



Configuring Unicast RPF

This chapter describes how to configure rate limits for egress traffic on Cisco NX-OS devices.

This chapter includes the following sections:

- [Information About Unicast RPF, page 1](#)
- [Licensing Requirements for Unicast RPF, page 3](#)
- [Guidelines and Limitations for Unicast RPF, page 3](#)
- [Default Settings for Unicast RPF, page 4](#)
- [Configuring Unicast RPF, page 4](#)
- [Configuration Examples for Unicast RPF, page 6](#)
- [Verifying the Unicast RPF Configuration, page 6](#)

Information About Unicast RPF

The Unicast RPF feature reduces problems that are caused by the introduction of malformed or forged (spoofed) IPv4 source addresses into a network by discarding IPv4 packets that lack a verifiable IP source address. For example, a number of common types of Denial-of-Service (DoS) attacks, including Smurf and Tribal Flood Network (TFN) attacks, can take advantage of forged or rapidly changing source IPv4 or IPv6 addresses to allow attackers to thwart efforts to locate or filter the attacks. Unicast RPF deflects attacks by forwarding only the packets that have source addresses that are valid and consistent with the IP routing table.

When you enable Unicast RPF on an interface, the switch examines all ingress packets received on that interface to ensure that the source address and source interface appear in the routing table and match the interface on which the packet was received. This examination of source addresses relies on the Forwarding Information Base (FIB).



Note

Unicast RPF is an ingress function and is applied only on the ingress interface of a switch at the upstream end of a connection.

Unicast RPF verifies that any packet received at a switch interface arrives on the best return path (return route) to the source of the packet by doing a reverse lookup in the FIB. If the packet was received from one of the best reverse path routes, the packet is forwarded as normal. If there is no reverse path route on the same

interface from which the packet was received, the source address might have been modified by the attacker. If Unicast RPF does not find a reverse path for the packet, the packet is dropped.

**Note**

With Unicast RPF, all equal-cost “best” return paths are considered valid, which means that Unicast RPF works where multiple return paths exist, if each path is equal to the others in terms of the routing cost (number of hops, weights, and so on) and as long as the route is in the FIB. Unicast RPF also functions where Enhanced Interior Gateway Routing Protocol (EIGRP) variants are being used and unequal candidate paths back to the source IP address exist.

Unicast RPF Process

Unicast RPF has several key implementation principles:

- The packet must be received at an interface that has the best return path (route) to the packet source (a process called *symmetric routing*). There must be a route in the FIB that matches the route to the receiving interface. Static routes, network statements, and dynamic routing add routes to the FIB.
- IP source addresses at the receiving interface must match the routing entry for the interface.
- Unicast RPF is an input function and is applied only on the input interface of a device at the upstream end of a connection.

You can use Unicast RPF for downstream networks, even if the downstream network has other connections to the Internet.

**Caution**

Be careful when using optional BGP attributes, such as weight and local preference, because an attacker can modify the best path back to the source address. Modification would affect the operation of Unicast RPF.

When a packet is received at the interface where you have configured Unicast RPF and ACLs, the Cisco NX-OS software performs the following actions:

Procedure

- Step 1** Checks the input ACLs on the inbound interface.
- Step 2** Uses Unicast RPF to verify that the packet has arrived on the best return path to the source, which it does by doing a reverse lookup in the FIB table.
- Step 3** Conducts a FIB lookup for packet forwarding.
- Step 4** Checks the output ACLs on the outbound interface.
- Step 5** Forwards the packet.

Global Statistics

Each time the Cisco NX-OS device drops a packet at an interface due to a failed unicast RPF check, that information is counted globally on the device on a per-forwarding engine (FE) basis. Global statistics on dropped packets provide information about potential attacks on the network, but they do not specify which interface is the source of the attack. Per-interface statistics on packets dropped due to a failed unicast RPF check are not available.

Licensing Requirements for Unicast RPF

Product	License Requirement
Cisco NX-OS	Unicast RPF requires no license. Any feature not included in a license package is bundled with the Cisco NX-OS system images and is provided at no extra charge to you. For an explanation of the Cisco NX-OS licensing scheme, see the <i>Cisco NX-OS Licensing Guide</i> .

Guidelines and Limitations for Unicast RPF

Unicast RPF has the following configuration guidelines and limitations:

- You must apply Unicast RPF at the interface downstream from the larger portion of the network, preferably at the edges of your network.
- The further downstream that you apply Unicast RPF, the finer the granularity you have in mitigating address spoofing and in identifying the sources of spoofed addresses. For example, applying Unicast RPF on an aggregation device helps to mitigate attacks from many downstream networks or clients and is simple to administer, but it does not help identify the source of the attack. Applying Unicast RPF at the network access server helps limit the scope of the attack and trace the source of the attack; however, deploying Unicast RPF across many sites does add to the administration cost of operating the network.
- The more entities that deploy Unicast RPF across Internet, intranet, and extranet resources, means that the better the chances are of mitigating large-scale network disruptions throughout the Internet community, and the better the chances are of tracing the source of an attack.
- Unicast RPF will not inspect IP packets that are encapsulated in tunnels, such as generic routing encapsulation (GRE) tunnels. You must configure Unicast RPF at a home gateway so that Unicast RPF processes network traffic only after the tunneling and encryption layers have been stripped off the packets.
- You can use Unicast RPF in any “single-homed” environment where there is only one access point out of the network or one upstream connection. Networks that have one access point provide symmetric routing, which means that the interface where a packet enters the network is also the best return path to the source of the IP packet.

- Do not use Unicast RPF on interfaces that are internal to the network. Internal interfaces are likely to have routing asymmetry, which means that multiple routes to the source of a packet exist. You should configure Unicast RPF only where there is natural or configured symmetry.
- Unicast RPF allows packets with 0.0.0.0 source and 255.255.255.255 destination to pass so that the Bootstrap Protocol (BOOTP) and the Dynamic Host Configuration Protocol (DHCP) can operate correctly.

**Note**

If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Default Settings for Unicast RPF

This table lists the default settings for Unicast RPF parameters.

Table 1: Default Unicast RPF Parameter Settings

Parameters	Default
Unicast RPF	Disabled

Configuring Unicast RPF

You can configure one the following Unicast RPF modes on an ingress interface:

Strict Unicast RPF mode

A strict mode check is successful when Unicast RPF finds a match in the FIB for the packet source address and the ingress interface through which the packet is received matches one of the Unicast RPF interfaces in the FIB match. If this check fails, the packet is discarded. You can use this type of Unicast RPF check where packet flows are expected to be symmetrical.

Loose Unicast RPF mode

A loose mode check is successful when a lookup of a packet source address in the FIB returns a match and the FIB result indicates that the source is reachable through at least one real interface. The ingress interface through which the packet is received is not required to match any of the interfaces in the FIB result.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.

	Command or Action	Purpose
Step 2	interface ethernet <i>slot/port</i> Example: <pre>switch(config)# interface ethernet 2/3 switch(config-if)#</pre>	Specifies an Ethernet interface and enters interface configuration mode.
Step 3	ip verify unicast source reachable-via {any [allow-default] rx} Example: <pre>switch(config-if)# ip verify unicast source reachable-via any</pre>	<p>Configures Unicast RPF on the interface for IPv4.</p> <p>The any keyword specifies loose Unicast RPF.</p> <p>If you specify the allow-default keyword, the source address lookup can match the default route and use that for verification.</p> <p>The rx keyword specifies strict Unicast RPF.</p>
Step 4	ipv6 verify unicast source reachable-via {any [allow-default] rx} Example: <pre>switch(config-if)# ipv6 verify unicast source reachable-via any</pre>	<p>Configures Unicast RPF on the interface for IPv6.</p> <p>The any keyword specifies loose Unicast RPF.</p> <p>If you specify the allow-default keyword, the source address lookup can match the default route and use that for verification.</p> <p>The rx keyword specifies strict Unicast RPF.</p>
Step 5	exit Example: <pre>switch(config-cmap)# exit switch(config)#</pre>	Exits class map configuration mode.
Step 6	show ip interface ethernet <i>slot/port</i> Example: <pre>switch(config)# show ip interface ethernet 2/3</pre>	<p>(Optional)</p> <p>Displays the IP information for an interface.</p>
Step 7	show running-config interface ethernet <i>slot/port</i> Example: <pre>switch(config)# show running-config interface ethernet 2/3</pre>	<p>(Optional)</p> <p>Displays the configuration for an interface in the running configuration.</p>
Step 8	copy running-config startup-config Example: <pre>switch(config)# copy running-config startup-config</pre>	<p>(Optional)</p> <p>Copies the running configuration to the startup configuration.</p>

Configuration Examples for Unicast RPF

The following example shows how to configure loose Unicast RPF for IPv4 packets:

```
interface Ethernet2/3
 ip address 172.23.231.240/23
 ip verify unicast source reachable-via any
```

The following example shows how to configure strict Unicast RPF for IPv4 packets:

```
interface Ethernet2/2
 ip address 172.23.231.240/23
 ip verify unicast source reachable-via rx
```

The following example shows how to configure loose Unicast RPF for IPv6 packets:

```
interface Ethernet2/1
 ipv6 address 2001:0DB8:c18:1::3/64
 ipv6 verify unicast source reachable-via any
```

The following example shows how to configure strict Unicast RPF for IPv6 packets:

```
interface Ethernet2/4
 ipv6 address 2001:0DB8:c18:1::3/64
 ipv6 verify unicast source reachable-via rx
```

Verifying the Unicast RPF Configuration

To display Unicast RPF configuration information, perform one of the following tasks:

Command	Purpose
show running-config interface ethernet <i>slot/port</i>	Displays the interface configuration in the running configuration.
show running-config ip [all]	Displays the IPv4 configuration in the running configuration.
show startup-config interface ethernet <i>slot/port</i>	Displays the interface configuration in the startup configuration.
show startup-config ip	Displays the IP configuration in the startup configuration.