WHITE PAPER

UNICAST REVERSE PATH FORWARDING ENHANCEMENTS FOR THE INTERNET SERVICE PROVIDER—INTERNET SERVICE PROVIDER NETWORK EDGE

HIGHLIGHTS

- New additions to Unicast Reverse Path Forwarding (uRPF) that would allow it to be used on the ISP-ISP edge of a network
- New DoS/DDoS reaction tool that would have BGP advertisements trigger drops on the edge of an ISP's network.

General questions on uRPF can be sent to unicast-rpf@cisco.com or cisco-nsp@puck.nether.net.

INTRODUCTION

Unicast Reverse Path Forwarding (uRPF) was a feature originally created to implement BCP 38/RFC 2827 *Network Ingress Filtering: Defeating Denial of Service Attacks Which Employ IP Source Address Spoofing*, by P. Ferguson and D. Senie. As such, uRPF was originally designed for the customer-to-Internet service provider (ISP) network edge (See Figure 1). The objective was to develop a feature that could be easily automated in the customer provisioning system, scale as new address blocks were allocated to the customer, and work with multiway trie- (MTRIE-) based Cisco® Express Forwarding switching. Unicast RPF met those objectives, even in situations where the customer was multihomed to one or more upstream ISPs.* Originally implemented in Cisco IOS® Software Release 11.1(17)CC, uRPF provided a new ISP security tool for BCP 38/RFC 2827 deployment.

Figure 1. Original URPF Deployment on the Customer-ISP Edge

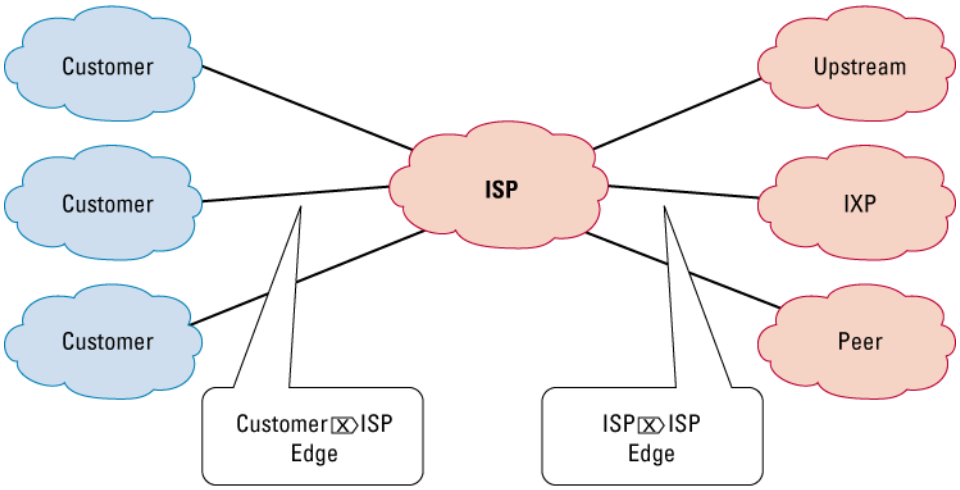

* Unicast RPF does work with asymmetrical routing on the customer-ISP edge. Detailed configurations and an explanation of the myth that uRPF does not work with asymmetrical routing is detailed in "Cisco ISP Essentials" at www.cisco.com/public/cons/isp/documents/IOSEssentialsPDF.zip.

Over time and with an increase in distributed denial-of-service (DDoS) attacks on the Internet, uRPF's functionality was reviewed as a tool that ISPs can use on the ISP-ISP edge to enable dynamic (Border Gateway Protocol [BGP]) triggered black hole filtering on the destination and the source address. The forwarding/routing characteristics on the ISP-ISP edge differ vastly from the forwarding characteristics on the customer-ISP edge. For example, the customer-ISP edge has relatively symmetrical flows (excluding some types of multihomed configurations). In this case, uRPF can use the best single path selection in the forwarding information base (FIB). However, the ISP-ISP edge can have several connections to different ISPs. Each ISP connection exchanges BGP routing information. With the original uRPF implementation (known as "strict mode"), the best single path selection that the BGP routing information base (RIB) sent to the FIB created issues when the FIB did not match the packet flows from the ISP connections. Strict mode uRPF would not work if the FIB and the packet flow did not match on an interface. Hence, the uRPF strict mode implementation was not an option for the ISP-ISP edge of a network.

In order to provide ISPs with a DDoS resistance tool on the ISP-ISP edge, uRPF was modified from its original "strict mode" implementation. Unicast RPF was modified to check the source addresses of each ingress packet without regard for the specific interface upon which it was received. This is known as "loose mode." This allowed uRPF the ability to automatically detect and drop packets such as the following:

- RFC 1918 source addresses
- Other documented special use addresses (DSUAs)* that should not appear in the source
- Unallocated addresses that have not been allocated by the regional Internet registries (RIRs)
- Source addresses that are routed to a null interface on the router

Any of these source addresses should not be traversing the Internet, are traversing the wrong path, or are defined to be dropped. Identifying and dropping these packets on the inter-ISP border are considered to be a very "OpEx-friendly" way of increasing resistance against many attacks. For this reason, uRPF functionality was reviewed and modified to enable it to work on the ISP-ISP edge of the network. The following outlines the first phase of enhancements made to uRPF to work on the ISP-ISP edge of the network.

WHAT HAS BEEN ENHANCED?

Unicast RPF is undergoing a complete review and overhaul. The first phase was to push the existing uRPF functionality to the limits of the original design. Specifically, the first phase documented how uRPF will work on the customer-ISP edge even with asymmetrical routing.** The second phase—committed through DDTS CSCdr93424—allowed a looser uRPF check and fixed some security holes. This "loose check" allows uRPF to be used on the ISP-ISP edge of the network. The compelling functionality that this second phase added to uRPF was an alternative to standard access control lists (ACLs) to drop packets on the edge of the network. A third phase is currently under way that will create a way to have strict enforcement of the uRPF check on individual ISP-ISP edges. Here, external BGP (eBGP) peer sessions will send specific prefixes to a dedicated Virtual routing and forwarding (VRF) table. This will allow uRPF to query the VRF table that contains all the routes for that specific eBGP peering session over the interface, thus verifying (authorizing) the source addresses of packets matching the advertised routes from the peered ISP. Together, these three phases will provide ISPs with the tools they need to increase the robustness and control of their networks.

* DSUAs are detailed in the Internet draft "Documenting Special Use IPv4 Address Blocks That Have Been Registered with IANA" by Bill Manning (http://www.ietf.org/rfc/rfc3330.txt?number=3330).

** Complete configuration details of how to use uRPF on the customer-ISP edge with multihoming and asymmetrical routing is documented in detail in the "ISP Essentials" white paper and presentation, located at ftp://ftp-eng.cisco.com/cons/isp/documents/.

This paper specifically addresses the second phase of the uRPF overhaul. Unicast RPF has been enhanced to allow it work on the ISP-ISP edge of the network. The new "loose check" enhancement removes the match requirement on the specific ingress interface, allowing uRPF to "loose" check packets. This allows an ISP peering router with multiple links to multiple ISPs to check the source IP address of ingress packets to see if they exist in the FIB. If they exist, then the packets are forwarded. If they do not exist in the FIB, then the packets fail and are dropped. This increases resistance against DoS/DDoS attacks that use spoofed source addresses based on RFC1918, Martian, and unallocated IP addresses.*

The uRPF enhancement also provides a new tool to drop packets based on source or destination IP addresses using BGP updates vs. ACL updates. DoS/DDoS attacks are often dynamic—they change their character over time. Hence, many ACL updates may be needed over the incident period. This new uRPF enhancement allows "drop" updates to be propagated by BGP to the edge of the network—triggering packet drops on the DoS/DDoS packets.

DDTS CSCdr93424 was committed to 12.0(13.06)S01 and 12.1E for the Cat6K/OSR support. The 7200, 7500, GSR Engine 0, and GSR Engine 1 were supported in the first CCO published version—12.0(14)S. GSR Engine 2 support was added in 12.0(17)S. This feature was also committed to 12.2(13)T for routers that support the IOS T train.

OBJECTIVES FOR THE UNICAST RPF ENHANCEMENT

Unicast RPF originally was designed to prevent source address spoofing at the customer-ISP edge. For example, in Figure 4, uRPF works well on the interface on router F leading from the ISP to the customer. It will also work if the customer is multihomed to the ISP or multiple ISPs. Unicast RPF will also work on links to Internet exchange points (IXPs; for example, routers A and D).** What does not work in the uRPF original "strict mode" implementation is if uRPF is applied on routers with multiple connections to multiple ISPs.

The new additions to uRPF were intended to achieve two goals:

- Create a new option for uRPF to work between ISPs—specifically, on routers with multiple links to multiple ISPs
- Create a rapid reaction tool that uses BGP to trigger filters on the edge of the network of an ISP—shut down attacks based on the source and destination IP addresses

UNICAST RPF BETWEEN INTERNET SERVICE PROVIDERS: WHAT IS THE PROBLEM?

The original uRPF "strict mode" implementation does not work between typical ISP router interconnections where there are multiple ISP peers over multiple interfaces. The commonly cited reason for why it will not work on the ISP-ISP edge is asymmetrical routing. However, the asymmetrical routing reason does not accurately describe the core reasons why it will not work. The core strength of uRPF was that it used the FIB of the router to validate the reverse path of the packet. This allows uRPF to use the same optimized MTRIE lookups to do its validation. At the same time, the use of the FIB is the core in-flexibility of the URPF deployment in a network. The best path selection algorithms for the RIB and the FIB will select one best path. However, there might be more than one best path. The way routing protocols and forwarding tables are built today will only allow one best path into the forwarding table (see Figure 2).

This best path forwarding/routing characteristic is the reason why we have asymmetrical routing on the Internet. It is also the reason why strict mode uRPF will not work on the ISP-ISP edge. Comprehending the crux of this limitation is key to understanding the deployment limitations of the original uRPF 'strict mode' implementation. Once one understands that the limitation is not asymmetrical routing on the Internet, but the best path selection of how routing/forwarding works on a router, new uRPF deployment options can be created for the ISP. Some of those new options include deployment options with the original uRPF with multihomed customers on the customer-ISP edge and special ISP-ISP peering options such as routers connected to IXPs.

* Unallocated IP addresses are those Internet Assigned Numbers Authority (IANA) reserved addresses that have not been delegated by the RIRs.

** Check out the "ISP Essentials" white paper for examples of uRPF with multihomed customers:
ISP Essentials white paper: www.cisco.com/public/cons/isp/documents/IOSEssentialsPDF.zip
ISP Essentials seminar slides: www.cisco.com/public/cons/isp/documents/IOSEssentials_Seminar.zip

Figure 2. Unicast RPF Limitation on the ISP-ISP Edge Has More to Do with RIB/FIB Best Path Selection than Asymmetrical Routing

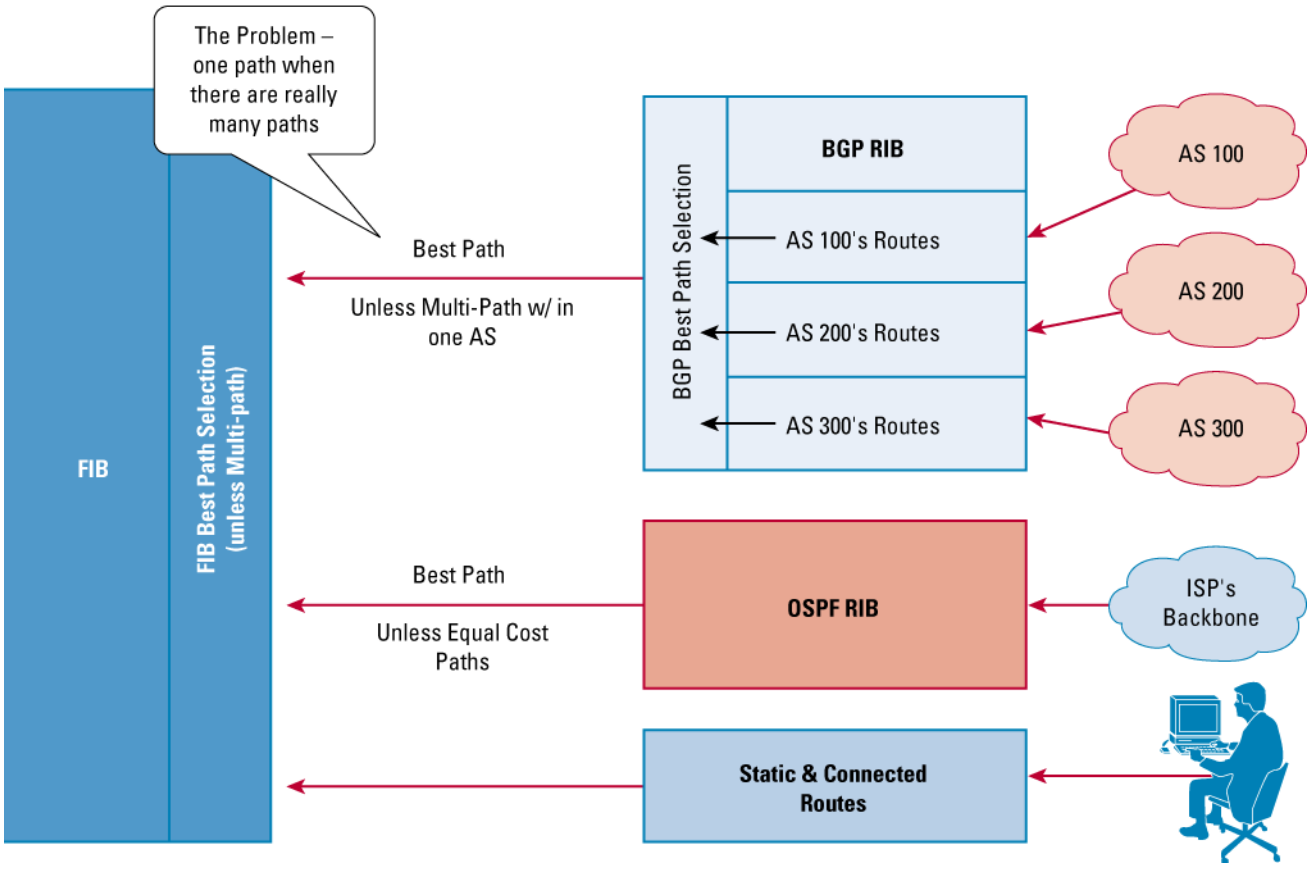

An IXP is a good example to demonstrate the original uRPF limitation. In Figure 3, the IXP border router has multiple ISP peers over one interface—the interface to the IXP switch. The BGP weight attribute is used on all the peers to keep the internal FIB symmetrically aligned so that any packet that arrives on the router goes out the best path of the interface connected to the IXP. With BGP weight and the single connection to the IXP switch, strict mode uRPF can be applied to this interface. Packets from each ISP that match their advertised policy will pass the uRPF check, since they will all have an adjacency equal to the one interface to the IXP.

The second diagram in Figure 3 illustrates what happens when multiple valid paths are fed to the router, yet the routing protocol and forwarding algorithm only allow one best path. In this example, the multiple ISP peers will result in different adjacencies for each route. Several ISPs might advertise the same prefix, with each being a valid path. BGP will pick one of them and insert it into the forwarding table. As a result, strict mode uRPF checks would fail on a valid packet sent from an ISP that is also advertising that route (since BGP has picked the route of another provider as the best path). Again, the problem with strict mode uRPF on the ISP-ISP edge has more to do with the character of how BGP works with forwarding tables than asymmetrical routing. Fixation on the term asymmetrical routing is misleading and fails to describe the many cases where uRPF will work and the few cases where it will not work.

Figure 3. Routers on the ISP-ISP Edge of the Network

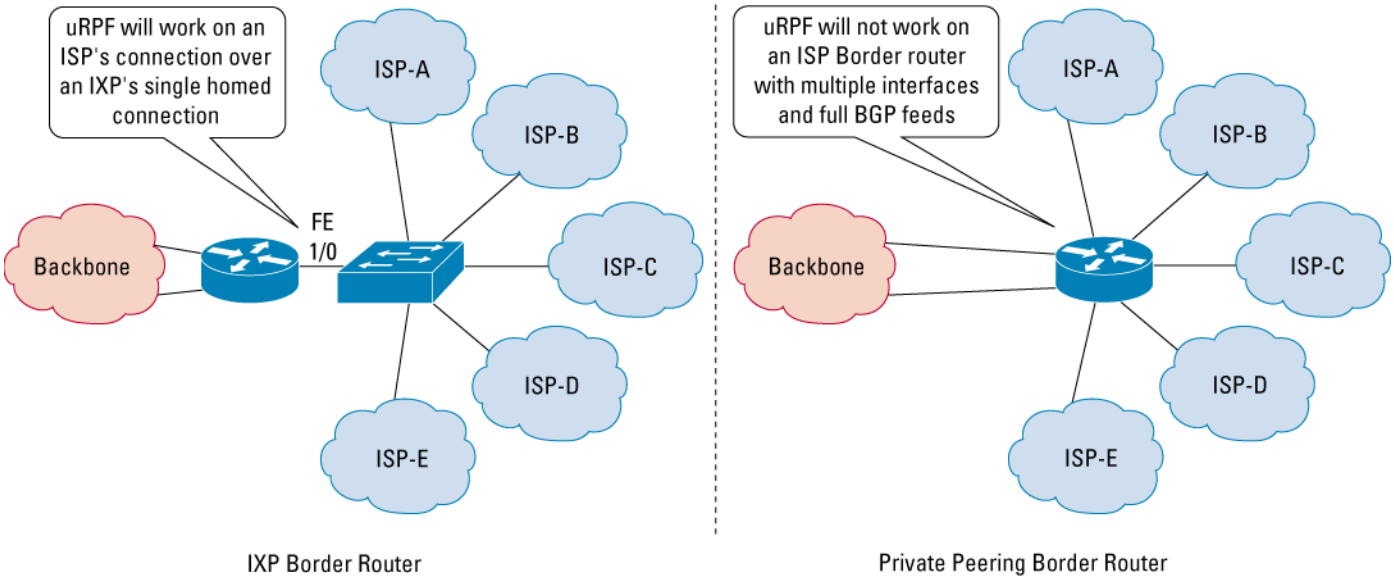

Another example that demonstrates how strict mode uRPF will work is on routers C and E in Figure 4. These routers are single homed between the two ISPs. The BGP weight attribute would be used to ensure that the router would always prefer the routes from the local eBGP session to all other BGP advertisements. This BGP weight trick is the key factor that will allow uRPF to work properly on the connection between the router and the single homed connection to the upstream ISP. This is not the case for routers A, B, and D. Each router has multiple BGP sessions to other ISPs over multiple links. BGP will take all the information from each of the BGP sessions and select the best path. This best path will be submitted to the forwarding table. Since there can only be one best path in this situation, the information in the forward table (FIB) will not necessarily match the traffic flow for any of the connections to other ISPs. As a result, strict mode uRPF will drop packets that should not be dropped.

As mentioned above, it is desirable to have uRPF work in all ISP-ISP edge scenarios. To achieve this, an enhancement was made, called “loose mode.”. This uRPF modification works around the requirement of best path selection in the forwarding/routing algorithms. It does this by only checking if there is an entry in the FIB. Nothing more. Nothing less. If an entry exists in the FIB, no matter what interface the packet arrived on, uRPF loose check would pass the packet. The loose check is not perfect, but there are no perfect solutions to Internet security—just more tools to make it more difficult to cause mischief on the Internet. The objective with uRPF loose check is to give ISPs a tool to make it more difficult to cause mischief.

Figure 4. ISP Network

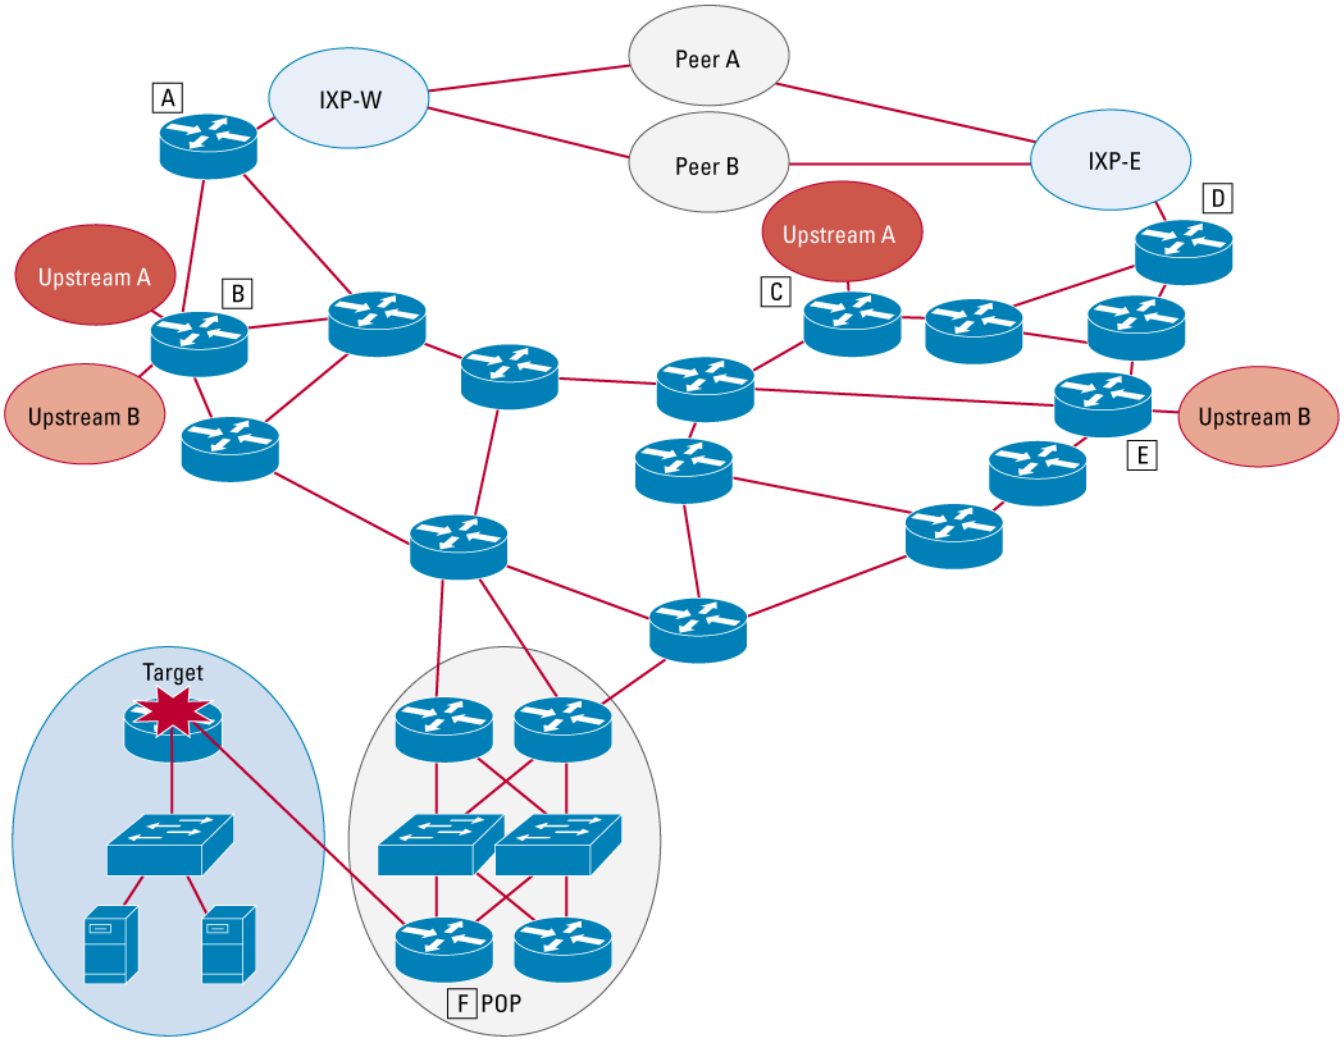

UNICAST RPF ALGORITHM

In the original uRPF code, the following command syntax was used:

```
ip verify unicast reverse-path <acl>
```

The original uRPF check logic for the above statement is the following:

```
    if the source address best path for a prefix is via the source interface
          pass the packet
    else
          if the source is 0.0.0.0 and destination is 255.255.255.255  /* BOOTP and DHCP */
             pass the packet
       else if destination is multicast
             pass the packet
```

```
      else if packet matches <acl>
            pass the packet
      else
            drop the packet
```

The new uRPF check logic for the loose mode implementation submitted through DDTS CSCdr93424 is the following:

```
look up source IP address
if entry found
    if ignore-default specified and entry is default route
         do drop logic & return
    if source of packet is different from FIB entry
        if exist-only specified  /* loose mode */
           count suppressed drop
      else
           do drop logic & return
    pass packet & return
else
      do drop logic & return

The drop logic is:

  if special addresses (e.g. DHCP/BOOTP)
       pass packet
  else if ACL permit
       count suppressed drop
       pass packet
  else
      count drop
      drop packet
```

NEW UNICAST RPF ENHANCEMENTS AND COMMAND-LINE INTERFACE

The new commands, enhancement, and fixes that were added via DDTS CSCdr93424 are as follows:

- **New mode of operation: “exists only” or “loose mode.”** In this mode, a source address need only be present in the FIB table and be resolved and reachable via any “real” interface.
- **Close ping DoS hole.** The original uRPF verification check allowed the router to ping its own interface. This created a potential DoS hole. The default behavior of URPF has been changed such that now, you must include the “allow-self-ping” option in the command to enable the router to ping its own interface. Without this option, ping packets generated by the router failed the RPF verification check. Caution should be used when enabling this feature, as this opens a potential DoS hole.
- **Allow secondary address pings.** A bug in the self-ping hole prevented the router from pinging a secondary address. This is fixed. Note that you must use the new allow-self-ping flag to make this work.

- **New command syntax.** A new, extendable syntax is used to support the new loose mode operation. It is as follows:

```
ip verify unicast source reachable-via (rx|any) [allow-default] [allow-self-ping] [<list>]

no ip verify unicast
```

Normally, sources found to be present in the FIB but only by way of the default will be dropped. The "allow-default" flag will override this default behavior and allow the lookup to match the default route and use it for verification. Note that this is current behavior, which matches (implicitly) with the old command format (see below):

```
ip verify unicast reverse-path [allow-self-ping] [<list>]
```

Note that the old command still works, and it is not converted to the new syntax (although it could be as "`ip verify unicast source reachable-via rx allow-default`").

These new enhancements allow uRPF to perform the following key functionality:

1. MTRIE checks on the source address to see if the route is in the FIB. If it is not, then the packet is dropped. The result is the elimination of any packet whose source is spoofed or from a restricted prefix (e.g., RFC1918 prefixes, or unallocated prefixes (those not allocated by the RIRs and reserved by IANA)). This will work on any interface, allowing an ISP border router with multiple links to multiple ISPs to have limited uRPF capability.
2. If the adjacency of a route equals Null0, then the packet is dropped. Unicast RPF will only pass a packet if the entry in the FIB points out of any real interface. This allows you to add a null0 route and cause packets to be dropped based on source IP address, since the FIB entry will point to the null interface descriptor block (IDB), which is not "real." "Real" interfaces do include loopback and tunnels.

USING THE UNICAST RPF LOOSE CHECK ENHANCEMENT AS A RAPID REACTION TOOL FOR DENIAL-OF-SERVICE/DISTRIBUTED-DENIAL-OF-SERVICE ATTACKS

When DoS/DDoS attacks occur, they can come from many sources. ISPs need the ability to rapidly apply filters on the edge of the network that will:

- Drop packets based on the source IP address.
- Be selective—marking valid packets as well as invalid packets.*
- Prevent frequent ACL updates on every edge router on the network. Filters should be passive and nonintrusive to the performance of the system. All routers are triggered at once. Applying ACLs to hundreds of routers is an operationally expensive process.

The objective is to empower the ISP to identify prefixes originating DoS/DDoS attacks and advertise these prefixes via BGP to all routers with preset filters as an operationally efficient mechanism for dropping attack packets at the edge of the ISP network based on source address (that is, pushing the problem to the edge of the network).

* Invalid packets are not part of a contiguous Transmission Control Protocol (TCP), Internet Control Message Protocol (ICMP), or User Datagram Protocol (UDP) flow. They are fragments or parts of a packet sequence.

Figure 5. URPF as a Rapid Reaction Tool

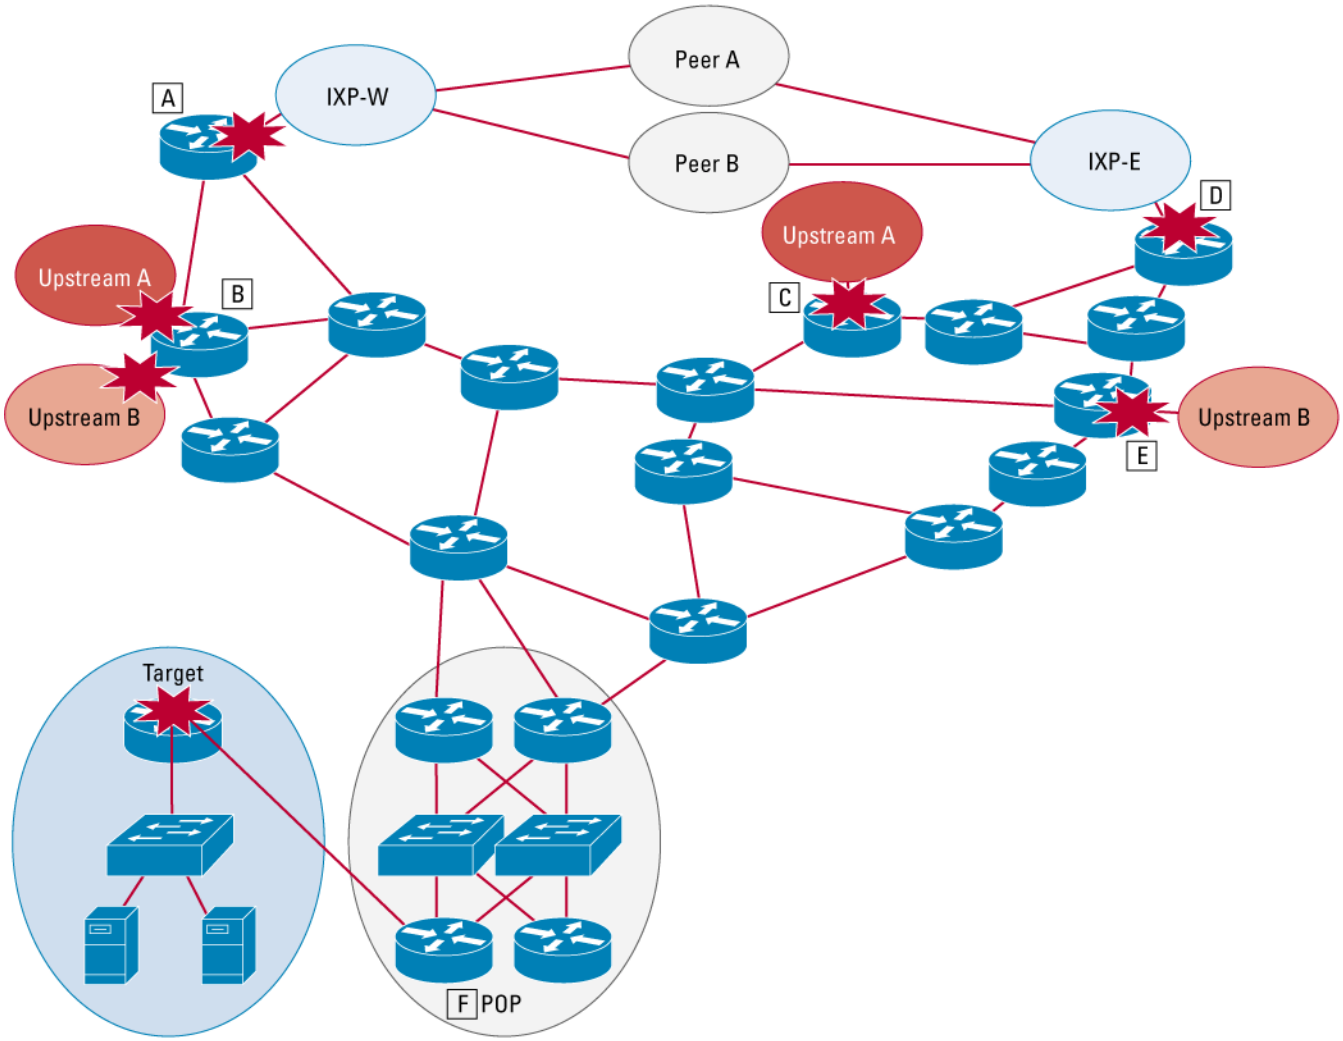

Motivation and Priority

DoS attacks are an increasing risk to the Internet. ISPs need passive tools that would make it more difficult to implement an attack and reactive tools to mitigate the effects of attacks when they happen. Maintaining anti-DoS ACLs on the edge of an ISP network is an operationally expensive process. Unicast RPF takes a different approach. It builds on the following assumptions of the current operational environment:

- Attacks are normal day-to-day events for an ISP.
- Multiple attacks against multiple customers of the ISP are normal.
- Attacks can shift in their character and type throughout the life cycles of the incident.
- Multiple updates will be required to the drop list through the life cycles of the incident.
- The drop list can be hundreds of lines long with frequent changes to respond to evolving or new incidents.
- ISPs will weigh the risk and take down (that is, drop all packets from) a specific source in order to mitigate the effects of the incident on the network of the customers or themselves

The drop list is the list of specific /32 prefixes that need to be dropped at the edge of an ISP network based on the source or destination of the packet.

Example of the Rapid Reaction Tool in Operation

One example of how this Rapid Reaction Tool can be used on the ISP-ISP edge of the network is via the use of a "tagged" route distributed through the ISP network. The route is tagged by setting its next hop to equal a specific prefix on each router. The specific prefix has a static route to Null0. With the new uRPF Null0-Check, all the traffic whose source IP address equals the "tagged" route would be dropped. The result is a way the ISP can classify the source IP addresses of the DoS packets and activate filters on the edge of the network through the insertion of a route in the BGP table.

For example, in Figure 7, router E would have uRPF applied to its ISP-ISP interfaces. BGP weight would be used with all the routers to ensure that local eBGP routes would be preferred over one advertised from another internal BGP (iBGP) source in the network. Router B would use the src-reachable-via any option (see above) to allow multiple eBGP interfaces on each router to be used. Routers A and D—while having multiple eBGP sessions—each only have one interface (a fast Ethernet interface) for all the eBGP sessions. So, uRPF would work with the BGP weight set to prefer local eBGP to iBGP.

With this new enhancement to uRPF, the ISP can preset the following in each router:

1. Configure Unicast RPF with the "'source reachable-via any" option (and without the "allow-default" option) on the interface connecting to the peering ISP or IXP.
2. Configure a static route for 192.0.2.1/32 to Null0 (that is, IP route 192.0.2.1 255.255.255.255 null0).

The 192.0.2.1/32 network is used as the next-hop specification; 192.0.2.0/24 is a network reserved by IANA for testing and documentation. It should not be routed on the Internet, so no valid packets with that destination address should be forward through the Internet. Adding a static route for 192.0.2.1/32 to Null0 will not harm any valid flows. This static route will be used to glue the Null0 adjacency to the prefix that needs to be dropped. For example, if 171.68.1.0/24 is the source of an ICMP "smurf" attack, we want to glue that prefix to Null0 without logging into all of our routers and adding statics to each one. A BGP advertisement is created on one router with the no-export BGP community (so the route does not get advertised outside the autonomous system). The next hop of that prefix is set to 192.0.2.1. When each edge router receives this advertisement, each router will glue 171.68.1.0/24 to a next hop of Null0 (see Figure 6).

Figure 6. Gluing the Network You Want to Drop to Null0

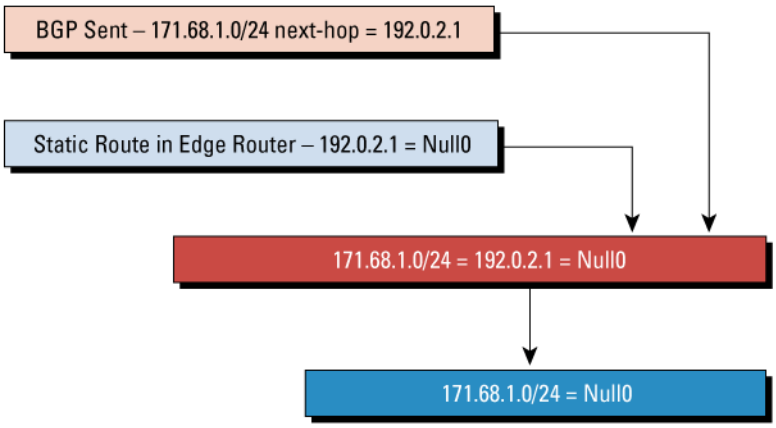

While finding the source addresses of the attacks is a topic for another paper (or an intrusion detection system [IDS] tool), once the source IP addresses of a specific DoS/DDoS attack are known, they are entered centrally via BGP at a convenient router or BGP-speaking workstation (gated, zebra, etc.) in the network operations center (NOC; router G in the example used in Figure 7). Since uRPF uses the FIB, all source addresses generating the attack can have their next hop set to an address in 192.0.2.1/32. For example, on a Cisco router, this would be accomplished via the network statement route map:

```
! Trigger Router Configuration…
router bgp XXXX
!
 redistribute static route-map Set-Next-Hop-To-TESTNET
!
route-map Set-Next-Hop-To-TESTNET permit 10
  set ip next-hop 192.0.2.1
  set community no-export
route-map Set-Next-Hop-To-TESTNET permit 20
!
ip route 192.0.2.1 255.255.255.255 Null0 250

interface Null0
  no ip unreachables
```

and

```
! Edge Router Configuration…

ip route 192.0.2.1 255.255.255.255 Null0 250

interface Null0
  no ip unreachables
```

For example, assume one wanted to drop all packets with a source IP address of 192.168.2.1/32 using URPF. On the trigger router then, one would apply a static route defining the next-hop for 192.168.2.1 as being 192.0.2.1. The result is that all routers on the edge of the network will have the following FIB/adjacency path for 192.168.2.1/32:

```
Edge_Router#sh ip route 192.168.2.1
Routing entry for 192.168.2.1/32
  Known via "bgp 100", distance 200, metric 0, type internal
  Last update from 192.0.2.1 00:09:25 ago
  Routing Descriptor Blocks:
  * 192.0.2.1, from 30.1.2.1, 00:09:25 ago
      Route metric is 0, traffic share count is 1
      AS Hops 0, BGP network version 4
```

```
Edge_Router#sh ip route 192.0.2.1
Routing entry for 192.0.2.1/32
  Known via "static", distance 1, metric 0 (connected)
  Routing Descriptor Blocks:
  * directly connected, via Null0
      Route metric is 0, traffic share count is 1

Edge_Router#sh ip cef 192.168.2.1
192.168.2.1/32, version 40
0 packets, 0 bytes
  via 192.0.2.1, 0 dependencies, recursive
    next hop 192.0.2.1, Null0 via 192.0.2.1/32
    valid null adjacency
```

Figure 7. Using Null0 as a Rapid Reaction Tool

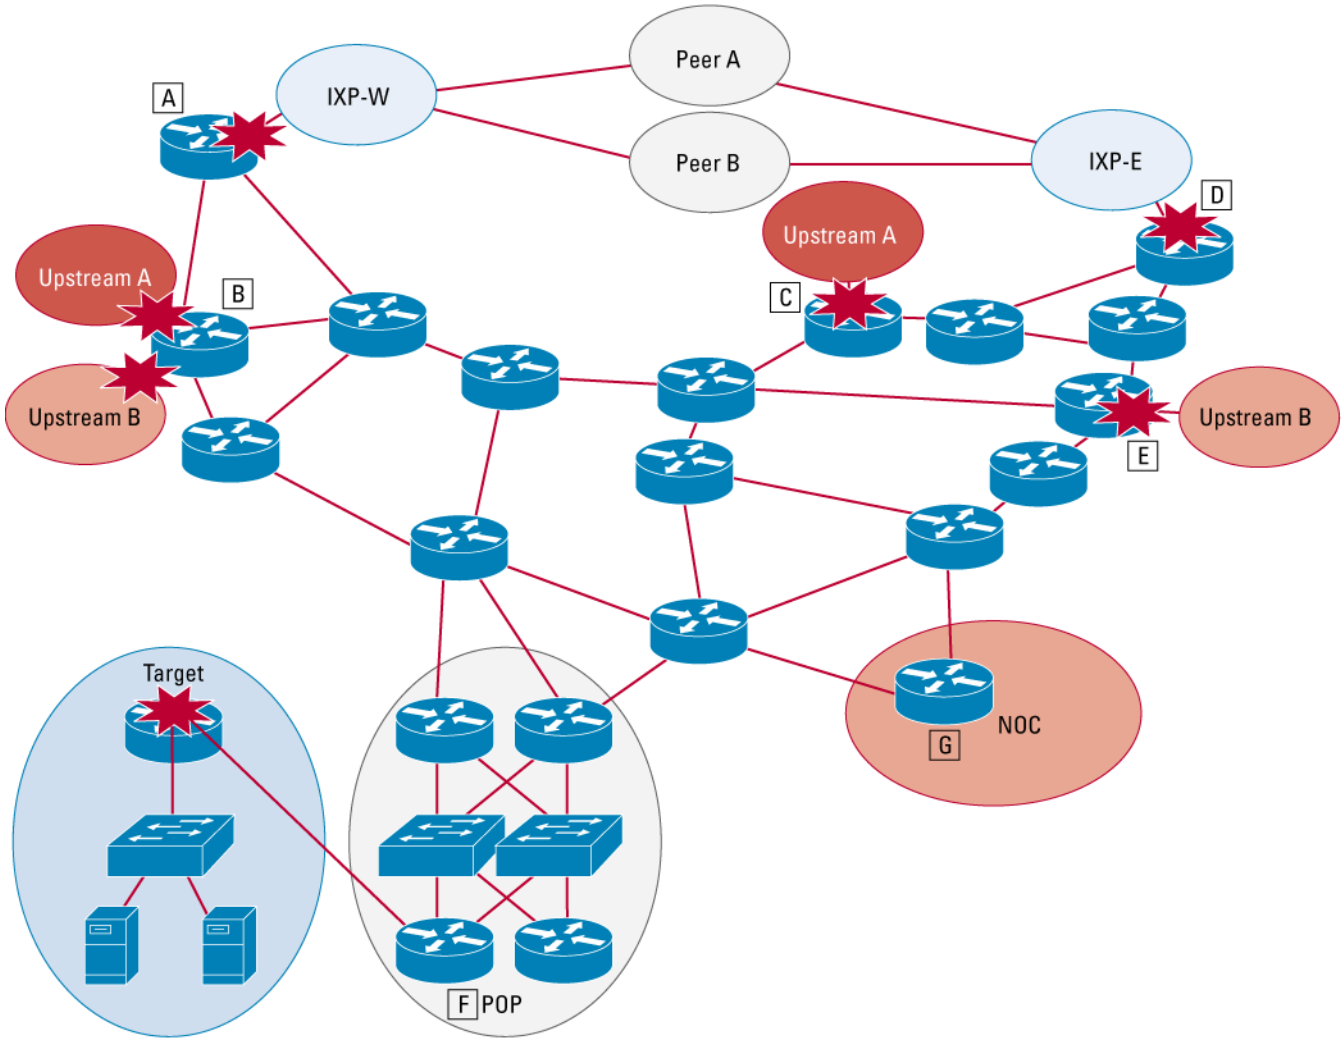

It is critical that the adjacency of 192.168.2.1/32 be set to Null0. Without uRPF applied, all packets with a destination address of 192.168.2.1/32 would be dropped via Null0. However, with uRPF and the Null0 check applied, all packets with **a source OR destination** equal to 192.168.2.1/32 would be dropped. In summary, this Unicast RPF enhancement allows the ISP to update its drop list for IP source and destination across its entire network, all based on a BGP routing update. Of course, since it does use BGP for the updates, when sources of the DoS/DDoS attacks shift in the middle of the incident, updates are just a matter of injecting more BGP advertisements, triggering more source addresses to be dropped on the edge.

Why Use Null0 and Not a Loopback Interface?

A router with Cisco Express Forwarding switching turned on handles packets sent to Null0 and the loopback interface differently. Null0 is considered to be a special Cisco Express Forwarding adjacency. Any packets with a next hop to Null0 will be dropped in the Cisco Express Forwarding path—on the line card—or via the ASIC. So packets that are black-holed to Null0 will have no performance impact when they are dropped. Loopback interfaces are valid virtual interfaces. Packets sent to the loopback interface are forwarded like any other packet bound for an interface. For routers such as the 7500 and GSR, the loopback interfaces are on the RP/GRP. Packets black-holed to a loopback will be sent from the line card to the RP/GRP to be processed by the interface. In some cases these are processed switched packets. In any case, packets that

are black-holed to a loopback interface with Cisco Express Forwarding switching turned on will have some performance impact on the router. The extent of the impact depends on the platform. It is recommended that any black-hole filtering techniques such as the uRPF one listed here use the Null0 interface.

WORKING WITH INTRUSION DETECTION SYSTEMS

The ability to trigger packet drops via an iBGP routing advertisement creates new options for new IDSs to integrate with the operations of an ISP. Currently IDS systems identify and classify attacks, then alert the operations team. They can then create ACLs to be uploaded to routers that will drop or rate-limit the attack. The new Unicast RPF enhancement would allow these tools to interface and update drop lists via BGP vs. trying to update the ACLs on hundreds of routers.

One of the key issues of deploying any type of IDS system is its ability to respond to attacks. Until now, ACLs have been created based on the information gathered from the IDS sensors in the network. These ACLs are then updated to the routers on the edge of the network. If the number of routers is small—say, one or two—then the ACL updates are straightforward. However, if the number of routers that need ACL updates is large—say, hundreds—then the ACL updates get very tricky and time consuming. In addition, updating ACLs on a live router during prime time traffic flows is a risk. Adding to this, the fact that several attacks may be occurring simultaneously results in complex and long ACLs, and then the ACL technique starts to get very cumbersome. What is needed is a technique where the IDS tool can trigger a packet filter without logging into the router. This new Unicast RPF technique provides that option.

Think of an IDS tool with gate D running and an iBGP peering session to one of the routers in the ISP network (see Figure 8). When the IDS sensors trigger an alarm, the following could happen:

1. An IDS tool alerts the NOC of a potential DoS/DDoS attack. The IDS tool posts the list of IP addresses that are generating the DoS/DDoS attack.
2. The IDS tool recommends dropping the attack. It creates a BGP advertisement to trigger uRPF to drop the attacks at the edge of the ISP network. The IDS tool waits for human approval.
3. The NOC team reviews the IDS recommendation and approves the BGP advertisement to drop the attack at the edge.
4. The IDS tool continues to monitor the attack.

Figure 8. IDS Systems Providing Feedback on the Attack to a Central NOC

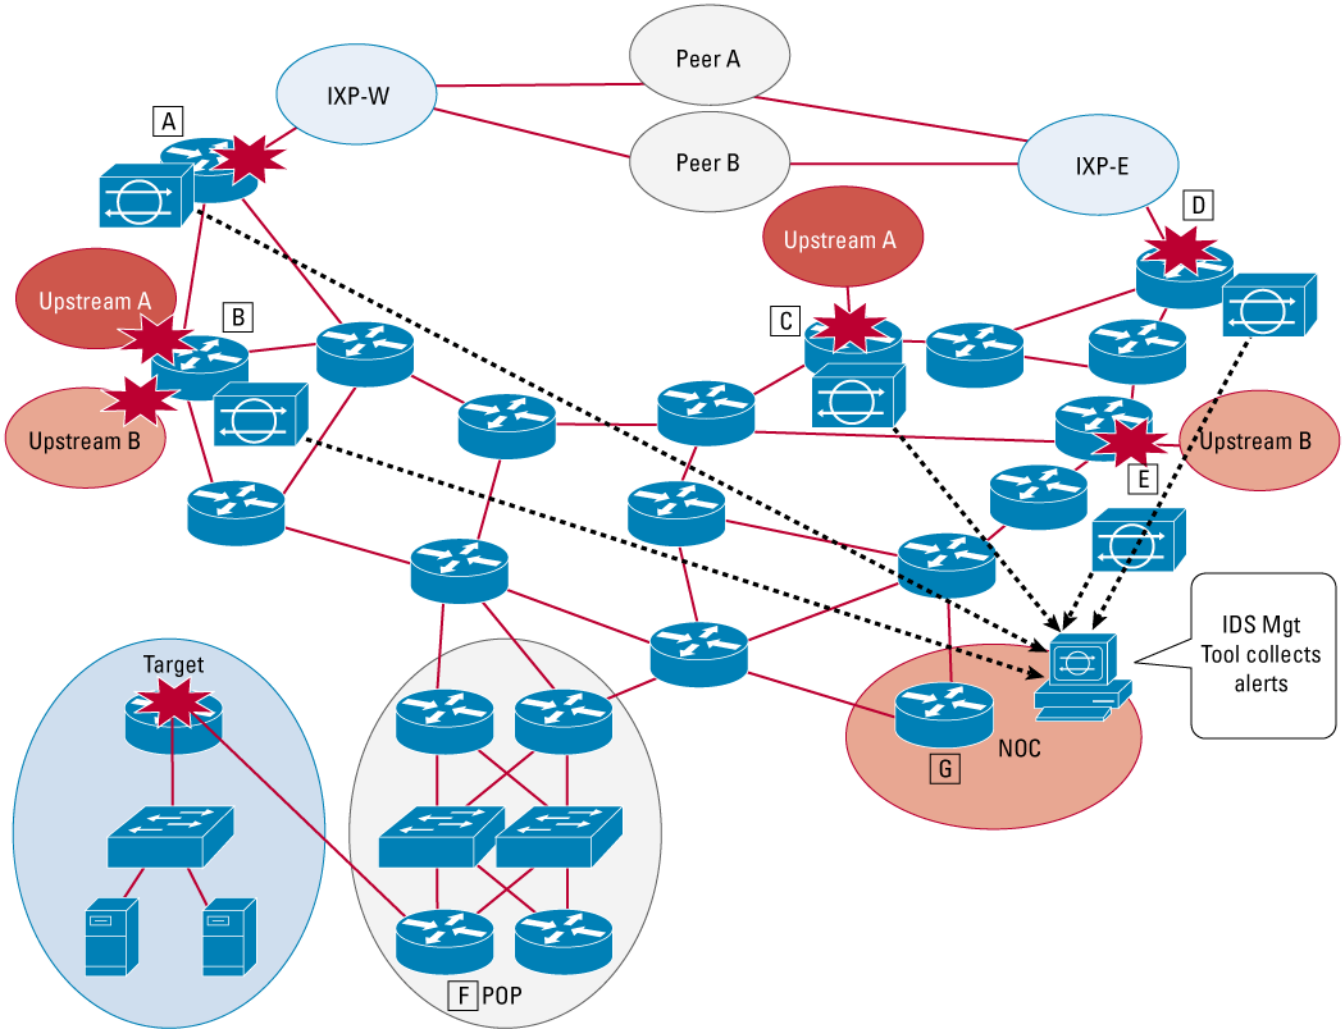

The IDS management tool generates a list of prefixes that would need to be dropped to fend off the attack. A BGP network advertisement is created to set the next hop for these prefixes to that of the shunt to Null0. Since the BGP advertisement will be sent to all the BGP speaking routers on the network, all the routes on the edge of the network would receive the update at the same time (see Figure 9). New attacks would be handled with additional advertisements. The result is a rapid response system integrated with IDS tools deployed throughout the network.

Figure 9. NOC Authorizes an iBGP Announcement to Trigger a Packet Drop

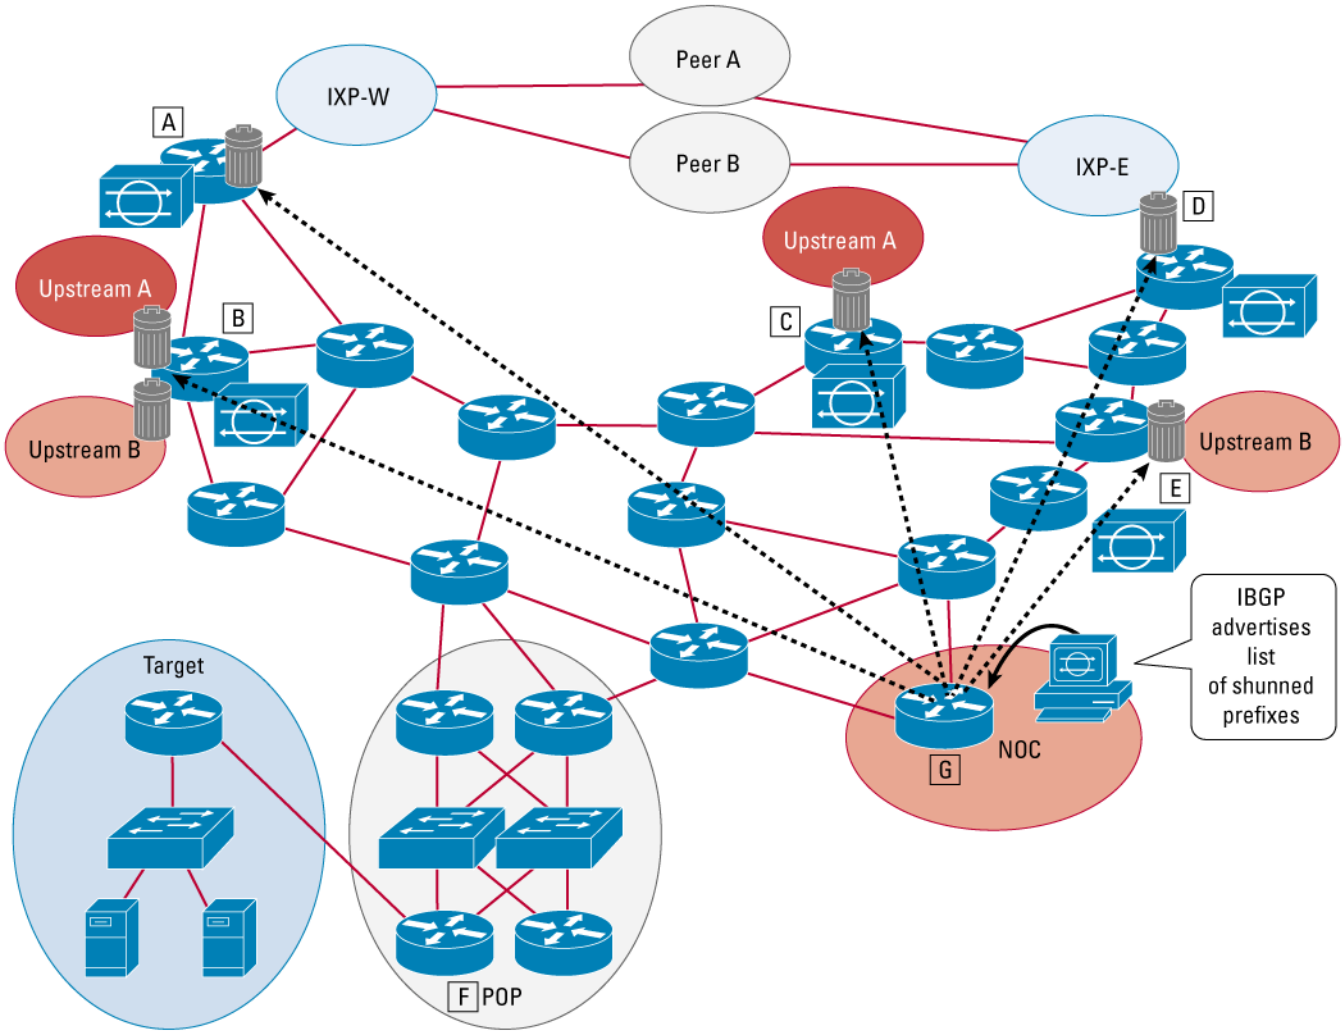

PERFORMANCE EXPECTATIONS

Minimal performance impact on the router is inherent in the design of uRPF. Since uRPF uses the same MTRIE lookups as the forwarding/switching process, the performance impact on the router and/or line card is expected to be minimal. Variance will occur in different ASIC architectures. For example, the GSR Engine 0, Engine 1, and Engine 2 line cards all have slightly different ASIC architectures. Hence, the performance impact with uRPF would be slightly different on each line card. Regardless, as shown in the lab test on an Engine 0 line card, 3 percent additional CPU while identifying and dropping packets (while under attack) is considered a minimal performance impact.

The lab test had four GSRs with Engine 0 line cards connected in a ring topology. The four routers are located in two different autonomous systems operating under hot potatoes routing. This architecture was used to simulate typical asymmetrical routing encountered on the Internet. This is a typical scenario in which two autonomous systems peer at two different exchange points that hand off the traffic at the closest exit, creating asymmetric routed traffic. The setup is shown in Figure 10.

Figure 10. Lab Test Architecture

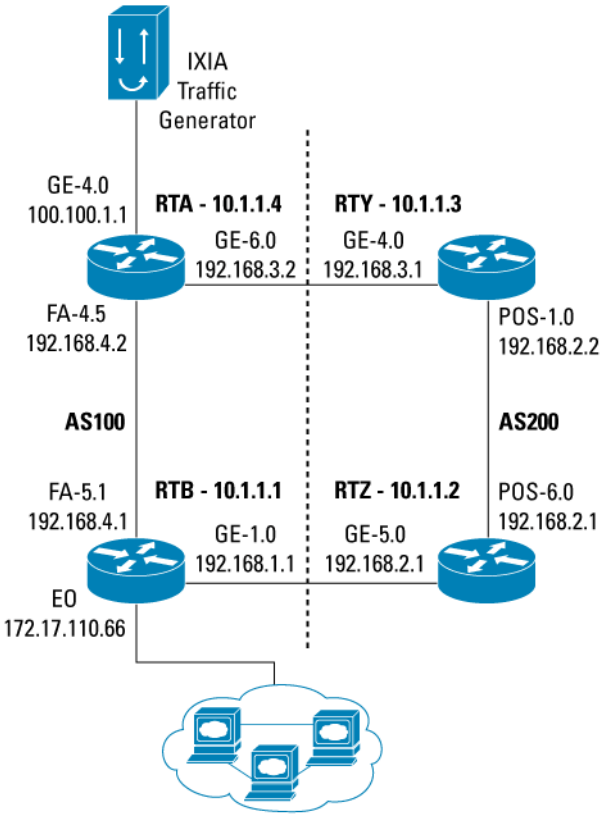

Lab Results

1. Normal operation without "ip verify unicast source reachable via any" on router RTY. Traffic destination is router RTZ. The result is obtained from the line card of intermediate router RTY. There is virtually no change to the GRP CPU utilization, hence the result is not recorded.

Attack Stream	Interface	Packet Size	CPU of LC	Packets per Second
0 Mbps	Gigabit 4/0	64 bytes	0%	0
30 Mbps	Gigabit 4/0	64 bytes	52%	61,000
50 Mbps	Gigabit 4/0	64 bytes	57%	100,000
100 Mbps	Gigabit 4/0	64 bytes	88%	210,000
160 Mbps	Gigabit 4/0	64 bytes	98%	330,000

2. Unicast RPF is turned on with spoofed source and valid destination address to router RTZ. The result is obtained from the intermediate router RTY line card.

Attack Stream	Interface	Packet Size	CPU of LC	Packets per Second
0 Mbps	Gigabit 4/0	64 bytes	3%	0
30 Mbps	Gigabit 4/0	64 bytes	55%	61,000
50 Mbps	Gigabit 4/0	64 bytes	60%	100,000
100 Mbps	Gigabit 4/0	64 bytes	90%	210,000
160 Mbps	Gigabit 4/0	64 bytes	100%	330,000

DEPLOYMENT OPTIONS WITH THE NEW ENHANCEMENTS

As mentioned throughout this paper, the new uRPF enhancements have created new deployment options for service providers. The following table reviews some of these deployment options.

Deployment Situation	Type of uRPF to Use	Config Notes
Leased Line Customer	Strict check	
Multihomed Leased Line Customer (Same ISP)	Strict check or loose check	Remember to use BGP weights on strict check
Multihomed Leased Line Customer (Different ISPs)	Strict check or loose check	Remember to use BGP weights on strict check
Dialup Customers	Strict check	
DSL Customers	Strict check	
Cable Modem Customers	Strict check	
IXP Connection—No Private Peering	Strict check	
IXP Connection with Private Peering	Loose check	
Private Peering—Dedicated Router	Strict check	Symmetry should be expected between the routes advertised and source addresses sent by the peering ISP
Private Peering with Several ISPs on the Same Router	Loose check	
Edge Router of Colocation Provider	Loose check	

REFERENCES

1. "Cisco ISP Essentials" white paper: www.cisco.com/public/cons/isp/documents/IOSEssentialsPDF.zip
2. "Cisco IOS Essentials – Best Practice" Power Session (presentation): www.cisco.com/public/cons/isp/documents/IOSEssentials_Seminar.zip
3. Unicast Reverse Path Forwarding Enhancements—Cisco IOS Software Release 12.1T documentation: www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121t/121t2/rpf_plus.htm
4. Cisco IOS Software Release 11.1CC Unicast RPF documentation: www.cisco.com/univercd/cc/td/doc/product/software/ios111/cc111/uni_rpf.htm
5. Strategies to Protect Against Distributed Denial of Service (DDoS) Attacks: www.cisco.com/warp/public/707/newsflash.html
6. Improving Security on Cisco Routers: www.cisco.com/warp/public/707/21.html
7. Cisco IOS Software Release 12.0S Unicast Reverse Path Forwarding Commands (original version): www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/secur_r/srprt6/srrpf.htm
8. Cisco Catalyst® 6000 Series Switch Support for the Original uRPF—release notes for MSFC Cisco IOS Software Release 12.1E: http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/relnotes/ol_2309.htm
9. Security Overview—Cisco IOS Software Release 12.1T documentation: www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/secur_c/scdoverv.htm
10. Cisco IOS Software Release 12.1 documentation—Configuring Unicast Reverse Path Forwarding (original version): www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/secur_c/scprt5/scdrpf.htm

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on **the Cisco Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

204170.h_ETMG_AE_2.05

Printed in the USA