooted. The information does not contain any private information; it contains restart summary information, in addition to daily heartbeat information. This allows the tracking of ProxySG unexpected restarts due to system issues, and allows Blue Coat to address system issues preemptively.

If the daily heartbeats setting is disabled, you can still send a heartbeat message by using the `send-heartbeat` command through the CLI (this feature is not available through the Management Console).

To Set Daily Heartbeats and/or Blue Coat Monitoring through the Management Console:

- Select Maintenance>Heartbeats.

The Heartbeats tab displays.

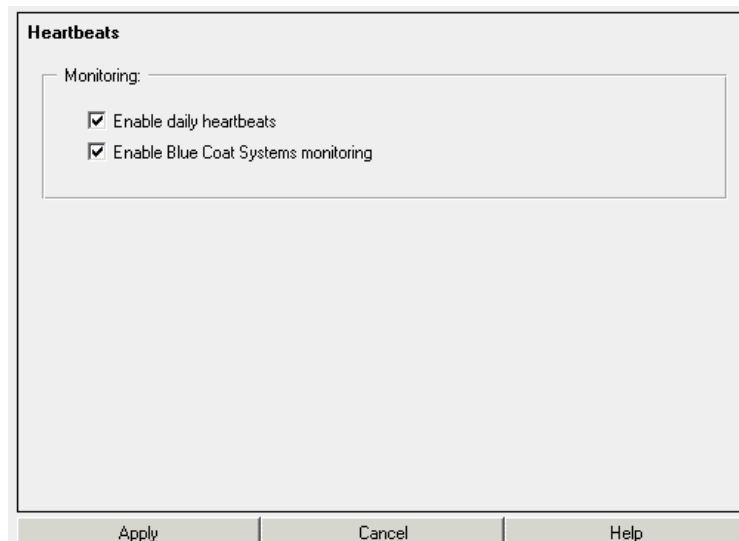


Figure E-10: Maintenance Heartbeats Tab

2. Select or deselect the Enable daily heartbeats or Enable Blue Coat monitoring checkbox.
3. Click Apply.

To Set Daily Heartbeats through the CLI:

At the (config) command prompt, enter the following command:

```
SGOS#(config) diagnostics  
SGOS#(config diagnostics) heartbeat enable
```

To Set Blue Coat Monitoring through the CLI:

At the (config) command prompt, enter the following command:

```
SGOS#(config) diagnostics  
SGOS#(config diagnostics) monitor enable
```

To Send an Immediate Heartbeat Message through the CLI:

At the (config) command prompt, enter the following command:

```
SGOS#(config) diagnostics  
SGOS#(config diagnostics) send-heartbeat
```

Note: This option is not available through the Management Console.

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 through the Management Console:

1. Select Statistics>Advanced.

A list of links to advanced statistic URLs displays.

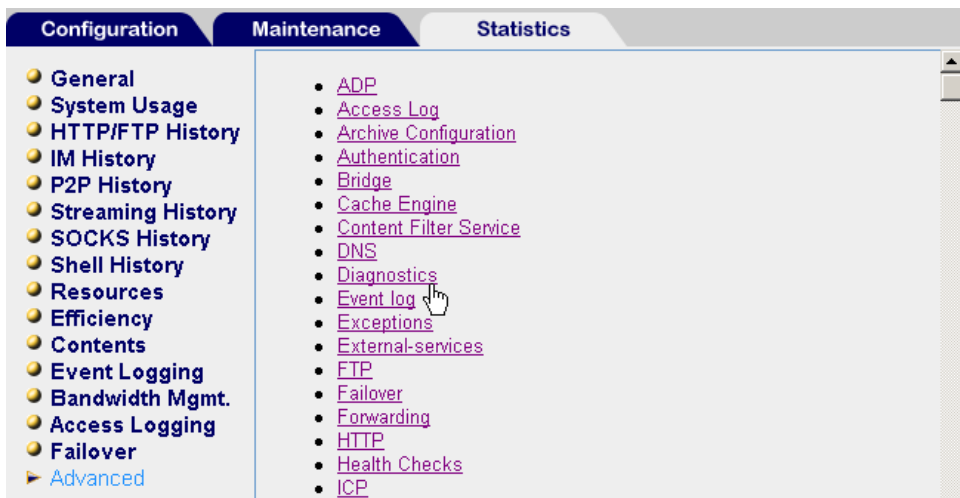


Figure E-11: Advanced Statistics Tab

2. Click the Diagnostics link.

A list of links to Diagnostic URLs displays.



Figure E-12: Diagnostic URL Links in the Advanced Statistics Tab

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.

To Configure and View CPU Monitoring through the CLI:

1. To enable or disable CPU monitoring, enter the following commands at the (config) command prompt:

```
SGOS#(config) diagnostics
SGOS#(config diagnostics) cpu-monitor {enable | disable}
```

2. To set the interval between CPU monitoring, enter the following command:

```
SGOS#(config diagnostics) cpu-monitor interval seconds
```

where *seconds* is a number from 1 to 59 that sets the frequency (in seconds) that the CPU monitor statistics are updated.

3. To view CPU monitoring results, enter the following command:

```
SGOS#(config diagnostics) view cpu-monitor
CPU Monitor:
Configured interval duration: 5 seconds
Current interval complete in: 3 seconds
CPU 0
      HTTP and FTP      36%
      Object Store     19%
      Miscellaneous     15%
      Miscellaneous     1%

CPU 1
      TCPIP            18%
      HTTP and FTP     15%
      HTTP and FTP     3%
```

If the CPU monitor is disabled, the `view` display will show the following message:

```
SGOS#(config diagnostics) view cpu-monitor
CPU Monitor is not running. Enable in diagnostics menu
```

Note: The total percentages 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 commands `SGOS#(config) show cpu` and `SGOS#(config diagnostics) view cpu-monitor` can sometimes display CPU statistics that differ by about 2-3%. This occurs because different measurement techniques are used for the two displays.

Appendix F: Using Blue Coat Director to Manage Multiple Appliances

If you are managing multiple ProxySG Appliances, you might find it easier to use Blue Coat Systems standalone product, Blue Coat Director, than to configure and control the appliances individually.

Director allows you to configure a ProxySG and then push that configuration out to as many ProxySG Appliances as you need. Director also allows you to delegate network and content control to multiple administrators and distribute user and content policy across a Content Delivery Network (CDN). With Director, you can:

- Reduce management costs by centrally managing all Blue Coat ProxySG Appliances.
- Eliminate the need to configure each remote ProxySG manually.
- Recover from system problems with configuration snapshots and recovery.

Configuration management specifically includes:

- Configure groups of ProxySG Appliances based on locations, applications, or other factors
- Delegate ProxySG administration by access level, group, or policy
- Rapidly deploy standardized configurations using profiles
- Manage the scheduling of policy and configuration changes
- Easily schedule incremental changes to one or more ProxySG Appliances
- Create and distribute policy across a system of ProxySG Appliances
- Automatically back up configuration snapshots
- Back up ProxySG backup files
- Compare backup files from different ProxySG Appliances
- Restore configuration backups to multiple ProxySG Appliances
- Automatically distribute software licenses
- Quickly monitor ProxySG status, statistics, and configurations
- Upgrade an entire content-smart network at once

How Director Works with ProxySG

Director consists of a management node (a *domain*) and the ProxySG Appliances that you want to manage. The appliances can be added to the domain through either Director's CLI or Management Console.

Note: Do not mix ProxySG versions within a domain; errors might result if you try to push the same configuration to machines that are running different versions of SGOS.

When a ProxySG is added to the domain, you provide connection information about the ProxySG: name (meaningful to you), IP address or full hostname, username/password, authentication method and credentials, and, optionally, a description.

Only the appliances added to the domain can be managed by the domain. Multiple domains can be created.

Once added to the domain, you can manage the ProxySG either individually, through the Quick View/Edit module, or you can manage multiple appliances through the Configuration Management module.

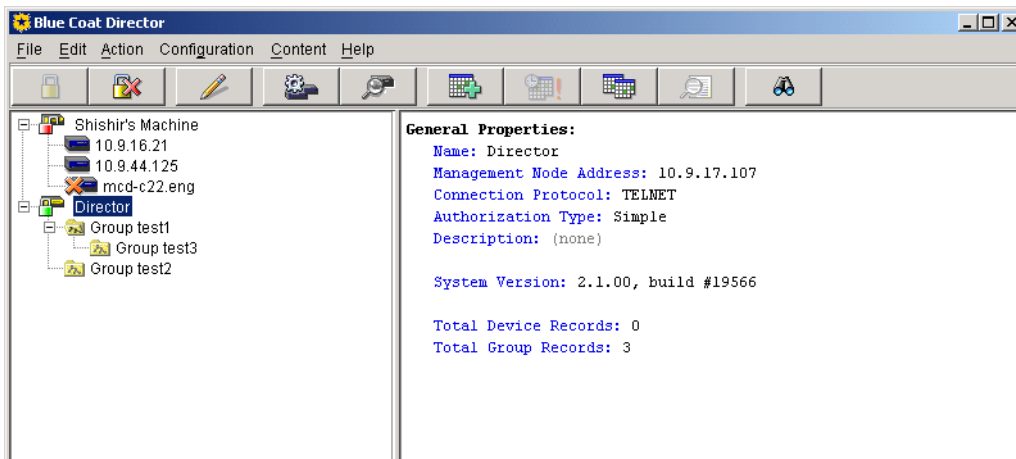


Figure F-1: Director Management Console

Communication Between Director and the ProxySG

Director and the ProxySG use SSHv2 as the default communication mode. To use SSHv1 or Telnet, you must do additional configuration on the ProxySG.

For Director to successfully manage multiple ProxySG Appliances, it must be able to communicate with a ProxySG using SSH/RSA and the Director's public key must be put on each ProxySG that Director manages. This creates a *golden profile*, meaning that the ProxySG is fully authenticated and can be used to push configurations to multiple ProxySGs using the same version of the software.

At initial set up of the ProxySG on Director, Director connects to the device using the authentication method established on the device: Telnet, SSH with simple authentication, or SSH/RSA. SSH/RSA is preferred, and must also be set up on Director before connecting to the ProxySG.

Note: You cannot connect to a ProxySG using Telnet without first enabling the Telnet-Console on the ProxySG.

Director can create an RSA keypair for a ProxySG to allow connections. However, for full functionality, Director's public key must be put on each ProxySG. You can put the key on the ProxySG two ways:

- Use Director to create and push the key.
- Use the `import-director-client-key` CLI command from the ProxySG.

Using Director to create and push client keys is the recommended method. The CLI command is provided for reference.

Complete the following steps to put Director's public key on the ProxySG using the CLI of the ProxySG. You must complete this procedure from the CLI. The Management Console is not available.

Note: For information on creating and pushing a SSH keypair on Director, refer to the *Blue Coat Director Installation Manual*.

Login to the ProxySG you want to manage from Director.

Using SGOS 2.x, complete the following steps:

1. From the (config) prompt, enter SSH mode.

```
SGOS#(config) ssh
```

The prompt changes to SGOS#(config sshd)

2. 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" (without quotes) must be the username, allowing you access to passwords in clear text.

```
SGOS#(config sshd) import director-client-key
```

```
Paste client key here, end with "... " (three periods)
```

```
1024 35
```

```
12999972424841580015120232422276923944398398687186826236351866926327873980944361
```

```
40021875567656584396224129186767257689612899184491932211183030490667098587892725
```

```
82600126521236261900215886326073496821771264041863915228768515743204386905918354
```

```
14418386150246458625744711564914544353588115604955636343894596239631
```

```
director@10.25.36.47-2.00e0.8105.d46b
```

```
...
```

```
ok
```

Using SGOS 4.x, complete the following steps:

1. From the (config) prompt, enter the services>ssh-console submode:

```
SGOS#(config) services
```

```
SGOS#(config services) ssh-console
```

```
SGOS#(config services ssh-console)
```

2. Import Director's key that was previously created on Director and copied to the clipboard.

Important: You must add the Director identification at the of the client key. The example shows the username, IP address, and MAC address of Director. "Director" (without quotes) must be the username, allowing you access to passwords in clear text.

```
SGOS#(config services ssh-console) import director-client-key
Paste client key here, end with "... " (three periods)
ssh-rsa
AAAAB3NzaC1yc2EAAAABIwAAAIEAvJIXt1ZauseE9qrcXem2IK/mC4dY8Cxxo1/B8th4KvedFY330ByO
/pvwucuhPZz+b1LETTY/zc3SL7jdVffq00KBN/ir4zu7L2XT68ML20Rwa9tXFedNmK1/iagI3/QZJ8T
8zQM6o7WnBzTvMC/ZE1MZZddAE3yPCv9+s2TR/Ipk=director@10.25.36.47-2.00e0.8105.d46b
...
ok
```

To View the Fingerprint of the Key:

```
SGOS#(config sshd) 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 sshd) delete director-client-key clientID
```

Importing VPM Policy

If you have your VPM policy stored locally and want to install it on a ProxySG, you can use SGOS inline commands to install them directly on the system. Note that VPM policy is stored in two files, `vpm-cpl` and `vpm-xml`. You must install both of them. (For more information on using VPM, see Chapter 14: “The Visual Policy Manager” on page 453.)

Note: VPM files are generally pulled from a specified ProxySG (reference device) and distributed to other ProxySG Appliances through the Director Management Console, and this is the recommended method. The procedure below is provided for reference.

For information on using VPM files with Director, refer to the *Blue Coat Director User Guide*.

Before you begin, copy the policy you are installing to the clipboard.

From the `(config)` prompt, enter the following commands:

```
SGOS#(config) inline policy vpm-cpl eof
<Proxy>
Deny url.domain="restricted"; Rule 1 eof
ok
```

where *eof* is the string you use to indicate to the system that you are beginning or ending. It can be any string of letters, but it should not be a string you type as part of the policy.

```
SGOS#(config) inline policy vpm-xml eof
<vpmap>
<conditionObjects>
destination-url name="URL1" port="-1" single="true" url="restricted" />
</conditionObjects>
<layers>
<layer layertype="com.bluecoat.sgos.vpm.WebAccessPolicyTable">
<name>Web Access Policy (1)</name>
<numRows>1</numRows>
<rowItem enabled="true" num="0">
<colItem col="0" value="1" />
<colItem col="1" name="Any" type="String" />
```

```
<colItem col="2" name="URL1" negate="false" type="Condition" />
<colItem col="3" name="Any" type="String" />
<colItem col="4" name="Deny" type="String" />
<colItem col="5" name="Any" type="String" />
<colItem col="6" name="" type="String" />
</rowItem>
</layer>
</layers>
</vpmapp>
eof
ok
```

Backing Up a ProxySG's SSL Settings

You can return to a previous ProxySG configuration by using a backup (a snapshot of the appliance at a point in time). Backups are either created explicitly by request or automatically prior to each profile push. They are stored on Director.

Backup configurations consist of specific configuration parameters related to a particular ProxySG. A backup saves all configuration settings. You can also back up SSL settings, including keyrings, CA-certificates, external-certificates, certificate signing-requests, certificate-lists, and cipher suites used.

When pushing a profile to another ProxySG, the SSL configuration needs to be pushed to the ProxySG before any other configuration is pushed to ensure that the required SSL keyrings are available for setting up the services.

Creating Profiles

A profile is a snapshot of a ProxySG configuration that can be used as a template to configure other ProxySGs.

When Director uses a profile, it takes the output of the ProxySG `show configuration` command from one system (*creating* the profile) and applies that configuration to the ProxySG appliances you specify (*pushing* the profile).

Note: Because the `show configuration` output is specific to one type of system, it is important to push profiles only to ProxySGs with similar platforms and versions.

It is also important to be sure that profile comes from a fully-authenticated ProxySG; that is, one that it is fully authenticated through SSH/RSA. Such systems are said to have *golden* profiles, profiles without invalid commands for other devices using the same SGOS version.

Keep in mind that if the configuration does not display in the `show configuration` output, it is not pushed to other systems as part of a profile. Specifically, keyrings configured with the `no-show` option are not part of the `show configuration` output. (Keyrings configured with the `show-director` option are part of the output only if Director issues the command using SSH-RSA.)

You can manipulate the `show configuration` output by

- creating configurations on different systems to provide profiles for different purposes.

- using the `restore-defaults keep-console` command: Restore to the factory defaults, restore to the factory defaults but keep the configured secure consoles

For more information on using the `restore-defaults` command, see "Restoring System Defaults" on page 787.

When a profile is created SSL configuration settings, such as self-signed certificates and certificate signing requests, are included if they were created with the non-interactive or inline form of the SSL commands. The interactive form of the SSL commands is never permitted in within a profile or overlay. (For information on using non-interactive SSL commands, see Chapter 7: "Using Secure Services" on page 203.)

For more information on using profiles, refer to Chapter 5, "Configuration Management," in the *Blue Coat Director User Guide*.

Creating Overlays

An overlay is one or more individual settings (such as time, SNMP, bandwidth gain, or SSL settings) that can be applied to one or a selected set of ProxySGs. An overlay is overlaid on a profile, changing specific settings created by the profile to fine-tune configuration specifics without having to create new profiles.

SSL configuration settings, such as self-signed certificates and certificate signing requests, must be created using the non-interactive commands to be used in Director overlays. To use SSL non-interactive settings:

- "To Create a Self-Signed Certificate Non-interactively Using Create Commands" on page 215
- "To Create a Signing Request Non-interactively Using Create Commands" on page 224
- "To Change the Cipher Suite of the SSL Client through the CLI" on page 233

SSL values can also be created using the SSL inline commands and can be used in overlays:

- "To Import a Keyring through the CLI" on page 210
- "To Import an External Certificate through the CLI Using Inline Commands" on page 220
- "To Import a CA Certificate through the CLI Using Inline Commands" on page 227

For more information on using overlays, refer to Chapter 5, "Configuration Management," in the *Blue Coat Director User Guide*.

Director Documentation

The following documentation is available:

- *Blue Coat Director Installation Manual*
- *Blue Coat Director User Guide*
- *Blue Coat Director Content Sync Module Guide*
- *Blue Coat Director Request Management Guide*

Blue Coat Director documentation can be found at <http://download.bluecoat.com/release/SGME/index.html>

Index

A

- accept-encoding request header modification, troubleshooting 184
- access control list
 - creating through CLI 249
 - creating through Management Console 248
 - restricting access with 247
- access lists
 - cache bypass
 - associating with service group 918
 - creating 917
 - creating 916
 - redirection
 - associating with service group 917
 - creating 917
 - syntax 916
- access logging
 - adding to log file 781
 - bandwidth management, setting 763
 - CLI default-logging command 757
 - commands
 - custom format 750
 - ELFF 750
 - continuous uploading 759
 - creating/editing log formats 747
 - custom
 - format, creating/editing 748
 - log formats 877
 - custom client
 - configuring through Management Console 773
 - editing through CLI 773
 - custom client port number 773
 - deleting log formats through CLI 750
 - disabling 756
 - ELFF
 - format, creating/editing 748
 - log formats 877
 - file compression, discussed 760
 - filename formats 882
 - FTP upload client
 - editing through CLI 767
 - editing through Management Console 766
 - port number 766
 - global settings 758
 - global settings commands 759
 - HTTP upload client
 - configuring through CLI 771
 - configuring through Management Console 769
 - port number 770
 - HTTPS upload client
 - CLI commands 771
 - ICAP 419
 - instant messaging format 747
 - log file
 - creating through CLI 752
 - creating through Management Console 751
 - deleting 755
 - edit commands 754
 - editing 752
 - log size, viewing statistics 854
 - log tail, viewing through Management Console 853
 - maximum log size, setting 758
 - NCSA/common format 747
 - NCSA/common log format
 - described 881
 - overriding 781
 - overview 743
 - P2P format 747
 - PASV, configuring for FTP client 767
 - policy, using with 781
 - protocol
 - disabling through CLI 758
 - disabling through Management Console 757
 - protocol association
 - configuring through CLI 757
 - configuring through Management Console 756
 - protocols, using with 755
 - remote max file size 752
 - resetting 781
 - scheduled uploading 759
 - show list of all logs 853
 - SQUID format 747
 - SQUID-compatible format 880
 - statistics
 - viewing 853
 - viewing through CLI 856
 - status statistics, viewing 855

- streaming format 748
- SurfControl client, editing through CLI 775
- SurfControl client, editing through Management Console 774
- tail options 853
- testing upload 780
- upload behavior 758
- upload client
 - configuring 759
 - configuring through CLI 763
 - configuring through Management Console 761
- upload client commands 763, 765
- upload client, configuring through CLI 765
- upload compression 762
- upload filename, configuring through Management Console 767
- upload schedule
 - configuring through CLI 778
 - configuring through Management Console 777
 - configuring, overview 776
- W3C format for Windows Media 615
- Websense client
 - editing through CLI 776
 - port number 776
- Websense client, editing through Management Console 775
- Windows Media 615
- access logs
 - digital signing
 - overview 761
 - verifying 765
- access restrictions
 - access control list for 247
 - configuring 247
- active client connections 823
- active content
 - and HTTPS tunneled connection 564
 - definition of 564
 - embed tags 566
 - JavaScript 565
 - object tags 566
 - script tags 564
 - stripping 564
 - types 564
 - types that can be removed or replaced 564
- active content, stripping 564
- Admin layer
 - actions 254
 - conditions 251
 - example 254
 - properties 253
- administration access policy, Visual Policy Manager reference 467
- administration authentication policy, Visual Policy Manager reference 467
- administrator
 - defining policies 250
 - read-only and read-write access 47
 - security levels 245
- advanced forwarding, *see* forwarding
- alternate hash table, creating 923
- alternate hashing
 - ProxySG example 932
 - router example 932
- AOL port service, creating 140
- ASX rewrite
 - command syntax 610
 - rules 609
 - setting up for Windows Media 608
- attack-detection
 - client
 - block-action, explained 97
 - connection-limit, explained 97
 - creating and editing 97
 - failure-limit, explained 97
 - global defaults 95
 - global defaults, changing 96
 - unblock-time, explained 97
 - warning-limit, explained 97
 - configuration, viewing 98
 - mode, entering 95
 - overview 95
 - server
 - add or remove server from group 99
 - configuration, viewing 99
 - configuring 99
 - creating 99
 - editing 99
 - hostname, explained 99
 - request-limit, explained 99
- authenticate.mode, NTLM, setting for 257
- Authenticate-401 130
- authentication
 - configuring transparent proxy authentication 258
 - definition 241
 - definition of 241
 - LDAP 279
 - policies 241, 271

- setting options for transparent proxy authentication through CLI 259
- setting options for transparent proxy authentication through Management Console 258
- authentication realm
 - definition of 271
 - in Visual Policy Manager 478
 - typical configuration 271
- authorization
 - definition of 241
 - LDAP 279
 - policies 205, 241
- automatic service information, enabling 942

B

- bandwidth gain
 - additional configurations affecting 175
 - byte-range support effects 175
 - revalidate pragma-no-cache effects 177
 - statistics 817
- bandwidth management
 - access logging, setting for 763
 - allocating bandwidth 377
 - allocation examples 388
 - class hierarchies 378
 - creating classes through CLI 383
 - creating classes through Management Console 382
 - definitions 375
 - deleting bandwidth-classes 386
 - editing classes through CLI 384
 - editing classes through Management Console 383
 - enabling or disabling through CLI 382
 - enabling or disabling through Management Console 381
 - flow classification 380
 - managing through policy 387
 - maximum bandwidth 378
 - minimum bandwidth 378
 - overview 376
 - policy examples 388
 - priority levels 378
 - viewing configurations through CLI 387
 - viewing configurations through Management Console 386
- bandwidth refresh, configuring 162
- base DN for a group in Visual Policy Manager 478
- BCAAA
 - COREid, using with 340

- event messages 870
- installation folder, selecting 863
- troubleshooting 869
- viewing the event log 868
- viewing the services 869
- WIDMS, configuring for 325
- blocking popup windows 561
- blocking Web content 635
- Blue Coat monitoring, enabling 958
- Blue Coat Web Filter
 - automatic download through CLI 650
 - automatic download through Management Console 649
 - configuring through CLI 647
 - configuring through Management Console 646
 - dynamic real-time rating, configuring 651
 - dynamic real-time rating, overview 645
 - selecting dynamic real-time rating settings 652
- bridging
 - about 75
 - configuring
 - failover 80
 - software bridge 78
 - failover 80
 - pass-through card 76
- bridging, transparent proxy, setting up 199
- browser
 - accessing the Management Console with 49
 - proxy, configuring for 151
 - setting for explicit proxies 151
 - supported 31
 - troubleshooting 791
 - viewing policy files with 449
- BWM, *see* bandwidth management
- bypass list
 - central 104
 - central, understanding 104
 - local
 - installing through CLI 108
 - installing through Management Console 105
 - understanding 100
 - policy based, understanding 108
- byte-range support
 - affecting bandwidth gain 175
 - configuring 176

C

- CAASNT, *see* BCAA
- CA-Certificates

- certificate signing request
 - creating through CLI 223
 - creating through Management Console 222
- error message 225
- lists
 - creating through CLI 229
 - creating through Management Console 228
- managing 225
- troubleshooting 225
- cache bypass list
 - associating with service group 918
 - creating 917
- cached objects by size 847
- CacheOS 4.x, logs, retrieving 859
- caching
 - clearing the system cache 790
 - efficiency statistics 843
 - freshness statistics 820
 - purging the DNS cache 790
 - restarting the ProxySG 785
- capturing packets, *see* packet capturing
- central bypass list 104
- Central policy file
 - automatic installation 447
 - configuring Management Console 447
 - configuring through CLI 447
 - e-mail notification
 - configuring through CLI 448
 - configuring through Management Console 448
 - managing 447
 - obtaining from Blue Coat Systems 444
 - update interval 448
 - configuring 448
 - updated, checking for 448
- Certificate Realm
 - authentication and authorization overview 313
 - configuring authentication and authorization 313
 - creating a realm through CLI 314
 - defining a realm through CLI 317
 - defining properties 315
 - defining realm server properties through Management Console 313
 - how it works 313
 - overview 313
 - policies, creating 318
- certificate realm
 - LDAP authorization, adding 315
 - local authorization, adding 315
 - requirements 313
 - results, viewing 317
- certificate signing request
 - creating 222
- certificates
 - chaining, about 221
 - challenge 214
 - commands
 - creating certificate 223
 - importing certificate 218
 - common name 214
 - country code 214
 - creating 212
 - CSA 204
 - commands 227
 - deleting through CLI 228
 - importing 225
 - deleting through CLI 221
 - encryption, using for 205
 - explained 204
 - external, using 205
 - importing existing 216
 - importing through CLI 218
 - importing through Management Console 218
 - self-signed 204
 - creating through CLI 214
 - creating through Management Console 213
 - troubleshooting 230
 - wildcard, using 205
 - X.509, using externally 205
- challenge type, explained 255
- cipher suites
 - changing through CLI 233
 - interactive mode, using 233
 - International Step-Up, working with 206
 - non-interactive mode, using 234
 - Server Gated Cryptography, working with 206
 - SGOS, supported by 205
- cipher suites shipped with ProxySG 205
- CLI
 - accessing 48
 - changing username and passwords in 52
- CLI configuration file, creating for ProxySG 926
- client map, *see* SSL client
- client-side bandwidth, enhancing 184
- commands
 - access-log
 - configuring upload client 763
 - configuring upload schedule 778
 - creating file 752

- customize format 750
- disabling upload 765
- editing custom-client 773
- editing FTP upload client 767
- editing HTTPS-client 771
- editing websense-client 776
- global settings 759
- acquire-UTC 64
- Blue Coat Web Filter content filtering
 - automatic download 650
 - configuring 647
- bypass-list local 108
- certificate realm
 - create-realm 317
 - edit-realm 317
- certificate, importing 218
- client-key, import 59
- DNS alternate 92
- DNS imputing 94
- DNS proxy service, creating and editing 132
- failover 116
- forwarding for explicit TCP-Tunnel 145
- FTP, create service 136
- host-keypair, create 56
- hostname 62
- HTTP-console, create port service 126
- HTTPS, configure origination 237
- HTTPS-console, create port service 124
- inline policy, empty 447
- interface, provide configuration instructions 152
- InterSafe content filtering, automatic download 657
- InterSafe content filtering, configuring 655
- keyring, create 210
- keyring, import 210
- keyring, view 211
- load balancing, configure multiple IP gateway 83
- local database content filtering, automatic downloads 645
- local database content filtering, configuring 642
- log file, edit 754
- NTP server, adding 65
- Optenet content filtering, automatic download 662
- Optenet content filtering, configuring 659
- policy poll-interval 448
- policy poll-now 448
- policy subscribe 447
- policy trace 442
- policy, e-mail notification 448
- policy, load 446
- Proventia content filtering, automatic download 667
- Proventia content filtering, configuring 664
- rip 90
- RIP settings, install 90
- security allowed-access 249
- security enable-password 53
- security enforce-ACL 249
- security front-panel-pin 53
- security hashed-enabled-password 53
- security hashed-front-panel 53
- security hashed-password 53
- security ldap edit-realm 283
- security ldap realm_name membership-attribute 290
- security ldap realm_name membership-type 290
- security local view 310
- security management display-realm 53
- security management login-timeout 54
- security NTLM create-realm 275
- security NTLM edit-realm 275
- security password 53
- security radius create-realm 299
- security radius edit-realm 299
- security sequence create-realm 360
- security sequence edit realm 360
- security siteminder create-realm 325
- security siteminder edit-realm 325
- security transparent-proxy-auth 259
- security username 53
- server certificate, verify 237
- services epmapper 134
- services http 137, 202
- services https 139
- services MMS 141
- services RTSP 141, 142
- services socks-proxy 191
- services, tcp-tunnel 145
- show configuration 450
- show policy 450
- show socks-proxy 190
- SmartFilter content filtering, automatic download 676
- SmartFilter content filtering, configuring 673
- socks 143
- socks-proxy 190
- ssh-console 127

- ssl ccl list-name 230
- ssl ccl list-name cert-name 230
- ssl ccl list-name view 230
- ssl client, create 232
- ssl create certificate 223
- ssl delete certificate 219
- ssl delete external certificate 221
- ssl import ca-certificate 227
- ssl no ca-certificate 228
- ssl set cipher-suite ssl-client 233
- ssl view summary ca-certificate 230
- static-routes, install 87
- SurfControl content filtering, automatic download 681
- SurfControl content filtering, configuring 678
- telnet-console 128
- view client-key 59
- view host-public-key 56
- virtual clear 113
- virtual IP address, creating 113
- virtual IP address, deleting 113
- WCCP, install 111
- Websense content filtering, automatic download 687
- Websense content filtering, configuring through CLI 684
- Websense integration service, configuring 686
- Webwasher content filtering, automatic download 692
- Webwasher content filtering, configuring 689
- common access log format 881
- community strings 805
- compression
 - behavior 179
 - boundary conditions 185
 - cache-control
 - no-transform directive ignored 185
 - cache-hit default behavior 180
 - cache-miss default behavior 180
 - client-side bandwidth settings 184
 - configuring 181
 - configuring through VPM 181
 - CPL, using with 183
 - CPU settings 184
 - exception pages issued 180
 - HTTP client compression object 181
 - HTTP server compression object 181
 - multiple content encoding 185
 - policy-based content transformation not stored 185
 - server-side bandwidth settings 184
 - variant served 185
- compression, overview 178
- configuration
 - archive running configuration 70
 - sharing between systems 66
- configuration file
 - creating with inline commands 926
 - creating with text editor 927
 - loading on ProxySG 927
- configuration mode, understanding 47
- CONNECT, using with origin-style redirection 258
- console account
 - minimum security 245
 - tab in Management Console 51
- console password, *see* password
- content filtering
 - blocking content 635
 - Blue Coat Web Filter, automatic download 649
 - Blue Coat Web Filter, automatic download through CLI 650
 - Blue Coat Web Filter, configuring through CLI 647
 - Blue Coat Web Filter, configuring through Management Console 646
 - definition of 635
 - example of category= 693
 - expired database, using 698
 - expired license, downloading a database with 698
 - InterSafe, automatic download through CLI 657
 - InterSafe, automatic download through Management Console 656
 - InterSafe, configuring through CLI 655
 - InterSafe, configuring through Management Console 653
 - local database, automatic download through CLI 645
 - local database, automatic download through Management Console 644
 - local database, configuring through CLI 642
 - local database, configuring through Management Console 640
 - Optenet, automatic download through CLI 662
 - Optenet, automatic download through Management Console 661
 - Optenet, configuring through CLI 659
 - Optenet, configuring through Management Console 658

- policy with vendor categories 695
- Proventia, automatic download through CLI 667
- Proventia, automatic download through Management Console 666
- Proventia, configuring through CLI 664
- Proventia, configuring through Management Console 663
- provider, selecting through CLI 638
- provider, selecting through Management Console 636
- SmartFilter, automatic download through CLI 676
- SmartFilter, automatic download through Management Console 675
- SmartFilter, configuring through CLI 673
- SmartFilter, configuring through Management Console 669
- SurfControl, automatic download through CLI 681
- SurfControl, automatic download through Management Console 680
- SurfControl, configuring through CLI 678
- SurfControl, configuring through Management Console 677
- Websense, automatic download through CLI 687
- Websense, automatic download through Management Console 687
- Websense, configuring through CLI 684
- Websense, configuring through Management Console 682
- Webwasher, automatic download through CLI 692
- Webwasher, automatic download through Management Console 691
- Webwasher, configuring through CLI 689
- Webwasher, configuring through Management Console 688
- content scanning
 - about 401
 - caching scanned objects 403
 - defined 23
 - ICAP service 403
 - policy for 402
- core image
 - restart options 957
- COREid
 - Access Server
 - specifying through CLI 343
 - specifying through Management Console 342
 - agents
 - configuring through CLI 341
 - agents, configuring 340
 - agents, configuring through Management Console 340
 - configuration overview 336
 - CPL, creating 346
 - forward proxy, using with 338
 - general settings
 - configuring 344
 - general settings, specifying through CLI 346
 - general settings, specifying through Management Console 345
 - ProxySG
 - challenges, avoiding 338
 - configuring 337
 - realm, creating 339
 - realm, creating through CLI 340
 - realm, creating through Management Console 339
 - SSO scheme, participating in 338
 - system, configuring 336
- CPL
 - Admin layer
 - actions 254
 - conditions 251
 - example 254
 - properties 253
 - Certificate Realm, policies, creating 318
 - creating through CLI 445
 - enabling ICP 736
 - generated by VPM 541
 - inline command 445
 - LDAP examples 294
 - local realm, creating policies 312
 - Netegrity SiteMinder policies, creating 335
 - NTLM policies, creating 278
 - policy overview 22
 - policy substitution realm, policies, creating 355
 - Proxy layer
 - actions 269
 - conditions 261
 - properties 267
 - RADIUS policies, creating 301
 - unloading policy files 446
- CPU
 - enhancing 184
 - utilization 816
- CPU monitoring
 - CLI commands 27
 - configuring 959

- overview 27
- CSR
 - interactive signing request mode, using 223
 - non-interactive create signing-request mode, using 224
- custom client
 - configuring for access logging 773
- custom format, creating/editing 748

D

- D range multicast address, explained 115
- data access pattern 845
- data allocation 842
- database
 - creating through ProxySG 308
 - local realm, setting up 305
 - viewing all users 309
- defaults, restoring system defaults 787
- deleting a ProxySG system 797
- deleting headers 567
- deleting objects from the ProxySG 808
- diagnostics
 - Blue Coat monitoring 958
 - core image restart options 957
 - CPU monitoring 959
 - heartbeats 958
 - packet capturing 951
 - sending service information 945
 - sending service information automatically 942
 - snapshot jobs 948
- digital signing
 - overview 761
 - verifying 765

Director

- client-key importing 965
- communicating with 964
- inline commands, using 966
- overlays, using with 968
- ProxySG, using with 963
- SSL non-interactive modes, using 214

disk

- multi-disk ProxySG 808
- reinitialization 807
- resource use 840
- single-disk ProxySG 808

DNS

- adding alternate server through CLI 92
- adding alternate server through Management Console 92

- adding primary through Management Console 91
- cache, purging 790
- negative caching, disabling 94
- negative caching, enabling 94
- to origin content server 239
- understanding 90

DNS servers

- addresses, specifying 90
- changing name imputing order 94
- changing order 93
- changing order of 92
- name imputing 93

DNS-Proxy

- commands 132
- configuring through CLI 132
- configuring through Management Console 131
- overview 131
- resolving name list, explained 131
- resource record, creating 132

document

- conventions 33
- organization 31

domain name for a group

- in Visual Policy Manager 479, 480

Do-Not-Fragment, *see* PMTU

dynamic bypass, troubleshooting 101

dynamic real-time rating

- categorize dynamically in real-time 652
- categorize dynamically in the background 652
- configuring 651
- do not categorize dynamically 652
- overview 645

E

ELFF

- access log formats 877
- creating/editing 748

embed tags 566

empty system 794

enable mode, understanding 47

Endpoint Mapper proxy

- commands 134
- configuring through CLI 134
- configuring through Management Console 133

error message, HTTPS Console 230

event log 848

event logging

- configuration, viewing through CLI 802
- contents, viewing through CLI 802

- event notification 800
 - log levels 798
 - log size 799
 - overview 797
 - event messages, BCAA 870
 - exceptions
 - built-in 568
 - compression 180
 - defining 568
 - definitions 572
 - hierarchy 573
 - installable list, about 574
 - installable list, install 577
 - user-defined 572
 - view 579
 - explicit proxy
 - browser settings 151
 - configuring 151
 - creating 151
 - definition 241
 - Internet Explorer, using with 186
 - overview 149
 - ProxySG, using as proxy server 151
 - explicit TCP-Tunnel, explained 144
 - Extended Log File Format, *see* ELFF 877
- F**
- failover
 - configuring through CLI 116
 - configuring through Management Console 114
 - configuring, overview 114
 - group secret 116
 - master 116
 - master, explained 114
 - multicast address, using 115
 - priority ranges 116
 - show failover configuration 117
 - statistics page, viewing 117, 857
 - statistics, viewing through CLI 117
 - VRRP, using with 114
 - filename formats, access logging 882
 - filter expressions for packet capturing 952
 - filtering, *see* content filtering
 - Finjan Vital Security scanning server 400
 - Firefox, versions supported 31
 - forms-based authentication
 - CPL substitutions for 363
 - CPL, using with 370
 - creating through CLI 367
 - creating, tips 363
 - creating/downloading through CLI 368
 - creating/editing form through Management Console 363
 - credentials sent in cleartext 371
 - customizing through ProxySG 367
 - editing through CLI 368
 - installing from local file 366
 - installing from remote URL 365
 - required values 362
 - storage options, setting through CLI 370
 - storage options, setting through Management Console 369
 - tips/boundary conditions 371
 - understanding 361
 - forward proxy, definition 241
 - forwarding
 - configuring through CLI 702
 - default sequence, creating 711
 - editing a host 705
 - host affinity, configuring 710
 - hosts/host groups, creating 702
 - load balancing, configuring 708
 - policy commands in forward layer 736
 - policy, managing with 736
 - using forwarding directives to create an installable list 713
 - fail open/closed 715
 - host timeout values 715
 - front panel PIN
 - clearing 243
 - creating 243
 - FTP
 - access logging, using with 756
 - content scanning 401
 - ProxySG, configuration for 931
 - router configuration for 930
 - spoofing 153
 - WCCP example 930
 - FTP clients, configuring 157
 - FTP port service
 - commands 136
 - creating 135
 - service defined 135
 - FTP proxy
 - configuring 152, 153
 - FTP upload client
 - editing through CLI 767
 - editing through Management Console 766

- troubleshooting 769

G

- gateways

- load balancing through CLI 83
- load balancing through Management Console 82
- switching to secondary 82
- understanding 81
- using multiple default IP gateways 81

- global configurations 61

- graph scale 811

H

- .htpasswd file

- creating password realm database 307
- loading 308
- uploading 308

- hash table, *see* redirection hash table

- hashed passwords, *see* passwords

- header

- policy substitution realm, using with 355

- headers

- request modification 402
- response modification 402

- health check

- creating forwarding 435
- creating general 431
- instant 435

- heartbeats, configuring 958

- home router

- mismatch errors 933
- ProxySG IP address 933
- troubleshooting 932
- version 1 usage 912
- version 2 configuration 913
- WCCP IP address 933

- hot spot, working with 923

- HTTP

- access logging, using with 755
- handoff, enabling 598
- persistent timeout, setting 66
- receive timeout, setting 66
- scanning HTTP objects 401
- timeout, configuring 66
- tolerant request parsing 167

- HTTP client compression object, using in VPM 181

- HTTP Console

- commands 126
- managing through CLI 126

- managing through Management Console 125

- HTTP port service

- commands 137, 202
- creating 136

- HTTP proxy

- acceleration profile 168
- bandwidth gain 175
- bandwidth gain profile 169
- byte-range support 175
- compression 178
- compression behavior 179
- compression boundary conditions 185
- compression, configuring 181
- normal profile 168
- portal profile 169
- profile settings, configuring 173
- profile settings, explained 170
- range request types 176
- revalidate pragma-no-cache 177
- traffic, controlling 168
- viewing settings 177

- HTTP redirection

- multicast address example 929
- multicast address router configuration 929
- password example 930
- ProxySG configuration 929
- ProxySG multicast address configuration 929
- ProxySG password example 930
- router configuration example 928
- router configuration for password 930

- HTTP server compression object, using in VPM 181

- HTTP upload client, configuring 769

- HTTPS

- content filtering, using with 698
- content scanning 401
- origination 236
- tunneled connection 564

- HTTPS Console

- certificate error message 230
- creating through CLI 123
- enabling 123
- IP address, selecting 123
- keyring, selecting 122
- managing through CLI 124
- managing through Management Console 123
- port service, creating 125
- troubleshooting certificate problems 230

- HTTPS port service

- commands 139

- creating 138, 236
- HTTPS termination
 - certificates 204
 - client map 206
 - configuring 207
 - DNS resolution to origin content server 239
 - keyring, creating 208
 - offloading SSL processing 203

I

ICAP

- access logging 419
- configuring the ProxySG for 403
- content scanning 401
- definition of 400
- Finjan Vital Security 400
- installing 403
- ISTags 400
- patience pages 405
- persistent connections 401
- refreshing content versions 403
- sense settings 400
- Symantec CarrierScan Server 400
- Trend Micro InterScan VirusWall 400
- WebWasher 400

ICMP broadcast echo

- configuring 118

ICMP error message

- ICMP host unreachable 119

ICMP timestamp echo

- configuring 119

ICP

- access logging, using with 756
- creating an installable list for 730
- enabling through CPL 736
- hierarchy 730
- icp_access_domain directive 733
- icp_access_ip directive 733
- installable list, creating through CLI 735
- installable list, creating through Management Console 734
- installing an ICP configuration 577
- restricting access 732

identification (Ident) protocol 729

imputing

- adding names through CLI 94
- adding names through Management Console 93
- changing name order 94
- changing suffix order 94

- definition of 93

- see also* DNS 90

- understanding 93

- inbound connections, rejecting 73

- inline commands

- creating policy with 442, 445

- Director, using 966

- using with forms-based authentication 368

- installable list

- ICP 730

- SOCKS 726

- instant messaging

- access log format 747

- access logging, using with 756

- AOL Messenger client configuration 632

- configuring clients 630

- configuring proxies 628

- creating 140

- defined 25

- MSN Messenger client configuration 632

- protocol policies 621

- proxy authentication 626

- securing 621

- statistics, IM clients tab 827

- statistics, IM data tab 826

- VPM 633

- Yahoo Messenger client configuration 631

- interface cards, configuring through CLI 152

- Internet Explorer, explicit proxy, using with 186

- Internet Explorer, versions supported 31

InterSafe

- automatic download through CLI 657

- automatic download through Management

- Console 656

- configuring through CLI 655

- configuring through Management Console 653

- IP address, configuring through CLI 73

- IP forwarding, enabling through Management Console 201

J

- JavaScript 565

- JREs

- supported 31

K

- keyring

- associating with certificate through CLI 218

- commands, create 210
- commands, import 210
- creating through CLI 210
- creating through Management Console 208
- importing through Management Console 216
- SSL client, associating 231
- view command 211

L

LDAP

- authentication and authorization overview 279
- authorization 288
- case-sensitive configuration 282
- certificate realm, adding to 315
- CPL examples 294
- defining Base DN through CLI 287
- defining Base DN through Management Console 286
- defining realm authorization properties and group information through Management Console 288, 290
- defining realm server properties through Management Console 280
- defining server properties through CLI 277, 282
- defining server properties through Management Console 281
- edit-realm commands 283
- group information 289
- membership-attribute command 290
- membership-type command 290
- policy-substitution realm, adding to 351
- search boundaries 289
- searching multiple base DN 285
- SSL, enabling 282
- v2/v3 support 279
- virtual URL, setting up 293

licensing

- about 35
- components 35
- expiration 37
- installing 40
- restore-default deletions 787
- trial period 37
- updating 44
- viewing 43

Lightweight Directory Access Protocol, *see* LDAP

link settings 74

load balancing

- assigning percentages 922

- gateways 82
 - understanding 922
 - using multiple default IP gateways 81
- local database
 - automatic download through CLI 645
 - automatic download through Management Console 644
 - clearing 643
 - configuring through CLI 642
 - configuring through Management Console 640
- local realm
 - authentication and authorization overview 302
 - certificate realm, adding to 315
 - changing properties 303
 - CPL, creating policies 312
 - database group, creating 308
 - database user, creating 308
 - database users, viewing 309
 - database, creating 306
 - database, creating through ProxySG 308
 - database, populated 306
 - database, setting up 305
 - defining realm server properties through Management Console 302
 - deleting groups 311
 - deleting users 310
 - groups, defined 307
 - groups, deleting 310
 - hashed passwords 307
 - policy substitution realm, adding to 351
 - results, viewing 305
 - user account, enabling 309
 - user name, defined 307
 - user password, creating 309
 - users, deleting 310
 - view all lists 310
 - virtual URL, setting up 304
- local user list
 - security settings, changing 311
- locking and unlocking ProxySG systems 796
- log file
 - creating 751
 - deleting 755
 - editing 752
- logging
 - event log 848
 - see* access logging and event logging
 - SNMP 804
 - syslog event monitoring 801

login parameters 49

logs

CacheOS 4.x, retrieving 859

SGOS 2.x, retrieving 859

M

management architecture, overview 22

Management Console

accessing 49

changing username and passwords in 51

configuring SSH 55

console account 51, 52

home page 49

HTTP Console 125

HTTPS Console 122

importing SSH client keypairs 57

logging in 49

logging out 50

managing 122, 129

SSH Console 126

Telnet Console 127

troubleshooting 791

menu bar in Visual Policy Manager 458

meta tags

parsing 167

MIBs 804

MMS

port service, creating 141

port services MMS commands 141, 142

modes, understanding 47

modifying headers 567

MSN port service, creating 140

multicast

D range address, explained 115

defined 584

failover, using with 115

unicast, converting by Windows Media 601

multicast address

configuring 916

ProxySG configuration 929

router configuration 929

syntax 916

multicast packet reception, enabling 919

N

name imputing, *see* imputing

name, configuring 61

NCSA, common access log format 747, 881

negate option, using in Visual Policy Manager 461

negative caching

disabling for DNS responses 94

enabling for DNS responses 94

netbios

using with policy substitution realm 354

Netegrity SiteMinder realm

agents, configuring 325

case-sensitive configuration 334

creating through CLI 325

creating through Management Console 324

defining server properties through CLI 332

display name, changing 334

making general settings through CLI 334

policies, creating 335

protected resource, entering 331

server mode, configuring 331

servers, configuring through Management Console 327

servers, editing through CLI 329

servers, editing through Management Console 328

SiteMinder agent, defining through CLI 326

SSO-only mode, enabling 331

viewing through CLI 330

Netscape, versions supported 31

network adapter

advanced configuration 73

configuring through CLI 73

link faults 75

link settings 74

rejecting inbound connections 73

Network Time Protocol server, *see* NTP

non-cacheable data 844

NTLM

authenticate.mode, setting 257

configuring authentication and authorization 273

creating a realm through CLI 275

defining realm server properties 273

defining realm server properties through Management Console 273

explicit proxy, using with Internet Explorer 186

force authentication

enabling through CPL 188

enabling through VPM 188

Internet Explorer, using with 186

overview 273

policies, creating 278

realm sequence position 360

single sign-on, configuring 278

- using with forms-based authentication 371

NTP

- adding server through CLI 65
- adding server through Management Console 65
- server order, changing 65
- time server, definition of 63
- understanding 64

O

- object tags 566
- objects
 - deleting from ProxySG 808
 - in Visual Policy Manager 463
 - served 821
 - served by size 847
- Oblix, *See* COREid
- one-time passwords, using with RADIUS 295
- Opnet content filtering
 - automatic download through CLI 662
 - automatic download through Management Console 661
 - configuring through CLI 659
 - configuring through Management Console 658
- optional negation syntax, using 924
- origination, HTTPS 236
- origin-style authentication
 - origin 255
 - origin-cookie 255
 - origin-cookie-redirect 255
 - origin-ip 255
 - origin-ip-redirect 255
- overlays
 - settings used in 968

P

P2P

- access log format 747
- access logging 582
- authentication 582
- managing 580
- policy 581

PAC file, defined 151

packet capturing

- about 951
- capturing 953
- common filter expressions 952
- file name format 951
- uploading data 956
- viewing current data 956

packet redirection

- enabling 918
- excluding 919

password

- changing through CLI 52
- changing through Management Console 51
- default for 51
- hashed, encrypted 244
- HTTP redirection example 930
- security, understanding 244
- see also* privileged-mode password
- with RIP 939

patience pages

- displaying 405
- troubleshooting 576

peer-to-peer

- access logging 582
- authentication 582
- managing 580
- policy 581

PMTU

- enabled by default 119
- overview 119

policy

- bandwidth management examples 388
- bypass list 108
- changing in Visual Policy Manager 544
- CLI inline command, using 445
- configuring policy evaluation order 441
- configuring the default policy proxy setting 441
- content scanning 402
- creating through CLI 445
- disabling 446
- disabling in Visual Policy Manager 544
- editing 442
- enabling in Visual Policy Manager 544
- example, limit access to certain web sites 695
- example, limit access to specified time of day 696
- files loading 442
- files, loading through CLI 446
- for maximum security 246
- for moderate security 246
- inline command 445
- inline commands, using 442
- layers in 542
- loading in Visual Policy Manager 543
- managing bandwidth 387
- overview 22
- policy editor 453

- saving in Visual Policy Manager 543
- source, viewing 450
- source, viewing through CLI 450
- statistics, viewing 451
- tabs for in Visual Policy Manager 459
- tracing 442
- tracing information 451
- unloading 446
- unloading/disabling files through CLI 447
- vendor categories, using with 695
- viewing through CLI 450
- viewing with browser 449
- Visual Policy Manager 453
- policy evaluation order
 - configuring through CLI 440
 - configuring through Management Console 440
- policy substitution realm
 - configuring 347
 - creating a realm through CLI 350
 - defining properties through CLI 351
 - defining properties through Management Console 350
 - defining realm server properties through Management Console 349
 - full usernames, constructing 351
 - general properties, defining through CLI 353
 - general properties, defining through Management Console 353
 - header, using with 355
 - how it works 347
 - LDAP authorization, adding 351
 - local authorization, adding 351
 - netbios, using with 354
 - policies, creating 355
 - results, viewing 354
 - user, username fields, explained 347
 - usernames, constructing 351
- pop-up ads, blocking 562
- port services
 - AOL, Yahoo, MSN, creating 140
 - attributes 130
 - attributes supported 130
 - creating/editing 129
 - FTP, creating 135
 - FTP, defined 135
 - HTTP, creating 136
 - HTTPS Console, creating 123, 125
 - HTTPS, creating 138, 236
 - instant messaging protocols 140
 - MMS port services commands 141, 142
 - MMS, creating 141
 - RTSP port services commands 141, 142
 - RTSP, creating 141
 - SOCKS, creating 142
 - SSH Console, creating 126
 - supported 129
 - TCP-Tunnel, creating 144
 - Telnet Console, creating 127
 - Telnet Console, explained 127
- privilege (enabled) mode, understanding 47
- privileged-mode password
 - changing through CLI 52
 - changing through Management Console 51
 - default for 51
- profiles
 - show configuration output, using 967
- prompt, customizing for Telnet 195
- Proventia
 - automatic download through CLI 667
 - automatic download through Management Console 666
 - configuring through CLI 664
 - configuring through Management Console 663
- proxies
 - configuring default settings through Management Console 441
 - definition 149, 241
 - explicit, browser settings 151
 - explicit, creating 151
 - interface settings 152
 - setting default HTTP proxy policy 164
 - setting up 149
 - SOCKS, configuring through CLI 190
 - SOCKS, configuring through Management Console 190
 - understanding 149
- Proxy layer
 - actions 269
 - conditions 261
 - properties 267
- proxy server, using ProxySG as 151
- ProxySG
 - accessing 48
 - alternate hashing example 932
 - browsers supported 31
 - configuration file, creating 923
 - configuration file, creating with text editor 927
 - configuration file, loading 927

- configuration file, quick start 914
- configuration file, syntax 924
- configuration file, using CLI 926
- deleting a system 797
- deleting objects from 808
- DNS server 90
- features 21
- FTP example 931
- home router IP address, verifying 933
- HTTP configuration example 929
- HTTP redirection multicast address example 929
- HTTP redirection with password example 930
- ICAP service configuration 403
- instant messaging, AOL Messenger client configuration 632
- instant messaging, configuring clients 630
- instant messaging, configuring proxies 628
- instant messaging, IM clients tab statistics 827
- instant messaging, IM data tab statistics 826
- instant messaging, MSN Messenger client configuration 632
- instant messaging, protocol policies 621
- instant messaging, proxy authentication 626
- instant messaging, securing 621
- instant messaging, VPM 633
- instant messaging, Yahoo Messenger client configuration 631
- IP address for 73
- load balancing 922
- locking and unlocking a system 796
- managing 794
- multi-disk 808
- optional negation syntax, using 924
- protocols supported 30
- read-only and read-write access 47, 245
- realm name, changing through CLI 53
- realm name, changing through Management Console 53, 54
- replacing a system 794, 797
- restarting 785
- reverse proxy example 931
- serial number, configuring 62
- setting the default system to boot 796
- simultaneous connections to, viewing 98
- single-disk 808
- subnet mask for 72
- system defaults 787
- time, configuring 63
- timeout, changing through CLI 54

- upgrading 791
- viewing details 795
- WCCP configuration, creating 921
- WCCP versions supported 911
- WCCP-known caches, displaying 928
- proxy-support header
 - disabling through CPL 187
 - disabling through VPM 186
 - Internet Explorer, using with 186
- purging the DNS cache 790

Q

- quick start
 - ProxySG, creating a configuration file 914
 - WCCP configuration 913
- QuickTime
 - access logging, using with 755

R

- RADIUS
 - authentication and authorization overview 295
 - case-sensitive usernames, setting 299
 - creating a realm through CLI 299
 - defining realm server properties through Management Console 295, 296
 - one-time passwords, using with 295
 - policies, creating 301
- range request types 176
- read-only access in ProxySG 47, 245
- read-write access in ProxySG 47, 245
- realm
 - COREid, creating 339
 - name, changing 53
 - timeout, changing 54
- realm banner, Telnet, customizing for 195
- realm sequence
 - creating 357
 - creating through CLI 359
 - managing through CLI 360
 - NLM realm position 360
 - promote/demote member realms 358
 - results, viewing through CLI 360
 - virtual URL 360
- RealMedia
 - access logging, using with 755
 - proxy authentication 590
- realms
 - COREid 336
 - COREid, access server 342

- COREid, agents
 - configuring 340
 - COREid, configuring general settings 344
 - COREid, CPL, crating 346
 - definition 241
 - promote/demote for sequence realms 358
 - sequence, troubleshooting 356
 - understanding 271
 - rebooting, *see* restarting 785
 - redirection access list, creating 917
 - redirection hash table
 - alternate, creating 923
 - assigning percentages 922
 - hot spot 923
 - understanding 922
 - refresh bandwidth statistics 820
 - refresh bandwidth, configuring 162
 - replacing a ProxySG system 797
 - reporting
 - event logging 797
 - syslog event monitoring 801
 - request modification 402
 - resolving name list, explained 131
 - resource use
 - disk 840
 - memory 840
 - response modification 402
 - restart
 - core image 957
 - restarting the ProxySG
 - restart options 785
 - setting the default system to boot 796
 - restoring system defaults 787
 - restricting access 247
 - revalidate pragma-no-cache
 - affects on bandwidth gain 177
 - configuring 177
 - reverse proxy
 - definition 241
 - ProxySG, configuration for 931
 - router configuration for 931
 - WCCP example 931
 - RFC-1323
 - configuring 118
 - RIP
 - configuring 87
 - definition of 87
 - installing configuration file through CLI 90
 - installing configuration file through Management Console 87
 - parameters 938
 - ProxySG-specific RIP parameters 939
 - using passwords with 939
 - routing
 - bypass list 100
 - central bypass list 104
 - policy-based bypass list 108
 - static routes 83
 - routing information protocol, *see* RIP
 - RTSP
 - port service creating 141
 - port services commands 141, 142
 - rules in policies
 - deleting in Visual Policy Manager 545
 - in Visual Policy Manager user interface 460
 - option menus for in Visual Policy Manager 460
 - ordering in Visual Policy Manager 541
- S**
- script tags 564
 - security
 - console account 245
 - local user list settings, changing 311
 - policies for 246, 250
 - self-signed certificate
 - interactive mode, using 214
 - non-interactive create mode, using 215
 - sequence realm
 - creating a realm through CLI 357
 - defining realm server properties through Management Console 356
 - sequences, troubleshooting 356
 - serial console, definition 242
 - serial number, configuring 62
 - serial port
 - password, creating 244
 - server-side bandwidth, enhancing 184
 - service information
 - enabling automatic 942
 - sending 945
 - set_aut.pl script, using with .htpasswd file 308
 - setup console
 - password, creating 244
 - SGOS 2.x, logs, retrieving 859
 - shell proxies
 - \$substitutions, using 193
 - boundary conditions for 194
 - CLI commands, using 197

- policy settings, customizing 193
- Telnet 195
- understanding 193
- show configuration
 - Director, using with 967
- Simple Network Management Protocol, *see* SNMP
- SiteMinder, *see* Netegrity SiteMinder
- SmartFilter
 - automatic download through CLI 676
 - automatic download through Management Console 675
 - configuring through CLI 673
 - configuring through Management Console 669
- snapshot jobs
 - creating and editing 948
- SNMP
 - community strings 805
 - enabling 804
 - MIB variables 804
 - MIBs 804
 - traps 806
- SOCKS
 - access logging, using with 756
 - commands 143
 - compression gain statistics 838
 - connections, viewing 837
 - creating an installable list for 726
 - enabling 191
 - gateway configuration 722
 - port service, creating 142
 - port services commands 191
 - SOCKS clients, viewing 837
 - statistics 837
- SOCKS gateway
 - default sequence, creating 725
 - HTTP, using with SOCKS 729
- SOCKS proxy
 - bind timeout on accept value 190
 - CLI commands 190
 - configuring through CLI 190
 - configuring through Management Console 190
 - connection timeout values 190
 - max-connection values 190
 - max-idle-timeout value 190
 - min-idle-timeout 190
 - show socks-proxy 190
- SQUID access log format 747, 880
- SSH
 - client, managing 56
 - configuring through Management Console 55
 - host connection, configuring 55
 - host keypairs, configuring through CLI 56
 - importing client keypairs through Management Console 57
 - password authentication 246
 - setting up 55
 - view client-key 59
 - view host-public-key 56
- SSH Console
 - port service commands 127
 - port services, creating 126
- SSH with RSA authentication, not controlled by policy 250
- SSL
 - authentication/authorization services, using with 260
 - caching behavior, SSL client 231
 - cipher suites interactive mode, using 233
 - cipher suites non-interactive mode, using 234
 - CSR interactive signing request mode, using 223
 - CSR non-interactive create signing-request mode, using 224
 - definition 241
 - interactive versus non-interactive modes 214
 - LDAP, enabling 282
 - no-show keyring option 209
 - self-signed certificate interactive mode, using 214
 - self-signed certificate non-interactive create mode, using 215
 - settings, backing up through Director 967
 - show keyring option 209, 218
 - show-director option 209, 218
 - timeout, configuring 236
- SSL certificates, *see* certificates.
- SSL client
 - cipher suite, changing 233
 - CLI commands 232
 - explained 206
 - keyring, associating 231
 - managing 231
- static routes 83
 - explained 83
 - loading 87
 - table, installing through CLI 87
 - table, installing through Management Console 84
- statistics
 - access logging log size 854

- access logging, status 855
 - access logging, viewing through CLI 856
 - access logging, viewing through Management Console 853
 - active client connections 823
 - bandwidth gain 817
 - cache efficiency 843
 - cache freshness 820
 - cached objects by size 847
 - CPU utilization 816
 - data access pattern 845
 - data allocation 842
 - event log 848
 - failover page 117
 - failover page, viewing 857
 - graph scale 811
 - HTTP/FTP bytes served 822
 - non-cacheable data 844
 - objects served 821
 - objects served by size 847
 - policy 451
 - refresh bandwidth 820
 - resource use 840
 - show list of all logs 853
 - SOCKS clients, viewing 837
 - system summary 811
 - total bytes served 846
 - streaming media
 - access log format 748
 - delivery type 584
 - live content defined 585
 - multicast defined 584
 - prepopulating content, description 592
 - unicast defined 584
 - streaming protocols, managing 141
 - stripping active content 564
 - subnet mask, configuring with the Management Console 72
 - SurfControl
 - automatic download through CLI 681
 - automatic download through Management Console 680
 - configuring for access logging 774
 - configuring through CLI 678
 - configuring through Management Console 677
 - Reporter, using with SGOS4.x 679
 - surrogate credentials, defined 255
 - Symantec CarrierScan server 400
 - syslog event monitoring 801
 - system cache
 - clearing 790
 - troubleshooting 791
 - system defaults, restoring 787
 - system summary 811
 - system time, *see* time 63
- ## T
- TCP NewReno
 - configuring 118
 - TCP-IP
 - configuration, showing 120
 - ICMP broadcast echo 118
 - ICMP timestamp echo 119
 - overview 117
 - PMTU, configuring 119
 - RFC-1323 118
 - TCP NewReno 118
 - TCP-Tunnel
 - access logging, using with 756
 - commands 145
 - commands, explicit 145
 - explicit 144
 - overview 143
 - port services, creating 144
 - Telnet
 - banner settings, configuring through CLI 197
 - banner settings, configuring through Management Console 195
 - boundary conditions for Telnet shell proxy 198
 - settings customizing 195
 - shell proxy, creating service 195
 - shell proxy, understanding 195
 - Telnet Console
 - commands 128
 - error message 127
 - port service, creating 127
 - port service, explained 127
 - troubleshooting 127
 - time, configuring in the ProxySG 63
 - timeout
 - configuring for SSL termination 236
 - HTTP, configuring 66
 - timeout, realm, changing 54
 - tolerant request parsing, setting through CLI 167
 - transforming active content tags 564
 - transparent proxy
 - CLI commands 259
 - definition 241

- hardware, configuring 199
- IP forwarding 201
- IP forwarding, enabling through CLI 201
- Layer-4 switch, using with 200
- overview 150
- pass-through card, setting up 199
- service, creating 201
- software bridging, setting up 199
- transparent proxy authentication
 - configuring 258
 - setting options for through CLI 259
 - setting options for through Management Console 258
- transparent redirection, using WCCP 911
- traps 806
- Trend Micro InterScan VirusWall 400
- troubleshooting
 - accept-encoding request header modification 184
 - BCAAA service 869
 - browsers 31, 791
 - CA-Certificates 225
 - cache-control no-transform directive ignored. 185
 - compression choices 185
 - CONNECT method 258
 - explicit proxy and Internet Explorer 186
 - forms-based authentication 371
 - FTP upload client, upload-now command 769
 - gzip, deflate formats with compression 185
 - HTTPS and content filtering 698
 - HTTPS Console 230
 - ICMP host unreachable error message 119
 - JREs 31
 - licenses disappear after restore-defaults command 787
 - multiple content encoding 185
 - patience pages 576
 - policy-based content transformations 185
 - show list of all logs 853
 - TCP_DENIED 256
 - Telnet Console 127
 - virtual IPs 113
 - WCCP, home router mismatch 935
 - XFTP users not prompted for proxy authentication 374

U

- unicast
 - defined 584
 - handling video on demand 585

- multicast, converting from by Windows Media 601
- Universal Time Coordinates, *see* UTC
- UNIX
 - creating a realm through CLI 303
 - creating a realm through Management Console 303
- upgrade enhancements 31
- upgrading
 - overview 791
 - system image from PC 792
 - through CLI 793
 - through Management Console 792
- upload client
 - configuring through Management Console 761
- username
 - changing through CLI 52
 - changing through Management Console 51
 - default for 51
- UTC time 63

V

- viewing changes 920
- virtual IPs
 - creating through CLI 113
 - creating through Management Console 112
 - flags 117
 - show virtual 113
 - understanding 112
- virtual URL
 - LDAP set up 293
 - realm sequence 360
- virus scanning
 - advanced configurations 418
 - managing 418
 - policies for in Visual Policy Manager 472
 - replacing the ICAP server 419
- virus, preventing 95
- Visual Policy Manager
 - administration access policy reference 467
 - administration authentication policy reference 467
 - changing policies 544
 - command reference 458
 - deleting a policy 545
 - Director, using inline commands 966
 - disabling a policy 544
 - downloading files for 545
 - enabling a policy 544

- files for 545
- generated CPL 541
- loading policies 543
- menu bar 458
- object naming 29
- objects 463
- overview 453
- policy layers 542
- policy-layer tabs 459
- rule options in the user interface 460
- rule order in 541
- rules in the user interface 460
- saving policies 543
- web access policy example 554
- web access policy reference 468, 469
- web authentication policy example 549
- web content policy reference 472

VRRP, failover, using with 114

W

W3C Extended Log File Format, *see* ELFF 877

WCCP

- access lists, creating 916
- alternate hash table, using 923
- alternate hashing example 932
- changes, viewing 920
- definition of 911
- examples 928
- global settings, syntax 915
- global settings, using 915
- home router mismatch, troubleshooting 935
- home router troubleshooting 932
- hot spot, working with 923
- HTTP redirection example 928
- installing settings through CLI 111
- installing settings through Management Console 109
- interface commands, syntax 918
- interface commands, using 918
- known caches, displaying 928
- load balancing, understanding 922
- multicast address, configuring 916
- multicast packet reception, enabling 919
- optional negation syntax, explained 924
- overview 911
- packet redirection, enabling 918
- packet redirection, excluding 919
- ProxySG configuration for 921
- quick start 913

- router configuration, initial 914
- saving changes 920
- service group, naming 916
- service group, setting up 915
- settings 108
- settings, understanding 108
- transparent redirection, using with 911
- version 1 overview 911
- version 1 rules 912
- version 2 overview 912
- version 2 router, configuring 914
- version 2, enabling 916

web access policy

- example in Visual Policy Manager 554
- Visual Policy Manager reference 468, 469

web access, content filtering 635

web authentication policy

- example in Visual Policy Manager 549

Web Cache Control Protocol, *see* WCCP

web interface, definition of 49

Websense

- automatic download through CLI 687
- automatic download through Management Console 687
- configuring through CLI 684
- configuring through Management Console 682
- integration service
 - configuring through CLI 686
- upload client, editing through CLI 776
- upload client, editing through Management Console 775

Webwasher content filtering

- automatic download through CLI 692
- automatic download through Management Console 691
- configuring through CLI 689
- configuring through Management Console 688

WebWasher scanning server 400

welcome banner, Telnet, customizing for 195

wildcard certificates, using 205

Windows Media

- .ASX-rewrite rules 609
- .nsc file 604
- access logging format 615
- access logging, using with 755
- ASX rewrite and NTLM incompatibility 611
- authentication limitations 589
- HTTP handoff enabling 598
- multicast station monitoring 605

multicast to unicast 601
Player 6.4 compatibility 608
prepopulating content description 592
setting up ASX rewrite 608

X

X.509 certificates
using for encryption 205

using with digital signing 761
wildcards, using with 205
XFTP users, not prompted for proxy authentication
374
XML validation 543

Y

Yahoo port service, creating 140