

# Zorp GPL 7 Reference Guide

Publication date March 04, 2024

## Abstract

This document is a detailed reference guide for Zorp GPL administrators.





Balasys

Copyright © 1996-2024 Balasys IT Zrt. (Private Limited Company)

This documentation and the product it describes are considered protected by copyright according to the applicable laws.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>). This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)

Linux™ is a registered trademark of Linus Torvalds.

Windows™ 10 is registered trademarks of Microsoft Corporation.

The Balasys™ name and the Balasys™ logo are registered trademarks of Balasys IT Zrt.

The Zorp™ name and the Zorp™ logo are registered trademarks of Balasys IT Zrt.

All other product names mentioned herein are the trademarks of their respective owners.

#### DISCLAIMER

Balasys is not responsible for any third-party websites mentioned in this document. Balasys does not endorse and is not responsible or liable for any content, advertising, products, or other material on or available from such sites or resources. Balasys will not be responsible or liable for any damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through any such sites or resources.



# Table of Contents

<b>Preface</b> .....	<b>xiii</b>
1. Summary of contents .....	xiii
2. Terminology .....	xiii
3. Target audience and prerequisites .....	xiv
4. Products covered in this guide .....	xv
5. Contact and support information .....	xv
5.1. Sales contact .....	xv
5.2. Support contact .....	xv
5.3. Training .....	xvi
6. About this document .....	xvi
6.1. Feedback .....	xvi
<b>1. How Zorp works</b> .....	<b>1</b>
1.1. Zorp startup and initialization .....	1
1.2. Handling incoming connections .....	1
1.2.1. Handling packet filtering services .....	2
1.2.2. Handling application-level services .....	2
1.3. Proxy startup and the server-side connection .....	3
<b>2. Configuring Zorp proxies</b> .....	<b>4</b>
2.1. Policies for requests and responses .....	4
2.1.1. Default actions .....	5
2.1.2. Response codes .....	6
2.2. Secondary sessions .....	7
2.3. Embedded protocol analysis .....	7
2.3.1. Proxy stacking .....	7
2.3.2. Program stacking .....	8
<b>3. The Zorp SSL framework</b> .....	<b>9</b>
3.1. The SSL protocol .....	9
3.1.1. The SSL handshake .....	9
3.2. Configuring TLS and SSL encrypted connections .....	10
3.2.1. Behavior of the SSL framework .....	10
3.2.2. Handshake callbacks .....	11
3.2.3. X.509 Certificates .....	12
3.2.4. Setting the allowed TLS protocol .....	14
3.2.5. SSL cipher selection .....	14
3.2.6. Enabling STARTTLS .....	15
3.2.7. Keybridging certificates .....	16
3.3. Related standards .....	18
3.4. SSL options reference .....	18
<b>4. Proxies</b> .....	<b>19</b>
4.1. General information on the proxy modules .....	19
4.2. Attribute values .....	19
4.3. Examples .....	20
4.4. Module AnyPy .....	20
4.4.1. Related standards .....	20
4.4.2. Classes in the AnyPy module .....	20



4.4.3. Class AbstractAnyPyProxy .....	20
4.4.4. Class AnyPyProxy .....	21
4.5. Module Finger .....	21
4.5.1. The Finger protocol .....	22
4.5.2. Proxy behavior .....	22
4.5.3. Related standards .....	22
4.5.4. Classes in the Finger module .....	23
4.5.5. Class AbstractFingerProxy .....	23
4.5.6. Class FingerProxy .....	25
4.6. Module Ftp .....	25
4.6.1. The FTP protocol .....	25
4.6.2. Proxy behavior .....	26
4.6.3. Related standards .....	29
4.6.4. Classes in the Ftp module .....	29
4.6.5. Class AbstractFtpProxy .....	30
4.6.6. Class FtpProxy .....	35
4.6.7. Class FtpProxyAnonRO .....	35
4.6.8. Class FtpProxyAnonRW .....	35
4.6.9. Class FtpProxyRO .....	36
4.6.10. Class FtpProxyRW .....	36
4.7. Module Http .....	36
4.7.1. The HTTP protocol .....	36
4.7.2. Proxy behavior .....	37
4.7.3. Related standards .....	48
4.7.4. Classes in the Http module .....	48
4.7.5. Class AbstractHttpProxy .....	49
4.7.6. Class HttpProxy .....	61
4.7.7. Class HttpProxyNonTransparent .....	61
4.7.8. Class HttpProxyURIFilter .....	61
4.7.9. Class HttpProxyURIFilterNonTransparent .....	61
4.7.10. Class HttpProxyURLCategoryFilter .....	61
4.7.11. Class HttpWebdavProxy .....	62
4.7.12. Class NontransHttpWebdavProxy .....	62
4.8. Module Plug .....	62
4.8.1. Proxy behavior .....	62
4.8.2. Related standards .....	63
4.8.3. Classes in the Plug module .....	63
4.8.4. Class AbstractPlugProxy .....	63
4.8.5. Class PlugProxy .....	65
4.9. Module Pop3 .....	65
4.9.1. The POP3 protocol .....	66
4.9.2. Proxy behavior .....	67
4.9.3. Related standards .....	70
4.9.4. Classes in the Pop3 module .....	70
4.9.5. Class AbstractPop3Proxy .....	70
4.9.6. Class Pop3Proxy .....	73
4.9.7. Class Pop3STLSProxy .....	73
4.10. Module Smtplib .....	73



4.10.1. The SMTP protocol	73
4.10.2. Proxy behavior	74
4.10.3. Related standards	76
4.10.4. Classes in the Smtplib module	76
4.10.5. Class AbstractSmtpProxy	76
4.10.6. Class SmtpProxy	79
4.11. Module Telnet	81
4.11.1. The Telnet protocol	81
4.11.2. Proxy behavior	82
4.11.3. Related standards	84
4.11.4. Classes in the Telnet module	84
4.11.5. Class AbstractTelnetProxy	84
4.11.6. Class TelnetProxy	85
4.11.7. Class TelnetProxyStrict	85
4.12. Module Whois	86
4.12.1. The Whois protocol	86
4.12.2. Proxy behavior	86
4.12.3. Related standards	86
4.12.4. Classes in the Whois module	86
4.12.5. Class AbstractWhoisProxy	86
4.12.6. Class WhoisProxy	87
<b>5. Core</b>	<b>88</b>
5.1. Module Auth	88
5.1.1. Authentication and authorization basics	88
5.1.2. Authentication and authorization in Zorp	88
5.1.3. Classes in the Auth module	89
5.1.4. Class AbstractAuthentication	90
5.1.5. Class AbstractAuthorization	91
5.1.6. Class AuthCache	91
5.1.7. Class AuthenticationPolicy	92
5.1.8. Class AuthorizationPolicy	93
5.1.9. Class BasicAccessList	94
5.1.10. Class InbandAuthentication	95
5.1.11. Class NEyesAuthorization	95
5.1.12. Class PairAuthorization	97
5.1.13. Class PermitGroup	97
5.1.14. Class PermitTime	98
5.1.15. Class PermitUser	99
5.1.16. Class SatyrAuthentication	99
5.1.17. Class ServerAuthentication	99
5.1.18. Class ZAAAuthentication	100
5.2. Module AuthDB	101
5.2.1. Classes in the AuthDB module	101
5.2.2. Class AbstractAuthenticationBackend	101
5.2.3. Class AuthenticationProvider	102
5.2.4. Class ZAS2AuthenticationBackend	102
5.3. Module Chainer	103
5.3.1. Selecting the network protocol	104



5.3.2. Classes in the Chainer module .....	104
5.3.3. Class AbstractChainer .....	105
5.3.4. Class AvailabilityChainer .....	105
5.3.5. Class ConnectChainer .....	106
5.3.6. Class FailoverChainer .....	107
5.3.7. Class MultiTargetChainer .....	108
5.3.8. Class RoundRobinAvailabilityChainer .....	109
5.3.9. Class RoundRobinChainer .....	110
5.3.10. Class SideStackChainer .....	110
5.3.11. Class StateBasedChainer .....	111
5.4. Module Detector .....	112
5.4.1. Classes in the Detector module .....	112
5.4.2. Class AbstractDetector .....	113
5.4.3. Class CertDetector .....	113
5.4.4. Class DetectorPolicy .....	114
5.4.5. Class HttpDetector .....	114
5.4.6. Class SniDetector .....	115
5.4.7. Class SshDetector .....	116
5.5. Module Encryption .....	116
5.5.1. SSL parameter constants .....	116
5.5.2. Classes in the Encryption module .....	118
5.5.3. Class AbstractVerifier .....	120
5.5.4. Class Certificate .....	122
5.5.5. Class CertificateCA .....	124
5.5.6. Class ClientCertificateVerifier .....	125
5.5.7. Class ClientNoneVerifier .....	128
5.5.8. Class ClientOnlyEncryption .....	128
5.5.9. Class ClientOnlyStartTLSEncryption .....	129
5.5.10. Class ClientSSLOptions .....	131
5.5.11. Class DHParam .....	135
5.5.12. Class DynamicCertificate .....	136
5.5.13. Class DynamicServerEncryption .....	137
5.5.14. Class EncryptionPolicy .....	139
5.5.15. Class FakeStartTLSEncryption .....	140
5.5.16. Class ForwardStartTLSEncryption .....	142
5.5.17. Class PrivateKey .....	144
5.5.18. Class SNIBasedCertificate .....	146
5.5.19. Class SSLOptions .....	147
5.5.20. Class ServerCertificateVerifier .....	151
5.5.21. Class ServerNoneVerifier .....	153
5.5.22. Class ServerOnlyEncryption .....	153
5.5.23. Class ServerSSLOptions .....	155
5.5.24. Class StaticCertificate .....	158
5.5.25. Class TwoSidedEncryption .....	159
5.6. Module Keybridge .....	161
5.6.1. Classes in the Keybridge module .....	161
5.6.2. Class X509KeyBridge .....	161
5.7. Module Matcher .....	163



5.7.1. Classes in the Matcher module .....	163
5.7.2. Class AbstractMatcher .....	164
5.7.3. Class CombineMatcher .....	164
5.7.4. Class DNSMatcher .....	165
5.7.5. Class MatcherPolicy .....	166
5.7.6. Class RegexpFileMatcher .....	166
5.7.7. Class RegexpMatcher .....	167
5.7.8. Class SmtplibInvalidRecipientMatcher .....	168
5.7.9. Class WindowsUpdateMatcher .....	170
5.8. Module NAT .....	170
5.8.1. Classes in the NAT module .....	170
5.8.2. Class AbstractNAT .....	171
5.8.3. Class GeneralNAT .....	172
5.8.4. Class HashNAT .....	173
5.8.5. Class NAT46 .....	174
5.8.6. Class NAT64 .....	174
5.8.7. Class NATPolicy .....	175
5.8.8. Class OneToOneMultiNAT .....	176
5.8.9. Class OneToOneNAT .....	177
5.8.10. Class RandomNAT .....	178
5.8.11. Class StaticNAT .....	178
5.9. Module Notification .....	179
5.9.1. Classes in the Notification module .....	179
5.9.2. Class AbstractNotificationMethod .....	179
5.9.3. Class EmailNotificationMethod .....	179
5.9.4. Class NotificationPolicy .....	180
5.10. Module Proxy .....	180
5.10.1. Functions in module Proxy .....	180
5.10.2. Classes in the Proxy module .....	180
5.10.3. Functions .....	180
5.10.4. Class Proxy .....	181
5.11. Module Resolver .....	184
5.11.1. Classes in the Resolver module .....	184
5.11.2. Class AbstractResolver .....	184
5.11.3. Class DNSResolver .....	184
5.11.4. Class HashResolver .....	185
5.12. Module Router .....	186
5.12.1. The source address used in the server-side connection .....	186
5.12.2. Classes in the Router module .....	187
5.12.3. Class AbstractRouter .....	187
5.12.4. Class DirectedRouter .....	188
5.12.5. Class InbandRouter .....	189
5.12.6. Class TransparentRouter .....	190
5.13. Module Rule .....	191
5.13.1. Evaluating firewall rules .....	191
5.13.2. Sample rules .....	193
5.13.3. Adding metadata to rules: tags and description .....	193
5.13.4. Classes in the Rule module .....	194



5.13.5. Class PortRange .....	194
5.13.6. Class Rule .....	195
5.14. Module Service .....	197
5.14.1. Naming services .....	197
5.14.2. Classes in the Service module .....	198
5.14.3. Class AbstractService .....	198
5.14.4. Class DenyService .....	199
5.14.5. Class PFService .....	200
5.14.6. Class Service .....	202
5.15. Module Session .....	207
5.15.1. Classes in the Session module .....	208
5.15.2. Class StackedSession .....	208
5.16. Module SockAddr .....	209
5.16.1. Classes in the SockAddr module .....	210
5.16.2. Class SockAddrInet .....	210
5.16.3. Class SockAddrInet6 .....	211
5.16.4. Class SockAddrInetHostname .....	211
5.16.5. Class SockAddrInetRange .....	212
5.16.6. Class SockAddrUnix .....	213
5.17. Module Stack .....	213
5.17.1. Classes in the Stack module .....	214
5.17.2. Class AbstractStackingBackend .....	214
5.17.3. Class RemoteStackingBackend .....	214
5.17.4. Class StackingProvider .....	215
5.18. Module Zone .....	216
5.18.1. Classes in the Zone module .....	217
5.18.2. Class Zone .....	217
5.19. Module Zorp .....	218
<b>6. Core-internal .....</b>	<b>219</b>
6.1. Module Cache .....	219
6.2. Module Core .....	219
6.3. Module Dispatch .....	219
6.3.1. Zone-based service selection .....	219
6.3.2. Classes in the Dispatch module .....	220
6.3.3. Class CSZoneDispatcher .....	220
6.3.4. Class Dispatcher .....	221
6.4. Module Globals .....	223
6.5. Module Stream .....	223
6.5.1. Classes in the Stream module .....	223
6.5.2. Class Stream .....	223
<b>Appendix A. Additional proxy information .....</b>	<b>225</b>
A.1. TELNET appendix .....	225
<b>Appendix B. Global options of Zorp .....</b>	<b>232</b>
B.1. Setting global options of Zorp .....	232
blob .....	233
audit .....	234
options .....	236
<b>Appendix C. Zorp manual pages .....</b>	<b>237</b>





instances.conf .....	238
policy.py .....	240
zorp .....	241
zorpctl .....	243
zorpctl.conf .....	245
<b>Appendix D. Zorp GPL End-User License Agreement .....</b>	<b>247</b>
D.1. 1. SUBJECT OF THE LICENSE CONTRACT .....	247
D.2. 2. DEFINITIONS .....	247
D.3. 3. LICENSE GRANTS AND RESTRICTIONS .....	248
D.4. 4. SUBSIDIARIES .....	250
D.5. 5. INTELLECTUAL PROPERTY RIGHTS .....	250
D.6. 6. TRADE MARKS .....	250
D.7. 7. NEGLIGENT INFRINGEMENT .....	250
D.8. 8. INTELLECTUAL PROPERTY INDEMNIFICATION .....	250
D.9. 9. LICENSE FEE .....	251
D.10. 10. WARRANTIES .....	251
D.11. 11. DISCLAIMER OF WARRANTIES .....	252
D.12. 12. LIMITATION OF LIABILITY .....	252
D.13. 13. DURATION AND TERMINATION .....	252
D.14. 14. AMENDMENTS .....	252
D.15. 15. WAIVER .....	253
D.16. 16. SEVERABILITY .....	253
D.17. 17. NOTICES .....	253
D.18. 18. MISCELLANEOUS .....	253
<b>Appendix E. Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd) License</b>	
.....	<b>254</b>
Index of Proxy attributes .....	259
Index of Core attributes .....	262
Index of all attributes .....	265



## List of Examples

2.1. Customizing FTP commands .....	5
2.2. Using the POLICY action .....	5
2.3. Default and explicit actions .....	5
2.4. Customizing response codes .....	6
2.5. Example PlugProxy allowing secondary sessions .....	7
2.6. HTTP proxy stacked into an HTTPS connection .....	8
2.7. Program stacking in HTTP .....	8
3.1. Disabling specific TLS protocols .....	14
3.2. Configuring FTPS support .....	15
4.1. Controlling the number of max hops .....	22
4.2. FTP protocol sample .....	26
4.3. Customizing FTP to allow only anonymous sessions .....	27
4.4. Configuring FTPS support .....	28
4.5. Example HTTP transaction .....	37
4.6. Proxy style HTTP query .....	38
4.7. Data tunneling with connect method .....	38
4.8. Implementing URL filtering in the HTTP proxy .....	39
4.9. 404 response filtering in HTTP .....	39
4.10. Header filtering in HTTP .....	40
4.11. URL redirection in HTTP proxy .....	41
4.12. Redirecting HTTP to HTTPS .....	41
4.13. Using parent proxies in HTTP .....	42
4.14. URL-filtering example .....	45
4.15. URL filtering HTTP proxy .....	61
4.16. POP3 protocol sample .....	66
4.17. Example for allowing only APOP authentication in POP3 .....	68
4.18. Example for converting simple USER/PASS authentication to APOP in POP3 .....	68
4.19. Rewriting the banner in POP3 .....	69
4.20. SMTP protocol sample .....	73
4.21. Example for disabling the Telnet X Display Location option .....	83
4.22. Rewriting the DISPLAY environment variable .....	84
4.23. Example WhoisProxy logging all whois requests .....	86
5.1. A simple authentication policy .....	92
5.2. Caching authentication decisions .....	92
5.3. A simple authorization policy .....	93
5.4. BasicAccessList example .....	94
5.5. A simple PairAuthorization policy .....	97
5.6. A simple PermitGroup policy .....	97
5.7. PermitTime example .....	98
5.8. A simple PermitUser policy .....	99
5.9. Outband authentication example .....	100
5.10. A sample authentication provider .....	102
5.11. A DirectedRouter using AvailabilityChainer .....	105
5.12. A sample ConnectChainer .....	106
5.13. A DirectedRouter using FailoverChainer .....	107



5.14. A DirectedRouter using RoundRobinAvailabilityChainer .....	109
5.15. A DirectedRouter using RoundRobinChainer .....	110
5.16. CertDetector example .....	113
5.17. HttpDetector example .....	114
5.18. SNIDetector example .....	115
5.19. SshDetector example .....	116
5.20. Loading a certificate .....	123
5.21. Loading DH parameters .....	135
5.22. Loading a private key .....	144
5.23. Whitelisting e-mail recipients .....	164
5.24. DNSMatcher example .....	165
5.25. RegexpFileMatcher example .....	166
5.26. RegexpMatcher example .....	167
5.27. SmtplInvalidMatcher example .....	168
5.28. WindowsUpdateMatcher example .....	170
5.29. GeneralNat example .....	172
5.30. Using Natpolicies .....	175
5.31. A simple DNSResolver policy .....	184
5.32. A simple HashResolver policy .....	185
5.33. DirectedRouter example .....	188
5.34. InbandRouter example .....	189
5.35. TransparentRouter example .....	190
5.36. Sample rule definitions .....	193
5.37. Tagging rules .....	194
5.38. A simple DenyService .....	199
5.39. PFService example .....	201
5.40. Service example .....	202
5.41. SockAddrInet example .....	210
5.42. SockAddrInet example .....	211
5.43. SockAddrInetHostname example .....	212
5.44. SockAddrUnix example .....	213
5.45. A simple StackingProvider class .....	215
5.46. Using a StackingProvider in an FTP proxy .....	215
5.47. Finding IP networks .....	216
5.48. Zone examples .....	216
5.49. Determining the zone of an IP address .....	217
6.1. CSZoneDispatcher example .....	220
6.2. Dispatcher example .....	221

# List of Procedures

1.1. Zorp startup and initialization .....	1
1.2.1. Handling packet filtering services .....	2
1.2.2. Handling application-level services .....	2
1.3. Proxy startup and the server-side connection .....	3
3.1.1. The SSL handshake .....	9
3.2.7.1. Configuring keybridging .....	16
B.1. Setting global options of Zorp .....	232



# Preface

Welcome to the Zorp GPL Reference Guide. This book contains reference documentation on the available Zorp proxies and their working environment, the Python framework.

This book contains information about the low-level proxy attributes available to customize proxy behavior and the low-level classes comprising Zorp's access control and service framework. Basic introduction to the various protocols is also provided for reference, but the detailed discussion of the protocols is beyond the scope of this book.

## 1. Summary of contents

*Chapter 1, How Zorp works (p. 1)* provides an overview of the internal working of Zorp, for example, how a connection is received.

*Chapter 2, Configuring Zorp proxies (p. 4)* describes the general concepts of configuring Zorp proxies.

*Chapter 3, The Zorp SSL framework (p. 9)* explains how to handle SSL-encrypted connections with Zorp.

*Chapter 4, Proxies (p. 19)* is a complete reference of the Zorp proxies, including their special features and options.

*Chapter 5, Core (p. 88)* is the reference of Zorp core modules which are directly used by gateway administrators, forming the access control and authentication framework.

*Appendix C, Zorp manual pages (p. 237)* is a collection of the command-line Zorp utilities.

*Appendix B, Global options of Zorp (p. 232)* is a reference the global options of Zorp.

## 2. Terminology

The following terms used throughout this documentation might require a brief explanation:

- *class*: A class is a set of attribute and method definitions performing certain specific functionality. Classes can inherit methods and attributes from one or more parent classes. Classes do not contain actual values for attributes; they only describe them.
- *instance*: An instance is a set of attribute values (as described by the class) and associated methods. Instances are also called objects. Instances are created from classes by "calling" the class, with arguments required by the constructor. For example, to create an instance of a class named "class" one would write `class(arg1, arg2 [, .. argN])` where `arg1` and `arg2` are arguments of the constructor.
- *method*: A function working in the context of an instance. It automatically receives a "self" argument which can be used to fetch or set attributes stored in the associated instance.
- *type*: Variables in Python are not strongly typed, meaning that it is possible to assign any kind of values to a variable; typing is assigned to the value.



- *attribute*: An attribute of an object is a variable holding some value, interpreted and manipulated by object methods. Although Python is not strongly typed, types were assigned to the variables in Zorp to indicate what kind of values they are supposed to hold.
- *actiontuple*: A tuple is a simple Python type defined as a list of values. An actiontuple is a special tuple defined by Zorp where the first value must be a value specifying what action to take, and trailing items specify arguments to the action. For example (`HTTP_REQ_REJECT`, "We don't like this request") is a tuple for rejecting HTTP requests and returning the message specified in the second value.

### 3. Target audience and prerequisites

This guide is intended for use by system administrators and consultants responsible for network security and whose task is the configuration and maintenance of Zorp firewalls. Zorp gives them a powerful and versatile tool to create full control over their network traffic and enables them to protect their clients against Internet-delinquency.

This guide is also useful for IT decision makers evaluating different firewall products because apart from the practical side of everyday Zorp administration, it introduces the philosophy behind Zorp without the marketing side of the issue.

The following skills and knowledge are necessary for a successful Zorp administrator.

Skill	Level/Description
Linux	At least a power user's knowledge is required.
Experience in system administration	Experience in system administration is certainly an advantage, but not absolutely necessary.
Programming language knowledge	It is not an explicit requirement to know any programming languages though being familiar with the basics of Python may be an advantage, especially in evaluating advanced firewall configurations or in troubleshooting misconfigured firewalls.
General knowledge on firewalls	A general understanding of firewalls, their roles in the enterprise IT infrastructure and the main concepts and tasks associated with firewall administration is essential. To fulfill this requirement a significant part of <i>Chapter 3, Architectural overview</i> in the <i>Zorp Administrator's Guide</i> is devoted to the introduction to general firewall concepts.
Knowledge on Netfilter concepts and IPTables	In-depth knowledge is strongly recommended; while it is not strictly required definitely helps understanding the underlying operations and also helps in shortening the learning curve.



Skill	Level/Description
Knowledge on TCP/IP protocol	High level knowledge of the TCP/IP protocol suite is a must, no successful firewall administration is possible without this knowledge.

Table 1. Prerequisites

## 4. Products covered in this guide

The Zorp Distribution DVD-ROM contains the following software packages:

- Current version of Zorp 7 packages.
- Current version of () 7.
- Current version of () 7 (GUI) for both Linux and Windows operating systems, and all the necessary software packages.
- Current version of () 7.
- Current version of the () 7, the client for both Linux and Windows operating systems.

For a detailed description of hardware requirements of Zorp, see .

For additional information on Zorp and its components visit the [Zorp website](#) containing white papers, tutorials, and online documentations on the above products.

## 5. Contact and support information

This product is developed and maintained by Balasys IT Zrt..

### Contact:

Balasys IT Zrt.  
 4 Alíz Street  
 H-1117 Budapest, Hungary  
 Tel: +36 1 646 4740  
 E-mail: <info@balasys.hu>  
 Web: <http://balasys.hu/>

### 5.1. Sales contact

You can directly contact us with sales related topics at the e-mail address <sales@balasys.hu>, or *leave us your contact information and we call you back.*

### 5.2. Support contact

To access the Balasys Support System, sign up for an account at *the Balasys Support System page.* Online support is available 24 hours a day.

Balasys Support System is available only for registered users with a valid support package.

Support e-mail address: <support@balasys.hu>.

### **5.3. Training**

Balasys IT Zrt. holds courses on using its products for new and experienced users. For dates, details, and application forms, visit the <https://www.balasys.hu/en/services#training> webpage.

## **6. About this document**

This guide is a work-in-progress document with new versions appearing periodically.

The latest version of this document can be downloaded from <https://docs.balasys.hu/>.

### **6.1. Feedback**

Any feedback is greatly appreciated, especially on what else this document should cover, including protocols and network setups. General comments, errors found in the text, and any suggestions about how to improve the documentation is welcome at <support@balasys.hu>.



# Chapter 1. How Zorp works

This chapter describes how Zorp works, and provides information about the core Zorp modules, explaining how they interoperate. For a detailed reference of the core modules, see the description of the particular in *Chapter 5, Core (p. 88)*.

- ***Zorp startup and initialization***: The main Zorp thread is started, and the rules listening for incoming connections are initialized.
- ***Handling incoming connections***: The client-side connection is established and the service to proxy the connection is selected.
- ***Proxy startup and server-side connections***: The proxy instance inspecting the traffic is created and connection to the server is established.

## 1.1. Procedure – Zorp startup and initialization

Step 1. The `zorpctl` utility loads the `instances.conf` file and starts the main zorp program. The `instances.conf` file stores the parameters of the configured Zorp instances.

Step 2. zorp performs the following initialization steps:

- Sets the stack limit.
- Creates its PID file.
- Changes the running user to the user and group specified for the instance.
- Initializes the handling of dynamic capabilities and sets the chroot directory.
- Loads the firewall policy from the `policy.py` file.

Step 3. The `init()` of Zorp initializes the ruleset defined for the Zorp instance.

Step 4. The `kzorp` kernel module uploads packet filtering services, rules, and zones into the kernel.

**Note**

Zorp creates four sockets (one for each type of traffic: TCP IPv6, TCP IPv4, UDP IPv6, UDP IPv4); the `kzorp` module directs the incoming connections to the appropriate socket.

## 1.2. Handling incoming connections

Incoming connections are first received by the `kzorp` kernel module, which is actually a netfilter table. The `kzorp` module determines the source and destination zones of the connection, and then tries to find a suitable firewall rule. If the rule points to a packet filtering service, the connection is processed according to *Procedure 1.2.1, Handling packet filtering services (p. 2)*; if it points to an application-level service, the connection is processed according to *Procedure 1.2.2, Handling application-level services (p. 2)*. If no suitable rule is found, the connection is rejected.



### 1.2.1. Procedure – Handling packet filtering services

Step 1. Zorp generates a session ID and creates a CONNTRACK entry for the connection. This ID is based on all relevant information about the connection, including the protocol (TCP/UDP) and the client's address.

The session ID uniquely identifies the connection and is included in every log message related to this particular connection.

Step 2. Based on the parameters of the connection, the Rule selects the service that will inspect the connection.

Step 3. The Router defined in the service determines the destination address of the server.

The Router performs the following actions:

- Determines the destination address of the server.
- Sets the source address of the server-side connection, according to the *forge\_address* settings of the router.

Step 4. If the client is permitted to access the selected service, the packet filter is instructed to let the connection pass Zorp.

Step 5. The kzorp module performs network address translation (NAT) on the connection, if needed.

### 1.2.2. Procedure – Handling application-level services

Step 1. For incoming connection requests that are processed on the application level, the main Zorp thread establishes the connection with the client. The connection is further processed in a separate thread; the main thread is listening for new connections.

Step 2. The *Dispatcher* creates the *MasterSession* object of the connection and generates the base session ID. This object stores all relevant information of the connection, including the protocol (TCP/UDP) and the client's address.

The session ID uniquely identifies the connection and is included in every log message related to this particular connection. Other components of Zorp add further digits to the session ID.

Step 3. For TCP-based connections, Zorp copies the Type of Service (ToS) value of the client-Zorp connection in the Zorp-client connection.

Step 4. The Rule selects the service that will inspect the connection.

Step 5. The Router defined in the service determines the destination address of the server. The result is stored in the Session object, where the Chainer can access it later.

The Router performs the following actions:

- Determines the destination address of the server.
- Sets the source address of the server-side connection (according to the *forge\_port*, *forge\_address* settings of the router).
- Sets the ToS value of the server-side connection.

Step 6. If the client is permitted to access the selected service, the `startInstance()` method of the service is started. The `startInstance()` method performs the following actions:

- Verifies that the new instance does not exceed the number of instances permitted for the service (*max\_instances* parameter).
- Creates the final session ID.
- Creates an instance of the proxy class associated with the service. This proxy instance is associated with a *StackedSession* object. The startup of the proxy is detailed in *Procedure 1.3, Proxy startup and the server-side connection (p. 3)*.

### 1.3. Procedure – Proxy startup and the server-side connection

- Step 1. To create an instance of the application-level proxy, the `__init__` constructor of the proxy class calls the `Proxy.__init__` function of the *Proxy* module. The proxy instance is created into a new thread from the *ZorpProxy* ancestor class.
- Step 2. From the new thread, the proxy loads its configuration.
- Step 3. The proxy initiates connection to the server.

**Note**

Some proxies connect the server only after receiving the first client request.

- Step 4. The `Proxy.connectServer()` method creates the server-side connection using the *Chainer* assigned to the service. The Chainer performs the following actions:
- Reads the parameters related to the server-side connection from the *Session* object. These parameters were set by the *Router* and the *Proxy*.
  - Performs source and destination network address translation. This may modify the addresses set by the *Router* and the *Proxy*.
  - Verifies that access to the server is permitted.
  - Establishes the connection using the *Attach* subsystem, and passes to the proxy the stream that represents the connection.

**Note**

The `Proxy.connectServer()` method connects stacked proxies with their parent proxies.



## Chapter 2. Configuring Zorp proxies

This chapter describes how Zorp proxies work in general, and how to configure them.

- For the details on configuring TLS/SSL connections, see *Chapter 2, Configuring Zorp proxies (p. 4)*.
- For a complete reference of the available Zorp proxies, see *Chapter 4, Proxies (p. 19)*.

### 2.1. Policies for requests and responses

Zorp offers great flexibility in proxy customization. Requests and commands, responses, headers, etc. can be managed individually in Zorp. This means that it is not only possible to enable/disable them one-by-one, but custom actions can be assigned to them as well. The available options are listed in the description of each proxy, but the general guidelines are discussed here.

All important events of a protocol have an associated policy hash: usually there is one for the requests or commands and one for the responses. Where applicable for a protocol, there are other policy hashes defined as well (e.g., for controlling the capabilities available in the IMAP protocol, etc.). The entries of the hash are the possible events of the protocol (e.g., the request hash of the FTP protocol contains the possible commands - RMD, DELE, etc.) and an action associated with the event - what Zorp should do when this event occurs. The available actions may slightly vary depending on the exact protocol and the hash, but usually they are the following:

Action	Description
ACCEPT	Enable the event; the command/response/etc. can be used and is allowed through the firewall.
REJECT	Reject the event and send an error message. The event is blocked and the client notified. The communication can continue, the connection is not closed.
DROP	Reject the event without sending an error message. The event is blocked but the client is not notified. The communication can continue, the connection is not closed. In some cases (depending on the protocol) this action is able to remove only a part of the message (e.g., a particular header in HTTP traffic) without rejecting the entire message.
ABORT	Reject the event and terminate the connection.
POLICY	Call a Python function to make a decision about the event. The final decision must be one of the above actions (i.e. POLICY is not allowed). The parameters received by the function are listed in the module

Action	Description
	descriptions. See the examples below and in the module descriptions for details.

Table 2.1. Action codes for protocol events

The use of the policy hashes and the action codes is illustrated in the following examples.



#### Example 2.1. Customizing FTP commands

In this example the 'RMD' command is rejected, and the connection is terminated if the user attempts to delete a file.

```
class MyFtp(FtpProxy):
    def config(self):
        self.request["RMD"] = (FTP_REQ_REJECT)
        self.request["DELE"] = (FTP_REQ_ABORT)
```



#### Example 2.2. Using the POLICY action

This example calls a function called pUser (defined in the example) whenever a USER command is received within an FTP session. All other commands are accepted. The parameter of the USER command (i.e. the username) is examined: if it is 'user1' or 'user2', the connection is accepted, otherwise it is rejected.

```
class MyFtp(FtpProxy):
    def config(self):
        self.request["USER"] = (FTP_REQ_POLICY, self.pUser)
        self.request["*"] = (FTP_REQ_ACCEPT)

    def pUser(self, command):
        if self.request_parameter == "user1" or self.request_parameter == "user2":
            return FTP_REQ_ACCEPT
        return FTP_REQ_REJECT
```

It must be noted that there is a difference between how Zorp processes the POLICY actions and all the other ones (e.g., ACCEPT, DROP, etc.). POLICY actions are evaluated on the policy (or Python) level of Zorp, while the other ones on the proxy (or C) level. Since the proxies of Zorp are written in C, and operate on the proxy level, the evaluation of POLICY actions is slightly slower, but this can be an issue only in very high-throughput environments with complex policy settings.

### 2.1.1. Default actions

Default actions for all events of a hash (e.g., all requests) can be set using the '\*' wildcard as the event. (Most hashes have default actions configured by default, these can be found in the description of the proxy classes.) It is important to note that setting the action using the '\*' wildcard does NOT override an action explicitly defined for an event, even if the explicit setting precedes the general one in the Python code. This feature is illustrated in the example below.



#### Example 2.3. Default and explicit actions

The following two proxy classes have the same effect, even though the order of the code lines is switched. The 'APPE' command is rejected, while all other commands are accepted.

```
class MyFtp1(FtpProxy):
    def config(self):
        self.request["APPE"] = (FTP_REQ_REJECT)
        self.request["*"] = (FTP_REQ_ACCEPT)
```

```
class MyFtp2(FtpProxy):
    def config(self):
        self.request["*"] = (FTP_REQ_ACCEPT)
        self.request["APPE"] = (FTP_REQ_REJECT)
```

**Warning**

If the relevant hash does not contain a received request or response, the '\*' entry is used which matches to every request/response. If there is no '\*' entry in the given hash, the request/response is denied.

## 2.1.2. Response codes

Responses in certain protocols include numeric response codes, e.g., in the FTP protocol responses start with a three-digit code. In Zorp it is possible to filter these codes as well, furthermore, to filter them based on the command to which the response arrives to. In these cases the hash contains both the command and the answer, and an action as well. The '\*' wildcard character can be used to match for every command or response code.

**Example 2.4. Customizing response codes**

The following example accepts the response '250' only to the 'DELE' command, but allows any response code to the 'LIST' command.

```
class MyFtp1(FtpProxy):
    def config(self):
        self.response["DELE", "250"] = (FTP_RSP_ACCEPT)
        self.response["*", "250"] = (FTP_RSP_REJECT)
        self.response["LIST", "*"] = (FTP_RSP_ACCEPT)
```

It is not necessary to specify the full response code, it is also possible to specify only the first, or the first two digits.

For example, all three response codes presented below are valid, but have different effects:

- "PWD", "200"  
Match exactly the answer 200 coming in a reply to a PWD command.
- "PWD", "2"  
Match every answer starting with '2' in a reply to a PWD command.
- "\*", "20"  
Match every answer between 200 and 209 in a reply to any command.

This kind of response code lookup is available in the following proxies: FTP, HTTP, NNTP, and SMTP. The precedence how the hash table entries are processed is the following:

1. Exact match. ("PWD", "200")
2. Exact command match, partial response matches ("PWD", "20"; "PWD", "2"; "PWD", "\*")
3. Wildcard command, with answer codes repeated as above. ("\*", "200"; "\*", "20"; "\*", "2")
4. Wildcard for both indexes. ("\*", "\*")

## 2.2. Secondary sessions

Certain proxies support the use of secondary sessions, i.e. several sessions using the same proxy instance (the same thread), effectively reusing proxy instances. As new sessions enter the proxy via a fastpath, using secondary sessions can significantly decrease the load on the firewall.

When a new connection is accepted, Zorp looks for the appropriate proxy instance which is willing to accept secondary sessions. If there is none, a new proxy instance is started. An already running proxy instance is appropriate if it is willing to accept secondary channels and the criteria about secondary sessions are met. (The criteria can be specified in the configuration of the proxy class.)

The criteria are set via the *secondary\_mask* attribute, while the number of secondary sessions allowed within the same instance is controlled by the *secondary\_sessions* attribute. The *secondary\_mask* attribute is an integer specifying which properties of an established session are considered to be important. If all important properties match, the connection can be handled as a secondary session by a proxy instance accepting secondary sessions, provided the new session does not exceed the limit set in *secondary\_sessions*. The *secondary\_mask* attribute is actually a bitfield interpreted as follows: bit 0 means source address; bit 1 means source port; bit 2 means destination address; bit 3 means destination port.

Currently the Plug supports the use of secondary sessions.



### Example 2.5. Example PlugProxy allowing secondary sessions

This example allows 100 parallel sessions in one proxy thread if the IP address and Port of the targets are the same.

```
class MyPlugProxy(PlugProxy):
    def config(self):
        PlugProxy.config(self)
        self.secondary_mask = 0xC
        self.secondary_sessions = 100
```

## 2.3. Embedded protocol analysis

Each protocol proxy available in Zorp inspects the traffic for conformance to the given protocol. Often further analysis of the data transferred via the protocol is required, this can be accomplished via stacking. Stacking is a method when the data transferred in the protocol is passed to another proxy or program. After performing the inspection, the stacked proxy or program returns the data to the original proxy, which resumes its transmission.

### 2.3.1. Proxy stacking

Proxy stacking is mainly used to inspect embedded protocols, or perform virus filtering: e.g., to inspect HTTPS traffic, the external SSL protocol is examined with a Pssl proxy, and then a HTTP proxy is stacked to inspect the internal protocol. It is possible to stack several layers of proxies into each other if needed, e.g., in the above example, a further virus filtering solution (like a module) could be stacked into the HTTP proxy.



#### Note

Starting with Zorp version 3.3FR1, every proxy is able to handle SSL/TLS-encrypted connection on its own, making the Pssl proxy redundant. This feature greatly decreases the need of proxy stacking, making it needed only in special cases, for example, to inspect HTTP traffic tunneled in SSH.

Stacking a proxy to inspect the embedded protocol is possible via the `self.request_stack` attribute; if another attribute has to be used, it is noted in the description of the given proxy. The HTTP proxy is special in the sense that it is possible to stack different proxies into the requests and the responses.

The parameters of the stack attribute has to specify the following:

- The protocol elements for which embedded inspection is required. This parameter can be used to specify if all received data should be passed to the stacked proxy ("\*"), or only the data related (sent or received) to specific protocol elements (e.g., only the data received with a GET request in HTTP).
- The mode how the data is passed to the stacked proxy. This parameter governs if only the data part should be passed to the stacked proxy (XXXX\_STK\_DATA, where XXXX depends on the protocol), or (if applicable) MIME header information should be included as well (XXXX\_STK\_MIME) to make it possible to process the data body as a MIME envelope. Please note that while it is possible to change the data part in the stacked proxy, it is not possible to change the MIME headers - they can be modified only by the upper level proxy. The available constants are listed in the respective protocol descriptions. The default value for this argument is XXXX\_STK\_NONE, meaning that no data is transferred to the stacked proxy. In some proxies it is also possible to call a function (using the XXXX\_STK\_POLICY action) to decide which part (if any) of the traffic should be passed to the stacked proxy.
- The proxy class that will perform inspection of the embedded protocol.

The use of proxy stacking is illustrated in the following example:



**Example 2.6. HTTP proxy stacked into an HTTPS connection**

The following proxy class stacks an Http proxy into a Pssl Proxy to inspect HTTPS traffic.

```
class HttpsPsslProxy(PsslProxy):
    def config(self):
        PsslProxy.config(self)
        self.stack_proxy=(Z_STACK_PROXY, HttpProxy)
```

For additional information on proxy stacking, see , and the various tutorials available at the [Balasys Documentation Page](#).

### 2.3.2. Program stacking

When stacking a program, the data received by a proxy within a protocol is directed to the standard input. Arbitrary commands (including command line scripts, or applications) working from the standard input can be run on this data stream. The original proxy obtains the processed data back from the standard output. When stacking a command, the command to be called has to be included in the proper stack attribute of the proxy between double-quotes. This is illustrated in the following example.



**Example 2.7. Program stacking in HTTP**

In this example a simple 'sed' (stream editor) command is stacked into the HTTP proxy to replace all occurrences of 'http' to 'https', thus securing the HTTP connections on one side of the firewall.

```
class MyHttp(HttpProxy):
    def config(self):
        HttpProxy.config(self)
        self.response_stack["GET"] = /
        (HTTP_STK_DATA, (Z_STACK_PROGRAM, "/bin/sed '/http:/s//https:/g'))
```



## Chapter 3. The Zorp SSL framework

This chapter describes the SSL protocol and the SSL framework available for every Zorp proxy.

### 3.1. The SSL protocol

Secure Socket Layer v3 (SSL) and Transport Layer Security v1 (TLS) are widely used crypto protocols guaranteeing data integrity and confidentiality in many PKI and e-commerce systems. They allow both the client and the server to authenticate each other. SSL/TLS use reliable TCP connection for data transmission and cooperate with any application level protocol. SSL/TLS guarantee that:

- Communication in the channel is private, only the other communicating party can decrypt the messages.
- The channel is authenticated, so the client can make sure that it communicates with the right server. Optionally, the server can also authenticate the client. Authentication is performed via certificates issued by a Certificate Authority (CA). Certificates identify the owner of an encryption keypair used in encrypted communication.
- The channel is reliable, which is ensured by message integrity verification using MAC.

SSL/TLS is almost never used in itself: it is used as a secure channel to transfer other, less secure protocols. The protocols most commonly embedded into SSL/TLS are HTTP and POP3 (i.e. these are the HTTPS and POP3S protocols).

#### 3.1.1. Procedure – The SSL handshake

As an initial step, both the client and the server collect information to start the encrypted communication.

Step 1. The client sends a CLIENT-HELLO message.

Step 2. The server answers with a SERVER-HELLO message containing the certificate of the server. At this point the parties determine if a new master key is needed.

**Note**

The server stores information (including the session ID and other parameters) about past SSL/TLS sessions in its session cache. Clients that have contacted a particular server previously can request to continue a session (by identifying its session ID); this can be used to accelerate the initialization of the connection. Zorp currently does not support this feature, but this does not cause any noticeable difference to the clients.

Step 3. The client verifies the server's certificate. If the certificate is invalid the client sends an ERROR message to the server.

**Note**

If a new master key is needed the client gets the server certificate from the SERVER-HELLO message and generates a master key, sending it to the server in a CLIENT-MASTER-KEY message.

- Step 4. The server sends a SERVER-VERIFY message, which authenticates the server itself.
- Step 5. Optionally, the server can also authenticate the client by requesting the client's certificate with a REQUEST-CERTIFICATE message.
- Step 6. The server verifies the certificate received from the client and finishes the handshake with a SERVER-FINISH message.

**Note**

In SSL two separate session keys are used, one for outgoing communication (which is of course incoming at the other end), and another key for incoming communication. These are known as SERVER/CLIENT-READ-KEY and SERVER/CLIENT-WRITE-KEY.

## 3.2. Configuring TLS and SSL encrypted connections

Zorp version 3.3FR1 introduces a common framework that allows every Zorp proxy to use SSL/TLS encryption, and also to support STARTTLS.

**Note**

In Zorp 7, the following proxies support STARTTLS: Ftp proxy (to start FTPS sessions), Smtplib proxy.

### 3.2.1. Behavior of the SSL framework

**Warning**

For the details of the attributes related to the SSL framework, see *Section 5.10.4, Class Proxy (p. 181)*.

The SSL framework was built for inspecting SSL/TLS connections, and also any other connections embedded into the encrypted SSL/TLS channel. SSL/TLS connections initiated from the client are terminated on the Zorp firewall; and two separate SSL/TLS connections are built: one between the client and the firewall, and one between the firewall and the server. If both connections are accepted by the local security policy (the certificates are valid, and only the allowed encryption algorithms are used), Zorp inspects the protocol embedded into the secure channel as well.

Several configuration examples and considerations are discussed in the Technical White Paper and Tutorial *Proxying secure channels - the Secure Socket Layer*, available at the [Balasys Documentation Page](#).

#### 3.2.1.1. General behavior

The SSL framework starts its operation by inspecting the values set in the `ssl.handshake_seq` attribute. When this attribute is set to `SSL_HSO_CLIENT_SERVER` the client side, otherwise (`SSL_HSO_SERVER_CLIENT`) the server side handshake is performed first.

As part of the handshake process the proxy checks if SSL is enabled on the given side (`ssl.client_connection_security` and `ssl.server_connection_security` attributes). It is not necessary for SSL to be enabled on both sides - Zorp can handle one-sided SSL connections as well (e.g., the



firewall communicates in an unencrypted channel with the client, but in a secure channel with the server). If SSL is not enabled, the handshake is skipped for that side.

When SSL is needed, the proxy will cooperate with the policy layer to have all required parameters (keys, certificates, etc.) set up. This is achieved using decision points in the hash named *ssl.handshake\_hash* which is explained later in detail.

The SSL handshake is slightly different for the client (in this case Zorp behaves as an SSL server) and the server (when Zorp behaves as an SSL client).

### 3.2.1.2. Client-side (SSL server) behavior

As an SSL server the first thing to present to an SSL client is a certificate/key pair, thus a call to the 'setup\_key' callback is made. It is expected that by the time this callback returns the attributes *ssl.client\_local\_privatekey* and *ssl.client\_local\_certificate* are filled appropriately.

If peer authentication is enabled (by setting the attribute *ssl.client\_verify\_type*) a list of trusted CA certificates must be set up (stored in the hash *ssl.client\_local\_ca\_list*). The list can be set up by the 'setup\_ca\_list' function call. Peer certificates are verified against the trusted CA list and their associated revocation lists. Revocations can be set up in the 'setup\_crl\_list' callback.

At the end of the verification another callback named 'verify\_cert' is called which can either ACCEPT or DENY the certificate possibly overriding the verification against the local CA database.

### 3.2.1.3. Server-side (SSL client) behavior

The server-side handshake is similar to the client-side handshake previously described. The difference is the order of certificate verification. On the server side Zorp verifies the server's certificate first and then sends its own certificate for verification. This is unlike the client side where the local certificate is sent first, and then the peer's certificate is verified.

So the callbacks are called in this order: 'setup\_ca\_list' and 'setup\_crl\_list' to set up CA and CRL information, 'verify\_cert' to finalize certificate validation, and 'setup\_key' to optionally provide a local certificate/key pair.

## 3.2.2. Handshake callbacks

As described earlier, the SSL framework provides a way to customize the SSL handshake process. This is done using the *ssl.client\_handshake* and *ssl.server\_handshake* hashes. These hashes are indexed by the keywords listed below.

The tuple can be separated to two parts: 1) tuple type, 2) parameters for the given type. For now only *SSL\_HS\_POLICY* is valid as tuple type, and it requires a function reference as parameter.

The following keys are accepted as indexes:

setup\_key

This function is called when the proxy needs the private key/certificate pair to be set up. All attributes filled in the earlier phases can be used to decide which key/certificate to use. The function expects two parameters: self, side.



<code>setup_ca_list</code>	This function is called when the proxy needs the trusted CA list to be set up. The function expects two parameters: <code>self</code> , <code>side</code> .
<code>setup_crl_list</code>	This function is called when the proxy needs the CRL list to be set up. This function gets a single string parameter which contains the name of the CA whose CRL is to be filled up. The function expects three parameters: <code>self</code> , <code>side</code> , <code>ca_name</code> .
<code>verify_cert</code>	This function is called to finalize the verification process. The function expects two parameters: <code>self</code> , <code>side</code> .

The function arguments as referenced above are defined as:

<code>self</code>	The proxy instance.
<code>side</code>	The side where handshake is being performed.
<code>ca_name</code>	Name of an X.509 certificate.

The functions returns one of the `SSL_HS_*` constants. Generally if the function returns `SSL_HS_ACCEPT` the handshake continues, otherwise the handshake is aborted. As an exception, `verify_cert` may return `SSL_HS_VERIFIED` in which case the certificate is accepted without further verification.

Name	Value
<code>SSL_HS_ACCEPT</code>	0
<code>SSL_HS_REJECT</code>	1
<code>SSL_HS_POLICY</code>	6
<code>SSL_HS_VERIFIED</code>	10

Table 3.1. Handshake policy decisions

### 3.2.3. X.509 Certificates

An X.509 certificate is a public key with a subject name specified as an X.500 DN (distinguished name) signed by a certificate issuing authority (CA). X.509 certificates are represented as Python policy objects having the following attributes:

<code>subject</code>	Subject of the certificate.
<code>issuer</code>	Issuer of the certificate (i.e. the CA that signed it).
<code>serial</code>	Serial number of the certificate.
<code>blob</code>	The certificate itself as a string in PEM format.

Zorp uses X.509 certificates to provide a convenient and efficient way to manage and distribute certificates and keys used by the various components and proxies of the managed firewall hosts. It is mainly aimed at providing certificates required for the secure communication between the different parts of the firewall system, e.g. firewall hosts and engine (the actual communication is realized by agents).

Certificates of trusted CAs (and their accompanying CRLs) are used in Zorp to validate the certificates of servers accessed by the clients. The hashes and structures below are used by the various certificate-related attributes of the Zorp Pssl proxy, particularly the ones of `certificate` type.

### 3.2.3.1. X.509 Certificate Names

A certificate name behaves as a string, and contains a DN in the following format (also known as one-line format):

```
/RDN=value/RDN=value/.../RDN=value/
```

The word RDN stands for relative distinguished name. For example, the DN *cn=Root CA, ou=CA Group, o=Foo Ltd, l=Bar, st=FooBar State, c=US* becomes */C=US/ST=FooBar State/L=Bar/O=Foo Ltd/OU=CA Group/CN=Root CA/*



#### Warning

The format and representation of certificate names may change in future releases.

### 3.2.3.2. X.509 Certificate Revocation List

A certifying authority may revoke the issued certificates. A revocation means that the serial number and the revocation date is added to the list of revoked certificates. Revocations are published on a regular basis. This list is called the Certificate Revocation List, also known as CRL. A CRL always has an issuer, a date when the list was published, and the expected date of its next update.

### 3.2.3.3. X.509 Certificate hash

The proxy stores trusted CA certificates in a Certificate hash. This hash can be indexed by two different types. If an integer index is used, the slot specified by this value is looked up; if a string index is used, it is interpreted as a one-line DN value, and the appropriate certificate is looked up. Each slot in this hash contains an X.509 certificate.

### 3.2.3.4. X.509 CRL hash

Similarly to the certificate hash, a separate hash for storing Certificate Revocation Lists was defined. A CRL contains revocation lists associated to CAs.

### 3.2.3.5. Certificate verification options

Zorp is able to automatically verify the certificates received. The types of accepted certificates can be controlled separately on the client and the server side using the *ssl.client\_verify\_type* and the *ssl.server\_verify\_type* attributes. These attributes offer an easy way to restrict encrypted access only to sites having trustworthy certificates. The available options are summarized in the following table.

Name	Value
TLS_TRUST_LEVEL_NONE	Accept invalid for example, expired certificates.
TLS_TRUST_LEVEL_UNTRUSTED	Both trusted and untrusted certificates are accepted.

Name	Value
TLS_TRUST_LEVEL_FULL	Only valid certificates signed by a trusted CA are accepted.

Table 3.2. Constants for trust level selection.

The `ssl.server_check_subject` can be used to compare the domain name provided in the `Subject` field of the server certificate to application level information about the server. Currently it can compare the `Subject` field to the domain name of the HTTP request in HTTPS communication. If the `ssl.server_check_subject` is set to `TRUE` and `ssl.server_verify_type` is `SSL_VERIFY_REQUIRED_UNTRUSTED` or `SSL_VERIFY_REQUIRED_TRUSTED`, the HTTP proxy using the SSL framework will deny access to the page and return an error if the `Subject` field does not match the domain name of the URL.

### 3.2.4. Setting the allowed TLS protocol

There are different and sometimes incompatible releases of the TLS protocol. TLS protocols can be set via the `ClientSSLOptions` and `ServerSSLOptions` classes, enabling all supported protocols by default. Set the appropriate `'disable_tls*'` parameters to disable the selected TLS protocols. Zorp currently supports the TLS v1, TLS v1.1, TLS v1.2 protocols.



#### Example 3.1. Disabling specific TLS protocols

The following example disables the TLSv1 protocol on the client and the server side.

```
EncryptionPolicy(
  name="MyTLSEncryption",
  encryption=TwoSidedEncryption(
    client_verify=None,
    server_verify=ServerCertificateVerifier(
      ca_directory="/etc/ca.d/certs/",
      crl_directory="/etc/ca.d/crls/",
    )
    client_ssl_options=ServerSSLOptions(disable_proto_tlsv1=TRUE),
    server_ssl_options=ServerSSLOptions(disable_proto_tlsv1=TRUE),
  )
)
```

### 3.2.5. SSL cipher selection

The cipher algorithms used for key exchange and mass symmetric encryption are specified by the attributes `ssl.client_ssl_ciphers` and `ssl.server_ssl_ciphers`. These attributes contain a cipher specification as specified by the OpenSSL manuals, see the manual page `ciphers(ssl)` for further details.

The default set of ciphers can be set by using the following predefined variables.

Name	Value
SSL_CIPHERS_HIGH	n/a
SSL_CIPHERS_MEDIUM	n/a
SSL_CIPHERS_LOW	n/a
SSL_CIPHERS_ALL	n/a



Name	Value
SSL_CIPHERS_CUSTOM	n/a

Table 3.3. Constants for cipher selection

Cipher specifications, as defined above, are sorted by key length. The cipher providing the best key length will be the most preferred one.

### 3.2.6. Enabling STARTTLS

Starting with version 3.3FR1, Zorp supports the STARTTLS method for encrypting connections. STARTTLS support can be configured separately for the client- and server side using the `ssl.client_connection_security` and `ssl.server_connection_security` parameters, respectively. The parameters have the following possible values:

Name	Value
SSL_NONE	Disable encryption between Zorp and the peer.
SSL_FORCE_SSL	Require encrypted communication between Zorp and the peer.
SSL_ACCEPT_STARTTLS	Permit STARTTLS sessions. Currently supported only in the Ftp, SmtP and Pop3 proxies.

Table 3.4. Client connection security type.

Name	Value
SSL_NONE	Disable encryption between Zorp and the peer.
SSL_FORCE_SSL	Require encrypted communication between Zorp and the peer.
SSL_FORWARD_STARTTLS	Forward STARTTLS requests to the server. Currently supported only in the Ftp, SmtP and Pop3 proxies.

Table 3.5. Server connection security type.

**Note**

In Zorp 7, the following proxies support STARTTLS: Ftp proxy (to start FTPS sessions), SmtP proxy.

**Example 3.2. Configuring FTPS support**

This example is a standard FtpProxy with FTPS support enabled.

```
class FtpsProxy(FtpProxy):
    def config(self):
        FtpProxy.config(self)
        self.max_password_length=64

        EncryptionPolicy(name="ForwardSTARTTLS",
encryption=ForwardStartTLSEncryption(client_verify=ClientCertificateVerifier(),
client_ssl_options=ClientSSLOptions(), server_verify=ServerCertificateVerifier(),
server_ssl_options=ServerSSLOptions(),
client_certificate_generator=DynamicCertificate(private_key=PrivateKey.fromFile(key_file_path="/etc/key.d/ZMS_Engine/key.pem"),
trusted_ca=Certificate.fromFile(certificate_file_path="/etc/ca.d/certs/my-trusted-ca-cert.pem"),
```

```
private_key=PrivateKey.fromFile("/etc/ca.d/keys/my-trusted-ca-cert.pem")),
untrusted_ca=Certificate.fromFile(certificate_file_path="/etc/ca.d/certs/my-untrusted-ca-cert.pem",
private_key=PrivateKey.fromFile("/etc/ca.d/keys/my-untrusted-ca-cert.pem")))))

def demo() :
    Service(name='demo/MyFTPSService', router=TransparentRouter(), chainer=ConnectChainer(),
proxy_class=FtpsProxy, max_instances=0, max_sessions=0, keepalive=Z_KEEPALIVE_NONE,
encryption_policy="ForwardSTARTTLS")

Rule(rule_id=2,
proto=6,
service='demo/MyFTPSService'
)
```

### 3.2.7. Keybridging certificates

Keybridging is a method to let the client see a copy of the server's certificate (or vice versa), allowing it to inspect it and decide about its trustworthiness. Because of proxying the SSL/TLS connection, the client is not able to inspect the certificate of the server directly, therefore Zorp generates a certificate based on the server's certificate on-the-fly. This generated certificate is presented to the client.

#### 3.2.7.1. Procedure – Configuring keybridging

##### Purpose:

To configure keybridging in a proxy, complete the following steps.

##### Steps:

Step 1. Create the required keys and CAs.

- Step a. Generate two local CA certificates. One of them will be used to sign bridging certificates for servers having trusted certificates, the other one for servers with untrusted or self-signed certificates. It is useful to reflect this difference somewhere in the CA's certificates, for example, in their common name (CA\_for\_Untrusted\_certs; CA\_for\_Trusted\_certs). These CA certificates can either be self-signed or signed by a local root CA. The certificate of the CA signing the trusted certificates should be imported to your clients to make the generated certificates 'trusted'. The other CA certificate should not be imported to the clients.



##### Warning

IMPORTANT: Do NOT set a password for these CAs, as they have to be accessible automatically by Zorp.

- Step b. Generate a new certificate. The private key of this keypair will be used in the on-the-fly generated certificates, the public part (DN and similar information) will not be used.
- Step c. Copy the generated certificate, the CA certificates, and the keys to the firewall, for example, into `/etc/zorp/sslbridge/`. This directory will be used in the `ssl.client_ca_directory` option.



**Note**

If you want to send the root CA of the CA certificates to the clients, also copy the root CA (and any intermediate CA certificates) to this directory.

Step d. Create a cache directory to store the keybridged certificates generated by Zorp, for example, `/var/lib/zorp/sslbridge/`, and make it writable for the `zorp` user.

**Note**

Zorp automatically creates a file called `serial.txt` in the cache directory. If you delete the certificates from the cache, do NOT delete this file. If you accidentally delete it, recreate it, and make sure that it is writable for the `zorp` user.

Step 2. Set up a proxy class (for example, a class derived from the `HttpProxy` class) and set the following attributes with the following values:

- Instruct Zorp to perform the handshake with the server first:  
`self.ssl.handshake_seq=SSL_HSO_SERVER_CLIENT`

```
class KeybrideHttpsProxy(HttpProxy):
    def config(self):
        self.ssl.handshake_seq=SSL_HSO_SERVER_CLIENT
```

- Enable keybridging. Depending on the direction the keybridging is performed, add the `self.ssl.client_keypair_generate` or the `self.ssl.server_keypair_generate` parameter, and set it to `TRUE`. When the generated certificates are shown to the clients, the `self.ssl.client_keypair_generate` parameter has to be used. (Actually, if a `keypair_generate` parameter is set, the proxy will request a keypair from the `key_generator` class. This class — discussed a bit later — returns either a newly generated keypair, or if its `key_file` parameter is set, a pregenerated keypair. In this example this latter option will be used.)

```
class KeybrideHttpsProxy(HttpProxy):
    def config(self):
        self.ssl.handshake_seq=SSL_HSO_SERVER_CLIENT
        self.ssl.client_keypair_generate = TRUE
```

- Configure the `key_generator` class. Note that the parameters of this class must be added to the proxy as a single line, for example:

```
class KeybrideHttpsProxy(HttpProxy):
    def config(self):
        self.ssl.handshake_seq=SSL_HSO_SERVER_CLIENT
        self.ssl.client_keypair_generate = TRUE
        self.ssl.key_generator=X509KeyBridge( \
            key_file="/etc/key.d/Keybridging_cert/key.pem", \
            key_passphrase="", cache_directory="/var/lib/zorp/sslbridge", \
            trusted_ca_files=("/etc/ca.d/certs/0000000070.pem", \
```



```
"/etc/ca.d/keys/0000000070.pem"), \  
untrusted_ca_files=("/etc/ca.d/certs/0000000069.pem", \  
"/etc/ca.d/keys/0000000069.pem")
```

Step 3. Create a service and a rule using the modified proxy class. Use the previously defined proxy class in your Service definition, set up service and access control properties as usual.

Step 4. Restart Zorp.

**Expected result:**

Every time the client connects to a previously unknown host, a new certificate will be generated, signed by one of the CAs specified above. This new certificate will be stored under `/var/lib/zorp/sslbridge` under a filename based on the original server certificate. It will also be shown to the client as the server certificate, and assuming the signer CA is trusted, the client (browser or other application) will not warn about untrusted certificates in any way. If the certificate is signed by the CA for untrusted certificates, the application will not recognize the issuer CA (since its certificate has not been imported to the client) and give a warning to the user. The user can then decide whether the certificate can be accepted or not.

(Actually, two files are stored on the firewall for each certificate: the original certificate received from the server, and the generated certificate. When a client connects to the server, the certificate provided by the server is compared to the stored one: if the two does not match, a new certificate is generated. This happens for example if the server certificate has been expired and refreshed.)

### 3.3. Related standards

- The SSL protocol is defined by Netscape Ltd. at <http://wp.netscape.com/eng/ssl3/ssl-toc.html>
- The TLS protocol is defined in RFC 2246.

### 3.4. SSL options reference

The SSL options are described in detail in the documentation of the Proxy class. See *Section 5.10.4, Class Proxy (p. 181)*.



## Chapter 4. Proxies

This chapter contains reference information on all the available Zorp proxies.

### 4.1. General information on the proxy modules

The sections discussing the available proxies are organized as follows. Overall introduction is followed by proxy class descriptions. Each module has an abstract class which is an interface between the policy and the proxy itself. Abstract classes are the point where the low-level attributes implemented by the proxy appear.

Each Python module contains an abstract proxy class (e.g., `AbstractFtpProxy`) and one or more preconfigured proxy classes derived from the abstract class (e.g., `FtpProxy`, `FtpProxyRO`, etc.). These abstract proxies are very low level classes which always require customization to operate at all, thus they are not directly usable. The preconfigured classes customize the base abstract proxy to perform actually useful functionality. These derived classes inherit all their attributes from the class they were derived from, but have some of their parameters set to default values. Consequently, they can be used for certain tasks without any (or only minimal) modification. Most default classes were derived directly from the abstract classes, but it is possible to derive a class from another derived class. In this case this new class inherits the attributes from its parent class and the abstract class as well. Abstract classes should not be used directly for configuring services in Zorp, always derive an own class and modify its attributes to suit the requirements.

### 4.2. Attribute values

The description of each abstract class includes a detailed list and definition of the attributes of the proxy class. The type and default value of the attribute is also provided. Most types of the attributes (e.g., integer, string, boolean, etc.) are self-explanatory; more complicated attributes (listed as complex type) are explained in their respective description or in the general proxy behavior section of the module.

Proxy attributes can be available and modified during configuration time, run time, or both. Configuration time attributes are set and modified when the proxy is configured, that is, when the session starts. Run time attributes are available when the connection is active, for example, information about the HTTP header being processed is available only when the header is processed. Access to the attributes is indicated in the header of the description of the attribute in the following format: *availability during configuration time : availability during run time*. The type of availability can be read (*r*) access, write (*w*) access, both, or not available (*n/a*). An attribute that is available for reading and writing during both configuration and run time is indicated as *rw:rw*, an attribute that is available only for reading during run time is indicated as *n/a:r*.

**Note**

Unless noted otherwise, default values related to lengths (e.g., line length, etc.) are in bytes.

Timeout values are always given in milliseconds. Setting a timeout to `-1` disables the timeout (i.e. it becomes unlimited).

The description of every proxy class includes a list or textual description of the attributes modified relative to their parent class. The values of the other attributes are inherited from the parent class.

### 4.3. Examples

A number of Python code samples is provided for the proxies to illustrate both their general operation and their capabilities. Most of the proxy configurations shown in the examples can be easily reproduced using the graphical interface. However, some of them utilize the advanced flexibility of Zorp and therefore require the use of configuration scripts written in Python. From these can be implemented, maintained and edited using the Class editor. (The Class editor is available under the **Proxies** tab of the **Zorp** component. When creating a new class, click on the **Class editor** button under the list of available classes.)

### 4.4. Module AnyPy

This module defines an interface to the AnyPy proxy implementation. AnyPy is basically a Python proxy which means that the proxy behaviour is defined in Python by the administrator.

#### 4.4.1. Related standards

#### 4.4.2. Classes in the AnyPy module

Class	Description
<i>AbstractAnyPyProxy</i>	Class encapsulating an AnyPy proxy.
<i>AnyPyProxy</i>	Class encapsulating the default AnyPy proxy.

Table 4.1. Classes of the AnyPy module

#### 4.4.3. Class AbstractAnyPyProxy

This class encapsulates AnyPy, a proxy module calling a Python function to do all of its work. It can be used for defining proxies for protocols not directly supported.



##### Warning

This proxy class is a basis for creating a custom proxy, and cannot be used on its own. Create a new proxy class using the AnyPyProxy as its parent, and implement the proxyThread method to handle the traffic.

Your code will be running as the proxy to transmit protocol elements. When writing your code, take care and be security conscious: do not make security vulnerabilities.

##### 4.4.3.1. Attributes of AbstractAnyPyProxy

<b>client_max_line_length (integer)</b>
Default: 4096
Size of the line buffer in the client stream in bytes. Default value: 4096
<b>server_max_line_length (integer)</b>
Default: 4096

<b>server_max_line_length (integer)</b>
Size of the line buffer in the server stream in bytes. Default value: 4096

#### 4.4.3.2. AbstractAnyPyProxy methods

Method	Description
<code>__init__(self, session)</code>	Constructor to initialize an AnyPy instance.
<code>proxyThread(self)</code>	Function called by the low-level proxy core to transfer requests.

Table 4.2. Method summary

#### Method `__init__(self, session)`

This constructor initializes a new AnyPy instance based on its arguments, and calls the inherited constructor.

#### Arguments of `__init__`

<b>session (unknown)</b>
Default: n/a
The session to be inspected with the proxy instance.

#### Method `proxyThread(self)`

This function is called by the proxy module to transfer requests. It can use the 'self.session.client\_stream' and 'self.session.server\_stream' streams to read data from and write data to.

### 4.4.4. Class AnyPyProxy

This class encapsulates AnyPy, a proxy module calling a Python function to do all of its work. It can be used for defining proxies for protocols not directly supported.

#### 4.4.4.1. Note

This proxy class can only be used as a basis for creating a custom proxy and cannot be used on its own. Please create a new proxy class with the AnyPyProxy as its parent and implement the proxyThread method for handling traffic.

Your code will be running as the proxy to transmit protocol elements, you'll have to take care and be security conscious not to make security vulnerabilities.

## 4.5. Module Finger

The Finger module defines the classes constituting the proxy for the Finger protocol.

### 4.5.1. The Finger protocol

Finger is a request/response based User Information Protocol using port TCP/79. The client opens a connection to the remote machine to initiate a request. The client sends a one line query based on the Finger query specification and waits for the answer. A remote user information program (RUIP) processes the query, returns the result and closes the connection. The response is a series of lines consisting of printable ASCII closed carriage return-line feed (CRLF, ASCII13, ASCII10). After receiving the answer the client closes the connection as well.

The following queries can be used:

- `<CRLF>` This is a simple query listing all users logged in to the remote machine.
- `USERNAME<CRLF>` A query to request all available information about the user USERNAME.
- `USERNAME@HOST1<CRLF>` Request the RUIP to forward the query to HOST1. The response to this query is all information about the user USERNAME available at the remote computer HOST1.
- `USERNAME@HOST1@HOST2<CRLF>` Request HOST1 to forward the query to HOST2. The response to this query is all information about the user USERNAME available at the remote computer HOST2.

### 4.5.2. Proxy behavior

Finger is a module built for parsing messages of the Finger protocol. It reads the QUERY at the client side, parses it and - if the local security policy permits - sends it to the server. When the RESPONSE arrives it processes the RESPONSE and sends it back to the client. It is possible to prepend and/or append a string to the response. Requests can also be manipulated in various ways using the *fingerRequest* function, which is called by the proxy if it is defined.

Length of the username, the line and the hostname can be limited by setting various attributes. Finger proxy also has the capability of limiting the number of hosts in a request, e.g.: `finger user@domain@server` normally results in fingering 'user@domain' performed by the host 'server'. By default, the proxy removes everything after and including the first '@' character. This behavior can be modified by setting the `max_hop_count` attribute to a non-zero value.



#### Example 4.1. Controlling the number of max hops

```
def MyFingerProxy(FingerProxy):
    def config(self):
        FingerProxy.config(self)
        self.max_hop_count = 2
        self.timeout = 30
```

### 4.5.3. Related standards

- The Finger User Information Protocol is described in RFC 1288.

#### 4.5.4. Classes in the Finger module

Class	Description
<i>AbstractFingerProxy</i>	Class encapsulating the abstract Finger proxy.
<i>FingerProxy</i>	Class encapsulating the default Finger proxy.

Table 4.3. Classes of the Finger module

#### 4.5.5. Class AbstractFingerProxy

This proxy implements the Finger protocol as specified in RFC 1288.

##### 4.5.5.1. Attributes of AbstractFingerProxy

<b>max_hop_count (integer, rw:r)</b>
Default: 0
Maximum number of '@' characters in the request. Any text after the last allowed '@' character is stripped from the request.

<b>max_hostname_length (integer, rw:r)</b>
Default: 30
Maximum number of characters in a single name of the hostname chain.

<b>max_line_length (integer, rw:r)</b>
Default: 132
Maximum number of characters in a single line in requests and responses.

<b>max_username_length (integer, rw:r)</b>
Default: 8
Maximum length of the username in a request.

<b>request_detailed (integer, n/a:rw)</b>
Default: n/a
Indicates if multi-line formatting request (/W prefix) was sent by the client (-l parameter). Request for multi-line formatting can be added/removed by the proxy during the <i>fingerRequest</i> event.

<b>request_hostnames (string, n/a:rw)</b>
Default: n/a
The hostname chain. The hostname chain can be modified by the proxy during the <i>fingerRequest</i> event.

<b>request_username (string, n/a:rw)</b>
Default: n/a
The username to be queried. The username can be modified by the proxy during the <i>fingerRequest</i> event.
<b>response_footer (string, rw:rw)</b>
Default:
String to be appended by the proxy to each finger response.
<b>response_header (string, n/a:rw)</b>
Default: ""
String to be prepended by the proxy to each finger response.
<b>strict_username_check (boolean, rw:r)</b>
Default: TRUE
If enabled (TRUE), only requests for usernames containing alphanumeric characters and underscore [a-zA-Z0-9_] are allowed.
<b>timeout (integer, rw:r)</b>
Default: n/a
Timeout value for the request in milliseconds.

#### 4.5.5.2. AbstractFingerProxy methods

Method	Description
<i>fingerRequest(self, username, hostname)</i>	Function processing finger requests.

Table 4.4. Method summary

#### Method `fingerRequest(self, username, hostname)`

This function is called by the Finger proxy to process requests. It can also modify request-specific attributes.

#### Arguments of `fingerRequest`

<b>hostname (unknown, n/a:n/a)</b>
Default: n/a
Destination hosts of the finger request.



<b>username (unknown, n/a:n/a)</b>
Default: n/a
Username to be fingered.

#### 4.5.6. Class FingerProxy

Simple FingerProxy based on AbstractFingerProxy.

### 4.6. Module Ftp

The Ftp module defines the classes constituting the proxy for the File Transfer Protocol (FTP).

#### 4.6.1. The FTP protocol

File Transfer Protocol (FTP) is a protocol to transport files via a reliable TCP connection between a client and a server. FTP uses two reliable TCP connections to transfer files: a simple TCP connection (usually referred to as the Control Channel) to transfer control information and a secondary TCP connection (usually referred to as the Data Channel) to perform the data transfer. It uses a command/response based approach, i.e. the client issues a command and the server responds with a 3-digit status code and associated status information in text format. The Data Channel can be initiated either from the client or the server; the Control Channel is always started from the client.

The client is required to authenticate itself before other commands can be issued. This is performed using the USER and PASS commands specifying username and password, respectively.

##### 4.6.1.1. Protocol elements

The basic protocol is as follows: the client issues a request (also called command in FTP terminology) and the server responds with the result. Both commands and responses are line based: commands are sent as complete lines starting with a keyword identifying the operation to be performed. A response spans one or more lines, each specifying the same 3-digit status code and possible explanation.

##### 4.6.1.2. Data transfer

Certain commands (for example RETR, STOR or LIST) also have a data attachment which is transferred to the peer. Data attachments are transferred in a separate TCP connection. This connection is established on-demand on a random, unprivileged port when a data transfer command is issued.

Endpoint information of this data channel is exchanged via the PASV and PORT commands, or their newer equivalents (EPSV and EPRT).

The data connection can either be initiated by the client (passive mode) or the server (active mode). In passive mode (PASV or EPSV command) the server opens a listening socket and sends back the endpoint information in the PASV response. In active mode (PORT or EPRT command) the client opens a listening socket and sends its endpoint information as the argument of the PORT command. The source port of the server is usually either 20, or the port number of the Command Channel minus one.

**Example 4.2. FTP protocol sample**

```

220 FTP server ready
USER account
331 Password required.
PASS password
230 User logged in.
SYST
215 UNIX Type: L8
PASV
227 Entering passive mode (192,168,1,1,4,0)
LIST
150 Opening ASCII mode data connection for file list
226-Transferring data in separate connection complete.
226 Quotas off
QUIT
221 Goodbye

```

## 4.6.2. Proxy behavior

FtpProxy is a module built for parsing commands of the Control Channel in the FTP protocol. It reads the REQUEST at the client side, parses it and - if the local security policy permits - sends it to the server. The proxy parses the arriving RESPONSES and sends them to the client if the policy permits that. FtpProxy uses a PlugProxy to transfer the data arriving in the Data Channel. The proxy is capable of manipulating commands and stacking further proxies into the Data Channel. Both transparent and non-transparent modes are supported.

The default low-level proxy implementation (AbstractFtpProxy) denies all requests by default. Different commands and/or responses can be enabled by using one of the several predefined proxy classes which are suitable for most tasks. Alternatively, use of the commands can be permitted individually using different attributes. This is detailed in the following two sections.

### 4.6.2.1. Configuring policies for FTP commands and responses

Changing the default behavior of commands can be done by using the hash attribute *request*, indexed by the command name (e.g.: USER or PWD). There is a similar attribute for responses called *response*, indexed by the command name and the response code. The possible values of these hashes are shown in the tables below. See *Section 2.1, Policies for requests and responses (p. 4)* for details. When looking up entries of the *response* attribute hash, the lookup precedence described in *Section 2.1.2, Response codes (p. 6)* is used.

Action	Description
FTP_REQ_ACCEPT	Allow the request to pass.
FTP_REQ_REJECT	Reject the request with the error message specified in the second optional parameter.
FTP_REQ_ABORT	Terminate the connection.

Table 4.5. Action codes for commands in FTP

Action	Description
FTP_RSP_ACCEPT	Allow the response to pass.
FTP_RSP_REJECT	Modify the response to a general failure with error message specified in the optional second parameter.

Action	Description
FTP_RSP_ABORT	Terminate the connection.

Table 4.6. Action codes for responses in FTP

**Example 4.3. Customizing FTP to allow only anonymous sessions**

This example calls a function called `pUser` (defined in the example) whenever a `USER` command is received. All other commands are accepted. The parameter of the `USER` command (i.e. the username) is examined: if it is 'anonymous' or 'Anonymous', the connection is accepted, otherwise it is rejected.

```
class AnonFtp(FtpProxy):
    def config(self):
        self.request["USER"] = (FTP_REQ_POLICY, self.pUser)
        self.request["*"] = (FTP_REQ_ACCEPT)

    def pUser(self, command):
        if self.request_parameter == "anonymous" or self.request_parameter == "Anonymous":
            return FTP_REQ_ACCEPT
        return FTP_REQ_REJECT
```

**4.6.2.2. Configuring policies for FTP features and FTPS support**

FTP servers send the list of supported features to the clients. For example, ProFTPD supports the following features: *LANG en*, *MDTM*, *UTF8*, *AUTH TLS*, *PBSZ*, *PROT*, *REST STREAM*, *SIZE*. The default behavior of FTP features can be changed using the hash attribute *features*, indexed by the name of the feature (e.g.: *UTF8* or *AUTH TLS*). The possible actions are shown in the table below. See *Section 2.1, Policies for requests and responses (p. 4)* for details.

The built-in FTP proxies permit the use of every feature by default.

Action	Description
FTP_FEATURE_ACCEPT	Forward the availability of the feature from the server to the client.
FTP_FEATURE_DROP	Remove the feature from the feature list sent by the server.
FTP_FEATURE_INSERT	Add the feature into the list of available features.

Table 4.7. Policy about enabling FTP features.

**Enabling FTPS connections**

For FTPS connections to operate correctly, the FTP server and client applications must comply to the *FTP Security Extensions (RFC 2228)* and *Securing FTP with TLS (RFC 4217)* RFCs.

For FTPS connections, the *AUTH TLS*, *PBSZ*, *PROT* features must be accepted. Also, *STARTTLS* support must be properly configured. See *Section 3.2, Configuring TLS and SSL encrypted connections (p. 10)* for details.

If the proxy is configured to disable encryption between Zorp and the client, the proxy automatically removes the *AUTH TLS*, *PBSZ*, *PROT* features from the list sent by the server.



If `STARTTLS` connections are accepted on the client side (`self.ssl.client_security=SSL_ACCEPT_STARTTLS`), but TLS-forwarding is disabled on the server side, the proxy automatically inserts the `AUTH TLS`, `PBSZ`, `PROT` features into the list sent by the server. These features are inserted even if encryption is explicitly disabled on the server side or the server does not support the `FEAT` command, making one-sided `STARTTLS` support feasible.

**Warning**

When using *inband routing* with the FTPS protocol, the server's certificate is compared to its hostname. The `subject_alt_name` parameter (or the `Common Name` parameter if the `subject_alt_name` parameter is empty) of the server's certificate must contain the hostname or the IP address (as resolved from the `Zorp` host) of the server (e.g., `ftp.example.com`).

Alternatively, the `Common Name` or the `subject_alt_name` parameter can contain a generic hostname, e.g., `*.example.com`.

Note that if the `Common Name` of the certificate contains a generic hostname, do not specify a specific hostname or an IP address in the `subject_alt_name` parameter.

**Note**

- The FTP proxy does not support the following FTPS-related commands: `REIN`, `CCC`, `CDC`.
- `STARTTLS` is supported in nontransparent scenarios as well.

**Example 4.4. Configuring FTPS support**

This example is a standard `FtpProxy` with FTPS support enabled.

```
class FtpsProxy(FtpProxy):
    def config(self):
        FtpProxy.config(self)
        self.max_password_length=64

        EncryptionPolicy(name="ForwardSTARTTLS",
encryption=ForwardStartTLSEncryption(client_verify=ClientCertificateVerifier(),
client_ssl_options=ClientSSLOptions(), server_verify=ServerCertificateVerifier(),
server_ssl_options=ServerSSLOptions(),
client_certificate_generator=DynamicCertificate(private_key=PrivateKey.fromFile(key_file_path="/etc/key.d/ZMS_Engine/key.pem"),
trusted_ca=Certificate.fromFile(certificate_file_path="/etc/ca.d/certs/my-trusted-ca-cert.pem",
private_key=PrivateKey.fromFile("/etc/ca.d/keys/my-trusted-ca-cert.pem")),
untrusted_ca=Certificate.fromFile(certificate_file_path="/etc/ca.d/certs/my-untrusted-ca-cert.pem",
private_key=PrivateKey.fromFile("/etc/ca.d/keys/my-untrusted-ca-cert.pem"))))

    def demo() :
        Service(name='demo/MyFTPSService', router=TransparentRouter(), chainer=ConnectChainer(),
proxy_class=FtpsProxy, max_instances=0, max_sessions=0, keepalive=Z_KEEPALIVE_NONE,
encryption_policy="ForwardSTARTTLS")

        Rule(rule_id=2,
proto=6,
service='demo/MyFTPSService'
)
```

**4.6.2.3. Stacking**

The available stacking modes for this proxy module are listed in the following table. For additional information on stacking, see *Section 2.3.1, Proxy stacking (p. 7)*.

Action	Description
FTP_STK_DATA	Pass the data to the stacked proxy or program.

Action	Description
FTP_STK_NONE	No proxy stacked.

Table 4.8. Stacking policy.

#### 4.6.2.4. Configuring inband authentication

The Ftp proxy supports *inband authentication* as well to use the built-in authentication method of the FTP and FTPS protocols to authenticate the client. The authentication itself is performed by the backend configured for the service.

If the client uses different usernames on and the remote server (e.g., he uses his own username to authenticate to , but anonymous on the target FTP server), the client must specify the usernames and passwords in the following format:

Username:

```
<ftp user>@<proxy user>@<remote site>[:<port>]
```

Password:

```
<ftp password>@<proxy password>
```

Alternatively, all the above information can be specified as the username:

```
<ftp user>@<proxy user>@<remote site>[:<port>]:<ftp password>@<proxy password>
```



##### Warning

When using *inband routing* with the FTPS protocol, the server's certificate is compared to its hostname. The `subject_alt_name` parameter (or the Common Name parameter if the `subject_alt_name` parameter is empty) of the server's certificate must contain the hostname or the IP address (as resolved from the Zorp host) of the server (e.g., `ftp.example.com`).

Alternatively, the Common Name or the `subject_alt_name` parameter can contain a generic hostname, e.g., `*.example.com`.

Note that if the Common Name of the certificate contains a generic hostname, do not specify a specific hostname or an IP address in the `subject_alt_name` parameter.

#### 4.6.3. Related standards

- The File Transfer Protocol is described in RFC 959.
- FTP Security Extensions including the FTPS protocol and securing FTP with TLS are described in RFC 2228 and RFC 4217.

#### 4.6.4. Classes in the Ftp module

Class	Description
<i>AbstractFtpProxy</i>	Class encapsulating the abstract FTP proxy.
<i>FtpProxy</i>	Default Ftp proxy based on AbstractFtpProxy.



Class	Description
<u>FtpProxyAnonRO</u>	FTP proxy based on AbstractFtpProxy, only allowing read-only access to anonymous users.
<u>FtpProxyAnonRW</u>	FTP proxy based on AbstractFtpProxy, allowing full read-write access, but only to anonymous users.
<u>FtpProxyRO</u>	FTP proxy based on AbstractFtpProxy, allowing read-only access to any user.
<u>FtpProxyRW</u>	FTP proxy based on AbstractFtpProxy, allowing full read-write access to any user.

Table 4.9. Classes of the Ftp module

#### 4.6.5. Class AbstractFtpProxy

This proxy implements the FTP protocol as specified in RFC 959. All traffic and commands are denied by default. Consequently, either customized Ftp proxy classes derived from the abstract class should be used, or one of the predefined classes (e.g.: FtpProxy, FtpProxyRO, etc.).

##### 4.6.5.1. Attributes of AbstractFtpProxy

<b>active_connection_mode (enum, rw:r)</b>
Default: FTP_ACTIVE_MINUSONE
In active mode the server connects the client. By default this must be from Command Channel port minus one (FTP_ACTIVE_MINUSONE). Alternatively, connection can also be performed either from port number 20 (FTP_ACTIVE_TWENTY) or from a random port (FTP_ACTIVE_RANDOM).
<b>auth_tls_ok_client (boolean, n/a:r)</b>
Default: ""
Shows whether the client-side authentication was performed over a secure channel.
<b>auth_tls_ok_server (boolean, n/a:r)</b>
Default: ""
Shows whether the server-side authentication was performed over a secure channel.
<b>buffer_size (integer, rw:r)</b>
Default: 4096
Buffer size for data transfer in bytes.
<b>data_mode (enum, rw:r)</b>
Default: FTP_DATA_KEEP

**data\_mode (enum, rw:r)**

The type of the FTP connection on the server side can be manipulated: leave it as the client requested (FTP\_DATA\_KEEP), or force passive (FTP\_DATA\_PASSIVE) or active (FTP\_DATA\_ACTIVE) connection.

**data\_port\_max (integer, rw:r)**

Default: 41000

On the proxy side, ports equal to or below the value of *data\_port\_max* can be allocated as the data channel.

**data\_port\_min (integer, rw:r)**

Default: 40000

On the proxy side, ports equal to or above the value of *data\_port\_min* can be allocated as the data channel.

**data\_protection\_enabled\_client (boolean, n/a:r)**

Default: ""

Shows whether the data channel is encrypted or not on the client-side.

**data\_protection\_enabled\_server (boolean, n/a:r)**

Default: ""

Shows whether the data channel is encrypted or not on the server-side.

**features (complex, rw:rw)**

Default:

Hash containing the filtering policy for FTP features.

**hostname (string, n/a:rw)**

Default:

The hostname of the FTP server to connect to, when inband routing is used.

**hostport (integer, n/a:rw)**

Default:

The port of the FTP server to connect to, when inband routing is used.

**masq\_address\_client (string, rw:r)**

Default: ""



<b>masq_address_client (string, rw:r)</b>
IP address of the firewall appearing on the client side. If its value is set, this IP is sent regardless of its true IP (where it is binded). This attribute may be used when network address translation is performed before Zorp.
<b>masq_address_server (string, rw:r)</b>
Default: ""
IP address of the firewall appearing on the server side. If its value is set, this IP is sent regardless of its true IP (where it is binded). This attribute may be used when network address translation is performed before Zorp.
<b>max_continuous_line (integer, rw:r)</b>
Default: 100
Maximum number of answer lines for a command.
<b>max_hostname_length (integer, rw:r)</b>
Default: 128
Maximum length of hostname. Used only in non-transparent mode.
<b>max_line_length (integer, rw:r)</b>
Default: 255
Maximum length of a line that the proxy is allowed to transfer. Requests/responses exceeding this limit are dropped.
<b>max_password_length (integer, rw:r)</b>
Default: 64
Maximum length of the password.
<b>max_username_length (integer, rw:r)</b>
Default: 32
Maximum length of the username.
<b>password (string, n/a:rw)</b>
Default:
The password to be sent to the server.



<b>permit_client_bounce_attack (boolean, rw:rw)</b>
Default: FALSE
If enabled the IP addresses of data channels will not need to match with the IP address of the control channel, permitting the use of FXP while increasing the security risks.
<b>permit_empty_command (boolean, rw:r)</b>
Default: TRUE
Enable transmission of lines without commands.
<b>permit_server_bounce_attack (boolean, rw:rw)</b>
Default: FALSE
If enabled the IP addresses of data channels will not need to match with the IP address of the control channel, permitting the use of FXP while increasing the security risks.
<b>permit_unknown_command (boolean, rw:r)</b>
Default: FALSE
Enable the transmission of unknown commands.
<b>proxy_password (string, n/a:rw)</b>
Default:
The password to be used for proxy authentication given by the user, when inband authentication is used.
<b>proxy_username (string, n/a:rw)</b>
Default:
The username to be used for proxy authentication given by the user, when inband authentication is used.
<b>request (complex, rw:rw)</b>
Default:
Normative policy hash for FTP requests indexed by command name (e.g.: "USER", "PWD" etc.). See also <i>Section 2.1, Policies for requests and responses (p. 4)</i> .
<b>request_command (string, n/a:rw)</b>
Default: n/a
When a request is evaluated on the policy level, this variable contains the requested command.



<b>request_parameter (string, n/a:rw)</b>
Default: n/a
When a request is evaluated on the policy level, this variable contains the parameters of the requested command.
<b>request_stack (complex, rw:rw)</b>
Default:
Hash containing the stacking policy for the FTP commands. The hash is indexed by the FTP command (e.g. RETR, STOR). See also <i>Section 2.3.1, Proxy stacking (p. 7)</i> .
<b>response (complex, rw:rw)</b>
Default:
Normative policy hash for FTP responses indexed by command name and answer code (e.g.: "USER", "331"; "PWD", "200" etc.). See also <i>Section 2.1, Policies for requests and responses (p. 4)</i> .
<b>response_parameter (string, n/a:rw)</b>
Default:
When a response is evaluated on the policy level, this variable contains answer parameters.
<b>response_status (string, n/a:rw)</b>
Default:
When a response is evaluated on the policy level, this variable contains the answer code.
<b>response_strip_msg (boolean, rw:r)</b>
Default: FALSE
Strip the response message and only send the response code.
<b>strict_port_checking (boolean, rw:rw)</b>
Default: TRUE
If enabled the foreign port is strictly checked: in active mode the server must be connected on port 20, while in any other situation the foreign port must be above 1023.
<b>target_port_range (string, rw:r)</b>
Default: "21"
The port where the client can connect through a non-transparent FtpProxy.

<b>timeout (integer, rw:r)</b>
Default: 300000
General I/O timeout in milliseconds. When there is no specific timeout for a given operation, this value is used.
<b>transparent_mode (boolean, rw:r)</b>
Default: TRUE
Specifies if the proxy works in transparent (TRUE) or non-transparent (FALSE) mode.
<b>username (string, n/a:rw)</b>
Default:
The username authenticated to the server.
<b>valid_chars_username (string, rw:r)</b>
Default: "a-zA-Z0-9._@"
List of the characters accepted in usernames.

#### 4.6.6. Class FtpProxy

A permitting Ftp proxy based on the AbstractFtpProxy, allowing all commands, responses, and features, including unknown ones. The connection is terminated if a response with the answer code 421 is received.

#### 4.6.7. Class FtpProxyAnonRO

FTP proxy based on AbstractFtpProxy, enabling read-only access (i.e. only downloading) to anonymous users (uploads and usernames other than 'anonymous' or 'ftp' are disabled). Commands and return codes are strictly checked, unknown commands and responses are rejected. Every feature is accepted.

The ABOR; ACCT; AUTH; CDUP; CWD; EPRT; EPSV; FEAT; LIST; MODE; MDTM; NLST; NOOP; OPTS; PASV; PASS; PORT; PWD; QUIT; REST; RETR; SIZE; STAT; STRU; SYST; TYPE; and USER commands are permitted, the CLNT; XPWD; MACB commands are rejected.

#### 4.6.8. Class FtpProxyAnonRW

FTP proxy based on AbstractFtpProxy, enabling full read-write access to anonymous users (the 'anonymous' and 'ftp' usernames are permitted). Commands and return codes are strictly checked, unknown commands and responses are rejected. Every feature is accepted.

The ABOR; ACCT; APPE; CDUP; CWD; DELE; EPRT; EPSV; LIST; MKD; MODE; MDTM; NLST; NOOP; OPTS; PASV; PASS; PORT; PWD; QUIT; RMD; RNFR; RNTD; REST; RETR; SIZE; STAT; STOR; STOU; STRU; SYST; TYPE; USER and FEAT commands are permitted, the AUTH; CLNT; XPWD; MACB commands are rejected.

#### 4.6.9. Class FtpProxyRO

FTP proxy based on AbstractFtpProxy, enabling read-only access to any user. Commands and return codes are strictly checked, unknown commands and responses are rejected. Every feature is accepted.

The ABOR; ACCT; AUTH; CDUP; CWD; EPRT; EPSV; FEAT; LIST; MODE; MDTM; NLST; NOOP; OPTS; PASV; PASS; PORT; PWD; QUIT; REST; RETR; SIZE; STAT; STRU; SYST; TYPE; and USER commands are permitted, the CLNT; XPWD; MACB commands are rejected.

#### 4.6.10. Class FtpProxyRW

FTP proxy based on AbstractFtpProxy, enabling full read-write access to any user. Commands and return codes are strictly checked, unknown commands and responses are rejected. Every feature is accepted.

The ABOR; ACCT; AUTH; CDUP; CWD; EPRT; EPSV; FEAT; LIST; MODE; MDTM; NLST; NOOP; OPTS; PASV; PASS; PORT; PWD; QUIT; REST; RETR; SIZE; STAT; STRU; SYST; TYPE; and USER commands are permitted, the CLNT; XPWD; MACB commands are rejected.

### 4.7. Module Http

The Http module defines the classes constituting the proxy for the HyperText Transfer Protocol (HTTP). HTTP is the protocol the Web is based on, therefore it is the most frequently used protocol on the Internet. It is used to access different kinds of content from the Web. The type of content retrieved via HTTP is not restricted, it can range from simple text files to hypertext files and multimedia formats like pictures, videos or audio files.

#### 4.7.1. The HTTP protocol

HTTP is an open application layer protocol for hypermedia information systems. It basically allows an open-ended set of methods to be applied to resources identified by Uniform Resource Identifiers (URIs).

##### 4.7.1.1. Protocol elements

HTTP is a text based protocol where a client sends a request comprising of a METHOD, an URI and associated meta information represented as MIME-like headers, and possibly a data attachment. The server responds with a status code, a set of headers, and possibly a data attachment. Earlier protocol versions perform a single transaction in a single TCP connection, HTTP/1.1 introduces persistency where a single TCP connection can be reused to perform multiple transactions.

An HTTP method is a single word - usually spelled in capitals - instructing the server to apply a function to the resource specified by the URI. Commonly used HTTP methods are "GET", "POST" and "HEAD". HTTP method names are not restricted in any way, other HTTP based protocols (such as WebDAV) add new methods to the protocol while keeping the general syntax intact.

Headers are part of both the requests and the responses. Each header consists of a name followed by a colon (':') and a field value. These headers are used to specify content-specific and protocol control information.

The response to an HTTP request starts with an HTTP status line informing the client about the result of the operation and an associated message. The result is represented by three decimal digits, the possible values are defined in the HTTP RFCs.

#### 4.7.1.2. Protocol versions

The protocol has three variants, differentiated by their version number. Version 0.9 is a very simple protocol which allows a simple octet-stream to be transferred without any meta information (e.g.: no headers are associated with requests or responses).

Version 1.0 introduces MIME-like headers in both requests and responses; headers are used to control both the protocol (e.g.: the "Connection" header) and to give information about the content being transferred (e.g.: the "Content-Type" header). This version has also introduced the concept of name-based virtual hosts.

Building on the success of HTTP/1.0, version 1.1 of the protocol adds persistent connections (also referred to as "connection keep-alive") and improved proxy control.

#### 4.7.1.3. Bulk transfer

Both requests and responses might have an associated data blob, also called an entity in HTTP terminology. The size of the entity is determined using one of three different methods:

1. The complete size of the entity is sent as a header (the Content-Length header).
2. The transport layer connection is terminated when transfer of the blob is completed (used by HTTP/0.9 and might be used in HTTP/1.1 in non-persistent mode).
3. Instead of specifying the complete length, smaller chunks of the complete blob are transferred, and each chunk is prefixed with the size of that specific chunk. The end of the stream is denoted by a zero-length chunk. This mode is also called chunked encoding and is specified by the Transfer-Encoding header.



##### Example 4.5. Example HTTP transaction

```
GET /index.html HTTP/1.1
Host: www.example.com
Connection: keep-alive
User-Agent: My-Browser-Type 6.0

HTTP/1.1 200 OK
Connection: close
Content-Length: 14

<html>
</html>
```

#### 4.7.2. Proxy behavior

The default low-level proxy implementation (*AbstractHttpProxy*) denies all requests by default. Different requests and/or responses can be enabled by using one of the several predefined proxy classes which are suitable for most tasks. Alternatively, a custom proxy class can be derived from *AbstractHttpProxy* and the requests and responses enabled individually using different attributes.

Several examples and considerations on how to enable virus filtering in the HTTP traffic are discussed in the Technical White Paper and Tutorial *Virus filtering in HTTP*, available at the BalaSys Documentation Page <http://www.balasy.hu/documentation/>.

#### 4.7.2.1. Transparent and non-transparent modes

HttpProxy is able to operate both in transparent and non-transparent mode. In transparent mode, the client does not notice (or even know) that it is communicating through a proxy. The client communicates using normal server-style requests.

In non-transparent mode, the address and the port of the proxy server must be set on the client. In this case the client sends proxy-style requests to the proxy.



##### Example 4.6. Proxy style HTTP query

```
GET http://www.example.com/index.html HTTP/1.1
Host: www.example.com
Connection: keep-alive
User-Agent: My-Browser-Type 6.0

HTTP/1.1 200 OK
Connection: close
Content-Length: 14

<html>
</html>
```

In non-transparent mode it is possible to request the use of the SSL protocol through the proxy, which means the client communicates with the proxy using the HTTP protocol, but the proxy uses HTTPS to communicate with the server. This technique is called data tunneling.



##### Example 4.7. Data tunneling with connect method

```
CONNECT www.example.com:443 HTTP/1.1
Host: www.example.com
User-agent: My-Browser-Type 6.0

HTTP/1.0 200 Connection established
Proxy-agent: My-Proxy/1.1
```

#### 4.7.2.2. Configuring policies for HTTP requests and responses

Changing the default behavior of requests is possible using the *request* attribute. This hash is indexed by the HTTP method names (e.g.: GET or POST). The *response* attribute (indexed by the request method and the response code) enables the control of HTTP responses. The possible actions are described in the following tables. See also *Section 2.1, Policies for requests and responses (p. 4)*. When looking up entries of the *response* attribute hash, the lookup precedence described in *Section 2.1.2, Response codes (p. 6)* is used.

Action	Description
HTTP_REQ_ACCEPT	Allow the request to pass.
HTTP_REQ_REJECT	Reject the request. The reason for the rejection can be specified in the optional second argument.
HTTP_REQ_ABORT	Terminate the connection.

Action	Description
HTTP_REQ_POLICY	Call the function specified to make a decision about the event. The function receives four arguments: self, method, url, version. See <i>Section 2.1, Policies for requests and responses (p. 4)</i> for details.

Table 4.10. Action codes for HTTP requests

Action	Description
HTTP_RSP_ACCEPT	Allow the response to pass.
HTTP_RSP_DENY	Reject the response and return a policy violation page to the client.
HTTP_RSP_REJECT	Reject the response and return a policy violation page to the client, with error information optionally specified as the second argument.
HTTP_RSP_POLICY	Call the function specified to make a decision about the event. The function receives five parameters: self, method, url, version, response. See <i>Section 2.1, Policies for requests and responses (p. 4)</i> for details.

Table 4.11. Action codes for HTTP responses

**Example 4.8. Implementing URL filtering in the HTTP proxy**

This example calls the filterURL function (defined in the example) whenever a HTTP GET request is received. If the requested URL is 'http://www.disallowedsite.com', the request is rejected and an error message is sent to the client.

```
class DmzHTTP(HttpProxy):
    def config(self):
        HttpProxy.config(self)
        self.request["GET"] = (HTTP_REQ_POLICY, self.filterURL)

    def filterURL(self, method, url, version):
        if (url == "http://www.disallowedsite.com"):
            self.error_info = 'Access of this content is denied by the local policy.'
            return HTTP_REQ_REJECT
        return HTTP_REQ_ACCEPT
```

**Example 4.9. 404 response filtering in HTTP**

In this example the 404 response code to GET requests is rejected, and a custom error message is returned to the clients instead.

```
class DmzHTTP(HttpProxy):
    def config(self):
        HttpProxy.config(self)
        self.response["GET", "404"] = (HTTP_RSP_POLICY, self.filter404)

    def filter404(self, method, url, version, response):
        self.error_status = 404
        self.error_info = "Requested page was not accessible."
        return HTTP_RSP_REJECT
```

### 4.7.2.3. Configuring policies for HTTP headers

Both request and response headers can be modified by the proxy during the transfer. New header lines can be inserted, entries can be modified or deleted. To change headers in the requests and responses use the *request\_header* hash or the *response\_header* hash, respectively.

Similarly to the request hash, these hashes are indexed by the header name (like "User-Agent") and contain an actiontuple describing the action to take.

By default, the proxy modifies only the "Host", "Connection", "Proxy-Connection" and "Transfer-Encoding" headers. "Host" headers need to be changed when the proxy modifies the URL; "(Proxy-)Connection" is changed when the proxy turns connection keep-alive on/off; "Transfer-Enconding" is changed to enable chunked encoding.

Action	Description
HTTP_HDR_ABORT	Terminate the connection.
HTTP_HDR_ACCEPT	Accept the header.
HTTP_HDR_DROP	Remove the header.
HTTP_HDR_POLICY	Call the function specified to make a decision about the event. The function receives three parameters: self, hdr_name, and hdr_value.
HTTP_HDR_CHANGE_NAME	Rename the header to the name specified in the second argument.
HTTP_HDR_CHANGE_VALUE	Change the value of the header to the value specified in the second argument.
HTTP_HDR_CHANGE_BOTH	Change both the name and value of the header to the values specified in the second and third arguments, respectively.
HTTP_HDR_INSERT	Insert a new header defined in the second argument.
HTTP_HDR_REPLACE	Remove all existing occurrences of a header and replace them with the one specified in the second argument.

Table 4.12. Action codes for HTTP headers



#### Example 4.10. Header filtering in HTTP

The following example hides the browser used by the client by replacing the value of the User-Agent header to Lynx in all requests. The use of cookies is disabled as well.

```
class MyHttp(HttpProxy):
    def config(self):
        HttpProxy.config(self)
        self.request_header["User-Agent"] = (HTTP_HDR_CHANGE_VALUE, "Lynx 2.4.1")
        self.request_header["Cookie"] = (HTTP_HDR_POLICY, self.processCookies)
        self.response_header["Set-Cookie"] = (HTTP_HDR_DROP, )

    def processCookies(self, name, value):
        # You could change the current header in self.current_header_name
        # or self.current_header_value, the current request url is
```



```
# in self.request_url
return HTTP_HDR_DROP
```

#### 4.7.2.4. Redirecting URLs

URLs or sets of URLs can be easily rejected or redirected to a local mirror by modifying some attributes during request processing.

When an HTTP request is received, normative policy chains are processed (*self.request*, *self.request\_header*). Policy callbacks for certain events can be configured with the HTTP\_REQ\_POLICY or HTTP\_HDR\_POLICY directives. Any of these callbacks may change the *request\_url* attribute, instructing the proxy to fetch a page different from the one specified by the browser. Please note that this is transparent to the user and does not change the URL in the browser.



##### Example 4.11. URL redirection in HTTP proxy

This example redirects all HTTP GET requests to the 'http://www.example.com/' URL by modifying the value of the requested URL.

```
class MyHttp(HttpProxy):
    def config(self):
        HttpProxy.config(self)
        self.request["GET"] = (HTTP_REQ_POLICY, self.filterURL)

    def filterURL(self, method, url, version):
        self.request_url = "http://www.example.com/"
        return HTTP_REQ_ACCEPT
```



##### Example 4.12. Redirecting HTTP to HTTPS

This example redirects all incoming HTTP connections to an HTTPS URL.

```
class HttpProxyHttpsredirect(HttpProxy):
    def config(self):
        HttpProxy.config(self)
        self.error_silent = TRUE
        self.request["GET"] = (HTTP_REQ_POLICY, self.reqRedirect)

    def reqRedirect(self, method, url, version):
        self.error_status = 301
        #self.error_info = 'HTTP/1.0 301 Moved Permanently'
        self.error_headers="Location: https://%s/" % self.request_url_host
        return HTTP_REQ_REJECT
```

#### 4.7.2.5. Request types

Zorp differentiates between two request types: server requests and proxy request.

- Server requests are sent by browsers directly communicating with HTTP servers. These requests include an URL relative to the server root (e.g.: /index.html), and a 'Host' header indicating which virtual server to use.
- Proxy requests are used when the browser communicates with an HTTP proxy. These requests include a fully specified URL (e.g.: http://www.example.com/index.html).

Zorp determines the type of the incoming request from the request URL, even if the Proxy-connection header exist. As there is no clear distinction between the two request types, the type of the request cannot always be accurately detected automatically, though all common cases are covered.

Requests are handled differently in transparent and non-transparent modes.

- A transparent HTTP proxy (*transparent\_mode* attribute is TRUE) is meant to be installed in front of a network where clients do not know about the presence of the firewall. In this case the proxy expects to see server type requests only. If clients communicate with a real HTTP proxy through the firewall, proxy type requests must be explicitly enabled using the *permit\_proxy\_requests* attribute, or transparent mode has to be used.
- The use of non-transparent HTTP proxies (*transparent\_mode* attribute is FALSE) must be configured in web browsers behind the firewall. In this case Zorp expects proxy requests only, and emits server requests (assuming *parent\_proxy* is not set).

#### 4.7.2.6. Using parent proxies

Parent proxies are non-transparent HTTP proxies used behind Zorp. Two things have to be set in order to use parent proxies. First, select a router which makes the proxy connect to the parent proxy, this can be either *InbandRouter()* or *DirectedRouter()*. Second, set the *parent\_proxy* and *parent\_proxy\_port* attributes in the *HttpProxy* class. Setting these attributes results in proxy requests to be emitted to the target server both in transparent and non-transparent mode.

The parent proxy attributes can be set both in the configuration phase (e.g.: *config()* event), or later on a per-request basis. This is possible because the proxy re-connects.



##### Example 4.13. Using parent proxies in HTTP

In this example the *MyHttp* proxy class uses a parent proxy. For this the domain name and address of the parent proxy is specified, and a service using an *InbandRouter* is created.

```
class MyHttp(HttpProxy):
    def config(self):
        HttpProxy.config(self)
        self.parent_proxy = "proxy.example.com"
        self.parent_proxy_port = 3128

    def instance():
        Service("http", MyHttp, router=InbandRouter())
        Listener(SocketAddrInet('10.0.0.1', 80), "http")
```

#### 4.7.2.7. FTP over HTTP

In non-transparent mode it is possible to let Zorp process *ftp://* URLs, effectively translating HTTP requests to FTP requests on the fly. This behaviour can be enabled by setting *permit\_ftp\_over\_http* to TRUE and adding port 21 to *target\_port\_range*. Zorp currently supports passive mode transfers only.

#### 4.7.2.8. Error messages

There are cases when the HTTP proxy must return an error page to the client to indicate certain error conditions. These error messages are stored as files in the directory specified by the *error\_files\_directory* attribute, and can be customized by changing the contents of the files in this directory.

Each file contains plain HTML text, but some special macros are provided to dynamically add information to the error page. The following macros can be used:

- *@INFO@* -- further error information as provided by the proxy

- `@VERSION@` -- Zorp version number
- `@DATE@` -- current date
- `@HOST@` -- hostname of Zorp

It is generally recommended not to display error messages to untrusted clients, as they may leak confidential information. To turn error messages off, set the `error_silent` attribute to TRUE, or strip error files down to a minimum.



**Note**

The language of the messages can be set using the `config.options.language` global option, or individually for every Http proxy using the `language` parameter. See *Appendix B, Global options of Zorp (p. 232)* for details.

#### 4.7.2.9. Stacking

HTTP supports stacking proxies for both request and response entities (e.g.: data bodies). This is controlled by the `request_stack` and `response_stack` attribute hashes. See also *Section 2.3.1, Proxy stacking (p. 7)*.

There are two stacking modes available: `HTTP_STK_DATA` sends only the data portion to the downstream proxy, while `HTTP_STK_MIME` also sends all header information to make it possible to process the data body as a MIME envelope. Please note that while it is possible to change the data part in the stacked proxy, it is not possible to change the MIME headers - they can be modified only by the HTTP proxy. The possible parameters are listed in the following tables.

Action	Description
<code>HTTP_STK_NONE</code>	No additional proxy is stacked into the HTTP proxy.
<code>HTTP_STK_DATA</code>	The data part of the HTTP traffic is passed to the specified stacked proxy.
<code>HTTP_STK_MIME</code>	The data part including header information of the HTTP traffic is passed to the specified stacked proxy.

Table 4.13. Constants for proxy stacking

Please note that stacking is skipped altogether if there is no body in the message.

#### 4.7.2.10. Webservers returning data in 205 responses

Certain webserver applications may return data entities in 205 responses. This is explicitly prohibited by the RFCs, but Zorp permits such responses for interoperability reasons.

#### 4.7.2.11. Session persistence in load balancing

Zorp's HTTP proxy offers the 'session persistence in load balancing' feature, further enhancing Zorp's load balancing capabilities by that.

With the help of this feature, the Round Robin chainer can identify connections by their session IDs and make sure that every connection with the same session ID is always addressed to the same server, so that the session persists.

For using the ‘session persistence in load balancing’ feature, the administrator has to configure the following three attributes for the HTTP proxy:

- *Enable\_session\_persistence*  
You can switch on or off the ‘Session persistence in load balancing’ feature with that parameter.
- *Session\_persistence\_cookie\_name*  
This parameter can only be configured if *enable\_session\_persistence* is set to TRUE. The administrator can provide the name of the cookie here: Zorpx directs all incoming requests to a web server and each web server sends a session ID back to Zorpx. The name of this session ID, that is, the cookie name, can be provided here to ensure that requests with the same session ID are directed to the same web server.
- *Session\_persistence\_cookie\_salt*  
This parameter can only be configured if *enable\_session\_persistence* is set to TRUE. The administrator can provide the salt here, with which the IP address of the web server can be hashed before the session ID. With the help of this hashed information Zorpx can next time identify to which server the next connection attempt of this session has to be directed.

#### 4.7.2.12. URL filtering in HTTP

Starting with version 3.3FR1, Zorpx supports category-based URL filtering using a regularly updated database.

- To configure URL-filtering, see *Section Configuring URL-filtering in HTTP (p. 44)*.
- For the list of categories available by default, see *Section List of URL-filtering categories (p. 45)*.
- To customize or expand the URL-database, see *Section Customizing the URL database (p. 48)*.

#### Configuring URL-filtering in HTTP

The URLs and domains in the database are organized into thematic categories like *adult*, *news*, *jobsearch*, etc.

To enable url-filtering, set the *enable\_url\_filter* and *enable\_url\_filter\_dns* options to *TRUE*. The *enable\_url\_filter\_dns* option is needed only to ensure that a domain or URL is correctly categorized even when it is listed in the database using its domain name, but the client tries to access it with its IP address (or vice-versa).



**Note**

URL-filtering is handled by the Zorpx Http proxy, without the need of using ZCV. The URL-filtering capability of Zorpx is available only after purchasing the *url-filter* license option.

Updates to the URL database are automatically downloaded daily from the BalaSys website using the *zavupdate* utility.

Access to specific categories can be set using the `url_category` option, which is a hash indexed by the name of the category. The following actions are possible:

Action	Description
HTTP_URL_ACCEPT	Permit access to the URL.
HTTP_URL_REJECT	Reject the request. The error code and reason for the rejection are specified in the second and third arguments. See <i>Section Configuring URL-filtering in HTTP</i> (p. 44) for details.
HTTP_URL_REDIRECT	Redirect the connection to the URL specified in the second argument.

Table 4.14. Action codes for URL filtering



#### Example 4.14. URL-filtering example

The following example blocks several categories and accepts the rest. For a complete list of categories, see *Section List of URL-filtering categories* (p. 45).

```
class MyHTTPUrlFilter(HttpProxy):
    def config(self):
        HttpProxy.config(self)
        self.enable_url_filter=TRUE
        self.enable_url_filter_dns=TRUE
        self.url_category['adult']=(HTTP_URL_REJECT, (403, "Adult website",))
        self.url_category['porn']=(HTTP_URL_REJECT, (403, "Porn website",))
        self.url_category['malware']=(HTTP_URL_REJECT, (403, "Site contains malware",))
        self.url_category['phishing']=(HTTP_URL_REJECT, (403, "Phishing site",))
        self.url_category['warez']=(HTTP_URL_REJECT, (403, "Warez site",))
        self.url_category['*']=(HTTP_URL_ACCEPT,)
```

The following example redirects access to online gaming sites to a dummy website.

```
class MyHTTPUrlFilter(HttpProxy):
    def config(self):
        HttpProxy.config(self)
        self.enable_url_filter=TRUE
        self.enable_url_filter_dns=TRUE
        self.url_category['onlinergames']=(HTTP_URL_REDIRECT, "http://example.com")
        self.url_category['*']=(HTTP_URL_ACCEPT,)
```

## List of URL-filtering categories

The Zorp URL database contains the following thematic categories by default.

- *abortion*: Abortion information excluding when related to religion
- *ads*: Advert servers and banned URLs
- *adult*: Sites containing adult material such as swearing but not porn
- *aggressive*: Similar to violence but more promoting than depicting
- *antispyware*: Sites that remove spyware
- *artnudes*: Art sites containing artistic nudity



- *astrology*: Astrology websites
- *audio-video*: Sites with audio or video downloads
- *banking*: Banking websites
- *beerliquorinfo*: Sites with information only on beer or liquors
- *beerliquorsale*: Sites with beer or liquors for sale
- *blog*: Journal/Diary websites
- *cellphones*: stuff for mobile/cell phones
- *chat*: Sites with chat rooms etc
- *childcare*: Sites to do with childcare
- *cleaning*: Sites to do with cleaning
- *clothing*: Sites about and selling clothing
- *contraception*: Information about contraception
- *culinary*: Sites about cooking et al
- *dating*: Sites about dating
- *desktopsillies*: Sites containing screen savers, backgrounds, cursers, pointers, desktop themes and similar timewasting and potentially dangerous content
- *dialers*: Sites with dialers such as those for pornography or trojans
- *drugs*: Drug related sites
- *ecommerce*: Sites that provide online shopping
- *entertainment*: Sites that promote movies, books, magazine, humor
- *filehosting*: Sites to do with filehosting
- *frencheducation*: Sites to do with french education
- *gambling*: Gambling sites including stocks and shares
- *games*: Game related sites
- *gardening*: Gardening sites
- *government*: Military and schools etc
- *guns*: Sites with guns
- *hacking*: Hacking/cracking information
- *homerepair*: Sites about home repair
- *hygiene*: Sites about hygiene and other personal grooming related stuff
- *instantmessaging*: Sites that contain messenger client download and web-based messaging sites
- *jewelry*: Sites about and selling jewelry
- *jobsearch*: Sites for finding jobs
- *kidstimewasting*: Sites kids often waste time on
- *mail*: Webmail and email sites
- *marketingware*: Sites about marketing products



- *medical*: Medical websites
- *mixed\_adult*: Mixed adult content sites
- *mobile-phone*: Sites to do with mobile phones
- *naturism*: Sites that contain nude pictures and/or promote a nude lifestyle
- *news*: News sites
- *onlineauctions*: Online auctions
- *onlinegames*: Online gaming sites
- *onlinepayment*: Online payment sites
- *personalfinance*: Personal finance sites
- *pets*: Pet sites
- *phishing*: Sites attempting to trick people into giving out private information
- *porn*: Pornography
- *proxy*: Sites with proxies to bypass filters
- *radio*: non-news related radio and television
- *religion*: Sites promoting religion
- *ringtones*: Sites containing ring tones, games, pictures and other
- *searchengines*: Search engines such as google
- *sect*: Sites about religious groups
- *sexuality*: Sites dedicated to sexuality, possibly including adult material
- *shopping*: Shopping sites
- *socialnetworking*: Social networking websites
- *sportnews*: Sport news sites
- *sports*: All sport sites
- *spyware*: Sites who run or have spyware software to download
- *updatesites*: Sites where software updates are downloaded from including virus sigs
- *vacation*: Sites about going on holiday
- *violence*: Sites containing violence
- *virusinfected*: Sites who host virus infected files
- *warez*: Sites with illegal pirate software
- *weather*: Weather news sites and weather related
- *weapons*: Sites detailing or selling weapons
- *webmail*: Just webmail sites
- *whitelist*: Contains site suitable for kids

### Customizing the URL database

To customize the database, you have to manually edit the relevant files of the database. The URL database is located on the Zorp hosts under the `/etc/zorp/urlfilter/` directory. Every thematic category is subdirectory containing two files called `domains` and `urls`. These files contain the list of domains (e.g., `example.com`) and URLs (e.g., `example.com/news/`) that fall into the specific category. Optionally, the subdirectory may contain a third file called `expressions`, where more complex rules can be defined using regular expressions.

- To to allow access (whitelist) to a domain or URL, add it to the `domains` or `urls` file of the `whitelist` category. Do not forget to configure your Http proxies to permit access to the domains of the `whitelist` category.



**Warning**

Deleting a domain from a category is not equivalent to whitelisting. Deleted domains will be re-added to their original category after the next database update.

- To add a new URL or domain to an existing category, create a new subdirectory under `/etc/zorp/urlfilter/`, create the `domains` and `urls` files for this new category, and add the domain or URL (without the `http://www.` prefix) to the `domains` or `urls` file. Zorp will automatically add these sites to the specific category after the next daily database update, or when the `zufupdate` command is executed.
- To create a new category, create a new subdirectory under `/etc/zorp/urlfilter/`, create the `domains` and `urls` files for this new category, and add domains and URLs to these files. Do not forget to configure your Http proxies to actually use the new category.



**Warning**

Manual changes to the URL database are not applied automatically, they become effective only after the next daily database update, or when the `zufupdate` command is executed.



**Note**

Manual changes are automatically merged with the original database during database updates.

If you are using the URL-filter database on several Zorp hosts and modify the database manually, make sure to copy your changes to the other hosts as well.

### 4.7.3. Related standards

- The Hypertext Transfer Protocol -- HTTP/1.1 protocol is described in RFC 2616.
- The Hypertext Transfer Protocol -- HTTP/1.0 protocol is described in RFC 1945.

### 4.7.4. Classes in the Http module

Class	Description
<code>AbstractHttpProxy</code>	Class encapsulating the abstract HTTP proxy.





Class	Description
<u><a href="#">HttpProxy</a></u>	Default HTTP proxy based on AbstractHttpProxy.
<u><a href="#">HttpProxyNonTransparent</a></u>	HTTP proxy based on HttpProxy, operating in non-transparent mode.
<u><a href="#">HttpProxyURIFilter</a></u>	HTTP proxy based on HttpProxy, with URI filtering capability.
<u><a href="#">HttpProxyURIFilterNonTransparent</a></u>	HTTP proxy based on HttpProxyURIFilter, with URI filtering capability and permitting non-transparent requests.
<u><a href="#">HttpProxyURLCategoryFilter</a></u>	HTTP proxy based on HttpProxy, with URL filtering capability based on categories.
<u><a href="#">HttpWebdavProxy</a></u>	HTTP proxy based on HttpProxy, allowing WebDAV extensions.
<u><a href="#">NontransHttpWebdavProxy</a></u>	HTTP proxy based on HttpProxyNonTransparent, allowing WebDAV extension in non-transparent requests.

Table 4.15. Classes of the Http module

#### 4.7.5. Class AbstractHttpProxy

This class implements an abstract HTTP proxy - it serves as a starting point for customized proxy classes, but is itself not directly usable. Service definitions should refer to a customized class derived from AbstractHttpProxy, or one of the predefined proxy classes, such as [HttpProxy](#) or [HttpProxyNonTransparent](#). AbstractHttpProxy denies all requests by default.

##### 4.7.5.1. Attributes of AbstractHttpProxy

<b>auth_by_cookie (boolean, rw:r)</b>
Default: FALSE
Authentication informations for one-time-password mode is organized by a cookie not the address of the client.

<b>auth_by_form (boolean, rw:r)</b>
Default: FALSE
When enabled, and the client tries to access an URL that requires authentication, a webpage where users can enter their authentication information is displayed. If the authentication is successful, the result is cached in a cookie.

<b>auth_cache_time (integer, rw:r)</b>
Default: 0

<b>auth_cache_time (integer, rw:r)</b>
Caching authentication information this amount of seconds.
<b>auth_cache_update (boolean, rw:r)</b>
Default: FALSE
Update authentication cache by every connection.
<b>auth_forward (boolean, rw:rw)</b>
Default: FALSE
Controls whether inband authentication information (username and password) should be forwarded to the upstream server. When a parent proxy is present, the incoming authentication request is put into a 'Proxy-Authorization' header. In other cases the 'WWW-Authentication' header is used.
<b>auth_realm (string, w:r)</b>
Default: "Zorp HTTP auth"
The name of the authentication realm to be presented to the user in the dialog window during inband authentication.
<b>buffer_size (integer, rw:r)</b>
Default: 1500
Size of the I/O buffer used to transfer entity bodies.
<b>connect_proxy (class, rw:rw)</b>
Default: PlugProxy
For CONNECT requests the HTTP proxy starts an independent proxy to control the internal protocol. The connect_proxy attribute specifies which proxy class is used for this purpose.
<b>connection_mode (enum, n/a:rw)</b>
Default: n/a
This value reflects the state of the session. If the value equals to 'HTTP_CONNECTION_CLOSE', the session will be closed after serving the current request. Otherwise, if the value is 'HTTP_CONNECTION_KEEPALIVE' another request will be fetched from the client. This attribute can be used to forcibly close a keep-alive connection.
<b>current_header_name (string, n/a:rw)</b>
Default: n/a

<b>current_header_name (string, n/a:rw)</b>
Name of the header. It is defined when the header is processed, and can be modified by the proxy to actually change a header in the request or response.
<b>current_header_value (string, n/a:rw)</b>
Default: n/a
Value of the header. It is defined when the header is processed, and can be modified by the proxy to actually change the value of the header in the request or response.
<b>default_port (integer, rw:rw)</b>
Default: 80
This value is used in non-transparent mode when the requested URL does not contain a port number. The default should be 80, otherwise the proxy may not function properly.
<b>enable_session_persistence (boolean, rw:rw)</b>
Default: FALSE
Allow persistent load balanced connections when accessing session-aware application servers.
<b>enable_url_filter (boolean, rw:r)</b>
Default: FALSE
Enables URL filtering in HTTP requests. See <i>Section 4.7.2.12, URL filtering in HTTP (p. 44)</i> for details.
<b>enable_url_filter_dns (boolean, rw:r)</b>
Default: FALSE
Enables DNS- and reverse-DNS resolution to ensure that a domain or URL is correctly categorized even when it is listed in the database using its domain name, but the client tries to access it with its IP address (or vice-versa). See <i>Section 4.7.2.12, URL filtering in HTTP (p. 44)</i> for details.
<b>error_files_directory (string, rw:rw)</b>
Default: "/usr/share/zorp/http"
Location of HTTP error messages.
<b>error_headers (string, n/a:rw)</b>
Default: n/a
A string included as a header in the error response. The string must be a valid header and must end with a " " sequence.

<b>error_info (string, n/a:rw)</b>
Default: n/a
A string to be included in error messages.
<b>error_msg (string, n/a:rw)</b>
Default: n/a
A string used as an error message in the HTTP status line.
<b>error_silent (boolean, rw:rw)</b>
Default: FALSE
Turns off verbose error reporting to the HTTP client (makes firewall fingerprinting more difficult).
<b>error_status (integer, rw:rw)</b>
Default: 500
If an error occurs, this value will be used as the status code of the HTTP response it generates.
<b>keep_persistent (boolean, rw:r)</b>
Default: FALSE
Try to keep the connection to the client persistent even if the server does not support it.
<b>language (string, rw:r)</b>
Default: "en"
Specifies the language of the HTTP error pages displayed to the client. English ( <i>en</i> ) is the default. Other supported languages: <i>de</i> (German); <i>hu</i> (Hungarian).
<b>max_auth_time (integer, rw:rw)</b>
Default: 0
Request password authentication from the client, invalidating cached one-time-passwords. If the time specified (in seconds) in this attribute expires, a new authentication from the client browser is requested even if it still has a password cached.
<b>max_body_length (integer, rw:rw)</b>
Default: 0
Maximum allowed length of an HTTP request or response body. The default "0" value means that the length of the body is not limited.

<b>max_chunk_length (integer, rw:rw)</b>
Default: 0
Maximum allowed length of a single chunk when using chunked transfer-encoding. The default "0" value means that the length of the chunk is not limited.
<b>max_header_lines (integer, rw:rw)</b>
Default: 50
Maximum number of header lines allowed in a request or response.
<b>max_hostname_length (integer, rw:rw)</b>
Default: 256
Maximum allowed length of the hostname field in URLs.
<b>max_keepalive_requests (integer, rw:rw)</b>
Default: 0
Maximum number of requests allowed in a single session. If the number of requests in the session reaches this limit, the connection is terminated. The default "0" value allows unlimited number of requests.
<b>max_line_length (integer, rw:r)</b>
Default: 4096
Maximum allowed length of lines in requests and responses. This value does not affect data transfer, as data is transmitted in binary mode.
<b>max_url_length (integer, rw:rw)</b>
Default: 4096
Maximum allowed length of an URL in a request. Note that this directly affects forms using the 'GET' method to pass data to CGI scripts.
<b>parent_proxy (string, rw:rw)</b>
Default: ""
The address or hostname of the parent proxy to be connected. Either DirectedRouter or InbandRouter has to be used when using parent proxy.
<b>parent_proxy_port (integer, rw:rw)</b>
Default: 3128
The port of the parent proxy to be connected.

<b>permit_ftp_over_http (boolean, rw:r)</b>
Default: FALSE
Allow processing FTP URLs in non-transparent mode.
<b>permit_http09_responses (boolean, rw:r)</b>
Default: TRUE
Allow server responses to use the limited HTTP/0.9 protocol. As these responses carry no control information, verifying the validity of the protocol stream is impossible. This does not pose a threat to web clients, but exploits might pass undetected if this option is enabled for servers. It is recommended to turn this option off for protecting servers and only enable it when Zorp is used in front of users.
<b>permit_invalid_hex_escape (boolean, rw:r)</b>
Default: FALSE
Allow invalid hexadecimal escaping in URLs (% must be followed by two hexadecimal digits).
<b>permit_null_response (boolean, rw:r)</b>
Default: TRUE
Permit RFC incompliant responses with headers not terminated by CRLF and not containing entity body.
<b>permit_proxy_requests (boolean, rw:r)</b>
Default: FALSE
Allow proxy-type requests in transparent mode.
<b>permit_server_requests (boolean, rw:r)</b>
Default: TRUE
Allow server-type requests in non-transparent mode.
<b>permit_unicode_url (boolean, rw:r)</b>
Default: FALSE
Allow unicode characters in URLs encoded as %u. This is an IIS extension to HTTP, UNICODE (UTF-7, UTF-8 etc.) URLs are forbidden by the RFC as default.
<b>request (complex, rw:rw)</b>
Default: empty
Normative policy hash for HTTP requests indexed by the HTTP method (e.g.: "GET", "PUT" etc.). See also <i>Section 4.7.2.2, Configuring policies for HTTP requests and responses (p. 38)</i> .

<b>request_count (integer, n/a:r)</b>
Default: 0
The number of keepalive requests within the session.
<b>request_header (complex, rw:rw)</b>
Default: empty
Normative policy hash for HTTP header requests indexed by the header names (e.g.: "Set-cookie"). See also <i>Section 4.7.2.3, Configuring policies for HTTP headers (p. 40)</i> .
<b>request_method (string, n/a:r)</b>
Default: n/a
Request method (GET, POST, etc.) sent by the client.
<b>request_mime_type (string, n/a:r)</b>
Default: n/a
The MIME type of the request entity. Its value is only defined when the request is processed.
<b>request_stack (complex, rw:rw)</b>
Default: n/a
Attribute containing the request stacking policy: the hash is indexed based on method names (e.g.: GET). See <i>Section 4.7.2.9, Stacking (p. 43)</i> .
<b>request_url (string, n/a:rw)</b>
Default: n/a
The URL requested by the client. It can be modified to redirect the current request.
<b>request_url_file (string, n/a:r)</b>
Default: n/a
Filename specified in the URL.
<b>request_url_host (string, n/a:r)</b>
Default: n/a
Remote hostname in the URL.
<b>request_url_passwd (string, n/a:r)</b>
Default: n/a

<b>request_url_passwd (string, n/a:r)</b>
Password in the URL (if specified).
<b>request_url_port (integer, n/a:r)</b>
Default: n/a
Port number as specified in the URL.
<b>request_url_proto (string, n/a:r)</b>
Default: n/a
Protocol specifier of the URL. This attribute is an alias for <i>request_url_scheme</i> .
<b>request_url_scheme (string, n/a:r)</b>
Default: n/a
Protocol specifier of the URL (http://, ftp://, etc.).
<b>request_url_username (string, n/a:r)</b>
Default: n/a
Username in the URL (if specified).
<b>request_version (string, n/a:r)</b>
Default: n/a
Request version (1.0, 1.1, etc.) used by the client.
<b>require_host_header (boolean, rw:r)</b>
Default: TRUE
Require the presence of the Host header. If set to FALSE, the real URL cannot be recovered from certain requests, which might cause problems with URL filtering.
<b>rerequest_attempts (integer, rw:rw)</b>
Default: 0
Controls the number of attempts the proxy takes to send the request to the server. In case of server failure, a reconnection is made and the complete request is repeated along with POST data.
<b>reset_on_close (boolean, rw:rw)</b>
Default: FALSE



**reset\_on\_close (boolean, rw:rw)**

Whenever the connection is terminated without a proxy generated error message, send an RST instead of a normal close. Causes some clients to automatically reconnect.

**response (complex, rw:rw)**

Default: empty

Normative policy hash for HTTP responses indexed by the HTTP method and the response code (e.g.: "PWD", "209" etc.). See also *Section 4.7.2.2, Configuring policies for HTTP requests and responses (p. 38)*.

**response\_header (complex, rw:rw)**

Default: empty

Normative policy hash for HTTP header responses indexed by the header names (e.g.: "Set-cookie"). See also *Section 4.7.2.3, Configuring policies for HTTP headers (p. 40)*.

**response\_mime\_type (string, n/a:r)**

Default: n/a

The MIME type of the response entity. Its value is only defined when the response is processed.

**response\_stack (complex, rw:rw)**

Default: n/a

Attribute containing the response stacking policy: the hash is indexed based on method names (e.g.: GET). See *Section 4.7.2.9, Stacking (p. 43)*.

**rewrite\_host\_header (boolean, rw:rw)**

Default: TRUE

Rewrite the Host header in requests when URL redirection is performed.

**session\_persistence\_cookie\_name (string, rw:rw)**

Default: "JSESSIONID"

The name of the cookie which will be used to persist load balanced connections when accessing session-aware application servers.

**session\_persistence\_cookie\_salt (string, rw:rw)**

Default: n/a

The salt to use when hashing the target server addresses in persistent load balanced connections. If session persistence is enabled, this parameter must be set.

<b>strict_header_checking (boolean, rw:r)</b>
Default: FALSE
Require RFC conformant HTTP headers.
<b>strict_header_checking_action (enum, rw:r)</b>
Default: HTTP_HDR_DROP
This attribute controls what should happen if a non-rfc conform or unknown header found in the communication. Only the HTTP_HDR_ACCEPT, HTTP_HDR_DROP and HTTP_HDR_ABORT can be used.
<b>target_port_range (string, rw:rw)</b>
Default: "80,443"
List of ports that non-transparent requests are allowed to use. The default is to allow port 80 and 443 to permit HTTP and HTTPS traffic. (The latter also requires the CONNECT method to be enabled).
<b>timeout (integer, rw:rw)</b>
Default: 300000
General I/O timeout in milliseconds. If there is no timeout specified for a given operation, this value is used.
<b>timeout_request (integer, rw:rw)</b>
Default: 10000
Time to wait for a request to arrive from the client.
<b>timeout_response (integer, rw:rw)</b>
Default: 300000
Time to wait for the HTTP status line to arrive from the server.
<b>transparent_mode (boolean, rw:r)</b>
Default: TRUE
Set the operation mode of the proxy to transparent (TRUE) or non-transparent (FALSE).
<b>url_category (complex, rw:rw)</b>
Default: empty
Normative policy hash for category-based URL-filtering. The hash is indexed by the name of the category.
<b>url_filter_uncategorized_action (enum, rw:rw)</b>
Default: HTTP_URL_ACCEPT

<b>url_filter_uncategorized_action (enum, rw:rw)</b>
The action applied to uncategorized (unknown) URLs when URL filtering is used. By default, uncategorized URLs are accepted: <code>self.url_filter_uncategorized_action=(HTTP_URL_ACCEPT, )</code> . Note that if you set this option to <code>HTTP_URL_REJECT</code> , you must add every URL on your intranet to a category and set an <code>HTTP_URL_ACCEPT</code> rule to this category, otherwise your clients will not be able to access your intranet sites. For details, see <i>Section Configuring URL-filtering in HTTP (p. 44)</i> .

<b>use_canonicalized_urls (boolean, rw:rw)</b>
Default: TRUE
This attribute enables URL canonicalization, which means to automatically convert URLs to their canonical form. This enhances security but might cause interoperability problems with some applications. It is recommended to disable this setting on a per-destination basis. URL filtering still sees the canonicalized URL, but at the end the proxy sends the original URL to the server.

<b>use_default_port_in_transparent_mode (boolean, rw:rw)</b>
Default: TRUE
Set the target port to the value of <code>default_port</code> in transparent mode. This ensures that only the ports specified in <code>target_port_range</code> can be used by the clients, even if InbandRouter is used.

#### 4.7.5.2. AbstractHttpProxy methods

Method	Description
<code>getRequestHeader(self, header)</code>	Function returning the value of a request header.
<code>getResponseHeader(self, header)</code>	Function returning the value of a response header.
<code>setRequestHeader(self, header, new_value)</code>	Function changing the value of a request header.
<code>setResponseHeader(self, header, new_value)</code>	Function changing the value of a response header.

Table 4.16. Method summary

#### Method `getRequestHeader(self, header)`

This function looks up and returns the value of a header associated with the current request.

#### Arguments of `getRequestHeader`

<b>header (unknown, n/a:n/a)</b>
Default: n/a
Name of the header to look up.

**Method getResponseHeader(self, header)**

This function looks up and returns the value of a header associated with the current response.

**Arguments of getResponseHeader**

<b>header (unknown, n/a:n/a)</b>
Default: n/a
Name of the header to look up.

**Method setRequestHeader(self, header, new\_value)**

This function looks up and changes the value of a header associated with the current request.

**Arguments of setRequestHeader**

<b>header (unknown, n/a:n/a)</b>
Default: n/a
Name of the header to change.

<b>new_value (unknown, n/a:n/a)</b>
Default: n/a
Change the header to this value.

**Method setResponseHeader(self, header, new\_value)**

This function looks up and changes the value of a header associated with the current response.

**Arguments of setResponseHeader**

<b>header (unknown, n/a:n/a)</b>
Default: n/a
Name of the header to change.

<b>new_value (unknown, n/a:n/a)</b>
Default: n/a
Change the header to this value.

#### 4.7.6. Class HttpProxy

HttpProxy is a default HTTP proxy based on AbstractHttpProxy. It is transparent, and enables the most commonly used HTTP methods: "GET", "POST" and "HEAD".

#### 4.7.7. Class HttpProxyNonTransparent

HTTP proxy based on HttpProxy. This class is identical to *HttpProxy* with the only difference being that it is non-transparent (*transparent\_mode = FALSE*). Consequently, clients must be explicitly configured to connect to this proxy instead of the target server and issue proxy requests. On the server side this proxy connects transparently to the target server.

For the correct operation the proxy must be able to set the server address on its own. This can be accomplished by using *InbandRouter*.

#### 4.7.8. Class HttpProxyURIFilter

HTTP proxy based on HttpProxy, having URL filtering capability. The matcher attribute should be initialized to refer to a Matcher object. The initialization should be done in the class body as shown in the next example.



Example 4.15. URL filtering HTTP proxy

```
class MyHttp(HttpProxyURIFilter):
    matcher = RegexpFileMatcher('/etc/zorp/blacklist.txt',
                               '/etc/zorp/whitelist.txt')
```

##### 4.7.8.1. Attributes of HttpProxyURIFilter

<b>matcher (class, rw:rw)</b>
Default: None
Matcher determining whether access to an URL is permitted or not.

#### 4.7.9. Class HttpProxyURIFilterNonTransparent

HTTP proxy based on HttpProxyURIFilter, but operating in non-transparent mode (*transparent\_mode = FALSE*).

#### 4.7.10. Class HttpProxyURLCategoryFilter

HTTP proxy based on HttpProxy with enabled URL filtering (with DNS and reverse-DNS resolution) and preconfigured default category actions.

The following categories have policy action *HTTP\_URL\_REJECT*:

- ads
- adult

- blacklist
- drugs
- gambling
- hacking
- phishing
- porn
- sexuality
- spyware
- violence
- virusinfected
- warez

The following categories have policy action *HTTP\_URL\_ACCEPT*:

- whitelist

#### 4.7.11. Class HttpWebdavProxy

HTTP proxy based on HttpProxy, also capable of inspecting WebDAV extensions of the HTTP protocol.

The following requests are permitted: PROPFIND; PROPPATCH; MKCOL; COPY; MOVE; LOCK; UNLOCK.

#### 4.7.12. Class NontransHttpWebdavProxy

HTTP proxy based on HttpProxyNonTransparent, operating in non-transparent mode (*transparent\_mode* = *FALSE*) and capable of inspecting WebDAV extensions of the HTTP protocol.

The following requests are permitted: PROPFIND; PROPPATCH; MKCOL; COPY; MOVE; LOCK; UNLOCK.

### 4.8. Module Plug

This module defines an interface to the Plug proxy. Plug is a simple TCP or UDP circuit, which means that transmission takes place without protocol verification.

#### 4.8.1. Proxy behavior

This class implements a general plug proxy, and is capable of optionally disabling data transfer in either direction. Plug proxy reads connection on the client side, then creates another connection at the server side. Arriving responses are sent back to the client. However, it is not a protocol proxy, therefore PlugProxy does not implement any protocol analysis. It offers protection to clients and servers from lower level (e.g.: IP) attacks. It is mainly used to allow traffic pass the firewall for which there is no protocol proxy available.

By default plug copies all data in both directions. To change this behavior, set the *copy\_to\_client* or *copy\_to\_server* attribute to *FALSE*.

Plug supports the use of secondary sessions. For details, see *Section 2.2, Secondary sessions (p. 7)*.



**Note**  
Copying of out-of-band data is not supported.

## 4.8.2. Related standards

Plug proxy is not a protocol specific proxy module, therefore it is not specified in standards.

## 4.8.3. Classes in the Plug module

Class	Description
<i>AbstractPlugProxy</i>	Class encapsulating the abstract Plug proxy.
<i>PlugProxy</i>	Class encapsulating the default Plug proxy.

Table 4.17. Classes of the Plug module

## 4.8.4. Class AbstractPlugProxy

An abstract proxy class for transferring data.

### 4.8.4.1. Attributes of AbstractPlugProxy

<b>bandwidth_to_client (integer, n/a:r)</b>
Default: n/a
Read-only variable containing the bandwidth currently used in server->client direction.
<b>bandwidth_to_server (integer, n/a:r)</b>
Default: n/a
Read-only variable containing the bandwidth currently used in client->server direction.
<b>buffer_size (integer, w:r)</b>
Default: 1500
Size of the buffer used for copying data.
<b>copy_to_client (boolean, w:r)</b>
Default: TRUE
Allow data transfer in the server->client direction.
<b>copy_to_server (boolean, w:r)</b>
Default: TRUE
Allow data transfer in the client->server direction.



<b>packet_stats_interval_packet (integer, w:r)</b>
Default: 0
The number of passing packages between two successive packetStats() events. It can be useful when the Quality of Service for the connection is influenced dynamically. Set to 0 to turn packetStats() off.
<b>packet_stats_interval_time (integer, w:r)</b>
Default: 0
The time in milliseconds between two successive packetStats() events. It can be useful when the Quality of Service for the connection is influenced dynamically. Set to 0 to turn packetStats() off.
<b>secondary_mask (secondary_mask, rw:r)</b>
Default: 0xf
Specifies which connections can be handled by the same proxy instance. See <i>Section 2.2, Secondary sessions (p. 7)</i> for details.
<b>secondary_sessions (integer, rw:r)</b>
Default: 10
Maximum number of allowed secondary sessions within a single proxy instance. See <i>Section 2.2, Secondary sessions (p. 7)</i> for details.
<b>shutdown_soft (boolean, w:r)</b>
Default: FALSE
If enabled, the two sides of a connection are closed separately. (E.g.: if the server closes the connection the client side connection is held until it is verified that no further data arrives, for example from a stacked proxy.) It is automatically enabled when proxies are stacked into the connection.
<b>stack_proxy (enum, w:r)</b>
Default: n/a
Proxy class to stack into the connection. All data is passed to the specified proxy.
<b>timeout (integer, w:r)</b>
Default: 600000
I/O timeout in milliseconds.



#### 4.8.4.2. AbstractPlugProxy methods

Method	Description
<i>packetStats(self, client_bytes, client_pkts, server_bytes, server_pkts)</i>	Function called when the packet_stats_interval is elapsed.

Table 4.18. Method summary

#### Method packetStats(self, client\_bytes, client\_pkts, server\_bytes, server\_pkts)

This function is called whenever the time interval set in packet\_stats\_interval elapses, or a given number of packets were transmitted. This event receives packet statistics as parameters. It can be used in managing the Quality of Service of the connections; e.g.: to terminate connections with excessive bandwidth requirements (for instance to limit the impact of a covert channel opened when using plug instead of a protocol specific proxy).

#### Arguments of packetStats

<b>client_bytes (unknown, n/a:n/a)</b>
Default: n/a
Number of bytes transmitted to the client.
<b>client_pkts (unknown, n/a:n/a)</b>
Default: n/a
Number of packets transmitted to the client.
<b>server_bytes (unknown, n/a:n/a)</b>
Default: n/a
Number of bytes transmitted to the server.
<b>server_pkts (unknown, n/a:n/a)</b>
Default: n/a
Number of packets transmitted to the server.

#### 4.8.5. Class PlugProxy

A default PlugProxy based on AbstractPlugProxy.

### 4.9. Module Pop3

The Pop3 module defines the classes constituting the proxy for the POP3 protocol.

### 4.9.1. The POP3 protocol

Post Office Protocol version 3 (POP3) is usually used by mail user agents (MUAs) to download messages from a remote mailbox. POP3 supports a single mailbox only, it does not support advanced multi-mailbox operations offered by alternatives such as IMAP.

The POP3 protocol uses a single TCP connection to give access to a single mailbox. It uses a simple command/response based approach, the client issues a command and a server can respond either positively or negatively.

#### 4.9.1.1. Protocol elements

The basic protocol is the following: the client issues a request (also called command in POP3 terminology) and the server responds with the result. Both commands and responses are line based, each command is sent as a complete line, a response is either a single line or - in case of mail transfer commands - multiple lines.

Commands begin with a case-insensitive keyword possibly followed by one or more arguments (such as RETR or DELE).

Responses begin with a status indicator ("+OK" or "-ERR") and a possible explanation of the status code (e.g.: "-ERR Permission denied.").

Responses to certain commands (usually mail transfer commands) also contain a data attachment, such as the mail body. See the *Section 4.9.1.3, Bulk transfers (p. 66)* for further details.

#### 4.9.1.2. POP3 states

The protocol begins with the server displaying a greeting message, usually containing information about the server.

After the greeting message the client takes control and the protocol enters the AUTHORIZATION state where the user has to pass credentials proving his/her identity.

After successful authentication the protocol enters TRANSACTION state where mail access commands can be issued.

When the client has finished processing, it issues a QUIT command and the connection is closed.

#### 4.9.1.3. Bulk transfers

Responses to certain commands (such as LIST or RETR) contain a long data stream. This is transferred as a series of lines, terminated by a "CRLF !" CRLF" sequence, just like in SMTP.



##### Example 4.16. POP3 protocol sample

```
+OK POP3 server ready
USER account
+OK User name is ok
PASS password
+OK Authentication successful
LIST
+OK Listing follows
1 5758
```



```

2 232323
3 3434
.
RETR 1
+OK Mail body follows
From: sender@sender.com
To: account@receiver.com
Subject: sample mail

This is a sample mail message. Lines beginning with
..are escaped, another '.' character is perpended which
is removed when the mail is stored by the client.
.
DELE 1
+OK Mail deleted
QUIT
+OK Good bye

```

## 4.9.2. Proxy behavior

Pop3Proxy is a module built for parsing messages of the POP3 protocol. It reads and parses COMMANDS on the client side, and sends them to the server if the local security policy permits. Arriving RESPONSEs are parsed as well, and sent to the client if the local security policy permits. It is possible to manipulate both the requests and the responses.

### 4.9.2.1. Default policy for commands

By default, the proxy accepts all commands recommended in RFC 1939. Additionally, the following optional commands are also accepted: USER, PASS, AUTH. The proxy understands all the commands specified in RFC 1939 and the AUTH command. These additional commands can be enabled manually.

### 4.9.2.2. Configuring policies for POP3 commands

Changing the default behavior of commands can be done using the hash named *request*. The hash is indexed by the command name (e.g.: USER or AUTH). See *Section 2.1, Policies for requests and responses (p. 4)* for details.

Action	Description
POP3_REQ_ACCEPT	Accept the request without any modification.
POP3_REQ_ACCEPT_MLINE	Accept multiline requests without modification. Use it only if unknown commands has to be enabled (i.e. commands not specified in RFC 1939 or RFC 1734).
POP3_REQ_REJECT	Reject the request. The second parameter contains a string that is sent back to the client.
POP3_REQ_POLICY	Call the function specified to make a decision about the event. See <i>Section 2.1, Policies for requests and responses (p. 4)</i> for details. This action uses two additional tuple items, which must be callable Python

Action	Description
	<p>functions. The first function receives two parameters: self and command.</p> <p>The second one is called with an answer, (if the answer is multiline, it is called with every line) and receives two parameters: self and response_param.</p>
POP3_REQ_ABORT	Reject the request and terminate the connection.

Table 4.19. Action codes for POP3 requests

**Example 4.17. Example for allowing only APOP authentication in POP3**

This sample proxy class rejects the USER authentication requests, but allows APOP requests.

```
class APop3(Pop3Proxy):
    def config(self):
        Pop3Proxy.config(self)
        self.request["USER"] = (POP3_REQ_REJECT)
        self.request["APOP"] = (POP3_REQ_ACCEPT)
```

**Example 4.18. Example for converting simple USER/PASS authentication to APOP in POP3**

The above example simply rejected USER/PASS authentication, this one converts USER/PASS authentication to APOP authentication messages.

```
class UToAPop3(Pop3Proxy):
    def config(self):
        Pop3Proxy.config(self)
        self.request["USER"] = (POP3_REQ_POLICY, self.DropUSER)
        self.request["PASS"] = (POP3_REQ_POLICY, self.UToA)

    def DropUSER(self, command):
        self.response_value = "+OK"
        self.response_param = "User ok Send Password"
        return POP3_REQ_REJECT

    def UToA(self, command):
        # Username is stored in self->username,
        # password in self->request_param,
        # and the server timestamp in self->timestamp,
        # consequently the digest can be calculated.
        # NOTE: This is only an example, calcdigest must be
        # implemented separately
        digest = calcdigest(self->timestamp+self->request_param)
        self->request_command = "APOP"
        self->request_param = name + " " + digest
        return POP3_REQ_ACCEPT
```

**4.9.2.3. Rewriting the banner**

As in many other protocols, POP3 also starts with a server banner. This banner contains the protocol version the server uses, the possible protocol extensions that it supports and, in many situations, the vendor and exact version number of the POP3 server.

This information is useful only if the clients connecting to the POP3 server can be trusted, as it might make bug hunting somewhat easier. On the other hand, this information is also useful for attackers when targeting this service.

To prevent this, the banner can be replaced with a neutral one. Use the *request* hash with the 'GREETING' keyword as shown in the following example.



**Example 4.19. Rewriting the banner in POP3**

```
class NeutralPop3(Pop3Proxy):
    def config(self):
        Pop3Proxy.config(self)
        self.request["GREETING"] = (POP3_REQ_POLICY, None, self.rewriteBanner)

    def rewriteBanner(self, response):
        self.response_param = "Pop3 server ready"
        return POP3_RSP_ACCEPT
```



**Note**

Some protocol extensions (most notably APOP) use random characters in the greeting message as salt in the authentication process, so changing the banner when APOP is used effectively prevents APOP from working properly.

#### 4.9.2.4. Stacking

The available stacking modes for this proxy module are listed in the following table. For additional information on stacking, see *Section 2.3.1, Proxy stacking (p. 7)*.

Action	Description
POP3_STK_POLICY	Call the function specified to decide which part (if any) of the traffic should be passed to the stacked proxy.
POP3_STK_NONE	No additional proxy is stacked into the POP3 proxy.
POP3_STK_MIME	The data part of the traffic including the MIME headers is passed to the specified stacked proxy.
POP3_STK_DATA	Only the data part of the traffic is passed to the specified stacked proxy.

Table 4.20. Action codes for proxy stacking

#### 4.9.2.5. Rejecting viruses and spam

When filtering messages for viruses or spam, the content vectoring modules reject infected and spam e-mails. In such cases the POP3 proxy notifies the client about the rejected message in a special e-mail.

To reject e-mail messages using the *ERR* protocol element, set the *reject\_by\_mail* attribute to *FALSE*. However, this is not recommended, because several client applications handle *ERR* responses incorrectly.

**Note**

Infected e-mails are put into the quarantine and deleted from the server.

### 4.9.3. Related standards

- Post Office Protocol Version 3 is described in RFC 1939.
- The POP3 AUTHentication command is described in RFC 1734.
- The POP3 STLS extension is described in RFC 2595.

### 4.9.4. Classes in the Pop3 module

Class	Description
<i>AbstractPop3Proxy</i>	Class encapsulating the abstract POP3 proxy.
<i>Pop3Proxy</i>	Default POP3 proxy based on AbstractPop3Proxy.
<i>Pop3STLSProxy</i>	POP3 proxy based on Pop3Proxy allowing Start TLS.

Table 4.21. Classes of the Pop3 module

### 4.9.5. Class AbstractPop3Proxy

This class implements an abstract POP3 proxy - it serves as a starting point for customized proxy classes, but is itself not directly usable. Service definitions should refer to a customized class derived from AbstractPop3Proxy, or a predefined Pop3Proxy proxy class. AbstractPop3Proxy denies all requests by default.

#### 4.9.5.1. Attributes of AbstractPop3Proxy

<b>max_authline_count (integer, rw:r)</b>
Default: 4
Maximum number of lines that can be sent during the authentication conversation. The default value is enough for password authentication, but might have to be increased for other types of authentication.
<b>max_password_length (integer, rw:r)</b>
Default: 16
Maximum allowed length of passwords.
<b>max_request_line_length (integer, rw:r)</b>
Default: 90



<b>max_request_line_length (integer, rw:r)</b>
Maximum allowed line length for client requests, without the CR-LF line ending characters.
<b>max_response_line_length (integer, rw:r)</b>
Default: 512
Maximum allowed line length for server responses, without the CR-LF line ending characters.
<b>max_username_length (integer, rw:r)</b>
Default: 8
Maximum allowed length of usernames.
<b>password (string, n/a:r)</b>
Default:
Password sent to the server (if any).
<b>permit_longline (boolean, rw:r)</b>
Default: FALSE
In multiline answer (especially in downloaded messages) sometimes very long lines can appear. Enabling this option allows the unlimited long lines in multiline answers.
<b>permit_unknown_command (boolean, rw:r)</b>
Default: FALSE
Enable unknown commands.
<b>reject_by_mail (boolean, rw:r)</b>
Default: TRUE
If the stacked proxy or content vectoring module rejects an e-mail message, reply with a special e-mail message instead of an <i>ERR</i> response. See <i>Section 4.9.2.5, Rejecting viruses and spam (p. 69)</i> for details.
<b>request (complex, rw:rw)</b>
Default:
Normative policy hash for POP3 requests indexed by the command name (e.g.: "USER", "UIDL", etc.). See also <i>Section 4.9.2.2, Configuring policies for POP3 commands (p. 67)</i> .
<b>request_command (string, n/a:rw)</b>
Default: n/a

<b>request_command (string, n/a:rw)</b>
When a command is passed to the policy level, its value can be changed to this value.
<b>request_param (string, n/a:rw)</b>
Default: n/a
When a command is passed to the policy level, the value of its parameters can be changed to this value.
<b>response_multiline (boolean, n/a:rw)</b>
Default: n/a
Enable multiline responses.
<b>response_param (string, n/a:rw)</b>
Default: n/a
When a command or response is passed to the policy level, the value its parameters can be changed to this value. (It has effect only if the return value is not POP3*_ACCEPT).
<b>response_stack (complex, rw:rw)</b>
Default:
Hash containing the stacking policy for multiline POP3 responses. The hash is indexed by the POP3 response. See also <i>Section 4.9.2.4, Stacking (p. 69)</i> .
<b>response_value (string, n/a:rw)</b>
Default: n/a
When a command or response is passed to the policy level, its value can be changed to this value. (It has effect only if the return value is not POP3*_ACCEPT).
<b>session_timestamp (string, n/a:r)</b>
Default: n/a
If the POP3 server implements the APOP command, with the greeting message it sends a timestamp, which is stored in this parameter.
<b>timeout (integer, rw:r)</b>
Default: 600000
Timeout in milliseconds. If no packet arrives within this interval, connection is dropped.



<b>username (string, n/a:r)</b>
Default: n/a
Username as specified by the client.

#### 4.9.6. Class Pop3Proxy

Pop3Proxy is the default POP3 proxy based on AbstractPop3Proxy, allowing the most commonly used requests.

The following requests are permitted: APOP; DELE; LIST; LAST; NOOP; PASS; QUIT; RETR; RSET; STAT; TOP; UIDL; USER; GREETING. All other requests (including CAPA) are rejected.

#### 4.9.7. Class Pop3STLSProxy

Pop3STLSProxy is based on Pop3Proxy, allowing the most commonly used requests.

The following requests are permitted: APOP; DELE; LIST; LAST; NOOP; PASS; QUIT; RETR; RSET; STAT; TOP; UIDL; USER; CAPA; STLS; GREETING. All other requests are rejected. The self.max\_request\_line\_length is set to 253.

### 4.10. Module Smtplib

Simple Mail Transport Protocol (SMTP) is a protocol for transferring electronic mail messages from Mail User Agents (MUAs) to Mail Transfer Agents (MTAs). It is also used for exchanging mails between MTAs.

#### 4.10.1. The SMTP protocol

The main goal of SMTP is to reliably transfer mail objects from the client to the server. A mail transaction involves exchanging the sender and recipient information and the mail body itself.

##### 4.10.1.1. Protocol elements

SMTP is a traditional command based Internet protocol; the client issues command verbs with one or more arguments, and the server responds with a 3 digit status code and additional information. The response can span one or multiple lines, the continuation is indicated by a '-' character between the status code and text.

The communication itself is stateful, the client first specifies the sender via the "MAIL" command, then the recipients using multiple "RCPT" commands. Finally it sends the mail body using the "DATA" command. After a transaction finishes the client either closes the connection using the "QUIT" command, or starts a new transaction with another "MAIL" command.



##### Example 4.20. SMTP protocol sample

```
220 mail.example.com ESMTP Postfix (Debian/GNU)
EHLO client.host.name
250-mail.example.com
250-PIPELINING
250-SIZE 50000000
250-VRFY
250-ETRN
250-XVERP
```

```
250 8BITMIME
MAIL From: <sender@sender.com>
250 Sender ok
RCPT To: <account@recipient.com>
250 Recipient ok
RCPT To: <account2@recipient.com>
250 Recipient ok
DATA
354 Send mail body
From: sender@sender.com
To: account@receiver.com
Subject: sample mail

This is a sample mail message. Lines beginning with
..are escaped, another '.' character is prepended which
is removed when the mail is stored by the client.
.
250 Ok: queued as BF47618170
QUIT
221 Farewell
```

#### 4.10.1.2. Extensions

Originally SMTP had a very limited set of commands (HELO, MAIL, RCPT, DATA, RSET, QUIT, NOOP) but as of RFC 1869, an extension mechanism was introduced. The initial HELO command was replaced by an EHLO command and the response to an EHLO command contains all the extensions the server supports. These extensions are identified by an IANA assigned name.

Extensions are used for example to implement inband authentication (AUTH), explicit message size limitation (SIZE) and explicit queue run initiation (ETRN). Each extension might add new command verbs, but might also add new arguments to various SMTP commands. The SMTP proxy has built in support for the most important SMTP extensions, further extensions can be added through customization.

#### 4.10.1.3. Bulk transfer

The mail object is transferred as a series of lines, terminated by the character sequence "CRLF" CRLF. When the '.' character occurs as the first character of a line, an escaping '.' character is prepended to the line which is automatically removed by the peer.

### 4.10.2. Proxy behavior

The Smtplib module implements the SMTP protocol as specified in RFC 2821. The proxy supports the basic SMTP protocol plus five extensions, namely: PIPELINING, SIZE, ETRN, 8BITMIME, and STARTTLS. All other ESMTP extensions are filtered by dropping the associated token from the EHLO response. If no connection can be established to the server, the request is rejected with an error message. In this case the proxy tries to connect the next mail exchange server.

#### 4.10.2.1. Default policy for commands

The abstract SMTP proxy rejects all commands and responses by default. Less restrictive proxies are available as derived classes (e.g.: SmtplibProxy), or can be customized as required.



#### 4.10.2.2. Configuring policies for SMTP commands and responses

Changing the default behavior of commands can be done by using the hash attribute *request*. These hashes are indexed by the command name (e.g.: MAIL or DATA). Policies for responses can be configured using the *response* attribute, which is indexed by the command name and the response code. The possible actions are shown in the tables below. See *Section 2.1, Policies for requests and responses (p. 4)* for details. When looking up entries of the *response* attribute hash, the lookup precedence described in *Section 2.1.2, Response codes (p. 6)* is used.

Action	Description
SMTP_REQ_ACCEPT	Accept the request without any modification.
SMTP_REQ_REJECT	Reject the request. The second parameter contains an SMTP status code, the third one an associated parameter which will be sent back to the client.
SMTP_REQ_ABORT	Reject the request and terminate the connection.

Table 4.22. Action codes for SMTP requests

Action	Description
SMTP_RSP_ACCEPT	Accept the response without any modification.
SMTP_RSP_REJECT	Reject the response. The second parameter contains an SMTP status code, the third one an associated parameter which will be sent back to the client.
SMTP_RSP_ABORT	Reject the response and terminate the connection.

Table 4.23. Action codes for SMTP responses

SMTP extensions can be controlled using the *extension* hash, which is indexed by the extension name. The supported extensions (SMTP\_EXT\_PIPELINING; SMTP\_EXT\_SIZE; SMTP\_EXT\_ETRN; SMTP\_EXT\_8BITMIME) can be accepted or dropped (SMTP\_EXT\_ACCEPT or SMTP\_EXT\_DROP) individually or all at once using the SMTP\_EXT\_ALL index value.

#### 4.10.2.3. Stacking

The available stacking modes for this proxy module are listed in the following table. For additional information on stacking, see *Section 2.3.1, Proxy stacking (p. 7)*.

Action	Description
SMTP_STK_NONE	No additional proxy is stacked into the SMTP proxy.

Action	Description
SMTP_STK_MIME	The data part including header information of the traffic is passed to the specified stacked proxy.

Table 4.24. Stacking options for SMTP

### 4.10.3. Related standards

- Simple Mail Transfer Protocol is described in RFC 2821.
- SMTP Service Extensions are described in the obsoleted RFC 1869.
- The STARTTLS extension is described in RFC 3207.

### 4.10.4. Classes in the Smtplib module

Class	Description
<i>AbstractSmtpProxy</i>	Class encapsulating the abstract SMTP proxy.
<i>SmtpProxy</i>	Default SMTP proxy based on AbstractSmtpProxy.

Table 4.25. Classes of the Smtplib module

### 4.10.5. Class AbstractSmtpProxy

This class implements an abstract SMTP proxy - it serves as a starting point for customized proxy classes, but is itself not directly usable. Service definitions should refer to a customized class derived from AbstractSmtpProxy, or one of the predefined proxy classes.

The following requests are permitted: HELO; MAIL; RCPT; DATA; RSET; QUIT; NOOP; EHLO; AUTH; ETRN. The following extensions are permitted: PIPELINING; SIZE; ETRN; 8BITMIME; STARTTLS.

#### 4.10.5.1. Attributes of AbstractSmtpProxy

<b>active_extensions (integer, n/a:r)</b>
Default: n/a
Active extension bitmask, contains bits defined by the constants 'SMTP_EXT_*'
<b>add_received_header (boolean, rw:rw)</b>
Default: FALSE
Add a Received: header into the email messages transferred by the proxy.
<b>append_domain (string, rw:rw)</b>
Default:

**append\_domain (string, rw:rw)**

Domain to append to email addresses which do not specify domain name. An address is rejected if it does not contain a domain and append\_domain is empty.

**autodetect\_domain\_from (enum, rw:rw)**

Default:

If you want to autodetect the domain name of the firewall and write it to the Received line, then set this. This attribute either set the method how the mailname should be detected. Only takes effect if add\_received\_header is TRUE.

**domain\_name (string, rw:rw)**

Default:

If you want to set a fix domain name into the added Receive line, set this. Only takes effect if add\_received\_header is TRUE.

**extensions (complex, rw:rw)**

Default:

Normative policy hash for ESMTP extension policy, indexed by the extension verb (e.g. ETRN). It contains an action tuple with the SMTP\_EXT\_\* values as possible actions.

**interval\_transfer\_noop (integer, rw:rw)**

Default: 600000

The interval between two NOOP commands sent to the server while waiting for the results of stacked proxies.

**max\_auth\_request\_length (integer, rw:r)**

Default: 256

Maximum allowed length of a request during SASL style authentication.

**max\_request\_length (integer, rw:r)**

Default: 256

Maximum allowed line length of client requests.

**max\_response\_length (integer, rw:r)**

Default: 512

Maximum allowed line length of a server response.



<b>permit_long_responses (boolean, rw:r)</b>
Default: FALSE
Permit overly long responses, as some MTAs include variable parts in responses which might get very long. If enabled, responses longer than <i>max_response_length</i> are segmented into separate messages. If disabled, such responses are rejected.
<b>permit_omission_of_angle_brackets (boolean, rw:r)</b>
Default: FALSE
Permit MAIL From and RCPT To parameters without the normally required angle brackets around them. They will be added when the message leaves the proxy anyway.
<b>permit_unknown_command (boolean, rw:r)</b>
Default: FALSE
Enable unknown commands.
<b>request (complex, rw:rw)</b>
Default:
Normative policy hash for SMTP requests indexed by the command name (e.g.: "USER", "UIDL", etc.). See also <i>Section 4.10.2.2, Configuring policies for SMTP commands and responses (p. 75)</i> .
<b>request_command (string, n/a:rw)</b>
Default: n/a
When a command is passed to the policy level, its value can be changed to this value.
<b>request_param (string, n/a:rw)</b>
Default: n/a
When a command is passed to the policy level, the value of its parameter can be changed to this value.
<b>request_stack (complex, rw:rw)</b>
Default:
Attribute containing the stacking policy for SMTP commands. See <i>Section 4.10.2.3, Stacking (p. 75)</i> .
<b>require_crlf (boolean, rw:r)</b>
Default: TRUE
Specifies whether the proxy should enforce valid CRLF line terminations.

<b>resolve_host (boolean, rw:rw)</b>
Default: FALSE
Resolve the client host from the IP address and add it to the Received line. Only takes effect if add_received_header is TRUE.
<b>response (complex, rw:rw)</b>
Default:
Normative policy hash for SMTP responses indexed by the command name and the response code. See also <i>Section 4.10.2.2, Configuring policies for SMTP commands and responses (p. 75)</i> .
<b>response_param (string, n/a:rw)</b>
Default: n/a
When a response is passed to the policy level, the value of its parameter can be changed to this value. (It has effect only when the return value is not SMTP_*_ACCEPT.)
<b>response_value (string, n/a:rw)</b>
Default: n/a
When a response is passed to the policy level, its value can be changed to this value. (It has effect only when the return value is not SMTP_*_ACCEPT.)
<b>timeout (integer, rw:r)</b>
Default: 600000
Timeout in milliseconds. If no packet arrives within this interval, the connection is dropped.
<b>tls_passthrough (boolean, rw:r)</b>
Default: FALSE
Change to passthrough mode after a successful STARTTLS request. The encrypted traffic is not processed or changed in any way, it is transported intact between the client and server.
<b>unconnected_response_code (integer, rw:rw)</b>
Default: 451
Error code sent to the client if connecting to the server fails.

#### 4.10.6. Class SmtProxy

SmtProxy implements a basic SMTP Proxy based on AbstractSmtProxy, with relay checking and sender/recipient check restrictions. (Exclamation marks and percent signs are not allowed in the e-mail addresses.)

#### 4.10.6.1. Attributes of SmtProxy

<b>error_soft (boolean, rw:rw)</b>
Default: FALSE
Return a soft error condition when recipient filter does not match. If enabled, the proxy will try to re-validate the recipient and send the mail again. This option is useful when the server used for the recipient matching is down.
<b>permit_exclamation_mark (boolean, rw:rw)</b>
Default: FALSE
Allow the '!' sign in the local part of e-mail addresses.
<b>permit_percent_hack (boolean, rw:rw)</b>
Default: FALSE
Allow the '%' sign in the local part of e-mail addresses.
<b>recipient_matcher (class, rw:rw)</b>
Default:
Matcher class (e.g.: SmtInvalidRecipientMatcher) used to check and filter recipient e-mail addresses.
<b>relay_check (boolean, rw:rw)</b>
Default: TRUE
Enable/disable relay checking.
<b>relay_domains (complex, rw:r)</b>
Default:
Domains mails are accepted for. Use Postfix style lists. (E.g.: '.example.com' allows every subdomain of example.com, but not example.com. To match example.com use 'example.com'.)
<b>relay_domains_matcher (class, rw:r)</b>
Default:
Domains mails are accepted for based on a matcher (e.g.: RegexpFileMatcher).
<b>relay_zones (complex, rw:r)</b>
Default:
Zones that are relayed. The administrative hierarchy of the zone is also used.





<b>sender_matcher (class, rw:rw)</b>
Default:
Matcher class (e.g.: SmtplibInvalidRecipientMatcher) used to check and filter sender e-mail addresses.

## 4.11. Module Telnet

The Telnet module defines the classes constituting the proxy for the TELNET protocol.

### 4.11.1. The Telnet protocol

The Telnet protocol was designed to remotely login to computers via the network. Although its main purpose is to access a remote standard terminal, it can be used for many other functions as well.

The protocol follows a simple scenario. The client opens a TCP connection to the server at the port 23. The server authenticates the client and opens a terminal. At the end of the session the server closes the connection. All data is sent in plain text format without any encryption.

#### 4.11.1.1. The network virtual terminal

The communication is based on the network virtual terminal (NVT). Its goal is to map a character terminal so neither the "server" nor "user" hosts need to keep information about the characteristics of each other's terminals and terminal handling conventions. NVT uses 7 bit code ASCII characters as the display device. An end of line is transmitted as a CRLF (carriage return followed by a line feed). NVT ASCII is used by many other protocols as well.

NVT defines three mandatory control codes which must be understood by the participants: NULL, CR (Carriage Return), which moves the printer to the left margin of the current line and LF (Line Feed), which moves the printer to the next line keeping the current horizontal position.

NVT also contains some optional commands which are useful. These are the following:

- *BELL* is an audible or visual sign.
- *BS* (Back Space) moves the printer back one position and deletes a character.
- *HT* (Horizontal Tab) moves the printer to the next horizontal tabular stop.
- *VT* Vertical Tab moves the printer to the next vertical tabular stop.
- *FF* (Form Feed) moves the printer to the top of the next page.

#### 4.11.1.2. Protocol elements

The protocol uses several commands that control the method and various details of the interaction between the client and the server. These commands can be either mandatory commands or extensions. During the session initialization the client and the server negotiates the connection parameters with these commands. Sub-negotiation is a process during the protocol which is for exchanging extra parameters of a command (e.g.: sending the window size). The commands of the protocol are:

Request/Response	Description
SE	End of sub-negotiation parameters.
NOP	No operation.
DM	Data mark - Indicates the position of Sync event within the data stream.
BRK	Break - Indicates that a break or attention key was hit.
IP	Suspend, interrupt or abort the process.
AO	Abort output - Run a command without sending the output back to the client.
AYT	Are you there - Request a visible evidence that the AYT command has been received.
EC	Erase character - Delete the character last received from the stream.
EL	Erase line - Erase a line without a CRLF.
GA	Go Ahead - Instruct the other machine to start the transmission.
SB	Sub-negotiation starts here.
WILL	Will (option code) - Indicates the desire to begin performing the indicated option, or confirms that it is being performed.
WONT	Will not (option code) - Indicates the refusal to perform, or continue performing, the indicated option.
DO	Do (option code) - Indicates the request that the other party perform, or confirmation that the other party is expected to perform, the indicated option.
DONT	Do not (option code) - Indicates the request that the other party stop performing the indicated option, or confirmation that its performing is no longer expected.
IAC	Interpret as command.

Table 4.26. Telnet protocol commands

#### 4.11.2. Proxy behavior

TelnetProxy is a module built for parsing TELNET protocol commands and the negotiation process. It reads and parses COMMANDs on the client side, and sends them to the server if the local security policy permits. Arriving RESPONSEs are parsed as well and sent to the client if the local security policy permits. It is possible to manipulate options by using TELNET\_OPT\_POLICY. It is also possible to accept or deny certain options and suboptions.

The Telnet shell itself cannot be controlled, thus the commands issued by the users cannot be monitored or modified.

#### 4.11.2.1. Default policy

The low level abstract Telnet proxy denies every option and suboption negotiation sequences by default. The different options can be enabled either manually in a derived proxy class, or the predefined `TelnetProxy` class can be used.

#### 4.11.2.2. Configuring policies for the TELNET protocol

The Telnet proxy can enable/disable the use of the options and their suboptions within the session. Changing the default policy can be done using the `option` multi-dimensional hash, indexed by the option and the suboption (optional). If the suboption is specified, the lookup precedence described in *Section 2.1.2, Response codes (p. 6)* is used. The possible action codes are listed in the table below.

Action	Description
TELNET_OPT_ACCEPT	Allow the option.
TELNET_OPT_DROP	Reject the option.
TELNET_OPT_ABORT	Reject the option and terminate the Telnet session.
TELNET_OPT_POLICY	Call the function specified to make a decision about the event. The function receives two parameters: <code>self</code> , and <code>option</code> (an integer). See <i>Section 2.1, Policies for requests and responses (p. 4)</i> for details.

Table 4.27. Action codes for Telnet options



#### Example 4.21. Example for disabling the Telnet X Display Location option

```
class MyTelnetProxy(TelnetProxy):
    def config(self):
        TelnetProxy.config(self)
        self.option[TELNET_X_DISPLAY_LOCATION] = (TELNET_OPT_REJECT)
```

Constants have been defined for the easier use of TELNET options and suboptions. These are listed in *Table A.1, TELNET options and suboptions (p. 225)*.

#### Policy callback functions

Policy callback functions can be used to make decisions based on the content of the suboption negotiation sequence. For example, the suboption negotiation sequences of the Telnet Environment option transfer environment variables. The low level proxy implementation parses these variables, and passes their name and value to the callback function one-by-one. These values can also be manipulated during transfer, by changing the `current_var_name` and `current_var_value` attributes of the proxy class.

**Example 4.22. Rewriting the DISPLAY environment variable**

```
class MyRewritingTelnetProxy(TelnetProxy):
    def config(self):
        TelnetProxy.config()
        self.option[TELNET_ENVIRONMENT, TELNET_SB_IS] = (TELNET_OPTION_POLICY, self.rewriteVar)

    def rewriteVar(self, option, name, value):
        if name == "DISPLAY":
            self.current_var_value = "rewritten_value:0"
        return TELNET_OPTION_ACCEPT
```

**Option negotiation**

In the Telnet protocol, options and the actual commands are represented on one byte. In order to be able to use a command in a session, the option (and its suboptions if there are any) corresponding to the command has to be negotiated between the client and the server. Usually the command and the option is represented by the same value, e.g.: the *TELNET\_STATUS* command and option are both represented by the value "5". However, this is not always the case. The *negotiation* hash is indexed by the code of the command, and contains the code of the option to be negotiated for the given command (or the *TELNET\_NEG\_NONE* when no negotiation is needed).

Currently the only command where the code of the command differs from the related option is `self.negotiation["239"] = int(TELNET_EOR)`.

**4.11.3. Related standards**

The Telnet protocol is described in RFC 854. The different options of the protocol are described in various other RFCs, listed in *Table A.1, TELNET options and suboptions (p. 225)*.

**4.11.4. Classes in the Telnet module**

Class	Description
<i>AbstractTelnetProxy</i>	Class encapsulating the abstract Telnet proxy.
<i>TelnetProxy</i>	Default Telnet proxy based on AbstractTelnetProxy.
<i>TelnetProxyStrict</i>	Telnet proxy based on AbstractTelnetProxy, allowing only the minimal command set.

*Table 4.28. Classes of the Telnet module*

**4.11.5. Class AbstractTelnetProxy**

This class implements the Telnet protocol (as described in RFC 854) and its most common extensions. Although not all possible options are checked by the low level proxy, it is possible to filter any option and suboption negotiation sequences using policy callbacks. AbstractTelnetProxy serves as a starting point for customized proxy classes, but is itself not directly usable. Service definitions should refer to a customized class derived from AbstractTelnetProxy, or one of the predefined TelnetProxy proxy classes. AbstractTelnetProxy denies all options by default.

#### 4.11.5.1. Attributes of AbstractTelnetProxy

<b>current_var_name (string, n/a:rw)</b>
Default: n/a
Name of the variable being negotiated.
<b>current_var_value (string, n/a:rw)</b>
Default: n/a
Value of the variable being negotiated (e.g.: value of an environment variable, an X display location value, etc.).
<b>enable_audit (boolean, w:r)</b>
Default: FALSE
Enable session auditing.
<b>negotiation (complex, rw:rw)</b>
Default:
Normative hash listing which options must be negotiated for a given command. See <i>Section Option negotiation (p. 84)</i> for details.
<b>option (complex, rw:rw)</b>
Default: n/a
Normative policy hash for Telnet options indexed by the option and (optionally) the suboption. See also <i>Section 4.11.2.2, Configuring policies for the TELNET protocol (p. 83)</i> .
<b>timeout (integer, rw:r)</b>
Default: 600000
I/O timeout in milliseconds.

#### 4.11.6. Class TelnetProxy

TelnetProxy is a proxy class based on AbstractTelnetProxy, allowing the use of all Telnet options.

#### 4.11.7. Class TelnetProxyStrict

TelnetProxyStrict is a proxy class based on AbstractTelnetProxy, allowing the use of the options minimally required for a useful Telnet session.

The following options are permitted: ECHO; SUPPRESS\_GO\_AHEAD; TERMINAL\_TYPE; NAWS; EOR; TERMINAL\_SPEED; X\_DISPLAY\_LOCATION; ENVIRONMENT. All other options are rejected.

## 4.12. Module Whois

WHOIS is a protocol providing information about domain and IP owners.

### 4.12.1. The Whois protocol

Whois is a netwide service to the Internet users maintained by DDN Network Information Center (NIC).

The protocol follows a very simple method. First the client opens a TCP connection to the server at the port 43 and sends a one line REQUEST closed with <CRLF>. This request can contain only ASCII characters. The server sends the result back and closes the connection.

### 4.12.2. Proxy behavior

WhoisProxy is a module build for parsing messages of the WHOIS protocol. It reads and parses the REQUESTs on the client side and sends them to the server if the local security policy permits. Arriving RESPONSEs are not parsed as they do not have any fixed structure or syntax.



Example 4.23. Example WhoisProxy logging all whois requests

```
class MyWhoisProxy(AbstractWhoisProxy):
    def whoisRequest(self, request):
        log(None, CORE_DEBUG, 3, "Whois request: '%s'" % (request))
        return ZV_ACCEPT
```

### 4.12.3. Related standards

- The NICNAME/WHOIS protocol is described in RFC 954.

### 4.12.4. Classes in the Whois module

Class	Description
<i>AbstractWhoisProxy</i>	Class encapsulating the abstract Whois proxy.
<i>WhoisProxy</i>	Default proxy class based on AbstractWhoisProxy.

Table 4.29. Classes of the Whois module

### 4.12.5. Class AbstractWhoisProxy

This class implements the WHOIS protocol as specified in RFC 954.

#### 4.12.5.1. Attributes of AbstractWhoisProxy

<b>max_line_length (integer, rw:r)</b>
Default: 132
Maximum number of characters allowed in a single line.

<b>max_request_length (integer, rw:r)</b>
Default: 128
Maximum allowed length of a Whois request.
<b>request (string, n/a:rw)</b>
Default:
The Whois request.
<b>response_footer (string, rw:rw)</b>
Default:
Append this string to each Whois response.
<b>response_header (string, rw:rw)</b>
Default:
Prepend this string to each Whois response.
<b>timeout (integer, rw:r)</b>
Default: 30000
I/O timeout in milliseconds.

#### 4.12.5.2. AbstractWhoisProxy methods

Method	Description
<i>whoisRequest(self, request)</i>	Function to process whois requests.

Table 4.30. Method summary

#### Method whoisRequest(self, request)

This function is called by the Whois proxy to process the requests. It can also be used to change specific attributes of the request.

#### 4.12.6. Class WhoisProxy

A default proxy class based on AbstractWhoisProxy.

## Chapter 5. Core

This chapter provides detailed description for the core modules of Zorp.

### 5.1. Module Auth

This module contains classes related to authentication and authorization. Together with the *AuthDB* module it implements the Authentication and Authorization framework.

User authentication verifies the identity of the user trying to access a particular network service. When performed on the connection level, that enables the full auditing of the network traffic. Authentication is often used in conjunction with authorization, allowing access to a service only to clients who have the right to do so.

#### 5.1.1. Authentication and authorization basics

Authentication is a method to ensure that certain services (access to a server, etc.) can be used only by the clients allowed to access the service. The process generally called as authentication actually consists of three distinct steps:

- *Identification*: Determining the clients identity (e.g.: requesting a username).
- *Authentication*: Verifying the clients identity (e.g.: requesting a password that only the real client knows).
- *Authorization*: Granting access to the service (e.g.: verifying that the authenticated client is allowed to access the service).



#### Note

It is important to note that although authentication and authorization are usually used together, they can also be used independently. Authentication verifies the identity of the client. There are situations where authentication is sufficient, because all users are allowed to access the services, only the event and the user's identity has to be logged. On the other hand, authorization is also possible without authentication, for example if access to a service is time-limited (e.g.: it can only be accessed outside the normal work-hours, etc.). In such situations authentication is not needed.

#### 5.1.2. Authentication and authorization in Zorp

Zorp can authenticate and authorize access to the services. The aim of authentication is to identify the user and the associated group memberships. When the client initiates a connection, it actually tries to use a service. Zorp checks if an *authentication policy* is associated to the service. If an authentication policy is present, Zorp contacts the *authentication provider* specified in the authentication policy. The type of authentication (the authentication class used, e.g., *InbandAuthentication*) is also specified in the authentication policy. The authentication provider connects to an *authentication backend* (e.g., a user database) to perform the authentication of the client - Zorp itself does not directly communicate with the database.

If the authentication is successful, the client is verified if it is allowed to access the service (by evaluating the *authorization policy* and the identity and group memberships of the client). If the client is authorized to access





the service, the server-side connection is built. The client is automatically authorized if no authorization policy is assigned to the service.

Currently only one authentication provider, the () is available via the *ZAS2AuthenticationBackend* class. Authentication providers are actually configured instances of the authentication backends, and it is independent from the database that the backend connects to. The authentication backend is that ties the authentication provider to the server storing the user data. For details on using , see the *Connection authentication and authorization* chapter of the *Zorp Administrator's Guide*.

The aim of authentication is to identify the user and resolve group memberships. The results are stored in the *auth\_user* and *auth\_groups* attributes of the *session* object. Note that apart from the information required for authentication, Zorp also sends session information (e.g., the IP address of the client) to the authentication provider.

Zorp provides the following authentication classes:

- *InbandAuthentication*: Use the built-in authentication of the protocol to authenticate the client on the Zorp.
- *ServerAuthentication*: Enable the client to connect to the target server, and extract its authentication information from the protocol.
- *ZAAuthentication*: Outband authentication using the .

If the authentication is successful, Zorp verifies that the client is allowed to access the service (by evaluating the authorization policy). If the client is authorized to access the service, the server-side connection is built. The client is automatically authorized if no authorization policy is assigned to the service.

Each service can use an authorization policy to determine whether a client is allowed to access the service. If the authorization is based on the identity of the client, it takes place only after a successful authentication - identity-based authorization can be performed only if the client's identity is known and has been verified. The actual authorization is performed by Zorp, based on the authentication information received from or extracted from the protocol.

Zorp provides the following authorization classes:

- *PermitUser*: Authorize listed users.
- *PermitGroup*: Authorize users belonging to the specified groups.
- *PermitTime*: Authorize connections in a specified time interval.
- *BasicAccessList*: Combine other authorization policies into a single rule.
- *PairAuthorization*: Authorize only user pairs.
- *NEyesAuthorization*: Have another client authorize every connection.

### 5.1.3. Classes in the Auth module

Class	Description
<i>AbstractAuthentication</i>	Class encapsulating the abstract authentication interface.

Class	Description
<u>AbstractAuthorization</u>	Class encapsulating the authorization interface.
<u>AuthCache</u>	Class encapsulating the authentication cache.
<u>AuthenticationPolicy</u>	A policy determining how the user is authenticated to access the service.
<u>AuthorizationPolicy</u>	A policy determining how the user is authorized to access the service.
<u>BasicAccessList</u>	Class encapsulating the authorization by access list.
<u>InbandAuthentication</u>	Class encapsulating the inband authentication interface.
<u>NEyesAuthorization</u>	Class encapsulating N eyes authorization.
<u>PairAuthorization</u>	Class encapsulating pair-based 4 eyes authorization.
<u>PermitGroup</u>	Class encapsulating the group membership based authorization.
<u>PermitTime</u>	Class encapsulating time based authorization.
<u>PermitUser</u>	Class encapsulating the user-name based authorization.
<u>SatyrAuthentication</u>	Class encapsulating the outband authentication interface using the Satyr application.
<u>ServerAuthentication</u>	Class encapsulating the server authentication interface.
<u>ZAAAuthentication</u>	Class encapsulating the outband authentication interface using the Zorp Authentication Agent.

Table 5.1. Classes of the Auth module

#### 5.1.4. Class AbstractAuthentication

This class encapsulates interfaces for inband and outband authentication procedures. Service definitions should refer to a customized class derived from AbstractAuthentication, or one of the predefined authentication classes, such as InbandAuthentication or ZAAAuthentication.

##### 5.1.4.1. AbstractAuthentication methods

Method	Description
<u>__init__</u> ( <i>self, authentication_provider, auth_cache</i> )	Constructor to initialize an AbstractAuthentication instance.

Table 5.2. Method summary

##### Method \_\_init\_\_ (*self, authentication\_provider, auth\_cache*)

This constructor initializes an instance of the AbstractAuthentication class.



### 5.1.5. Class AbstractAuthorization

This class encapsulates an authorization interface. Authorization determines whether the authenticated entity is in fact allowed to access a specific service. Service definitions should refer to a customized class derived from AbstractAuthorization, or one of the predefined authorization classes, such as *PermitUser* or *PermitGroup*.

### 5.1.6. Class AuthCache

This class encapsulates an authentication cache which associates usernames with client IP addresses. The association between a username and an IP address is valid only until the specified timeout. Caching the authentication results means that the users do not need to authenticate themselves for every request: it is assumed that the same user is using the computer within the timeout. E.g.: once authenticated for an HTTP service, the client can browse the web for **Timeout** period, but has to authenticate again to use FTP.

To use a single authorization cache for every service request of a client, set the *service\_equiv* attribute to *TRUE*. That way Zorp does not make difference between the different services (protocols) used by the client: after a successful authentication the user can use all available services without having to perform another authentication. E.g.: if this option is enabled in the example above, the client does not have to re-authenticate for starting an FTP connection.

#### 5.1.6.1. AuthCache methods

Method	Description
<i>__init__</i> ( <i>self</i> , <i>name</i> , <i>timeout</i> , <i>update_stamp</i> , <i>service_equiv</i> , <i>cleanup_threshold</i> )	Constructor to initialize an instance of the AuthCache class.

Table 5.3. Method summary

#### Method `__init__(self, name, timeout, update_stamp, service_equiv, cleanup_threshold)`

This constructor initializes and registers an AuthCache instance that can be referenced in authentication policies.

#### Arguments of `__init__`

<b>cleanup_threshold (integer)</b>
Default: 100
When the number of entries in the cache reaches the value of <i>cleanup_threshold</i> , old entries are automatically deleted.
<b>service_equiv (boolean)</b>
Default: FALSE
If enabled, then a single authentication of a user applies to every service from that client.

<b>timeout (integer)</b>
Default: 600
Timeout while an authentication is assumed to be valid.
<b>update_stamp (boolean)</b>
Default: TRUE
If set to <i>TRUE</i> , then cached authentications increase the validity period of the authentication cache. Otherwise, the authentication cache expires according to the timeout value set in <i>attribute timeout</i> (p. 92).

### 5.1.7. Class AuthenticationPolicy

Authentication policies determine how the user is authenticated to access the service. The *authentication\_policy* attribute of a service can reference an instance of the AuthenticationPolicy class.



#### Example 5.1. A simple authentication policy

The following example defines an authentication policy that can be referenced in service definitions. This policy uses inband authentication and references an *authentication provider*.

```
AuthenticationPolicy(name="demo_authentication_policy", cache=None,
authentication=InbandAuthentication(), provider="demo_authentication_provider")
```

To use the authentication policy, include it in the definition of the service:

```
Service(name="office_http_inter", proxy_class=HttpProxy,
authentication_policy="demo_authentication_policy", authorization_policy="demo_authorization_policy")
```



#### Example 5.2. Caching authentication decisions

The following example defines an authentication policy that caches the authentication decisions for ten minutes (600 seconds). For details on authentication caching, see see *Section 5.1.6, Class AuthCache* (p. 91).

```
AuthenticationPolicy(name="demo_authentication_policy", cache=AuthCache(timeout=600, update_stamp=TRUE,
service_equiv=TRUE, cleanup_threshold=100), authentication=InbandAuthentication(),
provider="demo_authentication_provider")
```

### 5.1.7.1. AuthenticationPolicy methods

Method	Description
<code>__init__(self, name, provider, authentication, cache)</code>	Constructor to initialize an instance of the AuthenticationPolicy class.

Table 5.4. Method summary

Method `__init__(self, name, provider, authentication, cache)`

Arguments of `__init__`

<b>authentication (class)</b>
Default: None
The authentication method used in the authentication process. See <i>Section 5.1.1, Authentication and authorization basics (p. 88)</i> for details.
<b>cache (class)</b>
Default: None
Caching method used to store authentication results.
<b>name (string)</b>
Default: n/a
Name identifying the AuthenticationPolicy instance.
<b>provider (class)</b>
Default: n/a
The authentication provider object used in the authentication process. See <i>Section 5.1.1, Authentication and authorization basics (p. 88)</i> for details.

### 5.1.8. Class AuthorizationPolicy

Authorization policies determine how the user is authorized to access the service. The `authorization_policy` attribute of a service can reference an instance of the AuthorizationPolicy class.



**Example 5.3. A simple authorization policy**

The following example defines an authorization policy that can be referenced in a service definition and permits only the members of the `admin` or `system` groups to access the service.

```
AuthorizationPolicy(name="demo_authorization_policy", authorization=PermitGroup(grouplist=("admin", "system")))
```

To use the authorization policy, include it in the definition of the service:

```
Service(name="office_http_inter", proxy_class=HttpProxy,
authentication_policy="demo_authentication_policy", authorization_policy="demo_authorization_policy")
```

### 5.1.8.1. AuthorizationPolicy methods

Method	Description
<code>__init__(self, name, authorization)</code>	

Table 5.5. Method summary

#### Method `__init__(self, name, authorization)`

#### Arguments of `__init__`

<b>authorization (class)</b>
Default: n/a
The authorization method (e.g., <i>PermitGroup</i> ) used in the instance. See <i>Section 5.1.8, Class AuthorizationPolicy (p. 93)</i> for examples.
<b>name (string)</b>
Default: n/a
Name of the AuthorizationPolicy instance. This name can be referenced in service definitions.

### 5.1.9. Class BasicAccessList

This class encapsulates an access list that uses any class derived from the *AbstractAuthorization* class. *BasicAccessList* allows to combine multiple access control requirements into a single decision.

*BasicAccessList* uses a list of rules. The rules are evaluated sequentially. Each rule can specify whether matching the current rule is *Sufficient* or *Required*. A connection is authorized if a *Sufficient* rule matches the connection, or all *Required* rules are fulfilled. If a *Required* rule is not met, the connection is refused.

Rules are represented as a list of Python tuples as the following example shows:



#### Example 5.4. BasicAccessList example

When referenced in a service definition, the following users can access the service:

- members of the *development* group;
- anyone with the *user1* username;
- anyone with the *user2* username.

```
AuthPolicy('intra',
authentication=ZAAAAuthentication
('zas2db', key_file='fwzaa.key', cert_file='fwzaa.crt'),
authorization=BasicAccessList(
((Z_BACL_SUFFICIENT, PermitUser('user1')),
(Z_BACL_SUFFICIENT, PermitUser('user2')),
(Z_BACL_REQUIRED, PermitGroup('development')))))
```

### 5.1.9.1. BasicAccessList methods

Method	Description
<code>__init__(self, acl)</code>	Constructor to initialize a BasicAccessList instance.

Table 5.6. Method summary

#### Method `__init__(self, acl)`

This constructor creates a new BasicAccessList instance which can be referenced in an authentication policy.

#### Arguments of `__init__`

<b>acl (complex)</b>
Default: n/a
Access control rules represented as a list of tuple.

### 5.1.10. Class InbandAuthentication

This class encapsulates inband authentication. Inband authentication is performed by the proxy using the rules of the application-level protocol. Only the authentication methods supported by the particular protocol can be used during inband authentication. *Authentication policies* can refer to instances of the InbandAuthentication class using the *auth* parameter.



#### Warning

Inband authentication is currently supported only for the Http, Ftp, and Socks proxy classes.

#### 5.1.10.1. InbandAuthentication methods

Method	Description
<code>__init__(self, authentication_provider, auth_cache)</code>	Constructor to initialize an InbandAuthentication instance.

Table 5.7. Method summary

#### Method `__init__(self, authentication_provider, auth_cache)`

This constructor initializes an instance of the InbandAuthentication class.

### 5.1.11. Class NEyesAuthorization

This class encapsulates an N-eyes based authorization method, which means that connections are authorized if other administrators authenticate themselves within the defined timelimits.

When *NEyesAuthorization* is used, the client trying to access the service has to be authorized by another (already authorized) client (this authorization chain can be expanded to multiple levels). *NEyesAuthorization* can only be used in conjunction with another *NEyesAuthorization* policy. One of them is the *authorizer* set to authorize the *authorized* policy.

In a simple 4-eyes scenario the *authorizer* policy points to the authorized policy in its *Authorization policy* parameter, and has its *wait\_authorization* parameter disabled. The *authorized* policy has an empty *Authorization policy* parameter (meaning that it is at lower the end of an N-eyes chain), and has its *wait\_authorization* parameter enabled, meaning that it has to be authorized by another policy.

For examples on using the *NEyesAuthorization* class, see the *Proxying secure channels - SSH tutorial* available from the BalaSys Documentation Page at <http://www.balasy.hu/documentation/>.

#### 5.1.11.1. NEyesAuthorization methods

Method	Description
<code>__init__(self, authorize_policy, wait_authorization, wait_timeout)</code>	Constructor to initialize a <i>NEyesAuthorization</i> instance.

Table 5.8. Method summary

#### Method `__init__(self, authorize_policy, wait_authorization, wait_timeout)`

This constructor initializes an *NEyesAuthorization* instance.

#### Arguments of `__init__`

<b>authorize_policy (class)</b>
Default: None
The authorization policy authorized by the current <i>NEyesAuthorization</i> policy.
<b>wait_authorization (boolean)</b>
Default: FALSE
Specifies whether the current authorization policy must wait for other authorization policies to finish. If this parameter is set, the client has to be authorized by another client. If set to <i>FALSE</i> , the current client is at the top of an authorizing chain.
<b>wait_timeout (integer)</b>
Default: 60000
The time (in milliseconds) Zorp will wait for the authorizing user to authorize the one accessing the service. If the other authorizations are not completed in time, the current authorization will fail.



### 5.1.12. Class PairAuthorization

This class encapsulates pair-based authorization method. Only two users simultaneously accessing the service are authorized, single users are not permitted to access the service. Set the time (in milliseconds) Zorp will wait for the second user to access the service using the `wait_timeout` parameter.



#### Example 5.5. A simple PairAuthorization policy

The following example permits access to the service only if two users having different usernames authenticate successfully within one minute.

```
AuthorizationPolicy(name="demo_pairauthorization_policy",
authorization=PairAuthorization(wait_timeout=60000))
```

For more detailed examples, see the *Proxying secure channels - SSH tutorial* available from the BalaSys Documentation Page at <http://www.balasyshu/documentation/>.

#### 5.1.12.1. PairAuthorization methods

Method	Description
<code>__init__(self, wait_timeout)</code>	Constructor to initialize a PairAuthorization instance.

Table 5.9. Method summary

#### Method `__init__(self, wait_timeout)`

This constructor initializes a PairAuthorization instance.

#### Arguments of `__init__`

<code>wait_timeout (integer)</code>
Default: 60000
The time (in milliseconds) Zorp will wait for the pair to complete the authorization. If the authorizations are not completed in time, the current authorization will fail.

### 5.1.13. Class PermitGroup

This class encapsulates an authorization decision based on group membership. Users who authenticate as a member of a usergroup specified in the policy receive access to the service. Otherwise access is denied.



#### Example 5.6. A simple PermitGroup policy

The following example permits only the members of the `admin` or `system` groups to access the service.

```
AuthorizationPolicy(name="demo_authorization_policy", authorization=PermitGroup(grouplist=("admin",
"system")))
```

### 5.1.13.1. PermitGroup methods

Method	Description
<code>__init__(self, grouplist)</code>	Constructor to initialize a PermitGroup instance.

Table 5.10. Method summary

#### Method `__init__(self, grouplist)`

This constructor initializes a PermitGroup instance.

#### Arguments of `__init__`

<b>grouplist (complex)</b>
Default: n/a
The list of authorized groups, represented as group names.

### 5.1.14. Class PermitTime

This class encapsulates an authorization decision based on the time when the connection is started. The connection is permitted if it is started in one of the permitted time periods (according to the system time of the host running Zorp).

Specify the permitted time intervals as a comma-separated list, where each element contains the beginning and ending time of the permitted interval in *HH:MM* format.



#### Example 5.7. PermitTime example

When used in the `intervals` attribute of a PermitTime instance, the following example permits access only from 07:00 to 09:00 and from 17:00 to 19:00.

```
(( "7:00", "9:00"), ("17:00", "19:00"))
```

The following is a complete authorization policy using the above intervals:

```
AuthorizationPolicy(name="demo_permittime_policy", authorization=PermitTime(intervals=(( "7:00", "9:00"), ("17:00", "19:00"))))
```

### 5.1.14.1. PermitTime methods

Method	Description
<code>__init__(self, intervals)</code>	Constructor to initialize a PermitTime instance.

Table 5.11. Method summary

#### Method `__init__(self, intervals)`

This constructor initializes a PermitTime instance.

### Arguments of `__init__`

<b>intervals (complex)</b>
Default: n/a
List of time intervals when connections are permitted (in <i>HH:MM</i> , <i>HH:MM</i> format).

### 5.1.15. Class PermitUser

This class encapsulates an authorization decision based on usernames. Users who authenticate using one of the usernames specified in the policy receive access to the service. Otherwise access is denied.



**Example 5.8. A simple PermitUser policy**

The following example permits only the *admin* and *root* users to access the service.

```
AuthorizationPolicy(name="demo_permituser", authorization=PermitUser(userlist=("admin", "root")))
```

#### 5.1.15.1. PermitUser methods

Method	Description
<code>__init__(self, userlist)</code>	Constructor to initialize a PermitUser instance.

Table 5.12. Method summary

#### Method `__init__(self, userlist)`

This constructor initializes a PermitUser instance.

### Arguments of `__init__`

<b>userlist (complex)</b>
Default: n/a
Comma-separated list of authorized usernames.

### 5.1.16. Class SatyrAuthentication

This class encapsulates outband authentication using the Satyr application. Satyr has been renamed to Zorp Authentication Agent, therefore this class is obsolete. Use ZAAAuthentication instead. See *Section 5.1.18, Class ZAAAuthentication (p. 100)* for details.

### 5.1.17. Class ServerAuthentication

This class encapsulates server authentication: Zorp authenticates the user based on the response of the server to the user's authentication request. Server authentication is a kind of inband authentication, it is performed within the application protocol, but the target server checks the credentials of the user instead of Zorp. This



authentication method is useful when the server can be trusted for authentication purposes, but you need to include an authorization decision in the service definition.

#### 5.1.17.1. ServerAuthentication methods

Method	Description
<code><u>init</u> (self)</code>	Constructor to initialize a ServerAuthentication instance.

Table 5.13. Method summary

#### Method `__init__(self)`

This constructor initializes an instance of the ServerAuthentication class.

### 5.1.18. Class ZAAAuthentication

This class encapsulates outband authentication using the Zorp Authentication Agent (ZAA). The Zorp Authentication Agent is an application that runs on the client computers and provides an interface for the users to authenticate themselves when Zorp requests authentication for accessing a service. This way any protocol, even those not supporting authentication can be securely authenticated. All communication between Zorp and ZAA is SSL-encrypted.



#### Example 5.9. Outband authentication example

The following authentication policy defines a class that uses outband authentication.

```
AuthenticationPolicy(name="demo_outbandauthentication_policy", cache=None,
authentication=ZAAAuthentication(port=1316, timeout=60000, connect_timeout=60000,
pki={"/etc/key.d/Zorp_certificate/cert.pem", "/etc/key.d/Zorp_certificate/key.pem"}),
provider="demo_authentication_provider")
```

#### 5.1.18.1. ZAAAuthentication methods

Method	Description
<code><u>init</u> (self, authentication_provider, pki, cert_file, key_file, port, timeout, connect_timeout, auth_cache)</code>	Constructor to initialize an instance of the ZAAAuthentication class.

Table 5.14. Method summary

#### Method `__init__(self, authentication_provider, pki, cert_file, key_file, port, timeout, connect_timeout, auth_cache)`

This constructor initializes an instance of the ZAAAuthentication authentication class that can be referenced in authentication policies to perform outband authentication.

## Arguments of `__init__`

<b>connect_timeout (integer)</b>
Default: 60000
Connection timeout (in milliseconds) to the Zorp Authentication Agent.
<b>pki (certificate)</b>
Default: None
A tuple containing the name of a certificate and a key file. Zorp uses this certificate to encrypt the communication with the Authentication Agents.
<b>port (integer)</b>
Default: 1316
The port number where the Zorp Authentication Agent is listening. Default value: 1316.
<b>timeout (integer)</b>
Default: 60000
Authentication timeout in milliseconds.

## 5.2. Module AuthDB

This module contains classes related to authentication databases. Together with the `Auth` module it implements the Authentication and Authorization framework. See *Section 5.1.1, Authentication and authorization basics (p. 88)* and *Section 5.1.2, Authentication and authorization in Zorp (p. 88)* for details.

### 5.2.1. Classes in the AuthDB module

Class	Description
<code><u>AbstractAuthenticationBackend</u></code>	Class encapsulating the abstract authentication backend like ZAS.
<code><u>AuthenticationProvider</u></code>	A database-independent class used by Zorp to connect to an authentication backend.
<code><u>ZAS2AuthenticationBackend</u></code>	Class encapsulating the ZAS authentication backend.

Table 5.15. Classes of the AuthDB module

### 5.2.2. Class `AbstractAuthenticationBackend`

This is an abstract class to encapsulate an authentication backend, which is responsible for checking authentication credentials against a backend database. In actual configurations, use one of the derived classes like `ZAS2AuthenticationBackend`.

The interface defined here is used by various authentication methods like [ZAAuthentication](#) and [InbandAuthentication](#).

### 5.2.3. Class AuthenticationProvider

The authentication provider is an intermediate layer that mediates between Zorp and the *authentication backend* (e.g., a user database) during connection authentication - Zorp itself does not directly communicate with the database.



**Example 5.10. A sample authentication provider**

The following example defines an authentication provider that uses the [ZAS2AuthenticationBackend](#) backend.

```
AuthenticationProvider(name="demo_authentication_provider",
backend=ZAS2AuthenticationBackend(serveraddr=SockAddrInet('192.168.10.10', 1317), use_ssl=TRUE,
ssl_verify_depth=3, pki_cert="/etc/key.d/ZAS_certificate/cert.pem",
"/etc/key.d/ZAS_certificate/key.pem"), pki_ca="/etc/ca.d/groups/demo_trusted_group/certs/",
"/etc/ca.d/groups/demo_trusted_group/crls/"))
```

#### 5.2.3.1. AuthenticationProvider methods

Method	Description
<code>__init__(self, name, backend)</code>	Constructor to initialize an AbstractAuthorizationBackend instance.

Table 5.16. Method summary

#### Method `__init__(self, name, backend)`

This constructor initializes an AbstractAuthorizationBackend instance.

#### Arguments of `__init__`

<b>backend (class)</b>
Default: n/a
Type of the database backend used by the ZAS instance.
<b>name (string)</b>
Default: n/a
Name of the ZAS instance.

### 5.2.4. Class ZAS2AuthenticationBackend

This class encapsulates a Zorp Authentication Server database and provides interface to other authentication classes to verify against users managed through ZAS. See [Section 5.2.3, Class AuthenticationProvider \(p. 102\)](#) for examples on using the ZAS2AuthenticationBackend class.

### 5.2.4.1. ZAS2AuthenticationBackend methods

Method	Description
<code>__init__(self, serveraddr, use_ssl, pki_cert, cert_file, key_file, pki_ca, ca_dir, crl_dir, ssl_verify_depth)</code>	Constructor to initialize a ZAS2AuthenticationProvider instance.

Table 5.17. Method summary

**Method `__init__(self, serveraddr, use_ssl, pki_cert, cert_file, key_file, pki_ca, ca_dir, crl_dir, ssl_verify_depth)`**

This constructor creates a new ZAS2AuthenticationProvider instance that can be used in authentication policies.

#### Arguments of `__init__`

<b>pki_ca (cagroup)</b>
Default: None
The name of a trusted CA group. When using SSL, ZAS must show a certificate signed by a CA that belongs to this group.
<b>pki_cert (certificate)</b>
Default: None
A tuple containing the name of a certificate and a key file. Zorp shows this certificate to ZAS when using SSL.
<b>serveraddr (sockaddr)</b>
Default: n/a
The IP address of this ZAS instance. ZAS accepts connections on this address.
<b>ssl_verify_depth (integer)</b>
Default: 3
Specifies the maximum number of CAs in the trust chain when verifying the certificate of Zorp.
<b>use_ssl (boolean)</b>
Default: FALSE
Enable this option if Zorp communicates with ZAS using SSL.

## 5.3. Module Chainer

Chainers establish a TCP or UDP connection between a proxy and a selected destination. The destination is usually a server, but the [SideStackChainer](#) connects an additional proxy before connecting the server.



### 5.3.1. Selecting the network protocol

The client-side and the server-side connections can use different networking protocols if needed. The *protocol* attribute of the chainer classes determines the network protocol used in the server-side connection. By default, the same protocol is used in both connections. The following options are available:

Name	Description
ZD_PROTO_AUTO	Use the protocol that is used on the client side.
ZD_PROTO_TCP	Use the TCP protocol on the server side.
ZD_PROTO_UDP	Use the UDP protocol on the server side.

Table 5.18. The network protocol used in the server-side connection

### 5.3.2. Classes in the Chainer module

Class	Description
<u><i>AbstractChainer</i></u>	Class encapsulating the abstract chainer.
<u><i>AvailabilityChainer</i></u>	This class enables establishing connection with multiple target addresses and using information from the Availability Checker daemon. <i>AvailabilityChainer</i> connects to target hosts in the order they have been specified.
<u><i>ConnectChainer</i></u>	Class to establish the server-side TCP/IP connection.
<u><i>FailoverChainer</i></u>	Class encapsulating the connection establishment with multiple target addresses and keeping down state between connects. <i>FailoverChainer</i> prefers connecting to target hosts in the order they were specified.
<u><i>MultiTargetChainer</i></u>	Class encapsulating connection establishment with multiple target addresses.
<u><i>RoundRobinAvailabilityChainer</i></u>	This class enables establishing connection with multiple target addresses and using information from the Availability Checker daemon.
<u><i>RoundRobinChainer</i></u>	Class encapsulating the connection establishment with multiple target addresses and keeping down state between connects.
<u><i>SideStackChainer</i></u>	Class to pass the traffic to another proxy.



Class	Description
<i>StateBasedChainer</i>	Class encapsulating connection establishment with multiple target addresses and keeping down state between connects.

Table 5.19. Classes of the Chainer module

### 5.3.3. Class AbstractChainer

AbstractChainer implements an abstract chainer that establishes a connection between the parent proxy and the selected destination. This class serves as a starting point for customized chainer classes, but is itself not directly usable. Service definitions should refer to a customized class derived from AbstractChainer, or one of the predefined chainer classes, such as *ConnectChainer* or *FailoverChainer*.

### 5.3.4. Class AvailabilityChainer

This class is based on the *MultiTargetChainer* class and encapsulates a real TCP/IP connection establishment. It is used when a top-level proxy wants to perform chaining. In addition to ConnectChainer, this class adds the capability to perform stateful failover HA functionality across a set of IP addresses.



#### Note

Use AvailabilityChainer if you want to connect to servers, the availability of which have been checked by the Availability Checker daemon monitoring them. Hosts which are in *Up* state are attempted to be connected.



#### Example 5.11. A DirectedRouter using AvailabilityChainer

The following service definition uses a DirectedRouter class with two possible destination addresses. The firewall uses these destinations in a failover fashion, targeting the second address only if the first one is marked unavailable by the Availability Checker daemon.

```
Service(name="intra_HTTP_inter", router=DirectedRouter(dest_addr=(SockAddrInet('192.168.55.55', 8080),
SockAddrInet('192.168.55.56', 8080)), forge_addr=FALSE, forge_port=Z_PORT_ANY, overrideable=FALSE),
chainer=AvailabilityChainer(protocol=ZD_PROTO_AUTO, timeout_connect=30000), max_instances=0,
proxy_class=HttpProxy,)
```

#### 5.3.4.1. AvailabilityChainer methods

Method	Description
<i>__init__</i> ( <i>self</i> , <i>protocol</i> , <i>timeout_connect</i> )	Constructor to initialize a AvailabilityChainer instance.

Table 5.20. Method summary

#### Method `__init__(self, protocol, timeout_connect)`

This constructor initializes a AvailabilityChainer class by filling arguments with appropriate values and calling the inherited constructor.

### Arguments of `__init__`

<b>protocol (enum)</b>
Default: <code>ZD_PROTO_AUTO</code>
Optional, specifies connection protocol ( <code>ZD_PROTO_TCP</code> or <code>ZD_PROTO_UDP</code> ), when not specified it defaults to the protocol used on the client side.
<b>timeout_connect (integer)</b>
Default: 30000
Specifies connection timeout to be used when connecting to the target server.

### 5.3.5. Class ConnectChainer

ConnectChainer is the default chainer class based on AbstractChainer. This class establishes a TCP or UDP connection between the proxy and the selected destination address.

ConnectChainer is used by default if no other chainer class is specified in the service definition.

ConnectChainer attempts to connect only a single destination address: if the connection establishment procedure selects multiple target servers (e.g., a *DNSResolver* with the `multi=TRUE` parameter or a *DirectedRouter* with multiple addresses), ConnectChainer will use the first address and ignore all other addresses. Use *FailoverChainer* to select from the destination from multiple addresses in a failover fashion, and *RoundRobinChainer* to distribute connections in a roundrobin fashion.



#### Example 5.12. A sample ConnectChainer

The following service uses a ConnectChainer that uses the UDP protocol on the server side.

```
Service(name="demo_service", proxy_class=HttpProxy, chainer=ConnectChainer(protocol=ZD_PROTO_UDP),
router=TransparentRouter(overrideable=FALSE, forge_addr=FALSE))
```

#### 5.3.5.1. ConnectChainer methods

Method	Description
<code>__init__(self, protocol, timeout_connect)</code>	Constructor to initialize an instance of the ConnectChainer class.

Table 5.21. Method summary

#### Method `__init__(self, protocol, timeout_connect)`

This constructor creates a new ConnectChainer instance which can be associated with a *Service*.

## Arguments of `__init__`

<b>protocol (enum)</b>
Default: <code>ZD_PROTO_AUTO</code>
Optional parameter that specifies the network protocol used in the connection protocol. By default, the server-side communication uses the same protocol that is used on the client side. See <i>Section 5.3.1, Selecting the network protocol (p. 104)</i> for details.
<b>timeout_connect (integer)</b>
Default: 30000
Specifies connection timeout to be used when connecting to the target server.

### 5.3.6. Class FailoverChainer

This class is based on the *StateBasedChainer* class and encapsulates a real TCP/IP connection establishment, and is used when a top-level proxy wants to perform chaining. In addition to *ConnectChainer* this class adds the capability to perform stateful, failover HA functionality across a set of IP addresses.



#### Note

Use *FailoverChainer* if you want to connect to the servers in a predefined order: i.e., connect to the first server, and only connect to the second if the first server is unavailable.

If you want to distribute connections between the servers (i.e., direct every new connection to a different server to balance the load) use *RoundRobinChainer*.



#### Example 5.13. A DirectedRouter using FailoverChainer

The following service definition uses a *DirectedRouter* class with two possible destination addresses. These destinations are used in a failover fashion, targeting the second address only if the first one is inaccessible.

```
Service(name="intra_HTTP_inter", router=DirectedRouter(dest_addr=(SockAddrInet('192.168.55.55', 8080),
SockAddrInet('192.168.55.56', 8080)), forge_addr=FALSE, forge_port=Z_PORT_ANY, overrideable=FALSE),
chainer=FailoverChainer(protocol=ZD_PROTO_AUTO, timeout_state=60000, timeout_connect=30000),
max_instances=0, proxy_class=HttpProxy,)
```

#### 5.3.6.1. FailoverChainer methods

Method	Description
<code>__init__(self, protocol, timeout, timeout_state, timeout_connect, round_robin)</code>	Constructor to initialize a <i>FailoverChainer</i> instance.

Table 5.22. Method summary

#### Method `__init__(self, protocol, timeout, timeout_state, timeout_connect, round_robin)`

This constructor initializes a *FailoverChainer* class by filling arguments with appropriate values and calling the inherited constructor.

### Arguments of `__init__`

<b>protocol (enum)</b>
Default: <code>ZD_PROTO_AUTO</code>
Optional, specifies connection protocol ( <code>ZD_PROTO_TCP</code> or <code>ZD_PROTO_UDP</code> ), when not specified it defaults to the protocol used on the client side.
<b>timeout_connect (integer)</b>
Default: 30000
Specifies connection timeout to be used when connecting to the target server.
<b>timeout_state (integer)</b>
Default: 60000
The down state of remote hosts is kept for this interval in milliseconds.

### 5.3.7. Class MultiTargetChainer

This class encapsulates a real TCP/IP connection establishment, and is used when a top-level proxy wants to perform chaining. In addition to `ConnectChainer`, this class adds the capability to perform stateless, simple load balance server connections among a set of IP addresses.

The same mechanism is used to set multiple server addresses as with a single destination address: the `Router` class sets a list of IP addresses in the `session.target_address` attribute.

#### 5.3.7.1. MultiTargetChainer methods

Method	Description
<code>__init__(self, protocol, timeout_connect)</code>	Constructor to initialize a <code>MultiTargetChainer</code> instance.

Table 5.23. Method summary

#### Method `__init__(self, protocol, timeout_connect)`

This constructor initializes a `MultiTargetChainer` class by filling arguments with appropriate values and calling the inherited constructor.

### Arguments of `__init__`

<b>protocol (enum)</b>
Default: <code>ZD_PROTO_AUTO</code>
Optional, specifies connection protocol (either <code>ZD_PROTO_TCP</code> or <code>ZD_PROTO_UDP</code> ), when not specified defaults to the same protocol as was used on the client side.

<b>self (class)</b>
Default: n/a
this instance

<b>timeout_connect (integer)</b>
Default: 30000
Specifies connection timeout to be used when connecting to the target server.

### 5.3.8. Class RoundRobinAvailabilityChainer

This class is based on the [AvailabilityChainer](#) class.



**Note**

Use RoundRobinAvailabilityChainer if you want to connect to servers, the availability of which have been checked by the Availability Checker daemon monitoring them. Hosts which are in *Up* state are attempted to be connected. In addition to AvailabilityChainer, this class adds the capability to perform stateful load balance server connections among a set of IP addresses.



**Example 5.14. A DirectedRouter using RoundRobinAvailabilityChainer**

The following service definition uses a DirectedRouter class with two possible destination addresses. The firewall uses these destinations in a failover fashion, targeting the second address only if the first one is marked unavailable by the Availability Checker daemon.

```
Service(name="intra_HTTP_inter", router=DirectedRouter(dest_addr=(SockAddrInet('192.168.55.55', 8080),
SockAddrInet('192.168.55.56', 8080)), forge_addr=FALSE, forge_port=Z_PORT_ANY, overrideable=FALSE),
chainer=RoundRobinAvailabilityChainer(protocol=ZD_PROTO_AUTO, timeout_connect=30000), max_instances=0,
proxy_class=HttpProxy,)
```

#### 5.3.8.1. RoundRobinAvailabilityChainer methods

Method	Description
<code>__init__(self, protocol, timeout_connect)</code>	Constructor to initialize a RoundRobinAvailabilityChainer instance.

Table 5.24. Method summary

#### Method `__init__(self, protocol, timeout_connect)`

This constructor initializes a RoundRobinAvailabilityChainer class by filling arguments with appropriate values and calling the inherited constructor.

#### Arguments of `__init__`

<b>protocol (enum)</b>
Default: ZD_PROTO_AUTO

**protocol (enum)**

Optional, specifies connection protocol ( *ZD\_PROTO\_TCP* or *ZD\_PROTO\_UDP* ), when not specified it defaults to the protocol used on the client side.

**timeout\_connect (integer)**

Default: 30000

Specifies connection timeout to be used when connecting to the target server.

### 5.3.9. Class RoundRobinChainer

This class is based on the *StateBasedChainer* class and encapsulates a real TCP/IP connection establishment, and is used when a top-level proxy wants to perform chaining. In addition to *ConnectChainer* this class adds the capability to perform stateful, load balance server connections among a set of IP addresses.


**Example 5.15. A DirectedRouter using RoundRobinChainer**

The following service definition uses a *RoundRobinChainer* class with two possible destination addresses. These destinations are used in a roundrobin fashion, alternating between the two destinations.

```
Service(name="intra_HTTP_inter", router=DirectedRouter(dest_addr=(SockAddrInet('192.168.55.55', 8080),
SockAddrInet('192.168.55.56', 8080)), forge_addr=FALSE, forge_port=Z_PORT_ANY, overrideable=FALSE),
chainer=RoundRobinChainer(protocol=ZD_PROTO_AUTO, timeout_state=60000, timeout_connect=30000),
max_instances=0, proxy_class=HttpProxy)
```

### 5.3.10. Class SideStackChainer

This class encapsulates a special chainer. Instead of establishing a connection to a server, it creates a new proxy instance and connects the server side of the current (parent) proxy to the client side of the new (child) proxy. The *right\_class* parameter specifies the child proxy.

It is possible to stack multiple proxies side-by-side. The final step of sidestacking is always to specify a regular chainer via the *right\_chainer* parameter that connects the last proxy to the destination server.

**Tip**

Proxy sidestacking is useful for example to create one-sided SSL connections. See the tutorials of the BalaSys Documentation Page available at <http://www.balasy.hu/documentation/> for details.

#### 5.3.10.1. Attributes of SideStackChainer

**right\_chainer (unknown)**

Default: n/a

The chainer used to connect to the destination of the side-stacked proxy class set in the *right\_class* attribute.

<b>right_class (unknown)</b>
Default: n/a
The proxy class to connect to the parent proxy. Both built-in and customized classes can be used.

### 5.3.10.2. SideStackChainer methods

Method	Description
<code>__init__(self, right_class, right_chainer)</code>	Constructor to initialize an instance of the SideStackChainer class.

Table 5.25. Method summary

#### Method `__init__(self, right_class, right_chainer)`

This constructor creates a new FailoverChainer instance which can be associated with a *Service*.

#### Arguments of `__init__`

<b>right_chainer (class)</b>
Default: None
The chainer used to connect to the destination of the side-stacked proxy class set in the <i>right_class</i> attribute.

<b>right_class (class)</b>
Default: n/a
The proxy class to connect to the parent proxy. Both built-in or customized classes can be used.

### 5.3.11. Class StateBasedChainer

This class encapsulates a real TCP/IP connection establishment, and is used when a top-level proxy wants to perform chaining. In addition to ConnectChainer, this class adds the capability to perform stateful, load balance server connections among a set of IP addresses.



#### Note

Both the *FailoverChainer* and *RoundRobinChainer* classes are derived from StateBasedChainer.

### 5.3.11.1. StateBasedChainer methods

Method	Description
<u><code>__init__(self, protocol, timeout_connect, timeout_state)</code></u>	Constructor to initialize a StateBasedChainer instance.

Table 5.26. Method summary

#### Method `__init__(self, protocol, timeout_connect, timeout_state)`

This constructor initializes a StateBasedChainer class by filling arguments with appropriate values and calling the inherited constructor.

#### Arguments of `__init__`

<b>protocol (enum)</b>
Default: ZD_PROTO_AUTO
Optional, specifies connection protocol ( <code>ZD_PROTO_TCP</code> or <code>ZD_PROTO_UDP</code> ), when not specified it defaults to the same protocol used on the client side.

<b>timeout_connect (integer)</b>
Default: 30000
Specifies connection timeout to be used when connecting to the target server.

<b>timeout_state (integer)</b>
Default: 60000
The down state of remote hosts is kept for this interval in milliseconds.

## 5.4. Module Detector

Detectors can be used to determine if the traffic in the incoming connection uses a particular protocol (for example, HTTP, SSH), or if it has other specific characteristics (for example, it uses SSL encryption with a specific certificate). Such characteristics of the traffic can be detected, and start a specific service to inspect the traffic (for example, start a specific HttpProxy for HTTP traffic, and so on).

### 5.4.1. Classes in the Detector module

Class	Description
<u><code>AbstractDetector</code></u>	Class encapsulating the abstract detector.
<u><code>CertDetector</code></u>	Class encapsulating a Detector that determines if an SSL/TLS-encrypted connection uses the specified certificate



Class	Description
<i>DetectorPolicy</i>	Class encapsulating a Detector which can be used by a name.
<i>HttpDetector</i>	Class encapsulating a Detector that determines if the traffic uses the HTTP protocol
<i>SniDetector</i>	Class encapsulating a Detector that determines whether a client targets a specific host in a SSL/TLS-encrypted connection.
<i>SshDetector</i>	Class encapsulating a Detector that determines if the traffic uses the SSHv2 protocol

Table 5.27. Classes of the Detector module

### 5.4.2. Class AbstractDetector

This abstract class encapsulates a detector that determines whether the traffic in a connection belongs to a particular protocol.

### 5.4.3. Class CertDetector

This Detector determines if an SSL/TLS-encrypted connection uses the specified certificate, and rejects any other protocols and certificates.



#### Example 5.16. CertDetector example

The following example defines a DetectorPolicy that detects if the traffic is SSL/TLS-encrypted, and uses the certificate specified.

```

mycertificate="-----BEGIN CERTIFICATE-----
MIIEdjCCA16gAwIBAgIIQ7Xu3Mwnk+4wDQYJKoZIhvcNAQEFBQAwSTELMAkGA1UE
BhMCVVMxZzARBgNVBAoTCKdvb2dsZS5BbmMxJTAjBgNVBAMTHEdvd2dsZS5BbnRl
cm5ldCBDbXRob3JpdHkgRzIwHhcNMjQwMTI5MTQwNTM3WhcNMjQwNTI5MDAwMDAw
WjBoMQswCQYDVQQGEwJVUzETMBEGA1UECAwKQ2FsaWZvcml5TEwMBQGA1UEBwwN
TW91bnRhaw4gVmllZDZETMBEGA1UECgwKR29vZ2x1IEIuYzEXMBUGA1UEAwOd3d3
Lmdvb2dsZS5jb20wggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQCkeHm
eYY7uMMRxKg14NPx8zFtD/VmUI2b4FdQYgD8AuRiFA+fqvxicEki7Td1SrZ4z1dn
AjbAS+fc0eQji8foJTosrkXgQgv5ds0+81U3dooVXoqemeJKUihzI/h+7cf1287/
7EbmI5RaDBUPThmZHeDtK38XUYsBrS93nICq4VDUAxy2BKsGSS219wRv14fhDDm
guQ5cRDkn/ppdYEqAqxFVEjamwjcUWSBsWlqSn37fI9s/MZDCzfMwz6AheFMrRNL
0oJ2Y3cVdBxiDvdqjGS+AG5qIUz/AsvHNL3JEsa550SRMFubCPCzYDMAVLKziqZX
5G25c0e/qh0bSK4/AgMBAAGjggFBMIIBPTAdBgNVHUSUEFjAUBggrBgEFBQcDAQYI
KwYBBQUHAwIwGQYDVR0RBBIwEII0d3d3Lmdvb2dsZS5jb20waAYIKwYBBQUHAQEE
XDBAMCSGCCsGAQUFBzAChh9odHRwOi8vcGtpLmdvb2dsZS5jb20vR01BRzIuY3J0
MCSGCCsGAQUFBzAChh9odHRwOi8vY2xpZW50czEuZ29vZ2x1LmNvbS9vY3NwMB0G
A1UdDgQWBRR1IOrR+bm3NNXp5DwKruhknMrpDAMBGNVHRMBAF8EAJAAMB8GA1Ud
IwYMBaAFerDbHbvPZotXb1gba7Yhq6WoEvMBCGA1UdIAQQMA4wDAYKKwYBBAHW
eQIFATAwBgNVHR8EKTAnMCwgI6Ahhh9odHRwOi8vcGtpLmdvb2dsZS5jb20vR01B
RzIuY3J0SMA0GCSqGSIb3DQEBAQUAA4IBAQA6j9oPKE5k/FX5sbLY4p7xsn1tndHD
N1oyzmb8+cmke6w/eFHsY0g+zUeUBW3zb0EMBNXWNTCB1aVICrGe8GUDDAnAZSX
MQBeBisNb69kn2untS7Rb1L83+8H787RsLeXucahr3kCoc61oTemI0HEI430DtVI
uFEDNjDE1wqsHkdZecNNS29IZySpK2skr3rH7qUkbP1lkzbFvsnFUyp3AJS4ib9+
4xPr65GQfUi/8vgoSVvOy5Y3rT/U3CtI9tPoDSZTYGTl64LDxJa8dEGYmTKHgjyJ
HmbKzes13N/BN18XUlvTnjEaifQXvJj9ypqcMHUFPjkqW1IHSyb1iRth
-----END CERTIFICATE-----"
DetectorPolicy(name="MyCertDetector", detector=CertDetector(certificate=mycertificate)

```

### 5.4.3.1. Attributes of CertDetector

<b>certificate (unknown)</b>
Default: n/a
The certificate to detect in PEM format. You can use the certificate directly, or store it in a file and reference the file with full path, for example, <code>DetectorPolicy(name="MyCertDetector", detector=CertDetector(certificate="/etc/key.d/mysite/cert.pem", ))</code>

### 5.4.3.2. CertDetector methods

Method	Description
<code>__init__(self, certificate)</code>	Constructor to initialize a CertDetector instance.

Table 5.28. Method summary

#### Method `__init__(self, certificate)`

This constructor initializes a CertDetector instance

#### Arguments of `__init__`

<b>certificate (certificate)</b>
Default: n/a
The certificate in PEM format. This must contain either the certificate as a string, or a full pathname to a file containing the certificate.

### 5.4.4. Class DetectorPolicy

DetectorPolicy instances are reusable detectors that contain configured instances of the detector classes (for example, HttpDetector, SshDetector) that detect if the traffic uses a particular protocol, or a particular certificate in an SSL/TLS connection. DetectorPolicy instances can be used in the `detect` option of firewall rules. For examples, see the specific detector classes.

### 5.4.5. Class HttpDetector

This Detector determines if the traffic uses the HTTP protocol, and rejects any other protocol.



#### Example 5.17. HttpDetector example

The following example defines a DetectorPolicy that detects HTTP traffic.

```
DetectorPolicy(name="http", detector=HttpDetector())
```

### 5.4.5.1. Attributes of HttpDetector

<b>ignore (unknown)</b>
Default: n/a
A list of compiled regular expressions which should be ignored when detecting the traffic type. By default, this list is empty.
<b>match (unknown)</b>
Default: n/a
A list of compiled regular expressions which result in a positive match. If the traffic matches this regular expression, it is regarded as HTTP traffic. Default value: <code>[OPTIONS GET HEAD POST PUT DELETE TRACE CONNECT] + ".*HTTP/1."</code>

### 5.4.5.2. HttpDetector methods

Method	Description
<code>__init__(self, **kw)</code>	Constructor to initialize a HttpDetector instance.

Table 5.29. Method summary

#### Method `__init__(self, **kw)`

This constructor initializes a HttpDetector instance

### 5.4.6. Class SniDetector

Class encapsulating a Detector that determines whether a client targets a specific host in a SSL/TLS-encrypted connection and rejects any other protocols and hostnames.



#### Example 5.18. SniDetector example

The following example defines a DetectorPolicy that detects if the traffic is SSL/TLS-encrypted, and uses targets the host `www.example.com`.

```
DetectorPolicy(name="MySniDetector",
detector=SniDetector(RegexMatcher(match_list=("www.example.com",))))
```

### 5.4.6.1. Attributes of SniDetector

<b>server_name_matcher (class)</b>
Default: n/a
Matcher class (e.g.: RegexMatcher) used to check and filter hostnames in Server Name Indication TLS extension, for example, <code>DetectorPolicy(name="MySniDetector", detector=SniDetector(RegexMatcher(match_list=("www.example.com",))))</code>

### 5.4.6.2. SniDetector methods

Method	Description
<code>__init__(self, server_name_matcher)</code>	Constructor to initialize a SniDetector instance.

Table 5.30. Method summary

#### Method `__init__(self, server_name_matcher)`

This constructor initializes a SniDetector instance

#### Arguments of `__init__`

<code>server_name_matcher</code> (class)
Default: n/a
Matcher class (e.g.: <code>RegexMatcher</code> ) used to check and filter hostnames in Server Name Indication TLS extension.

### 5.4.7. Class SshDetector

This Detector determines if the traffic uses the SSHv2 protocol, and rejects any other protocol.



#### Example 5.19. SshDetector example

The following example defines a DetectorPolicy that detects SSH traffic.

```
DetectorPolicy(name="ssh", detector=SshDetector())
```

## 5.5. Module Encryption

The SSL/TLS framework of the proxies is in a separate entity called Encryption policy. That way, you can easily share and reuse encryption settings between different services: you have to configure the Encryption policy once, and you can use it in multiple services. The SSL framework is described in *Chapter 3, The Zorp SSL framework* (p. 9).



#### Note

STARTTLS support is currently available only for the Ftp proxy to support FTPS sessions and for the SMTP and the Pop3 proxies.

### 5.5.1. SSL parameter constants

Name	Value
<code>SSL_CIPHERS_HIGH</code>	n/a
<code>SSL_CIPHERS_MEDIUM</code>	n/a
<code>SSL_CIPHERS_LOW</code>	n/a



Name	Value
SSL_CIPHERS_ALL	n/a
SSL_CIPHERS_CUSTOM	n/a

Table 5.31. Constants for cipher selection

Name	Value
TLSV1_3_CIPHERS_DEFAULT	n/a
TLSV1_3_CIPHERS_CUSTOM	n/a

Table 5.32. Constants for TLSv1.3 cipher selection

Name	Value
TLS_SHARED_GROUPS_DEFAULT	n/a
TLS_SHARED_GROUPS_CUSTOM	n/a

Table 5.33. Constants for shared group selection

Name	Value
SSL_HSO_CLIENT_SERVER	Perform the SSL-handshake with the client first.
SSL_HSO_SERVER_CLIENT	Perform the SSL-handshake with the server first.

Table 5.34. Handshake order.

Name	Value
SSL_NONE	Disable encryption between Zorp and the peer.
SSL_FORCE_SSL	Require encrypted communication between Zorp and the peer.
SSL_ACCEPT_STARTTLS	Permit STARTTLS sessions. Currently supported only in the Ftp, SmtP and Pop3 proxies.

Table 5.35. Client connection security type.

Name	Value
SSL_NONE	Disable encryption between Zorp and the peer.
SSL_FORCE_SSL	Require encrypted communication between Zorp and the peer.
SSL_FORWARD_STARTTLS	Forward STARTTLS requests to the server. Currently supported only in the Ftp, SmtP and Pop3 proxies.

Table 5.36. Server connection security type.

Name	Value
TLS_TRUST_LEVEL_NONE	Accept invalid for example, expired certificates.

Name	Value
TLS_TRUST_LEVEL_UNTRUSTED	Both trusted and untrusted certificates are accepted.
TLS_TRUST_LEVEL_FULL	Only valid certificates signed by a trusted CA are accepted.

Table 5.37. Constants for trust level selection.

Name	Value
TLS_INTERMEDIATE_REVOCATION_NONE	Ignore result of CA certificate revocation status check.
TLS_INTERMEDIATE_REVOCATION_SOFT_FAIL	Check every CA certificate revocation state in the certificate chain. Uncertainty is tolerated.
TLS_INTERMEDIATE_REVOCATION_HARD_FAIL	Check every CA certificate revocation state in the certificate chain. Uncertainty is not tolerated.

Table 5.38. Constants for intermediate certificates revocation check type.

Name	Value
TLS_LEAF_REVOCATION_NONE	Ignore result of leaf certificate revocation status check.
TLS_LEAF_REVOCATION_SOFT_FAIL	Check the revocation state of the leaf certificate. Uncertainty is tolerated.
TLS_LEAF_REVOCATION_HARD_FAIL	Check the revocation state of the leaf certificate. Uncertainty is not tolerated.

Table 5.39. Constants for leaf certificate revocation check type.

Name	Value
SSL_ERROR	n/a
SSL_DEBUG	n/a

Table 5.40. Verbosity level of the log messages

Name	Value
SSL_HS_ACCEPT	0
SSL_HS_REJECT	1
SSL_HS_POLICY	6
SSL_HS_VERIFIED	10

Table 5.41. Handshake policy decisions

### 5.5.2. Classes in the Encryption module

Class	Description
<i>AbstractVerifier</i>	Class encapsulating the abstract Certificate verifier.



<b>Class</b>	<b>Description</b>
<u><a href="#">Certificate</a></u>	Class encapsulating a certificate and its private key, and optionally the passphrase for the private key.
<u><a href="#">CertificateCA</a></u>	Class encapsulating the certificate of a Certificate Authority (CA certificate) and its private key, and optionally the passphrase for the private key.
<u><a href="#">ClientCertificateVerifier</a></u>	Class that can be used to verify the certificate of the client-side connection.
<u><a href="#">ClientNoneVerifier</a></u>	Disables certificate verification in client-side connection.
<u><a href="#">ClientOnlyEncryption</a></u>	The ClientOnlyEncryption class handles scenarios when only the client-Zorp connection is encrypted, the Zorp-server connection is not
<u><a href="#">ClientOnlyStartTLEncryption</a></u>	The client can optionally request STARTTLS encryption, but the server-side connection is always unencrypted.
<u><a href="#">ClientSSLOptions</a></u>	Class encapsulating a set of SSL options used in the client-side connection.
<u><a href="#">DHParam</a></u>	Class encapsulating DH parameters.
<u><a href="#">DynamicCertificate</a></u>	Class to perform SSL keybridging.
<u><a href="#">DynamicServerEncryption</a></u>	The DynamicServerEncryption class handles scenarios when both the client-firewall and the firewall-server connections could be encrypted but the server side encryption parameters set dynamically from proxies.
<u><a href="#">EncryptionPolicy</a></u>	Class encapsulating a named set of encryption settings.
<u><a href="#">FakeStartTLEncryption</a></u>	The client can optionally request STARTTLS encryption, but the server-side connection is always encrypted.
<u><a href="#">ForwardStartTLEncryption</a></u>	The ForwardStartTLEncryption class handles scenarios when the client can optionally request STARTTLS encryption.
<u><a href="#">PrivateKey</a></u>	Class encapsulating a private key.
<u><a href="#">SNIBasedCertificate</a></u>	Class to be used for Server Name Indication (SNI)
<u><a href="#">SSLOptions</a></u>	Class encapsulating the abstract SSL options.
<u><a href="#">ServerCertificateVerifier</a></u>	Class that can be used to verify the certificate of the server-side connection.
<u><a href="#">ServerNoneVerifier</a></u>	Disables certificate verification in server-side connection.

Class	Description
<u><a href="#">ServerOnlyEncryption</a></u>	The ServerOnlyEncryption class handles scenarios when only the Zorp-server connection is encrypted, the client-Zorp connection is not
<u><a href="#">ServerSSLOptions</a></u>	Class encapsulating a set of SSL options used in the server-side connection.
<u><a href="#">StaticCertificate</a></u>	Class encapsulating a static Certificate object.
<u><a href="#">TwoSidedEncryption</a></u>	The TwoSidedEncryption class handles scenarios when both the client-Zorp and the Zorp-server connections are encrypted.

Table 5.42. Classes of the Encryption module

### 5.5.3. Class AbstractVerifier

This class includes the settings and options used to verify the certificates of the peers in SSL and TLS connections. Note that you cannot use this class directly, use an appropriate derived class, for example, [ClientCertificateVerifier](#) or [ServerCertificateVerifier](#) instead.

#### 5.5.3.1. Attributes of AbstractVerifier

<b>intermediate_revocation_check_type (enum)</b>
Default: TLS_INTERMEDIATE_REVOCATION_SOFT_FAIL
Specify how intermediate certificates revocation status check should work.
<b>leaf_revocation_check_type (enum)</b>
Default: TLS_LEAF_REVOCATION_SOFT_FAIL
Specify how leaf certificate revocation status check should work.
<b>required (boolean)</b>
Default: trusted
If the <i>required</i> is TRUE, a certificate is required from the peer.
<b>trust_level (enum)</b>
Default: TLS_TRUST_LEVEL_FULL
Specify which certificate should be accepted as trusted.
<b>trusted_certs_directory (string)</b>
Default: ""





<b>trusted_certs_directory (string)</b>
A directory where trusted IP address - certificate assignments are stored. When a peer from a specific IP address shows the certificate stored in this directory, it is accepted regardless of its expiration or issuer CA. Each file in the directory should contain a certificate in PEM format. The filename must be the IP address.
<b>verify_ca_directory (string)</b>
Default: ""
Directory where the trusted CA certificates are stored. CA certificates are loaded on-demand from this directory when the certificate of the peer is verified.
<b>verify_crl_directory (string)</b>
Default: ""
Directory where the CRLs (Certificate Revocation Lists) associated with trusted CAs are stored. CRLs are loaded on-demand from this directory when the certificate of the peer is verified.
<b>verify_depth (integer)</b>
Default: 4
The length of the longest accepted CA verification chain. Longer CA chains are automatically rejected.

### 5.5.3.2. AbstractVerifier methods

Method	Description
<i><code>__init__(self, trust_level, intermediate_revocation_check_type, leaf_revocation_check_type, trusted_certs_directory, required, verify_depth, verify_ca_directory, verify_crl_directory)</code></i>	Constructor to initialize an AbstractVerifier instance.

Table 5.43. Method summary

**Method `__init__(self, trust_level, intermediate_revocation_check_type, leaf_revocation_check_type, trusted_certs_directory, required, verify_depth, verify_ca_directory, verify_crl_directory)`**

This constructor defines an AbstractVerifier with the specified parameters.

#### Arguments of `__init__`

<b>intermediate_revocation_check_type (enum)</b>
Default: TLS_INTERMEDIATE_REVOCATION_SOFT_FAIL
Specify how intermediate certificates revocation status check should work.



<b>leaf_revocation_check_type (enum)</b>
Default: TLS_LEAF_REVOCATION_SOFT_FAIL
Specify how leaf certificate revocation status check should work.
<b>required (boolean)</b>
Default: TRUE
If the <i>required</i> is TRUE, a certificate is required from the peer.
<b>trust_level (enum)</b>
Default: TLS_TRUST_LEVEL_FULL
Specify which certificate should be accepted as trusted.
<b>trusted_certs_directory (string)</b>
Default: ""
A directory where trusted IP address - certificate assignments are stored. When a peer from a specific IP address shows the certificate stored in this directory, it is accepted regardless of its expiration or issuer CA. Each file in the directory should contain a certificate in PEM format. The filename must be the IP address.
<b>verify_ca_directory (string)</b>
Default: ""
Directory where the trusted CA certificates are stored. CA certificates are loaded on-demand from this directory when the certificate of the peer is verified.
<b>verify_crl_directory (string)</b>
Default: ""
Directory where the CRLs (Certificate Revocation Lists) associated with trusted CAs are stored. CRLs are loaded on-demand from this directory when the certificate of the peer is verified.
<b>verify_depth (integer)</b>
Default: 4
The length of the longest accepted CA verification chain. Longer CA chains are automatically rejected.

#### 5.5.4. Class Certificate

The Certificate class stores a certificate, its private key, and optionally a passphrase for the private key. The certificate must be in PEM format.



When configuring Zorp manually using its configuration file, use the regular constructor of the Certificate class to load a certificate from a string. To load a certificate from a file, use the *Certificate.fromFile* method.



**Example 5.20. Loading a certificate**

The following example loads a certificate from the configuration file.

```
my_certificate = "-----BEGIN CERTIFICATE-----
MIICUTCCAfugAwIBAgIBADANBgkqhkiG9w0BAQQFADBXMQswCQYDVQQGEwJDTjEL
MAKGA1UECBMCUE4xCzAJBgNVBACtAkNOMQswCQYDVQQKEwJPTjELMAKGA1UECXM
VU4xFDASBgNVBAMTC0hlcm9uZyBZYW5nMB4XDTA1MDcxNTIxMTkON1oXDTA1MDgx
NDIxMTkON1owVzELMAKGA1UEBhMCQ04xCzAJBgNVBAGTA1B0MQswCQYDVQQHEwJD
TjELMAKGA1UEChMCT04xCzAJBgNVBAsTA1VOMRQwEgYDVQQDEwtIZXJvbmcgWwFu
ZzBcMAOGCSqGSIb3DQEBAQUAA0sAMEgCQQCp5hnG7ogBhtLynpOS21cBewKE/B7j
V14qeyslnr26xZUsSVko36Znhia0/zbM0oRcKK9vEcgMtcLFuQTWD13RagMBAAGj
gbEwga4wHQYDVR00BBYEFFXI70krXeQDxZgbaCQoR4jUDncEMH8GA1UdIwR4MHaa
FFXI70krXeQDxZgbaCQoR4jUDncEoVukWTBXMQswCQYDVQQGEwJDTjELMAKGA1UE
CBMCUE4xCzAJBgNVBACtAkNOMQswCQYDVQQKEwJPTjELMAKGA1UECXMVCU4xFDAS
BgNVBAMTC0hlcm9uZyBZYW5nggEAMAwGA1UdEwQFMAMBAF8wDQYJKoZIhvcNAQEE
BQADQQA/ugzBrjjK9jcwDvFgHlk3icNRq0oV7Ri32z/+HQX67aRfgZu7KwdI+Ju
Wm7DCfrPNGVvFWUQ0msPue9rZBgO
-----END CERTIFICATE-----"
my_certificate_object = Certificate(my_certificate, 'mypassphrase')
```

The following example loads a certificate from an external file.

```
my_certificate_object = Certificate.fromFile("/tmp/my_certificate.pem", 'mypassphrase')
```

**5.5.4.1. Attributes of Certificate**

<b>certificate_file_path (certificatechain)</b>
Default: n/a
The path and filename to the certificate file. The certificate must be in PEM format.
<b>private_key_password (string)</b>
Default: None
Passphrase used to access the private key of the certificate specified in <i>certificate_file_path</i> .

**5.5.4.2. Certificate methods**

Method	Description
<i>__init__(self, certificate, private key)</i>	Load a certificate from a string, and access it using its passphrase
<i>fromFile(certificate file path, private key)</i>	Load a certificate from a file, and access it using its passphrase

Table 5.44. Method summary

**Method `__init__(self, certificate, private_key)`**

Initializes a Certificate instance by loading a certificate from a string, and accesses it using its passphrase. To load a certificate from a file, use the *Certificate.fromFile* method.

### Arguments of `__init__`

<b>certificate_file_path (certificate)</b>
Default: n/a
The path and filename to the certificate file. The certificate must be in PEM format.
<b>private_key_password (string)</b>
Default: None
Passphrase used to access the private key of the certificate specified in <i>certificate_file_path</i> .

### Method `fromFile(certificate_file_path, private_key)`

Initializes a Certificate instance by loading a certificate from a file, and accesses it using its passphrase.

### Arguments of `fromFile`

<b>certificate_file_path (certificate)</b>
Default: n/a
The path and filename to the certificate file. The certificate must be in PEM format.
<b>passphrase (string)</b>
Default: None
Passphrase used to access the private key specified in <i>certificate_file_path</i> .

## 5.5.5. Class CertificateCA

The CertificateCA class stores a CA certificate, its private key, and optionally a passphrase for the private key. The certificate must be in PEM format.

### 5.5.5.1. Attributes of CertificateCA

<b>certificate_file_path (certificate)</b>
Default: n/a
The path and filename to the certificate file. The certificate must be in PEM format.
<b>private_key_password (string)</b>
Default: None
Passphrase used to access the private key of the certificate specified in <i>certificate_file_path</i> .



### 5.5.5.2. CertificateCA methods

Method	Description
<code>__init__(self, certificate, private_key)</code>	Load a CA certificate from a string, and access it using its passphrase

Table 5.45. Method summary

#### Method `__init__(self, certificate, private_key)`

Initializes a CertificateCA instance by loading a CA certificate, and accesses it using its passphrase.

#### Arguments of `__init__`

<b>certificate_file_path (certificate)</b>
Default: n/a
The path and filename to the CA certificate file. The certificate must be in PEM format.
<b>private_key_password (string)</b>
Default: None
Passphrase used to access the private key specified in <i>certificate_file_path</i> .

### 5.5.6. Class ClientCertificateVerifier

This class includes the settings and options used to verify the certificates of the peers in client-side SSL and TLS connections.

#### 5.5.6.1. Attributes of ClientCertificateVerifier

<b>ca_hint_directory (string)</b>
Default: ""
Set directory containing certificates to provide the client the list of CA certificates (subject names) that are used for verifying the client certificate.
<b>intermediate_revocation_check_type (enum)</b>
Default: TLS_INTERMEDIATE_REVOCATION_SOFT_FAIL
Specify how intermediate certificates revocation status check should work.
<b>leaf_revocation_check_type (enum)</b>
Default: TLS_LEAF_REVOCATION_SOFT_FAIL
Specify how leaf certificate revocation status check should work.



<b>required (boolean)</b>
Default: TRUE
If the <i>required</i> is TRUE, a certificate is required from the peer.
<b>trust_level (enum)</b>
Default: TLS_TRUST_LEVEL_FULL
Specify which certificate should be accepted as trusted.
<b>trusted_certs_directory (string)</b>
Default: ""
A directory where trusted IP address - certificate assignments are stored. When a peer from a specific IP address shows the certificate stored in this directory, it is accepted regardless of its expiration or issuer CA. Each file in the directory should contain a certificate in PEM format. The filename must be the IP address.
<b>verify_ca_directory (string)</b>
Default: ""
Directory where the trusted CA certificates are stored. CA certificates are loaded on-demand from this directory when the certificate of the peer is verified.
<b>verify_crl_directory (string)</b>
Default: ""
Directory where the CRLs (Certificate Revocation Lists) associated with trusted CAs are stored. CRLs are loaded on-demand from this directory when the certificate of the peer is verified.
<b>verify_depth (integer)</b>
Default: 4
The length of the longest accepted CA verification chain. Longer CA chains are automatically rejected.

### 5.5.6.2. ClientCertificateVerifier methods

Method	Description
<code>__init__(self, trust_level, intermediate_revocation_check_type, leaf_revocation_check_type, trusted_certs_directory, required, verify_depth, verify_ca_directory, verify_crl_directory, ca_hint_directory)</code>	Constructor to initialize a ClientCertificateVerifier instance.

Table 5.46. Method summary

**Method `__init__(self, trust_level, intermediate_revocation_check_type, leaf_revocation_check_type, trusted_certs_directory, required, verify_depth, verify_ca_directory, verify_crl_directory, ca_hint_directory)`**

This constructor defines a ClientCertificateVerifier with the specified parameters.

#### Arguments of `__init__`

<b><code>ca_hint_directory</code> (string)</b>
Default: ""
Set directory containing certificates to provide the client the list of CA certificates (subject names) that are used for verifying the client certificate.
<b><code>intermediate_revocation_check_type</code> (enum)</b>
Default: TLS_INTERMEDIATE_REVOCATION_SOFT_FAIL
Specify how intermediate certificates revocation status check should work.
<b><code>leaf_revocation_check_type</code> (enum)</b>
Default: TLS_LEAF_REVOCATION_SOFT_FAIL
Specify how leaf certificate revocation status check should work.
<b><code>required</code> (boolean)</b>
Default: TRUE
If the <i>required</i> is TRUE, a certificate is required from the peer.
<b><code>trust_level</code> (enum)</b>
Default: TLS_TRUST_LEVEL_FULL
Specify which certificate should be accepted as trusted.



<b>trusted_certs_directory (string)</b>
Default: ""
A directory where trusted IP address - certificate assignments are stored. When a peer from a specific IP address shows the certificate stored in this directory, it is accepted regardless of its expiration or issuer CA. Each file in the directory should contain a certificate in PEM format. The filename must be the IP address.
<b>verify_ca_directory (string)</b>
Default: ""
Directory where the trusted CA certificates are stored. CA certificates are loaded on-demand from this directory when the certificate of the peer is verified.
<b>verify_crl_directory (string)</b>
Default: ""
Directory where the CRLs (Certificate Revocation Lists) associated with trusted CAs are stored. CRLs are loaded on-demand from this directory when the certificate of the peer is verified.
<b>verify_depth (integer)</b>
Default: 4
The length of the longest accepted CA verification chain. Longer CA chains are automatically rejected.

### 5.5.7. Class ClientNoneVerifier

This class disables every certificate verification in client-side SSL and TLS connections.

### 5.5.8. Class ClientOnlyEncryption

The ClientOnlyEncryption class handles scenarios when only the client-Zorp connection is encrypted, the Zorp-server connection is not.

#### 5.5.8.1. Attributes of ClientOnlyEncryption

<b>client_certificate_generator (class)</b>
Default: n/a
The class that will generate the certificate that will be showed to the client. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.
<b>client_ssl_options (class)</b>
Default: ClientSSLOptions()
The protocol-level encryption settings used on the client side. This must be a <i>ClientSSLOptions</i> instance.



<b>client_verify (class)</b>
Default: ClientCertificateVerifierGroup()
The settings used to verify the certificate of the client. This must be a <i>ClientCertificateVerifier</i> instance.

### 5.5.8.2. ClientOnlyEncryption methods

Method	Description
<code>__init__(self, client_certificate_generator, client_verify, client_ssl_options)</code>	Initializes SSL/TLS connection on the client side.

Table 5.47. Method summary

#### Method `__init__(self, client_certificate_generator, client_verify, client_ssl_options)`

The ClientOnlyEncryption class handles scenarios when only the client-Zorp connection is encrypted, the Zorp-server connection is not.

#### Arguments of `__init__`

<b>client_certificate_generator (class)</b>
Default: n/a
The class that will generate the certificate that will be showed to the client. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.

<b>client_ssl_options (class)</b>
Default: ClientSSLOptions()
The protocol-level encryption settings used on the client side. This must be a <i>ClientSSLOptions</i> instance.

<b>client_verify (class)</b>
Default: ClientCertificateVerifierGroup()
The settings used to verify the certificate of the client. This must be a <i>ClientCertificateVerifier</i> instance.

### 5.5.9. Class ClientOnlyStartTLSEncryption

The ClientOnlyStartTLSEncryption class handles scenarios when the client can optionally request STARTTLS encryption. If the client sends a STARTTLS request, the client-side connection will use STARTTLS. The server-side connection will not be encrypted.



#### Warning

If the client does not send a STARTTLS request, the client-side communication will not be encrypted at all. The server-side connection will never be encrypted.

### 5.5.9.1. Attributes of ClientOnlyStartTLSEncryption

<b>client_certificate_generator (class)</b>
Default: n/a
The class that will generate the certificate that will be showed to the client. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.
<b>client_ssl_options (class)</b>
Default: ClientSSLOptions()
The protocol-level encryption settings used on the client side. This must be a <i>ClientSSLOptions</i> instance.
<b>client_verify (class)</b>
Default: ClientCertificateVerifierGroup()
The settings used to verify the certificate of the client. This must be a <i>ClientCertificateVerifier</i> instance.

### 5.5.9.2. ClientOnlyStartTLSEncryption methods

Method	Description
<i>__init__(self, client_certificate_generator, client_verify, client_ssl_options)</i>	The client can optionally request STARTTLS encryption, but the server-side connection is always unencrypted.

Table 5.48. Method summary

#### Method `__init__(self, client_certificate_generator, client_verify, client_ssl_options)`

The ClientOnlyStartTLSEncryption class handles scenarios when the client can optionally request STARTTLS encryption. If the client sends a STARTTLS request, the client-side connection will use STARTTLS. The server-side connection will not be encrypted.



**Warning**

If the client does not send a STARTTLS request, the client-side communication will not be encrypted at all. The server-side connection will never be encrypted.

#### Arguments of `__init__`

<b>client_certificate_generator (class)</b>
Default: n/a
The class that will generate the certificate that will be showed to the client. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.



<b>client_ssl_options (class)</b>
Default: ClientSSLOptions()
The protocol-level encryption settings used on the client side. This must be a <i>ClientSSLOptions</i> instance.

<b>client_verify (class)</b>
Default: ClientCertificateVerifier()
The settings used to verify the certificate of the client. This must be a <i>ClientCertificateVerifier</i> instance.

### 5.5.10. Class ClientSSLOptions

This class (based on the SSLOptions class) collects the TLS and SSL settings directly related to encryption, for example, the permitted protocol versions, ciphers, session reuse settings, and so on.

#### 5.5.10.1. Attributes of ClientSSLOptions

<b>cipher (enum)</b>
Default: n/a
Specifies the allowed ciphers. For details, see <i>Table 5.31, Constants for cipher selection (p. 116)</i> .

<b>cipher_server_preference (boolean)</b>
Default: FALSE
Use server and not client preference order when determining which cipher suite, signature algorithm or elliptic curve to use for an incoming connection.

<b>ciphers_tlsv1_3 (enum)</b>
Default: n/a
Specifies the allowed ciphers for TLSv1.3 connections. For details, see <i>Table 5.32, Constants for TLSv1.3 cipher selection (p. 117)</i> .

<b>dh_params (dhparams)</b>
Default: None
The DH parameter used by ephemeral DH key generation. Please be mind that this option is ignored in TLSv1.3 as it does not support custom DH parameters.

<b>disable_compression (boolean)</b>
Default: FALSE
Set this to TRUE to disable support for SSL/TLS compression even if it is supported. Please be mind that this option is ignored in TLSv1.3 as it does not support compression.



<b>disable_send_root_ca (boolean)</b>
Default: False
Inhibit sending Root CA to client, even if present in local certificate chain.
<b>disable_session_cache (boolean)</b>
Default: FALSE
Do not store session information in the session cache. Set this option to TRUE to disable SSL session reuse. Please be mind that this option is ignored in TLSv1.3 as it does not support session IDs.
<b>disable_ticket (boolean)</b>
Default: FALSE
Session tickets are a method for SSL session reuse, described in RFC 5077. Set this option to TRUE to disable SSL session reuse using session tickets.
<b>disable_tlsv1 (boolean)</b>
Default: TRUE
Do not allow using TLSv1 in the connection.
<b>disable_tlsv1_1 (boolean)</b>
Default: FALSE
Do not allow using TLSv1.1 in the connection.
<b>disable_tlsv1_2 (boolean)</b>
Default: FALSE
Do not allow using TLSv1.2 in the connection.
<b>disable_tlsv1_3 (boolean)</b>
Default: FALSE
Do not allow using TLSv1.3 in the connection.
<b>session_cache_size (integer)</b>
Default: 20480
The number of sessions stored in the session cache for SSL session reuse. Please be mind that this option is ignored in TLSv1.3 as it does not support session IDs.

<b>shared_groups (enum)</b>
Default: n/a
Specifies the allowed shared groups. For details, see <i>Table 5.33, Constants for shared group selection (p. 117)</i> .

### 5.5.10.2. ClientSSLOptions methods

Method	Description
<code>__init__(self, method, cipher, ciphers_tls1_3, shared_groups, cipher_server_preference, timeout, disable_sslv2, disable_sslv3, disable_tls1, disable_tls1_1, disable_tls1_2, disable_tls1_3, session_cache_size, disable_session_cache, disable_ticket, disable_compression, dh_params, disable_renegotiation, disable_send_root_ca)</code>	Constructor to initialize a ClientSSLOptions instance.

Table 5.49. Method summary

**Method `__init__(self, method, cipher, ciphers_tls1_3, shared_groups, cipher_server_preference, timeout, disable_sslv2, disable_sslv3, disable_tls1, disable_tls1_1, disable_tls1_2, disable_tls1_3, session_cache_size, disable_session_cache, disable_ticket, disable_compression, dh_params, disable_renegotiation, disable_send_root_ca)`**

This constructor defines a ClientSSLOptions with the specified parameters.

#### Arguments of `__init__`

<b>cipher (enum)</b>
Default: n/a
Specifies the allowed ciphers. For details, see <i>Table 5.31, Constants for cipher selection (p. 116)</i> .

<b>cipher_server_preference (boolean)</b>
Default: FALSE
Use server and not client preference order when determining which cipher suite, signature algorithm or elliptic curve to use for an incoming connection.

<b>ciphers_tls1_3 (enum)</b>
Default: n/a
Specifies the allowed ciphers for TLSv1.3 connections. For details, see <i>Table 5.32, Constants for TLSv1.3 cipher selection (p. 117)</i> .



<b>dh_param_file_path (string)</b>
Default: None
The path and filename to the DH parameter file. The DH parameter file must be in PEM format. Please be mind that this option is ignored in TLSv1.3 as it does not support custom DH parameters.
<b>disable_compression (boolean)</b>
Default: FALSE
Set this to TRUE to disable support for SSL/TLS compression even if it is supported. Please be mind that this option is ignored in TLSv1.3 as it does not support compression.
<b>disable_renegotiation (boolean)</b>
Default: TRUE
Set this to TRUE to disable client initiated renegotiation. Please be mind that this option is ignored in TLSv1.3 as it does not support renegotiation.
<b>disable_send_root_ca (boolean)</b>
Default: FALSE
Set this to TRUE to inhibit sending root ca to client, even if present in local chain.
<b>disable_session_cache (boolean)</b>
Default: FALSE
Do not store session information in the session cache. Set this option to TRUE to disable SSL session reuse. Please be mind that this option is ignored in TLSv1.3 as it does not support session IDs.
<b>disable_ticket (boolean)</b>
Default: FALSE
Session tickets are a method for SSL session reuse, described in RFC 5077. Set this option to TRUE to disable SSL session reuse using session tickets.
<b>disable_tlsv1 (boolean)</b>
Default: TRUE
Do not allow using TLSv1 in the connection.
<b>disable_tlsv1_1 (boolean)</b>
Default: FALSE
Do not allow using TLSv1.1 in the connection.

<b>disable_tlsv1_2 (boolean)</b>
Default: FALSE
Do not allow using TLSv1.2 in the connection.
<b>disable_tlsv1_3 (boolean)</b>
Default: FALSE
Do not allow using TLSv1.3 in the connection.
<b>session_cache_size (integer)</b>
Default: 20480
The number of sessions stored in the session cache for SSL session reuse. Please be mind that this option is ignored in TLSv1.3 as it does not support session IDs.
<b>shared_groups (enum)</b>
Default: n/a
Specifies the allowed shared groups. For details, see <i>Table 5.33, Constants for shared group selection (p. 117)</i> .
<b>timeout (integer)</b>
Default: 300
Drop idle connection if the timeout value (in seconds) expires.

### 5.5.11. Class DHParam

The DHParam class stores DH parameters. The DH parameters must be in PEM format.

When configuring Zorp manually using its configuration file, use the regular constructor of the DHParam class to load DH parameters key from a string. To load DH parameters key from a file, use the *DHParam.fromFile* method.



#### Example 5.21. Loading DH parameters

The following example loads DH parameters from the configuration file.

```
my_dh_params = "-----BEGIN DH PARAMETERS-----
MIIBCACkAQEAvvO8WguTntkDs33qe5u1T7Ij1l1mTrRnwFV4z7W4A0Du9j+prdRdD
UAb1HYBrQn30FsfG/6wDVTmUj8Lvgn9aFjwYTe6U3Ey7CQt4MBw2BhC03R19KDw7
Im8UdBbXuekuqZGifmKEEFzAcbiQepvBXiGMucDwgbLaaTY/FrKqb509DvoenSV
Aj/VNFnsefQTHXGo1Urg8ixawj7kTNhM3x7kj7BhK4ALfBuv/93aet2SQjU207C6
Oj3mku8CD93Xsbng6rIzmRd6pCANEFH0Rgo10X7+vMwwG5h5YDsF8cVAcRroZkxR
dyPdVNzY1z1X3Jx1n3It/6F2yyx/FOXAGwIBAg==
-----END DH PARAMETERS-----"
my_dh_params_object = DHParam(my_dh_params)
```

The following example loads DH parameters key from an external file.

```
my_dh_params_object = DHParam.fromFile("/tmp/my_dh_params.pem")
```

### 5.5.11.1. Attributes of DHParam

<b>params (string)</b>
Default: ""
The path and filename to the DH parameters file. The DH parameters must be in PEM format.

### 5.5.11.2. DHParam methods

Method	Description
<code>__init__(self, params)</code>	Load DH parameters key from a string
<code>fromFile(file_path)</code>	Load a DH parameters from a file

Table 5.50. Method summary

#### Method `__init__(self, params)`

Initializes a DHParam instance by loading DH parameters key from a string. To load a DH parameters from a file, use the `DHParam.fromFile` method.

#### Arguments of `__init__`

<b>params (certificate)</b>
Default: n/a
The path and filename to the DH parameters file. The DH parameters must be in PEM format.

#### Method `fromFile(file_path)`

Initializes a DHParam instance by loading a DH parameters from a file.

#### Arguments of `fromFile`

<b>file_path (dhparam)</b>
Default: n/a
The path and filename to the DH parameters file. The DH parameters must be in PEM format.

## 5.5.12. Class DynamicCertificate

This class is able to generate certificates mimicking another certificate, primarily used to transfer the information of a server's certificate to the client in keybridging. Can be used only in *TwoSidedEncryption*. For details on configuring keybridging, see *Section 3.2.7, Keybridging certificates (p. 16)*.



### 5.5.12.1. DynamicCertificate methods

Method	Description
<code>__init__(self, private_key, trusted_ca, untrusted_ca, cache_directory, extension_whitelist)</code>	Initializes a DynamicCertificate instance to use for keybridging

Table 5.51. Method summary

Method `__init__(self, private_key, trusted_ca, untrusted_ca, cache_directory, extension_whitelist)`

#### Arguments of `__init__`

<b>cache_directory (string)</b>
Default: None
The cache directory to store the keybridged generated certificates, for example, <code>/var/lib/zorp/sslbridge/</code> . The <code>zorp</code> user must have write privileges for this directory.
<b>extension_whitelist (complex)</b>
Default: None
<b>private_key (class)</b>
Default: n/a
The private key of the CA certificate set in <code>trusted_ca</code>
<b>trusted_ca (class)</b>
Default: n/a
The CA certificate that will used to sign the keybridged certificate of trusted peers.
<b>untrusted_ca (class)</b>
Default: n/a
The CA certificate that will used to sign the keybridged certificate of untrusted peers.

### 5.5.13. Class DynamicServerEncryption

The DynamicServerEncryption class handles scenarios when both the client-firewall and the firewall-server connections could be encrypted but the server side encryption parameters set dynamically from proxies.

### 5.5.13.1. Attributes of DynamicServerEncryption

<b>client_certificate_generator (class)</b>
Default: n/a
The class that will generate the certificate that will be showed to the client. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.
<b>client_security (enum)</b>
Default: n/a
Set security mode.
<b>client_ssl_options (class)</b>
Default: ClientSSLOptions()
The protocol-level encryption settings used on the client side. This must be a <i>ClientSSLOptions</i> instance.
<b>client_verify (class)</b>
Default: ClientCertificateVerifierGroup()
The settings used to verify the certificate of the client. This must be a <i>ClientCertificateVerifier</i> instance.

### 5.5.13.2. DynamicServerEncryption methods

Method	Description
<i>__init__(self, client_security, client_certificate_generator, client_verify, client_ssl_options)</i>	Initializes SSL/TLS connection on the client side.

Table 5.52. Method summary

#### Method `__init__(self, client_security, client_certificate_generator, client_verify, client_ssl_options)`

The DynamicServerEncryption class handles scenarios when both the client-firewall and the firewall-server connections could be encrypted but the server side encryption parameters set dynamically from proxies.

#### Arguments of `__init__`

<b>client_certificate_generator (class)</b>
Default: n/a
The class that will generate the certificate that will be showed to the client. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.

<b>client_ssl_options (class)</b>
Default: ClientSSLOptions()
The protocol-level encryption settings used on the client side. This must be a <i>ClientSSLOptions</i> instance.

<b>client_verify (class)</b>
Default: ClientCertificateVerifierGroup()
The settings used to verify the certificate of the client. This must be a <i>ClientCertificateVerifier</i> instance.

### 5.5.14. Class EncryptionPolicy

This class encapsulates a named set of encryption settings and an associated Encryption policy instance. Encryption policies provide a way to re-use encryption settings without having to define encryption settings for each service individually.

#### 5.5.14.1. Attributes of EncryptionPolicy

<b>encryption (class)</b>
Default: n/a
<p>An encryption scenario instance that will be used in the Encryption Policy. This describes the scenario and the settings how encryption is used in the scenario, for example:</p> <ul style="list-style-type: none"> <li>■ Both the client-side and the server-side connections are encrypted (<i>TwoSidedEncryption</i>)</li> <li>■ Only the client-side connection is encrypted (<i>ClientOnlyEncryption</i>)</li> <li>■ Only the server-side connection is encrypted (<i>ServerOnlyEncryption</i>)</li> <li>■ STARTTLS is enabled (<i>ClientOnlyStartTLSEncryption</i>, <i>FakeStartTLSEncryption</i>, or <i>ForwardStartTLSEncryption</i>)</li> </ul> <p>To customize the settings of a scenario (for example, to set the used certificates), derive a class from the selected scenario, set its parameters as needed for your environment, and use the customized class.</p>

<b>name (string)</b>
Default: n/a
Name identifying the Encryption policy.

### 5.5.14.2. EncryptionPolicy methods

Method	Description
<code>__init__(self, name, encryption)</code>	Constructor to create an Encryption policy.

Table 5.53. Method summary

#### Method `__init__(self, name, encryption)`

This constructor initializes an Encryption policy, based on the settings of the *encryption* parameter. This describes the scenario and the settings how encryption is used in the scenario, for example:

- Both the client-side and the server-side connections are encrypted (*TwoSidedEncryption*)
- Only the client-side connection is encrypted (*ClientOnlyEncryption*)
- Only the server-side connection is encrypted (*ServerOnlyEncryption*)
- STARTTLS is enabled (*ClientOnlyStartTLSEncryption*, *FakeStartTLSEncryption*, or *ForwardStartTLSEncryption*)

To customize the settings of a scenario (for example, to set the used certificates), derive a class from the selected scenario, set its parameters as needed for your environment, and use the customized class.

#### Arguments of `__init__`

<b>encryption (class)</b>
Default: n/a
An encryption scenario instance that will be used in the Encryption Policy.
<b>name (string)</b>
Default: n/a
Name identifying the Encryption policy.

### 5.5.15. Class FakeStartTLSEncryption

The FakeStartTLSEncryption class handles scenarios when the client can optionally request STARTTLS encryption. If the client sends a STARTTLS request, the client-side connection will use STARTTLS. The server-side connection will always be encrypted.

**Warning**

If the client does not send a STARTTLS request, the client-side communication will not be encrypted at all. The server-side connection will always be encrypted.

### 5.5.15.1. Attributes of FakeStartTLSEncryption

<b>client_certificate_generator (class)</b>
Default: n/a
The class that will generate the certificate that will be showed to the client. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.
<b>client_ssl_options (class)</b>
Default: ClientSSLOptions()
The protocol-level encryption settings used on the client side. This must be a <i>ClientSSLOptions</i> instance.
<b>client_verify (class)</b>
Default: ClientCertificateVerifierGroup()
The settings used to verify the certificate of the client. This must be a <i>ClientCertificateVerifier</i> instance.
<b>server_ssl_options (class)</b>
Default: ServerSSLOptions()
The protocol-level encryption settings used on the server side. This must be a <i>ServerSSLOptions</i> instance.
<b>server_verify (class)</b>
Default: ServerCertificateVerifierGroup()
The settings used to verify the certificate of the server. This must be a <i>ServerCertificateVerifier</i> instance.

### 5.5.15.2. FakeStartTLSEncryption methods

Method	Description
<i>__init__</i> ( <i>self</i> , <i>client_certificate_generator</i> , <i>client_verify</i> , <i>server_verify</i> , <i>client_ssl_options</i> , <i>server_ssl_options</i> )	Initializes a FakeStartTLSEncryption instance to handle scenarios when the client can optionally request STARTTLS encryption.

Table 5.54. Method summary

#### Method `__init__(self, client_certificate_generator, client_verify, server_verify, client_ssl_options, server_ssl_options)`

The FakeStartTLSEncryption class handles scenarios when the client can optionally request STARTTLS encryption. If the client sends a STARTTLS request, the client-side connection will use STARTTLS. The server-side connection will always be encrypted.

**Warning**

If the client does not send a STARTTLS request, the client-side communication will not be encrypted at all. The server-side connection will always be encrypted.

**Arguments of `__init__`**

<b>client_certificate_generator (class)</b>
Default: n/a
The class that will generate the certificate that will be showed to the client. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.
<b>client_ssl_options (class)</b>
Default: ClientSSLOptions()
The protocol-level encryption settings used on the client side. This must be a <i>ClientSSLOptions</i> instance.
<b>client_verify (class)</b>
Default: ClientCertificateVerifierGroup()
The settings used to verify the certificate of the client. This must be a <i>ClientCertificateVerifier</i> instance.
<b>server_ssl_options (class)</b>
Default: ServerSSLOptions()
The protocol-level encryption settings used on the server side. This must be a <i>ServerSSLOptions</i> instance.
<b>server_verify (class)</b>
Default: ServerCertificateVerifierGroup()
The settings used to verify the certificate of the server. This must be a <i>ServerCertificateVerifier</i> instance.

**5.5.16. Class ForwardStartTLSEncryption**

The ForwardStartTLSEncryption class handles scenarios when the client can optionally request STARTTLS encryption. If the client sends a STARTTLS request, the client-side connection will use STARTTLS, and the request will be forwarded to the server. If the server supports STARTTLS, the server-side connection will also use STARTTLS.

**Warning**

If the client does not send a STARTTLS request, the communication will not be encrypted at all. Both the client-Zorp and the Zorp-server connections will be unencrypted.

### 5.5.16.1. Attributes of ForwardStartTLSEncryption

<b>client_certificate_generator (class)</b>
Default: n/a
The class that will generate the certificate that will be showed to the client. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.
<b>client_ssl_options (class)</b>
Default: ClientSSLOptions()
The protocol-level encryption settings used on the client side. This must be a <i>ClientSSLOptions</i> instance.
<b>client_verify (class)</b>
Default: ClientCertificateVerifierGroup()
The settings used to verify the certificate of the client. This must be a <i>ClientCertificateVerifier</i> instance.
<b>server_ssl_options (class)</b>
Default: ServerSSLOptions()
The protocol-level encryption settings used on the server side. This must be a <i>ServerSSLOptions</i> instance.
<b>server_verify (class)</b>
Default: ServerCertificateVerifierGroup()
The settings used to verify the certificate of the server. This must be a <i>ServerCertificateVerifier</i> instance.

### 5.5.16.2. ForwardStartTLSEncryption methods

Method	Description
<i>__init__(self, client_certificate_generator, client_verify, server_verify, client_ssl_options, server_ssl_options)</i>	Initializes a ForwardStartTLSEncryption instance to handle scenarios when the client can optionally request STARTTLS encryption.

Table 5.55. Method summary

#### Method `__init__(self, client_certificate_generator, client_verify, server_verify, client_ssl_options, server_ssl_options)`

Initializes a ForwardStartTLSEncryption instance to handle scenarios when the client can optionally request STARTTLS encryption. If the client sends a STARTTLS request, the client-side connection will use STARTTLS, and the request will be forwarded to the server. If the server supports STARTTLS, the server-side connection will also use STARTTLS.

**Warning**

If the client does not send a STARTTLS request, the communication will not be encrypted at all. Both the client-Zorp and the Zorp-server connections will be unencrypted.

**Arguments of `__init__`**

<b>client_certificate_generator (class)</b>
Default: n/a
The class that will generate the certificate that will be showed to the client. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.
<b>client_ssl_options (class)</b>
Default: ClientSSLOptions()
The protocol-level encryption settings used on the client side. This must be a <i>ClientSSLOptions</i> instance.
<b>client_verify (class)</b>
Default: ClientCertificateVerifierGroup()
The settings used to verify the certificate of the client. This must be a <i>ClientCertificateVerifier</i> instance.
<b>server_ssl_options (class)</b>
Default: ServerSSLOptions()
The protocol-level encryption settings used on the server side. This must be a <i>ServerSSLOptions</i> instance.
<b>server_verify (class)</b>
Default: ServerCertificateVerifierGroup()
The settings used to verify the certificate of the server. This must be a <i>ServerCertificateVerifier</i> instance.

**5.5.17. Class PrivateKey**

The PrivateKey class stores a private key and optionally a passphrase for the private key. The private key must be in PEM format.

When configuring Zorp manually using its configuration file, use the regular constructor of the PrivateKey class to load a private key from a string. To load a private key from a file, use the *PrivateKey.fromFile* method.

**Example 5.22. Loading a private key**

The following example loads a private key from the configuration file.

```
my_private_key = "-----BEGIN RSA PRIVATE KEY-----
MIIEPgIBAAKCAQEAg9rbxqq+Zi70nRFAZe7SCTB6VgzP1PhkiUmOPmbwFmR0S1SSy
yMPSyIzaQqwELY0SQTZtsT3jhd6MCFPBZntym63/GwDuethGSjE9y8rt/9yr+T3I
zz+6ABnZXHJ38tdGYataF1Ndi3CsY5NXGszVFv1Is17P5mbYWQgJ7QzI/a5mPKa+
9pVxSDQthEV3BVUawIEJjNs0THD5XZQJ/MX6F4RPn+2MC9i/RbcAORVnLPmt2eiy
```



```
NV3+55sKdd7GpdMmEbRv9HZyW2xJNyU1xYbwU9YIP88dHCgvqo0gkAX2HLxCJ0y6
2gvsS8J7HEbohD98dxPJX7P8w9juORi6Hpsq0wIDAQABAOIBAQDXStIdJtuRC+GG
RXfXca/6iP3j3qV2KSzATRe+CkvAR0o1CC9T7z6zb+bPI5kLIblxwVPIJaW0nn4I
jj5JFhTVMaIagTeaz7yw5d2NR2r1SkZwW7Au2uePSv9ZIzL1IVLzzDnz/PW2xv5I
brOmT/Tr+N9GV8iIwNqu5sryp60FasKB/55LhCcKVYrkdy2WhJc8Y8TXUjF4n8Jn
Xuyd44N6uu5RUIEgn7bPsz01F1T8ujCICwDnNyuW91wSVvEC2EbTg841u2UcnE4k
grB7rCKLooDpY1KjXx/1o9Dj9Uv3hwLpStw2dYRoZS0kOFIKYACP1QcininrTGeL
cOPxyK6BAoGBAPvnBd7/U94Krp9Bp3jjxUEnlFrgf+B7QgRKpG7tN3RDRJmIVL8Z
mnxvbW6o4hsq4TzF/ratnRjqp+79Tw5wUz36G98ftwL1TUS620BznIkwImDGo+ysv
3QK8XUZ4Wg3EcnE5bG8Am0KoDRazc0g7UxopbHC+SNLRMZA/2dBvVh4zAoGBAPq6
UWIfcSnLyFYy7EPh3P7qmotBNPORGcX6aKdwR7pz6MqTADHxKvIP+eeDEWpF58T
RYBw7KxN4h6cNmglRZBbhED3h0NjKpYMGsq0hyczn40SIHHRf3iB07p35v7Eee82
2H/rT6BNrQF1fPIbz5spgT+eV5BuTAB7bsbwiuDhAoGBALVAgeT26y21mfhVkv9W
5LQA+qp5JworJlFYnAdtBx3M2StwASqQDazDsIYTVr4dmHvWk3Teb09iaPt5oMz0
3daWhd+D3VCv98FtM+r4FKGI/Zmd8Twd8HTrfGIcbw/A7mex3efxEHDkwqY28Rhk
N2N3sucNcx6GJjJQynVNxCRIpAoGBA0JyIEqUxyn0iP0BLm3osiXxUP7wN5i8FA7w
qFCBUecNt4uoCdiyk+fqbF10evT3UQQ07ZKJ71t3RAANAIZTU06buQjMBFMbAa90
4fP19BLtaQcAHh+HCcuX3I/+9rumS9JHIKX3qoTHYrdsMx03D/u9MqR4p/EkDLRq
xpQC9I9BAoGBAPZtXtEKc0xhYeuor4qIQbt1edr0+cfeZaxYUvjlLdG8rU3Yeh3
JLbYgcSNr4rMvEwhuvwbwgWJjed7TvqjKKEYYSW2ESwcmAjNIhDBVzX9oh1cY34
Ae/P630Ht89sWbb5oG2+fcb7xCwH3kYmVgT4/xPv0FQRspwErKYlCWg
-----END RSA PRIVATE KEY-----"
my_private_key_object = PrivateKey(my_private_key, 'mypassphrase')
```

The following example loads a private key from an external file.

```
my_private_key_object = PrivateKey.fromFile("/tmp/my_private.key", 'mypassphrase')
```

### 5.5.17.1. Attributes of PrivateKey

<b>key_file_path (string)</b>
Default: ""
The path and filename to the private key file. The private key must be in PEM format.
<b>passphrase (string)</b>
Default: None
Passphrase used to access the private key specified in <i>key_file_path</i> .

### 5.5.17.2. PrivateKey methods

Method	Description
<i>__init__(self, key, passphrase)</i>	Load a private key from a string, and access it using its passphrase
<i>fromFile(key_file_path, passphrase)</i>	Load a private key from a file, and access it using its passphrase

Table 5.56. Method summary

#### Method `__init__(self, key, passphrase)`

Initializes a `PrivateKey` instance by loading a private key from a string, and accesses it using its passphrase. To load a private key from a file, use the `PrivateKey.fromFile` method.



### Arguments of `__init__`

<b>key_file_path (certificate)</b>
Default: n/a
The path and filename to the private key file. The private key must be in PEM format.
<b>passphrase (string)</b>
Default: None
Passphrase used to access the private key specified in <i>key_file_path</i> .

### Method `fromFile(key_file_path, passphrase)`

Initializes a PrivateKey instance by loading a private key from a file, and accesses it using its passphrase.

### Arguments of `fromFile`

<b>key_file_path (certificate)</b>
Default: n/a
The path and filename to the private key file. The private key must be in PEM format.
<b>passphrase (string)</b>
Default: None
Passphrase used to access the private key specified in <i>key_file_path</i> .

## 5.5.18. Class SNIBasedCertificate

This class adds support for the Server Name Indication (SNI) TLS extension, as described in [RFC 6066](#). It stores a mapping between hostnames and certificates, and automatically selects the certificate to show to the peer if the peer has sent an SNI request.

### 5.5.18.1. Attributes of SNIBasedCertificate

<b>default (complex)</b>
Default: None
The certificate to show to the peer if no matching hostname is found in <i>hostname_certificate_map</i> .
<b>hostname_certificate_map (complex)</b>
Default: n/a

### hostname\_certificate\_map (complex)

A hash containing a matcher-certificate map. Each element of the hash contains a matcher and a certificate: if a matcher matches the hostname in the SNI request, the certificate is showed to the peer. You can use any matcher policy, though in most cases, RegexpMatcher will be adequate. Different elements of the hash can use different types of matchers, for example, RegexpMatcher and RegexpFileMatcher. For details on matcher policies, see *Section 5.7, Module Matcher (p. 163)*.

```
hostname_certificate_map={
  RegexpMatcher(
    match_list=("myfirstdomain.example.com", ): StaticCertificate(
      certificates=(Certificate.fromFile(
certificate_file_path="/etc/key.d/myfirstdomain/cert.pem",
      private_key=PrivateKey.fromFile(
        "/etc/key.d/myfirstdomain/key.pem")),),),}
```

#### 5.5.18.2. SNIBasedCertificate methods

Method	Description
<code>__init__(self, hostname_certificate_map, default)</code>	

Table 5.57. Method summary

Method `__init__(self, hostname_certificate_map, default)`

#### Arguments of `__init__`

default (complex)
Default: None
The certificate to show to the peer if no matching hostname is found in <code>hostname_certificate_map</code> .

hostname_certificate_map (complex)
Default: n/a
A matcher-certificate map that describes which certificate will be showed to the peer if the matcher part matches the hostname in the SNI request. For details on matcher policies, see <i>Section 5.7, Module Matcher (p. 163)</i> .

#### 5.5.19. Class SSLOptions

This class collects the TLS and SSL settings directly related to encryption, for example, the permitted protocol versions, ciphers, session reuse settings, and so on. Note that you cannot use this class directly, use an appropriate derived class, for example, *ClientSSLOptions* or *ServerSSLOptions* instead.



### 5.5.19.1. Attributes of SSLOptions

<b>cipher (complex)</b>
Default: n/a
Specifies the allowed ciphers. For details, see <i>Table 5.31, Constants for cipher selection (p. 116)</i> .
<b>ciphers_tlsv1_3 (complex)</b>
Default: n/a
Specifies the allowed ciphers for TLSv1.3 connections. For details, see <i>Table 5.32, Constants for TLSv1.3 cipher selection (p. 117)</i> .
<b>disable_compression (boolean)</b>
Default: FALSE
Set this to TRUE to disable support for SSL/TLS compression even if it is supported. Please be mind that this option is ignored in TLSv1.3 as it does not support compression.
<b>disable_session_cache (boolean)</b>
Default: FALSE
Do not store session information in the session cache. Set this option to TRUE to disable SSL session reuse. Please be mind that this option is ignored in TLSv1.3 as it does not support session IDs.
<b>disable_ticket (boolean)</b>
Default: FALSE
Session tickets are a method for SSL session reuse, described in RFC 5077. Set this option to TRUE to disable SSL session reuse using session tickets.
<b>disable_tlsv1 (boolean)</b>
Default: TRUE
Do not allow using TLSv1 in the connection.
<b>disable_tlsv1_1 (boolean)</b>
Default: FALSE
Do not allow using TLSv1.1 in the connection.
<b>disable_tlsv1_2 (boolean)</b>
Default: FALSE
Do not allow using TLSv1.2 in the connection.

<b>disable_tlsv1_3 (boolean)</b>
Default: FALSE
Do not allow using TLSv1.3 in the connection.
<b>session_cache_size (integer)</b>
Default: 20480
The number of sessions stored in the session cache for SSL session reuse. Please be mind that this option is ignored in TLSv1.3 as it does not support session IDs.
<b>shared_groups (complex)</b>
Default: n/a
Specifies the allowed shared groups. For details, see <i>Table 5.33, Constants for shared group selection (p. 117)</i> .

### 5.5.19.2. SSLOptions methods

Method	Description
<u><i>__init__(self, cipher, ciphers_tlsv1_3, shared_groups, timeout, disable_tlsv1, disable_tlsv1_1, disable_tlsv1_2, disable_tlsv1_3, session_cache_size, disable_session_cache, disable_ticket, disable_compression)</i></u>	Constructor to initialize an SSLOptions instance.

Table 5.58. Method summary

**Method `__init__(self, cipher, ciphers_tlsv1_3, shared_groups, timeout, disable_tlsv1, disable_tlsv1_1, disable_tlsv1_2, disable_tlsv1_3, session_cache_size, disable_session_cache, disable_ticket, disable_compression)`**

This constructor defines an SSLOptions with the specified parameters.

#### Arguments of `__init__`

<b>cipher (enum)</b>
Default: n/a
Specifies the allowed ciphers. For details, see <i>Table 5.31, Constants for cipher selection (p. 116)</i> .
<b>ciphers_tlsv1_3 (enum)</b>
Default: n/a
Specifies the allowed ciphers for TLSv1.3 connections. For details, see <i>Table 5.32, Constants for TLSv1.3 cipher selection (p. 117)</i> .

**disable\_compression (boolean)**

Default: FALSE

Set this to TRUE to disable support for SSL/TLS compression even if it is supported. Please be mind that this option is ignored in TLSv1.3 as it does not support compression.

**disable\_session\_cache (boolean)**

Default: FALSE

Do not store session information in the session cache. Set this option to TRUE to disable SSL session reuse. Please be mind that this option is ignored in TLSv1.3 as it does not support session IDs.

**disable\_ticket (boolean)**

Default: FALSE

Session tickets are a method for SSL session reuse, described in RFC 5077. Set this option to TRUE to disable SSL session reuse using session tickets.

**disable\_tlsv1 (boolean)**

Default: TRUE

Do not allow using TLSv1 in the connection.

**disable\_tlsv1\_1 (boolean)**

Default: FALSE

Do not allow using TLSv1.1 in the connection.

**disable\_tlsv1\_2 (boolean)**

Default: FALSE

Do not allow using TLSv1.2 in the connection.

**disable\_tlsv1\_3 (boolean)**

Default: FALSE

Do not allow using TLSv1.3 in the connection.

**session\_cache\_size (integer)**

Default: 20480

The number of sessions stored in the session cache for SSL session reuse. Please be mind that this option is ignored in TLSv1.3 as it does not support session IDs.



<b>shared_groups (enum)</b>
Default: n/a
Specifies the allowed shared groups. For details, see <i>Table 5.33, Constants for shared group selection (p. 117)</i> .

<b>timeout (integer)</b>
Default: 300
Drop idle connection if the timeout value (in seconds) expires.

## 5.5.20. Class ServerCertificateVerifier

This class includes the settings and options used to verify the certificates of the peers in server-side SSL and TLS connections. Note that the ServerCertificateVerifier class always requests a certificate from the server.

### 5.5.20.1. Attributes of ServerCertificateVerifier

<b>check_subject (boolean)</b>
Default: TRUE
If the <i>check_subject</i> parameter is TRUE, the Subject of the server-side certificate is compared with application-layer information (for example, it checks whether the Subject matches the hostname in the URL). For details, see <i>Section 3.2.3.5, Certificate verification options (p. 13)</i> .

<b>intermediate_revocation_check_type (enum)</b>
Default: TLS_INTERMEDIATE_REVOCATION_SOFT_FAIL
Specify how intermediate certificates revocation status check should work.

<b>leaf_revocation_check_type (enum)</b>
Default: TLS_LEAF_REVOCATION_SOFT_FAIL
Specify how leaf certificate revocation status check should work.

<b>trust_level (enum)</b>
Default: TLS_TRUST_LEVEL_FULL
Specify which certificate should be accepted as trusted.

<b>trusted_certs_directory (string)</b>
Default: ""
A directory where trusted IP address - certificate assignments are stored. When a peer from a specific IP address shows the certificate stored in this directory, it is accepted regardless of its expiration or issuer CA. Each file in the directory should contain a certificate in PEM format. The filename must be the IP address.



<b>verify_ca_directory (string)</b>
Default: ""
Directory where the trusted CA certificates are stored. CA certificates are loaded on-demand from this directory when the certificate of the peer is verified.

<b>verify_crl_directory (string)</b>
Default: ""
Directory where the CRLs (Certificate Revocation Lists) associated with trusted CAs are stored. CRLs are loaded on-demand from this directory when the certificate of the peer is verified.

<b>verify_depth (integer)</b>
Default: 4
The length of the longest accepted CA verification chain. Longer CA chains are automatically rejected.

#### 5.5.20.2. ServerCertificateVerifier methods

Method	Description
<code>__init__(self, trust_level, intermediate_revocation_check_type, leaf_revocation_check_type, trusted_certs_directory, verify_depth, verify_ca_directory, verify_crl_directory, check_subject)</code>	Constructor to initialize a ServerCertificateVerifier instance.

Table 5.59. Method summary

**Method `__init__(self, trust_level, intermediate_revocation_check_type, leaf_revocation_check_type, trusted_certs_directory, verify_depth, verify_ca_directory, verify_crl_directory, check_subject)`**

This constructor defines a ServerCertificateVerifier with the specified parameters.

#### Arguments of `__init__`

<b>check_subject (boolean)</b>
Default: TRUE
If the <code>check_subject</code> parameter is TRUE, the Subject of the server-side certificate is compared with application-layer information (for example, it checks whether the Subject matches the hostname in the URL). For details, see <i>Section 3.2.3.5, Certificate verification options (p. 13)</i> .

<b>intermediate_revocation_check_type (enum)</b>
Default: TLS_INTERMEDIATE_REVOCATION_SOFT_FAIL
Specify how intermediate certificates revocation status check should work.



<b>leaf_revocation_check_type (enum)</b>
Default: TLS_LEAF_REVOCATION_SOFT_FAIL
Specify how leaf certificate revocation status check should work.
<b>trust_level (enum)</b>
Default: TLS_TRUST_LEVEL_FULL
Specify which certificate should be accepted as trusted.
<b>trusted_certs_directory (string)</b>
Default: ""
A directory where trusted IP address - certificate assignments are stored. When a peer from a specific IP address shows the certificate stored in this directory, it is accepted regardless of its expiration or issuer CA. Each file in the directory should contain a certificate in PEM format. The filename must be the IP address.
<b>verify_ca_directory (string)</b>
Default: ""
Directory where the trusted CA certificates are stored. CA certificates are loaded on-demand from this directory when the certificate of the peer is verified.
<b>verify_crl_directory (string)</b>
Default: ""
Directory where the CRLs (Certificate Revocation Lists) associated with trusted CAs are stored. CRLs are loaded on-demand from this directory when the certificate of the peer is verified.
<b>verify_depth (integer)</b>
Default: 4
The length of the longest accepted CA verification chain. Longer CA chains are automatically rejected.

### 5.5.21. Class ServerNoneVerifier

This class disables every certificate verification in server-side SSL and TLS connections.

### 5.5.22. Class ServerOnlyEncryption

The ServerOnlyEncryption class handles scenarios when only the Zorp-server connection is encrypted, the client-Zorp connection is not.

### 5.5.22.1. Attributes of ServerOnlyEncryption

<b>server_certificate_generator (class)</b>
Default: None
The class that will generate the certificate that will be showed to the server. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.
<b>server_ssl_options (class)</b>
Default: ServerSSLOptions()
The protocol-level encryption settings used on the server side. This must be a <i>ServerSSLOptions</i> instance.
<b>server_verify (class)</b>
Default: ServerCertificateVerifierGroup()
The settings used to verify the certificate of the server. This must be a <i>ServerCertificateVerifier</i> instance.

### 5.5.22.2. ServerOnlyEncryption methods

Method	Description
<i>__init__(self, server_certificate_generator, server_verify, server_ssl_options)</i>	Initializes SSL/TLS connection on the server side.

Table 5.60. Method summary

#### Method `__init__(self, server_certificate_generator, server_verify, server_ssl_options)`

The ServerOnlyEncryption class handles scenarios when only the Zorp-server connection is encrypted, the client-Zorp connection is not.

#### Arguments of `__init__`

<b>server_certificate_generator (class)</b>
Default: None
The class that will generate the certificate that will be showed to the server. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.
<b>server_ssl_options (class)</b>
Default: ServerSSLOptions()
The protocol-level encryption settings used on the server side. This must be a <i>ServerSSLOptions</i> instance.

<b>server_verify (class)</b>
Default: ServerCertificateVerifierGroup()
The settings used to verify the certificate of the server. This must be a <i>ServerCertificateVerifier</i> instance.

### 5.5.23. Class ServerSSLOptions

This class (based on the SSLOptions class) collects the TLS and SSL settings directly related to encryption, for example, the permitted protocol versions, ciphers, session reuse settings, and so on.

#### 5.5.23.1. Attributes of ServerSSLOptions

<b>cipher (enum)</b>
Default: n/a
Specifies the allowed ciphers. For details, see <i>Table 5.31, Constants for cipher selection (p. 116)</i> .

<b>ciphers_tlsv1_3 (enum)</b>
Default: n/a
Specifies the allowed ciphers for TLSv1.3 connections. For details, see <i>Table 5.32, Constants for TLSv1.3 cipher selection (p. 117)</i> .

<b>disable_compression (boolean)</b>
Default: FALSE
Set this to TRUE to disable support for SSL/TLS compression even if it is supported. Please be mind that this option is ignored in TLSv1.3 as it does not support compression.

<b>disable_session_cache (boolean)</b>
Default: FALSE
Do not store session information in the session cache. Set this option to TRUE to disable SSL session reuse. Please be mind that this option is ignored in TLSv1.3 as it does not support session IDs.

<b>disable_ticket (boolean)</b>
Default: FALSE
Session tickets are a method for SSL session reuse, described in RFC 5077. Set this option to TRUE to disable SSL session reuse using session tickets.

<b>disable_tlsv1 (boolean)</b>
Default: TRUE
Do not allow using TLSv1 in the connection.



<b>disable_tlsv1_1 (boolean)</b>
Default: FALSE
Do not allow using TLSv1.1 in the connection.
<b>disable_tlsv1_2 (boolean)</b>
Default: FALSE
Do not allow using TLSv1.2 in the connection.
<b>disable_tlsv1_3 (boolean)</b>
Default: FALSE
Do not allow using TLSv1.3 in the connection.
<b>session_cache_size (integer)</b>
Default: 20480
The number of sessions stored in the session cache for SSL session reuse. Please be mind that this option is ignored in TLSv1.3 as it does not support session IDs.
<b>shared_groups (enum)</b>
Default: n/a
Specifies the allowed shared groups. For details, see <i>Table 5.33, Constants for shared group selection (p. 117)</i> .

### 5.5.23.2. ServerSSLOptions methods

Method	Description
<u><code>__init__(self, method, cipher, ciphers_tlsv1_3, shared_groups, timeout, disable_sslv2, disable_sslv3, disable_tlsv1, disable_tlsv1_1, disable_tlsv1_2, disable_tlsv1_3, session_cache_size, disable_session_cache, disable_ticket, disable_compression)</code></u>	Constructor to initialize a ServerSSLOptions instance.

Table 5.61. Method summary

**Method `__init__(self, method, cipher, ciphers_tlsv1_3, shared_groups, timeout, disable_sslv2, disable_sslv3, disable_tlsv1, disable_tlsv1_1, disable_tlsv1_2, disable_tlsv1_3, session_cache_size, disable_session_cache, disable_ticket, disable_compression)`**

This constructor defines a ServerSSLOptions with the specified parameters.

Arguments of `__init__`

<b>cipher (enum)</b>
Default: n/a
Specifies the allowed ciphers. For details, see <i>Table 5.31, Constants for cipher selection (p. 116)</i> .
<b>ciphers_tlsv1_3 (enum)</b>
Default: n/a
Specifies the allowed ciphers for TLSv1.3 connections. For details, see <i>Table 5.32, Constants for TLSv1.3 cipher selection (p. 117)</i> .
<b>disable_compression (boolean)</b>
Default: FALSE
Set this to TRUE to disable support for SSL/TLS compression even if it is supported. Please be mind that this option is ignored in TLSv1.3 as it does not support compression.
<b>disable_session_cache (boolean)</b>
Default: FALSE
Do not store session information in the session cache. Set this option to TRUE to disable SSL session reuse. Please be mind that this option is ignored in TLSv1.3 as it does not support session IDs.
<b>disable_ticket (boolean)</b>
Default: FALSE
Session tickets are a method for SSL session reuse, described in RFC 5077. Set this option to TRUE to disable SSL session reuse using session tickets.
<b>disable_tlsv1 (boolean)</b>
Default: TRUE
Do not allow using TLSv1 in the connection.
<b>disable_tlsv1_1 (boolean)</b>
Default: FALSE
Do not allow using TLSv1.1 in the connection.
<b>disable_tlsv1_2 (boolean)</b>
Default: FALSE
Do not allow using TLSv1.2 in the connection.



<b>disable_tlsv1_3 (boolean)</b>
Default: FALSE
Do not allow using TLSv1.3 in the connection.
<b>session_cache_size (integer)</b>
Default: 20480
The number of sessions stored in the session cache for SSL session reuse. Please be mind that this option is ignored in TLSv1.3 as it does not support session IDs.
<b>shared_groups (enum)</b>
Default: n/a
Specifies the allowed shared groups. For details, see <i>Table 5.33, Constants for shared group selection (p. 117)</i> .
<b>timeout (integer)</b>
Default: 300
Drop idle connection if the timeout value (in seconds) expires.

## 5.5.24. Class StaticCertificate

This class encapsulates a static Certificate that can be used in SSL/TLS connections.

### 5.5.24.1. Attributes of StaticCertificate

<b>certificates (complex)</b>
Default: n/a
List of certificate instances to show to the peer.

### 5.5.24.2. StaticCertificate methods

Method	Description
<code>__init__(self, certificates, certificate)</code>	Initializes a static Certificate object.

Table 5.62. Method summary

#### Method `__init__(self, certificates, certificate)`

A static Certificate that can be used in SSL/TLS connections.

**Arguments of `__init__`**

<b>certificates (complex)</b>
Default: n/a
List of certificate instances to show to the peer.

**5.5.25. Class TwoSidedEncryption**

The `TwoSidedEncryption` class handles scenarios when both the client-Zorp and the Zorp-server connections are encrypted. If you do not need encryption on the client- or the server-side, use the *ServerOnlyEncryption* or *ClientOnlyEncryption* classes, respectively. For a detailed example on keybridging, see *Section 3.2.7, Keybridging certificates (p. 16)*.

**5.5.25.1. Attributes of TwoSidedEncryption**

<b>client_certificate_generator (class)</b>
Default: n/a
The class that will generate the certificate that will be showed to the client. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.

<b>client_ssl_options (class)</b>
Default: <code>ClientSSLOptions()</code>
The protocol-level encryption settings used on the client side. This must be a <i>ClientSSLOptions</i> instance.

<b>client_verify (class)</b>
Default: <code>ClientCertificateVerifierGroup()</code>
The settings used to verify the certificate of the client. This must be a <i>ClientCertificateVerifier</i> instance.

<b>server_certificate_generator (class)</b>
Default: None
The class that will generate the certificate that will be showed to the server. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.

<b>server_ssl_options (class)</b>
Default: <code>ServerSSLOptions()</code>
The protocol-level encryption settings used on the server side. This must be a <i>ServerSSLOptions</i> instance.

<b>server_verify (class)</b>
Default: <code>ServerCertificateVerifierGroup()</code>



<b>server_verify (class)</b>
The settings used to verify the certificate of the server. This must be a <i>ServerCertificateVerifier</i> instance.

### 5.5.25.2. TwoSidedEncryption methods

Method	Description
<i>__init__</i> ( <i>self</i> , <i>client_certificate_generator</i> , <i>server_certificate_generator</i> , <i>client_verify</i> , <i>server_verify</i> , <i>client_ssl_options</i> , <i>server_ssl_options</i> )	Initializes SSL/TLS connection with both peers.

Table 5.63. Method summary

**Method `__init__` (*self*, *client\_certificate\_generator*, *server\_certificate\_generator*, *client\_verify*, *server\_verify*, *client\_ssl\_options*, *server\_ssl\_options*)**

The TwoSidedEncryption class handles scenarios when both the client-Zorp and the Zorp-server connections are encrypted.

#### Arguments of `__init__`

<b>client_certificate_generator (class)</b>
Default: n/a
The class that will generate the certificate that will be showed to the client. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.

<b>client_ssl_options (class)</b>
Default: ClientSSLOptions()
The protocol-level encryption settings used on the client side. This must be a <i>ClientSSLOptions</i> instance.

<b>client_verify (class)</b>
Default: ClientCertificateVerifierGroup()
The settings used to verify the certificate of the client. This must be a <i>ClientCertificateVerifier</i> instance.

<b>server_certificate_generator (class)</b>
Default: None
The class that will generate the certificate that will be showed to the server. You can use an instance of the <i>StaticCertificate</i> , <i>DynamicCertificate</i> , or <i>SNIBasedCertificate</i> classes.

<b>server_ssl_options (class)</b>
Default: ServerSSLOptions()



**server\_ssl\_options (class)**

The protocol-level encryption settings used on the server side. This must be a *ServerSSLOptions* instance.

**server\_verify (class)**

Default: `ServerCertificateVerifierGroup()`

The settings used to verify the certificate of the server. This must be a *ServerCertificateVerifier* instance.

## 5.6. Module Keybridge

Keybridging is a method to let the client see a copy of the server's certificate (or vice versa), allowing it to inspect it and decide about its trustworthiness. Because of proxying the SSL/TLS connection, the client is not able to inspect the certificate of the server directly, therefore a certificate based on the server's certificate is generated on-the-fly. This generated certificate is presented to the client.

For details on configuring keybridging, see *Section 3.2.7, Keybridging certificates (p. 16)*.

### 5.6.1. Classes in the Keybridge module

Class	Description
<i>X509KeyBridge</i>	Class to perform SSL keybridging.

Table 5.64. Classes of the Keybridge module

### 5.6.2. Class X509KeyBridge

This class is able to generate certificates mimicking another certificate, primarily used to transfer the information of a server's certificate to the client in keybridging. For details on configuring keybridging, see *Section 3.2.7, Keybridging certificates (p. 16)*.

#### 5.6.2.1. Attributes of X509KeyBridge

**cache\_directory (string)**

Default: ""

The directory where all automatically generated certificates are cached.

**key\_file (string)**

Default: ""

Name of the private key to be used for the newly generated certificates.

**key\_passphrase (string)**

Default: ""

<b>key_passphrase (string)</b>
Passphrase required to access the private key stored in <i>key_file</i> .
<b>trusted_ca_files (certificate)</b>
Default: None
A tuple of <i>cert_file</i> , <i>key_file</i> , <i>passphrase</i> ) for the CA used for keybridging trusted certificates.
<b>untrusted_ca_files (certificate)</b>
Default: None
A tuple of <i>cert_file</i> , <i>key_file</i> , <i>passphrase</i> ) for the CA used for keybridging untrusted certificates.

### 5.6.2.2. X509KeyBridge methods

Method	Description
<i>_old_init</i> (self, <i>key_file</i> , <i>cache_directory</i> , <i>trusted_ca_files</i> , <i>untrusted_ca_files</i> , <i>key_passphrase</i> , <i>extension_whitelist</i> )	None

Table 5.65. Method summary

Method *\_old\_init*(self, *key\_file*, *cache\_directory*, *trusted\_ca\_files*, *untrusted\_ca\_files*, *key\_passphrase*, *extension\_whitelist*)

n/a

#### Arguments of *\_old\_init*

<b>cache_directory (string)</b>
Default: <code>"/var/lib/zorp/keybridge-cache"</code>
The directory where all automatically generated certificates are cached.
<b>extension_whitelist (complex)</b>
Default: None
The following certificate extensions are transferred to the client side: <i>Key Usage</i> , <i>Subject Alternative Name</i> , <i>Extended Key Usage</i> . Other extensions will be automatically deleted during keybridging. This is needed because some certificate extensions contain references to the Issuer CA, which references become invalid for keybridged certificates. To transfer other extensions, list them in the <i>extension_whitelist</i> parameter. Note that modifying this parameter replaces the default values, so to extend the list of transferred extensions, include the <code>'keyUsage'</code> , <code>'subjectAltName'</code> , <code>'extendedKeyUsage'</code> list as well. For example:



<b>extension_whitelist (complex)</b>
<pre>self.extension_whitelist = ('keyUsage', 'subjectAltName', 'extendedKeyUsage', 'customExtension')</pre>
<b>key_file (certificate)</b>
Default: n/a
Name of the private key to be used for the newly generated certificates.
<b>key_passphrase (string)</b>
Default: ""
Passphrase required to access the private key stored in <i>key_file</i> .
<b>trusted_ca_files (certificate)</b>
Default: n/a
A tuple of <i>cert_file</i> , <i>key_file</i> , <i>passphrase</i> ) for the CA used for keybridging trusted certificates.
<b>untrusted_ca_files (certificate)</b>
Default: None
A tuple of <i>cert_file</i> , <i>key_file</i> , <i>passphrase</i> ) for the CA used for keybridging untrusted certificates.

## 5.7. Module Matcher

In general, matcher policies can be used to find out if a parameter is included in a list (or which elements of a list correspond to a certain parameter), and influence the behavior of the proxy class based on the results. Matchers can be used for a wide range of tasks, for example, to determine if the particular IP address or URL that a client is trying to access is on a black or whitelist, or to verify that a particular e-mail address is valid.

### 5.7.1. Classes in the Matcher module

Class	Description
<i>AbstractMatcher</i>	Class encapsulating the abstract string matcher.
<i>CombineMatcher</i>	Matcher for implementing logical expressions based on other matchers.
<i>DNSMatcher</i>	DNS matcher
<i>MatcherPolicy</i>	Class encapsulating a Matcher which can be used by a name.
<i>RegexpFileMatcher</i>	Class encapsulating Matcher which uses regular expressions stored in files for string matching.

Class	Description
<i>RegexpMatcher</i>	Class encapsulating a Matcher which uses regular expressions for string matching.
<i>SmtplibInvalidRecipientMatcher</i>	Class verifying the validity of the recipient addresses in E-mails.
<i>WindowsUpdateMatcher</i>	Windows Update matcher

Table 5.66. Classes of the Matcher module

### 5.7.2. Class AbstractMatcher

This abstract class encapsulates a string matcher that determines whether a given string is found in a backend database.

Specialized subclasses of AbstractMatcher exist such as 'RegexpFileMatcher' which use regular expressions stored in flat files to find matches.

### 5.7.3. Class CombineMatcher

This matcher makes it possible to combine the results of several matchers using logical operations. CombineMatcher uses prefix-notation in its expressions and uses the following format: the operand, a comma, first argument, a comma, second argument. For example, an AND expression should be formatted the following way: (*Z\_AND, matcher1, matcher2*). Expressions using more than one operands should be bracketed, e.g., (*Z\_OR (Z\_AND, matcher1, matcher2), matcher3*). The following operations are available:

- *Z\_AND* : Logical AND operation.
- *Z\_OR* : Logical OR operation.
- *Z\_XOR* : Logical XOR operation.
- *Z\_NOT* : Logical negation.
- *Z\_EQ* : Logical equation.



#### Example 5.23. Whitelisting e-mail recipients

A simple use for CombineMatcher is to filter the recipients of e-mail addresses using the following process:

1. An SmtplibInvalidMatcher (called *SmtplibCheckrecipient*) verifies that the recipient exists.
2. A RegexpMatcher (called *SmtplibWhitelist*) or RegexpFileMatcher is used to check if the address is on a predefined list (list of permitted addresses).
3. A CombineMatcher (called *SmtplibCombineMatcher*) sums up the results of the matchers with a logical AND operation.
4. An SmtplibProxy (called *SmtplibRecipientMatcherProxy*) references *SmtplibCombineMatcher* in its *recipient\_matcher* attribute.

Python:

```
class SmtplibRecipientMatcherProxy(SmtplibProxy):
    recipient_matcher="SmtplibCombineMatcher"
    def config(self):
        super(SmtplibRecipientMatcherProxy, self).config()

MatcherPolicy(name="SmtplibCombineMatcher", matcher=CombineMatcher (expr=(Z_AND, "SmtplibCheckrecipient",
"SmtplibWhitelist")))
MatcherPolicy(name="SmtplibWhitelist", matcher=RegexpMatcher (match_list=("info@example.com",),
ignore_list=None))
```

```
MatcherPolicy(name="Smtptestrecipient", matcher=SmtptestInvalidRecipientMatcher (server_port=25,
cache_timeout=60, attempt_delivery=FALSE, force_delivery_attempt=FALSE,
server_name="recipientcheck.example.com"))
```

### 5.7.4. Class DNSMatcher

DNSMatcher retrieves the IP addresses of domain names. This can be used in domain name based policy decisions, for example to allow encrypted connections only to trusted e-banking sites.

DNSMatcher operates as follows: it resolves the IP addresses stored in the list of domain names using the specified Domain Name Server, and compares the results to the IP address of the connection (i.e., the IP address of the server or the client). The matcher returns a true value if the IP addresses resolved from the list of domain names include the IP address of the connection.



#### Example 5.24. DNSMatcher example

The following DNSMatcher class uses the `dns.example.com` name server to resolve the `example2.com` and `example3.com` domain names.

```
MatcherPolicy(name="ExampleDomainMatcher", matcher=DNSMatcher (server="dns.example.com",
hosts=("example2.com", "example3.com")))
```

#### 5.7.4.1. DNSMatcher methods

Method	Description
<code>__init__(self, hosts, server, resolve_on_init)</code>	Constructor to initialize an instance of the DNSMatcher class.

Table 5.67. Method summary

#### Method `__init__(self, hosts, server, resolve_on_init)`

This constructor initializes an instance of the DNSMatcher class.

#### Arguments of `__init__`

<b>hosts (complex)</b>
Default: n/a
Hostnames to resolve.
<b>resolve_on_init (boolean)</b>
Default: FALSE
Resolve all hostnames on startup time. Otherwise, names will be resolved on-demand.
<b>server (string)</b>
Default: None

**server (string)**

IP address of the DNS server to query. Defaults to the servers set in the `resolv.conf` file.

**5.7.5. Class MatcherPolicy**

Matcher policies can be used to find out if a parameter is included in a list, or which elements of a list correspond to a certain parameter, and influence the behavior of the proxy class based on the results. Matchers can be used for a wide range of tasks, for example, to determine if the particular IP address or URL that a client is trying to access is on a black or whitelist, or to verify that a particular e-mail address is valid.

MatcherPolicy instances are reusable matchers that contain configured instances of the matcher classes (e.g., DNSMatcher, RegexpMatcher). For examples, see the specific matcher classes.

**5.7.6. Class RegexpFileMatcher**

This class is similar to *RegexpMatcher*, but stores the regular expressions to match and ignore in files. For example, this class can be used for URL filtering. The matcher itself stores only the paths and the filenames to the lists. The file is automatically monitored and reloaded when it is modified. Searches are case-insensitive.

**Example 5.25. RegexpFileMatcher example**

```
MatcherPolicy(name="demo_regexpfilematcher",
matcher=RegexpFileMatcher(match_fname="/tmp/match_list.txt", ignore_fname="/tmp/ignore_list.txt"))
```

**5.7.6.1. Attributes of RegexpFileMatcher****ignore\_date (unknown)**

Default: n/a

Date (in unix timestamp format) when the *ignore\_file* was loaded.

**ignore\_file (unknown)**

Default: n/a

Name of the file storing the patterns to ignore.

**match\_date (unknown)**

Default: n/a

Date (in unix timestamp format) when the *match\_file* was loaded.

**match\_file (unknown)**

Default: n/a

Name of the file storing the patterns for positive matches.

### 5.7.6.2. RegexpFileMatcher methods

Method	Description
<code>__init__(self, match_fname, ignore_fname)</code>	Constructor to initialize a RegexpFileMatcher instance.

Table 5.68. Method summary

#### Method `__init__(self, match_fname, ignore_fname)`

This constructor initializes an instance of the RegexpFileMatcher class.

#### Arguments of `__init__`

<b>ignore_fname (filename)</b>
Default: None
Name of the file storing the patterns to ignore.

<b>match_fname (filename)</b>
Default: None
Name of the file storing the patterns for positive matches.

### 5.7.7. Class RegexpMatcher

A simple regular expression based matcher with a match and an ignore list. Searches are case-insensitive.



#### Example 5.26. RegexpMatcher example

The following RegexpMatcher matches only the `smtp.example.com` string.

```
MatcherPolicy(name="Smtpdomains", matcher=RegexpMatcher (match_list=("smtp.example.com", ),
ignore_list=None))
```

#### 5.7.7.1. Attributes of RegexpMatcher

<b>ignore (unknown)</b>
Default: n/a
A list of compiled regular expressions defining the strings to be ignored even if <code>match</code> resulted in a positive match.

<b>match (unknown)</b>
Default: n/a
A list of compiled regular expressions which result in a positive match.

### 5.7.7.2. RegexpMatcher methods

Method	Description
<code>__init__(self, match_list, ignore_list, ignore_case)</code>	Constructor to initialize a RegexpMatcher instance.

Table 5.69. Method summary

#### Method `__init__(self, match_list, ignore_list, ignore_case)`

This constructor initializes a RegexpMatcher instance by setting the *match* and *ignore* attributes to an empty list.

#### Arguments of `__init__`

<b>ignore_list (filename)</b>
Default: None
The list of regular expressions to ignore.
<b>match_list (filename)</b>
Default: None
The list of regular expressions to match.

### 5.7.8. Class SmtplibInvalidRecipientMatcher

This class encapsulates a VRFY/RCPT based validity checker to transparently verify the existence of E-mail addresses. Instead of immediately sending the e-mail to the recipient SMTP server, an independent SMTP server is queried about the existence of the recipient e-mail address.

Instances of this class can be referred to in the *recipient\_matcher* attribute of the *SmtplibProxy* class. The SmtplibProxy will automatically reject unknown recipients even if the recipient SMTP server would accept them.



#### Example 5.27. SmtplibInvalidMatcher example

```
Python:
class SmtplibRecipientMatcherProxy(SmtplibProxy):
    recipient_matcher="SmtplibCheckrecipient"
    def config(self):
        super(SmtplibRecipientMatcherProxy, self).config()

MatcherPolicy(name="SmtplibCheckrecipient", matcher=SmtplibInvalidRecipientMatcher (server_port=25,
cache_timeout=60, attempt_delivery=FALSE, force_delivery_attempt=FALSE,
server_name="recipientcheck.example.com"))
```



## 5.7.8.1. SmtplibInvalidRecipientMatcher methods

Method	Description
<code>__init__(self, server_name, server_port, cache_timeout, attempt_delivery, force_delivery_attempt, sender_address, bind_name)</code>	

Table 5.70. Method summary

Method `__init__(self, server_name, server_port, cache_timeout, attempt_delivery, force_delivery_attempt, sender_address, bind_name)`

Arguments of `__init__`

<b>bind_name (string)</b>
Default: ""
Specifies the hostname to bind to before initiating the connection to the SMTP server.
<b>cache_timeout (integer)</b>
Default: 60
How long will the result of an address verification be retained (in seconds).
<b>force_delivery_attempt (boolean)</b>
Default: FALSE
Force a delivery attempt even if the autodetection code otherwise would use VRFY. Useful if the server always returns success for VRFY.
<b>sender_address (string)</b>
Default: "<>"
This value will be used as the mail sender for the attempted mail delivery. Mail delivery is attempted if the <code>force_delivery_attempt</code> is TRUE, or the recipient server does not support the VRFY command.
<b>server_name (string)</b>
Default: n/a
Domain name of the SMTP server that will verify the addresses.
<b>server_port (integer)</b>
Default: 25
Port of the target server.

### 5.7.9. Class WindowsUpdateMatcher

WindowsUpdateMatcher is actually a DNSMatcher used to retrieve the IP addresses currently associated with the `v5.windowsupdate.microsoft.nsatc.net`, `v4.windowsupdate.microsoft.nsatc.net`, and `update.microsoft.nsatc.net` domain names from the specified name server. Windows Update is running on a distributed server farm, using the DNS round robin method and a short TTL to constantly change the set of servers currently visible, consequently the IP addresses of the servers are constantly changing.



#### Example 5.28. WindowsUpdateMatcher example

```
MatcherPolicy(name="demo_windowsupdatematcher", matcher=WindowsUpdateMatcher())
```

#### 5.7.9.1. WindowsUpdateMatcher methods

Method	Description
<code>__init__(self, server)</code>	Constructor to initialize an instance of the WindowsUpdateMatcher class.

Table 5.71. Method summary

#### Method `__init__(self, server)`

This constructor initializes an instance of the WindowsUpdateMatcher class.

#### Arguments of `__init__`

<b>server (string)</b>
Default: None
The IP address of the name server to query.

## 5.8. Module NAT

Network Address Translation (NAT) is a technology that can be used to change source or destination addresses in a connection from one IP address to another one. This module defines the classes performing the translation for IP addresses.

Several different NAT methods are supported using different NAT classes, like *GeneralNAT* or *StaticNAT*. To actually perform network address translation in a service, you have to use a *NATPolicy* instance that contains a configured NAT class. NAT policies provide a way to re-use NAT instances without having to define NAT mappings for each service individually.

### 5.8.1. Classes in the NAT module

Class	Description
<i>AbstractNAT</i>	Class encapsulating the abstract NAT interface.

Class	Description
<u>GeneralNAT</u>	Class encapsulating a general subnet-to-subnet NAT.
<u>HashNAT</u>	Class which sets the address from a hash table.
<u>NAT46</u>	Class that performs translation from IPv4 to IPv6 addresses (NAT46)
<u>NAT64</u>	Class that performs translation from IPv6 to IPv4 addresses (NAT64)
<u>NATPolicy</u>	Class encapsulating named NAT instances.
<u>OneToOneMultiNAT</u>	Class translating addresses between two IP ranges.
<u>OneToOneNAT</u>	Class translating addresses between two IP ranges.
<u>RandomNAT</u>	Class generating a random IP address.
<u>StaticNAT</u>	Class that replaces the source or destination address with a predefined address.

Table 5.72. Classes of the NAT module

## 5.8.2. Class AbstractNAT

This class encapsulates an interface for application level network address translation (NAT). This NAT is different from the NAT used by packet filters: it modifies the outgoing source/destination addresses just before Zorp connects to the server.

Source and destination NATs can be specified when a Service is created.

The NAT settings are used by the ConnectChainer class just before connecting to the server.

### 5.8.2.1. AbstractNAT methods

Method	Description
<u>__init__(self)</u>	Constructor to initialize an AbstractNAT instance.
<u>performTranslation(self, session, addrs, nat_type)</u>	Function that performs the address translation.

Table 5.73. Method summary

#### Method \_\_init\_\_(self)

This constructor initializes an AbstractNAT instance. Currently it does nothing, but serves as a placeholder for future extensions.

#### Method performTranslation(self, session, addrs, nat\_type)

This function is called before connecting a session to the destination server. The function returns the address (a SockAddr instance) to bind to before establishing the connection.

## Arguments of performTranslation

<b>addrs (unknown)</b>
Default: n/a
tuple of (source, destination) address, any of them can be none in case of the other translation
<b>nat_type (unknown)</b>
Default: n/a
translation type, either NAT_SNAT or NAT_DNAT
<b>session (unknown)</b>
Default: n/a
Session which is about to connect the server.

### 5.8.3. Class GeneralNAT

This class encapsulates a general subnet-to-subnet NAT. It requires a list of *from*, *to*, *translated to* parameters:

- *from*: the source address of the connection.
- *to*: the destination address of the connection.
- *translated to*: the translated address.

If the NAT policy is used as SNAT, the translated address is used to translate the source address of the connection; if the NAT policy is used as DNAT, the translated address is used to translate the destination address of the connection. The translation occurs according to the first matching rule.



#### Example 5.29. GeneralNat example

The following example defines a simple GeneralNAT policy that maps connections coming from the *192.168.1.0/24* subnet and targeting the *192.168.10.0/24* subnet into the *10.70.0.0/24* subnet.

```
NATPolicy(name="Demo_GeneralNAT", nat=GeneralNAT(mapping=((InetSubnet("192.168.1.0/24"),
InetSubnet("192.168.10.0/24"), InetSubnet("10.70.0.0/24")))))
```

If the policy is used as SNAT, the *192.168.1.0/24* subnet is translated into the *10.70.0.0/24* subnet and used as the source address of the connection. If the policy is used as DNAT, the *192.168.10.0/24* subnet is translated into the *10.70.0.0/24* subnet and used as the target address of the connection.

### 5.8.3.1. GeneralNAT methods

Method	Description
<code>__init__(self, mapping)</code>	Constructor to initialize a GeneralNAT instance.

Table 5.74. Method summary

#### Method `__init__(self, mapping)`

This constructor initializes a GeneralNAT instance.

#### Arguments of `__init__`

<b>mapping (complex)</b>
Default: n/a
List of tuples of InetSubnets in (source domain, destination domain, mapped domain) format.

### 5.8.4. Class HashNAT

HashNAT statically maps an IP address to another using a hash table. The table is indexed by the source IP address, and the value is the translated IP address. Both IP addresses are stored in string format.

#### 5.8.4.1. HashNAT methods

Method	Description
<code>__init__(self, ip_hash, default_reject)</code>	Constructor to initialize a HashNAT instance.

Table 5.75. Method summary

#### Method `__init__(self, ip_hash, default_reject)`

This constructor initializes a HashNAT instance.

#### Arguments of `__init__`

<b>default_reject (boolean)</b>
Default: TRUE
Enable this parameter to reject all connections outside the specific source range.

<b>ip_hash (complex)</b>
Default: n/a
The hash storing the IP address.



### 5.8.5. Class NAT46

NAT46 embeds and IPv4 address into a specific portion of the IPv6 address according to the NAT46 specification as described in RFC6052 (<http://tools.ietf.org/html/rfc6052#section-2.2>).

#### 5.8.5.1. NAT46 methods

Method	Description
<code>__init__(self, prefix, prefix_mask, suffix)</code>	Constructor to initialize a NAT46 instance.

Table 5.76. Method summary

#### Method `__init__(self, prefix, prefix_mask, suffix)`

This constructor initializes a NAT46 instance.

#### Arguments of `__init__`

<b>prefix (string)</b>
Default: "64:ff9b::"
This parameter specifies the common leading part of the IPv6 address that the IPv4 address should map into. Bits that exceed the mask will be overwritten by the mapping.
<b>prefix_mask (integer)</b>
Default: 96
This parameter specifies the position to embed the IPv4 address to and must be one of 32, 40, 48, 56, 64, or 96.
<b>suffix (string)</b>
Default: "::"
This parameter specifies the common trailing part of the IPv6 address that the IPv4 address should map into. The length of the suffix must not exceed the empty bit count determined by the configured prefix mask.

### 5.8.6. Class NAT64

NAT64 maps specific bits of the IPv6 address to IPv4 addresses according to the NAT64 specification as described in RFC6052 (<http://tools.ietf.org/html/rfc6052#section-2.2>).

### 5.8.6.1. NAT64 methods

Method	Description
<code>__init__(self, prefix_mask)</code>	Constructor to initialize a NAT64 instance.

Table 5.77. Method summary

#### Method `__init__(self, prefix_mask)`

This constructor initializes a NAT64 instance.

#### Arguments of `__init__`

<b>prefix_mask (integer)</b>
Default: 96
This parameter specifies the length of the IPv6 address to consider and must be one of 32, 40, 48, 56, 64, or 96.

### 5.8.7. Class NATPolicy

This class encapsulates a name and an associated NAT instance. NAT policies provide a way to re-use NAT instances without having to define NAT mappings for each service individually.



#### Example 5.30. Using Natpolicies

The following example defines a simple NAT policy, and uses this policy for SNAT in a service.

```
NATPolicy(name="demo_natpolicy", nat=GeneralNAT(mapping=((InetSubnet(addr="10.0.1.0/24"),
InetSubnet(addr="192.168.1.0/24")),)))
Service(name="office_http_inter", proxy_class=HttpProxy, snat_policy="demo_natpolicy")
```

### 5.8.7.1. NATPolicy methods

Method	Description
<code>__init__(self, name, nat, cacheable)</code>	Constructor to initialize a NAT policy.

Table 5.78. Method summary

#### Method `__init__(self, name, nat, cacheable)`

This constructor initializes a NAT policy.

#### Arguments of `__init__`

<b>cacheable (boolean)</b>
Default: TRUE

<b>cacheable (boolean)</b>
Enable this parameter to cache the NAT decisions.
<b>name (string)</b>
Default: n/a
Name identifying the NAT policy.
<b>nat (class)</b>
Default: n/a
NAT object which performs address translation.

### 5.8.8. Class OneToOneMultiNAT



**Note**  
This class is obsolete, use *GeneralNAT* instead.

This class is similar to *OneToOneNAT* as it 1:1 address translation between the source and destination subnets. The difference is that the *OneToOneMultiNAT* class supports multiple mappings by using a list of mapping pairs.

If the source address is outside the given source address range, a *DACException* is raised. The source and destination subnets must have the same size.

#### 5.8.8.1. OneToOneMultiNAT methods

Method	Description
<code>__init__(self, mapping, default_reject)</code>	Constructor to initialize a <i>OneToOneMultiNAT</i> instance.

Table 5.79. Method summary

#### Method `__init__(self, mapping, default_reject)`

This constructor initializes an instance of the *OneToOneMultiNAT* class. Arguments must be *Subnet* instances specifying two non-overlapping IP subnets with the same size.

#### Arguments of `__init__`

<b>default_reject (boolean)</b>
Default: TRUE
Enable this parameter to reject all connections outside the specific source range.



<b>mapping (complex)</b>
Default: n/a
List of <i>Subnet</i> pairs in the <i>from</i> , <i>to</i> format.

### 5.8.9. Class OneToOneNAT



**Note**  
This class is obsolete, use *GeneralNAT* instead.

This class performs 1:1 address translation between the source and destination subnets. If the source address is outside the given source address range, a *DACException* is raised. The source and destination subnets must have the same size.



**Tip**  
Use *OneToOneNAT* to redirect a block of IP addresses to another block, for example, when the web servers located in the DMZ have dedicated IP aliases on the firewall.

#### 5.8.9.1. OneToOneNAT methods

Method	Description
<code>__init__(self, from_domain, to_domain, default_reject)</code>	Constructor to initialize a <i>OneToOneNAT</i> instance.

Table 5.80. Method summary

#### Method `__init__(self, from_domain, to_domain, default_reject)`

This constructor initializes a *OneToOneNAT* instance. Arguments must be *Subnet* instances specifying two non-overlapping IP subnets with the same size.

#### Arguments of `__init__`

<b>default_reject (boolean)</b>
Default: TRUE
Enable this parameter to reject all connections outside the specific source range.

<b>from_domain (class)</b>
Default: n/a
The source subnet ( <i>Subnet</i> instance).

<b>to_domain (class)</b>
Default: n/a
The destination subnet (Subnet instance).

### 5.8.10. Class RandomNAT

This class randomly selects an address from a list of IP addresses. This can be used for load-balancing several lines by binding each session to a different interface.

#### 5.8.10.1. RandomNAT methods

Method	Description
<code>__init__(self, addresses)</code>	Constructor to initialize a RandomNAT instance.

Table 5.81. Method summary

#### Method `__init__(self, addresses)`

This constructor initializes a RandomNAT instance.

#### Arguments of `__init__`

<b>addresses (complex)</b>
Default: n/a
List of the available interfaces. Each item of the list must be an instance of the <i>SockAddr</i> (or a derived) class.

### 5.8.11. Class StaticNAT

This class assigns a predefined value to the address of the connection.

#### 5.8.11.1. StaticNAT methods

Method	Description
<code>__init__(self, addr)</code>	Constructor to initialize a StaticNAT instance.

Table 5.82. Method summary

#### Method `__init__(self, addr)`

This constructor initializes a StaticNAT instance.



## Arguments of `__init__`

<b>addr (sockaddr)</b>
Default: n/a
The address that replaces all addresses.

## 5.9. Module Notification

### 5.9.1. Classes in the Notification module

Class	Description
<i>AbstractNotificationMethod</i>	Class encapsulating the abstract notification method.
<i>EmailNotificationMethod</i>	Class sending out notifications in e-mail.
<i>NotificationPolicy</i>	Class encapsulating a NotificationPolicy which describes how to send out notifications.

Table 5.83. Classes of the Notification module

### 5.9.2. Class AbstractNotificationMethod

This abstract class encapsulates a notification that is performed when a certain event occurs.

Specialized classes can be derived from AbstractNotification, such as the *EmailNotificationMethod* class.

### 5.9.3. Class EmailNotificationMethod

This class encapsulates a notification handler that sends an e-mail with the given mail properties.

#### 5.9.3.1. Attributes of EmailNotificationMethod

<b>recipient (string)</b>
Default: n/a
The e-mail address of the recipient.

#### 5.9.3.2. EmailNotificationMethod methods

Method	Description
<code>__init__(self, recipient)</code>	Constructor to initialize an EmailNotification instance.

Table 5.84. Method summary

#### Method `__init__(self, recipient)`

This constructor initializes an EmailNotification instance and sets the attributes of the outgoing e-mail.

### Arguments of `__init__`

<b>recipient (string)</b>
Default: n/a
The e-mail address of the recipient.

## 5.9.4. Class NotificationPolicy

### 5.10. Module Proxy

This module encapsulates the Proxy component. The Proxy module provides a common framework for protocol-specific proxies, implementing the functions that are used by all proxies. Protocol-specific proxy modules are derived from the Proxy module, and are described in *Chapter 4, Proxies (p. 19)*.

#### 5.10.1. Functions in module Proxy

Function	Description
<u><code>proxyLog</code></u>	Function to send a proxy-specific message to the system log.

Table 5.85. Function summary

#### 5.10.2. Classes in the Proxy module

Class	Description
<u><code>Proxy</code></u>	Class encapsulating the abstract proxy.

Table 5.86. Classes of the Proxy module

### 5.10.3. Functions

#### 5.10.3.1. Function `proxyLog(self, type, level, msg, args)`

This function sends a message into the system log. All messages start with the `session_id` that uniquely identifies the connection.

#### Arguments of `proxyLog`

<b>level (integer)</b>
Default: n/a
Verbosity level of the log message.

<b>msg (string)</b>
Default: n/a
The text of the log message.

<b>type (string)</b>
Default: n/a
The class of the log message.

### 5.10.4. Class Proxy

This class serves as the abstract base class for all proxies. When an instance of the Proxy class is created, it loads and starts a protocol-specific proxy. Proxies operate in their own threads, so this constructor returns immediately.

#### 5.10.4.1. Attributes of Proxy

<b>encryption_policy (class)</b>
Default: None
Name of the Encryption policy instance used to encrypt the sessions and verify the certificates used. For details, see <i>Section 5.5, Module Encryption (p. 116)</i> .

<b>language (string)</b>
Default: "en"
Determines the language used for user-visible error messages. Supported languages: <i>en</i> - English; <i>de</i> - German; <i>hu</i> - Hungarian.

#### 5.10.4.2. Proxy methods

Method	Description
<i>closedByAbort(self)</i>	Function called by the proxy core when an abort has been occurred.
<i>config(self)</i>	Function called by the proxy core to initialize the proxy instance.
<i>connectServer(self)</i>	Function called by the proxy instance to establish the server-side connection.
<i>getCredentials(self, method, username, domain, target, port)</i>	Function called when proxy requires credentials for server side authentication.
<i>invalidPolicyCall(self)</i>	Invalid policy function called.

Method	Description
<i>setServerAddress(self, host, port)</i>	Function called by the proxy instance to set the address of the destination server.
<i>setServerSideEncryption(self)</i>	Function called by the proxy instance to set up the server side encryption parameters dynamically.
<i>userAuthenticated(self, entity, groups, auth info)</i>	Function called when inband authentication is successful.

Table 5.87. Method summary

### Method closedByAbort(self)

This function is called when a callback gives abort or no result. It simply sets a flag that will be used for logging the reason of the proxy's ending.

### Method config(self)

This function is called during proxy startup. It sets the attributes of the proxy instance according to the configuration of the proxy.

### Method connectServer(self)

This function is called to establish the server-side connection. The function either connects a proxy to the destination server, or an embedded proxy to its parent proxy. The proxy may set the address of the destination server using the `setServerAddress` function.

The `connectServer` function calls the chainer specified in the service definition to connect to the remote server using the host name and port parameters.

The `connectServer` function returns the descriptor of the server-side data stream.

### Method getCredentials(self, method, username, domain, target, port)

The proxy instance calls this function to retrieve authentication credentials for authentication method *method* and the target user *username*.

### Arguments of getCredentials

<b>domain (string)</b>
Default: n/a
Domain the user name belongs to.
<b>method (string)</b>
Default: n/a
Method that will be used for authentication on target server.

<b>port (integer)</b>
Default: n/a
Target server port.

<b>target (string)</b>
Default: n/a
Target server hostname.

<b>username (string)</b>
Default: n/a
Username that will be used for authentication on target server.

#### Method `invalidPolicyCall(self)`

This function is called when invalid policy function has been called.

#### Method `setServerAddress(self, host, port)`

The proxy instance calls this function to set the address of the destination server. This function attempts to resolve the hostname of the server using the DNS; the result is stored in the `session.server_address` parameter. The address of the server may be modified later by the router of the service. See *Section 5.12, Module Router (p. 186)* for details.



#### Note

The `setServerAddress` function has effect only when *InbandRouter* is used.

#### Arguments of `setServerAddress`

<b>host (string)</b>
Default: n/a
The host name of the server.

<b>port (integer)</b>
Default: n/a
The Port number of the server.

### Method `setServerSideEncryption(self)`

Function called by the proxy instance when the encryption scenario is dynamic (eg.: `DynamicServerEncryption`) to set up the server side encryption parameters. It should return with a `DynamicServerEncryptionServerParams` if `DynamicServerEncryption` scenario used otherwise with `None`.

This method unconditionally raises a `NotImplementedError` exception to indicate that it must be overridden by descendant classes like 'Proxy'.

### Method `userAuthenticated(self, entity, groups, auth_info)`

The proxy instance calls this function to indicate that the inband authentication was successfully performed. The name of the client is stored in the `entity` parameter.

### Arguments of `userAuthenticated`

<b>entity (unknown)</b>
Default: n/a
Username of the authenticated client.

## 5.11. Module Resolver

This module defines the `AbstractResolver` interface and various derived classes to perform name lookups.

### 5.11.1. Classes in the Resolver module

Class	Description
<i><code>AbstractResolver</code></i>	Class encapsulating the abstract Resolver interface.
<i><code>DNSResolver</code></i>	Class encapsulating DNS-based name resolution.
<i><code>HashResolver</code></i>	Class encapsulating hash-based name resolution.

Table 5.88. Classes of the Resolver module

### 5.11.2. Class `AbstractResolver`

This class encapsulates an interface for application level name resolution.

### 5.11.3. Class `DNSResolver`

`DNSResolver` policies query the domain name server to resolve domain names.



#### Example 5.31. A simple `DNSResolver` policy

Below is a simple `DNSResolver` policy enabled to return multiple 'A' and 'AAAA' records from the nameserver 1.1.1.1 with 2s timeout.

```
ResolverPolicy(name="Mailservers", resolver=DNSResolver(name_server='1.1.1.1', timeout=2))
```



### 5.11.3.1. DNSResolver methods

Method	Description
<code>__init__(self, name_server, timeout, use_search_domain)</code>	Constructor to initialize a DNSResolver instance.

Table 5.89. Method summary

#### Method `__init__(self, name_server, timeout, use_search_domain)`

This constructor initializes a DNSResolver instance.

#### Arguments of `__init__`

<b>name_server (string)</b>
Default: None
IP address of the DNS server to query. Defaults to the servers set in the <code>resolv.conf</code> file.
<b>timeout (integer)</b>
Default: 2
Seconds to wait a response from a server.
<b>use_search_domain (boolean)</b>
Default: FALSE
Append the host's search domain to the query.

### 5.11.4. Class HashResolver

HashResolver policies are used to locally store the IP addresses belonging to a domain name. A domain name (Hostname) and one or more corresponding IP addresses (Addresses) can be stored in a hash. If the domain name to be resolved is not included in the hash, the name resolution will fail. The HashResolver can be used to direct incoming connections to specific servers based on the target domain name.



#### Example 5.32. A simple HashResolver policy

The resolver policy below associates the IP addresses `192.168.1.12` and `192.168.1.13` with the `mail.example.com` domain name.

```
ResolverPolicy(name="DMZ", resolver=HashResolver(mapping={"mail.example.com": ("192.168.1.12", "192.168.1.13"))})
```



### 5.11.4.1. HashResolver methods

Method	Description
<code>__init__(self, mapping)</code>	Constructor to initialize a HashResolver instance.

Table 5.90. Method summary

#### Method `__init__(self, mapping)`

This constructor initializes a HashResolver instance.

#### Arguments of `__init__`

<b>mapping (complex)</b>
Default: n/a
Mapping that describes hostname->IP address pairs.

## 5.12. Module Router

Routers define the target IP address and port of the destination server, based on information that is available before started. The simplest router (*DirectedRouter*) selects a preset destination as the server address, while the most commonly used *TransparentRouter* connects to the IP address requested by the client. Other routers may make more complex decisions. The destination address selected by the router may be overridden by the proxy and the DNAT classes used in the service.

### 5.12.1. The source address used in the server-side connection

Routers also define source address and port of the server-side connection. This is the IP address that is used to connect the server. The server sees that the connection originates from this address. The following two parameters determine the source address used in the server-side connection:

*forge\_addr*: If set to *TRUE*, the client's source address is used as the source of the server-side connection. Otherwise, the IP address of the interface connected to the server is used.

*forge\_port*: This parameter defines the source port that is used in the server-side connection. Specify a port number as an integer value, or use one of the following options:

Name	Description
Z_PORT_ANY	Selected a random port between 1024 and 65535. This is the default behavior of every router.
Z_PORT_GROUP	Select a random port in the same group as the port used by the client. The following groups are defined: 0-513, 514-1024, 1025-.
Z_PORT_EXACT	Use the same port as the client.



Name	Description
Z_PORT_RANDOM	Select a random port using a cryptographically secure function.

Table 5.91. Options defining the source port of the server-side connection

### 5.12.2. Classes in the Router module

Class	Description
<i>AbstractRouter</i>	Class encapsulating the abstract router.
<i>DirectedRouter</i>	Class encapsulating a Router which explicitly defines the target address.
<i>InbandRouter</i>	Class encapsulating the Router which extracts the destination address from the application-level protocol.
<i>TransparentRouter</i>	Class encapsulating a Router which provides transparent services.

Table 5.92. Classes of the Router module

### 5.12.3. Class AbstractRouter

AbstractRouter implements an abstract router that determines the destination address of the server-side connection. Service definitions should refer to a customized class derived from AbstractRouter, or one of the predefined router classes, such as *TransparentRouter* or *DirectedRouter*. Different implementations of this interface perform Transparent routing (directing the client to its original destination), and Directed routing (directing the client to a given destination).

A proxy can override the destination selected by the router using the the *setServerAddress* method.

#### 5.12.3.1. Attributes of AbstractRouter

<b>forge_addr (boolean)</b>
Default: n/a
If set to <i>TRUE</i> , the client's source address is used as the source of the server-side connection.
<b>forge_port (unknown)</b>
Default: n/a
Defines the source port that is used in the server-side connection. See <i>Section 5.12.1, The source address used in the server-side connection (p. 186)</i> for details.

### 5.12.4. Class DirectedRouter

This class implements directed routing, which means that the destination address is a preset address for each session.



#### Example 5.33. DirectedRouter example

The following service uses a DirectedRouter that redirects every connection to the `/var/sample.socket` Unix domain socket.

```
Service(name="demo_service", proxy_class=HttpProxy,
router=DirectedRouter(dest_addr=SockAddrUnix('/var/sample.socket'), overrideable=False,
forge_addr=False))
```

The following service uses a DirectedRouter that redirects every connection to the `192.168.2.24:8080` IP address.

```
Service(name="demo_service", proxy_class=HttpProxy,
router=DirectedRouter(dest_addr=SockAddrInet('192.168.2.24', 8080), overrideable=False,
forge_addr=False))
```

#### 5.12.4.1. Attributes of DirectedRouter

<b>dest_addr (unknown)</b>
Default: n/a
The destination address to connect to.

#### 5.12.4.2. DirectedRouter methods

Method	Description
<code>__init__(self, dest_addr, forge_addr, overrideable, forge_port)</code>	Constructor to initialize a DirectedRouter.

Table 5.93. Method summary

#### Method `__init__(self, dest_addr, forge_addr, overrideable, forge_port)`

This constructor initializes an instance of the DirectedRouter class.

#### Arguments of `__init__`

<b>dest_addr (complex)</b>
Default: n/a
The destination address to connect to.
<b>forge_addr (boolean)</b>
Default: FALSE
If set to <code>TRUE</code> , the client's source address is used as the source of the server-side connection.

**forge\_port (complex)**

Default: Z\_PORT\_ANY

Defines the source port that is used in the server-side connection. See *Section 5.12.1, The source address used in the server-side connection (p. 186)* for details.

**overrideable (boolean)**

Default: FALSE

If set to *TRUE*, the proxy may override the selected destination. Enable this option when the proxy builds multiple connections to the destination server, and the proxy knows the address of the destination server, for example, because it receives a redirect request. This situation is typical for the SQLNet proxy.

### 5.12.5. Class InbandRouter

This class implements inband routing, which means that the destination address will be determined by the protocol. Inband routing works only for protocols that can send routing information within the protocol, and is mainly used for non-transparent proxying. The InbandRouter class currently supports only the HTTP and FTP protocols.

**Example 5.34. InbandRouter example**

The following service uses an InbandRouter to extract the destination from the protocol.

```
Service(name="demo_service", proxy_class=HttpProxy, router=InbandRouter(forge_addr=FALSE))
```

#### 5.12.5.1. InbandRouter methods

Method	Description
<code>__init__(self, forge_addr, forge_port)</code>	Constructor to initialize a InbandRouter.

Table 5.94. Method summary

**Method `__init__(self, forge_addr, forge_port)`**

This constructor initializes an instance of the InbandRouter class.

**Arguments of `__init__`****forge\_addr (boolean)**

Default: FALSE

If set to *TRUE*, the client's source address is used as the source of the server-side connection.

**forge\_port (complex)**

Default: Z\_PORT\_ANY

Defines the source port that is used in the server-side connection. See *Section 5.12.1, The source address used in the server-side connection (p. 186)* for details.

**5.12.6. Class TransparentRouter**

This class implements transparent routing, which means that the destination server is the original destination requested by the client.

**Example 5.35. TransparentRouter example**

The following service uses a TransparentRouter that connects to the 8080 port of the server and uses the client's IP address as the source of the server-side connection.

```
Service(name="demo_service", proxy_class=HttpProxy, router=TransparentRouter(forced_port=8080,
overrideable=FALSE, forge_addr=TRUE))
```

**5.12.6.1. Attributes of TransparentRouter****forced\_port (unknown)**

Default: n/a

Defines the source port that is used in the server-side connection. See *Section 5.12.1, The source address used in the server-side connection (p. 186)* for details.

**forge\_addr (unknown)**

Default: n/a

If set to *TRUE*, the client's source address is used as the source of the server-side connection.

**5.12.6.2. TransparentRouter methods**

Method	Description
<code>__init__(self, forced_port, forge_addr, overrideable, forge_port)</code>	Constructor to initialize an instance of the TransparentRouter class.

Table 5.95. Method summary

**Method `__init__(self, forced_port, forge_addr, overrideable, forge_port)`**

This constructor creates a new TransparentRouter instance which can be associated with a *Service*.

## Arguments of `__init__`

<b>forced_port (integer)</b>
Default: 0
Modify the destination port to this value. Default value: 0 (do not modify the target port)
<b>forge_addr (boolean)</b>
Default: FALSE
If set to <i>TRUE</i> , the client's source address is used as the source of the server-side connection.
<b>forge_port (complex)</b>
Default: <i>Z_PORT_ANY</i>
Defines the source port that is used in the server-side connection. See <i>Section 5.12.1, The source address used in the server-side connection (p. 186)</i> for details.
<b>overrideable (boolean)</b>
Default: FALSE
If set to <i>TRUE</i> , the proxy may override the selected destination. Enable this option when the proxy builds multiple connections to the destination server, and the proxy knows the address of the destination server, for example, because it receives a redirect request. This situation is typical for the SQLNet proxy.

## 5.13. Module Rule

The Rule module defines the classes needed to create firewall rules.

### 5.13.1. Evaluating firewall rules

When Zorp receives a connection request from a client, it tries to select a rule matching the parameters of the connection. The following parameters are considered.

Name in	Name in <code>policy.py</code>
VPN	<i>reqid</i>
Source Interface	<i>src_iface</i>
Source Interface Group	<i>src_ifgroup</i>
Protocol	<i>proto</i>
Source Port	<i>src_port</i>
Destination Port	<i>dst_port</i>
Source Subnet	<i>src_subnet</i>



Name in	Name in policy.py
Source Zone	<i>src_zone</i>
Destination Subnet	<i>dst_subnet</i>
Destination Interface	<i>dst_iface</i>
Destination Interface Group	<i>dst_ifgroup</i>
Destination Zone	<i>dst_zone</i>

Table 5.96. Evaluated Rule parameters

Zorp selects the rule that most specifically matches the connection. Selecting the most specific rule is based on the following method.

- The order of the rules is not important.
- The parameters of the connection act as filters: if you do not set any parameters, the rule will match any connection.
- If multiple connections would match a connection, the rule with the most-specific match is selected. For example, you have configured two rules: the first has the *Source Zone* parameter set as the *office* (which is a zone covering all of your client IP addresses), the second has the *Source Subnet* parameter set as *192.168.15.15/32*. The other parameters of the rules are the same. If a connection request arrives from the *192.168.15.15/32* address, Zorp will select the second rule. The first rule will match every other client request.
- Zorp considers the parameters of a connection in groups. The first group is the least-specific, the last one is the most-specific. The parameter groups are listed below.
- The parameter groups are linked with a logical AND operator: if parameters of multiple groups are set in a rule, the connection request must match a parameter of every group. For example, if both the *Source Interface* and *Destination Port* are set, the connection must match both parameters.
- Parameters within the same group are linked with a logical OR operator: if multiple parameters of a group are set for a rule, the connection must match any one of the parameters. If there are multiple similar rules, the rule with the most specific parameter match for the connection will be selected.

**Note**

In general, avoid using multiple parameters of the same group in one rule, as it may lead to undesired side-effects. Use only the most specific parameter matching your requirements.

For example, suppose that you have a rule with the *Destination Zone* parameter set, and you want to create a similar rule for a specific subnet of this zone. In this case, create a new rule with the *Destination Subnet* parameter set, do not set the *Destination Zone* parameter in both rules. Setting the *Destination Zone* parameter in both rules and setting the *Destination Subnet* parameter in the second rule would work for connections targeting the specified subnet, but it would cause Zorp to reject the connections that target other subnets of the specified destination zone, because both rules would match for the connection.

- The parameter groups are the following from the least specific to the most specific ones. Parameters within each group are listed from left to right from the least specific to the most specific ones.



1. *Destination Zone > Destination Interface Group > Destination Interface > Destination Subnet*
  2. *Source Zone > Source Subnet*
  3. *Destination Port* (Note that port is more specific than port range.)
  4. *Source Port* (Note that port is more specific than port range.)
  5. *Protocol*
  6. *Source Interface Group > Source Interface > VPN*
- If no matching rule is found, Zorp rejects the connection.

**Note**

It is possible to create rules that are very similar, making debugging difficult.

### 5.13.2. Sample rules

**Example 5.36. Sample rule definitions**

The following rule starts the service called *MyPFService* for every incoming TCP connection (*proto=6*).

```
Rule(proto=6,
     service='MyPFService'
    )
```

The following rule starts a service for TCP or UDP connections from the *office* zone.

```
Rule(proto=(6,17),
     src_zone='office',
     service='MyService'
    )
```

The following rule permits connections from the *192.168.0.0/16* IPv4 and the *2001:db8:c001:ba80::/58* IPv6 subnets. Note that since the *src\_subnet* parameter has two values, they are specified as a Python tuple: (*'value1', 'value2'*).

```
Rule(proto=6,
     src_subnet=('192.168.0.0/16', '2001:db8:c001:ba80::/58'),
     service='MyService'
    )
```

The following rule has almost every parameter set:

```
Rule(src_iface=('eth0', ),
     proto=6,
     dst_port=443,
     src_subnet=('192.168.10.0/24', ),
     src_zone=('office', ),
     dst_subnet=('192.168.50.50/32', ),
     dst_zone=('finance', ),
     service='MyHttpsService'
    )
```

### 5.13.3. Adding metadata to rules: tags and description

To make the configuration file more readable and informative, you can add descriptions and tags to the rules. Descriptions can be longer texts, while tags are simple labels, for example, to identify rules that belong to the

same type of traffic. Adding metadata to rules is not necessary, but can be a great help when maintaining large configurations.

- To add a description to a rule, add the text of the description before the rule, enclosed between three double-quotes:

```
"""This rule is ..."""
```

- To tag a rule, add a comment line before the rule that contains the list of tags applicable to the rule, separated with commas.

```
#Tags: tag1, tag2
```



#### Example 5.37. Tagging rules

The following rule has two tags, marking the traffic type and the source zone: *http* and *office*.

```
#Tags: http, office
"""Description"""
Rule(proto=(6),
      src_zone='office',
      service='MyHttpService'
    )
```

### 5.13.4. Classes in the Rule module

Class	Description
<i>PortRange</i>	Specifies a port range for a rule
<i>Rule</i>	This class implements firewall rules

Table 5.97. Classes of the Rule module

### 5.13.5. Class PortRange

This class specifies a port range for a firewall rule. It can be used in the *src\_port* and *dst\_port* parameters of a rule. For example: *src\_port=PortRange(2000, 2100)*, or *src\_port=(PortRange(2000, 2100), PortRange(2500, 2600))*. When listing multiple elements, ports and port ranges can be mixed, for example: *src\_port=(4433, PortRange(2000, 2100), PortRange(2500, 2600))*

#### 5.13.5.1. Attributes of PortRange

<b>high (integer)</b>
Default: n/a
The higher value of the port range.
<b>low (integer)</b>
Default: n/a



<b>low (integer)</b>
The lower value of the port range.

### 5.13.6. Class Rule

This class implements firewall rules. For details, see *Section 5.13, Module Rule (p. 191)*.

#### 5.13.6.1. Rule methods

Method	Description
<code>__init__(self, **kw)</code>	Initializes a rule

Table 5.98. Method summary

#### Method `__init__(self, **kw)`

Initializes a rule

#### Arguments of `__init__`

<b>dst_iface (interface)</b>
Default: n/a
Permit traffic only for connections that target a configured IP address of the listed interfaces. This parameter can be used to provide nontransparent service on an interface that received its IP address dynamically. For example, <code>dst_iface='eth0'</code> , or <code>dst_iface=('eth0', 'tun1')</code> .

<b>dst_ifgroup (integer)</b>
Default: n/a
Permit traffic only for connections that target a configured IP address of the listed interface group. This parameter can be used to provide nontransparent service on an interface that received its IP address dynamically. For example, <code>dst_ifgroup=1</code> .

<b>dst_port (integer)</b>
Default: n/a
Permit traffic only if the client targets the listed port. For example, <code>dst_port=80</code> , or <code>dst_port=(80, 443)</code> . To specify port ranges, use the <code>PortRange</code> class, for example, <code>dst_port=PortRange(2000, 2100)</code> .

<b>dst_subnet (subnet)</b>
Default: n/a
Permit traffic only for connections targeting a listed IP address, or an address belonging to the listed subnet. The subnet can be IPv4 or IPv6 subnet. When listing multiple subnets, you can list both IPv4 and IPv6 subnets.



<b>dst_subnet (subnet)</b>
IP addresses are treated as subnets with a /32 (IPv4) or /128 (IPv6) netmask. If no netmask is set for a subnet, it is treated as a specific IP address. For example, <code>dst_subnet='192.168.10.16'</code> or <code>dst_subnet=('192.168.0.0/16', '2001:db8:c001:ba80::/58')</code> .
<b>dst_zone (zone)</b>
Default: n/a
Permit traffic only for connections targeting an address belonging to the listed zones. For example, <code>dst_zone='office'</code> or <code>dst_zone=('office', 'finance')</code> . Note that this applies to destination address of the client-side connection request: the actual address of the server-side connection can be different (for example, if a DirectedRouter is used in the service).
<b>proto (integer)</b>
Default: n/a
Permit only connections using the specified transport protocol. This is the transport layer (Layer 4) protocol of the OSI model, for example, TCP, UDP, ICMP, and so on. The protocol must be specified using a number: the decimal value of the "protocol" field of the IP header. This value is 6 for the TCP and 17 for the UDP protocol. For a list of protocol numbers, see the <a href="#">Assigned Internet Protocol Numbers page of IANA</a> . For example: <code>proto=(6,17)</code> . To permit any protocol, do not add the <code>proto</code> parameter to the rule.
<b>rule_id (integer)</b>
Default: n/a
A unique ID number for the rule. This parameter is optional, an ID number is automatically generated for the rule during startup.
<b>service (service)</b>
Default: n/a
The name of the service to start for matching connections. This is the only required parameter for the rule, everything else is optional. For example, <code>service='MyService'</code>
<b>src_iface (interface)</b>
Default: n/a
Permit traffic only for connections received on the listed interface. For example, <code>src_iface='eth0'</code> , or <code>src_iface=('eth0', 'tun1'),</code> .
<b>src_ifgroup (integer)</b>
Default: n/a



<b>src_ifgroup (integer)</b>
Permit traffic only for connections received on the listed interfacegroup. For example, <i>src_iface=1</i> . Interface groups can be defined in the <code>/etc/network/interfaces</code> file, for example:
<pre>iface eth0 inet dhcp     group 1     iface eth1 inet dhcp     group 1</pre>
<b>src_port (integer)</b>
Default: n/a
Permit traffic only if the client sends the connection request from the listed port. For example, <i>src_port=4455</i> . To specify port ranges, use the <i>PortRange</i> class, for example, <i>src_port=PortRange(2000, 2100)</i> .
<b>src_subnet (subnet)</b>
Default: n/a
Permit traffic only for the clients of the listed subnet or IP addresses. The subnet can be IPv4 or IPv6 subnet. When listing multiple subnets, you can list both IPv4 and IPv6 subnets. IP addresses are treated as subnets with a /32 (IPv4) or /128 (IPv6) netmask. If no netmask is set for a subnet, it is treated as a specific IP address. For example, <i>src_subnet='192.168.10.16'</i> or <i>src_subnet=('192.168.0.0/16', '2001:db8:c001:ba80::/58')</i> .
<b>src_zone (zone)</b>
Default: n/a
Permit traffic only for the clients of the listed zones. For example, <i>src_zone='office'</i> or <i>src_zone=('office', 'finance')</i> .

## 5.14. Module Service

This module defines classes encapsulating service descriptions. The services define how the incoming connection requests are handled. When a connection is accepted by a *Rule*, the service specified in the Rule creates an instance of itself. This instance handles the connection, and proxies the traffic between the client and the server. It also handles TLS and SSL encryption of the traffic if needed, as configured in the *encryption\_policy* parameter of the service. The instance of the selected service is created using the *'startInstance()'* method.

A service is not usable on its own, it needs a *Rule* to bind the service to a network interface of the firewall and activate it when a matching connection request is received. New instances of the service are started as the Rule accepts new connections.

### 5.14.1. Naming services

The name of the service must be a unique identifier; rules refer to this unique ID.

Use clear, informative, and consistent service names. Include the following information in the service name:

- Source zones, indicating which clients may use the service (e.g., *intranet*).
- The protocol permitted in the traffic (e.g., *HTTP*).
- Destination zones, indicating which servers may be accessed using the service (e.g., *Internet*).



**Tip**

Name the service that allows internal users to browse the Web *intra\_HTTP\_internet*. Use dots to indicate child zones, e.g., *intra.marketing\_HTTP\_inter*.

### 5.14.2. Classes in the Service module

Class	Description
<i>AbstractService</i>	Class encapsulating the abstract Service properties.
<i>DenyService</i>	DenyService prohibits access to certain services
<i>PFSERVICE</i>	Class encapsulating a packet-filter service definition.
<i>Service</i>	Class encapsulating a service definition.

Table 5.99. Classes of the Service module

### 5.14.3. Class AbstractService

AbstractService implements an abstract service. Service definitions should be based on a customized class derived from AbstractService, or on the predefined *Service* class.

#### 5.14.3.1. Attributes of AbstractService

<b>name (string)</b>
Default: n/a
The name of the service.

#### 5.14.3.2. AbstractService methods

Method	Description
<i>__init__(self, name)</i>	Constructor to initialize an instance of the AbstractService class.

Table 5.100. Method summary

#### Method *\_\_init\_\_(self, name)*

This constructor creates an AbstractService instance and sets the attributes of the instance according to the received arguments. It also registers the Service to the *services* hash so that rules can find the service instance.

## Arguments of `__init__`

name (string)
Default: n/a
The name of the service.

### 5.14.4. Class DenyService

The DenyService class is a type of service that rejects connections with a predefined error code. DenyServices can be specified in the `service` parameter of *Rules*. If the rule referencing the DenyService matches a connection request, the connection is rejected. DenyService is a replacement for the obsolete Umbrella zone concept.




#### Example 5.38. A simple DenyService

The following defines a DenyService and a rule to reject all traffic that targets port 5555.

```
def demo() :
    DenyService(name='DenyService', ipv4_setting=DenyIPv4.DROP, ipv6_setting=DenyIPv6.DROP)
    Rule(dst_port=5555,
        service='DenyService'
    )
```

#### 5.14.4.1. Attributes of DenyService

ipv4_setting (complex)
Default: n/a
Specifies how to reject IPv4 traffic. By default, the traffic is simply dropped without notifying the client ( <i>DenyIPv4.DROP</i> ). The following values are available: <i>DenyIPv4.DROP</i> , <i>DenyIPv4.TCP_RESET</i> , <i>DenyIPv4.ICMP_NET_UNREACHABLE</i> , <i>DenyIPv4.ICMP_HOST_UNREACHABLE</i> , <i>DenyIPv4.ICMP_PROTO_UNREACHABLE</i> , <i>DenyIPv4.ICMP_PORT_UNREACHABLE</i> , <i>DenyIPv4.ICMP_NET_PROHIBITED</i> , <i>DenyIPv4.ICMP_HOST_PROHIBITED</i> , <i>DenyIPv4.ICMP_ADMIN_PROHIBITED</i>
<div style="display: flex; align-items: flex-start;"> <div style="margin-right: 10px;">  </div> <div> <p><b>Note</b> When the <i>DenyIPv4.TCP_RESET</i> option is used, the TCP RESET packet is sent as if it was sent by the target server.</p> <p>When using an ICMP option, the appropriate ICMP packet is sent, just like a router would.</p> </div> </div>

ipv6_setting (complex)
Default: n/a
Specifies how to reject IPv6 traffic. By default, the traffic is dropped without notifying the client ( <i>DenyIPv6.DROP</i> ). The following values are available: <i>DenyIPv6.DROP</i> , <i>DenyIPv6.TCP_RESET</i> , <i>DenyIPv6.ICMP_NO_ROUTE</i> , <i>DenyIPv6.ICMP_ADMIN_PROHIBITED</i> , <i>DenyIPv6.ICMP_ADDR_UNREACHABLE</i> , <i>DenyIPv6.ICMP_PORT_UNREACHABLE</i>

<b>name (string)</b>
Default: n/a
The name of the service.

#### 5.14.4.2. DenyService methods

Method	Description
<code>__init__(self, name, logging, ipv4_setting, ipv6_setting, log_verbose, log_spec)</code>	Constructor to initialize a DenyService instance.

Table 5.101. Method summary

#### Method `__init__(self, name, logging, ipv4_setting, ipv6_setting, log_verbose, log_spec)`

This constructor defines a DenyService with the specified parameters.

#### Arguments of `__init__`

<b>log_spec (string)</b>
Default: None
Message filter expression.

<b>log_verbose (integer)</b>
Default: None
Default log verbosity level.

<b>name (string)</b>
Default: n/a
The name identifying the service.

#### 5.14.5. Class PFService

PFServices allow you to replace the FORWARD rules of iptables, and configure application-level and packet-filter rules from Zorp.



#### Note

The PFService class transfers packet-filter level services.

- To transfer connections on the packet-filter level, use the `PFService` class.
- To transfer connections on the application-level, use the `Service` class.



**Example 5.39. PFSERVICE example**

The following packet-filtering service transfers TCP connections that arrive to port 5555.

```
PFSERVICE(name="intranet_PF5555_internet", router=TransparentRouter())
```

The following example defines a few classes: the client and server zones, a simple services, and a rule that starts the service.

```
Zone('internet', ['0.0.0.0/0'])
Zone('intranet', ['192.168.0.0/16'])

def demo() :
    PFSERVICE(name="intranet_PF5555_internet", router=TransparentRouter())
    Rule(dst_port=5555,
         src_zone='intranet',
         dst_zone='internet',
         service='PFSERVICE'
        )
```

**5.14.5.1. Attributes of PFSERVICE**

<b>dnat_policy (class)</b>
Default: n/a
Name of the NAT policy instance used to translate the destination addresses of the sessions. See <i>Section 5.8, Module NAT (p. 170)</i> for details.
<b>router (class)</b>
Default: n/a
A router instance used to determine the destination address of the server. See <i>Section 5.12, Module Router (p. 186)</i> for details.
<b>snat_policy (class)</b>
Default: n/a
Name of the NAT policy instance used to translate the source addresses of the sessions. See <i>Section 5.8, Module NAT (p. 170)</i> for details.

**5.14.5.2. PFSERVICE methods**

<b>Method</b>	<b>Description</b>
<code>__init__(self, name, router, snat_policy, dnat_policy, log_verbose, log_spec)</code>	Constructor to initialize a PFSERVICE instance.

Table 5.102. Method summary

**Method `__init__(self, name, router, snat_policy, dnat_policy, log_verbose, log_spec)`**

This constructor defines a packetfilter-service with the specified parameters.

## Arguments of `__init__`

<b>log_spec (string)</b>
Default: None
Message filter expression.
<b>log_verbose (integer)</b>
Default: None
Default log verbosity level.

### 5.14.6. Class Service

A service is one of the fundamental objects. It stores the names of proxy-related parameters, and is also used for access control purposes to decide what kind of traffic is permitted.



#### Note

The Service class transfers application-level (proxy) services.

- To transfer connections on the packet-filter level, use the *PFService* class.
- To transfer connections on the application-level, use the *Service* class.



#### Example 5.40. Service example

The following service transfers HTTP connections. Every parameter is left at its default.

```
Service(name="demo_http, proxy_class=HttpProxy, router=TransparentRouter())
```

The following service handles HTTP connections. This service uses authentication and authorization, and network address translation on the client addresses (SNAT).

```
Service(name="demo_http", proxy_class=HttpProxy, authentication_policy="demo_authentication_policy",
authorization_policy="demo_permituser", snat_policy="demo_natpolicy", router=TransparentRouter())
```

The following example defines a few classes: the client and server zones, a simple services, and a rule that starts the service.

```
Zone('internet', ['0.0.0.0/0'])
Zone('office', ['192.168.1.0/32', '192.168.2.0/32'])

def demo_instance() :
Service(name="office_http_inter", proxy_class=HttpProxy, router=TransparentRouter())
Rule(src_zone='office',
proto=6,
dst_zone='internet',
service='office_http_inter'
)
```

#### 5.14.6.1. Attributes of Service

<b>auth_name (string)</b>
Default: n/a



<b>auth_name (string)</b>
Authentication name of the service. This string informs the users of the Zorp Authentication Agent about which service they are authenticating for. Default value: the name of the service.
<b>authentication_policy (class)</b>
Default: n/a
Name of the AuthenticationPolicy instance used to authenticate the clients. See <i>Section 5.1, Module Auth (p. 88)</i> for details.
<b>authorization_policy (class)</b>
Default: n/a
Name of the AuthorizationPolicy instance used to authorize the clients. See <i>Section 5.1, Module Auth (p. 88)</i> for details.
<b>chainer (class)</b>
Default: n/a
A chainer instance used to connect to the destination server. See <i>Section 5.3, Module Chainer (p. 103)</i> for details.
<b>dnat_policy (class)</b>
Default: n/a
Name of the NAT policy instance used to translate the destination addresses of the sessions. See <i>Section 5.8, Module NAT (p. 170)</i> for details.
<b>encryption_policy (class)</b>
Default: None
Name of the Encryption policy instance used to encrypt the sessions and verify the certificates used. For details, see <i>Section 5.5, Module Encryption (p. 116)</i> .
<b>instance_id (integer)</b>
Default: n/a
The sequence number of the last session started
<b>keepalive (integer)</b>
Default: Z_KEEPALIVE_NONE
The TCP keepalive option, one of the Z_KEEPALIVE_NONE, Z_KEEPALIVE_CLIENT, Z_KEEPALIVE_SERVER, Z_KEEPALIVE_BOTH values.



<b>max_instances (integer)</b>
Default: n/a
Permitted number of concurrent instances of this service. Usually each service instance handles one connection. The default value is 0, which allows unlimited number of instances.
<b>max_sessions (integer)</b>
Default: n/a
Maximum number of concurrent sessions handled by one thread.
<b>num_instances (integer)</b>
Default: n/a
The current number of running instances of this service.
<b>proxy_class (class)</b>
Default: n/a
Name of the proxy class instance used to analyze the traffic transferred in the session. See <i>Section 5.10, Module Proxy (p. 180)</i> for details.
<b>resolver_policy (unknown)</b>
Default: n/a
Name of the ResolvePolicy instance used to resolve the destination domain names. See <i>Section 5.11, Module Resolver (p. 184)</i> for details. Default value: <i>DNSResolver</i>
<b>router (class)</b>
Default: n/a
A router instance used to determine the destination address of the server. See <i>Section 5.12, Module Router (p. 186)</i> for details.
<b>snat_policy (class)</b>
Default: n/a
Name of the NAT policy instance used to translate the source addresses of the sessions. See <i>Section 5.8, Module NAT (p. 170)</i> for details.

### 5.14.6.2. Service methods

Method	Description
<code>__init__(self, name, proxy_class, router, chainer, snat_policy, snat, dnat_policy, dnat, authentication_policy, authorization_policy, max_instances, max_sessions, auth_name, resolver_policy, auth, auth_policy, keepalive, encryption_policy, limit_target_zones_to, detector_config, detector_default_service_name, session_counting)</code>	Constructor to initialize a Service instance.
<code>startInstance(self, session)</code>	Start a service instance.

Table 5.103. Method summary

**Method `__init__(self, name, proxy_class, router, chainer, snat_policy, snat, dnat_policy, dnat, authentication_policy, authorization_policy, max_instances, max_sessions, auth_name, resolver_policy, auth, auth_policy, keepalive, encryption_policy, limit_target_zones_to, detector_config, detector_default_service_name, session_counting)`**

This constructor defines a Service with the specified parameters.

#### Arguments of `__init__`

<b><code>auth_name</code> (string)</b>
Default: None
Authentication name of the service. This string informs the users of the Zorp Authentication Agent about which service they are authenticating for. Default value: the name of the service.
<b><code>authentication_policy</code> (class)</b>
Default: None
Name of the AuthenticationPolicy instance used to authenticate the clients. See <i>Section 5.1, Module Auth (p. 88)</i> for details.
<b><code>authorization_policy</code> (class)</b>
Default: None
Name of the AuthorizationPolicy instance used to authorize the clients. See <i>Section 5.1, Module Auth (p. 88)</i> for details.
<b><code>chainer</code> (class)</b>
Default: None



<b>chainer (class)</b>
Name of the chainer instance used to connect to the destination server. Defaults to <i>ConnectChainer</i> if no other chainer is specified.
<b>dnat_policy (class)</b>
Default: None
Name of the NAT policy instance used to translate the destination addresses of the sessions. See <i>Section 5.8, Module NAT (p. 170)</i> for details.
<b>encryption_policy (class)</b>
Default: None
Name of the Encryption policy instance used to encrypt the sessions and verify the certificates used. For details, see <i>Section 5.5, Module Encryption (p. 116)</i> .
<b>keepalive (integer)</b>
Default: Z_KEEPALIVE_NONE
The TCP keepalive option, one of the Z_KEEPALIVE_NONE, Z_KEEPALIVE_CLIENT, Z_KEEPALIVE_SERVER, Z_KEEPALIVE_BOTH values.
<b>limit_target_zones_to (complex)</b>
Default: None
A comma-separated list of zone names permitted as the target of the service. No restrictions are applied if the list is empty. Use this parameter to replace the obsolete <i>inbound_services</i> parameter of the Zone class.
<b>max_instances (integer)</b>
Default: 0
Permitted number of concurrent instances of this service. Usually each service instance handles one connection. Default value: 0 (unlimited).
<b>max_sessions (integer)</b>
Default: 0
Maximum number of concurrent sessions handled by one thread.
<b>name (string)</b>
Default: n/a
The name identifying the service.



<b>proxy_class (class)</b>
Default: n/a
Name of the proxy class instance used to analyze the traffic transferred in the session. See <i>Section 5.10, Module Proxy (p. 180)</i> for details.

<b>resolver_policy (class)</b>
Default: None
Name of the ResolvePolicy instance used to resolve the destination domain names. See <i>Section 5.11, Module Resolver (p. 184)</i> for details. Default value: <i>DNSResolver</i> .

<b>router (class)</b>
Default: None
Name of the router instance used to determine the destination address of the server. Defaults to <i>TransparentRouter</i> if no other router is specified.

<b>snat_policy (class)</b>
Default: None
Name of the NAT policy instance used to translate the source addresses of the sessions. See <i>Section 5.8, Module NAT (p. 170)</i> for details.

#### Method startInstance(self, session)

Called by the Rule to create an instance of this service.

#### Arguments of startInstance

<b>session (unknown)</b>
Default: n/a
The session object

### 5.15. Module Session

This module defines the abstract session interface in a class named *AbstractSession*, and two descendants *MasterSession* and *StackedSession*.

Sessions are hierarchically stacked into each other just like proxies. All sessions except the master session have a parent session from which child sessions inherit variables. Child sessions are stacked into their master sessions, so stacked sessions can inherit data from the encapsulating proxy instances. (Inheritance is implemented using a simple `getattr` wrapper.)



Instances of the Session classes store the parameters of the client-side and server-side connections in a session object (for example, the IP addresses and zone of the server and the client, and the username and group memberships of the user when authentication is used). Other components refer to this data when making various policy-based decisions.

### 5.15.1. Classes in the Session module

Class	Description
<u><i>StackedSession</i></u>	Class encapsulating a subsession.

Table 5.104. Classes of the Session module

### 5.15.2. Class StackedSession

This class represents a stacked session, e.g., a session within the session hierarchy. Every subsession inherits session-wide parameters from its parent.

#### 5.15.2.1. Attributes of StackedSession

<b>chainer (class)</b>
Default: n/a
The chainer used to connect to the parent proxy. If unset, the <i>server_stream</i> parameter must be set.
<b>owner (class)</b>
Default: n/a
The parent session of the current session.
<b>server_address (class)</b>
Default: n/a
The IP address to connect. Most often this is the IP address requested by the client, but the client requests can be redirected to different IPs.
<b>server_local (class)</b>
Default: n/a
The server is connected from this IP address. This is either the IP address of Zorp's external interface, or the IP address of the client (if Forge Port is enabled). The client's original IP address may be modified if SNAT policies are used.
<b>server_stream (class)</b>
Default: n/a
Server-side stream.



<b>server_zone (class)</b>
Default: n/a
Zone of the server.
<b>target_address (class)</b>
Default: n/a
The IP address to connect. Most often this is the IP address requested by the client, but the client requests can be redirected to different IPs.
<b>target_local (class)</b>
Default: n/a
The server is connected from this IP address. This is either the IP address of Zorp's external interface, or the IP address of the client (if Forge Port is enabled). The client's original IP address may be modified if SNAT policies are used.
<b>target_zone (class)</b>
Default: n/a
Zone of the server.

### 5.15.2.2. StackedSession methods

Method	Description
<i>setTargetAddress(self, addr)</i>	Set the target server address.

Table 5.105. Method summary

#### Method setTargetAddress(self, addr)

This is a compatibility function for proxies that override the routed target.

#### Arguments of setTargetAddress

<b>addr (unknown)</b>
Default: n/a
Server address

## 5.16. Module SockAddr

This module implements *inet\_ntoa* and *inet\_aton*. The module also provides an interface to the SockAddr services of the Zorp core. SockAddr is used for example to define the address of the ZAS server in *AuthenticationProvider* policies.

### 5.16.1. Classes in the SocketAddr module

Class	Description
<i>SocketAddrInet</i>	Class encapsulating an IPv4 address:port pair.
<i>SocketAddrInet6</i>	Class encapsulating an IPv6 address:port pair.
<i>SocketAddrInetHostname</i>	Class encapsulating a hostname:port pair.
<i>SocketAddrInetRange</i>	Class encapsulating an IPv4 address and a port range.
<i>SocketAddrUnix</i>	Class encapsulating a UNIX domain socket.

Table 5.106. Classes of the SocketAddr module

### 5.16.2. Class SocketAddrInet

This class encapsulates an IPv4 address:port pair, similarly to the *sockaddr\_in* struct in C. The class is implemented and exported by the Zorp core. The *SocketAddrInet* Python class serves only documentation purposes, and has no real connection to the behavior implemented in C.



#### Example 5.41. SocketAddrInet example

The following example defines an IPv4 address:port pair.

```
SocketAddrInet('192.168.10.10', 80)
```

The following example uses SocketAddrInet in a dispatcher.

```
Dispatcher(transparent=TRUE, bindto=DBSocketAddr(protocol=ZD_PROTO_TCP, sa=SocketAddrInet('192.168.11.11', 50080)), service="intra_HTTP_inter", backlog=255, rule_port="50080")
```

#### 5.16.2.1. Attributes of SocketAddrInet

<b>ip (unknown)</b>
Default: n/a
IP address (network byte order).
<b>ip_s (unknown)</b>
Default: n/a
IP address in string representation.
<b>port (unknown)</b>
Default: n/a
Port number (network byte order).

<b>type (string)</b>
Default: n/a
The <i>inet</i> value that indicates an address in the AF_INET domain.

### 5.16.3. Class SockAddrInet6

This class encapsulates an IPv6 address:port pair, similarly to the *sockaddr\_in* struct in C. The class is implemented and exported by the *Zorp* core. The *SockAddrInet* Python class serves only documentation purposes, and has no real connection to the behavior implemented in C.



#### Example 5.42. SockAddrInet example

The following example defines an IPv6 address:port pair.

```
SockAddrInet('fec0::1', 80)
```

The following example uses SockAddrInet in a dispatcher.

```
Dispatcher(transparent=TRUE, bindto=DBSocketAddr(protocol=ZD_PROTO_TCP, sa=SockAddrInet('fec0::1', 50080)), service="intra_HTTP_inter", backlog=255, rule_port="50080")
```

#### 5.16.3.1. Attributes of SockAddrInet6

<b>ip (unknown)</b>
Default: n/a
IP address (network byte order).

<b>ip_s (unknown)</b>
Default: n/a
IP address in string representation.

<b>port (unknown)</b>
Default: n/a
Port number (network byte order).

<b>type (string)</b>
Default: n/a
The <i>inet</i> value that indicates an address in the AF_INET domain.

### 5.16.4. Class SockAddrInetHostname

This class encapsulates a hostname:port or IPv4 address:port pair. Name resolution is only performed when creating the SockAddrInetHostname object (that is, during startup and reload). The class is implemented and

exported by the Zorp core. The *SockAddrInetHostname* Python class serves only documentation purposes, and has no real connection to the behavior implemented in C.



**Example 5.43. SockAddrInetHostname example**

The following example defines a hostname:port or IPv4 address:port pair.

```
SockAddrInetHostname('www.example.com', 80)
```

```
SockAddrInetHostname('192.168.10.10', 80)
```

The following example uses SockAddrInetHostname in a dispatcher.

```
Dispatcher(transparent=TRUE, bindto=DBSockAddr(protocol=ZD_PROTO_TCP,
sa=SockAddrInetHostname('www.example.com', 50080)), service="intra_HTTP_inter", backlog=255,
rule_port="50080")
```

### 5.16.4.1. Attributes of SockAddrInetHostname

<b>ip (unknown)</b>
Default: n/a
IP address (network byte order).
<b>ip_s (unknown)</b>
Default: n/a
IP address in string representation.
<b>port (unknown)</b>
Default: n/a
Port number (network byte order).
<b>type (string)</b>
Default: n/a
The <i>inet</i> value that indicates an address in the AF_INET domain.

### 5.16.5. Class SockAddrInetRange

A specialized SockAddrInet class which allocates a new port within the given range of ports when a dispatcher bounds to it. The class is implemented and exported by the Zorp core. The *SockAddrInetRange* Python class serves only documentation purposes, and has no real connection to the behavior implemented in C.

#### 5.16.5.1. Attributes of SockAddrInetRange

<b>ip (unknown)</b>
Default: n/a

<b>ip (unknown)</b>
IP address (network byte order).
<b>ip_s (unknown)</b>
Default: n/a
IP address in string representation.
<b>port (unknown)</b>
Default: n/a
Port number (network byte order).
<b>type (string)</b>
Default: n/a
The <i>inet</i> value that indicates an address in the AF_INET domain.

### 5.16.6. Class SockAddrUnix

This class encapsulates a UNIX domain socket endpoint. The socket is represented by a filename. The *SockAddrUnix* Python class serves only documentation purposes, and has no real connection to the behavior implemented in C.



#### Example 5.44. SockAddrUnix example

The following example defines a Unix domain socket.

```
SockAddrUnix('/var/sample.socket')
```

The following example uses SockAddrUnix in a DirectedRouter.

```
Service(name="demo_service", proxy_class=HttpProxy,
router=DirectedRouter(dest_addr=SockAddrUnix('/var/sample.socket'), overrideable=FALSE,
forge_addr=FALSE))
```

#### 5.16.6.1. Attributes of SockAddrUnix

<b>type (string)</b>
Default: n/a
The <i>unix</i> value that indicates an address in the UNIX domain.

### 5.17. Module Stack

Zorp is capable of stacking, that is, handing over parts of the traffic to other modules for further inspection (e.g., to other proxies to inspect embedded protocols, to content vectoring modules for virus filtering, etc.). The Stack module defines the classes required for this functionality.



Stacking in services is performed using *StackingProvider policies*, which reference the host that performs the stacked operations using the *RemoteStackingBackend* class.

### 5.17.1. Classes in the Stack module

Class	Description
<i>AbstractStackingBackend</i>	This is an abstract class, currently without any functionality.
<i>RemoteStackingBackend</i>	Constructor to initialize an instance of the RemoteStackingBackend class.
<i>StackingProvider</i>	This is a policy class that is used to reference a configured stacking provider in service definitions.

Table 5.107. Classes of the Stack module

### 5.17.2. Class AbstractStackingBackend

This is an abstract class, currently without any functionality.

### 5.17.3. Class RemoteStackingBackend

This class contains the address of the host that performs the stacked operations. It is typically used to access the Zorp Content Vectoring Server (ZCV) to perform virus filtering in the traffic. The remote backend can be accessed using the TCP protocol or a local socket, e.g., `RemoteStackingBackend(addr=(SockAddrInet('192.168.2.3', 1318),))` or `RemoteStackingBackend(addr=(SockAddrUnix('/var/run/zcv/zcv.sock'),))`.

#### 5.17.3.1. RemoteStackingBackend methods

Method	Description
<code>__init__(self, addr)</code>	

Table 5.108. Method summary

Method `__init__(self, addr)`

Arguments of `__init__`

<b>addr (complex)</b>
Default: n/a
The address of the remote backend in <i>SockAddrInet</i> or <i>SockAddrUnix</i> format. Separate addresses with commas to list more than one address for a backend. Zorp will connect to these addresses in a failover fashion.

### 5.17.4. Class StackingProvider

Instances of the `StackingProvider` class are policies that define which remote stacking backend a particular service uses to inspect the contents of the traffic.



#### Example 5.45. A simple `StackingProvider` class

The following class creates a simple stacking provider that can be referenced in service definitions. The remote host that provides the stacking services is located under the `192.168.12.12` IP address.

```
StackingProvider(name="demo_stackingprovider",
backend=RemoteStackingBackend(addr=(SockAddrInet('192.168.12.12', 1318),)))
```



#### Example 5.46. Using a `StackingProvider` in an FTP proxy

The following classes define a stacking provider that can be accessed a local ZCV instance using a domain socket. This service provider is then used to filter FTP traffic. The configuration of the ZCV (i.e., what modules it uses to filter the traffic is not discussed here).

```
class StackingFtpProxy(FtpProxy):
def config(self):
    super(StackingFtpProxy, self).config()
    self.request_stack["RETR"]=(FTP_STK_DATA, (Z_STACK_PROVIDER, "demo_stackingprovider",
"default_rulegroup"))

StackingProvider(name="demo_stackingprovider_socket",
backend=RemoteStackingBackend(addr=(SockAddrUnix('/var/run/zcv/zcv.sock'),)))
```

#### 5.17.4.1. `StackingProvider` methods

Method	Description
<code>__init__(self, name, backend)</code>	Constructor to initialize an instance of the <code>StackingProvider</code> class.

Table 5.109. Method summary

#### Method `__init__(self, name, backend)`

This constructor creates a `StackingProvider` instance and sets the attributes of the instance according to the received arguments.

#### Arguments of `__init__`

<b>backend (class)</b>
Default: n/a
A configured <code>RemoteStackingBackend</code> class containing the address of the remote stacking backend, e.g., <code>RemoteStackingBackend(addr=(SockAddrInet('192.168.2.3', 1318),))</code> or <code>RemoteStackingBackend(addr=(SockAddrUnix('/var/run/zcv/zcv.sock'),))</code> .
<b>name (string)</b>
Default: n/a

**name (string)**

Name of the Stacking provider policy. This name can be referenced in the service definitions.

## 5.18. Module Zone

This module defines the *Zone* class.

Zones are the basis of access control. A zone consists of a set of IP addresses, address ranges, or subnet. For example, a zone can contain an IPv4 or IPv6 subnet.

Zones are organized into a hierarchy created by the administrator. Child zones inherit the security attributes (set of permitted services etc.) from their parents. The administrative hierarchy often reflects the organization of the company, with zones assigned to the different departments.

When it has to be determined what zone a client belongs to, the most specific zone containing the searched IP address is selected. If an IP address belongs to two different zones, the most specific zone is selected.

**Example 5.47. Finding IP networks**

Suppose there are three zones configured: *Zone\_A* containing the *10.0.0.0/8* network, *Zone\_B* containing the *10.0.0.0/16* network, and *Zone\_C* containing the *10.0.0.25* IP address. Searching for the *10.0.44.0* network returns *Zone\_B*, because that is the most specific zone matching the searched IP address. Similarly, searching for *10.0.0.25* returns only *Zone\_C*.

This approach is used in the service definitions as well: when a client sends a connection request, the most specific zone containing the IP address of the client is looked up. Suppose that the clients in *Zone\_A* are allowed to use HTTP. If a client with IP *10.0.0.50* (thus belonging to *Zone\_B*) can only use HTTP if *Zone\_B* is the child of *Zone\_A*, or if a service definition explicitly permits *Zone\_B* to use HTTP.

**Example 5.48. Zone examples**

The following example defines a simple zone hierarchy. The following zones are defined:

- *internet*: This zone contains every possible IP addresses, if an IP address does not belong to another zone, than it belongs to the *internet* zone.
- *office*: This zone contains the *192.168.1.0/32* and *192.168.2.0/32* networks.
- *management*: This zone is separated from the *office* zone, because it contains an independent subnet *192.168.3.0/32*. But from the administrator's view, it is the child zone of the *office* zone, meaning that it can use (and accept) the same services as the *office* zone.
- *DMZ*: This is a separate zone.

```
Zone('internet', ['0.0.0.0/0', ':::0/0'])
Zone('office', ['192.168.1.0/32', '192.168.2.0/32'])
Zone('management', ['192.168.3.0/32'])
Zone('DMZ', ['10.50.0.0/32'])
```





### 5.18.1. Classes in the Zone module

Class	Description
<i>Zone</i>	Class encapsulating IP zones.

Table 5.110. Classes of the Zone module

### 5.18.2. Class Zone

This class encapsulates IPv4 and IPv6 zones.



**Example 5.49. Determining the zone of an IP address**

An IP address always belongs to the most specific zone. Suppose that *Zone A* includes the IP network *10.0.0.0/8* and *Zone B* includes the network *10.0.1.0/24*. In this case, a client machine with the *10.0.1.100/32* IP address belongs to both zones from an IP addressing point of view. But *Zone B* is more specific (in CIDR terms), so the client machine belongs to *Zone B*.

#### 5.18.2.1. Zone methods

Method	Description
<i>__init__</i> ( <i>self, name, addrs, hostnames, admin_parent, inbound_services, outbound_services</i> )	Constructor to initialize a Zone instance

Table 5.111. Method summary

**Method `__init__(self, name, addrs, hostnames, admin_parent, inbound_services, outbound_services)`**

This constructor initializes a Zone object.

#### Arguments of `__init__`

<b>addr (complex)</b>
Default: n/a
A string representing an address range interpreted by the domain class (last argument), *or* a list of strings representing multiple address ranges.
<b>admin_parent (string)</b>
Default: n/a
Name of the administrative parent zone. If set, the current zone inherits the lists of permitted inbound and outbound services from its administrative parent zone.
<b>hostnames (complex)</b>
Default: n/a

**hostnames (complex)**

A string representing a domain name, the addresses of its A and AAAA records are placed into the zone hierarchy \*or\* a list of domain names representing multiple domain names

**name (string)**

Default: n/a

Name of the zone.

## 5.19. Module Zorp

This module defines global constants (e.g., *TRUE* and *FALSE*) used by other components, and interface entry points to the core.

## Chapter 6. Core-internal

This chapter provides information about some of the internal Zorp modules.

### 6.1. Module Cache

Caching is used throughout the policy layer to improve performance. This module includes a couple of general caching classes used by various parts of the policy code.

### 6.2. Module Core

This module imports all public interfaces and makes it easy to use those from the user policy file by simply importing all symbols from Zorp.Core.

### 6.3. Module Dispatch

Dispatchers bind to a specific IP address and port of the Zorp firewall and wait for incoming connection requests. For each accepted connection, the Dispatcher creates a new service instance to handle the traffic arriving in the connection.

**Note**

Earlier product versions used different classes to handle TCP and UDP connections (Dispatchers, respectively). These classes have been merged into the Dispatcher module.

For each accepted connection, the Dispatcher creates a new service instance to handle the traffic arriving in the connection. The service started by the dispatcher depends on the type of the dispatcher:

- *Dispatchers* start the same service for every connection.
- *CSZoneDispatchers* start different services based on the zones the client and the destination server belong to.

**Note**

Only one dispatcher can bind to an IP address/port pair.

#### 6.3.1. Zone-based service selection

Dispatchers can start only a predefined service. Use CSZonedDispatchers to start different services for different connections. CSZoneDispatchers assign different services to different client-server zone pairs. Define the zones and the related services in the *services* parameter. The \* wildcard matches all client or server zones.

**Note**

The server zone may be modified by the proxy, the router, the chainer, or the NAT policy used in the service. To select the service, CSZoneDispatcher determines the server zone from the original destination IP address of the incoming client request. Similarly, the client zone is determined from the source IP address of the original client request.

To accept connections from the child zones of the selected client zones, set the *follow\_parent* attribute to *TRUE*. Otherwise, the dispatcher accepts traffic only from the client zones explicitly listed in the *services* attribute of the dispatcher.

### 6.3.2. Classes in the Dispatch module

Class	Description
<i>CSZoneDispatcher</i>	Class encapsulating the Dispatcher which starts a service by the client and server zone.
<i>Dispatcher</i>	Class encapsulating the Dispatcher which starts a service by the client and server zone.

Table 6.1. Classes of the Dispatch module

### 6.3.3. Class CSZoneDispatcher

This class is similar to a simple Dispatcher, but instead of starting a fixed service, it chooses one based on the client and the destined server zone.

It takes a mapping of services indexed by a client and the server zone name, with an exception of the '\*' zone, which matches anything.

NOTE: the server zone might change during proxy and NAT processing, therefore the server zone used here only matches the real destination if those phases leave the server address intact.

**Example 6.1. CSZoneDispatcher example**

The following example defines a CSZoneDispatcher that starts the service called *internet\_HTTP\_DMZ* for connections received on the *192.168.2.1* IP address, but only if the connection comes from the *internet* zone and the destination is in the *DMZ* zone.

```
CSZoneDispatcher(bindto=SockAddrInet('192.168.2.1', 50080), services={"internet",
"DMZ"}:"internet_HTTP_DMZ"}, transparent=TRUE, backlog=255, threaded=FALSE, follow_parent=FALSE)
```

#### 6.3.3.1. Attributes of CSZoneDispatcher

services (unknown)
Default: n/a
services mapping indexed by zone names

### 6.3.3.2. CSZoneDispatcher methods

Method	Description
<code>__init__(self, bindto, services, **kw)</code>	Constructor to initialize a CSZoneDispatcher instance.

Table 6.2. Method summary

#### Method `__init__(self, bindto, services, **kw)`

This constructor initializes a CSZoneDispatcher instance and sets its initial attributes based on arguments.

#### Arguments of `__init__`

<b>bindto (sockaddr)</b>
Default: n/a
An existing <i>socket address</i> containing the IP address and port number where the Dispatcher accepts connections.
<b>follow_parent (boolean)</b>
Default: n/a
Set this parameter to <i>TRUE</i> if the dispatcher handles also the connections coming from the child zones of the selected client zones. Otherwise, the dispatcher accepts traffic only from the explicitly listed client zones.
<b>services (complex)</b>
Default: n/a
Client zone - server zone - service name pairs using the <code>(("client_zone", "server_zone"): "service")</code> format; specifying the service to start when the dispatcher accepts a connection from the given client zone that targets the server zone.

### 6.3.4. Class Dispatcher

This class is the starting point of services. It listens on the given port, and when a connection is accepted it starts a session and the given service.



#### Example 6.2. Dispatcher example

The following example defines a transparent dispatcher that starts the service called *demo\_http\_service* for connections received on the *192.168.2.1* IP address.

```
Dispatcher(bindto=SockAddrInet('192.168.2.1', 50080), service="demo_http_service", transparent=TRUE,
backlog=255, threaded=FALSE)
```

### 6.3.4.1. Attributes of Dispatcher

<b>backlog (integer)</b>
Default: n/a
<i>Applies only to TCP connections.</i> This parameter sets the queue size (maximum number) of TCP connections that are established by the kernel, but not yet accepted by Zorp. This queue stores the connections that successfully performed the three-way TCP handshake with the Zorp host, until the dispatcher sends the <i>Accept</i> package.
<b>bindto (sockaddr)</b>
Default: n/a
An existing <i>socket address</i> containing the IP address and port number where the Dispatcher accepts connections.
<b>protocol (unknown)</b>
Default: n/a
the protocol we were bound to
<b>service (service)</b>
Default: n/a
Name of the service to start.
<b>threaded (boolean)</b>
Default: n/a
Set this parameter to <i>TRUE</i> to start a new thread for every client request. The proxy threads started by the dispatcher will start from the dispatcher's thread instead of the main thread. Incoming connections are accepted faster and optimizes queuing if this option is enabled. This improves user experience, but significantly increases the memory consumption of Zorp. Use it only if a very high number of concurrent connections have to be transferred.

### 6.3.4.2. Dispatcher methods

Method	Description
<code>__init__(self, bindto, service, **kw)</code>	Constructor to initialize a Dispatcher instance.

Table 6.3. Method summary

#### Method `__init__(self, bindto, service, **kw)`

This constructor creates a new Dispatcher instance which can be associated with a *Service*.



## Arguments of `__init__`

<b>bindto (sockaddr)</b>
Default: n/a
An existing <i>socket address</i> containing the IP address and port number where the Dispatcher accepts connections.
<b>service (service)</b>
Default: n/a
Name of the service to start.
<b>transparent (boolean)</b>
Default: n/a
Set this parameter to <i>TRUE</i> if the dispatcher starts a transparent service.

## 6.4. Module Globals

Global variables used by the policy layer.

## 6.5. Module Stream

This module defines the Stream class, encapsulating file descriptors and related functions.

### 6.5.1. Classes in the Stream module

Class	Description
<i>Stream</i>	Class encapsulating the file descriptor and related functions.

Table 6.4. Classes of the Stream module

### 6.5.2. Class Stream

This class encapsulates a full-duplex data tunnel, represented by a UNIX file descriptor. Proxies communicate with its peers through instances of this class. The *client\_stream* and *server\_stream* attributes of the *Session* class contain a Stream instance.

#### 6.5.2.1. Attributes of Stream

<b>bytes_recvd (integer)</b>
Default: n/a
The number of bytes received in the stream.



<b>bytes_sent (integer)</b>
Default: n/a
The number of bytes sent in the stream.

<b>fd (integer)</b>
Default: n/a
The file descriptor associated to the stream.

<b>name (string)</b>
Default: n/a
The name of the stream.

### 6.5.2.2. Stream methods

Method	Description
<code>__init__(self, fd, name)</code>	Constructor to initialize a stream.

Table 6.5. Method summary

#### Method `__init__(self, fd, name)`

This constructor initializes a Stream instance setting its attributes according to arguments.

#### Arguments of `__init__`

<b>fd (integer)</b>
Default: n/a
The file descriptor associated to the stream.

<b>name (string)</b>
Default: n/a
The name of the stream.



## Appendix A. Additional proxy information

### A.1. TELNET appendix

The constants defined for the easier use of TELNET options and suboptions are listed in the table below. Suboptions are listed directly under the option they refer to. All suboptions have the TELNET\_SB prefix. The RFC describing the given option is also shown in the table.

Name	Constant value of option/suboption	Detailed in RFC #
TELNET_BINARY	0	856
TELNET_ECHO	1	857
TELNET_SUPPRESS_GO_AHEAD	3	858
TELNET_STATUS	5	859
TELNET_SB_STATUS_SB_IS	0	
TELNET_SB_STATUS_SB_SEND	1	
TELNET_TIMING_MARK	6	860
TELNET_RCTE	7	726
TELNET_NAOCRD	10	652
TELNET_SB_NAOCRD_DR	0	
TELNET_SB_NAOCRD_DS	1	
TELNET_NAOHTS	11	653
TELNET_SB_NAOHTS_DR	0	
TELNET_SB_NAOHTS_DS	1	
TELNET_NAOHTD	12	654
TELNET_SB_NAOHTD_DR	0	
TELNET_SB_NAOHTD_DS	1	
TELNET_NAOFFD	13	655
TELNET_SB_NAOFFD_DR	0	
TELNET_SB_NAOFFD_DS	1	
TELNET_NAOVTS	14	656
TELNET_SB_NAOVTS_DR	0	
TELNET_SB_NAOVTS_DS	1	



Name	Constant value of option/suboption	Detailed in RFC #
TELNET_NAOVTD	15	657
TELNET_SB_NAOVTD_DR	0	
TELNET_SB_NAOVTD_DS	1	
TELNET_NAOLFD	16	658
TELNET_SB_NAOLFD_DR	0	
TELNET_SB_NAOLFD_DS	1	
TELNET_EXTEND_ASCII	17	698
TELNET_LOGOUT	18	727
TELNET_BM	19	735
TELNET_SB_BM_DEFINE	1	
TELNET_SB_BM_ACCEPT	2	
TELNET_SB_BM_REFUSE	3	
TELNET_SB_BM_LITERAL	4	
TELNET_SB_BM_CANCEL	5	
TELNET_DET	20	1043, 732
TELNET_SB_DET_DEFINE	1	
TELNET_SB_DET_ERASE	2	
TELNET_SB_DET_TRANSMIT	3	
TELNET_SB_DET_FORMAT	4	
TELNET_SB_DET_MOVE_CURSOR	5	
TELNET_SB_DET_SKIP_TO_LINE	6	
TELNET_SB_DET_SKIP_TO_CHAR	7	
TELNET_SB_DET_UP	8	
TELNET_SB_DET_DOWN	9	
TELNET_SB_DET_LEFT	10	
TELNET_SB_DET_RIGHT	11	
TELNET_SB_DET_HOME	12	
TELNET_SB_DET_LINE_INSERT	13	
TELNET_SB_DET_LINE_DELETE	14	
TELNET_SB_DET_CHAR_INSERT	15	



Name	Constant value of option/suboption	Detailed in RFC #
TELNET_SB_DET_CHAR_DELETE	16	
TELNET_SB_DET_READ_CURSOR	17	
TELNET_SB_DET_CURSOR_POSITION	18	
TELNET_SB_DET_REVERSE_TAB	19	
TELNET_SB_DET_TRANSMIT_SCREEN	20	
TELNET_SB_DET_TRANSMIT_UNPROTECTED	21	
TELNET_SB_DET_TRANSMIT_LINE	22	
TELNET_SB_DET_TRANSMIT_FIELD	23	
TELNET_SB_DET_TRANSMIT_REST_SCREEN	24	
TELNET_SB_DET_TRANSMIT_REST_LINE	25	
TELNET_SB_DET_TRANSMIT_REST_FIELD	26	
TELNET_SB_DET_TRANSMIT_MODIFIED	27	
TELNET_SB_DET_DATA_TRANSMIT	28	
TELNET_SB_DET_ERASE_SCREEN	29	
TELNET_SB_DET_ERASE_LINE	30	
TELNET_SB_DET_ERASE_FIELD	31	
TELNET_SB_DET_ERASE_REST_SCREEN	32	
TELNET_SB_DET_ERASE_REST_LINE	33	
TELNET_SB_DET_ERASE_REST_FIELD	34	
TELNET_SB_DET_ERASE_UNPROTECTED	35	
TELNET_SB_DET_FORMAT_DATA	36	
TELNET_SB_DET_REPEAT	37	
TELNET_SB_DET_SUPPRESS_PROTECTION	38	
TELNET_SB_DET_FIELD_SEPARATOR	39	
TELNET_SB_DET_FN	40	
TELNET_SB_DET_ERROR	41	
TELNET_SUPDUP	21	736, 734
TELNET_SUPDUP_OUTPUT	22	749
TELNET_SEND_LOCATION	23	779
TELNET_TERMINAL_TYPE	24	1091



Name	Constant value of option/suboption	Detailed in RFC #
TELNET_SB_TERMINAL_TYPE_IS	0	
TELNET_SB_TERMINAL_TYPE_SEND	1	
TELNET_EOR	25	885
TELNET_TUID	26	927
TELNET_OUTMRK	27	933
TELNET_TTYLOC	28 946	
TELNET_3270_REGIME	29	1041
TELNET_SB_3270_REGIME_IS	0	
TELNET_SB_3270_REGIME_ARE	1	
TELNET_X3_PAD	30	1053
TELNET_SB_X3_PAD_SET	0	
TELNET_SB_X3_PAD_RESPONSE_SET	1	
TELNET_SB_X3_PAD_IS	2	
TELNET_SB_X3_PAD_RESPONSE_IS	3	
TELNET_SB_X3_PAD_SEND	4	
TELNET_NAWS	31	1073
TELNET_TERMINAL_SPEED	32	1079
TELNET_SB_TERMINAL_SPEED_IS	0	
TELNET_SB_TERMINAL_SPEED_SEND	1	
TELNET_TOGGLE_FLOW_CONTROL	33	1372
TELNET_SB_TOGGLE_FLOW_CONTROL_CH	0	
TELNET_SB_TOGGLE_FLOW_CONTROL_CN	1	
TELNET_SB_TOGGLE_FLOW_CONTROL_RSPT_AY	2	
TELNET_SB_TOGGLE_FLOW_CONTROL_RSPT_XN	3	
TELNET_LINEMODE	34	1184
TELNET_SB_LINEMODE_MODE	1	
TELNET_SB_LINEMODE_FORWARDMASK	2	
TELNET_SB_LINEMODE_SLC	3	
TELNET_X_DISPLAY_LOCATION	35	1096
TELNET_SB_X_DISPLAY_LOCATION_IS	0	



Name	Constant value of option/suboption	Detailed in RFC #
TELNET_SB_X_DISPLAY_LOCATION_SEND	1	
TELNET_OLD_ENVIRONMENT	36	1408
TELNET_SB_OLD_ENVIRONMENT_IS	0	
TELNET_SB_OLD_ENVIRONMENT_SEND	1	
TELNET_SB_OLD_ENVIRONMENT_INFO	2	
TELNET_AUTHENTICATION	37	2941
TELNET_SB_AUTHENTICATION_IS	0	
TELNET_SB_AUTHENTICATION_SEND	1	
TELNET_SB_AUTHENTICATION_REPLY	2	
TELNET_SB_AUTHENTICATION_NAME	3	
TELNET_ENCRYPT	38	2946
TELNET_SB_ENCRYPT_IS	0	
TELNET_SB_ENCRYPT_SUPPORT	1	
TELNET_SB_ENCRYPT_REPLY	2	
TELNET_SB_ENCRYPT_START	3	
TELNET_SB_ENCRYPT_END	4	
TELNET_SB_ENCRYPT_REQUEST_START	5	
TELNET_SB_ENCRYPT_REQUEST_END	6	
TELNET_SB_ENCRYPT_ENC_KEYID	7	
TELNET_SB_ENCRYPT_DEC_KEYID	8	
TELNET_ENVIRONMENT	39	1572
TELNET_SB_ENVIRONMENT_IS	0	
TELNET_SB_ENVIRONMENT_SEND	1	
TELNET_SB_ENVIRONMENT_INFO	2	
TELNET_TN3270E	40	1647
TELNET_SB_TN3270E_ASSOCIATE	0	
TELNET_SB_TN3270E_CONNECT	1	
TELNET_SB_TN3270E_DEVICE_TYPE	2	
TELNET_SB_TN3270E_FUNCTIONS	3	
TELNET_SB_TN3270E_IS	4	



Name	Constant value of option/suboption	Detailed in RFC #
TELNET_SB_TN3270E_REASON	5	
TELNET_SB_TN3270E_REJECT	6	
TELNET_SB_TN3270E_REQUEST	7	
TELNET_SB_TN3270E_SEND	8	
TELNET_CHARSET	42	2066
TELNET_SB_CHARSET_REQUEST	1	
TELNET_SB_CHARSET_ACCEPTED	2	
TELNET_SB_CHARSET_REJECTED	3	
TELNET_SB_CHARSET_TTABLE_IS	4	
TELNET_SB_CHARSET_TTABLE_REJECTED	5	
TELNET_SB_CHARSET_TTABLE_ACK	6	
TELNET_SB_CHARSET_TTABLE_NAK	7	
TELNET_COM_PORT	44	2217
TELNET_SB_COM_PORT_CLI_SET_BAUDRATE	1	
TELNET_SB_COM_PORT_CLI_SET_DATASIZE	2	
TELNET_SB_COM_PORT_CLI_SET_PARITY	3	
TELNET_SB_COM_PORT_CLI_SET_STOPSIZE	4	
TELNET_SB_COM_PORT_CLI_SET_CONTROL	5	
TELNET_SB_COM_PORT_CLI_NOTIFY_LINESIZE	6	
TELNET_SB_COM_PORT_CLI_NOTIFY_MODEMSIZE	7	
TELNET_SB_COM_PORT_CLI_HOWCONTROL_SEND	8	
TELNET_SB_COM_PORT_CLI_HOWCONTROL_RESUME	9	
TELNET_SB_COM_PORT_CLI_SET_LINESIZE_MASK	10	
TELNET_SB_COM_PORT_CLI_SET_MODEMSIZE_MASK	11	
TELNET_SB_COM_PORT_CLI_PURGE_DATA	12	
TELNET_SB_COM_PORT_SVR_SET_BAUDRATE	101	
TELNET_SB_COM_PORT_SVR_SET_DATASIZE	102	
TELNET_SB_COM_PORT_SVR_SET_PARITY	103	
TELNET_SB_COM_PORT_SVR_SET_STOPSIZE	104	
TELNET_SB_COM_PORT_SVR_SET_CONTROL	105	



Name	Constant value of option/suboption	Detailed in RFC #
TELNET_SB_COMPFORSVRNOIFYINSAIE	106	
TELNET_SB_COMPFORSVRNOIFYMODEMSAIE	107	
TELNET_SB_COMPFORSRHOWCONFOSSEND	108	
TELNET_SB_COMPFORSRHOWCONFOSRESUME	109	
TELNET_SB_COMPFORSRSETINSAIEMASK	110	
TELNET_SB_COMPFORSRSETMODEMSAIEMASK	111	
TELNET_SB_COMPFORSVRPURGE_DATA	112	
TELNET_KERMIT	47	2840
TELNET_SB_KERMIT_START_SERVER	0	
TELNET_SB_KERMIT_STOP_SERVER	1	
TELNET_SB_KERMIT_REQ_START_SERVER	2	
TELNET_SB_KERMIT_REQ_STOP_SERVER	3	
TELNET_SB_KERMIT_SOP	4	
TELNET_SB_KERMIT_RESP_START_SERVER	8	
TELNET_SB_KERMIT_RESP_STOP_SERVER	9	
TELNET_EXOPL	255	861
TELNET_SUBLIMINAL_MSG	257	1097

Table A.1. TELNET options and suboptions

## Appendix B. Global options of Zorp

Zorp has a number of global options and variables that are used during the initialization of the engine, before any proxies or services are started. These options control the swapping of large data chunks (blobs) to disk, the handling of audit trails, and other miscellaneous parameters. To set these options, complete the following steps:

### B.1. Procedure – Setting global options of Zorp

- Step 1. Select the Zorp component, then select **Variables > New**.
- Step 2. Enter the name of the global option into the **Name** field, and select the type of the option in the **Type** field.
- Step 3. Select **OK**, then **Edit**.
- Step 4. Enter the desired value of the option, then select **OK**.



**Note**

Global options can be also set at the beginning of the `Config.py` file if managing the configuration of Zorp manually.



## blob

blob

### Description

These options control the handling of large data chunks (blobs), determine when they are swapped to disk, and also how much disk space and memory can be used by Zorp.

### Blob options

<code>config.blob.temp_directory</code>	The directory where the blobs are swapped to. Default value: <code>/var/lib/zorp/tmp/</code>
<code>config.blob.hiwat</code>	Zorp tries to store everything in the memory if possible. If the memory usage of Zorp reaches <code>hiwat</code> , it starts to swap the data onto the hard disk, until the memory usage reaches <code>lowat</code> . Default value: <code>128*0x100000</code> (128 MB)
<code>config.blob.lowat</code>	Global options can be also set at the beginning of the <code>Config.py</code> file if managing the configuration of Zorp manually. Lower threshold of data swapping. Default value: <code>96*0x100000</code> (96 MB)
<code>config.blob.max_disk_usage</code>	The maximum amount of hard disk space that Zorp is allowed to use. Default value: <code>1024*0x100000</code> (1 GB)
<code>config.blob.max_mem_usage</code>	The maximum amount of memory that Zorp is allowed to use. Default value: <code>256*0x100000</code> (256 MB)
<code>config.blob.noswap_max</code>	Objects smaller than this value (in bytes) are never swapped to hard disk. Default value: <code>16384</code>

## audit

audit

### Description

These options control the handling of audit trails in Zorp.

### Audit options

<code>config.audit.compress</code>	Enable the compression of audit trail files. The level of compression can be set via the <code>config.audit.compress_level</code> parameter. Default value: <i>TRUE</i>
<code>config.audit.compress_level</code>	The level of compression ranging from 1 (lowest, default) to 9 (highest). Please note that higher compression levels use significantly more CPU, therefore it is usually not recommended to set it to higher than 4. Default value: <i>1</i>
<code>config.audit.encrypt</code>	Encrypt the audit trail files using the key provided in the <code>config.audit.encrypt_certificate</code> parameter. Default value: <i>FALSE</i>
<code>config.audit.encrypt_certificate</code>	The X.509 PEM certificate used to encrypt the audit trail files. Default value: empty. The certificate should be placed in the following format: <pre>-----BEGIN CERTIFICATE----- insert key here -----END CERTIFICATE-----</pre>
<code>config.audit.encrypt_certificate_file</code>	Name and path of the file containing the X.509 PEM certificate used to encrypt the audit trail files. If this parameter is set, it overrides the settings of <code>config.audit.encrypt_certificate</code> . Default value: empty.
<code>config.audit.reopen_size_threshold</code>	The maximum size of a single audit trail file in bytes. Default value: <i>2000000000L</i> (2 GB)
<code>config.audit.per_session</code>	Store each session in its own audit file. Default value: <i>FALSE</i>
<code>config.audit.reopen_time_threshold</code>	The maximum time frame of a single audit file in seconds. Default value: <i>28800</i> (8 hours)
<code>config.audit.rate_limit</code>	Zorp considers it abnormal if the size of an audit trail is increasing faster than this value in byte/second. Default value: <i>2097152</i> (2 MB)
<code>config.audit.rate_notification_interval</code>	Time in seconds before repeating the notification about abnormally growing audit trails. Default value: <i>300</i> (5 minutes)



`config.audit.write_size_max` Maximum size of an audit trail file in bytes. Default value: *52428800* (50 MB)

`config.audit.terminate_on_max_size` If set to *TRUE*, Zorp terminates the connection if the corresponding audit trail file reaches the size limit set in *config.audit.write\_size\_max*. Default value: *FALSE*

## options

options

### Description

These options control various behavior of Zorp.

### Options

`config.options.dscp_prio_mapping` Priority mapping for transferring Differentiated Services Code Point (DSCP, also known as Type of Service or ToS). The low (0), normal (1), high (2), and urgent (3) priorities can be assigned to the DSCP classes. The assigned priority determines the priority of the Zorp thread that handles the connection. The mapping is actually a hash table consisting of the DSCP class ID, a colon (:), the priority of the class (0-3), and a comma (,) except for the last row. For example:

```
config.options.dscp_mapping = { 1: 3,  
                                2: 2,  
                                3: 2,  
                                4: 0 }
```

`config.options.language` The default language used to display user-visible messages, e.g., HTTP error pages. Default value: *en* (English). Other supported languages: *de* (German); *hu* (Hungarian).

`config.options.timeout_server_connect` The timeout (in milliseconds) used when establishing server-side connections. Default value: *30000* (30 sec)

### Cache options

Zorp caches certain data (e.g., to which zone a particular IP address belongs to) to decrease the time required to process a connection. The following parameters determine the size of these caches (the number of decisions stored in the cache). Adjusting these parameters is required only in environments having very complex zone structure and a large number of services. The following log message indicates that a cache is full: *Cache over shift-threshold, shifting*

`config.zone_cache_shift_threshold` Stores IP addresses and the zone they belong to. Default value: *1000*

`config.inbound_service_cache_threshold` Stores service-zone pairs, and if the service is permitted to enter the zone. Default value: *1000*

`config.outbound_service_cache_threshold` Stores service-zone pairs, and if the service is permitted to leave the zone. Default value: *1000*



## Appendix C. Zorp manual pages

## instances.conf

instances.conf — *zorp(8)* instances database

### Description

The `instances.conf` file describes the *zorp(8)* instances to be run on the system. It is processed by *zorpctl(8)* line by line, each line having the structure described below. Empty lines and lines beginning with '#' are comments ignored by *zorpctl*.

### Structure

`instance-name parameters [-- zorpctl-options]`

*instance-name* is the name of the Zorp instance to be started; it is passed to *zorp* with its `--as` parameter. Instance names may consist of the characters `[a-zA-Z0-9_]` and must begin with a letter.

*parameters* are space separated parameters entered into the *zorp* command-line. For details on these command-line parameters see *zorp(8)*.

*zorpctl-options* are space separated parameters control startup specific options. They are processed by *zorpctl* itself. The following *zorpctl* options are available:

<code>--auto-restart</code> or <code>-A</code>	Enable the automatic restart feature of <i>zorpctl</i> . When an instance is in auto-restart mode, it is restarted automatically in case the instance exits.
<code>--no-auto-restart</code> or <code>-a</code>	Disable automatic restart for this instance.
<code>--fd-limit &lt;number&gt;</code> or <code>-f &lt;number&gt;</code>	Set the file descriptor limit to <code>&lt;number&gt;</code> . The file descriptor limit defaults to the number of threads (specified by the <code>--threads</code> parameter of <i>zorp(8)</i> ) multiplied by 4.
<code>--num-of-processes &lt;number&gt;</code> or <code>-P &lt;number&gt;</code>	Run <code>&lt;number&gt;</code> of processes for the instance. <i>zorpctl</i> starts exactly one Zorp process in master mode and <code>&lt;number&gt;</code> of slave Zorp processes. This mode of operation is incompatible with old-style dispatchers, you must use the new rule-based policy with this option.

### Examples

```
zorp_ftp --policy /etc/zorp/policy.py --verbose 5
```

The line above describes a Zorp instance named *zorp\_ftp* using policy file `/etc/zorp/policy.py`, and having verbosity level 5.

```
zorp_intra -v4 -p /etc/zorp/policy.py --threads 500 --no-auto-restart --fd-limit 1024 --process-limit 512
```

This line describes a *zorp* instance named *zorp\_intra* using the policy file `/etc/zorp/policy.py`, verbosity level 4. The maximum number of threads is set to 500, file descriptor limit to 1024, process limit to 512.



## Files

The default location of `instances.conf` is `/etc/zorp/instances.conf`. Defaults for `zorpctl` tunables can be specified in `/etc/zorp/zorpctl`.

## Author

This manual page was written by the BalaSys Documentation Team <[documentation@balasys.hu](mailto:documentation@balasys.hu)>.

## Copyright

Copyright © 1996-2024 Balasys IT Zrt. All rights reserved.

## policy.py

policy.py — *zorp(8)* policy file.

### Description

The `policy.py` file is a Python module containing the zone and service definitions and other policy related settings used by *zorp(8)*. Empty lines and lines beginning with '#' are comments and are ignored.

The `policy.py` file is generated automatically by ZMC, the Zorp Management Console, or it can be edited manually.

**IMPORTANT:** Do not edit manually a file generated by ZMC, because the manual changes will not be retained by ZMC and will be lost when re-generating the file.

### Files

The default location of `policy.py` is `/etc/zorp/policy.py`.

### See Also

For further information on `policy.py` refer to the following sources:

A tutorial on manually editing the `policy.py` file can be found at <http://www.balasy.hu/documentation/>.

Additional information can also be found in the *Zorp Administrator's Guide*, the *Zorp Reference Guide*, and in the various tutorials available at the BalaSys Documentation Page at <http://www.balasy.hu/documentation>.

### Author

This manual page was written by the BalaSys Documentation Team <documentation@balasy.hu>.

### Copyright

Copyright © 1996-2024 Balasys IT Zrt. All rights reserved.



## zorp

zorp — Zorp Firewall Suite

### Synopsis

zorp [options]

### Description

The zorp command is the main entry point for a Zorp instance, and as such it is generally called by *zorpctl(8)* with command line parameters specified in *instances.conf(5)*.

### Options

<code>--version</code> or <code>-V</code>	Display version number and compilation information.
<code>--as &lt;name&gt;</code> or <code>-a &lt;name&gt;</code>	Set instance name to <name>. Instance names may consist of the characters [a-zA-Z0-9_] and must begin with a letter. Log messages of this instance are prefixed with this name.
<code>--no-syslog</code> or <code>-l</code>	Send log messages to the standard output instead of syslog. This option implies foreground mode, overriding the contradicting process options if present.
<code>--log-tags</code> or <code>-T</code>	Prepend log category and log level to each message.
<code>--log-escape</code>	Escape non-printable characters to avoid binary log files. Each character less than 0x20 and greater than 0x7F are escaped in the form <XX>.
<code>--log-spec &lt;spec&gt;</code> or <code>-s &lt;spec&gt;</code>	Set verbosity mask on a per category basis. Each log message has an assigned multi-level category, where levels are separated by a dot. For example, HTTP requests are logged under <i>http.request</i> . <spec> is a comma separated list of log specifications. A single log specification consists of a wildcard matching log category, a colon, and a number specifying the verbosity level of that given category. Categories match from left to right. E.g.: <code>--logspec 'http.*:5,core:3'</code> . The last matching entry will be used as the verbosity of the given category. If no match is found the default verbosity specified with <code>--verbose</code> is used.
<code>--threads &lt;num&gt;</code> or <code>-t &lt;num&gt;</code>	Set the maximum number of threads that can be used in parallel by this Zorp instance.
<code>--idle-threads &lt;num&gt;</code> or <code>-I</code>	Set the maximum number of idle threads; this option has effect only if threadpools are enabled (see the option <code>--threadpools</code> ).
<code>--threadpools</code> or <code>-0</code>	Enable the use of threadpools, which means that threads associated with sessions are not automatically freed, only if the maximum number of idle threads is exceeded.
<code>--user &lt;user&gt;</code> or <code>-u &lt;user&gt;</code>	Switch to the supplied user after starting up.

<code>--group &lt;group&gt; or -g &lt;group&gt;</code>	Switch to the supplied group after starting up.
<code>--chroot &lt;dir&gt; or -R &lt;dir&gt;</code>	Change root to the specified directory before reading the configuration file. The directory must be set up accordingly.
<code>--caps &lt;caps&gt; or -C &lt;caps&gt;</code>	Switch to the supplied set of capabilities after starting up. This should contain the required capabilities in the permitted set. For the syntax of capability description see the man page of <code>cap_from_text(3)</code> .
<code>--no-caps or -N</code>	Do not change capabilities at all.
<code>--crypto-engine &lt;engine&gt; or -E &lt;engine&gt;</code>	Set the OpenSSL crypto engine to be used for hardware accelerated crypto support.

## Files

`/etc/zorp/`

`/etc/zorp/policy.py`

`/etc/zorp/instances.conf`

## Author

This manual page was written by the BalaSys Documentation Team <[documentation@balasys.hu](mailto:documentation@balasys.hu)>.

## Copyright

Copyright © 1996-2024 Balasys IT Zrt. All rights reserved.

## zorpctl

zorpctl — Start and stop zorp instances.

### Synopsis

```
zorpctl command [options [instances/@instance-list-file]]
```

### Description

zorpctl starts and stops *zorp(8)* instances based on the contents of the *instances.conf(5)* file. Multiple instance names can be specified in the command-line or in a file to start or stop several instances. If an error occurs while stopping or starting an instance, an exclamation mark is appended to the instance name as zorpctl processes the request, and a summary is printed when the program exits. If no instance is specified, the command is executed on all instances. The instances to be controlled can be specified in a file instead of listing them in the command line, e.g.: `zorpctl command options instances.txt`. The `instances.txt` should contain every instance name in a new line.

### Commands

start	Starts the specified Zorp instance(s).
force-start	Starts the specified Zorp instance(s) even if they are disabled.
stop	Stops the specified Zorp instance(s).
force-stop	Forces the specified Zorp instance(s) to stop using the KILL signal.
restart	Restart the specified Zorp instance(s).
force-restart	Forces the specified Zorp instance(s) to restart by stopping them using the KILL signal.
reload	Reload the specified Zorp instance(s).
status	Display the status of the specified Zorp instance(s). --verbose or -v                      Display detailed status information.
gui-status	Display the status of the specified Zorp instance(s) in an internal format easily parsable by ZMC. NOTE: This command is mainly used internally within Zorp, and the structure of its output may change.
version	Display version information on Zorp.
inclog	Raise the verbosity (log) level of the specified Zorp instance(s) by one.
declog	Decrease the verbosity (log) level of the specified Zorp instance(s) by one.
log	Change various log related settings in the specified Zorp instance(s) using the following options:



	--vinc or -i	Increase verbosity level by one.
	--vdec or -d	Decrease verbosity level by one.
	--vset <verbosity> or -s <verbosity>	Set verbosity level to <verbosity>.
	--log-spec <spec> or -S <spec>	Set verbosity mask on a per category basis. The format of this value is described in <i>zorp(8)</i> .
	--help or -h	Display this help screen on the options of the log command.
szig		Display internal information from the specified Zorp instance(s). The information to be displayed can be specified with the following options:
	--walk or -w	Walk the specified tree.
	--root [node] or -r [node]	Set the root node of the walk operation to [node].
	--help or -h	Display a brief help on the options of the szig command.
help		Display a brief help message.

## Examples

```
zorpctl start zorp_ftp
```

The command above starts the zorp instance named *zorp-ftp* with parameters described in the `instances.conf` file.

## Files

The default location for `instances.conf` is `/etc/zorp/instances.conf`.

## Author

This manual page was written by the BalaSys Documentation Team <documentation@balasys.hu>.

## Copyright

Copyright © 1996-2024 Balasys IT Zrt. All rights reserved.

## zorpcctl.conf

zorpcctl.conf — *zorpcctl(8)* configuration file.

### Description

The *zorpcctl.conf* file describes various global options influencing the behavior of *zorpcctl(8)*. *zorpcctl(8)* processes the file line by line, each line having the structure described below. Empty lines and lines beginning with '#' are comments and are ignored.

### Structure

```
variable name = variable value
```

Each non-empty line specifies a variable name and its value separated by the equal sign ('='). The following variables are available:

AUTO_RESTART	Enable the automatic restart feature of <i>zorpcctl</i> . Instances in auto-restart mode are restarted automatically when they exit. Default value: 1 (TRUE).
STOP_CHECK_TIMEOUT	The number of seconds to wait for a stopping Zorp instance. Default value: 3.
START_CHECK_TIMEOUT	In <i>auto-restart</i> mode there is no real way to detect whether Zorp failed to load or not. <i>Zorpcctl</i> waits <i>START_CHECK_TIMEOUT</i> seconds and assumes that Zorp loaded successfully if it did not exit within this interval. Default value: 5 seconds.
START_WAIT_TIMEOUT	In <i>no-auto-restart</i> mode the successful loading of a Zorp instance can be verified by instructing Zorp to daemonize itself and waiting for the parent to exit. This parameter specifies the number of seconds to wait for Zorp to daemonize itself. Default value: 60 seconds.
ZORP_APPEND_ARGS	Zorp-specific arguments to be appended to the command line of each Zorp instance. Also recognised as <i>APPEND_ARGS</i> (deprecated). Default value: "".
ZORPCCTL_APPEND_ARGS	<i>Zorpcctl</i> -specific arguments to be appended to the command line of each instance. Default value: "".
CHECK_PERMS	Specifies whether to check the permissions of the Zorp configuration directory. If set, Zorp refuses to run if the <i>/etc/zorp</i> directory can be written by user other than <i>zorp</i> . Default value: 1 (TRUE).
CONFIG_DIR	The path to the Zorp configuration directory to check if <i>CHECK_PERMS</i> is enabled. NOTE: it does not change the Zorp policy file argument, this parameter is only used by the permission validating code. Default value: <code>\${prefix}/etc/zorp</code> .

CONFIG\_DIR\_OWNER,  
CONFIG\_DIR\_GROUP,  
CONFIG\_DIR\_MODE

The owner/group/permissions values considered valid for the configuration directory. `zorpctl` fails if the actual owner/group/permissions values conflict the ones set here. Default values: *root.zorp, 0750* .

PIDFILE\_DIR

The path to the Zorp pid file directory. The directory is created automatically prior to starting Zorp if it does not already exist. It is created if it does not exist, before NOTE: No `--pidfile` argument is passed to Zorp, only existence of the directory is verified. Default value: */var/run/zorp*.

PIDFILE\_DIR\_OWNER,  
PIDFILE\_DIR\_GROUP,  
PIDFILE\_DIR\_MODE

The owner/group/permission values the pidfile directory is created with if it does not exist. Default values: *root.root, 0700*.

## Files

The default location for `zorpctl.conf` is `/etc/zorp/zorpctl.conf`.

## Author

This manual page was written by the BalaSys Documentation Team <[documentation@balasys.hu](mailto:documentation@balasys.hu)>.

## Copyright

Copyright © 1996-2024 Balasys IT Zrt. All rights reserved.



## Appendix D. Zorp GPL End-User License Agreement

(c) Balasys IT Security Ltd.

### D.1. 1. SUBJECT OF THE LICENSE CONTRACT

1.1 This License Contract is entered into by and between Balasys and Licensee and sets out the terms and conditions under which Licensee and/or Licensee's Authorized Subsidiaries may use the Zorp GPL under this License Contract.

### D.2. 2. DEFINITIONS

In this License Contract, the following words shall have the following meanings:

#### 2.1 Balasys

Company name: Balasys IT Ltd.

Registered office: H-1117 Budapest, Alíz Str. 4.

Company registration number: 01-09-687127

Tax number: HU11996468-2-43

#### 2.2. Words and expressions

##### Annexed Software

Any third party software that is not a Balasys Product contained in the install media of the Balasys Product.

##### Authorized Subsidiary

Any subsidiary organization: (i) in which Licensee possesses more than fifty percent (50%) of the voting power and (ii) which is located within the Territory.

##### Balasys Product

Any software, hardware or service licensed, sold, or provided by Balasys including any installation, education, support and warranty services, with the exception of the Annexed Software.

##### License Contract

The present Zorp GPL License Contract.

##### Product Documentation



Any documentation referring to the Zorp GPL or any module thereof, with special regard to the reference guide, the administration guide, the product description, the installation guide, user guides and manuals.

#### Protected Hosts

Host computers located in the zones protected by Zorp GPL, that means any computer bounded to network and capable to establish IP connections through the firewall.

#### Protected Objects

The entire Zorp GPL including all of its modules, all the related Product Documentation; the source code, the structure of the databases, all registered information reflecting the structure of the Zorp GPL and all the adaptation and copies of the Protected Objects that presently exist or that are to be developed in the future, or any product falling under the copyright of Balasys.

#### Zorp GPL

Application software Balasys Product designed for securing computer networks as defined by the Product Description.

#### Warranty Period

The period of twelve (12) months from the date of delivery of the Zorp GPL to Licensee.

#### Territory

The countries or areas specified above in respect of which Licensee shall be entitled to install and/or use Zorp GPL.

#### Take Over Protocol

The document signed by the parties which contains

- a) identification data of Licensee;
- b) ordered options of Zorp GPL, number of Protected Hosts and designation of licensed modules thereof;
- c) designation of the Territory;
- d) declaration of the parties on accepting the terms and conditions of this License Contract; and
- e) declaration of Licensee that is in receipt of the install media.

### **D.3. 3. LICENSE GRANTS AND RESTRICTIONS**

3.1. For the Zorp GPL licensed under this License Contract, Balasys grants to Licensee a non-exclusive, non-transferable, perpetual license to use such Balasys Product under the terms and conditions of this License Contract and the applicable Take Over Protocol.

3.2. Licensee shall use the Zorp GPL in the in the configuration and in the quantities specified in the Take Over Protocol within the Territory.





3.3. On the install media all modules of the Zorp GPL will be presented, however, Licensee shall not be entitled to use any module which was not licensed to it. Access rights to modules and IP connections are controlled by an "electronic key" accompanying the Zorp GPL.

3.4. Licensee shall be entitled to make one back-up copy of the install media containing the Zorp GPL.

3.5. Licensee shall make available the Protected Objects at its disposal solely to its own employees and those of the Authorized Subsidiaries.

3.6. Licensee shall take all reasonable steps to protect Balasys's rights with respect to the Protected Objects with special regard and care to protecting it from any unauthorized access.

3.7. Licensee shall, in 5 working days, properly answer the queries of Balasys referring to the actual usage conditions of the

Zorp GPL, that may differ or allegedly differs from the license conditions.

3.8. Licensee shall not modify the Zorp GPL in any way, with special regard to the functions inspecting the usage of the software. Licensee shall install the code permitting the usage of the Zorp GPL according to the provisions defined for it by Balasys. Licensee may not modify or cancel such codes. Configuration settings of the Zorp GPL in accordance with the possibilities offered by the system shall not be construed as modification of the software.

3.9. Licensee shall only be entitled to analyze the structure of the Balasys Products (decompilation or reverse-engineering) if concurrent operation with a software developed by a third party is necessary, and upon request to supply the information required for concurrent operation Balasys does not provide such information within 60 days from the receipt of such a request. These user actions are limited to parts of the Balasys Product which are necessary for concurrent operation.

3.10. Any information obtained as a result of applying the previous Section

(i) cannot be used for purposes other than concurrent operation with the Balasys Product;

(ii) cannot be disclosed to third parties unless it is necessary for concurrent operation with the Balasys Product;

(iii) cannot be used for the development, production or distribution of a different software which is similar to the BalaSys Product

in its form of expression, or for any other act violating copyright.

3.11. For any Annexed Software contained by the same install media as the Balasys Product, the terms and conditions defined by its copyright owner shall be properly applied. Balasys does not grant any license rights to any Annexed Software.

3.12. Any usage of the Zorp GPL exceeding the limits and restrictions defined in this License Contract shall qualify as material breach of the License Contract.

3.13. The Number of Protected Hosts shall not exceed the amount defined in the Take Over Protocol.

3.14. Licensee shall have the right to obtain and use content updates only if Licensee concludes a maintenance contract that includes such content updates, or if Licensee has otherwise separately acquired the right to obtain



and use such content updates. This License Contract does not otherwise permit Licensee to obtain and use content updates.

#### **D.4. 4. SUBSIDIARIES**

4.1 Authorized Subsidiaries may also utilize the services of the Zorp GPL under the terms and conditions of this License Contract. Any Authorized Subsidiary utilising any service of the Zorp GPL will be deemed to have accepted the terms and conditions of this License Contract.

#### **D.5. 5. INTELLECTUAL PROPERTY RIGHTS**

5.1. Licensee agrees that Balasys owns all rights, titles, and interests related to the Zorp GPL and all of Balasys's patents, trademarks, trade names, inventions, copyrights, know-how, and trade secrets relating to the design, manufacture, operation or service of the Balasys Products.

5.2. The use by Licensee of any of these intellectual property rights is authorized only for the purposes set forth herein, and upon termination of this License Contract for any reason, such authorization shall cease.

5.3. The Balasys Products are licensed only for internal business purposes in every case, under the condition that such license does not convey any license, expressly or by implication, to manufacture, duplicate or otherwise copy or reproduce any of the Balasys Products.

No other rights than expressly stated herein are granted to Licensee.

5.4. Licensee will take appropriate steps with its Authorized Subsidiaries, as Balasys may request, to inform them of and assure compliance with the restrictions contained in the License Contract.

#### **D.6. 6. TRADE MARKS**

6.1. Balasys hereby grants to Licensee the non-exclusive right to use the trade marks of the Balasys Products in the Territory in accordance with the terms and for the duration of this License Contract.

6.2. Balasys makes no representation or warranty as to the validity or enforceability of the trade marks, nor as to whether these infringe any intellectual property rights of third parties in the Territory.

#### **D.7. 7. NEGLIGENT INFRINGEMENT**

7.1. In case of negligent infringement of Balasys's rights with respect to the Zorp GPL, committed by violating the restrictions and limitations defined by this License Contract, Licensee shall pay liquidated damages to Balasys. The amount of the liquidated damages shall be twice as much as the price of the Balasys Product concerned, on Balasys's current Price List.

#### **D.8. 8. INTELLECTUAL PROPERTY INDEMNIFICATION**

8.1. Balasys shall pay all damages, costs and reasonable attorney's fees awarded against Licensee in connection with any claim brought against Licensee to the extent that such claim is based on a claim that Licensee's authorized use of the Balasys Product infringes a patent, copyright, trademark or trade secret. Licensee shall notify Balasys in writing of any such claim as soon as Licensee learns of it and shall cooperate fully with Balasys



in connection with the defense of that claim. Balasys shall have sole control of that defense (including without limitation the right to settle the claim).

8.2. If Licensee is prohibited from using any Balasys Product due to an infringement claim, or if Balasys believes that any Balasys Product is likely to become the subject of an infringement claim, Balasys shall at its sole option, either: (i) obtain the right for Licensee to continue to use such Balasys Product, (ii) replace or modify the Balasys Product so as to make such Balasys Product non-infringing and substantially comparable in functionality or (iii) refund to Licensee the amount paid for such infringing Balasys Product and provide a pro-rated refund of any unused, prepaid maintenance fees paid by Licensee, in exchange for Licensee's return of such Balasys Product to Balasys.

8.3. Notwithstanding the above, Balasys will have no liability for any infringement claim to the extent that it is based upon:

- (i) modification of the Balasys Product other than by Balasys,
- (ii) use of the Balasys Product in combination with any product not specifically authorized by Balasys to be combined with the Balasys Product or
- (iii) use of the Balasys Product in an unauthorized manner for which it was not designed.

#### **D.9. 9. LICENSE FEE**

9.1. The number of the Protected Hosts (including the server as one host), the configuration and the modules licensed shall serve as the calculation base of the license fee.

9.2. Licensee acknowledges that payment of the license fees is a condition of lawful usage.

9.3. License fees do not contain any installation or post charges.

#### **D.10. 10. WARRANTIES**

10.1. Balasys warrants that during the Warranty Period, the optical media upon which the Balasys Product is recorded will not be defective under normal use. Balasys will replace any defective media returned to it, accompanied by a dated proof of purchase, within the Warranty Period at no charge to Licensee. Upon receipt of the allegedly defective Balasys Product, Balasys will at its option, deliver a replacement Balasys Product or Balasys's current equivalent to Licensee at no additional cost. Balasys will bear the delivery charges to Licensee for the replacement Product.

10.2. In case of installation by Balasys, Balasys warrants that during the Warranty Period, the Zorp GPL, under normal use in the operating environment defined by Balasys, and without unauthorized modification, will perform in substantial compliance with the Product Documentation accompanying the Balasys Product, when used on that hardware for which it was installed, in compliance with the provisions of the user manuals and the recommendations of Balasys. The date of the notification sent to Balasys shall qualify as the date of the failure. Licensee shall do its best to mitigate the consequences of that failure. If, during the Warranty Period, the Balasys Product fails to comply with this warranty, and such failure is reported by Licensee to Balasys within the Warranty Period, Balasys's sole obligation and liability for breach of this warranty is, at Balasys's sole option, either:

- (i) to correct such failure,



(ii) to replace the defective Balasys Product or

(iii) to refund the license fees paid by Licensee for the applicable Balasys Product.

#### **D.11. 11. DISCLAIMER OF WARRANTIES**

11.1. EXCEPT AS SET OUT IN THIS LICENSE CONTRACT, BALASYS MAKES NO WARRANTIES OF ANY KIND WITH RESPECT TO THE Zorp GPL. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, BALASYS EXCLUDES ANY OTHER WARRANTIES, INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF SATISFACTORY QUALITY, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS.

#### **D.12. 12. LIMITATION OF LIABILITY**

12.1. SOME STATES AND COUNTRIES, INCLUDING MEMBER COUNTRIES OF THE EUROPEAN UNION, DO NOT ALLOW THE LIMITATION OR EXCLUSION OF LIABILITY FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES AND, THEREFORE, THE FOLLOWING LIMITATION OR EXCLUSION MAY NOT APPLY TO THIS LICENSE CONTRACT IN THOSE STATES AND COUNTRIES. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW AND REGARDLESS OF WHETHER ANY REMEDY SET OUT IN THIS LICENSE CONTRACT FAILS OF ITS ESSENTIAL PURPOSE, IN NO EVENT SHALL BALASYS BE LIABLE TO LICENSEE FOR ANY SPECIAL, CONSEQUENTIAL, INDIRECT OR SIMILAR DAMAGES OR LOST PROFITS OR LOST DATA ARISING OUT OF THE USE OR INABILITY TO USE THE Zorp GPL EVEN IF BALASYS HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

12.2. IN NO CASE SHALL BALASYS'S TOTAL LIABILITY UNDER THIS LICENSE CONTRACT EXCEED THE FEES PAID BY LICENSEE FOR THE Zorp GPL LICENSED UNDER THIS LICENSE CONTRACT.

#### **D.13. 13. DURATION AND TERMINATION**

13.1. This License Contract shall come into effect on the date of signature of the Take Over Protocol by the duly authorized

representatives of the parties.

13.2. Licensee may terminate the License Contract at any time by written notice sent to Balasys and by simultaneously destroying all copies of the Zorp GPL licensed under this License Contract.

13.3. Balasys may terminate this License Contract with immediate effect by written notice to Licensee, if Licensee is in material or persistent breach of the License Contract and either that breach is incapable of remedy or Licensee shall have failed to remedy that breach within 30 days after receiving written notice requiring it to remedy that breach.

#### **D.14. 14. AMENDMENTS**

14.1. Save as expressly provided in this License Contract, no amendment or variation of this License Contract shall be effective unless in writing and signed by a duly authorised representative of the parties to it.

**D.15. 15. WAIVER**

15.1. The failure of a party to exercise or enforce any right under this License Contract shall not be deemed to be a waiver of that right nor operate to bar the exercise or enforcement of it at any time or times thereafter.

**D.16. 16. SEVERABILITY**

16.1. If any part of this License Contract becomes invalid, illegal or unenforceable, the parties shall in such an event negotiate in good faith in order to agree on the terms of a mutually satisfactory provision to be substituted for the invalid, illegal or unenforceable

provision which as nearly as possible validly gives effect to their intentions as expressed in this License Contract.

**D.17. 17. NOTICES**

17.1. Any notice required to be given pursuant to this License Contract shall be in writing and shall be given by delivering the notice by hand, or by sending the same by prepaid first class post (airmail if to an address outside the country of posting) to the address of the relevant party set out in this License Contract or such other address as either party notifies to the other from time to time. Any notice given according to the above procedure shall be deemed to have been given at the time of delivery (if delivered by hand) and when received (if sent by post).

**D.18. 18. MISCELLANEOUS**

18.1. Headings are for convenience only and shall be ignored in interpreting this License Contract.

18.2. This License Contract and the rights granted in this License Contract may not be assigned, sublicensed or otherwise transferred in whole or in part by Licensee without Balasys's prior written consent. This consent shall not be unreasonably withheld or delayed.

18.3. An independent third party auditor, reasonably acceptable to Balasys and Licensee, may upon reasonable notice to Licensee and during normal business hours, but not more often than once each year, inspect Licensee's relevant records in order to confirm that usage of the Zorp GPL complies with the terms and conditions of this License Contract. Balasys shall bear the costs of such audit. All audits shall be subject to the reasonable safety and security policies and procedures of Licensee.

18.4. This License Contract constitutes the entire agreement between the parties with regard to the subject matter hereof. Any modification of this License Contract must be in writing and signed by both parties.

## Appendix E. Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd) License

THE WORK (AS DEFINED BELOW) IS PROVIDED UNDER THE TERMS OF THIS CREATIVE COMMONS PUBLIC LICENSE ("CCPL" OR "LICENSE"). THE WORK IS PROTECTED BY COPYRIGHT AND/OR OTHER APPLICABLE LAW. ANY USE OF THE WORK OTHER THAN AS AUTHORIZED UNDER THIS LICENSE OR COPYRIGHT LAW IS PROHIBITED. BY EXERCISING ANY RIGHTS TO THE WORK PROVIDED HERE, YOU ACCEPT AND AGREE TO BE BOUND BY THE TERMS OF THIS LICENSE. TO THE EXTENT THIS LICENSE MAY BE CONSIDERED TO BE A CONTRACT, THE LICENSOR GRANTS YOU THE RIGHTS CONTAINED HERE IN CONSIDERATION OF YOUR ACCEPTANCE OF SUCH TERMS AND CONDITIONS.

### 1. Definitions

- a. "Adaptation" means a work based upon the Work, or upon the Work and other pre-existing works, such as a translation, adaptation, derivative work, arrangement of music or other alterations of a literary or artistic work, or phonogram or performance and includes cinematographic adaptations or any other form in which the Work may be recast, transformed, or adapted including in any form recognizably derived from the original, except that a work that constitutes a Collection will not be considered an Adaptation for the purpose of this License. For the avoidance of doubt, where the Work is a musical work, performance or phonogram, the synchronization of the Work in timed-relation with a moving image ("synching") will be considered an Adaptation for the purpose of this License.
- b. "Collection" means a collection of literary or artistic works, such as encyclopedias and anthologies, or performances, phonograms or broadcasts, or other works or subject matter other than works listed in Section 1(f) below, which, by reason of the selection and arrangement of their contents, constitute intellectual creations, in which the Work is included in its entirety in unmodified form along with one or more other contributions, each constituting separate and independent works in themselves, which together are assembled into a collective whole. A work that constitutes a Collection will not be considered an Adaptation (as defined above) for the purposes of this License.
- c. "Distribute" means to make available to the public the original and copies of the Work through sale or other transfer of ownership.
- d. "Licensor" means the individual, individuals, entity or entities that offer(s) the Work under the terms of this License.
- e. "Original Author" means, in the case of a literary or artistic work, the individual, individuals, entity or entities who created the Work or if no individual or entity can be identified, the publisher; and in addition (i) in the case of a performance the actors, singers, musicians, dancers, and other persons who act, sing, deliver, declaim, play in, interpret or otherwise perform literary or artistic works or expressions of folklore; (ii) in the case of a phonogram the producer being the person or legal entity who first fixes the sounds of a performance or other sounds; and, (iii) in the case of broadcasts, the organization that transmits the broadcast.

- f. "Work" means the literary and/or artistic work offered under the terms of this License including without limitation any production in the literary, scientific and artistic domain, whatever may be the mode or form of its expression including digital form, such as a book, pamphlet and other writing; a lecture, address, sermon or other work of the same nature; a dramatic or dramatico-musical work; a choreographic work or entertainment in dumb show; a musical composition with or without words; a cinematographic work to which are assimilated works expressed by a process analogous to cinematography; a work of drawing, painting, architecture, sculpture, engraving or lithography; a photographic work to which are assimilated works expressed by a process analogous to photography; a work of applied art; an illustration, map, plan, sketch or three-dimensional work relative to geography, topography, architecture or science; a performance; a broadcast; a phonogram; a compilation of data to the extent it is protected as a copyrightable work; or a work performed by a variety or circus performer to the extent it is not otherwise considered a literary or artistic work.
  - g. "You" means an individual or entity exercising rights under this License who has not previously violated the terms of this License with respect to the Work, or who has received express permission from the Licensor to exercise rights under this License despite a previous violation.
  - h. "Publicly Perform" means to perform public recitations of the Work and to communicate to the public those public recitations, by any means or process, including by wire or wireless means or public digital performances; to make available to the public Works in such a way that members of the public may access these Works from a place and at a place individually chosen by them; to perform the Work to the public by any means or process and the communication to the public of the performances of the Work, including by public digital performance; to broadcast and rebroadcast the Work by any means including signs, sounds or images.
  - i. "Reproduce" means to make copies of the Work by any means including without limitation by sound or visual recordings and the right of fixation and reproducing fixations of the Work, including storage of a protected performance or phonogram in digital form or other electronic medium.
2. *Fair Dealing Rights.* Nothing in this License is intended to reduce, limit, or restrict any uses free from copyright or rights arising from limitations or exceptions that are provided for in connection with the copyright protection under copyright law or other applicable laws.
3. *License Grant.* Subject to the terms and conditions of this License, Licensor hereby grants You a worldwide, royalty-free, non-exclusive, perpetual (for the duration of the applicable copyright) license to exercise the rights in the Work as stated below:
- a. to Reproduce the Work, to incorporate the Work into one or more Collections, and to Reproduce the Work as incorporated in the Collections; and,
  - b. to Distribute and Publicly Perform the Work including as incorporated in Collections.
- The above rights may be exercised in all media and formats whether now known or hereafter devised. The above rights include the right to make such modifications as are technically necessary to exercise the rights in other media and formats, but otherwise you have no rights to make Adaptations. Subject to 8(f), all rights not expressly granted by Licensor are hereby reserved, including but not limited to the rights set forth in Section 4(d).
4. *Restrictions.* The license granted in Section 3 above is expressly made subject to and limited by the following restrictions:

- a. You may Distribute or Publicly Perform the Work only under the terms of this License. You must include a copy of, or the Uniform Resource Identifier (URI) for, this License with every copy of the Work You Distribute or Publicly Perform. You may not offer or impose any terms on the Work that restrict the terms of this License or the ability of the recipient of the Work to exercise the rights granted to that recipient under the terms of the License. You may not sublicense the Work. You must keep intact all notices that refer to this License and to the disclaimer of warranties with every copy of the Work You Distribute or Publicly Perform. When You Distribute or Publicly Perform the Work, You may not impose any effective technological measures on the Work that restrict the ability of a recipient of the Work from You to exercise the rights granted to that recipient under the terms of the License. This Section 4(a) applies to the Work as incorporated in a Collection, but this does not require the Collection apart from the Work itself to be made subject to the terms of this License. If You create a Collection, upon notice from any Licensor You must, to the extent practicable, remove from the Collection any credit as required by Section 4(c), as requested.
- b. You may not exercise any of the rights granted to You in Section 3 above in any manner that is primarily intended for or directed toward commercial advantage or private monetary compensation. The exchange of the Work for other copyrighted works by means of digital file-sharing or otherwise shall not be considered to be intended for or directed toward commercial advantage or private monetary compensation, provided there is no payment of any monetary compensation in connection with the exchange of copyrighted works.
- c. If You Distribute, or Publicly Perform the Work or Collections, You must, unless a request has been made pursuant to Section 4(a), keep intact all copyright notices for the Work and provide, reasonable to the medium or means You are utilizing: (i) the name of the Original Author (or pseudonym, if applicable) if supplied, and/or if the Original Author and/or Licensor designate another party or parties (for example a sponsor institute, publishing entity, journal) for attribution ("Attribution Parties") in Licensor's copyright notice, terms of service or by other reasonable means, the name of such party or parties; (ii) the title of the Work if supplied; (iii) to the extent reasonably practicable, the URI, if any, that Licensor specifies to be associated with the Work, unless such URI does not refer to the copyright notice or licensing information for the Work. The credit required by this Section 4(c) may be implemented in any reasonable manner; provided, however, that in the case of a Collection, at a minimum such credit will appear, if a credit for all contributing authors of Collection appears, then as part of these credits and in a manner at least as prominent as the credits for the other contributing authors. For the avoidance of doubt, You may only use the credit required by this Section for the purpose of attribution in the manner set out above and, by exercising Your rights under this License, You may not implicitly or explicitly assert or imply any connection with, sponsorship or endorsement by the Original Author, Licensor and/or Attribution Parties, as appropriate, of You or Your use of the Work, without the separate, express prior written permission of the Original Author, Licensor and/or Attribution Parties.
- d. For the avoidance of doubt:
  - i. Non-waivable Compulsory License Schemes. In those jurisdictions in which the right to collect royalties through any statutory or compulsory licensing scheme cannot be waived, the Licensor reserves the exclusive right to collect such royalties for any exercise by You of the rights granted under this License;
  - ii. Waivable Compulsory License Schemes. In those jurisdictions in which the right to collect royalties through any statutory or compulsory licensing scheme can be waived, the Licensor reserves the exclusive right to collect such royalties for any exercise by You of the rights



granted under this License if Your exercise of such rights is for a purpose or use which is otherwise than noncommercial as permitted under Section 4(b) and otherwise waives the right to collect royalties through any statutory or compulsory licensing scheme; and,

- iii. Voluntary License Schemes. The Licensor reserves the right to collect royalties, whether individually or, in the event that the Licensor is a member of a collecting society that administers voluntary licensing schemes, via that society, from any exercise by You of the rights granted under this License that is for a purpose or use which is otherwise than noncommercial as permitted under Section 4(b).
  - e. Except as otherwise agreed in writing by the Licensor or as may be otherwise permitted by applicable law, if You Reproduce, Distribute or Publicly Perform the Work either by itself or as part of any Collections, You must not distort, mutilate, modify or take other derogatory action in relation to the Work which would be prejudicial to the Original Author's honor or reputation.
5. *Representations, Warranties and Disclaimer* UNLESS OTHERWISE MUTUALLY AGREED BY THE PARTIES IN WRITING, LICENSOR OFFERS THE WORK AS-IS AND MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND CONCERNING THE WORK, EXPRESS, IMPLIED, STATUTORY OR OTHERWISE, INCLUDING, WITHOUT LIMITATION, WARRANTIES OF TITLE, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT, OR THE ABSENCE OF LATENT OR OTHER DEFECTS, ACCURACY, OR THE PRESENCE OF ABSENCE OF ERRORS, WHETHER OR NOT DISCOVERABLE. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES, SO SUCH EXCLUSION MAY NOT APPLY TO YOU.
6. *Limitation on Liability*. EXCEPT TO THE EXTENT REQUIRED BY APPLICABLE LAW, IN NO EVENT WILL LICENSOR BE LIABLE TO YOU ON ANY LEGAL THEORY FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR EXEMPLARY DAMAGES ARISING OUT OF THIS LICENSE OR THE USE OF THE WORK, EVEN IF LICENSOR HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.
7. *Termination*
- a. This License and the rights granted hereunder will terminate automatically upon any breach by You of the terms of this License. Individuals or entities who have received Collections from You under this License, however, will not have their licenses terminated provided such individuals or entities remain in full compliance with those licenses. Sections 1, 2, 5, 6, 7, and 8 will survive any termination of this License.
  - b. Subject to the above terms and conditions, the license granted here is perpetual (for the duration of the applicable copyright in the Work). Notwithstanding the above, Licensor reserves the right to release the Work under different license terms or to stop distributing the Work at any time; provided, however that any such election will not serve to withdraw this License (or any other license that has been, or is required to be, granted under the terms of this License), and this License will continue in full force and effect unless terminated as stated above.
8. *Miscellaneous*
- a. Each time You Distribute or Publicly Perform the Work or a Collection, the Licensor offers to the recipient a license to the Work on the same terms and conditions as the license granted to You under this License.
  - b. If any provision of this License is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this License, and without further

action by the parties to this agreement, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

- c. No term or provision of this License shall be deemed waived and no breach consented to unless such waiver or consent shall be in writing and signed by the party to be charged with such waiver or consent.
- d. This License constitutes the entire agreement between the parties with respect to the Work licensed here. There are no understandings, agreements or representations with respect to the Work not specified here. Licensor shall not be bound by any additional provisions that may appear in any communication from You. This License may not be modified without the mutual written agreement of the Licensor and You.
- e. The rights granted under, and the subject matter referenced, in this License were drafted utilizing the terminology of the Berne Convention for the Protection of Literary and Artistic Works (as amended on September 28, 1979), the Rome Convention of 1961, the WIPO Copyright Treaty of 1996, the WIPO Performances and Phonograms Treaty of 1996 and the Universal Copyright Convention (as revised on July 24, 1971). These rights and subject matter take effect in the relevant jurisdiction in which the License terms are sought to be enforced according to the corresponding provisions of the implementation of those treaty provisions in the applicable national law. If the standard suite of rights granted under applicable copyright law includes additional rights not granted under this License, such additional rights are deemed to be included in the License; this License is not intended to restrict the license of any rights under applicable law.

## Index of Proxy attributes

### F

#### Finger

##### AbstractFingerProxy

- max\_hop\_count, 23
- max\_hostname\_length, 23
- max\_line\_length, 23
- max\_username\_length, 23
- request\_detailed, 23
- request\_hostnames, 23
- request\_username, 24
- response\_footer, 24
- response\_header, 24
- strict\_username\_check, 24
- timeout, 24

##### fingerRequest

- hostname, 24
- username, 25

#### Ftp

##### AbstractFtpProxy

- active\_connection\_mode, 30
- auth\_tls\_ok\_client, 30
- auth\_tls\_ok\_server, 30
- buffer\_size, 30
- data\_mode, 31
- data\_port\_max, 31
- data\_port\_min, 31
- data\_protection\_enabled\_client, 31
- data\_protection\_enabled\_server, 31
- features, 31
- hostname, 31
- hostport, 31
- masq\_address\_client, 32
- masq\_address\_server, 32
- max\_continuous\_line, 32
- max\_hostname\_length, 32
- max\_line\_length, 32
- max\_password\_length, 32
- max\_username\_length, 32
- password, 32
- permit\_client\_bounce\_attack, 33
- permit\_empty\_command, 33
- permit\_server\_bounce\_attack, 33
- permit\_unknown\_command, 33

- proxy\_password, 33
- proxy\_username, 33
- request, 33
- request\_command, 33
- request\_parameter, 34
- request\_stack, 34
- response, 34
- response\_parameter, 34
- response\_status, 34
- response\_strip\_msg, 34
- strict\_port\_checking, 34
- target\_port\_range, 34
- timeout, 35
- transparent\_mode, 35
- username, 35
- valid\_chars\_username, 35

### H

#### Http

##### AbstractHttpProxy

- auth\_by\_cookie, 49
- auth\_by\_form, 49
- auth\_cache\_time, 50
- auth\_cache\_update, 50
- auth\_forward, 50
- auth\_realm, 50
- buffer\_size, 50
- connection\_mode, 50
- connect\_proxy, 50
- current\_header\_name, 51
- current\_header\_value, 51
- default\_port, 51
- enable\_session\_persistence, 51
- enable\_url\_filter, 51
- enable\_url\_filter\_dns, 51
- error\_files\_directory, 51
- error\_headers, 51
- error\_info, 52
- error\_msg, 52
- error\_silent, 52
- error\_status, 52
- keep\_persistent, 52
- language, 52
- max\_auth\_time, 52
- max\_body\_length, 52
- max\_chunk\_length, 53
- max\_header\_lines, 53
- max\_hostname\_length, 53

- max\_keepalive\_requests, 53
- max\_line\_length, 53
- max\_url\_length, 53
- parent\_proxy, 53
- parent\_proxy\_port, 53
- permit\_ftp\_over\_http, 54
- permit\_http09\_responses, 54
- permit\_invalid\_hex\_escape, 54
- permit\_null\_response, 54
- permit\_proxy\_requests, 54
- permit\_server\_requests, 54
- permit\_unicode\_url, 54
- request, 54
- request\_count, 55
- request\_header, 55
- request\_method, 55
- request\_mime\_type, 55
- request\_stack, 55
- request\_url, 55
- request\_url\_file, 55
- request\_url\_host, 55
- request\_url\_passwd, 56
- request\_url\_port, 56
- request\_url\_proto, 56
- request\_url\_scheme, 56
- request\_url\_username, 56
- request\_version, 56
- require\_host\_header, 56
- rerequest\_attempts, 56
- reset\_on\_close, 57
- response, 57
- response\_header, 57
- response\_mime\_type, 57
- response\_stack, 57
- rewrite\_host\_header, 57
- session\_persistence\_cookie\_name, 57
- session\_persistence\_cookie\_salt, 57
- strict\_header\_checking, 58
- strict\_header\_checking\_action, 58
- target\_port\_range, 58
- timeout, 58
- timeout\_request, 58
- timeout\_response, 58
- transparent\_mode, 58
- url\_category, 58
- url\_filter\_uncategorized\_action, 59
- use\_canonicalized\_urls, 59
- use\_default\_port\_in\_transparent\_mode, 59

- getRequestHeader
  - header, 59
- getResponseHeader
  - header, 60
- HttpProxyURIFilter
  - matcher, 61
- setRequestHeader
  - header, 60
  - new\_value, 60
- setResponseHeader
  - header, 60
  - new\_value, 60

## P

### Plug

- AbstractPlugProxy
  - bandwidth\_to\_client, 63
  - bandwidth\_to\_server, 63
  - buffer\_size, 63
  - copy\_to\_client, 63
  - copy\_to\_server, 63
  - packet\_stats\_interval\_packet, 64
  - packet\_stats\_interval\_time, 64
  - secondary\_mask, 64
  - secondary\_sessions, 64
  - shutdown\_soft, 64
  - stack\_proxy, 64
  - timeout, 64
- packetStats
  - client\_bytes, 65
  - client\_pkts, 65
  - server\_bytes, 65
  - server\_pkts, 65

### Pop3

- AbstractPop3Proxy
  - max\_authline\_count, 70
  - max\_password\_length, 70
  - max\_request\_line\_length, 71
  - max\_response\_line\_length, 71
  - max\_username\_length, 71
  - password, 71
  - permit\_longline, 71
  - permit\_unknown\_command, 71
  - reject\_by\_mail, 71
  - request, 71
  - request\_command, 72
  - request\_param, 72
  - response\_multiline, 72

- response\_param, 72
- response\_stack, 72
- response\_value, 72
- session\_timestamp, 72
- timeout, 72
- username, 73

## S

### Smtp

#### AbstractSmtpProxy

- active\_extensions, 76
- add\_received\_header, 76
- append\_domain, 77
- autodetect\_domain\_from, 77
- domain\_name, 77
- extensions, 77
- interval\_transfer\_noop, 77
- max\_auth\_request\_length, 77
- max\_request\_length, 77
- max\_response\_length, 77
- permit\_long\_responses, 78
- permit\_omission\_of\_angle\_brackets, 78
- permit\_unknown\_command, 78
- request, 78
- request\_command, 78
- request\_param, 78
- request\_stack, 78
- require\_crlf, 78
- resolve\_host, 79
- response, 79
- response\_param, 79
- response\_value, 79
- timeout, 79
- tls\_passthrough, 79
- unconnected\_response\_code, 79

#### SmtpProxy

- error\_soft, 80
- permit\_exclamation\_mark, 80
- permit\_percent\_hack, 80
- recipient\_matcher, 80
- relay\_check, 80
- relay\_domains, 80
- relay\_domains\_matcher, 80
- relay\_zones, 80
- sender\_matcher, 81

## T

### Telnet

#### AbstractTelnetProxy

- current\_var\_name, 85
- current\_var\_value, 85
- enable\_audit, 85
- negotiation, 85
- option, 85
- timeout, 85

## W

### Whois

#### AbstractWhoisProxy

- max\_line\_length, 86
- max\_request\_length, 87
- request, 87
- response\_footer, 87
- response\_header, 87
- timeout, 87

## Index of Core attributes

### A

#### Auth

##### \_\_init\_\_

- acl, 95
- authentication, 93
- authorization, 94
- authorize\_policy, 96
- cache, 93
- cleanup\_threshold, 91
- connect\_timeout, 101
- grouplist, 98
- intervals, 99
- name, 93, 94
- pki, 101
- port, 101
- provider, 93
- service\_equiv, 91
- timeout, 92, 101
- update\_stamp, 92
- userlist, 99
- wait\_authorization, 96
- wait\_timeout, 96, 97

#### AuthDB

##### \_\_init\_\_

- backend, 102
- name, 102
- pki\_ca, 103
- pki\_cert, 103
- serveraddr, 103
- ssl\_verify\_depth, 103
- use\_ssl, 103

### C

#### Chainer

##### SideStackChainer

- right\_chainer, 110
- right\_class, 111

##### \_\_init\_\_

- protocol, 106, 107, 108, 110, 112
- right\_chainer, 111
- right\_class, 111
- self, 109
- timeout\_connect, 106, 107, 108, 109, 110, 112

- timeout\_state, 108, 112

### D

#### Dispatch

##### CSZoneDispatcher

- services, 220

##### Dispatcher

- backlog, 222
- bindto, 222
- protocol, 222
- service, 222
- threaded, 222

##### \_\_init\_\_

- bindto, 221, 223
- follow\_parent, 221
- service, 223
- services, 221
- transparent, 223

### M

#### Matcher

##### RegexpFileMatcher

- ignore\_date, 166
- ignore\_file, 166
- match\_date, 166
- match\_file, 166

##### RegexpMatcher

- ignore, 167
- match, 167

##### \_\_init\_\_

- bind\_name, 169
- cache\_timeout, 169
- force\_delivery\_attempt, 169
- hosts, 165
- ignore\_fname, 167
- ignore\_list, 168
- match\_fname, 167
- match\_list, 168
- resolve\_on\_init, 165
- sender\_address, 169
- server, 166, 170
- server\_name, 169
- server\_port, 169

### N

#### NAT

- performTranslation

- addrs, 172
- nat\_type, 172
- session, 172
- \_\_init\_\_
  - addr, 179
  - addresses, 178
  - cacheable, 176
  - default\_reject, 173, 176, 177
  - from\_domain, 177
  - ip\_hash, 173
  - mapping, 173, 177
  - name, 176
  - nat, 176
  - prefix, 174
  - prefix\_mask, 174, 175
  - suffix, 174
  - to\_domain, 178

## P

### Proxy

- getCredentials
  - domain, 182
  - method, 182
  - port, 183
  - target, 183
  - username, 183

### Proxy

- encryption\_policy, 181
- language, 181
- setServerAddress
  - host, 183
  - port, 183
- userAuthenticated
  - entity, 184

## R

### Resolver

- \_\_init\_\_
  - mapping, 186
  - name\_server, 185
  - timeout, 185
  - use\_search\_domain, 185

### Router

- AbstractRouter
  - forge\_addr, 187
  - forge\_port, 187
- DirectedRouter
  - dest\_addr, 188

### TransparentRouter

- forced\_port, 190
- forge\_addr, 190
- \_\_init\_\_
  - dest\_addr, 188
  - forced\_port, 191
  - forge\_addr, 188, 189, 191
  - forge\_port, 189, 190, 191
  - overrideable, 189, 191

## S

### Service

#### AbstractService

- name, 198

#### DenyService

- ipv4\_setting, 199
- ipv6\_setting, 199
- name, 200

#### PFService

- dnat\_policy, 201
- router, 201
- snat\_policy, 201

#### Service

- authentication\_policy, 203
- authorization\_policy, 203
- auth\_name, 203
- chainer, 203
- dnat\_policy, 203
- encryption\_policy, 203
- instance\_id, 203
- keepalive, 203
- max\_instances, 204
- max\_sessions, 204
- num\_instances, 204
- proxy\_class, 204
- resolver\_policy, 204
- router, 204
- snat\_policy, 204

#### startInstance

- session, 207

#### \_\_init\_\_

- authentication\_policy, 205
- authorization\_policy, 205
- auth\_name, 205
- chainer, 206
- dnat\_policy, 206
- encryption\_policy, 206
- keepalive, 206

- limit\_target\_zones\_to, 206
- log\_spec, 200, 202
- log\_verbose, 200, 202
- max\_instances, 206
- max\_sessions, 206
- name, 199, 200, 206
- proxy\_class, 207
- resolver\_policy, 207
- router, 207
- snat\_policy, 207

#### Session

- setTargetAddress
  - addr, 209

#### StackedSession

- chainer, 208
- owner, 208
- server\_address, 208
- server\_local, 208
- server\_stream, 208
- server\_zone, 209
- target\_address, 209
- target\_local, 209
- target\_zone, 209

#### SockAddr

##### SockAddrInet

- ip, 210
- ip\_s, 210
- port, 210
- type, 211

##### SockAddrInet6

- ip, 211
- ip\_s, 211
- port, 211
- type, 211

##### SockAddrInetHostname

- ip, 212
- ip\_s, 212
- port, 212
- type, 212

##### SockAddrInetRange

- ip, 213
- ip\_s, 213
- port, 213
- type, 213

##### SockAddrUnix

- type, 213

## Z

### Zone

- \_\_init\_\_
  - addr, 217
  - admin\_parent, 217
  - hostnames, 218
  - name, 218



## Index of all attributes

### A

acl, 95  
active\_connection\_mode, 30  
active\_extensions, 76  
addr, 179, 209, 217  
addresses, 178  
addrs, 172, 214  
add\_received\_header, 76  
admin\_parent, 217  
append\_domain, 77  
authentication, 93  
authentication\_policy, 203, 205  
authorization, 94  
authorization\_policy, 203, 205  
authorize\_policy, 96  
auth\_by\_cookie, 49  
auth\_by\_form, 49  
auth\_cache\_time, 50  
auth\_cache\_update, 50  
auth\_forward, 50  
auth\_name, 203, 205  
auth\_realm, 50  
auth\_tls\_ok\_client, 30  
auth\_tls\_ok\_server, 30  
autodetect\_domain\_from, 77

### B

backend, 102, 215  
backlog, 222  
bandwidth\_to\_client, 63  
bandwidth\_to\_server, 63  
bindto, 221, 222, 223  
bind\_name, 169  
buffer\_size, 30, 50, 63  
bytes\_recvd, 223  
bytes\_sent, 224

### C

cache, 93  
cacheable, 176  
cache\_directory, 137, 161, 162  
cache\_timeout, 169  
ca\_hint\_directory, 125, 127

certificate, 114  
certificates, 158, 159  
certificate\_file\_path, 123, 124, 125  
chainer, 203, 206, 208  
check\_subject, 151, 152  
cipher, 131, 133, 148, 149, 155, 157  
ciphers\_tlsv1\_3, 131, 133, 148, 149, 155, 157  
cipher\_server\_preference, 131, 133  
cleanup\_threshold, 91  
client\_bytes, 65  
client\_certificate\_generator, 128, 129, 130, 138, 141, 142, 143, 144, 159, 160  
client\_max\_line\_length, 20  
client\_pkts, 65  
client\_security, 138  
client\_ssl\_options, 128, 129, 130, 131, 138, 139, 141, 142, 143, 144, 159, 160  
client\_verify, 129, 130, 131, 138, 139, 141, 142, 143, 144, 159, 160  
connection\_mode, 50  
connect\_proxy, 50  
connect\_timeout, 101  
copy\_to\_client, 63  
copy\_to\_server, 63  
current\_header\_name, 51  
current\_header\_value, 51  
current\_var\_name, 85  
current\_var\_value, 85

### D

data\_mode, 31  
data\_port\_max, 31  
data\_port\_min, 31  
data\_protection\_enabled\_client, 31  
data\_protection\_enabled\_server, 31  
default, 146, 147  
default\_port, 51  
default\_reject, 173, 176, 177  
dest\_addr, 188  
dh\_params, 131  
dh\_param\_file\_path, 134  
disable\_compression, 131, 134, 148, 150, 155, 157  
disable\_renegotiation, 134  
disable\_send\_root\_ca, 132, 134  
disable\_session\_cache, 132, 134, 148, 150, 155, 157  
disable\_ticket, 132, 134, 148, 150, 155, 157  
disable\_tlsv1, 132, 134, 148, 150, 155, 157  
disable\_tlsv1\_1, 132, 134, 148, 150, 156, 157

disable\_tlsv1\_2, 132, 135, 148, 150, 156, 157  
disable\_tlsv1\_3, 132, 135, 149, 150, 156, 158  
dnat\_policy, 201, 203, 206  
domain, 182  
domain\_name, 77  
dst\_iface, 195  
dst\_ifgroup, 195  
dst\_port, 195  
dst\_subnet, 195  
dst\_zone, 196

## E

enable\_audit, 85  
enable\_session\_persistence, 51  
enable\_url\_filter, 51  
enable\_url\_filter\_dns, 51  
encryption, 139, 140  
encryption\_policy, 181, 203, 206  
entity, 184  
error\_files\_directory, 51  
error\_headers, 51  
error\_info, 52  
error\_msg, 52  
error\_silent, 52  
error\_soft, 80  
error\_status, 52  
extensions, 77  
extension\_whitelist, 137, 162

## F

fd, 224  
features, 31  
file\_path, 136  
follow\_parent, 221  
forced\_port, 190, 191  
force\_delivery\_attempt, 169  
forge\_addr, 187, 188, 189, 190, 191  
forge\_port, 187, 189, 190, 191  
from\_domain, 177

## G

grouplist, 98

## H

header, 59, 60  
high, 194  
host, 183

hostname, 24, 31  
hostnames, 218  
hostname\_certificate\_map, 147  
hostport, 31  
hosts, 165

## I

ignore, 115, 167  
ignore\_date, 166  
ignore\_file, 166  
ignore\_fname, 167  
ignore\_list, 168  
instance\_id, 203  
intermediate\_revocation\_check\_type, 120, 121, 125,  
127, 151, 152  
intervals, 99  
interval\_transfer\_noop, 77  
ip, 210, 211, 212, 213  
ipv4\_setting, 199  
ipv6\_setting, 199  
ip\_hash, 173  
ip\_s, 210, 211, 212, 213

## K

keepalive, 203, 206  
keep\_persistent, 52  
key\_file, 161, 163  
key\_file\_path, 145, 146  
key\_passphrase, 162, 163

## L

language, 52, 181  
leaf\_revocation\_check\_type, 120, 122, 125, 127, 151,  
153  
level, 180  
limit\_target\_zones\_to, 206  
log\_spec, 200, 202  
log\_verbose, 200, 202  
low, 195

## M

mapping, 173, 177, 186  
masq\_address\_client, 32  
masq\_address\_server, 32  
match, 115, 167  
matcher, 61  
match\_date, 166

match\_file, 166  
match\_fname, 167  
match\_list, 168  
max\_authline\_count, 70  
max\_auth\_request\_length, 77  
max\_auth\_time, 52  
max\_body\_length, 52  
max\_chunk\_length, 53  
max\_continuous\_line, 32  
max\_header\_lines, 53  
max\_hop\_count, 23  
max\_hostname\_length, 23, 32, 53  
max\_instances, 204, 206  
max\_keepalive\_requests, 53  
max\_line\_length, 23, 32, 53, 86  
max\_password\_length, 32, 70  
max\_request\_length, 77, 87  
max\_request\_line\_length, 71  
max\_response\_length, 77  
max\_response\_line\_length, 71  
max\_sessions, 204, 206  
max\_url\_length, 53  
max\_username\_length, 23, 32, 71  
method, 182  
msg, 181

## N

name, 93, 94, 102, 139, 140, 176, 198, 199, 200, 206, 216, 218, 224  
name\_server, 185  
nat, 176  
nat\_type, 172  
negotiation, 85  
new\_value, 60  
num\_instances, 204

## O

option, 85  
overrideable, 189, 191  
owner, 208

## P

packet\_stats\_interval\_packet, 64  
packet\_stats\_interval\_time, 64  
params, 136  
parent\_proxy, 53  
parent\_proxy\_port, 53

passphrase, 124, 145, 146  
password, 32, 71  
permit\_client\_bounce\_attack, 33  
permit\_empty\_command, 33  
permit\_exclamation\_mark, 80  
permit\_ftp\_over\_http, 54  
permit\_http09\_responses, 54  
permit\_invalid\_hex\_escape, 54  
permit\_longline, 71  
permit\_long\_responses, 78  
permit\_null\_response, 54  
permit\_omission\_of\_angle\_brackets, 78  
permit\_percent\_hack, 80  
permit\_proxy\_requests, 54  
permit\_server\_bounce\_attack, 33  
permit\_server\_requests, 54  
permit\_unicode\_url, 54  
permit\_unknown\_command, 33, 71, 78  
pki, 101  
pki\_ca, 103  
pki\_cert, 103  
port, 101, 183, 210, 211, 212, 213  
prefix, 174  
prefix\_mask, 174, 175  
private\_key, 137  
private\_key\_password, 123, 124, 125  
proto, 196  
protocol, 106, 107, 108, 110, 112, 222  
provider, 93  
proxy\_class, 204, 207  
proxy\_password, 33  
proxy\_username, 33

## R

recipient, 179, 180  
recipient\_matcher, 80  
reject\_by\_mail, 71  
relay\_check, 80  
relay\_domains, 80  
relay\_domains\_matcher, 80  
relay\_zones, 80  
request, 33, 54, 71, 78, 87  
request\_command, 33, 72, 78  
request\_count, 55  
request\_detailed, 23  
request\_header, 55  
request\_hostnames, 23  
request\_method, 55

request\_mime\_type, 55  
request\_param, 72, 78  
request\_parameter, 34  
request\_stack, 34, 55, 78  
request\_url, 55  
request\_url\_file, 55  
request\_url\_host, 55  
request\_url\_passwd, 56  
request\_url\_port, 56  
request\_url\_proto, 56  
request\_url\_scheme, 56  
request\_url\_username, 56  
request\_username, 24  
request\_version, 56  
required, 120, 122, 126, 127  
require\_crlf, 78  
require\_host\_header, 56  
rerequest\_attempts, 56  
reset\_on\_close, 57  
resolver\_policy, 204, 207  
resolve\_host, 79  
resolve\_on\_init, 165  
response, 34, 57, 79  
response\_footer, 24, 87  
response\_header, 24, 57, 87  
response\_mime\_type, 57  
response\_multiline, 72  
response\_param, 72, 79  
response\_parameter, 34  
response\_stack, 57, 72  
response\_status, 34  
response\_strip\_msg, 34  
response\_value, 72, 79  
rewrite\_host\_header, 57  
right\_chainer, 110, 111  
right\_class, 111  
router, 201, 204, 207  
rule\_id, 196

## S

secondary\_mask, 64  
secondary\_sessions, 64  
self, 109  
sender\_address, 169  
sender\_matcher, 81  
server, 166, 170  
serveraddr, 103  
server\_address, 208

server\_bytes, 65  
server\_certificate\_generator, 154, 159, 160  
server\_local, 208  
server\_max\_line\_length, 21  
server\_name, 169  
server\_name\_matcher, 115, 116  
server\_pkts, 65  
server\_port, 169  
server\_ssl\_options, 141, 142, 143, 144, 154, 159, 161  
server\_stream, 208  
server\_verify, 141, 142, 143, 144, 154, 155, 160, 161  
server\_zone, 209  
service, 196, 222, 223  
services, 220, 221  
service\_equiv, 91  
session, 21, 172, 207  
session\_cache\_size, 132, 135, 149, 150, 156, 158  
session\_persistence\_cookie\_name, 57  
session\_persistence\_cookie\_salt, 57  
session\_timestamp, 72  
shared\_groups, 133, 135, 149, 151, 156, 158  
shutdown\_soft, 64  
snat\_policy, 201, 204, 207  
src\_iface, 196  
src\_ifgroup, 197  
src\_port, 197  
src\_subnet, 197  
src\_zone, 197  
ssl\_verify\_depth, 103  
stack\_proxy, 64  
strict\_header\_checking, 58  
strict\_header\_checking\_action, 58  
strict\_port\_checking, 34  
strict\_username\_check, 24  
suffix, 174

## T

target, 183  
target\_address, 209  
target\_local, 209  
target\_port\_range, 34, 58  
target\_zone, 209  
threaded, 222  
timeout, 24, 35, 58, 64, 72, 79, 85, 87, 92, 101, 135, 151, 158, 185  
timeout\_connect, 106, 107, 108, 109, 110, 112  
timeout\_request, 58  
timeout\_response, 58

timeout\_state, 108, 112  
tls\_passthrough, 79  
to\_domain, 178  
transparent, 223  
transparent\_mode, 35, 58  
trusted\_ca, 137  
trusted\_ca\_files, 162, 163  
trusted\_certs\_directory, 121, 122, 126, 128, 151, 153  
trust\_level, 120, 122, 126, 127, 151, 153  
type, 181, 211, 212, 213

## U

unconnected\_response\_code, 79  
untrusted\_ca, 137  
untrusted\_ca\_files, 162, 163  
update\_stamp, 92  
url\_category, 58  
url\_filter\_uncategorized\_action, 59  
userlist, 99  
username, 25, 35, 73, 183  
use\_canonicalized\_urls, 59  
use\_default\_port\_in\_transparent\_mode, 59  
use\_search\_domain, 185  
use\_ssl, 103

## V

valid\_chars\_username, 35  
verify\_ca\_directory, 121, 122, 126, 128, 152, 153  
verify\_crl\_directory, 121, 122, 126, 128, 152, 153  
verify\_depth, 121, 122, 126, 128, 152, 153

## W

wait\_authorization, 96  
wait\_timeout, 96, 97