



## Configuring Unicast RPF

---

This chapter contains the following sections:

- [Information About Unicast RPF, on page 1](#)
- [Guidelines and Limitations for Unicast RPF, on page 2](#)
- [Default Settings for Unicast RPF, on page 4](#)
- [Configuring Unicast RPF, on page 4](#)
- [Configuration Examples for Unicast RPF, on page 6](#)
- [Verifying the Unicast RPF Configuration, on page 6](#)
- [Additional References for Unicast RPF, on 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 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).

Unicast RPF verifies that any packet received at a 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.

## Guidelines and Limitations for Unicast RPF

Unicast RPF (uRPF) has the following configuration guidelines and limitations:

- You must apply uRPF at the interface downstream from the larger portion of the network, preferably at the edges of your network.
- The further downstream that you apply uRPF, the finer the granularity you have in mitigating address spoofing and in identifying the sources of spoofed addresses. For example, applying uRPF 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 uRPF at the network access server helps limit the scope of the attack and trace the source of the attack; however, deploying uRPF across many sites does add to the administration cost of operating the network.
- The more entities that deploy uRPF 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.
- uRPF will not inspect IP packets that are encapsulated in tunnels, such as generic routing encapsulation (GRE) tunnels. You must configure uRPF at a home gateway so that uRPF processes network traffic only after the tunneling and encryption layers have been stripped off the packets.
- You can use uRPF 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 uRPF 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 uRPF only where there is natural or configured symmetry. Do not configure strict uRPF.
- uRPF 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.
- When uRPF is enabled, loose mode is applied for both IPv4 and IPv6. However, strict mode can be applied per protocol.
- For strict uRPF to work, you must enable it on both the ingress interface and the interface where the source IP address is learned.
- The switch hardware does not implement strict uRPF per the configured routing interface.
- Strict uRPF is implemented per learned route on strict uRPF-enabled interfaces.
- If a route is resolved as ECMP, strict uRPF will fall back to loose mode.
- Because of the hardware limitation on the trap resolution, uRPF might not be applied on supervisor-bound packets via inband.
- For IP traffic, both IPv4 and IPv6 configurations should be enabled simultaneously.
- Due to hardware limitations, the Cisco Nexus 3600 Series switches support only the following combinations:

| uRPF Configuration |         | Applied Traffic Check on Source IP Address |         |                     |
|--------------------|---------|--|---------|---------------------|
| IPv4               | IPv6    | IP Unipath                                 | IP ECMP | MPLS Encap/VPN/ECMP |
| Disable            | Disable | Allow                                      | Allow   | Allow               |

| uRPF Configuration |        | Applied Traffic Check on Source IP Address |            |            |
|--------------------|--------|--|------------|------------|
| Loose              | Loose  | uRPF loose                                 | uRPF loose | uRPF loose |
| Strict             | Strict | uRPF strict                                | uRPF loose | uRPF loose |

## 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 the Strict Unicast RPF mode or the Loose Unicast RPF mode on the ingress interface. For Strict Unicast mode, apply the configuration to interfaces where the source IP is attached. This allows you to configure the allowed list of specific sources.

### 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 isn't required to match any of the interfaces in the FIB result.

### Procedure

|               | Command or Action   | Purpose  |
|---------------|---|--|
| <b>Step 1</b> | <b>configure terminal</b><br><br><b>Example:</b><br><pre>switch# configure terminal switch(config)#</pre>                                 | Enters global configuration mode.  |
| <b>Step 2</b> | <b>interface ethernet <i>slot/port</i></b><br><br><b>Example:</b><br><pre>switch(config)# interface ethernet 2/3 switch(config-if)#</pre> | Specifies an ethernet interface and enters interface configuration mode. |

|               | Command or Action  | Purpose  |
|---------------|--|--|
| <b>Step 3</b> | <p><b>{ip   ipv6} verify unicast source reachable-via any</b></p> <p><b>Example:</b></p> <pre>switch(config-if)# ip verify unicast source reachable-via any</pre>      | <p>Configures unicast RPF on the interface for both IPv4 and IPv6.</p> <p><b>Note</b> Configure unicast RPF on each interface, since it's disabled by default. The configuration is shared across both IPv4 and IPv6. If you enable or disable on either IPv4 and IPv6, it affects all protocols on that interface.</p> <p><b>Note</b> When you enable uRPF for IPv4 or IPv6 (using the <b>ip</b> or <b>ipv6</b> keywords), unicast RPF is enabled for both IPv4 and IPv6.</p> <p><b>Note</b> You can configure only one version of the available IPv4 and IPv6 Unicast RPF command on an interface. When you configure one version, all the mode changes must be done by this version. The interface blocks all the other versions.</p> |
| <b>Step 4</b> | <p><b>exit</b></p> <p><b>Example:</b></p> <pre>switch(config-cmap)# exit switch(config)#</pre>   | Exits class map configuration mode.  |
| <b>Step 5</b> | <p>(Optional) <b>show ip interface ethernet slot/port</b></p> <p><b>Example:</b></p> <pre>switch(config)# show ip interface ethernet 2/3</pre>                         | Displays the IP information for an interface and verifies if the unicast RPF is enabled.   |
| <b>Step 6</b> | <p>(Optional) <b>show running-config interface ethernet slot/port</b></p> <p><b>Example:</b></p> <pre>switch(config)# show running-config interface ethernet 2/3</pre> | Displays the configuration for an interface in the running configuration.  |
| <b>Step 7</b> | <p>(Optional) <b>copy running-config startup-config</b></p> <p><b>Example:</b></p> <pre>switch(config)# copy running-config startup-config</pre>                       | Copies the running configuration to the startup configuration.   |

## Configuration Examples for Unicast RPF

The following examples show how to configure loose Unicast RPF for IPv4/IPv6 packets:

- ```
interface Ethernet2/3
ip address 172.23.231.240/23
ip verify unicast source reachable-via any
```
- ```
interface Ethernet2/3
ipv6 address 2001:0DB8:c18:1::3/64
ipv6 verify unicast source reachable-via any
```

The following examples show how to configure strict Unicast RPF for IPv4/IPv6 packets:

- ```
interface Ethernet2/2
ip address 172.23.231.240/23
ip verify unicast source reachable-via rx
```
- ```
interface Ethernet2/2
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  |
|--|--|
| <b>show running-config interface ethernet <i>slot/port</i></b> | Displays the interface configuration in the running configuration.                                   |
| <b>show running-config ip [all]</b>                            | Displays the IPv4 configuration in the running configuration.  |
| <b>show running-config ip6 [all]</b>                           | Displays the IPv6 configuration in the running configuration.  |
| <b>show startup-config interface ethernet <i>slot/port</i></b> | Displays the interface configuration in the startup configuration.                                   |
| <b>show ip interface ethernet <i>slot/port</i></b>             | Displays the IP information for an interface and verifies if the unicast RPF is enabled or disabled. |
| <b>show startup-config ip</b>                                  | Displays the IP configuration in the startup configuration.  |

## Additional References for Unicast RPF

This section includes additional information related to implementing unicast RPF.

**Related Documents**

| <b>Related Topic</b> | <b>Document Title</b>   |
|----------------------|---|
| MPLS VPN             | <a href="#">Cisco Nexus 3600 Series NX-OS Label Switching Configuration Guide</a> |

