CHAPTER 6

Network Policy Enforcement

Baseline network policy enforcement is primarily concerned with ensuring that traffic entering a network conforms to the network policy, including the IP address range and traffic types. Anomalous packets should be discarded as close to the edge of the network as possible, thereby minimizing the risk of exposure.

This section highlights the key steps to implementing baseline network policy enforcement, including:

- Access Edge Filtering
- IP Spoofing Protection

CSF Methodology Assessment

The results of applying the CSF methodology for baseline network policy enforcement are presented in Table 6-1 and highlight the technologies and features identified for enforcing baseline network policy enforcement and which are integrated in Network Security Baseline.

Total Visibility

Table 6-1 Network Policy Enforcement—Total Visibility

Identify	Monitor	Correlate
	• Logging – Syslog – SNMP	

Complete Control

Table 6-2 Network Policy Enforcement—Complete Control

Harden	Isolate	Enforcement
	• ACL	• Access Edge Filtering – iACLs • IP Spoofing Protection – uRPF – IP Source Guard – DHCP/ARP enforcement

Access Edge Filtering

The Network Security Baseline is focused on securing the network infrastructure itself, the control and management planes. Access edge filtering in this context is implemented to enforce policy on what traffic is permitted to be directed towards the network infrastructure devices themselves.

In Cisco IOS, access edge filtering for control and the data planes is achieved using ACLs developed to protect the infrastructure. These are referred to as infrastructure protection ACLs (iACLs).

Detailed information on the role, development and configuration of iACLs is provided in Infrastructure Protection ACLs (iACLs), page 8-9.

IP Spoofing Protection

Spoofing protection involves discarding traffic that has an invalid source address. Network Security Baseline includes source IP spoofing protection based on BCP38/RFC 2827 ingress traffic filtering.

Packets with spoofed source IP addresses represent a security risk as they are often used to conduct an attack, in order to evade traceability and bypass access controls. They may also be used to direct an attack at a spoofed source, something known as a “reflection attack”.

Table 6-3 IP Spoofing Protection

Type of Spoofing	Related Attacks	Possible Countermeasures
IP spoofing	ICMP unreachable storm	• ACLs • uRPF • IP Source Guard • DHCP Secured IP Address Assignment and DHCP Authorized ARP
	ISYN flood	
	SYN flood	
	Spoof trusted IP addresses to leverage trust relationship	
	DDoS	

Spoofed traffic with an invalid source IP address may include traffic from a:

- RFC1918, DSUA or non-allocated IP address range
- Valid IP network address range but not originating from the associated legitimate network

Implementing BCP38/RFC 2827 ingress traffic filtering to address source IP address spoofing renders the use of invalid source IP addresses ineffective, forcing attacks to be initiated from valid, reachable IP addresses. This is beneficial since it enables greater success in tracing the originator of an attack.

IP spoofing protection offers the following key benefits:

- Enables improved, clearer analysis of network telemetry data
- Increases traceback success
- Improves the traceability of the source of malicious behavior
- Eliminates bogon-sourced traffic from the peering edge
- Aides the effectiveness of iACLs

Cisco offers the following techniques for BCP38 ingress traffic filtering:

- Access Control Lists (ACLs)

  ACLs are the traditional technique for filtering forged source IP addresses. However, ACLs are not dynamic in nature, requiring manual configuration changes, and may have an impact on the performance of a device. It is thus recommended that ACLs are used only in a limited manner, as a complement to uRPF, for strict, static policies, such as filtering RFC 1918, DSUA and non-allocated IP addresses. They may also be used to complement uRPF loose mode for source IP address spoofing protection when uRPF strict mode is not possible. ACLs are covered in Appendix 4, "Device Resiliency and Survivability."

- uRPF

  uRPF offers a dynamic technique for enabling BCP38/RFC 2827 ingress traffic filtering, discarding packets with invalid source IP addresses based on a reverse-path look-up. Its dynamic nature provides the key advantages of offering minimal operational overhead and a scalable, timely enforcement technique. In addition, uRPF generally introduces minimal performance impact to a device on which it is enabled. uRPF is typically deployed as an edge technology in order to be most effective, minimizing the valid IP address space range and enforcing the discard of anomalous packets as close to their origin as possible.

- IP Source Guard

  IP Source Guard is used in switched environments to prohibit the use of forged MAC and source IP addresses. This feature is deployed on Layer 2 switching devices and is primarily designed for DHCP segments. Hosts with static address may also be supported, though additional operational complexity is introduced by doing so.

- DHCP Secured IP Address Assignment and DHCP Authorized ARP

  These Cisco IOS features are available on routers supported on the T-train and offer similar functionality in a routing environment as IP Source Guard in a switching environment. They are used in routed environments where the local router is also the local DHCP server to prohibit the use of forged MAC and source IP addresses

  http://www.cisco.com/en/US/docs/ios/12_2t/12_2t15/feature/guide/ftdsiaa.html
  http://www.cisco.com/en/US/docs/ios/12_3t/12_3t4/feature/guide/gtautarp.html

Network Security Baseline includes ACLs and uRPF to address IP spoofing protection, as these are the simplest techniques to implement and form a baseline upon which additional spoofing protection techniques may be deployed.

ACLs are covered in Chapter 4, "Device Resiliency and Survivability."

Unicast Reverse Path Forwarding (uRPF)

uRPF offers a dynamic technique for enabling BCP38/RFC 2827 ingress traffic filtering, discarding packets with invalid source IP addresses based on a reverse-path look-up. Its dynamic nature provides the following key advantages:

- Minimal operational overhead
- Scalable, timely enforcement technique
- Generally introduces minimal performance impact to a device on which it is enabled.

uRPF is thus a highly attractive alternative to traditional ACLs, which typically demand significant management overhead and have a greater impact on device performance.

uRPF is typically deployed as an edge technology in order to be most effective, minimizing the valid IP address space range and enforcing the discard of anomalous packets as close to their origin as possible.

The key function of uRPF is to verify that the path of an incoming packet is consistent with the local packet forwarding information. This is achieved by performing a reverse path look-up (hence the feature's name) using the source IP address of an incoming packet in order to determine the current path (adjacency) to that IP address. The validity of this path determines whether uRPF will pass or drop the packet.

If the path is valid, the packet will be passed. If the path is not valid, the packet will be silently discarded (unless an ACL exemption is configured). See Figure 6-1.

Figure 6-1 Unicast Reverse Path Forwarding (uRPF)

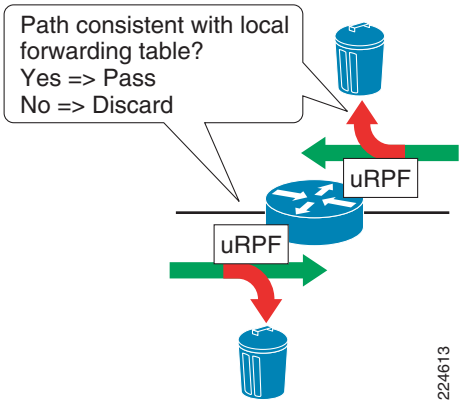

Once enabled on an interface, uRPF checks all IP packets (IPv4 and IPv6) on the input path of that interface. The key edge locations and the specific objectives of uRPF in each of these locations are shown Table 6-4.

Table 6-4 uRPF Key Edge Locations and Objectives

Network Edge Location	Specific uRPF Objectives
Access Layer Edge First Routed Hop[1]	• Discard incoming packets on the first routed hop of the access layer edge that do NOT have a source IP address within the access network IP range
Enterprise Internet Edge	• Discard incoming packets on the internal interface that do NOT have a source IP address within the internal network IP range • Discard incoming packets on the external Internet interface which have a source IP address within the internal network IP range

1. The term 'Access Layer Edge First Routed Hop' is used to refer to the first Layer 3 routed hop that is encountered in the access layer of a network. In an enterprise deployment, this is often referred to as the 'Distribution Layer Edge' and may be a router or a Layer2/Layer 3 switch

Access Layer First Routed Hop

In this location, ingress traffic filtering for source IP spoofing protection is applied at the first Layer 2-Layer 3 routed boundary of the access layer network (see Figure 6-2).

Figure 6-2 Access Layer First Routed Hop

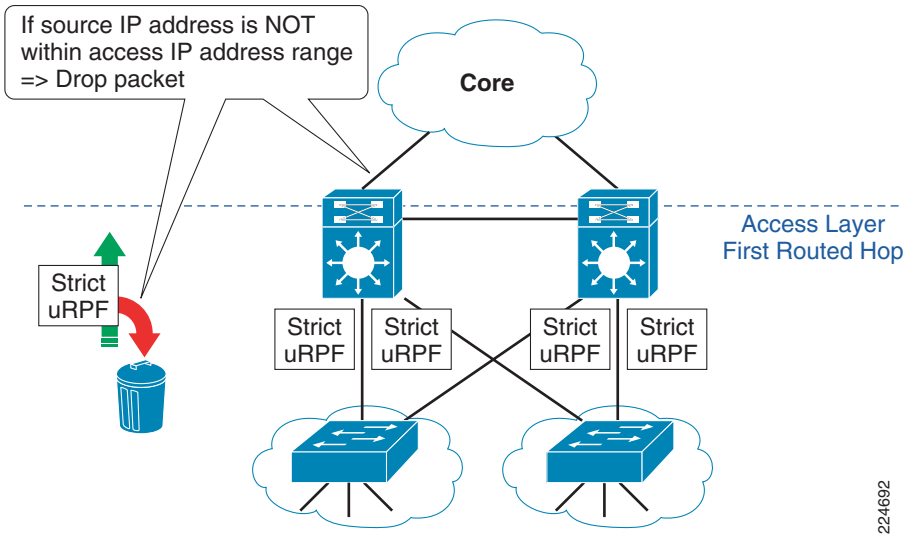

The key objective is:

Discard incoming packets on the first routed hop of the access layer edge that do NOT have a source IP address within the assigned access network IP range.

As illustrated in Figure 6-2, the general deployment approach is:

Enable uRPF strict mode on all first routed hop interfaces facing the Layer 2 access network downstream devices.

Deployment Considerations

In certain access layer edge scenarios, uRPF strict mode may either not be possible or may require some additional design work. Common scenarios requiring additional consideration include:

- Access layer topologies with downstream Layer 2 devices which are interconnected for redundancy, thereby creating the possibility of multiple paths to a particular IP prefix.
  This scenario typically creates a challenge with DHCP assignment and management applications.
- Dual-homed hosts which may be configured to send traffic to either gateway

IP Source Guard may be used to supplement uRPF but this is not included in Network Security Baseline.

Enterprise Internet Edge

In this location, ingress traffic filtering for source IP spoofing protection is applied on the SMB or Enterprise's Internet edge devices (see Figure 6-3).

Figure 6-3 Enterprise Internet Edge

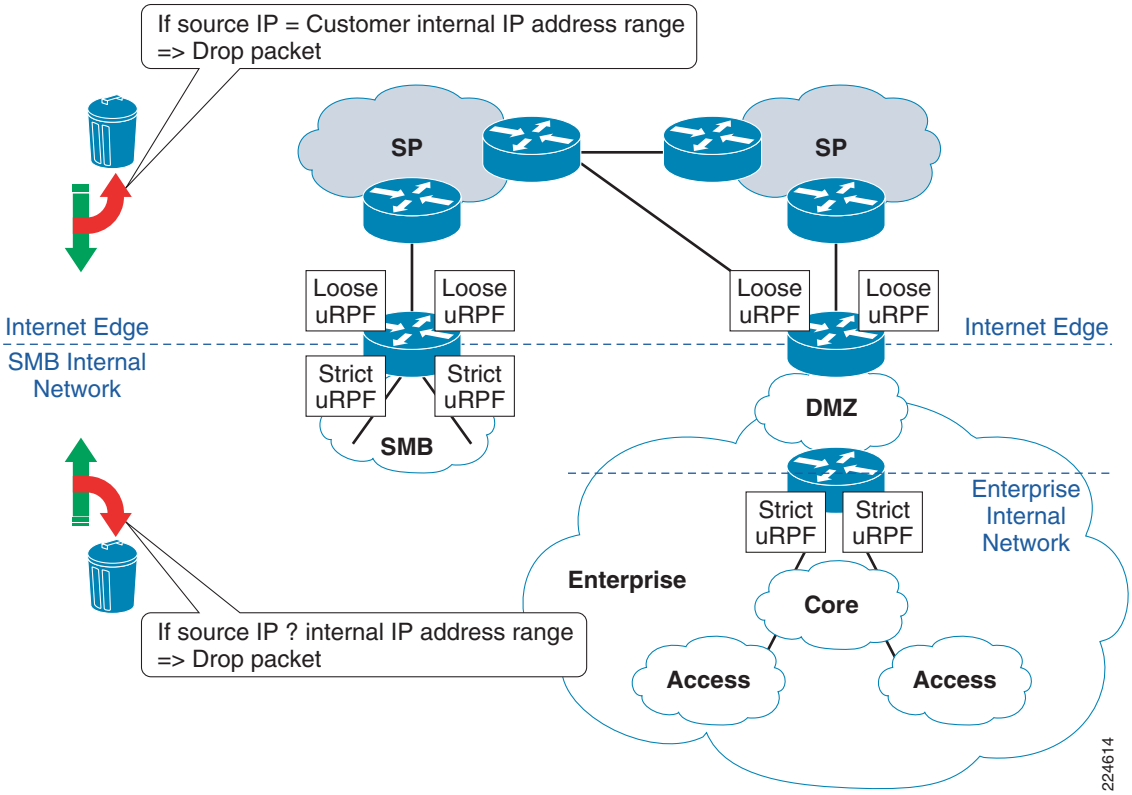

The key objectives being:

- Discard incoming packets on the internal interface that do NOT have a source IP address within the internal network IP range.
- Discard incoming packets on the external Internet interface which have a source IP address within the internal network IP range.

As illustrated in Figure 6-3, the general deployment approach is:

- Enable uRPF strict mode on the routed interfaces of first-hop devices facing the internal network
- Enable uRPF loose mode on external routed interfaces facing the Internet, along with source- based ACLs (since uRPF loose mode alone will not meet the design objective)

Deployment Considerations

- uRPF strict mode is only possible on the internal network edge interfaces if all paths to any internal IP prefix are of all equal cost, thereby allowing all valid paths to be present in the FIB. If this is not the case, as may be experienced in networks with multiple Internet connections from geographically disperse sites, uRPF loose mode must be deployed as an alternative.
- On networks with just a single external Internet connection, it may be possible to deploy uRPF strict mode. Review the FIB to verify.
- uRPF strict mode is typically not possible on the external Internet edge for networks which are multi-homed to one or more service providers, since multiple valid paths may exist to an IP prefix. The CEF FIB will typically only contain a subset of these valid paths, due to its ‘single best path’ selection algorithm, and thus uRPF strict mode may, inadvertently, drop valid packets.
- uRPF loose mode provides only limited source IP spoofing protection, since any path on the device will be valid. Consequently, only packets with source IP addresses not present in the FIB will be dropped.
- Since uRPF loose mode provides only limited source IP spoofing protection, source-based ACLs are required on the external Internet edge in order to supplement uRPF loose mode in achieving our objective of dropping incoming packets with a source IP address matching the internal network IP range.
- uRPF loose mode is extremely valuable in this location as a key enabler for the deployment of source IP based black hole and SRTBH rapid reaction attack tools.
- If the FIB contains a path to valid IP prefixes via a default route, the ‘allow-default’ uRPF option may be required to ensure valid packets are not inadvertently dropped. Note, however, that this may significantly reduce the effectiveness of uRPF in providing source IP spoofing protection.
- On external Internet interfaces running BGP and multi-homed to a single AS, it may be possible to deploy uRPF strict mode by using BGP parameters to force the FIB to be populated with all valid paths.
- A static route to null0 for bogon IP addresses may also be inserted into the routing table to enable uRPF to drop traffic from these invalid source IP addresses, thereby simplifying ACLs.

For more information on configuring uRPF, refer to the following URL:

http://www.cisco.com/en/US/docs/ios/security/configuration/guide/sec_cfg_unicast_rpf_ps6350_TSD_Products_Configuration_Guide_Chapter.html