

Converged Network Analyzer Command Reference Guide

for the Converged Network Analyzer (CNA), Adaptive Path Controller-Enterprise (APC-E), Adaptive Path Controller-Internet (APC-I), and the CNA Server Products

Version 4.1

14-300540
Issue 6
May 2008

Copyright 2008, Avaya Inc.
All Rights Reserved

This document contains information related to Converged Network Analyzer (CNA) software (as defined below) and Documentation ("Product"). "Documentation" means this document and Avaya's information manuals in printed or electronic form containing operating instructions and performance specifications that Avaya or its suppliers generally make available to users of its products, and which Avaya delivers to End User with the Products. "End User" means any customer of Avaya or its authorized resellers, or any end user of the Product. See the Software and Documentation DVD/CD inserts for additional legal and licensing information.

Notice

Changes and corrections to the information in this document may be incorporated in future releases.

Disclaimer

Avaya, its affiliates or subsidiaries ("Avaya") are not responsible for any modifications, additions or deletions to the original published version of the Documentation unless such modifications, additions or deletions were performed by Avaya. End User agrees to indemnify and hold harmless Avaya, Avaya's agents, servants, directors, officers, and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to the Documentation to the extent made by the End User.

Warranty

Avaya provides a limited warranty on the Product. Refer to your customer sales agreement to establish the terms of the limited warranty. In addition, Avaya's standard warranty language as well as information regarding support for the Product, while under warranty, is available through the following web site: http://www.avaya.com/support.

License

USE OR INSTALLATION OF THE PRODUCT INDICATES THE END USER'S ACCEPTANCE OF THE GENERAL LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE AT: http://www.avaya.com/support ("GENERAL LICENSE TERMS"). DO NOT USE THE PRODUCT IF YOU DO NOT WISH TO BE BOUND BY THE GENERAL LICENSE TERMS. IN ADDITION TO THE GENERAL LICENSE TERMS, THE FOLLOWING LICENSE TERMS AND RESTRICTIONS WILL APPLY TO THE PRODUCT.

Avaya grants End User a license within the scope of the license types described below. The applicable number of licenses and units of capacity for which the license is granted will be one (1), unless a different number of licenses or units of capacity is specified in the Documentation or other materials available to End User. "Designated Processor" means a single stand-alone computing device. "Server" means a Designated Processor that hosts a software application to be accessed by multiple users. "Software" means the computer programs in object code, originally licensed by Avaya and ultimately utilized by End User, whether as stand-alone products or pre-installed on Hardware. "Hardware" means the standard hardware products, originally sold by Avaya and ultimately utilized by End User.

If your system is running in a TDM environment, the following license restriction applies:

Designated System(s) License (DS). End User may install and use each copy of the Software on only one Designated Processor, unless a different number of Designated Processors is indicated in the Documentation or other materials available to End User. Avaya may require the Designated Processor(s) to be identified by type, serial number, feature key, location or other specific designation, or to be provided by End User to Avaya through electronic means established by Avaya specifically for this purpose.

If your system is running in an IP environment, the following license restriction applies:

Concurrent User License (CU). End User may install and use the Software on multiple Designated Processors or one or more Servers, so long as only the licensed number of Units are accessing and using the Software at any given time. A "Unit" means the unit on which Avaya, at its sole discretion, bases the pricing of its licenses and can be, without limitation, an agent, port or user, an e-mail or voice mail account in the name of a person or corporate function (e.g., webmaster or helpdesk), or a directory entry in the administrative database utilized by the Product that permits one user to interface with the Software. Units may be linked to a specific, identified Server.

For all systems, the following license restriction applies:

Shrinkwrap License (SR). With respect to Software that contains elements provided by third party suppliers, End User may install and use the Software in accordance with the terms and conditions of the "shrinkwrap" or "clickwrap" license accompanying the Software ("Shrinkwrap License"). The text of the Shrinkwrap License will be available from Avaya upon End User's request (see "Copyright" below for more information).

Copyright

Except where expressly stated otherwise, the Product is protected by copyright and other laws respecting proprietary rights. Unauthorized reproduction, transfer, and or use can be a criminal, as well as a civil, offense under the applicable law.

Certain Software programs or portions thereof included in the Product may contain software distributed under third party agreements ("Third Party Components"), which may contain terms that expand or limit rights to use certain portions of the Product ("Third Party Terms"). Information identifying Third Party Components and the Third Party Terms that apply to them is available on Avaya's web site at http://support.avaya.com/ThirdPartyLicense/.

The disclaimers of warranties and limitations of liability set forth in the Third Party Terms do not affect any express warranty or limitation of liability that may be provided to you by Avaya pursuant to the license terms covering the Product contained in a separate written agreement between you and Avaya. To the extent there is a conflict between the General License Terms or your customer sales agreement and any Third Party Terms, the Third Party Terms shall prevail solely for such Third Party Components.

Security and virus disclaimer

End User's decision to acquire products from third parties is End User's sole responsibility, even if Avaya helps End User identify, evaluate or select them. Avaya is not responsible for, and will not be liable for, the quality or performance of such third party products or their suppliers.

ALL INFORMATION IS BELIEVED TO BE CORRECT AT THE TIME OF PUBLICATION AND IS PROVIDED "AS IS". AVAYA DISCLAIMS ALL WARRANTIES, EITHER EXPRESS OR IMPLIED, INCLUDING THE WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE AND FURTHERMORE, AVAYA MAKES NO REPRESENTATIONS OR WARRANTIES THAT THE STEPS RECOMMENDED WILL ELIMINATE SECURITY OR VIRUS THREATS TO END USER' SYSTEMS. IN NO EVENT SHALL AVAYA BE LIABLE FOR ANY DAMAGES WHATSOEVER ARISING OUT OF OR IN CONNECTION WITH THE INFORMATION OR RECOMMENDED ACTIONS PROVIDED HEREIN, INCLUDING DIRECT, INDIRECT, CONSEQUENTIAL DAMAGES, LOSS OF BUSINESS PROFITS OR SPECIAL DAMAGES, EVEN IF AVAYA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Avaya does not warrant that this Product is immune from or will prevent unauthorized use of telecommunication services or facilities accessed through or connected to it. Avaya is not responsible for any damages or charges that result from either unauthorized uses or from incorrect installations of the security patches that are made available from time to time.

Suspected security vulnerabilities with Avaya products should be reported to Avaya by sending mail to securityalerts@avaya.com.

Trademarks

All trademarks identified by ® and ™ are registered trademarks or trademarks of Avaya Inc. All other trademarks are the property of their respective owners.

Contents

Chapter 1: Introduction to the Command Line Interface 17
Chapter contents 17
Intended audience 17
Documentation library 18
Typographic Conventions 19
Terminal Connection 21
Remote Connection 21
API 21
User access levels 21
Accessing and exiting command modes 22
Medic Services 24
Restarting the Decision Maker Process 26
Command Line Help 27
Legal Characters 27

Chapter 2: Commands 29
aaa authentication login 29
active-measurement avoid 30
active-measurement group 30
active-measurement max-probes-per-second 32
active-measurement name 32
active-measurement traceprobe avoid address (discontinued) 33
active-measurement-vip (deprecated) 34
alarm-absolute disable 35
alarm-absolute disable ping 36
alarm-absolute disable rtp 37
alarm-absolute disable tcp 39
alarm-absolute disable traceroute 40
alarm name 41
alarm-relative disable 42
alarm-relative disable ping 43
alarm-relative disable rtp 45
alarm-relative disable tcp 46
alarm-relative disable traceroute 47
alarm-threshold-exceeded 49
application 50

application age-out 51
application detection 52
application no-detection 52
application-model enterprise 53
application-model multimedia 56
application-model other 58
application-model streaming 60
application-model voice 62
application-model web 65
application priority 66
as-path 67
as-path-padding (discontinued) 69
assert-routes-matching-bgp 70
asserted-routes-timeout (discontinued) 70
auto-create-subnets 71
auto-create-subnets limit 71
auto-create-subnets mask-length 72
auto-create-subnets verify-frequency 73
backup 73
bgp 75
bgp-community 76
billing link 78
chatter assign 79
chatter join-hive 80
chatter leave-hive 80
chatter missed-heartbeats 81
chatter node add 82
chatter node split 83
chatter remove-cna-agent 83
cidr-lookup 84
clear active-measurements 85
clear arp 86
clear efc active-measurement 87
clear efc prefix (deprecated) 87
clear counters web-hits 88
clear prefix (deprecated) 88

clear snapshots (discontinued) . . . 88
clear statistics . . . 89
clear efc subnet . . . 90
clear subnet . . . 90
clear trending . . . 91
clear trending report . . . 92
clear tty . . . 93
cna-agent . . . 93
community . . . 94
configure memory . . . 95
configure subnet . . . 95
configure terminal . . . 96
contact . . . 96
content . . . 97
copy . . . 98
customer-prefix-file (deprecated) . . . 100
customer-subnet-file . . . 100
damped-mode . . . 101
data-source efc enable . . . 101
data-source efc action . . . 102
data-source efc sampling-period . . . 103
date . . . 103
debug . . . 104
debug display . . . 105
debug on . . . 106
decision-policy . . . 108
decision-policy-assignment . . . 110
delete . . . 111
delete report . . . 111
deny . . . 112
description . . . 114
diagnose target . . . 115
diagnose probe . . . 116
direction . . . 117
disable . . . 118
dump-measurements . . . 119

duplex 120
efc bytes minimum 120
efc bytes top 121
efc enable 123
efc filter 124
efc last-update-within 125
efc exclude-prefix (deprecated) 126
efc exclude-subnet 126
efc source 128
email 129
email format image 132
enable 132
enable password 133
end 134
erase startup-config 134
eventdump 135
exclude 136
exec-summary http-turns 137
exec-summary improvable min-hrtt-saving 138
exec-summary improvable min-percent-improvement 139
exec-summary min-bgp-hrtt 141
exit 142
export event 142
export report 143
export report (scheduled) 144
export trend 145
flap-control 146
group 148
halt 148
header 149
help 150
history 150
host 151
hostname 151
hrtt-smoothing-factor (deprecated) 152
ignore-measurements 153

interface 154
interval-between-packets 156
ip access-group 157
ip access-list 158
ip address 159
ip domain-list. 159
ip domain-lookup 160
ip domain-name 160
ip flow-export destination 161
ip ftp password. 162
ip ftp username 163
ip host 163
ip http password 164
ip http server header append (deprecated) 165
ip http server on 165
ip http server port 166
ip http server port ssl 167
ip http server ssl on 168
ip http username 169
ip http ustat-server map 169
ip http ustat-server redirect 170
ip mtu 171
ip name-server 172
ip route 172
ip ssh key generate 174
ip ssh key generate cna-agent 174
ip ssh on 175
ip telnet on 175
level 176
link 177
link-group 179
link-measurement-timeout 179
load evaluation-interval 180
load min-required-samples 182
load optimization 183
load threshold-table 184

load use-threshold-table 185
location 186
logging 186
logging buffered (deprecated) 187
logging console 188
logging facility 189
logging host 190
logging on 191
logging-server 192
logging trap 194
logout 195
ls 195
max-ttl 196
measurement intra-zone disable 197
medic call 198
medic set 198
min-required-load-samples 200
min-required-measurements 201
module 202
module chatter 203
module chatter or testplug 204
monitor enable 205
monitor interface 205
more 206
neighbor allow-next-hop-self (removed) 206
neighbor buffered (removed) 207
neighbor description 208
neighbor link 209
neighbor redirect-peering 210
neighbor remote-as 211
neighbor report-alias 211
neighbor shutdown 212
neighbor summarize-bgp-events 213
neighbor update-source 214
neighbor uses-next-hop-self-address 214
netflow enable 215

netflow filter interface 216
netflow interface 217
netflow port 218
netflow source 219
ntp broadcastclient 220
ntp broadcastdelay 220
ntp disable 221
ntp peer 222
ntp server 222
number-of-packets 223
outage-detection active-probe 224
outage-detection always-move-from-bgp (deprecated) 225
outage-detection brownout 226
outage-detection on (deprecated) 228
outage-detection silence 229
outage-detection silence always-move-from-bgp 230
outage-detection total-silence 231
outage 233
outbound-bgp-community (discontinued) 233
packet-length 234
penalty 235
performance-group 236
persistent-hosts 236
persistent-hosts timeout <hours> 238
neighbor uses-next-hop-self-address 240
performance-violation-detection metric (deprecated) 241
ping 241
polling bgp-status interface-ip-address 243
polling bgp-status interval 244
polling bgp-status on 244
polling bgp-status snmp-agent 245
polling bgp-status unknown-threshold 246
polling interface 247
polling interface-description (deprecated) 248
polling interface-index (deprecated) 248
polling interface-ip-address (deprecated) 249

polling interface-status down-threshold 250
polling interface-status interval 251
polling interface-status on 252
polling physical-link 253
polling snmp-agent 254
polling usage interval 255
polling usage on 256
polling usage unknown-threshold 257
polling usage use-64-bit-counter 258
port 259
port (Chatter module only) 260
prefer-in-winner-set bgp 261
prefer-in-winner-set most-reliable 262
prefix-cleanup schedule (deprecated) 263
prefix (deprecated) 263
prefix resync (deprecated) 263
prefix-update-rate (deprecated) 263
priority-subnets 263
probe-backoff 264
provider-as 266
publish report 266
publish trend 268
radius-server host 269
radius-server key 270
radius-server timeout 271
rate 272
rate (Chatter module only) 273
recover report 275
recover startup-config 275
refresh-time acquire 276
refresh-time host-disappeared 277
refresh-time proxy-disappeared 278
refresh-time reacquire 278
refresh-time verify 279
reliability half-life 280
reliability tolerance 281

reload 282
remark 282
remote-exec 283
rename 285
rename zone 285
repeat 286
report group 287
report prefix-alias (deprecated) 288
report prefix-limit (deprecated) 288
report subnet-alias 289
report subnet-limit 289
require-bgp-snmp-poll 290
require-routes-from-all-routers 291
restart decision-maker 292
restore 293
resync chatter chapId 295
resync chatter node-alias 295
resync efc filter 296
resync load threshold-table 296
resync prefix (deprecated) 297
resync subnet 297
retries 299
route-assert-filter 300
route-assert-mode 301
sample-collection 302
send-data 302
service password-encryption 303
service password-in-config 304
service timestamps log 305
set-alarm edge 306
set-alarm zone 307
set-application-model 308
set-decision-policy 309
set-link-group 310
show active-measurements 311
show active-measurements schedule 313

show arp 314
show bootvar. 315
show cell 315
show chatter cna-agent-device 320
show chatter master-key 322
show chatter ssl 322
show cidr-lookup 323
show cna-agent 324
show command-log 326
show configuration 329
show counts prefix (deprecated) 330
show counts subnet 330
show debugging 330
show diagbus 331
show efc 332
show event 333
show edge 335
show image 338
show interfaces 339
show ip bgp neighbor 340
show ip route. 341
show ip sockets 342
show link 342
show link billing 343
show local-debug 344
show logging. 345
show medic 347
show nodes 351
show ntp 352
show outage 354
show outage chatter 355
show outage-not-rescued 356
show outage-rescued 356
show penalty-setting link 357
show prefix (deprecated) 359
show reload 360

show report 360
show running-config 362
show snmp 363
show ssl 364
show startup-config 365
show statistics 366
show statistics chatter 367
show statistics dmaker 369
show subnet 370
show tech-support 375
show trending 376
show users 377
show zone 378
show version 381
shutdown 382
site 382
smtp from 383
smtp to 384
smtp to (Chatter module only) 384
smtp-server 385
smtp-server (Chatter module only) 386
snapshots clear-after (deprecated) 387
snapshots on 388
snapshots prefix-limit (deprecated) 388
snapshots schedule 388
snapshots subnet-limit 390
snmp-agent-ip-address 390
snmp host 391
snmp-server 392
snmp-server contact 393
snmp-server enable traps 394
snmp-server host 395
snmp-server informs 398
snmp-server location 399
snmpwalk 399
source-route 400

speed 400
ssl add 401
ssl delete 403
ssl export 403
ssl generate-csr 404
start-avoidance percent-below-max 406
subnet-cleanup schedule 407
subnet 407
subnet-cleanup schedule 408
subnet resync 409
subnet-update-rate 411
supernode 413
sysuptime 414
tacacs-server host 414
tacacs-server key 415
target 416
target (Chatter module only) 418
tcpdump 419
terminal 422
threshold host-disappeared 422
threshold host-link-fail 423
threshold ping 423
threshold proxy-disappeared 425
threshold proxy-link-fail 426
threshold rtp 426
threshold tcp 428
threshold traceroute 430
time 431
timeout 432
timeout (Chatter module only) 433
timezone 434
topological-analysis disable 452
tos byte 453
tos byte (Chatter module only) 454
traceprobe convergence-rule 455
traceprobe refresh 456

traceprobe retries-per-hop 457
traceprobe max-failing-hops 458
traceprobe hops 458
traceroute 459
traceroute-vips 461
tree 462
trending active-measurement group (discontinued) 462
trending clear-after (deprecated) 463
trending decision-policy 463
trending frequency (discontinued) 464
trending on 465
trending prefix-limit (deprecated). 465
trending subnet-limit 465
trending size (deprecated) 466
tunnel 466
type 467
usage contract (removed) 467
usage-reporting billing-period 468
usage-reporting cost 469
usage-reporting formula 470
usage-reporting percentile 472
use link-group 473
use named-vip 474
user-defined-application 475
username 476
version 477
vip 478
web-speedup http-turns 479
web-speedup improvable min-hrtt-saving 481
web-speedup improvable min-percent-improvement 482
web-speedup min-bgp-hrtt 483
web-speedup worst-bgp-hrtt-percent 484
winner-set-priority 485
winner-set-width 486
write 487
zone 488

Index 491

Chapter 1: Introduction to the Command Line Interface

Chapter contents

- Intended audience
- Documentation library
- Typographic Conventions
- Terminal Connection
- Remote Connection
- API
- User access levels
- Accessing and exiting command modes
- Medic Services
- Restarting the Decision Maker Process
- Command Line Help
- Legal Characters

Intended audience

This document is intended for use by IP administrators. You should have a detailed knowledge of how IP networks work, including an in-depth familiarity with the Border Gateway Protocol (BGP).

Documentation library

You can download the following documents at http://support.avaya.com:

Table 1: Documentation library

Document Title	Document No.	Content
CNA Install Guide	14-300539	• Hardware specifications • Installing a CNA server • Installing a CNA Agent Device
CNA Fundamentals Guide	14-601298	• Product overview • Configuration fundamentals • Configuring a new CNA server
CNA Monitoring Guide	14-601299	Current path monitoring • Collecting endpoint addresses • Security • Server-based measurement • Agent-based measurement • Decision policies and application models
Adaptive Path Controller-Enterprise Guide	14-601303	• Adaptive path control • High availability • WAN cost and load optimization
Adaptive Path Controller-Internet Guide	14-602423	• Configuring your routers • Polling routers • VPN integration
CNA Operator Guide	14-601304	• Web interface • Server-based measurement reports • Agent-based measurement reports
CNA Appendices	14-601329	• Diagnostic tools • Supported MIBs • Upgrade procedure • Migrating configuration files between hardware models • API

1 of 2

Table 1: Documentation library (continued)

Document Title	Document No.	Content
CNA Command Reference Guide	14-300540	Describes all CLI commands for the CNA server.
CNA Release Notes		• New features • Supported hardware platforms, browsers, and operating systems • Resolved issues • Known issues
		2 of 2

To set up user accounts for the web site's Support area, call:

- 1 (877) 733-5511 in the United States
- 001 (978) 552-0444 outside of the United States

Or you can send an e-mail message to *support@avaya.com.*

Typographic Conventions

Table 2 shows the typographical conventions that are used in this book:

Table 2: Typographic conventions

Convention	Meaning
`Courier type`	The Courier type face is used to denote one of the following: • Text as displayed on a monitor • Command syntax examples.
bold	Bold text denotes command keywords (words that must be entered exactly as shown).
,	A comma separates a list of values where one or more can be specified. For example: good, better, best.
	1 of 2

Table 2: Typographic conventions (continued)

Convention	Meaning
italics	Italics can mean any of the following: • filenames • book titles • command argument names and values • new terms not commonly understood
< >	Words enclosed in angle brackets should be replaced by user-specific data. In the following example, the words *userid* and *password* should be replaced by actual user ID and password values: `Login:` ***<userid>*** `Password:` ***<password>***
[ ]	Square brackets are used to denote optional command line arguments.
{ }	Curly braces enclose a range of choices for a required argument. When used within square brackets, curly braces indicate that if you choose to set the optional argument indicated by the square brackets, then you must choose one of the choices enclosed in the curly brackets. In the example below, the curly brackets indicate that either *in* or *out* is always required for the **`ip access-group`** command: **`ip access-group <name> {in | out}`** In the following example, the *version* argument is optional for the **`snmp-serverhost`** command, but if it is set it must be either *1* or *2c*: **`snmp-serverhost <host-id> [traps | inform] [version {1 | 2c}]`**
|	A vertical bar separates multiple choices among arguments. In the following example, the vertical bar indicates that the first argument of the **`clear subnet`** command can be either a subnet in standard, dotted-decimal format (1.2.3.4/*n*) or the keyword **all**: **`clear subnet {<subnet> | all} [routes-only]`** The second argument—the keyword **routes-only**—is enclosed in square brackets because it is optional.

2 of 2

Terminal Connection

To use the CLI, connect a terminal directly to the serial port of the device running the product, or use the telnet or ssh protocol to log in from a remote workstation.

If you connect your terminal to the serial port on the Management Processor module, the CLI can be used to execute commands on any of the product modules.

If you connect your terminal to the serial port on any module other than the Management Processor, the CLI will operate in local mode only and commands will apply only to that module. The number of commands that are available to you in local CLI mode are severely reduced from the number of commands that can be executed in system-wide mode.

Remote Connection

You can also establish a remote connection with the product through the ethernet port on the management module (top port only; the second ethernet port is not supported on any of the product modules).

The telnet protocol is enabled by default provided there is an **enable** password defined and a user account has been configured. See the **`enable password`** and **`username`** commands in Chapter 2: Commands.

If you prefer a more secure means of communication, you can disable the telnet protocol and enable the ssh protocol. See the **`ip telnet on`** and **`ip ssh on`** commands.

API

You can also configure the product from your client application, using the CNA5 API. For more information, see Appendix F: API in the *CNA Appendices*.

User access levels

There are two levels of user access to the CLI:

- Login—initial access level when you first log in to the system
- Privileged—requires the user to execute the **`enable`** command (which can be protected by a password)

Accessing and exiting command modes

The prompt that is displayed on the command line changes according to access level and operating mode.

The prompt always consists of the product's hostname followed by a description that indicates access level and mode:

- a right angle bracket (>) indicates login-level access
- a pound sign (#) indicates privileged-level access

Table 3 lists the command modes, how to access each command mode, and the prompt that each mode displays.

Table 3: Prompts displayed by the product

Access Level / Operating Mode	Prompt	Means of Entry
User level	`hostname>`	Default at login
System operations mode	`hostname#`	**`enable`** command
Remote-exec	`hostname(remote-exec)#`	**`remote-exec`** command (must be in login or enable modes)
Configuration mode	`hostname(config)#`	**`configure terminal`** command (must have privileged access)
Access-list configuration mode	`hostname(config-acl)#`	**`ip access-list`** command (must be in configuration mode)
Adaptive Routing Engine configuration mode	`hostname(config-engine)#`	**`module engine`** command (must be in configuration mode)
Active-measurement group configuration mode	`hostname(config-engine-active)`	**`active-measurement group`** command (must be in engine configuration mode)
BGP configuration mode	`hostname(config-engine-bgp)#`	**`bgp`** command (must be in engine configuration mode)
Link configuration mode	`hostname(config-engine-link)#`	**`link`** command (must be in engine configuration mode)

1 of 2

Table 3: Prompts displayed by the product (continued)

Access Level / Operating Mode	Prompt	Means of Entry
Policy configuration mode	`hostname(config-engine-policy)`	**`decision-policy`** command (must be in engine configuration mode)
Link policy configuration mode	`hostname(config-engine-policy-link)`	**`link`** command (must be in policy configuration mode)
Load threshold configuration mode	`hostname(config-engine-threshold)`	**`load threshold-table`** command (must be in engine configuration mode)
Interface configuration mode	`hostname(config-if)#`	**`interface`** command (must be in configuration mode)
Polling SNMP configuration mode	`hostname(config-polling-snmp)`	**`polling snmp-agent`** command (must be in configuration mode)
Analysis and Reporting configuration mode	`hostname(config-report)#`	**`module report`** command (must be in configuration mode)
USTAT configuration mode	`hostname(config-ustat)#`	**`module ustat`** command (must be in configuration mode)
Chatter Configuration	*`hostname`* `(config-chatter)#`	**`module chatter`** command (from Configuration mode)
Chatter Zone Configuration	*`hostname`* `(config-chatter-zone)#`	**`zone`** command (from Chatter Configuration mode)
Chatter Active Measurement Configuration	*`hostname`* `(config-chatter-active)#`	**`active-measurement name`** command. (from Chatter Configuration mode)
Chatter Alarm Configuration	*`hostname`* `(config-chatter-alarm)`	**`alarm name`** command (from Chatter Configuration mode)

Use either the **exit** or **end** command to return to the previous mode.

Use the **Control-Z** key combination to return to system operations mode.

For example, if you are in interface configuration mode (*config-if*), do either of the following to return to systems operations mode:

- **Ctrl-Z**
- **exit** to configuration mode and **exit** again (if you use the **exit** command one more time, you will be logged off the device)

Medic Services

Use the **show medic** command to check the status of various services. You can also use the **medic call** command to run certain routines and the **medic set** command to set parameters on some of the services.

Services, routines and parameters are itemized in Table 4.

If a service is identified as being on all modules, it means the service runs as a separate process on each module installed in your system. If you enter a slot number with a medic command, the command will apply to that module. If you enter the keyword **all**, the command will apply to each instance of the process. If you do not enter either a slot number or the keyword **all**, the command will apply to the management module.

If a service is identified as being on peripheral modules, it means the service runs on each module *except* the management module.

If a service is identified as being on the USTAT modules, it means the service runs as a separate process on each USTAT installed in your system:

Table 4: System services, routines, and parameters

Service	Module	Routines	Parameters (defaults)
active-measurement-process	USTATs	restart	--
active-measurement-scheduler	Engine	restart	--
data-reflector	Engine	--	--
decision-maker	Engine	restart	--
event-manager		restart	--

1 of 3

Table 4: System services, routines, and parameters (continued)

Service	Module	Routines	Parameters (defaults)
host-resources	Management	--	max-pcpu[1] (101) min-pcpu (50) max-pmem[2] (80) min-pmem (50) max-prdisk[3] (99) min-prdisk (80) max-phdisk[4] (80) min-phdisk (50)
host-resources	Engine	--	max-pcpu[1] (101) min-pcpu (50) max-pmem[2] (80) min-pmem (50) max-prdisk[3] (99) min-prdisk (80)
host-resources	Reporting	--	max-pcpu[1] (101) min-pcpu (50) max-pmem[2] (80) min-pmem (50) max-prdisk[3] (99) min-prdisk (80)
host-resources	USTATs	--	max-pcpu[1] (101) min-pcpu (50) max-pmem[2] (80) min-pmem (50) max-prdisk[3] (99) min-prdisk (80)
http-logger	USTATs	--	--
link-monitor	Engine	restart	--
logfilterd	All	--	--
management-web-server	Management	--	--
measurement-process	USTATs	restart	--
mount	Peripherals	--	--

Table 4: System services, routines, and parameters (continued)

Service	Module	Routines	Parameters (defaults)
mount-backup	Peripherals	--	--
mount-hd	Peripherals	--	--
object-model-proxy	Management	--	--
object-model-server	Management	--	--
reporting-datastore	Reporting		--
snmpd	All		--
snmptrapd	Management		--
ustat-web-server	USTATs	--	--
verify-flash	Peripherals	--	--

1 max-pcpu and min-pcpu parameters refer to % CPU usage

2 max-pmem and min-pmem parameters refer to % RAM usage

3 max-prdisk and min-prdisk parameters refer to % RAM disk usage

4 max-phdisk and min-phdisk parameters refer to % hard disk usage

3 of 3

Restarting the Decision Maker Process

Certain commands require a restart of the *decision-maker* process before they take effect:

```
bgp
link
module ustat
module engine
provider-as
neighbor allow-next-hop-self
neighbor link
neighbor remote-as
neighbor shutdown
vip
```

After entering any of these commands, including **no** forms, enter the **`restart decision-maker`** command, or you can use the **`medic call`** command:

```
medic call decision-maker restart
```

Command Line Help

Enter the global command **help** or a question mark (“?”) at any time to see a list of commands that are available to you in the current operating mode.

If you enter a command on the command line followed by the **help** or **?** command, a list of keywords and user-defined arguments will be displayed.

Note:

The keyword column of the help text is limited to a maximum of 20 characters. A few keywords exceed this limit and are truncated in the help text. For example, the **min-required-measurements** command is shown as **min-required-measure**. If you enter the keywords and command names on the command line as shown and press the TAB key, the CLI will complete the entry for you.

Legal Characters

- Text that you create as arguments to commands can consist of letters, numbers, hyphens, and the underscore character.
- Names that you create as arguments to commands can consist of letters, numbers, hyphens, and the underscore character, but they must begin with a letter, and not exceed 32 characters in length.
- Domain names and file URLs can contain periods, and URLs can contain a colon, forward slashes, and the @ character.
- Passwords can contain all printable characters. If a password contains a space ( ), backslash (\), question mark (?), or exclamation mark (!), then you must quote the entire password (either single (‘) or double (“) quotes can be used). Note that a quoted password can not use the same quote symbol as part of the password. E.g., “test 99” could not use the double (“) quote symbol as part of the actual password.

The exclamation mark (!) indicates comment text. The CNA software will ignore it and all text that follows until the end of the line. To use an exclamation mark in user-defined text (include description and remark blocks, user IDs and passwords), enclose the text block in double quote marks.

The CNA operating system will convert all alphabetical characters to lower case when interpreting keyword entries (but not when reading user-defined text, such as aliases and group names, which are case sensitive).

For more information about entering text, see Character Set in Chapter 2: Configuration Fundamentals in the *Fundamentals Guide*.

Chapter 2: Commands

This chapter describes all CLI commands for the product. The commands are presented in alphabetical order.

aaa authentication login

Sets the authentication method for users logging in to ANS6.

Syntax

```
aaa authentication login <method> [ <method> ]
no aaa authentication login
```

Arguments

- *method*—must be one of the following:
  - *radius*—authentication handled by RADIUS server
  - *tacacs+*—authentication handled by TACACS+ server (requires PAP, or password authorization protocol, in the server configuration)

Prerequisites

Must be in configuration mode (*config*) that requires privileged-level access.

Occurrence

Use once (but can specify multiple *method* arguments).

Default

The default authentication scheme is *local*.

Description

If multiple *method* values are specified, the device will try the first method in the list. If it fails due to a subnet failure condition, the device will attempt to authenticate by the next means in the list. If the authentication attempt fails for any other reason, the user will be denied access.

The method *local* (authentication handled by local device) is always appended to the list and therefore cannot be entered on the command line.

The **no** form of the command sets the authentication scheme to the default.

active-measurement avoid

Prevents ANS6 from probing a specified subnet.

Syntax

```
active-measurement avoid subnet <addr> [ traceprobe-convergence-only ]
no active-measurement avoid subnet <addr>
```

Arguments

- *addr*—subnet address in CIDR format (*a.b.c.d/n*)

Prerequisites

Must be in engine configuration mode (*config-engine*) that requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

none

Description

The **`active-measurement avoid`** command allows you to specify which addresses you want to prevent ANS6 from probing. Use the keyword **traceprobe-convergence-only** to prevent the address from being used as a convergence point, but still use it as an end point.

Use the **no** form of the command to remove the constraint from a subnet.

active-measurement group

Defines a group of active-measurement targets and sets configuration parameters for them.

Syntax

```
active-measurement group <name>
no active-measurement group <name>
```

Arguments

- *name*—name of the group

Prerequisites

Must be in engine configuration mode (*config-engine*) that requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

none

Description

The **`active-measurement group`** command defines a group of active measurement targets and places ANS6 into active-measurement group configuration mode (*config-engine-active*). In this mode, you may use the following commands (documented as separate command entries):

- description
- diagnose target
- disable
- efc bytes minimum ... efc source
- port
- probe-backoff
- rate
- refresh-time acquire ... refresh-time verify
- target
- threshold host-disappeared ... threshold proxy-link-fail
- timeout
- tos byte
- traceprobe convergence-rule ... traceprobe hops
- type
- use link-group, use named-vip

active-measurement max-probes-per-second

Specifies the maximum number of active measurement probes allowed per second, for all active probes.

Syntax

```
active-measurement max-probes-per-second <probes>
no active-measurement max-probes-per-second
```

Arguments

- *probes*—an integer, 0 or greater

Prerequisites

Must be in engine configuration mode (*config-engine*) that requires privileged-level access.

Occurrence

Use once.

Default

20

Description

The **`active-measurement max-probes-per-second`** command sets an upper limit on the number of probes ANS6 will attempt to send per second. If you have configured more probes than allowed by this limit, ANS6 will adjust the schedules of all probes on a pro-rated basis.

Set probes to 0 to disable all active probing.

The **no** form of the command reverts *probes* to its default value.

See the **`active-measurement target`** command.

active-measurement name

Enters Chatter Active Measurement Configuration command mode. Once in the configuration mode, enter the configuration commands for the specific Active Measurement test.

Syntax

active-measurement name {rtp | ping | tcp | traceroute}

no active-measurement {rtp | ping | tcp | traceroute}

Arguments

Argument	Description
{rtp | ping | tcp | traceroute}	Enter one of the following keywords: • **rtp**—RTP test • **ping**—ping test • **tcp**—TCP connect test • **traceroute**—traceroute test

Prerequisites

Must be in Chatter Configuration command mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrences

Available for use multiple times in the configuration.

Description

Defines a type of measurement test for the Chatter module to conduct.

Currently, Chatter permits one instance of a particular type of measurement test. For example, only one RTP test can be active in the Chatter hive.

The no command deletes the test that you specify.

active-measurement traceprobe avoid address (discontinued)

See active-measurement avoid. Prevents the CNA measuring processes from excessively probing an IP address.

Syntax

```
active-measurement traceprobe avoid address <addr>
no active-measurement traceprobe avoid address <addr>
```

Arguments

- *addr*—IP address (a.b.c.d format) to be avoided

Prerequisites

Must be in engine configuration mode (*config-engine*) that requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

The *traceprobe* type of active probe uses a **traceprobe retries-per-hop** process to get as close as possible to a designated target which is unreachable by conventional probing techniques.

It is expected that the probe will die before reaching the target, either at a firewall or a nearby host. The last interface to respond to the **traceroute** contact is the address that the CNA measurement process will continue to probe, for the duration of the interval specified by the **`traceprobe refresh`** command.

If continual probing of that address is unacceptable, you can flag it with the **`active-measurement traceprobe avoid address`** command, and the CNA measuring processes will not direct probes to it.

This is not a guarantee that the address will never be contacted again. At each refresh interval, the **traceroute** process will be invoked again and will fail again, likely at the same address as before. So, even though you flag an address to be avoided, it may still be contacted periodically. This cannot be prevented. There is no way to configure a **traceroute** packet to avoid intermediate hops.

See also the **`contact`** command.

---

active-measurement-vip (deprecated)

Defines a virtual IP address (VIP) for active measurements.

Syntax

```
active-measurement-vip <address>
no active-measurement-vip
```

Arguments

- *address*—IP address in dotted decimal format (a.b.c.d)

Prerequisites

Must be in USTAT configuration mode (*config-ustat*) that requires privileged-level access.

Occurrence

Use once.

Default

No active-measurement VIP address; active probes use the normal VIP address (see the **vip** command).

Description

The **active-measurement-vip** command has been incorporated into the **vip** command.

Both names are recognized in the current release of CNA, but the former name will be removed in a subsequent release.

If you have existing configurations, you should rename this command.

---

alarm-absolute disable

Disables generation of absolute alarms.

Syntax

alarm-absolute disable

no alarm-absolute disable

Prerequisites

Must be in Chatter Configuration command mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use once.

Default

Enabled.

Description

Disables generation of absolute alarms.

Absolute alarms are generated when a specific test measurement exceeds the configured threshold. To configure the absolute thresholds for specific test measurements, you create and define an alarm template. For information on how to create and define an alarm template, see alarm name.

Absolute alarms are forwarded as SNMP traps and to Syslog servers if configured to do so.

The no command restores the default setting of enabled.

alarm-absolute disable ping

Disables absolute alarms for Ping tests on any cell or edge that is assigned this alarm template.

You can control the generation of absolute alarms for each measurement variable individually or all measurements collectively.

Syntax

alarm-absolute disable ping (all | rtt | loss | inter-arrival-time)

no alarm-absolute disable ping

no alarm-absolute disable ping (all | rtt | loss | inter-arrival-time)

Arguments

Table 5: alarm-absolute disable ping arguments

Argument	Description
all	Disables absolute alarms for all Ping measurements: RTT, loss, and inter-arrival-time.
rtt	Disables absolute alarms for Ping round trip time (RTT).
loss	Disables absolute alarms for Ping loss. The loss threshold is a percentage of packets lost in transit.
inter-arrival-time	Disables absolute alarms for Ping inter-arrival-time. Inter arrival time is the time between tests, and can be used to trigger an alarm when a link is down or a set of endpoints becomes isolated.

Prerequisites

Must be in Chatter Alarm Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use multiple times.

Default

Absolute alarms are enabled for all Ping measurements.

Description

Disables absolute alarms for Ping tests on any cell or edge that is assigned this alarm template.

You can control the generation of absolute alarms for each measurement variable individually or all measurements collectively.

The no command restores the default value of enabled for all Ping measurements or the specified measurement variable.

---

alarm-absolute disable rtp

Disables absolute alarms for RTP tests on any cell or edge that is assigned this alarm template.

You can control the generation of absolute alarms for each measurement variable individually or all measurements collectively.

Syntax

alarm-absolute disable rtp (all | delay | jitter | loss | inter-arrival-time)

no alarm-absolute disable rtp

no alarm-absolute disable rtp (all | delay | jitter | loss | inter-arrival-time)

Arguments

Table 6: alarm-absolute disable rtp arguments

Argument	Description
all	Disables absolute alarms for all RTP measurements: delay, jitter, loss, and inter-arrival-time.
delay	Disables absolute alarms for RTP delay. Delay is how long (in milliseconds one-way) it takes the UDP packet stream to reach its destination.
jitter	Disables absolute alarms for RTP jitter. Jitter is the unevenness in packet delay (in milliseconds).
loss	Disables absolute alarms for RTP loss. The loss threshold is a percentage of packets lost in transit.
inter-arrival-time	Disables absolute alarms for RTP inter-arrival-time. Inter arrival time is the time between tests, and can be used to trigger an alarm when a link is down or a set of endpoints becomes isolated.

Prerequisites

Must be in Chatter Alarm Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Available for use multiple times in the configuration.

Default

Absolute alarms are enabled for all RTP measurements.

Description

Disables absolute alarms for RTP tests on any cell or edge that is assigned this alarm template.

You can control the generation of absolute alarms for each measurement variable individually or all measurements collectively.

The no command restores the default value of enabled for all RTP measurements or the specified measurement variable.

alarm-absolute disable tcp

Disables absolute alarms for TCP tests on any cell or edge that is assigned this alarm template.

You can control the generation of absolute alarms for each measurement variable individually or all measurements collectively.

Syntax

alarm-absolute disable tcp (all | delay | cerror | inter-arrival-time)

no alarm-absolute disable tcp

no alarm-absolute disable tcp (all | delay | cerror | inter-arrival-time)

Arguments

Table 7: alarm-absolute disable tcp arguments

Argument	Description
all	Disables absolute alarms for all TCP measurements: delay, cerror, and inter-arrival-time.
delay	Disables absolute alarms for TCP delay. Delay is how long (in milliseconds) the connection takes to complete.
cerror	Disables absolute alarms for TCP cerror. Cerror is the success or failure of the connection.
inter-arrival-time	Disables absolute alarms for TCP inter-arrival-time. Inter arrival time is the time between tests, and can be used to trigger an alarm when a link is down or a set of endpoints becomes isolated.

Prerequisites

Must be in Chatter Alarm Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Available for use multiple times in the configuration.

Default

Absolute alarms are enabled for all TCP measurements.

Description

Disables absolute alarms for TCP tests on any cell or edge that is assigned this alarm template.

You can control the generation of absolute alarms for each measurement variable individually or all measurements collectively.

The no command restores the default value of enabled for all TCP measurements or the specified measurement variable.

alarm-absolute disable traceroute

Disables absolute alarms for Traceroute tests on any cell or edge that is assigned this alarm template.

You can control the generation of absolute alarms for each measurement variable individually or all measurements collectively.

Syntax

alarm-absolute disable traceroute (all | hop-rtt | inter-arrival-time)

no alarm-absolute disable traceroute

no alarm-absolute disable traceroute (all | hop-rtt | inter-arrival-time)

Arguments

Table 8: alarm-absolute disable traceroute arguments

Argument	Description
all	Disables absolute alarms for all Traceroute measurements: hop-rtt and inter-arrival-time.

1 of 2

Table 8: alarm-absolute disable traceroute arguments (continued)

Argument	Description
hop-rtt <milliseconds>	Disables absolute alarms for Traceroute hop round trip time (RTT). Hop RTT is the round trip time for each hop along the route (UDP outbound and ICMP inbound).
inter-arrival-time <milliseconds>	Disables absolute alarms for Traceroute inter-arrival-time. Inter arrival time is the time between tests, and can be used to trigger an alarm when a link is down or a set of endpoints becomes isolated.
	2 of 2

Prerequisites

Must be in Chatter Alarm Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Available for use multiple times in the configuration.

Default

Absolute alarms are enabled for all TCP measurements.

Description

Disables absolute alarms for Traceroute tests on any cell or edge that is assigned this alarm template.

You can control the generation of absolute alarms for each measurement variable individually or all measurements collectively.

The no command restores the default value of enabled for all Traceroute measurements or the specified measurement variable.

alarm name

Creates a user-defined alarm template. Once the alarm template is created, you define it by assigning it a set of alarm thresholds and associating it with specific zone and edge objects.

This command also enters Chatter Alarm Configuration command mode.

Syntax

alarm name <name>

no alarm name <name>

Arguments

Table 9: alarm name arguments

Argument	Description
<name>	Name of the alarm template you are creating or configuring. Each alarm template must have a unique name.

Prerequisites

Must be in Chatter Configuration command mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Available for use multiple times in the configuration.

Description

Creates a user-defined alarm template. Once the alarm template is created, you define it by assigning it a set of alarm thresholds and associating it with specific zone and edge objects.

This command also enters Chatter Alarm Configuration command mode.

The no command deletes the specified alarm template.

alarm-relative disable

Disables generation of relative alarms.

Syntax

alarm-relative disable

no alarm-relative disable

Prerequisites

Must be in Chatter Configuration command mode.For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use once.

Default

Enabled.

Description

Disables generation of relative alarms.

Relative alarms are generated when the most recent results for a specific test show an abrupt increase or decrease. The thresholds for determining an “abrupt increase or decrease” are self-adjusting thresholds based on a statistical technique. Note that if the measurement remains near or outside the threshold, the threshold adjusts to reflect this changed norm. Once the threshold self-adjusts, the alarm clears even though the measurement itself has not changed. The steady state becomes the threshold value - this value can be higher or lower than configured threshold. Relative alarms are forwarded as SNMP traps and to Syslog servers if configured to do so.

The no command restores the default setting of enabled.

---

alarm-relative disable ping

Disables relative alarms for Ping tests on any cell or edge that is assigned this alarm template.

You can control the generation of relative alarms for each measurement variable individually or all measurements collectively.

Syntax

alarm-relative disable ping (all | rtt | loss | inter-arrival-time)

no alarm-relative disable ping

no alarm-relative disable ping (all | rtt | loss | inter-arrival-time)

Arguments

Table 10: alarm-relative disable ping arguments

Argument	Description
all	Disables relative alarms for all Ping measurements: RTT, loss, and inter-arrival-time.
rtt	Disables relative alarms for Ping round trip time (RTT).
loss	Disables relative alarms for Ping loss. The loss threshold is a percentage of packets lost in transit.
inter-arrival-time	Disables relative alarms for Ping inter-arrival-time. Inter arrival time is the time between tests, and can be used to trigger an alarm when a link is down or a set of endpoints becomes isolated.

Prerequisites

Must be in Chatter Alarm Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Available for use multiple times in the configuration.

Default

Relative alarms are enabled for all Ping measurements.

Description

Disables relative alarms for Ping tests on any cell or edge that is assigned this alarm template.

You can control the generation of relative alarms for each measurement variable individually or all measurements collectively.

The no command restores the default value of enabled for all Ping measurements or the specified measurement variable.

alarm-relative disable rtp

Disables relative alarms for RTP tests on any cell or edge that is assigned this alarm template.

You can control the generation of relative alarms for each measurement variable individually or all measurements collectively.

Syntax

alarm-relative disable rtp (all | delay | jitter | loss | inter-arrival-time)

no alarm-relative disable rtp

no alarm-relative disable rtp (all | delay | jitter | loss| inter-arrival-time)

Arguments

Table 11: alarm-relative disable rtp arguments

Argument	Description
all	Disables relative alarms for all RTP measurements: delay, jitter, loss, and inter-arrival-time.
delay	Disables relative alarms for RTP delay. Delay is how long (in milliseconds one-way) it takes the UDP packet stream to reach its destination.
jitter	Disables relative alarms for RTP jitter. Jitter is the unevenness in packet delay (in milliseconds).
loss	Disables relative alarms for RTP loss. The loss threshold is a percentage of packets lost in transit.
inter-arrival-time	Disables relative alarms for RTP inter-arrival-time. Inter arrival time is the time between tests, and can be used to trigger an alarm when a link is down or a set of endpoints becomes isolated.

Prerequisites

Must be in Chatter Alarm Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Available for use multiple times in the configuration.

Default

Relative alarms are enabled for all RTP measurements.

Description

Disables relative alarms for RTP tests on any cell or edge that is assigned this alarm template.

You can control the generation of relative alarms for each measurement variable individually or all measurements collectively.

The no command restores the default value of enabled for all RTP measurements or the specified measurement variable.

alarm-relative disable tcp

Disables relative alarms for TCP tests on any cell or edge that is assigned this alarm template.

You can control the generation of relative alarms for each measurement variable individually or all measurements collectively.

Syntax

alarm-relative disable tcp (all | delay | cerror | inter-arrival-time)

no alarm-relative disable tcp

no alarm-relative disable tcp (all | delay | cerror | inter-arrival-time)

Arguments

Table 12: alarm-relative disable tcp arguments

Argument	Description
all	Disables relative alarms for all TCP measurements: delay, cerror, and inter-arrival-time.
delay	Disables relative alarms for TCP delay. Delay is how long (in milliseconds) the connection takes to complete.

1 of 2

Table 12: alarm-relative disable tcp arguments (continued)

Argument	Description
cerror	Disables relative alarms for TCP cerror. Cerror is the success or failure of the connection.
inter-arrival-time	Disables relative alarms for TCP inter-arrival-time. Inter arrival time is the time between tests, and can be used to trigger an alarm when a link is down or a set of endpoints becomes isolated.
	2 of 2

Prerequisites

Must be in Chatter Alarm Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Available for use multiple times in the configuration.

Default

Relative alarms are enabled for all TCP measurements.

Description

Disables relative alarms for TCP tests on any cell or edge that is assigned this alarm template.

You can control the generation of relative alarms for each measurement variable individually or all measurements collectively.

The no command restores the default value of enabled for all TCP measurements or the specified measurement variable.

alarm-relative disable traceroute

Disables relative alarms for Traceroute tests on any cell or edge that is assigned this alarm template.

You can control the generation of relative alarms for each measurement variable individually or all measurements collectively.

Syntax

alarm-relative disable traceroute (all | hop-rtt | inter-arrival-time)

no alarm-relative disable traceroute

no alarm-relative disable traceroute (all | hop-rtt | inter-arrival-time)

Arguments

Table 13: alarm-relative disable traceroute arguments

Argument	Description
all	Disables relative alarms for all Traceroute measurements: hop-rtt and inter-arrival-time.
hop-rtt <milliseconds>	Disables relative alarms for Traceroute hop round trip time (RTT). Hop RTT is the round trip time for each hop along the route (UDP outbound and ICMP inbound).
inter-arrival-time <milliseconds>	Disables relative alarms for Traceroute inter-arrival-time. Inter arrival time is the time between tests, and can be used to trigger an alarm when a link is down or a set of endpoints becomes isolated.

Prerequisites

Must be in Chatter Alarm Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Available for use multiple times in the configuration.

Default

Relative alarms are enabled for all TCP measurements.

Description

Disables relative alarms for Traceroute tests on any cell or edge that is assigned this alarm template.

You can control the generation of relative alarms for each measurement variable individually or all measurements collectively.

The no command restores the default value of enabled for all Traceroute measurements or the specified measurement variable.

alarm-threshold-exceeded

Configures the number of times that an agent-based measurement threshold must be exceeded for an alarm to be generated.

Syntax

alarm-threshold-exceeded <count>

no alarm-threshold-exceeded

Arguments

Table 14: alarm-threshold-exceeded arguments

Argument	Description
<count>	Number of times the threshold must be exceeded for an alarm to be generated. The valid range is 1–255.

Prerequisites

Must be in Chatter Configuration command mode.For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use once.

Default

10 times.

Description

Configures the number of times that an agent-based measurement threshold must be exceeded for an alarm to be generated. This command applies to:

- Absolute alarms

  Absolute alarms are generated when tolerance thresholds are exceeded the number of times that you specify in this command.

- Relative alarms

  Relative alarms are generated when the most recent results for a specific test show an abrupt increase or decrease. The thresholds for determining an “abrupt increase or

decrease” are self-adjusting thresholds based on a statistical technique. Note that if the measurement remains near or outside the threshold, the threshold adjusts to reflect this changed norm. Once the threshold self-adjusts, the alarm clears even though the measurement itself has not changed. The steady state becomes the threshold value. This value can be higher or lower than the configured threshold.

The no command restores the default value of 10.

Note that you must balance the test rate with alarm sensitivity to meet your needs. For example, if the test rate is once per minute, and the zone hierarchy has 60 cells, all cells are tested once an hour. If you set the system to generate an alarm when the threshold is exceeded 10 times, the system will take 10 hours to generate an alarm for a cell.

To prevent a few spurious points from setting off an alarm, Avaya recommends that you set the system to generate an alarm when the threshold is exceeded no less than five times. Generally an alarm should indicate a persistent condition.

application

Associates an application with a group.

Syntax

```
application <name>

no application <name>
```

Arguments

- *name*—name of an application

Prerequisites

Must be in group configuration mode (*config-group*) that requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

none

Description

Use the **application** command to associate subnets with a group. When in group configuration mode (see the **group** command), you may add one application at a time to group. See also the **prefix (deprecated)** command.

Use the **no** form of the command to remove an application from a group.

---

application age-out

Sets the time required to remove the associates of an application with a subnet.

Syntax

```
application age-out <time>

no application age-out
```

Arguments

- *time*—time in seconds; range: 300 through 86400 (one day)

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*) that requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

3600

Description

Use the **application age-out** command to set the time, in seconds, required to elapse for the association of an application with a subnet to expire. The association is reestablished whenever the EFC modules detects the application on that subnet.

Use the **no** form of the command to remove an application from a group.

application detection

Explicitly sets which applications are to be detected by the EFC module.

Syntax

```
application detection
          ( name <app-name> ) | ( group <group> )

no application detection
no application detection name <app-name>
no application detection group <group>
```

Arguments

- *name*—name of an application (e.g., kazaa)
- *group*—name of an application group (see the **group** command)

Prerequisites

Must be in EFC configuration mode (*config-efc*) that requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

none

Description

Use the **application detection** command to enable the detection of the specified application (or group of applications) by the EFC module. To create new applications use the command: **user-defined-application**.

Use the **no** form of the command to stop detection of the specified application (or an application group).

application no-detection

Explicitly sets which applications are to be detected by the EFC module.

Syntax

```
application no-detection predefined | discovered | user-defined |
          ( name <app-name> ) | ( group <group> )

no application no-detection predefined
no application no-detection discovered
no application no-detection user-defined
no application no-detection name <app-name>
no application no-detection group <group>
```

Arguments

- *name*—name of an application (e.g., kazaa)
- *group*—name of an application group (see the **group** command)

Prerequisites

Must be in EFC configuration mode (*config-efc*) that requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

none

Description

Use the **application no-detection** command to disable the detection of the specified application, group, or category of applications (predefined, discovered, and user-defined) by the EFC module. For a quick way to see what is in the categories, use the following commands:

- show application predefined traffic-rate
- show application discovered traffic-rate
- show application user-defined traffic-rate

See also the command: **user-defined-application**.

Use the **no** form of the command to ignore no-detection rules.

application-model enterprise

Sets the characteristics that affect the behavior of the enterprise application model, which optimizes time-critical TCP transactions on your subnet, such as transactional web, file transfer, email, and gaming.

Syntax

```
application-model enterprise
           { ( application-delay metric-contribution <mcontr> ) |
             ( clip { half-life <hlife> |
                      metric-contribution <mcontr> |
                      threshold <threshold>
                    }
             ) |
             ( delay { half-life <hlife> } ) |
             ( loss { descent-rate <meas> <target-loss> } |
                        half-life <hlife> |
                    }
             ) |
             ( metric half-life <hlife> )
             }

no application-model enterprise
no application-model enterprise application-delay [ metric-contribution ]
no application-model enterprise clip [ half-life | metric-contribution
                                       threshold ]
no application-model enterprise delay [ half-life ]
no application-model enterprise loss [ descent-rate | half-life ]
no application-model enterprise metric [ half-life ]
```

Arguments

- *hlife*—(exact meaning depends upon which of the following keywords is used)
  - **clip half-life**: exponential decay of the effects of a clip; range: 0 through 15000 (default 20)
  - **delay half-life**: smoothing factor of average delay; range: 0 through 15000 (default 150)
  - **loss half-life**: smoothing factor of measured loss; range: 0 through 15000 (default 300)
  - **metric half-life**: internal parameter - do not change; range: 0 through 150000 (default 1500)
- *mcontr*—(exact meaning depends upon which of the following keywords is used)
  - **application-delay metric-contribution**: the weight given to delay, in milliseconds; range: 0 through 5000 (default is 1500)
  - **clip metric-contribution**: the weight given to a clip event; range: 0 through 1000 (default is 750)
- *threshold*—the number of consecutively lost packets that would be considered a clip event; range: 0 through 100 (default 4)
- *meas*—internal parameter - do not change; range: 0 through 5000 (default 250)
- *target-loss*—internal parameter - do not change; range: 1 through 100 (default 1)

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*) that requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **`application-model enterprise`** command to change the default behavior of this model to suit the particular needs of your application. You may create as many variants as you wish by assigning each to a specific decision policy. See Details on Application Models in Chapter 6: Decision Policies and Application Models of the *Monitoring Guide*.

Note:

When using the **enterprise** application model, you must use either the traceprobe or ICMP measurement modes. If you are able to employ only TCP, then you should use the **web** or **other** application model. The reason for this is because the **enterprise** model requires measurements of the loss characteristics in your subnet to a higher degree of precision than is possible with TCP.

The following keywords control the major components of the enterprise application model:

- **application-delay** - represents the average delay of a transaction and takes into account an average of three TCP turns for a single transaction. The metric-contribution affects the weighting of this delay in the decision-maker's calculations. Its value determines the mid-point of a curve that tells decision-maker how strictly it should regard application delay in its decisions. A higher number for this value will place greater importance on delay (i.e., an application is considered to be impacted more by a delay).
- **clip** - defines the impact and severity of the loss of consecutive packets. half-life sets the exponential decay of the effects of a clip on the decision-maker's calculations. A high number will cause decisions to be influenced by a clip event for a longer period. A low number will limit the effects of clip events to a shorter period, which can cause transitions to be less smooth (e.g., the decision-maker may change certain decisions more frequently). metric-contribution affects the weighting of a clip event's effect on the decision-maker's calculations. A higher value for this parameter will cause the decision-maker to treat a clip event more significantly. The threshold parameter sets the number of consecutively lost packets that determine a clip event.
- **delay** - determines how the moving average of delay is calculated. The half-life parameter controls the smoothness of the moving average: A high value results in smaller changes in the calculated average delay because the delay contribution of past measurements will be greater than when half-life is set low.
- **loss descent-rate** - this is an internal parameter, do not change it

- **loss** - controls the smoothing factor of measured loss. A high value results in a smaller contribution of past loss events; conversely, a low value will allow past loss events to have a greater influence on the decision-maker's calculations.
- **metric** - this is an internal parameter, do not change it

Use the **no** form of the command to revert to the default settings.

application-model multimedia

Sets the characteristics that affect the behavior of the **multimedia** application model, which optimizes real-time applications on your subnet. Examples include voice over IP, streaming audio, and video conferencing applications.

Syntax

```
application-model multimedia
             { ( clip { half-life <hlife> |
                      metric-contribution <mcontr> |
                      threshold <threshold>
                      }
               ) |
               ( delay { half-life <hlife> |
                      metric-contribution <mcontr> |
                      }
               ) |
               ( jitter min-buffer <mbuff> ) |
               ( loss { descent-rate <meas> <target-loss> } |
                        half-life <hlife> |
                        metric-contribution <mcontr>
                      }
               )
               ( metric half-life <hlife> )
             }
no application-model multimedia
no application-model multimedia clip [ half-life | metric-contribution
                                       threshold ]
no application-model multimedia delay [ half-life | metric-contribution ]
no application-model multimedia jitter [ min-buff ]
no application-model multimedia loss [ descent-rate | half-life |
                                       metric-contribution ]
no application-model multimedia metric [ half-life ]
```

Arguments

- *hlife*—(exact meaning depends upon which of the following keywords is used)
  - **clip half-life**: **exponential decay of the effects of a clip**; range: 0 through 15000 (default 20)
  - **delay half-life**: smoothing factor of average delay; range: 0 through 15000 (default 150)
  - **loss half-life**: smoothing factor of measured loss; range: 0 through 15000 (default 150)
  - **metric half-life**: internal parameter - do not change; range: 0 through 150000 (default 1500)
- *mcontr*—(exact meaning depends upon which of the following keywords is used)
  - **clip metric-contribution**: range: 0 through 1000 (default is 500)
  - **delay metric-contribution**: range: 0 through 1000 (default is 800)
  - **loss metric-contribution**: internal parameter - do not change; range: 0 through 1000 (default is 10)
- *threshold*—the number of consecutively lost packets that would be considered a clip event; range: 0 through 100 (default 2)
- *mbuff*—amount of buffer delay introduced by a client; range: 0 through 1000 (default 80)
- *meas*—internal parameter - do not change; range: 0 through 5000 (default 250)
- *target-loss*—internal parameter - do not change; range: 1 through 100 (default 1)

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **`application-model multimedia`** command to change the default behavior of this model to suit the particular needs of your subnet. You may create as many variants as you wish by assigning each to a specific decision policy. See Details on Application Models in Chapter 6: Decision Policies and Application Models of the *Monitoring Guide*.

Note:

When using the **multimedia** application model, you must use either the traceprobe or ICMP measurement modes. If you are able to employ only TCP, then you should use the **web** or **other** application model.

The following keywords control the major components of the multimedia application model:

- **clip** - defines the impact and severity of the loss of consecutive packets. half-life sets the exponential decay of the effects of a clip on the decision-maker's calculations. A high number will cause decisions to be influenced by a clip event for a longer period. A low number will limit the effects of clip events to a shorter period, which can cause transitions to be less smooth (e.g., the decision-maker may change certain decisions more frequently). metric-contribution affects the weighting of a clip event's effect on the decision-maker's calculations. A higher value for this parameter will cause the decision-maker to treat a clip event more significantly. The threshold parameter sets the number of consecutively lost packets that determine a clip event.
- **delay** - determines how the moving average of delay is calculated. The half-life parameter controls the smoothness of the moving average: A high value results in smaller changes in the calculated average delay because the delay contribution of past measurements will be greater than when half-life is set low.
- **jitter min-buffer** - sets the amount of buffer delay introduced by a client application (e.g., voice over IP, video conferencing), in milliseconds
- **loss descent-rate** - this is an internal parameter, do not change it
- **loss half-life** - controls the smoothing factor of measured loss. A high value results in a smaller contribution of past loss events; conversely, a low value will allow past loss events to have a greater influence on the decision-maker's calculations.
- **metric** - this is an internal parameter, do not change it

Use the **no** form of the command to revert to the default settings.

application-model other

Sets the characteristics that affect the behavior of the **other** application model, which optimizes web transactions on your subnet.

Syntax

```
application-model other
          {  ( delay { half-life <hlife> } ) |
             ( hrtt { half-life <hlife> |
                      metric-contribution <metric> |
                      threshold <threshold>
                    }
             ) |
             ( loss { descent-rate <meas> <target-loss> } |
                        half-life <hlife> |
                    }
             )
          }
```

```
no application-model other
no application-model other delay [ half-life ]
no application-model other hrtt [ half-life | metric-contribution | threshold ]
no application-model other loss [ descent-rate | half-life ]
```

Arguments

- *hlife*—(exact meaning depends upon which of the following keywords is used)
  - **delay half-life**: smoothing factor of average delay; range: 0 through 15000 (default 150)
  - **hrtt half-life**: smoothing factor of average handshake round trip time (hrtt); range: 0 through 15000 (default 50)
  - **loss half-life**: smoothing factor of measured loss; range: 0 through 15000 (default 300)
- *mcontr*—range: 0 through 5000 (default 915)
- *meas*—internal parameter - do not change; range: 0 through 5000 (default 250)
- *target-loss*—internal parameter - do not change; range: 1 through 100 (default 1)

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **`application-model other`** command to change the default behavior of this model to suit the particular needs of your subnet. You may create as many variants as you wish by assigning each to a specific decision policy. See Details on Application Models in Chapter 6: Decision Policies and Application Models of the *Monitoring Guide*.

Note:

The **other** application model is designed for TCP and user traffic test (UTT) measurements (i.e., single pixel gif). You can use ICMP and traceprobe measurements with the web model; however, the enterprise model has been designed to use the separate delay and loss signals from these measurement types. We recommend the **other** model if the majority of your measurements are TCP or UTT, and the enterprise model if you primarily use ICMP or traceprobe.

The following keywords control the major components of the other application model:

- **delay** - determines how the moving average of delay is calculated. The half-life parameter controls the smoothness of the delay moving average: A high value results in smaller changes in the calculated average delay because the delay contribution of past measurements will be greater than when half-life is set low. This parameter will not affect decision-maker's calculations, but is used only by the reporting module.
- **hrtt** - determines how handshake round trip time will influence decision-maker's calculations. The metric-contribution affects the weighting of the measured hrtt. Its value determines the mid-point of a curve that tells decision-maker how strictly it should regard measured hrtt in its decisions. A higher number for this value will place greater importance on it (i.e., an application is considered to be impacted more by longer handshake round trip times).
- **loss descent-rate** - this is an internal parameter, do not change it
- **loss half-life** - controls the smoothing factor of measured loss. A high value results in a smaller contribution of past loss events. This parameter will not affect decision-maker's calculations, and is used only by the reporting module.

Use the **no** form of the command to revert to the default settings.

application-model streaming

Sets the characteristics that affect the behavior of the **streaming** application model, which optimizes real-time applications on your subnet. Examples include voice over IP, streaming audio, and video conferencing applications.

Syntax

```
application-model streaming
             { ( clip { half-life <hlife> |
                      metric-contribution <mcontr> |
                      threshold <threshold>
                      }
               ) |
               ( delay { half-life <hlife> |
                      metric-contribution <mcontr> |
                      }
               ) |
               ( jitter min-buffer <mbuff> ) |
               ( loss { descent-rate <meas> <target-loss> } |
                        half-life <hlife> |
                        metric-contribution <mcontr>
                      }
               )
               ( metric half-life <hlife> )
             }
```

```
no application-model streaming
no application-model streaming clip [ half-life | metric-contribution
                                      threshold ]
no application-model streaming delay [ half-life | metric-contribution ]
no application-model streaming jitter [ min-buff ]
no application-model streaming loss [ descent-rate | half-life |
                                      metric-contribution ]
no application-model streaming metric [ half-life ]
```

Arguments

- *hlife*—(exact meaning depends upon which of the following keywords is used)
  - **clip half-life**: **exponential decay of the effects of a clip**; range: 0 through 15000 (default 20)
  - **delay half-life**: smoothing factor of average delay; range: 0 through 15000 (default 150)
  - **loss half-life**: smoothing factor of measured loss; range: 0 through 15000 (default 150)
  - **metric half-life**: internal parameter - do not change; range: 0 through 150000 (default 1500)
- *mcontr*—(exact meaning depends upon which of the following keywords is used)
  - **clip metric-contribution**: range: 0 through 1000 (default is 500)
  - **delay metric-contribution**: range: 0 through 1000 (default is 800)
  - **loss metric-contribution**: internal parameter - do not change; range: 0 through 1000 (default is 10)
- *threshold*—the number of consecutively lost packets that would be considered a clip event; range: 0 through 100 (default 2)
- *mbuff*—amount of buffer delay introduced by a client; range: 0 through 1000 (default 80)
- *meas*—internal parameter - do not change; range: 0 through 5000 (default 250)
- *target-loss*—internal parameter - do not change; range: 1 through 100 (default 1)

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **`application-model streaming`** command to change the default behavior of this model to suit the particular needs of your subnet. You may create as many variants as you wish by assigning each to a specific decision policy. See Details on Application Models in Chapter 6: Decision Policies and Application Models of the *Monitoring Guide*.

Note:

When using the **streaming** application model, you must use either the traceprobe or ICMP measurement modes. If you are able to employ only TCP, then you should use the **web** application model.

The following keywords control the major components of the streaming application model:

- **clip** - defines the impact and severity of the loss of consecutive packets. half-life sets the exponential decay of the effects of a clip on the decision-maker's calculations. A high number will cause decisions to be influenced by a clip event for a longer period. A low number will limit the effects of clip events to a shorter period, which can cause transitions to be less smooth (e.g., the decision-maker may change certain decisions more frequently). metric-contribution affects the weighting of a clip event's effect on the decision-maker's calculations. A higher value for this parameter will cause the decision-maker to treat a clip event more significantly. The threshold parameter sets the number of consecutively lost packets that determine a clip event.
- **delay** - determines how the moving average of delay is calculated. The half-life parameter controls the smoothness of the moving average: A high value results in smaller changes in the calculated average delay because the delay contribution of past measurements will be greater than when half-life is set low.
- **jitter min-buffer** - sets the amount of buffer delay introduced by a client application (e.g., voice over IP, video conferencing), in milliseconds
- **loss descent-rate** - this is an internal parameter, do not change it
- **loss half-life** - controls the smoothing factor of measured loss. A high value results in a smaller contribution of past loss events; conversely, a low value will allow past loss events to have a greater influence on the decision-maker's calculations.
- **metric** - this is an internal parameter, do not change it

Use the **no** form of the command to revert to the default settings.

application-model voice

Sets the characteristics that affect the behavior of the **voice** application model, which optimizes voice over IP traffic on your subnet.

Syntax

```
application-model voice
             { ( clip { half-life <hlife> |
                      metric-contribution <mcontr> |
                      threshold <threshold>
                      }
               ) |
               ( delay { half-life <hlife> |
                      metric-contribution <mcontr> |
                      }
               ) |
               ( jitter min-buffer <mbuff> ) |
               ( loss { descent-rate <meas> <target-loss> } |
                        half-life <hlife> |
                        metric-contribution <mcontr>
                      }
               )
               ( metric half-life <hlife> )
             }
no application-model voice
no application-model voice clip [ half-life | metric-contribution
                                    threshold ]
no application-model voice delay [ half-life | metric-contribution ]
no application-model voice jitter [ min-buff ]
no application-model multimvoiceedia loss [ descent-rate | half-life |
                                    metric-contribution ]
no application-model voice metric [ half-life ]
```

Arguments

- *hlife*—(exact meaning depends upon which of the following keywords is used)
  - **clip half-life**: **exponential decay of the effects of a clip**; range: 0 through 15000 (default 20)
  - **delay half-life**: smoothing factor of average delay; range: 0 through 15000 (default 150)
  - **loss half-life**: smoothing factor of measured loss; range: 0 through 15000 (default 150)
  - **metric half-life**: internal parameter - do not change; range: 0 through 150000 (default 1500)
- *mcontr*—(exact meaning depends upon which of the following keywords is used)
  - **clip metric-contribution**: range: 0 through 1000 (default is 500)
  - **delay metric-contribution**: range: 0 through 1000 (default is 800)
  - **loss metric-contribution**: internal parameter - do not change; range: 0 through 1000 (default is 10)
- *threshold*—the number of consecutively lost packets that would be considered a clip event; range: 0 through 100 (default 2)
- *mbuff*—amount of buffer delay introduced by a client; range: 0 through 1000 (default 80)
- *meas*—internal parameter - do not change; range: 0 through 5000 (default 250)
- *target-loss*—internal parameter - do not change; range: 1 through 100 (default 1)

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **`application-model voice`** command to change the default behavior of this model to suit the particular needs of your subnet. You may create as many variants as you wish by assigning each to a specific decision policy. See Details on Application Models in Chapter 6: Decision Policies and Application Models of the *Monitoring Guide*.

Note:

When using the **voice** application model, you must use either the traceprobe or ICMP measurement modes. If you are able to employ only TCP, then you should use the **web** application model.

The following keywords control the major components of the voice application model:

- **clip** - defines the impact and severity of the loss of consecutive packets. half-life sets the exponential decay of the effects of a clip on the decision-maker's calculations. A high number will cause decisions to be influenced by a clip event for a longer period. A low number will limit the effects of clip events to a shorter period, which can cause transitions to be less smooth (e.g., the decision-maker may change certain decisions more frequently). metric-contribution affects the weighting of a clip event's effect on the decision-maker's calculations. A higher value for this parameter will cause the decision-maker to treat a clip event more significantly. The threshold parameter sets the number of consecutively lost packets that determine a clip event.
- **delay** - determines how the moving average of delay is calculated. The half-life parameter controls the smoothness of the moving average: A high value results in smaller changes in the calculated average delay because the delay contribution of past measurements will be greater than when half-life is set low.
- **jitter min-buffer** - sets the amount of buffer delay introduced by a client application (e.g., voice over IP, video conferencing), in milliseconds
- **loss descent-rate** - this is an internal parameter, do not change it
- **loss half-life** - controls the smoothing factor of measured loss. A high value results in a smaller contribution of past loss events; conversely, a low value will allow past loss events to have a greater influence on the decision-maker's calculations.
- **metric** - this is an internal parameter, do not change it

Use the **no** form of the command to revert to the default settings.

application-model web

Sets the characteristics that affect the behavior of the **web** application model, which optimizes web transactions on your subnet.

Syntax

```
application-model web
          {  ( delay { half-life <hlife> } ) |
             ( hrtt { half-life <hlife> |
                      metric-contribution <metric> |
                      threshold <threshold>
                    }
             ) |
             ( loss { descent-rate <meas> <target-loss> } |
                        half-life <hlife> |
                    }
             )
          }

no application-model web
no application-model web delay [ half-life ]
no application-model web hrtt  [ half-life | metric-contribution | threshold ]
no application-model web loss  [ descent-rate | half-life ]
```

Arguments

- *hlife*—(exact meaning depends upon which of the following keywords is used)
  - **delay half-life**: smoothing factor of average delay; range: 0 through 15000 (default 150)
  - **hrtt half-life**: smoothing factor of average handshake round trip time (hrtt); range: 0 through 15000 (default 50)
  - **loss half-life**: smoothing factor of measured loss; range: 0 through 15000 (default 300)
- *mcontr*—range: 0 through 5000 (default 915)
- *meas*—internal parameter - do not change; range: 0 through 5000 (default 250)
- *target-loss*—internal parameter - do not change; range: 1 through 100 (default 1)

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **`application-model web`** command to change the default behavior of this model to suit the particular needs of your subnet. You may create as many variants as you wish by assigning each to a specific decision policy. See Details on Application Models in Chapter 6: Decision Policies and Application Models of the *Monitoring Guide*.

Note:

The web application model is designed for TCP and user traffic test (UTT) measurements (i.e., single pixel gif). You can use ICMP and traceprobe measurements with the web model; however, the enterprise model has been designed to use the separate delay and loss signals from these measurement types. We recommend the web model if the majority of your measurements are TCP or UTT, and the enterprise model if you primarily use ICMP or traceprobe.

The following keywords control the major components of the web application model:

- **delay** - determines how the moving average of delay is calculated. The half-life parameter controls the smoothness of the delay moving average: A high value results in smaller changes in the calculated average delay because the delay contribution of past measurements will be greater than when half-life is set low. This parameter will not affect decision-maker's calculations, but is used only by the reporting module.
- **hrtt** - determines how handshake round trip time will influence decision-maker's calculations. The metric-contribution affects the weighting of the measured hrtt. Its value determines the mid-point of a curve that tells decision-maker how strictly it should regard measured hrtt in its decisions. A higher number for this value will place greater importance on it (i.e., an application is considered to be impacted more by longer handshake round trip times).
- **loss descent-rate** - this is an internal parameter, do not change it
- **loss half-life** - controls the smoothing factor of measured loss. A high value results in a smaller contribution of past loss events. This parameter will not affect decision-maker's calculations, and is used only by the reporting module.

Use the **no** form of the command to revert to the default settings.

---

application priority

Sets the user preference for the ordering of application data in live trending reports.

Syntax

```
application priority
          ( name <app-name> ) | ( group <group-name> )

no application priority name <app-name>
no application priority group <group-name>
```

Arguments

- *name*—name of an application (e.g., kazaa)
- *group*—name of an application group (see the **group** command)

Prerequisites

Must be in report configuration mode (*config-report*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **application priority** command to ensure that the desired applications appear on the list for live trending. This is useful because of the system-imposed maximum of 100 live-trended applications; when there are more applications than this number, then the excess will not be live-trended.

Use the **no** form of the command to remove an application (or application group) from the priority list.

---

as-path

Configures the values to be used when a synthetic AS path is required.

Syntax

```
as-path {
          explicit <list> |
          synthetic <number> <length> |
          any-from-bgp
        }
no as-path
```

Arguments

- *list*—a list of specific AS numbers; can be:
  - “” for an empty list
  - list of space-separated integers in the range 1 through 65535, and may include square brackets, angle brackets or parenthesis—[ ], < > and ( )—plus a comma, for defining an AS set, a confed set, or a confed sequence; entire list must be enclosed in quotes unless list consists solely of a single integer
- *number*—Integer in the range 1 through 65535; AS number to be used as padding
- *length*—Integer in the range 1 through 30; number of pads to add to AS path

Prerequisites

Must be in engine or policy configuration mode (*config-eng* or *config-eng-policy*), which requires privileged-level access.

Default

synthetic

- *number*—AS number of the service provider represented by the link for which the AS path is being constructed
- *length*—number needed to make the AS path equal to the shortest BGP route to the subnet

Description

When CNA asserts a route to an edge router for a subnet, it includes the AS path that was received when the subnet was announced by the BGP peer.

However, CNA does not always have AS path data for every subnet through every link, such as after a system restart, or for subnets associated with a VPN which CNA knows about only because you added them to the configuration through the CLI. This no-available-path condition can also occur in a service provider environment where CNA is part of a mesh that includes upstream and downstream peerings.

When asserting such routes, CNA needs to construct an AS path using the values specified by the **`as-path`** command.

Use the **explicit** keyword and *list* argument if you know the precise AS path that should be used with the route. The AS path you specify will always be used. If a BGP-generated path exists, it will be overwritten.

Use the **synthetic** keyword with the *number* and *length* arguments if you want CNA to construct an AS path according to the following formula:

```
AS_path = link_AS + [ defined_AS [ defined_AS [ ... ] ] ] + origin_AS
```

...where:

- *link_AS* is the service provider's AS number (as set by the **`provider-as`** command when the link was configured)
- *origin_AS* is the AS number in which the subnet is located.

If *link_AS* is equal to *origin_AS*, the synthetic AS path will have a length of one and the AS number will be set equal to the *origin_AS* value.

If *link_AS* and *origin_AS* are different, the number of instances of *defined_AS* will be equal to the value of the *length* argument.

If you do not specify a value for the *length* argument, CNA will set it to whatever value is required to make the AS path equal in length to the shortest path in all BGP-provided routes that have been obtained for the subnet since the most recent restart of the *decision-maker* process.

For example, if the shortest BGP path has a length of 2, then the value of the *length* argument will be 0 and the synthetic AS path will be set to a length of 2; if the shortest BGP path has a length of 3, then the value of the *length* argument would be 1 and the synthetic AS path will be set to a length of 3.

There is an exception, however: if BGP's shortest AS path length for the subnet is 1 but CNA has determined that the BGP path is not the best performance route, then the synthetic AS path will *not* be set to the length of the shortest path. Instead, it will consist of *link_AS* and *origin_AS* and have a length of two.

The default length is intended to make CNA-asserted routes at least as attractive as the best BGP selection (other than the case where a BGP route has an AS path length of one and CNA asserts a different service provider).

This is useful when it is desired that the route propagates inside your subnet or to downstream networks.

Note:

However, such behavior is potentially dangerous. If the CNA-asserted route were to leak beyond your subnet boundaries into the Internet at large, unwanted traffic could be attracted to your subnet instead of the subnet in which the subnet is located. If there is a chance that asserted routes could leak into the Internet beyond your private subnet—if, for example, you have elected to not propagate the **no-export** community attribute that CNA attaches to all of its asserted routes—you may want to set the *number* value higher to make asserted routes unattractive to outsiders.

Use the **any-from-bgp** keyword if you want CNA to use the AS path it has learned from any BGP peer. If there are multiple BGP-generated AS paths available because of multiple peerings, the CNA *decision-maker* process will arbitrarily choose the first AS path it finds in the CNA database. If no BGP-generated AS path is available, the route will not be asserted.

Use the **no** form of the command, which takes no arguments, to reset the AS path construction method to its default.

as-path-padding (discontinued)

The **`as-path-padding`** command has been deprecated beginning in version 2.2, and removed in version 3.0. Use the **`as-path`** command instead.

assert-routes-matching-bgp

Forces CNA to assert a route even though it is also a BGP selected route.

Syntax

```
assert-routes-matching-bgp
no assert-routes-matching-bgp
```

Prerequisites

Must be in system-operations mode, with privileged-level access.

Default

Off

Description

When integrating CNA into some highly customized subnet topologies, it may be advantageous to force a route assertion on a router even though BGP has already selected the route for a given subnet. (This will have at least one major consequence: CNA will not be able to recognize subsequent changes in the natural BGP selection.)

If you are sure this behavior is correct for your subnet, use the **`assert-routes-matching-bgp`** command to enable *assert-matching-routes* mode.

Use the **no** form of the command to disable *assert-matching-routes* mode.

When this mode is disabled (the default condition), CNA will not assert a performance route to a router when the CNA choice and the BGP selected route for the subnet are the same.

Note:

If route assertion is not enabled (see the **`route-assert-mode`** command), the **`assert-routes-matching-bgp`** command will not have any affect on CNA operation.

asserted-routes-timeout (discontinued)

Sets maximum length of time performance routes will remain asserted for subnets that have no measurable traffic.

Syntax

```
asserted-routes-timeout <time>
no asserted-routes-timeout
```

Arguments

- *time*—number of seconds, in the range 3600 through 86400

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Default

86400 seconds (24 hours)

Description

The intended behavior of the asserted-routes-timeout command is now achieved by the **`outage-detection total-silence`** command.

---

auto-create-subnets

Enables the automatic create of subnets.

Syntax

```
auto-create-subnets { enable | disable }
no auto-create-subnets
```

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Description

Use this command when you want CNA to automatically create subnets based on observed addresses, either from active-measurements (see active-measurement group) or single-pixel.gif measurements (see ip http server on).

---

auto-create-subnets limit

Defines the maximum number of subnets to be automatically created (see auto-create-subnets).

Syntax

```
auto-create-subnets mask-length limit <number>
no auto-create-subnets mask-length limit
```

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Arguments

number—Integer in the range 1 through 65536

Default

1000

Description

Use this command to limit the number of subnets that can be automatically created by CNA. Setting it to a high limit may result in CPU or memory resource exhaustion.

---

auto-create-subnets mask-length

Defines the mask length of subnets that are automatically created.

Syntax

```
auto-create-subnets mask-length absolute <length>
no auto-create-subnets mask-length
```

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Arguments

length—Integer in the range 1 through 32; size of the mask length

Default

24

Description

Use this command to set the subnet mask length of the subnets that are to be automatically created. For example, if this mask-length is set to the default of 24 and CNA sees the address 100.1.2.3, then it will create the subnet 100.1.2.0/24.

auto-create-subnets verify-frequency

Defines the frequency at which automatically created subnets will be validated (see auto-create-subnets).

Syntax

```
auto-create-subnets verify-frequency <seconds>
no auto-create-subnets verify-frequency
```

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Arguments

- *seconds*—interval to wait, in seconds, before each validation run; range: 300 through 86400 (5 min - 1 day)

Default

3600

Description

Use this command to set the interval between each validation run. Automatically created subnets are removed from the system when they are determined by CNA not to be currently in use. A lower value for this interval means that the validation check will be done more often, which has the drawback of consuming more CPU resources, so be careful when you use this command.

backup

Creates a backup copy of the specified object.

Syntax

```
backup {image | image-config} [notes <text>] [<url>]
```

Arguments

- *text*—user-defined notes
- *url*—destination file location

Prerequisites

Must be in system-operations mode, with privileged-level access.

Description

The backup command can be used to create archival copies of the CNA image and configurations.

Use the keyword **image** to back up the CNA image.

Use the keyword **image-config** to back up the entire image plus startup configuration and, if they exist, SSL certification and SSH key files.

The keyword **data**, allowed in previous releases, has been discontinued.

Use the keyword **notes**, followed by a string of text, to add notes about the image being backed up. Enclose your text in double quotes if the text includes spaces. Your text will be displayed when you execute the **`show image`** command.

The backup file will contain a complete snapshot of the entire CNA operating system. If you do not have sufficient space on either the target disk (if backing up to an external location) or the internal hard disk, the system will advise you of the amount you need to clear.

Note:

Image files which you create using the **`backup`** command will be significantly smaller—by as much as 50 MB—than upgrade images obtained from the Avaya FTP server. These upgrade images contain installation-related files (the Java plug-in installers, for example) that are not included in a backup file.

Use the *url* argument to specify a target location either local to the device, or to a remote FTP or HTTP location. Data backups must be stored in a remote location. See the **`copy`** command for URL syntax.

The *url* argument is optional with the **image** and **image-config** keywords. If you do not specify a target file, the system will name the file for you.

See also **`ls`**, **`rename`**, and **`restore`** commands.

bgp

Places CNA into BGP configuration mode, so that you can configure peerings.

Syntax

```
bgp <num>
no bgp <num>
```

Arguments

- *as-num*—integer in the range 0 through 65,535 representing an Autonomous System (AS) number (64,512 through 65,535 are private AS numbers)

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

The bgp command places CNA in BGP configuration mode (*config-eng-bgp*). The *num* argument should be your subnet's AS number.

While in *config-engine-bgp* mode, you can enter any of the following commands to configure BGP peerings:

- **neighbor allow-next-hop-self (removed)**
- **neighbor buffered (removed)**
- **neighbor description**
- **neighbor link**
- **neighbor redirect-peering**
- **neighbor remote-as**
- **neighbor report-alias**
- **neighbor shutdown**
- **neighbor summarize-bgp-events**
- **neighbor update-source**

Use the **no** form of the command to remove the entire block of BGP commands from the CNA configuration.

bgp-community

Defines a BGP community string to be attached to selected route assertions.

Syntax

```
bgp-community all <list>
bgp-community assert if-winner <list>
bgp-community assert if-in-winner-set <list>
bgp-community assert if-not-in-winner-set <list>
bgp-community match (invalid-winner | force-winner) community <A:B>

no bgp-community all
no bgp-community assert if-winner
no bgp-community assert if-in-winner-set
no bgp-community assert if-not-in-winner-set
no bgp-community match invalid-winner
no bgp-community match force-winner
```

Arguments

- *list*—value of BGP community; can be one of the following:
  - one or more community numbers, each in *a:b* format where *a* and *b* are integers in the range 1 through 65534 separated by a colon
  - “”—two double quotes signifying an empty string; blocks inheritance of a BGP community value, resulting in no BGP community string
  - one of the well-known values defined in RFC 1997, in the following forms:
    - *no-export*
    - *no-advertise*
    - *no-export-subconfed*
- *A:B*—a community number, in *a:b* format where *a* and *b* are integers in the range 1 through 65534 separated by a colon

Prerequisites

For **all**, must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*).

For **assert**, must be in link mode at either the engine or policy level (*config-engine-link* or *config-engine-policy-link*).

For **match**, must be in link mode at the engine level (*config-engine-link*).

All modes require privileged-level access.

Occurrence

Use a command/keyword combination once per configuration mode (within each scope that will accept the command, there can be only one BGP community defined for a route).

Default

None, except that, within a decision policy, if no values are explicitly set the globally set values (if any) will apply

Description

The **`bgp-community`** command allows you to associate BGP community values with selected routes.

The values included in *list* will be applied to all routes that meet the criteria defined by the **keyword** used:

- **all**—values will be associated with all routes asserted by CNA (including BGP selected routes, if you have enabled **assert-routes-matching-bgp**), regardless of link
- **assert if-winner**—values will be associated with routes only when the specified link is the winner as determined by the CNA *decision-maker* process
- **assert if-in-winner-set**—values will be associated with routes only when the specified link is in the winner-set
- **assert if-not-in-winner-set**—values will be associate with routes when the specified link is not in the winner set
- **match**—when used with **invalid-winner**, will cause the link to be marked as invalid if the community matches a natural route

Setting a community value with the **`bgp-community all`** command will overwrite the default behavior, which is to associate the string *no-export* with all routes. If you want this behavior to remain in effect, you must explicitly include it in your *list* argument, at either the global level or within a decision policy.

If you override the *no-export* value at the global level, this behavior will be inherited by decision policies. If you want it added back for a specific policy, you must explicitly add it within that policy's configuration.

You can use the **`bgp-community all`** command without arguments to remove the no-export value from route assertions.

If a condition has been set at both global and policy levels, the more specific condition will apply. For example, if you specify a community value with the keyword **all** at both levels, the string specified in the policy level will override the global string.

If you set *list* equal to a null string (“”), no BGP community value will be associated with the subnet at this configuration scope. The empty string blocks inheritance of the value from the previous level in the hierarchy.

Each BGP community list can have multiple components. Separate each value by a space character and enclose the entire list in double quotes:

```
bgp-community all “65001:1 65002:2 65003:3”
```

The **bgp-community** command should be used in place of the **outbound-bgp-community (discontinued)** command, which will not be supported in future CNA releases.

Use the **no** form of the command to remove the specified form of the command from the scope in which it is entered.

In the following configuration example, the BGP community value is set to 100:00 for the link named *link_A* if it is the winning link selected by CNA and to 1000:1000 if it is not in the winner set.:

```
module engine 1
   link link_A
     bgp-community assert if-winner 100:100
     bgp-community assert if-not-in-winner-set 1000:1000
   end
```

You can change these values for selected subnets by creating a decision policy. For example, you can change the BGP community value to 1:1 when *link_A* is the winning link for a selected subnet and to nothing when it is not in the winner set for that subnet, by adding the following commands to your engine module configuration:

```
   decision-policy rule_A
     link link_A
       bgp-community assert if-winner 1:1
       bgp-community assert if-not-in-winner-set ""
     end
   end
   set-decision-policy 192.168.1.0/24
 end
```

In the example, the **bgp-community assert if-not-in-winner-set** command in the decision policy blocks inheritance of the **if-not-in-winner-set** value that was set at the *config-engine* level.

billing link

Creates a billing link configuration block to create a single billing link from multiple engine links.

Syntax

```
billing link <name>
no billing link <name>
```

Arguments

- *name*—name of the billing link

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Description

Use the **`billing link`** command to combine multiple engine links (see link.) into a single billing link. See usage-reporting billing-period, usage-reporting formula, usage-reporting cost, usage-reporting percentile.)

Use the **no** form of the command to remove the billing link.

---

chatter assign

Clears and recalculates the associations of cells and edges to Chatter modules.

Syntax

chatter assign

Prerequisites

Must be in Privileged command mode.For information on how to access this mode, see Accessing and exiting command modes.

Description

Clears and recalculates the associations of cells and edges to Chatter modules.

Note:

If this command changes the association of a cell or edge from one Chatter module to another, previously trended data for the cell or edge is lost. New data for the cell or edge is maintained by the Chatter module that the cell or edge is now associated with.

chatter join-hive

Adds the local Chatter module (on which you enter this command) to a remote Chatter module's hive.

Syntax

chatter join-hive ip-address <ip> master-key <key>

Arguments

Table 15: chatter join-hive arguments

Argument	Description
<ip>	IP address of the remote Chatter module whose hive you want the local Chatter to join.
<key>	Master key of the remote Chatter module whose hive you want the local Chatter to join. To obtain the master key of the remote Chatter module, enter the show chatter master-key command on the remote chatter module.

Prerequisites

Must be in Privileged command mode.For information on how to access this mode, see Accessing and exiting command modes.

Description

Adds the local Chatter module (on which you enter this command) to a remote Chatter module's hive.

To enter this command, you need the IP address and master key of the remote Chatter module. To obtain the master key of the remote Chatter module, enter the show chatter master-key command on the remote Chatter module.

chatter leave-hive

Removes the Chatter module from its current hive.

Syntax

chatter leave-hive

Prerequisites

Must be in Privileged command mode.For information on how to access this mode, see Accessing and exiting command modes.

Description

Removes the Chatter module from its current hive.

chatter missed-heartbeats

Defines the number of missed heartbeats for a Chatter module to be recognized as down by the hive.

Syntax

chatter missed-heartbeats <*value*>

no chatter missed-heartbeats

Arguments

Argument	Description
<*value*>	Number of missed heartbeats before the Chatter is considered down. The valid range is 1 to 255.

Prerequisites

Must be in Chatter Configuration command mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrences

Use once.

Default

10 missed heartbeats.

Description

Defines the number of missed heartbeats for a Chatter module to be recognized as down by the hive.

The Chatter modules in a hive send heartbeat messages between themselves to ensure that all Chatter modules are functional and performing the tests assigned to them. If the number of heartbeats specified in this command is missed from a Chatter module, it is recognized as down, and its tasks are redistributed between the remaining Chatter modules in the hive.

The default setting is 10.

The no command restores the default value.

chatter node add

Merges up to six hops in to a single router node.

Syntax

chatter node add <ip-list>

Arguments

Table 16: chatter node add arguments

Argument	Description
<ip-list>	List of IP addresses (hops) that you want merged to a router node. You can enter from 2 to 6 IP addresses. You must enter at least 2 IP addresses. IP addresses must be separated by a space.

Prerequisites

Must be in Privileged command mode.For information on how to access this mode, see Accessing and exiting command modes.

Description

Merges up to six hops in to a single router node.

Chatter automatically identifies and merges IP addresses that belong to a single router node. However, this functionality can fail for various reasons, and a node with multiple IP addresses will be displayed as more than one node. The **chatter node add** command makes it possible for you to manually merge intermediate hops into a node.

chatter node split

Splits a router node in to hops.

Syntax

chatter node split <ip-address>

Arguments

Table 17: chatter node split arguments

Argument	Description
<ip-address>	IP address of one of the hops in the node that you want split.

Prerequisites

Must be in Privileged command mode.For information on how to access this mode, see Accessing and exiting command modes.

Description

Splits a router node in to hops.

Chatter automatically identifies and merges IP addresses that belong to a single router node. However, this functionality can fail for various reasons. The **chatter node split** command makes it possible for you to manually split a router node in to intermediate hops.

Important:
Split nodes that were originally merged by Chatter may be remerged if Chatter determines that the IP addresses are the same device. The **chatter node split** command does not exclude the split node from Chatter's automatic merge functionality.

chatter remove-cna-agent

Manually removes a CNA agent from the agent-based measurement system.

Syntax

chatter remove-cna-agent [ip-address <ip> | not-assigned | all]

Arguments

Table 18: chatter remove-cna-agent arguments

Argument	Description
<ip>	IP address of the specific CNA agent to remove.
not-assigned	Removes all CNA agents that are not currently assigned to a zone.
all	Removes all CNA agents.

Prerequisites

Must be in Privileged command mode.For information on how to access this mode, see Accessing and exiting command modes.

Description

Once a CNA agent registers with a hive, it is not removed from the hive's database unless you do so manually. This functionality prevents the hive from losing track of agents that may be temporarily unavailable. If you are sure that you no longer want to use a specific CNA agent, manually remove it.

cidr-lookup

Enable substitution of subnets with the organization (when present in the CNA cidr lookup database.

Syntax

```
cidr-lookup enable
no cidr-lookup enable
```

Prerequisites

System operations mode

Description

Use the **cidr-lookup** command to cause CNA to use an organization's name, when it is present in the database, instead of a subnet (in a.b.c.d/n notation). The substitution will be applied to all reports.

Use the **no** form of the command to disable substitution (i.e., leave the subnet in a.b.c.d/n notation).

---

clear active-measurements

Clears some or all statistical data from the master table of active measurement statistics maintained by the RSOS software.

Syntax

```
clear active-measurements [ [ group <name>] | no-group]
                           [ target <probe> ]
                           [ address <addr> ]
                           [ range-address <range> ]
                           [ type { tcp | icmp | traceprobe } ]
                           [ ustat <uname> ]
                           [ successful ]
                           [ errors <list> ]
                           [ probing <address> ]
                           [ range-probing <range> ]
                           [ subnet <subnet> ]
                           [ subnet-probing <subnet> ]
```

Arguments

- *name*—name of an active measurement group
- *probe*—IP address or domain name of an active measurement probe, as entered in the configuration
- *addr*—IP address of probe, as resolved by DNS
- *range*—range in CIDR format (a.b.c.d /*n*)
- *uname*—USTAT module name
- *list*—error codes; can be any combination of the following:
  - *d*—DNS
  - *h*—Host unreachable
  - *n*—subnet unreachable
  - *p*—Port unreachable
  - *q*—Queue drop

  - *t*—Probe timeout
  - *o*—Other
- *subnet*—subnet address, in CIDR notation (a.b.c.d /*n*)
- *subnet-probing*—subnet address, in CIDR notation (a.b.c.d /*n*)

Prerequisites

System operations mode, with privileged-level access.

Description

The **`clear active-measurements`** command manages the tables displayed by the **`show active-measurements`** command.

CNA maintains a table of statistics for all currently configured active probes. When you execute the **`show active-measurements`** command, CNA locates and displays all rows of the master table that match the specified arguments.

Use the **`clear active-measurements`** command to clear selected rows in the master table.

Use the **range-address** keyword to clear addresses within a desired range.

Use the **range-probing** keyword to clear addresses within a desired traceprobing range.

Use the **no-group** keyword to clear data for all active measurements that are not associated with a group.

To clear data for all active measurements, do not specify either a group name or the **no-group** keyword.

---

clear arp

Clears the ARP table entry for a specified host on a specified CNA module.

Syntax

```
clear arp <address> [ <slot> ]
```

Arguments

- *address*—host name or IP address to be removed from the ARP table
- *slot*—slot number of module on which the ARP table resides

Prerequisites

System operations mode, privileged-level access.

Default

If *slot* not defined, defaults to the CNA management module.

clear efc active-measurement

Clears EFC data in the active-measurement table that is maintained by the CNA *efcserver* process.

Syntax

```
clear efc active-measurement { endpoint <address> | group <name> | all }
```

Arguments

- *address*—endpoint IP address (*a.b.c.d*)
- *name*—name of an active measurement group

Prerequisites

Must be in system operations mode, with privileged-level access.

Description

Use the **`clear efc active-measurement`** command to clear the measurement data for the specified endpoint, **group** of endpoints, or **all** endpoints, that are maintained in the EFC active-measurement table.

See the **`show efc`** and **`clear efc prefix (deprecated)`** commands.

See Chapter 2: Collecting Endpoint Addresses of the *Monitoring Guide*.

clear efc prefix (deprecated)

The clear efc prefix command has been deprecated. See the replacement command clear efc subnet on page 90.

clear counters web-hits

Clears the count of web-hits from the **show link** command output.

Syntax

```
clear counters web-hits
```

Prerequisites

Must be in system operations mode, with privileged-level access.

Description

The **clear counters web-hits** command clears the user-traffic-test (UTT) counters used in the in the output of the **show link** command.

The **clear counters web-hits** command clears the counter used by the **show link** command, while leaving the other counter (which is used for, among other things, the **Measurements** graph in the web-interface **Control Panel**, and for trend report data) untouched as long as the CNA *decision-maker* process continues to run.

clear prefix (deprecated)

The clear prefix command has been deprecated. See the replacement command clear subnet on page 90.

clear snapshots (discontinued)

Clears the snapshot database maintained by CNA.

Syntax

```
clear snapshots [ force ]
```

Prerequisites

System operations mode, privileged-level access.

Occurrence

Use once.

Description

The clear snapshots command has been discontinued in the 2.0 release. Use the **delete report** command instead.

---

clear statistics

Clears CNA counters for statistics which are visible in various show command output displays.

Syntax

```
clear statistics <process> [ <number> | all ] [ avoid | diagnose ]
```

Arguments

- *process*—name of the internal process; can be one of:
  - *am-scheduler*—active-measurement process on the engine module
  - *am-measurer*—active-measurement process on a USTAT module
  - *dmaker*—decision-maker process on the engine module
  - *dm-listener*—decision-maker listener process on the engine module
  - *edgestats*—user-traffic-test process on a USTAT module
  - *event-manager*—Event Manager process on the management module
  - *link-monitor*—SNMP polling process
  - *logging*—logging process on any module
  - *monitor*—monitor (SPAN listener) process on the EFC module
- *number*—slot or subsystem number for which statistics are to be cleared

Prerequisites

System operations mode, privileged-level access.

Description

The **clear statistics** command resets counters to 0 for statistics maintained by various CNA processes.

The keyword **diagnose** is valid only for the *am-measurer* and *am-scheduler* processes. Use it to clear all diagnostic statistics.

The keyword **avoid** is valid only for the *am-measurer* process. Use it to clear all entries in the probe avoidance table.

clear efc subnet

Clears EFC data in the subnet table maintained by the CNA *efcserver* process.

Syntax

```
clear efc subnet { <address> | all }
```

Arguments

- *address*—subnet address in CIDR format (*a.b.c.d/n*)

Prerequisites

System operations mode, log-in level access.

Description

Use the **`clear efc subnet`** command to clear the data that is maintained in the EFC subnet table.

Specify a subnet address in CIDR format (*a.b.c.d/n*) to clear data for that subnet, or use the keyword **all** to clear the entire table.

See the **`show efc`** and **`clear efc active-measurement`** commands.

See Chapter 2: Collecting Endpoint Addresses of the *Monitoring Guide*.

clear subnet

Causes CNA to withdraw best-path performance routes that have been asserted to routers.

Syntax

```
clear subnet { <address> | all | using-policy <name> }
             [ routes-only |
               measurement-policy |
               remove-auto-created ]
```

Arguments

- *address*—subnet address, in CIDR notation (a.b.c.d/*n*)
- *name*—name of policy

Prerequisites

System operations mode, privileged-level access.

Description

If the **all** keyword is used, all performance routes for all subnets will be withdrawn.

If you specify a subnet, performance routes for that subnet will be withdrawn.

If the **using-policy** keyword is used, only those subnets that are associated with a particular policy will be withdrawn.

If the **measurement-policy** keyword is used, subnets for which the decision maker process has automatically assigned a decision policy will have such policies removed.

If the **routes-only** keyword is specified, the subnet is withdrawn from the BGP tables but the counters which CNA uses to keep track of user-traffic test hits for that subnet are not reset.

If you want to start measuring traffic to the specified subnet as though your service provider links were new, allow the system to clear the traffic measurement counters.

If you want to remove a route temporarily but you want to maintain its history, use the **routes-only** keyword and keep the traffic measurement counters at their current values.

---

clear trending

Removes specified objects from the trending database.

Syntax

```
clear trending { [ link-usage { all [ force ] | link <lname> ] |
                 [ live-trending-subnet ] |
                 [ performance
                         { all [ force ] |
                           customer <name> |
                           link <lname> |
                           subnet <subnet> |
                           router <rname> |
                           summary
                         } ]
               }
```

Arguments

- *lname*—link name, as string (see **link** command)
- *name*—name of customer
- *subnet*—subnet in CIDR format (a.b.c.d /*n*)
- *rname*—IP address or alias of a router

Prerequisites

System operations mode, privileged-level access.

Occurrence

Use once.

Description

Use the performance keyword to clear data related to performance routing.

Use the link-usage keyword to clear data related to link bandwidth usage (*link-usage* data is obtained by polling routers to monitor individual link usage).

Use the performance keyword to refresh the list of subnets to be trended live.

Specify either all data, with the **all** keyword, or only the data associated with a specific subnet, router, or link *name* arguments.

Use the optional **force** keyword with the **all** keyword to cause the database to be cleared immediately; if the **force** keyword is not used, CNA will prompt the user before continuing.

See the **trending on** and **clear trending report** commands.

---

clear trending report

Removes all references to the specified object from all existing trend-report query files.

Syntax

```
clear trending report { subnet | router | link } <name>
```

Arguments

- *name*—can be one of:
  - subnet alias or address
  - router hostname or address
  - link name

Prerequisites

System operations mode, privileged-level access.

Description

When you change router names or addresses, or remove a link from a configuration, use the **`clear trending report`** command to purge references to the removed item from your trend-graph query files. You can also use this command to clear references to subnets for which trend data is no longer available, though this is expected to be a less common occurrence.

---

clear tty

Clears a specified tty session.

Syntax

```
clear tty <session>
```

Arguments

- *session*—session identifier

Prerequisites

System operations mode, privileged level access

Description

Use the **`clear tty`** command to terminate a remote session.

Use the **`show users`** command to display a list of all currently active sessions. The list will identify logged in users by user name, followed by a session identifier (such as *ttyp1*). Use the session identifier as the argument to the **`clear tty`** command.

---

cna-agent

Assigns a CNA agent to a zone.

Syntax

cna-agent <*ip-address*>

no cna-agent <*ip-address*>

Arguments

Table 19: cna-agent command arguments

Argument	Description
<ip-address>	IP address of the CNA agent. The format of the IP address is a.b.c.d.

Prerequisites

Must be in Chatter Zone Configuration command mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrences

Available for use multiple times in the configuration.

Description

Assigns a CNA agent to a zone or subzone. The no form of this command removes the CNA agent from the zone.

---

community

Provides an edge router's SNMP read-only community string to the CNA USTAT module that is associated with one of the router's service provider links.

Syntax

```
community <string>
no community
```

Arguments

- *string*—edge router's read-only (RO) community string

Prerequisites

Must be in polling-agent configuration mode (*config-polling-snmp*), which requires privileged-level access.

Description

The **`community`** command provides CNA with the router's read-only community string, which the USTAT module will use to periodically poll the router for usage data.

See the **`polling usage on`** command for details about polling.

Use the **no** form of the command to clear the community string from CNA.

---

configure memory

Resets the running configuration back to the startup configuration.

Syntax

```
configure memory
```

Prerequisites

System operations mode, privileged-level access.

---

configure subnet

Resets the running configuration from a configuration file stored somewhere on the subnet.

Syntax

```
configure subnet <url>
```

Arguments

- *url*—URL of configuration file

Prerequisites

Systems operation mode, privileged-level access.

Description

You can use FTP or HTTP to obtain a new configuration file, by specifying the appropriate URL.

If authentication is required to retrieve the file specified by *url*, you can either specify user name and password individually on the command line, or you can include them as part of the URL.

To specify authentication data on the command line, use the **ip ftp username** and **ip ftp password** commands or the **ip http username** and **ip http password** commands prior to executing the **configure subnet** command.

See the description of the **copy** command for information about how to specify authentication data as part of the URL.

configure terminal

Places CNA in configuration mode, allowing you to modify the running configuration from the command line.

Syntax

```
configure terminal
```

Prerequisites

System operations mode, privileged-level access.

Description

End configuration mode by entering the **end** or **exit** command, or Control-Z on the command line.

While you are in configuration mode, commands entered on the command line will be written to the running configuration.

contact

Customizes the identification string that is added to all **traceroute** UDP packets used for active measurements of type *traceprobe*.

Syntax

```
contact <text>
no contact
```

Arguments

- *text*—your data, up to 10 characters which can consist of letters, numbers, space, hyphen, period, comma, colon, underscore, parenthesis, forward slash, and “commercial at” - @.

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

none

Description

The CNA active-measurement system will embed identifying data in each **traceroute** UDP packet sent in support of an active measurement probe of type *traceprobe*.

You can change the value of part of this data to contain your own information (such as a name and e-mail address), by changing the value of *text*. The string “PathContrl” will also be sent with your text, along with host and domain information, if you have configured these. It will be in the form:

```
PathContrl <host>.<domain> <text>
```

If the data contains spaces, enclose the entire string in quote marks.

---

content

Customizes the content to be served by the USTAT web server.

Syntax

```
content <file>
no content
```

Arguments

- *file*—the file name of your content, which must end in .gif

Prerequisites

Must be in configuration mode (*ustat-server-map*), which requires privileged-level access.

Occurrence

Used once per map.

Default

none

Description

When the USTAT web server sees a request that matches the **location** string for this *ustat-server-map* configuration, then instead of the normal single-pixel .gif, it will serve the .gif specified by **content**.

Use the **no** form of the command to stop serving this content.

---

copy

Copies files from one location to another.

Syntax

```
copy <source> <dest>
copy /<data-type> rsync:[//<user>:<password>@<host>:<port>/<dir>]
copy rsync:[//<user>:<password>@<host>:<port>/<dir>] /<data-type>
```

Arguments

- *source*—can be one of the following:
  - URL of source file
  - *running-config*
  - *startup-config*
  - Name of a file in the CNA repository
- *dest*—can be one of the following:
  - URL of destination file
  - *running-config*
  - *startup-config*
  - Name of a file in the CNA repository
- *data-type*—type of data being replicated (must be *trends*)
- *user*—user name for logging on to remote **rsync** host
- *password*—password required for remote **rsync** host
- *host*—IP address or hostname of remote **rsync** host
- *port*—access port on remote **rsync** host
- *dir*—**rsync** root directory on remote **rsync** host

Prerequisites

System operations mode, privileged-level access.

Description—File Copy

Files are specified as URLs, in one of the following forms:

```
ftp://[<username>:<password>@]<host>[:<port>]/<directory>/<file>
http://[<username>:<password>@]<host>[:<port>]/<directory>/<file>
```

For ftp or http authentication, you can specify user name and password as part of the URL, as in the following examples:

```
ftp://user-1:password@my-ftp-host/directory-name/my-file-name
http://user-1:password@my-http-host/directory-name/my-file-name
```

Alternatively, you can specify them on the command line prior to executing the **copy** command. Use the **ip ftp username** and **ip ftp password** or **ip http username** and **ip http password** commands.

If username and password values are not specified, CNA will still attempt to copy the files from the subnet server.

If authentication is required and none is specified, the **copy** command will fail.

If the destination is specified as *running-config*, the commands in the source file will be merged with the currently running configuration.

If the destination is specified as *startup-config*, the current startup configuration will be copied to a backup location before the new one is written. The backup copy of the startup configuration can be obtained by executing the **recover startup-config** command.

You cannot copy an empty file; if the file has no contents, the command will fail.

When you copy a file containing multiple commands to the running configuration, the CLI will generate a status report indicating the number of commands contained in the file, the number of commands that were successfully applied, and the number and name of any commands that could not be applied because of some configuration conflict.

Description—rsync

The **copy** command can be used to force a one-time update of a remote repository of user data.

You can also use the **copy** command to restore the local user data repository from the remote site.

customer-prefix-file (deprecated)

The customer-prefix-file command has been deprecated. See the replacement command customer-subnet-file on page 100.

customer-subnet-file

This command enables the customer reports feature in CNA.

Syntax

```
customer-subnet-file on
no customer-subnet-file
```

Prerequisites

Must be in report configuration mode (*config-report*), which requires privileged-level access.

Occurrence

Use once.

Default

disabled

Description

Before you turn on the customer reports feature, you need to create a file containing customer-specific subnets. There must be only one subnet per line, and each entry in the file must use the “slash” notation in the form a.b.c.d/m. For example:

172.20.15.0/24

172.20.10.0/24

172.20.4.0/24

172.20.8.0/24

172.20.9.0/24

Assuming you have created the file on a host accessible to CNA, the only thing left to do is copy it to the user-data area (see *copy)*.

damped-mode

Select a new winner for a subnet only if the current winner is no longer in the winner set.

Syntax

```
damped-mode enable | disable }
no damped-mode
```

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once.

Default

Enabled

Description

Use the **damped-mode enable** command to tell the CNA decision maker process to select a new winner for a subnet only if the current winner is no longer in the winner set. This means that if a preferred link moves into the winner set, it will not be selected unless the current winner has moved out. Use the **disable** keyword to suspend this behavior without removing the command from the configuration.

Use the **no** form of the command to revert to default behavior.

data-source efc enable

This command enables the load optimization feature in CNA.

Syntax

```
data-source efc enable
no data-source efc enable
```

Prerequisites

Must be in configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Use once.

Default

disabled

Description

To use load optimization in its simplest form, you need only issue this command and then invoke restart decision-maker. See the commands, **`data-source efc action`** and **`data-source efc sampling-period`** for more information.

---

data-source efc action

Selects the type load optimization performed by CNA.

Syntax

```
data-source efc action [ simulate-measurements | compute-winner-set ]
no data-source efc action
```

Note:

Starting with CNA version 5, the keyword **simulate-measurements** has been deprecated and will no longer have any effect.

Prerequisites

Must be in configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Use once.

Default

simulate-measurements

Description

Use this command to change the way CNA will make load optimization decisions. The **simulate-measurements** keyword forces the CNA decision-maker process to treat all links as having the same measurement characteristics - a metric value of 999 is always applied to any subnets that are observed by the EFC module. This has the effect of providing a simple load optimization feature without requiring any measurements (note: all other types of

measurements, e.g. active probes, will be discarded). To take effect, you will have to restart decision-maker.

data-source efc sampling-period

Specifies how often EFC packets will be sent to the decision-maker process.

Syntax

```
data-source efc sampling period <period>
no data-source efc sampling period
```

Arguments

- *period*—the number of packets from a source address that should be dropped before sending another to the decision-maker; range: 1 to 4294967295

Prerequisites

Must be in configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Use once.

Default

1

Description

Use this command to down-sample each of the addresses sent by the EFC module to the decision-maker process by 1/<period>. For example, if you set the period to 5, then only one out of every five packets received from a particular address will be sent to the decision-maker. This can be useful if you expect a very high traffic load on the EFC module, which would in turn flood the decision-maker with more data than is useful.

date

Set or display the date.

Syntax

```
date [<value>]
```

Arguments

- *value*—date, in *mmddyyyy* format (use leading zeros for single-digit months and days; use four digits for year)

Prerequisites

System operations mode. Login-level access can view date; must have privileged-level access to set date.

Description

The date value must consist of eight digits; use leading zeros for single-digit months and days. Do not use separators between the *mm* and *dd* values, or between the *dd* and *yyyy* values.

If no date is specified on the command line, the current value will be displayed.

---

debug

Identifies an CNA subsystem to be debugged and causes CNA to begin generating debug messages for that system.

Syntax

```
debug <subsystem>
no debug [ <subsystem> ]
```

Arguments

- *subsystem*—name of the CNA subsystem to be debugged (can be any argument listed under **debug on** command, except *all* and *services*)

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Description

The **debug** command adds a line to the running configuration which causes the specified subsystem of the CNA operating system to begin generating debug messages. Under normal operating conditions, these messages are undesirable because of their volume and the frequency with which they are generated. But they can be valuable in debugging or understanding unexpected behavior.

The **debug** command by itself only initiates the generation of the messages. It does not cause the messages to be sent to a destination, or to be displayed. You need to tell the operating system where the messages should be directed. This allows you to open multiple terminal displays or CLI sessions, with debug messages going to some, while others show normal program output.

There are three related commands that regulate access to debug information generated by CNA. The commands are:

- **debug**—causes the specified subsystem to begin generating debug messages
- **debug on**—causes CNA to send log messages for the specified subsystem to the current terminal session
- **debug display**—determines whether debug messages are intermingled with normal program output

Use the **no** form of the **debug** command to deactivate generation of debug messages for the specified subsystem. If no subsystem is specified, all debug messages will be deactivated.

To determine whether debugging messages are being generated for a subsystem, you will need to inspect the running configuration for the presence of a **debug** command which specifies that subsystem. The **show debugging** command lists the display status for the current session only.

See also the **logging on** command.

debug display

Specifies how debug messages will interact with normal program output.

Syntax

```
debug display [ foreground | background ]
              [ | { include | exclude } <exp> ]
              [ > <file-url> ]
no debug display
```

Prerequisites

Must have privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

Debug messages are interspersed with regular CLI output.

Description

When the debugging process is enabled, messages are by default interspersed with normal command-line output.

Use the **debug display** command with the **foreground** keyword to force CNA to show debug messages exclusively. No other output will be displayed, and no other CLI commands can be issued in the terminal window in which the command was issued until you cancel the foreground display by pressing **Control-C**.

Use the **background** keyword to intermingle debug output with normal CLI output.

Use the pipe character (“|”) with either the **include** or **exclude** keyword and a regular expression to filter the messages that are displayed.

Use the **debug display** command with the redirection character (>) and a valid FTP, HTTP, or HTTPS URL to send debug messages to a file.

Use the **no** form of the command without either the **foreground** keyword or the > character and filename to return the CLI to its default behavior (debug text interspersed with other output).

Note:

The **debug display** command will not have any affect unless the **debug on** command was executed first.

---

debug on

Causes CNA to send log messages for the specified subsystem to the current terminal session.

Syntax

```
debug on <subsystem>
no debug on [ <subsystem> ]
```

Arguments

- *subsystem*—name of the CNA subsystem to be displayed

Prerequisites

Must have privileged-level access.

Description

The **debug on** command causes CNA to send all log messages that are being generated for the specified subsystem to be sent to the current terminal session.

You must have first enabled the generation of these messages with the **debug** command.

The subsystem identifier can be one of the following:

- *all*—enables all subsystems
- *am*—active measurements
- *am-measurer*—USTAT-resident active measurement manager
- *am-scheduler*—engine-resident active measurement scheduler
- *billing*—manages generation of monthly billing reports
- *builds*—manages backup/restore processes
- *coreman*—configuration, runtime management
- *dmaker-configd*—decision-maker bootstrap process
- *dmaker*—manages the *decision-maker* process
- *dmaker-decisionmaker*—process which decides subnet status
- *dmaker-efclistener*—engine-resident process, receives data from the EFC (endpoint flow collection) module
- *dmaker-main*—manages all *decision-maker* subsystems
- *dmaker-measurementlistener*—engine-resident process, receives measurements from USTATs
- *dmaker-peermgr*—manages BGP peering
- *dmaker-prefixscheduler*—generates route updates
- *dmaker-probefailurelistener*—alerts the *decision-maker* process to active-probe failures
- *dmaker-rsomd*—manages *decision-maker* database requests for snapshot reports
- *dmaker-showd*—provides information for engine decisions
- *dmaker-snmplistener*—keeps the *decision-maker* process informed about link load, and status of BGP peerings and interfaces
- *dmaker-trendprefixlive*—manages collection of trending data from the live database
- *dmaker-trendsummarylive*—manages summarizing of trending data from the live database
- *edgestats*—USTAT-resident process for measurement calculation
- *efc*—endpoint flow collector process
- *email*—automatic mailing of reports
- *event-manager*—automatic mailing of reports
- *event-reporter*—automatic mailing of reports
- *link-monitor*—link-usage polling process
- *medic*—CNA operating system monitor process
- *monitor*—SPAN port packet capture process

- *netflow*—netflow packet capture process
- *post-proc*—monitors post-processing stage of upgrade process
- *subnet-load*—subnet load computation
- *rrd-writer-main*—RRD writer daemon for trend records
- *rrd-writer-rrdpacket-receiver*—RRD writer daemon packet receiver
- *rrd-writer-rrdpacket-processor*—RRD writer daemon packet processor
- *services*—messages generated by third-party processes
- *snapshot*—monitors the creation of snapshot database objects
- *snapshot-main*—manages snapshot database and reports
- *snapshot-showd*—provides information for last hour's engine decisions
- *snapshot-trend*—manages collection of trending data from the hourly database snapshot
- *snapshot-update*—updates snapshot reports following an CNA upgrade
- *snmp*—monitors snmp requests to CNA

The *all* subsystem identifier allows you to toggle the display of all messages on and off. When on, this subsystem overrides your other settings. When *all* is subsequently turned off, your previous subsystem settings will be restored.

The *services* subsystem is made up of third-party processes that generate messages not written by Avaya.

Use the **no** form of the command to stop display of debug messages in the current display shell, for the specified subsystem. If no subsystem is specified, all debug messages are turned off in the current display shell.

Use the **`show debugging`** command to see if the current CLI session is configured to receive these messages. To see if CNA has been configured to generate the messages, you will need to inspect the running configuration for occurrences of the **`debug`** command.

See the **`debug`** and **`debug display`** commands.

decision-policy

Creates a decision-policy configuration block in which policies can be set to control how the CNA decision-maker process handles individual subnets.

Syntax

```
decision-policy <name>
no decision-policy <name>
```

Arguments

- *name*—arbitrary name of the policy

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Decision policy configuration blocks can be created to set parameters for individual subnets, with the following commands:

- **flap-control**
- **ignore-measurements**
- **outage-detection on (deprecated)**
- **outage-detection silence**
- **outage-detection always-move-from-bgp (deprecated)**
- **penalty**
- **prefix-update-rate (deprecated)**
- **priority**
- **priority-subnets**
- **route-assert-filter**
- **winner-set-priority**
- **winner-set-width**

When you execute the **decision-policy** command, the CNA software will enter policy configuration mode.

All commands entered before you enter the **end** command will be contained in the policy configuration block.

These policies will be applied to those subnets you associate with the block using the **set-decision-policy** command.

For more information about policies, see Chapter 6: Decision Policies and Application Models of the *Monitoring Guide*.

decision-policy-assignment

Define how quickly automatically assigned decision policy should expire.

Syntax

```
decision-policy-assignment measurement-based expire-after <duration>
no decision-policy-assignment measurement-based expire-after
```

Arguments

- *duration*—number of seconds after which an automatically assigned decision policy will expire

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access; range: 300 through 86400 (5 min - 1 day).

Occurrence

Use once.

Default

3600

Description

Use this command to set the length of time that the CNA *decision-maker* process will allow an automatically assigned decision policy to remain in effect. If no measurements are received from any subnets governed by this policy for this length of time, then the automatically assigned decision policy will no longer be in force. (See auto-create-subnets):

---

delete

Deletes a persistent object from the CNA storage repository.

Syntax

```
delete <name>
```

Arguments

- *name*—object to be deleted

Prerequisites

System operations mode, privileged-level access.

Description

Certain objects are stored in a persistent repository on CNA, including configuration files, snapshot report files, and persistent output from various **show** commands. Use the **delete** command to permanently remove the object; the **ls** command to see a list of object names; the **rename** command to rename an object; and the **copy** command to make a copy of an object.

---

delete report

Deletes the specified snapshot reports from the CNA hard drive.

Syntax

```
delete report [ <name> [ <object> ]]
              [ customer <customer-name> ]
              [ time { <timestamp> | last } ]
              [ force ]
```

Arguments

- *name*—name of the report; see **show report** command for valid report names
- *object*—name of router or link (required for selected reports)
- *timestamp*—timestamp of specific report, in *yyyy[-mm[-dd[_HH[-MM]]]]* format

Prerequisites

System operations mode, privileged-level access.

Description

Use the **delete report** command to permanently delete selected snapshot reports from CNA.

For information about snapshot reports, see Chapter 2: The Web interface of the *Operator Guide*.

Any incomplete entry of the **delete report** command will cause the deletion of all reports matching the partial data entered. See the **show report** command for more information about matching expressions.

CNA will prompt you for verification. To override this prompt, use the keyword **force**.

If you delete a report inadvertently, use the **recover report** command.

Deleted reports cannot be restored to CNA, even if you exported the reports first.

If you use the **last** keyword instead of a timestamp, the device will delete the most recently saved instances of the specified reports.

See also the **export report** command.

---

deny

Adds a filter condition to an access list.

Syntax

```
deny <protocol>
     {<src> <wildcard> | host <src-addr> | any}
     {<dest> <wildcard> | host <dest-addr> | any}

deny icmp
     {<src> <wildcard> | host <src-addr> | any}
     {<dest> <wildcard> | host <dest-addr> | any}
     [<icmp-type> [<icmp-code>]]
```

```
deny tcp
     {<src> <wildcard> | host <src-addr> | any}
     [<operator> <port> | range <port1> <port2>]
     {<dest> <wildcard> | host <dest-addr> | any}
     [<operator> <port> | range <port1> <port2>]
     [established | syn]

deny udp
     {<src> <wildcard> | host <src-addr> | any}
     [<operator> <port> | range <port1> <port2>]
     {<dest> <wildcard> | host <dest-addr> | any}
     [<operator> <port> | range <port1> <port2>]


no deny <protocol> <source> <source-wildcard> <dest> <dest-wildcard>

no deny icmp <src> <src-wildcard> <dest> <dest-wildcard> [<icmp-type> [<icmp-code>]


no deny tcp
        {<src> <wildcard> | host <src-addr> | any}
        [<operator> <port> | range <port1> <port2>]
        {<dest> <wildcard> | host <dest-addr> | any}
        [<operator> <port> | range <port1> <port2>] [established | syn]

no deny udp
        {<src> <wildcard> | host <src-addr> | any}
        [<operator> <port>|range<port1> <-port2>]
        {<dest> <wildcard> | host <dest-addr> | any}
        [<operator> <port> | range <port1> <port2>]
```

Arguments

- *protocol*—protocol name, or integer in the range 0 through 255, representing an IP protocol number
- *src*—IP address in dotted decimal format (a.b.c.d); source of packet
- *wildcard*—wildcard bits; dotted-decimal mask
- *src-addr*—IP address in dotted decimal format (a.b.c.d)
- *dest*—IP address in dotted-decimal format (a.b.c.d);destination of packet
- *dest-addr*—IP address in dotted-decimal format (ab.c.d)
- *icmp-type*—integer, in the range 0 through 255; ICMP message type
- *icmp-code*—integer, in the range 0 through 255; ICMP message code
- *operator*—can be one of the following string values:
  - *eq*—equal to
  - *lt*—less than
  - *gt*—greater than
  - *neq*—not equal to

- *port*—access port
- *port1*—low-end of range of ports
- *port2*—high-end of range of ports

Prerequisites

Access list configuration mode (*config-acl*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

IP protocol names and numbers are defined in RFC 1700.

The **no** form of the command deletes the matching element from the access list.

The **established** keyword indicates an established connection (ACK or RST bits must be set).

The **syn** keyword means no established connection (will match packets with the SYN bit set and the ACK and FIN bits cleared, which indicates a request to initiate a TCP connection).

The *src* and *dest* addresses are paired with wildcard masks; the *src-host-address* and *dest-host-address* values are complete IP addresses (with an implied mask of 0.0.0.0).

If you get an *insufficient or bad arguments* error message when you apply the access list to an interface, there may be a mismatch between IP address and wildcard mask in one of the rules. See Wildcard Masks in Chapter 3: Security of the *Monitoring Guide*.

If you add a malformed rule to an existing access list that is already applied to an interface, you may see the *insufficient or bad arguments* message whenever you attempt to execute any command related to the interface. To recover from this condition, remove the offending rule from the access list and re-enter the IP address and mask correctly.

See the **`permit`** command.

---

description

Sets the description attribute of an active measurement group, interface, service provider link, or Check Point VPN configuration block.

Syntax

```
description <text>
no description
```

Arguments

- *text*—text of the description, which can consist of letters, numbers, space, hyphen, period, comma, colon, underscore, parenthesis, forward slash, and “commercial at” - @.

Prerequisites

Must be in one of the following configuration modes, all of which require privileged-level access:

- active measurement group (*config-engine-active);* see **active-measurement group** command
- interface (*config-if*); see **interface** command
- link (*config-engine-link*); see **link** command

Occurrence

Use once per active measurement group, interface or link.

Default

Description field is an empty string.

Description

If the text includes spaces, enclose it in double quote characters (“...”).

Do not use single or double quotes within the text.

The **no** form of the command resets the description text to an empty string.

---

diagnose target

Enables logging diagnostics for a specified target.

Syntax

```
diagnose target <name> [ address <target-addr> ]
no diagnose <name>
```

Arguments

- *name*—name of target to be diagnosed:
- *target-addr*—IP address in dotted decimal format (a.b.c.d)

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once.

Default

none

Description

Use the **`diagnose target`** command to enable logging of diagnostics whenever the specified target is probed, or, with the **address** keyword, whenever the target resolves to that particular ip address.

Note:

If you configure more than one target to be diagnosed, only the first target in the config will actually be diagnosed.

Use the **no** form of the command to cancel logging of diagnostics.

---

diagnose probe

Enables logging diagnostics for a specified target.

Syntax

```
diagnose probe address <address> vip <vip-address>
                  type { icmp | tcp | traceroute } module <slot>
                  [ port <port> ] [ lsr <lsr-address> ]
```

Arguments

- *address*—IP address (*a.b.c.d* format)
- *vip-address*—source virtual IP address to use (*a.b.c.d* format)
- *slot*—slot number of the USTAT module
- *port*—tcp port number to use
- *lsr-address*—IP address (*a.b.c.d* format) to be loose source routed through

Prerequisites

System operations mode, privileged-level access.

Default

none

Description

Use the **diagnose probe** command to force the active measurement process to send an immediate probe to a given target and to record diagnostic information for that probe. Then you may view the results with the command **show stat am-measurer diagnose**. See show statistics.

---

direction

Specifies the type usage data to be used when evaluating link load.

Syntax

```
direction { inbound | outbound }
no direction
```

Prerequisites

Must be in load-threshold configuration mode (*config-engine-threshold*), which requires privileged-level access.

Occurrence

Use once per load threshold table.

Default

Outbound

Description

The **direction** command specifies the type of data to be used when evaluating load for links associated with the load threshold table in which the **direction** command is contained.

See the **load threshold-table** command.

disable

In system-operations mode, cancels privileged-level access (requires **enable** password to resume).

In an active-measurement group configuration block, suspends active probing for all targets in the group.

In a Check Point configuration block, suspends polling of the Check Point management system.

Syntax

```
disable
no disable
```

Prerequisites

Requires privileged-level access.

To suspend active probing, must be in active-measurement group configuration mode (*config-engine-active*).

Occurrence

Use once per active-measurement group or Check Point configuration block.

Description

If you execute the **disable** command at a privileged-level access prompt (not in *config* mode), privileged-level access will be cancelled. The **enable** password will be required to resume privileged-level access.

When the **disable** command is entered within an active-measurement group configuration block, active probing of all targets in the group is suspended. The active-measurement group configuration block remains in the configuration, and all targets remain defined, but no probes are conducted.

Disabled groups will be listed in the output of the **show statistics** command executed with the *am-scheduler* process.

When the **disable** command is enter within a Check Point configuration block, VPN polling is suspended.

Use the **no** form of the command to remove the disable command from the configuration block, thereby enabling the probes.

See the **active-measurement group** command for more information about configuring active measurement groups.

dump-measurements

Displays live measurements for a subnet as they are received the CNA.

Syntax

```
dump-measurements <subnet> | <address>
           [ count <number> ]
           [ | { include | exclude | begin } <exp> ]
           [ > <url> ]
```

Arguments

- *subnet*—subnet in CIDR format (a.b.c.d /*n*)
- *address*—IP address (*a.b.c.d*)
- *number*—show this number of measurements, then exit the command
- *exp*—regular expression
- *url*—destination of output; can be FTP, HTTP, or HTTPS URL

Prerequisites

System operations mode, login level access.

Description

Use the **dumpmeasurements** command to watch measurements as they are received by CNA.

You can pipe the command through one of the following filters (the vertical-line pipe character | is required when a filter is used):

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output after

If *url* is set to a valid FTP, HTTP or HTTPS URL (with the required pipe character, >), output will be sent to the specified destination. See the **copy** command for URL syntax.

To stop the **dumpmeasurements** process, enter **Control-C**.

duplex

Sets duplex mode for an interface.

Syntax

```
duplex [ full | half ]
no duplex
```

Prerequisites

Must be in interface configuration mode (*config-if*), which requires privileged-level access.

Occurrence

Use once per interface.

Default

Auto

Description

The **duplex** command with no keyword sets the mode to **auto**.

Both duplex and speed of an interface can be configured, but they are dependent on each other. If any one of them is set to **auto**, then both will be implemented as **auto** regardless of the setting of the other.

Changing either duplex or speed will bring the interface down and back up and you will lose Telnet or SSH connectivity. Changes to these modes should be done from a serial port connection.

See the **speed** command.

efc bytes minimum

Specifies a minimum amount of traffic volume for IP addresses sent to an active measurement group by the optional EFC module.

Syntax

```
efc bytes minimum <number>
no efc bytes minimum
```

Arguments

- *number*—number of bytes

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once per active-measurement group.

Default

Zero

Description

Use the **`efc bytes minimum`** command to specify a minimum amount of traffic volume for IP addresses added to an active measurement group by optional EFC module. Endpoints that did not transfer the number of bytes specified by *number* will not be added to the group.

The **`efc bytes minimum`** command is one of several commands which place restrictions on which endpoints obtained from the EFC module will be probed. See also:

- **`efc bytes top`**
- **`efc last-update-within`**
- **`efc exclude-subnet`**
- **`efc source`**

You can use them together in the same active measurement group.

Use the **no** form of the command to remove the minimum bytes restriction from the active-measurement group configuration block.

For more information about EFC, see Chapter 2: Collecting Endpoint Addresses of the *Monitoring Guide*.

---

efc bytes top

Specifies the number (based on traffic) of IP addresses obtained from the optional EFC module that will be probed as part of an active measurement group.

Syntax

```
efc bytes top <number> [ starting-from <start> ]
efc bytes top percent <number>
no efc bytes top
```

Arguments

- *number*—number of addresses (range: 1 through 1,000,000) or percentage (range: 1 through 100)
- *start*—exclude the top n addresses, where n is (*start - 1)*, range is 1 through 1,000,000

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once per active-measurement group.

Default

1000

Description

Use the **`efc bytes top`** command to use relative traffic volume as a limiting factor on the number of EFC-obtained IP addresses that will be probed as part of an active measurement group. Relative traffic volume is determined by the number of bytes that originated or terminated at the endpoint.

Without the **percent** keyword, the value of *number* means the number of the top (most-heavily trafficked) IP addresses collected by the EFC module that will be probed as part of the active measurement group.

If you include the **percent** keyword, CNA will probe the subnets that account for the specified percentage of traffic.

The **starting-from** keyword is useful when you want to exclude a range of addresses already covered by another active-measurement group.

The following version of the command would result in the 10 busiest addresses in the data stream being probed:

```
efc bytes top 10
```

With the **percent** keyword, the command would result in the number of subnets that account for the top specified percentage of traffic handled by the data stream being probed. The following example would result in probes being sent to the subnets responsible for 90 percent of the traffic:

```
efc bytes top 90 percent
```

The total traffic value used for determining the top endpoints (whether number of subnets or percentage of traffic) is based on traffic associated with endpoints already filtered for this active-measurement group.

The **`efc bytes top`** command is one of several commands which place restrictions on which endpoints obtained from the EFC module will be probed. See also:

- **`efc bytes minimum`**
- **`efc last-update-within`**
- **`efc exclude-subnet`**
- **`efc source`**

You can use them together in the same active measurement group.

Use the **no** form of the command to remove the top-traffic limitation.

For more information about EFC, see Chapter 2: Collecting Endpoint Addresses of the *Monitoring Guide*.

efc enable

Enables use of the optional Endpoint Flow Collector module as a source of probe targets for an active measurement group.

Syntax

```
efc enable
no efc enable
```

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once per active-measurement group.

Default

EFC collection of targets is off

Description

With the optional EFC module installed, the CNA can automatically create active-measurement probes from IP address received in a Netflow or SPAN data stream.

To enable this feature, use the **`efc enable`** command in the active-measurement group configuration block.

Use the **no** form of the command while in active-measurement group configuration mode to remove the command from the block and disable endpoint collection for the group.

For more information about EFC, see Chapter 2: Collecting Endpoint Addresses of the *Monitoring Guide*.

---

efc filter

Identifies an access list to be used by an active-measurement group for filtering targets received from the optional EFC module.

Syntax

```
efc filter <name> [ endpoint { source | destination } ]
no filter [ <name> ]
```

Arguments

- *name*—name of a configured access list

Prerequisites

Must be in active-measurement configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

Both source and destination endpoints are added to the group

Description

Use the **efc filter** command to specify an access list to be used to filter targets received from the optional EFC module.

Create the access list with the **ip access-list** command. If you modify an existing filter (one that has been referenced by the **efc filter** command), you must update the access list in the CNA *efcserver* process. See the **resync efc filter** command.

If you use the **endpoint source** keywords, only source addresses from the flows matching the access list criteria will be probed.

If you use the **endpoint destination** keywords, only destination addresses from flows matching the access list criteria will be probed.

By default, both source and destination addresses are probed.

If multiple filters are configured for the active measurement group, the access lists will be applied in the order in which they appear in the configuration.

For more information about EFC, see Chapter 2: Collecting Endpoint Addresses of the *Monitoring Guide*.

---

efc last-update-within

Isolates IP addresses within a flow received from the optional EFC module according to when they last sent or received data.

Syntax

```
efc last-update-within <interval>
                                  [
                                    seconds |
                                    minutes |
                                    hours |
                                    days |
                                    months |
                                    years
                                  ]
no efc last-update-within
```

Arguments

- *interval*—number of units of time

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once per active-measurement group.

Default

Interval units in seconds

Description

Use the **`efc last-update-within`** command to specify an interval of time in which IP addresses received from the optional EFC module must have sent or received data in order to be used as probe targets.

The **`efc last-update-within`** command is one of several commands which place restrictions on which endpoints obtained from the EFC module will be probed. See also:

- **`efc bytes minimum`**
- **`efc bytes top`**
- **`efc exclude-subnet`**
- **`efc source`**

You can use them together in the same active measurement group.

Use the **no** form of the command to remove the time restriction from the active-measurement group.

For more information about EFC, see Chapter 2: Collecting Endpoint Addresses of the *Monitoring Guide*.

efc exclude-prefix (deprecated)

The efc exclude-prefix command has been deprecated. See the replacement command, efc exclude-subnet on page 126.

efc exclude-subnet

Identifies specific subnets that are not to be automatically probed as a part of a data stream received by the optional EFC module.

Syntax

```
efc exclude-subnet <address>[/<mask>]
no efc exclude-subnet [ <address> ]
```

Arguments

- *address*—IP address (*a.b.c.d*)
- *mask*—mask portion of CIDR-formatted subnet address

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

None

Description

Use the **`efc exclude-subnet`** command to identify subnets that are not to be probed as a part of a data stream received by the optional EFC module.

If the /*mask* portion of the address is not specified, the address is assumed to be a /32 subnet (a single host).

The **`efc exclude-subnet`** command is one of several commands which place restrictions on which endpoints obtained from the EFC module will be probed. See also:

- **`efc bytes minimum`**
- **`efc bytes top`**
- **`efc last-update-within`**
- **`efc source`**

You can use them together in the same active measurement group.

Use the **no** form of the command without an address to remove all subnet exclusions from the active-measurement group configuration block.

Use the **no** form of the command with an address to remove a specific exclusion from the active-measurement group configuration block.

For more information about EFC, see Chapter 2: Collecting Endpoint Addresses of the *Monitoring Guide*.

efc source

Specify traffic-source type for IP addresses to be added to an active measurement group by the optional EFC module.

Syntax

```
efc source { netflow [ <address> ] | monitor}
no efc source [ netflow [ <address> ] | monitor ]
```

Arguments

- *address*—IP address (*a.b.c.d*) on the router where Netflow is enabled

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

Endpoints obtained from all input sources will be probed

Description

Use the **`efc source`** command to specify a type of input stream; endpoints gathered by the optional EFC module from the specified source will be probed by the active measurement group.

Specify an IP address with the **netflow** keyword if you want to identify a Netflow stream from a specific router interface.

Use the **monitor** keyword to identify SPAN data coming from a switch.

The **`efc source`** command is one of several commands which place restrictions on which endpoints obtained from the EFC module will be probed. See also:

- **`efc bytes minimum`**
- **`efc bytes top`**
- **`efc last-update-within`**
- **`efc exclude-subnet`**

You can use them together in the same active measurement group.

Use the **no** form of the command with no arguments to revert to the default behavior for the active-measurement group.

Use the **no** form of the command with the **netflow** keyword to remove the explicit **efc source netflow** configuration statement for all Netflow sources from the active measurement group configuration block.

Specify a specific address with the **no** form of the command and the **netflow** keyword to remove the explicit **efc source netflow** configuration statement for that particular Netflow source from the active measurement group configuration block.

Use the **no** form of the command with the **monitor** keyword to remove the explicit **efc source monitor** configuration statement from the active measurement group configuration block.

For more information about EFC, see Chapter 2: Collecting Endpoint Addresses of the *Monitoring Guide*.

email

Sets a schedule and target list for automatic mailings of reports.

Syntax

```
email { ( <report>
            [ <object> ]
            [ customer <name> | all ]
            [ subnet <subnet> ] )
        | ( trend <filename>
            [ application <app-name> ] ) }
      [ to <list> ]
      schedule <minute> [ <hour> [ <date> [ <month> [ <day> ]]]]

no email [ <report> [ <object> ]]
```

Arguments

- *report*—name of report to be e-mailed; can be one of the following (some reports, as indicated in the list, require a router name; the trend report requires a filename):
  - *application-delay*
  - *application-delay-blackouts*
  - *application-delay-poor*
  - *application-delay-very-poor*
  - *application-performance*
  - *application-performance-blackouts*

- *application-traffic*
- *asn-traffic*
- *billing*
- diagnostics
- diagnostics-brief
- *distribution-all <router>*
- *distribution-all-weighted <router>*
- *distribution-named <router>*
- *distribution-named-weighted <router>*
- *distribution-sufficient <router>*
- *distribution-sufficient-weighted <router>*
- *executive-summary*
- *executive-summary-bgp*
- *executive-summary-bgp-weighted*
- *executive-summary-weighted*
- *measurements*
- *measurements-blackouts*
- *subnet-traffic*
- *subnet-traffic-summary*
- *observed*
- *observed-blackouts*
- *observed-poor*
- *observed-very-poor*
- *optimized <router>*
- *optimized-blackouts <router>*
- *outage-application-summary*
- *outage-link-summary*
- *outage-subnet-summary*
- *outage-subnet-summary*
- *subnet-load*
- *subnet-load-summary*
- *trend <filename>*

- *object*—one of the following:
  - name or IP address of router
  - *.qry* or *.tmp* file for trend reports
- *name*—name of the customer for which the report is generated. The customer-specific subnet list will be used when generating such a report. If you use '**all**' for the name, then reports for all defined customers will be sent.
- *subnet*—subnet in CIDR format (a.b.c.d /*n*); generates a report specific to this subnet
- *list*—list of recipient user IDs supported by the SMTP server (enclose in quotes if the list includes spaces)
- *filename*—name of .qry or .tmpl file that defines the trend to be published
- *app-name*—name of an application (e.g., http, icmp, ftp, sap, etc.). One way to see what applications have been found by CNA is to use the command **show efc application**.
- *minute*—number of minutes after the hour when the report is to be mailed; range: 0 through 59 (or * for every minute, or a comma-separated list of minutes, such as 5,25,45)
- *hour*—hour of the day when the report is to be mailed; range: 0 through 23, or * for every hour (or * for every hour, or a comma-separated list of hours, such as 10,22)
- *date*—day of the month when the report is to be mailed; range: 1 through 31, or * for every day of the month (or * for each date of the month, or a comma-separated list of dates, such as 1,15)
- *month*—month report is to be mailed; range: 1 through 12 (or * for every month, or a comma-separated list of months, such as 1,7)
- *day*—day of week when the report is to be mailed; range 0 through 7 (or * for every day of the week, or a comma-separated list of days, such as 0,3); 0 and 7 both mean Sunday

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access, to set a schedule. One-time mailing of a report can be done with log-in level access.

Occurrence

Available for use multiple times in the configuration.

Default

* for all schedule increments (except *minute*)

Description

Use the **email** command to schedule automatic mailings of reports. You can set one schedule for each report; enter the command once for each.

You can also use this command directly from the command line (exec mode) to email reports immediately without needing to schedule it. In this case, you cannot use the **schedule** keyword.

Reports will be sent by the management module via the mail server defined by the **smtp-server** command. You should verify that the CNA management module's *eth0* interface has access to the mail server (SMTP servers typically listen on tcp port 25).

Use the **no** form of the command with no parameters to cancel all automatic e-mailing.

Use the **no** form of the command with a report name (and object type if required) to cancel a specific report schedule.

email format image

Sets the format for images when they are attached to email messages sent by CNA.

Syntax

```
email format image[ jpeg | png ]
no email format image
```

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Single

Default

png

enable

Enables a specific active-measurement task.

Syntax

enable

no enable

Prerequisites

Must be in Chatter Active Measurement Configuration command mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrences

Use once.

Default

Disabled.

Description

Enables a specific active-measurement task.

Chatter supports only one active measurement per particular measurement type (RTP, ping, TCP, and traceroute). If you attempt to enable an active-measurement that is already active, an error message is displayed.

By default, active-measurement tasks are disabled.

---

enable password

Sets or removes the password required for privileged-level access.

Syntax

```
enable password [encrypted] <pwd>
no enable password
```

Arguments

- *pwd*—password required for privileged-level access

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Default

No password is required until you set one.

Description

Use the **no** form of the command to remove password protection for privileged-level access.

If there is no **enable** password defined, anyone can execute the **`enable`** command and gain privileged-level access.

If the **encrypted** keyword is used, *pwd* must already have been encrypted by another CNA device.

end

Backs you out of the current operating mode.

Syntax

```
end
```

Description

This command is comparable to the **`exit`** command, except that the **`end`** command will not terminate the CNA session.

If you execute the **`end`** command while in system-operations mode, nothing will happen.

See the **`exit`** and **`logout`** commands.

erase startup-config

Deletes the startup configuration file.

Syntax

```
erase startup-config
```

Prerequisites

System operations mode, privileged-level access.

Description

When you execute the **`erase startup-config`** command, the device's startup configuration file is deleted from the system's internal storage. (The current startup configuration is backed up before it is erased.)

When you reboot the device, it will behave as it did when new. You will be prompted on the console for the basic information required to get the device connected to the local LAN.

eventdump

Displays CNA events on the screen as they are received by the CNA event manager process.

Syntax

```
eventdump
           [ detail ]
           [ count <number> ]
           [ subnet <subnet> | group <group-name> ]
           [ link <link-name> ]
           [ selected ]
           [ asserted ]
           [ confirmed ]
           [ bgp-announcement ]
           [ bgp-withdrawal ]
           [ bgp-summary ]
           [ link-status ]
           [ unlogged-action ]
           [ state-transition ]
           [ peer-status ]
           [ perf-violation ]
           [ outage ]
           [ not-rescued-outage ]
           [ | { include | exclude } <exp> ]
           [ > <url> ]
```

Arguments

- *number*—show this number of events, then exit the command
- *subnet*—subnet in CIDR format (a.b.c.d /*n*)
- *group-name*—name of a group
- *link-name*—name of a link
- *exp*—regular expression
- *url*—destination of output; can be FTP, HTTP, or HTTPS URL

Prerequisites

System operations mode, login level access.

Description

The **`eventdump`** command causes CNA to display data about events as they are sent to the CNA event manager.

The following additional keywords are supported:

- **detail**—displays all values for each event

- **subnet**—all events for a particular subnet
- **group**—all events for the subnets in a particular group
- **link**—all events for a particular link
- **selected**—route-selected events
- **asserted**—route-asserted events
- **confirmed**—route-confirmed events
- **bgp-announcement**—BGP announcement events
- **bgp-withdrawal**—BGP withdrawal events
- **bgp-summary**—BGP summary events
- **link-status**—link-status events
- **unlogged-action**—events generated when there are lost events due to an overflow of the event queue
- **state-transition**—decision maker state transition events
- **peer-status**—CNA BGP peering events
- **perf-violation**—performance violation events

You can pipe the command through one of the following filters (the vertical-line pipe character | is required when a filter is used):

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression

If *url* is set to a valid FTP, HTTP or HTTPS URL (with the required pipe character, >), output will be sent to the specified destination. See the **copy** command for URL syntax.

Note:

The keyword **raw** is no longer supported.

To stop the **eventdump** process, enter **Ctrl-C**.

See also the **show event** command.

exclude

Define number of HTTP turns on an average web page (used to show time savings in executive summary snapshot report).

Syntax

```
exclude ( subnet <subnet> ) | ( application <name> ) |
        ( filter <access-list> { source | destination } )
no exclude subnet <subnet>
no exclude application <name>
no exclude filter <access-list>
```

Arguments

- *subnet*—subnet in CIDR format (a.b.c.d /*n*)
- *name*—name of an application you want to exclude; user-defined names must start with an uppercase letter and can be up to 32 characters in length. Predefined applications will always start with a lowercase letter.
- *access-list*—name or number of an access list of the application you want to exclude

Prerequisites

Must be in EFC configuration mode (*efc-reporting*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

none

Description

Use the **exclude** command when you want the EFC module to completely ignore a particular flow. This means that the flow's statistics are not counted in traffic rates, and end points associated with that flow will not be used for active measurement.

Use the **no form of the** command to cause the EFC module to stop ignoring flows from a particular subnet or application.

---

exec-summary http-turns

Define number of HTTP turns on an average web page (used to show time savings in executive summary snapshot report).

Syntax

```
exec-summary http-turns <min> <max>
no exec-summary
no exec-summary http-turns
```

Arguments

- *min*—minimum number of turns on a page; range: 0 through 10000
- *max*—maximum number of turns on a page; range: 0 through10000

Prerequisites

Must be in report configuration mode (*config-reporting*), which requires privileged-level access.

Occurrence

Use once.

Default

min—20

max—150

Description

The **exec-summary http-turns** command specifies a range of HTTP turns, representing a page on your web server.

In order to download a complete web page, a client typically needs to contact the server several times, as it obtains graphics and other supporting files specified in the page definition. Each completed contact represents one HTTP turn.

These numbers will be used to calculate the *seconds saved* range of values shown on the *Executive Summary* snapshot report.

Use the **no exec-summary** command to reset all executive summary parameters to their defaults, including those set by each of the following commands:

- **exec-summary http-turns**
- **exec-summary improvable min-hrtt-saving**
- **exec-summary improvable min-percent-improvement**
- **exec-summary min-bgp-hrtt**
- **web-speedup worst-bgp-hrtt-percent**

Use the **no exec-summary http-turns** command to reset only the HTTP turn values to their defaults.

---

exec-summary improvable min-hrtt-saving

Specifies the minimum time savings required for a subnet to be considered improvable when generating the Executive Summary snapshot report.

Syntax

```
exec-summary improvable min-hrtt-saving <value>
no exec-summary
no exec-summary improvable min-hrtt-saving
```

Arguments

- *value*—minimum HRTT measurement (in milliseconds); range: 0 through 1000000

Prerequisites

Must be in report configuration mode (*config-report*), which requires privileged-level access.

Occurrence

Use once.

Default

value—0

Description

Use the **exec-summary improvable min-hrtt-saving** command to set the minimum time savings, in milliseconds, that must be achieved by switching from a BGP route to the CNA choice before a subnet is considered improvable in the executive summary snapshot report.

Time savings is one of two criteria used by the CNA report-generating process to determine if a subnet should be counted as improvable. The other is percentage of improvement. The subnet must meet conditions (minimum time savings and percentage speedup).

Use the **no exec-summary** command to reset all executive summary parameters to their defaults, including those set by each of the following commands:

- **exec-summary http-turns**
- **exec-summary improvable min-hrtt-saving**
- **exec-summary improvable min-percent-improvement**
- **exec-summary min-bgp-hrtt**
- **web-speedup worst-bgp-hrtt-percent**

Use the **no exec-summary http-turns** command to reset only minimum HRTT value to the default.

---

exec-summary improvable min-percent-improvement

Specifies the minimum percentage of speedup required for a subnet to be considered improvable when generating the Executive Summary snapshot report.

Syntax

```
exec-summary improvable min-percent-saving <value>
no exec-summary
no exec-summary improvable min-percent-saving
```

Arguments

- *value*—minimum percentage improvement required; range: 0 through 100

Prerequisites

Must be in report configuration mode (*config-reporting*), which requires privileged-level access.

Occurrence

Use once.

Default

value—10

Description

Use the **`exec-summary improvable min-percent-improvement`** command to set the minimum percentage of speedup that must be achieved by switching from a BGP route to the CNA choice before a subnet is considered improvable in the *Executive Summary* snapshot report. Percent improvement (time saved relative to final time) is calculated by subtracting the HRTT measurement across the CNA-chosen link from the HRTT measurement across the slowest BGP-selected link, dividing by the CNA-link's HRTT, and multiplying by 100:

$$\text{Percent Improvement} = \frac{\text{BGP HRTT} - \text{CNA HRTT}}{\text{CNA HRTT}} \times 100$$

If you have multiple edge routers, there will be multiple BGP-selected routes. The percent improvement shows how much better the CNA choice performs than the slowest of these BGP routes.

This percentage is one of two criteria used by the CNA report-generating process to determine if a subnet should be counted as improvable. The other is time savings. The subnet must meet both conditions (minimum time savings and percentage speedup).

Use the **`no exec-summary`** command to reset all executive summary parameters to their defaults, including those set by each of the following commands:

- **`exec-summary http-turns`**
- **`exec-summary improvable min-hrtt-saving`**
- **`exec-summary improvable min-percent-improvement`**

- **exec-summary min-bgp-hrtt**
- **web-speedup worst-bgp-hrtt-percent**

Use the **no exec-summary improvable min-percent- improvement** command to reset only the minimum percentage value to its default.

exec-summary min-bgp-hrtt

Defines one of the brownout conditions—speed of packet delivery—for the *Executive Summary* snapshot report.

Syntax

```
exec-summary min-bgp-hrtt <value_1> <value_2>
no exec-summary
no exec-summary min-bgp-hrtt
```

Arguments

- *value_1*—minimum HRTT value (milliseconds) required to be considered *poor*; range: 0 through 1000000
- *value_2*—minimum HRTT value required to be considered *very poor*; range: 0 through 1000000

Prerequisites

Must be in report configuration mode (*config-reporting*), which requires privileged-level access.

Occurrence

Use once.

Default

value_1—1000000

value_2—1000000

Description

Use the **exec-summary min-bgp-hrtt** command to set one of the two types of criteria used to define poor and very poor performance in the *Executive Summary* snapshot report. The other criteria is defined as “worst N percent of all subnets,” set by the **web-speedup worst-bgp-hrtt-percent** command.

The minimum BGP values set the threshold for HRTTs (handshake round trip times, in milliseconds) on the worst performing BGP route in your subnet:

- If a subnet measurement is at least as high as *value_1* and it qualifies as one of the worst N subnets, it will be factored into the *Poor Performance* column of the Executive Summary snapshot report.
- If the subnet measurement is at least as high as *value_2* and it qualifies as one of the worst N subnets, it will also be factored into the *Very Poor Performance* column.

Use the **no exec-summary** command to reset all executive summary parameters to their defaults, including those set by each of the following commands:

- **exec-summary http-turns**
- **exec-summary improvable min-hrtt-saving**
- **exec-summary improvable min-percent-improvement**
- **exec-summary min-bgp-hrtt**
- **web-speedup worst-bgp-hrtt-percent**

Use the **no exec-summary min-bgp-hrtt** command to reset only the minimum BGP HRTT values to their defaults.

exit

Exits the current operating mode.

Syntax

```
exit
```

Description

The **exit** command is comparable to the **end** command, except that the **exit** command will terminates the CNA session if you are in system operations mode.

See the **end** and **logout** commands.

export event

Exports events to a remote location.

Syntax

```
export event [ time [ <timestamp> | last ]
```

Arguments

- *timestamp*—timestamp of log message, in *yyyy-mm-dd_HH-MM* format

Prerequisites

System operations mode, login-level access.

Description

Use the **export event** command to export events to a remote location.

You can enter all or part of a time stamp. The command will export all messages that match the entry. Include hyphens to separate date and time elements, and the underscore character to separate date from time.

Use the keyword **last** instead of a timestamp to export the most recently logged event.

The event log on the device is limited to 500 MB of disk space. When this space is full, incoming messages overwrite the oldest existing messages.

See also the **show event** command.

---

export report

Exports the specified snapshot reports from the CNA hard drive to a remote location (or to a different local directory).

Syntax

```
export report [ <name> [ <object> ] ]
              [ customer <customer-name> ]
              [ time { <timestamp> | last } ]
              <url>
```

Arguments

- *name*—name of the report; see **show report** command for valid report names
- *customer-name*—name of the customer, defined by <customer>.subnet files (see the **customer-subnet-file** command)
- *object*—name of router or link (required for selected reports)
- *timestamp*—timestamp of specific report, in *yyyy[-mm[-dd[_HH[-MM]]]]* format

- *url*—destination; can be any valid FTP URL, but filename must have one of the following extensions:
  - *.tar*
  - *.gz*
  - *.tgz*
  - *.zip*

Prerequisites

System operations mode, privileged-level access.

Description

Use the **`export report`** command to export selected snapshot reports from CNA to a remote destination specified as an FTP URL. See the **`copy`** command for information on URL syntax.

The exported file will be formatted according to the extension of the specified filename in the URL. This file will contain the snapshot reports as HTML files individually compressed in *.gz* format.

Exporting reports does not remove them from the hard drive; to do so, you must also use the **`delete report`** command. You cannot restore deleted reports.

For information about snapshot reports, see Chapter 3: Reports of the *Operator Guide*.

Any incomplete entry of the **`export report`** command arguments will export all reports matching the partial data entered. See the **`show report`** command for more information about matching expressions.

If you use the **time last** keywords instead of **time** <*timestamp*>, the device will export the most recently saved instances of the specified reports.

export report (scheduled)

Schedules the export of specified snapshot reports from the CNA hard drive to a remote location (or to a different local directory).

Syntax

```
export report [ <name> [ <object> ] ]
              [ customer <customer-name> ]
              [ time { <timestamp> | last } ]
              <url>
```

Arguments

- *name*—name of the report; see **show report** command for valid report names
- *customer-name*—name of the customer, defined by <customer>.subnet files (see the **customer-subnet-file** command)
- *object*—name of router or link (required for selected reports)
- *timestamp*—timestamp of specific report, in *yyyy[-mm[-dd[_HH[-MM]]]]* format
- *url*—destination; can be any valid FTP URL, but filename must have one of the following extensions:
  - *.tar*
  - *.gz*
  - *.tgz*
  - *.zip*

Prerequisites

System operations mode, privileged-level access.

Description

Use the **export report** command to export selected snapshot reports from CNA to a remote destination specified as an FTP URL. See the **copy** command for information on URL syntax.

The exported file will be formatted according to the extension of the specified filename in the URL. This file will contain the snapshot reports as HTML files individually compressed in *.gz* format.

Exporting reports does not remove them from the hard drive; to do so, you must also use the **delete report** command. You cannot restore deleted reports.

For information about snapshot reports, see Chapter 3: Reports of the *CNA Reporting Guide*.

Any incomplete entry of the **export report** command arguments will export all reports matching the partial data entered. See the **show report** command for more information about matching expressions.

If you use the **time last** keywords instead of **time** <*timestamp*>, the device will export the most recently saved instances of the specified reports.

---

export trend

Export trend data file to a remote location.

Syntax

```
export trend <file> <url>
```

Arguments

- *file*—name of the trend file (*.qry*)
- *url*—destination; filename must end in one of the following:
  - *.gz*
  - *.tar*
  - *.tgz*
  - *.zip*

Prerequisites

System operations mode, log-in level access.

Occurrence

Available for use multiple times in the configuration.

Description

The **`export trend`** command retrieves the data required to build a trend graph from the trending database and exports it to a remote location. See the **copy** command for valid URL formats.

The exported file will be formatted according to the file extension specified for a destination. The contents of the file will be in XML format.

---

flap-control

Allows you to set values that control route flapping.

Syntax

```
flap-control <half-life> <reuse> <suppress>
no flap-control
```

Arguments

- *half-life*—integer in the range 1 through 45
- *reuse*—integer in the range 1 through 20,000
- *suppress*—integer in the range 1 through 20,000

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once.

Default

- *half-life*—15
- *reuse*—750
- *suppress*—2,000

Description

The **`flap-control`** command is used to achieve an acceptable interval between successive updates from CNA to an edge router for a single subnet, thereby avoiding flapping.

CNA measures flapping by use of an index derived from the elapsed time between updates and the user-configurable value of *half-life*.

The *half-life* value should represent the time (in minutes) that you want it to take for the index to decay by half.

The index is increased by 1,000 every time the subnet is updated. The *suppress* and *reuse* values serve as upper and lower update thresholds—if the index exceeds the value of *suppress*, CNA will not send updates to the router until the index has fallen below *reuse*.

If you find your CNA device to be flapping—sending updates for the same subnet too frequently—you should increase the value of *half-life*. Conversely, if your CNA device stops sending updates, or the sustainable rate of updates is too low, you should decrease the value of *half-life*.

The *suppress* and *reuse* values can also be adjusted, but you should do so cautiously.

Increasing the *suppress* value introduces more burstiness into the system.

If you adjust the rate of updates simply by increasing the *suppress* value, you may find the increase in frequency to be a temporary phenomenon.

If you do increase the *suppress* value, you should also increase the *reuse* value by a proportional amount. That is, the distances between 0 and *reuse* and between *reuse* and *suppress* should remain proportional to the distances between the comparable default values.

If the command is entered while in engine configuration mode, the parameters will be applied to all subnets globally.

If the command is entered while in policy configuration mode, the parameters will be applied only to those subnets associated with the policy.

For more information about policies, see the **decision-policy** command and Chapter 6: Decision Policies and Application Models of the *Monitoring Guide*.

group

Allows you to define a group of subnets.

Syntax

```
group <name> [ application <app-name> | subnet <subnet> ]
no group <name>
```

Arguments

- *name*—string; name of new group
- *subnet*—subnet in CIDR format (a.b.c.d /*n*)

Prerequisites

Must be in e configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

none

Description

Use the **group** command to create generic subnet groups or application groups which can then be referenced from within other groups in the configuration, such as report group and performance-group. After you create a new group, you will then be able to use the **prefix (deprecated)** or **application** command to add subnets or applications to the group.

Use the **no** form of the command to delete a group from the configuration.

halt

Halts CNA.

Syntax

```
halt
```

Description

Use this command to halt all processes on CNA to ready it for power shutdown.

---

header

Adds an HTTP header to the CNA web server's response string.

Syntax

```
header P3P [ add | append | set ] <text>
header P3P unset
no header append P3P
```

Arguments

- *text*—text to be sent as the header content, enclosed in single quote marks ('...'); should contain a URL to an *.xml* file describing your web site's privacy policy and, optionally, a Compact Policy statement, in the following form:
  - *'policyref="http://<domain>/<path to .xml file>", CP="<P3P policy codes>"'*

Prerequisites

The **header** command is only valid in the *ustat-server-map* configuration mode, which requires privileged-level access (see **ip http ustat-server map**).

Occurrence

Use once per map.

Description

Use the **header** command to add a header of the specified type to the response that is sent by a USTAT web server. Currently, CNA only allows you to append a P3P (Platform for Privacy Preferences) header, which requires the **P3P** keyword as part of the command:

The header text can consist of two parts:

- URL to a privacy policy statement in *.xml* format
- Compact Policy (CP) which summarizes your privacy policy in browser-readable codes

For more information about the structure of a policy statement, see the World Wide Web Consortium web site, at *http://www.w3.org*.

The value of *text* should be enclosed in single quotes, with the URL and the policy codes enclosed in double quotes. The entire string between the single quotes will be sent as the contents of a P3P header when the HTTP server responds to a request.

The header will be added to the response sent by the web server on the USTAT module being configured when it sends the single-pixel *.gif* file to a web surfer. In order to be fully P3P compliant, you must add the header to each USTAT on CNA.

Use the **add** keyword to create a new header.

Use the **append** keyword to append to the existing header.

Use the **set** keyword to overwrite the existing header.

Use the **unset** keyword to clear the existing header.

Use the **no** form of the command to remove the P3P header from the web server's response.

help

Displays all keywords that are valid as command-line entries.

Syntax

```
help
?
```

Prerequisites

Available in all modes, all levels of access.

Description

The **help** command is context sensitive—output will be a list of all commands/keywords that are valid entries in whatever mode you are currently in.

history

Displays a list of commands previously entered during the current session.

Syntax

```
history
```

Prerequisites

Available in all modes, all levels of access.

Description

To re-execute a listed command, enter the pound sign (#) followed by the list number.

---

host

Performs a DNS/reverse DNS lookup given a hostname or IP address end point.

Syntax

```
host { <name> | <address> }
```

Arguments

- *name*—hostname, which can consist of letters, numbers, and hyphens (cannot begin or end with a hyphen).
- *address*—IP address of the end point, in dotted decimal format (*a.b.c.d*)

Prerequisites

System operations mode, login-level access.

Default

none

Description

Use the **host** command to determine the IP address, if you know the host name, or to determine the host name if you know the IP address. If there are multiple ip addresses found by the lookup, CNA will show just the first ip address it receives.

---

hostname

Sets the hostname of CNA.

Syntax

```
hostname <name>
no hostname
```

Arguments

- *name*—hostname of the device, which can consist of letters, numbers, and hyphens (cannot begin or end with a hyphen).

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

Default hostname is *CNA*.

Description

Use the **no** form of the command to revert the hostname to its default.

---

hrtt-smoothing-factor (deprecated)

Sets a stabilizing factor used to modify raw traffic-measurement values before storing them in the database.

Syntax

```
hrtt-smoothing-factor <value>
no hrtt-smoothing-factor
```

Arguments

- *value*—integer; range: 0 through 1000

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), both of which require privileged-level access.

Occurrence

Use once globally or once per policy.

Default

global value—50

decision policy—none (uses the global setting)

Description

The value set by the **hrtt-smoothing-factor** command is used to compute the value that gets stored in the database whenever a traffic measurement is recorded. The acceptable range, 0 through 1000, is an integer that represents the half-life of a measurement.

Note:

If CNA is in route-assert mode, the smoothing factor should not be set to 0 except within a decision policy that also includes the **route-assert-filter deny** command. Routes should not be asserted to any subnet for which the smoothing factor is 0.

The specified value represents the number of new measurements required for a given measurement to contribute half of what it is currently contributing to the calculation of the *hrtt* value displayed in reports.

With the value of smoothing-factor set to 20, a given measurement's contribution to the HRTT value that is stored in the database will be cut in half after 20 additional measurements have been received.

If the factor is changed to 10, a newly received measurement's contribution to the HRTT value will be cut in half after 10 additional measurements have been received.

Adjusting this value allows you to determine whether newly made measurements should have more or less influence on CNA's decision-making process.

To give more weight to newer measurements, lower the value. This will make measurements decay more rapidly in the calculation of HRTTs. CNA will then be more sensitive to changes in traffic, resulting in more route assertions.

Use the **no** form of the command to reset the value to its default.

---

ignore-measurements

Prevents CNA decision-maker process from recording measurements for a subnet.

Syntax

```
ignore-measurements
no ignore-measurements
```

Prerequisites

Must be in policy configuration mode (*config-engine-policy*), which requires privileged-level access.

Occurrence

Use once per policy.

Default

Disabled

Description

The **ignore-measurements** command configures the CNA decision-maker process to ignore measurements for subnets associated with a decision policy.

This command can be used to prevent health checks (web hits initiated by you to verify server availability) from being recorded in the database. Health-check measurements can produce unreliable results in some CNA-generated reports and show command output.

Use the **no** form of the command to revert to the default condition.

---

interface

Places CNA into interface configuration mode (*config-if*) for the specified interface.

Syntax

```
interface { ethernet <num>/<port> | tunnel <id> }
no interface { ethernet <num>/<port> | tunnel <id> }
```

Arguments

- *num*—slot number in which module is inserted in a 5000 series device, or subsystem of module in a 3000 series device or generic installation of CNA
- *port*—port number of interface being configured
- *id*—integer identifying the GRE tunnel

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

The **interface** command is used to configure the address of a Ethernet or GRE tunnel interface.

The CNA architecture allows multiple modules to be configured on a single interface.

While in interface configuration mode, you can use any of the following commands to define the interface:

- **description**
- **duplex**
- **ip access-group**
- **ip address**
- **ip mtu**
- **ip route**
- **module**
- **speed**
- **tunnel**

Use the **no** form of the command to remove the interface definition from the configuration (use the **shutdown** command to temporarily disable the interface without removing commands from the configuration).

Ethernet

Ethernet interfaces are identified by slot or subsystem number—the physical interface between card and chassis—and ethernet port number.

Slot numbers correspond to the numbered slots in the chassis into which the removable modules are physically inserted.

GRE Tunnel

GRE tunnels provide a direct connection between a USTAT's Ethernet interface and an interface on a router, thereby avoiding intervening routing issues.

Tunnel IDs must be unique; no two tunnels can have the same *id* value. (In older versions of CNA, tunnels defined on a series 5000 removable module could reuse an ID number of a tunnel configured on a different module; but this method of creating tunnels is no longer possible and cannot be used in any longer. Tunnel IDs must not clash between tunnels created on a USTAT and tunnels created at the interface level.)

After you have created a tunnel with the **interface** command, you can associate it with a USTAT by specifying the tunnel ID as the *interface* argument of the **ip route** command within the USTAT module's configuration block.

Each tunnel must have a unique source/destination address pair. Use the **tunnel** command with the appropriate keyword to define these addresses.

You must define a destination address. If you don't specify a source address, the Ethernet of the tunnel's USTAT module will be assumed.

If multiple USTATs are configured on a single Ethernet interface, you should specify a unique private-space address as the tunnel source when you define the GRE tunnels. Otherwise, the tunnels would all assume the same Ethernet address as their respective source addresses, which is not allowed.

interval-between-packets

Defines the interval between packets that are sent between two endpoints of this active measurement test. The interval is specified in milliseconds.

Syntax

interval-between-packets <value>

no interval-between-packets

Arguments

Argument	Description
<*value*>	Interval in milliseconds between two consecutive packets sent during this active measurement test. Valid ranges are: • RTP test: 10–100 • Ping test: 10–1000 • Traceroute test: 10–100 • TCP test: Not applicable.

Prerequisites

Must be in Chatter Active Measurement Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrences

Use once.

Default

Table 20: Default settings for interval-between-packets command

Type of test	Default interval value
RTP	50
Ping	50
Traceroute	50

Description

Defines the interval between packets that are sent between two endpoints of this active measurement test. The interval is specified in milliseconds.

The no command restores the default value.

ip access-group

Applies an access list to the interface.

Syntax

```
ip access-group <name> {in | out}
no ip access-group [<name>] {in | out}
```

Arguments

- *name*—name of the access list to be applied to the interface

Prerequisites

Must be in interface configuration mode (*config-if*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

No access list.

Description

Use the **`ip access-group`** command to apply an access list (use the **`ip access-list`** command to create the list).

Use the keyword **in** if the access list is to be applied to packets inbound to the interface.

Use the keyword **out** if the access list is to be applied to packets outbound from the interface.

Use the **no** form of the command to remove the reference to an access list from an interface. If no *name* argument is specified, then all access list references will be removed from the interface.

---

ip access-list

Defines an access list and puts the CNA CLI into access-list configuration (*config-acl*) mode.

Syntax

```
ip access-list <name>
no ip access-list <name>
```

Arguments

- *name*—name of the access list

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

None.

Description

The **no** form of the command deletes the named list from the configuration.

ip address

Sets the address and subnet mask of an interface.

Syntax

```
ip address <address> <mask>
no ip address [<address> <mask>]
```

Arguments

- *address*—IP address of the interface, in dotted decimal format (*a.b.c.d*)
- *mask*—subnetwork mask for the interface, in dotted decimal format (*a.b.c.d*)

Prerequisites

Must be in interface configuration (*config-if*) mode, which requires privileged-level access.

Occurrence

Once per interface.

ip domain-list

Defines a list of domain names that will be used to complete unqualified host names.

Syntax

```
ip domain-list <name>
no ip domain-list [<name>]
```

Arguments

- *name*—domain name to be added to the list

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Do not include a leading dot in *name*. See the **ip domain-name** command for syntax.

If both **ip domain-name** and **ip domain-list** commands are used, both will show up in the configuration. But ip domain-name takes precedence. The list will be ignored.

To remove a name from the list, use the **no** form of the command with the domain name to be removed as the argument.

To remove the entire list, use the **no** form of the command with no argument.

---

ip domain-lookup

Enables/disables Domain Name System (DNS) capability, which translates host names into IP addresses.

Syntax

```
ip domain-lookup
no ip domain-lookup
```

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

DNS is enabled.

Description

The **ip domain-lookup** command enables DNS.

Use the **ip name-server** command to identify the DNS servers CNA will use.

---

ip domain-name

Defines the domain to which the device belongs and sets the default domain name to be appended to host names that are not fully qualified.

Syntax

```
ip domain-name <name>
no ip domain-name
```

Arguments

- *name*—name of the default domain

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

There is no default domain name.

Description

Use the **no** form of the command to delete the default domain name.

Do not include a leading dot (the operating system will append "*dot* <domain>" to an unqualified hostname). For example, if you declare the default domain name to be *mydomain.com* and the unqualified hostname is *host*, the CNA operating system will translate the hostname to the following:

```
host.mydomain.com
```

Allowable syntax for *name*, as defined in RFC 1034, is any alphanumeric character, plus the hyphen.

If both **`ip domain-name`** and **`ip domain-list`** commands are used, both will show up in the configuration. But ip domain-name takes precedence. The list will be ignored.

---

ip flow-export destination

Specifies the destination that is to receive a netflow fanout.

Syntax

```
ip flow-export destination <address> <port>
no ip flow-export destination <address>
```

Arguments

- *address*—IP address in the form a.b.c.d
- *port*— udp port

Prerequisites

Must be in EFC configuration mode (*config-efc*), which requires privileged-level access.

Occurrence

Use once or multiple times.

Default

none

Description

Use this command to enable netflow fanout to specific destinations. CNA's EFC server will copy all netflow packets it receives and send them to the destinations you specified.

Use the **no** form of the command to remove a specific destination from the fanout.

---

ip ftp password

Specifies a default password to be used when downloading a configuration file from an FTP server using the **`configure subnet`** or **`copy`** command.

Syntax

```
ip ftp password <pwd>
no ip ftp password
```

Arguments

- *pwd*—the password to be used with FTP

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Description

Use the **`ip ftp password`** command to establish a default password to be used when downloading files from a remote FTP server.

If you specify an FTP URL as an argument to the **configure subnet** or **copy** command without including a password and no default password has been configured with the **ip ftp password** command, CNA will attempt to download the file with no password. If authentication is required, the download attempt will fail.

See also the **ip ftp username** command.

Use the **no** form of the command to remove the default password from the configuration.

ip ftp username

Specifies a default FTP user name to be used when downloading a configuration file from an FTP server using the **configure subnet** or **copy** command.

Syntax

```
ip ftp username <name>
no ip ftp username
```

Arguments

- *name*—user name to be used with FTP

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Description

Use the **ip ftp username** command to establish a default user ID to be used when downloading files from a remote FTP server.

If you specify an FTP URL as an argument to the **configure subnet** or **copy** command without including a user ID and no default user ID has been configured with the **ip ftp username** command, CNA will attempt to download the file with no user ID. If authentication is required, the download attempt will fail.

See also the **ip ftp password** command.

Use the **no** form of the command to remove the default user ID from the configuration.

ip host

Defines host name aliases.

Syntax

```
ip host <name> <address>
no ip host <name>
```

Arguments

- *name*—hostname alias
- *address*—IP address in the form a.b.c.d

Prerequisites

Must be in configuration (*config*) mode, which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

---

ip http password

Specifies a default password to be sent to a remote http server when downloading a configuration file using the **`configure subnet`** or **`copy`** command.

Syntax

```
ip http password <pwd>
no ip http password
```

Arguments

- *pwd*—the password to be used for authentication

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Description

Use the **`ip http password`** command to establish a default password to be used when downloading files from a remote http server.

If you specify an HTTP URL as an argument to the **`configure subnet`** or **`copy`** command without including a password and no default password has been configured with the **`ip http password`** command, CNA will attempt to download the file with no password. If authentication is required, the download attempt will fail.

See also the **`ip http username`** command.

Note:

Unlike other **`ip http ...`** commands in the CNA CLI, which configure the internal CNA web servers, the **`ip http password`** and **`ip http username`** commands configure attributes used when communicating with remote servers.

Use the **no** form of the command to remove the default password from the configuration.

ip http server header append (deprecated)

This command has been replaced by **`ip http ustat-server map`**.

ip http server on

Enables an CNA HTTP server.

Syntax

```
ip http server on
no ip http server on
```

Prerequisites

Must be in configuration or USTAT configuration mode (*config* or *config-ustat*), which requires privileged-level access.

Occurrence

Use per server.

Default

The management module's HTTP server is disabled

Description

CNA has multiple HTTP servers.

The management module's server provides browser access to the management interface. If the server is enabled, you can open the web interface by entering *http://<PathControl hostname or management module eth0 IP address>* in your browser's location finder. The management interface requires the Java 1.4 runtime environment on your client workstation. This server is disabled by default. You must enable it in order to use the management interface.

Each USTAT module also contains an HTTP server, for serving the single-pixel *.gif* used to measure web traffic.

Use the **no** form of the **`ip http server on`** command to disable a server.

ip http server port

Allows you to change the web server port.

Syntax

```
ip http server port <num>
no ip http server port
```

Arguments

- *num*—new port number in the range 1024 through 56788

Prerequisites

Must be in configuration (*config*) mode, which requires privileged-level access.

Default

Port 80

Description

You can set the port to any valid port number in the specified range. Use the **no** form of the command to reset the port to its default value.

When you enter the **`ip http server port`** command in the CLI, CNA will reject any of the well-known ports (0 through 1023), as well as those ports which are reserved by the device for its exclusive use (all ports in the range 56789 through 65535).

Ports in the range 5000 through 49151 are known as “registered” ports and may be claimed by a third-party process used in some future implementation of the CNA operating system.

In addition, the ports in the range 1024 through 4999 are considered by the operating system to be available on demand to meet transient needs.

CNA has multiple processes exchanging data across multiple modules, and multiple users can log in simultaneously via Telnet or SSH and transfer files or initiate cross-module traffic. Each of these events requires one or more of these dynamically allocated ports for varying amounts of time.

You are free to use one of the ports in the dynamic pool, but there is a very small—though still present—chance that the port you choose will have been allocated at the moment the web

server attempts to bind to it. (Once the port is claimed by the HTTP server, it will be removed from the pool and there will be no further chance of conflict as long as the server remains up.)

You can also use one of the “registered” ports as long as it is available.

In the rare instance of a conflict, it is possible that the Medic process will still show the server to be running even though it is not. Currently, the only means by which you can detect a port conflict is to start the web server and attempt to log in to the web interface.

If the attempt fails, issue the **no ip http server on** command to tell Medic to release the server it thinks is running and then restart the server with the **ip http server on** command. Most dynamic port assignments will be temporary and short lived, so the requested port will in all probability be available this time (though this is not guaranteed).

If you change the web server port, you must also define a new access list which allows the management module to use not only the new port, but also all of the ports that were previously enabled on the interface by default.

For details about default access rules, see Access Lists in Chapter 3: Security of the *Monitoring Guide*.

Use the **no** form of the **ip http server port** command to reset the port to its default.

ip http server port ssl

Allows you to change the port number used by the web server for SSL traffic.

Syntax

```
ip http server ssl port <num>
no ip http server ssl port
```

Arguments

- *num*—new port number in the range 1024 through 56788

Prerequisites

Must be in configuration (*config*) mode, which requires privileged-level access.

Default

Port 443

Description

See the description of the **ip http server port** command. The same restrictions on port numbers apply, as does the requirement for an access list.

Use the **no** form of the command to reset the SSL port to its default.

ip http server ssl on

Enables or disables HTTPS on the CNA web server.

Syntax

```
ip http server ssl on
no ip http server ssl on
```

Prerequisites

Must be in configuration mode (*config*) or USTAT configuration mode (*config-ustat*), which requires privileged-level access.

Also, you must have already installed all required SSL certificate and key files, using the **ssl add** command.

Default

HTTPS server is disabled

Description

If you are in configuration mode (*config*), the **ip http server ssl on** command enables SSL on the management module's web server.

When HTTPS is enabled on the management module, the web server provides access to the web interface through a Secure Sockets Layer (SSL) connection at *https://<PathControl hostname>*. Data will be encrypted using SSL certificate and key files which you provide.

You can have both HTTP and HTTPS enabled simultaneously, and data will be transferred according to the appropriate protocol based on the form of URL entered. If you have only HTTPS enabled, the URL *http://<PathControl hostname>* will be redirected to the secure channel and data will be encrypted as though the URL had been entered with *https://*. (You can turn this behavior off with an access list that blocks the port specified for HTTP.)

If you are in USTAT configuration mode (*config-ustat*), the **ip http server ssl on** command enables SSL on a USTAT-SSL module. The command will be ignored by a standard (non-SSL equipped) USTAT module.

If SSL is not enabled, the USTAT-SSL module will only respond to *http://* requests.

If SSL is enabled, the USTAT-SSL module will respond to either *http://* or *https://* requests.

To disable HTTPS, use the **no** form of the command while in the appropriate configuration mode for the module of interest.

See the command **ip http server on** for information about enabling HTTP.

Use the **no** form of the **ip http server on** command to disable HTTP.

ip http username

Specifies a default username to be used when downloading a configuration file from an http server using the **configure subnet** or **copy** command.

Syntax

```
ip http username <name>
no ip http username
```

Arguments

- *name*—username to be used for authentication

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Description

Use the **ip http username** command to establish a default username to be used when downloading files from a remote http server.

If you specify an HTTP URL as an argument to the **configure subnet** or **copy** command without including a user name and no default user name has been configured with the **ip http password** command, CNA will attempt to download the file with no user name. If authentication is required, the download attempt will fail.

See also the **ip http password** command.

Note:

Unlike other **ip http ...** commands in the CNA CLI, which configure the internal CNA web servers, the **ip http password** and **ip http username** commands configure attributes used when communicating with remote servers.

Use the **no** form of the command to remove the default username from the configuration.

ip http ustat-server map

Places CNA into configuration mode for creating a ustat web server map.

Syntax

```
ip http ustat-server map <name>
no ip http ustat-server map <name>
```

Arguments

- *name*—used to identify this configuration map.

Prerequisites

Must be in configuration (*config*) mode, which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **`ip http ustat-server map`** command to create a ustat server map.

Once you execute the **`ip http ustat-server map`** command, you will be in ustat-server-map configuration mode. From there you will be able to set **location**, **content**, and **header**. See these commands for more information.

Use the **no** form of the command to remove the map.

---

ip http ustat-server redirect

Define redirection rules for content served by the ustat web servers.

Syntax

```
ip http ustat-server redirect [ server <server-name> | url <absolute-url> ]
no ip http ustat-server redirect
no ip http ustat-server redirect server <server-name>
no ip http ustat-server redirect url <absolute-url>
```

Arguments

- *server-name*—name of server to be redirected to
- *absolute-url*—redirect requests to an absolute url

Prerequisites

Must be in configuration (*config*) mode, which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

There are three forms of this command:

1. **`ip http ustat-server redirect`** enables dynamic redirection by CNA. Any url of the form: http://<ustat-ip-address>/redirect/<url> will be redirected to <url> by the USTAT.
2. **`ip http ustat-server redirect server <server-name>`** substitutes the <*server-name*> in the incoming request (via 302 response) and redirects it. For example, `ip http ustat-server redirect server www.avaya.com` will redirect an incoming http://www.speedtrak.com/something/test.gif to http://www.avaya.com/something/test.gif.
3. **`ip http ustat-server redirect url <url>`** redirects all requests to a USTAT to the specified absolute url.

Use the **no** form of the above commands to remove the redirection.

---

ip mtu

Sets the maximum transmission unit (MTU) size for the interface.

Syntax

```
ip mtu <bytes>
no ip mtu
```

Arguments

- *bytes*—MTU size in bytes, in the range 68 through 1,500 (see RFC 984)

Prerequisites

Must be in interface configuration mode (*config-if*), which requires privileged-level access.

Occurrence

Use once per interface.

Default

1,500

Description

Use the **no** form of the command to reset the MTU setting to its default.

ip name-server

Adds a DNS server to the list of servers used by CNA to resolve host names for an interface.

Syntax

```
ip name-server <ip-address>
no ip name-server [<ip-address>]
```

Arguments

- *ip-address*—hostname or IP address of the DNS server, in dotted decimal format (a.b.c.d)

Prerequisites

Must be in interface configuration mode (*config*), which requires privileged-level access.

Occurrence

May be used multiple times on each interface (more than one server can be specified for an interface).

Default

None.

Description

The **no** form of the command with an IP address as an argument removes the specified DNS server from the list.

The **no** form of the command with no argument deletes the entire list for the interface.

See the **`ip domain-lookup`** command.

ip route

Defines a static route.

Syntax

```
ip route <network_address> <network_mask> {<next_address> | <interface>}
no ip route <network_address> <network_mask>
            [<next_address> | <interface>]
```

Arguments

- *network_address*—IP address of the destination, in dotted decimal format (a.b.c.d)
- *network_mask*—subnet mask of the destination, in dotted decimal format (e.f.g.h.i
- *next_address*—IP address of the next hop to the destination
- *interface*—name of a specific interface to be used; can be either:
  - *ethernetN/0* (where *N* is slot number on which the interface is configured)
  - *tunnelN* (where *N* is the GRE tunnel number on a USTAT module)

Prerequisites

Must be in configuration, engine, ustat or interface configuration mode (*config*, *config-engine*, *config-ustat*, or *config-if*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **`ip route`** command to create a static route between CNA and a destination.

If you use the command in *config* mode on a 5000 series device, the route will be applied to the management module.

If you use the command while in module configuration mode (*config-engine* or *config-ustat*), the route will be applied to the module being configured.

If you use the command while in interface configuration mode (*config-if*), the route will be applied to the interface.

Note:

On a 3000 series device, you should always configure a static route on the interface assigned to the engine module if you have a globally defined route. This will avoid possible routing problems from a static route being shared between the management and engine modules. See Static Routes in Appendix E: Migrating Configurations Between Models of the *Appendices Guide*.

When interfaces or tunnels are created, you specify them as *ethernet N/N* or *tunnel N* with a space between the name and number portion of the value; when using them as values for the interface argument, they are specified as *ethernetN/N* or *tunnelN*—a single string with no spaces.

While CNA will allow you to define multiple static routes to a single destination, the device does not perform any round-robin, destination-based load optimization. It will use the first reachable default route in the list. If the next hop in that route subsequently goes down, CNA will lose connectivity to remote networks that rely on that route.

If you then bring the failed next-hop interface back up, CNA will then exercise correct fail-over behavior and move to the next route on the list.

Use the **no** form of the command to remove a static route. If neither *next_address* nor *interface* are specified, all matching static routes defined for the specified *network_address* and *network_mask* will be removed.

ip ssh key generate

Generates RSA public/private key pair for secure shell (SSH) communication.

Syntax

```
ip ssh key generate
```

Prerequisites

System operations mode, privileged-level access.

Description

Key generation must always be executed from the command line; commands are not stored in the configuration.

ip ssh key generate cna-agent

Generates a new SSH RSA public/private key pair for SSH connections between the CNA server and CNA Agent Devices.

Syntax

ip ssh key generate cna-agent

Prerequisites

System operations mode, privileged-level access.

Description

Generates a new SSH RSA public/private key pair for SSH connections between the CNA server and CNA Agent Devices. The CNA server uses SSH to remotely execute commands such as show chatter cna-agent-device on the CNA Agent Device.

The CNA server stores up to five SSH fingerprints and uses the most current one to establish SSH sessions with CNA Agent Devices. To view the fingerprints, use the show chatter cna-agent-device command.

ip ssh on

Enables the secure shell (SSH) server on CNA.

Syntax

```
ip ssh on
no ip ssh on
```

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

SSH is disabled.

Description

The **no ip ssh on** command disables SSH.

See also the **ip telnet on** command.

ip telnet on

Enables/disables the telnet protocol.

Syntax

```
ip telnet on
no ip telnet on
```

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Default

The telnet protocol is enabled.

Description

While the telnet protocol is enabled by default, you cannot use it until your configuration meets certain minimum requirements. For more information on the minimum requirements for subnet connectivity, see the *CNA Installation Guide*.

Use the **no** form of the command to disable the telnet protocol.

See also the **`ip ssh on`** command.

---

level

Creates a new row in the a load threshold table and specifies the total maximum load for all links in the row.

Syntax

```
level <max_load>
no level <max_load>
```

Arguments

- *max_load*—total maximum load of all links for the level, in megabits per second (mbps); range: 0 through 1,000,000

Prerequisites

Must be in threshold configuration mode (*config-engine-threshold*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

The **`level`** command establishes a new row in a load threshold table and puts CNA into threshold-level configuration mode (*config-engine-threshold-level*).

While in threshold-level configuration mode, use the **`link`** command to define thresholds for each configured link.

The value of *max_load* must be equal to the sum of the *max_avoidance* values that you specified for all of the links in the row.

You can specify as many levels as you need, but no two levels can have the same total load.

All levels of all threshold tables should contain the same set of links.

See the **load threshold-table** command.

link

Places CNA into service provider access link configuration mode to configure a link; or, associates an existing link with a USTAT module.

Syntax

```
link <name>
link <name>
     <max-avoidance>
     [ start-avoidance percent-below-max <value> ]
no link
no link <name>
```

Arguments

- *name*—name of link; can contain any combination of alpha-numerics, dash, and underscore.
- *max-avoidance*—(only when configuring load optimization threshold tables) level of usage (in megabits per second) at which the probability of a link being removed from the winner set is at its maximum; range: 0 through 1,000,000
- *value*—(only when configuring load optimization threshold tables) low end of avoidance threshold, expressed as a percentage below *max_avoidance*; range: 0 through 100

Prerequisites

The link command is valid on any of the following configuration modes:

- *config-engine*—create a new service provider link or change parameters of an existing link
- *config-engine-linkgroup*—add a configured link to a link group
- *config-engine-threshold-level*—configure parameters in a load threshold table
- *config-ustat*—associate a configured link with a USTAT module

All modes requires privileged-level access.

Occurrence

May be used multiple times, once per service provider access link.

Description

The **link** command is used in five different contexts:

- *config-engine* mode—defines a service provider access link in the CNA configuration
- *config-ustat* mode—associates a configured service provider access link with a USTAT module to associate the specified link with the USTAT module being configured.
- *config-engine-threshold-level* mode—defines load optimization parameters for a configured link.
- *config-engine-linkgroup* mode—associates a configured link with a link group, for active-measurement probes
- *config-engine-billing* mode—associates one or more configured links with a billing link

In *config-engine* mode, the **link** command puts CNA into link configuration mode (*config-engine-link*). Use the **provider-as** command to set the link's AS number and IP address and the optional **description** and **penalty** commands to set a descriptive test field and weight links relative to each other (see **winner-set-width** command).

When adding links to the CNA configuration, you must fully configure a link. The CNA *decision-maker* process will not start properly if your configuration includes a partially configured link. If the system is booted with an incomplete link, you should use the **medic call** command to manually restart the decision-maker process when the link has been fully configured.

In general, exercise the following precautions with respect to link configurations:

- No two link definitions in the engine module should have the same peering address.
- A link reference in a USTAT blade should always point to a well-defined link in the engine module.
- All USTAT blades should have a VIP defined.
- The link reference in a neighbor definition should point to a well-defined link in the engine module.
- The group of links that are referred to in all neighbor definitions in the engine module should be the same as the group of all links defined in the engine module. (Note that neighbors can still be linkless, that is, not all neighbors must have a link reference.)

When configuring load threshold tables, the following restrictions should be observed:

- Polling usage must be enabled for all links referenced in threshold tables.
- Polling usage must *not* be enabled for any configured links that are *not* referenced in the tables.
- All tables must reference the same set of links.

If multiple USTATs are configured to point to the same link (a means of increasing measurements for an important set of subnets, as in a VPN), the threshold configured by the **outage-detection active-probe** command will be divided among the USTATs. For example, if you configure the service provider link *provider_a* to be associated with both *ustat_x*

and *ustat_y* and you have set your probe-failure threshold to 20, the CNA *decision-maker* process will declare a subnet outage after 10 probe failures on each of the two USTATs.

Use the **no** form of the command without a link name argument while in *config-ustat* mode to remove the association between the USTAT module and its associated link.

Use the **no** form of the command with a link name argument while in *config-engine-threshold-level* mode to remove the link from the load threshold table row.

Use the **no** form of the command with a link name argument while in *config-engine-linkgroup* mode to remove the link from the link group.

link-group

Defines a link group and places CNA in *link-group* configuration mode.

Syntax

```
link-group <name>
no link-group <name>
```

Arguments

- *name*—name of the link group

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **`link-group`** command to define a new link group or edit the configuration of an existing group.

Use the **no** form of the command to delete the specified link-group from the CNA configuration. Specify the name of the link-group to be deleted.

link-measurement-timeout

Determines the maximum length of time links can be silent before being considered invalid.

Syntax

```
link-measurement-timeout <seconds>
no link-measurement-timeout
```

Arguments

- *seconds*—maximum time to wait, in seconds; range: 60 through 86400 (1-24 hours)

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Use once.

Default

300

Description

If no measurements are received for any subnet over a link for the length of time specified by seconds, CNA will consider the link to be invalid.

Invalid links are identified by the label *i* in **show prefix (deprecated)** output and web interface snapshot reports. You can also use the **`show link`** command to check the status of a link, and the **`show statistics am-measurer`** or **`show statistics edgestats`** command to see the status of CNA.

Use the **no** form of the command to reset the timeout limit to its default.

load evaluation-interval

Sets a required time interval between load samples used for load optimization.

Syntax

```
load evaluation-interval <min> <mean-cold> <mean-hot> <max>
no load-evaluation
```

Arguments

- *min*—minimum number of SNMP samples to skip before re-evaluating a link; range: 1 through 100

- *mean-cold*—multiplier used to calculate the random variable for cold links; range: 1 through 100
- *mean-hot*—multiplier used to calculate the random variable for hot links; range: 1 through 100
- *max*—maximum number of samples to skip; range: 1 through 100, with *max* – *min* >= 3

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once globally, or once per policy block.

Default

min—3

mean-cold—5

mean-hot—6

max—10

Description

The **`load evaluation-interval`** command sets multipliers which will be used by the CNA decision-maker process to spread out over time successive evaluations of the state of a link for load optimization. Forcing the decision maker process to choose a different interval at random for each subnet prevents flapping of the system (mass movements of subnets on or off a link in unison because calculations for each subnet are based on the same data points).

The interval is expressed as a multiplier of usage samples obtained by the CNA link-monitor process.

See the **`polling usage on`** command for more information about polling.

The actual multiplier used by the decision maker process is selected at random from an exponential distribution. The multiplier is selected at random in an effort to avoid clashes in time among subnets as the decision-maker process works through its queue of decisions to be made.

The *min* argument establishes a minimum amount of time—a multiple of the polling interval—to wait before re-evaluation of the link. You should not allow the decision-maker process to use consecutive load samples to evaluate a subnet. Using different load data points for different subnets helps avoid system flapping.

The *max* argument sets a maximum amount of time to be skipped before re-evaluation of the link, in order to avoid selection of an interval far in the future (because of the long tail of the exponential distribution from which the random interval is selected).

Be sure the difference between *min* and *max* is at least 3.

The *mean-cold* argument is a multiplier used by the decision maker process to compute a random interval for re-evaluating cold links.

The *mean-hot* argument is a multiplier used by the decision-maker process to compute a random interval for re-evaluating hot links.

The *mean-hot* argument should be larger than the *mean-cold* argument, to give hot links more time to cool off while still paying sufficient attention to cold links.

See also the following commands:

- **load min-required-samples**
- **load optimization**
- **load threshold-table**

Use the **no** form of the command to set all arguments to their defaults.

load min-required-samples

Sets the minimum number of usage samples that must be obtained from a link before load optimization will be used in the selection of a performance route.

Syntax

```
load min-required-samples <value>
no load min-required-samples
```

Arguments

- *value*—number of samples required; range: 1 through 100

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once globally, or once per decision policy block.

Default

3

Description

Use the **load min-required-samples** to set the number samples that must be obtained for a link before load optimization can be used after a system reboot, for example, or execution of the **clear prefix (deprecated)** command.

Use the **no** form of the command to reset *value* to its default.

---

load optimization

Enables load optimization, which factors a link's current usage load into the performance-route decision-making process.

Syntax

```
load optimization {enable | disable}
no load optimization
```

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Use once globally, or once per decision-policy block.

Default

Enabled

Description

When **load optimization** is enabled, current data rates through your service provider links will be evaluated against bandwidth thresholds during the performance route selection process. (If SNMP polling is not enabled, the minimum required load samples will never be reached and load optimization will never be used. See the **polling usage on** command for information about SNMP polling.)

Bandwidth thresholds are set with the **load threshold-table** command. Load optimization requires that a load threshold table be configured.

Minimum required samples is set with the **load min-required-samples** command.

The **no** form of the command reverts to the default status.

See also the **load evaluation-interval** command.

load threshold-table

Sets usage thresholds and direction of data flow for load optimization.

Syntax

```
load threshold-table <name>
no load threshold-table <name>
```

Arguments

- *name*—name of the table being defined

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **load threshold-table** command to create or modify a table of values used by the CNA decision maker listener process to compute the probability that a link will be removed from the winner set due to load optimization.

If you enter the **load threshold-table** command with a name not currently in use, a new table will be created and CNA will be placed in threshold configuration mode.

If you enter the **load threshold-table** command with the name of an existing table, CNA will be placed in threshold configuration mode for that table.

For more information about the load threshold table and the load optimization process, see Chapter 2: WAN Cost and Load Optimization of the *Adaptive Path Controller Guide*.

While in threshold configuration mode, you can use the following commands:

- **direction** - specifies which data to use for evaluating load (inboud or outbound)
- **level** - creates a new row in the threshold table
- **on-failure** - enters on-failure mode, allowing you to enter the links in the failover group
- **start-avoidance percent-below-max** - specifies the starting threshold

After creating or changing parameters in a threshold table, use the **resync load threshold-table** command to commit the table to the CNA *decision-maker* database.

Use the **load use-threshold-table** command in engine configuration mode (*config-engine*) to establish a table as the default global threshold table. This is required to use load optimization.

Use the **load use-threshold-table** command in policy configuration mode (*config-engine-policy*) to associate a table with a specific decision policy.

Use the **no** form of the command to remove a load-threshold table from the CNA configuration (requires a *decision-maker* process resynchronization with the **resync load threshold-table** command).

load use-threshold-table

Identifies a load threshold table to be used for a specific set of subnets.

Syntax

```
load use-threshold-table <name>
no load use-threshold-table
```

Arguments

- *name*—name of a configured load threshold table

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which require privileged-level access.

Occurrence

Use once globally, or once per decision policy.

Default

For decision policies, use the globally defined table.

Globally, there is no default—a configured table is required if load optimization is enabled.

Description

The **load use-threshold-table** command specifies which configured table to use globally or for a specific decision policy. (You must use the **resync load threshold-table** command first; an attempt to use a table that has not been committed will be rejected.)

For more information about the load optimization process, see Chapter 2: WAN Cost and Load Optimization in the *Adaptive Path Controller Guide*.

Also see the **load threshold-table** command.

location

Specifies the location that the map configuration will be applied to.

Syntax

```
location <string>
no location
```

Arguments

string—this will be matched to the URL being requested of CNA.

Prerequisites

The **`location`** command is only valid in the *ustat-server-map* configuration mode, which requires privileged-level access (see **ip http ustat-server map**).

Occurrence

Used once per map.

Description

Use the **`location`** command to specify the location string that CNA's USTAT web server will compare to incoming web requests. If a request matches the location string, then it will be handled as specified by the **`content`** and **`header`** commands.

The string must begin with a "/" and may contain the wildcard characters:

- * —will match any number of characters
- ? —will match a single character

Use the **no** form of the command to remove the location in the current *ustat-server-map*.

logging

Specifies hosts to receive log messages.

Syntax

```
logging <address>
no logging <address>
```

Arguments

- *address*—target machine's hostname or IP address in a.b.c.d format

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Messages generated by CNA are logged to a buffer. See the **logging buffered (deprecated)** command for a description of the buffer.

You can also direct log messages to one or more remote hosts that have been configured to receive Syslog messages, using the **logging** command.

Use the **no** form of the command to stop sending messages to a remote host.

Also see the **logging trap** command.

---

logging buffered (deprecated)

Note:

Starting with CNA version 4, CNA will automatically set the buffer size to 50000 (50megabytes).

Sets the size of the buffer used by CNA to store log messages.

Syntax

```
logging buffered <size>
no logging buffered
```

Arguments

- *size*—size of the buffer in kilobytes (KB=1024 bytes); range: 100 through 50000

Prerequisites

Must be in configuration mode (*config*), which requires privileged level access.

Occurrence

Use once.

Default

5,000 KB

Description

You can set the size of buffer used by CNA for holding log messages. The maximum size of the buffer is 50 MB, though you should not keep it that large for an extended period. If you increase the buffer to its maximum size for a specific purpose, you should lower it once the specific need has been met. Avaya recommends a buffer no larger than 100 MB for long-term use. Use a remote syslog host for storing large quantities of log messages.

Use the **logging** and **logging facility** commands to send log messages to a remote syslog server. The **logging on** command enables/disables logging.

The CNA log buffer is circular; when it is full, each new message will overwrite the oldest message in the buffer.

Use the **no** form of the command to revert to the default buffer size.

---

logging console

Sets a threshold level for controlling which messages are sent to the console.

Syntax

```
logging console <level>
no logging console
```

Arguments

- *level*—message level; must be one of the following (listed in descending order):
  - *emergency*
  - *alert*
  - *critical*
  - *error*
  - *warning*
  - *notice*
  - *informational*
  - *debug*

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

emergency

Description

When you set a message level, only messages at that level or above (in the order listed) are sent to the console. For example, if you set the level at *emergency*, only emergency messages will be sent to the console. If you set the level at *error*, then all *error*, *critical*, *alert*, and *emergency* level messages will be sent to the console.

Set the level to *debug* to have all messages that are sent to the buffer to also be sent to the console.

If you execute the **no** form of the command, *level* is reset to its default value.

---

logging facility

Sets the Syslog facility to be used by CNA when logging messages to remote hosts.

Syntax

```
logging facility <name>
no logging facility
```

Arguments

- *name*—facility to be used; can be one of the following strings:
  - *auth*—authorization system
  - *cron*—cron facility
  - *daemon*—system daemon
  - *kern*—kernel
  - *local<n>*—locally defined; *n* is an integer in the range 0 through 7, such as *local0* or *local7*
  - *lpr*—line printer

- *mail*—mail system
- *news*—USENET news
- *sys<n>*—system use; *n* is an integer in the range 9 through 14, such as *sys9* or *sys14*
- *syslog*—system log
- *user*—user process
- *uucp*—UNIX-to-UNIX copy system

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

Default facility is *local0*.

Description

The **no** form of the command resets the logging facility to its default value.

CNA uses its own logging process, logfilterd, to write logging messages. However, these messages conform to the Syslog standard and can be written to remote hosts for archiving.

See the **`logging`**, **`logging buffered`**, **`logging console`**, and **`logging on`** commands.

---

logging host

Specifies the Syslog servers to which Syslog messages for this alarm template are sent.

Syntax

logging host <host-id>

no logging host <host-id>

Argument

Table 21: logging host arguments

Argument	Description
<host-id>	IP address of the host or the host name of the Syslog server. Enter one of the hosts configured with the logging-server command.

Prerequisites

Must be in Alarm Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Up to four hosts can be configured per alarm template.

Description

Specifies the Syslog servers to which Syslog messages for this alarm template are sent. Syslog messages are sent for alarms originating on a zone or edge that is assigned this alarm template. You can specify any or all of the hosts configured with the logging-server command.

The no command disables Syslog messages to this host for this alarm template.

logging on

Enables logging of messages.

Syntax

```
logging on
no logging on
```

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

Logging is enabled.

Description

When logging is enabled, all messages will be logged to Syslog and can be forwarded to a remote host (see the **logging** and **logging trap** commands).

Use the **no** form of the command to disable logging.

logging-server

Forwards a Syslog message to the host specified when Chatter generates an alarm.

Syntax

logging-server host <host-id> [port <port-number>] [facility <facility-name>] [severity <severity-level>]

no logging-server host <host-id>

Arguments

Table 22: logging-server arguments

Argument	Description
<host-id>	IP address of the Syslog server to which Chatter sends Syslog messages. Enter the IP address in the format of a.b.c.d.
<port-number>	The UDP port number that the Syslog server receives Syslog messages on. The default value is UDP port 514.

1 of 2

Table 22: logging-server arguments (continued)

Argument	Description
<facility-name>	The Syslog facility that Chatter includes in its Syslog messages. Options are: • auth—authorization system • cron—cron facility • daemon—system daemon • kern—kernel • local<n>—locally defined; n is an integer in the range 0–7, such as local0 or local7 • lpr—line printer • mail—mail system • news—USENET news • sys<n>—system use; n is an integer in the range 9–14, such as sys9 or sys14 • syslog—system log • user—user process • uucp—UNIX-to-UNIX copy system The default facility is local0.
<severity-level>	Severity level to include in the message.
	2 of 2

Prerequisites

Must be in Chatter Configuration command mode.For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Up to four logging server hosts can be configured for the Chatter module.

Default

Table 23: Default settings for logging-server command

Argument	Default value
<port-number>	514
<facility-name>	local0
<severity-level>	warning

Description

Forwards a Syslog message to the host specified when Chatter generates an alarm.

The no command disables Syslog messages to the specified host.

---

logging trap

Sets limit on level of log messages sent to remote hosts.

Syntax

```
logging trap <level>
no logging trap
```

Arguments

- *level*—level of log messages; can be one of the following (listed in descending order):
  - *emergency*
  - *alert*
  - *critical*
  - *error*
  - *warning*
  - *notice*
  - *informational*
  - *debug*

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

Messages at the warning level and above are sent to remote hosts.

Description

When logging is on, all messages at or above the specified level in the list defined under the **Arguments** heading will be sent to designated remote hosts.

Specify *debug* level to send all log messages to the remote host.

Use the **no** form of the command to reset the logging level to the default.

logout

End the current CNA session.

Syntax

```
logout
```

Description

The **logout** command logs you out of the current CLI session, regardless of which operating level or configuration mode you are in.

See also the **end** and **exit** commands.

ls

Lists persistent objects in the CNA storage repository.

Syntax

```
ls [<extension>]
```

Arguments

- *extension*—filename extension of files to list

Prerequisites

System operations mode, login-level access.

Description

Certain objects are stored as files in a persistent repository on CNA, including configuration files, snapshot report files, and persistent output from various **show** commands. Use the **ls** command to see a list of filenames; the **delete** command to permanently remove an object; the **rename** command to rename an object; and the **copy** command to make a copy.

The **ls** command by itself will list all objects, or you can filter the list by specifying one of the following as *extension*:

- *cfg*—configuration files
- *gif*—gif images
- *html*—Internet files
- *img*—CNA images
- *name*—subnets alias mapping files
- *subnet*—bgp subnets (used for customer-based reports)
- *qry*—trend-report query files (defines data to be retrieved for trend-report graphs)
- *tmpl*—template files for trend reports
- *txt*—text files

Trend report filenames displayed on the command line (file extension *.qry*) may differ from the filenames displayed in the web interface, which allows the use of spaces and the slash character in trend report names (but not when naming other file types).

These characters will be translated into a commonly accepted URL encoding scheme when the file is written to disk. The web interface translates the filename back to an unencoded format when displaying the filename, but the filenames are shown in their encoded form on the command line.

For example, a trend report listed as *trend report/1* in the list of trend reports in the web interface will appear as *trend_report=2F1.qry* on the command line. The template file listed as *Bandwidth (5 Day) Template* in the web interface appears as *Bandwidth_=285_Day=29Template.tmpl* on the command line.

max-ttl

Defines the maximum TTL (in hops) for traceroute tests.

Syntax

max-ttl <value>

no max-ttl

Arguments

Table 24: max-ttl arguments

Argument	Description
<*value*>	Maximum TTL (in hops) for traceroute tests. The valid range is 1–31.

Prerequisites

Must be in Chatter Active Measurement Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use once.

Default

9 hops.

Description

Defines the maximum TTL (in hops) for traceroute tests.

The no command restores the default value of 9 hops.

measurement intra-zone disable

Disables intra-zone tests for agent-based measurement.

Syntax

measurement intra-zone disable

no measurement intra-zone disable

Prerequisites

Must be in Chatter Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Default

Enabled.

Description

Disables intra-zone tests for agent-based measurement. The no form of this command enables intra-zone tests.

Intra-zone tests are tests run between two CNA agents in the same zone or subzone. These tests have the same zone or subzone for their source and destination. By default, intra-zone testing is enabled.

---

medic call

Force immediate execution of the specified Medic routine.

Syntax

```
medic call [<slot> | all] <service> <routine>
```

Arguments

- *slot*—CNA blade slot number
- *service*—name of Medic service
- *routine*—name of the Medic routine

Prerequisites

System operations mode, privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

See Table 4: System services, routines, and parameters in the section Medic Services of Chapter 1: Introduction to the Command Line Interface for names of services and parameters, and default values.

---

medic set

Sets or clears a Medic parameter.

Syntax

```
medic set [<slot> | all] <service> <parameter> <value>
no medic set [<slot> | all] <service> <parameter>
```

Arguments

- *slot*—slot number of CNA module
- *service*—name of Medic service
- *parameter*—name of parameter
- *value*—value of parameter

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

The **no** form of the command resets the specified parameter to its default value.

If *slot* is specified, the values will apply only to the module that occupies that slot number.

If the keyword **all** is specified, the values will apply to each module in CNA.

If neither *slot* nor the keyword **all** is specified, the values will apply to the management module.

See Table 4: System services, routines, and parameters in the section Medic Services of Chapter 1: Introduction to the Command Line Interface for names of services and parameters, and default values.

Parameters that establish minimums for ranges can be set between -1 and the range's current maximum.

Parameters that establish maximums can be set between the range's current minimum and 101.

For example, the default range for the CPU resource on the management module is 50 through 80%.

That means you can reset the *max-pcpu* parameter on the management module to a value in the range 50 through 101, and the *min-pcpu* parameter can be set to a value of -1 through 80.

Once you change either of these values, the allowable range for the other parameter will change accordingly. If you attempt to set a value outside of its allowed range, an error message will inform you of the correct range.

Since these parameters represent percentages, the actual maximum range of each is 0 through 100%. However, the allowable range is -1 through 101.

If you never want an alarm to be triggered, set the maximum value to 101. Resource usage cannot exceed 100 percent; the maximum will never be reached and the alarm will never be set off.

And, if you never want the alarm trigger to be reset, set the minimum value to –1. Resource usage cannot be less than 0; the minimum will never be reached and the alarm will never be reset.

min-required-load-samples

Sets minimum sample threshold for load optimization.

Syntax

```
min-required-load-samples <value>
no min-required-load-samples
```

Arguments

- *value*—minimum number of samples required before load optimization can be applied to route selection process; range: 1 through 100

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Use once.

Default

3

Description

The **`min-required-load-samples`** command sets the minimum number of usage samples required before the CNA decision-maker process can apply load optimization in evaluating links.

Use the **no** form of the command to revert to the default number of samples.

See the **`load optimization`** command.

min-required-measurements

Sets the minimum number of measurements required by the Decision Maker process before a subnet can be considered *sufficiently measured* to be a candidate for updating.

Syntax

```
min-required-measurement <value>
no min-required-measurement
```

Arguments

- *value*—number measurements; range: 1 through 65535

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Use once.

Default

12

Description

The value you set establishes a minimum number of traffic test measurements required before CNA will determine a best path for a subnet.

If too many subnets are labeled I (insufficient data) when you execute the **`show subnet`** command, decrease the **`min-required-measurements`** value. This will allow CNA to make more best-path determinations.

Because of the way the CNA software calculates the HRTT value that is stored in the database, a very low minimum measurement value could incorrectly give too much influence to the first few measurements. Avaya recommends you set this value no lower than 4, except in test or demonstration environments.

The **no** form of the command resets the minimum to its default value.

module

Places CNA into configuration mode for the specified CNA module, or associates a named USTAT with an interface.

Syntax

```
module <type> [ <name> ]
no module <type> [ <name> ]
```

Arguments

- *type*—an CNA module; can be one of the following:
  - *engine*
  - *reporting*
  - *ustat*
  - efc
- *name*—name of the module (USTATs only, in configurations written for version 2.2 or later)

Prerequisites

Must be in configuration or interface configuration mode (*config* or *config-if*), which requires privileged-level access.

Occurrence

One instance per module in a configuration or one instance per interface, but you can use the command multiple times in order to enter module-configuration mode to edit a module's configuration block.

Description

Use the **module** command in configuration mode (*config*) to identify an CNA module to be configured.

Use the **module** command in interface configuration mode (*config-if*) to associate a named USTAT with an interface.

Once you execute the **module** command, you will be in a specific configuration mode (such as *config-engine* or *config-ustat*). You must exit that mode before issuing the **no** command, which requires you to be in configuration mode (*config*).

There is no *management* module type; the management module is implied.

Use the *name* argument with USTATs only. The name may contain dashes, underscores and alpha-numeric characters, but the first character must be a letter of the alphabet (a-z or A-Z).

Note:

In versions earlier than 2.2, all modules required a numerical argument matching the 5000-series chassis slot in which the removable module is installed. In order to provide backward compatibility with existing configurations, slot numbers are still allowed in **module engine** and **module reporting** commands, but the argument will be ignored.

On a 5000 series device, names can be numbers matching the slot in which the removable USTAT is inserted. This is a legacy feature to allow compatibility with configurations that predate version 2.2.

If you choose to identify USTATs by names instead of by slot numbers on 5000 series devices, you must use the **module** command in interface configuration mode in order to associate the USTAT with an interface.

See the **interface** command.

In spite of this required additional step to associate USTATs with interfaces, you should consider migrating your USTATs to names instead of slot numbers for upward compatibility with future CNA releases and CNA hardware upgrades.

If you have created a GRE tunnel, you can associate it with a USTAT module by specifying the tunnel ID as the interface argument of the **ip route** command while in config-ustat mode. When done, your configuration should contain a command block similar to this:

```
module ustat myUstat
  ...
  ip route 0.0.0.0 0.0.0.0 tunnel 1
  ...
  end
```

Use the **show version** or **show medic all** commands to see current slot assignments.

Use the **no** form of the command to remove all custom settings on the module (all settings will revert to defaults).

module chatter

Enters Chatter Configuration command mode.

Syntax

module chatter

Prerequisites

Must be in Configuration command mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Available for use multiple times in the configuration.

Description

Opens the module block to include all configuration statements related to Chatter.

Note:

The previous **module chatter** command (see module chatter or testplug) is entered from Interface Configuration mode and associates an interface with the Chatter module or CNA agent.

---

module chatter or testplug

Associates an interface with the CNA server's Chatter module or embedded CNA agent.

Syntax

module chatter | testplug

no module chatter

no module testplug

Prerequisites

Must be in Interface Configuration command mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Available for use multiple times in the configuration.

Default

Interface **Ethernet 0/0** is the default interface for Chatter and CNA server's embedded CNA agent.

Description

Enables the CNA server's Chatter module or embedded CNA agent on a specific interface.

Important:

The Chatter module and embedded CNA agent must be associated with an Ethernet or Fast Ethernet interface. You cannot associate the Chatter module or embedded CNA agent with a tunnel interface.

monitor enable

Enable the monitor process on the EFC to receive SPAN data.

Syntax

```
monitor enable
no monitor enable
```

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

Monitor is not enabled

Description

Use the **`monitor enable`** command to allow the optional EFC module to receive SPAN data.

Use the **no** form of the command to restore the default.

monitor interface

Specifies the interfaces on which SPAN data will be received.

Syntax

```
monitor interface [ ethernet0 | ethernet1 | gigabitEthernet0 | gigabitEthernet1 ]
no monitor interface [ ethernet0 | ethernet1 | gigabitEthernet0 | gigabitEthernet1]
```

Prerequisites

Requires the optional EFC module. Must be in EFC configuration mode (*config-efc*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

none

Description

Use the **monitor interface** command to set the interfaces used by the optional EFC module to receive SPAN data. If this is not specified, then all interfaces will be monitored by the EFC.

Use the **no** form of the command to remove an interface previously specified.

See the **interface** command for information about assigning an IP address to the interface.

See the **netflow interface** command for information about receiving Netflow data.

---

more

Displays the contents of a specified file.

Syntax

```
more <url>
```

Arguments

- *url*—URL identifying the file

Prerequisites

System operations mode, privileged-level access.

Description

The URL can refer to a file anywhere on the subnet. The URL can contain a username and password if authentication is required. See the description of the **copy** command for details.

---

neighbor allow-next-hop-self (removed)

Note:

Starting with CNA version 5, this command has been removed completely (see neighbor uses-next-hop-self-address).

Provides CNA with a next-hop address when routers are configured with *next-hop-self*.

Syntax

```
neighbor <address> allow-next-hop-self <next-hop>
no neighbor <address> allow-next-hop-self
```

Arguments

- *address*—IP address of the BGP neighbor, in dotted decimal format (a.b.c.d)
- *next-hop*—next-hop address of the edge router, in dotted decimal format (a.b.c.d)

Prerequisites

Must be in BGP configuration mode (*config-engine-bgp*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

Edge routers cannot be configured with next-hop-self

Description

Use the **`neighbor allow-next-hop-self`** command to identify the IP address to be used by CNA when asserting routes to edge routers configured with next-hop-self.

When *next-hop-self* is configured on a router, CNA will attempt to match the router's advertised next hop to a *next-hop* argument configured by a **`neighbor allow-next-hop-self`** command in one of the configuration's BGP neighbor blocks.

When a match is found, CNA will set the next-hop address to the facing address of the link whose AS number (declared when the link was created with the **`link`** command) matches the first AS number in the route's AS path.

To be sure these facing addresses are included in the routing tables on each of your edge routers, you can redistribute them into IGP or you can create static routes on each edge router for all non-connected service providers.

Use the **no** form of the command to remove the *next-hop-self* mapping from the BGP neighbor configuration block.

---

neighbor buffered (removed)

Note:

Starting with CNA release 5.0, this command has been removed.

Creates a buffer for receipt of BGP update tables.

Syntax

```
neighbor <address> buffered [ <size> ]
no neighbor <address> buffered
```

Arguments

- *address*—IP address of the BGP peer
- *size*—size of buffer, in megabytes; range: 1 through 100 (recommended minimum is 10)

Prerequisites

Must be in BGP configuration mode (*config-engine-bgp*), which requires privileged-level access.

Occurrence

Use once per BGP peer.

Default

30

Description

Use the **neighbor buffered** command to force CNA to place all BGP update data into a special CNA-configured buffer before processing the table.

This buffer should not be necessary under most circumstances and will not speed up a BGP peering that is functioning correctly.

You should consult a Avaya support engineer before using this command.

Use the **no** form of the command to disable this CNA-configured BGP peering buffer.

---

neighbor description

Sets descriptive text for a BGP peer.

Syntax

```
neighbor <address> description <text>
no neighbor <address> description
```

Arguments

- *address*—IP address of the BGP peer, in dotted decimal form (a.b.c.d)
- *text*—text of the description, which can consist of letters, numbers, space, hyphen, period, comma, colon, underscore, parenthesis, forward slash, and “commercial at” - @.

Prerequisites

Must be in BGP configuration mode (*config-engine-bgp*), which requires privileged-level access.

Occurrence

Use once per BGP peer

Default

No text.

Description

Use the **no** form of the command to delete the descriptive text.

---

neighbor link

Binds a service provider access link to a BGP neighbor.

Syntax

```
neighbor <address> link <name>
no neighbor <address> link <name>
```

Arguments

- *address*—IP address of the BGP neighbor, in dotted decimal format (a.b.c.d)
- *name*—name of the service provider link (see **`link`** and **`show link`** commands)

Prerequisites

Must be in BGP configuration mode (*config-engine-bgp*), which requires privileged-level access.

Occurrence

Use once.

Description

Use the **no** form of the command to remove the *neighbor link* line from the configuration.

Link names are used as database keys when retrieving data from the trending database. Removal of links from the configuration, or renaming them, after the trending database has accumulated data can result in confusing trend-report graphics. See Y Axis Key Lists in Chapter 3: Reports of the *Operator Guide*.

---

neighbor redirect-peering

Redirects a BGP peering to a different IP address and port, for testing the CNA image without multiple routers.

Syntax

```
neighbor <address> redirect-peering <new-address> <port>
no neighbor <address> redirect-peering
```

Arguments

- *address*—IP address of the BGP peer
- *new-address*—new IP address to be used for testing
- *port*—port on *new-address* device to be used for testing; range: 1024 through 65535

Prerequisites

Must be in BGP configuration mode (*config-engine-bgp*), which requires privileged-level access.

Occurrence

Use once per BGP peer.

Description

When testing an CNA image, it is often useful to test multiple BGP peerings without dedicating multiple routers to the test.

Avaya engineers use the **`neighbor redirect-peerings`** command to redirect the configured peerings to multiple ports on a single interface.

The **no** form of the command removes the redirection from the peering.

neighbor remote-as

Adds an entry to the BGP neighbor table.

Syntax

```
neighbor <address> remote-as <number>
no neighbor <address>
```

Arguments

- *address*—IP address of the BGP neighbor, in dotted decimal format (a.b.c.d)
- *number*—neighbor's Autonomous System (AS) number; integer in the range 1 through 65,535

Prerequisites

Must be in BGP configuration mode (*config-engine-bgp*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

The **`neighbor remote-as`** command creates a new BGP neighbor block to the running configuration.

Use the **no** form of the command to clear all lines in the block from the configuration.

neighbor report-alias

Associates an alias with a BGP peer.

Syntax

```
neighbor <address> report-alias <alias>
no neighbor <address> report-alias
```

Arguments

- *address*—neighbor's IP address, in dotted decimal form (a.b.c.d)
- *alias*—alias to be used for neighbor

Prerequisites

Must be in BGP configuration mode (*config-engine-bgp*), which requires privileged-level access.

Occurrence

Use once.

Description

Use the **no** form of the command to remove the alias.

The alias is used by CNA when generating reports.

---

neighbor shutdown

Disables a BGP neighbor association.

Syntax

```
neighbor <address> shutdown
no neighbor <address> shutdown
```

Arguments

- *address*—neighbor's IP address, in dotted decimal format (a.b.c.d)

Prerequisites

Must be in BGP configuration mode (*config-engine-bgp*), which requires privileged-level access.

Occurrence

Use once.

Description

The **shutdown** command renders the interface inactive without removing it form the configuration.

Use the **no** form of the command to remove the **neighbor shutdown** command from the running configuration, thereby re-enabling the BGP neighbor association.

neighbor summarize-bgp-events

Defines the rate at which BGP announcement events are sent to the CNA *event-manager* process.

Syntax

```
neighbor <address> summarize-bgp-events <interval>
no neighbor <address> summarize-bgp-events
```

Arguments

- *address*—IP address of the BGP neighbor
- *interval*—interval between events, in seconds; range: 0 through 86400

Prerequisites

Must be in BGP configuration mode (*config-engine-bgp*), which requires privileged-level access.

Occurrence

Use once per BGP neighbor.

Default

300

Description

The **`neighbor summarize-bgp-events`** command determines how frequently BGP-announcement events are sent to the CNA *event-manager* process, which in turn sends event packets to the **Control Panel** for display in the web interface **Event Log**.

If the interval is set to 0, all BGP announcements are sent to the *event-manager* process as they occur. When this happens, the volume of BGP announcements can overwhelm events from the *decision-maker* process in the **Event Log**.

When the interval is set at some number greater than zero, the *event-manager* process will receive (and pass on to the **Control Panel**) summary events, which report the number of BGP announcements that were sent during the interval.

For example, if the **neighbor summarize-bgp-events** *interval* argument is 60, then the **Event Log** in the **Control Panel** will display one BGP summary event at the end of every minute, instead of potentially hundreds of individual BGP announcement events being logged as the minute passes.

A summary event will always be sent at the end of the interval, even if there were no BGP announcements.

The **no** form of the command will revert to the default.

neighbor update-source

Associates the named interface with a BGP session endpoint.

Syntax

```
neighbor <address> update-source <interface>
neighbor <address> update-source
```

Arguments

- *address*—IP address of the BGP neighbor
- *interface*—name of the CNA-configured interface; can be:
  - *ethernet0*
  - *ethernet1*

Prerequisites

Must be in BGP configuration mode (*config-engine-bgp*), which requires privileged-level access.

Occurrence

Use once per BGP neighbor.

Description

The **`neighbor update-source`** command associates the IP address of *interface* with the BGP neighbor's IP address.

You should also use the CNA **`ip route`** command to direct packets that are bound for the BGP neighbor through the named interface.

Use the **no** form of the command to disable the command for the BGP peering.

neighbor uses-next-hop-self-address

Provides CNA with a next-hop address when routers are configured with *next-hop-self*.

Syntax

```
neighbor <address> uses-next-hop-self-address <next-hop>
no neighbor <address> uses-next-hop-self-address
```

Arguments

- *address*—IP address of the BGP neighbor, in dotted decimal format (a.b.c.d)
- *next-hop*—next-hop address of the edge router, in dotted decimal format (a.b.c.d)

Prerequisites

Must be in BGP configuration mode (*config-engine-bgp*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

Edge routers cannot be configured with next-hop-self

Description

Use the **`neighbor uses-next-hop-self-address`** command to identify the IP address to be used by CNA when asserting routes to edge routers configured with next-hop-self.

When *next-hop-self* is configured on a router, CNA will attempt to match the router's advertised next hop to a *next-hop* argument configured by a **`uses-next-hop-self-address`** command in one of the CNA configuration's BGP neighbor blocks.

When a match is found, CNA will set the next-hop address to the facing address of the link whose AS number (declared when the link was created with the **`link`** command) matches the first AS number in the route's AS path.

To be sure these facing addresses are included in the routing tables on each of your edge routers, you can redistribute them into IGP or you can create static routes on each edge router for all non-connected service providers.

Use the **no** form of the command to remove the *next-hop-self* mapping from the BGP neighbor configuration block.

---

netflow enable

Enable the netflow process on the EFC to receive netflow data.

Syntax

```
netflow enable
no netflow enable
```

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

Netflow is not enabled

Description

Use the **`netflow enable`** command to allow the optional EFC module to receive Netflow data.

Use the **no** form of the command to restore the default.

---

netflow filter interface

Enables port-based netflow packet filtering.

Syntax

```
netflow filter interface <name> <address> permit
        [ input-interface | output-interface ]
        [ <ifIndex> | <ifDescr> | <interface-address>

no netflow interface name> <address> permit
        [ input-interface | output-interface ]
        [ <ifIndex> | <ifDescr> | <interface-address>

no netflow interface
```

Arguments

- *name*— name of SNMP agent (see polling snmp-agent to see how you create this name on CNA)
- *address*—IP address (*a.b.c.d* format) of the device sending netflow to CNA
- *ifIndex*— interface index
- *ifDescr*— interface name
- *interface-address*— ip address (*a.b.c.d* format) of the interface on the sending device

Prerequisites

Must be in EFC configuration mode (*config-efc*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

none

Description

Use the **`netflow filter interface`** command to create netflow packet filters which will tell CNA to accept netflow packets from specific router/switch interfaces. If there are no filters specified for a particular device, then all packets will be accepted from that device. If there is at least one filter configured for a device, then netflow packets will be accepted from only the interfaces on that device which you have specified with this command.

If you specify *ifIndex* instead of *ifDescr* or *interface-address* then CNA will not need to query the sending device to determine the actual numbering of the interface. In this case, the SNMP agent name will not be used, and the output from show stat link-monitor will not include this filter since there is no need to SNMP poll to acquire the index.

Use the **no** form of the command to restore the default.

---

netflow interface

Specifies the interface on which Netflow data will be received.

Syntax

```
netflow interface [ ethernet0 | ethernet 1 ]
no netflow interface
```

Prerequisites

Must be in EFC configuration mode (*config-efc*), which requires privileged-level access.

Occurrence

Use once.

Default

ethernet0

Description

Use the **netflow interface** command to change the interface used by the optional EFC module to receive Netflow data.

Use the **no** form of the command to restore the default.

See the **interface** command for information about assigning an IP address to the interface.

See the **monitor interface** command for information about receiving SPAN data.

---

netflow port

Specifies the UDP port on which Netflow data will be received by the EFC module.

Syntax

```
netflow port <number>
no netflow port
```

Arguments

- *number*—UDP port number; range: 1025 through 65535

Prerequisites

Must be in EFC configuration mode (*config-efc*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

Ports 9995 and 9996 will be open by default

Description

Use the **netflow port** command to specify the UDP port number on which the optional EFC module will receive Netflow data. This port number must match the Netflow target port number configured on the router.

Use the **no** form of the command to remove the port designation (which will shutdown the receipt of Netflow data).

See the **interface** command for information about assigning an IP address to the interface.

See the **monitor interface** command for information about receiving SPAN data.

Prerequisites

Must be in EFC configuration mode (*config-efc*), which requires privileged-level access.

netflow source

Specifies the IP address of the router interface from which Netflow data originates.

Syntax

```
netflow source <address> [ version <string> ]
no netflow source [ <address> ]
```

Arguments

- *address*—IP address (*a.b.c.d* format)
- *string*—'5' and '7' are currently supported

Prerequisites

Must be in EFC configuration mode (*config-efc*), which requires privileged-level access.

Occurrence

Use once.

Description

Use the **`netflow source`** command to specify the IP address on the router from which Netflow data originates. You can specify multiple Netflow sources.

CAUTION:

If you do not use this command, then the EFC will accept netflow data from all sources.

Use the **version** keyword to specify which netflow version (either '5' or '7')will be accepted by CNA. If you omit this keyword, then both version 5 and version 7 will be accepted.

Use the **no** form of the command with an IP address to remove a specific source. Use the **no** form of the command with no address to remove all sources.

ntp broadcastclient

Configures the Network Time Protocol (NTP) process to listen for broadcast packets instead of polling for updates.

Syntax

```
ntp broadcastclient
no ntp broadcastclient
```

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

Listener mode is disabled; CNA must poll for updates.

ntp broadcastdelay

Specify the default propagation delay from the broadcast/multicast server and CNA.

Syntax

```
ntp broadcastdelay <microseconds>
no ntp broadcastdelay
```

Arguments

- *microseconds*—time, in microseconds, to be used as delay value; integer in the range 1 through 9,999,999

Prerequisites

Must be in configuration mode, which requires privileged level access.

Occurrence

Use once.

Default

microseconds—0

Description

The delay valued specified is used only if the delay cannot be computed automatically by the protocol.

Use the **no** form of the command to reset microseconds to its default.

---

ntp disable

Disables (or enables) the Network Time Protocol (NTP) on CNA.

Syntax

```
ntp disable
no ntp disable
```

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Default

NTP is enabled.

Description

To enable NTP, use the **no** form of the **`ntp disable`** command.

Use the following commands to configure CNA as an NTP client:

- **`ntp server`**
- **`ntp peer`**
- **`ntp broadcastclient`**
- **`ntp broadcastdelay`**

CNA cannot be configured as an NTP server.

ntp peer

Defines a Network Time Protocol (NTP) peer.

Syntax

```
ntp peer <address> [version <number>] [prefer]
no ntp peer [<address>]
```

Arguments

- *address*—hostname or IP address in dotted decimal form (a.b.c.d), of peer
- *number*—NTP version

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

number—4

Description

Version number can be in the range 1 through 4.

Use the **no** form of the command with *address* to disable the peer relationship with the specified host.

Use the **no** form of the command without an argument to disable all NTP peer relationships.

Use the keyword **prefer** to indicate a preference (and reduce switching among servers).

NTP peers synchronize with each other and exchange NTP packets. All peripheral CNA modules are automatically synchronized with the management module.

ntp server

Defines the Network Time Protocol (NTP) servers to be polled for updates.

Syntax

```
ntp server <address> [version <number>] [prefer]
no ntp server <address>
```

Arguments

- *address*—hostname or IP address of NTP server in dotted decimal format (a.b.c.d)
- *number*—NTP protocol version

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

number—4

Description

Version number can be in the range 1 through 4.

Use the **no** form of the command to remove an NTP server from the list of servers. If no host is specified, all servers will be removed.

Use the keyword **prefer** to indicate a preference (and reduce switching among servers).

All peripheral CNA modules are automatically synchronized with the management module.

---

number-of-packets

Defines the number of packets that are sent between endpoints for one instance of this active measurement test.

Syntax

number-of-packets <value>

no number-of-packets

Arguments

Argument	Description
<*value*>	Number of packets that are sent between endpoints for one instance of this test. Valid ranges are: • RTP test: 1–100 • Ping test: 1–100 • Traceroute test: 1–10 • TCP test: Not applicable.

Prerequisites

Must be in Chatter Active Measurement Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use once.

Default

Table 25: Default settings for number-of-packets command

Type of test	Default number of packets
RTP	10
Ping	10
Traceroute	1

Description

Defines the number of packets that are sent between endpoints for one instance of this active measurement test.

The no command restores the default value.

outage-detection active-probe

Defines the number of consecutive active-probe failures that should be interpreted by the CNA *decision-maker* process as an outage for the subnet.

Syntax

```
outage-detection active-probe { enable | disable |
       failure-threshold <fail-num> | measurement-threshold <meas-num> }
no outage-detection active-probe { enable | disable |
       failure-threshold | measurement-threshold }
```

Arguments

- *fail-num*—number of consecutive failures that constitute an outage; range: 1 through 20; default 5
- *meas-num*—number of successes that constitute the end of an outage; range: 1 through 255; default 5

Prerequisites

Must be in engine or decision policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once globally, and once per decision policy.

Description

Use the **`outage-detection active-probe`** command to define an outage as measured by active probing.

If an active probe fails *fail-num* consecutive times over the same link, an outage will be declared for the targeted subnet over that link. (Failure means no response from the target for a length of time specified as the timeout value in the active probe configuration.)

The current value of *num* is displayed as *probe-failure-threshold* in the output generated by the **`show statistics`** command with *dmaker* as the argument.

The outage will be considered to have ended after *meas-num* active probe successes for that subnet over the same link.

Use the **no** form of the command to revert to the defaults, either globally or within a decision policy.

---

outage-detection always-move-from-bgp (deprecated)

Enables *always-move-from-bgp* behavior when a service outage occurs on a BGP-selected route.

Syntax

```
outage-detection always-move-from-bgp
no outage-detection always-move-from-bgp
```

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once.

Default

always-move-from-bgp behavior is disabled (BGP's selection prevails, even though CNA has determined there is a service outage)

Description

By default, CNA is reluctant to move traffic away from a BGP-selected route for any reason other than a measured improvement in performance.

Use the **`outage-detection always-move-from-bgp`** command to remove this reluctance. When this behavior is enabled, CNA will remove the BGP-selected route from the winner set used in recomputing a performance route after a service outage has been detected on a BGP-selected route.

When configured from engine mode, the command is applied globally, to all subnets.

When configured from policy mode, the command is applied only to those subnets that are associated with the policy.

Use the **no** form of the command to disable the *always-move-from-bgp* behavior, thereby reinstating the reluctance to move away from a BGP route.

---

outage-detection brownout

Specifies the parameters used by CNA for detecting subnet brownout conditions.

Syntax

```
outage-detection brownout
        { enable | disable |
         absolute threshold <abs-rating> |
         relative { sample-period <seconds> | threshold <rel-rating> } |
         tolerance <tolerance> }
no outage-detection brownout
no outage-detection brownout absolute threshold
```

```
no outage-detection brownout tolerance
no outage-detection brownout relative threshold
```

Arguments

- *abs-rating*—decimal value in the range, 0.0 to 4.9
- *rel-rating*—decimal value in the range, 0.1 to 5.0
- *time*—sample period in seconds (60 through 86400)
- *tolerance*—decimal value in the range, 0.0 to 4.9

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once.

Default

- *abs-rating*—5.0
- *rel-rating*—0.5
- *time*—14440 (4 hours)
- *tolerance*—0.1

Description

Note:

A subnet brownout, sometimes called a soft outage, occurs when there is a significant reduction in application performance to end users as a result of subnet problems such as congestion, packet loss, link or equipment failures, or configuration errors. Brownouts differ from blackouts (also known as hard outages) which occur when a subnet or portion of a subnet fails completely. Brownouts occur when connectivity theoretically still exists but excessive latency and loss make the subnet effectively unusable. Standard routing protocols only detect hard outages, so while blackouts can eventually be overcome, normal routing will not detect or correct a brownout, allowing the condition to persist.

CNA is able to detect and take corrective action for subnet brownouts. Use the **`outage-detection brownout`** command to change CNA's sensitivity to subnet brownout conditions to suit the specific needs of your subnet. In order detect a brownout, CNA will establish the baseline performance of each subnet, ranging from 0.0 (worst) to 5.0 (best), over time. The **relative threshold** is subtracted from this value and compared to the **absolute threshold**. The lesser of these values yields the *effective threshold*, which is used to determine if there is a brownout condition. If the subnet is currently performing equal to or worse than this

computed value minus the **tolerance**, then there is a brownout. The brownout condition ends when the current performance is better than the *effective threshold* **plus** the **tolerance**. Note that the **tolerance** parameter provides hysteresis for brownout transitions.

Use the **relative sample-period** keyword to change the time interval between each measurement used for calculating the baseline performance of a subnet.

Use the **no** form of the **`outage-detection brownout`** command to revert to the default settings.

outage-detection on (deprecated)

Enables outage-detection, the ability to detect when a previously asserted link is no longer able to provide coverage for a subnet.

Syntax

```
outage-detection on
no outage-detection on
```

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once.

Default

Outage detection is enabled.

Description

CNA keeps track of the time intervals between HRTT measurements for each subnet, through each service provider link. Whenever the device detects an unacceptable increase in this interval when measured against a rolling-average of intervals, the link will be considered no longer able to serve that subnet.

See the **`outage-detection silence`** command for information about unacceptable intervals.

When configured from engine mode, the command is applied globally, to all subnets.

When configured from policy mode, the command is applied only to those subnets which are associated with the policy.

If the link on which an outage is detected is the current winner for the subnet, CNA will recompute the performance route.

If the link is the route that has been selected by natural BGP, the link will remain in the winner set from which the new performance route is selected, unless the **`outage-detection always-move-from-bgp`** command has been enabled (it is disabled by default). Otherwise, the invalidated link will be removed from the winner set.

Use the **no** form of the **`outage detection on`** command to disable outage-detection.

outage-detection silence

Specifies the maximum amount of time allowed between HRTT measurements before a subnet is determined to be unreachable through a service provider link.

Syntax

```
outage-detection silence { enable | disable |
        tslm_threshold <time> <multiplier> }
no outage-detection silence { enable | disable | tslm_threshold }
```

Arguments

- *time*—interval in milliseconds (100 through 10,000,000)
- *multiplier*—integer (1 through 1,000,000)

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once.

Default

- *time*—10000 (10 seconds)
- *multiplier*—10

Description

CNA will consider a link no longer able to reach a subnet if both of the following conditions are met:

- time since the last recorded measurement through the link exceeds the value of *time*

- time since the last recorded measurement through the link exceeds the value of *multiplier* times the maximum rolling-average interval between measurements for the subnet across all of your service provider links

An excessive number of subnets categorized as I (insufficient measurements) in output from the **show prefix (deprecated)** command can be a symptom of premature declaration of outages for low-traffic subnets and at least one link categorized as o (outage detected), especially in systems with few service provider links. To avoid this condition, increase the value of *time* to a value larger than the typical inter-measurement time for the subnet.

When configured from engine mode, the command is applied globally, to all subnets.

When configured from policy mode, the command is applied only to those subnets which are associated with the policy.

For active probes, you should create a decision policy and set the multiplier value to 3. Setting the value any lower can cause premature declarations of outages and is not recommended.

Note:

The per-link outage-detection algorithm compares measurement times across each of your service provider links to determine if there is an outage on any one of them. But if the currently asserted route experiences an outage, the single-pixel *.gif* that is used in the passive user traffic test will never be served to users in the subnet, and therefore no measurements will be collected. CNA must use a different mechanism to detect that condition. See the **outage-detection total-silence** command.

Use the **no** form of the command to revert to default values for *time* and *multiplier*.

outage-detection silence always-move-from-bgp

Enables *always-move-from-bgp* behavior when a service outage occurs on a BGP-selected route.

Syntax

```
outage-detection always-move-from-bgp { enable | disable }
no outage-detection always-move-from-bgp { enable | disable }
```

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once.

Default

always-move-from-bgp behavior is disabled (BGP's selection prevails, even though CNA has determined there is a service outage)

Description

By default, CNA is reluctant to move traffic away from a BGP-selected route for any reason other than a measured improvement in performance.

Use the **`outage-detection always-move-from-bgp`** command to remove this reluctance. When this behavior is enabled, CNA will remove the BGP-selected route from the winner set used in recomputing a performance route after a service outage has been detected on a BGP-selected route.

When configured from engine mode, the command is applied globally, to all subnets.

When configured from policy mode, the command is applied only to those subnets which are associated with the policy.

Use the **no** form of the command to disable the *always-move-from-bgp* behavior, thereby reinstating the reluctance to move away from a BGP route.

---

outage-detection total-silence

Controls the way CNA detects an outage (when all HRTT measurements for a subnet on all links cease).

Syntax

```
outage-detection total-silence { enable | disable |
        tslm_threshold <min> <max> <multiplier> }
no outage-detection total-silence { enable | disable | tslm_threshold }
```

Arguments

- *min*—minimum time interval between measurements before an outage can be declared; in milliseconds (100 through 2,147,483,647)
- *max*—maximum time between measurements before an outage will be declared; in milliseconds (100 through 2,147,483,647)
- *multiplier*—integer (1 through 10,000,000)

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once.

Default

Global level:

- *min*—60000 (60 seconds)
- *max*—86,400,000 (24 hours)
- *multiplier*—100

Values within a decision policy default to the settings currently in effect at the global level.

Description

The **`outage-detection total-silence`** command sets the parameters used by CNA to determine that a subnet has been totally silent—no UTT measurements recorded for any of your links—for an unacceptable period of time.

Each time a measurement is received for a sufficiently measured subnet, CNA will compute a deadline for the next measurement. If no measurements are received for the subnet across any of your links before the deadline arrives, CNA will declare an outage for that subnet across all links.

When an outage is declared across all links, CNA will withdraw its route assertion for the subnet, unless the subnet is user-defined (not learned from BGP). For non-BGP subnets, CNA will take no action as a result of the outage.

The deadline is calculated as follows:

- *interval* = *multiplier* times *maxIMT*
- if *interval* < *min*, then *interval* = *min*
- if *interval* > than *max*, then interval = *max*
- *deadline* = *interval*+ time last measurement received

Note:

The total-silence detection process runs once a minute in the background, as a backstop for subnet outages that might otherwise go undetected. There may be as much as a two-minute delay between the actual and reported times of an outage.

See the **`outage-detection silence`** command.

For more information on the outage-detection mechanisms, see Subnet Outage Control in Chapter 3: Adaptive Path Control of the *Adaptive Path Controller Guide*.

Use the **no** form of the command to revert to default values for *min, max* and *multiplier*.

outage

Include additional events in the outage reports.

Syntax

```
outage { perf-violation | probe-failure | silence } enable
no outage perf-violation enable
no outage probe-failure enable
no outage silence enable
```

Note:

Starting with CNA version 5, the keywords **perf-violation** and **probe-failure** have been deprecated and will no longer have any effect.

Prerequisites

Must be in report configuration mode (*config-reporting*), which requires privileged-level access.

Occurrence

Use once.

Default

Disabled

Description

Use the **perf-violation** keyword to include performance violation events in the outage reports.

Use the **probe-failure** keyword to include probe failure events in the outage reports.

Use the **silence** keyword to include performance violation events in the outage reports.

Use the **no** form of the command to prevent the additional events from appearing in the outage reports.

outbound-bgp-community (discontinued)

Defines lists of BGP community values that will be attached to CNA route assertions.

Syntax

```
outbound-bgp-community <value>
no outbound-bgp-community
```

Arguments

- *value*—string of community numbers, each in *a:b* format where *a* and *b* are integers in the range 0 through 65535 separated by a colon

Prerequisites

Must be in link configuration mode (*config-engine-link* or *config-engine-policy-link*), which requires privileged-level access.

Occurrence

Use once per link configuration block.

Description

Use the **bgp-community** command. The **outbound-bgp-community** command will be discontinued in a future CNA release.

packet-length

Defines the packet length in bytes for packets that are sent between endpoints of the particular active measurement test.

Syntax

packet-length <value>

no packets-length

Arguments

Argument	Description
<*value*>	Packet length in bytes for this active measurement test Valid ranges are: • RTP test: 88–1000 • Ping test: 10–1000 This command is applicable only for Ping and RTP tests.

Prerequisites

Must be in Chatter Active Measurement Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use once.

Default

Table 26: Default settings for packet-length command

Type of test	Default packet length (in bytes)
RTP	256
Ping	64

Description

Defines the packet length in bytes for packets that are sent between endpoints of the particular active measurement test.

This command is applicable only for Ping and RTP tests.

The no command restores the default value.

---

penalty

Assigns a numerical penalty to a service provider access link.

Syntax

```
penalty <value>
no penalty
```

Arguments

- *value*—integer in the range 0 through 1,000

Prerequisites

Must be in link configuration mode (*config-engine-link* or *config-engine-policy-link*), which requires privileged-level access.

Occurrence

Use once per link configuration block.

Default

Penalty *value* is 0.

Description

The **no** form of the command resets *value* to the default.

The penalty value will be used when calculating the performance metric which CNA uses to evaluate a service provider link against other links for a given subnet.

Increasing a penalty decreases the likelihood that a service provider link will be asserted as best path.

If mode is *config-engine-link*, the **`penalty`** command applies to the link for all subnets.

If mode is *config-engine-link-policy*, the **`penalty`** command applies to the link for those subnets associated with the policy.

Values set globally, in *config-engine-link* mode, will apply to all links not explicitly set in *config-engine-policy-link* mode.

For more information about policies, see the **`decision-policy`** command and Chapter 6: Decision Policies and Application Models in the *Monitoring Guide*.

performance-group

This command has been deprecated and will be translated to the site command.

persistent-hosts

Description

Controls whether an unreachable subnet affects or does not affect the availability score of service-provider links. A subnet is unreachable when probes on all service-provider links fail.

The no form of this command restores the default value:

- When entered in *Adaptive Routing Engine Configuration mode*, the command restores the default value for the *global setting*.
- When entered in *Policy Configuration mode*, the command restores the default value for the *policy decision*.

For more information on the command modes for the **`persistent-hosts`** command, see Command modes.

Syntax

persistent-hosts {enable | disable}

no persistent-hosts

Arguments

Table 27: persistent-hosts arguments

Argument	Description
enable	Marks a subnet *not-rescued* when probes to it fail. A subnet is not-rescued when probes to targets on the subnet fail, and CNA cannot find an alternate path to the targets (all service provider links have outages for the subnet). Not-rescued subnets affect the availability score of service-provider links. Use this option for subnets that contain persistent hosts such as routers, gateways, and so forth.
disable	Marks a subnet *insufficient data* when probes on all service-provider links fail. Subnets that are marked insufficient data do not affect the availability score of service-provider links. If CNA is operating in route assert mode, asserted routes for the subnet are withdrawn. Use this option for subnets that contain non-persistent hosts such as PCs, laptops, and so forth.

Command modes

Table 28: persistent-hosts command modes

When entered in this command mode...	The command applies to...
Adaptive Routing Engine Configuration mode	the global setting. The global setting applies to all subnets that: • are *not* assigned a policy decision • are assigned a decision policy that does not explicitly set this parameter.
Policy Configuration mode	the policy decision. The policy decision applies to all subnets to which it is assigned. This setting overrides the global setting. **Note:** Any parameter not explicitly set in a decision policy assumes the globally set value for that parameter.

For information on how to access these modes, see Accessing and exiting command modes.

Default

Disabled.

persistent-hosts timeout <hours>

Description

Sets the time after which an unreachable subnet no longer affects availability scores of service-provider links. A subnet is unreachable when probes on all service-provider links fail.

Use this command for subnets that have persistent-hosts enabled.

The no form of this command restores the default value:

- When entered in *Adaptive Routing Engine Configuration mode*, the command restores the default value for the *global setting*.
- When entered in *Policy Configuration mode*, the command restores the default value for the *policy decision*.

For more information on the command modes for the **`persistent-hosts timeout`** command, see Command modes.

Syntax

persistent-hosts timeout <hours>

no persistent-hosts timeout

Arguments

Table 29: persistent-hosts timeout arguments

Argument	Description
<hours>	Number of hours after which an unreachable subnet no longer affects availability scores of service-provider links.

Command modes

Table 30: persistent-hosts command modes

When entered in this command mode...	The command applies to...
Adaptive Routing Engine Configuration mode	the global setting. The global setting applies to all subnets that: • are *not* assigned a policy decision • are assigned a decision policy that does not explicitly set this parameter.
Policy Configuration mode	the policy decision. The policy decision applies to all subnets to which it is assigned. This setting overrides the global setting. **Note:** Any parameter not explicitly set in a decision policy assumes the globally set value for that parameter.

For information on how to access these modes, see Accessing and exiting command modes.

Default

72 hours (3 days)

neighbor uses-next-hop-self-address

Provides CNA with a next-hop address when routers are configured with *next-hop-self*.

Syntax

```
neighbor <address> uses-next-hop-self-address <next-hop>
no neighbor <address> uses-next-hop-self-address
```

Arguments

- *address*—IP address of the BGP neighbor, in dotted decimal format (a.b.c.d)
- *next-hop*—next-hop address of the edge router, in dotted decimal format (a.b.c.d)

Prerequisites

Must be in BGP configuration mode (*config-engine-bgp*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

Edge routers cannot be configured with next-hop-self

Description

Use the **`neighbor uses-next-hop-self-address`** command to identify the IP address to be used by CNA when asserting routes to edge routers configured with next-hop-self.

When *next-hop-self* is configured on a router, CNA will attempt to match the router's advertised next hop to a *next-hop* argument configured by a **`uses-next-hop-self-address`** command in one of the CNA configuration's BGP neighbor blocks.

When a match is found, CNA will set the next-hop address to the facing address of the link whose AS number (declared when the link was created with the **`link`** command) matches the first AS number in the route's AS path.

To be sure these facing addresses are included in the routing tables on each of your edge routers, you can redistribute them into IGP or you can create static routes on each edge router for all non-connected service providers.

Use the **no** form of the command to remove the *next-hop-self* mapping from the BGP neighbor configuration block.

performance-violation-detection metric (deprecated)

Note:

Starting with CNA, this command has been deprecated and will be translated to the **`outage-detection brownout`** command.

Sets metrics used by the decision-maker process to determine when a performance violation occurs (or ceases).

Syntax

```
performance-violation-detection metric { enable | disable |
               absolute { threshold | tolerance } < num > }
no performance-violation-detection metric enable | disable
no performance-violation-detection metric threshold
no performance-violation-detection metric tolerance
```

Arguments

- *num*—for threshold: (range is 100 through 1000); for tolerance: (range is 0 through 50)

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **`performance-violation-detection`** command to set the **threshold** and **tolerance**, which affect the way the decision-maker process will decide if there is a performance violation. With the respective default values of 300 and 10, the decision-maker process will consider a violation to have occurred when a link's quality-metric exceeds 310. It will consider the violation to have ceased when that link's quality-metric falls below 290. For more information on quality-metric, see Details on Application Models in Chapter 6: Decision Policies and Application Models of the *Monitoring Guide*.

ping

Tests connectivity between devices.

Syntax

```
ping [ count <packets> ]
     [ wait <seconds> ]
     [ record-route ]
     [ bypass-routing-table ]
     [ numeric ]
     [ source <source-ip-address> ]
     [ packet-size <bytes> ]
     [ module <slot> ]
     [ link <name> ]
     <target>
```

Arguments

- *packets*—number of packets to be counted
- *seconds*—number of seconds to wait between tests
- *source-ip-address*—IP address in dotted decimal format (a.b.c.d) to be used as the source address of the ping
- *bytes*—number of data bytes to be sent
- *slot*—CNA slot or subsystem number; origin of test
- *name*—perform the ping from the named link
- *target*—hostname or IP address in dotted decimal format (a.b.c.d)

Prerequisites

System operations mode, either login-level or privileged-level access.

Default

slot—0

packets—continuous stream of packets

bytes—56 (total of 84 bytes sent to target, including 8-byte ICMP header and 20-byte IP header)

seconds—1

Description

You can specify any CNA module as *target* by using its internally defined host name *moduleN*, where *N* is the slot or subsystem number, as follows:

```
ping module 1 module3
```

In the example, the keyword **module** and its argument *1* follows the **`ping`** command (note that the keyword **module** and its argument are separated by a space).

There is no space in *module3*, however. This entire string is the value of the *target* argument.

To ping subsystem 1 from subsystem 0 on a 3300 base unit, the command would be:

```
ping module1
```

Use the **count** keyword to restrict the test to a specified number of packets.

Use the **wait** keyword to specify the amount of time to wait between sending packets.

Use the **record-route** keyword to display the routes taken by the packet (maximum of nine)

Use the **numeric** keyword to suppress symbolic name lookup for host addresses.

Use the **bypass-routing-table** keyword to communicate directly with a host on an attached subnet.

Use the **packet-size** keyword to define packet size.

Use the **module** keyword to specify the CNA module to be source of the ping.

polling bgp-status interface-ip-address

Identifies the IP address of the interface to be polled for BGP status.

Syntax

```
polling bgp-status interface-ip-address <address>
no polling bgp-status interface-ip-address
```

Arguments

- *address*—IP address (in *a.b.c.d* format) of the interface to be used when polling an edge router's SNMP agent for BGP status

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Occurrence

Use once per link.

Description

When polling edge routers, you can specify one SNMP agent for BGP status and a different one for link usage information and interface status.

Use the **`polling bgp-status interface-ip-address`** command to identify the interface to be used when polling for BGP status.

See the **`polling bgp-status snmp-agent`** command.

polling bgp-status interval

Determines how often CNA will poll edge routers for status of the BGP peering relationship.

Syntax

```
polling bgp-status interval <seconds>
no polling bgp-status interval
```

Arguments

seconds—polling frequency; range: 5 through 3600

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Default

30

Occurrence

Use once per link.

Description

When BGP status polling behavior is enabled, CNA will contact your edge router to evaluate the status of the BGP peering relationship at the frequency set by the *seconds* argument.

This behavior is enabled by the **`polling bgp-status on`** command.

Use the **no** form of the **`polling bgp-status interval`** command to reset the frequency to the default.

polling bgp-status on

Enables polling of edge routers to determine status of BGP peering relationship.

Syntax

```
polling bgp-status on
no polling bgp-status on
```

Prerequisites

You must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Default

Polling for BGP status is disabled.

Description

When polling for BGP status is enabled, CNA will contact the edge router at regular intervals to evaluate the status of the BGP peering relationship. CNA's decision-maker process will use this information to help determine whether a subnet is experiencing a subnet outage.

The frequency at which the edge router will be polled for BGP status is determined by the **`polling bgp-status interval`** command.

Use the no form of the command to disable polling for BGP status.

---

polling bgp-status snmp-agent

Identifies an SNMP agent on an edge router to be used for BGP status polling.

Syntax

```
polling bgp-status snmp-agent <name>
no polling bgp-status snmp-agent
```

Arguments

- *name*—name of the SNMP agent

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Occurrence

Use once per link.

Description

Use the **`polling bgp-status snmp-agent`** command to associate a block of SNMP configuration commands with a link.

These SNMP commands are created in *config* mode with the **polling snmp-agent** command, to provide CNA with the information it needs to communicate with an SNMP agent on an edge router.

If you do not use the **polling bgp-status snmp-agent** command, CNA will poll the same SNMP agent for BGP status that it polls for link usage.

You will also need to add the IP address of the SNMP agent to the link configuration.

See the **polling bgp-status interface-ip-address** and **polling snmp-agent** commands.

polling bgp-status unknown-threshold

Sets a threshold for number of times a status check of a BGP peering can fail before the service provider access link is declared down.

Syntax

```
polling bgp-status unknown-threshold <num>
no polling bgp-status unknown-threshold
```

Arguments

- *num*—number of times status check can fail before status is declared unknown (range: 1 through 100)

Prerequisites

You must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Default

3

Occurrence

Once for each link.

Description

If the **polling bgp-status on** and **require-bgp-snmp-poll** commands are both enabled, the value set by the **polling bgp-status unknown-threshold** command determines how many times the CNA link monitor process must report the peering as unknown before the link is considered to be unreliable.

This provides a buffer against prematurely declaring the link to be down as a result of transient, non-fatal SNMP errors.

See Link Outage Control in Chapter 3: Adaptive Path Control of the *Adaptive Path Controller Guide* for more information.

Use the **no** form of the command to reset the *number* argument to the default.

polling interface

Tells CNA which interface to poll, using the *ifDescr* MIB variable as the key.

Syntax

```
polling interface-description { <address> | <text> | <index> }
no polling interface
```

Arguments

address—IP address of the interface, in dotted decimal format (a.b.c.d)

text—value of the interface's *ifDescr* MIB variable on the router (maximum length: 255 characters)

index—interface's *ifIndex* MIB variable on router (integer in the range 1 through 2,147,483,647)

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Occurrence

Use once per link.

Description

When polling for usage data, CNA needs to know where to look. The device will use whichever of the above methods you have set (use only one in the CNA configuration):

Use the **no** form of the command to remove the interface identifier.

polling interface-description (deprecated)

Note:

See the new command: **polling interface**.Tells CNA which interface to poll, using the *ifDescr* MIB variable as the key.

Syntax

```
polling interface-description <text>
no polling interface-description
```

Arguments

text—value of the interface's *ifDescr* MIB variable on the router (maximum length: 255 characters).

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Occurrence

Use once per link.

Description

When polling for usage data, CNA needs to know where to look. The device will use whichever of the following you have set (use only one in the CNA configuration):

- the interface's *ifDescr* MIB variable (use the CNA command **polling interface-description**)
- the interface's *ifIndex* MIB variable (use the CNA command **polling interface-index (deprecated)**)
- the edge router's ISP-facing IP address (use the CNA command **polling interface-ip-address (deprecated)**)

Use the **no** form of the command to clear the association between the link and an SNMP configuration block.

polling interface-index (deprecated)

See the new command: **polling interface**. Tells CNA which interface to poll, using the *ifIndex* MIB variable as the key.

Syntax

```
polling interface-index <index>
no polling interface-index
```

Arguments

index—interface's *ifIndex* MIB variable on router (integer in the range 1 through 2,147,483,647)

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Occurrence

Use once per link.

Description

When polling for usage data, CNA needs to know where to look. The device will use whichever of the following you have set (use only one in the CNA configuration):

- the interface's *ifDescr* MIB variable (use the CNA command **`polling interface-description`**)
- the interface's *ifIndex* MIB variable (use the CNA command **`polling interface-index`**)
- the edge router's ISP-facing IP address (use the CNA command **`polling interface-ip-address`**)

Use the **no** form of the command to clear the association between the link and an SNMP configuration block.

---

polling interface-ip-address (deprecated)

See the new command: **`polling interface`**. Tells CNA which interface to poll, using the IP address of the edge router's ISP-facing interface.

Syntax

```
polling interface-ip-address <address>
no polling interface-ip-address
```

Arguments

address—edge router's ISP-facing IP address

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Occurrence

Use once per link.

Description

When SNMP agents on edge routers, CNA needs to know which interface to contact. The device will use whichever of the following you have set (use only one in the CNA configuration):

- the interface's *ifDescr* MIB variable (use the CNA command **`polling interface-description`**)
- the interface's *ifIndex* MIB variable (use the CNA command **`polling interface-index`**)
- the edge router's ISP-facing IP address (use the CNA command **`polling interface-ip-address`**)

Use the **no** form of the command to clear the association between the link and an SNMP configuration block.

You can also configure a separate interface for polling a different SNMP agent for BGP status. See the **`polling bgp-status snmp-agent`** command.

---

polling interface-status down-threshold

Sets the number of consecutive times an interface can be reported down by the CNA link monitor process before the link is considered down by the decision maker process.

Syntax

```
polling interface-status down-threshold <number>
no polling interface-status down-threshold
```

Arguments

- *number*—threshold value (range: 1 through 100)

Prerequisites

You must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Default

3

Description

The CNA software will evaluate the status of the ISP-facing interface on the edge router to determine if the link is up or down if either of the following are true:

- **`polling bgp-status on`** command is enabled and the CNA link monitor has reported the BGP peering to be down
- BGP status is reported as unknown and the **`require-bgp-snmp-poll`** command is not enabled

The number set by the **`polling interface-status down-threshold`** command determines the number of consecutive times the interface must be reported down before the link is considered to be down.

This provides a buffer against prematurely declaring the link to be down as a result of transient, non-fatal SNMP errors.

See Link Outage Control in Chapter 3: Adaptive Path Control of the *Adaptive Path Controller Guide* for more information.

Use the **no** form of the command to restore *number* to its default.

---

polling interface-status interval

Interval for checking the status of an edge-router interface.

Syntax

```
polling interface-status interval <seconds>
no polling interface-status interval <seconds>
```

Arguments

- *seconds*—interval (range: 5 through 3600)

Prerequisites

You must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Default

30 second

Description

CNA will check the interface configured for the link at the specified interval. The ip address that will be checked is the address configured for the link by the **`polling interface-ip-address`** command.

CNA will not consider a link to be down unless it fails this check a user-specified number of times. See the **`polling interface-status down-threshold`** command for information about setting the consecutive-times number.

Use the **no** form of the **`polling interface-status interval`** command to reset *seconds* to the default.

---

polling interface-status on

Enables polling of interfaces to determine link status.

Syntax

```
polling interface-status on
no polling interface-status on
```

Prerequisites

You must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Default

Interface-status polling is disabled.

Occurrence

Use once per link.

Description

If interface polling is enabled, CNA will contact the edge router at regular intervals to determine the status of the link.

The frequency of contact is determined by the **`polling interface-status interval`** command.

Use the **no** form of the command to disable interface status polling.

polling physical-link

Used for combining multiple physical links into a single engine link.

Syntax

```
polling physical-link <name> snmp-agent <agent-name>
         interface { <address> | <text> | <index> }
         [ bandwidth <mbps> ]
         [ passive ]
         [ use-64-bit-counter interval-greater-than <seconds> ]
no physical-link <name>
```

Arguments

name—name of physical link

agent-name—name of SNMP agent to be used for obtaining link usage information

address—IP address of the interface, in dotted decimal format (a.b.c.d)

text—value of the interface's *ifDescr* MIB variable on the router (maximum length: 255 characters)

index—interface's *ifIndex* MIB variable on router (integer in the range 1 through 2,147,483,647)

mbps—maximum capacity (mbps) of the physical link (any non-negative number)

seconds—when present, indicates that a 64-bit counter should only be used for those objects that have a polling interval greater than this (integer in the range 5 through 3600)

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Default

none

Occurrence

Once for each physical link.

Description

The **`polling physical-link`** command enables you to group physical links into a single engine link. This means that multiple physical links will be polled to determine the link loading for the engine link.

Use the **passive** keyword for links that the decision maker process cannot control (its link usage will still be used for load optimization. For more information about the load threshold table and the load optimization process, see Chapter 2: WAN Cost and Load Optimization in the *Adaptive Path Controller Guide*.

Use the **bandwidth** keyword to specify the maximum physical capacity of a link.

Note:

Such a link must be declared as a physical link; instead of **`polling interface`**, you need to use the **`polling physical-link`** command.

When using the **use-64-bit-counter** keyword, you may also constrain the use of 64-bit counters to only those objects which are polled at a slower interval (specified with the **interval-greater-than** keyword). This is to work around an update problem exhibited by some router models.

Use the **no** form of the command to remove the polling configuration from a particular physical link.

polling snmp-agent

Provides the CNA link monitor with contact information needed to poll an SNMP agent on a remote device for link usage and BGP and interface status.

Syntax

```
polling snmp-agent <name>
no polling snmp-agent <name>
```

Arguments

- *name*—name of the SNMP agent

Prerequisites

To create an SNMP configuration block, you must be in configuration mode (*config*), which requires privileged-level access.

To associate the configuration block with a service provider link, you must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Description

You will need to add a block of SNMP configuration commands for each device that is to be polled, then you will need to associate the block with the edge router on which the link is configured. You can also have a separate block of configuration commands for an SNMP agent dedicated to BGP status polling.

Enter configuration mode (*config*) to create the configuration block using the **`polling snmp-agent`** command. The *name* argument of the **`polling snmp-agent`** command should be a unique string that meets the criteria described in the section Legal Characters in Chapter 1: Introduction to the Command Line Interface.

This will place CNA into SNMP-polling configuration mode (*config-polling-snmp*), which will allow you to use the following commands to set the following parameters:

- **`community`**
- **`port`**
- **`retries`**
- **`snmp-agent-ip-address`**
- **`timeout`**

When you have configured the parameters, exit out of *config-polling-snmp* mode and enter link configuration mode (*config-engine-link*) for the link that represents the same service provider.

Enter the **`polling snmp-agent`** command again, with the same *name* argument.

Then continue configuring the link with the following commands:

- **`polling bgp-status interface-ip-address`**
- **`polling usage interval`**
- **`polling usage on`**

See the **`polling bgp-status snmp-agent`** command.

polling usage interval

Sets the sampling interval to be used when storing polled data in the trending database.

Syntax

```
polling usage interval <collection-rate> <billing-rate>
no polling usage interval
```

Arguments

- *collection-rate*—polling rate; frequency of data collection from routers in seconds (range: 5 through 300)
- *billing-rate*—ISP's sampling rate, for billing, in seconds (range: 5 through 3600)

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Default

collection-rate—30

billing-rate—300

Occurrence

Once for each service provider link

Description

The **`polling usage interval`** command sets the frequency at which the CNA link monitor process will gather link-usage data from the edge router and tells CNA the usage-sampling interval used by your ISP to calculate your bill.

Changing the interval will delete existing usage data from the trending database.

Set the value of the *billing-rate* to the interval that your ISP uses for its billing calculations. This value must be a multiple of the interval specified by the *collection-rate* argument. If it is not, the command will not be accepted and an error message will be displayed.

Collected data will be stored in the CNA's internal trending database so that you can generate trend report charts showing actual usage as well as graphs which approximate your ISP's billing samples.

See the **`polling usage on`** command for additional polling details.

Use the **no** form of the **`polling usage interval`** command to revert to default *collection-rate* and *billing-rate* values.

---

polling usage on

Enables polling of the edge router for link-usage data.

Syntax

```
polling usage on
no polling usage on
```

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Default

Polling is off

Occurrence

Once for each service provider link.

Description

The **polling usage on** command enables the polling of an edge router for ISP link-usage data. Polling is enabled on a per-link basis.

You must also provide the following information to CNA:

- the link's ISP-facing address on your edge router, using the **polling interface-ip-address** command
- the address to use for obtaining SNMP information from your edge router, using the **snmp-agent-ip-address** command
- the read-only community string for obtaining SNMP information from your edge router, using the **community** command
- how frequently you want CNA to poll the edge router, using the **polling usage interval** command
- the sampling interval you want CNA to use in calculating a projected bill for that link, and the ISP's billing period, using the **polling usage interval**

If a load threshold table has been defined, only those links that are contained in the table should be polled for usage. See the **load threshold-table** command.

Use the **no** form of the polling usage on command to disable the polling feature.

---

polling usage unknown-threshold

Short description

Syntax

```
polling usage unknown-threshold <value>
no polling usage unknown-threshold
```

Arguments

- *value*—threshold value; range: 1 through 100

Prerequisites

Must be in link configuration mode (*config-engine-link* or *config-engine-policy-link*), which requires privileged-level access.

Occurrence

Use once per link configuration block.

Default

3

Description

When SNMP polling is enabled, the value set by the **`polling usage unknown-threshold`** command establishes a threshold for the number of consecutive times the CNA link monitor process can fail to obtain load data from the remote SNMP agent.

When this threshold is exceeded, the CNA decision maker process will conclude that there is no load information available on the link and will automatically drop out of route-assert mode if that mode is enabled.

---

polling usage use-64-bit-counter

Configures the CNA *link-monitor* process to use a 64-bit counter for all polled values.

Syntax

```
polling usage use-64-bit-counter [ interval-greater-than <seconds> ]
no polling usage use-64-bit-counter
```

Arguments

seconds—minute of the hour; range: 5 through 3600

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Occurrence

Use once per link.

Default

Disabled (uses 32 bit counters)

Description

The **`polling usage use-64-bit-counters`** command configures the CNA *link-monitor* process to use 64-bit counters when polling edge routers for link usage data.

By default, the *link-monitor* process obtains usage data from the router's smaller, but more reliable, 32-bit counters. However, with a Gigabit Ethernet interface, the 32-bit counter will fill up and roll over to 0 in about 32 seconds.

Your edge router may not update the 64-bit counter as frequently as the 32-bit counter. You should configure a polling interval of 60 seconds if you use the 64-bit counter.

Use the **interval-greater-than** keyword to prevent the use of 64-bit counters when the collection-rate (see **polling usage interval**) has been set to a relatively low number which can cause some routers to give incorrect results. For example, if you have set the collection-rate (see **polling usage interval**) to 60 seconds, and configured **interval-greater-than** to a value less than 60, then 64-bit counter will NOT be used.

port

Tells CNA what port to use on an external device that is being polled or probed.

Syntax

port <number>

no port

Arguments

- *number*—a valid port number on the device to be polled; range: 1 through 65535

Prerequisites

Must be in polling-agent (*config-polling-snmp*) or active-measurement group (*config-engine-active*) configuration mode, both of which require privileged-level access.

Default

SNMP polling—161

Active-measurement group—no default

Description

The **`port`** command tells CNA what port to contact on a remote device.

You can use the **`port`** command in either of the following configuration modes:

- polling-agent (*config-polling-snmp*) mode—when polling for link usage is enabled, the CNA SNMP link-monitor agent will contact an external device (usually a router) using the port number specified by *number*

- active-measurement group (*config-engine-active*) mode—if specified for an active-measurement group, CNA will use this port to contact a target of an active-measurement probe; the group-wide value is superseded by a port specified for an individual target

  Note:

  If you are configuring a TCP-type active probe, a *port* value must be specified either for the individual probe or the group, or the probe will fail. The output of the **show statistics** command for the *am-scheduler* process includes a configuration errors table which will note the lack of a port assignment.

See the **polling usage on** command for more information about polling for link usage data. See the **polling snmp-agent** command for more information about polling parameters.

See the **active-measurement group** command for general information about active measurement groups. See the **target** command information about group-wide versus specific port assignments.

Use the **no** form of the **port** command to reset the port number to the default. (This will remove the **port** command from an active measurement group.)

port (Chatter module only)

Defines the TCP port to which the TCP active measurement test is performed.

Syntax

port <value>

no port

Arguments

Table 31: port arguments

Argument	Description
<value>	Destination TCP port number for the active measurement test. The valid range is 1–65535.

Prerequisites

Must be in Chatter Active Measurement Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use once.

Default

TCP port 80.

Description

Defines the TCP port to which the TCP active measurement test is performed.

This command is applicable only to TCP tests.

The no command restores the default value of TCP port 80.

prefer-in-winner-set bgp

Tells the CNA decision-maker process that it should consider BGP status when selecting a performance route from the winner set.

Syntax

```
prefer-in-winner-set bgp { enable | disable }
no prefer-in-winner-set bgp
```

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once.

Default

Not enabled (CNA decision maker will follow default selection rules)

Description

Use the **`prefer-in-winner-set bgp enable`** command as one of the tie-breaker criteria to require that the winning link be at least as good as the BGP selected route. Use the **disable** keyword to suspend this behavior without removing the command from the configuration.

Enabling BGP-status evaluation ensures that the link chosen from the winner set will perform at least as well as a BGP-selected link that is also in the same winner set. Or, put another way, it

ensures that the winning link will not perform any worse than a BGP route in the winner set. It does not, however, guarantee that a BGP selection in the winner set is an automatic winner.

For information about performance route selection criteria, see Winner Set in Chapter 3: Adaptive Path Control of the *Adaptive Path Controller Guide*.

Use the **no** form of the command to disable BGP-status evaluation and remove the command from the configuration.

prefer-in-winner-set most-reliable

Tells the CNA decision-maker process that it should consider a link's reliability when selecting a performance route from the winner set.

Syntax

```
prefer-in-winner-set most-reliable { enable | disable }
no prefer-in-winner-set most-reliable
```

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once.

Default

Not enabled (CNA decision maker will follow default selection rules)

Description

Use the **`prefer-in-winner-set most-reliable`** command to allow the CNA decision-maker process to consider link reliability when determining a winner in a tie-breaker situation. Use the **disable** keyword to suspend this behavior without removing the command from the configuration.

Enabling reliability evaluation ensures that the link chosen from the winner set will have the best reliability score, which is determined by the frequency that a link moves out of the winner set (for example, a link which rarely moves out of the winner set will tend to have better reliability).

For information about performance route selection criteria, see Winner Set in Chapter 3: Adaptive Path Control of the *Adaptive Path Controller Guide*.

Use the **no** form of the command to disable reliability status evaluation and remove the command from the configuration.

prefix-cleanup schedule (deprecated)

The prefix-cleanup schedule command has been deprecated. See the replacement command, subnet-cleanup schedule on page 407.

prefix (deprecated)

The prefix command has been deprecated. See the replacement command, subnet on page 407.

prefix resync (deprecated)

The prefix resync command has been deprecated. See the replacement command, subnet resync on page 409.

prefix-update-rate (deprecated)

The prefix-update-rate command has been deprecated. See the replacement command, subnet-update-rate on page 411.

priority-subnets

Enables priority handling in a decision-policy configuration block.

Syntax

```
priority-subnet
no priority-subnet
```

Prerequisites

Must be in policy configuration mode (*config-engine-policy*), which requires privileged-level access.

Occurrence

Use once per policy block.

Default

Disabled

Description

When priority handling is enabled in a decision-policy configuration block, all subnets associated with the policy will be given expedited handling by the decision-maker process. Non-priority subnets will be handled at a lower priority.

See the **`prefix-update-rate (deprecated)`** command for information about queues.

For more information about policies, see the **`decision-policy`** command and Chapter 6: Decision Policies and Application Models of the *Monitoring Guide*.

for information about decision-policy configuration blocks.

Use the **no** form of the command to disable priority handling for the policy block.

---

probe-backoff

Defines the length of timeout for probing and sets criteria for CNA's response to a timeout for probes with an active-measurement group.

Syntax

```
probe-backoff [ trigger <time> ] [ wait <delay> ]
no probe-backoff
```

Arguments

- *time*—number of seconds to wait for a subnet to respond to probing; range: 1 through 31556736
- *delay*—number of seconds to wait after a timeout has been triggered before resumption of probes; range: 0 through 31556736

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once.

Default

timeout—3600

delay—0

Description

Use the **probe-backoff** command to define the length of time the CNA *am-scheduler* process should wait for a response to before it stops probing an unreachable or non-responsive target over an individual link, and the length of time the *am-scheduler* process should wait to resume probing that target again.

The **probe-backoff** command sets values which will apply only to active measurement targets that are associated with the active-measurement group. Trigger and wait times for probes that are not in a group are set by the **active-measurement probe-backoff** command.

You can set either or both (or none) of the values. Any value that you do not set when the command is added to the configuration will revert to the default. Entering the command without specifying a value is the equivalent of the **no** form of the command, which also reverts the values to their defaults.

You should be careful not to set the trigger interval smaller than the timeout value in the probe configuration. This will cause probing to be suspended after a single packet times out.

The trigger time should also take into account the number of probe failures that define a blackout (a user-configured number; see the **outage-detection active-probe** command). The default trigger time assumes a very low probing rate. The trigger time should be at least as large as the configured number of probe failures multiplied by the time between individual probes. You may want to set it larger than this calculated value to accommodate periodic fluctuations in subnet activity.

Note:

If you set *delay* to 0 (the default value), the *am-scheduler* process will not wait any length of time before resumption of probing after a timeout has been triggered, effectively turning off the probe-backoff feature.

See Chapter 4: Server-Based Measurement of the *Monitoring Guide*.

provider-as

Sets the Autonomous System number on a service provider link.

Syntax

```
provider-as <number> <address>
no provider-as
```

Arguments

- *number*—integer in the range 0 through 65,535; a valid AS number (top of range is reserved for private AS numbers)
- *address*—hostname or IP address of the service provider link, in dotted decimal format (a.b.c.d)

Prerequisites

Must be in link configuration mode (*config-engine-link*—see **link** command), which requires privileged-level access.

Occurrence

Use once.

Description

The **no** form of the command removes the link AS from the configuration.

publish report

Sets a schedule for automatic uploading of reports (may be used in config as well as executed directly).

Syntax

```
publish report <report>
     [ <object> ]
     [ customer { <name> | all } ]
     <file-url>
     schedule <minute> [ <hour> [ <date> [ <month> [ <day> ]]]]
no publish report <report> [ <object> ]
```

Arguments

- *report*—name of report to be published; see the **email** command for a current list of reports (some reports, as indicated in that list, require a router name.
- *object*—one of the following:
  - name or IP address of router
- *name*—name of the customer for which the report is generated. The customer-specific subnet list will be used when generating such a report. If you use '**all**' for the name, then reports for all defined customers will be sent.
- *file-url*—you may specify the resulting file with either of these two methods:
  - directory name—when you end the name with a slash "/" character, then the file name will be derived from the time it was generated. This will ensure that the remote file will have a unique name and thus will not be overwritten.
  - file name—use this (without the ending slash "/" character) instead of a directory name when you want each scheduled upload to overwrite the previous one.
- *minute*—number of minutes after the hour when the report is to be mailed; range: 0 through 59 (or * for every minute, or a comma-separated list of minutes, such as 5,25,45)
- *hour*—hour of the day when the report is to be mailed; range: 0 through 23, or * for every hour (or * for every hour, or a comma-separated list of hours, such as 10,22)
- *date*—day of the month when the report is to be mailed; range: 1 through 31, or * for every day of the month (or * for each date of the month, or a comma-separated list of dates, such as 1,15)
- *month*—month report is to be mailed; range: 1 through 12 (or * for every month, or a comma-separated list of months, such as 1,7)
- *day*—day of week when the report is to be mailed; range 0 through 7 (or * for every day of the week, or a comma-separated list of days, such as 0,3); 0 and 7 both mean Sunday

Prerequisites

Must be in report configuration mode (*config-report*), which requires privileged-level access, to set a schedule. One-time publishing of a report can be done with log-in level access (exec mode).

Occurrence

Available for use multiple times in the configuration.

Default

* for all schedule increments (except *minute,* which must be specified)

Description

Use the **`publish report`** command to schedule automatic publishing of reports. You may set one schedule for each report; enter the command once for each.

You may also use this command directly from the command line (exec mode) to publish reports immediately without needing to schedule it. In this case, you cannot use the **schedule** keyword.

Reports will be sent by the management module via FTP.

Use the **no** form of the command with no parameters to cancel all automatic publishing.

Use the **no** form of the command with a report name (and object type if required) to cancel a specific publish schedule.

---

publish trend

Sets a schedule for automatic publishing of trends (may be used in config as well as executed directly).

Syntax

```
publish trend <filename>
     [ <object> ]
     [ customer { <name> | all } ]
     [ application <app-name> ]
     [ subnet <subnet> ]
     [ to <list> ]
     schedule <minute> [ <hour> [ <date> [ <month> [ <day> ]]]]
no publish [ <report> [ <object> ]]
```

Arguments

- *filename*—name of .qry or .tmpl file that defines the trend to be published.
- *subnet*—subnet in CIDR format (a.b.c.d /*n*); generates a trend specific to this subnet.
- *app-name*—name of an application (for example, http, icmp, ftp, sap, etc.). One way to see what applications have been found by CNA is to use the command **`show efc application`**.
- *file-url*—you may specify the resulting file with either of these two methods:
  - directory name—when you end the name with a slash “/” character, then the file name will be derived from the time it was generated. This will ensure that the remote file will have a unique name and thus will not be overwritten.
  - file name—use this (without the ending slash “/” character) instead of a directory name when you want each scheduled upload to overwrite the previous one.

- *minute*—number of minutes after the hour when the report is to be mailed; range: 0 through 59 (or * for every minute, or a comma-separated list of minutes, such as 5,25,45)
- *hour*—hour of the day when the report is to be mailed; range: 0 through 23, or * for every hour (or * for every hour, or a comma-separated list of hours, such as 10,22)
- *date*—day of the month when the report is to be mailed; range: 1 through 31, or * for every day of the month (or * for each date of the month, or a comma-separated list of dates, such as 1,15)
- *month*—month report is to be mailed; range: 1 through 12 (or * for every month, or a comma-separated list of months, such as 1,7)
- *day*—day of week when the report is to be mailed; range 0 through 7 (or * for every day of the week, or a comma-separated list of days, such as 0,3); 0 and 7 both mean Sunday

Prerequisites

Must be in report configuration mode (*config-report*), which requires privileged-level access, to set a schedule. One-time publishing of a report can be done with log-in level access (exec mode).

Occurrence

Available for use multiple times in the configuration.

Default

* for all schedule increments (except *minute*)

Description

Use the **`publish trend`** command to schedule automatic publishing of trends. You may set one schedule for each trend; enter the command once for each.

You may also use this command directly from the command line (exec mode) to publish reports immediately without needing to schedule it. In this case, you cannot use the **schedule** keyword.

Trends will be sent by the management module via FTP.

Use the **no** form of the command with no parameters to cancel all automatic publishing.

Use the **no** form of the command with a trend name (and object type if required) to cancel a specific publish schedule.

---

radius-server host

Defines the RADIUS server to be used by CNA for authentication.

Syntax

```
radius-server host <address>
                  [auth-port <port>]
                  [timeout <seconds>]
                  [key <string>]

no radius-server-host <address>
```

Arguments

- *address*—the RADIUS server's hostname or IP address in dotted decimal format (a.b.c.d)
- *port*—authentication port number
- *seconds*—number of seconds before connection times out (range: 1 through 30)
- *string*—RADIUS encryption key

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

port—1812

Description

The **no** form of the command disables the definition for the specified RADIUS server.

Values set for *port*, *seconds* and *string* apply only to the current server definition. See **`radius-server key`** and **`radius-server timeout`** commands for information about setting global values.

---

radius-server key

Defines the global encryption key used to exchange information with all RADIUS servers for which a key was not explicitly set.

Syntax

```
radius-server key <string>
no radius-server key
```

Arguments

- *string*—RADIUS encryption key

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Description

The **no** form of the command deletes the global RADIUS encryption key.

If an encryption key was explicitly set for a particular RADIUS server (see the **`radius-server host`** command), that key will override the global key for that RADIUS server host.

---

radius-server timeout

Defines a global RADIUS server timeout value—the amount of time in seconds CNA will wait for a response from the RADIUS server before retransmitting the authentication request.

Syntax

```
radius-server timeout <seconds>
no radius-server timeout
```

Arguments

- *seconds*—number of seconds to wait; range: 1 through 30

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

3

Description

The **no** form of the command deletes the global timeout value.

If a timeout value is explicitly set for a particular RADIUS server (see the **`radius-server host`** command), that value will override the global timeout value for that RADIUS server.

---

rate

Defines the rate at which active probes should be sent.

Syntax

```
rate <value> <unit>
no rate
```

Arguments

- *value*—number portion of the rate
- *unit*—unit portion of the rate; must be one of:
  - *per-second*
  - *per-minute*
  - *per-hour*
  - *per-day*

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once per active-measurement group.

Default

3 per minute

Description

Use the **`rate`** command while in active-measurement group configuration mode (*config-engine-active*) to set a group-wide rate at which probes will be contacted.

This rate will be used for all targets that are members of the group but for which an individual rate was not specified.

See the **`target`** command for information about setting individual probe rates.

See the **active-measurement group** command for an overview of active-measurement groups.

Use the **no** form of the command to remove a rate from a group.

rate (Chatter module only)

Defines the test frequency for this type of active measurement test.

Important:
This measurement rate applies to all CNA agents in the Chatter subnet performing this type of test. It is not the measurement rate between two specific CNA agents.

Syntax

rate <value> {per-second | per-minute | per-hour | per-day}

no rate

Arguments

Table 32: rate arguments

Argument	Description
<*value*>	Number of tests conducted during the time interval that you specify. Enter a positive, non-zero integer. The range is from 1 to 49,000.
{per-second | per-minute | per-hour | per-day}	Time interval for the test frequency.

Prerequisites

Must be in Chatter Active Measurement Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrences

Use once.

Default

5 tests per minute

Description

Defines the test frequency for this type of active measurement test.

Important:
This measurement rate applies to all CNA agents in the Chatter subnet performing this type of test. It is not the measurement rate between two specific CNA agents.

Defines the overall measurement rate for this type of active measurement test conducted across the Chatter subnet.

The no command restores the default measurement rate of 5 tests per minutes.

There is a trade-off between setting the test rate:

- High enough to quickly detect down links and other problems

  AND

- Low enough so that all of the tests together do not generate a load that exceeds the hive or end point capacity.

In the ideal case, with unlimited capacity, set the rate based on the zone hierarchy as follows:

1. Follow down the levels of the hierarchy. At each branch point, add the number of paths as follows:
   - For each two-way branch, add 2 paths.
   - For each three-way branch, add 6 paths.
   - For each four-way branch, add 12 paths.
   - For each five-way branch, add 20 paths.
   - Do not add any for zones that contain endpoints instead of subzones.
2. Divide the total number of paths by the number of minutes desired between tests.

   For example, 100 paths to be checked every minute means a rate of 100 tests per minute. If two minutes between tests is acceptable, then 50 tests per minute is sufficient.

To find the best rate for your system, Avaya recommends the following steps:

1. Start the system with the default setting (5 tests per minute).
2. Wait a while and then look at several links and note the Inter Arrival Time (time between tests).
3. If the Inter Arrival Time is too long, restart the system with a higher rate.
4. Look at the link test spacing and the drops to see if tests are being dropped. If tests are being dropped, the total frequencies (for all types of tests) is too high.
5. Raise the frequency if the system can handle it.

recover report

Recovers from the most recently executed **delete report command**.

Syntax

```
recover report
```

Prerequisites

System operations mode, privileged-level access.

Description

Use the **recover report** command to recover snapshot reports inadvertently deleted.

The **recover report** command will undo the most recently executed **delete report** command.

When a report or group of reports is deleted, the CNA software moves the deleted material to a user-created data space on the hard disk.

The system will only hold the results of a single delete operation, which will be shown as *recycled-reports* in a listing of files by the **ls** command. You can delete this file with the **delete** command.

There is only one user-data area on the hard disk. All data created by all users in all CLI or web sessions is written to this single area.

A user deleting reports in a separate CLI session can cause reports deleted by another user to be irretrievably lost. Only the most recent delete operation is recoverable.

In addition, the system will not retrieve a report from this discard directory if a copy of the report already exists—if, for example, an interim billing report has been regenerated after the delete but prior to the recovery attempt, or if after deleting the reports you restore an image which contains one of the deleted reports. The recover operation will not overwrite the existing report.

recover startup-config

Restores startup configuration from back-up copy.

Syntax

```
recover startup-config
```

Prerequisites

System operations mode, privileged-level access.

Occurrence

Use once.

Description

Each time you replace the CNA startup configuration, a backup copy of the current startup configuration is made before the new one is written. The **`recover startup-config`** command retrieves this backed-up copy and makes it the current startup configuration.

There is only one backup configuration. If you use the **`copy`** command to overwrite startup configuration A with the contents of configuration file B, the system first saves A as the backup and then makes B the current startup configuration. If you then execute the **`recover startup-config`** command, B becomes the backup and A becomes the current startup configuration.

---

refresh-time acquire

Sets the time to wait between attempts to find a convergence point for a target for the first time.

Syntax

```
refresh-time acquire [ seconds ]
no refersh-time acquire
```

Arguments

- *seconds*—time in seconds, in the range 0 through 31,556,736 (1 year).

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once.

Default

21600

Description

Use the **refresh-time acquire** command to set the number of seconds to wait between each attempt to find a new convergence point. This applies only until a convergence point for a particular target has been found for the first time. After a convergence point has been found for the first time, but subsequently becomes unreachable, the time specified with the **reacquire** key word will apply.

Use the **no** form of the command to reset the value to its default.

---

refresh-time host-disappeared

Sets the time to wait before searching for a new convergence point after a host has disappeared.

Syntax

```
refresh-time host-disappeared [ seconds ]
no refresh-time host-disappeared
```

Arguments

- *seconds*—time in seconds, in the range 0 through 31,556,736 (1 year).

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once.

Default

0

Description

Use the **refresh-time host-disappeared** command to set the number of seconds to wait (after a *host* point is unreachable over all links) before searching for a new convergence point.

Use the **no** form of the command to reset the value to its default.

refresh-time proxy-disappeared

Sets the time to wait after a convergence point has disappeared before searching for a new one.

Syntax

```
refresh-time proxy-disappeared [ seconds ]
no refresh-time proxy-disappeared
```

Arguments

- *seconds*—time in seconds, in the range 0 through 31,556,736 (1 year).

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once.

Default

0

Description

Use the **`refresh-time proxy-disappeared`** command to set the number of seconds to wait (after a *convergence* point is unreachable over all links) before searching for a new convergence point.

Use the **no** form of the command to reset the value to its default.

refresh-time reacquire

Sets the time to wait between subsequent attempts to find a convergence point for a target.

Syntax

```
refresh-time reacquire [ seconds ]
no refersh-time reacquire
```

Arguments

- *seconds*—time in seconds, in the range 0 through 31,556,736 (1 year).

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once.

Default

1800

Description

Use the **`refresh-time reacquire`** command to set the number of seconds to wait between each attempt to find a new convergence point. This applies after a convergence point for a particular target had been found but was subsequently lost or invalid. For targets where a convergence point has never been found, the time specified with the **acquire** key word will apply.

Use the **no** form of the command to reset the value to its default.

---

refresh-time verify

Sets the time to wait between attempts to verify that a convergence point is still on the routing path to a particular target.

Syntax

```
refresh-time verify [ seconds ]
no refersh-time verify
```

Arguments

- *seconds*—time in seconds, in the range 0 through 31,556,736 (1 year).

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once.

Default

3600

Description

Use the **`refresh-time verify`** command to set the number of seconds to wait between each attempt to verify that a convergence point is still on the routing path to a particular target.

Use the **no** form of the command to reset the value to its default.

---

reliability half-life

Sets the half-life parameter which influences link reliability.

Syntax

```
reliability half-life <value>
no reliability half-life
```

Arguments

- *value*—half-life in the range 600 through 14400

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once.

Default

1440

Description

Use the **`reliability half-life`** command to change the half-life (in minutes) of the reliability metric. The metric is calculated for each subnet on a link as follows: The first time it appears in the winner-set it is given a value of zero (this is the best score). Then, each time it moves out of the winner-set, this score is increased by one. Thus, subnets on a link which move

out (and then eventually back in) of the winner set more often than on other links will be regarded as less reliable than ones which tend to stay in the winner set. The score is reduced by half when the subnet stays in the winner set for the number of minutes specified by half-life.

Use the **no** form of the command to revert to the default value.

reliability tolerance

Sets the tolerance parameter which influences link reliability.

Syntax

```
reliability tolerance <value>
no reliability tolerance
```

Arguments

- *value*—tolerance in the range 0 through 100

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once.

Default

10

Description

Use the **`reliability tolerance`** command to set the tolerance of the reliability metrics between links in a winner set where they will still be considered to be equal. For example, if there are two links in a winner set that the CNA decision maker process calculates to be equal (after it has looked at other, higher-priority tie-breaks and still considers them equal) then it will use link with the better reliability metric as long as the difference between the two is greater than the **tolerance**. If it is not greater, then the links will still be tied and further tie-breaks will be used.

Refer to the command **`reliability half-life`** to see how link metrics are determined.

Use the **no** form of the command to revert to the default value.

For information about performance route selection criteria, see Winner Set in Chapter 3: Adaptive Path Control of the *Adaptive Path Controller Guide*.

reload

Restarts CNA, using the configuration saved as the startup-config.

Syntax

```
reload [ erase ]
       [ cancel ]
       [ cna-agent-device <address> ]
       [ { in <hhh:mm> } | { at <hh:mm> [ <mm-dd> ] } [ text ] ]
```

Arguments

- *hhh:mm*—specifies the number of hours and minutes from now that the reload is to occur. (*hhh* in the range 0 through 999; *mm* in the range 0 through 59; 000:00 is not allowed)
- *hh:mm*—specifies the exact time of day (24-hour clock) at which the reload is to occur. (*hh* in the range 0 through 23; *mm* in the range 0 through 59
- *mm-dd*—specifies the month and day when the reload is to occur. (*mm* in the range 1 through 12; dd in the range 1 through 31
- *text*—specifies the text to be displayed when a **`show reload`** command is issued.

Prerequisites

System operations mode, privileged-level access.

Description

When the system is rebooted by the **`reload`** command, all performance and subnet data will be lost. The system will need to rebuild its database, including all traffic test results, after rebooting. This can take a significant amount of time.

Use the **erase** keyword to erase all data from the device. When the reload is complete, the device will be in an as-new, out-of-the-box state. Because of the severity of this action, you will be prompted twice for confirmation.

Use the **in** keyword to schedule a reload to occur some number of hours and minutes from now.

Use the **at** keyword to schedule a reload to occur at a specified time in the future.

Use the **`show reload`** command to see the details of a scheduled reload, if any.

remark

Associates a block of descriptive text with an access list.

Syntax

```
remark <text>
no remark <text>
```

Arguments

- *text*—text of the remark

Prerequisites

Must be in access-list configuration mode (*config-acl*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **no** form of the command to remove a remark.

---

remote-exec

Places CNA in remote-exec mode, allowing you to manage multiple devices from a single location using the command line.

Syntax

```
remote-exec
```

Prerequisites

Available in login and enable modes.

Occurrence

Available for use multiple times in the configuration.

Description

The **remote-exec** command allows exec commands entered in this mode to be executed on all servers in the hive, including the local server.

Exec commands are available in both login and enable modes. Configuration commands are not allowed and can only be executed on the local server. In login mode, **remote-exec** will only access login commands, which do not require the enable password. When **remote-exec** is used in enable mode, most exec commands are available. However, there are some

commands that unavailable with **remote-exec**. Following is a list of the commands that are not available:

- **ssl generate-csr**
- **generate -config**
- **chatter join-hive**
- **chatter leave-hive**
- **chatter assign**
- **copy scp**
- **copy chatter-config**
- **email trend**
- **publish trend**
- **ssl add**
- **ssl export**
- **show debugging**
- **debug on**
- **debug display**
- **help**
- **history**
- **terminal**
- **tree**
- **tcpdump (without the count parameter spcified)**
- **ping (without the count parameter specified)**
- **eventdump**
- **dump-measurements**
- **traceroute-vips continuous**

Note:

If previous commands are attempted with **remote-exec**, the software will generate a user error in the CLI.

Devices can be grouped in hives and exchange information using the chatter module. All servers in the hive are aware of each other and have the ability to view information and run commands on any other server in the hive. For more information about hives, see the *Monitoring Guide*.

If the command returns output, the output information will be concatenated with one server information after the other. The **show version** command and **restore** command have

custom output. This output is collected, processed, and presented to the user in a form that would be useful for the particular command.

The complete results for the last **remote-exec** command executed. The results are saved to print normal or custom outputs. The results are named `remote-exec-last.txt`, and are saved in the user's space. The results can be accessed with the **more remote-exec-last.txt** command.

The **remote-exec** command can take a parameter, a comma separated list of servers, to execute the commands on. Using a parameter allows the command to be executed on specific servers instead of the default, all servers. This list of servers must be a subset of the servers in the hive. The list of servers in the hive can be obtained by using the **show hive** command.

rename

Renames objects in the CNA repository.

Syntax

```
rename <old> <new>
```

Arguments

- *old*—current name of the object
- *new*—new name of the object

Prerequisites

System operations mode, privileged-level access.

Description

Use the **ls** command to obtain a list of objects.

rename zone

Renames a specific zone.

Syntax

rename zone <orig-name> <new-name>

Arguments

Table 33: rename zone command arguments

Argument	Description
<orig-name>	The specific, existing name of the zone or subzone that you are renaming. Enter just the name, not the fully qualified name (x.y.z).
<new-name>	New name for the zone or subzone. Enter just the name, not the fully qualified name (x.y.z). Restrictions on the name are: • The name cannot be an integer. • The name must start with a letter (upper or lower case) and be followed by a letter (upper or lower case), digit, dash or underscore. • All zone and subzone names must be unique.

Prerequisites

None.

Description

Renames a specific zone.

Note:

The hierarchy of a zone cannot be changed. You cannot use this command to move a subzone from one zone to another. To change the hierarchy of a zone, you must delete it and then recreate it with the new hierarchy.

The **rename** command is designed primarily to correct an error in a specific zone name. To enforce this requirement, the zone name in the command cannot contain dots. The zone name that you enter must be a single zone name, not a fully qualified name.

repeat

Defines the number of times that this particular active measurement test is conducted.

Syntax

repeat <value>

no repeat

Arguments

Table 34: repeat arguments

Argument	Description
<*value*>	Number of times this test is conducted. Zero indicates that the test is conducted forever. The valid range for this variable is 0 (zero) to 1000.

Prerequisites

Must be in Chatter Active Measurement Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use once.

Default

0, which indicates forever.

Description

Defines the number of times that this particular active measurement test is conducted. A value of 0 (zero) indicates that the test is conducted forever.

The no command restores the default value of 0.

report group

Assigns subnets to a named group for generating CNA reports.

Syntax

```
report group <name> { group <generic-group> | subnet <subnet> }
no report group <name>
no report group <name> subnet <subnet>
no report group <name> group <generic-group>
```

Arguments

- *name*—name of the report
- *subnet*—subnet in CIDR format (a.b.c.d /*n*)
- *generic-group-name*—name of a previously defined group (see the **group** command)

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration (for multiple subnets in a given report, and for multiple report names).

Description

Use the **`report group`** command to define a collection of subnets that will be labeled by the specified group name. This will allow you to sort by group name in snapshot reports, and gather trend data for subnets of interest. You can also associate previously created groups of subnets with the report group when you use the second **group** keyword (instead of **subnet**). Refer to the **`group`** command see how to create generic groups.

A group is created whenever you associate one or more subnets with a name using the **`report group`** command.

If you add a group called *trend-live* to your configuration, CNA will gather trend data for any subnet associated with the group at a fixed sampling rate of every 30 seconds (subnet-related trend data is normally sampled once every hour; see the command **`trending on`**). The trending database will maintain the number of data points specified by the *performance* family for subnet-related data.

The group column in a snapshot reports will not show more than one group for a subnet. If a subnet is in both the *trend-live* group as well as another user-defined group, the snapshot report will show only the user-defined group name.

The **no** form of the command without a subnet specification deletes the named report group.

The **no** form of the command with a subnet specification deletes the subnet from the named report group.

report prefix-alias (deprecated)

The report prefix-alias command has been deprecated. See the replacement command, report subnet-alias on page 289.

report prefix-limit (deprecated)

The report prefix-limit command has been deprecated. See the replacement command, report subnet-limit on page 289.

report subnet-alias

Associates a subnet with an alias to be used by CNA when generating reports.

Syntax

```
report subnet-alias <alias> <subnet>
no report subnet-alias <alias>
```

Arguments

- *alias*—name of the alias
- *subnet*—subnet address, in CIDR format (a.b.c.d /*n*)

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration. Each subnet can only be used once. An alias can be assigned to more than one subnet.

Description

The **no** form of the command removes the specified alias from the configuration.

The **`report subnet-alias`** command associates subnets and aliases in an internal configuration file.

The alias is used only for generating reports.

In contrast, the **`performance-group`** command associates subnets under a common name for internal group processing by CNA.

report subnet-limit

Sets a limit on the number of rows displayed when the **`show prefix (deprecated)`** command is executed.

Syntax

```
report subnet-limit <limit>
no report subnet-limit
```

Arguments

- *limit*—number of rows to display (range: 1 through 200,000)

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

500

Description

The **no** form of the command resets *limit* to the default.

Use the **`report subnet-limit`** command to establish a new default value for the *limit* argument of the **`show prefix (deprecated)`** command. This value determines the number of rows displayed when you execute the **`show prefix (deprecated)`** command without an explicitly set limit. You can also specify a number each time you execute the **`show`** command.

---

require-bgp-snmp-poll

Sets a flag used by the CNA decision maker process when determining the status of a link.

Syntax

```
require-bgp-snmp-poll
no require-bgp-snmp-poll
```

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Use once.

Default

Disabled

Description

When **require-bgp-snmp-poll** behavior is enabled by adding this command to the CNA configuration, the CNA decision-maker process will immediately stop route-assert mode if the CNA link monitor process reports the status of the BGP peering as unknown consecutively a user-specified number of times. This number is set by the **`polling bgp-status unknown-threshold`** command.

The behavior is disabled by default because of its dependence on SNMP, which must be running on your edge router.

Once you have enabled SNMP and configured the CNA link monitor, you can add the **`require-bgp-snmp-poll`** command to your CNA configuration. For additional information about this behavior, see Chapter 3: Adaptive Path Control of the *Adaptive Path Controller Guide*.

See the **`polling bgp-status on`** command for information about CNA link monitor and SNMP polling.

Use the **no** form of the command to disable **require-bgp-snmp-poll** behavior.

---

require-routes-from-all-routers

Enables and disables *require* behavior, which determines whether CNA will assert a route before hearing about the subnet from all edge routers.

Syntax

```
require-routes-from-all-routers
no require-routes-from-all-routers
```

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Use once.

Default

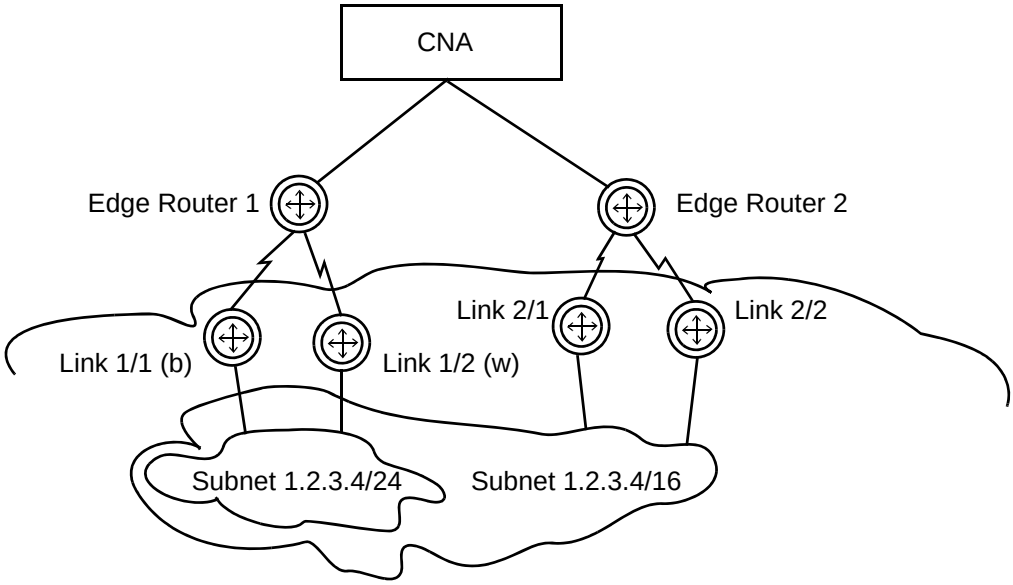
Require behavior is enabled.

Description

When *require* behavior is enabled, CNA will not assert a route until all of its BGP peers have reported routes to the subnet. This behavior prevents CNA from being the source of routing information previously unknown to the router.

In the following example, there are two links on each of two routers: 1/1, 1/2, 2/1 and 2/2.

Figure 1: Asserting routes with multiple routers



Router 1 has received the minimum required number of measurements for subnet 1.2.3.4/24 across both Link1/1 (the BGP selected route to the subnet, labeled “b”) and Link 1/2, and CNA has determined that Link 1/2 is the performance-route winner (labeled “w”).

If you have *disabled* require behavior, CNA will assert Link 1/2 as the winner route to subnet 1.2.3.4/24 to both routers, even though Router 2 has reported that it knows only about subnet 1.2.3.4/16.

With *require* behavior enabled (the default behavior), CNA will not assert any route to subnet 1.2.3.4/24 until Edge Router 2 has sent notice that it knows about the subnet. This will result in Edge Router 1 continuing to use Link 1/1, the BGP selected route, to subnet 1.2.3.4/24, even though CNA has determined that Link 1/2 is the better one.

Use the **no** form of the **`require-routes-from-all-routers`** command to disable this behavior.

restart decision-maker

Restarts the CNA *decision-maker* process.

Syntax

```
restart decision-maker
```

Prerequisites

Must be in system-operations mode, with privileged-level access.

Description

The **restart decision-maker** command restarts the CNA *decision-maker* process. You will need to do this whenever you make a topological change to your CNA configuration, such as adding or removing routers and service provider links.

Restarting the *decision-maker* process will bring down all BGP peerings and clear all performance measurements from the internal database.

BGP peerings will have to be reestablished and subnets relearned; additionally, routes will not be selected for any subnet until it has a minimum number of measurements in the database (see the **min-required-measurements** command).

---

restore

Retrieves a specified object from a backup.

Syntax

```
restore {image | image-config | license } <file-url>
```

Arguments

- *file-url*—location of the backup copy

Prerequisites

Must be in system-operations mode, with privileged-level access.

Description

The **restore** command restores a license file, or an image from an archival copy that was previously created with the **backup** command (the archive can contain the entire CNA image or the entire image plus startup configuration).

For example, if your Avaya support user ID were *rs_customer* and your password were *p_wrD*, you could use the following command to restore the 2.1.2 image on your CNA device from the image stored on the Avaya FTP server at *ftp.avaya.com*:

```
restore image ftp://rs_customer:p_wrD@ftp.avaya.com/releases/fcs-2.1.2.img
```

If you specify the **image** keyword when restoring a backup file that contains both operating system image and configurations, only the image will be extracted.

If you specify the **image-config** keyword (instead of **image**, in the example) and the URL identifies an image you previously backed up with the **backup** command and **image-config** keyword, the configuration that was stored with the image when it was backed up will overwrite your existing startup config.

If you specify **image-config** when restoring a backup file that contains only an image, the restore will succeed (with no action taken regarding the missing config). See the **backup** command description for information about saving a config with images during backup.

Use the **license** keyword to restore a license file (*.lic*) to the device.

When upgrading to an image downloaded from the Avaya FTP server, be aware that those images do contain a startup config with minimal content. If you upgrade with the **image-config** keyword, CNA will be restored to an as-new condition.

The archive can be located either locally on CNA, or a remote location accessible by FTP or HTTP. See the **copy** command for URL syntax.

If you want to halt a restoration in progress, use **Control-C** (you may have to wait a few minutes and retry if the process is in the midst of a critical step and ignores the halt signal.)

To restore an image from a remote device, you will need free disk space on CNA equal to twice the size of the stored object. Use the **ls** command to display the repository's current contents, including both the name of the object and the amount of space it occupies.

To restore an image stored locally on CNA, you need free space equal to the size of the image.

When you attempt to restore a file, the device will calculate the amount of space required and warn you if there is not enough space currently available.

See the note about image size in the description of the **backup** command.

During the restoration process, progress is indicated by a display of dots that appear at intervals during the process, as shown in the following example:

```
Checking local disk space          ..                        [OK]
Verifying image                    .. ..                     [OK]
Upgrading management               .. .. .. ..               [OK]
Preparing for peripheral upgrades  .. .. .. .. ..            [OK]
Upgrading engine module in slot 2  .. ..                     [OK]
Upgrading ustat module in slot 3   .. ..                     [OK]
Upgrading ustat module in slot 4   .. ..                     [OK]
Upgrading ustat module in slot 7   .. ..                     [OK]
Upgrading ustat module in slot 12  .. ..                     [OK]
Local image (0616.img) was successfully restored
```

Some advisory messages may be displayed between steps but, as along as the line ends with the text *OK*, the restore will continue. If any other text is displayed inside the brackets, check the log for additional information.

Dot pairs indicate completion of individual steps in the process. You can interrupt the restore process with the **Control-C** keys, but be aware that the intervals between display of dot pairs is not uniform and can take from a few seconds to several minutes. Also, while the number of dot

pairs displayed for any one step is fixed, neither the number of dots nor the steps themselves are guaranteed to remain static from release to release.

resync chatter chapld

Manually resynchronizes the local Chatter module's data with that of other Chatter modules in the hive.

Syntax

resync chatter chapld

Prerequisites

Must be in privileged command mode.

Description

Manually resynchronizes the local Chatter module's data with that of other Chatter modules in the hive. This data is periodically synchronized automatically. But this command makes it possible for you to force a resynchronization.

resync chatter node-alias

Refreshes the list of node names that are known to Chatter.

Syntax

resync chatter node-alias

Prerequisites

Must be in Privileged command mode.

Description

Refreshes the list of node names that are known to Chatter.

resync efc filter

Updates an access list that is used as a filter for EFC data flows.

Syntax

```
resync efc filter [ <name> ]
```

Arguments

- *name*—name of a configured access list being used as an efc filter

Prerequisites

System-operations mode, with privileged-level access.

Description

After making changes to an access list that has been referenced by the **efc filter** within an active measurement group configuration block, use the **`resync efc filter`** command to update the list in the CNA *efcserver* process.

If you specify a *name* argument, only the named list will be updated. If you specify a list that has not been referenced as a filter, the command will have no effect.

If you omit the *name* argument, all configured access lists that are known to the *efcserver* process will be updated.

resync load threshold-table

Commits the specified load threshold table to the CNA *decision-maker* database.

Syntax

```
resync load threshold-table <name>
```

Arguments

- *name*—name of the table

Prerequisites

Must be in systems-operation mode with privileged-level access.

Occurrence

Use once.

Description

Newly configured load threshold tables, or changes to existing tables, are not implemented until you execute the **`resync load threshold-table`** command.

See the **`load threshold-table`** command.

resync prefix (deprecated)

The resync prefix command has been deprecated. See the replacement command, resync subnet on page 297.

resync subnet

Forces the CNA *decision-maker* process to immediately reevaluate specified subnets.

Syntax

```
resync subnet { all |
                confirmed |
                asserted |
                selected |
                adequate |
                measured |
                using-policy <name> |
                quiet <qtime> |
                <addr>
              }
              {
                compute-winner-set |
                clear |
                clear-routes-only |
                reassert-route |
                verify-auto-created
              }
              [ limit <num> ]
```

Arguments

- *name*—name of a decision policy
- *qtime*—quiet time (number of seconds since last measurement); range: 60 through 86400
- *addr*—subnet address in CIDR format (*a.b.c.d/n*)
- *num*—maximum number of subnets to re-evaluate; range: 1 through 1000000

Prerequisites

Must be in module engine mode with privileged level access.

Default

num—60

Description

The **`resync subnet`** command is the *exec* (immediately executing) form of the **`subnet resync configuration`** command, which results in a periodic reevaluation of subnets.

Use one of the following keywords with either version of the command:

- **all**—all subnets
- **confirmed**—all subnets that have been asserted by CNA and confirmed as in use by the routers
- **asserted**—all subnets that have been asserted (includes those that have been confirmed)
- **selected**—all subnets that have been selected as the winning route by CNA (includes those that have been asserted)
- **adequate**—all subnets that have a sufficient number of measurements in the CNA database to be considered adequately measured
- **measured**—all subnets that have at least one measurement in the database
- **using-policy**—all subnets associated with the named decision policy
- **quiet**—all subnets for which no measurements have been received for the specified period of time

Or, you can specify an individual subnet by its address in CIDR format (a.b.c.d/n).

You can impose a limit on the number of subnets that will be evaluated regardless of the size of the set, by setting the *limit* argument.

The *limit* argument can be used with both **`subnet resync`**, the *config* form of this command, and **`resync subnet`**, the *exec* form. But its greatest value is with the config form, when it can be used to avoid thrashing (nonstop reassertions of routes) if your interval is small and the set large. In some cases, it could take at least as long as the interval to re-evaluate all of the subnets in the set, causing the *decision-maker* process to immediately start a new **resync** cycle immediately upon completion of the previous evaluation.

You must also choose one of the following actions to be taken by the CNA *decision-maker* process:

- **compute-winner-set**—the *decision-maker* process will recompute the winner-set for the specified set of subnets and reassert the current winner if it is in the winner set
- **clear**—the *decision-maker* process will withdraw any CNA-asserted routes for the specified subnets; this is the equivalent of you executing the **clear prefix (deprecated)** command for the specified subnets
- **clear-routes-only**—the *decision-maker* process will withdraw any CNA-asserted routes for the specified subnets without resetting the counters which CNA uses to keep track of user-traffic test hits; this is the equivalent of you executing the **clear prefix (deprecated)** command for the specified subnets with **routes-only** keyword
- **reassert-route**—the *decision-maker* process will reassert the route that was most recently selected as the winning route by the normal subnet evaluation process
- **verify-auto-created**—the *decision-maker* process will immediately run a validation check on automatically created subnets (see auto-create-subnets verify-frequency)

Use the keyword **limit** to restrict subnets that will be re-evaluated to the specified number that have gone the longest time without an evaluation.

The following example forces an immediate recomputation of winner set of the 50 asserted-and-confirmed subnets that have gone the longest without a *decision-maker* evaluation:

```
resync subnet confirmed compute-winner-set limit 50
```

retries

Tells CNA how many times to retry polling an external device

Syntax

retries <number>

no retries

Arguments

- *number*—number of times to retry polling; range: 0 through 10

Prerequisites

Must be in polling-agent configuration mode (*config-polling-snmp*), which requires privileged-level access.

Default

1

Description

When link-usage polling is on, CNA will attempt to obtain data from the external device as many times as is specified by *number*.

See the **`polling usage on`** command for more information about polling for link usage data; see the **`polling snmp-agent`** command for more information about polling parameters.

Use the **no** form of the **`retries`** command to reset the number of tries to the default.

---

route-assert-filter

Allows the user to determine conditions under which the CNA decision maker will assert routes known to BGP.

Syntax

```
route-assert-filter permit-if-in-bgp | force | deny
no route-assert-filter
```

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), which requires privileged-level access.

Occurrence

Use once globally, and once per policy configuration block.

Default

permit-if-in-bgp

Description

When **route-assert-filter permit-if-in-bgp** is enabled (the default condition), the CNA decision maker process will assert routes to a subnet if the subnet is present in the BGP feed CNA receives from its BGP peer.

If **force** is enabled, the CNA decision maker process will assert routes to a subnet regardless of whether it is in the BGP feed.

If **deny** is enabled, the CNA decision maker process will *not* assert routes to a subnet regardless of whether it is in the BGP feed.

When configured in engine configuration mode, the command is global and applies to all subnets.

When configured in policy configuration mode, the command applies only to those subnets associated with the policy.

For more information about policies, see the **decision-policy** command and Chapter 6: Decision Policies and Application Models of the *Monitoring Guide*.

In *config-engine* mode, the **no** form of the command reverts the behavior of CNA to the default.

In *config-engine-policy* mode, the **no** form of the command removes the command from the policy configuration block, which means the global setting will apply to subnets associated with the policy.

route-assert-mode

Places CNA in route-assert mode.

Syntax

```
route-assert-mode
no route-assert-mode
```

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Default

Route-assert mode is off.

Description

When route-assert mode is on, CNA will assert its preferred route (based on its analysis of user traffic tests) to its BGP peers as the best path for a given subnet.

When route-assert mode is off, best-path decisions are not asserted to routers, but are calculated for reporting purposes only.

Use the **no** form of the command to revert to the default.

Whenever you disable route-assert mode, all asserted routes are withdrawn but measurements and measurement counts will be retained in the database. When you re-enable route-assert mode, CNA will resume the decision making process immediately.

Use the **clear prefix (deprecated)** command to withdraw existing route assertions without disabling route assertion.

sample-collection

Sets the time when CNA will record a snapshot of the database.

Syntax

```
sample-collection <minutes>
no sample-collection
```

Arguments

- *minutes*—number of minutes, in the range 0 through 59

Prerequisites

Must be in report configuration mode (*config-report*), which requires privileged-level access.

Occurrence

Use once.

Default

0 (snapshot taken on the hour)

Description

CNA records a snapshot of the database once every hour.

The **`sample-collection`** command allows you to determine how many minutes past the hour the recording occurs.

This hourly snapshot is used to generate snapshot and trending reports (see **`snapshots schedule`** and **`trending on`** commands).

Use the no form of the command to set *minutes* to its default value.

send-data

Turns on a stream of outgoing data to a remote receiver.

Syntax

```
send-data <type> <address> <port>
no send data <type> <address> <port>
```

Arguments

- *type*—type of data stream; can be one of the following:
  - *decision*—send CNA *decision-maker* process data
  - *http-request*—send web-hit/HTTP request data
  - *utt*—send CNA user traffic test (UTT) data
- *address*—IP address of receiver
- *port*—port on receiver; range: 1024 through 65535

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

No data streams are enabled

Description

Use the **send-data** command to initiate an outgoing stream of data to a remove receiver.

If *type* is *utt*, the data stream will consist of measurement data from user traffic tests and active probes.

If *type* is *http-request*, the data stream will consist of logs of HTTP requests received from user traffic tests.

Multiple receivers can listen to each data stream.

For more information about the **send-data** command, contact Avaya support.

Use the **no** form of the command to turn off a specified data stream.

---

service password-encryption

When this command is executed, passwords which provide access to CNA will be shown as encrypted strings in all output (**show** and **copy** commands, for example).

Syntax

```
service password-encryption
no service password-encryption
```

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

Passwords are shown in clear text.

Description

If you add the **`service password-encryption`** command to your configuration, passwords entered in clear text will be encrypted by the CNA software as soon as you hit the **Enter** key, and the **encrypted** keyword will be added to the **`username`** command.

Passwords entered with the **ip ftp password** or **ip http password** will not be encrypted, so that they can be passed to external protocols which require clear text.

Use the **no** form of the command to disable password encryption. New passwords added to the configuration will be displayed in clear text. However, already encrypted passwords will remain encrypted. The **no** command will not toggle back to clear text.

---

service password-in-config

Determines whether passwords will be displayed in configurations, or be replaced by asterisks.

Syntax

```
service password-in-config
no service password-in-config
```

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

Passwords are displayed.

Description

Use the **no** form of the command to obscure passwords in all configurations; passwords (including both user account passwords and the **enable** password) will be displayed as a string of eight asterisks: ********.

This will cause the following behavior when you use the **copy** command, which will copy the asterisks as simple characters (without special meaning):

- If you copy the running configuration to the startup configuration, passwords will be displayed as asterisks (********) when you view the startup configuration; but the correct passwords will be correctly configured upon reboot
  - If you delete users from the running configuration and copy the startup configuration back to the running configuration, the user definitions will not be written into the running-config and the users will not have access to the device; but they will still be configured properly upon reboot
- If you copy the running configuration to a file, passwords will be exported as ********.
- If you copy a configuration from a file URL to the running configuration, passwords that are displayed as asterisks in the source file will be ignored by the **copy** command. For this reason, the string ******** is not allowed as a password.
- If you copy a configuration from a file URL to the startup configuration, passwords are copied as is; those that were displayed as asterisks in the source file will be ignored when the system is rebooted.
- When you copy to the running configuration, the **copy** command performs a merge operation, so that all passwords defined in the running configuration remain intact; after the following commands, all users would continue to be defined in the running configuration:

  ```
  copy running-config <some_file>
  copy <some_file> running_config
  ```

  However, if you delete users between copy commands, users will not remain intact; after the following sequence of actions, users would be lost:

  ```
  copy running-config <some_file>
  <delete users from running configuration>
  copy <some_file> running_config
  ```

service timestamps log

Enables timestamps on log messages.

Syntax

```
service timestamps log [uptime | datetime [msec] [show-timezone] ]
```

```
no service timestamps log
```

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

Timestamps are enabled as:

```
service timestamps log datetime
```

Description

Use the **uptime** keyword to record the number of elapsed seconds since the last reboot.

Use the **datetime** keyword to record the date and time of the message. If you include the **msec** modifier, the timestamp will specify time to the millisecond.

If you use the **show-timezone** keyword, the time zone will be identified in the timestamp.

If neither **uptime** nor **datetime** keywords are specified, **datetime** is assumed.

The **`no service timestamps log`** command resets timestamps to the default.

---

set-alarm edge

Assigns an alarm template to a Chatter edge. Each edge can have one alarm template assigned.

Syntax

set-alarm edge <ip-address1> <ip-address2> alarm <alarm-name>

no set-alarm edge <ip-address1> <ip-address2>

Arguments

Table 35: set-alarm edge arguments

Argument	Description
<ip-address1>	IP address of the source end point.
<ip-address2>	IP address of the destination end point.
<alarm-name>	Alarm template that you are assigning to the zone.

Prerequisites

Must be in Chatter Configuration command mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Available for use multiple times in the configuration. Each edge can have one alarm template assigned.

Description

Assigns an alarm template to a Chatter edge. Each edge can have one alarm template assigned.

An *edge* is the IP link that connects a pair of subnet devices, for example, routers, layer 3 switches, or Chatter end points. CNA displays Chatter edges as colored lines.

You can also assign a default set of absolute alarm settings to all zones or edges to which you have not assigned a specific alarm template.

The no command assigns the default alarm template to the edge.

set-alarm zone

Assigns an alarm template to a Chatter zone. Each zone can have one alarm template assigned.

Syntax

set-alarm zone <name> alarm <alarm-name>

no set-alarm zone <name>

Arguments

Table 36: set-alarm zone arguments

Argument	Description
<name>	Zone to which you are assigning an alarm-template.
<alarm-name>	Alarm template that you are assigning to the zone.

Prerequisites

Must be in Chatter Configuration command mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Available for use multiple times in the configuration. Each zone can have one alarm template assigned.

Description

Assigns an alarm template to a Chatter zone. Each zone can have one alarm template assigned.

You can also assign a default set of absolute alarm settings to all zones or edges to which you have not assigned a specific alarm template.

The no command reverts the alarm setting back to the system default alarm template for the zone.

set-application-model

Chooses which application model for particular decision policies CNA should use when making decisions.

Syntax

```
set-application-model { enterprise |
                        other |
                        streaming |
                        voice |
                        multimedia |
                        web }
no set-application-model
```

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), both of which require privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **set-application-model** command to associate an application model with a named policy. If you use it in *config-engine* mode, then this will set the model globally, to be applied to decision policies where you don't specify the application model (in *config-engine-policy* mode).

For more information and examples about policies and application models, see the **decision-policy** command and Chapter 6: Decision Policies and Application Models of the *Monitoring Guide*.

See the commands: **application-model enterprise**, **application-model multimedia**, **application-model other**, **application-model streaming**, **application-model voice**, and **application-model web**.

---

set-decision-policy

Associates a subnet with a specified decision policy.

Syntax

```
set-decision-policy <name> { subnet <address> |
        group <group-name> |
        active-measurement-group <am-group-name> }
no set-decision-policy <name>
no set-decision-policy <name> group <group-name>
no set-decision-policy <name> active-measurement-group <am-group-name>
```

Arguments

- *name*—name of the policy
- *address*—address in CIDR format (a.b.c.d/n)
- *group-name*—name of the group of subnets to associate this policy with
- *am-group-name*—name of the active-measurement group to associate this policy with

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **`set-decision-policy`** command once for each subnet, group of subnets, or active-measurement group to be associated with a named policy.

For more information about policies, see the **`decision-policy`** command and Chapter 6: Decision Policies and Application Models of the *Monitoring Guide*.

---

set-link-group

Associates a subnet with a link-group.

Syntax

```
set-link-group <link-group-name> subnet <subnet>
no set-link-group <link-group-name> subnet <subnet>
```

Arguments

- *link-group-name*—name of the link group
- *subnet*—address in CIDR format (a.b.c.d/n)

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **`set-link-group`** command assign a specific subnet to a link group. This will cause the decision maker process to assert routes to such subnets only to links in the specified link-group.

Use the no form of the command to remove a subnet from a link-group

show active-measurements

Displays a table of statistics about currently configured active probes.

Syntax

```
show active-measurements [ traceprobe | detail | schedule ]
                      [ [ group <name> ] | no-group ]
                      [ target <probe> ]
                      [ address <address> ]
                      [ range-address <range> ]
                      [ type { tcp | icmp | traceprobe } ]
                      [ ustat <uname> ]
                      [ successful ]
                      [ errors <list> ]
                      [ probing <address> [ range-probing <range> ] ]
                      [ subnet <subnet> ]
                      [ subnet-probing <subnet> ]
                      [ errors <list> ]
                      [ | { include | exclude | begin } <exp> ]
                      [ > <url> ]
```

Arguments

- *name*—name of an active-measurement group
- *probe*—IP address or domain name of a probe, as entered in CNA configuration
- *address*—DNS-resolved IP address of a probe
- *range*—range in CIDR format (a.b.c.d /*n*)
- *uname*—USTAT module name
- *list*—error codes; can be any combination of the following:
  - *d*—DNS
  - *h*—Host unreachable
  - *n*—subnet unreachable
  - *p*—Port unreachable
  - *q*—Queue drop
  - *t*—Probe timeout
  - *o*—Other
- *subnet*—subnet address in CIDR format (*a.b.c.d/n*)
- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System configuration mode, login-level access.

Description

The CNA software maintains a table of statistics for all currently configured active probes. When you execute the **`show active-measurements`** command, the CNA software locates and displays all rows of the master table that match the specified arguments.

The **traceprobe** keyword changes the display format to show a new column in the output (labeled “probing”). The address in this column is that of the host that the traceprobe has found and is using instead of the actual target address; this happens because traceprobes run a traceroute towards the target and will use the farthest responding host in the traceroute path. When using this command, you don’t need to type the keywords **type traceprobe** - this will be automatic.

Use the **schedule** keyword to show the active probing schedule for each target.

Use the **detail** keyword to get detailed information on active probes.

Use the **range-address** keyword to limit output to addresses within a desired range.

Use the **range-probing** keyword to limit output to addresses within a desired traceprobing range.

Use the **no-group** keyword to display data for all active measurements not associated with a group.

To see data for all active measurements, do not specify either a group name or the **no-group** keyword.

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **`copy`** command for URL syntax.

Use the **`cidr-lookup`** command to clear selected rows in the master table.

show active-measurements schedule

Displays the active probing schedule for various targets.

Syntax

```
show active-measurements schedule
                         [ group <name> ]
                         [ target <probe> ]
                         [ type { tcp | icmp | traceprobe } ]
                         [ | { include | exclude | begin } <exp> ]
                         [ > <url> ]
```

Arguments

- *name*—name of an active-measurement group
- *probe*—IP address or domain name of a probe, as entered in CNA configuration
- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System configuration mode, login-level access.

Description

Use the **`show active-measurements schedule`** command to display the scheduling information about each active measurement target.

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **`copy`** command for URL syntax.

show arp

Displays contents of CNA ARP tables.

Syntax

```
show arp [ | { include | exclude | begin } <exp> ]
         [ > <url> ]
```

Arguments

- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System configuration mode, login-level access.

Description

The **show arp** command works in both local and system-wide modes. If you are connected to a peripheral module's console port, the data displayed will be local to that module only. The *slot* argument, if specified, will be ignored.

If you are connected to the management module, or you use telnet, you can specify a module by using the *slot* argument and the data displayed will be local to that module. If no *slot* argument is specified, the data will be obtained from the management module.

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show bootvar

Displays values of CNA environment variables.

Syntax

```
show bootvar [ | { include | exclude | begin } <exp> ] [ > <url> ]
```

Arguments

- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System configuration mode, login-level access.

Description

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show cell

Displays a tabular report of the test results between cells.

Syntax

show cell [test {rtp | ping | tcp} [variable {cerror | delay | drops | jitter | inter-arrival-time | loss | rtt}]] [from <zone-name> to <zone-name> [all-datapoints]]

Arguments

Table 37: show cell arguments

Option or argument	Description
test {rtp | ping | tcp}	Test for which you want to view test measurements. Options are: RTP, Ping, or TCP.
variable {cerror | delay | drops | jitter | inter-arrival-time | loss | rtt}	Specific test variable for which you want to view test measurements. Options are: • RTP: - Delay - Drops - Jitter - Loss - Inter arrival time • Ping - Round trip time - Loss - Inter arrival time - Drops • TCP - Delay - Cerror - Inter arrival time - Drops
from <*zone-name*>	Source zone of the cell for which you want to view test measurements. You can enter a fully qualified zone name or a specific zone name.
to <*zone-name*>	Destination zone of the cell for which you want to view test measurements. You can enter a fully qualified zone name or a specific zone name.
	1 of 2

Table 37: show cell arguments (continued)

Option or argument	Description
[all-datapoints]	**Important:** To use the **all-datapoints** option, you must specify both a test variable and cell. You cannot display all data points for all cells or all tests in one report Displays test results for all available CNA agents in the cell. This option displays the data used to generate the historical data plot. See all-datapoints option.
	2 of 2

Prerequisites

Must be in Privileged command mode.

Description

Displays a tabular report of the test results between cells. If you do not specify a cell, test results are displayed for all cells in the agent-based measurement subnet. For a definition and examples of cells, see Agent-based measurement tests in Chapter 5: Configuring Agent-Based Measurement of the *Monitoring Guide*.

To display the results for a specific test variable, specify the test variable. By default the command displays test results for all tests.

Report display

The report consists of one row per test per cell. For each test, the following criteria are displayed:

- Absolute alarm state
- Relative alarm state
- Minimum measurement
- Median measurement
- Maximum measurement
- Average measurement

Sample reports

Table 38 shows a sample report for the **show cell** command. This command displays data for all currently active tests for all cells.

Table 38: sample output for show cell command

Parent Zone	From Zone	To Zone	Test Name	Test Variable	Absolute	Relative	Min	Median	Max	Avg
US	APAC	EMEA	RTP	Delay (ms)	Above	Rising	47	140	342	120
US			RTP	Jitter (ms)	Below					
			RTP	Loss (%)	Above	Rising				
			RTP	Drops (%)	Below	Steady				
			RTP	Inter Arrival Time (secs)	Steady					
			Ping	RTT (ms)	Steady					
			Ping	Drops (%)	Below					
			Ping	Loss (%)	Above					
			Ping	Inter Arrival Time (secs)	Steady					
US	EMEA	Europe	. . .							
	. . .	. . .								
California	Milpitas	Concord	. . .							

Table 39 shows a sample report for the **show cell test ping-rtt** command. This command displays Ping RTT test measurements for all cells.

Table 39: sample output for show cell test ping-rtt command

Parent Zone	From Zone	To Zone	Test Name	Test Variable	Absolute	Relative	Min	Median	Max	Avg
US	APAC	EMEA	Ping	RTT	Above	Rising	47	140	342	120
					Below					
					Above	Rising				
					Below	Steady				
					Steady					
					Steady					
					Below					
					Above					
					Steady					
US	EMEA	Europe	Ping	RTT						
	. . .	. . .								
California	Milpitas	Concord	Ping	RTT						

Table 40 shows a sample report for the **show cell from APAC to EMEA** command. This command displays data for all currently active tests between the APAC and EMEA zones.

Table 40: sample output for show cell from APAC to EMEA command

Parent Zone	From Zone	To Zone	Test Name	Test Variable	Absolute	Relative	Min	Median	Max	Avg
US	APAC	EMEA	RTP	Delay (ms)	Above	Rising	47	140	342	120
US	APAC	EMEA	RTP	Jitter (ms)	Below					
US	APAC	EMEA	RTP	Loss (%)	Above	Rising				
US	APAC	EMEA	RTP	Drops (%)	Below	Steady				
US	APAC	EMEA	RTP	Inter Arrival Time (secs)	Steady					
US	APAC	EMEA	Ping	RTT (ms)	Steady					
US	APAC	EMEA	Ping	Drops (%)	Below					
US	APAC	EMEA	Ping	Loss (%)	Above					
US	APAC	EMEA	Ping	Inter Arrival Time (secs)	Steady					

Table 41 shows a sample report for the **show cell test rtp-jitter from APAC to EMEA** command. This command displays RTP jitter test measurements between the APAC and EMEA zones.

Table 41: sample output for show cell test rtp-jitter from APAC to EMEA command

Parent Zone	From Zone	To Zone	Test Name	Test Variable	Absolute	Relative	Min	Median	Max	Avg
US	APAC	EMEA	RTP	Jitter	Above	Rising	47	140	342	120

Table 42 shows a sample report for the **show cell test rtp-jitter** command. This command displays RTP jitter test measurements for all cells.

Table 42: sample output for the show cell test rtp-jitter command

Parent Zone	From Zone	To Zone	Test Name	Test Variable	Absolute	Relative	Min	Median	Max	Avg
US	APAC	EMEA	RTP	Jitter	Above	Rising	47	140	342	120
US	EMEA	Europe	RTP	Jitter	Below					
	. . .	. . .	RTP	Jitter						
California	Milpitas	Concord	RTP	Jitter	Stale	Stale	23	30	44	35

all-datapoints option

Use the **all-datapoints** option to display test results for all available CNA agents in the cell.This option displays the data used to generate the historical data plot in the Web interface. When you use this option, the resulting report contains the following columns:

- Time Since
- Value
- From CNA Agent
- To CNA Agent
- Absolute
- Absolute Strikes
- Relative
- Relative Strikes

For example, the following command displays the data used to generate the scatter plot for a cell: **show cell test rtp-delay from Milpitas to Concord all-datapoints**

show chatter cna-agent-device

Displays information about the stand-alone CNA Agent Devices in your agent-based measurement subnet.

Syntax

show chatter cna-agent-device {fingerprint | log | information}

Arguments

Table 43: show chatter cna-agent-device

Argument	Description
fingerprint	Displays the SSH fingerprints for CNA Agent Devices. The CNA server generates the fingerprints and uses them to establish SSH connections with CNA Agent Devices. The CNA server stores up to five SSH fingerprints and uses the most current one to establish SSH sessions with CNA Agent Devices. The CNA Agent Device uses the fingerprint to authenticate the CNA server. Once the server is authenticated, the CNA Agent Device: • Allows the CNA Server to remotely manage it. • Retrieves CNA agent software upgrades from the CNA server when it is upgraded. To generate a new SSH RSA public/private key pair and fingerprint, use the ip ssh key generate cna-agent command.
log	Displays the contents of the syslog file from CNA Agent Devices.
information	Displays the device's status of registration with a CNA server, date and time of registration, and software version number.

Prerequisites

Must be in Privileged command mode. For information on how to access this mode, see Accessing and exiting command modes.

Description

Displays information about the stand-alone CNA Agent Devices in your agent-based measurement subnet.

Sample report

```
homeless-a15# sho cna-agent fingerprint

Date       Time     Fingerprint Order   Fingerprint Data
09/05/2005 08:58:42 current             fingerprint5
09/05/2005 08:58:43 fingerprint5        41:50:96:cf:e8:92:15:9a:ad:b9:fb:87:98:a7:70:88
09/05/2005 07:57:27 fingerprint4        24:90:f8:2e:52:cc:f6:fd:64:b6:dd:68:5d:98:c6:45
09/05/2005 08:55:31 fingerprint3        b5:21:c1:be:0a:bb:4a:e4:d3:5b:e5:7b:c2:2b:e8:58
09/05/2005 08:48:01 fingerprint2        98:bf:9e:94:b2:0b:60:02:09:67:11:d9:7c:2d:fc:5e
09/05/2005 08:35:42 fingerprint1        24:f3:6d:c3:1c:d9:19:a2:c0:6a:be:c8:a6:31:55:6b

Current Fingerprint: 41:50:96:cf:e8:92:15:9a:ad:b9:fb:87:98:a7:70:88
```

show chatter master-key

Displays the master key for the local Chatter module (on which you enter this command).

Syntax

show chatter master-key

Prerequisites

Must be in Privileged command mode.

Description

Displays the master key for the local Chatter module (on which you enter this command).

You need this key to add a remote Chatter module to the local Chatter module's hive. To add a remote Chatter module to the local Chatter module's hive, enter the chatter join-hive command on the remote Chatter module.

show chatter ssl

Displays the SSL certificate that CNA test agents use to authenticate and register with the CNA server.

Syntax

show chatter ssl

Prerequisites

Must be in Privileged command mode.

Description

Displays the SSL certificate that CNA test agents use to authenticate and register with the CNA server. The certificate is displayed in Privacy Enhanced Mail (PEM) format.

---

show cidr-lookup

Displays the owners of specified subnets.

Syntax

```
show cidr-lookup <list>
                  [ | { include | exclude | begin } <exp> ]
                  [ > <url> ]
```

Arguments

- *list*—list of subnets, in CIDR format (a.b.c.d/n)
- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Description

The **`show cidr-lookup`** command identifies owners of specified subnets.

You can see multiple subnet listings by entering multiple subnets using CIDR syntax (*a.b.c.d/n*) on the command line. Separate each by a space character and enclose the list in double quotation marks:

```
show cidr-lookup “172.10.3.0/24 172.16.0.0/16”
```

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show cna-agent

Displays information about the CNA test agents in your agent-based measurement subnet.

Syntax

show cna-agent [assigned [zone <zone-name>] | unassigned | detail | measurement [ip-address <addr>] [duration <s>]]

Arguments

Table 44: show cna-agent

Option or argument	Description
[assigned [zone *<zone-name>*]	Displays all CNA agents that are assigned to a zone. To see the CNA agents that are assigned to one specific zone, use the **zone *<zone-name>*** option.
unassigned	Displays all CNA agents that are not assigned to a zone. Use this option to see new CNA agents that you need to assign to a zone.
	1 of 2

Table 44: show cna-agent (continued)

Option or argument	Description
detail	Displays the following information for each CNA agent: • IP address • CNA server with which it registered • Date and time of registration • Version of CNA agent software • Software version • Type of device (whether a stand-alone CNA Agent device, or an embedded agent) • MAC address • subnet • Call server (for IP phones) • Port on the CNA agent to which the CNA Server sends test requests. • Port out which the CNA agent sends RTP test requests to a destination CNA agent. • Hardware model • Firmware version
measurement	Displays the following information for each CNA agent: • IP address • CNA server with which it registered • Duration in seconds that the report covers • Number of test requests sent to a destination CNA agent. • Number of test results (ok or busy) received from a destination CNA agent • Number of CNA server test requests to which the CNA agent sends a busy response. When busy, the CNA agent does not send test requests to destination CNA agents. • Number of test requests received from a source CNA agent. • Number of test results (ok or busy) sent to a source CNA agent • Time last test request was sent to destination CNA agent. • Time last test result was sent to source CNA agent.
[ip-address <addr>]	IP address of a CNA agent for which you want to view measurement data.
[duration <s>]	Time period in seconds for which you want to see measurement data. The default duration is 0: all available measurement data is displayed.

2 of 2

Prerequisites

Must be in Privileged command mode.

Description

Displays information about the CNA test agents in your agent-based measurement subnet.

The report lists one row per CNA agent. If an agent is assigned to a particular zone, it is indicated in the report.

Sample report

```
IP Address        Name                Assigned To        Device Type
148.147.3.112     Milpitas-TP         California         CNA Server => IBM
148.147.12.46     Basking-Ridge-TP    New Jersey         CNA Agnet => Avaya
12.14.16.10       Pune-TP             India              IP Phone
135.10.13.66      -                   -                  IP Phone
```

---

show command-log

Displays a list of commands that have been executed or configured on CNA during the specified period.

Syntax

```
show command-log [ historic ]
                 [ all | exec | config ]
                 [ start-time <begin> ]
                 [ end-time <end> ]
                 [ detail ]
                 [ reverse ]
                 [ | { include | exclude | begin } <exp> ]
                 [ > <url> ]
```

Arguments

- *begin*—first time stamp of the display range, in *mm-dd-yy_HH-MM* format
- *end*—last time stamp of the display range, in *mm-dd-yy_HH-MM* format

Prerequisites

System operations mode, privileged-level access.

Description

The **show command-log** command displays a list of all commands of the specified type that have been invoked on CNA during the specified time period.

When you execute the command with no type keyword (**all**, **exec** or **config**), the display will include all **exec** and **config** commands that the **show command-log** process can identify by command name.

When you invoke the command with the keyword **all**, the display will include a listing for every command executed during the specified period, including those that cannot be identified or categorized by the **show command-log** process.

The keyword **historic** will show commands prior to the last system boot.

Commands invoked by internal processes, as well as many commands invoked by the web client, are passed to the *coreman* process as XML tags, which do not use the same command names as the CLI. Those commands cannot be mapped to their CLI names by the **show command-log** process, and they will be listed as N/A (not available) in the output for both command name and type. Even though you will not be able to identify the command from the listing, you will be able to locate commands that take long periods of time to complete or consume large amounts of memory. By matching the start time with general CNA activity as recorded in log messages, you may be able to deduce the command name.

If you invoke the **show command-log** command with the **exec** keyword, the output will be limited to just those commands that the CNA software executes immediately upon entry without adding them to the configuration (such as **show** commands, for example).

The **config** keyword will limit the display to just those commands that the CNA software adds to the configuration. Config commands are also executed immediately.

The timestamps entered with the **start-time** and **end-time** keywords define the time period to be covered by the **show command-log** output. You must enter complete time stamps. Unlike the **show report** command, the **show command-log** command will not attempt to match a partial time stamp.

If you don't specify a time range, the list will include all commands that have been issued since the most recent reboot, except that there are limits on the total number of commands that will be kept in the log. These limits, which are displayed at the end of the command output, are: 10,000 **config** commands entered through the CLI, 2,000 **exec** commands entered through the CLI, and 1,000 commands invoked by a process other than the CLI.

Output will include the following information about each command that was invoked during the specified time:

- **Number**—a sequential number showing the order in which commands were invoked since the last reboot; there may be gaps in the sequence as commands exceed the command log's buffer capacity for that type of command and are discarded from the log
- **Start Time**—time stamp showing when the command was invoked
- **Mode**—whether it was an **exec** or **config** command
- **Command**—the command name

The **detail** keyword causes the display to include more details about each command entry, including the following additional columns:

- **Duration**—length of time it took to complete or, in the case of commands not yet completed (marked by an asterisk), how long the command has been running; time spans are identified as follows:

  d—days

  h—hours

  M—minutes

  s—seconds

  m—milliseconds

- **Start Memory**—the total amount of memory in use by the CNA *coreman* process at the time the command was invoked
- **Mem Diff**—the difference between the amount of memory in use by the CNA when the command was invoked and when the command completed; may be negative if memory was released while the command was in progress

If you have configured an **enable** password, the **`enable`** command will appear once in the log for each time the command was executed. The CNA software allows you to enter privileged mode from a console connection without requiring that an **enable** password be configured. So it must first check to see if a password is required. If it is, the **`enable`** command will be rejected the first time and the CLI will prompt the user for a password. Then the **`enable`** command will be passed back to the *CNA* process a second time.

The **reverse** keyword causes the commands to be listed in reverse order (starting with the earliest command in the log).

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **`copy`** command for URL syntax.

show configuration

Note:

This command behaves exactly the same as the command **show startup-config**. It is provided for historical reasons only.

Displays the contents of the startup configuration.

Syntax

```
show startup-config [ | { include | exclude | begin | begin } <exp> ]
                    [ > <url> ]
```

Arguments

- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, privileged-level access. Can also be executed in configuration (*config*) mode.

Description

Use the **show startup-config** command to display the contents of the startup configuration.

You can pipe the command through one of the following filters:

- **include**—shows all commands present in the startup configuration which contain a match to the specified regular expression
- **exclude**—shows all commands present in the startup configuration except those which contain a match to the specified regular expression
- **begin**—shows all commands present in the startup configuration beginning with the first match to the specified regular expression

The vertical-line pipe character is required when a filter is used:

```
PathControl8# show startup-config | include interface
interface Ethernet 0/0
interface Ethernet 3/0
interface Ethernet 4/0
interface Ethernet 9/0
  polling interface-ip-address 172.30.100.54
  polling interface-ip-address 172.29.100.54
```

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show counts prefix (deprecated)

The show counts prefix command has been deprecated. See the replacement command, show counts subnet on page 330.

show counts subnet

Displays information on a scheduled reload.

Syntax

```
show counts subnet category | origin | policy | outage
```

Prerequisites

System operations mode.

Description

The **category** keyword specifies a summary report for the first 500 subnets. (The command **show prefix (deprecated)** allows you to get more detailed information on subnets)

The **origin** keyword shows how many subnets were user-defined, bgp-defined, system-defined, or user-defined AND bgp-defined.

The **policy** keyword shows how many subnets have a given policy and how such were assigned (global/bgp, user-specified, inherited, measurement-based).

The **outage** keyword shows, on a per-link basis, a summary table with counts for outages, silence outages, and rescued networks.

show debugging

Displays a list of all subsystems for which the **debug on** command can be issued, along with the status for the current CLI session.

Syntax

```
show debugging
```

Arguments

- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Description

The **show debugging** command displays a list of subsystems and labels each as either *off* or *ON*, depending on whether you have enabled display of a subsystem's debug messages.

See the **debug** and **debug on** commands.

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

---

show diagbus

Displays diagnostic bus and boot log information for the specified CNA module.

Syntax

```
show diagbus [ detail ]
             [ <num> | all ]
             [ | { include | exclude | begin } <exp> ]
             [ > <url> ]
```

Arguments

- *num*—slot or subsystem number of CNA module, or *all*
- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Description

The **show diagbus** command generates diagnostic information about the state of CNA.

An Avaya technician may ask you to send this information to an FTP or HTTP server; use the **> <url>** option (a right angle bracket followed by an URL) to divert the **show** command output to a remote file location.

Use the **detail** keyword to obtain extensive diagnostic information; omit the **detail** keyword to obtain less extensive information.

Specify a slot or subsystem number to see data about the specified system, or the keyword **all** to see data about all systems. If no system is specified, the output will refer to the management module.

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

---

show efc

Displays data gathered by the EFC module and maintained by the CNA *efcserver* process.

Syntax

```
show efc {
           active-measurement [ summary ] [ group <g-name> | all ] |
           application |
           application-endpoint |
         }
```

Note:

Starting with CNA version 5: the keyword **link** has been deprecated and will no longer have any effect; the keyword **subnet** has been deprecated and will be automatically translated to **show subnet traffic-rate**.

Arguments

- *g-name*—name of an active measurement group in the CNA configuration
- *address*—subnet address in CIDR format (*a.b.c.d/n*)

Prerequisites

System operations mode, log-in level access.

Description

Use the **show efc** command to view data obtained by the EFC module.

Use the **active-measurement** keyword to see data maintained in the active-measurement table, either per specified address or for all addresses. With the **summary** keyword, you can produce summary numbers per group without the list of EFC end points.

Use the **application** keyword to see data on each application (e.g., oracle, sap, telnet) gathered by the EFC module.

Use the **application-endpoint** keyword to see data for each endpoint with applications recognized by the EFC module.

Use the **subnet** keyword to see data maintained in the subnet table, either per specified subnet address or for all subnets.

Use the **link** keyword to see traffic distribution for endpoints gathered by the EFC module.

Data in the active-measurement and subnet tables can be reset using the **clear efc active-measurement** and **clear efc prefix (deprecated)** commands.

See Chapter 2: Collecting Endpoint Addresses of the *Monitoring Guide*.

---

show event

Displays the contents of the event log.

Syntax

```
show event
           [ <start> ]
           [ end <end> ]
           [ detail ]
           [ count <number> ]
           [ subnet <subnet> | group <group-name> ]
           [ link <link-name> ]
           [ selected ]
           [ asserted ]
           [ confirmed ]
           [ bgp-announcement ]
           [ bgp-withdrawal ]
```

```
[ bgp-summary ]
[ link-status ]
[ unlogged-action ]
[ state-transition ]
[ peer-status ]
[ em-start ]
[ outage ]
[ not-rescued-outage ]
[ | { include | exclude } <exp> ]
[ > <url> ]
```

Arguments

- *start*—show events from this time, in *yyyy-mm-dd_HH-MM-SS* format
- *end*—show events until this time, in *yyyy-mm-dd_HH-MM-SS* format
- *number*—show this number of events, then exit the command
- *subnet*—subnet in CIDR format (a.b.c.d /*n*)
- *group-name*—name of a group
- *link-name*—name of a link
- *exp*—regular expression
- *url*—destination of output; can be FTP, HTTP, or HTTPS URL

Prerequisites

System operations mode, login level access.

Description

Use the **`show event`** command to view a list of event log messages, beginning at *start* time and ending at *end* time (if specified).

Specify all or part of a timestamp for *start* and *end*. Events shown will be those that come closest to matching the specified timestamps. For example, if you specify only a year for *start*, all events for that year will be displayed. (You must specify everything to the left of the desired time element. If you specify month you must include a year; if you specify day, you must include month and year, and so on.)

Note:

The keyword **raw** is no longer supported.

The following additional keywords are supported:

- **detail**—displays all values for each event
- **subnet**—all events for a particular subnet
- **group**—all events for the subnets in a particular group
- **link**—all events for a particular link

- **selected**—route-selected events
- **asserted**—route-asserted events
- **confirmed**—route-confirmed events
- **bgp-announcement**—BGP announcement events
- **bgp-withdrawal**—BGP withdrawal events
- **bgp-summary**—BGP summary events
- **link-status**—link-status events
- **unlogged-action**—events generated when there are lost events due to an overflow of the event queue
- **state-transition**—decision maker state transition events
- **peer-status**—CNA BGP peering events
- **em-start**—event manager start events
- **outage**—outage events
- **not-rescued-outage**—not rescued outage events

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

The event log is limited to 500 MB of disk space. When this space is full, incoming messages overwrite the oldest existing messages.

See also the **eventdump** command.

show edge

Displays a tabular report of the test results for a specific edge or all edges in your agent-based measurement subnet.

Syntax

show edge [test {rtp | ping | tcp | traceroute} [variable {delay | jitter | loss | inter-arrival-time | rtt | cerror | hop-rtt}]] [from <ip-address> to <ip-address> [all-datapoints]]

Arguments

Table 45: show edge arguments

Argument	Description
test {rtp | ping | tcp | traceroute}	Test for which you want to view test measurements. Options are: RTP, Ping, TCP, or Traceroute.
[variable {delay | jitter | loss | inter-arrival-time | rtt | cerror | hop-rtt}]	Specific test variable for which you want to view test measurements. Options are: • RTP: - Delay - Jitter - Loss - Inter arrival time • Ping - Round trip time - Loss - Inter arrival time • TCP - Delay - Cerror - Inter arrival time • Traceroute - Hop round trip time - Inter arrival time
from <ip-address>	Source IP address of the edge for which you want to view test measurements.
to <ip-address>	Destination IP address of the edge for which you want to view test measurements.
[all-datapoints]	**Important:** To use the **all-datapoints** option, you must specify both a test variable and an edge. You cannot display all data points for all edges or all tests in one report Displays test results for all available CNA agents on the edge. This command displays the data used to generate the historical data plot. For more information, see all-datapoints option.

Prerequisites

Must be in Privileged command mode.For information on how to access this mode, see Accessing and exiting command modes.

Description

Displays a tabular report of the test results for an edge. If you do not specify an edge, test results are displayed for all edges in your agent-based measurement subnet.

To display the results for a specific test variable, specify the test variable. By default the command displays test results for all tests.

Report display

The report consists of one row per test per edge. For each test, the following criteria are displayed:

- Absolute alarm state
- Relative alarm state
- Minimum measurement
- Median measurement
- Maximum measurement

Sample report

From IP	From Name	To IP	To Name	Test Name	Test Variable	Absolute	Relative	Min	Median	Max	Avg
12.4.5.43	ATT	147.16.23.12	Yahoo	RTP	Jitter	Above	Rising	47	140	342	120
135.148.13.56	Avaya		…	RTP	Jitter	Below					
	. . .		. . .	RTP	Jitter						
	Milpitas		Concord	RTP	Jitter	Stale	Stale	23	30	44	35

all-datapoints option

Use the **all-datapoints** option to display test results for all available CNA agents on the edge. This option displays the data used to generate the historical data plot in the Web interface. When you use this option, the resulting report contains the following columns:

- Time Since
- Value
- From CNA Agent
- To CNA Agent
- From Zone
- To Zone

- Absolute
- Absolute Strikes
- Relative
- Relative Strikes

For example, the following command displays the data used to generate the scatter plot for a cell:

show edge test rtp variable delay from 10.10.10.1 to 10.10.11.1 all-datapoints

show image

Displays information about an archived copy of the CNA image.

Syntax

```
show image <file>
           [ | { include | exclude | begin } <exp> ]
           [ > <url> ]
```

Arguments

- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, requires privileged-level access.

Description

The **`show image`** command will display data contained in an information block in an archived copy of the CNA image that was created by the **`backup`** command.

If you added notes during the backup, your text will be displayed in the **show image** output.

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show interfaces

Displays information about interfaces.

Syntax

```
show interfaces [ <type> ]
                [ detail ]
                [ | { include | exclude | begin } <exp> ]
                [ > <url> ]
```

Arguments

- *type*—type of interface; can be:
  - *Ethernet*
  - *tunnel*
  - *loopback*
- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Description

The **show interfaces** command displays information about interfaces configured on CNA modules.

If the **detail** keyword is not specified, each interface included in the output will be listed on a single line along with its IP address and its current state.

If the **detail** keyword is specified, more information about each interface will be displayed.

If *slot* number is not specified, the management module is assumed.

If *type* is specified, only interfaces of that type will be included in the output; otherwise, information about all interfaces configured on the specified module will be included.

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression

- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show ip bgp neighbor

Displays information about all bgp neighbors and their CNA-configured links.

Syntax

```
show ip bgp neighbor [ | { include | exclude | begin } <exp> ]
                     [ > <url> ]
```

Arguments

- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Occurrence

Use once.

Description

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show ip route

Shows the IP routing table for the specified module.

Syntax

```
show ip route [<num>] [ | { include | exclude | begin } <exp> ]
              [ > <url> ]
```

Arguments

- *num*—slot or subsystem number of an CNA module
- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Description

The **`show ip route`** command works in both local and system-wide modes. If you are connected to a peripheral module's console port, the data displayed will be local to that module only. The *slot* argument, if specified, will be ignored.

If you are connected to the management module, or you use telnet, you can specify a module by using the *slot* argument and the data displayed will be local to that module. If no *slot* argument is specified, the data will be obtained from the management module.

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **`copy`** command for URL syntax.

show ip sockets

Displays all open IP sockets by remote host IP address.

Syntax

```
show ip sockets [<num>]
                [ | { include | exclude | begin } <exp> ]
                [ > <url> ]
```

Arguments

- *num*—CNA slot or subsystem number
- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

Must be in system-operations mode, with privileged-level access.

Description

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show link

Displays information about service provider links.

Syntax

```
show link [<name>]
          [ | { include | exclude | begin } <exp> ]
          [ > <url> ]
```

Arguments

- *name*—name of the link
- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Description

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

---

show link billing

Displays an estimated bill to date, based on usage from the beginning of the current billing period.

Syntax

```
show link billing [<name>]
                  [ | { include | exclude | begin } <exp> ]
                  [ > <url> ]
```

Arguments

- *name*—name of a link
- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Description

The **`show link billing`** command uses the service provider billing configured for a link to display an estimate of what the current bill would be based on usage from the beginning of the billing period to the current date.

See the **`usage-reporting billing-period`**, **`usage-reporting cost`**, **`usage-reporting formula`** and **`usage-reporting percentile`** commands for more information about billing.

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **`copy`** command for URL syntax.

---

show local-debug

Displays data about the local CNA module, for troubleshooting. It cannot be used with the management module (see show tech-support).

Syntax

```
show local-debug [ | { include | exclude | begin } <exp> ] [ > <url> ]
```

Arguments

- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

Must be physically connected to the module of interest. System operations mode, login-level access.

Description

The **show local-debug** command displays information about the local module. Output is similar to the output from the **show tech-support** command, except it is obtained from the local module only.

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

---

show logging

Displays contents of logging buffer.

Syntax

```
show logging
   [ [ { first | last } <count> ] | all ]
   [ level <name> ]
   [ | { include | exclude | begin } <exp> ]
   [ > <url> ]
```

Arguments

- *count*—number of lines to display
- *name*—logging level; output will include messages whose level is at or higher in the following hierarchy list:
  - *emergency*
  - *alert*
  - *critical*

  - *error*
  - *warning*
  - *notice*
  - *informational*
  - *debug*
- *exp*—regular expression
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Occurrence

Use once.

Default

last100

debug and above (all levels)

Description

The **`show logging`** command displays all messages up to a specified number, or only those messages that match a regular expression. If *exp* contains spaces, it should be enclosed in quotes (“ ”).

Use the *count* argument with the keyword **first** to specify the number of lines in the log to display, beginning with the oldest message in the buffer.

Use *count* with the **last** keyword to specify a number of messages beginning with the most recent message in the buffer.

Use the keyword **all** to specify all lines at the specified level.

When you specify **level**, include an equals sign as the first character to limit the display to only the named level.Otherwise, the output will contain all messages whose level is at or higher than the level specified.

For example, if you specify **level** as *alert*, the output will contain messages whose level is either *alert* or *emergency*. If you specify **level** as =*alert*, the output will contain only those messages whose level is *alert*.

See the **`logging buffered`** command.

See the **`debug on`** command for subsystem names, which can be used for match expressions.

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show medic

Displays information about CNA.

Syntax

```
show medic [ detail ]
           [ <service> [ <routine> | <parameter> ] ]
           [ | { include | exclude | begin } <exp> ]
           [ > <url> ]
```

Arguments

- *service*—name of Medic service
- *routine*—name of Medic routine
- *parameter*—name of Medic parameter
- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Description

Omit the **detail** keyword to see a summary of conditions.

Use the keyword **detail** to see a more comprehensive report.

If you are connected to a peripheral CNA module (anything other than the management module), the **show medic** command operates in local mode and the data displayed refers only to the local module.

The **Status** column can contain one of the following:

- **up**—module is operating correctly
- **no connectivity**—internal subnet connectivity between management module and peripheral is down (no *ping*)
- **verifying flash**—connectivity established with peripheral but peripheral has not yet verified that the flash image is correct
- **(no entry)**—module was not identified by the system (could be an electrical problem or bad ROM)

The **Uptime** column tells how long the module has been operating (**d** for days, **h** for hours, **m** for minutes, **s** for seconds; the numbers following a **d** entry refer to additional hours and minutes—1d, 6:05 means the module has been operating for one day, six hours and five minutes).

The **Services** column identifies how many services are currently running, out of the total number of services on that module.

The **CPU** and **Mem** columns show percentage of use (an entry of 17 means 17% of total capacity is currently being used).

If you enter the **show medic** command with the keyword **all** and the keyword **detail**, the status of each module will be expanded into a block of detailed text. Output will look similar to the following block of text for a USTAT module:

```
CNA> show medic detail all
 1   ustat       up 19:28            srv: 13/13  cpu: 2%   mem: 34%
       PASS  active-measurement-process         eval: 21:58:29 (5s ago)
               [ Running ]
               routine: restart INIT (never run)
       PASS  environment-monitor                eval: 21:58:20 (14s ago)
       PASS  host-resources                     eval: 21:58:20 (14s ago)
               max-pcpu = 80 (def = 80) SNMP trap max % CPU utilized
               max-pmem = 80 (def = 80) SNMP trap max % memory utilized
               max-prdisk = 99 (def = 99) SNMP trap max % root disk utilized
               min-pcpu = 50 (def = 50) SNMP trap min % CPU utilized
               min-pmem = 50 (def = 50) SNMP trap min % memory utilized
               min-prdisk = 80 (def = 80) SNMP trap min % root disk utilized
       PASS  interface-resources                eval: 21:58:31 (3s ago)
       PASS  klogd                              eval: 21:58:29 (5s ago)
       PASS  logfilterd                         eval: 21:58:29 (5s ago)
       PASS  measurement-process                eval: 21:58:30 (4s ago)
               [ Running ]
               routine: restart INIT (never run)
       PASS  mount                              eval: Thu 02:32 (19h26m ago)
       PASS  mount-hd                           eval: Thu 02:32 (19h26m ago)
       PASS  ping-management                    eval: 21:58:31 (3s ago)
       PASS  snmpd                              eval: 21:58:33 (1s ago)
       PASS  ustat-web-server                   eval: 21:58:31 (3s ago)
```

```
        [ Running ]
        routine: restart INIT (never run)
PASS  verify-flash                    eval: Thu 02:32 (19h26m ago)
```

The first line in each block identifies the slot number and module name. Two asterisks following the slot number (*N***) indicate a problem in one of the routines.

The next column in the first line tells you the status of the module.

In the example, the USTAT module in slot 1 has been operating (status is *up*) for 19 hours and 28 minutes. Operating time is shown in UNIX/Linux **uptime** format. If the system has been operating for more than 24 hours, the time will be presented in number of days plus hours and minutes (for example, *up 1 day, 19:28* would mean the module had been operating for one day, 19 hours and 28 minutes).

The next column tells you how many Medic services are running out of how many total. In the example shown, *srv: 13/13* means that 13 of the 13 services on the module were successfully initialized and are currently available.

The next two columns tell you what percentage of total CPU capacity and RAM are currently being used.

Services, and their routines and parameters, are then itemized.

One of the following status labels precedes the service name:

- PASS—Medic was able to complete the service's startup routines with no critical errors
- FAIL—indicates a serious error was encountered when Medic attempted to run the service's startup routines; CNA must be rebooted in order to start the service
- EVAL—the service is in the process of starting up or attempting to recover from a failure; may be waiting for a dependent service to start (reissue the **`show medic`** command to see if the service has completed its startup routines)
- INIT—Medic has not yet tried to start the service (this state should never be visible unless the CNA startup sequence has encountered a serious error)
- WARN—the service is not currently running, possibly due to user action or configuration error; Medic may be in the process of restarting a service that halted unexpectedly (if Medic is unable to restart a service after several attempts, the status will be changed to FAIL)

The column following the service's name gives the date and time the service was last evaluated by the Medic process, with an interval in parentheses telling how long ago that was.

In the following example, the service was last evaluated by Medic on Thursday at 4:01 p.m., which was 3 days, 18 hours ago from the present:

```
eval: Thu 16:01 (3d18h ago)
```

If the service is currently running (not all will be running at all times), the next line will show the following:

```
[ Running ]
```

If the service has a user-accessible routine, it will be identified, along with its status:

```
routine: restart INIT (never run)
```

This line shows that the service's **restart** routine is in the INIT state, which means it has never been run. Possible routine states are:

- IDLE—routine is inactive (no tasks currently scheduled)
- INIT—routine has never been run
- RETRY—routine is waiting a specified time to repeat a failed task
- SLEEP—routine is waiting to perform an action at a specified time
- SUSPEND—routine is either waiting for a required service/routine to finish running, or the routine is paused while a system resource clears (if the database has been shut down or restarted, for example, it may take a few seconds before a related routine can proceed)

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show nodes

Displays the nodes known to the Chatter subnet and information about each node.

Syntax

show nodes [topology]

Arguments

Table 46: show nodes arguments

Argument	Description
[topology]	Displays the x and y coordinates. Use this option when you enter the **show nodes** command from the Web interface.

Prerequisites

Must be in Privileged command mode.

Description

Displays the nodes known to the Chatter subnet and information about each node.

The report displays one row per node. Additional IP addresses that are associated with the node are also listed one per row. These IP addresses are associated to the node using the Merge test.

Sample report

```
IP Address          Name                Device Type         Additional
                                                            Interfaces
148.147.3.1         Milpitas-Switch     Switch
                                                            148.147.12.43
198.152.2.30        Basking-Ridge-Route  Router
                    r
144.3.12.150        Pune-Access         Router
                                                            144.4.12.63
```

show ntp

Displays information about all of the NTP servers currently peered with CNA.

Syntax

```
show ntp [ association | status ]
         [ | { include | exclude | begin } <exp> ]
         [ > <url> ]
```

Arguments

- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, login-level privileges.

Description

Output from the **show ntp** command is generated by a call to the **ntpq** query utility. More information about this utility and its output is available at *http://www.ntp.org*.

Use the **status** keyword to display the current timing status with the synchronized peer. Output, which is generated by calling the **ntpq** utility with a **-p** option, will look similar to the following:

```
   remote     refid     st t when poll reach   delay   offset    disp
==============================================================================
*LOCAL(1)  LOCAL(1)     10 l   49   64  377     0.00    0.000   10.01
```

Columns identify the following:

- **remote**—NTP server and peers defined in the CNA configuration; in the example, there are no NTP peerings configured, so time synchronization is provided by the management module, shown as LOCAL
- **refid**—current source of synchronization, again LOCAL for the management module.
- **st**—NTP stratum.
- **t**—NTP type; can be one of the following:
  - *u*—unicast
  - *m*—multicast
  - *l*—local
  - - —don't know
- **poll**—polling interval, in seconds

- **when**—number of seconds since peer was last heard
- **reach**—octal value representing status of *reachability register* (see RFC-1305)
- **delay**—latest delay, in milliseconds (ms)
- **offset**—latest offset in ms
- **disp**—latest jitter (dispersion) in ms

The **association** keyword displays the output of a call to the **ntpq** utility with the **-c association** option. This output will look similar to the following:

```
ind assID status  conf reach auth condition  last_event cnt
===========================================================
  1 33732  9614   yes   yes  none  sys.peer   reachable  1
```

Each line in the **association** table corresponds to a line in the **status** table, where each line represents an NTP server or peer.

Columns identify the following:

- **ind**—index number
- **assID**—unique identifier for each mobilized association
- **status**—hex number indicating peer status, defined in the NTP specification
- **conf, reach, auth, condition, last_event,** and **cnt**—shows data decoded from the status word, as defined in RFC-1305 (*sys.peer* in the **condition** column identifies the *system peer*, the server which is currently providing synchronization; this should correspond to the line identified by an asterisk, *, in the status table)

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show outage

Displays the current outages.

Syntax

```
show outage [ subnet <subnet> ]
            [ link <link> ]
            [ trended-live ]
            [ limit <num> ]
            [ | { include | exclude | begin } <exp> ]
            [ > <url> ]
```

Arguments

- *subnet*—subnet in CIDR format (a.b.c.d /*n*)
- *link-name*—name of a link
- *num*—maximum number of networks to be displayed in the report
- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, any level access.

Description

The **`show outage`** command displays outages within the last 24 hours. To see outages beyond 24 hours, refer to the summary outage reports (see show report **outage-subnet-summary** and show report **outage-link-summary**.

Use the **subnet** keyword to filter the outage report by subnet.

Use the **link** keyword to filter the outage report by the specified link.

Use the **trended-live** keyword to allow in the outage report only those networks that are currently being trended live.

Use the **limit** keyword to restrict the number of networks that will be displayed in the report.

For descriptions of individual reports, see Chapter 3: Reports of the *Operator Guide*.

show outage chatter

Displays the cells that are currently in a fault state based on their Chatter measurements. The fault states are:

- Above or Rapidly Rising, typically colored red in Web interface
- Stale, typically colored yellow in Web interface

Syntax

show outage chatter

Prerequisites

Must be in Privileged command mode.

Description

Displays a list of cells that are currently in a fault state based on their Chatter measurements. The fault states are:

- Above or Rapidly Rising, typically colored red in Web interface
- Stale, typically colored yellow in Web interface

A cell is considered in a fault state if any of its test variables has a fault state.

The report contains the following columns:

- Zone From
- Zone To
- Absolute State
- Relative State
- Test Name
- Test Variable

Sample report

Zone From	Zone To	Absolute	Relative	Test Variable
US	APAC	Above	Rising	tcp-cerror
Milpitas	Westminster	Above	Rising	tcp-cerror

show outage-not-rescued

Displays the rescued outages.

Syntax

```
show outage-not-rescued [ subnet <subnet> ]
            [ subnet-pair ]
            [ link <link> ]
            [ link-down <link> ]
```

Arguments

- *subnet*—subnet in CIDR format (a.b.c.d /*n*)
- *link*—name of a link

Prerequisites

System operations mode, any level access.

Description

The **show outages** command displays unrescued outages (outages that CNA was **NOT** able to fix) within the last 24 hours.

Use the **subnet** keyword to filter the outage report by subnet.

Use the **link** keyword to filter the outage report by the specified link.

Use the **link-down** keyword to display only link down outages for the specified link.

For descriptions of individual reports, see Chapter 3: Reports of the *Operator Guide*.

show outage-rescued

Displays the rescued outages.

Syntax

```
show outage-rescued [ subnet <subnet> ]
            [ link <link> ]
            [ link-down <link> ]
            [ limit <limit> ]
            [ trend-live <trend-live> ]
```

Arguments

- *subnet*—subnet in CIDR format (a.b.c.d /*n*)
- *link*—name of a link
- limit—
- trend-live—

Prerequisites

System operations mode, any level access.

Description

The **show outages** command displays rescued outages (outages that CNA was able to fix) within the last 24 hours.

Use the **subnet** keyword to filter the outage report by subnet.

Use the **link** keyword to filter the outage report by the specified link.

Use the **link-down** keyword to display only link down outages for the specified link.

For descriptions of individual reports, see Chapter 3: Reports of the *Operator Guide*.

---

show penalty-setting link

Displays the penalty values that would be required to move a percentage of subnets off a specified link.

Syntax

```
show [ last ] penalty-setting link <name>
     [ move-percentage <percent> ]
     [ load]
     [ | { include | exclude | begin } <exp> ]
     [ > <url> ]
```

Arguments

- *name*—name of service provider link
- *percent*—percentage of traffic to be moved
- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, any level access.

Description

The **show penalty-setting link** command displays the penalty values required to move various percentages of subnets off the link specified by the *name* argument. See the **penalty** command for information about assigning penalties to service provider links.

Note:

Output of the **show penalty-setting link** command is valid as long as the CNA decision-maker process uses the same winner-set-width and penalty values when selecting routes, for all subnets. If you override these values for some subnets using the **decision-policy** command, the suggested penalty settings will not be accurate.

If you use the **last** keyword, data used for the calculations will be obtained from the most recent snapshot image of the database.

If you use the **move-percentage** keyword with a *percent* value, the output will show the penalty value required to move the specified percentage of subnets. If you omit the **move-percentage** keyword, the output will show the values required to move subnets in 10% increments, up to 100%.

If you use the **load** keyword, data for each subnet will be weighted by an estimate of its load, and percentages will approximate traffic volume for each link instead of just showing numbers of subnets.

For example, if you have a service provider link named *link_1*, you might get output similar to the following:

```
show penalty-setting link link_1

% of Total % moved Penalty req'd
        25       0             0
        22      10            35
        20      20            42
        17      32            46
        13      47            48
        11      56            49
         7      69            50
         5      76            51
         4      81            52
         2      90            56
         0     100           291
```

If you wanted to know how high to set the penalty in order to move 50 percent of the subnets currently served by the link off the specified link, the command and output would look similar to the following:

```
CNA# show penalty-setting link ISP-1 move-percentage 50
Penalty setting information for link: ISP-1
```

```
Current penalty setting: 0


% of Total % moved Penalty req'd
        25       0             0
        11      56            49
```

The first line of numbers shows the current distribution on the link, at the current penalty setting.

The second line shows the penalty required to achieve the specified percentage of movement.

Note:

Since neither subnets nor penalty values can be fractional, the numbers in the table will be adjusted as necessary in order to be able to move whole subnets or to display penalties as integers.This also means that the values in the middle column only approximate multiples of 10. And, if you specify a percentage that would require a fractional penalty, the output will be adjusted to the next highest penalty and its equivalent percentages. In the example, we asked for the penalty required to move 50 percent of the subnets. Since that would require a penalty between 48 and 49, the output shows the higher integer, 49, which will move 56 percent of the subnets.

The first column, **% of Total**, shows what percentage of the total subnets known to CNA are currently being served by the link. The second column, **% moved**, tells you what percentage of those subnets would be moved if the penalty were set to the value in the third column, **Penalty Req'd**.

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show prefix (deprecated)

The show prefix command has been deprecated. See the replacement command, show subnet on page 370.

show reload

Displays information on a scheduled reload.

Syntax

```
show reload
```

Prerequisites

Enable mode.

Description

Use this command to see if and when a scheduled reload will take place. See the **reload** command.

show report

Output snapshot report to monitor.

Syntax

```
show report [ <name> [ <object> ] ]
            [ customer <customer-name>
            [ time { <timestamp> | last | all } ]
```

Arguments

- *name*—name of report; can be one of:
  - *billing*
  - *distribution-all*
  - *distribution-all-weighted*
  - *distribution-named*
  - *distribution-named-weighted*
  - *distribution-sufficient*
  - *distribution-sufficient-weighted*
  - *executive-summary*
  - *executive-summary-weighted*

  - *hrtt*
  - *measurements*
  - *metric*
  - *observed*
  - *optimized*
  - *outage-link-summary*
  - *outage-subnet-summary*
  - *summary-of-improvements*
  - *summary-of-improvements-weighted*
- *object*—name of router or link (required for selected reports)
- *customer-name*—name of the customer reports that are defined by <customer>.subnet files (see the **`customer-subnet-file`** command)
- *timestamp*—timestamp of specific report, in *yyyy-mm-dd_HH-MM_TZ* format
- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Description

The **`show report`** command will cause a specified snapshot report to be displayed on the monitor.

For descriptions of individual reports, see Chapter 3: Reports of the *Operator Guide*.

If you enter the **`show report`** command with partial data and multiple reports match the entry, the output will display a list of the reports that match the data.

For example, entering just **show report** will cause CNA to display a list of all reports.

If you enter **show report** *executive-summary*, the device will display a list of all *executive-summary* reports.

Specify all or part of a date for *timestamp.* Reports shown will be those that come closest to matching the specified time.

For example, if you specify only a year, a list of all reports for that year will be displayed. (You must specify everything to the left of the desired time element. If you specify month you must include a year; if you specify day, you must include month and year, and so on.)

The timezone element of *timestamp* (TZ) should be the three-letter code displayed in the list of reports, not the full timezone name entered with the **`timezone`** command.

If you enter **show report** *executive-summary 2002-06*, the device will display a list of all *executive-summary* reports currently on the hard drive with a timestamp any time in June of 2002. Note the use of hyphens as separators between elements of the date and elements of the time, and the use of the underscore to separate date from time.

If only one report matches the partial data, the report itself will be shown.

If you use the **last** keyword instead of a timestamp, the device will display the most recently saved instance of the specified report.

See also the **`delete report`** and **`export report`** commands.

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **`copy`** command for URL syntax.

show running-config

Displays CNA's running configuration.

Syntax

```
show running-config [ | { include | exclude | begin } <exp> ]
                    [ > <url> ]
```

Arguments

- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, privileged-level access. Can also be executed in configuration (*config*) mode.

Description

Use the **`show running-config`** command to display the contents of the running configuration.

You can pipe the command through one of the following filters:

- **include**—shows all commands present in the running configuration which contain a match to the specified regular expression
- **exclude**—shows all commands present in the running configuration except those which contain a match to the specified regular expression
- **begin**—shows all commands present in the running configuration beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

```
PathControl8# show running-config | include interface
interface Ethernet 0/0
interface Ethernet 3/0
interface Ethernet 4/0
interface Ethernet 9/0
  polling interface-ip-address 172.30.100.54
  polling interface-ip-address 172.29.100.54
```

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **`copy`** command for URL syntax.

---

show snmp

Displays the status of the SNMP polling configuration.

Syntax

```
show snmp [ | { include | exclude | begin } <exp> ] [ > <url> ]
```

Arguments

- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

Must be in system-operations mode, with login-level access.

Description

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

---

show ssl

Displays information about the Secure Sockets Layer certificates currently installed on CNA.

Syntax

```
show ssl [ | { include | exclude | begin } <exp> ] [ > <url> ]
```

Arguments

- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

Must have privileged-level access.

Description

Once installed, SSL certificates remain in memory as long as SSL is enabled. If you have used the **`ssl delete`** command to remove files from a module on which SSL is still enabled, the deleted certificates will still exist in memory and therefore will appear in the output from the **`show ssl`** command.

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression

- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show startup-config

Displays the contents of the startup configuration.

Syntax

```
show startup-config [ | { include | exclude | begin | begin } <exp> ]
                    [ > <url> ]
```

Arguments

- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, privileged-level access. Can also be executed in configuration (*config*) mode.

Description

Use the **show startup-config** command to display the contents of the startup configuration.

You can pipe the command through one of the following filters:

- **include**—shows all commands present in the startup configuration which contain a match to the specified regular expression
- **exclude**—shows all commands present in the startup configuration except those which contain a match to the specified regular expression
- **begin**—shows all commands present in the startup configuration beginning with the first match to the specified regular expression

The vertical-line pipe character is required when a filter is used:

```
PathControl8# show startup-config | include interface
interface Ethernet 0/0
interface Ethernet 3/0
interface Ethernet 4/0
interface Ethernet 9/0
  polling interface-ip-address 172.30.100.54
```

```
polling interface-ip-address 172.29.100.54
```

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show statistics

Displays information about specified internal operating system processes.

Syntax

```
show statistics
     <process>
     [ chatter ]
     [ database ]
     [ | { include | exclude | begin } <exp> ]
```

Arguments

- *process*—name of the internal process; can be one of:
  - *am-measurer*—active-measurement process on a USTAT module
  - *am-scheduler*—active-measurement process on the engine module
  - *dm-listener*—decision-maker listener process on the engine module
  - *dmaker*—decision-maker process on the engine module
  - *edgestats*—user-traffic-test process on a USTAT module
  - *efc*—endpoint collection process on the EFC module
  - *event-manager*—Event Manager process on the management module
  - *link-monitor*—SNMP polling process
  - *logging*—logging process on any module
  - *monitor*—monitor (SPAN listener) process on the EFC module
- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

Must have privileged-level access.

Description

The **show statistics** command displays data related to a specified CNA operating-system process.

For additional options available to you for *dmaker* process output, see the **show statistics dmaker** command description.

You can only use the **detail** keyword with the dmaker and efc processes.

Use the *slot* argument or the keyword **all** only for processes that can reside on multiple modules, such as USTAT-based processes or logging. Engine module process do not take a module-position argument.

The keyword **avoid** is valid only with the *am-measurer* process; use it to display a list of addresses that are not to be probed, either because you have named them as arguments to the **active-measurement traceprobe avoid address (discontinued)** command or because they have returned an ICMP Administratively Unreachable message.

The keyword **diagnose** is valid only with the *am-measurer* and *am-measurer* processes; use it to display diagnostic statistics captured with the **diagnose probe** command.

You can pipe the command through one of the following filters:

- **include**—display all lines in the **show statistics** output which contain a match to the specified regular expression
- **exclude**—display all lines in the **show statistics** output except those which contain a match to the specified regular expression
- **begin**—display all lines in the **show statistics** output beginning with the first match to the specified regular expression

The vertical-line pipe character is required when a filter is used:

```
CNA# show statistics am-measurer | include Restart
Restart Time (Since Epoch): 1031084562
Restart Time: Tue Sep  3 13:22:42 2002
Restart Time (Since Epoch): 1031084571
Restart Time: Tue Sep  3 13:22:51 2002
```

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show statistics chatter

Displays a CNA server's agent-based measurement configuration.

Syntax

show statistics chatter [admin | alarms | cells | classes | destinations | edges | load | tasks | zones]

Arguments

Table 47: show statistics chatter arguments

Argument	Description
admin	Displays the following data: • Chatter missed heartbeats • Intra-zone testing status (enabled or disabled) • Relative alarm status (enabled or disabled) • Absolute alarm status (enabled or disabled) • Alarm threshold exceeded • Alarm control resistance
alarms	Displays all alarm templates and their configurations.
cells	Displays all cells, their source and destination, and owner. The cell owner is the CNA server that stores a cell's data.
classes	
destinations	Shows all configured alarm destinations and their configurations: • E-mail recipients • SNMP hosts • Syslog servers
edges	Displays data for each edge in the hive: • Source and destination IP addresses • Owner (CNA server that stores the edge's data) • Alarm template assigned to it
load	Displays data about the state of load optimization on CNA.
tasks	Displays configuration for each of four tests: RTP, TCP, Ping, and Traceroute.
zones	Displays all zones in the hive, the alarm template assigned to each zone, and the number of CNA agents assigned to each zone.

Prerequisites

Must be in Privileged command mode.

Description

Displays a CNA server's agent-based measurement configuration.

---

show statistics dmaker

Displays detailed information about the CNA decision-maker process.

Syntax

```
show statistics dmaker [ decision-making |
                         dynamic-subnets |
                         event-queue |
                         group |
                         link-dump [ <link> ]
                         load |
                         observed-applications
                         peerings |
                         performance-group
                         policy [ <policy-name> ] |
                         subnet-dump <subnet> |
                         subnet-scheduling |
                         scheduler-queue |
                       ] [ detail ]
```

Arguments

- *policy-name*—name of a decision policy in the CNA configuration
- *subnet*—subnet address, in CIDR format (*a.b.c.d/n*)
- *link*—name of link in the CNA configuration

Description

The **`show statistics dmaker`** command is an extension of the **`show statistics`** command.

Use one of the following key words to specify the part of the CNA *decision-maker* process for which you want to see statistics:

- **decision-making**—displays statistical data about decisions made by the CNA *decision-maker* process

- **dynamic-subnets**—displays statistics about subnets that were assigned a dynamic policy by CNA.
- **event-queue**—displays statistical tables about the subnet event queue
- **link-dump**—displays configuration parameters, SNMP data, and load data for the specified link (called SPAL in the output, for *service provider access link*)
- **load**—displays data about the state of load optimization on CNA
- **observed-applications**—displays statistical data about applications (e.g., bgp, oracle, kazaa, etc.) that CNA has encountered.
- **peerings**—displays statistical data about BGP announcements and withdrawals
- **policy**—displays configuration parameters for the specified decision policy
- **subnet-dump**—displays data about the subnet as currently known to the CNA system, and configuration parameters currently being applied to the subnet
- **subnet-scheduling**—summarizes what the CNA system knows about subnets in general
- **scheduler-queue**—displays data about the queue used by the CNA *decision-maker* process to determine the order in which subnets will be evaluated
- **detail**—displays detailed information about the requested data (this does not include **policy**, **subnet-dump**, or **link-dump**).

show subnet

Displays information about a subnet or group of subnets.

Syntax

```
show [ last ] subnet
                [ all |
                 sufficient |
                 asserted |
                 <subnet> |
                 <address> |
                 <domain-name> |
                 alias <alias-name> |
                 group <group_name> |
                 using-policy <policy-name>
                ]
                [ trended-live ]
                [ limit <number> ]
                { application |
                 application-loss |
                 application-performance |
                 application-delay |
                 measurements |
```

```
raw-measurements |
observed |
optimized |
policy |
info |
hrtt |
metric |
subnet-loss |
jitter |
quality-metric |
reliability |
rtt
summary
traffic-counter
traffic-rate
}
[ > <url> ]
```

Arguments

- *subnet*—individual subnet address, in CIDR format (a.b.c.d/*n*)
- *address*—individual IP address, in dotted decimal format (a.b.c.d)
- *alias-name*—subnet alias (see **report prefix-alias (deprecated)** command)
- *group_name*—name of subnet group (see **report group** command)
- *policy_name*—name of policy (see **decision-policy** command)
- *number*—number of subnets to include in output (range 0 through 1,000,000; default is 500; specifying 0 indicates no limit)
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Default

none

Description

The **show subnet** command will display a variety of data related to subnets.

Use the **last** keyword to obtain the data from the snapshot database maintained on the reporting module.

If you omit the **last** keyword, the data displayed will be taken from the live database maintained on the engine module. Avaya strongly recommends that whenever you specify **all**, you also use the **last** keyword. Generating data for all subnets can consume an enormous amount of system resources.

Use the keyword **all** to display data for all subnets in the database.

Use the **trended-live** keyword to show only those subnets that are currently being trended live.

The keyword **limit** followed by a number limits the number of subnets in the display. Set **limit** to 0 for no restriction on quantity (all subnets meeting the specified criteria will be included in the output).

Use the keyword **sufficient** to limit the display to subnets which have met the criteria for route assertion (minimum required measurements have been received for the subnet, and the subnet is reachable by at least two service provider links).

You can filter subnets by substituting values for any one of the following optional arguments: *number*, *subnet*, *address*, *domain-name*, or *group*.

You can specify any of the following data-type keywords:

- **application**—output will display information for each application (or group of applications) on all networks.
- **application-loss**—output will display loss statistics for all links associated with the specified subnets.
- **application-performance**—output will display performance ratings from 1.0 (worst) to 5.0 (best) for all links associated with the specified subnets.
- **application-delay**—output will display an estimate (given the raw measurements observed for a subnet) of the application delay that would be experienced for all links associated with the specified subnets.
- **measurements**—display will show a table of measurement data for each service provider link on each subnet, including number of measurements made, average inter-measurement time (IMT), and time since last measurement (TSLM).
- **raw measurements**—display will show actual measurement values for all links associated with the specified subnets. These data include round trip time (RTT), jitter, loss, and TCP-based round trip time (HRTT).
- **observed**—display will show a comparison of HRTT scores for the best-performing, worst-performing link, and BGP-selected link on each edge router. For details about the contents of the display, see Observed report in Chapter 3: Reports of the *Operator Guide.*
- **optimized**—all data associated with the specified subnets obtained from the specified router (identified by address or alias defined by the **`neighbor report-alias`** command)
- **policy**—display will show information on the decision policy of the specified subnets.
- **info**—display will show summary information for the specified subnet.
- **hrtt**—output will display handshake round trip time (HRTT) values obtained from CNA's user traffic tests (UTT) and TCP probes, in milliseconds for all links associated with the specified subnets.

- **metric**—display will show an index for each subnet listed (a number in the range 0 through 1,000 calculated from HRTTs and user-defined penalty) depicting the quality of a web transaction as perceived by the end user; CNA uses this metric as a basis for making routing decisions (smaller is better)
- **subnet-loss**—output will display the average loss for all links associated with the specified subnets. These values are available only for subnets receiving measurements from ICMP or traceprobe active-probing.
- **jitter**—displays the average jitter values in milliseconds for all links associated with the specified subnets. Jitter is calculated as the difference between consecutive round trip time (RTT) values. Jitter is available only for subnets receiving measurements from ICMP or traceprobe active-probing.
- **quality-metric**—output will display a quality index for all links associated with the given subnets. The index, which is calculated form raw measurements, ranges from 0 (best) to 1000 (worst) and is indicative of the quality of the performance of each link per subnet.
- **reliability**—the average time a link is in the winner set for each subnet
- **rtt**—output will display round trip time (RTT) values, in milliseconds, for all links associated with the specified subnets. RTT is available only for subnets receiving measurements from ICMP or traceprobe active-probing.
- **traffic-counter**—output will display application since the start of CNA.
- **traffic-rate**—output will display average traffic rate l.

subnet categories are:

- A—advised winner (CNA has asserted a route for this subnet that is better than the one selected by BGP but the router has not yet confirmed that it has switched)
- B—BGP winner (CNA has determined that the route chosen by BGP is in fact the best path to this subnet)
- C—confirmed winner (CNA has asserted a route for this subnet that is better than the one selected by BGP and the router has confirmed that it has switched)
- D—determined winner (CNA has determined that a route for this subnet is better than the one selected by BGP but has not yet asserted this route to the router)
- I—insufficient data (CNA has not yet received enough measurement tests from this subnet; or the BGP next hop is unknown to CNA)
- N—not used; subnet is in the BGP table, but CNA considers it to be an invalid route
- R—remote BGP; traffic to the subnet is routed through a service provider link that is not defined in the CNA configuration
- U—unavailable data (subnet is in the BGP table but CNA has no measurement data for it)

The category identifier may also show one or both of the following symbols:

- ~—new data (measurement data has been received for this subnet but nothing has not yet been processed)

- ^—non-bgp subnet; CNA learned of this subnet from some source other than BGP (a private, VPN-related subnet added by the **set-decision-policy** command, for example)

When load optimization is enabled, the subnet category will also be accompanied by one of the following labels, indicating how the winner was chosen:

- r—winner was selected from the regular winner set
- e—winner was selected from the extended winner set
- h—all links hot; winner is “least hot”
- j—no eligible winner in the extended winner set at the threshold-table level appropriate for current usage; the decision-maker process had to jump to a higher level (indicates a narrow extended winner set; you should consider increasing the size)
- l—no eligible winners could be selected by the load optimization algorithm; winner selected by load optimization (current usage has exceeded the maximum total load configured at the highest level of your load threshold table; it means that either the table is not configured correctly or you have insufficient capacity to meet demand)

The level indicator label will be accompanied by the total load value configured for that level.

Service provider links may be labeled by one or more of the following codes in some displays:

- a—CNA has asserted link to be best path
- b—link selected by BGP (multiple links may show b label for a single subnet, if there are multiple routers)
- h—link has been categorized as hot by CNA *decision-maker* process. See Determining the Winner Set in Chapter 2: WAN Cost and Load Optimization of the *Adaptive Path Controller Guide.*
- i—data is invalid; link no longer considered reliable for this subnet
- o—outage detected; subnet is no longer generating measurement traffic
- s—link is in winner set (see **winner-set-width** command)
- w—link is CNA-determined winner (best path to subnet)

Output of the **show subnet** command can be diverted to a file, using the > character as a keyword, followed by a file name in URL format (see the **copy** command for information about supported protocols and URL syntax).

Note that a space is required between the > character and the URL.

Note:

Occasionally, the **show subnet** command will fail with an “unknown error” message. This could be due to a conflict with a database task being executed by CNA concurrently with your **show subnet** request. Wait a few minutes and reissue the **show prefix (deprecated)** command.

See also the **report prefix-limit (deprecated)** command.

show tech-support

Displays diagnostic information about the current state of CNA.

Syntax

```
show tech-support [ detail ]
                  [ no-config ]
                  [ | { include | exclude | begin } <exp> ]
                  [ > <url> ]
```

Arguments

- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Description

If you use the **detail** keyword, output will include the output generated by the **`show diagbus`** command.

Use the **no-config** keyword to exclude the running-config from the output.

Output of the **`show tech-support`** command can be diverted to a file, using the > character as a keyword, followed by a file name in URL format (see the **`copy`** command for information about supported protocols and URL syntax).

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **`copy`** command for URL syntax.

show trending

Displays a list of configured sizes for trending database files, and a list of all objects that are being tracked in the trending database.

Syntax

```
show trending
              [ | { include | exclude | begin } <exp> ]
              [ > <url> ]
```

Arguments

- *family*—trend size family; must be one of:
  - *performance*
  - *event-manager (deprecated)*
  - *link-usage*
- *address*—subnet address in CIDR format (a.b.c.d/*n*)
- *lname*—link name, as string (see **`link`** command)
- *rname*—IP address of a router
- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Description

The **`show trending`** command will display current configured sizes of the trending database files for the specified trending size family, plus the globally configured sizes.

The output also shows the total disk space used by the trending database, in bytes; whether trending is enabled (by the **`trending on`** command); and the number of minutes after the hour when the database snapshot will be made (configured by the **`snapshots schedule`** command).

If you use the keyword **subnet**, the display will show a list of subnets for which trend data is available. If you also include a specific subnet address in CIDR format, the output will include a list of all files related to that subnet.

If you use the keyword **link**, the display will show a list of links for which trend data is available. If you also include a specific CNA-configured link name, the display will include a list of all files in the trend database related to that link.

If you use the keyword **router**, the display will show a list of routers for which trend data is available. If you also include a specific router's IP address or hostname, the display will include a list of all files in the trend database related to that router.

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show users

Displays a list of users currently logged on to CNA.

Syntax

```
show users [ | { include | exclude | begin } <exp> ] [ > <url> ]
```

Arguments

- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Description

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

show zone

Displays a tabular report of test results between subzones of a zone.

Syntax

show zone
[name <zone-name>] [test {rtp-delay | rtp-jitter | rtp-loss | rtp-inter-arrival-time | rtp-drops | ping-rtt | ping-loss | ping-inter-arrival-time | ping-drops | tcp-delay | tcp-cerror | tcp-inter-arrival-time | tcp-drops}]
[criterion {relative | absolute | median | minimum | maximum}]

Arguments

Table 48: show zone arguments

Argument	Description
<zone-name>	Zone for which you want to view test measurements. You can enter a fully qualified zone name or a specific zone name. The default value is the root zone.

Table 48: show zone arguments (continued)

Argument	Description
[test {rtp-delay | rtp-jitter | rtp-loss | rtp-inter-arrival-time | rtp-drops | ping-rtt | ping-loss | ping-inter-arrival-time | ping-drops | tcp-delay | tcp-cerror | tcp-inter-arrival-time | tcp-drops}]	Specific test variable for which you want to view test measurements. Options are: • RTP: - Delay - Jitter - Loss - Inter arrival time - Drops • Ping - Round trip time - Loss - Inter arrival time - Drops • TCP - Delay - Cerror - Inter arrival time - Drops
relative	Displays status of test measurements in context of the test's relative thresholds. For example, if the measurements are consistently falling within the upper and lower thresholds, the report displays `Steady`. **Note:** Relative alarms are generated when the most recent results for a specific test show an abrupt increase or decrease. The thresholds for determining an "abrupt increase or decrease" are self-adjusting thresholds based on a statistical technique. Note that if the measurement remains near or outside the threshold, the threshold adjusts to reflect this changed norm. Once the threshold self-adjusts, the alarm clears even though the measurement itself has not changed. The steady state becomes the threshold value. This value can be higher or lower than the configured threshold.

2 of 3

Table 48: show zone arguments (continued)

Argument	Description
absolute	Displays status of test measurements in context of the test's absolute threshold. For example, if the measurements are consistently below the absolute threshold, the report displays `Ok`.
median	Displays median of test measurements.
minimum	Displays minimum of test measurements.
maximum	Displays maximum of test measurements.
	3 of 3

Prerequisites

Must be in Privileged command mode.

Default

Table 49: Default settings show zone command

Argument	Default setting
<zone-name>	Root zone name
criterion	relative

Description

Displays a tabular report of test results between subzones of a particular zone. To display the results for a specific test variable, specify the test variable. By default the command displays test results for the first active test measurements in the list.

Table 50 shows a sample report for zone AV, RTP delay median.

Table 50: Sample output for show zone

	APAC	AVBackbone	CALA	EMEA	US
APAC					
AVBackbone	125				
CALA		87			
EMEA					59
US					

show version

Displays information about the CNA software as well as the hardware installed on each module.

Syntax

```
show version [ | { include | exclude | begin } <exp> ] [ > <url> ]
```

Arguments

- *exp*—a regular expression, to be used as a filter
- *url*—destination for output

Prerequisites

System operations mode, login-level access.

Description

The **show version** command shows version information about CNA, as well as information about CPU type and speed, and size of RAM and memory cache, for each module installed in CNA.

If you execute this command during an upgrade process, the output will list any upgrade tasks that are still running in the background, as well as any that will be completed only after the next time the **reload** command is executed. (A reload is the last step in the upgrade process; it should not be done while tasks are shown as still running.)

You can pipe the command through one of the following filters:

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression
- **begin**—shows all lines in the output beginning with the first match to the specified regular expression

The vertical-line pipe character | is required when a filter is used.

You can send the output to a remote file location, by piping the display to a remote URL (leading > required). See the **copy** command for URL syntax.

shutdown

Places a **shutdown** command in the interface definition block of the running configuration.

Syntax

```
shutdown
no shutdown
```

Prerequisites

Must be in interface configuration mode (*config-if*), which requires privileged-level access.

Occurrence

Use once per interface.

Default

Interface is not shutdown.

Description

The **shutdown** command renders the interface inactive without removing it form the configuration.

The **no** form of the command removes the **shutdown** command from the interface block of the running configuration, thereby restoring the previously shut-down interface to active status.

See also the **neighbor shutdown** command.

site

Associates a subnet with a group name.

Syntax

```
site <site-name>
        { subnet <subnet>| site <generic-site-name> }
no site <site-name>
no site <site-name> subnet <subnet>
no site <site-name> site <generic-site-name>
```

Arguments

- *site-name*—string; name of group

- *subnet*—subnet in CIDR format (a.b.c.d /*n*)
- *generic-site-name*—string; name of generic site

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

The **site** command associates subnets under a common name for internal group processing by CNA.

In contrast, the **report subnet-alias** command simply associates subnets and aliases in an internal configuration file, and the alias is used only for generating reports.

If the **no** form is executed with no subnet, all subnets are removed from the group.

If the **no** form of the command is executed with a subnet, the specified subnet is deleted from the group.

If the **no** form of the command is executed with a generic group name, then you will remove the association of the subnets in the generic group from the subnets in this group.

---

smtp from

Sets the default text for the **from** field in CNA's web client e-mail messages.

Syntax

```
smtp from <text>
no smtp from
```

Arguments

- *text*—name

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Description

CNA's web client provides the ability to send e-mail messages.

Use the **`smtp from`** command to set default text for the from field in the e-mail message (can be changed by the user when sending mail).

Use the **no** form of the command to remove default text.

---

smtp to

Sets the default text of the **to** field in CNA's web client e-mail messages.

Syntax

```
smtp to <text>
no smtp to
```

Arguments

- *text*—addressee's name

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Description

Use the **no** form of the command to remove the default text.

---

smtp to (Chatter module only)

Specifies the e-mail addresses to which alarm notifications for this alarm template are sent.

Syntax

smtp to <recipient>

no smtp to <recipient>

Argument

Table 51: smtp to arguments

Argument	Description
<recipient>	E-mail address to which you want alarm notifications sent. Enter one of the e-mail addresses configured with the smtp-server (Chatter module only) command. Enter the e-mail address in the format a@b.c, or use an e-mail alias that is defined in the directory.

Prerequisites

Must be in Chatter Alarm Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Up to four e-mail recipients can be configured per alarm template.

Description

Specifies the e-mail addresses to which alarm notifications for this alarm template are sent. Alarm notifications are sent for alarms originating on a zone or edge that is assigned this alarm template. You can specify any or all of the e-mail addresses configured with the smtp-server (Chatter module only) command.

The sender of the e-mail message is as configured by the smtp-server (Chatter module only) command.

The no command disables alarm notifications to the specified e-mail address for this alarm template.

smtp-server

Sets the SMTP (Simple Mail Transfer Protocol) server to be used when sending e-mail from the CNA web client.

Syntax

```
smtp-server <name> [ <port> ]
no smtp-server
```

Arguments

- *name*—SMTP server domain name (i.e., mail.Avaya.com)
- *port*—access port on the server

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

port—25

Description

Use the **no** form of the command to delete the SMTP server definition.

smtp-server (Chatter module only)

Forwards alarms to an e-mail address.

Syntax

smtp-server to <recipient> host <ip-address> [port <port-number>] [from <sender>]

no smtp-server to <recipient> host <ip-address>

Arguments

Table 52: smtp-server arguments

Argument	Description
<recipient>	E-mail address to which you want the alarms forwarded. Enter the e-mail address in the format a@b.c, or use an e-mail alias that is defined in the directory.
<ip-address>	IP address of the SMTP server.

Table 52: smtp-server arguments

Argument	Description
<port-number>	Port number of the SMTP server.
<sender>	E-mail address from which you want the e-mail sent. Enter the e-mail address in the format a@b.c, or use an e-mail alias that is defined in the directory.

Prerequisites

Must be in Chatter Configuration command mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Up to four instances of this command can be configured for the Chatter module.

Default

Table 53: Default settings for smtp-server command

Argument	Default value
<port-number>	25
<sender>	CNA

Description

Forwards alarms to an e-mail address.

snapshots clear-after (deprecated)

The snapshots clear-after command has been deprecated.

snapshots on

Causes CNA to generate snapshot reports.

Syntax

```
snapshot on
no snapshot on
```

Prerequisites

Must be in report configuration mode (*config-report*), which requires privileged-level access. Snapshot reporting mode requires CNA's optional Analysis and Reporting Module.

Occurrence

Use once.

Default

Snapshot reporting mode is on (reports are generated) if the system includes the optional Analysis and Reporting Module.

Description

Use the **no** form of the command to turn snapshot reporting mode off.

When snapshot reporting mode is on, CNA will generate snapshot reports according to the schedule defined by the **`snapshots schedule`** command.

This process may take a few minutes to complete. As a safety mechanism to prevent reporting errors, the process will timeout after 90 minutes.

snapshots prefix-limit (deprecated)

The snapshots prefix-limit command has been deprecated. See the replacement command, snapshots subnet-limit on page 390.

snapshots schedule

Sets a schedule for generating snapshot reports by the Analysis and Reporting module.

Syntax

```
snapshots schedule <hour> [ <day-of-month> [ <month> [ <day-of-week> ]]]
no snapshots schedule
```

Arguments

- *hour*—integer in range 0 through 23, or * (for every hour), or comma separated list of hours (1,2)
- *day-of-month*—integer in range 1 through 31, or * (for every day of the month), or comma-separated list of days (1,15)
- *month*—integer in range 1 through 12, or * (for every month of the year), or comma-separated list of months (1,4,7,10)
- *day-of-week*—integer in range 0 through 7, or * (for every day of the week; both 0 and 7 designate Sunday), or comma-separated list of days (1,2,3,4,5)

Prerequisites

Must be in report configuration mode (*config-report*), which requires privileged-level access.

Occurrence

Use once.

Default

Reports are generated every 12 hours, at 00:00 and 12:00:

- *hour*—0,12
- *day-of-month*—*
- *month*—*
- *day-of-week*—*

Description

Scheduling is done by the **cron** utility. Arguments are specified in limited **crontab** format (only those forms listed in the *Arguments* section of this description; do not include spaces in comma-separated lists).

If you specify both a day of the month and a day of the week, both will be honored. That is, if you specify the third day of the week (Tuesday) and the fourth day of the month, reports will be generated every Tuesday and on the fourth day of the month regardless of the day of week.

Use the **no** form of the command to reset the schedule to the default.

Snapshot reports are generated from the hourly snapshot image of the database. The schedule you set with the **`snapshots schedule`** command determines how often the reporting module will extract data from the database to. This extracted data will be used for the following purposes:

- generate snapshot reports

- update the trending database

See the **sample-collection** command for information about the hourly snapshot image.

See the **trending size** command for information about the trending database.

snapshots subnet-limit

Sets the number of subnets to be included in snapshot reports.

Syntax

```
snapshots subnet-limit <number>
no snapshots subnet-limit
```

Arguments

- *number*—number of subnets to include in snapshot reports; range: 1 through 1000

Prerequisites

Must be in report configuration mode (*config-report*), which requires privileged-level access.

Occurrence

Use once.

Default

100

Description

When snapshot reports are enabled, CNA automatically generates snapshot reports each time a snapshot image of the database is made. These reports show data for a certain number of the most-active subnets (those which have generated the most measurements), plus named subnets and active-probe targets. The **snapshots subnet-limit** command determines how many subnets are in the most-active grouping.

Use the **no** form of the command to reset the limit to the default.

snmp-agent-ip-address

Provides CNA with the address of an SNMP-enabled interface on an edge router, for polling.

Syntax

```
snmp-agent-ip-address <address>
```

Arguments

- *address*—the IP address of an SNMP-enabled interface on the edge router

Prerequisites

Must be in polling agent configuration mode (*config-polling-snmp*), which requires privileged-level access.

Occurrence

Use once per polling agent

Description

Use the **no** form of the command to remove the IP address from the polling agent.

See the **`polling snmp-agent`** and **`polling usage on`** commands for more information about polling edge routers.

---

snmp host

Specifies the hosts to which SNMP traps for this alarm template are sent.

Syntax

snmp host <host-id>

no snmp-server host <host-id>

Argument

Table 54: snmp host arguments

Argument	Description
<host-id>	IP address of the host or the host name to which traps are sent. Enter one of the hosts configured with the snmp-server command.

Prerequisites

Must be in Chatter Alarm Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Up to four hosts can be configured per alarm template.

Description

Specifies the hosts to which SNMP traps for this alarm template are sent. SNMP traps are sent for alarms originating on a zone or edge that is assigned this alarm template.You can specify any or all of the hosts configured with the snmp-server command.

The no command disables traps to this host for this alarm template.

snmp-server

Defines the host to which Chatter sends SNMP traps or informs.

Syntax

snmp-server host <host-id> community <community-string> [udp-port <port>]

no snmp-server host <host-id>

Arguments

Table 55: snmp-server arguments

Argument	Description
<host-id>	IP address of the host to which Chatter sends SNMP traps or informs.
<community-string>	Community string to be sent with the SNMP trap or inform.
<port>	UDP port number on the host where Chatter sends SNMP traps or informs.

Prerequisites

Must be in Chatter Configuration command mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Up to four SNMP server hosts can be configured for the Chatter module.

Default

No SNMP notifications are generated by Chatter unless you use this command to configure them.

Only version 1 traps are generated.

By default notifications are sent to UDP port 162.

Description

Defines the host to which Chatter sends SNMP traps.

The command specifies whether the notification is sent as a trap.

The no command disables Chatter traps or informs to the specified host.

---

snmp-server contact

Sets the SNMP MIB's *sysContact* string.

Syntax

```
snmp-server contact <text>
no snmp-server contact
```

Arguments

- *text*—contact data

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

No system-contact information available.

Description

Use the **no** form of the command to remove the contact information from the MIB.

snmp-server enable traps

Configures traps or informs that will be generated by CNA.

Syntax

```
snmp-server enable traps <family>
no snmp-server enable traps [ <family> ]
```

Arguments

- *family*—trap family to be enabled/disabled; can be one of the following:
  - ifmib
  - v2mib
  - pcmib
  - pcmib-verbose
  - outage

Prerequisites

Must be in configuration (*config*) mode, which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

All trap families are enabled by default.

Description

The **`snmp-server enable traps`** command works in conjunction with the **`snmp-server host`** command to configure SNMP traps and informs.

See the **`snmp-server host`** command for details and a list of traps within each family.

See also the **`snmp-server informs`** command.

The **no** form of the command disables the specified trap generated by CNA. If you use the **no** form without a *family* argument, all traps will be disabled.

snmp-server host

Defines hosts to which SNMP traps and informs should be sent.

Syntax

```
snmp-server host <address>
                 [ traps | informs ]
                 [ version {1 | 2c} ]
                 <community-string>
                 [ udp-port <port> ]
                 [ <family> ... ]

no snmp-server host <address> <community-string>
```

Arguments

- *address*—hostname or IP address in dotted decimal format (a.b.c.d) of destination
- *community-string*—access string on **host**
- *port*—UDP port on host
- *family*—trap family; can be one or more of the following:
  - ifmib
  - v2mib
  - pcmib
  - pcmib-verbose

Prerequisites

Must be in configuration (*config*) mode, which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

No notifications unless configured.

If **informs** keyword not specified, **traps** is assumed.

version 1

port—162

If *family* is not specified, all trap families that have been enabled are sent to the specified host

Description

The **snmp-server host** command works in conjunction with the **snmp-server enable traps** and **snmp-server informs** commands to send traps and informs to remote hosts.

If *community-string* is set to a value that has not been defined with the **snmp-server contact** command, it will be defined automatically by CNA.

Multiple *family* values can be specified, but only one host *address* per command. You specify which traps are to be sent by naming one or more families. You cannot set individual traps.

Traps (and their descriptions) are:

- **v2mib family:**
  - *authenticationFailure*—CNA's SNMP agent has received a request without proper authentication; Avaya has disabled access to the standard MIB variable *snmpEnableAuthenTraps* to prevent users from turning off the *authenticationFailure* trap.
  - *coldStart*—The CNA SNMP agent has been re-initialized, signifying a potential change in the configuration.
- **ifmib family:**
  - *linkUp*—An ethernet interface configured on CNA has been brought up; in this context, link does not refer to a service provider link as used elsewhere in Avaya terminology, but is any configured ethernet interface.
  - *linkDown*—An ethernet interface configured on CNA has been brought down, either by physically disconnecting it or by configuration command; in this context, link does not refer to a service provider link as used elsewhere in Avaya terminology, but is any configured ethernet interface.
- **pcmib family:**
  - *rsBgpPeerDown*—an CNA-configured BGP peering has gone down.
  - *rsBgpPeerUp*—an CNA-configured BGP peering has been established.
  - *rsConfigurationChangeTrap*—an CNA CLI user has entered *config* mode.
  - *rsDeviceShutdownTrap*—CNA is being shut down.
  - *rsEnvMonVoltageTrap*—A voltage threshold has been crossed on a module. See snmptrapd in Appendix A: Diagnostics of the *Appendices Guide*.
  - *rsEnvonTemperatureTrap*—A temperature threshold has been crossed on a module. See snmptrapd in Appendix A: Diagnostics of the *Appendices Guide*.
  - *rsImageCopyCompletionTrap*—Conveys status at the completion of an image-copy operation; see the **copy** command.
  - *rsLinkMeasurementTimeoutTrap*—No measurements have been received by a link for the interval specified by the **link-measurement-timeout** command.

- *rsLinkSnmpTrap*—The CNA *event-manager* process has received a service provider link status update from the link polling process. See Chapter 3: Polling Routers of the *Adaptive Path Controller-Internet Guide*.
- *rsMedicCpuTrap*—A CPU resource threshold has been crossed. See host-resources in Appendix A: Diagnostics of the *Appendices Guide*.
- *rsMedicHardDiskTrap*—A threshold for hard disk usage has been crossed. See host-resources in Appendix A: Diagnostics of the *Appendices Guide*.
- *rsMedicMemoryTrap*—A RAM usage threshold has been crossed. See host-resources in Appendix A: Diagnostics of the *Appendices Guide*.
- *rsMedicRamDiskTrap*—A threshold for RAM disk usage has been crossed. See host-resources in Appendix A: Diagnostics of the *Appendices Guide*.
- *rsMedicServiceTrap*—A service of the CNA Medic process has received a utilization-threshold trap or is attempting a restart.
- *rsStateTransitionTrap*—CNA has changed state to or from route-assert mode. Can be due to user configuration, a BGP status failure as defined in Link Outage Control of Chapter 3: Adaptive Path Control in the *Adaptive Path Controller Guide*, or lack of load data as defined by the **`polling usage unknown-threshold`** command.

- **pcmib-verbose family:**
  - *rsPrefixLinkActiveProbeBasedOutage*—A subnet outage has been detected by the CNA active-probe-based outage-detection process or a previously detected subnet outage has been cleared; see the **`outage-detection active-probe`** command.
  - *rsPrefixLinkSilenceBasedOutageDetectionTrap*—A subnet outage has been detected on one of the configured links by the CNA outage-detection process or a previously detected subnet outage has been cleared; see the **`outage-detection on (deprecated)`** command.
  - *rsPrefixTotalSilenceBasedOutageDetectionTrap*—A subnet outage has been detected over all links by the CNA outage-detection process or a previously detected subnet outage has been cleared; see the **`outage-detection total-silence`** command.

The traps in the **v2mib** and **ifmib** families are defined in standard MIBs, *RFC1213-MIB-2.txt* and *interfaces-MIB.txt* files respectively. The traps in the **pcmib** and **pcmib-verbose** family are defined in the Avaya enterprise MIBs.

The **no** form of the command disables CNA traps to the specified host.

See the *Appendices Guide* for more information about SNMP:

- Appendix A: Diagnostics, host-resources and snmptrapd for information about when certain traps are thrown
- Appendix B: SNMP for a list of all MIBs supported and the complete text of Avaya MIBs

snmp-server informs

Sets options for SNMP informs.

Syntax

```
snmp-server informs [ retries <num-tries> ]
                    [ timeout <seconds> ]
                    [ pending <num-pending> ]
no snmp-server informs
```

Arguments

- *num-tries*—number of retries
- *seconds*—number of seconds to wait for the host to respond
- *num-pending*—maximum size of queue of unacknowledged notifications

Prerequisites

Must be in configuration (*config*) mode, which requires privileged-level access.

Occurrence

Use once.

Default

num-tries—3

seconds—30

pending—25

Description

Use the **retries** keyword to specify the number of times an unacknowledged inform will be sent out.

Use the **timeout** keyword to specify how long CNA should wait for an acknowledgement before resending the notification.

CNA will maintain a queue of notifications waiting for acknowledgement. If the number in the queue exceeds the value of *num-pending*, older notifications will be discarded.

Use the **no** form of the command to reset all arguments to their defaults.

See the **`snmp-server host`** command for additional details.

snmp-server location

Sets the SNMP MIB's *sysLocation* variable for CNA.

Syntax

```
snmp-server location <text>
no snmp-server location
```

Arguments

- *text*—location information

Prerequisites

Must be in configuration mode, which requires privileged-level access.

Occurrence

Use once.

Default

Value of *sysLocation* variable is an empty string.

Description

The **no** form of the command resets the value of the *sysLocation* variable to an empty string.

snmpwalk

Displays contents of the a router's SNMP MIB.

Syntax

```
snmpwalk <name> <oid>
```

Arguments

- *name*—Name of the SNMP agent, as configured on CNA by the **polling snmp-agent** command
- *oid*—identifier of the MIB object where the walk is to begin

Prerequisites

Must be in system operations mode, with either login or privileged level access.

Description

Use the **snmpwalk** command to view the contents of a MIB on a router.

---

source-route

Enables Loose Source Routing (LSR) for obtaining measurement data.

Syntax

```
source-route active-measurement
no source-route active-measurement
```

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

Disabled

Description

The **source-route active-measurement** command enables or disables Loose Source Routing (LSR) for active measurement probes.

Before enabling Loose Source Routing, you must be certain that all routers on all routes to all active-probe targets will accept and process source-routed IP datagrams. (Web-surfer initiated User Traffic Tests are not affected by this command.)

Use the **no** form of the command to revert to the default condition.

---

speed

Set the speed of an interface.

Syntax

```
speed [ 10 | 100 ]
no speed
```

Prerequisites

Must be in interface configuration mode (*config-if*), which requires privileged-level access.

Occurrence

Use once per interface.

Default

Auto

Description

The **speed** command with no keyword sets the speed to **auto**.

Both speed and duplex of an interface can be configured, but they are dependent on each other. If any one of them is set to **auto**, then both will be implemented as **auto** regardless of the setting of the other.

Changing either speed or duplex will bring the interface down and back up and you will lose Telnet or SSH connectivity. Changes to these modes should be done from a serial port connection.

See the **duplex** command.

---

ssl add

Installs Secure Sockets Layer certificate and private key files on an CNA module.

Syntax

```
ssl add <url> [ encrypted ] [ ustat <name> ]
```

Arguments

- *url*—FTP or HTTP location of the files to be added (see **copy** command for URL syntax)
- *slot*—chassis slot of the USTAT module for which the certificate and key are valid
- *name*—name of the USTAT module for which the certificate and key are valid

Prerequisites

Must have privileged-level access.

Description

Each web server needs its own SSL certificate and private key files, as well as any required intermediate files. If your Certificate Authority (such as VeriSign or Thawte, for example) provided you with one or more Intermediate CAs when you obtained your certificate and key, then you must install all of them on CNA.

Use the **encrypted** keyword if the key is encrypted. CNA will prompt you for the pass phrase required to decrypt the key.

If you are adding certificate files for SSL-enabled USTAT modules, you can use either the *slot* parameter to specify module number, or the keyword **ustat** followed by the *name* of the ustat. Each USTAT-SSL requires the installation of its own certificate and key. See SSL Certificate/ Key Files in Chapter 3: Security of the *Monitoring Guide*.

If you are enabling SSL for the CNA web interface, do not specify a module number.

The certificate file should be named *server.crt* and any intermediate files should be named *chainca.crt*.

To add files one at a time, specify the filename as the *url* argument. This file must be in must be in Base64-encoded format.

Use this method if you have already added a *.key* file using the **`ssl generate-csr`** command.

If you have not yet installed the *.key* file, you must install it with the *server.crt* file. Package them together in *.tar*, *.tar.gz*, *.tgz*, or *.zip* format and specify that file as the *url* argument.

You can also include one or more *chainca.crt* files if desired (may be included in any order).

This will replace any existing SSL *server.key* and *server.crt* files.

Note:

The *server.key* and *server.crt* files must be issued together by your CA. If they don't match, they cannot work together. Additionally, the certificate file must be named *server.crt*.

Each of the SSL files must meet the follow requirements:

- *server.crt*—X.509 certificate in PEM (Base64-encoded DER) format
- *server.key*—the private RSA key in PEM (Base64-encoded DER) format; for general browser compatibility, the key size should be 1024 bits or less
- *chainca.crt*—intermediate X.509 certificate in PEM (Base64-encoded DER) format

If the common name (CN) in the server certificate is a fully qualified domain name, not an IP address, your USTAT-SSL VIPs should be reachable through DNS.

After installing the files on the module, you must enable, or re-enable, SSL with the **`ip http server ssl on`** command. (If the web server is already running, the command will restart it.)

You must restart the web server on each module, one at a time.

However, each server can be restarted with a single command, using the following shorthand notation (assuming USTAT modules named one, two and three):

```
CNA# config t module ustat one ip http server ssl on
CNA# config t module ustat two ip http server ssl on
CNA# config t module ustat three ip http server ssl on
```

ssl delete

Remove SSL certificates and key files from CNA.

Syntax

```
ssl delete [ ustat <name> ]
```

Arguments

- *name*—name of the USTAT module for which the certificate and key are valid

Prerequisites

Must have privileged-level access.

Description

You must disable SSL before CNA will allow the files to be deleted. See the **`ip http server ssl on`** command for information about disabling SSL.

Use the *slot* parameter to specify module number if you are removing files from a USTAT module.

If you are removing certification files for the web UI, do not specify a slot number.

Removal of the certification files does not automatically disable SSL. You must explicitly disable SSL with the **no** form of the **`ip http server ssl on`** command.

ssl export

Sends text of installed SSL files to a file URL or the screen.

Syntax

```
ssl export [ ustat <name> ] <url>
```

Arguments

- *name*—name of the USTAT module for which the certificate and key are valid
- *url*—destination file

Prerequisites

Must have privileged-level access.

Description

The **`ssl export`** command sends the contents of the *.crt* and *.key* files that are installed on CNA to a remote location. See the **`copy`** command for valid URL forms.

When you enter the command on the command line, CNA will prompt for a passkey. This passkey will be used to lock the file; you will need this passkey in order to decrypt the exported file.

If you do not enter a passkey and just press the **Enter** (or **Return**) key at each prompt, the *.key* file will not be encrypted in the exported file, but will be exported in clear text.

Use the *slot* parameter to specify module number if you are exporting files from a USTAT module.

If you are exporting files for the web UI, do not specify a slot number.

Files can be exported in any of the following formats, depending on the file type named in the URL. It can be one of the following:

- *.tar*
- *.tar.gz*
- *.tgz*
- *.zip*

If you don't specify a filename in the URL, the exported file will be named *server.tgz*.

The URL must end in either a filename with one of the allowed file extensions or a directory name with terminating slash character.

If a file with the same name as the one you have specified exists, CNA will overwrite it unless the FTP server prevents it.

ssl generate-csr

Generates a Certificate Signing Request for obtaining an SSL certificate.

Syntax

```
ssl generate-csr [ ustat <name> ] [ with-certificate ] [ > <url> ]
```

Arguments

- *name*—name of the USTAT module for which the certificate and key are valid
- *url*—remote destination for the request

Prerequisites

Must have privileged-level access.

Description

The **ssl generate-csr** command will create a Certificate Signing Request (in the form of a.*csr* file), which you can then send to a Certificate Authority (CA) in order to obtain the SSL certificate file (*server.crt*) required by the **ssl add** command.

When you submit a certificate request to your CA, you should specify that certificates be in the same format used by an Apache web server running SSLeay. Specifically, CNA does *not* support PKCS7 files.

The **ssl generate-csr** command also creates a *server.key* file and automatically installs it (unencrypted) into the specified module.

If you are generating a new CSR for a module that has already had SSL files installed, use the **ssl delete** command first. Otherwise, the new key will overwrite the existing key but the old certificate will still exist, and the module's web server will not function correctly.

When you enter the **ssl generate-csr** command on the command line, CNA will prompt you for the following information, which will be used by your CA to create the certificate (field names are defined in the X.509 standard):

- **Common Name (CN)**—domain name of the CNA module for which certificate files are sought; must be an exact match to the hostname portion of the URL used to reach the module (if a USTAT, use the domain name in the *single-pixel.gif* image tag, such as *www.one-pixel-gif.yourdomain.com*; if the management module, use the domain name you use to reach the web interface, such as *CNA.yourcompany.com*)
- **Organizational Unit (OU)**—your unit in your organization, such as department or business unit
- **Organization (O)**—your company
- **City/Locality (L)**—city in which CN is located
- **State/Province (ST)**—state or province in which L is located
- **Country (C)**—country in which ST is located
- **Contact e-mail address**—e-mail address associated with the web-server

Some of the fields may be optional; some may require a specific syntax. Abbreviations in names may be prohibited in some places and required in others (you may be required to enter Saint Louis instead of St. Louis, but US for United States and GB for Great Britain, for example). You should check with your CA before generating this *.csr* file.

Use the **with-certificate** keyword to also generate a self-signed crt file on CNA. See ip http server ssl on for more information on enabling SSL.

When you execute the **`ssl generate-csr`** command, the Certificate Signing Request file will be sent to the screen, or, if you use the redirection character >, the location specified by the command-line argument *url*. See the **`copy`** command for valid URL formats.

start-avoidance percent-below-max

Sets the low end of the threshold for load optimization.

Syntax

```
start-avoidance percent-below-max <value>
no start-avoidance percent-below-max
```

Arguments

- *value*—threshold, expressed as a percentage below maximum load; range: 0 through 100

Prerequisites

Must be in threshold configuration mode (*config-engine-threshold*), which requires privileged-level access.

Occurrence

Use once per threshold table.

Default

value—10

Description

The **`start-avoidance percent-below-max`** command sets the minimum threshold at which the probability of a link being removed from the winner set due to load optimization. See Determining the Winner Set in Chapter 2: WAN Cost and Load Optimization of the *Adaptive Path Controller Guide*.

The value set by the **`start-avoidance percent-below-max`** command is used for all levels in a load threshold table, unless you explicitly set a different percentage for a specific link

at a specific level of the table using the **start-avoidance percent-below-max** keyword of the link command.

See the **load threshold-table** command.

subnet-cleanup schedule

Sets the time of day when unreachable subnets will be pruned from CNA's database.

Syntax

```
subnet-cleanup schedule <hour> [<minute>]
no subnet cleanup schedule
```

Arguments

- *hour*—hour of the day; range: 0 through 23
- *minute*—minute of the hour; range: 0 through 59

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access

Default

0 hour, 0 minute (midnight)

Description

subnets which have been declared unreachable by BGP for a minimum of 24 hours are deleted from the live database maintained on the engine module by a garbage collection routine that runs once every day.

Use the **subnet-cleanup schedule** command to set the time of day when this occurs.

Use the **no** form of the command to revert to the default.

subnet

Associates a subnet with a group.

Syntax

```
subnet <address>
```

```
no subnet <address>
```

Arguments

- *addr*—subnet address in CIDR format (*a.b.c.d/n*)

Prerequisites

Must be in group configuration mode (*config-group*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

none

Description

Use the **subnet** command to associate subnets with a group. When in group configuration mode (see the **group** command), you can add one subnet at a time to group.

Use the **no** form of the command to remove a subnet from a group.

---

subnet-cleanup schedule

Sets the time of day when unreachable subnets will be pruned from CNA's database.

Syntax

```
subnet-cleanup schedule <hour> [<minute>]
no subnet cleanup schedule
```

Arguments

- *hour*—hour of the day; range: 0 through 23
- *minute*—minute of the hour; range: 0 through 59

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access

Default

0 hour, 0 minute (midnight)

Description

subnets which have been declared unreachable by BGP for a minimum of 24 hours are deleted from the live database maintained on the engine module by a garbage collection routine that runs once every day.

Use the **subnet-cleanup schedule** command to set the time of day when this occurs.

Use the **no** form of the command to revert to the default.

---

subnet resync

Causes the CNA *decision-maker* process to reevaluate specified subnets.

Syntax

```
subnet resync { all |
                confirmed |
                asserted |
                selected |
                adequate |
                measured |
                using-policy <name> |
                quiet <qtime> |
                <addr>
              }
              {
                compute-winner-set |
                clear |
                clear-routes-only |
                reassert-route |
                verify-auto-created
              }
              [ limit <num> ]
```

Arguments

- *name*—name of a decision policy
- *qtime*—quiet time (number of seconds since last measurement); range: 60 through 86400
- *addr*—subnet address in CIDR format (*a.b.c.d/n*)
- *num*—maximum number of subnets to re-evaluate; range: 1 through 1000000

Prerequisites

Must be in system operations mode with privileged level access.

Default

num—60

Description

The **subnet resync** command is the configuration (periodic reevaluation of subnets) form of the command, **resync subnet,** which results in immediate reevaluation of subnets.

Use one of the following keywords with either version of the command:

- **all**—all subnets
- **confirmed**—all subnets that have been asserted by CNA and confirmed as in use by the routers
- **asserted**—all subnets that have been asserted (includes those that have been confirmed)
- **selected**—all subnets that have been selected as the winning route by CNA (includes those that have been asserted)
- **adequate**—all subnets that have a sufficient number of measurements in the CNA database to be considered adequately measured
- **measured**—all subnets that have at least one measurement in the database
- **using-policy**—all subnets associated with the named decision policy
- **quiet**—all subnets for which no measurements have been received for the specified period of time

Or, you can specify an individual subnet by its address in CIDR format (a.b.c.d/n).

You can impose a limit on the number of subnets that will be evaluated regardless of the size of the set, by setting the *limit* argument.

The *limit* argument can be used with both **subnet resync**, the *config* form of this command, and **resync subnet**, the *exec* form. But its greatest value is with the config form, when it can be used to avoid thrashing (nonstop reassertions of routes) if your interval is small and the set large. In some cases, it could take at least as long as the interval to re-evaluate all of the subnets in the set, causing the *decision-maker* process to immediately start a new **resync** cycle immediately upon completion of the previous evaluation.

You must also choose one of the following actions to be taken by the CNA *decision-maker* process:

- **compute-winner-set**—the *decision-maker* process will recompute the winner-set for the specified set of subnets and reassert the current winner if it is in the winner set
- **clear**—the *decision-maker* process will withdraw any CNA-asserted routes for the specified subnets; this is the equivalent of you executing the **clear prefix (deprecated)** command for the specified subnets
- **clear-routes-only**—the *decision-maker* process will withdraw any CNA-asserted routes for the specified subnets without resetting the counters which CNA uses to keep track of user-traffic test hits; this is the equivalent of you executing the **clear prefix (deprecated)** command for the specified subnets with **routes-only** keyword
- **reassert-route**—the *decision-maker* process will reassert the route that was most recently selected as the winning route by the normal subnet evaluation process
- **verify-auto-created**—the *decision-maker* process will immediately run a validation check on automatically created subnets (see auto-create-subnets verify-frequency)

Use the keyword **limit** to restrict subnets that will be re-evaluated to the specified number that have gone the longest time without an evaluation.

The following example forces an immediate recomputation of winner set of the 50 asserted-and-confirmed subnets that have gone the longest without a *decision-maker* evaluation:

```
subnet resync confirmed compute-winner-set limit 50
```

subnet-update-rate

Sets the rate at which subnet update requests are taken out of the queue by CNA.

Syntax

```
subnet-update-rate <rate> [ <priority_rate> ] <unit>
no subnet-update-rate
```

Arguments

- *rate*—representing number of updates; range depends upon *unit*
  - *per-second*: 0 through 100
  - *per-minute*: 0 through 6000
  - *per-hour*: 0 through 360000
  - *per-day*: 0 through 8640000

- *priority_rate*—representing number of priority-subnet updates. It must be equal to or greater than *rate.* If you choose not to set this value, it will be set to *rate*; range: same as *rate*
- *unit*—interval; can be one of the following strings:
  - *per-second*
  - *per-minute*
  - *per-hour*
  - *per-day*

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Default

rate—10 per second

priority_rate—same as rate, if you do not set this value

unit—per-second

Description

Whenever CNA determines that it is necessary to assert a change in routes, it places an update request into a queue. (Changes in response to critical conditions—loss of coverage for a particular subnet, or link failure—bypass this queue and are dealt with immediately.)

The update rate determines how fast update requests are processed out of the queue.

The priority rate determines how fast updates for priority subnets are processed. See the **`priority-subnets`** command.

In general, the faster the rate of updates the faster the system can react to general performance deteriorations. The actual rate should be based on your subnet's traffic patterns, link bandwidths, and number of subnets. The more resources your subnet has, the faster you can set the rate.

Setting the rate to 0 causes CNA to stop asserting routes, an alternative to using the **`route-assert-mode`** command to shift temporarily to reporting mode (which has the side effect of clearing the CNA subnet database and measurement counters). You should not rely on a rate of 0 for long-term suspension of route assertions, since it will prevent CNA from responding to brownout or blackout conditions.

With an update rate of 0, no decisions will be allowed out of the decision-maker's queue, thereby preventing assertions; however, route withdrawals may still occur, if, for example, **assert-routes-matching-bgp** behavior is disabled or all natural BGP links are in the winner set.

The **no** form of the **`subnet-update-rate`** command resets update rates to their defaults.

supernode

Configures a single server in a hive to be designated as the supernode. Configure the supernode in a shared configuration. Configuring a second server as a supernode will overwrite the previous supernode configuration.

Syntax

configure terminal supernode <address>

Arguments

Table 56: supernode arguments

Argument	Description
<*address*>	IP address of the server to be designated as the supernode. Enter the IP address in the format **1.2.3.4**

Prerequisites

Must be in configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrences

Use once.

Description

Configures a single server in a hive to be designated as the supernode. Only one supernode can be configured in a hive. Configuring a second server as a supernode will overwrite the previous supernode configuration.

This is a share command and within a few minutes the configuration will appear on all servers, and all hive members will know this address as the new supernode. This command is propagated through chatter.

By default, when servers are initially started they are automatically configured as supernodes. When the server joins a hive, the server must accept the existing supernode in the hive.

sysuptime

Sets the unit for sysuptime used by a particular router.

Syntax

```
sysuptime unit <unit>
no sysuptime unit
```

Arguments

- *unit*—unit portion of sysuptime; must be one of the following:
  - *seconds*
  - *tenths-of-seconds*
  - *hundredths-of-seconds*
  - *milliseconds*

Prerequisites

Must be in polling agent configuration mode (*config-polling-snmp)*, which requires privileged-level access.

Occurrence

Once per SNMP-agent

Default

unit—hundredths-of-seconds

Description

The **`sysuptime unit`** command specifies the units used by a router which is being polled for link-usage data.

tacacs-server host

Defines the TACACS+ server to be used by CNA for user authentication.

Syntax

```
tacacs-server host <address> [ key <string> ]
no tacacs-server host <address>
```

Arguments

- *address*—IP address in dotted decimal format of the TACACS+ server, or host name
- *string*—string containing TACACS+ server's encryption key

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

The key applies only to the TACACS+ server currently being defined. If *string* is not specified when the TACACS+ server is defined, the global key defined by the **`tacacs-server key`** command is used.

CNA connects to the TACACS+ server on port 49.

The **no** form of the command removes the specified TACACS+ server from the configuration.

---

tacacs-server key

Defines a global encryption key for communicating with TACACS+ servers.

Syntax

```
tacacs-server key <string>
no tacacs-server key
```

Arguments

- *string*—string containing TACACS+ server's encryption key

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Description

The key defined by **`tacacs-server key`** command will be used when communicating with TACACS+ servers for which no encryption key was specified; see the **`tacacs-server host`** command.

---

target

Defines an active-measurement target.

Syntax

```
target <address>
     [ <type> ]
     [ [ <rate> <unit> [ <port> ] ]
     [ measure-type <mtype> ]
     [ timeout <value> ]
     [ description <text> ]

no target <address> <type>
```

Arguments

- *address*—hostname or IP address of the host to be probed
- *type*—protocol to be used; can be *tcp*, *icmp*, or *traceprobe*
- *rate*—value portion of probe frequency; positive, non-zero integer
- *unit*—unit portion of probe frequency; must be one of the following:
  - *per-second*
  - *per-minute*
  - *per-hour*
  - *per-day*
- *port*—port to be contacted on host
- *mtype*—how CNA should parcel out measurements among USTATs; must be one of the following:
  - *simultaneous*
  - *round-robin*
- *value*—number of milliseconds to wait for a response; range: 0 through 780000 (default: 3000)
- *text*—text of the description; can consist of any combination of alphanumerical characters, dash, underscore and space

Prerequisites

Must be in active-measurement configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

May be used multiple times within an active-measurement group configuration block.

Description

Active measurements represent one of the means by which CNA obtains subnet traffic data. The system actively establishes contact with a remote host in order to obtain measurement data by opening and immediately closing a TCP session on the specified port.

This is in contrast with the User Traffic Test (UTT), where CNA waits for a web surfer to initiate contact and then measures the time it takes to respond to the user's request for an image.

The values you set for the *rate* and *unit* parameters of the **`active-measurement target`** command determine the frequency of the active probes.

Note:

If you are configuring a TCP probe, a *port* value must be specified either for the individual probe or the group, or the probe will fail. The output of the **`show statistics`** command for the *am-scheduler* process includes a configuration errors table which will note the lack of a port assignment.

The timeout *value* sets the number of seconds CNA will wait for a response from the target. If *value* is set to 0, CNA will not timeout on its own (though the underlying kernel will timeout in approximately 13 minutes). The timeout value is used by the **`outage-detection active-probe`** command in determining whether a subnet is experiencing an outage.

If you set *mtype* to *simultaneous*, CNA will contact the target through all service provider links simultaneously at each scheduled probe time. Round-trip traffic measurements can be compared directly and the fastest route observed empirically.

If you set *mtype* to *round-robin*, CNA will contact the target through a single service provider link at each scheduled probe time, proceeding through all links in round-robin fashion. This method places less demand on the target host, but it also produces fewer measurement points, and the measurements do not all represent the same moment in time.

For example, the following command, entered while in engine configuration mode, would create a probe using ICMP aimed at a target whose IP address is 192.168.121.10:

```
CNA (config-engine)# active-measurement target 192.168.121.10 icmp 2 per-second
```

Default values were accepted for measurement type and timeout (ICMP probes don't take a port assignment). This will cause CNA to send a probe to the target through every configured link twice every second (*simultaneous* measurement type).

See also the **`active-measurement group`** command.

Use the **no** form of the command to delete an active measurement probe.

target (Chatter module only)

Configures the test to be conducted from CNA test agents in the agent-based measurement subnet to an external IP address.

Syntax

target <*ip-address*>

no target

Arguments

Table 57: target arguments

Argument	Description
<*ip-address*>	IP address of the end point to which you want the test conducted. Enter the IP address in the format **a.b.c.d**

Prerequisites

Must be in Chatter Active Measurement Configuration command mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrences

Use once.

Default

0.0.0.0, which conducts the test between CNA agents in the hive.

Description

Configures the test to be conducted from CNA agents in the agent-based measurement subnet to an external IP address.

Only one instance of a test is conducted at any time. If the external target is specified and this test was already active between CNA agents in the Chatter subnet, the active test will be discontinued. The test is reinitiated with all CNA agents in the Chatter subnet conducting the test on the specified external target.

The no command restores the default setting, 0.0.0.0, which conducts the test between CNA agents in the hive.

Note:

To respond to TCP tests, the device at the specified IP address must be running TCP on the port to which the test is sent. The default port is 80. To change this port number, use the port (Chatter module only) command.

Only CNA agents can serve as targets for RTP tests. No other devices support RTP tests.

tcpdump

Displays headers from packets arriving on the specified interface.

Syntax

```
tcpdump [ count <packets> ]
        [ interface <name> ]
        [ snarf <bytes> ]
        [ numeric ]
        [ absolute ]
        [ hexadecimal ]
        [ quiet | verbose | more-verbose ]
        [ no-timestamp ]
        [ timeout <minutes> ]
        [ raw-packets ]
        [ module <slot> ]
        [ <expression> ]
        [ | { include | exclude } <regexp> ]
        [ > <url> ]
```

Arguments

- *packets*—number of packets to count before exiting
- *name*—name of the interface; can be one of the following:
  - *ethernet0*
  - *tunnelN* (where N is number assigned to the tunnel in the configuration)
  - *lo*
  - *lo:0*
  - lo:1
- *bytes*—number of bytes in each packet to be analyzed (equivalent to *snaplen*)
- *minutes*—number of minutes to wait
- *slot*—chassis position of target module
- *expression*—filter for which packets to be dumped

- *regexp*—regular expression
- *url*—URL of remote file where output is to be sent (cannot be local to CNA)

Prerequisites

System operations mode, login-level access.

Occurrence

Use once.

Default

count—1 when run on a peripheral module in local mode, otherwise none

name—name of lowest numbered interface

bytes—68

timeout—2 minutes

If no interface is specified, default is *Ethernet0*

Time stamp is included

Description

The **`tcpdump`** command is available both as a system-wide command and a local command.

When you are connected to the system through the management module's ethernet port, the **`tcpdump`** command will work on any module that you specify with the *slot* argument.

When you are connected to a peripheral module's console port, the command operates as a local command on the module to which you are physically connected. Not all arguments are supported in local mode.

All but one of the listed arguments are optional (*expression* is required).

Order is important—if you specify multiple arguments, they must appear in the order shown:

- Use the **count** keyword to specify the number of packets to count before exiting. When you use this keyword, be aware that it may take some time for the command to complete (or timeout). During this time, there are no on-screen activity indicators.
- Use the **interface** keyword to specify an CNA interface (ethernet interfaces will be placed in promiscuous mode and all traffic, not just those packets that are addressed to the interface, will be included in the output).
- Use the **snarf** keyword to specify the number of bytes in a packet.
- Use the **numeric** keyword to keep addresses in numeric form, instead of converting them to symbolic names (not available in local mode).
- Use the **absolute** keyword to display absolute (as opposed to relative) TCP sequence numbers.

- Use the **hexadecimal** keyword to display the output in hexadecimal format.
- Use the **quiet**, **verbose**, or **more-verbose** keyword to determine how much data is output.
- Use the **no-timestamp** keyword to suppress timestamps in the output.
- Use the **timeout** keyword to specify the number of minutes to wait for a reply (not available in local mode).
- Use the **raw-packets** keyword with the > character and a file URL to divert raw packets to a file (not available in local mode). If you are specifying a module number or an expression, these arguments must follow the **raw-packets** keyword but the URL must be last. Additionally, the URL must point to a remote location; you cannot direct data to a file local to CNA. Packets written to the file follow the regular **tcpdump** binary format compatible with the **-w** and **-r** options defined by the **tcpdump** utility.
- Use the **module** keyword to specify an CNA slot number.
- The *expression* argument should follow regular **tcpdump** expression semantics. It can consist of one or more primitives (such as an ID name or number), preceded by one or more qualifiers. The entire expression (primitive and qualifier) must be enclosed in double quotes. Qualifiers can be one of the following:
  - type—kind of object identified by primitive (possible types include *host*, or *net*, or *port*); if none specified, *host* is assumed
  - dir—direction of transfer; to or from the primitive; can be *src*, *dst*, *src or dst*, or *src and dst*, for example
  - proto—protocol, such as *ether*, *fddi*, *ip*, *arp*, *tcp*, or *udp*, for example; if none specified, all protocols consistent with *type* are assumed

  Each qualifier/primitive pair must be separated by *and, or, not*:

```
tcpdump raw-packets “host 1.2.3.4 and port 80”
                                        > ftp://mydomain.com/myfile
```

- **include**—shows all lines in the output which contain a match to the specified regular expression
- **exclude**—shows all lines in the output except those which contain a match to the specified regular expression

If *url* is set to a valid FTP, HTTP or HTTPS URL (with the required pipe character, >), output will be sent to the specified destination. See the **copy** command for URL syntax.

See *http://www.tcpdump.org* for more information about the **tcpdump** utility.

terminal

Sets output parameters. Parameter terminal hardwrap to set whether the CLI tool will add a hard linefeed character when wrapping lines. Default is OFF.

Syntax

```
no terminal [ length <num_rows> ] [ width <num_cols> ] no terminal [ hardwrap ]
no terminal [ status ]
```

Arguments

- *num_rows*—number of lines to be displayed
- *num_cols*—number of characters per row

Description

Use the **length** argument to set the number of lines to be displayed per screen.

Use the **width** argument to set the number of characters to be included in each line.

Use the **status** keyword to see current settings.

threshold host-disappeared

Sets the number of probe failures to a host that will cause it to be considered unreachable.

Syntax

```
threshold host-disappeared [ count ]
no threshold host-disappeared
```

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once.

Default

3

Description

Use the **threshold host-disappeared** command to set the number of probe failures from all USTATs to a particular host that will cause CNA to consider that host to be unreachable.

Use the **no** form of the command to reset the value to its default.

threshold host-link-fail

Sets the number of probe failures over a single link to a host that will cause the host to be considered unreachable.

Syntax

```
threshold host-link-fail [ count ]
no threshold host-link-fail
```

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once.

Default

3

Description

Use the **threshold host-link-fail** command to set the number of probe failures over a particular link to a host that will cause CNA to consider that host to be unreachable.

Use the **no** form of the command to reset the value to its default.

threshold ping

Defines the alarm thresholds for Ping tests on any cell or edge that is assigned this alarm template.

Syntax

threshold ping (rtt <milliseconds>) | (loss <percent>) | (inter-arrival-time <milliseconds>)

no threshold ping {rtt | loss | inter-arrival-time}

Arguments

Table 58: threshold ping arguments

Argument	Description
rtt <milliseconds>	The round trip time (RTT) threshold in milliseconds for Ping tests. The valid range is 0 to 780,000
loss <percent>	The loss threshold as a percentage of loss for Ping tests. The valid range is 0 to 100. The loss threshold is a percentage of packets lost in transit.
inter-arrival-time <milliseconds>	The inter-arrival-time threshold in milliseconds for Ping tests. The valid range is 0 to 2147483647. Inter arrival time is the time between tests, and can be used to trigger an alarm when a link is down or a set of endpoints becomes isolated.

Prerequisites

Must be in Chatter Alarm Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use once.

Default

Table 59: Default settings for threshold ping command

Argument	Default value
RTT threshold	300 milliseconds
loss threshold	1 percent
inter arrival time threshold	864 seconds

Description

Defines the alarm thresholds for Ping tests on any cell or edge that is assigned this alarm template. Thresholds are defined for each variable measured by Chatter for the ping test.

These threshold values are absolute thresholds and used to generate an alarm if absolute alarms are enabled. These threshold values are also used for coloring the absolute topology.

The no command restores the default value for the specified threshold.

---

threshold proxy-disappeared

Sets the number of probe failures to a convergence point that will cause it to be considered unreachable.

Syntax

```
threshold proxy-disappeared [ count ]
no threshold proxy-disappeared
```

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once.

Default

3

Description

Use the **`threshold proxy-disappeared`** command to set the number of probe failures to a convergence point that will cause CNA to consider that host to be unreachable.

Use the **no** form of the command to reset the value to its default.

threshold proxy-link-fail

Sets the number of probe failures over a single link to a convergence point that will cause it to be considered unreachable.

Syntax

```
threshold proxy-link-fail [ count ]
no threshold proxy-link-fail
```

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once.

Default

3

Description

Use the **`threshold proxy-link-fail`** command to set the number of probe failures set the number of probe failures over a particular link to a convergence point that will cause CNA to consider that host to be unreachable.

Use the **no** form of the command to reset the value to its default.

threshold rtp

Defines the alarm thresholds for RTP tests on any cell or edge that is assigned this alarm template.

Syntax

threshold rtp {(delay <milliseconds>) | (jitter <milliseconds>) (loss <percent>) | (inter-arrival-time <milliseconds>)}

no threshold rtp {delay | jitter | loss | inter-arrival-time}

Arguments

Table 60: threshold rtp arguments

Argument	Description
delay <milliseconds>	The delay threshold in milliseconds for RTP tests. The valid range is 0 to 780,000. Delay is how long (in milliseconds one-way) it takes the UDP packet stream to reach its destination.
jitter <milliseconds>	The jitter threshold in milliseconds for RTP tests. The valid range is 0 to 780,000 Jitter is the unevenness in packet delay (in milliseconds).
loss <percent>	The loss threshold for RTP tests. The valid range is 0 to 100. The loss threshold is a percentage of packets lost in transit.
inter-arrival-time <milliseconds>	The inter-arrival-time threshold in milliseconds for RTP tests. The valid range is 0 to 2147483647. Inter arrival time is the time between tests, and can be used to trigger an alarm when a link is down or a set of endpoints becomes isolated.

Prerequisites

Must be in Chatter Alarm Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use once.

Default

Table 61: Default settings for threshold-rtp command

Argument	Default value
delay threshold	300 milliseconds
jitter threshold	30 milliseconds
	1 of 2

Table 61: Default settings for threshold-rtp command (continued)

Argument	Default value
loss threshold	1 percent
inter arrival time threshold	864 seconds
	2 of 2

Description

Defines the alarm thresholds for RTP tests on any cell or edge that is assigned this alarm template. Thresholds are defined for each variable measured by Chatter for the RTP test.

These threshold values are absolute thresholds and used to generate an alarm if absolute alarms are enabled. The threshold values are also used for coloring the absolute topology.

The no command restores the default value for the specified threshold.

threshold tcp

Defines the alarm thresholds for TCP tests on any cell or edge that is assigned this alarm template.

Syntax

threshold tcp {(delay <milliseconds>) | (cerror <percent>) | (inter-arrival-time <milliseconds>)}

no threshold tcp delay | cerror | inter-arrival-time

Arguments

Table 62: threshold tcp arguments

Argument	Description
delay <milliseconds>	The delay threshold in milliseconds for TCP tests. The valid range is 0 to 780,000 Delay is how long (in milliseconds) the connection takes to complete.
	1 of 2

Table 62: threshold tcp arguments (continued)

Argument	Description
cerror <percent>	The cerror threshold as a percentage of cerror for TCP tests. The valid range is 0 to 100. Cerror is the success or failure of the connection.
inter-arrival-time <milliseconds>	The inter-arrival-time threshold in milliseconds for TCP tests. The valid range is 0 to 2147483647. Inter arrival time is the time between tests, and can be used to trigger an alarm when a link is down or a set of endpoints becomes isolated.

2 of 2

Prerequisites

Must be in Chatter Alarm Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use once.

Default

Table 63: Default settings for threshold tcp command

Argument	Default value
delay threshold	300 milliseconds
cerror threshold	1 percent
inter arrival time threshold	864 seconds

Description

Defines the alarm thresholds for TCP tests on any cell or edge that is assigned this alarm template. Thresholds are defined for each variable measured by Chatter for the ping test.

These threshold values are absolute thresholds and used to generate an alarm if absolute alarms are enabled. The threshold values are also used for coloring the absolute topology.

The no command restores the default value for the specified threshold.

threshold traceroute

Defines the alarm thresholds for Traceroute tests on any cell or edge that is assigned this alarm template.

Syntax

threshold traceroute {(hop-rtt <milliseconds>) | (inter-arrival-time <milliseconds>)}

no threshold traceroute {hop-rtt | inter-arrival-time}

Arguments

Table 64: threshold traceroute arguments

Argument	Description
hop-rtt <milliseconds>	The hop round trip time (RTT) threshold in milliseconds for traceroute tests. The valid range is 0 to 780,000. Hop RTT is the round trip time for each hop along the route (UDP outbound and ICMP inbound).
inter-arrival-time <milliseconds>	The inter-arrival-time threshold in milliseconds for traceroute tests. The valid range is 0 to 2147483647. Inter arrival time is the time between tests, and can be used to trigger an alarm when a link is down or a set of endpoints becomes isolated.

Prerequisites

Must be in Chatter Alarm Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use once.

Default

Table 65: Default settings for threshold traceroute command

Argument	Default value
Hop RTT threshold	300 milliseconds
inter arrival time threshold	864 seconds

Description

Defines the alarm thresholds for Traceroute tests on any cell or edge that is assigned this alarm template. Thresholds are defined for each variable measured by Chatter for the traceroute test.

These threshold values are absolute thresholds and used to generate an alarm if absolute alarms are enabled. These threshold values are also used for coloring the absolute topology.

The no command restores the default value for the specified threshold.

time

Displays or sets the system time.

Syntax

```
time [<value>]
```

Arguments

- *value*—time to set, in HHMMSS format (HH is a 24-hour value)

Prerequisites

System operations mode; requires login-level access to display, privileged-level access to set.

Description

All peripheral modules in CNA take their NTP feed from the Management Processor module.

There must always be six digits in *value*. Do not use colons, but do use leading zeros for single-digit hours, minutes or seconds.

You cannot reset the time while NTP is enabled.

Resetting the system time to an earlier date/time may result in a gap in your snapshot reports as well as polling by the link-monitor process.

The snapshot reports problem is self-correcting; snapshots will resume when the system time and the snapshot schedule become synchronized again.

In order to resume polling immediately, you should restart the *link-monitor* process, using the **`medic call`** command.

CNA must be restarted before time changes will be implemented by some database-related CNA processes.

Changing the system time while active probes are enabled can disrupt active probing. See Troubleshooting Agent-Based Measurement in Chapter 4: Server-Based Measurement of the *Monitoring Guide*.

timeout

Sets the amount of time CNA will wait for a response from a remote device.

Syntax

```
timeout <time>
no timeout
```

Arguments

- *time*—amount of time to wait, as follows:
  - for SNMP, specify time in seconds; range: 1 through 300
  - for active-measurement probes, specify time in milliseconds; range: 0 through 780000

Prerequisites

Must be in one of the following:

- SNMP polling configuration mode (*config-polling-snmp*)
- Active measurement group configuration mode (*config-engine-active*)

Both require privileged-level access.

Default

1 second for SNMP polling

3000 ms for active-measurement probes

Description

You can use the **timeout** command in either of the following configuration modes:

- SNMP polling agent (*config-polling-snmp*)—determines how long the CNA link monitor will wait for an edge router to respond when polling for usage data
- Active measurement group (*config-engine-active*)—determines how long CNA will wait for active probe targets to respond (regardless of probe type—*tcp* or *icmp*)

See the **polling snmp-agent** command for information about configuring the link monitor for SNMP polling.

See the **active-measurement group** command for information about configuring active-measurement groups. The group wide **timeout** command will apply to targets that don't have their individual **timeout** arguments set.

Use the **no** form of the command to reset the value of *time* to its default.

timeout (Chatter module only)

Defines the time (in milliseconds) at which this active measurement test times out if no response is received from the remote device.

Syntax

timeout <value>

no timeout

Arguments

Table 66: timeout arguments

Argument	Description
<*value*>	Number of milliseconds at which the test times out if no response is received from the remote device. Valid ranges are: • RTP test: 100–5000 • Ping test: 1–1000 • TCP test: 1–1000 • Traceroute test: 1000–20000

Prerequisites

Must be in Chatter Active Measurement Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use once.

Default

Table 67: Default settings for timeout command

Type of test	Default timeout value (in milliseconds)
RTP	500
Ping	500
TCP	500
Traceroute	5000

Description

Defines the time (in milliseconds) at which this active measurement test times out if no response is received from the remote device.

The no command restores the default value.

timezone

Sets the time zone to be used when displaying time on CNA.

Syntax

```
timezone <name>
no timezone
```

Arguments

- *name*—name of the timezone

Prerequisites

System operations mode; requires privileged-level access.

Occurrence

Use once.

Default

UTC (Universal, or Greenwich, Time)

Description

Use the **timezone** command to set the time zone to be used by CNA when displaying dates. This time zone will be used for all CLI output and for all times displayed in the management web user interface unless over-ridden by user selection in the web UI session.

You must reboot CNA for timezone changes to take effect, which means it must be present in the startup configuration. If the **timezone** command is present only in the running configuration, CNA will remain at its default upon reboot.

Identify the zone to be used by setting *name* equal to an entry in the Time Zone column in Table 68. Use the **no** form of the command to reset the timezone back to the default.

Table 68: Time Zones

Country Code	Time Zone	Description
AD	Europe/Andorra	
AE	Asia/Dubai	
AF	Asia/Kabul	
AG	America/Antigua	
AI	America/Anguilla	
AL	Europe/Tirane	
AM	Asia/Yerevan	
AN	America/Curacao	
AO	Africa/Luanda	
AQ	Antarctica/Casey	Casey Station, Bailey Peninsula
AQ	Antarctica/Davis	Davis Station, Vestfold Hills
		1 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
AQ	Antarctica/ DumontDUrville	Dumont-d'Urville Base, Terre Adelie
AQ	Antarctica/Mawson	Mawson Station, Holme Bay
AQ	Antarctica/McMurdo	McMurdo Station, Ross Island
AQ	Antarctica/Palmer	Palmer Station, Anvers Island
AQ	Antarctica/South_Pole	Amundsen-Scott Station, South Pole
AQ	Antarctica/Syowa	Syowa Station, E Ongul I
AR	America/Buenos_Aires	E Argentina (BA, DF, SC, TF)
AR	America/Catamarca	Catamarca (CT)
AR	America/Cordoba	W Argentina (CB, SA, TM, LR, SJ, SL, NQ, RN)
AR	America/Jujuy	Jujuy (JY)
AR	America/Mendoza	Mendoza (MZ)
AR	America/Rosario	NE Argentina (SF, ER, CN, MN, CC, FM, LP, CH)
AS	Pacific/Pago_Pago	
AT	Europe/Vienna	
AU	Australia/Adelaide	South Australia
AU	Australia/Brisbane	Queensland—most locations
AU	Australia/Broken_Hill	New South Wales—Broken Hill
AU	Australia/Darwin	Northern Territory
AU	Australia/Hobart	Tasmania
AU	Australia/Lindeman	Queensland—Holiday Islands
		2 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
AU	Australia/Lord_Howe	Lord Howe Island
AU	Australia/Melbourne	Victoria
AU	Australia/Perth	Western Australia
AU	Australia/Sydney	New South Wales—most locations
AW	America/Aruba	
AZ	Asia/Baku	
BA	Europe/Sarajevo	
BB	America/Barbados	
BD	Asia/Dacca	
BE	Europe/Brussels	
BF	Africa/Ouagadougou	
BG	Europe/Sofia	
BH	Asia/Bahrain	
BI	Africa/Bujumbura	
BJ	Africa/Porto-Novo	
BM	Atlantic/Bermuda	
BN	Asia/Brunei	
BO	America/La_Paz	
BR	America/Araguaina	Tocantins
BR	America/Belem	Amapa, E Para
BR	America/Boa_Vista	Roraima
BR	America/Cuiaba	Mato Grosso, Mato Grosso do Sul
		3 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
BR	America/Fortaleza	NE Brazil (MA, PI, CE, RN, PR, PE)
BR	America/Maceio	Alagoas, Sergipe
BR	America/Manaus	Amazonas
BR	America/Noronha	Fernando de Noronha
BR	America/Porto_Acre	Acre
BR	America/Porto_Velho	W Para, Rondonia
BR	America/Sao_Paulo	S and SE Brazil (BA, GO, DF, MG, ES, RJ, SP, PR, SC, RS)
BS	America/Nassau	
BT	Asia/Thimbu	
BW	Africa/Gaborone	
BY	Europe/Minsk	
BZ	America/Belize	
CA	America/Cambridge_Bay	Central Time—west Nunavut
CA	America/Dawson	Pacific Time—north Yukon
CA	America/Dawson_Creek	Mountain Standard Time—Dawson Creek & Fort Saint John, British Columbia
CA	America/Edmonton	Mountain Time—Alberta, east British Columbia & west Saskatchewan
CA	America/Glace_Bay	Atlantic Time—Nova Scotia—places that did not observe DST 1966-1971
CA	America/Goose_Bay	Atlantic Time—E Labrador
CA	America/Halifax	Atlantic Time—Nova Scotia (most places), NB, W Labrador, E Quebec & Prince Edward Island
		4 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
CA	America/Inuvik	Mountain Time—west Northwest Territories
CA	America/Iqaluit	Central Time—east Nunavut
CA	America/Montreal	Eastern Time—Ontario & Quebec—most locations
CA	America/Nipigon	Eastern Time—Ontario & Quebec—places that did not observe DST 1967-1973
CA	America/Pangnirtung	Central Time—Pangnirtung, Nunavut
CA	America/Rainy_River	Central Time—Rainy River & Fort Frances, Ontario
CA	America/Rankin_Inlet	Central Time—central Nunavut
CA	America/Regina	Central Standard Time—Saskatchewan—most locations
CA	America/St_Johns	Newfoundland Island
CA	America/Swift_Current	Central Standard Time—Saskatchewan—midwest
CA	America/Thunder_Bay	Eastern Time—Thunder Bay, Ontario
CA	America/Vancouver	Pacific Time—west British Columbia
CA	America/Whitehorse	Pacific Time—south Yukon
CA	America/Winnipeg	Central Time—Manitoba & west Ontario
CA	America/Yellowknife	Mountain Time—central Northwest Territories
CC	Indian/Cocos	
CD	Africa/Kinshasa	west Dem. Rep. of Congo
CD	Africa/Lubumbashi	east Dem. Rep. of Congo

5 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
CF	Africa/Bangui	
CG	Africa/Brazzaville	
CH	Europe/Zurich	
CI	Africa/Abidjan	
CK	Pacific/Rarotonga	
CL	America/Santiago	mainland
CL	Pacific/Easter	Easter Island
CM	Africa/Douala	
CN	Asia/Chungking	China mountains
CN	Asia/Harbin	north Manchuria
CN	Asia/Kashgar	Eastern Turkestan
CN	Asia/Shanghai	China coast
CN	Asia/Urumqi	Tibet & Xinjiang
CO	America/Bogota	
CR	America/Costa_Rica	
CU	America/Havana	
CV	Atlantic/Cape_Verde	
CX	Indian/Christmas	
CY	Asia/Nicosia	
CZ	Europe/Prague	
DE	Europe/Berlin	
DJ	Africa/Djibouti	
		6 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
DK	Europe/Copenhagen	
DM	America/Dominica	
DO	America/Santo_Domingo	
DZ	Africa/Algiers	
EC	America/Guayaquil	mainland
EC	Pacific/Galapagos	Galapagos Islands
EE	Europe/Tallinn	
EG	Africa/Cairo	
EH	Africa/El_Aaiun	
ER	Africa/Asmera	
ES	Africa/Ceuta	Ceuta & Melilla
ES	Atlantic/Canary	Canary Islands
ES	Europe/Madrid	mainland
ET	Africa/Addis_Ababa	
FI	Europe/Helsinki	
FJ	Pacific/Fiji	
FK	Atlantic/Stanley	
FM	Pacific/Kosrae	Kosrae
FM	Pacific/Ponape	Ponape (Pohnpei)
FM	Pacific/Truk	Truk (Chuuk)
FM	Pacific/Yap	Yap
FO	Atlantic/Faeroe	
		7 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
FR	Europe/Paris	
GA	Africa/Libreville	
GB	Europe/Belfast	Northern Ireland
GB	Europe/London	Great Britain
GD	America/Grenada	
GE	Asia/Tbilisi	
GF	America/Cayenne	
GH	Africa/Accra	
GI	Europe/Gibraltar	
GL	America/Godthab	southwest Greenland
GL	America/Scoresbysund	east Greenland
GL	America/Thule	northwest Greenland
GM	Africa/Banjul	
GN	Africa/Conakry	
GP	America/Guadeloupe	
GQ	Africa/Malabo	
GR	Europe/Athens	
GS	Atlantic/South_Georgia	
GT	America/Guatemala	
GU	Pacific/Guam	
GW	Africa/Bissau	
GY	America/Guyana	
		8 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
HK	Asia/Hong_Kong	
HN	America/Tegucigalpa	
HR	Europe/Zagreb	
HT	America/Port-au-Prince	
HU	Europe/Budapest	
ID	Asia/Jakarta	Java & Sumatra
ID	Asia/Jayapura	Irian Jaya & the Moluccas
ID	Asia/Ujung_Pandang	Borneo & Celebes
IE	Europe/Dublin	
IL	Asia/Jerusalem	
IN	Asia/Calcutta	
IO	Indian/Chagos	
IQ	Asia/Baghdad	
IR	Asia/Tehran	
IS	Atlantic/Reykjavik	
IT	Europe/Rome	
JM	America/Jamaica	
JO	Asia/Amman	
JP	Asia/Tokyo	
KE	Africa/Nairobi	
KG	Asia/Bishkek	
KH	Asia/Phnom_Penh	
		9 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
KI	Pacific/Enderbury	Phoenix Islands
KI	Pacific/Kiritimati	Line Islands
KI	Pacific/Tarawa	Gilbert Islands
KM	Indian/Comoro	
KN	America/St_Kitts	
KP	Asia/Pyongyang	
KR	Asia/Seoul	
KW	Asia/Kuwait	
KY	America/Cayman	
KZ	Asia/Almaty	east Kazakhstan
KZ	Asia/Aqtau	west Kazakhstan
KZ	Asia/Aqtobe	central Kazakhstan
LA	Asia/Vientiane	
LB	Asia/Beirut	
LC	America/St_Lucia	
LI	Europe/Vaduz	
LK	Asia/Colombo	
LR	Africa/Monrovia	
LS	Africa/Maseru	
LT	Europe/Vilnius	
LU	Europe/Luxembourg	
LV	Europe/Riga	
		10 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
LY	Africa/Tripoli	
MA	Africa/Casablanca	
MC	Europe/Monaco	
MD	Europe/Chisinau	most locations
MD	Europe/Tiraspol	Transdniestria
MG	Indian/Antananarivo	
MH	Pacific/Kwajalein	Kwajalein
MH	Pacific/Majuro	most locations
MK	Europe/Skopje	
ML	Africa/Bamako	southwest Mali
ML	Africa/Timbuktu	northeast Mali
MM	Asia/Rangoon	
MN	Asia/Hovd	Bayan-Olgiy, Hovd, Uvs
MN	Asia/Ulaanbaatar	most locations
MO	Asia/Macao	
MP	Pacific/Saipan	
MQ	America/Martinique	
MR	Africa/Nouakchott	
MS	America/Montserrat	
MT	Europe/Malta	
MU	Indian/Mauritius	
MV	Indian/Maldives	
		11 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
MW	Africa/Blantyre	
MX	America/Cancun	Central Time—Quintana Roo
MX	America/Chihuahua	Mountain Time—Chihuahua
MX	America/Hermosillo	Mountain Standard Time—Sonora
MX	America/Mazatlan	Mountain Time—S Baja, Nayarit, Sinaloa
MX	America/Mexico_City	Central Time—most locations
MX	America/Tijuana	Pacific Time
MY	Asia/Kuala_Lumpur	peninsular Malaysia
MY	Asia/Kuching	Sabah & Sarawak
MZ	Africa/Maputo	
NA	Africa/Windhoek	
NC	Pacific/Noumea	
NE	Africa/Niamey	
NF	Pacific/Norfolk	
NG	Africa/Lagos	
NI	America/Managua	
NL	Europe/Amsterdam	
NO	Europe/Oslo	
NP	Asia/Katmandu	
NR	Pacific/Nauru	
NU	Pacific/Niue	
NZ	Pacific/Auckland	most locations
		12 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
NZ	Pacific/Chatham	Chatham Islands
OM	Asia/Muscat	
PA	America/Panama	
PE	America/Lima	
PF	Pacific/Gambier	Gambier Islands
PF	Pacific/Marquesas	Marquesas Islands
PF	Pacific/Tahiti	Society Islands
PG	Pacific/Port_Moresby	
PH	Asia/Manila	
PK	Asia/Karachi	
PL	Europe/Warsaw	
PM	America/Miquelon	
PN	Pacific/Pitcairn	
PR	America/Puerto_Rico	
PS	Asia/Gaza	
PT	Atlantic/Azores	Azores
PT	Atlantic/Madeira	Madeira Islands
PT	Europe/Lisbon	mainland
PW	Pacific/Palau	
PY	America/Asuncion	
QA	Asia/Qatar	
RE	Indian/Reunion	

13 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
RO	Europe/Bucharest	
RU	Asia/Anadyr	Moscow+10—Bering Sea
RU	Asia/Irkutsk	Moscow+05—Lake Baikal
RU	Asia/Kamchatka	Moscow+09—Kamchatka
RU	Asia/Krasnoyarsk	Moscow+04—Yenisei River
RU	Asia/Magadan	Moscow+08—Magadan & Sakhalin
RU	Asia/Novosibirsk	Moscow+03—Novosibirsk
RU	Asia/Omsk	Moscow+03—west Siberia
RU	Asia/Vladivostok	Moscow+07—Amur River
RU	Asia/Yakutsk	Moscow+06—Lena River
RU	Asia/Yekaterinburg	Moscow+02—Urals
RU	Europe/Kaliningrad	Moscow-01—Kaliningrad
RU	Europe/Moscow	Moscow+00—west Russia
RU	Europe/Samara	Moscow+01—Caspian Sea
RW	Africa/Kigali	
SA	Asia/Riyadh	
SB	Pacific/Guadalcanal	
SC	Indian/Mahe	
SD	Africa/Khartoum	
SE	Europe/Stockholm	
SG	Asia/Singapore	
SH	Atlantic/St_Helena	
		14 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
SI	Europe/Ljubljana	
SJ	Arctic/Longyearbyen	Svalbard
SJ	Atlantic/Jan_Mayen	Jan Mayen
SK	Europe/Bratislava	
SL	Africa/Freetown	
SM	Europe/San_Marino	
SN	Africa/Dakar	
SO	Africa/Mogadishu	
SR	America/Paramaribo	
ST	Africa/Sao_Tome	
SV	America/El_Salvador	
SY	Asia/Damascus	
SZ	Africa/Mbabane	
TC	America/Grand_Turk	
TD	Africa/Ndjamena	
TF	Indian/Kerguelen	
TG	Africa/Lome	
TH	Asia/Bangkok	
TJ	Asia/Dushanbe	
TK	Pacific/Fakaofo	
TM	Asia/Ashkhabad	
TN	Africa/Tunis	
		15 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
TO	Pacific/Tongatapu	
TP	Asia/Dili	
TR	Europe/Istanbul	
TT	America/Port_of_Spain	
TV	Pacific/Funafuti	
TW	Asia/Taipei	
TZ	Africa/Dar_es_Salaam	
UA	Europe/Kiev	most locations
UA	Europe/Simferopol	central Crimea
UA	Europe/Uzhgorod	Ruthenia
UA	Europe/Zaporozhye	Zaporozh'ye, E Lugansk
UG	Africa/Kampala	
UM	Pacific/Johnston	Johnston Atoll
UM	Pacific/Midway	Midway Islands
UM	Pacific/Wake	Wake Island
US	America/Adak	Aleutian Islands
US	America/Anchorage	Alaska Time
US	America/Boise	Mountain Time—south Idaho & east Oregon
US	America/Chicago	Central Time
US	America/Denver	Mountain Time
US	America/Detroit	Eastern Time—Michigan—most locations
		16 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
US	America/Indiana/Knox	Eastern Standard Time—Indiana—Starke County
US	America/Indiana/Marengo	Eastern Standard Time—Indiana—Crawford County
US	America/Indiana/Vevay	Eastern Standard Time—Indiana—Switzerland County
US	America/Indianapolis	Eastern Standard Time—Indiana—most locations
US	America/Juneau	Alaska Time—Alaska panhandle
US	America/Los_Angeles	Pacific Time
US	America/Louisville	Eastern Time—Louisville, Kentucky
US	America/Menominee	Central Time—Michigan—Wisconsin border
US	America/New_York	Eastern Time
US	America/Nome	Alaska Time—west Alaska
US	America/Phoenix	Mountain Standard Time—Arizona
US	America/Shiprock	Mountain Time—Navajo
US	America/Yakutat	Alaska Time—Alaska panhandle neck
US	Pacific/Honolulu	Hawaii
UY	America/Montevideo	
UZ	Asia/Samarkand	west Uzbekistan
UZ	Asia/Tashkent	east Uzbekistan
VA	Europe/Vatican	
VC	America/St_Vincent	

17 of 18

Table 68: Time Zones (continued)

Country Code	Time Zone	Description
VE	America/Caracas	
VG	America/Tortola	
VI	America/St_Thomas	
VN	Asia/Saigon	
VU	Pacific/Efate	
WF	Pacific/Wallis	
WS	Pacific/Apia	
YE	Asia/Aden	
YT	Indian/Mayotte	
YU	Europe/Belgrade	
ZA	Africa/Johannesburg	
ZM	Africa/Lusaka	
ZW	Africa/Harare	
		18 of 18

topological-analysis disable

Disables or enables topological analysis for this type of active measurement test.

Topological analysis issues a traceroute between the CNA agents to discover intermediate hops and devices. This enables Chatter to identify potential hops where subnet disruption is present.

Syntax

topological-analysis disable

no topological-analysis disable

Prerequisites

Must be in Chatter Active Measurement Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use once.

Default

Enabled.

Description

Disables or enables topological analysis for this type of active measurement test.

Topological analysis issues a traceroute between the CNA agents to discover intermediate hops and devices. This traceroute makes it possible for Chatter to identify potential hops where subnet disruption is present.

---

tos byte

Defines the type of service (tos) byte to be used when sending active probes. It is assigned on a per packet basis and influences the packet priority along the subnet path on devices with quality of service (QoS) implemented.

Syntax

```
tos byte <value>
no tos
```

Arguments

- *value*—byte value (0 through 255)

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once per active-measurement group.

Default

0

Description

Use the **tos** command while in active-measurement group configuration mode (*config-engine-active*) to set the value of the type of service field in the IP header of the probes to be sent to targets in that group.

See the **active-measurement group** command for an overview of active-measurement groups.

Use the **no** form of the command to revert to the default value.

---

tos byte (Chatter module only)

Defines the type-of-service (TOS) byte to use when performing this active measurement test.

Syntax

tos byte <value>

no tos byte

Arguments

Table 69: tos byte arguments

Argument	Description
<*value*>	TOS value for this test. The valid range is 0–255.

Prerequisites

Must be in Chatter Active Measurement Configuration mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Use once.

Default

TOS value of 0.

Description

Defines the type-of-service (TOS) byte to use when performing the test. Use this command to mark agent-based measurement traffic.

Define a TOS byte for agent-based measurement traffic if you want CNA to measure the performance of one (or more) specific TOS markings. Applicable for RTP, Ping, and Traceroute tests. This command is not applicable for TCP tests.

The no command restores the default value of 0.

---

traceprobe convergence-rule

Sets the convergence algorithm for traceprobes.

Syntax

```
traceprobe convergence-rule [ normal | strict ]
no traceprobe convergence-rule
```

Prerequisites

Must be in engine configuration mode (*config-engine*), which requires privileged-level access.

Occurrence

Use once.

Default

normal

Description

If you have more than two links, you should use `normal`.

If your subnet has two links, use `strict` along with **traceprobe hops min** and **traceprobe avoid** (to ensure that the traceprobe system cannot converge on any host that is within your subnet; otherwise, the system may not be able to detect a link failure). Note: future changes to your subnet configuration may require you to update parameters for **traceprobe hops min** and **traceprobe avoid**.

Use the **no** form of the command to revert to the default value.

traceprobe refresh

Sets a time interval for updating the **traceroute** data for a **traceprobe**.

Syntax

```
traceprobe refresh <time>
no traceprobe refresh
```

Arguments

- *time*—interval between **`traceroute`** commands, in seconds; range: 0 through 31,556,736

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once per active-measurement group

Default

21600 seconds

Description

The **`traceprobe refresh`** command sets the interval between invocations of the **traceprobe retries-per-hop** process on behalf of a *traceprobe* type of active measurement probe.

Measurement data for a probe of type *traceprobe* is derived from the replies received from a **traceroute** packet. When you configure a probe of this type, the CNA am-*scheduler* process sends a **traceroute** packet to the probe target address. The HRTT that goes into the database is based on packet transit times between CNA and the closest address reached by the **traceroute** packet. (The *traceprobe* type of probe is designed to provide measurements for targets that are unreachable by *tcp* or *icmp* probes. See Chapter 4: Server-Based Measurement of the *Monitoring Guide*.)

But the **`traceroute`** command is costly—it can take several minutes to complete one probe. The CNA software caches the address of the last endpoint reached and begins probing that address directly. This is not reliable over an extended period of time, however, so the **traceroute** process must be invoked periodically, at the interval specified by the *time* parameter.

Use the **no** form of the command to reset the value of *time* to its default.

See also the **`active-measurement traceprobe avoid address (discontinued)`** command.

traceprobe retries-per-hop

Sets the retry count for each hop of the **traceroute**. The total number of attempts will be one plus the retry count.

Syntax

```
traceprobe retries-per-hop <count>
no traceprobe retries-per-hop
```

Arguments

- *count*—the maximum number of retries to attempt for each hop of the **traceroute**

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once per active-measurement group

Default

0 retries

Description

The **`traceprobe retries-per-hop`** command sets the retry count for each hop of the **traceroute**. A retry is attempted whenever a probe fails to reach a hop, but not when an *unreachable* (e.g., *host unreachable*) message is received. The total number of attempts will be one (the initial probe) plus the retry count.

Use the **no** form of the command to reset the value of *count* to its default.

traceprobe max-failing-hops

The **traceprobe** will abort if it fails to receive a response from this number of consecutive hops.

Syntax

```
traceprobe max-failing-hops <count>
no traceprobe max-failing-hops
```

Arguments

- *count*—the number of failed hops which will cause the **traceprobe to abort.**

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once per active-measurement group

Default

3 failed hops

traceprobe hops

Sets the minimum and maximum hop count for the **traceroute**.

Syntax

```
traceprobe hops min <min> max <max>
no traceprobe hops
```

Arguments

- *min*— 1 through 255 the minimum number of hops for a destination to be a valid target of a **traceroute**
- *max*—1 through 255 the maximum number of hops allowed by a **traceroute** to a destination

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once per active-measurement group

Default

min = 2

max = 31

Description

The **`traceprobe retries-per-hop`** command sets the minimum and maximum hop count for the **traceroute**. The minimum number of hops determines the distance, in hops, from that **traceroute** will start its probe (for example, the default value of 2 typically excludes the border router between and the Internet). The maximum number of hops limits the distance, in hops, from that **traceroute** will probe.

Use the **no** form of the command to reset the value of *count* to its default.

---

traceroute

Displays the route which packets travel between CNA and their destination.

Syntax

```
traceroute [ do-not-fragment ]
           [ first-ttl <first_seconds> ]
           [ source { <interface> | <address> } ]
           [ icmp-echo ]
           [ max-ttl <max_hops> ]
           [ numeric ]
           [ base-udp-port] <port> ]
           [ bypass-routing-table ]
           [ tos <TOS> ]
           [ verbose ]
           [ wait <wait_seconds> ]
           [ packet-size <bytes> ]
           [ module <module> ]
           [ link <name> ]
           <host>
```

Arguments

- *first_seconds*—initial time-to-live, in seconds
- *interface*— interface to obtain the source ip address for the outgoing probe packet
- *address*—source IP address
- *max_hops*—maximum number of hops, 1 through 100
- *port*—UDP port number
- *TOS*—integer, 0 through 255 (type of service bits)
- *wait_seconds*—number of seconds to wait for response
- *bytes*—size of probe packet, in bytes
- *module*—module number
- *name*—perform the traceroute from the named link
- *host*—address of **`traceroute`** command target

Prerequisites

System operations mode, any level access.

Defaults

- *max_hops*—30
- *port*—33434
- *TOS*—0
- *wait_seconds*—5
- *bytes*—40

Description

Order of arguments is important. If multiple arguments are used, they must follow in the order listed here:

- **do-not-fragment**—do not break up large packets into smaller units
- **first-ttl**—initial time-to-live value for the first outgoing probe packet
- **source**—identifies the address or CNA interface to be used as the source address in outgoing probe packets
- **icmp-echo**—use ICMP ECHO instead of UDP datagrams
- **max-ttl**—maximum number of hops (time-to-live) for outgoing probe packets
- **numeric**—print hop addresses numerically, rather than both numerically and symbolically
- **base-udp-port**—set base UDP port number

- **bypass-routing-table**—send directly to a host on an attached subnet
- **tos**—type of service in probe packets
- **verbose**—verbose output
- The *bytes* value of the **packet-size** keyword specifies the length of a probe's datagram

traceroute-vips

Performs a simultaneous traceroute with the source as each configured vip on CNA. See traceroute.

Syntax

```
traceroute [ continuous ]
           [ icmp-echo ]
           [ max-ttl <max_hops> ]
           [ numeric ]
           [ base-udp-port] <port> ]
           [ wait <wait_seconds> ]
           <host>
```

Arguments

- *max_hops*—maximum number of hops, 1 through 100
- *port*—UDP port number
- *wait_seconds*—number of seconds to wait for a response
- *host*—address of **traceroute** target

Prerequisites

System operations mode, any level access.

Defaults

- *ports*—33434
- *wait_seconds*—5

Description

Order of arguments is important. If multiple arguments are used, they must follow in the order listed here:

- **continuous**—will force the system to do traceroute-vips until you terminate it with control-c

- **icmp-echo**—use ICMP ECHO instead of UDP datagrams
- **max-ttl**—maximum number of hops (time-to-live) for outgoing probe packets
- **numeric**—print hop addresses numerically, rather than both numerically and symbolically
- **base-udp-port**—set base UDP port number
- **wait**—number of seconds to wait for a response

tree

Displays a hierarchical tree of commands.

Syntax

```
tree
```

Description

The tree that is displayed will show all of the command-line keywords, in nested hierarchical form, that are available to you at the time you execute the **`tree`** command.

trending active-measurement group (discontinued)

Enables gathering of trending data for subnets in an active-measurement group.

Syntax

```
trending active-measurement group <name>
no trending active-measurement group [ <name> ]
```

Arguments

- *name*—name of the group

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

If you want trending data to be gathered for subnets that are targeted by active probes, identify the group in which they are defined as the *name* argument of the **`trending active-measurement group`** command.

Collection of trending data cannot be enabled for individual probes not defined in a group (unless they qualify for trending for some other reason. See the *CNA Reporting Guide*).

Trending data will be gathered for all subnets in the group.

Use the **no** form of the command with the *name* argument to disable collection of data for subnets that are in a specific group, or without an argument to disable trending of all subnets that are being trended because they are in an active-measurement group.

If a subnet that is included in an active-measurement group qualifies for trend-data collection for some other reason, the data will continue to be collected.

trending clear-after (deprecated)

trending decision-policy

Enables gathering of trending data for subnets belonging to a decision policy.

Syntax

```
trending decision-policy <name>
no trending decision-policy <name>
```

Arguments

- *name*—name of a decision policy

Prerequisites

Must be in report configuration mode (*config-reporting*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **`trending decision-policy`** command if you want trending data to be gathered for subnets that belong to a particular decision policy.

Use the **no** form of the command with the *name* argument to disable collection of data for subnets in that decision policy.

trending frequency (discontinued)

Sets the frequency for gathering trend data for specified data types.

Syntax

```
trending frequency <family> <type> <interval>
no trending frequency <family> <type>
```

Arguments

- *family*—trending family; must be *performance*
- *type*—type of data; must be *summary*
- *interval*—time in seconds; range: 30 through 300

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

interval—120 seconds

Description

The **`trending frequency`** command sets the time interval for gathering data for selected data types.

Currently, the command will only accept a user-defined interval for *live-summary* data for *performance* family types.

Use the **no** form of the command to restore an interval to its default.

trending on

Enables or disables data collection for trend reports.

Syntax

```
trending on
no trending on
```

Prerequisites

Must be in report configuration mode (*config-report*), which requires privileged-level access.

Occurrence

Use once.

Default

Data collection for trend reports is on.

Description

Use the **no** form of the command to turn trending data collection off (does not delete data that has been collected).

See the **`clear trending`** command.

This process may take a few minutes to complete. As a safety mechanism to prevent reporting errors, the process will timeout after 90 minutes.

trending prefix-limit (deprecated)

The trending prefix-limit command has been deprecated. See the replacement command, trending subnet-limit on page 465.

trending subnet-limit

Sets the number of subnets for which trending data will be extracted from the snapshot database.

Syntax

```
trending subnet-limit <number>
no trending subnet-limit
```

Arguments

- *number*—number of subnets, range: 1 through 1000

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Use once.

Default

500

Description

When trending is enabled, CNA will extract data from the periodic snapshot image of the database for a certain number of the most active subnets (those that have generated the most number of measurements), plus all named subnets and all active probe targets.

Use the **`trending subnet-limit`** command to determine the number of most-active subnets.

Use the **no** form of the command to reset the limit to the default.

trending size (deprecated)

The trending size command has been deprecated. This command is longer need to be used as the trending sizes is automatically determined by CNA.

tunnel

This command has been removed starting from version 3.0. Use the **`interface`** command to create a GRE tunnel.

type

Sets type of probe in an active-measurement group.

Syntax

```
type [ tcp | icmp | traceprobe ]
no type
```

Prerequisites

Must be in active-measurement-group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once per active-measurement group.

Default

none

Description

While in active-measurement group configuration mode, use the **type** command to set the type of probes that will be launched to targets.

See the **active-measurement group** command.

See also the **contact** command, for information about identification data that is included in all *traceprobe* type packets.

usage contract (removed)

Specifies the contracted usage rate for a service provider access link.

Syntax

```
usage contract <bandwidth> <percentile>
no usage contract
```

Arguments

- *bandwidth*—contracted bandwidth, in kilobits per second
- *percentile*—usage percentile for billing

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Occurrence

Use once per link.

Description

This command was discontinued in version 2.0 and will be removed from the CLI in a subsequent release. Use the following commands, instead:

- **usage-reporting billing-period**
- **usage-reporting cost**
- **usage-reporting percentile**
- **usage-reporting formula**

---

usage-reporting billing-period

Configures the CNA link with the service provider's billing cycle.

Syntax

```
usage billing-period <day> [ <hour> ]
no usage billing-period
```

Arguments

- *day*—day of month on which billing cycle begins; range: 01 through 31
- *hour*—hour of day when cycle begins; range 00 through 23

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Occurrence

Use once per link.

Default

day—01

hour—00 (midnight)

Description

Use the **usage-reporting billing-period** command to change the beginning of a service provider's billing cycle from its default of day 01, hour 00 (midnight before the first day of the month).

The billing cycle is the period used to generate billing reports from link-usage data obtained by polling your edge routers. See the command **polling usage on**.

The **usage-reporting billing-period** command is one of several related usage parameter commands:

- **usage-reporting cost**
- **usage-reporting percentile**
- **usage-reporting formula**

Use the **no** form of the command to delete the billing period from the link configuration block.

---

usage-reporting cost

Configures the CNA link with the price charged and bandwidth allowed by the service provider.

Syntax

```
usage-reporting cost <bandwidth>
                     <dollars>
                     [
                       per-mbps |
                       per-mbps-ceiling |
                       per-mbps-round |
                       per-mbps-truncate
                     ]
                     [ offset <offset-dollars> ]
no usage-reporting cost <bandwidth>
```

Arguments

- *bandwidth*—maximum dataflow permitted for this usage tier, in megabits per second
- *dollars*—contracted dollar amount for this usage tier (can be a fraction, to two decimal points)
- *offset-dollars*—amount to be added to total cost if usage enters this tier (can be a fraction, to two decimal points); number specified can be either positive or negative

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

The *bandwidth* and *dollars* values entered with the **usage-reporting cost** command are used by the CNA software to calculate monthly billing reports from link usage data gathered by the SNMP polling agent. See the **polling usage on** command.

Some contracts which may have a special bonus or penalty clause that changes the cost of service when a tier threshold is broken. To accommodate those contracts, the value specified by the optional *offset-dollars* argument will be added (or subtracted, if you configure a negative number) to the total cost if usage enters the tier in which the value is configured.

If the optional keyword **per-mbps** is used, the *dollars* amount will be considered the cost of one megabit per second.

Use the keyword **per-mbps-ceiling** to round all fractional values up to the next highest integer, **per-mbps-round** to round up if the fraction is .5 or greater or down if the fraction is .49 or less, and **per-mbps-truncate** to drop all fractional values.

If the keyword **per-mbps** or one of its variants is omitted, the *dollars* amount will be considered a fixed cost for the tier.

If you have a flat rate contract (any fixed-tier contract that ties you to one price regardless of bandwidth actually used), enter the **usage cost** command once.

If you have a burstable contract, enter the **usage-reporting cost** command multiple times, specifying the bandwidth and dollar amounts for each tier of your multi-tiered contract.

The **usage-reporting cost** command is one of several related usage parameter commands:

- **usage-reporting billing-period**
- **usage-reporting percentile**
- **usage-reporting formula**

Use the **no** form of the command to delete a cost tier from the link configuration block.

---

usage-reporting formula

Configures link with the formula used by the service provider to determine billable usage rate.

Syntax

```
usage formula <method> [ { sum-of-each-sample | sum-over-billing-period } ]
no usage formula
```

Arguments

- *method*—the formula used to determine billable usage rate; can be:
  - *percentile-inbound*—rate determined by Nth percentile of all incoming data samples
  - *percentile-outbound*—rate determined by Nth percentile of all outgoing data samples
  - *percentile-sum-inbound-outbound*—rate determined by Nth percentile of sum of all incoming and outgoing data samples
  - *sum-percentile-inbound-outbound*—rate based on sum of Nth percentile of all incoming data samples and Nth percentile of all outgoing data samples
  - *percentile-inbound-or-outbound*—rate based on Nth percentile of set of all incoming and outgoing data samples
  - *percentile-max-inbound-outbound*—rate based on Nth percentile of set of larger of incoming or outgoing samples
  - *max-avg- percentile-daily-in-avg- percentile-daily-out*—plots the monthly maximum of the following two averages: 1) the daily Nth percentile of all incoming data samples; 2) the daily Nth percentile of all outgoing data samples
  - *max-percentile-inbound-outbound*—rate based on either Nth percentile of all incoming data samples or Nth percentile of all outgoing data samples, whichever is larger
  - *percentile-avg-inbound-outbound*—rate based on Nth percentile of average of all incoming and outgoing data samples
  - *avg-percentile-inbound-outbound*—rate based on average of Nth percentile of all incoming data and average of Nth percentile of all outgoing data
  - *avg-max-inbound-outbound*—rate based on the average of the daily maximum of all incoming data and outgoing data

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Occurrence

Use once per link.

Default

percentile-outbound

Description

The method set by the **`usage-reporting formula`** command determines how CNA calculates the billing report that estimates the service provider's bill based on link usage data obtained from your edge router. See the **`polling usage on`** command.

The value of *N* in the method descriptions (*Nth percentile*) is determined by the **usage-reporting percentile** command.

See the discussion of formulas in the Chapter 3: Polling Routers of the *Adaptive Path Controller-Internet Guide*.

The **usage-reporting formula** command is one of several related usage parameter commands:

- **usage-reporting billing-period**
- **usage-reporting cost**
- **usage-reporting percentile**

Use the **sum-of-each-sample** or **sum-over-billing-period** keywords to combine multiple physical links into one billing link. The first will sum each sample from all the physical links; then, at the end of the billing period, the formula will be applied to this value to obtain the cost based on the billing tier. Conversely, with the **sum-over-billing-period** keyword, the formula will be applied to all the samples of each physical link over the billing period, and then summed to determine the total bandwidth usage for the billing link to find the corresponding billing tier to obtain the cost.

Use the **no** form of the command to revert to the default value.

usage-reporting percentile

Configures a link with the **usage-reporting percentile** value used by the service provider to determine billable bandwidth.

Syntax

```
usage percentile <number> [ plus-one ]
no usage percentile
```

Arguments

- *number*—percentile; range: 1 through 100

Prerequisites

Must be in link configuration mode (*config-engine-link*), which requires privileged-level access.

Occurrence

Use once per link.

Default

number—95

no **plus-one** (round down)

Description

The **usage-reporting percentile** command specifies the percentile used by the service provider to determine billable bandwidth. The CNA software will discard measurements above this percentile when calculating billing reports from link-usage data. See the command **polling usage on**.

Use the **plus-one** keyword with the usage percentile command to direct the CNA software to always round up to the next higher data point when rounding is required.

See the discussion of percentile in Chapter 3: Polling Routers of the *Adaptive Path Controller-Internet Guide*.

The **usage-reporting percentile** command is one of several related usage parameter commands:

- **usage-reporting billing-period**
- **usage-reporting cost**
- **usage-reporting formula**

Use the **no** form of the command to remove the percentile value from a link configuration block.

---

use link-group

Associates a link-group with an active-measurement group.

Syntax

```
use link-group <name>
no use link-group
```

Arguments

- *name*—name of the link group

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once per active-measurement group.

Description

The **use link-group** command associates a link group with the active measurement group.

Probes in the group will only be sent through links in the link group. See Using Link-Groups in Chapter 4: Server-Based Measurement of the *Monitoring Guide*.

Use the **link-group** command to create or edit a link group.

---

use named-vip

Associates probes within an active-measurement group to a specified named VIP.

Syntax

```
use { named-vip <name> | link-group <name> }
no use named-vip
```

Arguments

- *name*—name of the VIP to be used as the probe source

Prerequisites

Must be in active-measurement group configuration mode (*config-engine-active*), which requires privileged-level access.

Occurrence

Use once per active-measurement group.

Description

Use the **use named-vip** command to associate probes with a named VIP (virtual IP address) that has previously been created on each USTAT.

When probes are launched for the current active-measurement group, they will only be sent from USTATs on which there is a VIP with the specified name. See the **vip** command.

If the active measurement group was configured with both the **use named-vip** and **use link-group** commands, only those USTATs which meet both conditions will be used for the probes.

There is no comparable command for globally configured probes (not in an active measurement group).

Use the **no** form of the command to remove the **`use named-vip`** command from the active measurement group configuration.

See also the **`active-measurement group`** command.

user-defined-application

Allows the user to expand the list of applications recognized by the EFC module.

Syntax

```
user-defined-application name <name>
         [ efc-filter <access-list> endpoint { source | destination } ]
no user-defined-application name <name>
no user-defined-application name <name> efc-filter
no user-defined-application name <name> efc-filter <access-list>
```

Arguments

- *name*—name of the application you want to define; user-defined names must start with an uppercase letter and can be up to 32 characters in length. Predefined applications will always start with a lowercase letter.
- *access-list*—name or number of an access list of the application you want to define; user-defined names must start with an uppercase letter and can be up to 32 characters in length. Predefined applications will always start with a lowercase letter.

Prerequisites

Must be in EFC configuration mode (*config-efc*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

Use the **user-defined-application** command to extend the EFC module's capability to recognize applications beyond the list of predefined ones. After you create a new application, you will then be able to use the **efc-filter** command to add references to access list names or numbers which define the rules for identifying an application. Currently, these rules may be any combination of:

- protocol
- source ip address
- destination ip address
- source port number
- destination port number

Use the **no** form of the command with the **efc-filter** keyword to remove all of the filters from the application definition. If you specify this with the <access-list> then only that filter will be removed.

---

username

Defines a local login account, or changes a user's password.

Syntax

```
username <name> password [ encrypted ] <pwd > [ role dbaccess]
no username [ <name> ]
```

Arguments

- *name*—username
- *pwd*—password

Prerequisites

Must be in configuration mode (*config*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Default

No password.

Description

Use the **no** form of the command with an existing username to remove that username account.

Use the **no** form of the command without an argument to remove all local login accounts.

Use the **encrypted** keyword if the value of *pwd* has already been encrypted by CNA (copied from an existing configuration).

When you boot a new installation of CNA, no user account is needed to obtain a CLI prompt, although you can only use a console port connection. Once you create a local account, however, this userless access is no longer allowed. You must log in to obtain a command prompt. (For information about obtaining a command prompt without a password, see Password Recovery in Chapter 2: Configuration Fundamentals of the *Fundamentals Guide*.)

Local login accounts provide users with access to CNA even when you lose connectivity with your authentication servers.

Without a local login account, you will be locked out of CNA unless an authentication server is online.

You may have multiple user accounts.

If the *name* value specifies an existing user account, that account's password will be changed to the value of *pwd* currently being specified.

For example, create the *this_user* account with the password *abc*, as follows:

```
username this_user password abc
```

The password for the *this_user* account is now *abc*.

Now re-enter the user name command with the same user ID but a different password, *xyz*:

```
username this_user password xyz
```

The password for the *this_user* account is now *xyz*.

---

version

Set version of SNMP used by the SNMP agent on an edge router.

Syntax

```
version <value>
```

Arguments

- *value*—version of SNMP on the edge router; can be one of:
  - *1*
  - *2c*

Prerequisites

Must be in SNMP configuration mode (*config-polling-snmp*), which requires privileged-level access.

Occurrence

Use once per edge router SNMP agent.

Default

2c

Description

Use the **`version`** command to configure CNA to use the correct version of SNMP when the CNA link monitor communicates with the SNMP agent on the an edge router.

If you change the version to 1 and then reset the version to 2c, the SNMP version will be reset to the default value but the command will not appear in the configuration.

---

vip

Specifies a virtual IP address (VIP) for a USTAT module.

Syntax

```
vip [ name <name> | active-measurement ] <address> [ device loopback ]
no vip [ <address> | name <name> | active-measurement ]
```

Arguments

- *name*—name of VIP
- *address*—IP address, in dotted decimal form (a.b.c.d)

Prerequisites

Must be in USTAT configuration mode (*config-ustat*), which requires privileged-level access.

Occurrence

Available for use multiple times in the configuration.

Description

The **`vip`** command is used to designate an IP address for the USTAT web server.

You can have only one web-server VIP on a USTAT.

The *name* argument should be a unique string that meets the criteria described in the section Legal Characters in Chapter 1: Introduction to the Command Line Interface.

In addition to a web-server VIP, you can have one unnamed active-measurement VIP and several named active-measurement VIPs on the same USTAT:

- **Unnamed active-measurement VIPs**—Use the **active-measurement** keyword to create a VIP (virtual IP address) for the exclusive use of active probes, so that acknowledgement packets can bypass firewalls and load balancers which might otherwise drop them. Use an IP address that is reachable from the Internet but is unknown to your load balancer.
- **Named active-measurement VIPs**—Use the **name** keyword to create a named active-measurement VIP. Once you have created a named VIP on each of your USTATs, use the **`use named-vip`** command in an active measurement group to direct probes from that group to the appropriately named VIP. You can have multiple named VIPs on a single USTAT in order to selectively change the IP address used for different probe groups.

If neither a named nor an unnamed active-measurement VIP is defined, the USTAT will send active probes through the regular VIP.

VIP addresses are by default associated with the USTAT's Ethernet interface. In releases prior to 2.3, VIPs were associated with the loopback interface by default.

Because of the new implementation, the USTAT will now respond to an ARP request for the VIP address.

If you prefer the VIP to be associated with the USTAT's loopback interface as in earlier versions, include the **device loopback** option in the command.

Use the **no** form of the command with an address argument to remove a web-server VIP from a USTAT.

Use the **no** form of the command with the name argument (and its value) to remove a specified named VIP from a USTAT.

Use the **no** form of the command with the active-measurement keyword to remove an unnamed active-measurement VIP from a USTAT.

web-speedup http-turns

Define number of HTTP turns on an average web page (used to show time savings in web speedup snapshot report).

Syntax

```
web-speedup http-turns <min> <max>
no web-speedup
no web-speedup http-turns
```

Arguments

- *min*—minimum number of turns on a page; range: 0 through 10000
- *max*—maximum number of turns on a page; range: 0 through 10000

Prerequisites

Must be in report configuration mode (*config-reporting*), which requires privileged-level access.

Occurrence

Use once.

Default

min—20

max—150

Description

The **web-speedup http-turns** command specifies a range of HTTP turns, representing a page on your web server.

In order to download a complete web page, a client typically needs to contact the server several times, as it obtains graphics and other supporting files specified in the page definition. Each completed contact represents one HTTP turn.

These numbers will be used to calculate the *seconds saved* range of values shown on the *Web Speedup* snapshot report.

Use the **no web-speedup** command to reset all executive summary parameters to their defaults, including those set by each of the following commands:

- **web-speedup http-turns**
- **web-speedup improvable min-hrtt-saving**
- **web-speedup improvable min-percent-improvement**
- **web-speedup min-bgp-hrtt**
- **web-speedup worst-bgp-hrtt-percent**

Use the **no web-speedup http-turns** command to reset only the HTTP turn values to their defaults.

web-speedup improvable min-hrtt-saving

Specifies the minimum time savings required for a subnet to be considered improvable when generating the Web Speedup snapshot report.

Syntax

```
web-speedup improvable min-hrtt-saving <value>
no web-speedup
no web-speedup improvable min-hrtt-saving
```

Arguments

- *value*—minimum HRTT measurement (in milliseconds); range: 0 through 1000000

Prerequisites

Must be in report configuration mode (*config-report*), which requires privileged-level access.

Occurrence

Use once.

Default

value—0

Description

Use the **web-speedup improvable min-hrtt-saving** command to set the minimum time savings, in milliseconds, that must be achieved by switching from a BGP route to the CNA choice before a subnet is considered improvable in the web speedup snapshot report.

Time savings is one of two criteria used by the CNA report-generating process to determine if a subnet should be counted as improvable. The other is percentage of improvement. The subnet must meet conditions (minimum time savings and percentage speedup).

Use the **no web-speedup** command to reset all web speedup parameters to their defaults, including those set by each of the following commands:

- **web-speedup http-turns**
- **web-speedup improvable min-hrtt-saving**
- **web-speedup improvable min-percent-improvement**
- **web-speedup min-bgp-hrtt**
- **web-speedup worst-bgp-hrtt-percent**

Use the **no web-speedup http-turns** command to reset only minimum HRTT value to the default.

web-speedup improvable min-percent-improvement

Specifies the minimum percentage of speedup required for a subnet to be considered improvable when generating the Web Speedup snapshot report.

Syntax

```
web-speedup improvable min-percent-saving <value>
no web-speedup
no web-speedup improvable min-percent-saving
```

Arguments

- *value*—minimum percentage improvement required; range: 0 through 100

Prerequisites

Must be in report configuration mode (*config-reporting*), which requires privileged-level access.

Occurrence

Use once.

Default

value—10

Description

Use the **`web-speedup improvable min-percent-improvement`** command to set the minimum percentage of speedup that must be achieved by switching from a BGP route to the CNA choice before a subnet is considered improvable in the *Web Speedup* snapshot report. Percent improvement (time saved relative to final time) is calculated by subtracting the HRTT measurement across the CNA-chosen link from the HRTT measurement across the slowest BGP-selected link, dividing by the CNA-link's HRTT, and multiplying by 100:

$$\text{Percent Improvement} = \frac{\text{BGP HRTT} - \text{CNA HRTT}}{\text{CNA HRTT}} \times 100$$

If you have multiple edge routers, there will be multiple BGP-selected routes. The percent improvement shows how much better the CNA choice performs than the slowest of these BGP routes.

This percentage is one of two criteria used by the CNA report-generating process to determine if a subnet should be counted as improvable. The other is time savings. The subnet must meet both conditions (minimum time savings and percentage speedup).

Use the **no web-speedup** command to reset all web speedup parameters to their defaults, including those set by each of the following commands:

- **web-speedup http-turns**
- **web-speedup improvable min-hrtt-saving**
- **web-speedup improvable min-percent-improvement**
- **web-speedup min-bgp-hrtt**
- **web-speedup worst-bgp-hrtt-percent**

Use the **no web-speedup improvable min-percent-improvement** command to reset only the minimum percentage value to its default.

web-speedup min-bgp-hrtt

Defines one of the brownout conditions—speed of packet delivery—for the *Web Speedup* snapshot report.

Syntax

```
web-speedup min-bgp-hrtt <value_1> <value_2>
no web-speedup
no web-speedup min-bgp-hrtt
```

Arguments

- *value_1*—minimum HRTT value (milliseconds) required to be considered *poor*; range: 0 through 1000000
- *value_2*—minimum HRTT value required to be considered *very poor*; range: 0 through 1000000

Prerequisites

Must be in report configuration mode (*config-reporting*), which requires privileged-level access.

Occurrence

Use once.

Default

value_1—1000000

value_2—1000000

Description

Use the **web-speedup min-bgp-hrtt** command to set one of the two types of criteria used to define poor and very poor performance in the *Web Speedup* snapshot report. The other criteria is defined as "worst N percent of all subnets," set by the **web-speedup worst-bgp-hrtt-percent** command.

The minimum BGP values set the threshold for HRTTs (handshake round trip times, in milliseconds) on the worst performing BGP route in your subnet:

- If a subnet measurement is at least as high as *value_1* and it qualifies as one of the worst N subnets, it will be factored into the *Poor Performance* column of the Executive Summary snapshot report.
- If the subnet measurement is at least as high as *value_2* and it qualifies as one of the worst N subnets, it will also be factored into the *Very Poor Performance* column.

Use the **no web-speedup** command to reset all executive summary parameters to their defaults, including those set by each of the following commands:

- **web-speedup http-turns**
- **web-speedup improvable min-hrtt-saving**
- **web-speedup improvable min-percent-improvement**
- **web-speedup min-bgp-hrtt**
- **web-speedup worst-bgp-hrtt-percent**

Use the **no web-speedup min-bgp-hrtt** command to reset only the minimum BGP HRTT values to their defaults.

---

web-speedup worst-bgp-hrtt-percent

Defines one of the brownout conditions—worst N percent of all subnets—for the *Executive Summary* snapshot report.

Syntax

```
web-speedup worst-bgp-hrtt-percent <value_1> <value_2>
no web-speedup
no web-speedup worst-bgp-hrtt-percent
```

Arguments

- *value_1*—percentage of subnets considered *poor* based on HRTT measurements; range: 0 through 100
- *value_2*—percentage of subnets considered *very poor* based on HRTT measurements; range: 0 through 100

Prerequisites

Must be in report configuration mode (*config-reporting*), which requires privileged-level access.

Occurrence

Use once.

Default

value_1—20

value_2—3

Description

Use the **web-speedup worst-bgp-hrtt-percent** command to set one of the two criteria used to define poor and very poor performance in the *Executive Summary* snapshot report. The other piece of information is speed of packet delivery, set by the **exec-summary min-bgp-hrtt** command.

The values set by this command, plus the minimum HRTT threshold, determine which subnets are used to calculate the *Poor Performance* and *Very Poor Performance* columns of the Executive Summary report.

All subnets are sorted by HRTT score, and the worst *value_1* and *value_2* percent are included in the poor and very poor performers columns.

Use the **no web-speedup** command to reset all executive summary parameters to their defaults, including those set by each of the following commands:

- **web-speedup http-turns**
- **web-speedup improvable min-hrtt-saving**
- **web-speedup improvable min-percent-improvement**
- **web-speedup min-bgp-hrtt**
- **web-speedup worst-bgp-hrtt-percent**

Use the **no web-speedup worst-bgp-hrtt-percent** command to reset only the worst HRTT percentage value to its default.

winner-set-priority

Specifies a priority value for a service provider access link, to be used by the CNA decision maker process when narrowing the winner set in order to pick one link to assert as the best performing route to a subnet.

Syntax

```
winner-set-priority { prefer | avoid | <value> }
no winner-set-priority
```

Arguments

- *value*—priority value; integer in the range: –10 through 10

Prerequisites

Must be in link configuration mode (*config-engine-link* or *config-engine-policy-link*), which requires privileged-level access.

Occurrence

Use once per link configuration block.

Default

0

Description

After a winner set has been determined (see the **winner-set-width** command), the **winner-set-priority** value is applied to further narrow the field.

Values set globally, in *config-engine-link* mode, will apply to all links not explicitly set in *config-engine-policy-link* mode.

The keyword **prefer** is the equivalent of the highest possible priority value.

The keyword **avoid** is the equivalent of the lowest possible priority value.

Use the **no** form of the command to reset a link's priority to the default.

For information about how the CNA decision maker process selects the best performance route from the winner set, see Route Assertion in Chapter 3: Adaptive Path Control of the *Adaptive Path Controller Guide.*

---

winner-set-width

Sets a *winner-set* value, which establishes an acceptable range for preferred-route performance scores—to be in the winner-set (the set of service provider links from which the best performing route will be selected), a route's performance score must be within this range.

Syntax

```
winner-set-width <value> [ <val_2> ]
no winner-set-width
```

Arguments

- *value*—maximum allowable difference in metrics among links considered to be in the winner set; range: 1 through 1000
- *val_2*—maximum allowable difference in metrics among links considered to be in the winner set if winner set contains no links because of load optimization; range: 1 through 1000

Prerequisites

Must be in engine or policy configuration mode (*config-engine* or *config-engine-policy*), both of which require privileged-level access.

Occurrence

Use once for global value, and once per policy configuration.

Default

value—20

val_2—will be set to *value*

Description

All routes whose performance metrics are lower than the lowest performance score plus the **winner-set-width** *value* are considered to fall within the *winner-se*t. (Penalties assigned using the **`penalty`** command are factored into the performance score calculation.)

For information about how the CNA decision maker process selects the best performance route from the winner set, see Route Assertion in Chapter 3: Adaptive Path Control of the *Adaptive Path Controller Guide.*

Use the no form of the command to reset *value* to its default.

See also the **`winner-set-priority`** command.

---

write

Writes the current running configuration to either the hard disk (memory) or the screen (terminal).

Syntax

```
write { memory | terminal }
```

Prerequisites

System operations mode, privileged level access.

Description

When executed as **write memory**, the command copies the current running config as the startup config.

When executed as **write terminal**, the command displays the current running config on the screen.

---

zone

Creates a zone or subzone. Also enters Chatter Zone Configuration command mode.

Syntax

zone <*name*> [define]

no zone <*name*>

Arguments

Table 70: zone command arguments

Argument	Description
<*name*>	Fully qualified name of the zone or subzone that you are creating. The name consists of a zone and multiple subzones that are separated by a period (x.y.z format). Restrictions on the name are: • The name cannot be an integer. • The name must start with a letter (upper or lower case) and be followed by a letter (upper or lower case), digit, dash or underscore. • All zone and subzone names must be unique.
[define]	Define a block of information for this zone. Use [define] to assign a CNA agent to a specific zone.

Prerequisites

Must be in Chatter configuration command mode. For information on how to access this mode, see Accessing and exiting command modes.

Occurrence

Available for use multiple times in the configuration.

Description

Creates a zone or subzone. Zones contain subzones, which in turn contain either lower level subzones or CNA agents. The zone name in the following example is a fully qualified name. Zone names at each level are separated by a dot.

When you enter a fully qualified subzone name, the higher level subzones and zone are automatically created, if not already created.

When you enter the name of an existing zone or subzone for <*name*>, this command enters Chatter Zone Configuration command mode for the zone or subzone that you specify.

The no form of this command deletes the zone or subzone and all of its lower subzones or CNA agents.

Example

Zone AV contains subzones:

- APAC
- AVBackbone
- CALA
- EMEA
- US

Each of these zones contains further subzones. For example:

- APAC contains subzones:
  - Beijing
  - Malaysia
  - Pune
  - Sydney
- AVBackbone contains subzones:
  - APACBB
  - EMEABB
  - USBB

Leaf zones (the lowest level subzones) contain CNA agents.

In the following configuration example, zone US and subzone Beijing have CNA agents assigned to them.

```
zone AV.APAC.Beijing define
  cna-agent ...
end
zone AV.APAC.Malaysia
zone AV.APAC.Pune
zone AV.APAC.Sydney
zone AV.AVBackbone.APACBB
zone AV.AVBackbone.EMEABB
zone AV.AVBackbone.USBB
zone AV.CALA
zone AV.EMEA
zone AV.US define
  cna-agent ...
end
```

To delete everything contained in the AVBackbone zone, enter **no zone AV. AVBackbone**. This command removes all the subzones and CNA agents contained in the AVBackbone zone.

Index

A

aaa . . . 29
access list . . . 157, 158
active probe . . . 32
active-measurement name . . . 32
advised winner . . . 373
agent-based measurement
  active-measurement name . . . 32
  alarm name . . . 41
  alarm-absolute disable . . . 35
  alarm-absolute disable ping . . . 36
  alarm-absolute disable rtp . . . 37
  alarm-absolute disable tcp . . . 39
  alarm-absolute disable traceroute . . . 40
  alarm-relative disable . . . 42
  alarm-relative disable ping . . . 43
  alarm-relative disable rtp . . . 45
  alarm-relative disable tcp . . . 46
  alarm-relative disable traceroute . . . 47
  alarm-threshold-exceeded . . . 49
  chatter assign . . . 79
  chatter join-hive . . . 80
  chatter leave-hive . . . 80
  chatter missed-heartbeats . . . 81
  chatter node add . . . 82
  chatter node split . . . 83
  chatter remove-cna-agent . . . 83
  cna-agent . . . 93
  enable . . . 132
  interval-between-packets . . . 156
  ip ssh key generate cna-agent . . . 174
  logging host . . . 191
  logging-server . . . 192
  max-ttl . . . 196
  measurement intra-zone disable . . . 197
  module chatter . . . 203
  module chatter or testplug . . . 204
  number-of-packets . . . 223
  packet-length . . . 234
  port . . . 260
  rate . . . 273
  rename zone . . . 285
  repeat . . . 286
  resync chatter chapld . . . 295
  resync chatter node-alias . . . 295
  set-alarm edge . . . 306
  set-alarm zone . . . 307
  show cell . . . 315
  show chatter cna-agent-device . . . 320
  show chatter master-key . . . 322
  show chatter ssl . . . 322
  show cna-agent . . . 324
  show edge . . . 335
  show nodes . . . 351
  show outage chatter . . . 355
  show statistics chatter . . . 367
  show zone . . . 378
  smtp to . . . 384
  smtp-server . . . 386
  snmp host . . . 391
  snmp-server . . . 392
  target . . . 418
  threshold ping . . . 423
  threshold rtp . . . 426
  threshold tcp . . . 428
  threshold traceroute . . . 430
  timeout . . . 433
  topological-analysis disable . . . 452
  tos byte . . . 454
  zone . . . 488
alarm name . . . 41
alarm-absolute disable . . . 35
alarm-absolute disable ping . . . 36
alarm-absolute disable rtp . . . 37
alarm-absolute disable tcp . . . 39
alarm-absolute disable traceroute . . . 40
alarm-relative disable . . . 42
alarm-relative disable ping . . . 43
alarm-relative disable rtp . . . 45
alarm-relative disable tcp . . . 46
alarm-relative disable traceroute . . . 47
alarm-threshold-exceeded . . . 49
am . . . 107
am-measurer . . . 89, 107, 366
am-scheduler . . . 89, 107, 366
application . . . 50
application age-out . . . 51
application detection . . . 52
application no-detection . . . 52
application priority . . . 66
application-model other . . . 58
application-model streaming . . . 60
application-model voice . . . 62
arp . . . 86
authentication . . . 29

B

bgp . . . 75
BGP winner . . . 373
billing . . . 107
billing-rate . . . 256
builds . . . 107

C

case . . . 27
chatter assign . . . 79
chatter join-hive . . . 80
chatter leave-hive . . . 80
chatter missed-heartbeats . . . 81
chatter node add . . . 82
chatter node split . . . 83
chatter remove-cna-agent . . . 83
Check Point . . . 114, 118
cidr-lookup . . . 84
clear active-measurements . . . 85
clear counters web-hits . . . 88
clear efc active-measurement . . . 87
cna-agent . . . 93
command mode
  Chatter Active Measurement Configuration . . . 32
  Chatter Alarm Configuration . . . 41
  Chatter Configuration . . . 203
  Chatter Zone Configuration . . . 488
confirmed winner . . . 373
content . . . 97
coreman . . . 107

D

damped-mode . . . 101
data collection . . . 465
date . . . 103
Decision Maker . . . 201
description . . . 114
determined winner . . . 373
diagnose . . . 115, 116
diagnose target . . . 115, 116
diagnostic . . . 331
dmaker . . . 89, 107, 366
dmaker-configd . . . 107
dmaker-decisionmaker . . . 107
dmaker-main . . . 107
dmaker-measurementlistener . . . 107
dmaker-peermgr . . . 107
dmaker-prefixscheduler . . . 107
dmaker-probefailurelistener . . . 107
dmaker-snmplistener . . . 107
dmaker-trendsummarylive . . . 107
dm-listener . . . 89, 366
DNS . . . 160, 172
Domain Naming System . . . 160
dump-measurements . . . 119

E

edgestats . . . 89, 107, 366
efc . . . 107, 366
efc enable . . . 129
email . . . 107, 129
email format image . . . 132
enable . . . 132
erase . . . 134
event-manager . . . 89, 366
exclude . . . 136
export report scheduled . . . 144
export trend . . . 145

F

filenames . . . 195, 196
flapping . . . 146
ftp . . . 99, 162, 163

H

handshake round trip time . . . 372
help . . . 27
hierarchical tree . . . 462
host . . . 151
HRTT . . . 153
hrtt . . . 372
http . . . 99, 169

I

insufficient data . . . 373
interval-between-packets . . . 156
ip ssh key generate cna-agent . . . 174

L

link-monitor . . . 89, 107, 366
local . . . 29
local login . . . 476
local mode . . . 21
logging . . . 89, 366
logging host . . . 191
logging-server . . . 192
login . . . 21

M

Management Processor . . . 21
max-ttl . . . 196
measurement intra-zone disable . . . 197
medic . . . 107
mode
  local . . . 21
module chatter . . . 203
module chatter or testplug . . . 204
monitor . . . 89, 107, 366
more . . . 206
MTU . . . 171
multiple static routes . . . 173

N

neighbor . . . 240
neighbor uses-next-hop-self-address . . . 240
netflow . . . 108
Network Time Protocol . . . 220
next-hop-self . . . 206, 214, 240
not used . . . 373
NTP . . . 220
number-of-packets . . . 223

O

outage-detection brownout . . . 226

P

P3P . . . 149
packet-length . . . 234
passwords . . . 303
persistent-hosts . . . 236
persistent-hosts timeout . . . 238
Platform for Privacy Preferences . . . 149
port . . . 260
post-proc . . . 108
privileged . . . 21
probe . . . 32
publish . . . 266
publish report . . . 266
publish trend . . . 268

R

RADIUS . . . 29, 269
rate . . . 273
reboot . . . 134
refresh-time . . . 276
refresh-time acquire . . . 278
refresh-time host-disappeared . . . 277
refresh-time proxy-disappeared . . . 278
refresh-time reacquire . . . 278
refresh-time verify . . . 279
reliability half-life . . . 280
reliability tolerance . . . 281
remote BGP . . . 373
remote connection . . . 21
remote-exec . . . 283
rename zone . . . 285
repeat . . . 286
resync chatter chapld . . . 295
resync chatter node-alias . . . 295
route flapping . . . 146
rrd-writer-main . . . 108
rrd-writer-rrdpacket-receiver . . . 108
rsync . . . 108
running config . . . 96, 99

S

services . . . 108
set-alarm edge . . . 306
set-alarm zone . . . 307
set-link-group . . . 310
show . . . 360
show cell . . . 315
show chatter cna-agent-device . . . 320
show chatter master-key . . . 322
show chatter ssl . . . 322
show cna-agent . . . 324
show edge . . . 335
show nodes . . . 351
show outage chatter . . . 355
show outage-not-rescued . . . 356
show outage-rescued . . . 356
show penalty-setting link . . . 357
show statistics chatter . . . 367
show zone . . . 378
showd-engine . . . 107
showd-reporting . . . 108
smtp to . . . 384
smtp-server . . . 386
snapshot . . . 88, 108, 111
snapshot reports . . . 388
snapshot-main . . . 108
snapshots clear-after . . . 387
snapshot-update . . . 108
snmp . . . 108
snmp host . . . 391
snmp-server . . . 392
snmp-server enable traps . . . 394
ssh . . . 21, 174, 175
startup config . . . 99, 134
static route . . . 172

sufficiently measured 201

T

TACACS+ 29, 414
target (agent-based measurement) 418
telnet 21, 175
terminal 96
terminate 142
threshold host-disappeared 422, 423
threshold host-link-fail 423
threshold ping 423
threshold proxy-disappeared 425
threshold proxy-link-fail 426
threshold rtp 426
threshold tcp 428
threshold traceroute 430
timeout (agent-based measurement) 433
timestamps 305
topological-analysis disable 452
tos byte 453
tos byte (agent-based measurement) 454
traceprobe convergence-rule 455
trend-live 107
trend-snapshot 108

U

unknown error 374
User Traffic Test 417
user-defined-application 475
UTT 372, 417

V

virtual IP address 478

W

web-speedup http-turns 479
web-speedup improvable min-hrtt-saving 481
web-speedup improvable min-percent-improvement . 482
web-speedup min-bgp-hrtt 483
web-speedup worst-bgp-hrtt-percent 484
winner-set 486

Z

zone 488