Management Interfaces

Table of contents

Management interfaces are used in order to provide access to management user interfaces. NVIDIA switches support out-of-band (OOB) dedicated interfaces (e.g., mgmt0, mgmt1) and in-band dedicated interfaces. In addition, most systems feature a serial port that provides access to the CLI only.On systems with two OOB management ports, both of them may be configured on the same VLAN if needed. In this case, ARP replies to the IP of those management interfaces is answered from either of them.

Configuring Management Interfaces with Static IP Addresses

If the system was set during initialization to obtain dynamic IP addresses through DHCP and you wish to switch to static assignments, perform the following steps:

1. Enter Config configuration mode. Run:

```
switch > enable
switch # configure terminal
```

2. Disable setting IP addresses using the DHCP using the following command:

```
switch (config) # no interface <ifname> dhcp
```

3. Define your interfaces statically using the following command:

```
switch (config) # interface <ifname> ip address <IP address>
<netmask>
```

Configuring IPv6 Address on the Management Interface

1. Enable IPv6 on this interface.

```
switch (config) # interface mgmt0 ipv6 enable
```

2. Set the IPv6 address to be configured automatically.

```
switch (config) # interface mgmt0 ipv6 address autoconfig
```

3. Verify the IPv6 address is configured correctly.

```
switch (config) # show interfaces mgmt0 brief
```

Dynamic Host Configuration Protocol (DHCP)

DHCP is used for automatic retrieval of management IP addresses.

For all other systems (and software versions) DHCP is disabled by default.

Note

If a user connects through SSH, runs the wizard and turns off DHCP, the connection is immediately terminated as the management interface loses its IP address.

```
<localhost># ssh admin@<ip-address>
NVIDIA MLNX-OS Switch Management
Password:
NVIDIA switch
NVIDIA configuration wizard
Do you want to use the wizard for initial
configuration? yes
Step 1: Hostname? [my-switch]
```

```
Step 2: Use DHCP on mgmt0 interface? [yes] no
<localhost>#
```

In this case the serial connection should be used.

Default Gateway

To configure manually the default gateway, use the “ip route” command, with “0.0.0.0” as prefix and mask. The next-hop address must be within the range of one of the IP interfaces on the system.

```
switch (config)# ip route 0.0.0.0 0.0.0.0 10.10.0.2
switch (config)# show ip route
Destination       Mask              Gateway          Interface
Source       Distance/Metric
default           0.0.0.0           10.10.0.2        mgmt0          static
0/0
10.10.0.0         255.255.254.0     0.0.0.0          mgmt0          direct
0/0
```

Configuring Hostname via DHCP (DHCP Client Option 12)

This feature, also known as the DHCP Client Option 12, is enabled by default and assigns the switch system a hostname via DHCP as long as network manager configures hostname to the management interfaces’ (i.e. mgmt0, mgmt1) MAC address. If a network manager configures the hostname manually through any of the user interfaces, the hostname is not retrieved from the DHCP server.

To enable fetching hostname from DHCP server, run the following:

```
switch (config interface mgmt0) # dhcp hostname
```

To disable fetching hostname from DHCP server, run the following:

```
switch (config interface mgmt0) # no dhcp hostname
```

> **Note**
>
> Getting the hostname through DHCP is enable by default and will change the switch hostname if the hostname is not set by the user. Therefore, if a switch is part of an HA cluster the user would need to make sure the HA master has the same HA node names as the DHCP server.

Management Interface Commands

Interface

interface

	interface {mgmt0 | mgmt1 | lo | vlan<id> | ib0} Enters a management interface context.}	
Syntax Description	mgmt0	Management port 0 (out of band).
	mgmt1	Management port 1 (out of band).
	lo	Loopback interface.
	vlan<id>	In-band management interface (e.g., vlan10).
	ib0	IPoIB in-band management.
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config)# interface mgmt0 switch (config interface mgmt0)#	
Related Commands	show interfaces <ifname>	
Notes		

ip address

	ip address <IP address> <netmask> no ip address

	Sets the IP address and netmask of this interface. The no form of the command clears the IP address and netmask of this interface.	
Syntax Description	IP address	IPv4 address
	netmask	Subnet mask of IP address
Default	0.0.0.0/0	
Configuration Mode	config interface management	
History	3.1.0000	
Example	switch (config interface mgmt0)# ip address 10.10.10.10 255.255.255.0	
Related Commands	show interfaces <ifname>	
Notes	If DHCP is enabled on the specified interface, then the DHCP IP assignment will hold until DHCP is disabled	

ip default-gateway

	ip default-gateway <next-hop-IP-address> <interface-name> no default-gateway <next-hop-IP-address> <interface-name> Configures a default route. The no form of the command removes the current default route.	
Syntax Description	next hop IP address	gateway IP address
	interface name	default gateway interface name
Default	N/A	
Configuration Mode	config interface management	
History	3.1.0000 3.8.1000: Updated Command & Syntax description	
Example	switch (config interface mgmt0)# ip default-gateway mgmt1	
Related Commands		
Notes		

alias

	alias <index> ip address < IP address> <netmask> no alias <index> Adds an additional IP address to the specified interface. The secondary address will appear in the output of “show interface” under the data of the primary interface along with the alias. The no form of the command removes the secondary address to the specified interface.	
Syntax Description	index	A number that is to be aliased to (associated with) the secondary IP.
	IP address	Additional IP address.
	netmask	Subnet mask of the IP address.
Default	N/A	
Configuration Mode	config interface management	
History	3.1.0000	
Example	switch (config interface mgmt0)# alias 2 ip address 9.9.9.9 255.255.255.255	
Related Commands	show interfaces <ifname>	
Notes	• If DHCP is enabled on the specified interface, then the DHCP IP assignment will hold until DHCP is disabled • More than one additional IP address can be added to the interface	

mtu

	mtu <bytes> no mtu <bytes> Sets the Maximum Transmission Unit (MTU) of this interface. The no form of the command resets the MTU to its default.	
Syntax	bytes	The entry range is 68-1500.

Description		
Default	1500	
Configuration Mode	config interface management	
History	3.6.3004	
Example	switch (config interface mgmt0)# mtu 1500	
Related Commands	show interfaces <ifname>	
Notes		

duplex

	duplex <duplex> no duplex Sets the interface duplex. The no form of the command resets the duplex setting for this interface to its default value.	
Syntax Description	duplex	Sets the duplex mode of the interface. The following are the possible values: • half—half duplex • full—full duplex • auto—auto duplex sensing (half or full)
Default	auto	
Configuration Mode	config interface management	
History	3.1.0000	
Example	switch (config interface mgmt0)# duplex auto	
Related Commands	show interfaces <ifname>	
Notes	• Setting the duplex to “auto” also sets the speed to “auto” • Setting the duplex to one of the settings “half” or “full” also sets the speed to a manual setting which is determined by	

	querying the interface to find out its current auto-detected state

speed

	speed <speed> no speed Sets the interface speed. The no form of the command resets the speed setting for this interface to its default value.	
Syntax Description	speed	Sets the speed of the interface. The following are the possible values: • 10—fixed to 10Mbps • 100—fixed to 1000Mbps • 1000—fixed to 1000Mbps • auto—auto speed sensing (10/100/1000Mbps)
Default	auto	
Configuration Mode	config interface management	
History	3.1.0000	
Example	switch (config interface mgmt0)# speed auto	
Related Commands	show interfaces <ifname>	
Notes	• Setting the speed to “auto” also sets the duplex to “auto” • Setting the speed to one of the manual settings (generally “10”, “100”, or “1000”) also sets the duplex to a manual setting which is determined by querying the interface to find out its current auto-detected state	

dhcp

	dhcp [renew]

	no dhcp Enables DHCP on the specified interface. The no form of the command disables DHCP on the specified interface.	
Syntax Description	renew	Forces a renewal of the IP address. A restart on the DHCP client for the specified interface will be issued.
Default	Could be enabled or disabled (per part number) manufactured with 3.2.0500	
Configuration Mode	config interface management	
History	3.1.0000 3.9.1900: Added note	
Example	switch (config interface mgmt0)# dhcp	
Related Commands	show interfaces <ifname> configured	
Notes	• When enabling DHCP, the IP address and netmask are received via DHCP hence, the static IP address configuration is ignored • Enabling DHCP disables zeroconf and vice versa • Setting a static IP address and netmask does not disable DHCP. DHCP is disabled using the “no” form of this command, or by enabling zeroconf. • When static IP is configured, DHCP will not run.	

dhcp hostname

	dhcp hostname no dhcp hostname Enables fetching the hostname from DHCP for this interface. The no form of the command disables fetching the hostname from DHCP for this interface.
Syntax Description	N/A
Default	Enabled
Configuration Mode	config interface management
History	3.5.1000

Example	switch (config interface mgmt0)# dhcp hostname
Related Commands	hostname <hostname> show interfaces <ifname> configured
Notes	• If a hostname is configured manually by the user, that configuration would override the “dhcp hostname” configuration • When a default hostname is not configured, the DHCP server assigns the new hostname for your machine (after upgrading to version 3.5.1000) • These commands do not work on in-band interfaces

shutdown

	shutdown no shutdown Disables the specified interface. The no form of the command enables the specified interface.
Syntax Description	N/A
Default	no shutdown
Configuration Mode	config interface management
History	3.1.0000
Example	switch (config interface mgmt0)# no shutdown
Related Commands	show interfaces <ifname> configured
Notes	

zeroconf

	zeroconf no zeroconf Enables zeroconf on the specified interface. It randomly chooses a unique link-local IPv4 address from the 169.254.0.0/16 block. This command is an alternative to DHCP.

	The no form of the command disables the use of zeroconf on the specified interface.
Syntax Description	N/A
Default	no zeroconf
Configuration Mode	config interface management
History	3.1.0000
Example	switch (config interface mgmt0)# zeroconf
Related Commands	show interfaces <ifname> configured
Notes	Enabling zeroconf disables DHCP and vice versa.

comment

	comment <comment> no comment Adds a comment for an interface. The no form of the command removes a comment for an interface.	
Syntax Description	comment	A free-form string that has no semantics other than being displayed when the interface records are listed.
Default	no comment	
Configuration Mode	config interface management	
History	3.1.0000	
Example	switch (config interface mgmt0)# comment my-interface	
Related Commands		
Notes		

ipv6 enable

	ipv6 enable no ipv6 enable

	Enables all IPv6 addressing for this interface. The no form of the command disables all IPv6 addressing for this interface.	
Syntax Description	N/A	
Default	IPv6 addressing is disabled	
Configuration Mode	config interface management	
History	3.1.0000	
Example	switch (config interface mgmt0)# ipv6 enable	
Related Commands	ipv6 address show interface <ifname>	
Notes	• The interface identifier is a 64-bit long modified EUI-64, which is based on the MAC address of the interface • If IPv6 is enabled on an interface, the system will automatically add a link-local address to the interface. Link-local addresses can only be used to communicate with other hosts on the same link, and packets with link-local addresses are never forwarded by a router. • A link-local address, which may not be removed, is required for proper IPv6 operation. The link-local addresses start with “fe80::”, and are combined with the interface identifier to form the complete address.	

ipv6 address

	ipv6 address {<IPv6 address/netmask> | autoconfig [default | privacy]} no ipv6 {<IPv6 address/netmask> | autoconfig [default | privacy]} Configures IPv6 address and netmask to this interface, static or autoconfig options are possible. The no form of the command removes the given IPv6 address and netmask or disables the autoconfig options.	
Syntax Description	IPv6 address/netmask	Configures a static IPv6 address and netmask. Format example: 2001:db8:1234::5678/64.
	autoconfig	Enables IPv6 stateless address auto configuration (SLAAC) for this interface. An address will be

		automatically added to the interface based on an IPv6 prefix learned from router advertisements, combined with an interface identifier.
	autoconfig default	Enables default learning routes. The default route will be discovered automatically, if the autoconfig is enabled.
	autoconfig privacy	Uses privacy extensions for SLAAC to construct the autoconfig address, if the autoconfig is enabled.
Default	No IP address available, auto config is enabled	
Configuration Mode	config interface management	
History	3.1.0000	
Example	switch (config interface mgmt0)# ipv6 fe80::202:c9ff:fe5e:a5d8/64	
Related Commands	ipv6 enable show interface <ifname>	
Notes	• On a given interface, up to 16 addresses can be configured • For Ethernet, the default interface identifier is a 64-bit long modified EUI-64, which is based on the MAC address of the interface	

ipv6 dhcp primary-intf

	ipv6 dhcp primary-intf <if-name> no ipv6 dhcp primary-intf Sets the interface from which non-interface-specific (resolver) configuration is accepted via DHCPv6. The no form of the command resets non-interface-specific (resolver) configuration.	
Syntax Description	if-name	Interface name: • lo • mgmt0 • mgmt1
Default	N/A	

Configuration Mode	config
History	3.1.0000
Example	switch (config)# ipv6 dhcp primary-intf mgmt0
Related Commands	ipv6 enable ipv6 address show interface <ifname>
Notes	

ipv6 dhcp stateless

	ipv6 dhcp stateless no ipv6 dhcp stateless Enables stateless DHCPv6 requests. The no form of the command disables stateless DHCPv6 requests.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	switch (config)# ipv6 dhcp stateless
Related Commands	ipv6 enable ipv6 address show interface <ifname>
Notes	• This command only gets DNS configuration, not an IPv6 address • The no form of the command requests all information, including an IPv6 address

ipv6 dhcp client enable

	ipv6 dhcp client enable

	no ipv6 dhcp client enable Enables DHCPv6 on this interface. The no form of the command disables DHCPv6 on this interface.
Syntax Description	N/A
Default	ipv6 dhcp client enable
Configuration Mode	config interface management
History	3.7.11xx 3.9.1900: Added note
Example	switch (config interface mgmt0)# ipv6 dhcp client enable
Related Commands	ipv6 dhcp client renew show ipv6 dhcp
Notes	When static IP is configured, DHCP will not run.

ipv6 dhcp client renew

	ipv6 dhcp client renew Renews DHCPv6 lease for this interface.
Syntax Description	N/A
Default	N/A
Configuration Mode	config interface management
History	3.7.11xx
Example	switch (config interface mgmt0)# ipv6 dhcp client renew
Related Commands	ipv6 dhcp client enable show ipv6 dhcp
Notes	

show interfaces mgmt0

	show interface mgmt0

	Displays information on the management interface configuration and status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.8008: Updated example 3.9.1900: Updated example—added new output option of "no (Static IP is configured)"

Example	switch (config)# show interfaces mgmt0
	Interface mgmt0 status:
	Comment :
	Admin up : yes
	Link up : yes
	DHCP running : no (Static IP is configured)
	IP address : 10.12.67.33
	Netmask : 255.255.255.128
	IPv6 enabled : yes
	Autoconf enabled: no
	Autoconf route : yes
	Autoconf privacy: no
	DHCPv6 running : no (Static IP is configured)
	IPv6 addresses : 2
	`IPv6 address: 1::1/64 fe80::7efe:90ff:fe65:dea8/64`
	`Speed : UNKNOWN Duplex : full Interface type : ethernet Interface source: bridge Bonding master : vrf_vrf-default MTU : 1500 HW address : 7C:FE:90:65:DE:A8`
	`Rx: 13840892 bytes 58605 packets 0 mcast packets 2 discards 0 errors 0 overruns 0 frame`
	`Tx: 3796 bytes 38 packets 0 discards 0 errors 0 overruns 0 carrier 0 collisions 1000 queue len`

Related Commands	

Notes	

show interfaces mgmt0 brief

	show interface mgmt0 brief Displays brief information on the management interface configuration and status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.8008: Updated example
Example	switch (config)# show interfaces mgmt0 brief Interface mgmt0 status: Comment : Admin up : yes Link up : yes DHCP running : yes IP address : 10.12.67.33 Netmask : 255.255.255.128 IPv6 enabled : yes Autoconf enabled: no Autoconf route : yes Autoconf privacy: no DHCPv6 running : yes (but no valid lease) IPv6 addresses : 1 IPv6 address: fe80::268a:7ff:fe53:3d8e/64 Speed : 1000Mb/s (auto) Duplex : full (auto) Interface type : ethernet Interface source: bridge MTU : 1500 HW address : 24:8a:07:53:3d:8e
Related Commands	
Notes	

show interfaces mgmt0 configured

	show interface mgmt0 configured Displays configuration information about the specified interface.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.5.1000: Updated example with “DHCP Hostname” 3.6.8008: Updated example
Example	switch (config)# show interfaces mgmt0 configured Interface mgmt0 configuration: Comment : Enabled : yes DHCP : yes DHCP Hostname : yes Zeroconf : no IP address : Netmask : IPv6 enabled : yes Autoconf enabled: no Autoconf route : yes Autoconf privacy: no DHCPv6 enabled : yes IPv6 addresses : 0 Speed : auto Duplex : auto MTU : 1500
Related Commands	
Notes	

Hostname Resolution

hostname

	hostname <hostname> no hostname Sets a static system hostname. The no form of the command clears the system hostname.	
Syntax Description	hostname	A free-form string
Default	Default hostname	
Configuration Mode	config	
History	3.1.0000 3.6.3004: Added support for the character “.”	
Example	switch (config)# hostname my-switch-hostname	
Related Commands	show hosts	
Notes	• Hostname may contain letters, numbers, periods (‘.’), and hyphens (‘-’), in any combination • Hostname may be 1-63 characters long • Hostname may not begin with a hyphen • Hostname may not contain other characters, such as “%”, “_” etc. • Hostname may not be set to one of the valid logging commands (i.e. debug-files, fields, files, format, level, local, monitor, receive, trap) • Changing the hostname stamps a new HTTPS certificate	

ip name-server

	ip name-server <IPv4/IPv6 address> no ip name-server <IPv4/IPv6 address> Sets the static name server. The no form of the command clears the name server.	
Syntax Description	IPv4/IPv6 address	IPv4 or IPv6 address.
Default	No server name	
Configuration Mode	config	

History	3.1.0000
Example	switch (config)# ip name-server 9.9.9.9
Related Commands	show hosts
Notes	

ip domain-list

	ip domain-list <domain-name> no ip domain-list <domain-name> Sets the static domain name. The no form of the command clears the domain name.	
Syntax Description	domain-name	The domain name in a string form. A domain name is an identification string that defines a realm of administrative autonomy, authority, or control in the Internet.
Default	No static domain name	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config)# ip domain-list mydomain.com	
Related Commands	show hosts	
Notes		

ip/ipv6 host

	{ip | ipv6} host <hostname> <ip-address> no {ip | ipv6} host <hostname> <ip-address> Configures the static hostname IPv4 or IPv6 address mappings. The no form of the command clears the static mapping.	
Syntax Description	hostname	The hostname in a string form.

	IP Address	The IPv4 or IPv6 address.
Default	No static domain name	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config)# ip host my-host 2.2.2.2 switch (config)# ipv6 host my-ipv6-host 2001::8f9	
Related Commands	show hosts	
Notes		

ip/ipv6 map-hostname

	{ip |ipv6} map-hostname no {ip | ipv6} map-hostname Maps between the currently-configured hostname and the loopback address 127.0.0.1. The no form of the command clears the mapping.
Syntax Description	N/A
Default	IPv4 mapping is enabled by default IPv6 mapping is disabled by default
Configuration Mode	config
History	3.1.0000
Example	switch (config)# ip map-hostname
Related Commands	show hosts
Notes	• If no mapping is configured, a mapping between the hostname and the IPv4 loopback address 127.0.0.1 will be added • The no form of the command maps the hostname to the IPv6 loopback address if there is no statically configured mapping from the hostname to an IPv6 address (disabled by default) • Static host mappings are preferred over DNS results. As a result, with this option set, you will not be able to look up your

	hostname on your configured DNS server; but without it set, some problems may arise if your hostname cannot be looked up in DNS.

show hosts

	show hosts Displays hostname, DNS configuration, and static host mappings.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.8.1000: Updated example
Example	switch (config)# show hosts Hostname: switch1 Name servers: 10.7.77.192 dynamic (DHCP on mgmt0) 10.7.77.135 dynamic (DHCP on mgmt0) 10.198.0.169 dynamic (DHCP on mgmt0) (*) 10.211.0.124 dynamic (DHCP on mgmt0) Domain names: mtl.labs.mlnx dynamic (DHCP on mgmt0) (*) Inactive due to system limits on name servers and domain names. Static IPv4 host mappings: 10.7.144.133 --> switch1 127.0.0.1 --> localhost Static IPv6 host mappings: ::1 --> localhost6 Automatically map hostname to loopback address : yes Automatically map hostname to IPv6 loopback address: no
Related Commands	
Notes	

Routing

IP route

	{ip | ipv6} route {<network-prefix> <netmask> | <network-prefix>/<masklen>} <next-hop> no ip route {<network-prefix> <netmask> | <network-prefix>/<masklen>} <next-hop> Sets a static route for a given IP. The no form of the command deletes the static route.	
Syntax Description	network-prefix	IPv4 or IPv6 network prefix
	netmask	IPv4 netmask formats are: • /24 • 255.255.255.0 IPv6 netmask format is: • /48 (as a part of the network prefix)
	nexthop-address	The IPv4 or IPv6 address of the next hop router for this route
	ifname	The interface name (e.g., mgmt0, mgmt1)
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config)# ip route 20.20.20.0 255.255.255.0 mgmt0	
Related Commands	show ip route	
Notes		

ipv6 default-gateway

	ipv6 default-gateway {<ip-address> | <ifname>} no ipv6 default-gateway Sets a static default gateway.

	The no form of the command deletes the default gateway.	
Syntax Description	ip address	The default gateway IP address (IPv6)
	ifname	The interface name (e.g.,, mgmt0, mgmt1)
Default	N/A	
Configuration Mode	config	
History	3.1.0000 3.2.0500: Removed IPv4 configuration option	
Example	switch (config)# ipv6 default-gateway ::1	
Related Commands	show ip/ipv6 route show ipv6 default-gateway	
Notes	• The configured default gateway will not be used if DHCP is enabled • In order to configure ipv4 default-gateway use 'ip route' command.	

show ip/ipv6 route

	show {ip | ipv6} route [static] Displays the routing table in the system.	
Syntax Description	static	Filters the table with the static route entries
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example		

switch (config)# show ip route

Destination	Mask	Gateway	Interface	Source
default	0.0.0.0	172.30.0.1	mgmt0	DHCP
10.10.10.10	255.255.255.255	0.0.0.0	mgmt0	static
20.10.10.10	255.255.255.255	172.30.0.1	mgmt0	static
20.20.20.0	255.255.255.0	0.0.0.0	mgmt0	static
172.30.0.0	255.255.0.0	0.0.0.0	mgmt0	interface

switch (config)# show ipv6 route Destination prefix Gateway Interface Source --- ::/0 :: mgmt0 static ::1/128 :: lo local 2222:2222:2222::/64 :: mgmt1 interface	
Related Commands	ip route
Notes	

show ipv6 default-gateway

	show ipv6 default-gateway [static] Displays the default gateway.	
Syntax Description	static	Displays the static configuration of the default gateway
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example	switch (config)# show ipv6 default-gateway Active default gateways: 172.30.0.1 (interface: mgmt0) switch (config)# show ipv6 default-gateway static Configured default gateway: 10.10.10.10	
Related Commands	ipv6 default-gateway	
Notes	The configured IPv4 default gateway will not be used if DHCP is enable	

Network to Media Resolution (ARP & NDP)

IPv4 network use Address Resolution Protocol (ARP) to resolve IP address to MAC address, while IPv6 network uses Network Discovery Protocol (NDP) that performs basically the same as ARP.

ipv6 neighbor

	ipv6 neighbor <ipv6-address> <ifname> <mac-address> no ipv6 neighbor <ipv6-address> <ifname> <mac-address> Adds a static neighbor entry. The no form of the command deletes the static entry.	
Syntax Description	ipv6-address	The IPv6 address
	ifname	The management interface (i.e. mgmt0, mgmt1)
	mac-address	The MAC address
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config)# ipv6 neighbor 2001:db8:701f::8f9 mgmt0 00:11:22:33:44:55	
Related Commands	show ipv6 neighbor ipv6 route arp clear ipv6 neighbors	
Notes	• ARP is used only with IPv4. In IPv6 networks, Neighbor Discovery Protocol (NDP) is used similarly. • Use The no form of the command to remove static entries. Dynamic entries can be cleared via the “clear ipv6 neighbors” command.	

clear ipv6 neighbors

	clear ipv6 neighbors Clears the dynamic neighbors cache.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000 3.6.4110: Updated command
Example	switch (config)# clear ipv6 neighbors
Related Commands	ipv6 neighbor show ipv6 neighbor arp
Notes	• Clearing Neighbor Discovery Protocol (NDP) cache removes only the dynamic entries learned and not the static entries configured • Use the no form of the command to remove static entries

show ipv6 neighbors

	show ipv6 neighbors [static] Displays the Neighbor Discovery Protocol (NDP) table.	
Syntax Description	static	Filters only the table of the static entries.
Default	N/A	
Configuration Mode	Any command mode	
History	3.1.0000	
Example		
switch (config)# show ipv6 neighbors IPv6 Address Age MAC Address State Interf ---------- ----- ---------- ---------- ------ 2001::2 9428 aa:aa:aa:aa:aa:aa permanent mgmt0		
Related Commands	ipv6 neighbor clear ipv6 neighbor show ipv6	

Notes	

DHCP

ip dhcp

	ip dhcp {default-gateway yield-to-static | hostname <hostname>| primary-intf <ifname> | send-hostname} no ip dhcp {default-gateway yield-to-static| hostname | | primary-intf | send-hostname} Sets global DHCP configuration. The no form of the command deletes the DHCP configuration.	
Syntax Description	yield-to-static	Does not allow you to install a default gateway from DHCP if there is already a statically configured one.
	hostname	Specifies the hostname to be sent during DHCP client negotiation if send-hostname is enabled.
	primary-intf <ifname>	Sets the interface from which a non-interface-specific configuration (resolver and routes) will be accepted via DHCP. Default: "primary-intf mgmt0"
	send-hostname	Enables the DHCP client to send a hostname during negotiation.
Default	Disabled	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config)# ip dhcp default-gateway yield-to-static	
Related Commands	show ip dhcp dhcp [renew]	
Notes	DHCP is supported for IPv4 networks only	

show ip dhcp

	show ip dhcp Displays the DHCP configuration and status.
Syntax Description	N/A
Default	N/A
Configuration Mode	Any command mode
History	3.1.0000 3.6.5000: Updated example
Example	switch (config)# show ip dhcp ---------- Interface DHCP DHCP Valid Enabled Running lease ---------- dummy0 no no no lo no no no mgmt0 yes yes yes mgmt1 no no no mgmts0 no no no mgmts1 no no no vif1 no no no IPv4 dhcp default gateway yields to static configuration: no DHCP primary interface: Configured: mgmt0 Active: mgmt0 DHCP client options: Send Hostname: no Client Hostname: 1.1.1.1
Related Commands	ip dhcp dhcp [renew]
Notes	

IP Diagnostic Tools

ping

	ping [-LRUbdfnqrvVaA] [-c count] [-i interval] [-w deadline] [-p pattern] [-s packetsize] [-t ttl] [-I interface or address] [-M mtu

	discovery hint] [-S sndbuf] [-T timestamp option] [-Q tos] [hop1 ...] destination Sends ICMP echo requests to a specified host.	
Syntax Description	Linux Ping options	https://www.lifewire.com/uses-of-command-ping-2201076
Default	N/A	
Configuration Mode	config	
History	3.1.0000	
Example	switch (config)# ping 172.30.2.2 PING 172.30.2.2 (172.30.2.2) 56(84) bytes of data. 64 bytes from 172.30.2.2: icmp_seq=1 ttl=64 time=0.703 ms 64 bytes from 172.30.2.2: icmp_seq=2 ttl=64 time=0.187 ms 64 bytes from 172.30.2.2: icmp_seq=3 ttl=64 time=0.166 ms 64 bytes from 172.30.2.2: icmp_seq=4 ttl=64 time=0.161 ms 64 bytes from 172.30.2.2: icmp_seq=5 ttl=64 time=0.153 ms 64 bytes from 172.30.2.2: icmp_seq=6 ttl=64 time=0.144 ms ... --- 172.30.2.2 ping statistics --- 6 packets transmitted, 6 received, 0% packet loss, time 5004ms rtt min/avg/max/mdev = 0.144/0.252/0.703/0.202 ms	
Related Commands	traceroutes	
Notes		

traceroute

	traceroute [-46dFITUnrAV] [-f first_ttl] [-g gate,...] [-i device] [-m max_ttl] [-N squeries] [-p port] [-t tos] [-l flow_label] [-w waittime] [-q nqueries] [-s src_addr] [-z sendwait] host [packetlen] Traces the route packets take to a destination.	
Syntax Description	-4	Uses IPv4
	-6	Uses IPv6
	-d	Enables socket level debugging
	-F	Sets DF (do not fragment bit) on
	-I	Uses ICMP ECHO for tracerouting

	-T	Uses TCP SYN for tracerouting
	-U	Uses UDP datagram (default) for tracerouting
	-n	Does not resolve IP addresses to their domain names
	-r	Bypasses the normal routing and send directly to a host on an attached network
	-A	Performs AS path lookups in routing registries and print results directly after the corresponding addresses
	-V	Prints version info and exit
	-f	Starts from the first_ttl hop (instead from 1)
	-g	Routes packets through the specified gateway (maximum 8 for IPv4 and 127 for IPv6)
	-i	Specifies a network interface with which to operate
	-m	Sets the max number of hops (max TTL to be reached). Default is 30.
	-N	Sets the number of probes to be tried simultaneously (default is 16)
	-p	Uses destination port. It is an initial value for the UDP destination port (incremented by each probe, default is 33434), for the ICMP seq number (incremented as well, default from 1), and the constant destination port for TCP tries (default is 80).
	-t	Sets the TOS (IPv4 type of service) or TC (IPv6 traffic class) value for outgoing packets
	-l	Uses specified flow_label for IPv6 packets
	-w	Sets the number of seconds to wait for response to a probe (default is 5.0). Non-integer (float point) values allowed too.
	-s	Uses source src_addr for outgoing packets.
	-q	Sets the number of probes per each hop. Default is 3.
	-z	Sets minimal time interval between probes (default is 0). If the value is more than 10, then it specifies a number in milliseconds, else it is a number of seconds (float point values allowed too).

Default	N/A
Configuration Mode	config
History	3.1.0000
Example	
switch (config)# traceroute 192.168.10.70 traceroute to 192.168.10.70 (192.168.10.70), 30 hops max, 40 byte packets 1 172.30.0.1 (172.30.0.1) 3.632 ms 2.849 ms 3.544 ms 2 10.222.128.46 (10.222.128.46) 3.176 ms 3.289 ms 3.656 ms 3 10.158.128.30 (10.158.128.30) 15.331 ms 15.819 ms 16.388 ms 4 10.158.128.65 (10.158.128.65) 20.468 ms 7.893 ms 12.27 ms 5 10.7.34.115 (10.7.34.115) 16.405 ms 11.985 ms 12.264 ms6 192.168.10.70 (192.168.10.70) 16.377 ms 16.091 ms 20.475 ms	
Related Commands	ping
Notes	

tcpdump

	tcpdump [-aAdDefILnNOpqRStuUvxX] [-c count] [-C file_size] [-E algo:secret] [-F file] [-i interface] [-M secret] [-r file] [-s snaplen] [-T type] [-w file] [-W filecount] [-y datalinktype] [-Z user] [-D list possible interfaces] [expression] Invokes standard binary, passing command line parameters straight through. Runs in foreground, printing packets as they arrive, until the user hits Ctrl+C.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.1.0000
Example	switch (config)# tcpdump ..... 09:37:38.678812 IP 192.168.10.7.ssh > 192.168.10.1.54155: P 1494624:1494800(176) ack 625 win 90 <nop,nop,timestamp 5842763 858672398>

	09:37:38.678860 IP 192.168.10.7.ssh > 192.168.10.1.54155: P 1494800:1495104(304) ack 625 win 90 <nop,nop,timestamp 5842763 858672398> ... 9141 packets captured 9142 packets received by filter 0 packets dropped by kernel
Related Commands	
Notes	

Control Plane Policing (CoPP)

Control Plane Policing or Policies (CoPP) ensures the CPU and control plane are not over-utilized which is essential for the robustness of the switch. CoPP limits the number of control plane packets.

This software implements several CoPP mechanisms:

- ACLs may be used to limit the rate of packets or bytes of a certain type, including L3 control packets (L2 control packets are forwarded to the CPU before the ACL)
- Policers on traffic going to the CPU—these policers are configured by the operating system and cannot be modified by the user
- IP filter tables limit the traffic to the CPU coming in from the management ports

IP Table Filtering

IP table filtering is a mechanism that allows the user to apply actions to a specific control packet flow identified by a certain flow key.

This mechanism is used in order to protect switch control traffic against attacks. For example, it could allow traffic coming from a specific trusted management subnet only, block the SNMP UDP port from receiving traffic, and force ping rate to be lower than a specific threshold.

Each IP table rule is defined by key, priority, and action:

- Key—the key is a combination of physical port and layer 3 parameters (e.g. SIP, DIP, SPORT, DPORT, etc.), and other fields. Each part of the key, can be set to a specific value or masked.
- Priority—each rule in the IP table is assigned a priority, and the rule with the highest priority whose key matches the packet executes the action.
- Action—the action describes the behavior of packets which match the key. The action type may be drop, accept, rate limit, etc.

An IP-table rule is bound to an IP interface that can be a management out-of-band interface, VLAN interface, or router port interface. Once bound, all traffic received (ingress

rule) or transmitted (egress rule) in this direction is being verified with all bounded rules.

Once a match was found, the rule action is executed. If no match is found, the default policy of the chain shall apply.

> **Note**
>
> IP table rules get a lower priority than ACL mechanism.

> **Note**
>
> In the rare case that IP filter is used while the input policy is "drop" (i.e., ip filter chain input policy drop) and an NTP server or an InfiniBand switch with SM HA enabled is used, then the following rule needs to be added that allows src-ip 127.0.0.1 (which is a requirement for any clustered application (e.g., sm-ha) and NTP):
>
> ip filter chain input rule append tail target accept dup-delete source-addr 127.0.0.1 /32

Configuring IP Table Filtering

Prerequisite for IPv6:

```
switch (config) # ipv6 enable
```

To configure IPv4 table filtering:

1. Select the policy that applies to the input/output chain (default is “accept”).

   ```
   switch (config)# ip filter chain input policy drop
   ```

```
switch (config)# ip filter chain output policy accept
```

2. Append filtering rules to the list or set a specific rule number, select a target, and (optional) any additional filter conditions. For example:

```
switch (config) # ip filter chain input rule append tail target
rate-limit 2 protocol udp
switch (config) # ip filter chain input rule set 2 target drop
protocol icmp in-intf mgmt1
switch (config) # ip filter chain output rule append tail
target drop protocol icmp
```

3. Enable IP table filtering.

```
switch (config) # ip filter enable
```

4. Verify IP table filtering configuration.

```
switch (config) # show ip filter configured

Packet filtering for IPv4: enabled

IPv4 configuration:
  Chain 'input' Policy 'accept':
    Rule 1:
      Target      : rate-limit 2 pps
      Protocol    : udp
      Source      : all
      Destination : all
      Interface   : all
      State       : any
```

```
      Other Filter:  -

    Rule 2:
      Target      : drop
      Protocol    : icmp
      Source      : all
      Destination : all
      Interface   : mgmt1 (ingress)
      State       : any
      Other Filter:  -

  Chain 'output' Policy 'accept':
    Rule 1:
      Target      : drop
      Protocol    : icmp
      Source      : all
      Destination : all
      Interface   : all
      State       : any
      Other Filter: -
```

Modifying IP Table Filtering

To modify IP table filtering configuration:

```
switch (config) # ip filter chain input rule modify 3 target
reject-with icmp6-adm-prohibited source-addr 10::0 /126
```

To delete an existing IP table filtering rule:

```
switch (config) # no ip filter chain input rule 2
```

To delete all existing IP table filtering rules:

```
switch (config) # no ip filter chain output rule all
```

To insert an IP table filtering rule in a chain:

```
switch (config) # ip filter chain input rule 2 set target drop
protocol tcp dest-port 22 in-intf mgmt1
```

Rate-Limit Rule Configuration

Using a rate-limit target allows to create a rule to limit the rate of certain traffic types. The limit is specified in packets per second (pps) and can be anywhere between 1-1000 pps. When enabled, the system takes the user specified rate and converts it into units of 1/10000 of a second. Therefore, any value greater than 100 can have a slight difference when the rule is displayed using the show command.

Rate limits can be set using the parameter "rate-limit-above" in order to drop packets whenever traffic is above the set limit. For example: ip filter chain input rule append tail target drop rate-limit-above 1/second source-addr 1.1.1.1 /32.

Another option is to use the parameter "rate-limit". This should be followed by a rule that drops additional packets of the same “type”. Alternatively, this can be implicitly achieved by setting the chain policy to “drop” so that it drops packets not processed by matching rules. Otherwise, no effect of the rule is observed as the remaining traffic simply gets accepted.

Note

Rate-limit is implemented with an average rate and a burst-limit. Rate values are specified in pps and take a range from 1-1000 pps. For rate values in the range 1-100, the burst value is set equal to the rate value. For rate values in the range 101-1000, the burst limit is set to 100.

IP Table Filtering Default Rules

IP table filtering is enabled on both ipv4 and ipv6 and Firewall default IP filter rules are applied.

- To reset/apply default rules on system, run the command “ip filter reset-to-default-rules” for ipv4 **or** "ipv6 filter reset-to-default-rules" for IPv6.
- To enable IP Filter, run the command “ip filter enable”, "ipv6 filter enable".
- To list the default firewall rules, run the command “show ip filter”, "show ipv6 filter".
- Note when touching a default rule (delete/move/modify) all IP Filter rules will be reflected on “show running-config”, to restore default rules, run the command “ip filter reset-to-default-rules” **or "**ipv6 filter reset-to-default-rules**"**
- Restoring factory default configuration will reset the default rules and enable the feature

IPv4 Firewall Default Rules

Prerouting-Mangle Chain Rules
• ip filter chain prerouting-mangle rule append tail target drop in-intf mgmt0 protocol tcp conntrack new tcp-op-mss mss-not-in-range 536:65535 not-dest-port 22-23
Input Chain Rules
• ip filter chain input rule append tail target accept in-intf lo • ip filter chain input rule append tail target drop dup-delete dest-addr 127.0.0.0 /8 in-intf mgmt0 • ip filter chain input rule append tail target accept dup-delete in-intf mgmt0 state established,related

Prerouting-Mangle Chain Rules
• ip filter chain input rule append tail target drop dup-delete not-dest-port 22-23 in-intf mgmt0 protocol tcp state new tcp-op syn match-not-syn • ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 fragment enable • ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 protocol tcp tcp-op flags all • ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 protocol tcp tcp-op flags none • ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 state invalid • ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 protocol tcp tcp-op flags reset rate-limit-above 2/second burst-limit-above 2 • ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 protocol tcp state new rate-limit-above 50/second burst-limit-above 50 • ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 protocol tcp conntrack new rate-limit-above 60/second burst-limit-above 20 • ip filter chain input rule append tail target drop dup-delete in-intf mgmt0 recent name portscan recent rcheck-sec 86400 • ip filter chain input rule append tail target none dup-delete in-intf mgmt0 recent name portscan recent remove • ip filter chain input rule append tail target none dup-delete dest-port 22 in-intf mgmt0 protocol tcp conntrack new recent set • ip filter chain input rule append tail target drop dup-delete dest-port 22 in-intf mgmt0 protocol tcp conntrack new recent update-sec 60 recent hitcount 50 • ip filter chain input rule append tail target none dup-delete dest-port 23 in-intf mgmt0 protocol tcp conntrack new recent set • ip filter chain input rule append tail target drop dup-delete dest-port 23 in-intf mgmt0 protocol tcp conntrack new recent update-sec 60 recent hitcount 50 • ip filter chain input rule append tail target none dup-delete dest-port 443 in-intf mgmt0 protocol tcp conntrack new recent set • ip filter chain input rule append tail target drop dup-delete dest-port 443 in-intf mgmt0 protocol tcp conntrack new recent update-sec 60 recent hitcount 150 • ip filter chain input rule append tail target none dup-delete dest-port 80 in-intf mgmt0 protocol tcp conntrack new recent set • ip filter chain input rule append tail target drop dup-delete dest-port 80 in-intf mgmt0 protocol tcp conntrack new recent update-sec 60 recent hitcount 150 • ip filter chain input rule append tail target none dup-delete dest-port 23108 in-intf mgmt0 protocol tcp conntrack new recent set • ip filter chain input rule append tail target drop dup-delete dest-port 23108 in-intf mgmt0 protocol tcp conntrack new recent update-sec 60 recent hitcount 150 • ip filter chain input rule append tail target none dup-delete dest-port 161 in-intf mgmt0 protocol udp conntrack new recent set

Prerouting-Mangle Chain Rules
• ip filter chain input rule append tail target drop dup-delete dest-port 161 in-intf mgmt0 protocol udp conntrack new recent update-sec 60 recent hitcount 100 • ip filter chain input rule append tail target accept dup-delete dest-port 22 in-intf mgmt0 protocol tcp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 23 in-intf mgmt0 protocol tcp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 443 in-intf mgmt0 protocol tcp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 80 in-intf mgmt0 protocol tcp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 23108 in-intf mgmt0 protocol tcp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 179 in-intf mgmt0 protocol tcp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 68 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 122 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 161 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 6306 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 69 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 389 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 389 in-intf mgmt0 protocol tcp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 1812-1813 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 49 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 49 in-intf mgmt0 protocol tcp conntrack new,established • ip filter chain input rule append tail target accept dup-delete source-port 53 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete source-port 53 in-intf mgmt0 protocol tcp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 500 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 4500 in-intf mgmt0 protocol udp conntrack new,established

Prerouting-Mangle Chain Rules
• ip filter chain input rule append tail target accept dup-delete dest-port 1293 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 1293 in-intf mgmt0 protocol tcp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 1707 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 1707 in-intf mgmt0 protocol tcp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 3786 in-intf lo protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 33000 in-intf lo protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete in-intf mgmt0 protocol icmp • ip filter chain input rule append tail target accept dup-delete source-port 5353 dest-port 5353 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target reject-with icmp-port-unreachable dup-delete dest-port 33434-33523 in-intf mgmt0 protocol udp • ip filter chain input rule append tail target accept dup-delete dest-port 123 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 514 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete dest-port 67 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dup-delete comment "Feature HA port" dest-port 60102 in-intf mgmt0 protocol tcp conntrack new,established • ip filter chain input rule append tail target accept dest-port 636 in-intf mgmt0 protocol udp conntrack new,established • ip filter chain input rule append tail target accept dest-port 636 in-intf mgmt0 protocol tcp conntrack new,established • ip filter chain input rule append tail target logging dup-delete in-intf mgmt0
Output Chain Rules
• ip filter chain output rule append tail target drop out-intf mgmt0 state invalid • ip filter chain output rule append tail target accept out-intf mgmt0
Logging Chain Rules
• ip filter chain logging rule append tail target nflog in-intf mgmt0 rate-limit 1/minute logging-options prefix "IPTables-Dropped-<Domain>: " logging-options group 3

Prerouting-Mangle Chain Rules
• ip filter chain logging rule append tail target drop in-intf mgmt0

IPv6 Firewall Default Rules

Prerouting-Mangle Chain Rules
• ipv6 filter chain prerouting-mangle rule append tail target drop dup-delete not-dest-port 22-23 in-intf mgmt0 protocol tcp conntrack new tcp-op-mss mss-not-in-range 1280:65535
Input Chain Rules
• ipv6 filter chain input rule append tail target accept dup-delete in-intf lo • ipv6 filter chain input rule append tail target drop dup-delete dest-addr ::1 /128 in-intf mgmt0 • ipv6 filter chain input rule append tail target accept dup-delete in-intf mgmt0 state established,related • ipv6 filter chain input rule append tail target drop dup-delete not-dest-port 22-23 in-intf mgmt0 protocol tcp state new tcp-op syn match-not-syn • ipv6 filter chain input rule append tail target drop dup-delete in-intf mgmt0 fragment enable • ipv6 filter chain input rule append tail target drop dup-delete in-intf mgmt0 protocol tcp tcp-op flags all • ipv6 filter chain input rule append tail target drop dup-delete in-intf mgmt0 protocol tcp tcp-op flags none • ipv6 filter chain input rule append tail target drop dup-delete in-intf mgmt0 state invalid • ipv6 filter chain input rule append tail target drop dup-delete in-intf mgmt0 protocol tcp tcp-op flags reset rate-limit-above 2/second burst-limit-above 2 • ipv6 filter chain input rule append tail target drop dup-delete in-intf mgmt0 protocol tcp state new rate-limit-above 50/second burst-limit-above 50 • ipv6 filter chain input rule append tail target drop dup-delete in-intf mgmt0 protocol tcp conntrack new rate-limit-above 60/second burst-limit-above 20 • ipv6 filter chain input rule append tail target drop dup-delete in-intf mgmt0 recent name portscan recent rcheck-sec 86400 • ipv6 filter chain input rule append tail target none dup-delete in-intf mgmt0 recent name portscan recent remove • ipv6 filter chain input rule append tail target none dup-delete dest-port 22 in-intf mgmt0 protocol tcp conntrack new recent set • ipv6 filter chain input rule append tail target drop dup-delete dest-port 22 in-intf mgmt0 protocol tcp conntrack new recent update-sec 60 recent hitcount 50

Prerouting-Mangle Chain Rules
• ipv6 filter chain input rule append tail target none dup-delete dest-port 23 in-intf mgmt0 protocol tcp conntrack new recent set • ipv6 filter chain input rule append tail target drop dup-delete dest-port 23 in-intf mgmt0 protocol tcp conntrack new recent update-sec 60 recent hitcount 50 • ipv6 filter chain input rule append tail target none dup-delete dest-port 443 in-intf mgmt0 protocol tcp conntrack new recent set • ipv6 filter chain input rule append tail target drop dup-delete dest-port 443 in-intf mgmt0 protocol tcp conntrack new recent update-sec 60 recent hitcount 150 • ipv6 filter chain input rule append tail target none dup-delete dest-port 80 in-intf mgmt0 protocol tcp conntrack new recent set • ipv6 filter chain input rule append tail target drop dup-delete dest-port 80 in-intf mgmt0 protocol tcp conntrack new recent update-sec 60 recent hitcount 150 • ipv6 filter chain input rule append tail target none dup-delete dest-port 23108 in-intf mgmt0 protocol tcp conntrack new recent set • ipv6 filter chain input rule append tail target drop dup-delete dest-port 23108 in-intf mgmt0 protocol tcp conntrack new recent update-sec 60 recent hitcount 150 • ipv6 filter chain input rule append tail target none dup-delete dest-port 161 in-intf mgmt0 protocol udp conntrack new recent set • ipv6 filter chain input rule append tail target drop dup-delete dest-port 161 in-intf mgmt0 protocol udp conntrack new recent update-sec 60 recent hitcount 100 • ipv6 filter chain input rule append tail target drop dup-delete in-intf mgmt0 routing-header-type 0 • ipv6 filter chain input rule append tail target drop dup-delete in-intf mgmt0 hop-by-hop-header enable • ipv6 filter chain input rule append tail target accept dup-delete dest-port 22 in-intf mgmt0 protocol tcp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 23 in-intf mgmt0 protocol tcp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 443 in-intf mgmt0 protocol tcp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 80 in-intf mgmt0 protocol tcp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 23108 in-intf mgmt0 protocol tcp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 179 in-intf mgmt0 protocol tcp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 547 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 122 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 161 in-intf mgmt0 protocol udp conntrack new,established

Prerouting-Mangle Chain Rules
• ipv6 filter chain input rule append tail target accept dup-delete dest-port 6306 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 69 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 389 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 389 in-intf mgmt0 protocol tcp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 1812-1813 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 49 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 49 in-intf mgmt0 protocol tcp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete source-port 53 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete source-port 53 in-intf mgmt0 protocol tcp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 500 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 4500 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 1293 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 1293 in-intf mgmt0 protocol tcp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 1707 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 1707 in-intf mgmt0 protocol tcp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 3786 in-intf lo protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 33000 in-intf lo protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete source-port 5353 dest-port 5353 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target reject-with icmp6-port-unreachable dup-delete dest-port 33434-33523 in-intf mgmt0 protocol udp • ipv6 filter chain input rule append tail target accept dup-delete dest-port 123 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 514 in-intf mgmt0 protocol udp conntrack new,established

Prerouting-Mangle Chain Rules
• ipv6 filter chain input rule append tail target accept dup-delete dest-port 546 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete comment "Feature HA port" dest-port 60102 in-intf mgmt0 protocol tcp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 636 in-intf mgmt0 protocol udp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete dest-port 636 in-intf mgmt0 protocol tcp conntrack new,established • ipv6 filter chain input rule append tail target accept dup-delete in-intf mgmt0 protocol-icmp-type destination-unreachable • ipv6 filter chain input rule append tail target accept dup-delete in-intf mgmt0 protocol-icmp-type packet-too-big • ipv6 filter chain input rule append tail target accept dup-delete in-intf mgmt0 protocol-icmp-type time-exceeded • ipv6 filter chain input rule append tail target accept dup-delete in-intf mgmt0 protocol-icmp-type parameter-problem • ipv6 filter chain input rule append tail target accept dup-delete in-intf mgmt0 protocol-icmp-type echo-request • ipv6 filter chain input rule append tail target accept dup-delete in-intf mgmt0 protocol-icmp-type echo-reply • ipv6 filter chain input rule append tail target accept dup-delete in-intf mgmt0 protocol-icmp-type router-advertisement • ipv6 filter chain input rule append tail target accept dup-delete in-intf mgmt0 protocol-icmp-type neighbor-solicitation • ipv6 filter chain input rule append tail target accept dup-delete in-intf mgmt0 protocol-icmp-type neighbor-advertisement • ipv6 filter chain input rule append tail target logging dup-delete in-intf mgmt0
Output Chain Rules
• ipv6 filter chain output rule append tail target accept dup-delete out-intf mgmt0
Logging Chain Rules
• ipv6 filter chain logging rule append tail target nflog dup-delete in-intf mgmt0 logging-options prefix IP6Tables-Dropped: logging-options group 3 • ipv6 filter chain logging rule append tail target drop dup-delete in-intf mgmt0

Control Plane Policing Commands

ip filter enable | ipv6 filter enable

	{ip | ipv6} filter enable no {ip | ipv6} filter enable Enables IP filtering. The no form of the command disables IP filtering.
Syntax Description	N/A
Default	ip Enabled ip6 Disabled
Configuration Mode	config
History	3.5.1000 3.10.3000 IP Filter is enabled by default
Example	switch (config) # ip filter enable
Related Commands	
Notes	It is recommended to run this command only after configuring all of the IP table filter parameters.

ip filter chain policy | ipv6 filter chain policy

	{ip | ipv6} filter chain <chain_name> policy {accept | drop} no {ip | ipv6} filter chain <chain_name> policy Configures default policy for a specific chain (if no rule matches this default policy action shall apply). The no form of the command resets default policy for a specific chain.	
Syntax Description	chain_name	Selects a chain for which to add or modify a filter: • input – input chain or ingress interfaces • output – output chain or egress interfaces
	accept	Accepts all traffic by default for this chain
	drop	Drops all traffic by default for this chain

Default	Accept for input and output chains
Configuration Mode	config
History	3.5.1000
Example	switch (config) # ipv6 filter chain input policy accept
Related Commands	
Notes	

ip filter chain rule target | ipv6 filter chain rule target

	no {ip | ipv6} filter chain <chain_name> rule {<number> | all} Inserts rule before specified rule number. The no form of the command deletes rule for a specific chain.	
Syntax Description	chain_name	A chain to which to add or modify a filter: • input – input chain or ingress interfaces • output – output chain or egress interfaces
	rule	• append tail – appends operation to the bottom of operation list • insert <oper_num> – inserts operation at specified position (existing operation at that position moves back in the list) • modify <oper_num> – modifies existing operation at specified position. Only the parameters specified in this invocation are altered; everything else is left untouched. • move <oper_num1> to <oper_num2> – moves one operation to another place in the operation list • set <oper_num> – sets operation at specified position (overwrites existing)
	target	• accept – allows the packets that match the rule into the management plane • drop – drops packets that match the rule • rate-limit – allows with rate limiting in packets per sec (PPS) • reject-with – drops the packet and replies with an ICMP error message

	param	• rate-limit /[second | minute | hour] - matches is traffic is less than the set limit • rate-limit-above /[second | minute | hour] - matches is traffic is more than the set limit • burst-limit - Maximum initial number of packets to match when setting rate-limit. The default is 5. • burst-limit-above - Maximum initial number of packets to match when setting rate-limit-above. The default is 5. • comment <text> – specifies description string for this rule (60 chars max) • dest-addr <ip> – IP matching a specific destination address or address range. A specific IPv4 address can be provided or an entire subnet by giving an address along with netmask in dot notation or as a CIDR notation (e.g. /24). • not-dest-addr <ip> – IP not matching a specific destination address range • dest-port <port(s)> – matching a specific destination port or port range • not-dest-port <port(s)> – port not matching a specific destination port or port range • dup-delete – deletes any preexisting duplicates of this rule • in-intf – interface matching a specific inbound interface • not-in-intf <if_name> – interface not matching a specific inbound interface • out-intf <if_name> – matches a specific outbound interface • not-out-intf <if_name> – interface not matching a specific outbound interface

	param4 (cont.)	• protocol <if_name> – matches a specific protocol
		  ◦ tcp
		  ◦ udp
		  ◦ icmp
		  ◦ all
		• not-protocol <protocol> – does not match a specific protocol
		  ◦ tcp
		  ◦ udp
		  ◦ icmp
		  ◦ all
		• source-addr <ip> – matches a specific source address range

		• not-source-addr <ip> – does not match a specific source address range • source-port <port(s)> – matches a specific source port or port range • not-source-port <port(s)> – does not match a specific source port or port range • state – matches packets in a particular state. Possible values: ◦ established – packet associated with an established connection which has seen traffic in both directions ◦ related – packet that starts a new connection but is related to an existing connection ◦ new – packet that starts a new, unrelated connection ◦ A combination can be entered separated by commas
	param (cont.)	• routing-header-type <type> – matches IPv6 routing header type. (only supported for ipv6) • hop-by-hop-header enable - matches IPv6 packet with hop by hop header. (only supported for ipv6)
Default	N/A	
Configuration Mode	config	
History	3.5.1000	
Example	switch (config) # ipv6 filter enable chain input rule append tail target drop state related protocol all dup-delete	
Related Commands		
Notes	• The source and destination ports may each be either a single number, or a range specified as “<low>-<high>”. For example: “10-20” would specify ports 10 through 20 (inclusive). • The port parameter only works in conjunction with TCP and UDP • Setting a “positive” rule removes any corresponding “not-” rules, and vice-versa • The “state” parameter is a classification of the packet relative to existing connections	

	• If TCP or UDP are selected for the “protocol” parameter, source and/or destination ports may be specified. If ICMP is selected, these options are either ignored, or an error is produced.

ip filter options include-bridges

	{ip | ipv6} filter options include-bridges no {ip | ipv6} filter options include-bridges Applies IP filters to bridges
Syntax Description	N/A
Default	Disabled
Configuration Mode	config
History	3.5.1000
Example	switch (config) # ip filter options include-bridges
Related Commands	
Notes	

ip filter reset-to-default-rules | ipv6 filter reset-to-default-rules

	{ip | ipv6} filter reset-to-default-rules Deletes all configured IP filter rules and add the default rules defined in the user manual under section "IP Table Filtering Default Rules", above.
Syntax Description	N/A
Default	N/A
Configuration Mode	config

History	3.10.3000 3.11.4002: Added support for support IPv6
Example	switch (config) # ip filter reset-to-default-rules switch (config) # ipv6 filter reset-to-default-rules
Related Commands	
Notes	

show ip filter

	show ip filter Displays IPv4 filtering state.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000

Example	switch (config) # show ip filter
	Packet filtering for IPv4: enabled
	Active IPv4 filtering rules (omitting any not from configuration):
	  Chain 'input' Policy 'accept':
	  Rule 1:
	  Target : accept
	  Protocol : all
	  Source : all
	  Destination : 1.1.1.0/24
	  Interface : all
	  State : any
	  Other Filter: -
	Chain 'output' Policy 'accept':
	  Rule 1:
	  Target : reject-with icmp-net-unreachable
	  Protocol : tcp
	  Source : all
	  Destination : all
	  Interface : all
	  State : any
	  Other Filter: dest-port 1000

Related Commands	
Notes	

show ip filter all

	show ip filter all Displays IPv4 filtering state (including un-configured rules).
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000
Example	switch (config) # show ip filter all Destination : 1.1.1.0/24 Interface : all State : any Other Filter: - Chain 'output' Policy 'accept': Rule 1: Target : reject-with icmp-net-unreachable Protocol : tcp Source : all Destination : all Interface : all State : any Other Filter: dest-port 1000
Related Commands	
Notes	

show ip filter configured

	show ip filter configured Displays IPv4 filtering configuration.

Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000
Example	switch (config) # show ip filter configured Packet filtering for IPv4: enabled IPv4 configuration: Chain 'input' Policy 'accept': Rule 1: Target : accept Protocol : all Source : all Destination : 1.1.1.0/24 Interface : all State : any Other Filter: - Chain 'output' Policy 'accept': Rule 1: Target : reject-with icmp-net-unreachable Protocol : tcp Source : all Destination : all Interface : all State : any Other Filter: dest-port 1000
Related Commands	
Notes	

show ipv6 filter

	show ipv6 filter Displays IPv6 filtering state.
Syntax Description	N/A
Default	N/A
Configuration	config

Mode	
History	3.6.6000

Example	switch (config) # show ipv6 filter
	Packet filtering for IPv6: enables
	Active IPv6 filtering rules (omitting any not from configuration):
	  Chain 'input' Policy 'accept':
	  Rule 1:
	  Target : accept
	  Protocol : all
	  Source : all
	  Destination : 1.1.1.0/24
	  Interface : all
	  State : any
	  Other Filter: -
	  Chain 'output' Policy 'accept':
	  Rule 1:
	  Target : reject-with icmp-net-unreachable
	  Protocol : tcp
	  Source : all
	  Destination : all
	  Interface : all
	  State : any
	  Other Filter: dest-port 1000

Related Commands	
Notes	

show ipv6 filter all

	show ipv6 filter all Displays IPv6 filtering state (including un-configured rules).
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000
Example	switch (config) # show ipv6 filter all Packet filtering for IPv6: enables All active IPv6 filtering rules:

	Chain 'input' Policy 'accept': Rule 1: Target : accept Protocol : all Source : all Destination : 1.1.1.0/24 Interface : all State : any Other Filter: - Chain 'output' Policy 'accept': Rule 1: Target : reject-with icmp-net-unreachable Protocol : tcp Source : all Destination : all Interface : all State : any Other Filter: dest-port 1000
Related Commands	
Notes	

show ipv6 filter configured

	show ipv6 filter configured Displays IPv6 filtering configuration.
Syntax Description	N/A
Default	N/A
Configuration Mode	config
History	3.6.6000
Example	switch (config) # show ipv6 filter configured Packet filtering for IPv6: enables IPv6 configuration: Chain 'input' Policy 'accept': Rule 1: Target : accept Protocol : all Source : all Destination : 1.1.1.0/24 Interface : all

	State : any Other Filter: - Chain 'output' Policy 'accept': Rule 1: Target : reject-with icmp-net-unreachable Protocol : tcp Source : all Destination : all Interface : all State : any Other Filter: dest-port 1000
Related Commands	
Notes	

LLDP Over Managament Interface

management-lldp

management-lldp enable

New command	management-lldp enable no management-lldp enable Enables lldp on management interfaces. The no form of the command disables lldp on management interfaces.
Syntax Description	Enable lldp on management interfaces.
Default	Disabled.
Configuration Mode	Configure terminal.
History	
Role	Admin
Examples	switch (config) # management-lldp enable switch (config) # no management-lldp enable
Related Commands	show management-lldp status
Notes	Enables/disables RX and TX of lldp packets on management interface.

show management-lldp status

New command	show management-lldp status
Syntax Description	Show lldp over management interfaces status

Default	N/A
Configuration Mode	Any command mode
History	
Role	Admin
Examples	switch (config) # show management-lldp status lldp status: enabled
Related Commands	management-lldp enable
Notes	

show management-lldp neighbors

New command	show management-lldp neighbors
Syntax Description	Show lldp neighbors over management interfaces
Default	N/A
Configuration Mode	Any command mode
History	
Examples	switch (config) # show management-lldp neighbors --- Interface Device ID Port ID System Name Capability TTL --- mgmt0 7C:FE:90:65:DE:A8 16 switch22 Bridge 120
Related Commands	management-lldp enable
Notes	• To view active lldp neighbors, lldp must be enabled using the cli "management-lldp enable" cmd. • The switch will not show the lldp neighbor when lldp is disabled on that switch.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation (“NVIDIA”) makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer’s own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, “MATERIALS”) ARE BEING PROVIDED “AS IS.” NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF

ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA and the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.