

Highlights

- New additions to 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.

Unicast Reverse Path Forwarding (uRPF) Enhancements for the ISP-ISP Edge

Version 1.5

Contacts: Barry Raveendran Greene bgreene@cisco.com and Neil Jarvis njarvis@cisco.com
General questions on uRPF can be sent to unicast-rpf@cisco.com or cisco-nsp@puck.nether.net

Introduction

Unicast RPF is 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, D. Senie. As such, Unicast RPF was designed for the customer ↔ ISP edge. The objective was a feature that can be easily automated in the customer provisioning system, scale as new addresses blocks were allocated to the customer, and work with the MTRIE based CEF switching. Unicast RPF meets these objectives, even in situations where the customer was multihomed to one or more upstream ISPs.¹ Originally implemented in 11.1(17)CC, uRPF provided a new ISP Security tool for BCP 38/ RFC 2827 deployment.

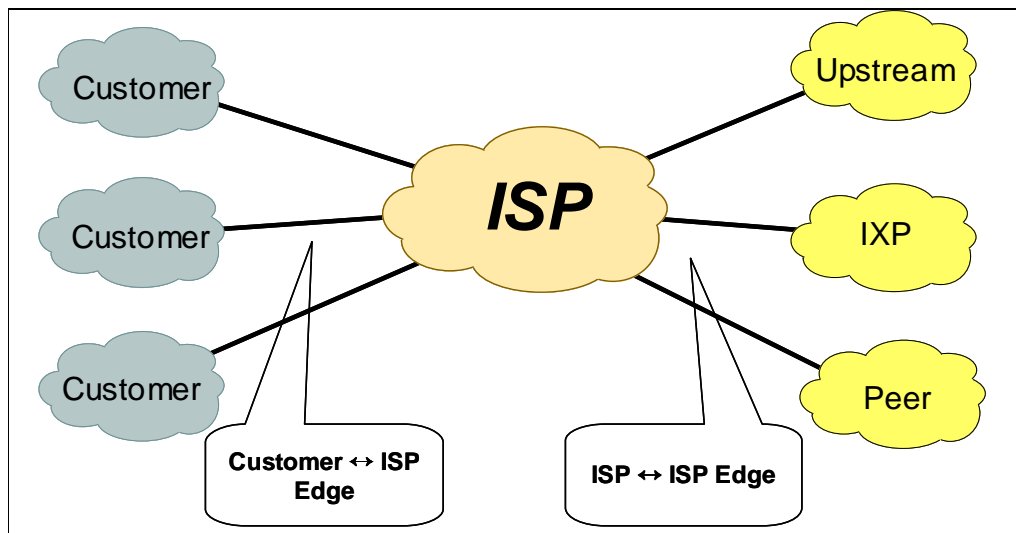


Figure 1 - Original uRPF deployment was on the Customer ↔ ISP Edge.

¹ Unicast RPF does work with asymmetrical routing on the Customer ↔ ISP Edge. Detailed configurations and a explanation of the myth that uRPF does not work with asymmetrical routing is details in ISP/IOS Essentials at <http://www.cisco.com/public/cons/isp/documents/IOSEssentialsPDF.zip>.

Sunday, February 18, 2001

Over time and with an increase in DDOS attacks on the Internet, uRPF's functionality was reviewed as one of several ways ISP Security tools that could be implemented the ISP ↔ ISP. The forwarding/routing characteristics on the ISP ↔ ISP are vastly different from the forwarding characteristics. For example, Customer ↔ ISP edge has relatively symmetrical flows (excluding some types of multihomed configurations). So uRPF can use the best single path selection the RIBs send to the FIB. On the other hand, the ISP ↔ ISP can have several connections to ISPs. Each ISP connection would be exchanging BGP routing information. The best single path selection the BGP RIB sends to the FIB will create issues where the FIB will not match the packet flows from the ISP connections. Unicast RPF will not work if the FIB and the packet flow does not match on an interface. Hence, uRPF's originally implementation is not an option for the ISP ↔ ISP edge of a network.

Yet, if uRPF could be modified to work on the ISP edge, it would provide the ISPs with a new DDOS resistance tool. For example, I uRPF would check the source addresses of each ingress packet inbound from another ISP, it would be able to automatically detect and drop packets like:

- RFC 1918 source addresses.
- Other DUSA² addresses that should not appear in the source.
- Un-Allocated addresses that have not been allocated by the RIRs.
- Sources addresses that do not match the BGP advertisement of the peering ISP.

Each of these source addresses either should not be traversing the Internet or are traversing the wrong path. Identifying and dropping these packets on the inter-ISP border is considered to be a way to increase resistance against all sorts of security attacks. For this reason, uRPF was reopened and reviewed to find ways to have it 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. The first phase of the enhancement will filter the first three of four listed items above. The final spoofed source is left to future uRPF enhancements.

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*.³ The second phase – committed through DDTS CSCdr93424 – will allow a looser uRPF check and fix some security holes. This will allow uRPF to be used on the ISP ↔ ISP edge of the network. The compelling functionality this second phase adds to uRPF is an alternative to standard ACLs to drop packets on the edge of the network (see below). The third phase will create a way to have strict enforcement of the uRPF check on the ISP ↔ ISP edge. It is hoped that an eBGP peer session can be send to a dedicated VRF table. This will allow uRPF to query the VFR table that contains all the routes for the eBGP peering session over the interface. This will allow uRPF to verify is the source address of a packet matches the advertised routes from the peered ISP. Together, it is hoped that these three phases will provide ISPs tools to increase the difficulty to perpetrators of security incidences on the Net.

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 will remove the match on the specific interface – allowing uRPF to “loose” check. This would allow an ISP's peering router with multiple links to multiple ISPs to check source packets to see if they exist in the FIB. If they exist, then they pass. If they do not exist, they drop. This would build resistance against DOS/DDOS attacks that use spoofed source addresses based on RFC1918, Martin, and un-allocated IP addresses.⁴

² Documented Special Use Addresses 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/internet-drafts/draft-manning-dsua-06.txt>)

³ 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* whitepaper and presentation located at <http://www.cisco.com/public/cons/isp/documents/>.

⁴ Un-allocated IP addresses are those IANA reserved addresses that have not been delegated by the Regional Internet Registries (RIRs)

Sunday, February 18, 2001

The Unicast RPF enhancement will also allow for a new tool to drop packets based on BGP updates vs ACL updates. The hypotheses being that DOS/DDOS attacks are dynamic – changing their character over the time of the incident. Hence, several ACL updates will be needed over the incident period. This new enhancement will allow these updates to be propagated by BGP to the edge of the network – triggering packet drops on the DOS/DDOS packets.

DDTS CSCdr93424 has been committed to 12.0(13.06)S01 and will be committed to 12.1E for the Cat6K/OSR support. 7200, 7500, GSR Engine 0, and GSR Engine 1 are supported in the first CCO published version – 12.0(14)S. GSR Engine 2 support is scheduled around 12.0(17)S.

Objectives for the Unicast RPF Enhancement

Unicast RPF (uRPF) originally was designed to prevent source address spoofing between 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 was multihomed to the ISP or multiple ISPs. Unicast RPF will also work on links to IXPs (for example Routers A and D).⁵ What does not work is if uRPF was applied on routers with multiple connections to multiple ISPs.

The proposed additions to uRPF are 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 would use BGP to trigger filters on the edge of an ISP's network - shut down attacks based on the source and destination IP address.

Unicast RPF between ISPs – What is the problem?

Currently uRPF will not work between typical ISP router where there are *multiple* ISP peers over multiple interfaces. The common reason for why it will not work on the ISP ↔ ISP edge is '*asymmetrical routing*'. The problem is that 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 router's FIB to validate the reverse path of the packet. This allows uRPF to use the same optimized MTRIE look-ups to do its validation. At the same time, the use of the FIB is the core in-flexibility of uRPF's deployment in a network. The RIB's and FIB's best path selection algorithms will select one best path. There might be more than one *best paths*, but the way routing protocols and forwarding 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 Unicast RPF will not work on the ISP ↔ ISP edge. Comprehending the crux of this is Internet limitation is key to understanding the deployment limitations of the original Unicast RPF. Once one grasps that the limitation is not *asymmetrical routing on the Internet*, but the *best path selection* of how routing/forwarding works on a router, new Unicast RPF deployment options are created for the ISP. Some of those new options include deployment options with the original uRPF with multihomed customers on the Customer ↔ ISP and special ISP ↔ ISP peering options – like routers connected to Internet eXchange Points (IXPs).

⁵Check out the "IOS Essentials" whitepaper for examples of uRPF with multihomed customers:

ISP Essentials Whitepaper <http://www.cisco.com/public/cons/isp/documents/IOSEssentialsPDF.zip>
ISP Essentials Seminar Slides http://www.cisco.com/public/cons/isp/documents/IOSEssentials_Seminar.zip

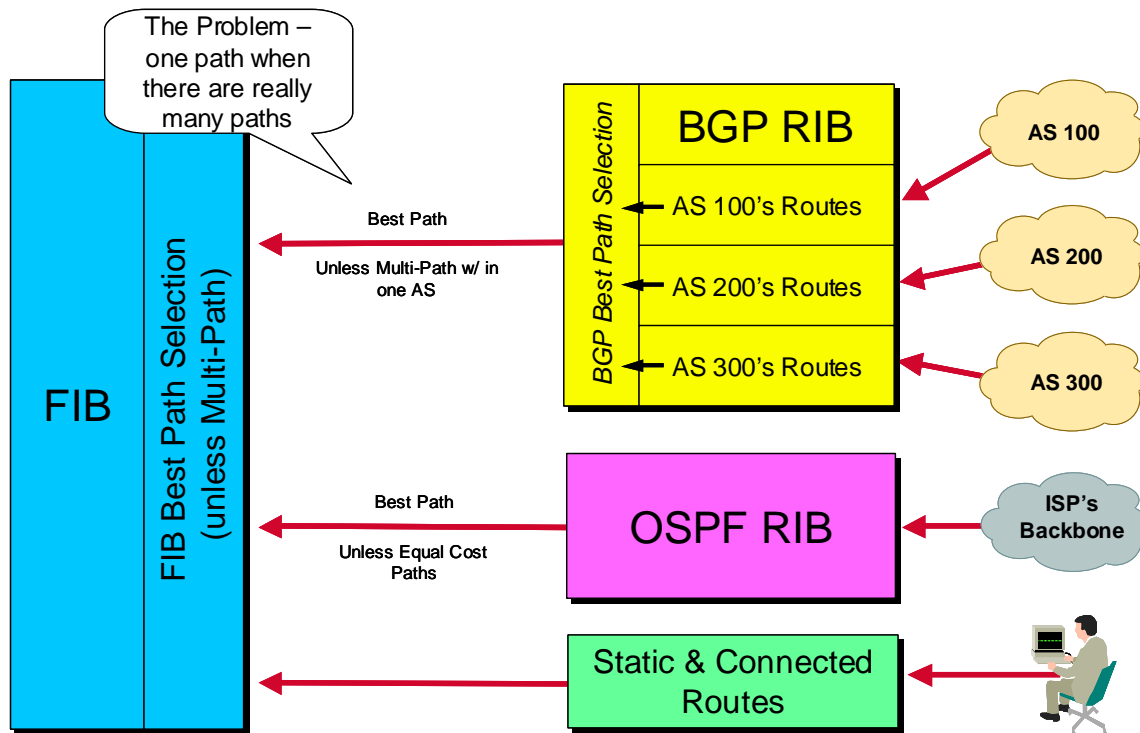


Figure 2 - uRPF's Limitation on the ISP ↔ ISP edge has more to do with RIB/FIB best path selection than asymmetrical routing.

An IXP is a good example to demonstrate the original uRPF's limitation. In Figure 3, the IXP Border router has multiple ISP peers, but they are all over one interface - the interface to the IXP's switch. BGP Weight 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, Unicast RPF 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 presents a illustrates what happens when multiple valid paths are feed to the router, yet the routing protocol and forwarding algorithm only allow one best path. Here there 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 forward table. As a result, uRPF checks would fail on a valid packet sent from an ISP that is also advertising that route (since BGP has picked another provider's route as the *best path*). Again, the problem with uRPF on the ISP ↔ ISP edge has more to do with the character of how BGP works with forward 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.

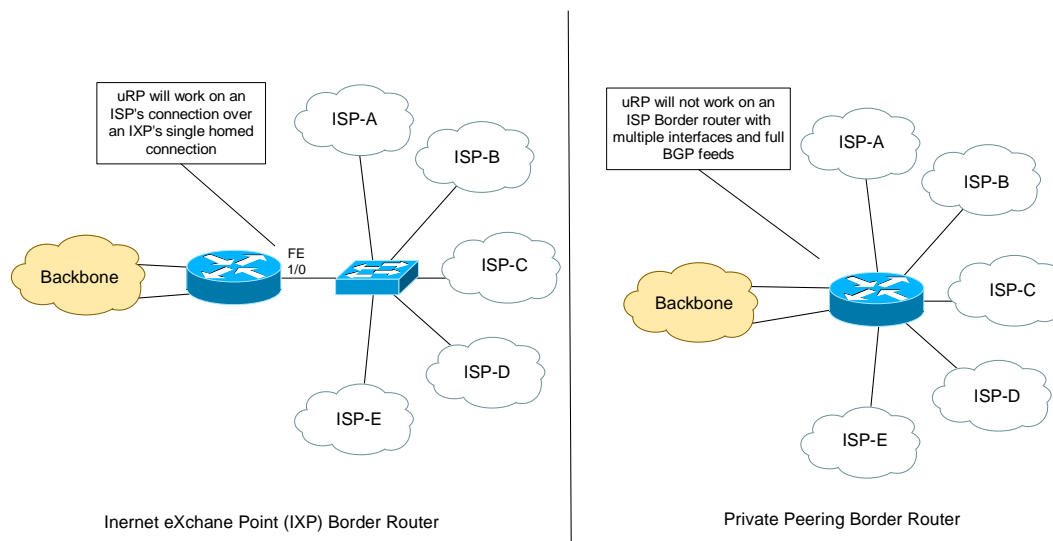


Figure 3 - Routers on the ISP-ISP Edge of the Network

Another example that demonstrates how this uRPF will work is on Routers C and E in Figure 4. These routers are single homed between the two ISPs. The BGP weight command would be used to insure 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 will 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 forward 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, uRPF will drop packets that should not be dropped.

As mentioned before, it is desirable to have Unicast RPF work in all ISP ↔ ISP edge scenarios. To do this, we looked a various enhancements that could be easily added to uRPF. Enhancements that would work across the ASIC families used on all the equipment sold to ISPs. The simplest technique was selected as the first enhancement. This technique would work around the limitations of the best path selection in the forwarding/routing algorithms. It does so by just 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's *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 Net. The objective with uRPF Loose Check is to give ISPs a tool to make it more difficult to cause mischief.

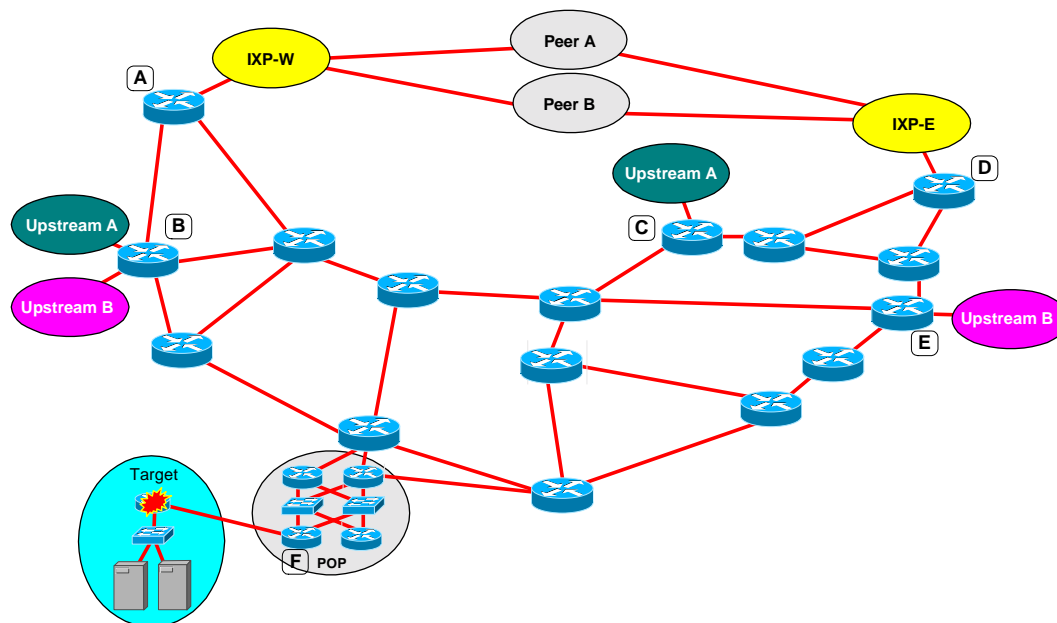


Figure 4 - ISP's Network

Existing uRPF CLI and Algorithm

In the existing Unicast RPF code, the following command syntax is used:

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

The original Unicast RPF algorithm with for the above statement is the following:

```
if the source address's 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 a 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 checking logic submitted through DDTS CSCdr93424:

Sunday, February 18, 2001

```
lookup 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
            count suppressed drop
        else
            do drop logic & return
    pass packet & return
else
    do drop logic & return
```

The drop logic is:

```
if special addresses
    pass packet
else if ACL permit
    count suppressed drop
    pass packet
else
    count drop
    drop packet
```

New Unicast RPF Enhancements, CLI, and Algorithm

The new commands, enhancement, and fixes that will be added via DDTS - CSCdr93424 :

- **New mode of operation - "exists-only".** In this mode, a source address needs only be present in the FIB table, be resolved and reachable via a "real" interface to be verified. The new command is

```
ip verify unicast source reachable-via any [allow-default]
```

The allow-default flag means allow the lookup to match the default route and use it for verification. Note this is today's behavior, so is implicit with the old command format (see below).

- **Close ping DoS hole.** There is a hole in the verification check to allow the router to ping its own interface. This is a denial-of-service hole. You must now specify allow-self-ping in the command to enable this hole.
- **Allow secondary address pings.** There was a bug in the self-ping hole, which prevented the router pinging a secondary address. This is fixed. Note 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 modes of operation. It is:

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

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

```
no ip verify unicast
```

Note that the old command works; also it is not re-nvgened in the new syntax (although it could be as 'ip verify unicast source reachable-via rx allow-default').

Sunday, February 18, 2001

- **Config vs. hardware support check.** To support implementing uRPF on hardware platforms in the future (and at different times in the future for different platforms), there is now a LOOP registry that the command code calls to verify that the uRPF configuration is supported on the specified interface. If any registered function returns TRUE (e.g. config not supported), the config attempt aborts with an error message. Part of this commit ensures that RPF in any form cannot be configured on GSR engine 2 (and later) linecards. When the functionality is added, changes can be made to the `grp_rpf_config_unsupported()` function.

This update to Unicast RPF will allow 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 gets dropped. The result is elimination of any packet whose source is spoofed from a restricted prefix (i.e. RFC1918 prefixes) and any un-allocated prefixes (i.e. not allocation by the RIRs and reserved by IANA). This will work on any interface - allowing an ISP's border router with multiple links to multiple ISPs to have limited Unicast RPF capability.
2. If the route's adjacency equals Null0, then the packet is dropped. The Null0 check is in addition to the new exists-only check to 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, since the FIB entry will point to the nullidb, which is not "real". "Real" interfaces include loopback and tunnels.

Using the uRPF Loose Check Enhancement as a Rapid Reaction Tool for DOS/DDOS Attacks

When DOS/DDOS attacks happen, they can come from any direction. 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.⁶
- Prevent frequent ACL updates on every edge router on the network. Filters should be passive and non-intrusive to the performance of the system. All routers are triggered at once. Applying ACLs to hundreds of routers is an operational hindrance.

The objective is to empower the ISP to identify the networks originating the DOS/DDOS attacks, advertise the networks via BGP to all routers with the pre-set filters, and start dropping the packets based on the source address at the edge of the ISP's network (i.e. pushing the problem to the edge of the Network).

⁶ Invalid packets are not part of a contiguous TCP, ICMP, or UDP flow. They are fragment or parts of a packet sequence.




Figure 5 - uRPF as a Rapid Reaction Tool

Motivation and Priority

Denials of Service (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 the attack(s) when they happen. Maintaining *anti-DOS ACLs* on the edge of an ISP's network is an operational hindrance. The uRPF takes a different approach. It builds on the following assumptions of today's operational environment:

- Attacks are normal day to day events on an ISP.
- Multiple attacks against multiple customers of the ISP are normal.
- Attacks can shift in their character and type through out 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 weight the risk and take down (i.e. drop all packets) a specific source in order to mitigate the effects of the incident on their customer's or their own network.

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

Once reviewed, these additions to uRPF will be a priority featurette for the 12.0S and 12.1E code trains (covering the 7200, 7500, ESR, GSR, and 6500/OSR).

Example of the Rapid Reaction Tool in Operation

One example of how this can be used on the ISP – ISP edge of the network is via the use of a “tagged” route distributed through the ISP’s 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, the all 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.

Sunday, February 18, 2001

For example, in Figure 7 Router E each router would have uRPF applied to their ISP – ISP interfaces. BGP Weight would be used with all the routers to insure that local eBGP routes would be preferred over one advertised from another iBGP source in the network. Router B would use the *exempt-interface* option (see above) to allow multiple eBGP interfaces on each router to be used. Routers A and D – while having multiple eBGP sessions – 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 'src-reachable-via any' and 'ignore-default' options on the interface connecting to the peering ISP or IXP.
2. Configure a static route for 192.0.2.1/32 to Null0.

We are using the network 192.0.2.1/32 as the next-hop trigger. 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 a destination address should be forward through the Internet. Putting a static route for 192.0.2.1/32 to Null0 will not harm any valid flows. We will use this static route to *glue* the Null0 adjacency to the network we want to drop. For example, if 171.68.1.0/24 is the source of an ICMP SMURF attack, you want to glue that prefix to Null0 with out logging into the routers adding statics to each one. So you advertise create a BGP advertisement on one router with a *local-AS* BGP Community (so the route does not get advertised out side of the AS). You set the *next-hop* of that prefix to 192.0.2.1. When each edge router receives the prefix, the router will *glue* 171.68.1.0/24 to a *next-hop* of Null0 (see Figure 6).

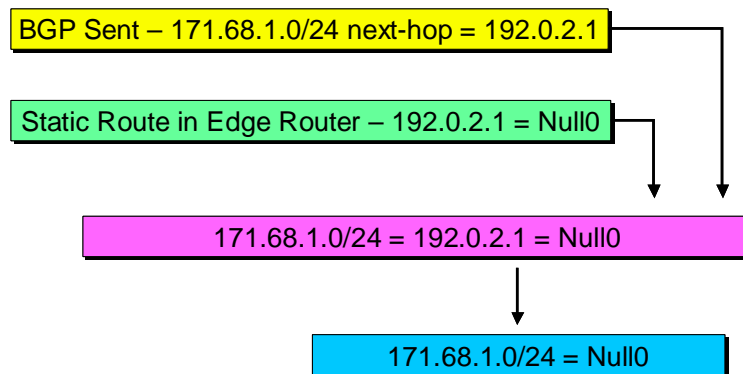


Figure 6 - Gluing the network you want to drop to Null0.

While finding the source addresses of the attacks is a topic for another paper (or an IDS tool), once you do find the source IP addresses of a specific DOS/DDOS attack(s), they are entered centrally via BGP at a convent router or gated workstation in the NOC (Router G in the example using Figure 7). Since Unicast RPF uses the FIB, a whole bunch of source addresses generating the attack can have their next-hop is set to an address in 192.0.2.1/32. For example, if you generate this on a Cisco Router, you would do so via the network statement's route-map:

Sunday, February 18, 2001

```
router bgp XXXX
!
network 171.68.1.0 mask 255.255.255.0 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 XXXX:66 local-AS
!
route-map Set-Next-Hop-To-TESTNET permit 20
!
!
ip route 171.68.1.0 255.255.255.0 Null0 250
ip route 192.0.2.0 255.255.255.0 Null0 250
```

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 (the simulated network to be set to Null0):

```
Excalabur#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

Excalabur#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

Excalabur#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
Excalabur#
```

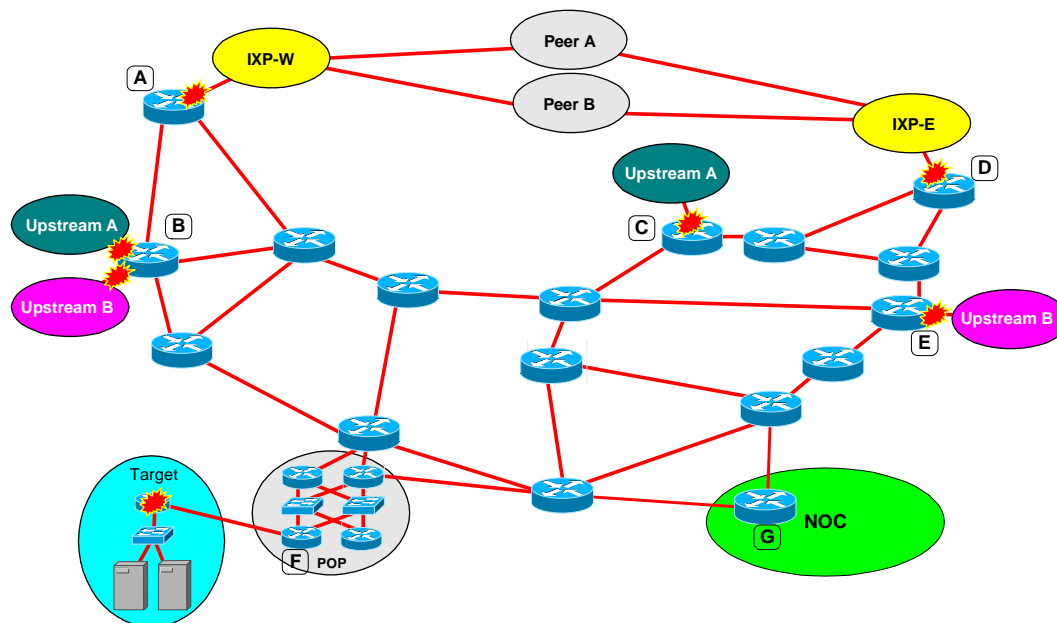


Figure 7 - Using Null0 as a Rapid Reaction Tool

It is critical that 192.168.2.1/32's adjacency is set to Null0. Without uRPF applied, all packets with a destination address to 192.168.2.1/32 would be dropped via Null0. With the 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 would allow the ISP to update 'drop list' for IP source and destination across their 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 CEF switching turned on handles packets sent to Null0 and the Loopback interface differently. Null0 is considered to be a special CEF adjacency. Any packets with a next-hop to Null0 will be dropped in the CEF 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. In routers like 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 CEF 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 like 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 Intrusion Detection Systems (IDS) to integrate with an ISP's operations. Currently IDS systems identify and classify attacks then alerts the operations team. They can then create ACLs to be uploaded to routers that will drop or rate limit the attack. The new

Sunday, February 18, 2001

Unicast RPF enhancement would allow these tools to interface and update drop list 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 are 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 was small – say one or two – then the ACL updates are straightforward. But, if the numbers of routers that need ACL updates are large – say ten to fifty – then the ACL updates get very tricky. In addition, ACL updates on a live router during prime time traffic is a risk. When you add the fact that there are several attacks per day resulting in several ACL update per day, 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 gateD running and an iBGP peering session to one of the routers in the ISP's network (see Figure 8). When the IDS sensors trigger an alarm, the following could happen:

1. IDS Tool Alerts the NOC of a potential DOS/DDOS attack. The IDS Tool post the list of IP addresses that are generating the DOS/DDOS attack.
2. IDS Tool Recommend Dropping the attack. It creates a BGP Advertisement to trigger uRPF to drop the attacks at the edge of the ISP's network. The IDS tool waits for human approval.
3. NOC Team reviews the IDS Recommendation and approves the BGP advertisement to drop the attack at the edge.
4. IDS Tool continues to monitor the attack.

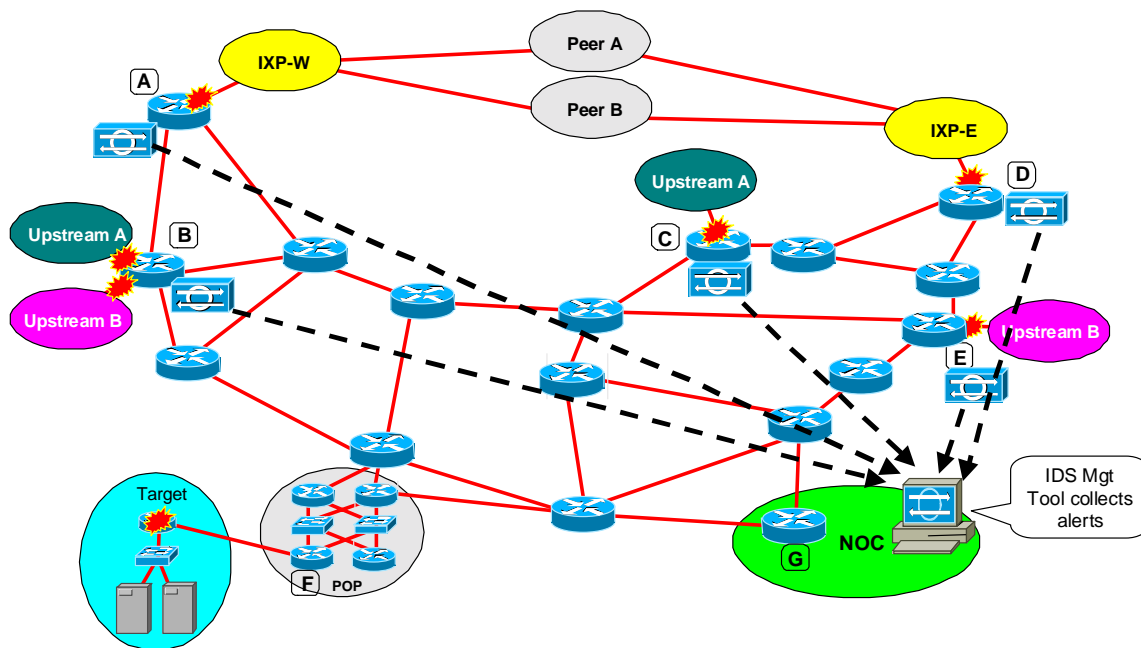


Figure 8 - IDS Systems providing feedback on the attack to a central NOC

What is new is that the IDS management tool generates a list of BGP networks that would need to be dropped to fend off the attack. The BGP network advertisement would set the next-hop to equal a route that is shunted to Null0 on each of the border routers. 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.

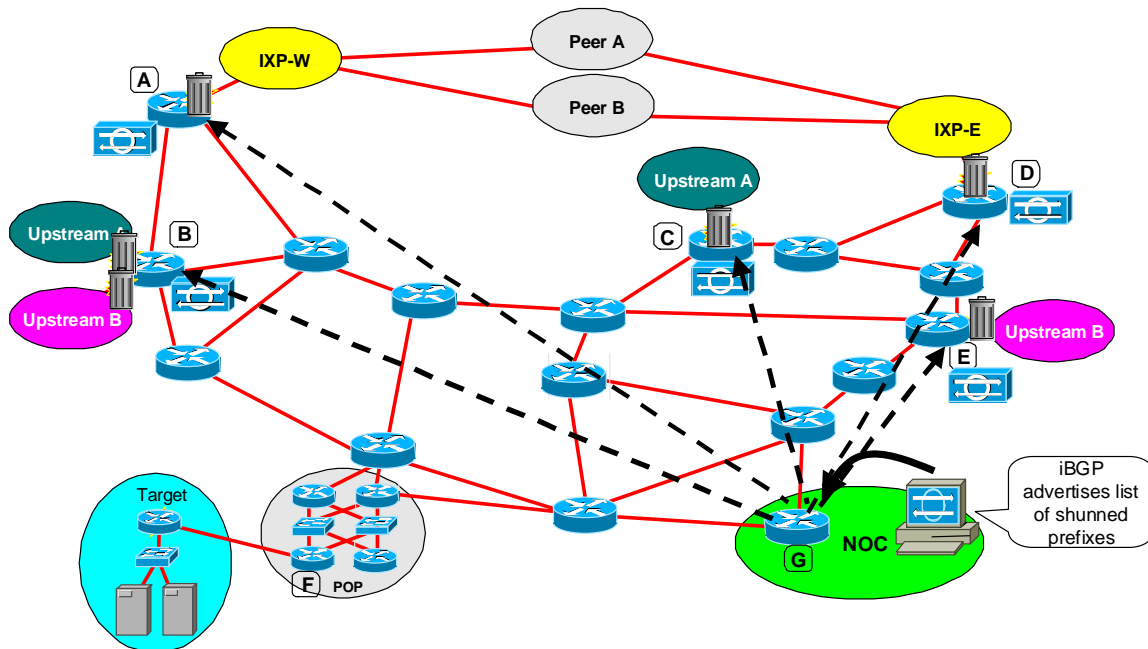


Figure 9 - NOC authorizes an iBGP announcement to trigger a packet drop.

Performance Expectations

Minimal performance impact on the router is inherent in the design of Unicast RPF. Since uRPF uses the same MTRIE look-ups 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 slight different on each line card. Regard less, as show in the lab test on an Engine 0 line card, 3% additional CPU while identifying and dropping packets (i.e. under attack) is considered a minimal performance impact.

The lab test had four GSR with Engine 0 line cards connected in a ring topology. The four routers are located in two different AS 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 AS peer at two different exchange points that hands-off the traffic at the closest exit creating asymmetric routed traffic. The setup is as below:

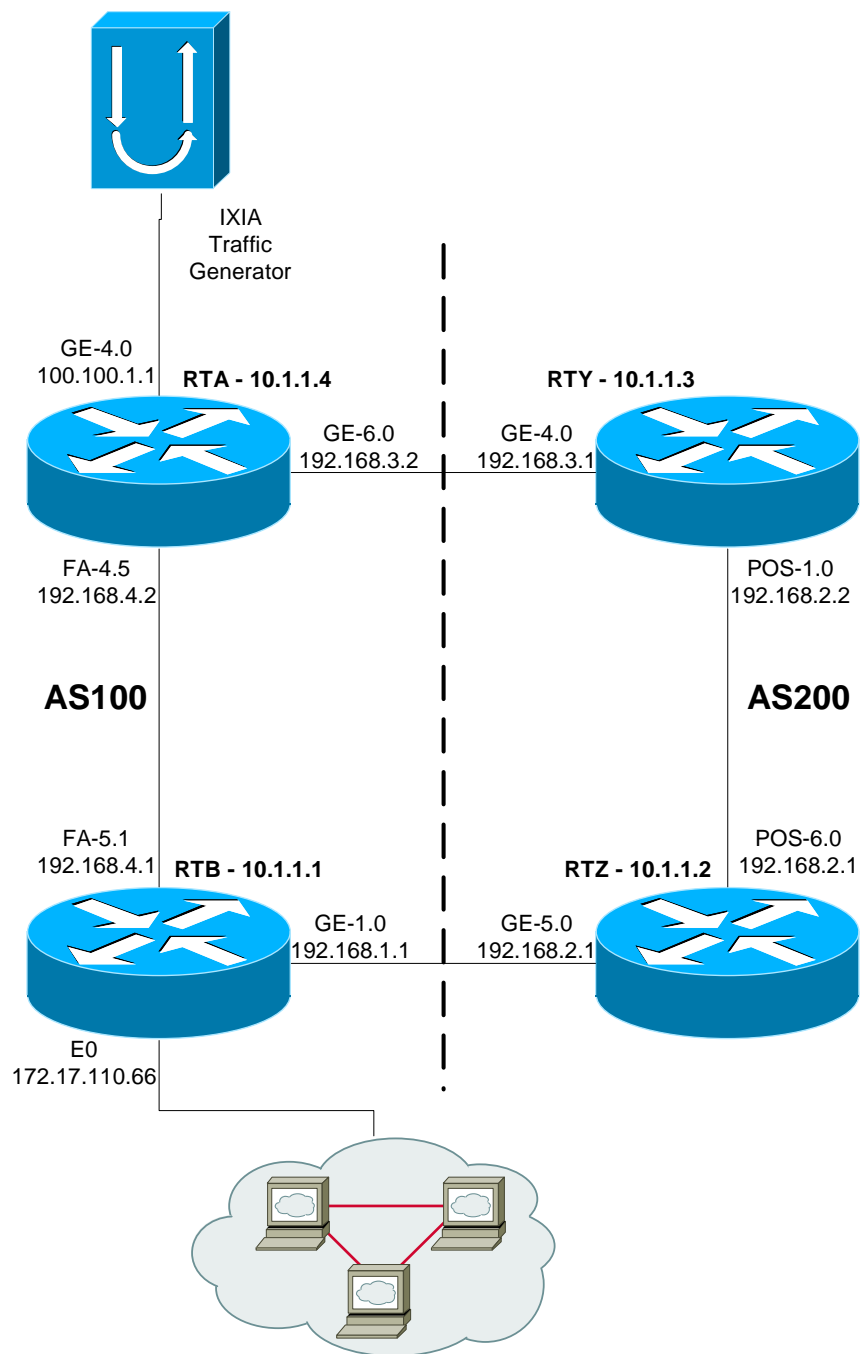


Figure 10 - Lab Test Architecture

Lab Results

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

Attack Stream	Interface	Packet Size	CPU of LC	Packet/Sec
0Mbps/s	Gigabit 4/0	64bytes	0%	0
30Mbps/s	Gigabit 4/0	64bytes	52%	61K
50Mbps/s	Gigabit 4/0	64bytes	57%	100K
100Mbps/s	Gigabit 4/0	64bytes	88%	210K
160Mbps/s	Gigabit 4/0	64bytes	98%	330K

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

Attack Stream	Interface	Packet Size	CPU of LC	Packet/Sec
0Mbps/s	Gigabit 4/0	64bytes	3%	0
30Mbps/s	Gigabit 4/0	64bytes	55%	61K
50Mbps/s	Gigabit 4/0	64bytes	60%	100K
100Mbps/s	Gigabit 4/0	64bytes	90%	210K
160Mbps/s	Gigabit 4/0	64bytes	100%	330K

Deployment Options with the new Enhancements

As mention through out this paper, the new Unicast RPF enhancements has create new deployment options for service providers. The following table reviews some of these deployment options.

Deployment Situation	Type of uRPF to use	Config Notes
Lease Line Customer	Strict Check	
Multi-homed Lease Line Customer (same ISP)	Strict Check or Loose Check	Remember to use BGP Weights on Strict Check
Multi-homed Lease Line Customer (different ISPs)	Strict Check or Loose Check	Remember to use BGP Weights on Strict Check
Dial-up Customers	Strict Check	
DSL Customers	Strict Check	
Cable Modem Customers	Strict Check	
IXP Connection – no private peering	Strict Check	
IXP Connection w/ private peering	Loose Check	
Private Peering – Dedicated Router	Strict Check	Symmetry should be expect between the routes advertised and source addresses sent by the peering ISP.
Private Peering /w several ISPs on the same router	Loose Check	
Co-Location Provider’s Edge Routers	Loose Check	

Sunday, February 18, 2001

References

- ISP/IOS Essentials Whitepaper
<http://www.cisco.com/public/cons/isp/documents/IOSEssentialsPDF.zip>
- ISP/IOS Essentials Power Sessions (Presentation)
http://www.cisco.com/public/cons/isp/documents/IOSEssentials_Seminar.zip
- Unicast Reverse Path Forwarding Enhancements – 12.1T Documentation
http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121newft/121t/121t2/rpf_plus.htm
- 11.1CC Unicast RPF Documentation
http://www.cisco.com/univercd/cc/td/doc/product/software/ios111/cc111/uni_rpf.htm
- Strategies to Protect Against Distributed Denial of Service (DDoS) Attacks
<http://www.cisco.com/warp/public/707/newsflash.html>
- Improving Security on Cisco Routers
<http://www.cisco.com/warp/public/707/21.html>
- 12.0S Unicast Reverse Path Forwarding Commands (original version)
http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/secur_r/srprt6/srrpf.htm
- CAT 6000 Support for the Original uRPF - Release Notes for MSFC IOS Release 12.1(3a)E4
http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/relnotes/78_11385.htm
- Security Overview – 12.1T Documentation
http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/secur_c/scdoverv.htm
- 12.1 Documentation - Configuring Unicast Reverse Path Forwarding (original version)
http://www.cisco.com/univercd/cc/td/doc/product/software/ios121/121cgcr/secur_c/scprt5/scdrpf.htm