Cisco Expo 2009

# ISP Infrastructure & Operational Security

**Dirk Schroetter, Consulting Systems Engineer**

Welcome to the Human Network.

  

# Agenda

- Security Overview and Process
- Securing the Router
- Securing the Management Plane
- Securing the Control Plane
- Securing the Data Plane
- Staying informed
- Q&A
- References

# Credits

- This presentation is based on material from a number of esteemed colleagues in Cisco and in the field at large. My gratitude to them for their excellent material.
- Special thanks to:

  Michael Behringer

  Gregg Schudel

  Dawit Birhanu

  David Barak

  Barry Raveendran Greene

  Mark Prior

# Security Overview & Process

Welcome to the Human Network.

  

# Everything you need to know on 1 slide

- Security is a process as well as a craft & discipline
- Defense at the edge is dead -> defense in the deep
- There are many ressources out there -> stay infromed
- Expect things to break -> be paranoid

# Defence in the deep

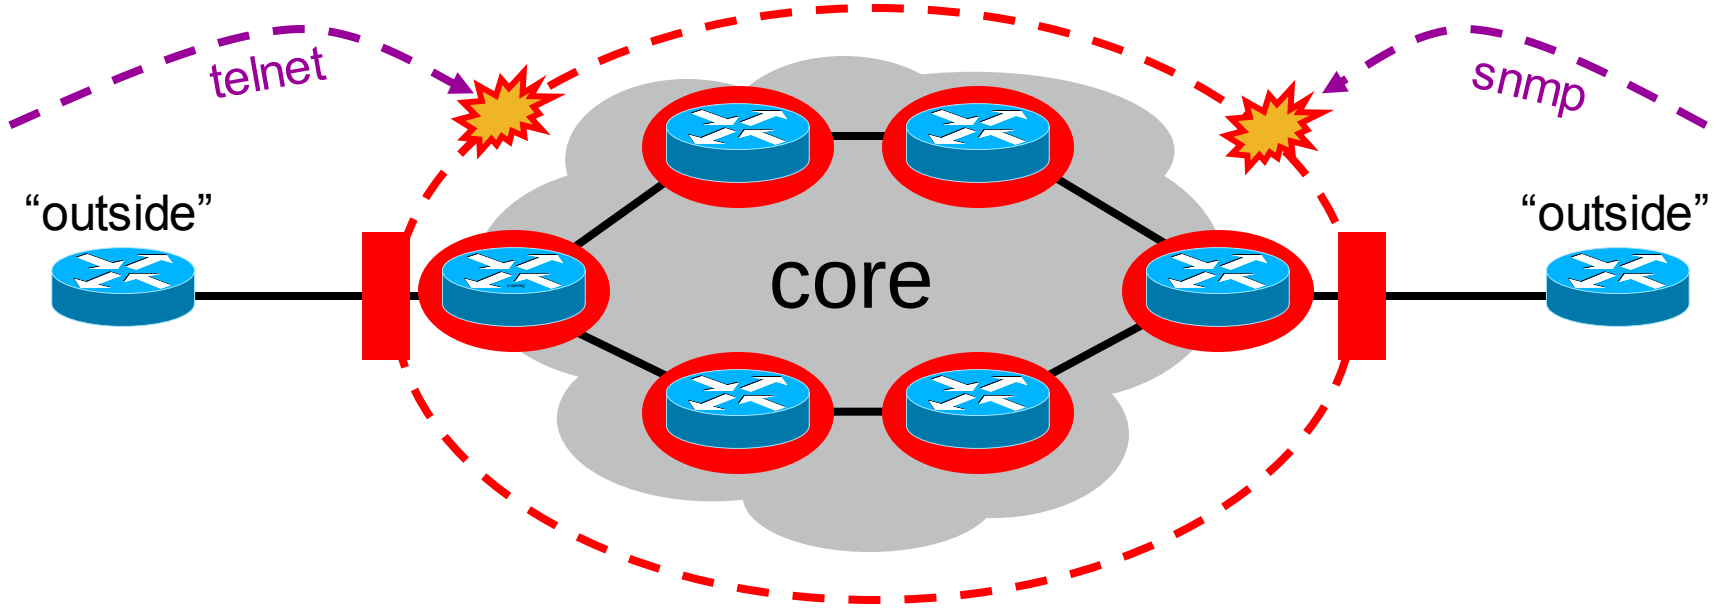


- Core routers individually secured PLUS
- Infrastructure protection
- Routers generally NOT accessible from outside

# Securing the Router

Welcome to the Human Network.

  

# Securing the Router

1. Basic Security

   AAA, SSH, SNMPv3, rACL, CoPP, etc…

2. Don't let packets into (!) the core

   → No way to attack core, except through routing, thus:

3. Secure the routing protocol

   Neighbor authentication, maximum routes, dampening, GTSM, …

4. Design for transit traffic

   Correct Core Design

   Capacity / QoS

   Choose correct router for bandwidth

5. Operate Securely

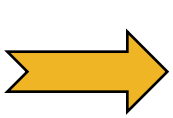

Individual router security

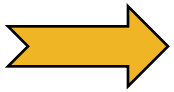

Still "open": routing protocol

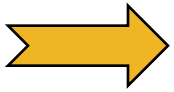

Only attack vector: Transit traffic

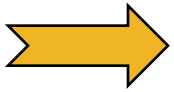

Now only insider attacks possible

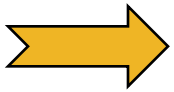

Avoid insider attacks

# Securing the Management Plane

Welcome to the Human Network.

  

# Control Plane Policing

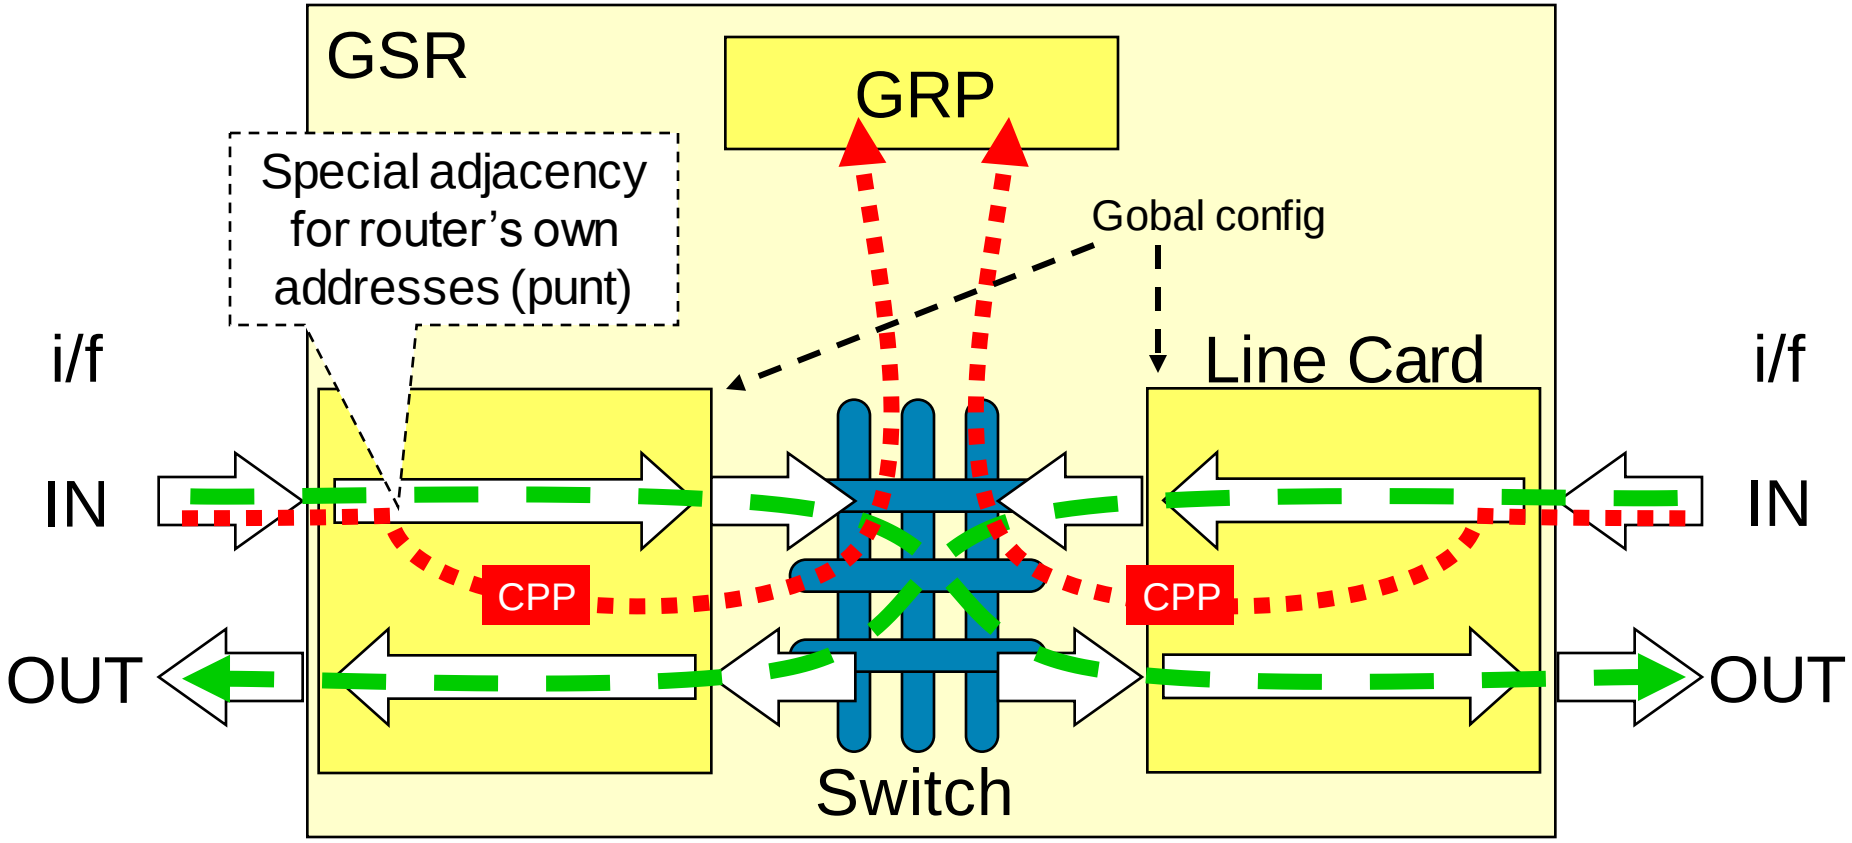

# Configuring CPP

Control Plane Policing

```
Router(config)# access-list 140 deny tcp host 10.1.1.1 any eq telnet
Router(config)# access-list 140 permit tcp any any eq telnet
```

Traffic to be rate limited: All telnet but not from host 10.1.1.1

```
Router(config)# class-map telnet-class
Router(config-cmap)# match access-group 140
Router(config-cmap)# exit
```

Define class-map for this traffic

```
Router(config)# policy-map control-plane-policy
Router(config-pmap)# class telnet-class
Router(config-pmap-c)# police 80000 conform transmit exceed drop
Router(config-pmap-c)# exit
Router(config-pmap)# exit
```

Define the policy for this class map: up to 80 kbps: transmit, else drop

```
Router(config)# control-plane
Router(config-cp)# service-policy input control-plane-policy
Router(config-cp)# exit
```

Apply policy: to control-plane

# Local Packet Transport Services IOS-XR

- Why is LPTS needed?
  - –LPTS enables distributed applications to reside on any or all RPs, DRPs, or LCs
  - –Filters and polices local ‘for-us’ packets and sends them only to the nodes that need them
  - –Re-assembling fragments
  - –High Availability for NSR (Non-Stop Routing)
- LPTS has HW policers on line cards to limit traffic sent to local or remote nodes
  - –LPTS entries in TCAM classifies packets to select a policer
  - –Polices on protocol (BGP, OSPF, SSH) and flow state (BGP established flows, BGP listen)
  - –Policing done on the LC ASIC before packets hit RP/LC CPU
  - –All filters are automatically and dynamically installed by the IOS XR infrastructure

# Local Packet Transport Services

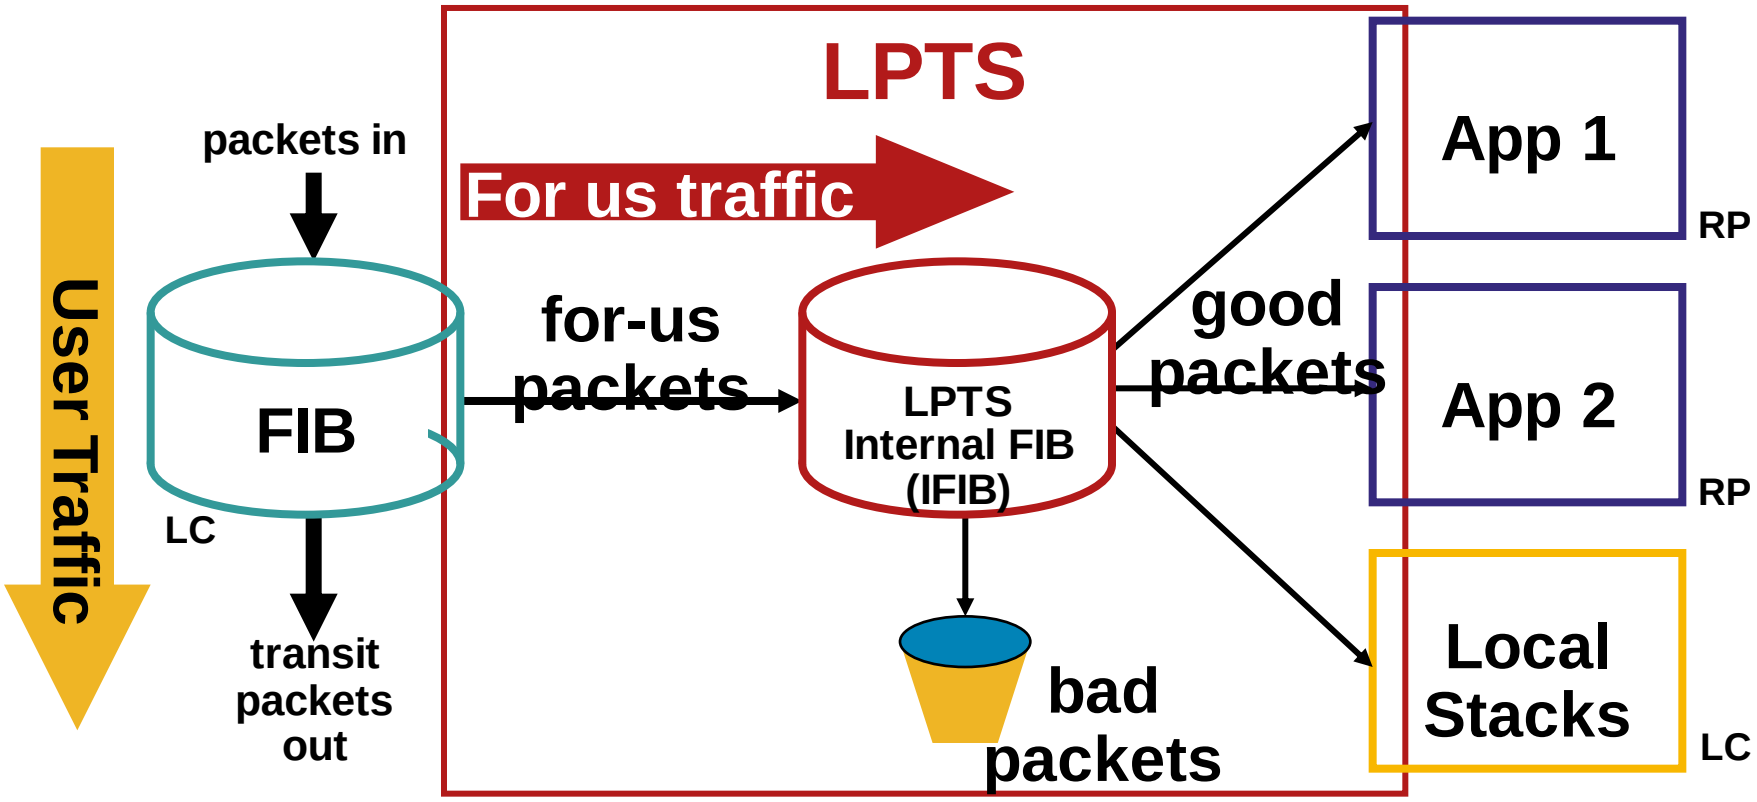


- LPTS enables applications to reside on any or all RPs, DRPs, or LCs
  - Active/Standby, Distributed Applications, Local processing
- IFIB forwarding is based on matching control plane flows
  - Built in firewall for control plane traffic.
- LPTS is transparent and automatic



  

# Securing the Control Plane

Welcome to the Human Network.

 

# Control Plane Security

Data plane
Control plane
Management plane
Services plane

## *Control Plane Security Features*

| | |
|---|---|
| Receive-Path Access List (rACL) | IOS 12.0S-only ACL applied to “receive” path packets |
| Control Plane Policing (aCoPP and dCoPP) | IOS-wide MQC-based policing applied to all punt-path packets |
| Selective Packet Discard (SPD) | IOS-wide “process” level queuing and packet prioritization into CPU |
| BGP Security-Related Commands | |
| IP Prefix List | BGP prefix filtering mechanism using a prefix-list |
| IP Community List | BGP prefix filtering mechanism using a community B |
| IP AS-Path Access Lists | BGP prefix filtering mechanism using an AS path list |
| Route Map | Method for attaching policies to BGP neighbor process |
| Class Map | MQC mechanism for expressing policy |
| Policy Map | MQC mechanism for expressing policy |
| Static/Null0 Routes | Static route definition mechanism |
| ISIS Security-Related Commands | ISIS related security commands |
| OSPF Security-Related Commands | OSPF related security commands |
| EIGRP Security-Related Commands | EIGRP related security commands |
| RIPv2 Security-Related Commands | RIPv2 related security commands |
| LDP Security-Related Commands | LDP related security commands |

# Control Plane Security

- Data plane
- Control plane
- Management plane
- Services plane

## *Control Plane Security Features*

| | |
|---|---|
| RSVP Security-Related Commands | RSVP related security commands |
| PIM Security-Related Commands | PIM related security commands |
| IGMP Security-Related Commands | IGMP related security commands |
| IP icmp rate-limits | Rate limit ICMP error message generation |
| IP redirects | Toggle generation of ICMP redirect messages |
| IP unreachables | Toggle generation of ICMP unreachable messages |
| IP mask-reply | Toggle generation of ICMP mask-reply messages |
| IP information-reply | Toggle generation of ICMP information-reply messages |
| IP proxy-arp | Toggle ip proxy-arp functionality |
| Key Chain | Define key-chains (used for some routing protocol authentication mechanisms) |
| | |
| | |
| | |
| | |
| | |
| | |

# Securing the Data Plane

Welcome to the Human Network.

 

# uRPF Overview

- uRPF uses information in the Forwarding Information Base (FIB) created by the Cisco Express Forwarding (CEF) switching path to perform reverse path resolution on the source IP address of an incoming packet.
  - The FIB table is dynamically populated from the routing information of all active routing protocols and static routes.
  - There is an entry for each known prefix, along with associated path(s) information.
- uRPF is superior to ingress ACLs for ingress anti-spoofing protection
  - It's dynamic and automatically updated
  - It's operationally simple to maintain and scale
  - Minimal performance impact (it's in the fast path)

# uRPF Basic Operations

- uRPF "Strict Mode" (aka "version 1")

  Requires the existence of a valid FIB path entry via the exact same interface as that on which the packet arrived for the source IP address of an incoming packet.

  If the FIB return path does not exist, or refers to a different interface than the one on which the packets arrived, the packet is dropped.

- uRPF "Loose Mode" (aka "version 2")

  Requires the existence of a valid FIB path entry via any interface for the source IP address of an incoming packet (i.e. the strict interface adherence is no longer enforced).

  If a valid FIB entry does not exist, or if the FIB entry refers to the Null0 interface, the packet is dropped.

- uRPF "VRF Mode" (aka "version 3")

  Requires that the source IP address of the incoming packet exist or not exist (depending on the mode selected – white list or black list mode) within the prefix list contained within the designated VRF.

  [Note: this uRPF v3 availability is limited to certain C12K images and only Engine 0 Line Cards to date.]

# uRPF Basic Operations "strict mode" vs. "loose mode"

router(config-if)# ip verify unicast source reachable-via **rx**

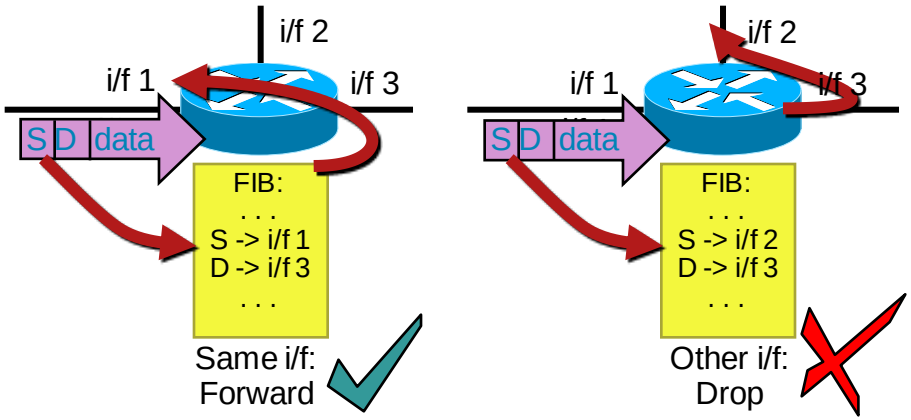


"Strict Mode"
(aka "v1")

router(config-if)# ip verify unicast source reachable-via **any**

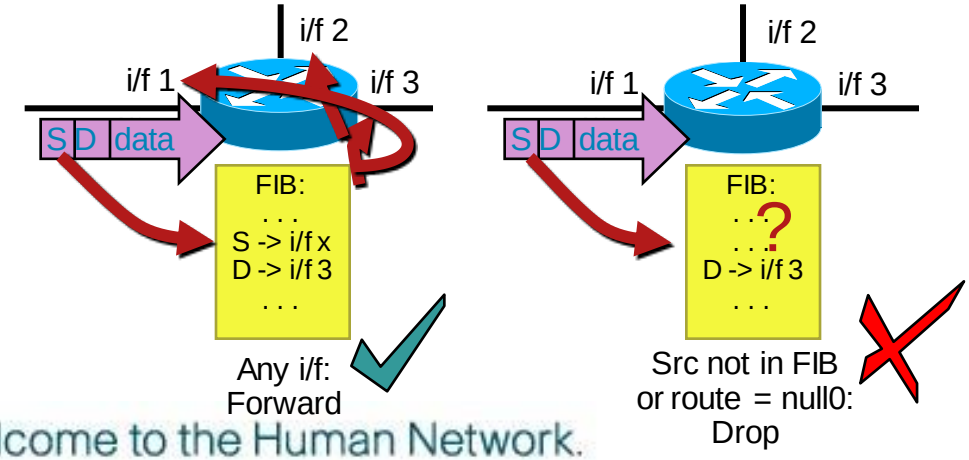


"Loose Mode"
(aka "v2")

Welcome to the Human Network.

# uRPF Overview

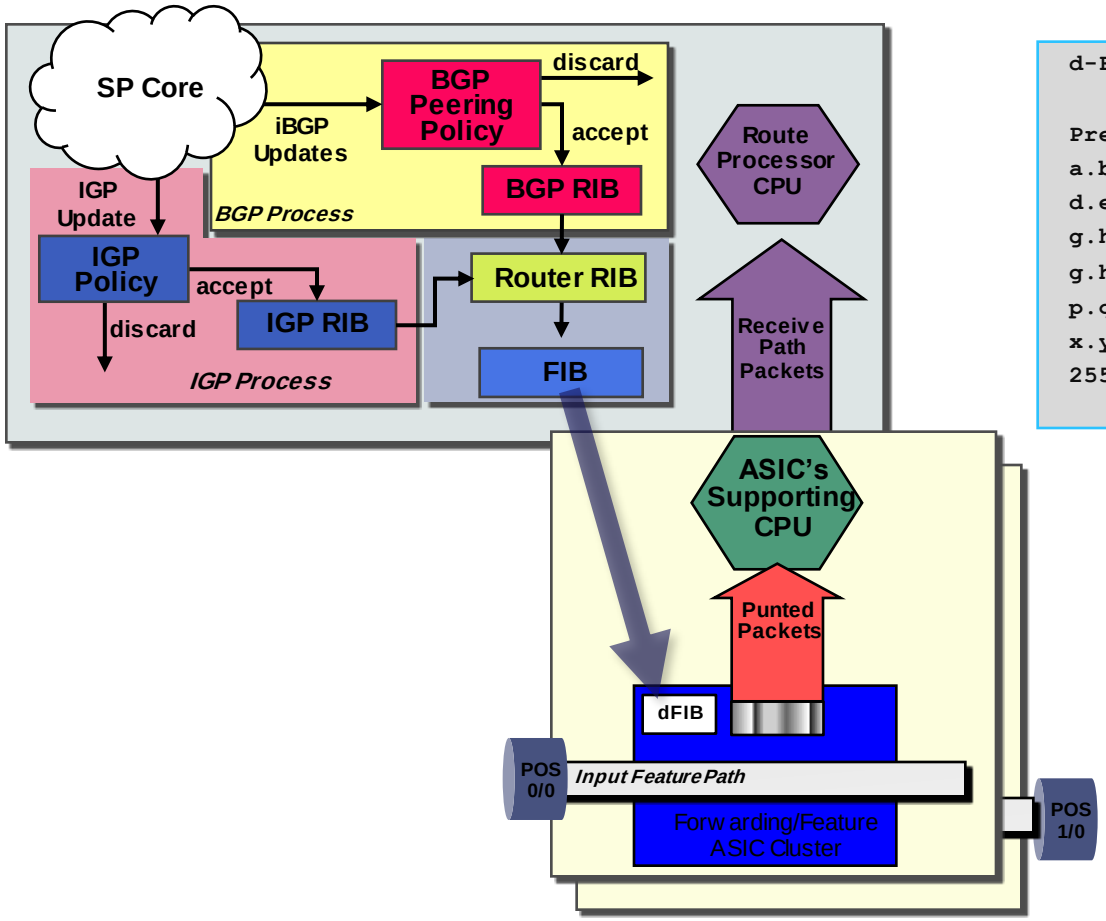


d-FIB

| Prefix | Next Hop | Interface |
|---|---|---|
| a.b.c.0/24 | 1.1.1.1 | POS0/0 |
| d.e.f.0/24 | 2.2.2.2 | POS1/0 |
| g.h.i.0/30 | attached | POS0/0 |
| g.h.i.1/32 | receive | |
| p.q.r.s/32 | attached | Null0 |
| x.y.z.0/24 | 2.2.2.2 | POS1/0 |
| 255.255.255.255/32 | receive | |

# uRPF "Strict Mode"

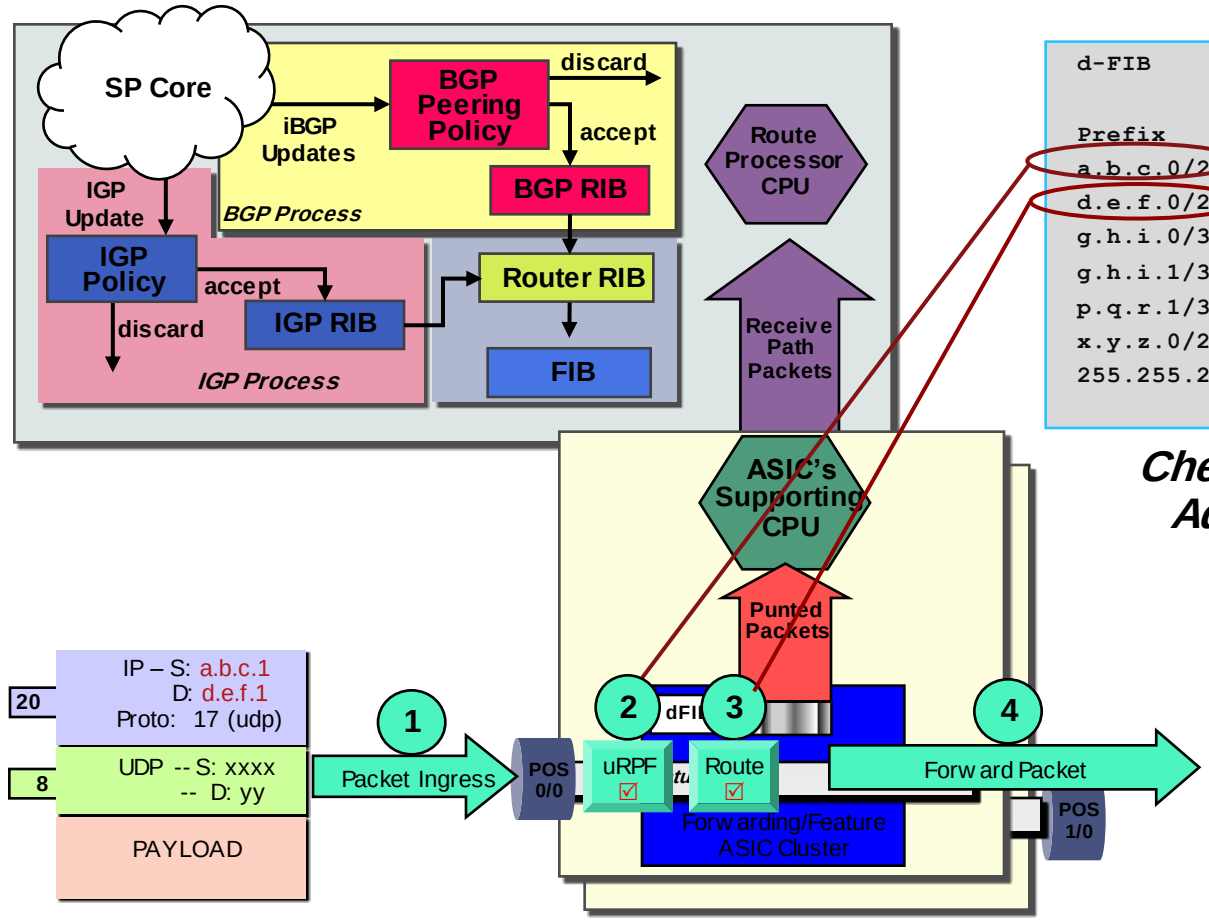


d-FIB

| Prefix | Next Hop | Interface |
|---|---|---|
| a.b.c.0/24 | 1.1.1.1 | POS0/0 |
| d.e.f.0/24 | 2.2.2.2 | POS1/0 |
| g.h.i.0/30 | attached | POS0/0 |
| g.h.i.1/32 | receive | |
| p.q.r.1/32 | attached | Null0 |
| x.y.z.0/24 | 2.2.2.2 | POS1/0 |
| 255.255.255.255/32 | receive | |

***Check FIB – Does Source Address and next-hop adjacency match?***

```
router(config)# int pos0/0
router(config-if)# ip verify unicast source reachable-via rx
```

Welcome to the Human Network.

# uRPF "Loose Mode"

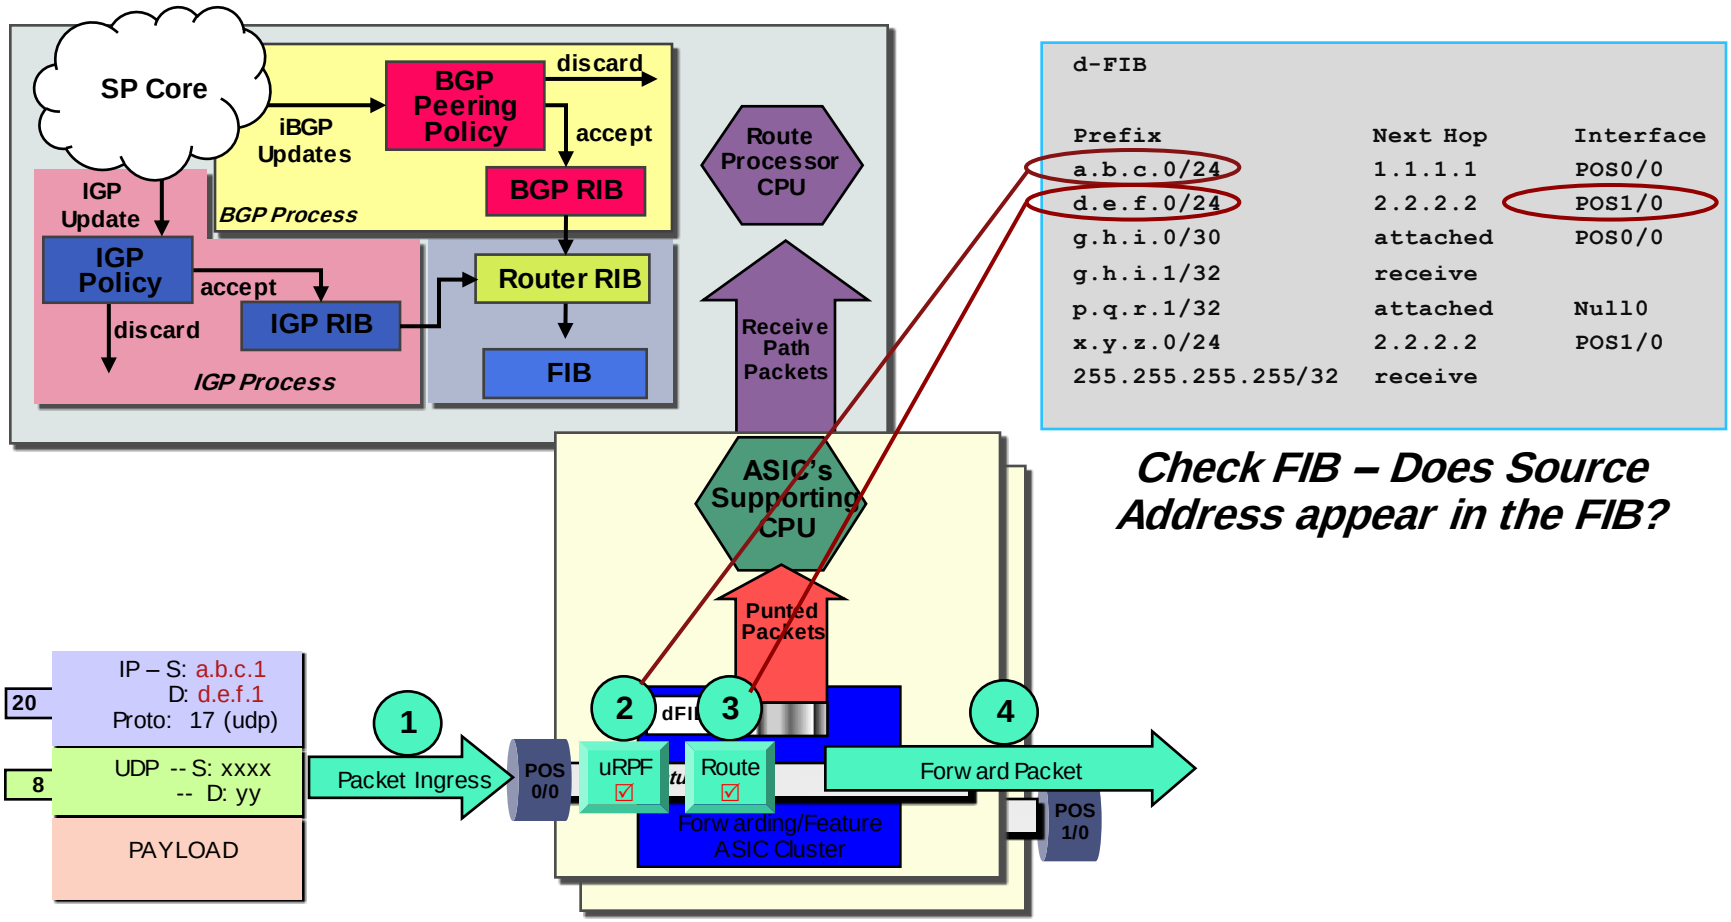


```
router(config)# int pos0/0
router(config-if)# ip verify unicast source reachable-via any
```

Welcome to the Human Network.

# Where To Deploy uRPF? Service Provider Configuration

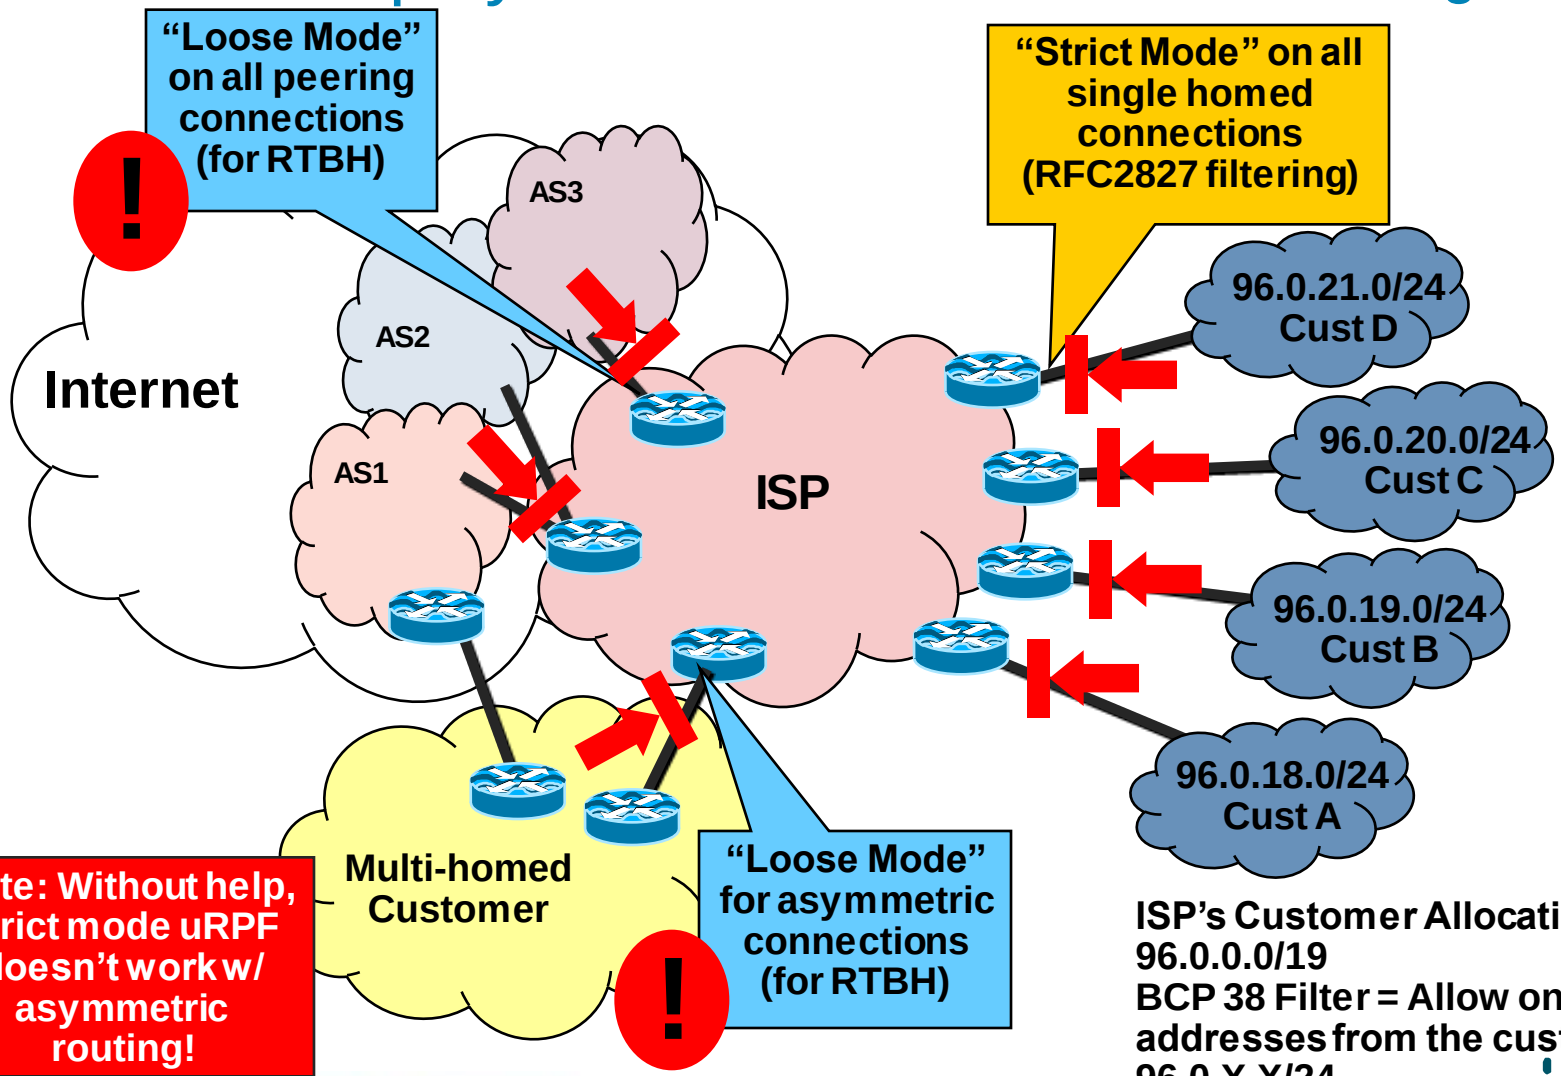


ISP's Customer Allocation Block: 96.0.0.0/19
BCP 38 Filter = Allow only source addresses from the customer's 96.0.X.X/24

Welcome to the Human Network.

# Remote Triggered Black Hole (RTBH) Filtering

## What Is It?

- Remotely Triggered Black Hole (RTBH) filtering is an **SP network foundation tool**

  RTBH provides a reaction technique for security events (e.g. DoS/DDoS attacks) that enables network-wide destination and source IP address –based drop capabilities

- RTBH Uses BGP to trigger network-wide attack flow responses

  Simple pre-configured static route allows the SP to trigger network-wide destination-based black holes at iBGP update speeds

  Coupling with Unicast Reverse Path Forwarding (uRPF) enables network-wide source-based triggered black holes, including at the ISP—ISP edge

  Effective against spoofed and valid source addresses

- Preparation does not adversely impact SP operations or router performance

# Reacting to an Attack with RTBH

## What's Needed?

- Pre-configured static route to Null0 for a /32 address from TEST-NET (192.0.2.0/24) on all border/edge routers
- Unicast Reverse Path Forwarding (uRPF) strict or loose mode on all border/edge routers (optional – for source-based RTBH)
- Method to inject a BGP advertisements into the network with a tag or community to trigger the drop

  Include the no-export, no-advertise communities, and good egress route-filters to prevent leakage
- A way to quickly identify and classify the attack traffic

  NetFlow and Narus, Arbor, or other

# RTBH: Global and Regional Deployment

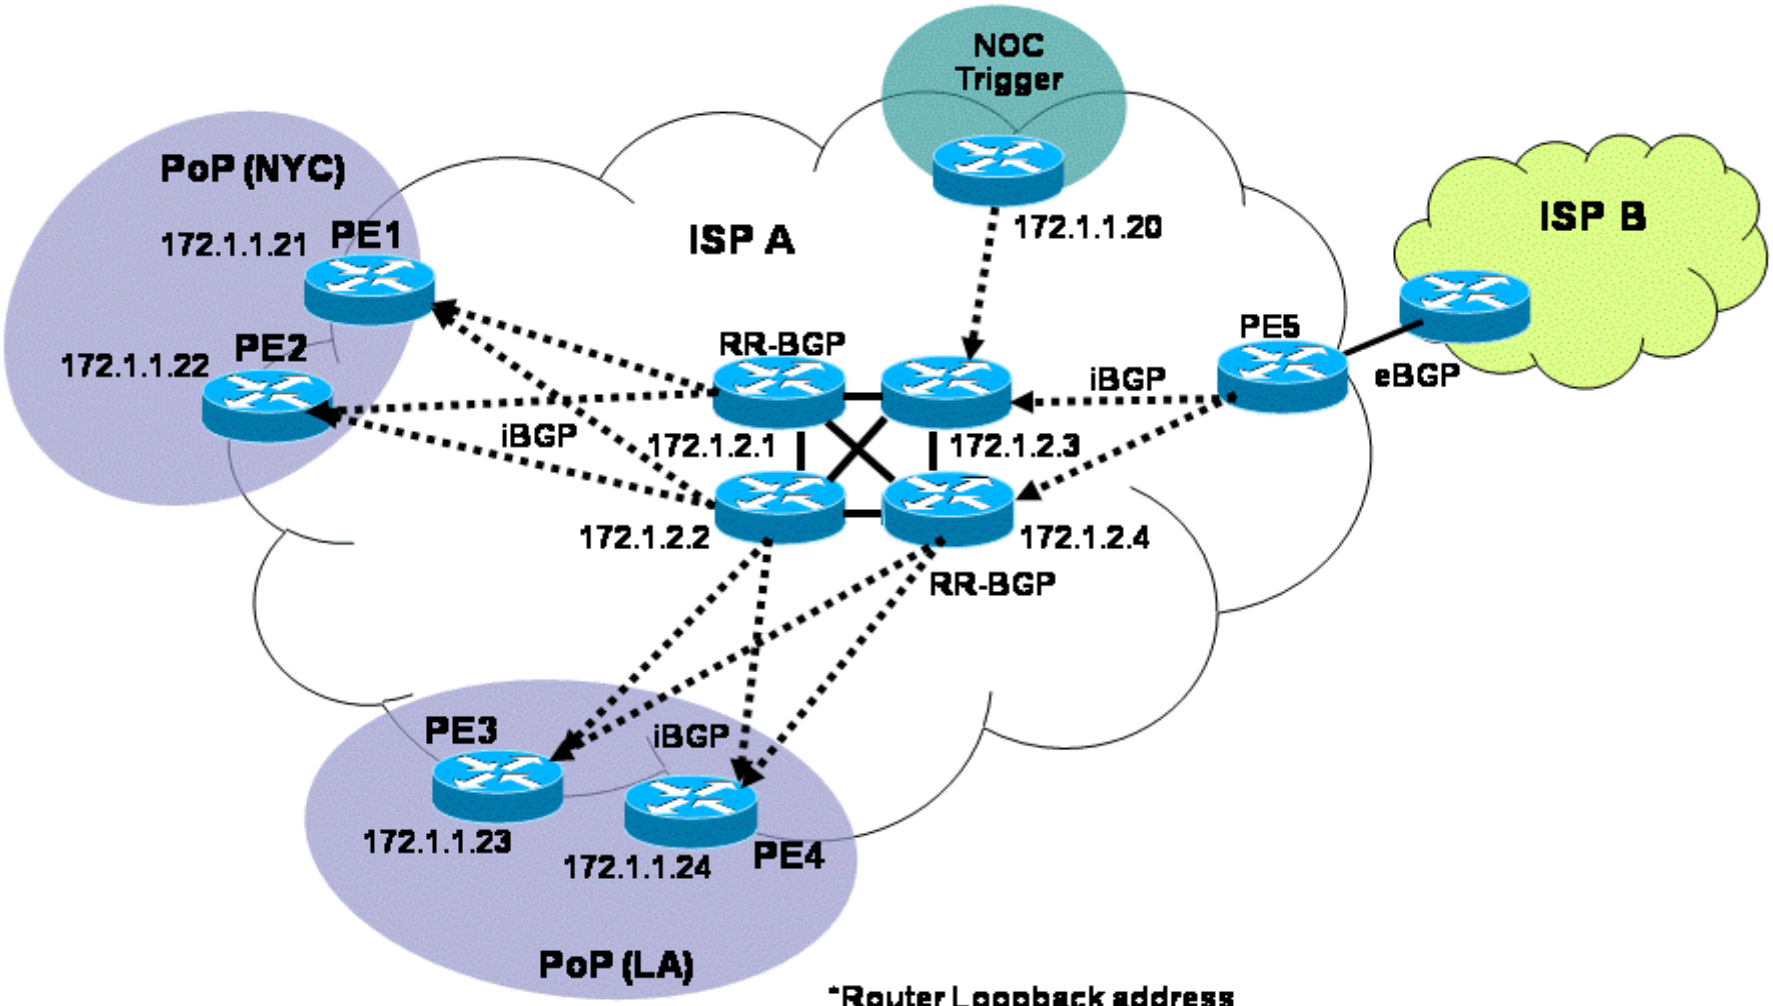

# RTBH: Basic Tag-Based Dst Drops

## Trigger router pre-configuration

```
!
hostname trigger
ip bgp-community new-format
!
interface Loopback0
 ip address 172.1.1.20 255.255.255.255
!
router bgp 1
 no synchronization
 bgp router-id 172.1.1.20
 bgp log-neighbor-changes
 redistribute static route-map RTBH-trigger
 neighbor 172.1.2.3 remote-as 1
 neighbor 172.1.2.3 update-source Loopback0
 neighbor 172.1.2.3 version 4
 neighbor 172.1.2.3 route-map no-prefix-in in
 no auto-summary
!
ip prefix-list no-in seq 5 deny 0.0.0.0/0
ip route 192.0.2.1 255.255.255.255 Null0
!
route-map no-prefix-in permit 10
 match ip address prefix-list no-in
!
route-map RTBH-trigger permit 10
 match tag 666
 set ip next-hop 192.0.2.1
 set local-preference 200
 set origin igp
 set community no-export
!
```

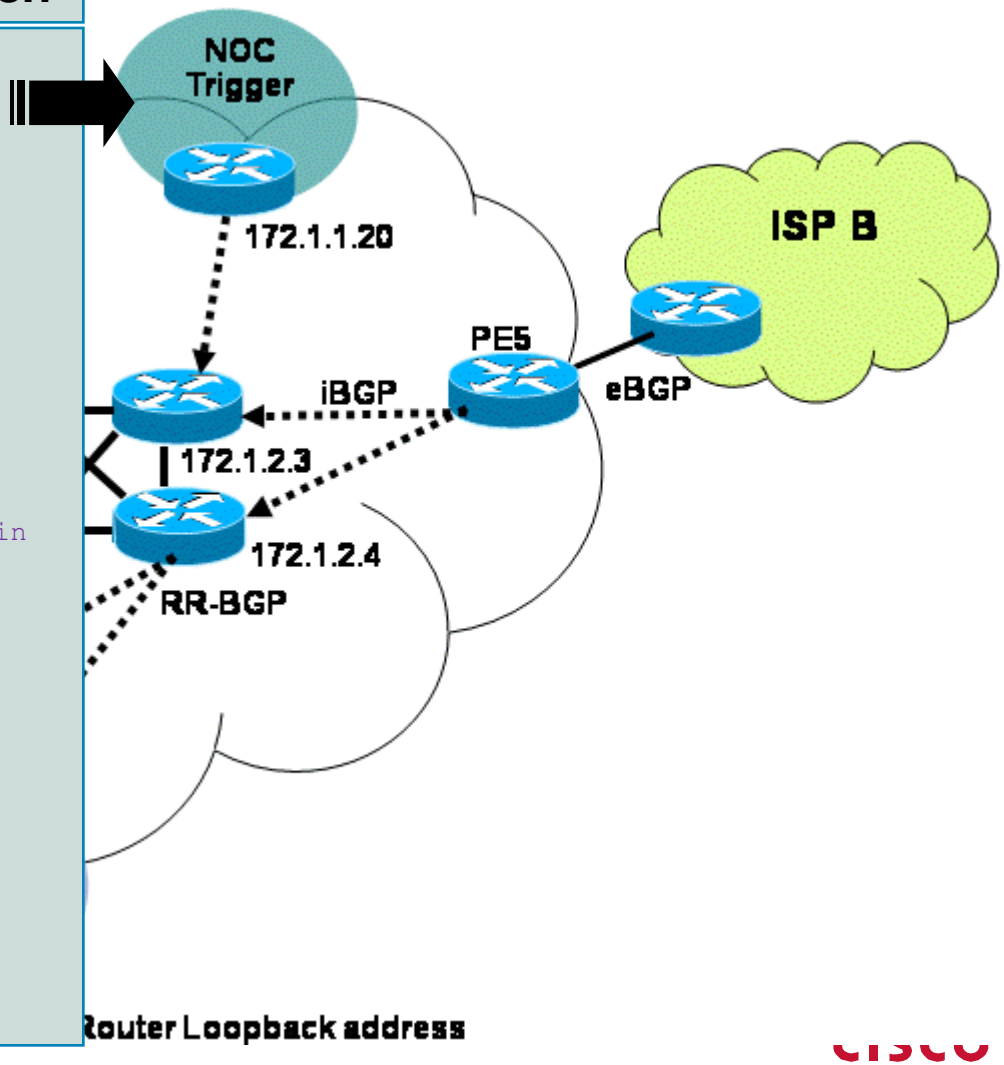

# RTBH: Basic Tag-Based Dst Drops

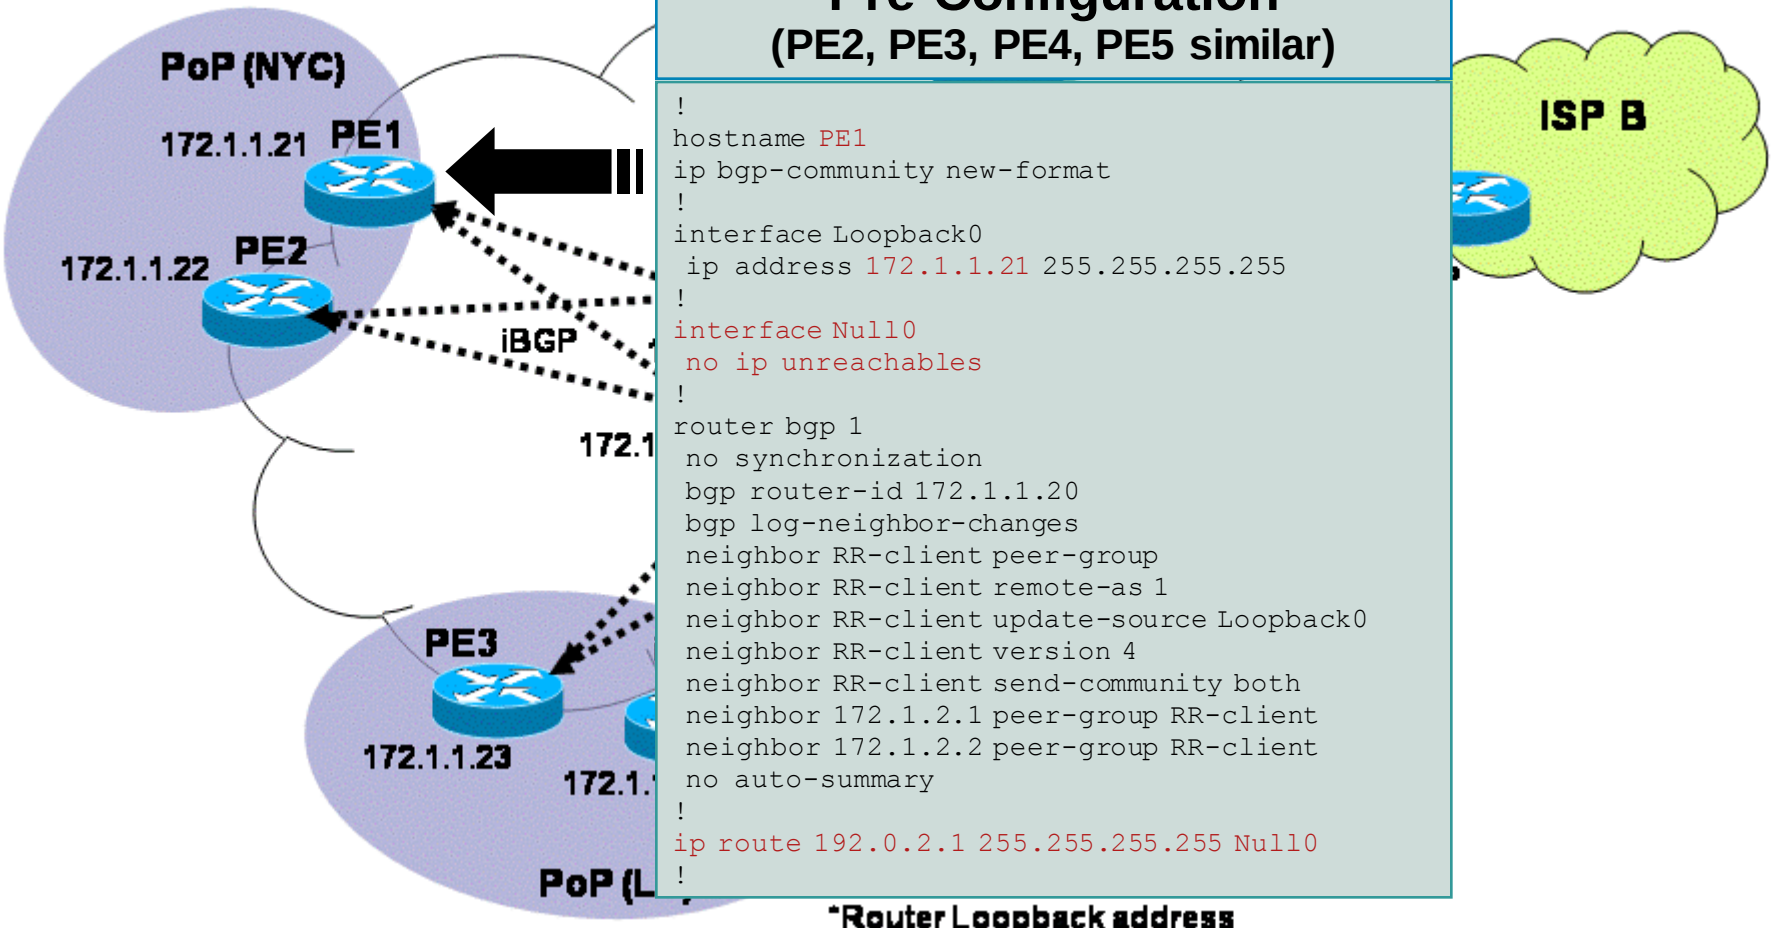


## PE1 Edge Router Pre-Configuration
### (PE2, PE3, PE4, PE5 similar)

```
!
hostname PE1
ip bgp-community new-format
!
interface Loopback0
 ip address 172.1.1.21 255.255.255.255
!
interface Null0
 no ip unreachables
!
router bgp 1
 no synchronization
 bgp router-id 172.1.1.20
 bgp log-neighbor-changes
 neighbor RR-client peer-group
 neighbor RR-client remote-as 1
 neighbor RR-client update-source Loopback0
 neighbor RR-client version 4
 neighbor RR-client send-community both
 neighbor 172.1.2.1 peer-group RR-client
 neighbor 172.1.2.2 peer-group RR-client
 no auto-summary
!
ip route 192.0.2.1 255.255.255.255 Null0
!
```

*Router Loopback address

# RTBH: Global and Regional Deployment Example: Basic Tag-Based Dst Drops

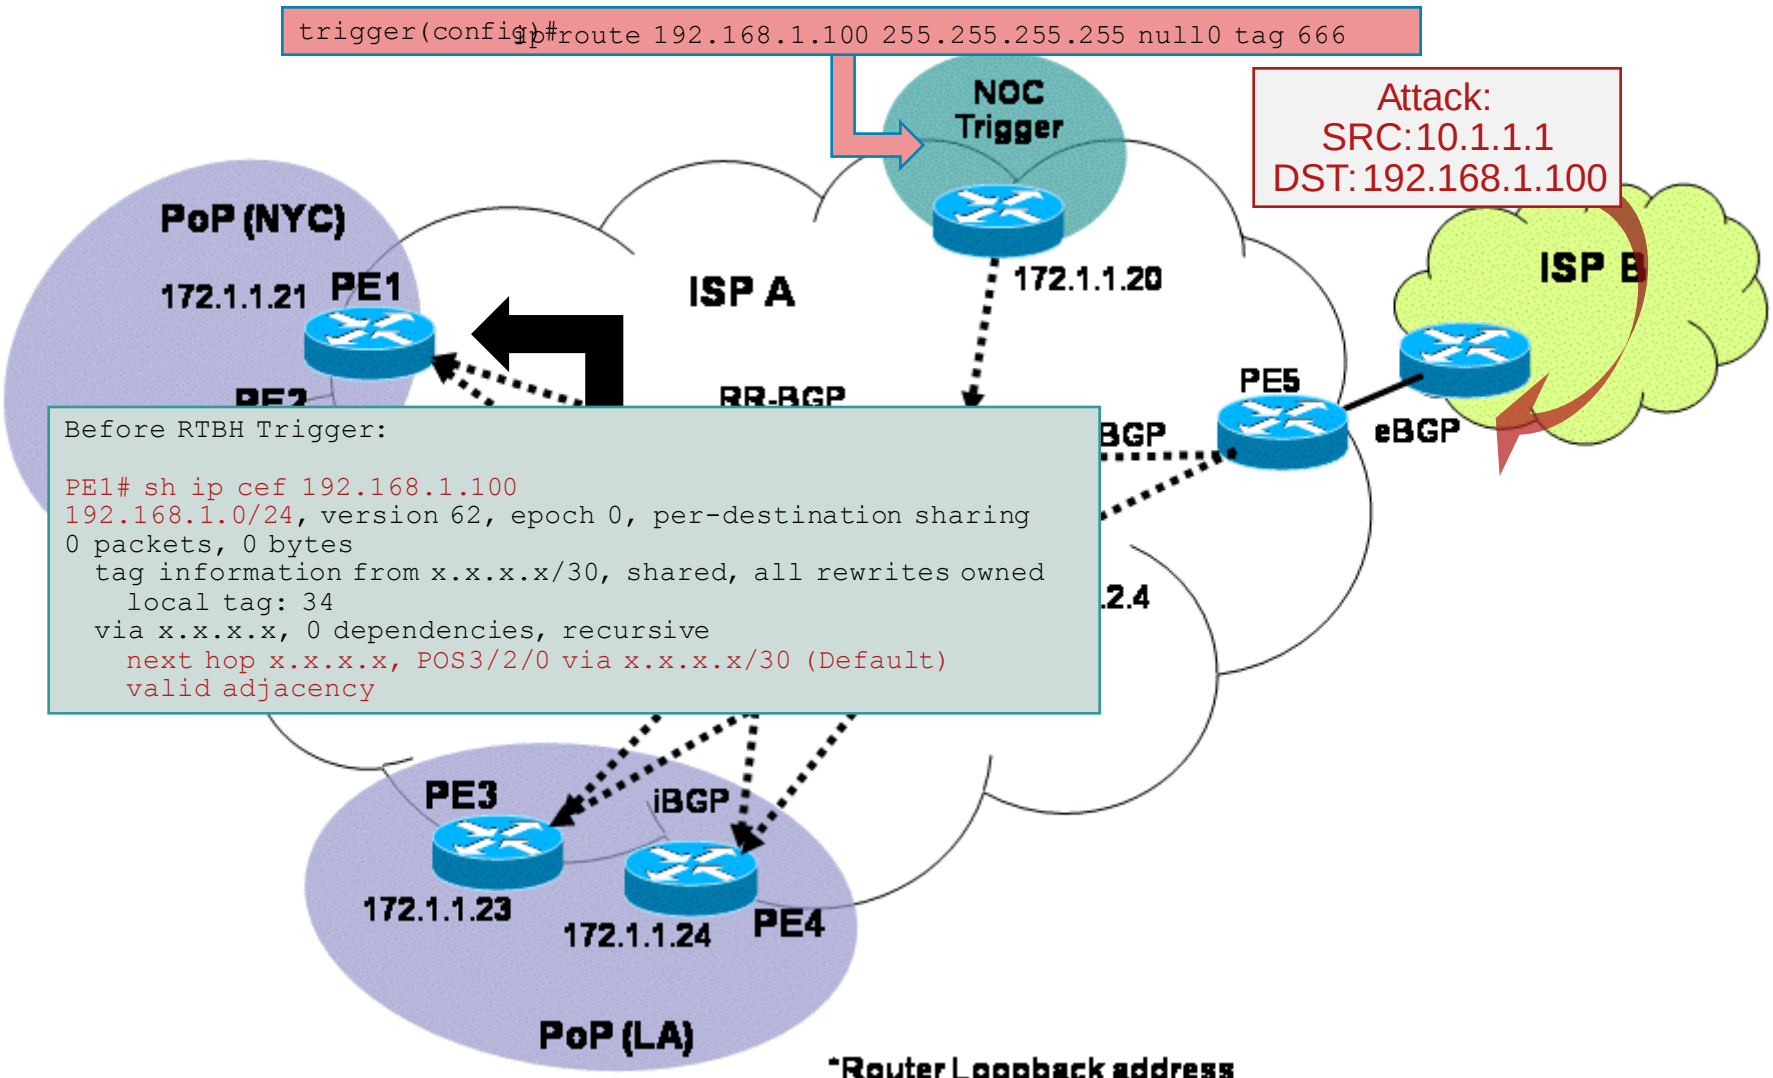

# RTBH: Global and Regional Deployment Example: Basic Tag-Based Dst Drops

```
trigger(config)#ip route 192.168.1.100 255.255.255.255 null0 tag 666
```

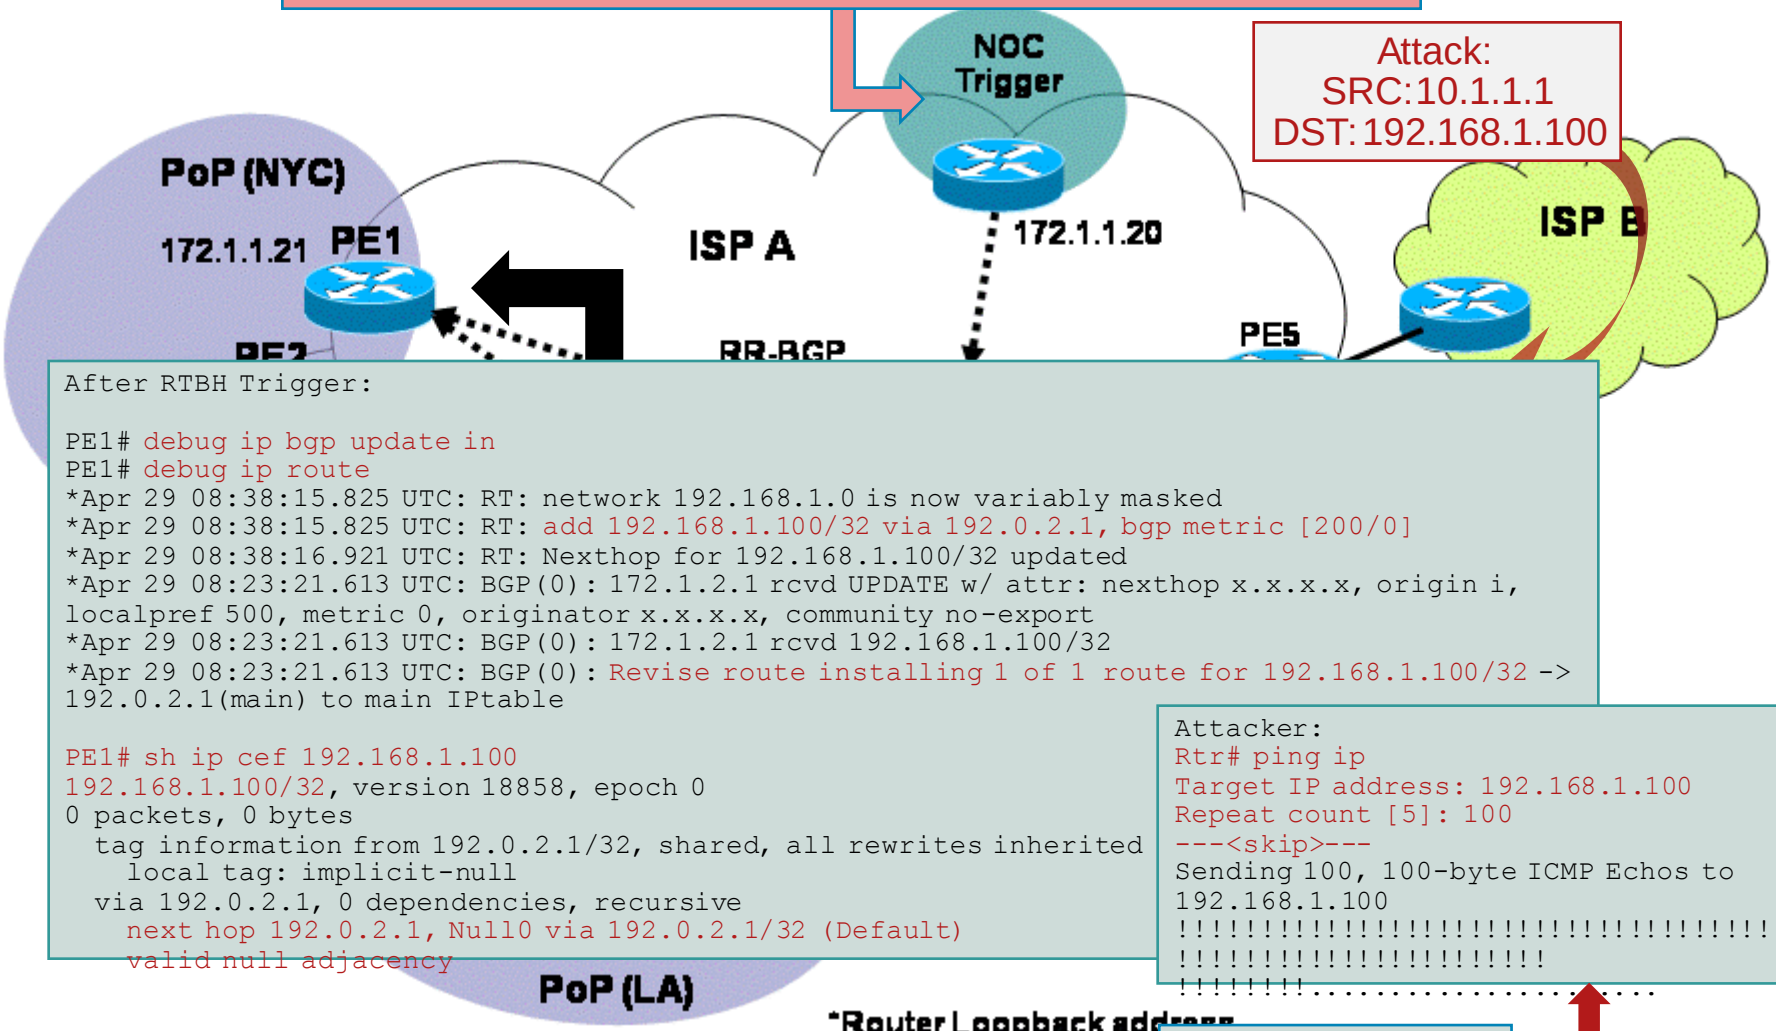


```
After RTBH Trigger:

PE1# debug ip bgp update in
PE1# debug ip route
*Apr 29 08:38:15.825 UTC: RT: network 192.168.1.0 is now variably masked
*Apr 29 08:38:15.825 UTC: RT: add 192.168.1.100/32 via 192.0.2.1, bgp metric [200/0]
*Apr 29 08:38:16.921 UTC: RT: Nexthop for 192.168.1.100/32 updated
*Apr 29 08:23:21.613 UTC: BGP(0): 172.1.2.1 rcvd UPDATE w/ attr: nexthop x.x.x.x, origin i,
localpref 500, metric 0, originator x.x.x.x, community no-export
*Apr 29 08:23:21.613 UTC: BGP(0): 172.1.2.1 rcvd 192.168.1.100/32
*Apr 29 08:23:21.613 UTC: BGP(0): Revise route installing 1 of 1 route for 192.168.1.100/32 ->
192.0.2.1(main) to main IPtable

PE1# sh ip cef 192.168.1.100
192.168.1.100/32, version 18858, epoch 0
0 packets, 0 bytes
  tag information from 192.0.2.1/32, shared, all rewrites inherited
    local tag: implicit-null
  via 192.0.2.1, 0 dependencies, recursive
    next hop 192.0.2.1, Null0 via 192.0.2.1/32 (Default)
    valid null adjacency
```

```
Attacker:
Rtr# ping ip
Target IP address: 192.168.1.100
Repeat count [5]: 100
---<skip>---
Sending 100, 100-byte ICMP Echos to
192.168.1.100
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!
!!!!!!!!.....................
```

Verification…

# Filter, filter, filter ...

- Always filter all the BGP advertisements you send and receive
  - That includes your peers and customers
  - Filter all advertisements of DUSA (see RFC 3300)
- Positively, allow only allocated IPv4 blocks
  - Observe the minimum practical allocations for each block
- Be paranoid, your peers and customers will send you bad advertisements.

# Stay informed

Welcome to the Human Network.

 

# Reference: ISP Security Essentials

- The “bible” for Core Security
- Available as book, and on FTP:

  ftp://ftp-eng.cisco.com/cons/isp/security

- How to secure the core

  Security for devices, routing, traffic, management, …

# Stay informed

- NANOG
- RIPE
- Cymru:
  - Up to date information about SP core security, secure routing, bogon address ranges, etc… [www.cymru.com]
- nsp-security:
  - Closed list with the operations contacts of the big ISPs. Only for trusted contacts, very controlled. [https://puck.nether.net/mailman/listinfo/nsp-security]
- nsp-security-discuss:
  - Closed list with the operations contacts of the big ISPs. Only for trusted contacts, a bit easier to access than nsp-security [https://puck.nether.net/mailman/listinfo/nsp-security-discuss]

# References

# References

Product Security:

- Cisco's Product Vulnerabilities; A page that every SE MUST know!!! [http://www.cisco.com/warp/public/707/advisory.html]
- Security Reference Information: Various white papers on DoS attacks and how to defeat them [http://www.cisco.com/warp/public/707/ref.html]

ISP Essentials:

- Technical tips for ISPs every ISP should know [ftp://ftp-eng.cisco.com/cons/isp/]

Technical tips:

- Troubleshooting High CPU Utilization on Cisco Routers [http://www.cisco.com/warp/public/63/highcpu.html]
- The "show processes" command [http://www.cisco.com/warp/public/63/showproc_cpu.html]
- NetFlow Performance White Paper [http://www.cisco.com/warp/public/cc/pd/iosw/prodlit/ntfo_wp.htm]

Mailing lists:

- cust-security-announce: All customers should be on this list.
- cust-security-discuss: For informal discussions.

# Recommended Reading

## Router Security Strategies: Securing IP Network Traffic Planes

*Router Security Strategies: Securing IP Network Traffic Planes provides a comprehensive approach to understand and implement IP traffic plane separation and protection on IP routers. This book details the distinct traffic planes of IP networks and the advanced techniques necessary to operationally secure them. This includes the data, control, management, and services planes that provide the infrastructure for IP networking.*

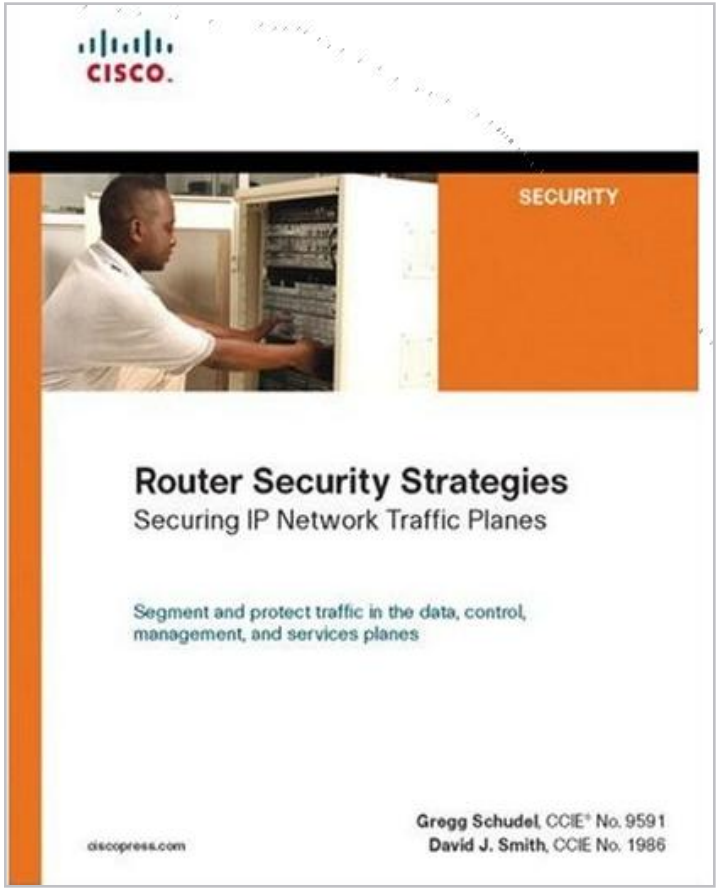


**http://www.ciscopress.com/bookstore/product.asp?isbn=1587053365**

Welcome to the Human Network.


CISCO