Security Configuration Guide: Unicast Reverse Path Forwarding, Cisco IOS XE Everest 16.6

Americas Headquarters

CONTENTS

CHAPTER 1

Read Me First

Important Information about Cisco IOS XE 16

Effective Cisco IOS XE Release 3.7.0E for Catalyst Switching and Cisco IOS XE Release 3.17S (for Access and Edge Routing) the two releases evolve (merge) into a single version of converged release—the Cisco IOS XE 16—providing one release covering the extensive range of access and edge products in the Switching and Routing portfolio.

Feature Information

Use Cisco Feature Navigator to find information about feature support, platform support, and Cisco software image support. An account on Cisco.com is not required.

Related References

- Cisco IOS Command References, All Releases

Obtaining Documentation and Submitting a Service Request

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.

- To get the business impact you're looking for with the technologies that matter, visit Cisco Services.

- To submit a service request, visit Cisco Support.

- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco Marketplace.

- To obtain general networking, training, and certification titles, visit Cisco Press.

- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

CHAPTER 2

Unicast Reverse Path Forwarding Strict Mode

The Unicast Reverse Path Forwarding feature limits the malicious traffic on a network. This feature enables devices to verify the reachability of the source address in packets that are being forwarded and limit the appearance of spoofed or malformed addresses on a network. If the source IP address is not valid, Unicast Reverse Path Forwarding (RPF) discards the packet.

This module describes the Unicast Reverse Path Forwarding feature.

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see Bug Search Tool and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Unicast Reverse Path Forwarding

- Unicast Reverse Path Forwarding (RPF) requires Cisco Express Forwarding to function properly on a device.

- Prior to configuring Unicast RPF, you must configure the following access control lists (ACLs):

  - Configure standard or extended ACL to mitigate the transmission of invalid IP addresses (by performing egress filtering). Configuring standard or extended ACLs permit only valid source addresses to leave your network and enter the Internet.

• Configure standard or extended ACL entries to drop (deny) packets that have invalid source IP addresses (by performing ingress filtering). Invalid source IP addresses include the following types:

- Broadcast addresses (including multicast addresses)
- Loopback addresses
- Private addresses (RFC 1918, *Address Allocation for Private Internets*)
- Reserved addresses
- Source addresses that fall outside the range of valid addresses that are associated with the protected network

Restrictions for Unicast Reverse Path Forwarding

• Unicast RPF does not support access control list (ACL) templates.

The following basic restrictions apply to multihomed clients:

• Clients should not be multihomed on the same device because multihoming defeats the purpose of creating a redundant service for a client.

• Ensure that packets that flow up the link (out to the Internet) match the route advertised out of the link. Otherwise, Unicast RPF filters these packets as malformed packets.

Information About Unicast Reverse Path Forwarding

Overview of Unicast Reverse Path Forwarding

The Unicast Reverse Path Forwarding feature helps to mitigate problems that are caused by the introduction of malformed or forged (spoofed) IP source addresses into a network by discarding IP packets that lack verifiable IP source addresses. For example, a number of common types of denial-of-service (DoS) attacks, including Smurf and Tribal Flood Network (TFN), can take advantage of forged or rapidly changing source IP addresses to allow attackers to thwart efforts to locate or filter these attacks. For ISPs that provide public access, Unicast RPF deflects such attacks by forwarding only packets that have source addresses that are valid and consistent with the IP routing table, thereby protecting the network of the ISP, ISP customers, and the Internet.

Unicast RPF Operation

When Unicast RPF is enabled on an interface of a device, the device examines all packets received as input on that interface to ensure that the source address and source interface information appears in the routing table and matches the interface on which packets are received. This ability to "look backwards" is available only when Cisco Express Forwarding is enabled on a device because the lookup relies on the presence of a Forwarding Information Base (FIB). Cisco Express Forwarding generates a FIB as part of its operation.

Note Unicast RPF is an input function and is applied only on the input interface of a device at the upstream end of a connection.

Unicast RPF does a reverse lookup in the Cisco Express Forwarding table to check if any packet received at the interface of a device arrives on the best return path (or return route) to the source of the packet. If the packet was received from one of the best reverse path routes, the packet is forwarded as normal. No reverse path route on the interface from which the packet was received can mean that the source address was modified.

Note With Unicast RPF, all equal-cost "best" return paths are considered valid. Unicast RPF supports multiple return paths, provided that each path is equal to the others in terms of the routing cost (such as number of hops, weights, and so on) and the route is available in the FIB. Unicast RPF also functions where Enhanced Interior Gateway Routing Protocol (EIGRP) variants are used.

Before forwarding a packet that is received at the interface on which Unicast RPF and ACLs have been configured, Unicast RPF does the following checks:

1. If input ACLs are configured on the inbound interface.

2. If the packet has arrived on the best return path to the source by doing a reverse lookup in the FIB table.

3. Does a lookup of the Cisco Express Forwarding table for packet forwarding.

4. Checks output ACLs on the outbound interface.

5. Forwards the packet.

Access Control Lists and Logging

When you configure an access control list (ACL) and a packet fails the Unicast RPF check, the Unicast RPF checks the ACL to see if the packet should be dropped (by using a deny statement in the ACL) or forwarded (by using a permit statement in the ACL). Regardless of whether the packet is dropped or forwarded, the packet is counted in the global IP traffic statistics for Unicast RPF drops and in the interface statistics for Unicast RPF.

If no ACL is configured, the device drops the forged or malformed packet immediately, and no ACL logging occurs. The device and the interface Unicast RPF logging counters are updated.

To log Unicast RPF events, specify the logging option for ACL entries. Using the log information, administrators can view source addresses that are used in an attack, the time at which packets arrived at an interface, and so on.

Caution Logging requires CPU and memory resources. Logging Unicast RPF events for attacks that have a high rate of forged packets can degrade the performance of a device.

Per-Interface Statistics

Each time a packet is dropped or forwarded at an interface, that information is counted two ways: globally on the router and at each interface where you have applied Unicast RPF. Global statistics on dropped packets provide information about potential attacks on the network; however, these global statistics do not help to specify which interface is the source of the attack.

Per-interface statistics allow network administrators to track two types of information about malformed packets: Unicast RPF drops and Unicast RPF suppressed drops. Statistics on the number of packets that Unicast RPF drops help to identify the interface that is the entry point of the attack. The Unicast RPF drop count tracks the number of drops at the interface. The Unicast RPF suppressed drop count tracks the number of packets that failed the Unicast RPF check but were forwarded because of the permit permission set up in the ACL. Using the drop count and suppressed drop count statistics, a network administrator can takes steps to isolate the attack at a specific interface.

Note Judicious use of ACL logging can further identify the address or addresses that are being dropped by Unicast RPF.

The figure below illustrates how Unicast RPF and CEF work together to validate IP source addresses by verifying packet return paths. In this example, a customer has sent a packet having a source address of 192.168.1.1 from interface FDDI 2/0/0. Unicast RPF checks the FIB to see if 192.168.1.1 has a path to FDDI 2/0/0. If there is a matching path, the packet is forwarded. If there is no matching path, the packet is dropped.

Figure 1: Unicast RPF Validating IP Source Addresses
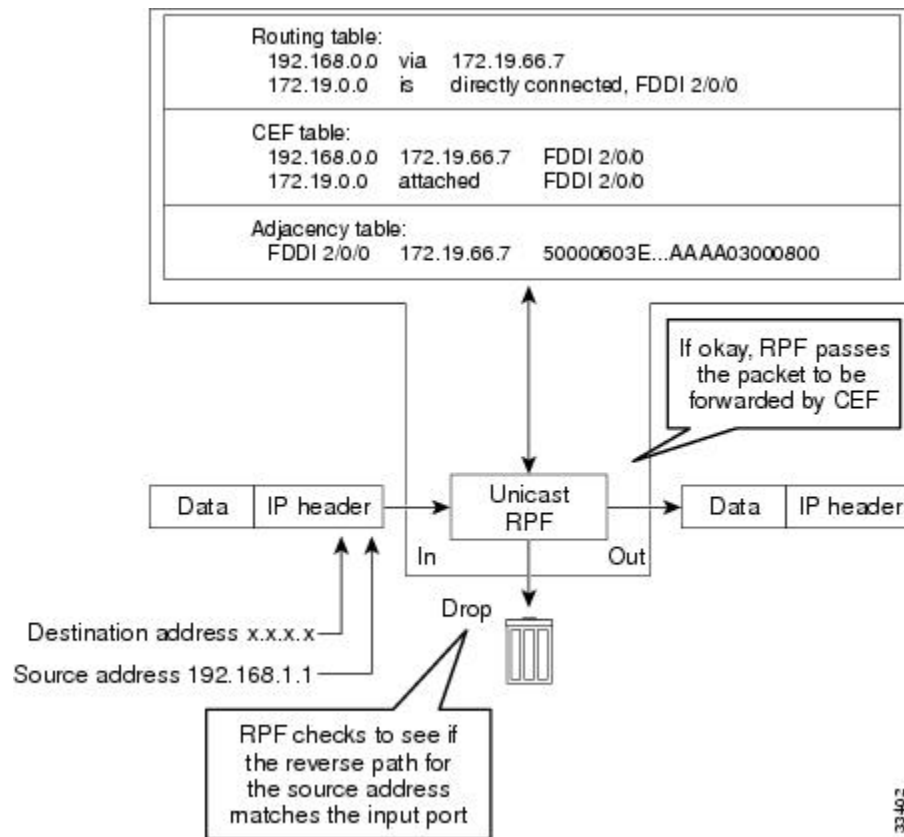
The figure below illustrates how Unicast RPF drops packets that fail validation. In this example, a customer has sent a packet having a source address of 209.165.200.225, which is received at interface FDDI 2/0/0. Unicast RPF checks the FIB to see if 209.165.200.225 has a return path to FDDI 2/0/0. If there is a matching path, the packet is forwarded. In this case, there is no reverse entry in the routing table that routes the customer packet back to source address 209.165.200.225 on interface FDDI 2/0/0, and so the packet is dropped.

Figure 2: Unicast RPF Dropping Packets That Fail Verification

Rules for Implementing Unicast RPF

The following rules apply when implementing Unicast Reverse Path Forwarding (RPF):

- Packets must be received at an interface that has the best return path (route) to the packets' source. This process is called symmetric routing. A route in the Forwarding Information Base (FIB) must match the route to the receiving interface. Add a route in the FIB through dynamic or static routing or by using a network statement.

- IP source addresses at the receiving interface must match the routing entry for the interface.

- Unicast RPF is an input function and can be applied at the input interface of a device at the upstream end of a connection.

Network administrators can use Unicast RPF for their customers and also for their downstream network or ISP, even if the downstream network or ISP has other connections to the Internet.

⚠

Caution Using optional Border Gateway Protocol (BGP) attributes, such as weight and local preference, the best path back to source addresses can be modified. The best path modification will affect the operation of Unicast RPF.

The following sections provides information about the implementation of Unicast RPF:

Security Policy and Unicast RPF

When determining how to deploy Unicast Reverse Path Forwarding (RPF), consider the following points:

- Apply Unicast RPF at the downstream interface, away from the larger portion of the network, preferably at the edges of your network. The further you apply Unicast RPF, the finer the granularity you have in mitigating address spoofing and in identifying 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 Unicast RPF does not help in identifying the source of the attack. Applying Unicast RPF at the network access server helps to limit the scope of the attack and trace the source of the attack. However, deploying Unicast RPF across many sites adds to the administration cost of operating a network.

- When you deploy Unicast RPF on many entities on a network (for example, across the Internet, intranet, and extranet resources), you have better chances of mitigating large-scale network disruptions throughout the Internet community, and of tracing the source of an attack.

- Unicast RPF does not inspect IP packets that are encapsulated in tunnels, such as the generic routing encapsulation (GRE), Layer 2 Tunneling Protocol (L2TP), or Point-to-Point Tunneling Protocol (PPTP). Configure Unicast RPF on a home gateway so that Unicast RPF processes network traffic only after tunneling and encryption layers are stripped off from the packets.

Ingress and Egress Filtering Policy for Unicast RPF

Unicast Reverse Path Forwarding (RPF) can be more effective at mitigating spoofing attacks when combined with a policy of ingress and egress filtering by using access control lists (ACLs).

Ingress filtering applies filters to traffic that is received at a network interface from either internal or external networks. With ingress filtering, packets that arrive from other networks or the Internet and that have a source address that matches a local network or private or broadcast addresses are dropped. For example, in ISP environments, ingress filtering can be applied to traffic that is received at a device from either a client (customer) or the Internet.

Egress filtering applies filters to the traffic that exits a network interface (the sending interface). By filtering packets on devices that connect your network to the Internet or to other networks, you can permit only packets with valid source IP addresses to leave your network.

For more information on network filtering, refer to RFC 2267, *Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing*.

Where to Use Unicast RPF

Unicast Reverse Path Forwarding (RPF) can be used in any "single-homed" environment where there is essentially only one access point out of the network, which means that there is only one upstream connection to the network. Networks having one access point offer the best example of 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. Unicast RPF is best used at the network perimeter for Internet, intranet, or extranet environments, or in ISP environments for customer network terminations.

The following sections describe two sample network environments in which Unicast RPF is implemented:

Enterprise Networks with a Single Connection to an ISP

In enterprise networks, you can use Unicast Reverse Path Forwarding (RPF) to filter traffic at the input interface (a process called ingress filtering) to protect from malformed packets that arrive from the Internet.

Traditionally, local networks that have one connection to the Internet use access control lists (ACLs) at the receiving interface to prevent spoofed packets from entering their local network.

ACLs work well for single-homed customers. However, when ACLs are used as ingress filters, the following two commonly referenced limitations apply:

- Packet-per-second (PPS) performance at very high packet rates

- ACL maintenance (whenever there are new addresses added to the network)

Unicast RPF addresses both the limitations described above. With Unicast RPF, ingress filtering is done at Cisco Express Forwarding PPS rates. Because Unicast RPF uses the Forwarding Information Base (FIB), ACL maintenance is not required, and thus, the administration overhead of traditional ACLs is reduced. The following figure and example demonstrate how Unicast RPF is configured for ingress filtering.

The figure below illustrates an enterprise network that has a single link to an upstream ISP. In this example, Unicast RPF is applied at GigabitEthernet interface 1/0/2 on the enterprise device for protection from malformed packets arriving from the Internet. Unicast RPF is also applied at GigabitEthernet interface 1/0/2 on the ISP device for protection from malformed packets arriving from the enterprise network.

Figure 3: Enterprise Network Using Unicast RPF for Ingress Filtering

A typical configuration on an ISP device that uses the topography in the figure above would be as follows:

```
ip cef
interface loopback 0
  description Loopback interface on Gateway Device 2
  ip address 192.168.3.1 255.255.255.255
  no ip redirects
  no ip directed-broadcast
  no ip proxy-arp
!
interface GigabitEthernet 1/0/2
  description 128K HDLC link to ExampleCorp WT50314E  R5-0
  bandwidth 128
  ip unnumbered loopback 0


  no ip redirects
  no ip directed-broadcast
  no ip proxy-arp
!
ip route 192.168.10.0 255.255.252.0 GigabitEthernet 1/0/2
```

The gateway device configuration of the enterprise network will be similar to the following:

```
ip cef
interface FastEthernet 0/0/0
 description ExampleCorp LAN
 ip address 192.168.10.1 255.255.252.0
 no ip redirects
 no ip directed-broadcast
 no ip proxy-arp
!
interface GigabitEthernet 1/0/2
```

```
    description 128K HDLC link to ExampleCorp Internet Inc WT50314E  C0
    bandwidth 128
    ip unnumbered FastEthernet 0/0/0


    no ip redirects
    no ip directed-broadcast
    no ip proxy-arp
    !
    ip route 0.0.0.0 0.0.0.0 GigabitEthernet 1/0/2
```

Notice that Unicast RPF works with a single default route. There are no additional routes or routing protocols. Network 192.168.10.0/22 is a connected network. Hence, packets coming from the Internet with a source address in the network 192.168.10.0/22 will be dropped by Unicast RPF.

Applying Unicast RPF to Network Access Servers

If a network access server supports Cisco Express Forwarding, Unicast RPF will work on that network. A network access server (NAS) allows users to access a network by checking the credentials of the users accessing the network. Aggregation devices support Unicast RPF with single-homed clients. Unicast RPF works well on leased lines or on a digital subscriber line (DSL), ISDN, or public switched telephone network (PSTN) customer connections that are connected to the Internet. Dialup connections are a big source of denial of service (Dos) attacks that use forged IP addresses.

Aggregation devices need routing prefixes information (IP address block) for routing traffic. In the topology described below, aggregation devices do not have a full Internet routing table, and as a result, Unicast RPF uses the information configured or redistributed by the Interior Gateway Protocol (IGP) or Internal Border Gateway Protocol (IBGP) (based on how customer routes are added to the network) to route traffic. Unicast RPF is applied upstream on the customer dialup connection device that is on the receiving (input) interfaces of ISP aggregation devices.

The figure below illustrates how Unicast RPF is applied to aggregation and access devices for an ISP or point of presence (PoP) with ISP devices providing dialup connections.

Figure 4: Unicast RPF Applied to PSTN/ISDN Customer Connections

Routing Table Requirements

Unicast Reverse Path Forwarding (RPF) uses the routing information in Cisco Express Forwarding tables for routing traffic. The amount of routing information that must be available in Cisco Express Forwarding tables depends on the device where Unicast RPF is configured and the functions the device performs in the network. For example, in an ISP environment where a device is a leased-line aggregation device for customers, the information about static routes that are redistributed into the Interior Gateway Protocol (IGP) or Internal Border Gateway Protocol (IBGP) (depending on which technique is used in the network) is required in the routing table. Because Unicast RPF is configured on customer interfaces, only minimal routing information is required. If a single-homed ISP configures Unicast RPF on the gateway to the Internet, the full Internet routing table information is required by Unicast RPF to help protect the ISP from external denial of service (DoS) attacks that use addresses that are not in the Internet routing table.

Where Not to Use Unicast RPF

Do not use Unicast Reverse Path Forwarding (RPF) on interfaces that are internal to a network. Internal interfaces are likely to have routing asymmetry (see the figure below), which means that there can be multiple routes to the source of a packet. Unicast RPF is applied only where there is a natural or configured symmetry.

For example, devices at the edge of an ISP network are more likely to have symmetrical reverse paths than devices that are in the core of an ISP network. The best forwarding path to forward packets from devices that are at the core of an ISP network may not be the best forwarding path that is selected for packets that are returned to the device.

We recommend that you do not apply Unicast RPF where there is a chance of asymmetric routing, unless you configure access control lists (ACLs) to allow the device to accept incoming packets. ACLs permit the use of Unicast RPF when packets arrive through specific, less-optimal asymmetric input paths.

The figure below illustrates how Unicast RPF can block legitimate traffic in an asymmetric routing environment.

Figure 5: Unicast RPF Blocking Legitimate Traffic in an Asymmetric Routing Environment

Unicast RPF with BOOTP and DHCP

Unicast RPF allows packets with 0.0.0.0 as the source IP address and 255.255.255.255 as the destination IP address to pass through a network to enable Bootstrap Protocol (BOOTP) and DHCP functions to work properly when Unicast RPF is configured.

How to Configure Unicast Reverse Path Forwarding

Configuring Unicast RPF

Before you begin

To use Unicast Reverse Path Forwarding, you must configure a device for Cisco Express Forwarding switching or distributed Cisco Express Forwarding switching. If Cisco Express Forwarding is not enabled globally on a device, Unicast RPF will not work on that device. If Cisco Express Forwarding is running on a device, individual interfaces on the device can be configured with other switching modes. Unicast RPF is an input-side function that is enabled on an interface or subinterface that supports any type of encapsulation, and Unicast RPF operates on IP packets that are received by the device.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip cef distributed**
4. **interface** *slot/subslot/port*
5. **exit**
6. **end**
7. **show cef interface** [*type number*]

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** Device> enable	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** Device# configure terminal	Enters global configuration mode.
Step 3	**ip cef distributed** **Example:** Device(config)# ip cef distributed	Enables Cisco Express Forwarding or distributed Cisco Express Forwarding on a device.
Step 4	**interface** *slot/subslot/port* **Example:** Device(config)# interface GigabitEthernet 0/0	Selects the input interface on which you want to apply Unicast Reverse Path Forwarding and enters interface configuration mode. • The interface that is configured is the receiving interface, which allows Unicast RPF to verify the best return path before forwarding a packet to the next destination.
Step 5	**exit** **Example:** Device(config-if)# exit	Exits interface configuration mode.
Step 6	**end** **Example:** Device(config-if)# end	Exits interface configuration mode and enters privileged EXEC mode.
Step 7	**show cef interface** [*type number*] **Example:** Device# show cef interface	Displays detailed Cisco Express Forwarding information for a specified interface or for all interfaces.

Example:

Troubleshooting Tips

HSRP Failure

The failure to disable Unicast RPF before disabling Cisco Express Forwarding can cause a Hot Standby Router Protocol (HSRP) failure. If you want to disable Cisco Express Forwarding on a device, you must first disable Unicast RPF.

Configuration Examples for Unicast Reverse Path Forwarding

Example: Configuring Unicast RPF

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Command List, All Releases
Unicast RPF command descriptions	• Cisco IOS Security Command Reference: Commands A to C • Cisco IOS Security Command Reference: Commands D to L • Cisco IOS Security Command Reference: Commands M to R • Cisco IOS Security Command Reference: Commands S to Z
Cisco Express Forwarding commands	Cisco IOS IP Switching Command Reference

Standards & RFCs

Standard/RFC	Title
RFC 1918	*Address Allocation for Private Internets*
RFC 2267	*Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing*

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for Unicast Reverse Path Forwarding

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for Unicast Reverse Path Forwarding

Feature Name	Releases	Feature Information
Unicast Reverse Path Forwarding	Cisco IOS XE Release 2.1	The Unicast Reverse Path Forwarding feature limits the malicious traffic on a network. This feature enables devices to verify the reachability of the source address in packets that are being forwarded and limit the appearance of spoofed or malformed addresses on a network. If the source IP address is not valid, Unicast Reverse Path Forwarding (RPF) discards the packet.

Unicast Reverse Path Forwarding ACL Support

The Unicast Reverse Path Forwarding feature helps to mitigate problems that are caused by malformed or forged IP source addresses that pass through a device. The Unicast Reverse Path Forwarding ACL Support feature adds the access control list (ACL) support to the Unicast Reverse Path Forwarding feature. With the ACL support, Unicast Reverse Path Forwarding (RPF) can determine whether to drop or to forward data packets that have malformed or forged IP source addresses.

This module describes the ACL support for Unicast RPF.

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see Bug Search Tool and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for Unicast Reverse Path Forwarding ACL Support

- Unicast RPF requires Cisco Express Forwarding to function properly on a device.

- Prior to configuring Unicast RPF, you must configure the following ACLs:

  - Configure standard or extended ACLs to mitigate the transmission of invalid IP addresses (by performing egress filtering). Configuring standard or extended ACLS, permit only valid source addresses to leave your network and enter the Internet.

- Configure standard or extended ACL entries to drop (deny) packets that have invalid source IP addresses (by performing ingress filtering). Invalid source IP addresses include the following types:

    - Broadcast addresses (including multicast addresses)
    - Loopback addresses
    - Private addresses (RFC 1918, *Address Allocation for Private Internets*)

    - Reserved addresses
    - Source addresses that fall outside the range of valid addresses associated with a protected network

- Configure standard or extended ACL entries to forward (permit) packets that fail the Unicast RPF checks and allow specific traffic from known asymmetric routed sources.

- Configure ACLs to track Unicast RPF events to provide additional information about network attacks.

Restrictions for Unicast Reverse Path Forwarding ACL Support

ACL templates are not supported.

Information About Unicast Reverse Path Forwarding ACL Support

Unicast RPF Operation

When Unicast RPF is enabled on an interface of a device, the device examines all packets received as input on that interface to ensure that the source address and source interface information appears in the routing table and matches the interface on which packets are received. This ability to "look backwards" is available only when Cisco Express Forwarding is enabled on a device because the lookup relies on the presence of a Forwarding Information Base (FIB). Cisco Express Forwarding generates a FIB as part of its operation.

Note   Unicast RPF is an input function and is applied only on the input interface of a device at the upstream end of a connection.

Unicast RPF does a reverse lookup in the Cisco Express Forwarding table to check if any packet received at the interface of a device arrives on the best return path (or return route) to the source of the packet. If the packet was received from one of the best reverse path routes, the packet is forwarded as normal. No reverse path route on the interface from which the packet was received can mean that the source address was modified.

![Note icon]

Note   With Unicast RPF, all equal-cost "best" return paths are considered valid. Unicast RPF supports multiple return paths, provided that each path is equal to the others in terms of the routing cost (such as number of hops, weights, and so on) and the route is available in the FIB. Unicast RPF also functions where Enhanced Interior Gateway Routing Protocol (EIGRP) variants are used.

Before forwarding a packet that is received at the interface on which Unicast RPF and ACLs have been configured, Unicast RPF does the following checks:

1. If input ACLs are configured on the inbound interface.

2. If the packet has arrived on the best return path to the source by doing a reverse lookup in the FIB table.

3. Does a lookup of the Cisco Express Forwarding table for packet forwarding.

4. Checks output ACLs on the outbound interface.

5. Forwards the packet.

Access Control Lists and Logging

When you configure an access control list (ACL) and a packet fails the Unicast RPF check, the Unicast RPF checks the ACL to see if the packet should be dropped (by using a deny statement in the ACL) or forwarded (by using a permit statement in the ACL). Regardless of whether the packet is dropped or forwarded, the packet is counted in the global IP traffic statistics for Unicast RPF drops and in the interface statistics for Unicast RPF.

If no ACL is configured, the device drops the forged or malformed packet immediately, and no ACL logging occurs. The device and the interface Unicast RPF logging counters are updated.

To log Unicast RPF events, specify the logging option for ACL entries. Using the log information, administrators can view source addresses that are used in an attack, the time at which packets arrived at an interface, and so on.

![Caution icon]

Caution   Logging requires CPU and memory resources. Logging Unicast RPF events for attacks that have a high rate of forged packets can degrade the performance of a device.

Per-Interface Statistics

Each time a packet is dropped or forwarded at an interface, that information is counted two ways: globally on the router and at each interface where you have applied Unicast RPF. Global statistics on dropped packets provide information about potential attacks on the network; however, these global statistics do not help to specify which interface is the source of the attack.

Per-interface statistics allow network administrators to track two types of information about malformed packets: Unicast RPF drops and Unicast RPF suppressed drops. Statistics on the number of packets that Unicast RPF drops help to identify the interface that is the entry point of the attack. The Unicast RPF drop count tracks the number of drops at the interface. The Unicast RPF suppressed drop count tracks the number of packets that failed the Unicast RPF check but were forwarded because of the permit permission set up in the ACL.

Using the drop count and suppressed drop count statistics, a network administrator can takes steps to isolate the attack at a specific interface.

✎

Note Judicious use of ACL logging can further identify the address or addresses that are being dropped by Unicast RPF.

The figure below illustrates how Unicast RPF and CEF work together to validate IP source addresses by verifying packet return paths. In this example, a customer has sent a packet having a source address of 192.168.1.1 from interface FDDI 2/0/0. Unicast RPF checks the FIB to see if 192.168.1.1 has a path to FDDI 2/0/0. If there is a matching path, the packet is forwarded. If there is no matching path, the packet is dropped.

Figure 6: Unicast RPF Validating IP Source Addresses
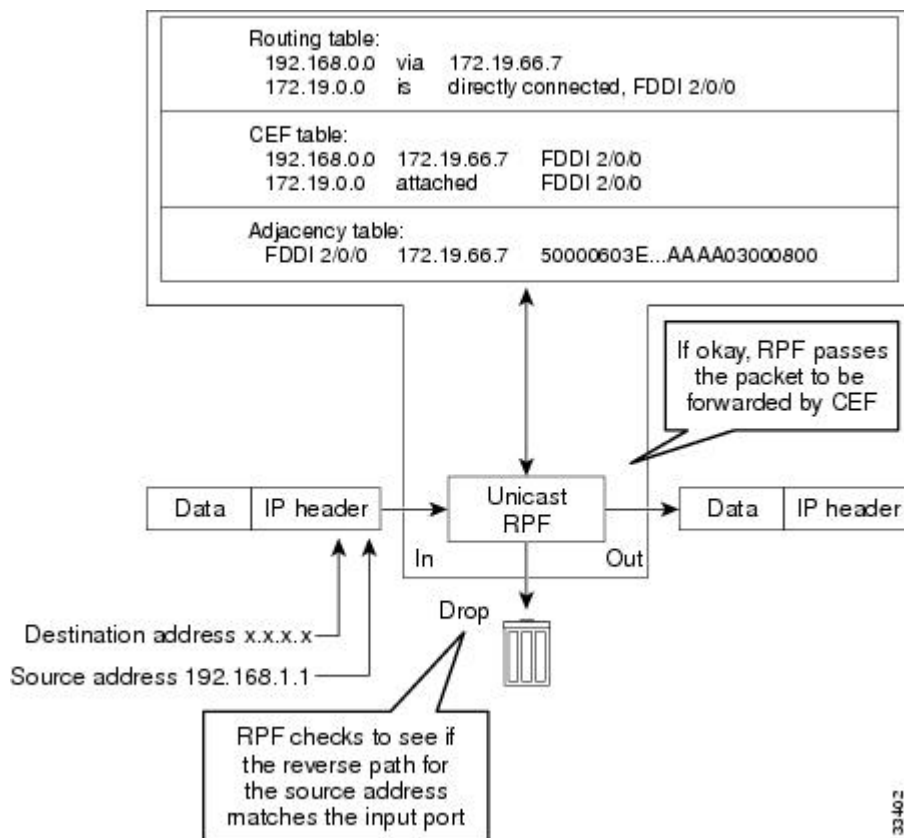
The figure below illustrates how Unicast RPF drops packets that fail validation. In this example, a customer has sent a packet having a source address of 209.165.200.225, which is received at interface FDDI 2/0/0. Unicast RPF checks the FIB to see if 209.165.200.225 has a return path to FDDI 2/0/0. If there is a matching path, the packet is forwarded. In this case, there is no reverse entry in the routing table that routes the customer packet back to source address 209.165.200.225 on interface FDDI 2/0/0, and so the packet is dropped.

Figure 7: Unicast RPF Dropping Packets That Fail Verification

How to Configure Unicast Reverse Path Forwarding ACL Support

Configuring Unicast RPF with ACL Support

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *type number*
4. **ipv6 address** *ipv6-address/prefix-length*
5. **ipv6 verify unicast source reachable-via** {**rx** | **any**} [*access-list*]
6. **end**
7. **show cef interface** [*type number*]
8. **show ipv6 traffic**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *type number* **Example:** `Device(config)# interface gigabitethernet 0/0/1`	Configures an interface and enters interface configuration mode.
Step 4	**ipv6 address** *ipv6-address/prefix-length* **Example:** `Device(config-if)# ipv6 address 2001:DB8:1::1/96`	Configures an IPv6 address based on an IPv6 general prefix and enables IPv6 processing on an interface.

Step 5	**ipv6 verify unicast source reachable-via** {**rx** | **any**} [*access-list*]
	
	Example:
	
	`Device(config-if)# ipv6 verify unicast source`
	`reachable-via any acl1`
		Verifies that a source address exists in the FIB table and enables Unicast RPF.

Step 6	**end** **Example:** `Device(config-if)# end`	Exits interface configuration mode and enters privileged EXEC mode.
Step 7	**show cef interface** [*type number*] **Example:** `Device# show cef interface gigabitethernet 0/0/1`	Displays detailed Cisco Express Forwarding information for a specified interface or for all interfaces.
Step 8	**show ipv6 traffic** **Example:** `Device# show ipv6 traffic`	Displays statistics about IPv6 traffic.

Example:

The following is sample output from the **show cef interface gigabitethernet 0/0/1** command:

```
Device# show cef interface gigabitethernet 0/0/1

 GigabitEthernet0/0/1 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::C67D:4FFF:FEB6:E410
  No Virtual link-local address(es):
  Global unicast address(es):
    2001::1, subnet is 2001::/64
```

```
   Joined group address(es):
     FF02::1
     FF02::2
     FF02::1:FF00:1
     FF02::1:FFB6:E410
   MTU is 1500 bytes
   ICMP error messages limited to one every 100 milliseconds
   ICMP redirects are enabled
   ICMP unreachables are sent
   Input features: Verify Unicast Reverse-Path
  IPv6 verify source reachable-via rx, ACL test
    0 verification drop(s) (process), 0 (CEF)
    0 suppressed verification drop(s) (process), 0 (CEF)
   ND DAD is enabled, number of DAD attempts: 1
   ND reachable time is 30000 milliseconds (using 30000)
   ND advertised reachable time is 0 (unspecified)
   ND advertised retransmit interval is 0 (unspecified)
   ND router advertisements are sent every 200 seconds
   ND router advertisements live for 1800 seconds
   ND advertised default router preference is Medium
   Hosts use stateless autoconfig for addresses.
```

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

```
Device# show ipv6 traffic

 IPv6 statistics:
       Rcvd:  6 total, 0 local destination
              0 source-routed, 0 truncated
              0 format errors, 0 hop count exceeded
              0 bad header, 0 unknown option, 0 bad source
              0 unknown protocol, 0 not a router
              0 fragments, 0 total reassembled
              0 reassembly timeouts, 0 reassembly failures
       Sent:  34 generated, 28 forwarded
              0 fragmented into 0 fragments, 0 failed
              0 encapsulation failed, 0 no route, 0 too big
              0 RPF drops, 0 RPF suppressed drops
       Mcast: 6 received, 34 sent

ICMP statistics:
       Rcvd: 6 input, 0 checksum errors, 0 too short
             0 unknown info type, 0 unknown error type
             unreach: 0 routing, 0 admin, 0 neighbor, 0 address, 0 port
                      0 sa policy, 0 reject route
             parameter: 0 error, 0 header, 0 option
             0 hopcount expired, 0 reassembly timeout,0 too big
             0 echo request, 0 echo reply
             0 group query, 0 group report, 0 group reduce
             0 router solicit, 0 router advert, 0 redirects
             0 neighbor solicit, 0 neighbor advert
       Sent: 34 output, 0 rate-limited
             unreach: 0 routing, 0 admin, 0 neighbor, 0 address, 0 port
                      0 sa policy, 0 reject route
             parameter: 0 error, 0 header, 0 option
             0 hopcount expired, 0 reassembly timeout,0 too big
             0 echo request, 0 echo reply
             0 group query, 0 group report, 0 group reduce
             0 router solicit, 18 router advert, 0 redirects
             2 neighbor solicit, 2 neighbor advert
```

Configuration Examples for Unicast Reverse Path Forwarding ACL Support

Example: Configuring Unicast RPF with ACL Support

```
Device# configure terminal
Device(config)# interface gigabitethernet 0/0/1
Device(config-if)# ipv6 address 2001:DB8:1::1/96
Device(config-if)# ipv6 verify unicast source reachable-via any acl1
Device(config-if)# end
```

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Command List, All Releases
Unicast RPF command descriptions	• Cisco IOS Security Command Reference: Commands A to C • Cisco IOS Security Command Reference: Commands D to L • Cisco IOS Security Command Reference: Commands M to R • Cisco IOS Security Command Reference: Commands S to Z
Cisco Express Forwarding commands	Cisco IOS IP Switching Command Reference

Standards & RFCs

Standard/RFC	Title
RFC 1918	*Address Allocation for Private Internets*
RFC 2267	*Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing*

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for Unicast Reverse Path Forwarding ACL Support

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 2: Feature Information for Unicast Reverse Path Forwarding ACL Support

Feature Name	Releases	Feature Information
Unicast Reverse Path Forwarding ACL Support	Cisco IOS XE Release 3.7S	The Unicast Reverse Path Forwarding feature helps to mitigate problems that are caused by malformed or forged IP source addresses that pass through a device. The Unicast Reverse Path Forwarding ACL support feature adds the ACL support to the Unicast Reverse Path Forwarding feature. With the ACL support, Unicast RPF can determine whether to drop or to forward data packets that have malformed or forged IP source addresses. The following commands were introduced or modified: **ip verify unicast source reachable-via** and **ipv6 verify unicast source reachable-via**.

CISCO-IP-URPF-MIB Support

The CISCO-IP-URPF-MIB support provides Simple Network Management Protocol (SNMP) notification when a specified drop-rate threshold on a managed device is exceeded. You can use the IP Unicast Reverse Path Forwarding (RPF) feature to avert denial of service (DoS) attacks by verifying the validity of the source IP of an incoming packet. You can configure the Unicast RPF drop-rate threshold globally for a device or per interface.

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see Bug Search Tool and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for CISCO-IP-URPF-MIB Support

Before you configure CISCO-IP-URPF-MIB, you must configure the following features:

- Cisco Express Forwarding switching

- IP routing

- SNMP

- Unicast RPF

Restrictions for CISCO-IP-URPF-MIB Support

- Because Cisco IOS software does not support VPN routing and forwarding (VRF)-specific Unicast RPF counters, it does not support the following MIB objects related to VRF:

    - cipUrpfIfVrfName
    - cipUrpfVrfName
    - cipUrpfVrfIfDrops
    - cipUrpfVrfIfDiscontinuityTime

- This implementation of the CISCO-IP-URPF MIB supports only IPv4.

Information About CISCO-IP-URPF-MIB Support

Implementation of Unicast RPF Notification

Unicast RPF is a security feature that verifies the validity of the source IP of an incoming packet. When a packet arrives at an interface and its source IP is unknown in the routing table or is a known bad source address, Unicast RPF drops the packet. IP verification of the source is done to prevent the DoS attacks by detecting problems with the incoming packets on an interface. However, deploying Unicast RPF without some automated monitoring capability is a challenge.

The CISCO-IP-URPF-MIB lets you specify a Unicast RPF drop-rate threshold on interfaces of a managed device that will send an SNMP notification when the threshold is exceeded. The MIB includes objects for specifying global and per-interface drop counts and drop rates and a method to generate SNMP traps when the drop rate exceeds a configurable per-interface threshold.

Although you can configure some parameters globally, you must configure the CISCO-IP-URPF-MIB on individual interfaces.

Elements of Unicast RPF Notification

The elements described in the following sections make Unicast RPF drop-rate notification work:

Drop-Rate Computation

Whenever Unicast RPF is configured on an interface, the drop-rate calculation is done periodically (at intervals specified by the cipUrpfComputeInterval object). Drop rates are computed over a constantly sliding window,

whose period starts at the configured number of seconds before the calculation and ends with the performance of the calculation.

Global Scalars

The following global scalars affect how the MIB agent computes all drop rates and generates notifications:

- cipUrpfDropRateWindow--This object specifies the window of time in the recent past over which the drop rate computation occurs. If there was no window (that is, the window is the epoch since booting up), an identical drop count burst at a later time would produce a lower drop rate than the one occurring earlier.

- cipUrpfComputeInterval--This object specifies how often the drop-rate computation occurs.

- cipUrpfDropNotifyHoldDownTime--This object specifies the minimum time between notifications for a particular packet flow on an interface.

Global Tables

The CISCO-IP-URPF-MIB includes the following global tables:

- cipUrpfTable--This table contains the global drop count and drop-rate objects per packet flow. These global rates are useful for quickly determining whether the managed device had Unicast RPF activity at a specific time.

- cipUrpfVrfTable--This table contains the index drop counters by VRF (if a VRF routing table is used to determine Unicast RPF checking). The table provides a method for VRF to index all the Unicast RPF-enabled interfaces.

Per-Interface Configuration

The following MIB objects enable per-interface configuration:

- cipUrpfIfDropRateNotifyEnable--This object specifies whether the system produces the cipUrpfIfDropRateNotify notification because Unicast RPF has dropped version cipUrpfIfIpVersion IP packets on the specified interface.

- cipUrpfIfNotifyDropRateThreshold--This object specifies the drop-rate threshold value above which a notification is generated.

Per-Interface Statistics

The following MIB objects track per-interface statistics:

- cipUrpfIfMonTable--This table contains the statistics for a particular packet flow on an interface.

- cipUrpfIfDrops--This object accumulates Unicast RPF drops on an interface. Snapshots of this value are used in the drop-rate computation. The computed drop rate is specified in the cipUrpfIfDropRate object. If Unicast RPF is configured on a subinterface, drop rates are computed.

How to Configure Unicast RPF Drop-Rate Notification

Configuring Unicast RPF Drop-Rate Notification via Syslog

Perform this task to configure the Unicast RPF drop-rate threshold and computation parameters for notification via syslog.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip verify drop-rate compute window** *seconds*
4. **ip verify drop-rate compute interval** *seconds*
5. **ip verify drop-rate notify hold-down** *seconds*
6. **interface** *type number*
7. **ip verify unicast notification threshold** *packets-per-second*
8. **end**
9. **show ip interface** *type number*
10. **debug ip verify mib**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**ip verify drop-rate compute window** *seconds* **Example:** `Router(config)# ip verify drop-rate compute window 60`	Configures the period of time, in seconds, over which the Unicast RPF drop count used in the drop-rate computation is collected. • The range is from 30 to 300. The default is 300. **Note** The value for the compute window must be greater than or equal to that entered using the **ip verify drop-rate compute interval** command.
Step 4	**ip verify drop-rate compute interval** *seconds* **Example:**	Configures the interval of time, in seconds, between Unicast RPF drop-rate computations.

	Command or Action	Purpose
| | ```
Router(config)# ip verify drop-rate compute
interval 60
``` | • The range is from 30 to 300. The default is 30.<br><br>**Note** The value for the compute interval must be less than or equal to that entered using the **ip verify drop-rate compute window** command. |
| Step 5 | **ip verify drop-rate notify hold-down** *seconds*<br><br>**Example:**<br><br>```
Router(config)# ip verify drop-rate notify
hold-down 60
``` | Configures the minimum time, in seconds, between Unicast RPF drop-rate notifications.<br><br>• The range is from 30 to 300. The default is 300. |
| Step 6 | **interface** *type number*<br><br>**Example:**<br><br>```
Router(config)# interface ethernet 3/0
``` | Configures an interface and enters interface configuration mode. |
| Step 7 | **ip verify unicast notification threshold** *packets-per-second*<br><br>**Example:**<br><br>```
Router(config-if)# ip verify unicast notification
 threshold 750
``` | Configures the threshold value, in packets per second, which determines whether to send a Unicast RPF drop-rate notification.<br><br>• The range is from 0 to 2147483647. The default is 1000.<br><br>**Note** If you configure the threshold as 0, every packet drop triggers a notification. |
| Step 8 | **end**<br><br>**Example:**<br><br>```
Router(config-if)# end
``` | Returns to privileged EXEC mode. |
| Step 9 | **show ip interface** *type number*<br><br>**Example:**<br><br>```
Router# show ip interface ethernet 2/3
``` | (Optional) Displays the verification drop rate and the number of verification drops when Unicast RPF is configured for an interface. |
| Step 10 | **debug ip verify mib**<br><br>**Example:**<br><br>```
Router# debug ip verify mib
``` | (Optional) Displays output that is useful for troubleshooting Unicast RPF notification. |

Configuring Unicast RPF Drop-Rate Notification via SNMP

Perform this task to configure the Unicast RPF drop-rate threshold and computation parameters for notification via SNMP.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip verify drop-rate compute window** *seconds*
4. **ip verify drop-rate compute interval** *seconds*
5. **ip verify drop-rate notify hold-down** *seconds*
6. **interface** *type number*
7. **ip verify unicast notification threshold** *packets-per-second*
8. **snmp trap ip verify drop-rate**
9. **end**
10. **show ip interface** *type number*
11. **debug ip verify mib**

DETAILED STEPS

	Command or Action	**Purpose**
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**ip verify drop-rate compute window** *seconds* **Example:** `Router(config)# ip verify drop-rate compute window 60`	Configures the period of time, in seconds, over which the Unicast RPF drop count used in the drop-rate computation is collected. • The range is from 30 to 300. The default is 300. **Note** The value for the compute window must be greater than or equal to that entered using the **ip verify drop-rate compute interval** command.
Step 4	**ip verify drop-rate compute interval** *seconds* **Example:** `Router(config)# ip verify drop-rate compute interval 60`	Configures the interval of time, in seconds, between Unicast RPF drop-rate computations. • The range is from 30 to 300. The default is 30. **Note** The value for the compute interval must be less than or equal to that entered using the **ip verify drop-rate compute window** command.
Step 5	**ip verify drop-rate notify hold-down** *seconds* **Example:**	Configures the minimum time, in seconds, between Unicast RPF drop-rate notifications.

	Command or Action	Purpose
	`Router(config)# ip verify drop-rate notify hold-down 60`	• The range is from 30 to 300. The default is 300.
Step 6	**interface** *type number* **Example:** `Router(config)# interface ethernet 3/0`	Configures an interface and enters interface configuration mode.
Step 7	**ip verify unicast notification threshold** *packets-per-second* **Example:** `Router(config-if)# ip verify unicast notification threshold 750`	Configures the threshold value, in packets per second, which determines whether to send a Unicast RPF drop-rate notification. • The range is from 0 to 2147483647. The default is 1000. **Note** If you configure the threshold to be 0, every packet drop triggers a notification.
Step 8	**snmp trap ip verify drop-rate** **Example:** `Router(config-if)# snmp trap ip verify drop-rate`	Configures the router to send an SNMP notification when the Unicast RPF drop rate exceeds the configured threshold.
Step 9	**end** **Example:** `Router(config-if)# end`	Returns to privileged EXEC mode.
Step 10	**show ip interface** *type number* **Example:** `Router# show ip interface ethernet 2/3`	(Optional) Displays the verification drop rate and the number of verification drops when Unicast RPF is configured for an interface.
Step 11	**debug ip verify mib** **Example:** `Router# debug ip verify mib`	(Optional) Displays output that is useful for troubleshooting Unicast RPF notification.

Configuration Examples for CISCO-IP-URPF-MIB Support

Example Configuring Unicast RPF Drop-Rate Notification via Syslog

The following example shows how to configure Unicast RPF drop-rate notification via syslog:

```
Router> enable
Router# configure terminal
Router(config)# ip verify drop-rate compute window 60
Router(config)# ip verify drop-rate compute interval 60
Router(config)# ip verify drop-rate notify hold-down 60
Router(config)# i
nterface ethernet 3/0
Router(config-if)# ip verify unicast notification threshold 750
Router(config-if)# end
```

Example Configuring Unicast RPF Drop-Rate Notification via SNMP

The following example shows how to configure Unicast RPF drop-rate notification via SNMP:

```
Router> enable
Router# configure terminal
Router(config)# ip verify drop-rate compute window 60
Router(config)# ip verify drop-rate compute interval 60
Router(config)# ip verify drop-rate notify hold-down 60
Router(config)# interface ethernet 3/0
Router(config-if)# ip verify unicast notification threshold 750
Router(config-if)# snmp trap ip verify drop-rate
Router(config-if)# end
```

Example Verifying and Troubleshooting the Unicast RPF Configuration

The following is sample output from the **show ip interface** command. The output displays the verification drop rate and the number of verification drops when Unicast RPF is configured for an interface. The last five lines in the following example show the output of the **show ip interface**command when Unicast RPF is configured:

```
Router# show ip interface ethernet 2/3
Ethernet2/3 is up, line protocol is up
  Internet address is 10.10.5.4/16
  Broadcast address is 255.255.255.255
  Address determined by non-volatile memory
  MTU is 1500 bytes
  Helper address is not set
  Directed broadcast forwarding is disabled
  Outgoing access list is not set
  Inbound  access list is not set
  Proxy ARP is enabled
  Local Proxy ARP is disabled
  Security level is default
  Split horizon is enabled
  ICMP redirects are always sent
  ICMP unreachables are always sent
  ICMP mask replies are never sent
  IP fast switching is disabled
  IP Flow switching is disabled
  IP CEF switching is disabled
  IP Null turbo vector
  IP Null turbo vector
  IP multicast fast switching is disabled
  IP multicast distributed fast switching is disabled
  IP route-cache flags are No CEF
  Router Discovery is disabled
```

```
   IP output packet accounting is disabled
   IP access violation accounting is disabled
   TCP/IP header compression is disabled
   RTP/IP header compression is disabled
   Probe proxy name replies are disabled
   Policy routing is disabled
   Network address translation is disabled
   WCCP Redirect outbound is disabled
   WCCP Redirect inbound is disabled
   WCCP Redirect exclude is disabled
   BGP Policy Mapping is disabled
   Input features: uRPF
   IP verify source reachable-via RX, allow default
    0 verification drops
    0 suppressed verification drops
    0 verification drop-rate
Router#
```

The following is sample output from the **debug ip verify mib** command. The command displays output that is useful for troubleshooting Unicast RPF notification:

```
Router# debug ip verify mib
01:29:45: cipUrpfScalar_get, searchType 161
01:29:45: ipurpfmib_get_scalars
01:29:45: cipUrpfScalar_get, searchType 161
01:29:45: cipUrpfScalar_get, searchType 161
01:29:45: ipurpfmib_get_scalars
01:29:45: cipUrpfScalar_get, searchType 161
01:29:45: cipUrpfScalar_get, searchType 161
01:29:45: ipurpfmib_get_scalars
01:29:45: cipUrpfScalar_get, searchType
161ipurpfmib_get_urpf_entryipurpfmib_get_urpf_entryipurpfmib_get_urpf_entryipurpfmib_get_
urpf_entry
01:29:45: cipUrpfIfMonEntry_get, searchType 161
01:29:45: ipurpfmib_get_urpf_ifmon_entry entry: ST 161, if 1, ip 1
01:29:45: cipUrpfIfMonEntry_get, searchType 161
01:29:45: ipurpfmib_get_urpf_ifmon_entry entry: ST 161, if 1, ip 1
01:29:45: cipUrpfIfMonEntry_get, searchType 161
01:29:45: ipurpfmib_get_urpf_ifmon_entry entry: ST 161, if 1, ip 1
01:29:45: cipUrpfIfMonEntry_get, searchType 161
01:29:45: ipurpfmib_get_urpf_ifmon_entry entry: ST 161, if 1, ip 1
01:29:45: cipUrpfIfMonEntry_get, searchType 161
01:29:45: ipurpfmib_get_urpf_ifmon_entry entry: ST 161, if 1, ip 1
01:29:45: cipUrpfIfMonEntry_get, searchType 161
01:29:45: ipurpfmib_get_urpf_ifmon_entry entry: ST 161, if 1, ip 1
01:29:45: cipUrpfIfMonEntry_get, searchType 161
01:29:45: ipurpfmib_get_urpf_ifmon_entry entry: ST 161, if 1, ip 1
01:29:45: cipUrpfIfMonEntry_get, searchType 161
01:29:45: ipurpfmib_get_urpf_ifmon_entry entry: ST 161, if 1, ip 1
01:29:45: cipUrpfIfMonEntry_get, searchType 161
01:29:45: ipurpfmib_get_urpf_ifmon_entry entry: ST 161, if 1, ip 1
```

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Commands List, All Releases
Security commands	*Cisco IOS Security Command Reference*
Configuring Unicast RPF	"Configuring Unicast Reverse Path Forwarding" module in the *Cisco IOS Security Configuration Guide: Securing the Data Plane*
Configuring SNMP	"Configuring SNMP Support" module in the *Network Management Configuration Guide*

MIBs

MIB	MIBs Link
CISCO-IP-URPF-MIB	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFC	Title
None	--

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for CISCO-IP-URPF-MIB Support

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 3: Feature Information for CISCO-IP-URPF-MIB Support

Feature Name	Releases	Feature Information
CISCO-IP-URPF-MIB Support	Cisco IOS XE Release 3.9S	The CISCO-IP-URPF-MIB provides SNMP notification when a specified drop-rate threshold on a managed device is exceeded. You can use the IP Unicast RPF feature to avert DoS attacks by verifying the validity of the source IP of an incoming packet. You can configure the Unicast RPF drop-rate threshold globally for a device or per interface. The following commands were introduced or modified: **debug ip verify mib**, **ip verify drop-rate compute interval**, **ip verify drop-rate compute window**, **ip verify drop-rate notify hold-down**, **ip verify unicast notification threshold**, **show ip interface**, **snmp trap ip verify drop-rate**