

show b – show cq

- show backup-package, on page 3
- show bfd drops, on page 5
- show bfd map, on page 7
- show bfd neighbors, on page 9
- show bfd summary, on page 11
- show bgp, on page 13
- show bgp all community, on page 20
- show bgp all neighbors, on page 23
- show bgp cidr-only, on page 28
- show bgp community, on page 30
- show bgp community-list, on page 32
- show bgp filter-list, on page 35
- show bgp injected-paths, on page 37
- show bgp ipv4, on page 39
- show bgp ipv6, on page 41
- show bgp ipv6 community, on page 44
- show bgp ipv6 community-list, on page 47
- show bgp ipv6 filter-list, on page 50
- show bgp ipv6 inconsistent-as, on page 53
- show bgp ipv6 neighbors, on page 55
- show bgp ipv6 paths, on page 63
- show bgp ipv6 prefix-list, on page 65
- show bgp ipv6 quote-regexp, on page 67
- show bgp ipv6 regexp, on page 69
- show bgp ipv6 route-map, on page 71
- show bgp ipv6 summary, on page 73
- show bgp neighbors, on page 75
- show bgp paths, on page 86
- show bgp policy-list, on page 88
- show bgp prefix-list, on page 89
- show bgp regexp, on page 90
- show bgp replication, on page 92
- show bgp rib-failure, on page 94

- show bgp summary, on page 96
- show bgp system-config, on page 100
- show blocks, on page 101
- show bootvar, on page 110
- show bridge-group, on page 112
- show call-home, on page 114
- show call-home registered-module status, on page 119
- show capture, on page 120
- show chardrop, on page 126
- show checkheaps, on page 127
- show checksum, on page 128
- show chunkstat, on page 129
- show class, on page 131
- show clns, on page 132
- show clock, on page 142
- show cluster, on page 144
- show cluster history, on page 147
- show cluster info, on page 150
- show cluster user-identity, on page 159
- show cluster vpn-sessiondb distribution, on page 161
- show compression, on page 163
- show configuration, on page 165
- show configuration session, on page 169
- show conn, on page 171
- show console-output, on page 184
- show context, on page 185
- show controller, on page 189
- show coredump filesystem, on page 195
- show coredump log, on page 197
- show counters, on page 199
- show cpu, on page 203

show backup-package

To display back-up package status and summary information on the Cisco ISA 3000, use the **show backup-package** command in privileged EXEC or global configuration mode.

show backup-package { **status** { **backup | restore** } **| summary** }

Note This command applies only to the Cisco ISA 3000 appliance.

Syntax Description

backup | restore	Specifies the type of **status** information to be displayed.
status	Displays mode, location, passphrase, and most-recent time information for either back-up or restore operations.
summary	Displays status information for both back-up and restore operations.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	—	—

Command History

Release	Modification
9.7(1)	This command was added.

Usage Guidelines

The **show backup-package** commands are also available in global configuration mode.

Examples

The following example shows backup-package summary statistics:

```
ciscoasa# show backup-package summary
 backup mode      : auto
 backup location  : disk3:
 backup passphrase: cisco
 last backup time : Mar 23 2014 22:05:52
 restore mode     : auto
 restore location  : disk3:
```

```
restore passphrase: cisco
Last restore time : Mar 24 2014 05:07:32
```

show bfd drops

To display the numbered of dropped packets in BFD, use the show bfd drops command in global configuration mode.

show bfd drops

Syntax Description

This command has no arguments or keywords.

Command Default

This command has no default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Global configuration	• Yes	—	• Yes	• Yes	—

Command History

Release	Modification
9.6(2)	This command was added.

Examples

The following example displays the BFD dropped packets.

```
ciscoasa# show bfd drops
BFD Drop Statistics
                        IPV4 IPV6 IPV4-M IPV6-M
Invalid TTL             0    0    0      0
BFD Not Configured      0    0    0      0
No BFD Adjacency        0    0    0      0
Invalid Header Bits     0    0    0      0
Invalid Discriminator   0    0    0      0
Session AdminDown       0    0    0      0
Authen invalid BFD ver  0    0    0      0
Authen invalid len      0    0    0      0
Authen invalid seq      0    0    0      0
Authen failed           0    0    0      0
```

Related Commands

Command	Description
authentication	Configures authentication in a BFD template for single-hop and multi-hop sessions.
bfd echo	Enables BFD echo mode on the interface,
bfd interval	Configures the baseline BFD parameters on the interface.

Command	Description
bfd map	Configures a BFD map that associates addresses with multi-hop templates.
bfd slow-timers	Configures the BFD slow timers value.
bfd template	Binds a single-hop BFD template to an interface.
bfd-template single-hop | multi-hop	Configures the BFD template and enters BFD configuration mode.
clear bfd counters	Clears the BFD counters.
echo	Configures echo in the BFD single-hop template.
neighbor	Configures BFD support for BGP so that BGP is registered to receive forwarding path detection failure messages from BFD.
show bfd map	Displays the configured BFD maps.
show bfd neighbors	Displays a line-by-line listing of existing BFD adjacencies.
show bfd summary	Displays summary information for BFD.

show bfd map

To display the configured BFD maps, use the show bfd map command in global configuration mode.

show bfd map

Syntax Description

This command has no arguments or keywords.

Command Default

This command has no default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Global configuration	• Yes	—	• Yes	• Yes	—

Command History

Release	Modification
9.6(2)	This command was added.

Examples

The following example displays the BFD maps.

```
ciscoasa# show bfd map
Destination: 40.40.40.2/24
 Source: 50.50.50.2/24
 Template: mh
 Authentication(Type): sha-1
```

Related Commands

Command	Description
authentication	Configures authentication in a BFD template for single-hop and multi-hop sessions.
bfd echo	Enables BFD echo mode on the interface,
bfd interval	Configures the baseline BFD parameters on the interface.
bfd map	Configures a BFD map that associates addresses with multi-hop templates.
bfd slow-timers	Configures the BFD slow timers value.
bfd template	Binds a single-hop BFD template to an interface.

Command	Description
bfd-template single-hop | multi-hop	Configures the BFD template and enters BFD configuration mode.
clear bfd counters	Clears the BFD counters.
echo	Configures echo in the BFD single-hop template.
neighbor	Configures BFD support for BGP so that BGP is registered to receive forwarding path detection failure messages from BFD.
show bfd drops	Displays the numbered of dropped packets in BFD.
show bfd neighbors	Displays a line-by-line listing of existing BFD adjacencies.
show bfd summary	Displays summary information for BFD.

show bfd neighbors

To display a line-by-line listing of existing BFD adjacencies, use the show bfd neighbors command in global configuration mode.

show bfd neighbors [**client** { **bgp** } | **details** | **interface** *interface-name* | **ipv4** *ip-address* | **ipv6** *ipv6-address* | **multihop-ipv4** *ip-address* | **multihop-ipv6** *ipv6-address*]

Syntax Description

client	(Optional) Displays the neighbors of a specific client.
bgp	(Optional) Displays a BGP client.
details	(Optional) Displays all BFD protocol parameters and timers for each neighbor.
interface interface-name	(Optional) Displays neighbors at the specified interface.
ipv4 ip-address	(Optional) Displays specified single-hop IP neighbors.
ipv6 ipv6-address	(Optional) Displays specified single-hop IPv6 neighbors.
multihop-ipv4 ip-address	(Optional) Displays specified multi-hop IP neighbors.
multihop-ipv6 ipv6-address	(Optional) Displays specified multi-hop IPv6 neighbors.

Command Default

This command has no default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Global configuration	• Yes	—	• Yes	• Yes	—

Command History

Release	Modification
9.6(2)	This command was added.

Usage Guidelines

Use this command to troubleshoot BFD issues.

Examples

The following example displays the BFD neighbors.

```
ciscoasa# show bfd neighbors
OurAddr     NeighAddr    LD/RD  RH     Holdown(mult)     State Int
172.16.10.1 172.16.10.2  1/6    1      260 (3 )          Up    Fa0/1
```

Related Commands

Command	Description
authentication	Configures authentication in a BFD template for single-hop and multi-hop sessions.
bfd echo	Enables BFD echo mode on the interface,
bfd interval	Configures the baseline BFD parameters on the interface.
bfd map	Configures a BFD map that associates addresses with multi-hop templates.
bfd slow-timers	Configures the BFD slow timers value.
bfd template	Binds a single-hop BFD template to an interface.
bfd-template single-hop | multi-hop	Configures the BFD template and enters BFD configuration mode.
clear bfd counters	Clears the BFD counters.
echo	Configures echo in the BFD single-hop template.
neighbor	Configures BFD support for BGP so that BGP is registered to receive forwarding path detection failure messages from BFD.
show bfd drops	Displays the numbered of dropped packets in BFD.
show bfd map	Displays the configured BFD maps.
show bfd summary	Displays summary information for BFD.

show bfd summary

To display summary information for BFD, use the show bfd summary command in global configuration mode.

show bfd summary [**client | host | session**]

Syntax Description

client	(Optional) Displays the BFD summary for clients.
host	(Optional) Displays the BFD summary for sessions.
session	(Optional) Displays the BFD summary for protocols.

Command Default

This command has no default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Global configuration	• Yes	—	• Yes	• Yes	—

Command History

Release	Modification
9.6(2)	This command was added.

Usage Guidelines

Use this command to display summary information about BFD, BFD clients, or BFD sessions. When a BFD client launches a session with a peer, BFD sends periodic BFD control packets to the peer. Information about the following states of a session are included in the output of this command:

- Up—When another BFD interface acknowledges the BFD control packets, the session moves into an Up state.
- Down—The session and the data path are declared down if a data path failure occurs and BFD does not receive a control packet within the configured amount of time. When a session is down, BFD notifies the BFD client so that the client can perform necessary actions to reroute the traffic.

Examples

The following example displays the BFD summaries.

```
ciscoasa# show bfd summary

         Session         Up      Down
Total    1               1       0
ciscoasa# show bfd summary session
Protocol                  Session     Up  Down
IPV4                      1     1     0
```

```
Total                      1     1     0
ciscoasa# show bfd summary client
Client                    Session      Up     Down
BGP                       1       1    0
EIGRP                     1       1    0
Total                     2       2    0
```

Related Commands

Command	Description
authentication	Configures authentication in a BFD template for single-hop and multi-hop sessions.
bfd echo	Enables BFD echo mode on the interface,
bfd interval	Configures the baseline BFD parameters on the interface.
bfd map	Configures a BFD map that associates addresses with multi-hop templates.
bfd slow-timers	Configures the BFD slow timers value.
bfd template	Binds a single-hop BFD template to an interface.
bfd-template single-hop | multi-hop	Configures the BFD template and enters BFD configuration mode.
clear bfd counters	Clears the BFD counters.
echo	Configures echo in the BFD single-hop template.
neighbor	Configures BFD support for BGP so that BGP is registered to receive forwarding path detection failure messages from BFD.
show bfd drops	Displays the numbered of dropped packets in BFD.
show bfd map	Displays the configured BFD maps.
show bfd neighbors	Displays a line-by-line listing of existing BFD adjacencies.

show bgp

To display entries in the Border Gateway Protocol (BGP) routing table, use the show bgp command in user EXEC or privileged EXEC mode.

show bgp [*ip-address* [*mask* [**longer-prefixes** [**injected**] | **shorter-prefixes** [*length*] | **bestpath | multipaths | subnets**] | **bestpath | multipaths**] | **all | prefix-list name | pending-prefixes | route-map** *name*]]

Syntax Description

ip-address	(Optional) Specifies the AS path access list name..
mask	(Optional) Mask to filter or match hosts that are part of the specified network.
longer-prefixes	(Optional) Displays the specified route and all more specific routes.
injected	(Optional) Displays more specific prefixes injected into the BGP routing table.
shorter-prefixes	(Optional) Displays the specified route and all less specific routes.
length	(Optional) The prefix length. The value for this argument is a number from 0 to 32.
bestpath	(Optional) Displays the bestpath for this prefix
multipaths	(Optional) Displays multipaths for this prefix.
subnets	(Optional) Displays the subnet routes for the specified prefix.
all	(Optional) Displays all address family information in the BGP routing table.
prefix-list name	(Optional) Filters the output based on the specified prefix list.
pending-prefixes	(Optional) Displays prefixes that are pending deletion from the BGP routing table.
route-map name	(Optional) Filters the output based on the specified route map.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	—	• Yes	• Yes	—

Command History

Release	Modification
9.2(1)	This command was added

Usage Guidelines

The show bgp command is used to display the contents of the BGP routing table. The output can be filtered to display entries for a specific prefix, prefix length, and prefixes injected through a prefix list, route map, or conditional advertisement.

In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, the Cisco implementation of 4-byte autonomous system numbers uses asplain—65538 for example—as the default regular expression match and output display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain format and the asdot format as described in RFC 5396. To change the default regular expression match and output display of 4-byte autonomous system numbers to asdot format, use the bgp asnotation dot command followed by the clear bgp * command to perform a hard reset of all current BGP sessions.

Examples

The following sample output shows the BGP routing table:

```
Router# show bgp
BGP table version is 22, local router ID is 10.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.1.1/32      0.0.0.0                  0         32768 i
*>i10.2.2.2/32      172.16.1.2               0    100      0 i
*bi10.9.9.9/32      192.168.3.2              0    100      0 10 10 i
*>                  192.168.1.2                           0 10 10 i
* i172.16.1.0/24    172.16.1.2               0    100      0 i
*>                  0.0.0.0                  0         32768 i
*> 192.168.1.0      0.0.0.0                  0         32768 i
*>i192.168.3.0      172.16.1.2               0    100      0 i
*bi192.168.9.0      192.168.3.2              0    100      0 10 10 i
*>                  192.168.1.2                           0 10 10 i
*bi192.168.13.0     192.168.3.2              0    100      0 10 10 i
*>                  192.168.1.2                           0 10 10 i
```

Table 4-1 shows each field description.

Table 1: show bgp Fields

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.

Field	Description
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: • s—The table entry is suppressed. • d—The table entry is dampened. • h—The table entry history. • *—The table entry is valid. • >—The table entry is the best entry to use for that network. • i—The table entry was learned via an internal BGP (iBGP) session. • r—The table entry is a RIB-failure. • S—The table entry is stale. • m—The table entry has multipath to use for that network. • b—The table entry has backup path to use for that network. • x—The table entry has best external route to use for the network.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: • i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. • e—Entry originated from an Exterior Gateway Protocol (EGP). • ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	IP address of a network entity.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.
Metric	If shown, the value of the interautonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.
(stale)	Indicates that the following path for the specified autonomous system is marked as "stale" during a graceful restart process.

Examples

show bgp (4-Byte Autonomous System Numbers): Example

The following sample output shows the BGP routing table with 4-byte autonomous system numbers, 65536 and 65550, shown under the Path field. This example requires Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, or a later release.

```
RouterB# show bgp
BGP table version is 4, local router ID is 172.17.1.99
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.1.0/24      192.168.1.2              0             0 65536  i
*> 10.2.2.0/24      192.168.3.2              0             0 65550  i
*> 172.17.1.0/24    0.0.0.0                  0         32768 i
```

show bgp ip-address: Example

The following sample output displays information about the 192.168.1.0 entry in the BGP routing table:

```
Router# show bgp 192.168.1.0
BGP routing table entry for 192.168.1.0/24, version 22
Paths: (2 available, best #2, table default)
  Additional-path
  Advertised to update-groups:
     3
  10 10
    192.168.3.2 from 172.16.1.2 (10.2.2.2)
      Origin IGP, metric 0, localpref 100, valid, internal, backup/repair
  10 10
    192.168.1.2 from 192.168.1.2 (10.3.3.3)
      Origin IGP, localpref 100, valid, external, best , recursive-via-connected
```

The following sample output displays information about the 10.3.3.3 255.255.255.255 entry in the BGP routing table:

```
Router# show bgp 10.3.3.3 255.255.255.255
BGP routing table entry for 10.3.3.3/32, version 35
Paths: (3 available, best #2, table default)
Multipath: eBGP
Flag: 0x860
  Advertised to update-groups:
     1
  200
    10.71.8.165 from 10.71.8.165 (192.168.0.102)
      Origin incomplete, localpref 100, valid, external, backup/repair
      Only allowed to recurse through connected route
  200
    10.71.11.165 from 10.71.11.165 (192.168.0.102)
      Origin incomplete, localpref 100, weight 100, valid, external, best
      Only allowed to recurse through connected route
  200
    10.71.10.165 from 10.71.10.165 (192.168.0.104)
      Origin incomplete, localpref 100, valid, external,
      Only allowed to recurse through connected route
```

Table 4-2 shows each field description.

Table 2: show bgp (4 byte autonomous system numbers) Fields

Field	Description
BGP routing table entry fo	IP address or network number of the routing table entry.
version	Internal version number of the table. This number is incremented whenever the table changes.
Paths	The number of available paths, and the number of installed best paths. This line displays "Default-IP-Routing-Table" when the best path is installed in the IP routing table.
Multipath	This field is displayed when multipath loadsharing is enabled. This field will indicate if the multipaths are iBGP or eBGP.
Advertised to update-groups	The number of each update group for which advertisements are processed.
Origin	Origin of the entry. The origin can be IGP, EGP, or incomplete. This line displays the configured metric (0 if no metric is configured), the local preference value (100 is default), and the status and type of route (internal, external, multipath, best).
Extended Community	This field is displayed if the route carries an extended community attribute. The attribute code is displayed on this line. Information about the extended community is displayed on a subsequent line.

Examples

show bgp all: Example

The following is sample output from the show bgp command entered with the all keyword. Information about all configured address families is displayed.

```
Router# show bgp all
For address family: IPv4 Unicast   *****
BGP table version is 27, local router ID is 10.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.1.0/24      0.0.0.0                  0         32768 ?
*> 10.13.13.0/24    0.0.0.0                  0         32768 ?
*> 10.15.15.0/24    0.0.0.0                  0         32768 ?
*>i10.18.18.0/24    172.16.14.105         1388  91351      0 100 e
*>i10.100.0.0/16    172.16.14.107          262    272      0 1 2 3 i
*>i10.100.0.0/16    172.16.14.105         1388  91351      0 100 e
*>i10.101.0.0/16    172.16.14.105         1388  91351      0 100 e
*>i10.103.0.0/16    172.16.14.101         1388    173    173 100 e
*>i10.104.0.0/16    172.16.14.101         1388    173    173 100 e
*>i10.100.0.0/16    172.16.14.106         2219  20889      0 53285 33299 51178 47751 e
*>i10.101.0.0/16    172.16.14.106         2219  20889      0 53285 33299 51178 47751 e
*  10.100.0.0/16    172.16.14.109         2309             0 200 300 e
*>                  172.16.14.108         1388             0 100 e
*  10.101.0.0/16    172.16.14.109         2309             0 200 300 e
*>                  172.16.14.108         1388             0 100 e
*> 10.102.0.0/16    172.16.14.108         1388             0 100 e
*> 172.16.14.0/24   0.0.0.0                  0         32768 ?
```

```
*> 192.168.5.0      0.0.0.0                  0         32768 ?
*> 10.80.0.0/16     172.16.14.108         1388             0 50 e
*> 10.80.0.0/16     172.16.14.108         1388             0 50 e
```

show bgp longer-prefixes: Example

The following is sample output from the show bgp command entered with the longer-prefixes keyword:

```
Router# show bgp 10.92.0.0 255.255.0.0 longer-prefixes
BGP table version is 1738, local router ID is 192.168.72.24
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Metric LocPrf Weight Path
*> 10.92.0.0         10.92.72.30         8896         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.1.0         10.92.72.30         8796         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.11.0        10.92.72.30        42482         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.14.0        10.92.72.30         8796         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.15.0        10.92.72.30         8696         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.16.0        10.92.72.30         1400         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.17.0        10.92.72.30         1400         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.18.0        10.92.72.30         8876         32768 ?
*                    10.92.72.30                          0 109 108 ?
*> 10.92.19.0        10.92.72.30         8876         32768 ?
*                    10.92.72.30                          0 109 108 ?
```

show bgp shorter-prefixes: Example

The following is sample output from the show bgp command entered with the shorter-prefixes keyword. An 8-bit prefix length is specified.

```
Router# show bgp 172.16.0.0/16 shorter-prefixes 8
*> 172.16.0.0         10.0.0.2                               0 ?
*                     10.0.0.2                 0             0 200 ?
```

show bgp prefix-list: Example

The following is sample output from the show bgp command entered with the prefix-list keyword:

```
Router# show bgp prefix-list ROUTE
BGP table version is 39, local router ID is 10.0.0.1
Status codes:s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes:i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 192.168.1.0      10.0.0.2                               0 ?
*                    10.0.0.2                 0             0 200 ?
```

show bgp route-map: Example

The following is sample output from the show bgp command entered with the route-map keyword:

```
Router# show bgp route-map LEARNED_PATH
BGP table version is 40, local router ID is 10.0.0.1
Status codes:s suppressed, d damped, h history, * valid, > best, i -
```

```
internal
Origin codes:i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 192.168.1.0      10.0.0.2                             0 ?
*                   10.0.0.2                 0           0 200 ?
```

show bgp all community

To display routes for all address families belonging to a particular Border Gateway Protocol (BGP) community, use the show bgp all community command in user EXEC or privileged EXEC configuration mode.

show bgp all community [*community-number....* [*community-number*]] [**local-as**] [**no-advertise**] [**no-export**] [**exact-match**]

Syntax Description

community-number.	(Optional) Displays the routes pertaining to the community numbers specified. You can specify multiple community numbers. The range is from 1 to 4294967295 or AA:NN (autonomous system:community number, which is a 2-byte number).
local-as	(Optional) Displays only routes that are not sent outside of the local autonomous system (well-known community).
no-advertise	(Optional) Displays only routes that are not advertised to any peer (well-known community).
no-export	(Optional) Displays only routes that are not exported outside of the local autonomous system (well-known community).
exact-match	(Optional) Displays only routes that match exactly with the BGP community list specified. **Note** The availability of keywords in the command depends on the command mode. The exact-match keyword is not available in user EXEC mode.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Usage Guidelines

User can enter the local-as, no-advertise and no-export keywords in any order. When using the bgp all community command, be sure to enter the numerical communities before the well-known communities.

.For example, the following string is not valid:

```
ciscoasa# show bgp all community local-as 111:12345
```

Use the following string instead:

```
ciscoasa# show bgp all community 111:12345 local-as
```

Examples

The following is sample output from the show bgp all community command, specifying communities of 1, 2345, and 6789012:

```
ciscoasa# show bgp all community 1 2345 6789012 no-advertise local-as no-export exact-match
For address family: IPv4 Unicast
BGP table version is 5, local router ID is 30.0.0.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network    Next Hop            Metric LocPrf Weight  Path
*> 10.0.3.0/24       10.0.0.4                               0 4 3 ?
*> 10.1.0.0/16       10.0.0.4                 0             0 4 ?
*> 10.12.34.0/24     10.0.0.6                 0             0 6 ?
```

Table 4-26 shows each field description.

Table 3: show bgp all community Fields

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes
local router ID	The router ID of the router on which the BGP communities are set to display. A 32-bit number written as 4 octets separated by periods (dotted-decimal format).
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened.h—The table entry is history.*—The table entry is valid.>—The table entry is the best entry to use for that network.i—The table entry was learned via an internal BGP session.
Origin codes	Indicates the origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from the Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from the Exterior Gateway Protocol (EGP).?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	The network address and network mask of a network entity. The type of address depends on the address family.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. The type of address depends on the address family
Metric	The value of the inter autonomous system metric. This field is not used frequently.

Field	Description
LocPrf	Local preference value as set with the set local-preference command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

show bgp all neighbors

To display information about Border Gateway Protocol (BGP) connections to neighbors of all address families, use the show bgp all neighbors command in user EXEC or privileged EXEC mode.

show bgp all neighbors [*ip-address*] [**advertised-routes | paths** [*reg-exp*] **| policy** [**detail**] **| received prefix-filter | received-routes | routes**]

Syntax Description

ip-address	(Optional) IP address of a neighbor. If this argument is omitted, information about all neighbors is displayed.
advertised-routes	Optional) Displays all routes that have been advertised to neighbors.
paths reg-exp	(Optional) Displays autonomous system paths learned from the specified neighbor. An optional regular expression can be used to filter the output.
policy	(Optional) Displays the policies applied to neighbor per address family.
detail	(Optional) Displays detailed policy information such as route maps, prefix lists, community lists, Access Control Lists (ACLs), and autonomous system path filter lists.
received prefix-filter	(Optional) Displays the prefix-list (outbound route filter [ORF]) sent from the specified neighbor.
received-routes	(Optional) Displays all received routes (both accepted and rejected) from the specified neighbor.
routes	(Optional) Displays all routes that are received and accepted. The output displayed when this keyword is entered is a subset of the output displayed by the received-routes keyword.

Command Default

The output of this command displays information for all neighbors.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Usage Guidelines

Use the show bgp all neighbors command to display BGP and TCP connection information for neighbor sessions specific to address families such as IPv4.

Examples

The following example shows output of the show bgp all neighbors command:

```
ciscoasa# show bgp all neighbors
For address family: IPv4 Unicast
BGP neighbor is 172.16.232.53, remote AS 100, external link
 Member of peer-group internal for session parameters
  BGP version 4, remote router ID 172.16.232.53
  BGP state = Established, up for 13:40:17
  Last read 00:00:09, hold time is 180, keepalive interval is 60 seconds
  Message statistics:
    InQ depth is 0
    OutQ depth is 0
                         Sent       Rcvd
    Opens:                  3          3
    Notifications:          0          0
    Updates:                0          0
    Keepalives:           113        112
    Route Refresh:          0          0
    Total:                116        11
  Default minimum time between advertisement runs is 5 seconds
  Connections established 22; dropped 21
  Last reset 13:47:05, due to BGP Notification sent, hold time expired
External BGP neighbor may be up to 2 hops away.
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)
Event Timers (current time is 0x1A0D543C):
Timer          Starts    Wakeups            Next
Retrans          1218          5             0x0
TimeWait            0          0             0x0
AckHold          3327       3051             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0
iss: 1805423033  snduna: 1805489354  sndnxt: 1805489354     sndwnd:  15531
irs:  821333727  rcvnxt:  821591465  rcvwnd:      15547  delrcvwnd:    837
SRTT: 300 ms, RTTO: 303 ms, RTV: 3 ms, KRTT: 0 ms
minRTT: 8 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: higher precedence, nagle
Datagrams (max data segment is 1420 bytes):
Rcvd: 4252 (out of order: 0), with data: 3328, total data bytes: 257737
Sent:4445 (retransmit: 5), with data: 4445, total data bytes;244128
```

Table 4-4 shows each field description.

Table 4: show bgp all neighbor Fields

Field	Description
For address family	Address family to which the following fields refer.
BGP neighbor	IP address of the BGP neighbor and its autonomous system number.
remote AS	Autonomous system number of the neighbor.
external link	External Border Gateway Protocol (eBGP) peerP.

Field	Description
BGP version	BGP version being used to communicate with the remote router.
remote router ID	IP address of the neighbor.
BGP state	State of this BGP connection
up for	Time, in hh:mm:ss, that the underlying TCP connection has been in existence.
Last read	Time, in hh:mm:ss, since BGP last received a message from this neighbor.
hold time	Time, in seconds, that BGP will maintain the session with this neighbor without receiving messages.
keepalive interval	Time interval, in seconds, at which keepalive messages are transmitted to this neighbor.
Message statistics	Statistics organized by message type.
InQ depth is	Number of messages in the input queue.
OutQ depth is	Number of messages in the output queue.
Sent	Total number of transmitted messages.
Rcvd	Total number of received messages.
Opens	Number of open messages sent and received.
Notifications	Number of notification (error) messages sent and received.
Updates	Number of update messages sent and received.
Keepalives	Number of keepalive messages sent and received.
Route Refresh	Number of route refresh request messages sent and received.
Total	Total number of messages sent and received.
Default minimum time between...	Time, in seconds, between advertisement transmissions.
Connections established	Number of times a TCP and BGP connection has been successfully established.
dropped	Number of times that a valid session has failed or been taken down.
Last reset	Time, in hh:mm:ss, since this peering session was last reset. The reason for the reset is displayed on this line.
External BGP neighbor may be...	Indicates that the BGP Time-to-live (TTL) security check is enabled. The maximum number of hops that can separate the local and remote peer is displayed on this line.
Connection state	Connection status of the BGP peer.

Field	Description
Local host, Local	IP address of the local BGP speaker and the port number.
Foreign host, Foreign port	Neighbor address and BGP destination port number.
Enqueued packets for retransmit:	Packets queued for retransmission by TCP.
Event Timers	TCP event timers. Counters are provided for starts and wakeups (expired timers).
Retrans	Number of times a packet has been retransmitted.
TimeWait	Time waiting for the retransmission timers to expire.
AckHold	Acknowledgment hold timer.
SendWnd	Transmission (send) window.
KeepAlive	Number of keepalive packets.
GiveUp	Number times a packet is dropped due to no acknowledgment.
PmtuAger	Path MTU discovery timer.
DeadWait	Expiration timer for dead segments.
iss:	Initial packet transmission sequence number.
snduna:	Last transmission sequence number that has not been acknowledged
sndnxt:	Next packet sequence number to be transmitted.
sndwnd:	TCP window size of the remote host.
irs:	Initial packet receives sequence number.
rcvnxt:	Last receive sequence number that has been locally acknowledged.
rcvwnd:	TCP window size of the local host.
delrcvwnd:	Delayed receive window—data the local host has read from the connection, but has not yet subtracted from the receive window the host has advertised to the remote host. The value in this field gradually increases until it is larger than a full-sized packet, at which point it is applied to the rcvwnd field.
SRTT:	A calculated smoothed round-trip timeout.
RTTO:	Round-trip timeout.
RTV:	Variance of the round-trip time.
KRTT:	New round-trip timeout (using the Karn algorithm). This field separately tracks the round-trip time of packets that have been re-sent.

Field	Description
minRTT:	Smallest recorded round-trip timeout (hard-wire value used for calculation).
maxRTT:	Largest recorded round-trip timeout.
ACK hold	Length of time the local host will delay an acknowledgment to carry (piggyback) additional data.
IP Precedence value	IP precedence of the BGP packets.
Datagrams	Number of update packets received from a neighbor.
Rcvd:	Number of received packets.
with data	Number of update packets sent with data.
total data bytes	Total amount of data received, in bytes.
Sent	Number of update packets sent.
with data	Number of update packets received with data.
total data bytes	Total amount of data sent, in bytes.

show bgp cidr-only

To display routes with classless inter domain routing (CIDR), use the show bgp cidr-only command in EXEC mode.

show bgp cidr-only

Syntax Description

This command has no arguments or keywords.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Examples

The following is sample output from the show bgp cidr-only command:

```
 ciscoasa# show bgp cidr-only

BGP table version is 220, local router ID is 172.16.73.131
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Metric LocPrf Weight Path
*> 192.168.0.0/8    172.16.72.24                         0 1878 ?
*> 172.16.0.0/16    172.16.72.30                         0 108 ?
```

Table 4-5 shows each field description.

Table 5: show bgp cidr-only Fields

Field	Description
BGP table version is 220	Internal version number of the table. This number is incremented whenever the table changes..
local router ID	IP address of the router.

Field	Description
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP..

show bgp community

To display routes that belong to specified BGP communities, use the show bgp community command in EXEC mode.

show bgp community *community-number* [**exact**]

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Examples

The following is sample output from the show bgp community command in privileged EXEC mode:

```
ciscoasa# show bgp community 111:12345 local-as
 BGP table version is 10, local router ID is 224.0.0.10
 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
 Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop          Metric LocPrf Weight Path
 *> 172.16.2.2/32    10.43.222.2           0             0 222 ?
 *> 10.0.0.0         10.43.222.2           0             0 222 ?
 *> 10.43.0.0        10.43.222.2           0             0 222 ?
 *> 10.43.44.44/32   10.43.222.2           0             0 222 ?
 *  10.43.222.0/24   10.43.222.2           0             0 222 i
 *> 172.17.240.0/21  10.43.222.2           0             0 222 ?
 *> 192.168.212.0    10.43.222.2           0             0 222 i
 *> 172.31.1.0       10.43.222.2           0             0 222 ?
```

Table 4-6 shows each field description.

Table 6: show bgp community Fields

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.

Field	Description
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP.

show bgp community-list

To display routes that are permitted by the Border Gateway Protocol (BGP) community list, use the show bgp community-list command in user or privileged EXEC mode.

show bgp community-list { *community-list-number* | *community-list-name* [**exact-match**] }

Syntax Description

community-list-number	A standard or expanded community list number in the range from 1 to 500.
community-list-name	Community list name. The community list name can be standard or expanded.
exact-match	(Optional) Displays only routes that have an exact match.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Usage Guidelines

This command requires you to specify an argument when used. The exact-match keyword is optional.

Examples

The following is sample output of the show bgp community-list command in privileged EXEC mode:

```
ciscoasa# show bgp community-list 20
BGP table version is 716977, local router ID is 192.168.32.1
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Metric LocPrf Weight Path
* i10.3.0.0         10.0.22.1             0    100      0 1800 1239 ?
*>i                 10.0.16.1             0    100      0 1800 1239 ?
* i10.6.0.0         10.0.22.1             0    100      0 1800 690 568 ?
*>i                 10.0.16.1             0    100      0 1800 690 568 ?
* i10.7.0.0         10.0.22.1             0    100      0 1800 701 35 ?
*>i                 10.0.16.1             0    100      0 1800 701 35 ?
*                   10.92.72.24                         0 1878 704 701 35 ?
* i10.8.0.0         10.0.22.1             0    100      0 1800 690 560 ?
*>i                 10.0.16.1             0    100      0 1800 690 560 ?
*                   10.92.72.24                         0 1878 704 701 560 ?
* i10.13.0.0        10.0.22.1             0    100      0 1800 690 200 ?
*>i                 10.0.16.1             0    100      0 1800 690 200 ?
*                   10.92.72.24                         0 1878 704 701 200 ?
```

```
* i10.15.0.0        10.0.22.1             0    100      0 1800 174 ?
*>i                 10.0.16.1             0    100      0 1800 174 ?
* i10.16.0.0        10.0.22.1             0    100      0 1800 701 i
*>i                 10.0.16.1             0    100      0 1800 701 i
*                   10.92.72.24                        0 1878 704 701 i
```

Table 4-7 shows each field description.

Table 7: show bgp community-list Fields

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.

Field	Description
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP..

show bgp filter-list

To display routes that conform to a specified filter list, use the show bgp filter-list command in EXEC mode.

show bgp filter-list *access-list-name*

Syntax Description

access-list-name	Name of an autonomous system path access list.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Examples

The following is sample output of the show bgp filter-list command in privileged EXEC mode:

```
ciscoasa# show bgp filter-list filter-list-acl
BGP table version is 1738, local router ID is 172.16.72.24
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Metric LocPrf Weight Path
*  172.16.0.0       172.16.72.30                           0 109 108 ?
*  172.16.1.0       172.16.72.30                           0 109 108 ?
*  172.16.11.0      172.16.72.30                           0 109 108 ?
*  172.16.14.0      172.16.72.30                           0 109 108 ?
*  172.16.15.0      172.16.72.30                           0 109 108 ?
*  172.16.16.0      172.16.72.30                           0 109 108 ?
*  172.16.17.0      172.16.72.30                           0 109 108 ?
*  172.16.18.0      172.16.72.30                           0 109 108 ?
*  172.16.19.0      172.16.72.30                           0 109 108 ?
*  172.16.24.0      172.16.72.30                           0 109 108 ?
*  172.16.29.0      172.16.72.30                           0 109 108 ?
*  172.16.30.0      172.16.72.30                           0 109 108 ?
*  172.16.33.0      172.16.72.30                           0 109 108 ?
*  172.16.35.0      172.16.72.30                           0 109 108 ?
*  172.16.36.0      172.16.72.30                           0 109 108 ?
*  172.16.37.0      172.16.72.30                           0 109 108 ?
*  172.16.38.0      172.16.72.30                           0 109 108 ?
*  172.16.39.0      172.16.72.30                           0 109 108 ?
```

Table 4-8 shows each field description.

Table 8: show bgp filter-list Fields

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP..

show bgp injected-paths

To display all the injected paths in the Border Gateway Protocol (BGP) routing table, use the show bgp injected-paths command in user or privileged EXEC mode.

show bgp injected-paths

Syntax Description

This command has no arguments or keywords.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Examples

The following is sample output from the show bgp injected-paths command in EXEC mode:

```
ciscoasa# show bgp injected-paths
BGP table version is 11, local router ID is 10.0.0.1
Status codes:s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes:i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 172.16.0.0       10.0.0.2                              0 ?
*> 172.17.0.0/16    10.0.0.2                              0 ?
```

Table 4-9 shows each field description.

Table 9: show bgp injected-path Fields

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.

Field	Description
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP.

show bgp ipv4

To display entries in the IP version 4 (IPv4) Border Gateway Protocol (BGP) routing table, use the show bgp ipv4 command in privileged EXEC mode.

show bgp ipv4

Syntax Description

This command has no arguments or keywords.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Examples

The following is sample output from the show bgp ipv4 unicast command:

```
ciscoasa# show bgp ipv4 unicast
 BGP table version is 4, local router ID is 10.0.40.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.10.10.0/24    172.16.10.1              0             0  300 i
*> 10.10.20.0/24    172.16.10.1              0             0  300 i
*  10.20.10.0/24    172.16.10.1              0             0  300 i
```

The following is sample output from the show bgp ipv4 multicast command:

```
Router# show bgp ipv4 multicast
 BGP table version is 4, local router ID is 10.0.40.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.10.10.0/24    172.16.10.1              0             0  300 i
*> 10.10.20.0/24    172.16.10.1              0             0  300 i
*  10.20.10.0/24    172.16.10.1              0             0  300 i
```

Table 4-10 shows each field description.

Table 10: show bgp ipv4 Fields

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP..

show bgp ipv6

To display entries in the IPv6 Border Gateway Protocol (BGP) routing table, use the show bgp ipv6 command in user EXEC or privileged EXEC mode.

show bgp ipv6 unicast [*ipv6-prefix/prefix-length*] [**longer-prefixes**] [**labels**]

Syntax Description

unicast	*Specifies IPv6 unicast address prefixes.*
ipv6-prefix	(Optional) IPv6 network number, entered to display a particular network in the IPv6 BGP routing table. This argument must be in the form documented in RFC 2373 where the address is specified in hexadecimal using 16-bit values between colons.
/prefix-length	(Optional) The length of the IPv6 prefix. A decimal value that indicates how many of the high-order contiguous bits of the address comprise the prefix (the network portion of the address). A slash mark must precede the decimal value.
longer-prefixes	(Optional) Displays the route and more specific routes.
labels	(Optional) Displays the policies applied to this neighbor per address family.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.3(2)	This command was added

Examples

The following is sample output from the show bgp ipv6 command:

```
ciscoasa# show bgp ipv6 unicast
BGP table version is 12612, local router ID is 172.16.7.225
Status codes: s suppressed, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop            Metric LocPrf Weight Path
*> 10.10.10.0/24   172.16.10.1             0            0  300 i
*> 10.10.20.0/24   172.16.10.1             0            0  300 i
*  10.20.10.0/24   172.16.10.1             0            0  300 i
```

The following is sample output from the show bgp ipv4 multicast command:

```
Router# show bgp ipv4 multicast
 BGP table version is 4, local router ID is 10.0.40.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
Network            Next Hop            Metric LocPrf Weight Path
*                   3FFE:C00:E:C::2                        0 3748 4697 1752 i
*                   3FFE:1100:0:CC00::1
                                                           0 1849 1273 1752 i
*  2001:618:3::/48  3FFE:C00:E:4::2          1             0 4554 1849 65002 i
*>                  3FFE:1100:0:CC00::1
                                                           0 1849 65002 i
*  2001:620::/35    2001:0DB8:0:F004::1
                                                           0 3320 1275 559 i
*                   3FFE:C00:E:9::2                        0 1251 1930 559 i
*                   3FFE:3600::A                           0 3462 10566 1930 559 i
*                   3FFE:700:20:1::11
                                                           0 293 1275 559 i
*                   3FFE:C00:E:4::2          1             0 4554 1849 1273 559 i
*                   3FFE:C00:E:B::2                        0 237 3748 1275 559 i
```

Table 4-10 shows each field description.

Table 11: show bgp ipv6 Fields

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.

Field	Description
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP.

The following is sample output from the show bgp ipv6 command, showing information for prefix 3FFE:500::/24:

```
ciscoasa# show bgp ipv6 unicast 3FFE:500::/24
BGP routing table entry for 3FFE:500::/24, version 19421
Paths: (6 available, best #1)
  293 3425 2500
    3FFE:700:20:1::11 from 3FFE:700:20:1::11 (192.168.2.27)
      Origin IGP, localpref 100, valid, external, best
  4554 293 3425 2500
    3FFE:C00:E:4::2 from 3FFE:C00:E:4::2 (192.168.1.1)
      Origin IGP, metric 1, localpref 100, valid, external
  33 293 3425 2500
    3FFE:C00:E:5::2 from 3FFE:C00:E:5::2 (209.165.18.254)
      Origin IGP, localpref 100, valid, external
  6175 7580 2500
    3FFE:C00:E:1::2 from 3FFE:C00:E:1::2 (209.165.223.204)
      Origin IGP, localpref 100, valid, external
1849 4697 2500, (suppressed due to dampening)
    3FFE:1100:0:CC00::1 from 3FFE:1100:0:CC00::1 (172.31.38.102)
      Origin IGP, localpref 100, valid, external
237 10566 4697 2500
    3FFE:C00:E:B::2 from 3FFE:C00:E:B::2 (172.31.0.3)
      Origin IGP, localpref 100, valid, external
ciscoasa# show bgp ipv6 unicast
BGP table version is 28, local router ID is 172.10.10.1
Status codes:s suppressed, h history, * valid, > best, i -
internal,
              r RIB-failure, S Stale
Origin codes:i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*>i4004::/64        ::FFFF:172.11.11.1
                                             0    100      0 ?
* i                 ::FFFF:172.30.30.1
                                             0    100      0 ?
```

show bgp ipv6 community

To display entries in the IPv6 Border Gateway Protocol (BGP) routing table, use the show bgp ipv6community command in user EXEC or privileged EXEC mode.

show bgp ipv6 unicast community [*community-number*] [**exact-match**] [**local-as | no-advertise | no-export**]

Syntax Description

unicast	*Specifies IPv6 unicast address prefixes.*
community-number	(Optional) Valid value is a community number in the range from 1 to 4294967295 or AA:NN (autonomous system-community number:2-byte number).
exact-match	(Optional) Displays only routes that have an exact match.
local-as	(Optional) Displays only routes that are not sent outside of the local autonomous system (well-known community).
no-advertise	(Optional) Displays only routes that are not advertised to any peer (well-known community).
no-export	(Optional) Displays only routes that are not exported outside of the local autonomous system (well-known community).

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.3(2)	This command was added

Examples

The show bgp ipv6 community command provides output similar to the show ip bgp community command, except it is IPv6-specific.

Communities are set with the set community route-map configuration command. You must enter the numerical communities before the well-known communities. For example, the following string is not valid:

```
ciscoasa# show ipv6 bgp unicast community local-as 111:12345
```

Use following strings instead:

```
ciscoasa# show ipv6 bgp unicast community 111:12345 local-as
```

Examples

The following is sample output from the show bgp ipv6 community command:

```
BGP table version is 69, local router ID is 10.2.64.5
Status codes:s suppressed, h history, * valid, > best, i - internal
Origin codes:i - IGP, e - EGP, ? - incomplete
       Network              Next Hop              Metric LocPrf Weight Path
*> 2001:0DB8:0:1::1/64       ::                                   0 32768 i
*> 2001:0DB8:0:1:1::/80      ::                                   0 32768 ?
*> 2001:0DB8:0:2::/64        2001:0DB8:0:3::2                        0 2 i
*> 2001:0DB8:0:2:1::/80      2001:0DB8:0:3::2                        0 2 ?
* 2001:0DB8:0:3::1/64        2001:0DB8:0:3::2                        0 2 ?
*>                         ::                                   0 32768 ?
*> 2001:0DB8:0:4::/64        2001:0DB8:0:3::2                        0 2 ?
*> 2001:0DB8:0:5::1/64       ::                                   0 32768 ?
*> 2001:0DB8:0:6::/64        2000:0:0:3::2                        0 2 3 i
*> 2010::/64                 ::                                   0 32768 ?
*> 2020::/64                 ::                                   0 32768 ?
*> 2030::/64                 ::                                   0 32768 ?
*> 2040::/64                 ::                                   0 32768 ?
*> 2050::/64                 ::                                   0 32768 ?
```

Table 12: show bgp ipv6 community fields

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	A 32-bit number written as 4 octets separated by periods (dotted-decimal format).
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.

Field	Description
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP.

show bgp ipv6 community-list

To display routes that are permitted by the IPv6 Border Gateway Protocol (BGP) community list, use the show bgp ipv6 community-list command in user EXEC or privileged EXEC mode.

show bgp ipv6 unicast community-list { *number* | *name* } [**exact-match**]

Syntax Description

unicast	*Specifies IPv6 unicast address prefixes.*
number	Community list number in the range from 1 to 199.
name	Community list name.
exact-match	(Optional) Displays only routes that have an exact match.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.3(2)	This command was added

Examples

The show bgp ipv6 unicast community-list command provide output similar to the show ip bgp community-list command, except they are IPv6-specific.

Examples

The following is sample output of the show bgp ipv6 community-list command for community list number 3:

```
ciscoasa# show bgp ipv6 unicast community-list 3
BGP table version is 14, local router ID is 10.2.64.6
Status codes:s suppressed, h history, * valid, > best, i - internal
Origin codes:i - IGP, e - EGP, ? - incomplete
    Network                Next Hop            Metric LocPrf Weight Path
*> 2001:0DB8:0:1::/64      2001:0DB8:0:3::1                         0 1 i
*> 2001:0DB8:0:1:1::/80    2001:0DB8:0:3::1                         0 1 i
*> 2001:0DB8:0:2::1/64     ::                                   0 32768 i
*> 2001:0DB8:0:2:1::/80    ::                                   0 32768 ?
* 2001:0DB8:0:3::2/64      2001:0DB8:0:3::1                         0 1 ?
*>                       ::                                   0 32768 ?
*> 2001:0DB8:0:4::2/64     ::                                   0 32768 ?
```

```
*> 2001:0DB8:0:5::/64       2001:0DB8:0:3::1                        0 1 ?
*> 2010::/64                2001:0DB8:0:3::1                        0 1 ?
*> 2020::/64                2001:0DB8:0:3::1                        0 1 ?
*> 2030::/64                2001:0DB8:0:3::1                        0 1 ?
*> 2040::/64                2001:0DB8:0:3::1                        0 1 ?
*> 2050::/64                2001:0DB8:0:3::1                        0 1 ?
```

Table below describes the significant fields shown in the display.

Table 13: show bgp ipv6 community-list fields

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	A 32-bit number written as 4 octets separated by periods (dotted-decimal format).
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.

Field	Description
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP.

show bgp ipv6 filter-list

To display routes that conform to a specified IPv6 filter list, use the show bgp ipv6 filter-list command in user EXEC or privileged EXEC mode.

show bgp ipv6 unicast filter-list *access-list-number*

Syntax Description

unicast	*Specifies IPv6 unicast address prefixes.*
>*access-list-number*	Number of an IPv6 autonomous system path access list. It can be a number from 1 to 199.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.3(2)	This command was added

Examples

The show bgp ipv6 filter-list command provides output similar to the show ip bgp filter-list command, except that it is IPv6-specific.

Examples:

The following is sample output from the show bgp ipv6 filter-list command for IPv6 autonomous system path access list number 1:

```
ciscoasa# show bgp ipv6 unicast filter-list 1
BGP table version is 26, local router ID is 192.168.0.2
Status codes:s suppressed, h history, * valid, > best, i - internal
Origin codes:i - IGP, e - EGP, ? - incomplete
    Network              Next Hop                          Metric LocPrf Weight Path
*> 2001:0DB8:0:1::/64        2001:0DB8:0:4::2                         0 2 1 i
*> 2001:0DB8:0:1:1::/80      2001:0DB8:0:4::2                         0 2 1 i
*> 2001:0DB8:0:2:1::/80      2001:0DB8:0:4::2                         0 2 ?
*> 2001:0DB8:0:3::/64        2001:0DB8:0:4::2                         0 2 ?
*> 2001:0DB8:0:4::/64        ::                                         32768  ?
*                            2001:0DB8:0:4::2                         0 2 ?
*> 2001:0DB8:0:5::/64        ::                                         32768  ?
*                            2001:0DB8:0:4::2                         0 2 1 ?
*> 2001:0DB8:0:6::1/64       ::                                         32768  i
*> 2030::/64                 2001:0DB8:0:4::2                         0 1
```

```
*> 2040::/64                2001:0DB8:0:4::2                        0 2 1 ?
*> 2050::/64                2001:0DB8:0:4::2                        0 2 1 ?
Table below describes the significant fields shown in the display.
```

Table 14: show bgp ipv6 community-list fields

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	A 32-bit number written as 4 octets separated by periods (dotted-decimal format).
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.

Field	Description
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP.

show bgp ipv6 inconsistent-as

To display IPv6 Border Gateway Protocol (BGP) routes with inconsistent originating autonomous systems, use the show bgp ipv6 inconsistent-as command in user EXEC or privileged EXEC mode.

show bgp ipv6 unicast inconsistent-as

Syntax Description

unicast	*Specifies IPv6 unicast address prefixes.*

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.3(2)	This command was added

Examples

The show bgp ipv6 unicast inconsistent-as command provide output similar to the show ip bgp inconsistent-as command, except they are IPv6-specific.

Examples

The following is sample output from the show bgp ipv6 inconsistent-as command:

```
ciscoasa# show bgp ipv6 unicast inconsistent-as
BGP table version is 12612, local router ID is 192.168.7.225
Status codes: s suppressed, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*  3FFE:1300::/24   2001:0DB8:0:F004::1                  0 3320 293 6175 ?
*                   3FFE:C00:E:9::2                     0 1251 4270 10318 ?
*                   3FFE:3600::A                        0 3462 6175 ?
*                   3FFE:700:20:1::11                   0 293 6175 ?
Table 4-15 below describes the significant fields shown in the display.
```

Table 15: show bgp ipv6 community-list fields

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.

Field	Description
local router ID	A 32-bit number written as 4 octets separated by periods (dotted-decimal format).
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP.

show bgp ipv6 neighbors

To display information about IPv6 Border Gateway Protocol (BGP) connections to neighbors, use the **show bgp ipv6 neighbors** command in user EXEC or privileged EXEC mode.

show bgp ipv6 unicast neighbors [*ipv6-address*] [**received-routes | routes | advertised-routes | paths** *regular-expression*]

Syntax Description

unicast	*Specifies IPv6 unicast address prefixes.*
ipv6-address	(Optional) Address of the IPv6 BGP-speaking neighbor. If you omit this argument, all IPv6 neighbors are displayed. This argument must be in the form documented in RFC 2373 where the address is specified in hexadecimal using 16-bit values between colons.
received-routes	(Optional) Displays all received routes (both accepted and rejected) from the specified neighbor.
routes	(Optional) Displays all routes received and accepted. This is a subset of the output from the **received-routes** keyword.
advertised-routes	(Optional) Displays all the routes the networking device advertised to the neighbor.
paths *regular-expression*	(Optional) Regular expression used to match the paths received.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.3(2)	This **command was added.**

Examples

The **show bgp ipv6 unicast neighbors** provide output similar to the **show ip bgp neighbors** command, except they are IPv6-specific.

Examples

The following is sample output from the **show bgp ipv6 neighbors** command:

```
ciscoasa# show bgp ipv6 unicast neighbors
BGP neighbor is 3FFE:700:20:1::11,  remote AS 65003, external link
  BGP version 4, remote router ID 192.168.2.27
  BGP state = Established, up for 13:40:17
  Last read 00:00:09, hold time is 180, keepalive interval is 60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 31306 messages, 20 notifications, 0 in queue
  Sent 14298 messages, 1 notifications, 0 in queue
  Default minimum time between advertisement runs is 30 seconds
 For address family: IPv6 Unicast
  BGP table version 21880, neighbor version 21880
  Index 1, Offset 0, Mask 0x2
  Route refresh request: received 0, sent 0
  Community attribute sent to this neighbor
  Outbound path policy configured
  Incoming update prefix filter list is bgp-in
  Outgoing update prefix filter list is aggregate
  Route map for outgoing advertisements is uni-out
  77 accepted prefixes consume 4928 bytes
  Prefix advertised 4303, suppressed 0, withdrawn 1328
  Number of NLRIs in the update sent: max 1, min 0
  1 history paths consume 64 bytes
  Connections established 22; dropped 21
  Last reset 13:47:05, due to BGP Notification sent, hold time expired
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 3FFE:700:20:1::12, Local port: 55345
Foreign host: 3FFE:700:20:1::11, Foreign port: 179
Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)
Event Timers (current time is 0x1A0D543C):
Timer          Starts    Wakeups            Next
Retrans          1218          5             0x0
TimeWait            0          0             0x0
AckHold          3327       3051             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0
iss: 1805423033  snduna: 1805489354  sndnxt: 1805489354     sndwnd:  15531
irs:  821333727  rcvnxt:  821591465  rcvwnd:      15547  delrcvwnd:    837
SRTT: 300 ms, RTTO: 303 ms, RTV: 3 ms, KRTT: 0 ms
minRTT: 8 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: higher precedence, nagle
Datagrams (max data segment is 1420 bytes):
Rcvd: 4252 (out of order: 0), with data: 3328, total data bytes: 257737
Sent: 4445 (retransmit: 5), with data: 4445, total data bytes: 244128
```

The table below describes the significant fields shown in the display.

Table 16: show bgp ipv6 community-list fields

Field	Description
BGP neighbor	IP address of the BGP neighbor and its autonomous system number. If the neighbor is in the same autonomous system as the router, then the link between them is internal; otherwise, it is considered external.
remote AS	Autonomous system of the neighbor.

Field	Description
internal link	Indicates that this peer is an interior Border Gateway Protocol (iBGP) peer.
BGP version	BGP version being used to communicate with the remote router; the router ID (an IP address) of the neighbor is also specified.
remote router ID	A 32-bit number written as 4 octets separated by periods (dotted-decimal format).
BGP state	Internal state of this BGP connection.
up for	Amount of time that the underlying TCP connection has been in existence.
Last read	Time that BGP last read a message from this neighbor.
hold time	Maximum amount of time that can elapse between messages from the peer.
keepalive interval	Time period between sending keepalive packets, which help ensure that the TCP connection is up.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor.
Route refresh	Indicates that the neighbor supports dynamic soft reset using the route refresh capability.
Address family IPv6 Unicast	Indicates that BGP peers are exchanging IPv6 reachability information.
Received	Number of total BGP messages received from this peer, including keepalives.
notifications	Number of error messages received from the peer .
Sent	Total number of BGP messages that have been sent to this peer, including keepalives.
notifications	Number of error messages the router has sent to this peer.
advertisement runs	Value of the minimum advertisement interval.
For address family	Address family to which the following fields refer.
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
neighbor version	Number used by the software to track the prefixes that have been sent and those that must be sent to this neighbor.
Route refresh request	Number of route refresh requests sent and received from this neighbor. .

Field	Description
Community attribute (not shown in sample output)	Appears if the neighbor send-community command is configured for this neighbor. .
Inbound path policy (not shown in sample output)	Indicates whether an inbound filter list or route map is configured.
Outbound path policy (not shown in sample output)	Indicates whether an outbound filter list, route map, or unsuppress map is configured.
bgp-in (not shown in sample output)	Name of the inbound update prefix filter list for the IPv6 unicast address family.
aggregate (not shown in sample output)	Name of the outbound update prefix filter list for the IPv6 unicast address family.
uni-out (not shown in sample output)	Name of the outbound route map for the IPv6 unicast address family.
accepted prefixes	Number of prefixes accepted.
Prefix advertised	Number of prefixes advertised.
suppressed	Number of prefixes suppressed
withdrawn	Number of prefixes withdrawn.
history paths (not shown in sample output)	Number of path entries held to remember history.
Connections established	Number of times the router has established a TCP connection and the two peers have agreed to speak BGP with each other.
dropped	Number of times that a good connection has failed or been taken down.
Last reset	Elapsed time (in hours:minutes:seconds) since this peering session was last reset.
Connection state	State of the BGP Peer
unread input bytes	Number of bytes of packets still to be processed.
Local host, Local port	Peering address of the local router, plus the port.
Foreign host, Foreign port	Peering address of the neighbor.
Event Timers	Table that displays the number of starts and wakeups for each timer.
snduna	Last send sequence number for which the local host sent but has not received an acknowledgment.
sndnxt	Sequence number the local host will send next.
sndwnd	TCP window size of the remote host.

Field	Description
irs	Initial receive sequence number.
rcvnxt	Last receive sequence number the local host has acknowledged.
rcvwnd	TCP window size of the local host.
delrecvwnd	Delayed receive window--data the local host has read from the connection, but has not yet subtracted from the receive window the host has advertised to the remote host. The value in this field gradually increases until it is larger than a full-sized packet, at which point it is applied to the rcvwnd field.
SRTT	A calculated smoothed round-trip timeout (in milliseconds).
RTTO	Round-trip timeout (in milliseconds).
RTV	Variance of the round-trip time (in milliseconds).
KRTT	New round-trip timeout (in milliseconds) using the Karn algorithm. This field separately tracks the round-trip time of packets that have been re-sent.
minRTT	Smallest recorded round-trip timeout (in milliseconds) with hard wire value used for calculation.
maxRTT	Largest recorded round-trip timeout (in milliseconds).
ACK hold	Time (in milliseconds) the local host will delay an acknowledgment in order to "piggyback" data on it.
Flags	IP precedence of the BGP packets.
Datagrams: Rcvd	Number of update packets received from neighbor.
with data	Number of update packets received with data.
total data bytes	Total number of bytes of data.
Sent	Number of update packets sent.
with data	Number of update packets with data sent.
total data bytes	Total number of data bytes.

The following is sample output from the **show bgp ipv6 neighbors** command with the **advertised-routes** keyword:

```
ciscoasa# show bgp ipv6 unicast neighbors 3FFE:700:20:1::11 advertised-routes
BGP table version is 21880, local router ID is 192.168.7.225
Status codes: s suppressed, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 2001:200::/35    3FFE:700:20:1::11                       0 293 3425 2500 i
```

```
*> 2001:208::/35    3FFE:C00:E:B::2                          0 237 7610 i
*> 2001:218::/35    3FFE:C00:E:C::2                          0 3748 4697 i
```

b

The following is sample output from the **show bgp ipv6 neighbors** command with the **routes** keyword:

```
ciscoasa# show bgp ipv6 unicast neighbors 3FFE:700:20:1::11 routes
BGP table version is 21885, local router ID is 192.168.7.225
Status codes: s suppressed, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 2001:200::/35    3FFE:700:20:1::11                        0 293 3425 2500 i
*  2001:208::/35    3FFE:700:20:1::11                        0 293 7610 i
*  2001:218::/35    3FFE:700:20:1::11                        0 293 3425 4697 i
*  2001:230::/35    3FFE:700:20:1::11                        0 293 1275 3748 i
Table below describes the significant fields shown in the display.
```

Table 17: show bgp ipv6 neighbors advertised-routes and routes fields

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	A 32-bit number written as 4 octets separated by periods (dotted-decimal format).
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.

Field	Description
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP.

The following is sample output from the **show bgp ipv6 neighbors** command with the **paths** keyword:

```
ciscoasa# show bgp ipv6 unicast neighbors 3FFE:700:20:1::11 paths ^293
Address    Refcount Metric Path
0x6131D7DC        2      0 293 3425 2500 i
0x6132861C        2      0 293 7610 i
0x6131AD18        2      0 293 3425 4697 i
0x61324084        2      0 293 1275 3748 i
0x61320E0C        1      0 293 3425 2500 2497 i
0x61326928        1      0 293 3425 2513 i
0x61327BC0        2      0 293 i
0x61321758        1      0 293 145 i
0x61320BEC        1      0 293 3425 6509 i
0x6131AAF8        2      0 293 1849 2914 ?
0x61320FE8        1      0 293 1849 1273 209 i
0x613260A8        2      0 293 1849 i
0x6132586C        1      0 293 1849 5539 i
0x6131BBF8        2      0 293 1849 1103 i
0x6132344C        1      0 293 4554 1103 1849 1752 i
0x61324150        2      0 293 1275 559 i
0x6131E5AC        2      0 293 1849 786 i
0x613235E4        1      0 293 1849 1273 i
0x6131D028        1      0 293 4554 5539 8627 i
0x613279E4        1      0 293 1275 3748 4697 3257 i
0x61320328        1      0 293 1849 1273 790 i
0x6131EC0C        2      0 293 1275 5409 i
The table below describes the significant fields shown in the display.
```

show bgp ipv6 neighbors paths fields

Field	Description
Address	Internal address where the path is stored.
Refcount	Number of routes using that path.
Metric	The Multi Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.)
Path	The autonomous system path for that route, followed by the origin code for that route.

The following sample output from the **show bgp ipv6 neighbors** command shows the received routes for IPv6 address 2000:0:0:4::2:

```
ciscoasa# show bgp ipv6 unicast neighbors 2000:0:0:4::2 received-routes
BGP table version is 2443, local router ID is 192.168.0.2
Status codes:s suppressed, h history, * valid, > best, i - internal
Origin codes:i - IGP, e - EGP, ? - incomplete
Network               Next Hop           Metric LocPrf Weight Path
*> 2000:0:0:1::/64     2000:0:0:4::2                        0 2 1 i
*> 2000:0:0:2::/64     2000:0:0:4::2                        0 2 i
*> 2000:0:0:2:1::/80   2000:0:0:4::2                        0 2 ?
*> 2000:0:0:3::/64     2000:0:0:4::2                        0 2 ?
*  2000:0:0:4::1/64    2000:0:0:4::2                        0 2 ?
```

show bgp ipv6 paths

To display all the IPv6 Border Gateway Protocol (BGP) paths in the database, use the **show bgp ipv6 paths** command in user EXEC or privileged EXEC mode.

show bgp ipv6 unicast paths *regular-expression*

Syntax Description

unicast	Specifies IPv6 unicast address prefixes.
regular-expression	Regular expression that is used to match the received paths in the database.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.3(2)	This command was added

Examples

The **show bgp ipv6 unicast paths** command provide output similar to the **show ip bgp paths** command, except they are IPv6-specific.

Examples

The following is sample output from the **show bgp ipv6 paths** command:

```
ciscoasa# show bgp ipv6 unicast paths
Address    Hash Refcount Metric Path
0x61322A78    0        2      0 i
0x6131C214    3        2      0 6346 8664 786 i
0x6131D600   13        1      0 3748 1275 8319 1273 209 i
0x613229F0   17        1      0 3748 1275 8319 12853 i
0x61324AE0   18        1      1 4554 3748 4697 5408 i
0x61326818   32        1      1 4554 5609 i
0x61324728   34        1      0 6346 8664 9009 ?
0x61323804   35        1      0 3748 1275 8319 i
0x61327918   35        1      0 237 2839 8664 ?
0x61320504   38        2      0 3748 4697 1752 i
0x61320988   41        2      0 1849 786 i
0x6132245C   46        1      0 6346 8664 4927 i
Table below describes the significant fields shown in the display.
```

Field	Description
Address	Internal address where the path is stored.
Refcount	Number of routes using that path.
Metric	The Multi Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.)
Path	The autonomous system path for that route, followed by the origin code for that route.

show bgp ipv6 prefix-list

To display routes that match a prefix list, use the **show bgp ipv6 prefix-list** command in user EXEC or privileged EXEC mode.

show bgp ipv6 unicast prefix-list *name*

Syntax Description

unicast	Specifies IPv6 unicast address prefixes.
name	The specified prefix-list

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.3(2)	This command was added

Examples

The specified prefix list must be an IPv6 prefix list, which is similar in format to an IPv4 prefix list.

Example

```
The following is sample output from the show bgp ipv6 prefix-list
command:
Router# show bgp ipv6 unicast prefix-list pin
ipv6 prefix-list pin:
   count:4, range entries:3, sequences:5 - 20, refcount:2
   seq 5 permit 747::/16 (hit count:1, refcount:2)
   seq 10 permit 747:1::/32 ge 64 le 64 (hit count:2, refcount:2)
   seq 15 permit 747::/32 ge 33 (hit count:1, refcount:1)
   seq 20 permit 777::/16 le 124 (hit count:2, refcount:1)
The ipv6 prefix-list match the following prefixes:
   seq 5: matches the exact match 747::/16
   seq 10:first 32 bits in prefix must match with a prefixlen of /64
   seq 15:first 32 bits in prefix must match with any prefixlen up to /128
   seq 20:first 16 bits in prefix must match with any prefixlen up to /124
Table below describes the significant fields shown in the display.
```

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.

Field	Description
local router ID	A 32-bit number written as 4 octets separated by periods (dotted-decimal format).
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP.

show bgp ipv6 quote-regexp

To display IPv6 Border Gateway Protocol (BGP) routes matching the autonomous system path regular expression as a quoted string of characters, use the **show bgp ipv6 quote-regexp** command in user EXEC or privileged EXEC mode.

show bgp ipv6 unicast quote-regexp *regular expression*

Syntax Description

unicast	Specifies IPv6 unicast address prefixes.
regular expression	Regular expression that is used to match the BGP autonomous system paths

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.3(2)	This command was added

Examples

The **show bgp ipv6 unicast quote-regexp** command provide output similar to the **show ip bgp quote-regexp** command, except they are IPv6-specific.

Example

The following is sample output from the **show bgp ipv6 quote-regexp** command that shows paths beginning with 33 or containing 293:

```
Router# show bgp ipv6 unicast quote-regexp ^33|293
BGP table version is 69964, local router ID is 192.31.7.225
Status codes: s suppressed, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*  2001:200::/35    3FFE:C00:E:4::2          1             0 4554 293 3425 2500 i
*                   2001:0DB8:0:F004::1
                                                           0 3320 293 3425 2500 i
*  2001:208::/35    3FFE:C00:E:4::2          1             0 4554 293 7610 i
*  2001:228::/35    3FFE:C00:E:F::2                        0 6389 1849 293 2713 i
*  3FFE::/24        3FFE:C00:E:5::2                        0 33 1849 4554 i
*  3FFE:100::/24    3FFE:C00:E:5::2                        0 33 1849 3263 i
*  3FFE:300::/24    3FFE:C00:E:5::2                        0 33 293 1275 1717 i
* 3FFE:C00:E:F::2                        0 6389 1849 293 1275
Table below describes the significant fields shown in the display.
```

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	A 32-bit number written as 4 octets separated by periods (dotted-decimal format).
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP.

show bgp ipv6 regexp

To display IPv6 Border Gateway Protocol (BGP) routes matching the autonomous system path regular expression, use the **show bgp ipv6 regexp** command in user EXEC or privileged EXEC mode.

show bgp ipv6 unicast regexp *regular-expression*

Syntax Description

unicast	Specifies IPv6 unicast address prefixes.
regular-expression	Regular expression that is used to match the BGP autonomous system paths

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.3(2)	This command was added

Examples

The **show bgp ipv6 unicast regexp** command provide output similar to the **show ip bgp regexp** command, except they are IPv6-specific.

Example

The following is sample output from the **show bgp ipv6 regexp** command that shows paths beginning with 33 or containing 293:

```
Router# show bgp ipv6 unicast regexp ^33|293
BGP table version is 69964, local router ID is 192.168.7.225
Status codes: s suppressed, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*  2001:200::/35    3FFE:C00:E:4::2          1             0 4554 293 3425 2500 i
*                   2001:0DB8:0:F004::1
                                                           0 3320 293 3425 2500 i
*  2001:208::/35    3FFE:C00:E:4::2          1             0 4554 293 7610 i
*  2001:228::/35    3FFE:C00:E:F::2                        0 6389 1849 293 2713 i
*  3FFE::/24        3FFE:C00:E:5::2                        0 33 1849 4554 i
*  3FFE:100::/24    3FFE:C00:E:5::2                        0 33 1849 3263 i
*  3FFE:300::/24    3FFE:C00:E:5::2                        0 33 293 1275 1717 i
*                   3FFE:C00:E:F::2                        0 6389 1849 293 1275
Table below describes the significant fields shown in the display.
```

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	A 32-bit number written as 4 octets separated by periods (dotted-decimal format).
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP.

show bgp ipv6 route-map

To display IPv6 Border Gateway Protocol (BGP) routes that failed to install in the routing table, use the **show bgp ipv6 route-map** command in user EXEC or privileged EXEC mode.

show bgp ipv6 unicast route-map *name*

Syntax Description

unicast	Specifies IPv6 unicast address prefixes.
name	A specified route map to match.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.3(2)	This command was added

Examples

The following is sample output from the **show bgp ipv6 route-map** command for a route map named rmap:

```
Router# show bgp ipv6 unicast route-map rmap
BGP table version is 16, local router ID is 172.30.242.1
Status codes:s suppressed, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes:i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*>i12:12::/64       2001:0DB8:101::1         0    100     50 ?
*>i12:13::/64       2001:0DB8:101::1         0    100     50 ?
*>i12:14::/64       2001:0DB8:101::1         0    100     50 ?
*>i543::/64         2001:0DB8:101::1         0    100     50 ?
```

The table below describes the significant fields shown in the display:

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	A 32-bit number written as 4 octets separated by periods (dotted-decimal format).

Field	Description
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP.

show bgp ipv6 summary

To display the status of all IPv6 Border Gateway Protocol (BGP) connections, use the **show bgp ipv6 summary** command in user EXEC or privileged EXEC mode.

show bgp ipv6 unicast summary

Syntax Description

unicast	Specifies IPv6 unicast address prefixes.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.3(2)	This command was added

Examples

The **show bgp ipv6 unicast summary** command provides output similar to the **show ip bgp summary** command, except they are IPv6-specific.

Examples

The following is sample output from the **show bgp ipv6 summary** command:

```
ciscoasa# show bgp ipv6 unicast summary
BGP device identifier 172.30.4.4, local AS number 200
BGP table version is 1, main routing table version 1
Neighbor           V    AS  MsgRcvd  MsgSent   TblVer  InQ  OutQ  Up/Down   State/PfxRcd
2001:0DB8:101::2   4    200    6869     6882        0    0     0  06:25:24  Active
```

The table below describes the significant fields shown in the display.

Field	Description
BGP device identifier	IP address of the networking device.
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
main routing table version	Last version of BGP database that was injected into the main routing table.

Field	Description
Neighbor	IPv6 address of a neighbor.
V	BGP version number spoken to that neighbor.
AS	Autonomous System
MsgRcvd	BGP messages received from that neighbor.
MsgSent	BGP messages sent to that neighbor
TblVer	Last version of the BGP database that was sent to that neighbor.
InQ	Number of messages from that neighbor waiting to be processed.
OutQ	Number of messages waiting to be sent to that neighbor.
Up/Down	The length of time that the BGP session has been in state Established, or the current state if it is not Established.
State/PfxRcd	Current state of the BGP session/the number of prefixes the device has received from a neighbor. When the maximum number (as set by the **neighbor maximum-prefix** command) is reached, the string "PfxRcd" appears in the entry, the neighbor is shut down, and the connection is Idle. An (Admin) entry with Idle status indicates that the connection has been shut down using the **neighbor shutdown** command.

show bgp neighbors

To display information about Border Gateway Protocol (BGP) and TCP connections to neighbors, use the show bgp neighbors command in user or privileged EXEC mode.

show bgp neighbors [**slow** | *ip-address* [**advertised-routes | | paths** [*reg-exp* **| policy** [**detail**] **| received prefix-filter | received-routes | routes**]]

Syntax Description

slow	(Optional) Displays information about dynamically configured slow peers
ip-address	(Optional) Displays information about the IPv4 neighbor. If this argument is omitted, information about all neighbors is displayed.
advertised-routes	(Optional) Displays all routes that have been advertised to neighbors.
paths reg-exp	(Optional) Displays autonomous system paths learned from the specified neighbor. An optional regular expression can be used to filter the output.
policy	(Optional) Displays the policies applied to this neighbor per address family.
detail	(Optional) Displays detailed policy information such as route maps, prefix lists, community lists, access control lists (ACLs), and autonomous system path filter lists.
received prefix-filter	(Optional) Displays the prefix-list (outbound route filter [ORF]) sent from the specified neighbor.
received-routes	(Optional) Displays all received routes (both accepted and rejected) from the specified neighbor.
routes	(Optional) Displays all routes that are received and accepted. The output displayed when this keyword is entered is a subset of the output displayed by the received-routes keyword.

Command Default

The output of this command displays information for all neighbors.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Usage Guidelines

Use the show bgp neighbors command to display BGP and TCP connection information for neighbor sessions. For BGP, this includes detailed neighbor attribute, capability, path, and prefix information. For TCP, this includes statistics related to BGP neighbor session establishment and maintenance.

Prefix activity is displayed based on the number of prefixes that are advertised and withdrawn. Policy denials display the number of routes that were advertised but then ignored based on the function or attribute that is displayed in the output.

Cisco implementation of 4-byte autonomous system numbers uses asplain—65538 for example—as the default regular expression match and output display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain format and the asdot format as described in RFC 5396. To change the default regular expression match and output display of 4-byte autonomous system numbers to asdot format, use the bgp asnotation dot command followed by the clear bgp * command to perform a hard reset of all current BGP sessions.

Examples

Example output is different for the various keywords available for the show bgp neighbors command. Examples using the various keywords appear in the following sections:

show bgp neighbors: Example

The following example shows output for the BGP neighbor at 10.108.50.2. This neighbor is an internal BGP (iBGP) peer. This neighbor supports the route refresh and graceful restart capabilities.

```
ciscoasa# show bgp neighbors 10.108.50.2
BGP neighbor is 10.108.50.2,  remote AS 1, internal link
  BGP version 4, remote router ID 192.168.252.252
  BGP state = Established, up for 00:24:25
  Last read 00:00:24, last write 00:00:24, hold time is 180, keepalive interval is
   60 seconds
  Neighbor capabilities:
    Route refresh: advertised and received(old & new)
    MPLS Label capability: advertised and received
    Graceful Restart Capability: advertised
    Address family IPv4 Unicast: advertised and received
  Message statistics:
    InQ depth is 0
    OutQ depth is 0
                         Sent       Rcvd
    Opens:                  3          3
    Notifications:          0          0
    Updates:                0          0
    Keepalives:           113        112
    Route Refresh:          0          0
    Total:                116        115
  Default minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP additional-paths computation is enabled
  BGP advertise-best-external is enabled
  BGP table version 1, neighbor version 1/0
 Output queue size : 0
  Index 1, Offset 0, Mask 0x2
  1 update-group member
                               Sent       Rcvd
  Prefix activity:             ----       ----
    Prefixes Current:             0          0
    Prefixes Total:               0          0
    Implicit Withdraw:            0          0
    Explicit Withdraw:            0          0
    Used as bestpath:           n/a          0
    Used as multipath:          n/a          0
```

```
                                  Outbound    Inbound
  Local Policy Denied Prefixes:   --------    -------
    Total:                               0          0
  Number of NLRIs in the update sent: max 0, min 0
  Connections established 3; dropped 2
  Last reset 00:24:26, due to Peer closed the session
External BGP neighbor may be up to 2 hops away.
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled
Local host: 10.108.50.1, Local port: 179
Foreign host: 10.108.50.2, Foreign port: 42698
Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)
Event Timers (current time is 0x68B944):
Timer          Starts    Wakeups            Next
Retrans            27          0             0x0
TimeWait            0          0             0x0
AckHold            27         18             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0
iss: 3915509457  snduna: 3915510016  sndnxt: 3915510016     sndwnd:  15826
irs:  233567076  rcvnxt:  233567616  rcvwnd:      15845  delrcvwnd:    539
SRTT: 292 ms, RTTO: 359 ms, RTV: 67 ms, KRTT: 0 ms
minRTT: 12 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: passive open, nagle, gen tcbs
IP Precedence value : 6
Datagrams (max data segment is 1460 bytes):
Rcvd: 38 (out of order: 0), with data: 27, total data bytes: 539
Sent: 45 (retransmit: 0, fastretransmit: 0, partialack: 0, Second Congestion: 08
```

Below table describes the significant fields shown in the display. Fields that are preceded by the asterisk character (*) are displayed only when the counter has a nonzero value.

Table 4-10 shows each field description.

Table 18: show bgp ipv4 Fields

Field	Description
BGP neighbor	IP address of the BGP neighbor and its autonomous system number.
remote AS	Autonomous system number of the neighbor.
local AS 300 no-prepend (not shown in display)	Verifies that the local autonomous system number is not prepended to received external routes. This output supports the hiding of the local autonomous systems when migrating autonomous systems.
internal link	"internal link" is displayed for iBGP neighbors. "external link" is displayed for external BGP (eBGP) neighbors.
BGP version	BGP version being used to communicate with the remote router.
remote router ID	IP address of the neighbor.
BGP state	Finite state machine (FSM) stage of session negotiation.

Field	Description
up for	Time, in hhmmss, that the underlying TCP connection has been in existence.
Last read	Time, in hhmmss, since BGP last received a message from this neighbor.
last write	Time, in hhmmss, since BGP last sent a message to this neighbor.
hold time	Time, in seconds, that BGP will maintain the session with this neighbor without receiving a messages.
keepalive interval	Time interval, in seconds, at which keepalive messages are transmitted to this neighbor.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor. "advertised and received" is displayed when a capability is successfully exchanged between two routers
Route Refresh	Status of the route refresh capability.
Graceful Restart Capability	Status of the graceful restart capability.
Address family IPv4 Unicast	IP Version 4 unicast-specific properties of this neighbor.
Message statistics	Statistics organized by message type.
InQ depth is	Number of messages in the input queue.
OutQ depth is	Number of messages in the output queue.
Sent	Total number of transmitted messages.
Received	Total number of received messages.
Opens	Number of open messages sent and received.
notifications	Number of notification (error) messages sent and received.
Updates	Number of update messages sent and received.
Keepalives	Number of keepalive messages sent and received.
Route Refresh	Number of route refresh request messages sent and received.
Total	Total number of messages sent and received.
Default minimum time between...	Time, in seconds, between advertisement transmissions.
For address family:	Address family to which the following fields refer.
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.

Field	Description
neighbor version	Number used by the software to track prefixes that have been sent and those that need to be sent.
update-group	Number of update-group member for this address family
Prefix activity	Prefix statistics for this address family.
Prefixes current	Number of prefixes accepted for this address family.
Prefixes total	Total number of received prefixes.
Implicit Withdraw	Number of times that a prefix has been withdrawn and readvertised.
Explicit Withdraw	Number of times that prefix has been withdrawn because it is no longer feasible.
Used as bestpath	Number of received prefixes installed as bestpaths.
Used as multipath	Number of received prefixes installed as multipaths.
* Saved (soft-reconfig)	Number of soft resets performed with a neighbor that supports soft reconfiguration. This field is displayed only if the counter has a nonzero value.
* History paths	This field is displayed only if the counter has a nonzero value.
* Invalid paths	Number of invalid paths. This field is displayed only if the counter has a nonzero value.
Local Policy Denied Prefixes	Prefixes denied due to local policy configuration. Counters are updated for inbound and outbound policy denials. The fields under this heading are displayed only if the counter has a nonzero value.
* route-map	Displays inbound and outbound route-map policy denials.
* filter-list	Displays inbound and outbound filter-list policy denials.
* prefix-list	Displays inbound and outbound prefix-list policy denials.
* AS_PATH too long	Displays outbound AS-path length policy denials.
* AS_PATH loop	Displays outbound AS-path loop policy denials.
* AS_PATH confed info	Displays outbound confederation policy denials.
* AS_PATH contains AS 0	Displays outbound denials of autonomous system (AS) 0.
* NEXT_HOP Martian	Displays outbound martian denials.
* NEXT_HOP non-local	Displays outbound non-local next-hop denials.
* NEXT_HOP is us	Displays outbound next-hop-self denials.
* CLUSTER_LIST loop	Displays outbound cluster-list loop denials.

Field	Description
* ORIGINATOR loop	Displays outbound denials of local originated routes.
* unsuppress-map	Displays inbound denials due to an unsuppress-map.
* advertise-map	Displays inbound denials due to an advertise-map.
* Well-known Community	Displays inbound denials of well-known communities.
* SOO loop	Displays inbound denials due to site-of-origin.
* Bestpath from this peer	Displays inbound denials because the bestpath came from the local router.
* Suppressed due to dampening	Displays inbound denials because the neighbor or link is in a dampening state.
* Bestpath from iBGP peer	Deploys inbound denials because the bestpath came from an iBGP neighbor.
* Incorrect RIB for CE	Deploys inbound denials due to RIB errors for a CE router.
* BGP distribute-list	Displays inbound denials due to a distribute list.
Number of NLRIs...	Number of network layer reachability attributes in updates.
Connections established	Number of times a TCP and BGP connection has been successfully established.
dropped	Number of times that a valid session has failed or been taken down.
Last reset	Time since this peering session was last reset. The reason for the reset is displayed on this line.
External BGP neighbor may be... (not shown in the display)	Indicates that the BGP TTL security check is enabled. The maximum number of hops that can separate the local and remote peer is displayed on this line.
Connection state	Connection status of the BGP peer.
Connection is ECN Disabled	Explicit congestion notification status (enabled or disabled).
Local host: 10.108.50.1, Local port: 179	IP address of the local BGP speaker. BGP port number 179.
Foreign host: 10.108.50.2, Foreign port: 42698	Neighbor address and BGP destination port number.
Enqueued packets for retransmit:	Packets queued for retransmission by TCP.
Event Timers	TCP event timers. Counters are provided for starts and wakeups (expired timers).
Retrans	Number of times a packet has been retransmitted.

Field	Description
TimeWait	Time waiting for the retransmission timers to expire.
AckHold	Acknowledgment hold timer.
SendWnd	Transmission (send) window.
KeepAlive	Number of keepalive packets.
GiveUp	Number times a packet is dropped due to no acknowledgment.
PmtuAger	Path MTU discovery timer
DeadWait	Expiration timer for dead segments.
iss:	Initial packet transmission sequence number.
snduna	Last transmission sequence number that has not been acknowledged.
sndnxt:	Next packet sequence number to be transmitted.
sndwnd:	TCP window size of the remote neighbor.
irs:	Initial packet receive sequence number.
rcvnxt:	Last receive sequence number that has been locally acknowledged.
rcvwnd:	TCP window size of the local host.
delrcvwnd:	Delayed receive window—data the local host has read from the connection, but has not yet subtracted from the receive window the host has advertised to the remote host. The value in this field gradually increases until it is larger than a full-sized packet, at which point it is applied to the rcvwnd field.
SRTT:	A calculated smoothed round-trip timeout.
RTTO:	Round-trip timeout.
RTV:	Variance of the round-trip time.
KRTT:	New round-trip timeout (using the Karn algorithm). This field separately tracks the round-trip time of packets that have been re-sent.
minRTT:	Smallest recorded round-trip timeout (hard-wire value used for calculation).
maxRTT:	Largest recorded round-trip timeout.
ACK hold:	Length of time the local host will delay an acknowledgment to carry (piggyback) additional data.
IP Precedence value:	IP precedence of the BGP packets.
Datagrams	Number of update packets received from a neighbor.

Field	Description
Rcvd:	Number of received packets.
with data	Number of update packets sent with data.
total data bytes	Total amount of data received, in bytes.
Sent	Number of update packets sent.
Second Congestion	Number of second retransmissions sent due to congestion.
Datagrams: Rcvd	Number of update packets received from a neighbor.
out of order:	Number of packets received out of sequence.
with data	Number of update packets received with data.
Last reset	Elapsed time since this peering session was last reset.
unread input bytes	Number of bytes of packets still to be processed.
retransmit	Number of packets retransmitted.
fastretransmit	Number of duplicate acknowledgments retransmitted for an out of order segment before the retransmission timer expires.
partialack	Number of retransmissions for partial acknowledgments (transmissions before or without subsequent acknowledgments).

show bgp neighbors advertised-routes: Example

The following example displays routes advertised for only the 172.16.232.178 neighbor:

```
ciscoasa# show bgp neighbors 172.16.232.178 advertised-routes
BGP table version is 27, local router ID is 172.16.232.181
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
Network          Next Hop          Metric LocPrf Weight Path
*>i10.0.0.0      172.16.232.179         0    100      0 ?
*> 10.20.2.0     10.0.0.0               0         32768 i
```

Table 4-19 shows each field description.

Table 19: show bgp neighbors advertised routes Fields

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.

Field	Description
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	Internet address of the network the entry describes.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the access server has some non-BGP route to this network.
Metric	If shown, the value of the inter autonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path. At the end of the path is the origin code for the path: i—The entry was originated with the IGP and advertised with a network router configuration command. e—The route originated with EGP. ?—The origin of the path is not clear. Usually this is a path that is redistributed into BGP from an IGP..

Examples

show bgp neighbors paths: Example

The following is example output from the show bgp neighbors command entered with the paths keyword:

```
ciscoasa# show bgp neighbors 172.29.232.178 paths ^10
Address    Refcount Metric Path
0x60E577B0        2     40 10 ?
```

Table 4-20 shows each field description.

Table 20: show bgp neighbors paths Fields

Field	Description
Address	Internal address where the path is stored.
Refcount	Number of routes using that path..
Metric	Multi Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.).
Path	Autonomous system path for that route, followed by the origin code for that route..

Examples

show bgp neighbors received prefix-filter: Example

The following example shows that a prefix-list that filters all routes in the 10.0.0.0 network has been received from the 192.168.20.72 neighbor:

```
ciscoasa# show bgp neighbors 192.168.20.72 received prefix-filter
Address family:IPv4 Unicast
ip prefix-list 192.168.20.72:1 entries
   seq 5 deny 10.0.0.0/8 le 32
```

Table 4-21 shows each field description.

Table 21: show bgp neighbors received prefix filter Fields

Field	Description
Address family	Address family mode in which the prefix filter is received.
ip prefix-list	Prefix list sent from the specified neighbor.

Examples

show bgp neighbors policy: Example

The following sample output shows the policies applied to the neighbor at 192.168.1.2. The output displays policies configured on the neighbor device.

```
ciscoasa# show bgp neighbors 192.168.1.2 policy
Neighbor: 192.168.1.2, Address-Family: IPv4 Unicast
Locally configured policies:
 route-map ROUTE in
Inherited polices:
 prefix-list NO-MARKETING in
 route-map ROUTE in
 weight 300
 maximum-prefix 10000
```

show bgp neighbors: Example

The following is sample output from the show bgp neighbors command that verifies that BGP TCP path maximum transmission unit (MTU) discovery is enabled for the BGP neighbor at 172.16.1.2:

```
ciscoasa# show bgp neighbors 172.16.1.2
BGP neighbor is 172.16.1.2,  remote AS 45000, internal link
  BGP version 4, remote router ID 172.16.1.99
.
.
.
 For address family: IPv4 Unicast
  BGP table version 5, neighbor version 5/0
.
.
.
  Address tracking is enabled, the RIB does have a route to 172.16.1.2
  Address tracking requires at least a /24 route to the peer
  Connections established 3; dropped 2
  Last reset 00:00:35, due to Router ID changed
  Transport(tcp) path-mtu-discovery is enabled
.
.
.
SRTT: 146 ms, RTTO: 1283 ms, RTV: 1137 ms, KRTT: 0 ms
minRTT: 8 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: higher precedence, retransmission timeout, nagle, path mtu capable
```

The following is partial output from the show bgp neighbors command that verifies the status of the BGP graceful restart capability for the external BGP peer at 192.168.3.2. Graceful restart is shown as disabled for this BGP peer.

```
ciscoasa# show bgp neighbors 192.168.3.2
BGP neighbor is 192.168.3.2,  remote AS 50000, external link
 Inherits from template S2 for session parameters
  BGP version 4, remote router ID 192.168.3.2
  BGP state = Established, up for 00:01:41
  Last read 00:00:45, last write 00:00:45, hold time is 180, keepalive intervals
  Neighbor sessions:
    1 active, is multisession capable
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Address family IPv4 Unicast: advertised and received
.
.
.
Address tracking is enabled, the RIB does have a route to 192.168.3.2
  Connections established 1; dropped 0
  Last reset never
  Transport(tcp) path-mtu-discovery is enabled
  Graceful-Restart is disabled
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
```

show bgp paths

To display all the BGP paths in the database, use the show bgp paths command in EXEC mode.

show bgp paths
Cisco 10000 Series Router
show bgp paths *regexp*

Syntax Description

regexp	Regular expression to match the BGP autonomous system paths.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Examples

The following is sample output from the show bgp paths command in privileged EXEC mode:

```
ciscoasa# show bgp paths
Address    Hash Refcount Metric Path
0x60E5742C    0        1      0 i
0x60E3D7AC    2        1      0 ?
0x60E5C6C0   11        3      0 10 ?
0x60E577B0   35        2     40 10 ?
```

Table 4-22 shows each field description.

Table 22: show bgp paths Fields

Field	Description
Address	Internal address where the path is stored.
Hash	Hash bucket where path is stored.
Refcount	Number of routes using that path.
Metric	The Multi Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.)

Field	Description
Path	The autonomous system path for that route, followed by the origin code for that route.

show bgp policy-list

To display information about a configured policy list and policy list entries, use the show bgp policy-list command in user EXEC mode.

show bgp policy-list [*policy-list-name*]

Syntax Description

policy-list-name	(Optional) Displays information about the specified policy list with this argument.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Examples

The following is sample output from the show bgp policy-list command. The output of this command will display the policy-list name and configured match clauses. The following sample output is similar to the output that will be displayed:

```
ciscoasa# show bgp policy-list
policy-list POLICY-LIST-NAME-1 permit
  Match clauses:
    metric 20
policy-list POLICY-LIST-NAME-2 permit
  Match clauses:
    as-path (as-path filter): 1
```

show bgp prefix-list

To display information about a prefix list or prefix list entries, use the show bgp prefix-list command in user or privileged EXEC mode

show bgp prefix-list [**detail | summary**] [*prefix-list-name* [**seq** *sequence-number* | *network/length* [**longer**|**first-match**]]]

Syntax Description

detail | summary	(Optional) Displays detailed or summarized information about all prefix lists.
first-match	(Optional) Displays the first entry of the specified prefix list that matches the given network/length.
longer	(Optional) Displays all entries of the specified prefix list that match or are more specific than the given network/length.
network/length	(Optional) Displays all entries in the specified prefix list that use this network address and netmask length (in bits).
prefix-list-name	(Optional) Displays the entries in a specific prefix list.
seq sequence-number	(Optional) Displays only the prefix list entry with the specified sequence number in the specified prefix-list.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Examples

The following example shows the output of the show bgp prefix-list command with details about the prefix list named test:

```
ciscoasa# show bgp prefix-list detail test
ip prefix-list test:
Description: test-list
 count: 1, range entries: 0, sequences: 10 - 10, refcount: 3
 seq 10 permit 10.0.0.0/8 (hit count: 0, refcount: 1)
```

show bgp regexp

To display routes matching the autonomous system path regular expression, use the show bgp regexp command in EXEC mode.

show bgp regexp *regexp*

Syntax Description

regexp	Regular expression to match the BGP autonomous system paths. For more details about autonomous system number formats, see the router bgp command.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Usage Guidelines

Cisco implementation of 4-byte autonomous system numbers uses asplain—65538 for example—as the default regular expression match and output display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain format and the asdot format as described in RFC 5396. To change the default regular expression match and output display of 4-byte autonomous system numbers to asdot format, use the bgp asnotation dot command followed by the clear bgp * command to perform a hard reset of all current BGP sessions.

To ensure a smooth transition we recommend that all BGP speakers within an autonomous system that is identified using a 4-byte autonomous system number, are upgraded to support 4-byte autonomous system numbers.

Examples

The following is sample output from the show bgp regexp command in privileged EXEC mode:

```
Router# show bgp regexp 108$
BGP table version is 1738, local router ID is 172.16.72.24
Status codes: s suppressed, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Metric LocPrf Weight Path
*  172.16.0.0       172.16.72.30                           0 109 108 ?
*  172.16.1.0       172.16.72.30                           0 109 108 ?
*  172.16.11.0      172.16.72.30                           0 109 108 ?
*  172.16.14.0      172.16.72.30                           0 109 108 ?
*  172.16.15.0      172.16.72.30                           0 109 108 ?
```

```
*  172.16.16.0      172.16.72.30                           0 109 108 ?
*  172.16.17.0      172.16.72.30                           0 109 108 ?
*  172.16.18.0      172.16.72.30                           0 109 108 ?
*  172.16.19.0      172.16.72.30                           0 109 108 ?
*  172.16.24.0      172.16.72.30                           0 109 108 ?
*  172.16.29.0      172.16.72.30                           0 109 108 ?
*  172.16.30.0      172.16.72.30                           0 109 108 ?
*  172.16.33.0      172.16.72.30                           0 109 108 ?
*  172.16.35.0      172.16.72.30                           0 109 108 ?
*  172.16.36.0      172.16.72.30                           0 109 108 ?
*  172.16.37.0      172.16.72.30                           0 109 108 ?
*  172.16.38.0      172.16.72.30                           0 109 108 ?
*  172.16.39.0      172.16.72.30                           0 109 108 ?
```

After the bgp asnotation dot command is configured, the regular expression match format for 4-byte autonomous system paths is changed to asdot notation format. Although a 4-byte autonomous system number can be configured in a regular expression using either asplain or asdot format, only 4-byte autonomous system numbers configured using the current default format are matched. In the first example, the show bgp regexp command is configured with a 4-byte autonomous system number in asplain format. The match fails because the default format is currently asdot format and there is no output. In the second example using asdot format, the match passes and the information about the 4-byte autonomous system path is shown using the asdot notation.

Note The asdot notation uses a period which is a special character in Cisco regular expressions. to remove the special meaning, use a backslash before the period.

```
Router# show bgp regexp ^65536$
Router# show bgp regexp ^1\.0$
BGP table version is 2, local router ID is 172.17.1.99
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.1.0/24      192.168.1.2              0             0 1.0 i
```

The following is sample output from the show bgp regexp command after the bgp asnotation dot command has been entered to display 4-byte autonomous system numbers

Note The asdot notation uses a period which is a special character in Cisco regular expressions. to remove the special meaning, use a backslash before the period.

```
Router# show bgp regexp ^1\.14$
BGP table version is 4, local router ID is 172.17.1.99
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.1.1.0/24      192.168.1.2              0             0 1.14   i
```

show bgp replication

To display update replication statistics for Border Gateway Protocol (BGP) update groups, use the show bgp replication command in EXEC mode.

show bgp replication [*index-group* | *ip-address*]

Syntax Description

index-group	(Optional) Displays update replication statistics for the update group with the corresponding index number. The range of update-group index numbers is from 1 to 4294967295.
ip-address	(Optional) Displays update replication statistics for this neighbor.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Usage Guidelines

The output of this command displays BGP update-group replication statistics.

When a change to outbound policy occurs, the router automatically recalculates update-group memberships and applies the changes by triggering an outbound soft reset after a 3-minute timer expires. This behavior is designed to provide the network operator with time to change the configuration if a mistake is made. You can manually enable an outbound soft reset before the timer expires by entering the clearbgp ip-address soft out command.

Examples

The following sample output from the show bgp replication command shows update-group replication information for all neighbors:

```
ciscoasa# show bgp replication
BGP Total Messages Formatted/Enqueued : 0/0
     Index     Type  Members          Leader   MsgFmt  MsgRepl  Csize  Qsize
         1 internal        1       10.4.9.21        0        0      0      0
         2 internal        2        10.4.9.5        0        0      0      0
The following sample output from the show bgp replication command shows update-group
statistics for the 10.4.9.5 neighbor:
Router# show bgp replication 10.4.9.5
     Index     Type  Members          Leader   MsgFmt  MsgRepl  Csize  Qsize
         2 internal        2        10.4.9.5        0        0      0      0
```

Table 4-23 shows each field description.

Table 23: show bgp replication Fields

Field	Description
Index	Index number of the update group.
Type	Type of peer (internal or external).
Members	Number of members in the dynamic update peer group.
Leader	First member of the dynamic update peer group.

show bgp rib-failure

To display Border Gateway Protocol (BGP) routes that failed to install in the Routing Information Base (RIB) table, use the show bgp rib-failure command in privileged EXEC mode.

show bgp rib-failure

Syntax Description

This command has no keywords or arguments.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Examples

The following is a sample output from the show bgp rib-failure command:

```
ciscoasa# show bgp rib-failure
Network          Next Hop                    RIB-failure   RIB-NH Matches
10.1.15.0/24     10.1.35.5          Higher admin distance             n/a
10.1.16.0/24     10.1.15.1          Higher admin distance             n/a
```

Table 4-24 shows each field description.

Table 24: show bgp rib-failure Fields

Field	Description
Network	IP address of a network entity
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.
RIB-failure	Cause of RIB failure. Higher admin distance means that a route with a better (lower) administrative distance such as a static route already exists in the IP routing table.

Field	Description
RIB-NH Matches	Route status that applies only when Higher admin distance appears in the RIB-failure column and bgp suppress-inactive is configured for the address family being used. There are three choices: • Yes—Means that the route in the RIB has the same next hop as the BGP route or next hop recurses down to the same adjacency as the BGP nexthop. • No—Means that the next hop in the RIB recurses down differently from the next hop of the BGP route. • n/a—Means that bgp suppress-inactive is not configured for the address family being used.

show bgp summary

To display the status of all Border Gateway Protocol (BGP) connections, use the show bgp summary command in user EXEC or privileged EXEC mode.

show bgp summary

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
9.2(1)	This command was added

Usage Guidelines

The show bgp summary command is used to display BGP path, prefix, and attribute information for all connections to BGP neighbors.

A prefix is an IP address and network mask. It can represent an entire network, a subset of a network, or a single host route. A path is a route to a given destination. By default, BGP will install only a single path for each destination. If multipath routes are configured, BGP will install a path entry for each multipath route, and only one multipath route will be marked as the bestpath.

BGP attribute and cache entries are displayed individually and in combinations that affect the bestpath selection process. The fields for this output are displayed when the related BGP feature is configured or attribute is received. Memory usage is displayed in bytes.

The Cisco implementation of 4-byte autonomous system numbers uses asplain—65538 for example—as the default regular expression match and output display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain format and the asdot format as described in RFC 5396. To change the default regular expression match and output display of 4-byte autonomous system numbers to asdot format, use the bgp asnotation dot command followed by the clear bgp * command to perform a hard reset of all current BGP sessions.

Examples

The following is sample output from the show bgp summary command in privileged EXEC mode:

```
Router# show bgp summary
BGP router identifier 172.16.1.1, local AS number 100
BGP table version is 199, main routing table version 199
37 network entries using 2850 bytes of memory
59 path entries using 5713 bytes of memory
18 BGP path attribute entries using 936 bytes of memory
```

```
2 multipath network entries and 4 multipath paths
10 BGP AS-PATH entries using 240 bytes of memory
7 BGP community entries using 168 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
90 BGP advertise-bit cache entries using 1784 bytes of memory
36 received paths for inbound soft reconfiguration
BGP using 34249 total bytes of memory
Dampening enabled. 4 history paths, 0 dampened paths
BGP activity 37/2849 prefixes, 60/1 paths, scan interval 15 secs
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down State/PfxRcd
10.100.1.1      4   200      26      22      199    0    0 00:14:23 23
10.200.1.1      4   300      21      51      199    0    0 00:13:40 0
```

Table 4-25 shows each field description.

Table 25: show bgp summary Fields

Field	Description
BGP router identifier	In order of precedence and availability, the router identifier specified by the bgp router-id command, a loopback address, or the highest IP address.
BGP table version	Internal version number of BGP database.
main routing table version	Last version of BGP database that was injected into the main routing table.
...network entries	Number of unique prefix entries in the BGP database.
...using ... bytes of memory	Amount of memory, in bytes, that is consumed for the path, prefix, or attribute entry displayed on the same line.
...path entries using	Number of path entries in the BGP database. Only a single path entry will be installed for a given destination. If multipath routes are configured, a path entry will be installed for each multipath route.
...multipath network entries using	Number of multipath entries installed for a given destination.
* ...BGP path/bestpath attribute entries using	Number of unique BGP attribute combinations for which a path is selected as the bestpath.
* ...BGP rrinfo entries using	Number of unique ORIGINATOR and CLUSTER_LIST attribute combinations.
...BGP AS-PATH entries using	Number of unique AS_PATH entries.
...BGP community entries using	Number of unique BGP community attribute combinations.
*...BGP extended community entries using	Number of unique extended community attribute combinations.
BGP route-map cache entries using	Number of BGP route-map match and set clause combinations. A value of 0 indicates that the route cache is empty.
...BGP filter-list cache entries using	Number of filter-list entries that match an AS-path access list permit or deny statements. A value of 0 indicates that the filter-list cache is empty.

Field	Description
BGP advertise-bit cache entries using	(Cisco IOS Release 12.4(11)T and later releases only) Number of advertised bitfield entries and the associated memory usage. A bitfield entry represents a piece of information (one bit) that is generated when a prefix is advertised to a peer. The advertised bit cache is built dynamically when required
...received paths for inbound soft reconfiguration	Number paths received and stored for inbound soft reconfiguration.
BGP using...	Total amount of memory, in bytes, used by the BGP process.
Dampening enabled...	Indicates that BGP dampening is enabled. The number of paths that carry an accumulated penalty and the number of dampened paths are displayed on this line.
BGP activity...	Displays the number of times that memory has been allocated or released for a path or prefix.
Neighbor	IP address of the neighbor.
V	BGP version number spoken to the neighbor.
AS	Autonomous system number.
MsgRcvd	Number of messages received from the neighbor.
MsgSent	Number of messages sent to the neighbor.
TblVer	Last version of the BGP database that was sent to the neighbor.
InQ	Number of messages queued to be processed from the neighbor.
OutQ	Number of messages queued to be sent to the neighbor.
Up/Down	The length of time that the BGP session has been in the Established state, or the current status if not in the Established state.
State/PfxRcd	Current state of the BGP session, and the number of prefixes that have been received from a neighbor or peer group. When the maximum number (as set by the neighbor maximum-prefix command) is reached, the string "PfxRcd" appears in the entry, the neighbor is shut down, and the connection is set to Idle. An (Admin) entry with Idle status indicates that the connection has been shut down using the neighbor shutdown command.

Examples

The following output from the show bgp summary command shows that the BGP neighbor 192.168.3.2 was dynamically created and is a member of the listen range group, group192. The output also shows that the IP prefix range of 192.168.0.0/16 is defined for the listen range group named group192. In Cisco IOS Release 12.2(33)SXH and later releases, the BGP dynamic neighbor feature added the ability to support the dynamic creation of BGP neighbor peers using a subnet range associated with a peer group (listen range group).

```
ciscoasa# show bgp summary
BGP router identifier 192.168.3.1, local AS number 45000
BGP table version is 1, main routing table version 1
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
*192.168.3.2    4 50000       2       2        0    0    0 00:00:37        0
* Dynamically created based on a listen range command
Dynamically created neighbors: 1/(200 max), Subnet ranges: 1
BGP peergroup group192 listen range group members:
  192.168.0.0/16
```

The following output from the show bgp summary command shows two BGP neighbors, 192.168.1.2 and 192.168.3.2, in different 4-byte autonomous system numbers, 65536 and 65550. The local autonomous system 65538 is also a 4-byte autonomous system number and the numbers are displayed in the default asplain format.

```
Router# show bgp summary
BGP router identifier 172.17.1.99, local AS number 65538
BGP table version is 1, main routing table version 1
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  Statd
192.168.1.2     4       65536       7       7        1    0    0 00:03:04      0
192.168.3.2     4       65550       4       4        1    0    0 00:00:15      0
```

The following output from the show bgp summary command shows the same two BGP neighbors, but the 4-byte autonomous system numbers are displayed in asdot notation format. To change the display format the bgp asnotation dot command must be configured in router configuration mode.

```
Router# show bgp summary
BGP router identifier 172.17.1.99, local AS number 1.2
BGP table version is 1, main routing table version 1
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  Statd
192.168.1.2     4         1.0       9       9        1    0    0 00:04:13      0
192.168.3.2     4        1.14       6       6        1    0    0 00:01:24      0
```

The following example displays sample output of the show bgp summary slow command:

```
ciscoasa> show bgp summary slow
BGP router identifier 2.2.2.2, local AS number 100
BGP table version is 37, main routing table version 37
36 network entries using 4608 bytes of memory
36 path entries using 1872 bytes of memory
1/1 BGP path/bestpath attribute entries using 124 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
1 BGP extended community entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 6700 total bytes of memory
BGP activity 46/0 prefixes, 48/0 paths, scan interval 60 secs

Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down State/PfxRcd
6.6.6.6 4 100 11 10 1 0 0 00:44:20 0
```

show bgp system-config

To display running configuration for bgp of system context in user context, use the show bgp system-config command in user or privileged EXEC mode.

show bgp system-config

Syntax Description

This command has no arguments or keywords.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC, User EXEC	• Yes	• Yes	• Yes	—	—

Command History

Release	Modification
9.2(1)	This command was added

Usage Guidelines

This command can be used only in user context without any arguments or keywords. This command can be useful for checking the running configuration enforced on user context by system context.

Examples

The following sample output is similar to the output that will be displayed when the show bgp system-config command is entered in user EXEC mode:

```
ciscoasa/c1(config)# show bgp system-config
router bgp 1
 bgp log-neighbor-changes
 no bgp always-compare-med
 no bgp asnotation dot
 no bgp bestpath med
 no bgp bestpath compare-routerid
 bgp default local-preference 100
 no bgp deterministic-med
 bgp enforce-first-as
 bgp maxas-limit 0
 bgp transport path-mtu-discovery
 timers bgp 60 180 0
 address-family ipv4 unicast
  bgp scan-time 0
  bgp nexthop trigger enable
  bgp nexthop trigger delay 5
 exit-address-family
```

show blocks

To show the packet buffer utilization, use the **show blocks** command in privileged EXEC mode.

show blocks [**core | export-failed | interface**] [{ **address** *hex* **| all | assigned | free | old | pool** *size* [**summary**] } [**diagnostics | dump | header | packet] | queue history |** [**exhaustion snapshot | history** [**list**] [*1-MAX_NUM_SNAPSHOT* | *index*] [**detail**]] [**depleted** *size*]

Syntax Description

address *hex*	(Optional) Shows a block corresponding to this address, in hexadecimal.
all	(Optional) Shows all blocks.
assigned	(Optional) Shows blocks that are assigned and in use by an application.
core	(Optional) Shows core-specific buffers.
depleted	(Optional) Shows the depleted block details for the specified block size. Valid sizes are 0, 4, 80, 256, 1550, 2560, 2048, 4096, 8192, 9344, 16384 and 65536/65664.
detail	(Optional) Shows a portion (128 bytes) of the first block for each unique queue type.
dump *	(Optional) Shows the entire block contents, including the header and packet information. The difference between dump and packet is that dump includes additional information between the header and the packet.
diagnostics	(Optional) Shows block diagnostics.
exhaustion snapshot	(Optional) Prints the last x number (x is currently 10) of snapshots that were taken and the time stamp of the last snapshot). After a snapshot is taken, another snapshot is not taken if less than 5 minutes has passed.
export-failed	(Optional) Show system buffer export failure counters.
free	(Optional) Shows blocks that are available for use.
header	(Optional) Shows the header of the block.
history *1-MAX_NUM_SNAPSHOT* **history** *index* **history list**	The **history** option displays recent and all snapshots in the history. The **history list** option displays a summary of snapshots in the history. The **history** *index* option displays the index of snapshots in the history. The **history** *1-MAX_NUM_SNAPSHOT* option displays only one snapshot in the history.
interface	(Optional) Show buffers attached to interfaces.
old *	(Optional) Shows blocks that were assigned more than a minute ago.
packet	(Optional) Shows the header of the block as well as the packet contents.

pool *size* *	(Optional) Shows blocks of a specific size.
queue history	(Optional) Shows where blocks are assigned when the ASA runs out of blocks. Sometimes, a block is allocated from the pool but never assigned to a queue. In that case, the location is the code address that allocated the block.
summary	(Optional) Shows detailed information about block usage sorted by the program addresses of applications that allocated blocks in this class, program addresses of applications that released blocks in this class, and the queues to which valid blocks in this class belong.

*When these commands are integrated in scripts and run within aggresive intervals, it might overload the system. Therefore, use these commands after verifying the system capacity to withstand the load.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
7.0(1)	The **pool summary** option was added.
8.0(2)	The dupb block uses 0 length blocks now instead of 4 byte blocks. An additional line was added for 0 byte blocks.
9.1(5)	The **exhaustion snapshot** , **history list** , **history** *index*, and **history** *1-MAX_NUM_SNAPSHOT* options were added.
9.14(1)	The **depleted** keyword was added to the command to display the depleted block details.
9.16(2)	The output of this command was enhanced to include the failed count.

Usage Guidelines

The **show blocks** command helps you determine if the ASA is overloaded. This command lists preallocated system buffer utilization. A full memory condition is not a problem as long as traffic is moving through the ASA. You can use the **show conn** command to see if traffic is moving. If traffic is not moving and the memory is full, there may be a problem.

You can also view this information using SNMP.

The information shown in a security context includes the system-wide information as well as context-specific information about the blocks in use and the high water mark for block usage.

See the “Examples” section for a description of the display output.

Examples

The following is sample output from the **show blocks** command in single mode:

```
ciscoasa# show blocks
  SIZE    MAX    LOW    CNT    FAILED
     0    100     99    100       0
     4   1600   1598   1599       0
    80    400    398    399       0
   256   3600   3540   3542       0
  1550   4716   3177   3184       0
 16384     10     10     10       0
  2048   1000   1000   1000       0
```

The following table shows each field description.

Table 26: show blocks Fields

Field	Description
SIZE	Size, in bytes, of the block pool. Each size represents a particular type.
0	Used by dupb blocks. Duplicating existing blocks - This includes storing the original (for re-transmission) and sending a copy, or for passing copies of blocks between subsystems. Block zeros allocated by function [1] deals with both encrypted and normal traffic. These blocks are recovered on acknowledgment or reply from peer. Delay of acknowledgment or reply from peer leads to blocks depletion.
4	Duplicates existing blocks in applications such as DNS, ISAKMP, URL filtering, uauth, TFTP, and TCP modules. Also, this sized block can be used normally by code to send packets to drivers, etc.
80	Used in TCP intercept to generate acknowledgment packets and for failover hello messages.
256	Used for Stateful Failover updates, syslogging, and other TCP functions. These blocks are mainly used for Stateful Failover messages. The active ASA generates and sends packets to the standby ASA to update the translation and connection table. In bursty traffic, where high rates of connections are created or torn down, the number of available blocks might drop to 0. This situation indicates that one or more connections were not updated to the standby ASA. The Stateful Failover protocol catches the missing translation or connection the next time. If the CNT column for 256-byte blocks stays at or near 0 for extended periods of time, then the ASA is having trouble keeping the translation and connection tables synchronized because of the number of connections per second that the ASA is processing. Syslog messages sent out from the ASA also use the 256-byte blocks, but they are generally not released in such quantity to cause a depletion of the 256-byte block pool. If the CNT column shows that the number of 256-byte blocks is near 0, ensure that you are not logging at Debugging (level 7) to the syslog server. This is indicated by the logging trap line in the ASA configuration. We recommend that you set logging at Notification (level 5) or lower, unless you require additional information for debugging purposes. There are 100 blocks of 256 byte blocks have been reserved for HA FHELLO/FHELLO_TEST communication.

Field	Description
1550	Used to store Ethernet packets for processing through the ASA. General packet processing, buffering URL filter packets, HA Failover. When a packet enters an ASA interface, it is placed on the input interface queue, passed up to the operating system, and placed in a block. The ASA determines whether the packet should be permitted or denied based on the security policy and processes the packet through to the output queue on the outbound interface. If the ASA is having trouble keeping up with the traffic load, the number of available blocks will hover close to 0 (as shown in the CNT column of the command output). When the CNT column is zero, the ASA attempts to allocate more blocks. The maximum can be greater than 8192 for 1550-byte blocks if you issue this command. If no more blocks are available, the ASA drops the packet.
2048	Control or guided frames used for control updates. See the description for 1550 for more information about Ethernet packets.
2560	Used for IKE Messages.
4096	Used for QoS Metrics.
8192	Used for QoS Metrics.
9344	Interface blocks for jumbo frames when "jumbo-frame reservation" is enabled.
16384	Only used for the 64-bit, 66-MHz Gigabit Ethernet cards (i82543). Only used for jumbo frames in 5580.
65536	Used for QoS Metrics.
MAX	Maximum number of blocks available for the specified byte block pool. The maximum number of blocks are carved out of memory at bootup. Typically, the maximum number of blocks does not change. The exception is for the 256- and 1550-byte blocks, where the ASA can dynamically create more when needed. The maximum can be greater than 8192 for 1550-byte blocks if you issue this command.
LOW	Low-water mark. This number indicates the lowest number of this size blocks available since the ASA was powered up, or since the last clearing of the blocks (with the **clear blocks** command). A zero in the LOW column indicates a previous event where memory was full. **Note** To reset this value to MAX, you must reboot ASA.
CNT	Current number of blocks available for that specific size block pool. A zero in the CNT column means memory is full now.
FAILED	When the memory count for a block size is completely exhausted (LOW and CNT value is zero), the corresponding FAILED column is incremented with the number of allocation request for the same block size received thereafter. Eventually, when memory space is feed, the current available blocks for allocation increments. However, the FAILED value does not decrease, as it is a record of the number of failures that have occurred. If CNT and FAILED values increase, it indicates an issue and must be resolved.

[1] proxyi_fill_tx_pkt()

Examples

The following is sample output from the **show blocks all** command:

```
ciscoasa# show blocks all
Class 0, size 4
     Block   allocd_by    freed_by  data size    alloccnt     dup_cnt  oper location
0x01799940  0x00000000  0x00101603          0           0           0 alloc not_specified
0x01798e80  0x00000000  0x00101603          0           0           0 alloc not_specified
0x017983c0  0x00000000  0x00101603          0           0           0 alloc not_specified
...
    Found 1000 of 1000 blocks
    Displaying 1000 of 1000 blocks
```

Table 4-27 shows each field description.

Table 27: show blocks all Fields

Field	Description
Block	The block address.
allocd_by	The program address of the application that last used the block (0 if not used).
freed_by	The program address of the application that last released the block.
data size	The size of the application buffer/packet data that is inside the block.
alloccnt	The number of times this block has been used since the block came into existence.
dup_cnt	The current number of references to this block if used: 0 means 1 reference, 1 means 2 references.
oper	One of the four operations that was last performed on the block: alloc, get, put, or free.
location	The application that uses the block, or the program address of the application that last allocated the block (same as the allocd_by field).

Examples

The following is sample output from the **show blocks** command in a context. In multiple context mode, the output includes information on the number of blocks currently in use by the context (INUSE), and the high-water mark for blocks used by the context (HIGH).

```
ciscoasa/contexta# show blocks
  SIZE    MAX    LOW    CNT  INUSE   HIGH
     4   1600   1599   1599      0      0
    80    400    400    400      0      0
   256   3600   3538   3540      0      1
  1550   4616   3077   3085      0      0
```

The following is sample output from the **show blocks queue history** command:

```
ciscoasa# show blocks queue history
Each Summary for User and Queue_type is followed its top 5 individual queues
Block Size: 4
Summary for User "http", Queue "tcp_unp_c_in", Blocks 1595, Queues 1396
Blk_cnt Q_cnt Last_Op Queue_Type        User      Context
    186     1 put                               contexta
     15     1 put                               contexta
      1     1 put                               contexta
```

```
      1     1 put                                 contextb
      1     1 put                                 contextc
Summary for User "aaa", Queue "tcp_unp_c_in", Blocks 220, Queues 200
Blk_cnt Q_cnt Last_Op Queue_Type        User      Context
     21     1 put                                 contexta
      1     1 put                                 contexta
      1     1 put                                 contexta
      1     1 put                                 contextb
      1     1 put                                 contextc
Blk_cnt Q_cnt Last_Op Queue_Type        User      Context
    200     1 alloc   ip_rx             tcp       contexta
    108     1 get     ip_rx             udp       contexta
     85     1 free    fixup             h323_ras  contextb
     42     1 put     fixup             skinny    contextb
Block Size: 1550
Summary for User "http", Queue "tcp_unp_c_in", Blocks 1595, Queues 1000
Blk_cnt Q_cnt Last_Op Queue_Type        User      Context
    186     1 put                                 contexta
     15     1 put                                 contexta
      1     1 put                                 contexta
      1     1 put                                 contextb
      1     1 put                                 contextc
...
```

The following is sample output from the **show blocks queue history detail** command:

```
ciscoasa# show blocks queue history detail
History buffer memory usage: 2136 bytes (default)
Each Summary for User and Queue type is followed its top 5 individual queues
Block Size: 4
Summary for User "http", Queue_Type "tcp_unp_c_in", Blocks 1595, Queues 1396
Blk_cnt Q_cnt Last_Op Queue_Type        User      Context
    186     1 put                                 contexta
     15     1 put                                 contexta
      1     1 put                                 contexta
      1     1 put                                 contextb
      1     1 put                                 contextc
 First Block information for Block at 0x.....
  dup_count 0, flags 0x8000000, alloc_pc 0x43ea2a,
  start_addr 0xefb1074, read_addr 0xefb118c, write_addr 0xefb1193
  urgent_addr 0xefb118c, end_addr 0xefb17b2
  0efb1150: 00 00 00 03 47 c5 61 c5 00 05 9a 38 76 80 a3 00  |  ....G.a....8v...
  0efb1160: 00 0a 08 00 45 00 05 dc 9b c9 00 00 ff 06 f8 f3  |  ....E...........
  0efb1170: 0a 07 0d 01 0a 07 00 50 00 17 cb 3d c7 e5 60 62  |  .......P...=..`b
  0efb1180: 7e 73 55 82 50 18 10 00 45 ca 00 00 2d 2d 20 49  |  ~sU.P...E...-- I
  0efb1190: 50 20 2d 2d 0d 0a 31 30 2e 37 2e 31 33 2e 31 09  |  P --..10.7.13.1.
  0efb11a0: 3d 3d 3e 09 31 30 2e 37 2e 30 2e 38 30 0d 0a 0d  |  ==>.10.7.0.80...
Summary for User "aaa", Queue "tcp_unp_c_in", Blocks 220, Queues 200
Blk_cnt Q_cnt Last_Op Queue_Type        User      Context
     21     1 put                                 contexta
      1     1 put                                 contexta
      1     1 put                                 contexta
      1     1 put                                 contextb
      1     1 put                                 contextc
 First Block information for Block at 0x.....
  dup_count 0, flags 0x8000000, alloc_pc 0x43ea2a,
  start_addr 0xefb1074, read_addr 0xefb118c, write_addr 0xefb1193
  urgent_addr 0xefb118c, end_addr 0xefb17b2
  0efb1150: 00 00 00 03 47 c5 61 c5 00 05 9a 38 76 80 a3 00  |  ....G.a....8v...
  0efb1160: 00 0a 08 00 45 00 05 dc 9b c9 00 00 ff 06 f8 f3  |  ....E...........
  0efb1170: 0a 07 0d 01 0a 07 00 50 00 17 cb 3d c7 e5 60 62  |  .......P...=..`b
  0efb1180: 7e 73 55 82 50 18 10 00 45 ca 00 00 2d 2d 20 49  |  ~sU.P...E...-- I
  0efb1190: 50 20 2d 2d 0d 0a 31 30 2e 37 2e 31 33 2e 31 09  |  P --..10.7.13.1.
  0efb11a0: 3d 3d 3e 09 31 30 2e 37 2e 30 2e 38 30 0d 0a 0d  |  ==>.10.7.0.80...
```

```
...
total_count: total buffers in this class
```

The following is sample output from the **show blocks pool summary** command:

```
ciscoasa# show blocks pool 1550 summary
Class 3, size 1550
================================================
         total_count=1531    miss_count=0
Alloc_pc        valid_cnt       invalid_cnt
0x3b0a18        00000256        00000000
         0x01ad0760 0x01acfe00 0x01acf4a0 0x01aceb40 00000000 0x00000000
0x3a8f6b        00001275        00000012
         0x05006aa0 0x05006140 0x050057e0 0x05004520 00000000
0x00000000
================================================
         total_count=9716    miss_count=0
Freed_pc        valid_cnt       invalid_cnt
0x9a81f3        00000104        00000007
         0x05006140 0x05000380 0x04fffa20 0x04ffde00 00000000 0x00000000
0x9a0326        00000053        00000033
         0x05006aa0 0x050057e0 0x05004e80 0x05003260 00000000 0x00000000
0x4605a2        00000005        00000000
         0x04ff5ac0 0x01e8e2e0 0x01e2eac0 0x01e17d20 00000000 0x00000000
...
================================================
         total_count=1531    miss_count=0
Queue   valid_cnt       invalid_cnt
0x3b0a18        00000256        00000000  Invalid Bad qtype
         0x01ad0760 0x01acfe00 0x01acf4a0 0x01aceb40 00000000 0x00000000
0x3a8f6b        00001275        00000000  Invalid Bad qtype
         0x05006aa0 0x05006140 0x050057e0 0x05004520 00000000
0x00000000
================================================
free_cnt=8185  fails=0  actual_free=8185  hash_miss=0
   03a8d3e0  03a8b7c0  03a7fc40  03a6ff20  03a6f5c0  03a6ec60 kao-f1#
```

The following is sample output from the **show blocks exhaustion history list** command:

```
ciscoasa# show blocks exhaustion history list
1 Snapshot created at 18:01:03 UTC Feb 19 2014:
   Snapshot created due to 16384 blocks running out
2 Snapshot created at 18:02:03 UTC Feb 19 2014:
   Snapshot created due to 16384 blocks running out
3 Snapshot created at 18:03:03 UTC Feb 19 2014:
   Snapshot created due to 16384 blocks running out
4 Snapshot created at 18:04:03 UTC Feb 19 2014:
   Snapshot created due to 16384 blocks running out
```

Table 4-28 shows each field description.

Table 28: show blocks pool summary Fields

Field	Description
total_count	The number of blocks for a given class.
miss_count	The number of blocks not reported in the specified category due to technical reasons.
Freed_pc	The program addresses of applications that released blocks in this class.

Field	Description
Alloc_pc	The program addresses of applications that allocated blocks in this class.
Queue	The queues to which valid blocks in this class belong.
valid_cnt	The number of blocks that are currently allocated.
invalid_cnt	The number of blocks that are not currently allocated.
Invalid Bad qtype	Either this queue has been freed and the contents are invalid or this queue was never initialized.
Valid tcp_usr_conn_inp	The queue is valid.

The following is sample output from the **show blocks depleted** command:

```
ciscoasa# show blocks depleted 8192

Block Class: 8,   Class Size: 8192,   Count: 100

1 Depletion created at 11:47:48 UTC Feb 17 2022

First Snapshot Details:

 Block             allocd_by          freed_by            alloccnt    dup_cnt  timestamp
  oper location

0x00007f117bd5f9c0 0x0000560e84d1236b 0x0000560e822144e4     1           0     246610
     alloc 0x0000560e84d1236b

0x00007f117bd5d300 0x0000560e84d1236b 0x0000560e822144e4     1           0     246610
     alloc 0x0000560e84d1236b

0x00007f117bd5ac40 0x0000560e84d1236b 0x0000560e822144e4     1           0     246610
     alloc 0x0000560e84d1236b

0x00007f117bd58580 0x0000560e84d1236b 0x0000560e822144e4     1           0     246610
     alloc 0x0000560e84d1236b

0x00007f117bd55ec0 0x0000560e84d1236b 0x0000560e822144e4     1           0     246610
     alloc 0x0000560e84d1236b

0x00007f117bd53800 0x0000560e84d1236b 0x0000560e822144e4     1           0     246610
     alloc 0x0000560e84d1236b

0x00007f117bd51140 0x0000560e84d1236b 0x0000560e822144e4     1           0     246610
     alloc 0x0000560e84d1236b

0x00007f117bd4ea80 0x0000560e84d1236b 0x0000560e822144e4     1           0     246610
     alloc 0x0000560e84d1236b

0x00007f117bd4c3c0 0x0000560e84d1236b 0x0000560e822144e4     1           0     246610
     alloc 0x0000560e84d1236b

0x00007f117bd49d00 0x0000560e84d1236b 0x0000560e822144e4     1           0     246610
     alloc 0x0000560e84d1236b

0x00007f117bd47640 0x0000560e84d1236b 0x0000560e822144e4     1           0     246610
     alloc 0x0000560e84d1236b
```

```
0x00007f117bd44f80 0x0000560e84d1236b 0x0000560e822144e4     1          0     246610
     alloc 0x0000560e84d1236b

0x00007f117bd428c0 0x0000560e84d1236b 0x0000560e822144e4     1          0     246610
     alloc 0x0000560e84d1236b

0x00007f117bd40200 0x0000560e84d1236b 0x0000560e822144e4     1          0     246610
     alloc 0x0000560e84d1236b

0x00007f117bd3db40 0x0000560e84d1236b 0x0000560e822144e4     1          0     246620
     alloc 0x0000560e84d1236b

0x00007f117bd3b480 0x0000560e84d1236b 0x0000560e822144e4     1          0     246620
     alloc 0x0000560e84d1236b

<--- More --->

0x00007f117bc85a40 0x0000560e84d1236b 0x0000560e822144e4     1          0     263390
     alloc 0x0000560e84d1236b

0x00007f117bc83380 0x0000560e84d1236b 0x0000560e822144e4     1          0     263390
     alloc 0x0000560e84d1236b

0x00007f117bc80cc0 0x0000560e84d1236b 0x0000560e822144e4     1          0     263390
     alloc 0x0000560e84d1236b

0x00007f117bc7e600 0x0000560e84d1236b 0x0000560e822144e4     1          0     263390
     alloc 0x0000560e84d1236b

0x00007f117bc7bf40 0x0000560e84d1236b 0x0000560e822144e4     1          0     263390
     alloc 0x0000560e84d1236b

0x00007f117bc79880 0x0000560e84d1236b 0x0000560e822144e4     1          0     263390
     alloc 0x0000560e84d1236b

<--- More --->

. . . . .
. . . . .

0x00007f117bc771c0 0x0000560e84d1236b 0x0000560e822144e4     1          0     263390
     alloc 0x0000560e84d1236b

0x00007f117bc74b00 0x0000560e84d1236b 0x0000560e822144e4     1          0     263390
     alloc 0x0000560e84d1236b

0x00007f117bc72440 0x0000560e84d1236b 0x0000560e822144e4     1          0     263390
     alloc 0x0000560e84d1236b

0x00007f117bc6fd80 0x0000560e84d1236b 0x0000560e822144e4     1          0     263390
     alloc 0x0000560e84d1236b
```

Related Commands

Command	Description
blocks	Increases the memory assigned to block diagnostics
clear blocks	Clears the system buffer statistics.
show conn	Shows active connections.

show bootvar

To show the boot file and configuration properties, use the **show bootvar** command in privileged EXEC mode.

show bootvar

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
7.2(1)	This command was added.

Usage Guidelines

The BOOT variable specifies a list of bootable images on various devices. The CONFIG_FILE variable specifies the configuration file used during system initialization. Set these variables with the **boot system** command and **boot config** command, respectively.

Examples

The BOOT variable contains disk0:/f1_image, which is the image booted when the system reloads. The current value of BOOT is disk0:/f1_image; disk0:/f1_backupimage. This value means that the BOOT variable has been modified with the **boot system** command, but the running configuration has not been saved with the **write memory** command. When the running configuration is saved, the BOOT variable and current BOOT variable will both be disk0:/f1_image; disk0:/f1_backupimage. Assuming that the running configuration is saved, the boot loader will try to load the contents of the BOOT variable, starting with disk0:/f1image, but if that is not present or invalid, the boot loader will try to boot disk0:1/f1_backupimage.

The CONFIG_FILE variable points to the system startup configuration. In this example it is not set, so the startup configuration file is the default specified with the **boot config** command. The current CONFIG_FILE variable may be modified with the **boot config** command and saved with the **write memory** command.

The following is sample output from the **show bootvar** command:

```
ciscoasa# show bootvar
BOOT variable = disk0:/f1_image
Current BOOT variable = disk0:/f1_image; disk0:/f1_backupimage
```

```
CONFIG_FILE variable =
Current CONFIG_FILE variable =
ciscoasa#
```

Related Commands

Command	Description
boot	Specifies the configuration file or image file used at startup.

show bridge-group

To show bridge group information such as interfaces assigned, MAC addresses, and IP addresses, use the **show bridge-group** command in privileged EXEC mode.

show bridge-group *bridge_group_number*

Syntax Description

bridge_group_number	Specifies the bridge group number as an integer between 1 and 100.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	• Yes	—

Command History

Release	Modification
8.4(1)	This command was added.
9.7(1)	We added support in routed mode for Integrated Routing and Bridging.

Examples

The following is sample output from the **show bridge-group** command with IPv4 addresses:

```
ciscoasa# show bridge-group 1
 Interfaces: GigabitEthernet0/0.101, GigabitEthernet0/0.201
 Management System IP Address: 10.0.1.1 255.255.255.0
 Management Current IP Address: 10.0.1.1 255.255.255.0
 Management IPv6 Global Unicast Address(es):
    N/A
 Static mac-address entries: 0
 Dynamic mac-address entries: 2
```

The following is sample output from the **show bridge-group** command with IPv4 and IPv6 addresses:

```
ciscoasa# show bridge-group 1
 Interfaces: GigabitEthernet0/0.101, GigabitEthernet0/0.201
 Management System IP Address: 10.0.1.1 255.255.255.0
 Management Current IP Address: 10.0.1.1 255.255.255.0
 Management IPv6 Global Unicast Address(es):
    2000:100::1, subnet is 2000:100::/64
    2000:101::1, subnet is 2000:101::/64
    2000:102::1, subnet is 2000:102::/64
```

```
Static mac-address entries: 0
Dynamic mac-address entries: 2
```

Related Commands

Command	Description
bridge-group	Groups transparent firewall interfaces into a bridge group.
clear configure interface bvi	Clears the bridge group interface configuration.
interface	Configures an interface.
interface bvi	Creates a bridge virtual interface.
ip address	Sets the management IP address for a bridge group.
show running-config interface bvi	Shows the bridge group interface configuration.

show call-home

To display the configured Call Home information, use the **show call-home** command in privileged EXEC mode.

[**cluster exec**] **show call-home** [**alert-group | detail | events | mail-server | status | profile** { *profile_name* | **all** } | **statistics**]

Syntax Description

alert-group	(Optional) Displays the available alert group.
cluster exec	(Optional) In a clustering environment, enables you to issue the **show call-home** command in one unit and run the command in all the other units at the same time.
detail	(Optional) Displays the Call Home configuration in detail.
events	(Optional) Displays current detected events.
mail-server status	(Optional) Displays the Call Home mail server status information.
profile *profile _name* **all**	(Optional) Displays configuration information for all existing profiles.
statistics	(Optional) Displays the Call Home statistics.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	—	• Yes

Command History

Release	Modification
8.2(2)	This command was added.
9.1(3)	A new type of Smart Call Home message has been added to include the output of the **show cluster history** command and **show cluster info** command.

Examples

The following is sample output from the show call-home command and displays the configured Call Home settings:

```
ciscoasa# show call-homeCurrent Smart Call-Home settings:Smart Call-Home feature : enableSmart
```

```
 Call-Home message's from address: from@example.comSmart Call-Home message's reply-to
address: reply-to@example.comcontact person's email address: example@example.comcontact
person's phone: 111-222-3333street address: 1234 Any Street, Any city, Any state,
12345customer ID: ExampleCorpcontract ID: X123456789site ID: SantaClaraMail-server[1]:
Address: smtp.example.com Priority: 1Mail-server[2]: Address: 192.168.0.1 Priority:
10Rate-limit: 60 message(s) per minute
Available alert groups:Keyword                  State
----------------------- -------
Syslog Enable
diagnostic Enableenvironmental Enableinventory Enableconfiguration Enablefirewall
Enabletroubleshooting Enablereport Enable
Profiles:Profile Name: CiscoTAC-1Profile Name: prof1Profile Name: prof2
```

The following is sample output from the show call-home detail command and displays detailed Call Home configuration information:

```
ciscoasa# show call-home detailDescription: Show smart call-home configuration in
detail.Supported Modes: single mode and system context in multi mode,
routed/transparent.Output:
Current Smart Call-Home settings:Smart Call-Home feature: enableSmart Call-Home message's
from address: from@example.example.comSmart Call-Home message's reply-to address:
reply-to@example.example.comcontact person's email address: abc@example.comcontact person's
 phone: 111-222-3333street address: 1234 Any Street, Any city, Any state, 12345customer ID:
 111111contract ID: 123123site ID: SantaClaraMail-server[1]: Address: example.example.com
Priority: 1Mail-server[2]: Address: example.example.com Priority: 10Rate-limit: 60 message(s)
 per minute
Available alert groups:Keyword State----------------------- -------syslog Enablediagnostic
 Enableenvironmental Enableinventory Enableconfiguration Enablefirewall Enabletroubleshooting
 Enablereport Enable
Profiles:
Profile Name: CiscoTAC-1Profile status: ACTIVE Preferred Message Format: xmlMessage Size
Limit: 3145728 BytesEmail address(es): anstage@cisco.comHTTP address(es):
https://tools.cisco.com/its/service/oddce/services/DDCEService
Periodic inventory message is scheduled monthly at 01:00
Alert-group Severity----------------------- ------------inventory n/a
Profile Name: prof1Profile status: ACTIVE Preferred Message Format: xmlMessage Size Limit:
 3145728 BytesEmail address(es): example@example.comHTTP address(es):
https://kafan-lnx-01.cisco.com:8443/sch/sch.jsp
Periodic configuration message is scheduled daily at 01:00
Periodic inventory message is scheduled every 60 minutes
Alert-group Severity----------------------- ------------configuration n/ainventory n/a
Profile Name: prof2Profile status: ACTIVE Preferred Message Format: short-textMessage Size
 Limit: 1048576 BytesEmail address(es): example@example.comHTTP address(es):
https://example.example.com:8443/sch/sch.jsp
Periodic configuration message is scheduled every 1 minutes
Periodic inventory message is scheduled every 1 minutes
Alert-group Severity----------------------- ------------configuration n/ainventory n/a
```

The following is sample output from the show call-home events command and displays available Call Home events:

```
ciscoasa# show call-home eventsDescription: Show current detected events.Supported Modes:
single mode and system context in multi mode, routed/transparent.Output:Active event
list:Event client alert-group severity active
(sec)-------------------------------------------------------------------Configuration
Client configuration none 5Inventory inventory none 15
```

The following is sample output from the show call-home mail-server status command and displays available Call Home mail-server status:

```
ciscoasa# show call-home mail-server statusDescription: Show smart call-home configuration,
 status, and statistics.Supported Modes: single mode and system context in multi mode,
```

```
routed/transparent.Output:Mail-server[1]: Address: example.example.com Priority: 1
[Available]Mail-server[2]: Address: example.example.com Priority: 10 [Not Available]
```

The following is sample output from the show call-home alert-group command and displays the available alert groups:

```
ciscoasa# show call-home alert-groupDescription: Show smart call-home alert-group
states.Supported Modes: single mode and system context in multi mode,
routed/transparent.Output:Available alert groups:Keyword State-----------------------
-------syslog Enablediagnostic Enableenvironmental Enableinventory Enableconfiguration
Enablefirewall Enabletroubleshooting Enablereport Enable
```

The following is sample output from the show call-home profile profile-name | all command and displays information for all predefined and user-defined profiles:

```
ciscoasa# show call-home profile {profile-name
 | all}Description: Show smart call-home profile configuration.Supported Modes: single mode
 and system context in multi mode, routed/transparent.Output:Profiles:
Profile Name: CiscoTAC-1Profile status: ACTIVE Preferred Message Format: xmlMessage Size
Limit: 3145728 BytesEmail address(es): anstage@cisco.comHTTP address(es):
https://tools.cisco.com/its/service/oddce/services/DDCEService
Periodic inventory message is scheduled monthly at 01:00
Alert-group Severity----------------------- ------------inventory n/a
Profile Name: prof1Profile status: ACTIVE Preferred Message Format: xmlMessage Size Limit:
 3145728 BytesEmail address(es): example@example.comHTTP address(es):
https://example.example.com:8443/sch/sch.jsp
Periodic configuration message is scheduled daily at 01:00
Periodic inventory message is scheduled every 60 minutes
Alert-group Severity----------------------- ------------configuration n/ainventory n/a
Profile Name: prof2Profile status: ACTIVE Preferred Message Format: short-textMessage Size
 Limit: 1048576 BytesEmail address(es): example@example.comHTTP address(es):
https://example.example.com:8443/sch/sch.jsp
Periodic configuration message is scheduled every 1 minutes
Periodic inventory message is scheduled every 1 minutes
Alert-group Severity----------------------- ------------configuration n/ainventory n/a
```

The following is sample output from the show call-home statistics command and displays the call-home statistics:

```
ciscoasa# show call-home statisticsDescription: Show smart call-home statistics.Supported
Modes: single mode and system context in multi mode, routed/transparent.Output:Message Types
 Total Email HTTP-------------------- ---------------- ---------------- ----------------Total
 Success 0 0 0
Total In-Queue 0 0 0
Total Dropped 5 4 1Tx Failed 5 4 1inventory 3 2 1configuration 2 2 0
Event Types Total-------------------- ----------------Total Detected 2inventory 1configuration
 1
Total In-Queue 0
Total Dropped 0
Last call-home message sent time: 2009-06-17 14:22:09 GMT-07:00
```

The following is sample output from the show call-home status command and displays the call-home status:

```
ciscoasa# show call-home mail-server status
Description: Show smart call-home configuration, status, and statistics.Supported Modes:
single mode and system context in multi mode, routed/transparent.Output:Mail-server[1]:
Address: kafan-lnx-01.cisco.com Priority: 1 [Available]Mail-server[2]: Address:
kafan-lnx-02.cisco.com Priority: 10 [Not Available]37. ciscoasa# show call-home events
Description: Show current detected events.Supported Modes: single mode and system context
in multi mode, routed/transparent.Output:Active event list:Event client alert-group severity
```

```
 active (sec)-------------------------------------------------------------------Configuration
 Client configuration none 5Inventory inventory none 15
```

The following is sample output from the **cluster exec show call-home statistics** command and displays call-home statistics for a cluster:

```
ciscoasa(config)# cluster exec show call-home statistics
A(LOCAL):*******************************************************************
Message Types           Total            Email             HTTP
--------------------   ---------------- ---------------- ----------------
       Total Success                  3                3                0
                test                  3                3                0
 Total In-Delivering                  0                0                0
      Total In-Queue                  0                0                0
Total Dropped                         8                8                0
           Tx Failed                  8                8                0
       configuration                  2                2                0
                test                  6                6                0
Event Types             Total
--------------------   ----------------
      Total Detected                10
       configuration                 1
                test                 9
 Total In-Processing                 0
      Total In-Queue                 0
Total Dropped                        0
Last call-home message sent time: 2013-04-15 05:37:16 GMT+00:00
B:*******************************************************************
Message Types           Total            Email             HTTP
--------------------   ---------------- ---------------- ----------------
       Total Success                  1                1                0
                test                  1                1                0
 Total In-Delivering                  0                0                0
      Total In-Queue                  0                0                0
Total Dropped                         2                2                0
           Tx Failed                  2                2                0
       configuration                  2                2                0
Event Types             Total
--------------------   ----------------
      Total Detected                 2
       configuration                 1
                test                 1
 Total In-Processing                 0
      Total In-Queue                 0
Total Dropped                        0
Last call-home message sent time: 2013-04-15 05:36:16 GMT+00:00
C:*******************************************************************
Message Types           Total            Email             HTTP
--------------------   ---------------- ---------------- ----------------
       Total Success                  0                0                0
 Total In-Delivering                  0                0                0
      Total In-Queue                  0                0                0
Total Dropped                         2                2                0
           Tx Failed                  2                2                0
       configuration                  2                2                0
Event Types             Total
--------------------   ----------------
      Total Detected                 1
       configuration                 1
 Total In-Processing                 0
      Total In-Queue                 0
Total Dropped                        0
Last call-home message sent time: n/a
```

```
D:*********************************************************************
Message Types           Total            Email             HTTP
--------------------   ---------------- ---------------- ----------------
       Total Success                  1                1                0
                test                  1                1                0
 Total In-Delivering                  0                0                0
      Total In-Queue                  0                0                0
Total Dropped                         2                2                0
           Tx Failed                  2                2                0
       configuration                  2                2                0
Event Types             Total
--------------------   ----------------
      Total Detected                  2
       configuration                  1
                test                  1
 Total In-Processing                  0
      Total In-Queue                  0
Total Dropped                         0
Last call-home message sent time: 2013-04-15 05:35:34 GMT+00:00
ciscoasa(config)#
```

Related Commands

Command	Description
call-home	Enters call home configuration mode.
call-home send alert-group	Sends a specific alert group message.
service call-home	Enables or disables Call Home.

show call-home registered-module status

To display the registered module status, use the **show call-home** registered-module status command in privileged EXEC mode.

show call-home registered-module status [**all**]

Note The [**all**] option is only valid in system context mode.

Syntax Description

all	Displays module status based on the device, not per context. In multiple context mode, if a module is enabled in at least one context, it is displayed as enabled if the "**all**" option is included.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	—	• Yes

Command History

Release	Modification
8.2(2)	This command was added.

Examples

The following example displays the **show call-home** registered-module status **all** output:

```
Output:
Module Name  Status
------------------------------------- --------------------Smart Call-Home enabledFailover
 Standby/Active
```

Related Commands

Command	Description
call-home	Enters call-home configuration mode.
call-home send alert-group	Sends a specific alert group message.
service call-home	Enables or disables Call Home.

show capture

To display the capture configuration when no options are specified, use the **show capture** command in privileged EXEC mode.

[**cluster exec**] **show capture** [*capture_name*] [**access-list** *access_list_name*] [**count** *number*] [**decode**] [**detail**] [**dump**] [**packet-number** *number*] [**trace**]

Syntax Description

access-list *access_list_name*	(Optional) Displays information for packets that are based on IP or higher fields for the specific access list identification .
capture_name	(Optional) Specifies the name of the packet capture.
cluster exec	(Optional) In a clustering environment, enables you to issue the **show capture** command in one unit and run the command in all the other units at the same time.
count *number*	(Optional) Displays the number of packets specified data.
decode	This option is useful when a capture of type isakmp is applied to an interface. All ISAKMP data flowing through that interface will be captured after decryption and shown with more information after decoding the fields.
detail	(Optional) Displays additional protocol information for each packet.
dump	(Optional) Displays a hexadecimal dump of the packets that are transported over the data link.
packet-number *number*	Starts the display at the specified packet number.
trace	Displays extended trace information for each packet.

Command Default

This command has no default settings.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
7.0(1)	This command was added.
8.4(2)	Detailed information in the output for IDS was added.

Release	Modification
9.0(1)	The **cluster exec** option was added.
9.2(1)	The **vpn-user** domain name was changed to **filter-aaa** in the output.
9.3(1)	Output for SGT plus Ethernet Tagging was added.
9.10(1)	IP decode support for GRE and IPinIP encapsulation was added.
9.13(1)	The show capture for asp-drop capture type was enhanced to include location details of the drop.
9.20(2)	The show capture detail for the physical port displays the drop configuration (disable or mac-filter).

Usage Guidelines

If you specify the *capture_name* , then the capture buffer contents for that capture are displayed.

The **dump** keyword does not display MAC information in the hexadecimal dump.

The decoded output of the packets depend on the protocol of the packet. Typically, this command supports IP decode for the ICMP, UDP, and TCP protocols. From version 9.10(1), this command is enhanced to support the display of the IP decode output for GRE and IPinIP encapsulation on ICMP, UDP, and TCP.

In Table 4-29, the bracketed output is displayed when you specify the **detail** keyword.

Table 29: Packet Capture Output Formats

Packet Type	Capture Output Format
802.1Q	*HH:MM:SS.ms* [ether-hdr] *VLAN-info encap-ether-packet*
ARP	*HH:MM:SS.ms* [ether-hdr] *arp-type arp-info*
IP/ICMP	*HH:MM:SS.ms* [ether-hdr] *ip-source* > *ip-destination:* icmp: *icmp-type icmp-code* [checksum-failure]
IP/UDP	*HH:MM:SS.ms* [ether-hdr] *src-addr* . *src-port dest-addr* . *dst-port* : [checksum-info] udp *payload-len*
IP/TCP	*HH:MM:SS.ms* [ether-hdr] *src-addr* . *src-port* d *est-addr* . *dst-port* : *tcp-flags* [header-check] [checksum-info] *sequence-number ack-number tcp-window urgent-info tcp-options*
IP/GRE	**ICMP encapsulated with GRE:** *HH:MM:SS.ms* [ether-hdr] *carrier-ip-source> carrier-ip-destination: gre:* [gre-flags] *ip-source > ip-destination: icmp: icmp-type icmp-code* [checksum-failure] **UDP encapsulated with GRE:** *HH:MM:SS.ms* [ether-hdr] *carrier-ip-source> carrier-ip-destination: gre:* [gre-flags] *src-addr.src-port dest-addr.dst-port:* [checksum-info] *udp payload-len* **TCP encapsulated with GRE:** *HH:MM:SS.ms* [ether-hdr] *carrier-ip-source> carrier-ip-destination: gre:* [gre-flags] *src-addr.src-port dest-addr.dst-port: tcp-flags* [header-check] [checksum-info] *sequence-number ack-number tcp-window urgent-info tcp-options*

Packet Type	Capture Output Format
IP/IPinIP	**ICMP encapsulated with IPinIP:** *HH:MM:SS.ms* [ether-hdr] *carrier-ip-source> carrier-ip-destination: ipip-proto-4: ip-source > ip-destination: icmp: icmp-type icmp-code* [checksum-failure] **UDP encapsulated with IPinIP:** *HH:MM:SS.ms* [ether-hdr] *carrier-ip-source> carrier-ip-destination: ipip-proto-4: src-addr.src-port dest-addr.dst-port:* [checksum-info] *udp payload-len* **TCP encapsulated with IPinIP:** *HH:MM:SS.ms* [ether-hdr] *carrier-ip-source> carrier-ip-destination: ipip-proto-4: src-addr.src-port dest-addr.dst-port: tcp-flags* [header-check] [checksum-info] *sequence-number ack-number tcp-window urgent-info tcp-options*
IP/Other	*HH:MM:SS.ms* [ether-hdr] *src-addr dest-addr* : *ip-protocol ip-length*
Other	*HH:MM:SS.ms ether-hdr* : *hex-dump*

Usage Guidelines

If the ASA receives packets with an incorrectly formatted TCP header and drops them because of the ASP drop reason *invalid-tcp-hdr-length* , the **show capture** command output on the interface where those packets are received does not show those packets.

From version 9.13(1), to facilitate troubleshooting, the show capture output is enhanced to include the drop location information when showing the capture type of asp-drop. While troubleshooting using ASP drop counters, the exact location of the drop is unknown, especially when the same ASP drop reason is used in many different places. This information is critical in finding root cause of the drop. With this enhancement, the ASP drop details such as the build target, ASA release number, hardware model, and ASLR memory text region (to facilitate the decode of drop location) are shown.

Examples

This example shows how to display the capture configuration:

```
ciscoasa(config)# show capture
capture arp ethernet-type arp interface outside
capture http access-list http packet-length 74 interface inside
```

This example shows how to display the packets that are captured by an ARP capture:

```
ciscoasa(config)# show capture arp
2 packets captured
19:12:23.478429 arp who-has 171.69.38.89 tell 171.69.38.10
19:12:26.784294 arp who-has 171.69.38.89 tell 171.69.38.10
2 packets shown
```

The following example shows how to display the packets that are captured on a single unit in a clustering environment:

```
ciscoasa(config)# show capture
capture 1 cluster type raw-data interface primary interface cluster [Buffer Full - 524187
bytes]
capture 2 type raw-data interface cluster [Capturing - 232354 bytes]
```

The following example shows how to display the packets that are captured on all units in a clustering environment:

```
ciscoasa(config)# cluster exec
            show capture
mycapture (LOCAL):--------------------------------------------------------------
capture 1 type raw-data interface primary [Buffer Full - 524187 bytes]
capture 2 type raw-data interface cluster [Capturing - 232354 bytes]
yourcapture:--------------------------------------------------------------------
capture 1 type raw-data interface primary [Capturing - 191484 bytes]
capture 2 type raw-data interface cluster [Capturing - 532354 bytes]
```

The following example shows the packets that are captured on the cluster control link in a clustering environment after the following commands are entered:

```
ciscoasa (config)# capture a interface cluster
ciscoasa (config)# capture cp interface cluster match udp any eq 49495 any
ciscoasa (config)# capture cp interface cluster match udp any any eq 49495
ciscoasa (config)# access-list cc1 extended permit udp any any eq 4193
ciscoasa (config)# access-list cc1 extended permit udp any eq 4193 any
ciscoasa (config)# capture dp interface cluster access-list cc1
ciscoasa (config)# capture lacp type lacp interface gigabitEthernet 0/0
ciscoasa(config)# show capture
capture a type raw-data interface cluster [Capturing - 970 bytes]
capture cp type raw-data interface cluster [Capturing - 26236 bytes]
  match udp any eq 49495 any
capture dp type raw-data access-list cc1 interface cluster [Capturing - 4545230 bytes]
capture lacp type lacp interface gigabitEthernet0/0 [Capturing - 140 bytes]
```

The following example shows the packets that are captured when SGT plus Ethernet tagging has been enabled on an interface:

```
ciscoasa(config)# show capture my-inside-capture
1: 11:34:42.931012 INLINE-TAG 36 10.0.101.22 > 11.0.101.100: icmp: echo request
2: 11:34:42.931470 INLINE-TAG 48 11.0.101.100 > 10.0.101.22: icmp: echo reply
3: 11:34:43.932553 INLINE-TAG 36 10.0.101.22 > 11.0.101.100: icmp: echo request
4: 11.34.43.933164 INLINE-TAG 48 11.0.101.100 > 10.0.101.22: icmp: echo reply
```

When SGT plus Ethernet tagging has been enabled on an interface, the interface can still receive tagged or untagged packets. The example shown is for tagged packets, which have INLINE-TAG 36 in the output. When the same interface receives untagged packets, the output remains unchanged (that is, no “INLINE-TAG 36” entry is included in the output).

The following example shows the GRE, IPinIP, and other packets that are generated by packet tracer and the subsequent capture output on interface inside:

GRE packets:

```
packet-tracer input inside gre ipv4 31.1.1.6 32.1.1.6 tcp 1.1.1.1 1234 2.2.2.2 80
packet-tracer input inside gre ipv4 31.1.1.6 32.1.1.6 udp 1.1.1.1 1234 2.2.2.2 80
packet-tracer input inside gre ipv4 31.1.1.6 32.1.1.6 icmp 1.1.1.1 8 0 2.2.2.2
```

IPinIP Packets:

```
packet-tracer input inside ipip 31.1.1.6 32.1.1.6 tcp 1.1.1.1 1234 2.2.2.2 80
```

```
packet-tracer input inside ipip 31.1.1.6 32.1.1.6 udp 1.1.1.1 1234 2.2.2.2 80
packet-tracer input inside ipip 31.1.1.6 32.1.1.6 icmp 1.1.1.1 8 0 2.2.2.2
```

Regular tcp/udp/icmp packets:

```
packet-tracer input inside tcp 1.1.1.1 1234 2.2.2.2 80
packet-tracer input inside udp 1.1.1.1 1234 2.2.2.2 80
packet-tracer input inside icmp 1.1.1.1 8 0 2.2.2.2


ciscoasa(config)# show capture inside
12:10:37.523746       31.1.1.6 > 32.1.1.6: gre: 1.1.1.1.1234 > 2.2.2.2.80: S
2145492151:2145492151(0) win 8192
12:10:37.623624       31.1.1.6 > 32.1.1.6: gre: 1.1.1.1.1234 > 2.2.2.2.80:  udp 0
12:10:37.714471       31.1.1.6 > 32.1.1.6: gre: 1.1.1.1 > 2.2.2.2 icmp: echo request
12:10:37.806690       31.1.1.6 > 32.1.1.6: ipip-proto-4: 1.1.1.1.1234 > 2.2.2.2.80: S
1501131661:1501131661(0) win 8192
12:10:37.897673       31.1.1.6 > 32.1.1.6: ipip-proto-4: 1.1.1.1.1234 > 2.2.2.2.80:  udp 0

12:10:41.974604      31.1.1.6 > 32.1.1.6: ipip-proto-4: 1.1.1.1 > 2.2.2.2 icmp: echo request
12:16:14.957225       1.1.1.1.1234 > 2.2.2.2.80: S 2091733697:2091733697(0) win 8192
12:16:15.023909       1.1.1.1.1234 > 2.2.2.2.80:  udp 0
12:16:15.090449       1.1.1.1 > 2.2.2.2 icmp: echo request
```

Note The GRE and IPinIP packets are decoded using the TCP/UDP/ICMP decode functionality to display the inner packet.

The following example shows the enhancement made to the output of this command. The drop location program counter or current instruction (which is later decoded), the build target, ASA release number, hardware model, and ASLR memory text region are captured and displayed to facilitate the decode of drop location:

```
ciscoasa(config)# capture gtp_drop type asp-drop inspect-gtp
ciscoasa(config)# show capture gtp_drop [trace]
Target:         SSP
Hardware:   FPR4K-SM-12
Cisco Adaptive Security Appliance Software Version 9.13.1
ASLR enabled, text region 55cd421df000-55cd47530ea9
1 packets captured

1: 15:55:58.522983       192.168.108.12.41245 > 171.70.168.183.2123:  udp 27 Drop-reason:
(inspect-gtp) GTP inspection, Drop-location: frame 0x000055cd43db4019 flow (NA)/NA

ciscoasa(config)# show capture gtp_drop trace detail
Target:         SSP
Hardware:   FPR4K-SM-12
Cisco Adaptive Security Appliance Software Version 9.13.1
ASLR enabled, text region 55cd421df000-55cd47530ea9
1 packets captured

1: 15:55:58.522983 0050.56b0.bd99 5057.a884.2beb 0x0800 Length: 69
192.168.108.12.41245 > 171.70.168.183.2123:  [udp sum ok] udp 27 (ttl 64, id 53384)
Drop-reason: (inspect-gtp) GTP inspection, Drop-location: frame 0x000055cd43db4019 flow
(NA)/NA
```

The following example shows the packet captured with the mac-filter drop enabled:

```
ciscoasa(config)# show capture test detail
Packet Capture info
Name:test
Session: 3
Admin State: disabled
OperState:down
OperState Reason: Session_Admin_Shut
Config Success:yes
Config Fail Reason:
Append Flag: overwrite
Session Mem Usage: 256
Session PcapSnap Len: 1518
Error Code:0
Drop Count:0
Total Physical ports involved in Packet Capture: 1

Physical port:
Slot Id: 1
Port Id: 1
Pcapfile:/mnt/disk0/packet-capture/sess-3-test-ethernet-1-1-0.
pcapPcapsize:0
Drop:mac-filter
Filter:test-1-1
Packet Capture Filter Info
Name:test-1-1
Protocol:0
Ivlan: 0
Ovlan: 3178
SrcIp:0.0.0.0
DestIp: 0.0.0.0
SrcIpv6:::
DestIpv6: ::
SrcMAC: 00:00:00:00:00:00
DestMAC:00:00:00:00:00:00
SrcPort:0
DestPort: 0
Ethertype: 0
Total Physical breakout ports involved in Packet Capture: 0
0 packet captured on disk using switch capture
```

Related Commands

Command	Description
capture	Enables packet capture capabilities for packet sniffing and network fault isolation.
clear capture	Clears the capture buffer.
copy capture	Copies a capture file to a server.

show chardrop

To display the count of characters dropped from the serial console, use the **show chardrop** command in privileged EXEC mode.

show chardrop

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
7.2(1)	This command was added.

Examples

The following is sample output from the **show chardrop** command:

```
ciscoasa# show chardrop
Chars dropped pre-TxTimeouts: 0, post-TxTimeouts: 0
```

Related Commands

Command	Description
show running-config	Shows the current operating configuration.

show checkheaps

To show the checkheaps statistics, use the **show checkheaps** command in privileged EXEC mode. Checkheaps is a periodic process that verifies the sanity of the heap memory buffers (dynamic memory is allocated from the system heap memory region) and the integrity of the code region.

show checkheaps

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	—	• Yes

Command History

Release	Modification
7.0(1)	This command was added.

Examples

The following is sample output from the **show checkheaps** command:

```
ciscoasa# show checkheaps
Checkheaps stats from buffer validation runs
--------------------------------------------
Time elapsed since last run     : 42 secs
Duration of last run            : 0 millisecs
Number of buffers created       : 8082
Number of buffers allocated     : 7808
Number of buffers free          : 274
Total memory in use             : 43570344 bytes
Total memory in free buffers    : 87000 bytes
Total number of runs            : 310
```

Related Commands

Command	Description
checkheaps	Sets the checkheap verification intervals.

show checksum

To display the configuration checksum, use the **show checksum** command in privileged EXEC mode.

show checksum

Syntax Description

This command has no arguments or keywords.

Command Default

This command has no default settings.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	• Yes	—

Command History

Release	Modification
7.2(1)	This command was added.

Usage Guidelines

The show checksum command allows you to display four groups of hexadecimal numbers that act as a digital summary of the configuration contents. This checksum is calculated only when you store the configuration in flash memory.

If a dot (“.”) appears before the checksum in the show config or show checksum command output, the output indicates a normal configuration load or write mode indicator (when loading from or writing to the ASA flash partition). The “.” shows that the ASA is preoccupied with the operation but is not “hung up.” This message is similar to a “system processing, please wait” message.

Examples

This example shows how to display the configuration or the checksum:

```
ciscoasa(config)# show checksum
Cryptochecksum: 1a2833c0 129ac70b 1a88df85 650dbb81
```

show chunkstat

To display the chunk statistics, use the **show chunkstat** command in privileged EXEC mode.

show chunkstat

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	—	• Yes

Command History

Release	Modification
7.0(1)	This command was added.

Usage Guidelines

The output of this command is primarily for internal development usage, and it is not meaningful for customers.

Examples

This example shows how to display the chunk statistics:

```
ciscoasa# show chunkstat
Global chunk statistics: created 181, destroyed 34, siblings created 94, siblings destroyed
 34
Per-chunk statistics: siblings created 0, siblings trimmed 0
Dump of chunk at 01edb4cc, name "Managed Chunk Queue Elements", data start @ 01edbd24, end
 @ 01eddc54
next: 01eddc8c, next_sibling: 00000000, prev_sibling: 00000000
flags 00000001
maximum chunk elt's: 499, elt size: 16, index first free 498
# chunks in use: 1, HWM of total used: 1, alignment: 0
Per-chunk statistics: siblings created 0, siblings trimmed 0
Dump of chunk at 01eddc8c, name "Registry Function List", data start @ 01eddea4, end @
01ede348
next: 01ede37c, next_sibling: 00000000, prev_sibling: 00000000
flags 00000001
maximum chunk elt's: 99, elt size: 12, index first free 42
# chunks in use: 57, HWM of total used: 57, alignment: 0
```

Related Commands

Command	Description
show counters	Displays the protocol stack counters.
show cpu	Displays the CPU utilization information.

show class

To show the contexts assigned to a class, use the **show class** command in privileged EXEC mode.

show class *name*

Syntax Description

name	Specifies the name as a string up to 20 characters long. To show the default class, enter **default** for the name.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	—	—	• Yes

Command History

Release	Modification
7.2(1)	This command was added.

Examples

The following is sample output from the **show class default** command:

```
ciscoasa# show class default
Class Name           Members    ID   Flags
default                All       1    0001
```

Related Commands

Command	Description
class	Configures a resource class.
clear configure class	Clears the class configuration.
context	Configures a security context.
limit-resource	Sets the resource limit for a class.
member	Assigns a context to a resource class.

show clns

To show Connectionless-mode Network Service (CLNS) information for IS-IS, use the **show clns** command in privileged EXEC mode.

show clns { **filter-set | interface** [*interface_name*] | **is-neighbors** [*interface_name*] [**detail**] | **neighbors** [**areas**] [*interface_name*] [**detail**] | **protocol** [*domain*] | **traffic** [**since** { **bootup | show** }] }

Syntax Description

areas	(Optional) Shows CLNS multiarea adjacencies.
bootup	Shows CLNS protocol statistics since bootup.
detail	(Optional) Shows the areas associated with the intermediate systems. Otherwise, a summary display is provided.
domain	(Optional) Shows routing protocol process information for a CLNS domain.
filter-set	Shows CLNS filter sets.
interface	Shows CLNS interface status and configuration.
interface_name	(Optional) Specifies the interface name.
is-neighbors	Shows IS neighbor adjacencies. Neighbor entries are sorted according to the area in which they are located.
neighbors	Displays end system (ES), intermediate system (IS), and multitopology Integrated Intermediate System-to-Intermediate System (M-ISIS) neighbors
protocol	Shows CLNS routing protocol process information. There will always be at least two routing processes, a Level 1 and a Level 2, and there can be more.
show	Shows CLNS protocol statistics since the last time you used this **show** command.
since	(Optional) Shows CLNS protocol statistics since either bootup or the last time you used this **show** command.
traffic	Lists the CLNS packets that this router has seen.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	—	• Yes	• Yes	—

Command History

Release	Modification
9.6(1)	We introduced this command.

Usage Guidelines

This command shows CLNS information for IS-IS.

Examples

The following display assumes filter sets have been defined with the following commands:

```
ciscoasa(config)# clns filter-set US-OR-NORDUNET 47.0005..
.
ciscoasa(config)# clns filter-set US-OR-NORDUNET 47.0023...

ciscoasa(config)# clns filter-set LOCAL 49.0003...
```

The following is a sample output from the **show clns filter-set** command:

```
ciscoasa# show clns filter-set
CLNS filter set US-OR-NORDUNET
      permit 47.0005...
      permit 47.0023...
CLNS filter set LOCAL
      permit 49.0003...
```

The following is sample output from the **show clns interface** command that includes information for Token Ring and serial interfaces:

```
ciscoasa# show clns interface
GigabitEthernet0/1 is up, line protocol is up
  Checksums enabled, MTU 1500
  ERPDUs enabled, min. interval 10 msec.
  DEC compatibility mode OFF for this interface
  Next ESH/ISH in 0 seconds
  Routing Protocol: IS-IS
    Circuit Type: level-1-2
    Interface number 0x0, local circuit ID 0x1
    Level-1 Metric: 10, Priority: 64, Circuit ID: c2.01
    DR ID: c2.01
    Level-1 IPv6 Metric: 10
    Number of active level-1 adjacencies: 3
    Level-2 Metric: 10, Priority: 64, Circuit ID: c2.01
    DR ID: c2.01
    Level-2 IPv6 Metric: 10
    Number of active level-2 adjacencies: 3
    Next IS-IS LAN Level-1 Hello in 1 seconds
    Next IS-IS LAN Level-2 Hello in 1 seconds
```

Table 30: show clns interface fields

Field	Description
GigabitEthernet0/1 is up, line protocol is up	Shown to be up, and the line protocol is up.
Checksums enabled	Can be enabled or disabled.
MTU	The number following maximum transmission unit (MTU) is the maximum transmission size for a packet on this interface.
ERPDUs	Displays information about the generation of error protocol data units (ERPDUs). They can be either enabled or disabled. If they are enabled, they are sent out no more frequently than the specified interval.
RDPDUs	Provides information about the generation of redirect protocol data units (RDPDUs). They can be either enabled or disabled. If they are enabled, they are sent out no more frequently than the specified interval. If the address mask is enabled, redirects are sent out with an address mask.
Congestion Experienced	Tells when CLNS will turn on the congestion experienced bit. The default is to turn this bit on when there are more than four packets in a queue.
DEC compatibility mode	Indicates whether Digital Equipment Corporation (DEC) compatibility has been enabled.
Next ESH/ISH	Displays when the next end system (ES) hello or intermediate system (IS) hello will be sent on this interface.
Routing Protocol	Lists the areas that this interface is in. In most cases, an interface will be in only one area.
Circuit Type	Indicates whether the interface has been configured for local routing (level 1), area routing (level 2), or local and area routing (level 1-2).
Interface number, local circuit ID; Level-1 Metric; DR ID; Level-1 IPv6 Metric; Number of active level-1 adjacencies; Level-2 Metric; DR ID; Level-2 IPv6 Metric; Number of active level-2 adjacencies; Next IS-IS LAN Level-1; Next IS-IS LAN Level-2	Last series of fields displays information pertaining to Intermediate System-to-Intermediate System (IS-IS). For IS-IS, the Level 1 and Level 2 metrics, priorities, circuit IDs, and number of active Level 1 and Level 2 adjacencies are specified.
BFD enabled	BFD has been enabled on the interface.

The following is sample output from the **show clns is-neighbors** command:

```
ciscoasa# show clns is-neighbors
System Id      Interface   State  Type Priority  Circuit Id         Format
```

```
CSR7001        inside      Up     L1L2 64/64    ciscoasa.01         Phase V
CSR7002        inside      Up     L1L2 64/64    ciscoasa.01         Phase V
```

Table 31: show clns is-neighbors Fields

Field	Description
System Id	Identification value of the system.
Interface	Interface on which the router was discovered.
State	Adjacency state. Up and Init are the states. See the show clns neighbors description.
Type	L1, L2, and L1L2 type adjacencies. See the **show clns neighbors** description.
Priority	IS-IS priority that the respective neighbor is advertising. The highest priority neighbor is elected the designated IS-IS router for the interface.
Circuit Id	Neighbor's idea of what the designated IS-IS router is for the interface.
Format	Indicates if the neighbor is either a Phase V (OSI) adjacency or Phase IV (DECnet) adjacency.

The following is sample output from the **show clns is-neighbors detail** command:

```
ciscoasa# show clns is-neighbors detail
System Id      Interface   State  Type Priority  Circuit Id         Format
CSR7001        inside      Up     L1L2 64/64    ciscoasa.01         Phase V
  Area Address(es): 49.0001
  IP Address(es):  1.3.3.3*
  Uptime: 00:12:49
  NSF capable
  Interface name: inside
CSR7002        inside      Up     L1L2 64/64    ciscoasa.01         Phase V
  Area Address(es): 49.0001
  IP Address(es):  20.3.3.3*
  Uptime: 00:12:50
  NSF capable
  Interface name: inside
```

The following is sample output from the **show clns neighbors detail** command:

```
ciscoasa# show clns neighbors detail
System Id      Interface   SNPA            State  Holdtime  Type Protocol
CSR7001        inside      000c.2921.ff44  Up          26   L1L2
  Area Address(es): 49.0001
  IP Address(es):  1.3.3.3*
  Uptime: 01:16:33
  NSF capable
  Interface name: inside
CSR7002        inside      000c.2906.491c  Up          27   L1L2
  Area Address(es): 49.0001
  IP Address(es):  20.3.3.3*
  Uptime: 01:16:33
  NSF capable
  Interface name: inside
```

The following is sample output from the **show clns neighbors** command:

```
ciscoasa# show clns neighbors
```

```
System Id      Interface   SNPA            State  Holdtime  Type Protocol
CSR7001        inside      000c.2921.ff44  Up          29   L1L2
CSR7002        inside      000c.2906.491c  Up          27   L1L2
```

Table 32: show clns neighbors Fields

Field	Description
System Id	Six-byte value that identifies a system in an area.
Interface	Interface name from which the system was learned.
SNPA	Subnetwork Point of Attachment. This is the data-link address.
State	State of the ES, IS, or M-ISIS. • Init—System is an IS and is waiting for an IS-IS hello message. IS-IS regards the neighbor as not adjacent. • Up—Believes the ES or IS is reachable.
Holdtime	Number of seconds before this adjacency entry times out.
Type	The adjacency type. Possible values are as follows: • ES—End-system adjacency either discovered via the ES-IS protocol or statically configured. • IS—Router adjacency either discovered via the ES-IS protocol or statically configured. • M-ISIS—Router adjacency discovered via the multitopology IS-IS protocol. • L1—Router adjacency for Level 1 routing only. • L1L2—Router adjacency for Level 1 and Level 2 routing. • L2—Router adjacency for Level 2 only.
Protocol	Protocol through which the adjacency was learned. Valid protocol sources are ES-IS, IS-IS, ISO IGRP, Static, DECnet, and M-ISIS.

The following is sample output from the **show clns protocol** command:

```
ciscoasa# show clns protocol
IS-IS Router
  System Id: 0050.0500.5008.00  IS-Type: level-1-2
  Manual area address(es):
        49.0001
  Routing for area address(es):
        49.0001
  Interfaces supported by IS-IS:
        outside - IP
  Redistribute:
    static (on by default)
  Distance for L2 CLNS routes: 110
  RRR level: none
  Generate narrow metrics: level-1-2
  Accept narrow metrics:   level-1-2
  Generate wide metrics:   none
  Accept wide metrics:     none
```

The following is sample output from the **show clns traffic** command:

```
ciscoasa# show clns traffic
CLNS:  Time since last clear: never
CLNS & ESIS Output: 0, Input: 8829
CLNS Local: 0, Forward: 0
CLNS Discards:
  Hdr Syntax: 0, Checksum: 0, Lifetime: 0, Output cngstn: 0
  No Route: 0, Discard Route: 0, Dst Unreachable 0, Encaps. Failed: 0
  NLP Unknown: 0, Not an IS: 0
CLNS Options: Packets 0, total 0 , bad 0, GQOS 0, cngstn exprncd 0
CLNS Segments:  Segmented: 0, Failed: 0
CLNS Broadcasts: sent: 0, rcvd: 0
Echos: Rcvd 0 requests, 0 replies
      Sent 0 requests, 0 replies
ESIS(sent/rcvd): ESHs: 0/0, ISHs: 0/0, RDs: 0/0, QCF: 0/0
Tunneling (sent/rcvd): IP: 0/0, IPv6: 0/0
Tunneling dropped (rcvd) IP/IPV6:  0
ISO-IGRP: Querys (sent/rcvd): 0/0 Updates (sent/rcvd): 0/0
ISO-IGRP: Router Hellos: (sent/rcvd): 0/0
ISO-IGRP Syntax Errors: 0
IS-IS: Time since last clear: never
IS-IS: Level-1 Hellos (sent/rcvd): 1928/1287
IS-IS: Level-2 Hellos (sent/rcvd): 1918/1283
IS-IS: PTP Hellos     (sent/rcvd): 0/0
IS-IS: Level-1 LSPs sourced (new/refresh): 7/13
IS-IS: Level-2 LSPs sourced (new/refresh): 7/14
IS-IS: Level-1 LSPs flooded (sent/rcvd): 97/2675
IS-IS: Level-2 LSPs flooded (sent/rcvd): 73/2628
IS-IS: LSP Retransmissions: 0
IS-IS: Level-1 CSNPs (sent/rcvd): 642/0
IS-IS: Level-2 CSNPs (sent/rcvd): 639/0
IS-IS: Level-1 PSNPs (sent/rcvd): 0/554
IS-IS: Level-2 PSNPs (sent/rcvd): 0/390
IS-IS: Level-1 DR Elections: 1
IS-IS: Level-2 DR Elections: 1
IS-IS: Level-1 SPF Calculations: 9
IS-IS: Level-2 SPF Calculations: 8
IS-IS: Level-1 Partial Route Calculations: 0
IS-IS: Level-2 Partial Route Calculations: 0
IS-IS: LSP checksum errors received: 0
IS-IS: Update process queue depth: 0/200
IS-IS: Update process packets dropped: 0
```

Table 33: show clns traffic Fields

Field	Description
CLNS & ESIS Output	Total number of packets that this router has sent.
Input	Total number of packets that this router has received.
CLNS Local	Lists the number of packets that were generated by this router.
Forward	Lists the number of packets that this router has forwarded.
CLNS Discards	Lists the packets that CLNS has discarded, along with the reason for the discard.
CLNS Options	Lists the options seen in CLNS packets.

Field	Description
CLNS Segments	Lists the number of packets segmented and the number of failures that occurred because a packet could not be segmented.
CLNS Broadcasts	Lists the number of CLNS broadcasts sent and received.
Echos	Lists the number of echo request packets and echo reply packets received. The line following this field lists the number of echo request packets and echo reply packets sent.
ESIS (sent/rcvd)	Lists the number of End System Hello (ESH), Intermediate System Hello (ISH), and redirects sent and received.
ISO IGRP	Lists the number of ISO Interior Gateway Routing Protocol (IGRP) queries and updates sent and received.
Router Hellos	Lists the number of ISO IGRP router hello packets sent and received.
IS-IS: Level-1 hellos (sent/rcvd)	Lists the number of Level 1 IS-IS hello packets sent and received.
IS-IS: Level-2 hellos (sent/rcvd)	Lists the number of Level 2 IS-IS hello packets sent and received.
IS-IS: PTP hellos (sent/rcvd)	Lists the number of point-to-point IS-IS hello packets sent and received over serial links.
IS-IS: Level-1 LSPs (sent/rcvd)	Lists the number of Level 1 link-state Protocol Data Unit (PDUs) sent and received.
IS-IS: Level-2 LSPs (sent/rcvd)	Lists the number of Level 2 link-state PDUs sent and received.
IS-IS: Level-1 CSNPs (sent/rcvd)	Lists the number of Level 1 Complete Sequence Number Packets (CSNP) sent and received.
IS-IS: Level-2 CSNPs (sent/rcvd)	Lists the number of Level 2 CSNPs sent and received.
IS-IS: Level-1 PSNPs (sent/rcvd)	Lists the number of Level 1 Partial Sequence Number Packets (PSNP) sent and received.
IS-IS: Level-2 PSNPs (sent/rcvd)	Lists the number of Level 2 PSNPs sent and received.
IS-IS: Level-1 DR Elections	Lists the number of times Level 1 designated router election occurred.
IS-IS: Level-2 DR Elections	Lists the number of times Level 2 designated router election occurred.
IS-IS: Level-1 SPF Calculations	Lists the number of times the Level 1 shortest-path-first (SPF) tree was computed.
IS-IS: Level-2 SPF Calculations	Lists the number of times the Level 2 SPF tree was computed.

Related Commands

Command	Description
advertise passive-only	Configures the ASA to advertise passive interfaces.
area-password	Configures an IS-IS area authentication password.
authentication key	Enables authentication for IS-IS globally.
authentication mode	Specifies the type of authentication mode used in IS-IS packets for the IS-IS instance globally.
authentication send-only	Configure the IS-IS instance globally to have authentication performed only on IS-IS packets being sent (not received).
clear isis	Clears IS-IS data structures.
default-information originate	Generates a default route into an IS-IS routing domain.
distance	Defines the administrative distance assigned to routes discovered by the IS-IS protocol.
domain-password	Configures an IS-IS domain authentication password.
fast-flood	Configures IS-IS LSPs to be full.
hello padding	Configures IS-IS hellos to the full MTU size.
hostname dynamic	Enables IS-IS dynamic hostname capability.
ignore-lsp-errors	Configures the ASA to ignore IS-IS LSPs that are received with internal checksum errors rather than purging the LSPs.
isis adjacency-filter	Filters the establishment of IS-IS adjacencies.
isis advertise-prefix	Advertises IS-IS prefixes of connected networks in LSP advertisements on an IS-IS interface.
isis authentication key	Enables authentication for an interface.
isis authentication mode	Specifies the type of authentication mode used in IS-IS packets for the IS-IS instance per interface
isis authentication send-only	Configure the IS-IS instance per interface to have authentication performed only on IS-IS packets being sent (not received).
isis circuit-type	Configures the type of adjacency used for the IS-IS.
isis csnp-interval	Configures the interval at which periodic CSNP packets are sent on broadcast interfaces.
isis hello-interval	Specifies the length of time between consecutive hello packets sent by IS-IS.
isis hello-multiplier	Specifies the number of IS-IS hello packets a neighbor must miss before the ASA declares the adjacency as down.

Command	Description
isis hello padding	Configures IS-IS hellos to the full MTU size per interface.
isis lsp-interval	Configures the time delay between successive IS-IS LSP transmissions per interface.
isis metric	Configures the value of an IS-IS metric.
isis password	Configures the authentication password for an interface.
isis priority	Configures the priority of designated ASAs on the interface.
isis protocol shutdown	Disables the IS-IS protocol per interface.
isis retransmit-interval	Configures the amount of time between retransmission of each IS-IS LSP on the interface.
isis retransmit-throttle-interval	Configures the amount of time between retransmissions of each IS-IS LSP on the interface.
isis tag	Sets a tag on the IP address configured for an interface when the IP prefix is put into an LSP.
is-type	Assigns the routing level for the IS-IS routing process.
log-adjacency-changes	Enables the ASA to generate a log message when an NLSP IS-IS adjacency changes state (up or down).
lsp-full suppress	Configures which routes are suppressed when the PDU becomes full.
lsp-gen-interval	Customizes IS-IS throttling of LSP generation.
lsp-refresh-interval	Sets the LSP refresh interval.
max-area-addresses	Configures additional manual addresses for an IS-IS area.
max-lsp-lifetime	Sets the maximum time that LSPs persist in the ASA's database without being refreshed.
maximum-paths	Configures multi-path load sharing for IS-IS.
metric	Globally changes the metric value for all IS-IS interfaces.
metric-style	Configures an ASA running IS-IS so that it generates and only accepts new-style, length, value objects (TLVs).
net	Specifies the NET for the routing process.
passive-interface	Configures a passive interface.
prc-interval	Customizes IS-IS throttling of PRCs.
protocol shutdown	Disables the IS-IS protocol globally so that it cannot form any adjacency on any interface and will clear the LSP database.

Command	Description
redistribute isis	Redistributes IS-IS routes specifically from Level 1 into Level 2 or from Level 2 into Level 1.
route priority high	Assigns a high priority to an IS-IS IP prefix.
router isis	Enables IS-IS routing.
set-attached-bit	Specifies constraints for when a Level 1-Level 2 router should set its attached bit.
set-overload-bit	Configures the ASA to signal other routers not to use it as an intermediate hop in their SPF calculations.
show clns	Shows CLNS-specific information.
show isis	Shows IS-IS information.
show route isis	Shows IS-IS routes.
spf-interval	Customizes IS-IS throttling of SPF calculations.
summary-address	Creates aggregate addresses for IS-IS.

show clock

To view the time on the ASA, use the **show clock** command in user EXEC mode.

show clock [**detail**]

Syntax Description

detail	(Optional) Indicates the clock source (NTP or user configuration) and the current summer-time setting (if any).

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
User EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
7.0(1)	This command was added.

Examples

The following is sample output from the **show clock** command:

```
ciscoasa# show clock
12:35:45.205 EDT Tue Jul 27 2004
```

The following is sample output from the **show clock detail** command:

```
ciscoasa# show clock detail
12:35:45.205 EDT Tue Jul 27 2004
Time source is user configuration
Summer time starts 02:00:00 EST Sun Apr 4 2004
Summer time ends 02:00:00 EDT Sun Oct 31 2004
```

Related Commands

Command	Description
clock set	Manually sets the clock on the ASA.
clock summer-time	Sets the date range to show daylight saving time.
clock timezone	Sets the time zone.
ntp server	Identifies an NTP server.

Command	Description
show ntp status	Shows the status of the NTP association.

show cluster

To view aggregated data for the entire cluster or other information, use the **show cluster** command in privileged EXEC mode.

show cluster [**chassis**] { **access-list** [*acl_name*] | **conn** [**count**] | **context** [*context_name*] | **cpu** [**usage**] **interface-mode | memory | resource | service-policy | traffic | xlate count** }

Syntax Description

access-list [*acl_name*]	Shows hit counters for access policies. To see the counters for a specific ACL, enter the *acl_name* .
chassis	For the Firepower 9300 ASA security module, shows the cluster information for the chassis.
conn [**count**]	Shows the aggregated count of in-use connections for all units. If you enter the **count** keyword, only the connection count is shown.
context [*context_name*]	Shows usage per security context in multiple context mode.
cpu [**usage**]	Shows CPU usage information.
interface-mode	Shows the cluster interface mode, either spanned or individual.
memory	Shows system memory utilization and other information.
resource usage	Shows system resources and usage.
service-policy	Shows the MPF service policy statistics.
traffic	Shows traffic statistics.
xlate count	Shows current translation information.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	—	• Yes

Command History

Release	Modification
9.0(1)	This command was added.

Release	Modification
9.4(1)	The **service-policy** keyword was added.
9.4(1.152)	The **chassis** keyword was added.
9.9(1)	The **context** keyword was added.

Usage Guidelines

See also the **show cluster info** and **show cluster user-identity** commands.

Examples

The following is sample output from the **show cluster access-list** command:

```
ciscoasa# show cluster access-list
hitcnt display order: cluster-wide aggregated result, unit-A, unit-B,  unit-C, unit-D
access-list cached ACL log flows: total 0, denied 0 (deny-flow-max 4096) alert-interval 300
access-list 101; 122 elements; name hash: 0xe7d586b5
access-list 101 line 1 extended permit tcp 192.168.143.0 255.255.255.0 any eq www (hitcnt=0,
 0, 0, 0, 0) 0x207a2b7d
access-list 101 line 2 extended permit tcp any 192.168.143.0 255.255.255.0 (hitcnt=0, 0,
0, 0, 0) 0xfe4f4947
access-list  101 line 3 extended permit tcp host 192.168.1.183 host 192.168.43.238 (hitcnt=1,
 0, 0, 0, 1) 0x7b521307
access-list 101 line 4 extended permit tcp host 192.168.1.116 host 192.168.43.238 (hitcnt=0,
 0, 0, 0, 0) 0x5795c069
access-list 101 line 5 extended permit tcp host 192.168.1.177 host 192.168.43.238 (hitcnt=1,
 0, 0, 1, 0) 0x51bde7ee
access list 101 line 6 extended permit tcp host 192.168.1.177 host 192.168.43.13 (hitcnt=0,
 0, 0, 0, 0) 0x1e68697c
access-list 101 line 7 extended permit tcp host 192.168.1.177 host 192.168.43.132 (hitcnt=2,
 0, 0, 1, 1) 0xc1ce5c49
access-list 101 line 8 extended permit tcp host 192.168.1.177 host 192.168.43.192 (hitcnt=3,
 0, 1, 1, 1) 0xb6f59512
access-list 101 line 9 extended permit tcp host 192.168.1.177 host 192.168.43.44 (hitcnt=0,
 0, 0, 0, 0) 0xdc104200
access-list 101 line 10 extended permit tcp host 192.168.1.112 host 192.168.43.44 (hitcnt=429,
 109, 107, 109, 104)
0xce4f281d
access-list 101 line 11 extended permit tcp host 192.168.1.170 host 192.168.43.238 (hitcnt=3,
 1, 0, 0, 2) 0x4143a818
access-list 101 line 12 extended permit tcp host 192.168.1.170 host 192.168.43.169 (hitcnt=2,
 0, 1, 0, 1) 0xb18dfea4
access-list 101 line 13 extended permit tcp host 192.168.1.170 host 192.168.43.229 (hitcnt=1,
 1, 0, 0, 0) 0x21557d71
access-list 101 line 14 extended permit tcp host 192.168.1.170 host 192.168.43.106 (hitcnt=0,
 0, 0, 0, 0) 0x7316e016
access-list 101 line 15 extended permit tcp host 192.168.1.170 host 192.168.43.196 (hitcnt=0,
 0, 0, 0, 0) 0x013fd5b8
access-list 101 line 16 extended permit tcp host 192.168.1.170 host 192.168.43.75 (hitcnt=0,
 0, 0, 0, 0) 0x2c7dba0d
```

To display the aggregated count of in-use connections for all units, enter:

```
ciscoasa# show cluster conn count
Usage Summary In Cluster:*********************************************
  200 in use (cluster-wide aggregated)
     cl2(LOCAL):************************************************************
  100 in use, 100 most used
  cl1:************************************************************
  100 in use, 100 most used
```

Related Commands

Command	Description
show cluster info	Shows cluster information.
show cluster user-identity	Shows cluster user identity information and statistics.

show cluster history

To view event history for the cluster, use the **show cluster history** command in privileged EXEC mode.

show cluster history [**brief**] [**latest** [*number*]] [**reverse**] [**time** [*year month day*] *hh* **:** *mm* **:** *ss*]

Syntax Description

brief	Shows cluster history without generic events.
latest [*number*]	Displays the latest events. By default, the device shows the last 512 events. You can limit the *number* of events, between 1 and 512.
reverse	Shows events in reverse order.
time [*year month day*] *hh***:***mm***:***ss*	Shows events before a specified date and time.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	—	• Yes

Command History

Release	Modification
9.0(1)	This command was added.
9.14(1)	The output was enhanced to show more details for troubleshooting.
9.16(1)	We added the **brief**, **latest**, **reverse**, **time** keywords.
9.19(1)	Command output was changed from **Master** and **Slave** to **Control_Node** and **Data_Node**.

Usage Guidelines

The following is sample output from the **show cluster history time** command:

```
asa1(cfg-cluster)# show cluster history time august 26 10:10:05
================================================================================
From State            To State              Reason
================================================================================

10:08:49 UTC Aug 26 2020
DISABLED              DISABLED              Disabled at startup
```

```
10:09:43 UTC Aug 26 2020
DISABLED              ELECTION              Enabled from CLI

10:10:01 UTC Aug 26 2020
ELECTION              ONCALL                Event: Cluster unit A state is CONTROL_NODE

10:10:02 UTC Aug 26 2020
ONCALL                DATA_NODE_COLD            Data Node proceeds with configuration sync

10:10:02 UTC Aug 26 2020
DATA_NODE_COLD            DATA_NODE_CONFIG          Client progression done

10:10:04 UTC Aug 26 2020
DATA_NODE_CONFIG          DATA_NODE_FILESYS         Configuration replication finished

10:10:05 UTC Aug 26 2020
DATA_NODE_FILESYS         DATA_NODE_BULK_SYNC       Client progression done
```

The following is sample output from the **show cluster history brief** command:

```
asa1(cfg-cluster)# show cluster history brief
==========================================================================
From State            To State              Reason
==========================================================================

10:08:49 UTC Aug 26 2020
DISABLED              DISABLED              Disabled at startup

10:09:43 UTC Aug 26 2020
DISABLED              ELECTION              Enabled from CLI



10:10:02 UTC Aug 26 2020
ONCALL                DATA_NODE_COLD            Data Node proceeds with configuration sync

10:10:02 UTC Aug 26 2020
DATA_NODE_COLD            DATA_NODE_CONFIG          Client progression done

10:10:04 UTC Aug 26 2020
DATA_NODE_CONFIG          DATA_NODE_FILESYS         Configuration replication finished

10:10:05 UTC Aug 26 2020
DATA_NODE_FILESYS         DATA_NODE_BULK_SYNC       Client progression done
```

The following is sample output from the **show cluster history latest** command:

```
asa1(cfg-cluster)# show cluster history latest 3
==========================================================================
From State            To State              Reason
==========================================================================
```

```
10:10:05 UTC Aug 26 2020
DATA_NODE_FILESYS         DATA_NODE_BULK_SYNC       Client progression done

10:10:04 UTC Aug 26 2020
DATA_NODE_CONFIG          DATA_NODE_FILESYS         Configuration replication finished

10:10:02 UTC Aug 26 2020
DATA_NODE_COLD            DATA_NODE_CONFIG          Client progression done
```

Related Commands

Command	Description
show cluster	Shows aggregated data for the entire cluster and other information.
show cluster info	Shows cluster information.
show cluster user-identity	Shows cluster user identity information and statistics.

show cluster info

To view cluster information, use the **show cluster info** command in privileged EXEC mode.

show cluster info [**auto-join** | **clients** | **conn-distribution** | **counters** | **flow-mobility** | **goid** [*options*] | **health** [**details**] | **incompatible-config** | **instance-type** | **loadbalance** | **load-monitor** | **old-members** | **packet-distribution** | **trace** [*options*] | **transport** { **asp** | **cp** [**detail**] } | **unit-join-acceleration incompatible-config**]

Syntax Description

auto-join	(Optional) Shows whether the cluster node will automatically rejoin the cluster after a time delay and if the failure conditions (such as waiting for the license, chassis health check failure, and so on) are cleared. If the node is permanently disabled, or if the node is already in the cluster, then this command will not show any output.
clients	(Optional) Shows the version of register clients.
conn-distribution	(Optional) Shows the connection distribution in the cluster.
flow-mobility counters	(Optional) Shows EID movement and flow owner movement information.
goid [*options*]	(Optional) Shows the global object ID database. Options include: classmap conn-set hwidb idfw-domain idfw-group interface policymap virtual-context
health [**details**]	(Optional) Shows health monitoring information. The details keyword shows the number heartbeat message failures.
incompatible-config	(Optional) Shows commands that are incompatible with clustering in the current running configuration. This command is useful before you enable clustering.
instance-type	(Optional) Shows the ASA Virtual CPU cores and RAM per cluster node.
loadbalance	(Optional) Shows load balancing information.
load-monitor	(Optional) Monitors the traffic load for cluster nodes, including total connection count, CPU and memory usage, and buffer drops. If the load is too high, you can choose to manually disable clustering on the node if the remaining nodes can handle the load, or adjust the load balancing on the external switch. This feature is enabled by default with the **load-monitor** command.

old-members	(Optional) Shows former nodes of the cluster.
packet-distribution	(Optional) Shows packet distribution in the cluster.
trace [*options*]	(Optional) Shows the clustering control module event trace. Options include: • **latest** [*number*]—Displays the latest *number* events, where the number is from 1 to 2147483647. The default is to show all. • **level** *level* —Filters events by level where the *level* is one of the following: **all** , **critical** , **debug** , **informational** , or **warning** . • **module** *module* —Filters events by module where the *module* is one of the following: **ccp** , **datapath** , **fsm** , **general** , **hc** , **license** , **rpc** , **or transport** . • **time** {[*month day*] [*hh* **:** *mm* **:** *ss*]}—Shows events before the specified time or date.
transport { **asp** | **cp** [**detail**]}	(Optional) Show transport related statistics for the following: • **asp** —Data plane transport statistics. • **cp** —Control plane transport statistics. If you enter the **detail** keyword, you can view cluster reliable transport protocol usage so you can identify packet drop issues when the buffer is full in the control plane.
unit-join-acceleration incompatible-config	(Optional) When the **unit join-acceleration** command is enabled (the default), some configuration commands are not compatible with accelerated cluster joining; if these commands are present on the node, even if accelerated cluster joining is enabled, configuration syncing will always occur. You must remove the incompatible configuration for accelerated cluster joining to work. Use the **show cluster info unit-join-acceleration incompatible-config** command to view incompatible configuration.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	—	• Yes

Command History

Release	Modification
9.0(1)	This command was added.
9.3(1)	Improved support for modules in the **show cluster info health** command was added.
9.5(1)	Site ID information was added to the output.
9.5(2)	The **flow-mobility counters** keywords were added.
9.8(1)	The **health details** keyword was added.
9.9(2)	Added the **auto-join** keyword
9.9(2)	Added the **detail** keyword for **transport cp** .
9.13(1)	Added **load-monitor** and **unit-join-acceleration incompatible-config** keywords.
9.17(1)	Added the **instance-type** keyword for the ASA Virtual, and added ASA Virtual information to the **show cluster info** output.
9.19(1)	Command output was changed from **Master** and **Slave** to **Control_Node** and **Data_Node**.
9.24(1)	The **show cluster info trace** command output is enhanced to display priority-polling status.

Usage Guidelines

If you do not specify any options, the **show cluster info** command shows general cluster information including the cluster name and status, the cluster nodes, the node states, and so on.

Clear statistics using the **clear cluster info** command.

See also the **show cluster** and **show cluster user-identity** commands.

Examples

The following is sample output from the **show cluster info** command for a hardware platform:

```
ciscoasa# show cluster info
Cluster stbu: On
  This is "C" in state DATA_NODE
      ID        : 0
      Site ID : 1
      Version   : 9.5(1)
      Serial No.: P3000000025
      CCL IP    : 10.0.0.3
      CCL MAC   : 000b.fcf8.c192
      Last join : 17:08:59 UTC Sep 26 2011
      Last leave: N/A
Other members in the cluster:
  Unit "D" in state DATA_NODE
      ID        : 1
      Site ID : 1
      Version   : 9.5(1)
      Serial No.: P3000000001
      CCL IP    : 10.0.0.4
      CCL MAC   : 000b.fcf8.c162
      Last join : 19:13:11 UTC Sep 23 2011
      Last leave: N/A
  Unit "A" in state CONTROL_NODE
      ID        : 2
      Site ID : 2
```

```
      Version   : 9.5(1)
      Serial No.: JAB0815R0JY
      CCL IP    : 10.0.0.1
      CCL MAC   : 000f.f775.541e
      Last join : 19:13:20 UTC Sep 23 2011
      Last leave: N/A
  Unit "B" in state DATA_NODE
      ID        : 3
      Site ID : 2
      Version   : 9.5(1)
      Serial No.: P3000000191
      CCL IP    : 10.0.0.2
      CCL MAC   : 000b.fcf8.c61e
      Last join : 19:13:50 UTC Sep 23 2011
      Last leave: 19:13:36 UTC Sep 23 2011
```

The following is sample output from the **show cluster info** command for the ASA Virtual:

```
Cluster ASAv-cluster: On
    Interface mode: individual
Cluster Member Limit : 16
    This is "A" in state CONTROL_NODE
        ID        : 0
        Version   : 9.17(1)79
        Serial No.: 9A3GVQ1EL7W
        CCL IP    : 10.1.1.24
        CCL MAC   : 0050.56a8.7d4f
        Module    : ASAv
        Resource  : 2 cores / 4096 MB RAM
        Last join : 16:27:46 UTC Feb 18 2021
        Last leave: N/A
Other members in the cluster:
    Unit "B" in state DATA_NODE
        ID        : 1
        Version   : 9.17(1)79
        Serial No.: 9ACE28176EE
        CCL IP    : 10.1.1.25
        CCL MAC   : 0050.56a8.883e
        Module    : ASAv
        Resource  : 2 cores / 4096 MB RAM
        Last join : 20:29:25 UTC Feb 19 2021
        Last leave: 16:31:46 UTC Feb 19 2021
```

The following is sample output from the **show cluster info incompatible-config** command:

```
ciscoasa(cfg-cluster)# show cluster info incompatible-config
INFO: Clustering is not compatible with following commands which given a user's confirmation
 upon enabling clustering, can be removed automatically from running-config.
policy-map global_policy
 class scansafe-http
  inspect scansafe http-map fail-close
policy-map global_policy
 class scansafe-https
  inspect scansafe https-map fail-close
INFO: No manually-correctable incompatible configuration is found.
```

The following is sample output from the **show cluster info trace** command:

```
ciscoasa# show cluster info trace
 Feb 02 14:19:47.456 [DBUG]Receive CCP message: CCP_MSG_LOAD_BALANCE
 Feb 02 14:19:47.456 [DBUG]Receive CCP message: CCP_MSG_LOAD_BALANCE
```

```
 Feb 02 14:19:47.456 [DBUG]Send CCP message to all: CCP_MSG_KEEPALIVE from 80-1 at
CONTROL_NODE
```

The following is sample output from the **show cluster info health** command on the ASA 5500-X:

```
ciscoasa# show cluster info health
Member ID to name mapping:
  0 - A   1 - B(myself)
                         0        1
GigabitEthernet0/0       up       up
Management0/0            up       up
ips (policy off)         up       None
sfr (policy off)         None     up
Unit overall             healthy   healthy
Cluster overall          healthy
```

The above output lists both ASA IPS (ips) and ASA FirePOWER (sfr) modules, and for each module the ASA shows “policy on” or “policy off” to show if you configured the module in the service policy. For example:

```
class-map sfr-class
match sfr-traffic
policy-map sfr-policy
class sfr-class
sfr inline fail-close
service-policy sfr interface inside
```

With the above configuration, the ASA FirePOWER module (“sfr”) will be displayed as “policy on”. If one cluster node has a module as “up”, and the other node has the module as “down” or “None”, then the node with the down module will be kicked out of the cluster. However, if the service policy is not configured, then the cluster node would not be kicked out of the cluster; the module status is only relevant if the module is running.

The following is sample output from the **show cluster info health** command on the ASA 5585-X:

```
ciscoasa# show cluster info health
spyker-13# sh clu info heal
Member ID to name mapping:
  0 - A(myself) 1 - B
                       0 1
GigabitEthernet0/0       upup
SSM Card (policy off)    upup
Unit overall             healthyhealth
Cluster overall         healthyhealth
```

If you configure the module in the service policy, then the output shows “policy on”. If you do not configure the service policy, then the output shows “policy off”, even if a module is present in the chassis.

The following is sample output from the **show cluster info flow-mobility counters** command:

```
ciscoasa# show cluster info flow-mobility counters
EID movement notification received  : 0
EID movement notification processed : 0
Flow owner moving requested         : 0
```

The following is sample output from the **show cluster info auto-join** command:

```
ciscoasa(cfg-cluster)# show cluster info auto-join
Unit will try to join cluster in 253 seconds.
```

```
Quit reason: Received control message DISABLE

ciscoasa(cfg-cluster)# show cluster info auto-join

Unit will try to join cluster when quit reason is cleared.
Quit reason: Control Node has application down that data node has up.

ciscoasa(cfg-cluster)# show cluster info auto-join

Unit will try to join cluster when quit reason is cleared.
Quit reason: Chassis-blade health check failed.

ciscoasa(cfg-cluster)# show cluster info auto-join

Unit will try to join cluster when quit reason is cleared.
Quit reason: Service chain application became down.

ciscoasa(cfg-cluster)# show cluster info auto-join
Unit will try to join cluster when quit reason is cleared.
Quit reason: Unit is kicked out from cluster because of Application health check failure.

ciscoasa(cfg-cluster)# show cluster info auto-join
Unit join is pending (waiting for the smart license entitlement: ent1)

ciscoasa(cfg-cluster)# show cluster info auto-join
Unit join is pending (waiting for the smart license export control flag)
```

See the following output for the **show cluster info transport cp detail** command:

```
ciscoasa# show cluster info transport cp detail
Member ID to name mapping:
  0 - unit-1-1   2 - unit-4-1   3 - unit-2-1
Legend:
  U    - unreliable messages
  UE   - unreliable messages error
  SN   - sequence number
  ESN  - expecting sequence number
  R    - reliable messages
  RE   - reliable messages error
  RDC  - reliable message deliveries confirmed
  RA   - reliable ack packets received
  RFR  - reliable fast retransmits
  RTR  - reliable timer-based retransmits
  RDP  - reliable message dropped
  RDPR - reliable message drops reported
  RI   - reliable message with old sequence number
  RO   - reliable message with out of order sequence number
  ROW  - reliable message with out of window sequence number
  ROB  - out of order reliable messages buffered
  RAS  - reliable ack packets sent
This unit as a sender
--------------------------
        all       0         2         3
  U     123301    3867966   3230662   3850381
  UE    0         0         0         0
  SN    1656a4ce  acb26fe   5f839f76  7b680831
  R     733840    1042168   852285    867311
  RE    0         0         0         0
  RDC   699789    934969    740874    756490
  RA    385525    281198    204021    205384
  RFR   27626     56397     0         0
  RTR   34051     107199    111411    110821
  RDP   0         0         0         0
```

```
  RDPR  0         0         0         0
This unit as a receiver of broadcast messages
-------------------------------------------
        0         2         3
  U     111847    121862    120029
  R     7503      665700    749288
  ESN   5d75b4b3  6d81d23   365ddd50
  RI    630       34278     40291
  RO    0         582       850
  ROW   0         566       850
  ROB   0         16        0
  RAS   1571      123289    142256
This unit as a receiver of unicast messages
-------------------------------------------
        0         2         3
  U     1         3308122   4370233
  R     513846    879979    1009492
  ESN   4458903a  6d841a84  7b4e7fa7
  RI    66024     108924    102114
  RO    0         0         0
  ROW   0         0         0
  ROB   0         0         0
  RAS   130258    218924    228303
Gated Tx Buffered Message Statistics
------------------------------------
    current sequence number: 0
    total:                   0
    current:                 0
    high watermark:          0
    delivered:               0
    deliver failures:        0
    buffer full drops:       0
    message truncate drops:  0
    gate close ref count:    0
    num of supported clients:45
MRT Tx of broadcast messages
============================
Message high watermark: 3%
  Total messages buffered at high watermark: 5677
  [Per-client message usage at high watermark]
  ------------------------------------------------------------------
  Client name                          Total messages  Percentage
  Cluster Redirect Client                        4153         73%
  Route Cluster Client                            419          7%
  RRI Cluster Client                             1105         19%
Current MRT buffer usage: 0%
  Total messages buffered in real-time: 1
  [Per-client message usage in real-time]
  Legend:
        F - MRT messages sending when buffer is full
        L - MRT messages sending when cluster node leave
        R - MRT messages sending in Rx thread
  ------------------------------------------------------------------------------
  Client name                          Total messages  Percentage    F   L   R
  VPN Clustering HA Client                          1        100%    0   0   0
MRT Tx of unitcast messages(to member_id:0)
===========================================
Message high watermark: 31%
  Total messages buffered at high watermark: 4059
  [Per-client message usage at high watermark]
  ------------------------------------------------------------------
  Client name                          Total messages  Percentage
  Cluster Redirect Client                        3731         91%
  RRI Cluster Client                              328          8%
```

```
Current MRT buffer usage: 29%
  Total messages buffered in real-time: 3924
  [Per-client message usage in real-time]
  Legend:
        F - MRT messages sending when buffer is full
        L - MRT messages sending when cluster node leave
        R - MRT messages sending in Rx thread
  ------------------------------------------------------------------------------
  Client name                          Total messages  Percentage    F   L   R
  Cluster Redirect Client                        3607         91%    0   0   0
  RRI Cluster Client                              317          8%    0   0   0
MRT Tx of unitcast messages(to member_id:2)
==========================================
Message high watermark: 14%
  Total messages buffered at high watermark: 578
  [Per-client message usage at high watermark]
  ---------------------------------------------------------------
  Client name                          Total messages  Percentage
  VPN Clustering HA Client                        578        100%
Current MRT buffer usage: 0%
  Total messages buffered in real-time: 0
MRT Tx of unitcast messages(to member_id:3)
==========================================
Message high watermark: 12%
  Total messages buffered at high watermark: 573
  [Per-client message usage at high watermark]
  ---------------------------------------------------------------
  Client name                          Total messages  Percentage
  VPN Clustering HA Client                        572         99%
  Cluster VPN Unique ID Client                      1          0%

Current MRT buffer usage: 0%
  Total messages buffered in real-time: 0
```

The following is sample output from the **show cluster info load-monitor** command:

```
ciscoasa(cfg-cluster)# show cluster info load-monitor
ID  Unit Name
0  B
1  A_1
Information from all units with 50 second interval:
Unit     Connections    Buffer Drops   Memory Used   CPU Used
Average from last 1 interval:
   0           0            0            14              25
   1           0            0            16              20
Average from last 25 interval:
   0           0            0            12              28
   1           0            0            13              27
```

The following is sample output from the **show cluster info unit-join-acceleration incompatible-config** command:

ciscoasa# show cluster info unit-join-acceleration incompatible-config

INFO: Clustering is not compatible with following commands. User must manually remove them to activate the cluster unit join-acceleration feature.

```
zone sf200 passive
```

The following is sample output from the show cluster info instance-type command for an ASA Virtual cluster:

```
ciscoasa# show cluster info instance-type
      Unit            Module Type     CPU Cores            RAM (MB)
      A                   ASAv           2                  4096
      B                   ASAv           2                  4096
```

Related Commands

Command	Description
show cluster	Displays aggregated data for the entire cluster.
show cluster user-identity	Shows cluster user identity information and statistics.

show cluster user-identity

To view cluster-wide user identity information and statistics, use the **show cluster user-identity** command in privileged EXEC mode.

show cluster user-identity [**statistics** [**user** *name* | **user-group** *group_name*] | **user** [**active** [**domain** *name*] | **user** *name* | **user-group** *group_name*] [**list** [*detail*] | **all** [**list** [*detail*] | **inactive** { **domain** *name* | **user-group** *group_name*] [**list** [*detail*]] }

Syntax Description

active	Shows users with active IP-user mappings.
all	Shows all users in the user database.
domain *name*	Shows user info for a domain.
inactive	Shows users with inactive IP-user mappings.
list [*detail*]	Shows a list of users.
statistics	Shows cluster user identity statistics.
user	Shows the user database.
user *name*	Show information for a specific user.
user-group *group_name*	Shows information for each user of a specific group.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	—	• Yes

Command History

Release	Modification
9.0(1)	This command was added.

Usage Guidelines

See also the **show cluster info** and **show cluster** commands.

Examples

The following is sample output from the << **command** >> command:

Related Commands

Command	Description
show cluster	Displays aggregated data for the entire cluster.
show cluster info	Shows cluster information.

show cluster vpn-sessiondb distribution

To view how active and backup sessions are distributed across the cluster, execute this command in privileged EXEC mode.

show cluster vpn-sessiondb distribution

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	—	• Yes	—	• Yes

Command History

Release	Modification
9.9(1)	Command added for the Firepower 9300.
9.23(1)	Support added for the Secure Firewall 4200.

Usage Guidelines

This show command provides a quick view of the sessions, rather than having to execute **show vpn-sessiondb summary** on each member.

Each row contains the member id, member name, number of active sessions, and on which members the backup sessions reside.

Examples

For example, if the output of show cluster vpn-sessiondb distribution was:

Member 0 (unit-1-1): active: 209; backups at: 1(111), 2(98)

Member 1 (unit-1-3): active: 204; backups at: 0(108), 2(96)

Member 2 (unit-1-2): active: 0

One would read the information as:

- Member 0 has 209 active sessions, 111 sessions are backed up on member 1, 98 sessions are backed up on member 2
- Member 1 has 204 active sessions, 108 sessions are backed up on member 0, 96 sessions are backed up on member 2
- Member 2 has NO active sessions, therefore, no cluster members are backing up sessions for this node

Related Commands

Command	Description
cluster redistribute vpn-sessiondb	Redistribute the active VPN sessions on a distributed VPN cluster.

show compression

To view compression statistics on the ASA, use the **show compression** command from privileged EXEC mode.

show compression [**all | anyconnect-ssl | http-comp**]

Command Default

There is no default behavior for this command.

Syntax Description

all	Show all (anyconect-ssl, http-comp) compression statistics
anyconnect-ssl	Show AnyConnect SSL Compression statistics.
http-comp	Show HTTP-COMP Compression statistics

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Global configuration	• Yes	—	• Yes	—	—

Command History

Release	Modification
7.1(1)	This command was added.
9.0(1)	Support for multiple context mode was added.

Examples

The following types of statistics are displayed for show compression all:

```
Compression AnyConnect Client Sessions              0
Compressed Frames                                   0
Compressed Data In (bytes)                          0
Compressed Data Out (bytes)                         0
Expanded Frames                                     0
Compression Errors                                  0
Compression Resets                                  0
Compression Output Buf Too Small                    0
Compression Ratio                                   0
Decompressed Frames                                 0
Decompressed Data In                                0
Decompressed Data Out                               0
Decompression CRC Errors                            0
Decompression Errors                                0
Decompression Resets                                0
Decompression Ratio                                 0
Block Allocation Failures                           0
```

```
Compression Skip Percent                              0%
Time Spent Compressing (peak)                       0.0%
Time Spent Decompressing (peak)                     0.0%
Number of http bytes in                               0
Number of http gzipped bytes out                      0
```

Related Commands

Command	Description
compression	Enables compression for all SVC and WebVPN connections.

show configuration

To display the configuration that is saved in flash memory on the ASA, use the **show configuration** command in privileged EXEC mode.

show configuration

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
7.0(1)	This command was modified.

Usage Guidelines

The show configuration command displays the saved configuration in flash memory on the ASA. Unlike the **show running-config** command, the **show configuration** command does not use many CPU resources to run.

To display the active configuration in memory (including saved configuration changes) on the ASA, use the **show running-config** command.

Examples

The following is sample output from the **show configuration** command:

```
ciscoasa# show configuration
: enable password 8Ry2YjIyt7RRXU24 encrypted
names
dns-guard
!
interface Ethernet0/0
 nameif inside
 security-level 100
 ip address 192.168.2.5 255.255.255.0
!
interface Ethernet0/1
 nameif outside
 security-level 0
 ip address 10.132.12.6 255.255.255.0
!
interface Ethernet0/2
```

```
 nameif dmz
 security-level 50
 ip address 10.0.0.5 255.255.0.0
!
interface Ethernet0/3
 shutdown
 no nameif
 no security-level
 no ip address
!
interface Management0/0
 nameif management
 security-level 100
 ip address 192.168.1.1 255.255.255.0
 management-only
!
passwd 2KFQnbNIdI.2KYOU encrypted
boot system disk0:/newImage
ftp mode passive
access-list acl1 extended permit ip any any
access-list mgcpacl extended permit udp any any eq 2727
access-list mgcpacl extended permit udp any any eq 2427
access-list mgcpacl extended permit udp any any eq tftp
access-list mgcpacl extended permit udp any any eq 1719
access-list permitIp extended permit ip any any
pager lines 25
logging enable
logging console debugging
logging buffered debugging
logging asdm informational
mtu inside 1500
mtu outside 1500
mtu dmz 1500
mtu management 1500
icmp unreachable rate-limit 1 burst-size 1
icmp permit any inside
icmp permit any outside
icmp permit any dmz
asdm image disk0:/pdm
no asdm history enable
arp timeout 14400
global (outside) 1 10.132.12.50-10.132.12.52
global (outside) 1 interface
global (dmz) 1 interface
nat (inside) 1 0.0.0.0 0.0.0.0
access-group permitIp in interface inside
access-group permitIp in interface outside
access-group mgcpacl in interface dmz
!
router ospf 1
 network 10.0.0.0 255.255.0.0 area 192.168.2.0
 network 192.168.2.0 255.255.255.0 area 192.168.2.0
 log-adj-changes
 redistribute static subnets
 default-information originate
!
route outside 0.0.0.0 0.0.0.0 10.132.12.1 1
route outside 10.129.0.0 255.255.0.0 10.132.12.1 1
route outside 88.0.0.0 255.0.0.0 10.132.12.1 1
timeout xlate 3:00:00
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00 h225 1:00:00 mgcp 0:05:00 mgcp-pat 0:05:00
timeout sip 0:30:00 sip_media 0:02:00 sip-invite 0:03:00 sip-disconnect 0:02:00
timeout uauth 0:05:00 absolute
```

```
dynamic-access-policy-record DfltAccessPolicy
aaa authentication ssh console LOCAL
http server enable
http 10.132.12.0 255.255.255.0 outside
http 192.168.2.0 255.255.255.0 inside
http 192.168.1.0 255.255.255.0 management
no snmp-server location
no snmp-server contact
snmp-server enable traps snmp authentication linkup linkdown coldstart
telnet 192.168.2.0 255.255.255.0 inside
telnet 10.132.12.0 255.255.255.0 outside
telnet timeout 5
ssh 192.168.2.0 255.255.255.0 inside
ssh timeout 5
console timeout 0
dhcpd address 192.168.1.2-192.168.1.254 management
dhcpd enable management
!
threat-detection basic-threat
threat-detection statistics access-list
!
class-map inspection_default
 match default-inspection-traffic
!
!
policy-map type inspect dns preset_dns_map
 parameters
  message-length maximum 512
policy-map global_policy
 class inspection_default
  inspect dns preset_dns_map
  inspect ftp
  inspect h323 h225
  inspect h323 ras
  inspect rsh
  inspect rtsp
  inspect esmtp
  inspect sqlnet
  inspect skinny
  inspect sunrpc
  inspect xdmcp
  inspect sip
  inspect netbios
  inspect tftp
  inspect mgcp
policy-map type inspect mgcp mgcpapp
 parameters
  call-agent 150.0.0.210 101
  gateway 50.0.0.201 101
  gateway 100.0.0.201 101
  command-queue 150
!
service-policy global_policy global
webvpn
 memory-size percent 25
 enable inside
 internal-password enable
 onscreen-keyboard logon
username snoopy password /JcYsjvxHfBHc4ZK encrypted
prompt hostname context
Cryptochecksum:62bf8f5de9466cdb64fe758079594635:
end
```

Related Commands

Command	Description
configure	Configures the ASA from the terminal.

show configuration session

To display the current configuration sessions and the changes within the sessions, use the **show configuration session** command in privileged EXEC mode.

show configuration session [*session_name*]

Syntax Description

session_name	The name of an existing configuration session. If you omit this parameter, all existing sessions are shown.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Global configuration	• Yes	• Yes	• Yes	• Yes	—

Command History

Release	Modification
9.3(2)	This command was added.

Usage Guidelines

Use this command in conjunction with the **configure session** command, which creates isolated sessions for editing ACLs and their objects. This command shows the names of the sessions, and all of the configuration changes that have been made in the sessions.

If a session shows as committed, you can open the session and revert the changes if you decide they did not work as expected.

Examples

The following example shows all available sessions:

```
ciscoasa# show configuration session

config-session abc (un-committed)
 access-list abc permit ip any any
 access-list abc permit tcp any any

config-session abc2 (un-committed)
 object network test
 host 1.1.1.1
 object network test2
 host 2.2.2.2

ciscoasa#
```

Related Commands

Command	Description
clear configuration session	Deletes a configuration session and its contents.
clear session	Clears the contents of a configuration session or resets its access flag.
configure session	Creates or opens a session.

show conn

To display the connection state for the designated connection type, use the show **conn** command in privileged EXEC mode. This command supports IPv4 and IPv6 addresses.

show conn [**count** | [**all**] [**detail** [**data-rate-filter** { **lt | eq | gt** } *value*]] [**long**] [**state** *state_type*] [**protocol** { **tcp | udp | sctp** }] [**scansafe**] [**address** *src_ip* [*-src_ip*] [**netmask** *mask*]] [**port** *src_port* [*-src_port*]] [**address** *dest_ip* [*-dest_ip*] [**netmask** *mask*]] [**port** *dest_port* [*-dest_port*]] [**user-identity | user** [*domain_nickname* \ ] *user_name* **| user-group** [*domain_nickname* \ \ ] *user_group_name*] | **security-group**] [**zone** *zone_name* [**zone** *zone_name*] [**...**]] [**data-rate**]

Syntax Description

address	(Optional) Displays connections with the specified source or destination IP address.
all	(Optional) Displays connections that are to the device or from the device, in addition to through-traffic connections.
count	(Optional) Displays the number of active connections.
dest_ip	(Optional) Specifies the destination IP address (IPv4 or IPv6). To specify a range, separate the IP addresses with a dash (-). For example: `10.1.1.1-10.1.1.5`
dest_port	(Optional) Specifies the destination port number. To specify a range, separate the port numbers with a dash (-). For example: `1000-2000`
detail	(Optional) Displays connections in detail, including translation type and interface information.
data-rate-filter { **lt	eq**	**gt** } *value*	(Optional) Displays connections that are filtered based on a data-rate value (bytes per second). For example: `data-rate-filter gt 123`
long	(Optional) Displays connections in long format.
netmask *mask*	(Optional) Specifies a subnet mask for use with the given IP address.
port	(Optional) Displays connections with the specified source or destination port.
protocol { **tcp	udp**	**sctp** }	(Optional) Specifies the connection protocol.
scansafe	(Optional) Shows connections being forwarded to the Cloud Web Security server.
security-group	(Optional) Specifies that all connections displayed belong to the specified security group.

src_ip	(Optional) Specifies the source IP address (IPv4 or IPv6). To specify a range, separate the IP addresses with a dash (-). For example: `10.1.1.1-10.1.1.5`
src_port	(Optional) Specifies the source port number. To specify a range, separate the port numbers with a dash (-). For example: `1000-2000`
state *state_type*	(Optional) Specifies the connection state type. See <xref> for a list of the keywords available for connection state types.
user [*domain_nickname* \ ] *user_name*	(Optional) Specifies that all connections displayed belong to the specified user. When you do not include the *domain_nickname* argument, the ASA displays information for the user in the default domain.
user-group [*domain_nickname* \\ ] *user_group_name*	(Optional) Specifies that all connections displayed belong to the specified user group. When you do not include the *domain_nickname* argument, the ASA displays information for the user group in the default domain.
user-identity	(Optional) Specifies that the ASA display all connections for the Identity Firewall feature. When displaying the connections, the ASA displays the user name and IP address when it identifies a matching user. Similarly, the ASA displays the host name and an IP address when it identifies a matching host.
zone [*zone_name*]	(Optional) Displays connections for a zone. The **long** and **detail** keywords show the primary interface on which the connection was built and the current interface used to forward the traffic.
data-rate	(Optional) Displays whether data-rate tracking status is enabled or disabled.

Command Default

All through connections are shown by default. You need to use the **all** keyword to also view management connections to the device.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	• Yes	—

Command History

Release	Modification
7.0(8)/7.2(4)/8.0(4)	The syntax was simplified to use source and destination concepts instead of "local" and "foreign." In the new syntax, the source address is the first address entered and the destination is the second address. The old syntax used keywords like **foreign** and **fport** to determine the destination address and port.
7.2(5)/8.0(5)/8.1(2)/8.2(4)/8.3(2)	The **tcp_embryonic** state type was added. This type shows all TCP connections with the i flag (incomplete connections); i flag connections for UDP are not shown.
8.2(1)	The b flag was added for TCP state bypass.
8.4(2)	The **user-identity** , **user** , and **user-group** keywords were added to support the Identity Firewall.
9.0(1)	Support for clustering was added. We added the **scansafe** and **security-group** keywords.
9.3(2)	The **zone** keyword was added.
9.5(2)	The L flag was added for traffic subject to LISP flow-mobility.
9.5(2)	The Q flag for detailed output was added for Diameter connections. The **protocol sctp** keyword was added. The o flag for detailed output was added for off-loaded flows.
9.6(2)	The u flag for detailed output was added for STUN connections. The v flag was added for M3UA connections.
9.7(1)	The l flag was added to indicate the stub flow is local director Yl or local backup yl when using cluster director localization.
9.9(1)	VPN Stub at the end of the detail output, indicating that the connection is playing the role of a VPN encryption stub flow in addition to its cluster role.
9.13(1)	Dead Connection Detection (DCD) initiator/responder probe counts were added to the **show conn detail** output for DCD-enabled connections.
9.14(1)	Connection data-rate tracking status was added. The **data-rate-filter** keyword was added to the **show conn detail** command to filter the connections by user-specified data rate value.
9.16(1)	Multicast data connection entries were no longer displayed in the output. The entries were moved to the **show local-host** output.
9.23(1)	Added the Yo, yo, and zo flow types for conns that are asymmetric and offloaded in a cluster.

Usage Guidelines

The show conn command displays the number of active TCP and UDP connections, and provides information about connections of various types. Use the **show conn all** command to see the entire table of connections.

Note When the ASA creates a pinhole to allow secondary connections, this is shown as an incomplete conn by the **show conn** command. To clear this incomplete conn use the **clear conn** command.

The connection types that you can specify using the **show conn state** command are defined in Table 4-34. When specifying multiple connection types, use commas without spaces to separate the keywords.

Table 34: Connection State Types

Keyword	Connection Type Displayed
up	Connections in the up state.
conn_inbound	Inbound connections.
ctiqbe	CTIQBE connections
data_in	Inbound data connections.
data_out	Outbound data connections.
finin	FIN inbound connections.
finout	FIN outbound connections.
h225	H.225 connections
h323	H.323 connections
http_get	HTTP get connections.
mgcp	MGCP connections.
nojava	Connections that deny access to Java applets.
rpc	RPC connections.
service_module	Connections being scanned by an SSM.
sip	SIP connections.
skinny	SCCP connections.
smtp_data	SMTP mail data connections.
sqlnet_fixup_data	SQL*Net data inspection engine connections.
tcp_embryonic	TCP embryonic connections.
vpn_orphan	Orphaned VPN tunneled flows.

Usage Guidelines

When you use the detail option, the system displays information about the translation type and interface information using the connection flags defined in Table 4-34. Also, VPN stub may be shown at the end of

the output of this command indicating that the connection is playing the role of a VPN encryption stub flow in addition to its cluster role. A VPN stub is used to encrypt clear text packets in an asymmetric VPN traffic scenario or hub-n-spoke scenario.

Note In Firepower 3100 and Cisco Secure Firewall 1200 model devices, though the priority queue is disabled and no connection is established, the output of the **show conn details** displays a note that mentions an invalid connection due to an invalid internal-data/Rx-ring number.

Table 35: Connection Flags

Flag	Description
a	awaiting outside ACK to SYN
A	awaiting inside ACK to SYN
b	TCP state bypass
B	initial SYN from outside
C	Computer Telephony Interface Quick Buffer Encoding (CTIQBE) media connection
d	dump
D	DNS
E	outside back connection. This is a secondary data connection that must be initiated from the inside host. For example, using FTP, after the inside client issues the PASV command and the outside server accepts, the ASA preallocates an outside back connection with this flag set. If the inside client attempts to connect back to the server, then the ASA denies this connection attempt. Only the outside server can use the preallocated secondary connection.
f	inside FIN
F	outside FIN
g	Media Gateway Control Protocol (MGCP) connection
G	connection is part of a group[2]
h	H.225
H	H.323
i	incomplete TCP or UDP connection
I	inbound data
k	Skinny Client Control Protocol (SCCP) media connection
K	GTP t3-response

Flag	Description
l	local director/backup stub flow
L	traffic subject to LISP flow-mobility
m	SIP media connection
M	SMTP data
N	Inspected by Snort. If the system is configured to preserve connections if Snort goes down (this is enabled by default), the N flag includes a number. • 1—This connection will be preserved if Snort goes down. • 2—Snort did go down, and this connection was preserved. The connection will no longer be inspected by Snort.
o	Off-loaded flow.
O	outbound data
p	replicated (unused)
P	inside back connection. This is a secondary data connection that must be initiated from the inside host. For example, using FTP, after the inside client issues the PORT command and the outside server accepts, the ASA preallocates an inside back connection with this flag set. If the outside server attempts to connect back to the client, then the ASA denies this connection attempt. Only the inside client can use the preallocated secondary connection.
q	SQL*Net data
Q	Diameter connection
r	inside acknowledged FIN
R	outside acknowledged FIN for TCP connection
R	UDP RPC[3]
s	awaiting outside SYN
S	awaiting inside SYN
t	SIP transient connection[4]
T	SIP connection[5]
u	STUN connection
U	up
v	M3UA connection

Flag	Description
V	VPN orphan
w	For inter-chassis clustering on the Firepower 9300, identifies a flow on a backup owner on a separate chassis.
W	WAAS
X	Inspected by the service module, such as a CSC SSM.
y	For clustering, identifies a backup owner flow.
yo	For clustering, indicates that packets are (forwarded and) offloaded by director backup of the conn. For this case, owner is also director of the conn.
Y	For clustering, identifies a director flow.
Yo	For clustering, indicates that packets are (forwarded and) offloaded by director of the connection.
z	For clustering, identifies a forwarder flow.
zo	For clustering, indicates that packets are offloaded by a forwarder of the connection.
Z	Cloud Web Security

[2] The G flag indicates the connection is part of a group. It is set by the GRE and FTP Strict inspections to designate the control connection and all its associated secondary connections. If the control connection terminates, then all associated secondary connections are also terminated.

[3] Because each row of show conn command output represents one connection (TCP or UDP), there will be only one R flag per row.

[4] For UDP connections, the value t indicates that it will timeout after one minute.

[5] For UDP connections, the value T indicates that the connection will timeout according to the value specified using the timeout sip command.

Note For connections using a DNS server, the source port of the connection may be replaced by the *IP address of DNS server* in the **show conn** command output.

A single connection is created for multiple DNS sessions, as long as they are between the same two hosts, and the sessions have the same 5-tuple (source/destination IP address, source/destination port, and protocol). DNS identification is tracked by app_id, and the idle timer for each app_id runs independently.

Because the app_id expires independently, a legitimate DNS response can only pass through the ASA within a limited period of time and there is no resource build-up. However, when you enter the **show conn** command, you will see the idle timer of a DNS connection being reset by a new DNS session. This is due to the nature of the shared DNS connection and is by design.

Note When there is no TCP traffic for the period of inactivity defined by the **timeout conn** command (by default, 1:00:00), the connection is closed and the corresponding conn flag entries are no longer displayed.

If a LAN-to-LAN/Network-Extension Mode tunnel drops and does not come back, there might be a number of orphaned tunnel flows. These flows are not torn down as a result of the tunnel going down, but all the data attempting to flow through them is dropped. The **show conn** command output shows these orphaned flows with the **V** flag.

When the following TCP connection directionality flags are applied to connections between same-security interfaces (see the **same-security permit** command), the direction in the flag is not relevant because for same-security interfaces, there is no "inside" or "outside." Because the ASA has to use these flags for same-security connections, the ASA may choose one flag over another (for example, f vs. F) based on other connection characteristics, but you should ignore the directionality chosen.

- B—Initial SYN from outside
- a—Awaiting outside ACK to SYN
- A—Awaiting inside ACK to SYN
- f—Inside FIN
- F—Outside FIN
- s—Awaiting outside SYN
- S—Awaiting inside SYN

To display information for a specific connection, include the **security-group** keyword and specify a security group table value or security group name for both the source and destination of the connection. The ASA displays the connection matching the specific security group table values or security group names.

When you specify the **security-group** keyword without specifying a source and destination security group table value or a source and destination security group name, the ASA displays data for all SXP connections.

The ASA displays the connection data in the format *security_group_name* (*SGT_value*) or just as the *SGT_value* when the security group name is unknown.

Note Security group data is not available for stub connections because stub connection do not go through the slow path. Stub connections maintain only the information necessary to forward packets to the owner of the connection.

You can specify a single security group name to display all connections in a cluster; for example, the following example displays connections matching security-group mktg in all units of the cluster:

```
ciscoasa# show cluster conn security-group name mktg
```

Use the **data-rate** keyword to view the current state of the connection data rate tracking feature—enabled or disabled. Use the **data-rate filte** r keyword to filter the connections based on the data-rate value in bytes per second. Use the relational operators (lesser than, equal to, or greater than) to filter the connections data.

The output displays the active connections along with two data rate values—instantaneous one-second and maximum data rate, for both forward and reverse flows.

Examples

When specifying multiple connection types, use commas without spaces to separate the keywords. The following example displays information about RPC, H.323, and SIP connections in the Up state:

```
ciscoasa# show conn state up,rpc,h323,sip
```

The following is sample output from the **show conn count** command:

```
ciscoasa# show conn count
54 in use, 123 most used
```

The following is sample output from the**show conn** command. This example shows a TCP session connection from inside host 10.1.1.15 to the outside Telnet server at 10.10.49.10. Because there is no B flag, the connection is initiated from the inside. The “U”, “I”, and “O” flags denote that the connection is active and has received inbound and outbound data.

```
ciscoasa# show conn
54 in use, 123 most used
TCP out 10.10.49.10:23 in 10.1.1.15:1026 idle 0:00:22, bytes 1774, flags UIO
UDP out 10.10.49.10:31649 in 10.1.1.15:1028 idle 0:00:14, bytes 0, flags D-
TCP dmz 10.10.10.50:50026 inside 192.168.1.22:5060, idle 0:00:24, bytes 1940435, flags UTIOB
TCP dmz 10.10.10.50:49764 inside 192.168.1.21:5060, idle 0:00:42, bytes 2328346, flags UTIOB
TCP dmz 10.10.10.51:50196 inside 192.168.1.22:2000, idle 0:00:04, bytes 31464, flags UIB
TCP dmz 10.10.10.51:52738 inside 192.168.1.21:2000, idle 0:00:09, bytes 129156, flags UIOB
TCP dmz 10.10.10.50:49764 inside 192.168.1.21:0, idle 0:00:42, bytes 0, flags Ti
TCP outside 192.168.1.10(20.20.20.24):49736 inside 192.168.1.21:0, idle 0:01:32, bytes 0,
flags Ti
TCP dmz 10.10.10.50:50026 inside 192.168.1.22:0, idle 0:00:24, bytes 0, flags Ti
TCP outside 192.168.1.10(20.20.20.24):50663 inside 192.168.1.22:0, idle 0:01:34, bytes 0,
flags Ti
TCP dmz 10.10.10.50:50026 inside 192.168.1.22:0, idle 0:02:24, bytes 0, flags Ti
TCP outside 192.168.1.10(20.20.20.24):50663 inside 192.168.1.22:0, idle 0:03:34, bytes 0,
flags Ti
TCP dmz 10.10.10.50:50026 inside 192.168.1.22:0, idle 0:04:24, bytes 0, flags Ti
TCP outside 192.168.1.10(20.20.20.24):50663 inside 192.168.1.22:0, idle 0:05:34, bytes 0,
flags Ti
TCP dmz 10.10.10.50:50026 inside 192.168.1.22:0, idle 0:06:24, bytes 0, flags Ti
TCP outside 192.168.1.10(20.20.20.24):50663 inside 192.168.1.22:0, idle 0:07:34, bytes 0,
flags Ti
```

The following is sample output from the **show conn**command, whcih includes the “X” flag to indicate that the connection is being scanned by the SSM.

```
ciscoasa# show conn address 10.0.0.122 state service_module
TCP out 10.1.0.121:22 in 10.0.0.122:34446 idle 0:00:03, bytes 2733, flags UIOX
```

The following is sample output from the **show conn detail**command. This example shows a UDP connection from outside host 10.10.49.10 to inside host 10.1.1.15. The D flag denotes that this is a DNS connection. The number 1028 is the DNS ID over the connection.

`Uptime` represents the total time the SCTP association has been active on the device. The counter begins at the initial handshake (INIT/INIT-ACK) of the association. `Idle` indicates the duration since the last packet (including heartbeats) was processed for a specific multi-homing connection.

Examples

```
ciscoasa# show conn detail
54 in use, 123 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed,
       C - CTIQBE media, c - cluster centralized,
       D - DNS, d - dump, E - outside back connection, e - semi-distributed,
       F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, L - LISP triggered flow owner mobility
       l - local director/backup stub flow
       M - SMTP data, m - SIP media, n - GUP
       N - inspected by Snort
       O - outbound data, o - offloaded,
       P - inside back connection,
       Q - Diameter, q - SQL*Net data,
       R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up, u - STUN,
       V - VPN orphan, v - M3UA W - WAAS,
       w - secondary domain backup,
       X - inspected by service module,
       x - per session, Y - director stub flow, y - backup stub flow,
       Z - Scansafe redirection, z - forwarding stub flow
Cluster units to ID mappings:
  ID 0: asa1
  ID 255: The default cluster member ID which indicates no ownership or affiliation
          with an existing cluster member
TCP outside:10.10.49.10/23 inside:10.1.1.15/1026,
    flags UIO, idle 39s, uptime 1D19h, timeout 1h0m, bytes 1940435
UDP outside:10.10.49.10/31649 inside:10.1.1.15/1028,
    flags dD, idle 39s, uptime 1D19h, timeout 1h0m, bytes 1940435
TCP dmz:10.10.10.50/50026 inside:192.168.1.22/5060,
    flags UTIOB, idle 39s, uptime 1D19h, timeout 1h0m, bytes 1940435
TCP dmz:10.10.10.50/49764 inside:192.168.1.21/5060,
    flags UTIOB, idle 56s, uptime 1D19h, timeout 1h0m, bytes 2328346
TCP dmz:10.10.10.51/50196 inside:192.168.1.22/2000,
    flags UIB, idle 18s, uptime 1D19h, timeout 1h0m, bytes 31464
TCP dmz:10.10.10.51/52738 inside:192.168.1.21/2000,
    flags UIOB, idle 23s, uptime 1D19h, timeout 1h0m, bytes 129156
TCP outside:10.132.64.166/52510 inside:192.168.1.35/2000,
    flags UIOB, idle 3s, uptime 1D21h, timeout 1h0m, bytes 357405
TCP outside:10.132.64.81/5321 inside:192.168.1.22/5060,
    flags UTIOB, idle 1m48s, uptime 1D21h, timeout 1h0m, bytes 2083129
TCP outside:10.132.64.81/5320 inside:192.168.1.21/5060,
    flags UTIOB, idle 1m46s, uptime 1D21h, timeout 1h0m, bytes 2500529
TCP outside:10.132.64.81/5319 inside:192.168.1.22/2000,
    flags UIOB, idle 31s, uptime 1D21h, timeout 1h0m, bytes 32718
TCP outside:10.132.64.81/5315 inside:192.168.1.21/2000,
    flags UIOB, idle 14s, uptime 1D21h, timeout 1h0m, bytes 358694
TCP outside:10.132.64.80/52596 inside:192.168.1.22/2000,
    flags UIOB, idle 8s, uptime 1D21h, timeout 1h0m, bytes 32742
TCP outside:10.132.64.80/52834 inside:192.168.1.21/2000,
    flags UIOB, idle 6s, uptime 1D21h, timeout 1h0m, bytes 358582
TCP outside:10.132.64.167/50250 inside:192.168.1.35/2000,
    flags UIOB, idle 26s, uptime 1D21h, timeout 1h0m, bytes 375617
```

The following is sample output from the **show conn** command when an orphan flow exists, as indicated by the **V** flag:

```
ciscoasa# show conn
16 in use, 19 most used
```

```
TCP out 192.168.110.251:7393 in 192.168.150.252:21 idle 0:00:00, bytes 1048, flags UOVB
TCP out 192.168.110.251:21137 in 192.168.150.252:21 idle 0:00:00, bytes 1048, flags UIOB
```

To limit the report to those connections that have orphan flows, add the **vpn_orphan** option to the **show conn state** command, as in the following example:

```
ciscoasa# show conn state vpn_orphan
14 in use, 19 most used
TCP out 192.168.110.251:7393 in 192.168.150.252:5013, idle 0:00:00, bytes 2841019, flags
UOVB
```

For clustering, to troubleshoot the connection flow, first see connections on all units by entering the **cluster exec show conn** command on the master unit. Look for flows that have the following flags: director (Y), backup (y), and forwarder (z). The following example shows an SSH connection from 172.18.124.187:22 to 192.168.103.131:44727 on all three ASAs; ASA 1 has the z flag showing it is a forwarder for the connection, ASA3 has the Y flag showing it is the director for the connection, and ASA2 has no special flags showing it is the owner. In the outbound direction, the packets for this connection enter the inside interface on ASA2 and exit the outside interface. In the inbound direction, the packets for this connection enter the outside interface on ASA 1 and ASA3, are forwarded over the cluster control link to ASA2, and then exit the inside interface on ASA2.

```
ciscoasa/ASA1/master# cluster exec show conn
ASA1(LOCAL):**************************************************************
18 in use, 22 most used
Cluster stub connections: 0 in use, 5 most used
TCP outside  172.18.124.187:22 inside  192.168.103.131:44727, idle 0:00:00, bytes 37240828,
 flags z
ASA2:**************************************************************
12 in use, 13 most used
Cluster stub connections: 0 in use, 46 most used
TCP outside  172.18.124.187:22 inside  192.168.103.131:44727, idle 0:00:00, bytes 37240828,
 flags UIO
ASA3:**************************************************************
10 in use, 12 most used
Cluster stub connections: 2 in use, 29 most used
TCP outside  172.18.124.187:22 inside  192.168.103.131:44727, idle 0:00:03, bytes 0, flags
  Y
```

The output of **show conn detail** on ASA2 shows that the most recent forwarder was ASA1:

```
ciscoasa/ASA2/slave# show conn detail
12 in use, 13 most used
Cluster:
        fwd connections: 0 in use, 0 most used
        dir connections: 0 in use, 0 most used
        centralized connections: 1 in use, 61 most used
Flags: A - awaiting inside ACK to SYN, a - awaiting outside ACK to SYN,
       B - initial SYN from outside, b - TCP state-bypass or nailed,
       C - CTIQBE media, c - cluster centralized,
       D - DNS, d - dump, E - outside back connection, e - semi-distributed,
       F - outside FIN, f - inside FIN,
       G - group, g - MGCP, H - H.323, h - H.225.0, I - inbound data,
       i - incomplete, J - GTP, j - GTP data, K - GTP t3-response
       k - Skinny media, L - LISP triggered flow owner mobility
       l - local director/backup stub flow
       M - SMTP data, m - SIP media, n - GUP
       N - inspected by Snort
       O - outbound data, o - offloaded,
       P - inside back connection,
       Q - Diameter, q - SQL*Net data,
```

```
       R - outside acknowledged FIN,
       R - UDP SUNRPC, r - inside acknowledged FIN, S - awaiting inside SYN,
       s - awaiting outside SYN, T - SIP, t - SIP transient, U - up, u - STUN,
       V - VPN orphan, v - M3UA W - WAAS,
       w - secondary domain backup,
       X - inspected by service module,
       x - per session, Y - director stub flow, y - backup stub flow,
       Z - Scansafe redirection, z - forwarding stub flow
Cluster units to ID mappings:
  ID 0: asa1
  ID 1: asa2
  ID 255: The default cluster member ID which indicates no ownership or affiliation
          with an existing cluster member
TCP outside: 172.18.124.187/22 inside: 192.168.103.131/44727,
    flags UIO , idle 0s, uptime 25s, timeout 1h0m, bytes 1036044, cluster sent/rcvd bytes
0/1032983, cluster sent/rcvd total bytes 0/1080779, owners (1,255)
Traffic received at interface outside
        Locally received: 0 (0 byte/s)
        From most recent forwarder ASA1: 1032983 (41319 byte/s)
Traffic received at interface inside
        Locally received: 3061 (122 byte/s)
```

The following examples show how to display connections for the Identity Firewall feature:

```
ciscoasa# show conn user-identity
1219 in use, 1904 most used
UDP inside (www.yahoo.com))10.0.0.2:1587 outside (user1)192.0.0.2:30000, idle 0:00:00, bytes
 10, flags -
UDP inside (www.yahoo.com)10.0.0.2:1586 outside (user2)192.0.0.1:30000, idle 0:00:00, bytes
 10, flags -
UDP inside 10.0.0.34:1586 outside 192.0.0.25:30000, idle 0:00:00, bytes 10, flags -
...
ciscoasa# show conn user user1
2 in use
UDP inside (www.yahoo.com))10.0.0.2:1587 outside (user1)192.0.0.2:30000, idle 0:00:00, bytes
 10, flags -
```

See the following output for the **show conn long zone** command:

```
ciscoasa# show conn long zone zone-inside zone zone-outside
TCP outside-zone:outside1(outside2): 10.122.122.1:1080 inside-zone:inside1(inside2):
10.121.121.1:34254, idle 0:00:02, bytes 10, flags UO
```

When you use the **detail** keyword, you can see information about Dead Connection Detection (DCD) probing, which shows how often the connection was probed by the initiator and responder. For example, the connection details for a DCD-enabled connection would look like the following:

```
TCP dmz: 10.5.4.11/5555 inside: 10.5.4.10/40299,
    flags UO , idle 1s, uptime 32m10s, timeout 1m0s, bytes 11828, cluster sent/rcvd bytes
0/0, owners (0,255)
  Traffic received at interface dmz
        Locally received: 0 (0 byte/s)
  Traffic received at interface inside
        Locally received: 11828 (6 byte/s)
  Initiator: 10.5.4.10, Responder: 10.5.4.11
  DCD probes sent: Initiator 5, Responder 5
```

The following is sample output from the **show conn detail**command, which includes the Initiator and Responder field values for a UDP traffic flow.

```
ciscoasa# UDP BulkInlineSetNo1:interface2(interface2): 172.16.2.255/138
```

```
BulkInlineSetNo1:interface1(interface1): 172.16.2.1/138,
    flags - N(egress-optimization FR), idle 11s, uptime 11s, timeout 2m0s, bytes 437, flow
 id 12123, Snort id 12, rule id 268434436, Rx-RingNum 23, Internal-Data0/1
  Initiator: 172.16.2.1, Responder: 172.16.2.255
  Connection lookup keyid: 7021890
```

The following example shows how to view the status of connection data-rate tracking feature:

```
ciscoasa# show conn data-rate
Connection data rate tracking is currently enabled.
```

The following example shows how to filter the connection based on a specified data-rate:

```
ciscoasa# show conn detail data-rate-filter ?
eq  Enter this keyword to show conns with data-rate equal to specified value
gt   Enter this keyword to show conns with data-rate greater than specified
       value
lt    Enter this keyword to show conns with data-rate less than specified value
ciscoasa# show conn detail data-rate-filter gt ?
<0-4294967295>  Specify the data rate value in bytes per second
ciscoasa# show conn detail data-rate-filter gt 123 | grep max rate
    max rate:     3223223/399628 bytes/sec
    max rate:     3500123/403260 bytes/sec
```

Related Commands

Commands	Description
clear conn	Clears connections.
clear conn data-rate	Clears the current maximum data-rate stored.

show console-output

To display the currently captured console output, use the **show console-output** command in privileged EXEC mode.

show console-output

Syntax Description

This command has no arguments or keywords.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
7.0(1)	This command was added.

Examples

The following is sample output from the **show console-output** command, which displays the following message when there is no console output:

```
ciscoasa# show console-output
Sorry, there are no messages to display
```

Related Commands

Command	Description
clear configure console	Restores the default console connection settings.
clear configure timeout	Restores the default idle time durations in the configuration.
console timeout	Sets the idle timeout for a console connection to the ASA.
show running-config console timeout	Displays the idle timeout for a console connection to the ASA.

show context

To show context information including allocated interfaces and the configuration file URL, the number of contexts configured, or from the system execution space, a list of all contexts, use the **show context** command in privileged EXEC mode.

show context [*name* | **detail** | **count**]

Syntax Description

count	(Optional) Shows the number of contexts configured.
detail	(Optional) Shows additional detail about the context(s) including the running state and information for internal use.
name	(Optional) Sets the context name. If you do not specify a name, the ASA displays all contexts. Within a context, you can only enter the current context name.

Command Default

In the system execution space, the ASA displays all contexts if you do not specify a name.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	—	• Yes	• Yes

Command History

Release	Modification
7.0(1)	This command was added.
8.0(2)	Information about assigned IPS virtual sensors was added.

Usage Guidelines

See the "Examples" section for a description of the display output.

Examples

The following is sample output from the **show context** command. The following sample display shows three contexts:

```
ciscoasa# show context
Context Name      Interfaces                  URL
*admin            GigabitEthernet0/1.100      flash:/admin.cfg
                  GigabitEthernet0/1.101
contexta          GigabitEthernet0/1.200      flash:/contexta.cfg
                  GigabitEthernet0/1.201
contextb          GigabitEthernet0/1.300      flash:/contextb.cfg
                  GigabitEthernet0/1.301
Total active Security Contexts: 3
```

Table 4-36 shows each field description.

Table 36: show context Fields

Field	Description
Context Name	Lists all context names. The context name with the asterisk (*) is the admin context.
Interfaces	The interfaces assigned to the context.
URL	The URL from which the ASA loads the context configuration.

Examples

The following is sample output from the **show context detail** command in the system execution space:

```
ciscoasa# show context detail
Context "admin", has been created, but initial ACL rules not complete
  Config URL: flash:/admin.cfg
  Real Interfaces: Management0/0
  Mapped Interfaces: Management0/0
  Real IPS Sensors: ips1, ips2
  Mapped IPS Sensors: highsec, lowsec
  Flags: 0x00000013, ID: 1
Context "ctx", has been created, but initial ACL rules not complete
  Config URL: ctx.cfg
  Real Interfaces: GigabitEthernet0/0.10, GigabitEthernet0/1.20,
     GigabitEthernet0/2.30
  Mapped Interfaces: int1, int2, int3
  Real IPS Sensors: ips1, ips3
  Mapped IPS Sensors: highsec, lowsec
  Flags: 0x00000011, ID: 2
Context "system", is a system resource
  Config URL: startup-config
  Real Interfaces:
  Mapped Interfaces: Control0/0, GigabitEthernet0/0,
     GigabitEthernet0/0.10, GigabitEthernet0/1, GigabitEthernet0/1.10,
     GigabitEthernet0/1.20, GigabitEthernet0/2, GigabitEthernet0/2.30,
     GigabitEthernet0/3, Management0/0, Management0/0.1
  Flags: 0x00000019, ID: 257
Context "null", is a system resource
  Config URL: ... null ...
  Real Interfaces:
  Mapped Interfaces:
  Flags: 0x00000009, ID: 258
```

Table 4-37 shows each field description.

Table 37: Context States

Field	Description
Context	The context name. The null context information is for internal use only. The system context represents the system execution space.
State Message:	The context state. See the possible messages below.

Field	Description
Has been created, but initial ACL rules not complete	The ASA parsed the configuration but has not yet downloaded the default ACLs to establish the default security policy. The default security policy applies to all contexts initially, and includes disallowing traffic from lower security levels to higher security levels, enabling application inspection, and other parameters. This security policy ensures that no traffic can pass through the ASA after the configuration is parsed but before the configuration ACLs are compiled. You are unlikely to see this state because the configuration ACLs are compiled very quickly.
Has been created, but not initialized	You entered the **context** *name* command, but have not yet entered the **config-url** command.
Has been created, but the config hasn't been parsed	The default ACLs were downloaded, but the ASA has not parsed the configuration. This state might exist because the configuration download might have failed because of network connectivity issues, or you have not yet entered the **config-url** command. To reload the configuration, from within the context, enter **copy startup-config running-config** . From the system, reenter the **config-url** command. Alternatively, you can start configuring the blank running configuration.
Is a system resource	This state applies only to the system execution space and to the null context. The null context is used by the system, and the information is for internal use only.
Is a zombie	You deleted the context using the **no context** or **clear context** command, but the context information persists in memory until the ASA reuses the context ID for a new context, or you restart.
Is active	This context is currently running and can pass traffic according to the context configuration security policy.
Is ADMIN and active	This context is the admin context and is currently running.
Was a former ADMIN, but is now a zombie	You deleted the admin context using the **clear configure context** command, but the context information persists in memory until the ASA reuses the context ID for a new context, or you restart.
Real Interfaces	The interfaces assigned to the context. If you mapped the interface IDs in the **allocate-interface** command, this display shows the real name of the interface.
Mapped Interfaces	If you mapped the interface IDs in the **allocate-interface** command, this display shows the mapped names. If you did not map the interfaces, the display lists the real names again.
Real IPS Sensors	The IPS virtual sensors assigned to the context if you have an AIP SSM installed. If you mapped the sensor names in the **allocate-ips** command, this display shows the real name of the sensor.
Mapped IPS Sensors	If you mapped the sensor names in the **allocate-ips** command, this display shows the mapped names. If you did not map the sensor names, the display lists the real names again.

Field	Description
Flag	For internal use only.
ID	An internal ID for this context.

Examples

The following is sample output from the **show context count** command:

```
ciscoasa# show context count
Total active contexts: 2
```

Related Commands

Command	Description
admin-context	Sets the admin context.
allocate-interface	Assigns interfaces to a context.
changeto	Changes between contexts or the system execution space.
config-url	Specifies the location of the context configuration.
context	Creates a security context in the system configuration and enters context configuration mode.

show controller

To view controller-specific information of all interfaces present, use the **show controller** command in privileged EXEC mode.

show controller [**slot**] [*physical_interface*] [**pci** [**bridge** [*bridge-id* [*port-num*]]]] [**detail**]
show controller Internal-Data 0/1 qos [**statistics** | **rules**]

Syntax Description

bridge	(Optional) Displays PCI bridge-specific information for the ASA 5585-X.
bridge-id	(Optional) Displays each unique PCI bridge identifier for the ASA 5585-X.
detail	(Optional) Shows additional detail about the controller.
Internal-Data 0/1 qos	(Optional) Shows information about internal quality of service queues for control traffic in the NIC. This feature is available on Secure Firewall 1200 devices. This information is primarily for internal TAC support personel. You cannot configure the MCAM QoS rules, the system creates and manages these rules automatically. Add one of the following keywords: • **statistics**—Display the hit counters for MCAM rules. • **rules**—Display the raw hexdump of the MCAM rules.
pci	(Optional) Displays a summary of PCI devices along with their first 256 bytes of PCI configuration space for the ASA 5585-X.
physical_interface	(Optional) Identifies the interface ID.
port-num	(Optional) Displays the unique port number within each PCI bridge for the ASA 5585-X adaptive ASA.
slot	(Optional) Displays PCI-e bus and slot information for the ASA 5580 only.

Command Default

If you do not identify an interface, this command shows information for all interfaces.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
7.2(1)	This command was added.

Release	Modification
8.0(2)	This command now applies to all platforms, and not just the ASA 5505. The **detail** keyword was added.
8.1(1)	The **slot** keyword was added for the ASA 5580.
8.2(5)	The **pci** , **bridge** , *bridge-id,* and *port-num* options were added for the ASA 5585-X with an IPS SSP installed. In addition, support for sending pause frames to enable flow control on 1 GigabitEthernet interfaces has been added for all ASA models.
8.6(1)	Support was added for the **detail** keyword for the ASA 5512-X through ASA 5555-X Internal-Control0/0 interface, used for control traffic between the ASA and the software module, and for the Internal-Data0/1 interface used for data traffic to the ASA and the software module.
9.22(1)	The **Internal-Data 0/1 qos** keywords were added.

Usage Guidelines

This command helps Cisco TAC gather useful debug information about the controller when investigating internal and customer found defects. The actual output depends on the model and Ethernet controller. The command also displays information about all the PCI bridges of interest in the ASA 5585-X with an IPS SSP installed. For the ASA Services Module, the **show controller** command output does not show any PCIe slot information.

Examples

The following is sample output from the **show controller** command:

```
ciscoasa# show controller
Ethernet0/0:
   Marvell 88E6095 revision 2, switch port 7
     PHY Register:
         Control:       0x3000  Status:        0x786d
         Identifier1:   0x0141  Identifier2:   0x0c85
         Auto Neg:      0x01e1  LP Ability:    0x40a1
         Auto Neg Ex:   0x0005  PHY Spec Ctrl: 0x0130
         PHY Status:    0x4c00  PHY Intr En:   0x0400
         Int Port Sum:  0x0000  Rcv Err Cnt:   0x0000
         Led select:    0x1a34
         Reg 29:        0x0003  Reg 30:        0x0000
     Port Registers:
         Status:        0x0907  PCS Ctrl:      0x0003
         Identifier:    0x0952  Port Ctrl:     0x0074
         Port Ctrl-1:   0x0000  Vlan Map:      0x077f
         VID and PRI:   0x0001  Port Ctrl-2:   0x0cc8
         Rate Ctrl:     0x0000  Rate Ctrl-2:   0x3000
         Port Asc Vt:   0x0080
         In Discard Lo: 0x0000  In Discard Hi: 0x0000
         In Filtered:   0x0000  Out Filtered:  0x0000
     Global Registers:
         Control:       0x0482
--------------------------------------------------------------------
Number of VLANs: 1
--------------------------------------------------------------------
Vlan[db]\Port|  0 |  1 |  2 |  3 |  4 |  5 |  6 |  7 |  8 |  9 | 10 |
--------------------------------------------------------------------
  <0001[01]> | EUT| EUT| EUT| EUT| EUT| EUT| EUT| EUT| EUM| NM | NM |
--------------------------------------------------------------------
....
Ethernet0/6:
```

```
   Marvell 88E6095 revision 2, switch port 1
     PHY Register:
         Control:       0x3000  Status:        0x7849
         Identifier1:   0x0141  Identifier2:   0x0c85
         Auto Neg:      0x01e1  LP Ability:    0x0000
         Auto Neg Ex:   0x0004  PHY Spec Ctrl: 0x8130
         PHY Status:    0x0040  PHY Intr En:   0x8400
         Int Port Sum:  0x0000  Rcv Err Cnt:   0x0000
         Led select:    0x1a34
         Reg 29:        0x0003  Reg 30:        0x0000
     Port Registers:
         Status:        0x0007  PCS Ctrl:      0x0003
         Identifier:    0x0952  Port Ctrl:     0x0077
         Port Ctrl-1:   0x0000  Vlan Map:      0x07fd
         VID and PRI:   0x0001  Port Ctrl-2:   0x0cc8
         Rate Ctrl:     0x0000  Rate Ctrl-2:   0x3000
         Port Asc Vt:   0x0002
         In Discard Lo: 0x0000  In Discard Hi: 0x0000
         In Filtered:   0x0000  Out Filtered:  0x0000
   ----Inline power related counters and registers----
   Power on fault: 0  Power off fault: 0
   Detect enable fault: 0  Detect disable fault: 0
   Faults: 0
   Driver counters:
   I2C Read Fail: 0   I2C Write Fail: 0
   Resets: 1  Initialized: 1
   PHY reset error: 0
   LTC4259 registers:
   INTRPT STATUS = 0x88  INTRPT MASK   = 0x00  POWER EVENT   = 0x00
   DETECT EVENT  = 0x03  FAULT EVENT   = 0x00  TSTART EVENT  = 0x00
   SUPPLY EVENT  = 0x02  PORT1 STATUS  = 0x06  PORT2 STATUS  = 0x06
   PORT3 STATUS  = 0x00  PORT4 STATUS  = 0x00  POWER STATUS  = 0x00
   OPERATE MODE  = 0x0f  DISC. ENABLE  = 0x30  DT/CLASS ENBL = 0x33
   TIMING CONFIG = 0x00  MISC. CONFIG  = 0x00
...
Internal-Data0/0:
   Y88ACS06 Register settings:
    rap                          0xe0004000 = 0x00000000
    ctrl_status                  0xe0004004 = 0x5501064a
    irq_src                      0xe0004008 = 0x00000000
    irq_msk                      0xe000400c = 0x00000000
    irq_hw_err_src               0xe0004010 = 0x00000000
    irq_hw_err_msk               0xe0004014 = 0x00001000
    bmu_cs_rxq                   0xe0004060 = 0x002aaa80
    bmu_cs_stxq                  0xe0004068 = 0x01155540
    bmu_cs_atxq                  0xe000406c = 0x012aaa80
   Bank 2: MAC address registers:

....
```

The following is sample output from the **show controller detail** command:

```
ciscoasa# show controller gigabitethernet0/0 detail
GigabitEthernet0/0:
   Intel i82546GB revision 03
     Main Registers:
         Device Control:             0xf8260000 = 0x003c0249
         Device Status:              0xf8260008 = 0x00003347
         Extended Control:           0xf8260018 = 0x000000c0
         RX Config:                  0xf8260180 = 0x0c000000
         TX Config:                  0xf8260178 = 0x000001a0
         RX Control:                 0xf8260100 = 0x04408002
         TX Control:                 0xf8260400 = 0x000400fa
         TX Inter Packet Gap:        0xf8260410 = 0x00602008
```

```
    RX Filter Cntlr:            0xf8260150 = 0x00000000
    RX Chksum:                  0xf8265000 = 0x00000300
RX Descriptor Registers:
    RX Descriptor 0 Cntlr:      0xf8262828 = 0x00010000
    RX Descriptor 0 AddrLo:     0xf8262800 = 0x01985000
    RX Desccriptor 0 AddrHi:    0xf8262804 = 0x00000000
    RX Descriptor 0 Length:     0xf8262808 = 0x00001000
    RX Descriptor 0 Head:       0xf8262810 = 0x00000000
    RX Descriptor 0 Tail:       0xf8262818 = 0x000000ff
    RX Descriptor 1 Cntlr:      0xf8262828 = 0x00010000
    RX Descriptor 1 AddrLo:     0xf8260138 = 0x00000000
    RX Descriptor 1 AddrHi:     0xf826013c = 0x00000000
    RX Descriptor 1 Length:     0xf8260140 = 0x00000000
    RX Descriptor 1 Head:       0xf8260148 = 0x00000000
    RX Descriptor 1 Tail:       0xf8260150 = 0x00000000
TX Descriptor Registers:
    TX Descriptor 0 Cntlr:      0xf8263828 = 0x00000000
    TX Descriptor 0 AddrLo:     0xf8263800 = 0x01987000
    TX Descriptor 0 AddrHi:     0xf8263804 = 0x00000000
    TX Descriptor 0 Length:     0xf8263808 = 0x00001000
    TX Descriptor 0 Head:       0xf8263810 = 0x00000000
    TX Descriptor 0 Tail:       0xf8263818 = 0x00000000
RX Address Array:
    Ethernet Address 0:         0012.d948.ef58
    Ethernet Address 1:         Not Valid!
    Ethernet Address 2:         Not Valid!
    Ethernet Address 3:         Not Valid!
    Ethernet Address 4:         Not Valid!
    Ethernet Address 5:         Not Valid!
    Ethernet Address 6:         Not Valid!
    Ethernet Address 7:         Not Valid!
    Ethernet Address 8:         Not Valid!
    Ethernet Address 9:         Not Valid!
    Ethernet Address a:         Not Valid!
    Ethernet Address b:         Not Valid!
    Ethernet Address c:         Not Valid!
    Ethernet Address d:         Not Valid!
    Ethernet Address e:         Not Valid!
    Ethernet Address f:         Not Valid!
PHY Registers:
    Phy Control:                0x1140
    Phy Status:                 0x7969
    Phy ID 1:                   0x0141
    Phy ID 2:                   0x0c25
    Phy Autoneg Advertise:      0x01e1
    Phy Link Partner Ability:   0x41e1
    Phy Autoneg Expansion:      0x0007
    Phy Next Page TX:           0x2801
    Phy Link Partnr Next Page:  0x0000
    Phy 1000T Control:          0x0200
    Phy 1000T Status:           0x4000
    Phy Extended Status:        0x3000
Detailed Output - RX Descriptor Ring:
rx_bd[000]: baddr      = 0x019823A2, length = 0x0000, status  = 0x00
            pkt chksum = 0x0000,     errors = 0x00,   special = 0x0000
rx_bd[001]: baddr      = 0x01981A62, length = 0x0000, status  = 0x00
            pkt chksum = 0x0000,     errors = 0x00,   special = 0x0000
```

.....

The following is sample output from the **show controller detail** command for the Internal interfaces on the ASA 5512-X through ASA 5555-X:

```
ciscoasa# show controller detail
Internal-Control0/0:
   ASA IPS/VM Back Plane TunTap Interface , port id 9
     Major Configuration Parameters
         Device Name           : en_vtun
         Linux Tun/Tap Device  : /dev/net/tun/tap1
         Num of Transmit Rings : 1
         Num of Receive Rings  : 1
         Ring Size             : 128
         Max Frame Length      : 1550
         Out of Buffer         : 0
         Reset                 : 0
         Drop                  : 0
     Transmit Ring [0]:
         tx_pkts_in_queue      : 0
         tx_pkts               : 176
         tx_bytes              : 9664
     Receive Ring [0]:
         rx_pkts_in_queue      : 0
         rx_pkts               : 0
         rx_bytes              : 0
         rx_drops              : 0
Internal-Data0/1:
   ASA IPS/VM Management Channel TunTap Interface , port id 9
     Major Configuration Parameters
         Device Name           : en_vtun
         Linux Tun/Tap Device  : /dev/net/tun/tap2
         Num of Transmit Rings : 1
         Num of Receive Rings  : 1
         Ring Size             : 128
         Max Frame Length      : 1550
         Out of Buffer         : 0
         Reset                 : 0
         Drop                  : 0
     Transmit Ring [0]:
         tx_pkts_in_queue      : 0
         tx_pkts               : 176
         tx_bytes              : 9664
     Receive Ring [0]:
         rx_pkts_in_queue      : 0
         rx_pkts               : 0
         rx_bytes              : 0
         rx_drops              : 0
```

The following is sample output from the **show controller slot** command:

```
Slot  Card Description                              PCI-e Bandwidth Cap.
----  ----------------                             ----------------------
3.    ASA 5580 2 port 10GE SR Fiber Interface Card   Bus: x4, Card: x8
4.    ASA 5580 4 port GE Copper Interface Card       Bus: x4, Card: x4
5.    ASA 5580 2 port 10GE SR Fiber Interface Card   Bus: x8, Card: x8
6.    ASA 5580 4 port GE Fiber Interface Card        Bus: x4, Card: x4
7.    empty                                          Bus: x8
8.    empty                                          Bus: x8
```

The following is sample output from the **show controller pci** command:

```
ciscoasa# show controller
           pci
PCI Evaluation Log:
  ------------------------------------------------------------------------------
    Empty
```

```
PCI Bus:Device.Function (hex): 00:00.0 Vendor ID: 0x8086 Device ID: 0x3406
-------------------------------------------------------------------------------
  PCI Configuration Space (hex):
  0x00: 86 80 06 34 00 00 10 00 22 00 00 06 10 00 00 00
  0x10: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  0x20: 00 00 00 00 00 00 00 00 00 00 00 00 86 80 00 00
  0x30: 00 00 00 00 60 00 00 00 00 00 00 00 05 01 00 00
  0x40: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  0x50: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  0x60: 05 90 02 01 00 00 00 00 00 00 00 00 00 00 00 00
  0x70: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  0x80: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  0x90: 10 e0 42 00 20 80 00 00 00 00 00 00 41 3c 3b 00
  0xa0: 00 00 41 30 00 00 00 00 c0 07 00 01 00 00 00 00
  0xb0: 00 00 00 00 3e 00 00 00 09 00 00 00 00 00 00 00
  0xc0: 01 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  0xd0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  0xe0: 01 00 03 c8 08 00 00 00 00 00 00 00 00 00 00 00
  0xf0: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  Link Capabilities: x4, Gen1
  Link Status: x4, Gen1
```

The following example is from ASA Virtual. In this case, the rx_dropped_packets indicates that packets are being dropped at the VM level before entering the ASA Virtual, possibly due to lack of bandwidth. One possible cause is that there is a blast/burst of traffic destined for the VM beyond what the VM can handle.

```
ciscoasa# show controller TenGigabitEthernet 0/2

TenGigabitEthernet0/2:
   DPDK Statistics
                rx_good_packets : 13186640462
                tx_good_packets : 3225386
                  rx_good_bytes : 12526548356100
                  tx_good_bytes : 383943970
                      rx_errors : 0
                      tx_errors : 0
       rx_mbuf_allocation_errors : 0
                   rx_q0packets : 0
                     rx_q0bytes : 0
                    rx_q0errors : 0
                   tx_q0packets : 0
                     tx_q0bytes : 0
                       rx_bytes : 12526548273860
             rx_unicast_packets : 13186630349
           rx_multicast_packets : 10025
           rx_broadcast_packets : 0
             rx_dropped_packets : 15357499
    rx_unknown_protocol_packets : 0
                       tx_bytes : 383943970
             tx_unicast_packets : 3224181
           tx_multicast_packets : 1205
           tx_broadcast_packets : 0
             tx_dropped_packets : 0
               tx_error_packets : 0
```

Related Commands

Command	Description
show interface	Shows the interface statistics.
show tech-support	Shows information so Cisco TAC can diagnose problems.

show coredump filesystem

To show the contents of the coredump filesystem, enter the show coredump filesystem command.

show coredump filesystem

Syntax Description

This command has no arguments or keywords.

Command Default

By default, coredumps are not enabled.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Global configuration	• Yes	• Yes	• Yes	• Yes	—

Command History

Release	Modification
8.2(1)	This command was added.

Usage Guidelines

This command shows the contents of the coredump filesystem.

Examples

To show the contents of any recent coredumps generated, enter the show coredump filesystem command.

```
ciscoasa(config)# show coredump filesystem
Coredump Filesystem Size is 100 MB
Filesystem type is FAT for disk0
Filesystem 1k-blocks Used Available Use% Mounted on
/dev/loop0 102182 75240 26942 74% /mnt/disk0/coredumpfsys
Directory of disk0:/coredumpfsys/
246 -rwx 20205386 19:14:53 Nov 26 2008 core_lina.2008Nov26_191244.203.11.gz
247 -rwx 36707919 19:17:27 Nov 26 2008 core_lina.2008Nov26_191456.203.6.gz
```

Related Commands

Command	Description
coredump enable	Enables the coredump feature.
clear configure coredump	Removes any coredumps currently stored on the coredump filesystem and clears the coredump log. Does not touch the coredump filesystem itself and does not change or affect the coredump configuration.

Command	Description
clear coredump	Removes any coredumps currently stored on the coredump filesystem and clears the coredump log. Does not touch the coredump filesystem itself and does not change/effect the coredump configuration.
show coredump log	Shows the coredump log.

show coredump log

To show the contents of the coredump log, newest first, enter the **show coredump log** command. To show the contents of the coredump log, oldest first, enter the **show coredump log reverse** command.

show coredump log
show coredump log [**reverse**]

Syntax Description

reverse	Shows the oldest coredump log.

Command Default

By default, coredumps are not enabled.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Global configuration	• Yes	• Yes	• Yes	• Yes	—

Command History

Release	Modification
8.2(1)	This command was added.

Usage Guidelines

This command displays the contents of the coredump log. The logs should reflect what is currently on the disk.

Examples

The following example shows the output from these commands:

```
ciscoasa(config)# show coredump log
[ 1 ] Wed Feb 18 22:12:09 2009: Coredump completed for module 'lina', coredump file
'core_lina.2009Feb18_221032.203.6.gz', size 971722752 bytes, compressed size 21293688
[ 2 ] Wed Feb 18 22:11:01 2009: Filesystem full on 'disk0', removing module coredump record
 'core_lina.2009Feb18_213558.203.11.gz'
[ 3 ] Wed Feb 18 22:10:32 2009: Coredump started for module 'lina', generating coredump
file 'core_lina.2009Feb18_221032.203.6.gz' on 'disk0'
[ 4 ] Wed Feb 18 21:37:35 2009: Coredump completed for module 'lina', coredump file
'core_lina.2009Feb18_213558.203.11.gz', size 971722752 bytes, compressed size 21286383
[ 5 ] Wed Feb 18 21:35:58 2009: Coredump started for module 'lina', generating coredump
file 'core_lina.2009Feb18_213558.203.11.gz' on 'disk0'
```



Note The older coredump file is deleted to make room for the new coredump. This is done automatically by the ASA in the event the coredump filesystem fills and room is needed for the current coredump. This is why it is imperative to archive coredumps as soon as possible, to insure they don't get overwritten in the event of a crash.

ciscoasa(config)# show coredump log reverse

```
[ 1 ] Wed Feb 18 21:35:58 2009: Coredump started for module 'lina', generating coredump
file 'core_lina.2009Feb18_213558.203.11.gz' on 'disk0''
[ 2 ] Wed Feb 18 21:37:35 2009: Coredump completed for module 'lina', coredump file
'core_lina.2009Feb18_213558.203.11.gz', size 971722752 bytes, compressed size 21286383
[ 3 ] Wed Feb 18 22:10:32 2009: Coredump started for module 'lina', generating coredump
file 'core_lina.2009Feb18_221032.203.6.gz' on 'disk0'
[ 4 ] Wed Feb 18 22:11:01 2009: Filesystem full on 'disk0', removing module coredump record
 'core_lina.2009Feb18_213558.203.11.gz'
[ 5 ] Wed Feb 18 22:12:09 2009: Coredump completed for module 'lina', coredump file
'core_lina.2009Feb18_221032.203.6.gz', size 971722752 bytes, compressed size 21293688
```

Related Commands

Command	Description
coredump enable	Enables the coredump feature.
clear configure coredump	Removes any coredumps currently stored on the coredump filesystem and clears the coredump log. Does not touch the coredump filesystem itself and does not change/effect the coredump configuration.
clear coredump	Removes any coredumps currently stored on the coredump filesystem and clears the coredump log. Does not touch the coredump filesystem itself and does not change or affect the coredump configuration.
show coredump filesystem	Shows the contents of the coredump filesystem.

show counters

To display the protocol stack counters, use the **show counters** command in privileged EXEC mode.

show counters [**all** | **context** *context-name* | **summary** | `top` *N*][**detail**][**protocol** *protocol_name* [**:** *counter_name*]][**threshold** *N*]

Syntax Description

all	Displays the filter details.
context *context-name*	Specifies the context name.
:*counter_name*	Specifies a counter by name.
detail	Displays additional counters information.
protocol *protocol_name*	Displays the counters for the specified protocol.
summary	Displays a counter summary.
threshold *N*	Displays only those counters at or above the specified threshold. The range is 1 through 4294967295.
top *N*	Displays the counters at or above the specified threshold. The range is 1 through 4294967295.

Command Default

show counters summary detail threshold 1

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
7.2(1)	This command was added.
9.2(1)	Counters for the event manager were added.
9.13(1)	A new counter was added for the Firepower 1000 and 2100 in Appliance mode: HTTPERR: the number of HTTP request message timeouts to FXOS.

Release	Modification
9.12(1)	Five new counters were added for ACL search levels: • OBJGRP_SEARCH_THRESHOLD (Exceeding threshold 10000 search count) • OBJGRP_SEARCH_THRESHOLD_LEVEL4 (Between 7500 to 10000 searches) • OBJGRP_SEARCH_THRESHOLD_LEVEL3 (Between 5000 to 7500 searches) • OBJGRP_SEARCH_THRESHOLD_LEVEL2 (Between 2500 to 5000 searches) • OBJGRP_SEARCH_THRESHOLD_LEVEL1 (Between 1 to 2500 searches)
9.20(3)	Three new counters were added for rate-limited preauthenticated SSL connections: • SOCK_PRE_AUTH_COUNT_EXCEEDED: An increment by one indicates that the number of simultaneous preauthenticated SSL connections has exceeded the VPN limit. New connection attempts are blocked until the SOCK_PRE_AUTH_COUNT counter is 0. • SOCK_COUNT_RATE_LIMIT: Indicates the connections that the ASA has reset after exceeding therate limit. • SOCK_PRE_AUTH_COUNT: Number of concurrent preauthentication connections at any given time.
9.24(1)	New counters were introduced for PKI operation: • All lookaside asymmetric cipher counters are denoted with 'LAS_' prefix to the counter. • Enqueue and dequeue counters are identified by a suffix '_ENQ' and '_DEQ' respectively.

Examples

The following example shows how to display all counters:

```
ciscoasa#

show counters all
Protocol     Counter           Value   Context
IOS_IPC      IN_PKTS               2   single_vf
IOS_IPC      OUT_PKTS              2   single_vf
ciscoasa# show counters
Protocol     Counter           Value   Context
NPCP         IN_PKTS            7195   Summary
NPCP         OUT_PKTS           7603   Summary
IOS_IPC      IN_PKTS             869   Summary
IOS_IPC      OUT_PKTS            865   Summary
IP           IN_PKTS             380   Summary
IP           OUT_PKTS            411   Summary
IP           TO_ARP              105   Summary
IP           TO_UDP                9   Summary
UDP          IN_PKTS               9   Summary
UDP          DROP_NO_APP           9   Summary
FIXUP        IN_PKTS             202   Summary
UAUTH        IPV6_UNSUPPORTED     27   Summary
IDFW         HIT_USER_LIMIT        2   Summary
```

The following example shows how to display a summary of counters:

```
ciscoasa#

show counters summary
Protocol     Counter           Value   Context
IOS_IPC      IN_PKTS               2   Summary
IOS_IPC      OUT_PKTS              2   Summary
```

The following example shows how to display counters for a context:

```
ciscoasa# show counters context single_vf
Protocol     Counter           Value   Context
IOS_IPC      IN_PKTS               4   single_vf
IOS_IPC      OUT_PKTS              4   single_vf
```

The following example shows how to display counters for the event manager:

```
ciscoasa# show counters protocol eem
Protocol     Counter           Value   Context
EEM          SYSLOG            22      Summary
EEM          COMMANDS          6       Summary
EEM          FILES             3       Summary
```

The following example shows how to display counters for ACL search levels:

```
ciscoasa# show counters
Protocol     Counter                                                                 Value
   Context
ACL          OBJGRP_SEARCH_THRESHOLD                          1582   Summary
ACL          OBJGRP_SEARCH_THRESHOLD_LEVEL4                   534   Summary
ACL          OBJGRP_SEARCH_THRESHOLD_LEVEL3                   524   Summary
ACL          OBJGRP_SEARCH_THRESHOLD_LEVEL2                   307   Summary
ACL          OBJGRP_SEARCH_THRESHOLD_LEVEL1                   216   Summary
```

The following example shows how to display counters for rate-limited preauthenticated SSL connections:

```
ciscoasa# show counters
MIDPATH      SOCK_PRE_AUTH_COUNT_EXCEEDED                       3   Summary
MIDPATH      SOCK_COUNT_RATE_LIMIT                           4260   Summary
MIDPATH      SOCK_PRE_AUTH_COUNT                              335   Summary
```

Each PKI command is made up of an enqueue and dequeue operation. The same number of enqueue and dequeue counts are expected for a given PKI command type. In the output, all asymmetric cipher counters have 'LAS_' prefix to denote lookaside; the enqueue counters have '_ENQ' suffix and dequeue counters have a '_DEQ' suffix. The ENQs and DEQs should have matching numbers.

```
csf6170# show counter | grep LAS
CRYPTO               LAS_RSA_PRIV_ENC_ENQ                        1   Summary
CRYPTO               LAS_RSA_PRIV_DEC_ENQ                        1   Summary
CRYPTO               LAS_RSA_PUB_ENC_ENQ                         1   Summary
CRYPTO               LAS_RSA_PUB_DEC_ENQ                         1   Summary
CRYPTO               LAS_DH_KEY_PAIR_ENQ                         1   Summary
CRYPTO               LAS_RSA_PRIV_ENC_DEQ                        1   Summary
CRYPTO               LAS_RSA_PRIV_DEC_DEQ                        1   Summary
CRYPTO               LAS_RSA_PUB_ENC_DEQ                         1   Summary
CRYPTO               LAS_RSA_PUB_DEC_DEQ                         1   Summary
CRYPTO               LAS_DH_KEY_PAIR_DEQ                         1   Summary
CRYPTO               LAS_OUTBOUND_BUF_ALLOC                      5   Summary
CRYPTO               LAS_OUTBOUND_BUF_FORWARD                    5   Summary
CRYPTO               LAS_INBOUND_BUF_RCV                         5   Summary
```

```
CRYPTO                LAS_INBOUND_BUF_FREE                          5    Summary
```

Related Commands

Command	Description
clear counters	Clears the protocol stack counters.

show cpu

To display the CPU utilization information, use the show cpu command in privileged EXEC mode.

[**cluster exec**] **show cpu** [**usage** *core-id* **| profile | dump | detailed**]

From the system configuration in multiple context mode:

[**cluster exec**] **show cpu** [**usage**] [**context** { **all** | *context_name* }]

Syntax Description

all	Specifies that the display show all contexts.
cluster exec	(Optional) In a clustering environment, enables you to issue the **show cpu** command in one unit and run the command in all the other units at the same time.
context	Specifies that the display show a context.
context_name	Specifies the name of the context to display.
core-id	Specifies the number of the processor core.
detailed	(Optional) Displays the CPU usage internal details.
dump	(Optional) Displays the dump profiling data to the TTY.
profile	(Optional) Displays the CPU profiling data.
usage	(Optional) Displays the CPU usage.

Command Default

No default behavior or values.

Command Modes

The following table shows the modes in which you can enter the command:

Command Mode	Firewall Mode		Security Context		
	Routed	Transparent	Single	Multiple	
				Context	System
Privileged EXEC	• Yes	• Yes	• Yes	• Yes	• Yes

Command History

Release	Modification
7.0(1)	This command was added.
8.6(1)	The *core-id* option was added to support the ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X.
9.1(2)	The output was updated for the **show cpu profile** and **show cpu profile dump** commands.
9.2(1)	Virtual platform CPU usage was added to the output for the ASA Virtual.

Usage Guidelines

The CPU usage is computed using an approximation of the load every five seconds, and by further feeding this approximation into two, following moving averages.

You can use the **show cpu** command to find process related loads (that is, activity on behalf of items listed by the output of the **show process** command in both single mode and from the system configuration in multiple context mode).

Further, you can request, when in multiple context mode, a breakdown of the process related load to CPU consumed by any configured contexts by changing to each context and entering the **show cpu** command or by entering the **show cpu context** command.

While process related load is rounded to the nearest whole number, context related loads include one additional decimal digit of precision. For example, entering the **show cpu** command from the system context produces a different number than from entering the **show cpu context system** command. The former is an approximate summary of everything that appears in the **show cpu context all** command, and the latter is only a portion of that summary.

You can use the show cpu profile dump command in conjunction with the **cpu profile activate** command to collect information for TAC use in troubleshooting CPU issues. The **show cpu profile dump** command output is in hexadecimal format.

If the CPU profiler is waiting for a starting condition to occur, the **show cpu profile** command displays the following output:

```
CPU profiling started: 12:45:57.209 UTC Wed Nov 14 2012
CPU Profiling waiting on starting condition.
Core 0: 0 out of 10 samples collected.
Core 1: 0 out of 10 samples collected.
Core 2: 0 out of 10 samples collected.
Core 3: 0 out of 10 samples collected.
CP
0 out of 10 samples collected.
```

For the ASA Virtual, note the following licensing guidelines:

- The number of allowed vCPUs is determined by the vCPU platform license installed.
  - If the number of licensed vCPUs matches the number of provisioned vCPUs, the state is Compliant.
  - If the number of licensed vCPUs is less than the number of provisioned vCPUs, the state is Noncompliant: Over-provisioned.
  - If the number of licensed vCPUs is more than the number of provisioned vCPUs, the state is Compliant: Under-provisioned.
- The memory limit is determined by the number of vCPUs provisioned.
  - If the provisioned memory is at the allowed limit, the state is Compliant.
  - If the provisioned memory is above the allowed limit, the state is Noncompliant: Over-provisioned.
  - If the provisioned memory is below the allowed limit, the state is Compliant: Under-provisioned.
- The Frequency Reservation limit is determined by the number of vCPUs provisioned.
  - If the frequency reservation memory is at or above the required minimum (1000 MHz), the state is Compliant.

- If the frequency reservation memory is below the required minimum (1000 MHz), the state is Compliant: Under-provisioned.

For example, the following output shows that no license has been applied. The number of allowed vCPUs refers to the number licensed, and Noncompliant: Over-provisioned indicates that the product is running with more resources than have been licensed.

```
Virtual platform CPU resources
------------------------------
Number of vCPUs              :        1
Number of allowed vCPUs      :        0
vCPU Status                  :        Noncompliant: Over-provisioned
```

Copy this information and provide it to the TAC for decoding.

Note When ASA is running on FXOS chassis, the number of CPU cores displayed in the **show cpu** command outputs may be less than the number displayed in the **show version** command output on some platforms, including Firepower 4100 and 9300 (FXOS-based) platforms.

The **show cpu** command output in Firepower 4100 and 9300 platforms has been modified due to the introduction of dynamic hyper-threading support. If the traffic throughput is low, the output in the **show cpu [detailed | core | external]** CLI is different, as seen with the standalone ASA output. If the CPU hyper-threading feature is disabled, the later part of the CPU core usage output is low. When the ASA traffic throughput is above the threshold limit, enabling the CPU hyper-threading feature results in the **show cpu** command displaying the same output as the standalone ASA.

Examples

The following example shows how to display the CPU utilization:

```
ciscoasa# show cpu usage
CPU utilization for 5 seconds = 18%; 1 minute: 18%; 5 minutes: 18%
```

The following example shows how to display detailed CPU utilization information. Note that the per core information is a sum of the (data path usage + control plane usage), as shown in parentheses.

```
ciscoasa# show cpu detailed
Break down of per-core data path versus control point cpu usage:
Core         5 sec            1 min            5 min
Core 0       0.0 (0.0 + 0.0)  3.3 (0.0 + 3.3)  2.4 (0.0 + 2.4)
Current control point elapsed versus the maximum control point elapsed for:
      5 seconds = 99.0%; 1 minute: 99.8%; 5 minutes: 95.9%
CPU utilization of external processes for:
      5 seconds = 0.2%; 1 minute: 0.0%; 5 minutes: 0.0%
Total CPU utilization for:
      5 seconds = 0.2%; 1 minute: 3.3%; 5 minutes: 2.5%
```

Note The “Current control point elapsed versus the maximum control point elapsed for” statement means that the current control point load is compared to the maximum load seen within the defined time period. This is a ratio instead of an absolute number. The figure of 99% for the 5-second interval means that the current control point load is at 99% of the maximum load that is visible over this 5-second interval. If the load continues to increase all the time, then it will always remain at 100%. However, the actual CPU may still have a lot of free capacity because the maximum absolute value has not been defined. Note that this number is not the sum of the "control plane usage" numbers for the cores.

The following example shows how to display the CPU utilization for the system context in multiple mode:

```
ciscoasa# show cpu context system
CPU utilization for 5 seconds = 9.1%; 1 minute: 9.2%; 5 minutes: 9.1%
```

The following example shows how to display the CPU utilization for all contexts:

```
ciscoasa# show cpu usage context all
5 sec  1 min  5 min  Context Name
9.1%   9.2%   9.1%  system
0.0%   0.0%   0.0%  admin
5.0%   5.0%   5.0%  one
4.2%   4.3%   4.2%  two
```

The following example shows how to display the CPU utilization for a context named “one”:

```
ciscoasa/one# show cpu usage
CPU utilization for 5 seconds = 5.0%; 1 minute: 5.0%; 5 minutes: 5.0%
```

The following example activates the profiler and instructs it to store 1000 samples.

```
ciscoasa# cpu profile activate
Activated CPU profiling for 1000 samples.
Use "show cpu profile" to display the progress or "show cpu profile dump" to interrupt
profiling and display the incomplete results.
```

The following examples show the status of the profiling (in-progress and completed):

```
ciscoasa# show cpu profile
CPU profiling started: 13:45:10.400 PST Fri Nov 16 2012
CPU profiling currently in progress:
Core 0: 209 out of 1000 samples collected.
Use "show cpu profile dump" to see the results after it is complete or to interrupt profiling
 and display the incomplete results.
ciscoasa# show cpu profile dump
Cisco Adaptive Security Appliance Software Version 9.1(2)
Hardware:   ASA5555
CPU profiling started: 09:13:32.079 UTC Wed Jan 30 2013
No CPU profiling process specified.
No CPU profiling trigger specified.
cores: 2
Process virtual address map:
---------------------------
...
---------------------------
```

```
End of process map
Samples for core 0 - stopped
{0x00000000007eadb6,0x000000000211ee7e} ...
```

The following example shows CPU usage for the ASA Virtual:

```
ciscoasa# show cpu
CPU utilization for 5 seconds = 0%; 1 minute: 0%; 5 minutes: 0%
Virtual platform CPU resources
------------------------------
Number of vCPUs              :     2
Number of allowed vCPUs      :     2
vCPU Status                  :  Compliant
Frequency Reservation        :  1000 MHz
Minimum required             :  1000 MHz
Frequency Limit              :  4000 MHz
Maximum allowed              :  56000 MHz
Frequency Status             :  Compliant
Average Usage (30 seconds)   :   136 MHz
```

The following example shows details of CPU usage for the ASA Virtual:

```
Break down of per-core data path versus control point cpu usage:
Core         5 sec            1 min            5 min
Core 0       0.0 (0.0 + 0.0)  0.0 (0.0 + 0.0)  0.0 (0.0 + 0.0)
Core 1       0.0 (0.0 + 0.0)  0.2 (0.2 + 0.0)  0.0 (0.0 + 0.0)
Core 2       0.0 (0.0 + 0.0)  0.0 (0.0 + 0.0)  0.0 (0.0 + 0.0)
Core 3       0.0 (0.0 + 0.0)  0.1 (0.0 + 0.1)  0.0 (0.0 + 0.0)
Current control point elapsed versus the maximum control point elapsed for:
     5 seconds = 0.0%; 1 minute: 0.0%; 5 minutes: 0.0%
CPU utilization of external processes for:
     5 seconds = 0.0%; 1 minute: 0.0%; 5 minutes: 0.0%
Total CPU utilization for:
     5 seconds = 0.1%; 1 minute: 0.1%; 5 minutes: 0.1%
Virtual platform CPU resources
------------------------------
Number of vCPUs              :     4
Number of allowed vCPUs      :     4
vCPU Status                  :  Compliant
Frequency Reservation        :  1000 MHz
Minimum required             :  1000 MHz
Frequency Limit              :  20000 MHz
Maximum allowed              :  20000 MHz
Frequency Status             :  Compliant
Average Usage (30 seconds)   :    99 MHz
```

From ASA version 9.6.1, two or four cores are choosen for control point (CP) processing to limit the number of cores CP can run on, instead of letting CP float around all the available cores. Even if there is no traffic load, selected cores for CP processing has constant load for CPU pinning, which polls data path (DP) on each core for checking DP thread. This load is included in the **show cpu core** output, but excluded in the **show cpu detail** output because show cpu detail checks for CP and DP load.

Note On Secure Firewall 4200 series devices, core 0 is dedicated for control point, while the other cores are used to execute the data path processes.

Examples

The following example shows different CPU utilization values (Core 0 and Core 2) in the output of **show cpu core** and **show cpu detail** commands:

```
ciscoasa(config)# show cpu core
Core 5 sec 1 min 5 min
Core 0 18.0% 18.0% 18.0%
Core 1 0.0% 0.0% 0.0%
Core 2 18.6% 18.5% 18.6%
Core 3 0.0% 0.0% 0.0%
ciscoasa(config)# show cpu detail
Break down of per-core data path versus control point cpu usage:
Core 5 sec 1 min 5 min
Core 0 1.6 (0.0 + 1.6) 1.6 (0.0 + 1.6) 1.6 (0.0 + 1.6)
Core 1 0.0 (0.0 + 0.0) 0.0 (0.0 + 0.0) 0.0 (0.0 + 0.0)
Core 2 1.6 (0.0 + 1.6) 1.6 (0.0 + 1.6) 1.6 (0.0 + 1.6)
Core 3 0.0 (0.0 + 0.0) 0.0 (0.0 + 0.0) 0.0 (0.0 + 0.0)
```

Related Commands

Command	Description
show counters	Displays the protocol stack counters.
cpu profile activate	Activates CPU profiling.