



# ***Securing the Network and Data Plane***

***Seo Boon NG <sbng@cisco.com>***



# What is Ingress and Egress?

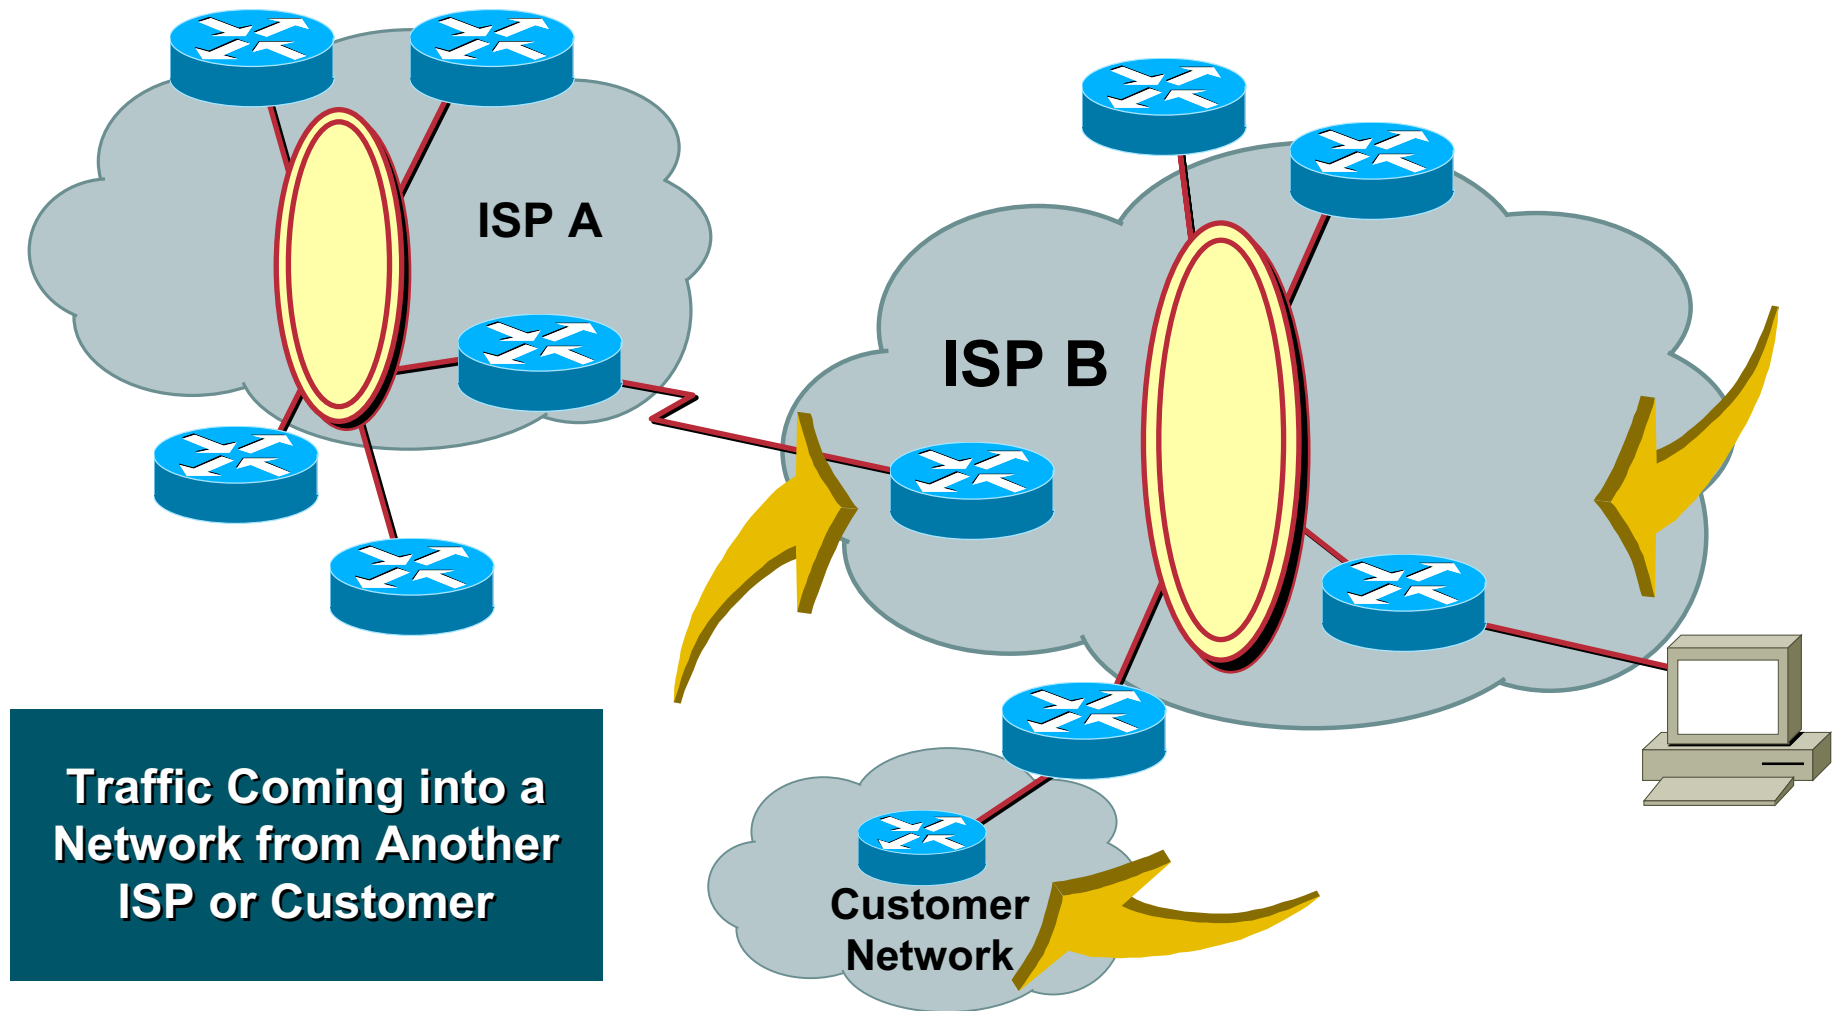
# Securing the Network

Cisco.com

- **Packet filtering - Dropping**
- **Traffic Rate limiting – Degrading**

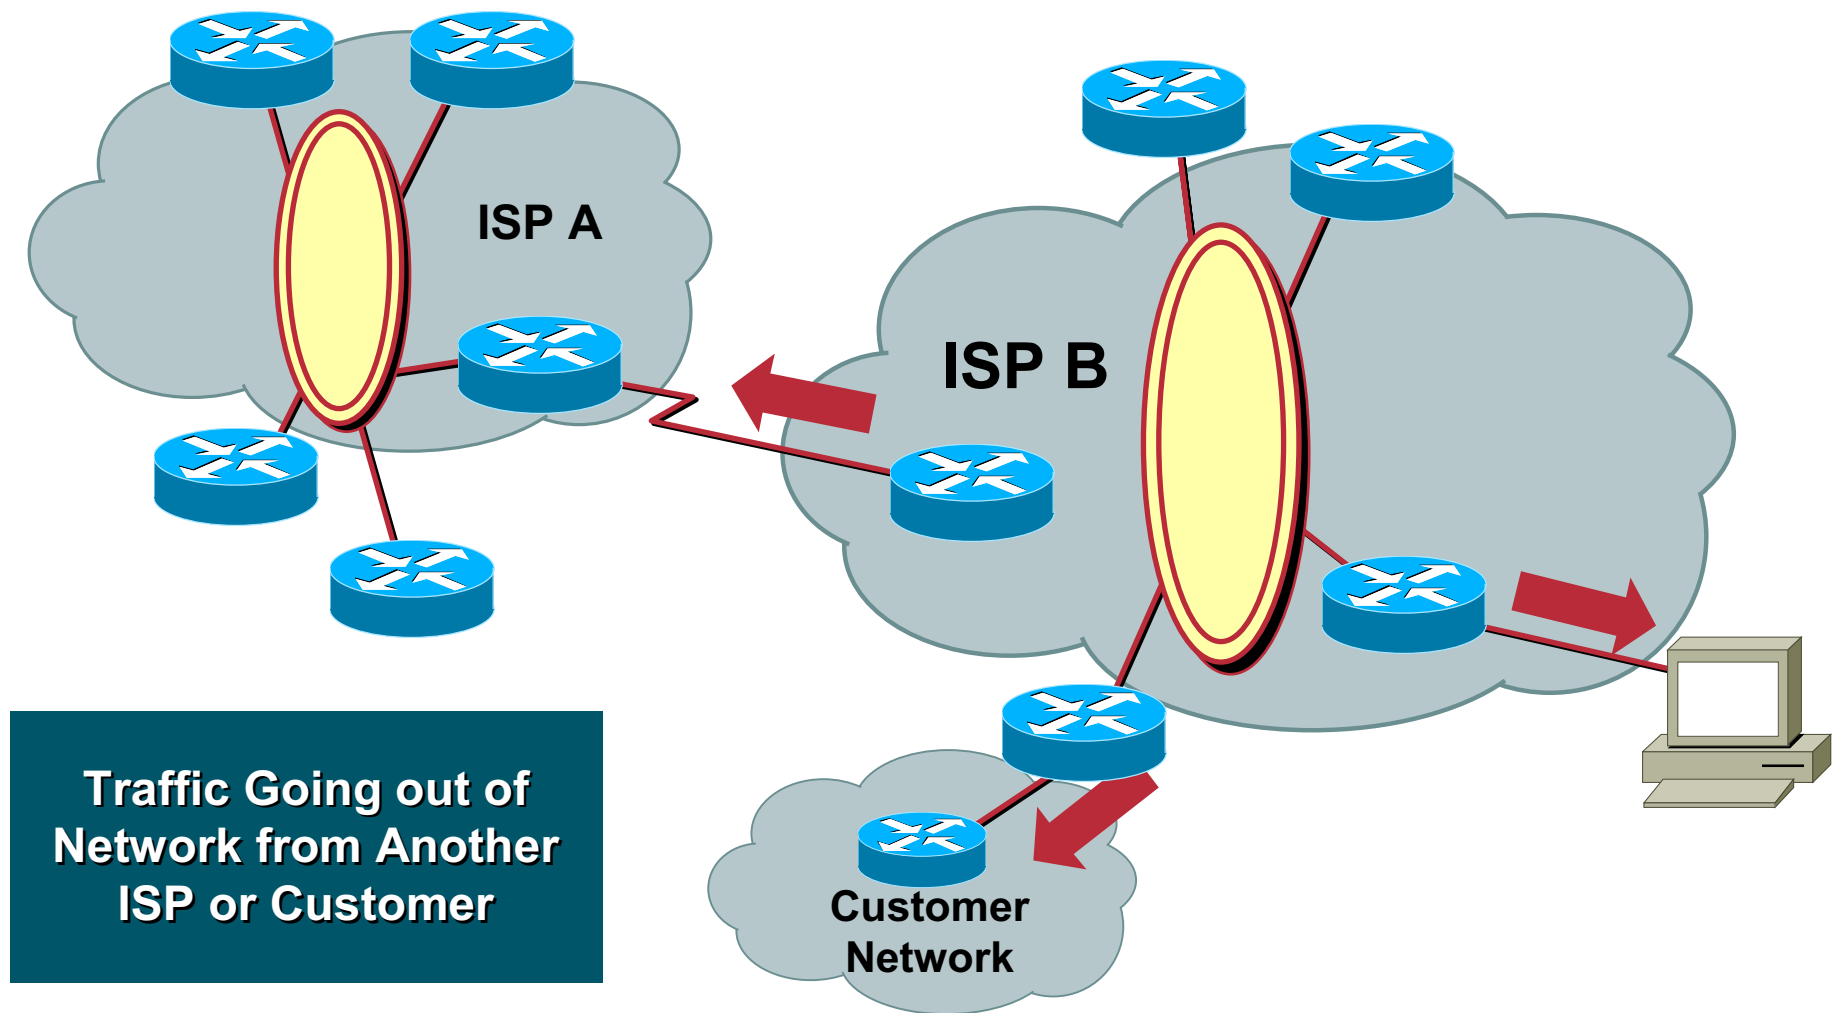
# Ingress Filters—Inbound Traffic

Cisco.com



# Egress Filters—Outbound Traffic

Cisco.com



# BCP 38 Ingress Packet Filtering

# BCP 38 Ingress Packet Filtering

Cisco.com

**Your customers should not be sending **any** IP packets out to the Internet with a source address other than the address you have allocated to them!**

# BCP 38 Ingress Packet Filtering

Cisco.com

- **BCP 38/ RFC 2827**
- **Title: Network Ingress Filtering: Defeating Denial of Service Attacks which Employ IP Source Address Spoofing**
- **Author(s): P. Ferguson, D. Senie**

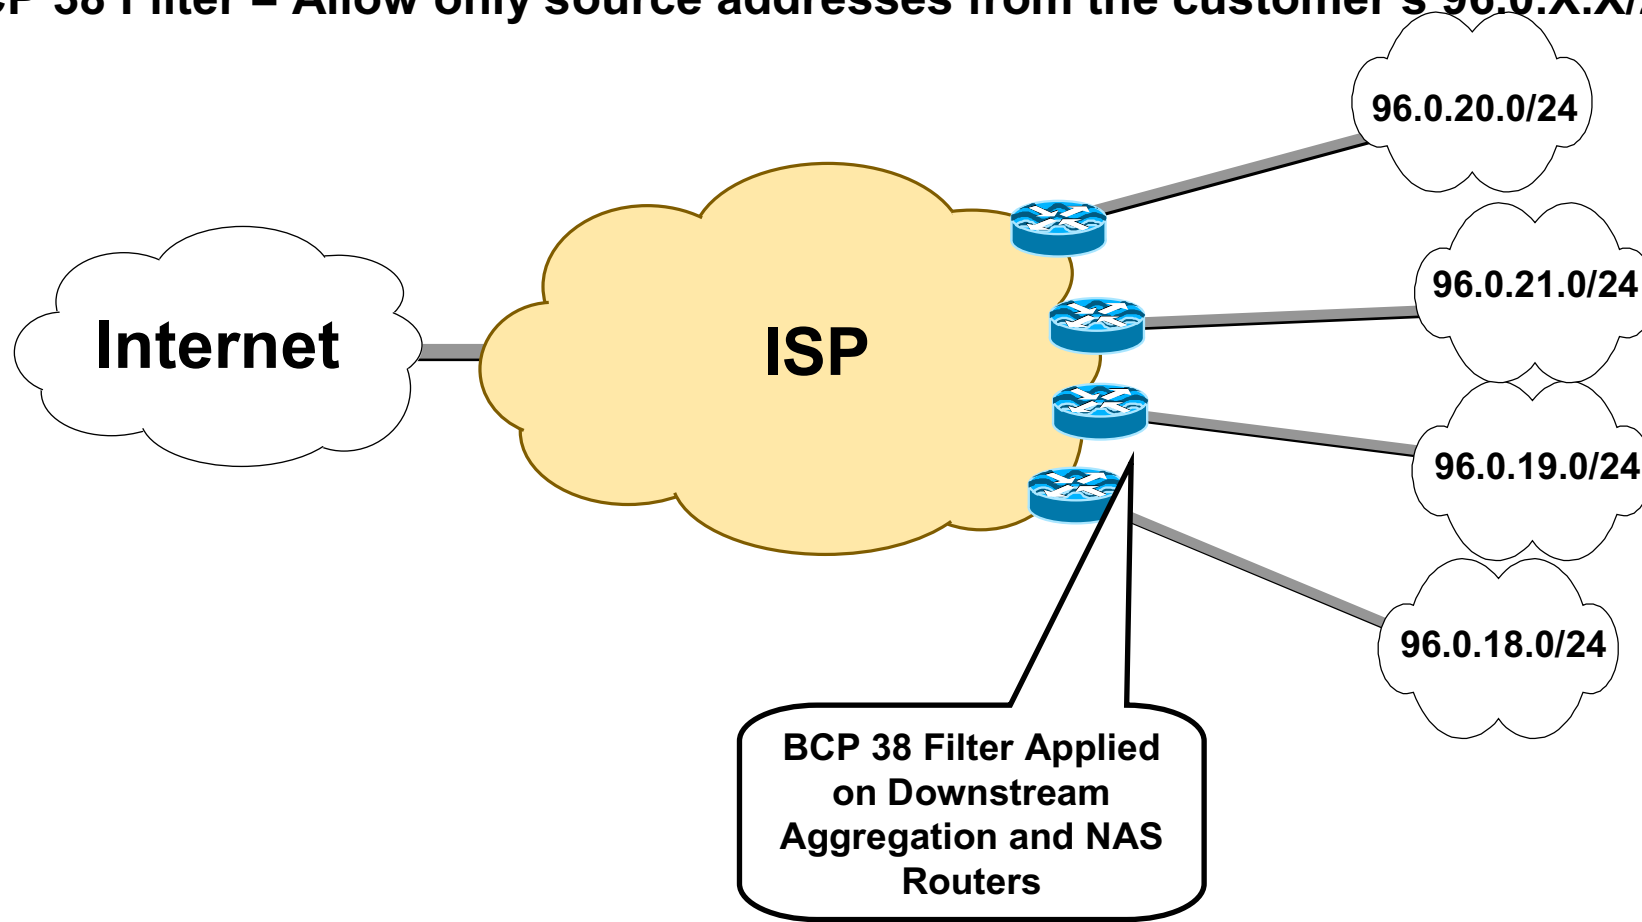


# BCP 38 Ingress Packet Filtering

Cisco.com

**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**



# BCP 38 Packet Filtering: Principles

Cisco.com

- **Filter as close to the edge as possible**
- **Filter as precisely as possible**
- **Filter both source and destination where possible**

# Techniques for BCP 38 Ingress Packet Filtering

# Three Techniques for BCP 38 Filtering

Cisco.com

- **Static access list on the edge of the network**
- **Dynamic access list with AAA profiles**
- **Unicast RPF**

# Techniques for BCP 38 Ingress Packet Filtering

## Static ACLs

# Static BCP 38 ACLs on Customer Links

Cisco.com


- **Static ACLs are the traditional mode of insuring the customer's source address matches their allocation block.**
  - ✓ **Permit All Traffic whose source address equals the allocation block.**
  - ✓ **Deny any other packet (might include logging).**

# Static BCP 38 Ingress Packet Filtering

Cisco.com

**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**



# Guidelines – Example

Cisco.com

- **Access-list 100:**
  - ✓ **Permit icmp**
  - ✓ **Permit established tcp connections (ie block TCP-SYN)**
  - ✓ **Permit SecureShell**
  - ✓ **Allow WWW to Webserver**
  - ✓ **Allow SMTP to Mailserver**
  - ✓ **Allow DNS to Nameserver**
  - ✓ **Allow NTP for time synchronisation**
  - ✓ **Block NFS**
  - ✓ **Permit only unprivileged UDP ports**
  - ✓ **Block everything else, and log it**
- **Access-list 101:**
  - ✓ **Permit only packets from my address block out**
  - ✓ **Block everything else, and log it**



# Guidelines – Example

Cisco.com

```
interface serial 0
  description Connection to Planet ISP
  ip unnumbered Ethernet 0
  ip access-group 100 in
  ip access-group 101 out
  no ip directed-broadcast
!
access-list 100 permit icmp any any
access-list 100 permit tcp any any established
access-list 100 permit tcp any any eq 22
access-list 100 permit tcp any host 221.4.0.1 eq www
access-list 100 permit tcp any host 221.4.0.2 eq smtp
access-list 100 permit udp any host 221.4.0.3 eq domain
access-list 100 permit tcp any host 221.4.0.3 eq domain
access-list 100 permit udp any any eq ntp
access-list 100 deny    udp any any eq 2049
access-list 100 permit udp any any gt 1023
access-list 100 deny    ip any any log
!
access-list 101 permit ip 221.4.0.0 0.0.3.255 any
access-list 101 deny    ip any any log
!
```

- **Lab Module 5 Exercise 1**

# Techniques for BCP 38 Ingress Packet Filtering

## Dynamic ACLs



# Techniques for BCP 38 Ingress Packet Filtering

## ***Strict Mode Unicast Reverse Path Forwarding (uRPF)***

# uRPF *Strict Mode*

Cisco.com

- **What problems is it solving?**
  - ✓ **Scaling BCP 38!**
  - ✓ **How do you manage BCP 38 ACLs for over 10,000 lease line customers?**
  - ✓ **What is needed is a command that automatically executes BCP 38 filtering.**
  - ✓ **It would be really nice if the line engineer who first brings up the customer interface can configure this feature without needing to create ACLs or touch the routing protocols.**
  - ✓ **It would be nice if the *filter* could be automatically updated.**

# uRPF *Strict Mode*

- **The answer is here .... uRPF Strict Mode.**
  - ✓ **One line config on the customer's interface.**
  - ✓ **Uses the CEF table as the certification authority for which packets pass and which are dropped – which means the the “drop table” is updated through routing.**
  - ✓ **Works on all single homed customers – which is 80% - 90% of your customers.**
  - ✓ **Works on most multi-homed customers – even with asymmetrical packet flows.**



# Unicast RPF Details

# Unicast Reverse Path Forwarding

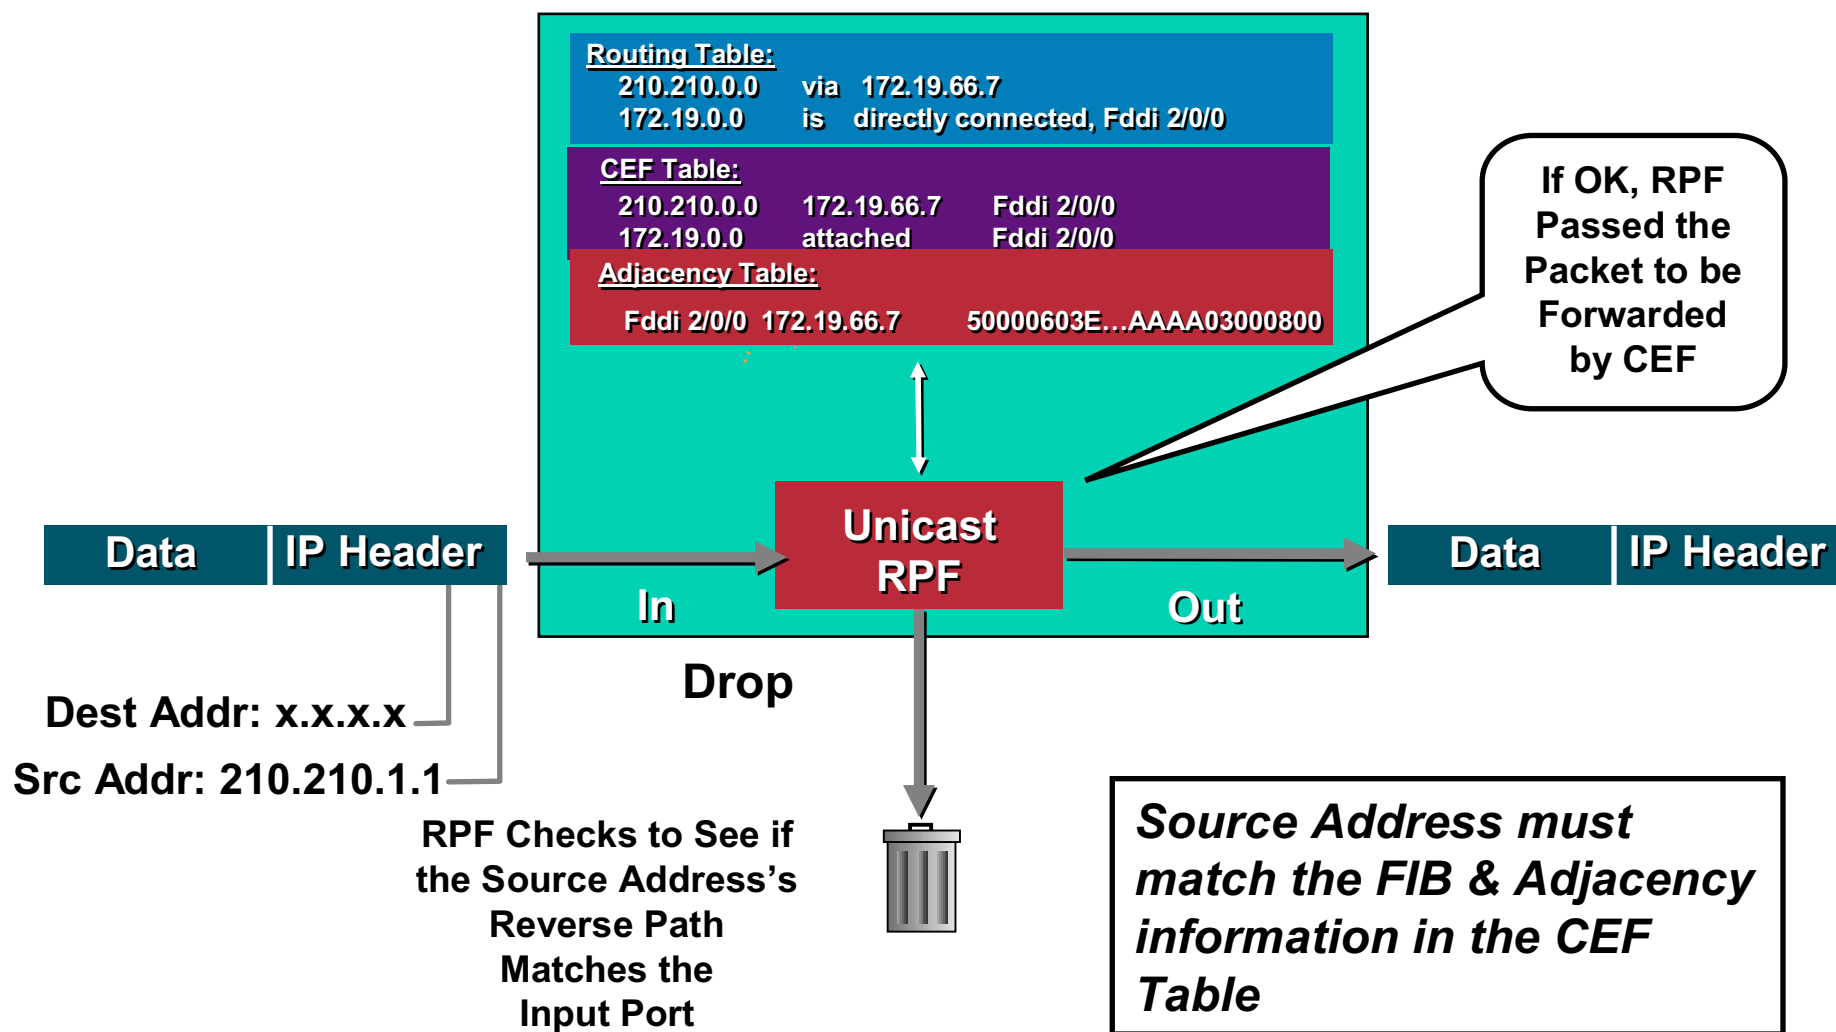
Cisco.com

- Supported from 11.1(17)CC images
- CEF switching must be enabled
- Source IP packets are checked to ensure that the route back to the source uses the same interface
- Care required in multihoming situations
- Two Flavors of uRPF:
  - ✓ Strict Mode for BCP 38/ RFC 2827 Filters on Customer Ingress Edge
  - ✓ Loose Mode for ISP ✎ ISP Edge for Remote Triggered Black Hole Filtering



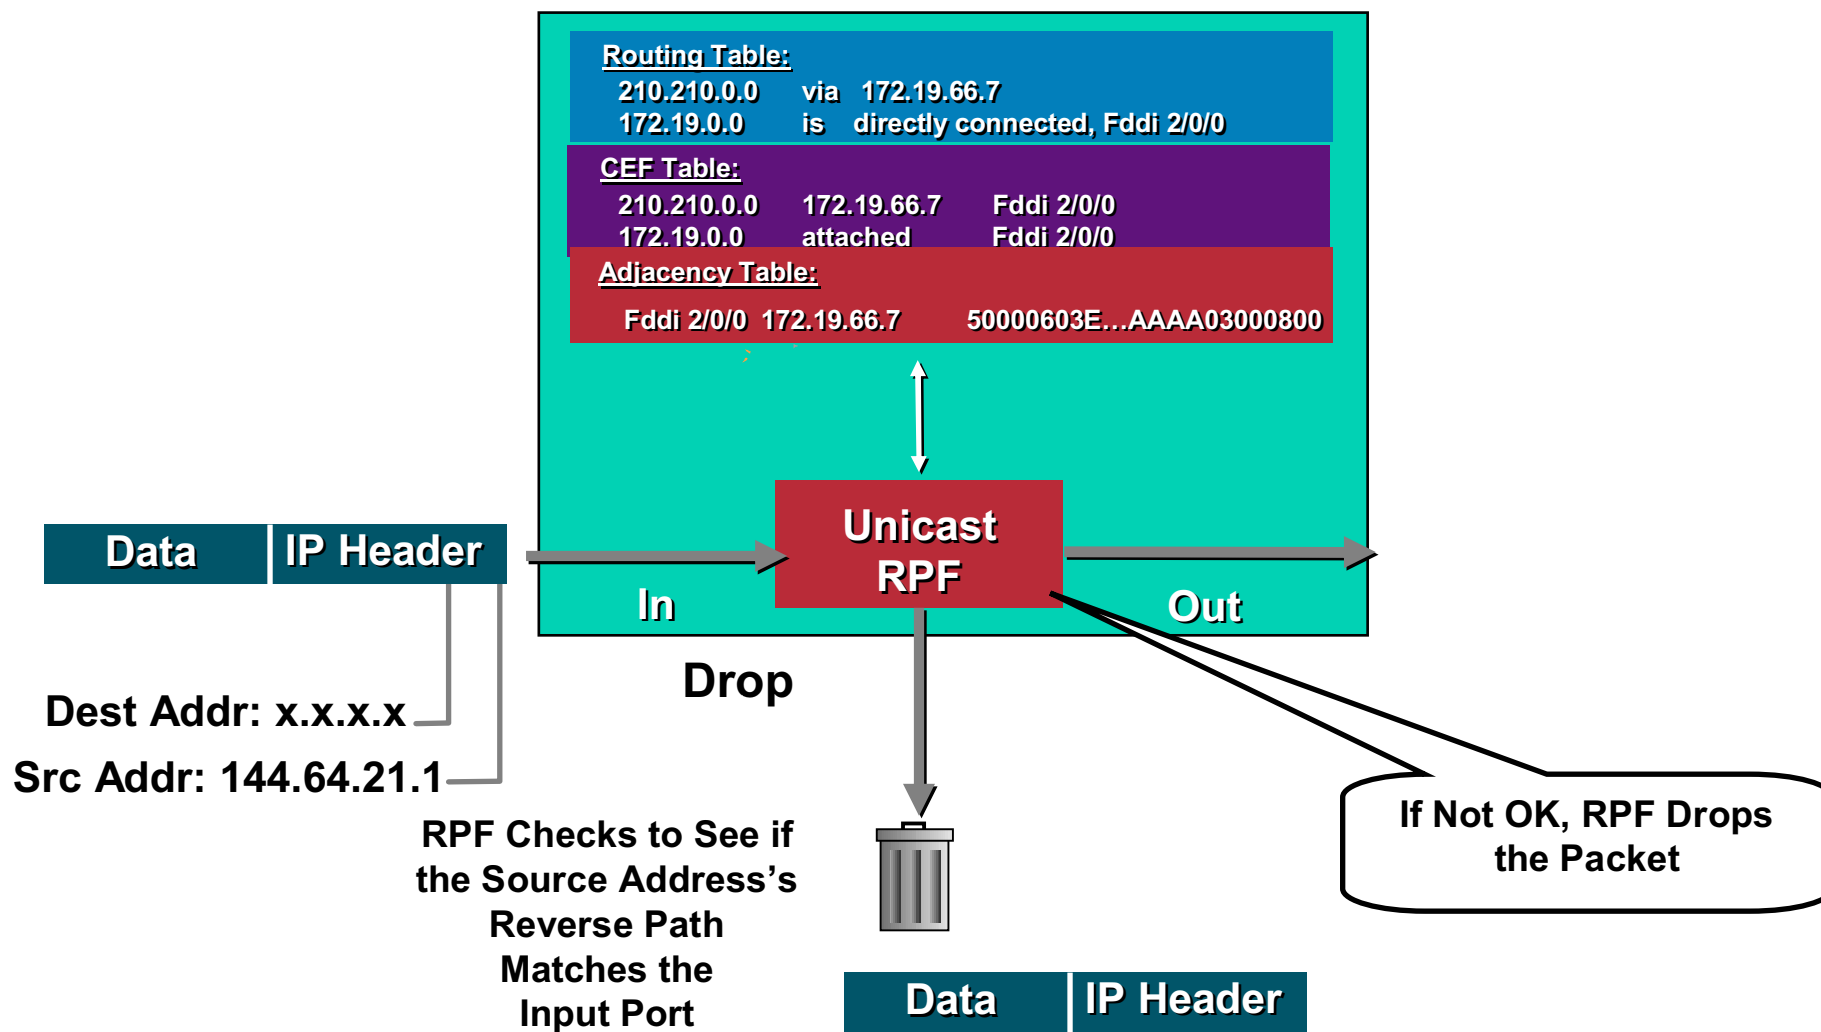
# CEF Unicast RPF (Strict Mode)

Cisco.com



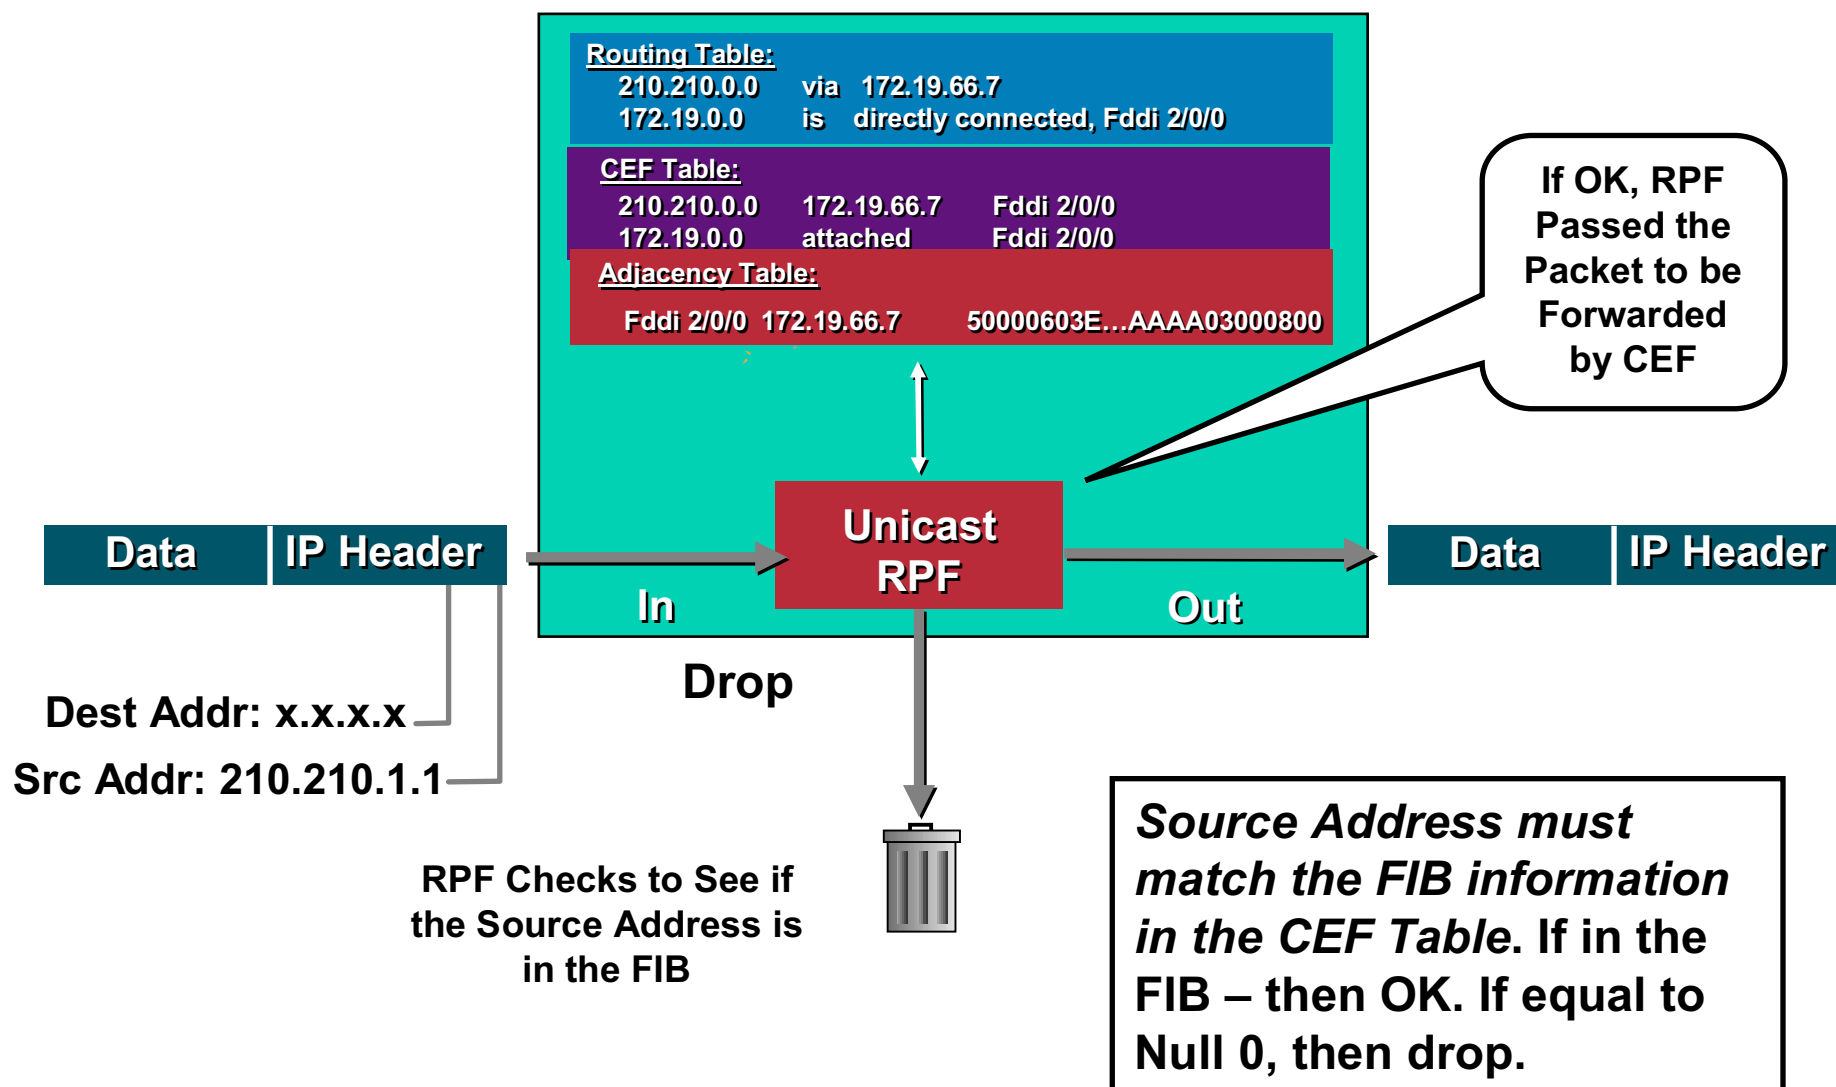
# CEF Unicast RPF (Strict Mode)

Cisco.com



# CEF Unicast RPF (Loose Check Mode)


Cisco.com



# uRPF Originally Designed for the Customer→ISP Edge

Cisco.com

- Unicast RPF was originally designed for deployment on the customer→ISP edge
- New enhancements allow it to work on the ISP→ISP edge



# Unicast RPF Commands (Strict Mode)

Cisco.com

- **Configure RPF on the interface using the following interface command syntax:**

```
[no] ip verify unicast reverse-path [<ACL>]
```

- **For example on a leased line aggregation router:**

```
ip cef ! or "ip cef distributed" for an RSP+VIP based  
box
```

```
!
```

```
interface serial 5/0/0
```

```
    ip verify unicast reverse-path
```

- ***Interface group-async* command for dial-up ports:**

```
ip cef
```

```
!
```

```
interface Group-Async1
```

```
    ip verify unicast reverse-path
```

# Unicast RPF Drop Logic (Strict Mode)

Cisco.com

- **Exceptions to RPF**

```
lookup source address in forwarding database
```

```
if the source address is reachable 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
```

```
    drop the packet
```

# Unicast RPF (Strict Mode) —Simple Single Homed Customer Example

Cisco.com


```
interface loopback 0
  description Loopback interface on Gateway Router 2
  ip address 215.17.3.1 255.255.255.255
  no ip redirects
  no ip directed-broadcast
  no ip proxy-arp
!
interface Serial 5/0
  description 128K HDLC link to Galaxy Publications Ltd [galpub1] R5-0
  bandwidth 128
  ip unnumbered loopback 0
  ip verify unicast reverse-path ! Unicast RPF activated here
  no ip redirects
  no ip directed-broadcast
  no ip proxy-arp
!
ip route 215.34.10.0 255.255.252.0 Serial 5/0
```



# Unicast RPF (Strict Mode) —Simple Single Homed Customer Example

Cisco.com

```
interface Ethernet 0
  description Galaxy Publications LAN
  ip address 215.34.10.1 255.255.252.0
  no ip redirects
  no ip directed-broadcast
  no ip proxy-arp
!
interface Serial 0
  description 128K HDLC link to Galaxy Internet Inc WT50314E C0
  bandwidth 128
  ip unnumbered ethernet 0
  ip verify unicast reverse-path ! Unicast RPF activated here
  no ip redirects
  no ip directed-broadcast
  no ip proxy-arp
!
ip route 0.0.0.0 0.0.0.0 Serial 0
```





# CEF Unicast RPF (Strict Mode)

Cisco.com

- **Unicast RPF provides**
  - ✓ **Automatic Ingress filtering based on routing information**
  - ✓ **Can be part of the default configuration**
  - ✓ **Packet drops at CEF—Before the router processes spoofed packets**
- **If this feature is so great, why is it not used?**

# Why Is Unicast RPF (Strict Mode) Not Widely Deployed?

Cisco.com

- The **myth**

- ✓ What people say:


**Unicast RPF will not work with asymmetrical routing; since the Internet has a lot of asymmetrical routing, it will not work**

- ✓ The real reason:

**ISP network engineers have not given the feature enough thought!**

# What is Asymmetrical Routing?

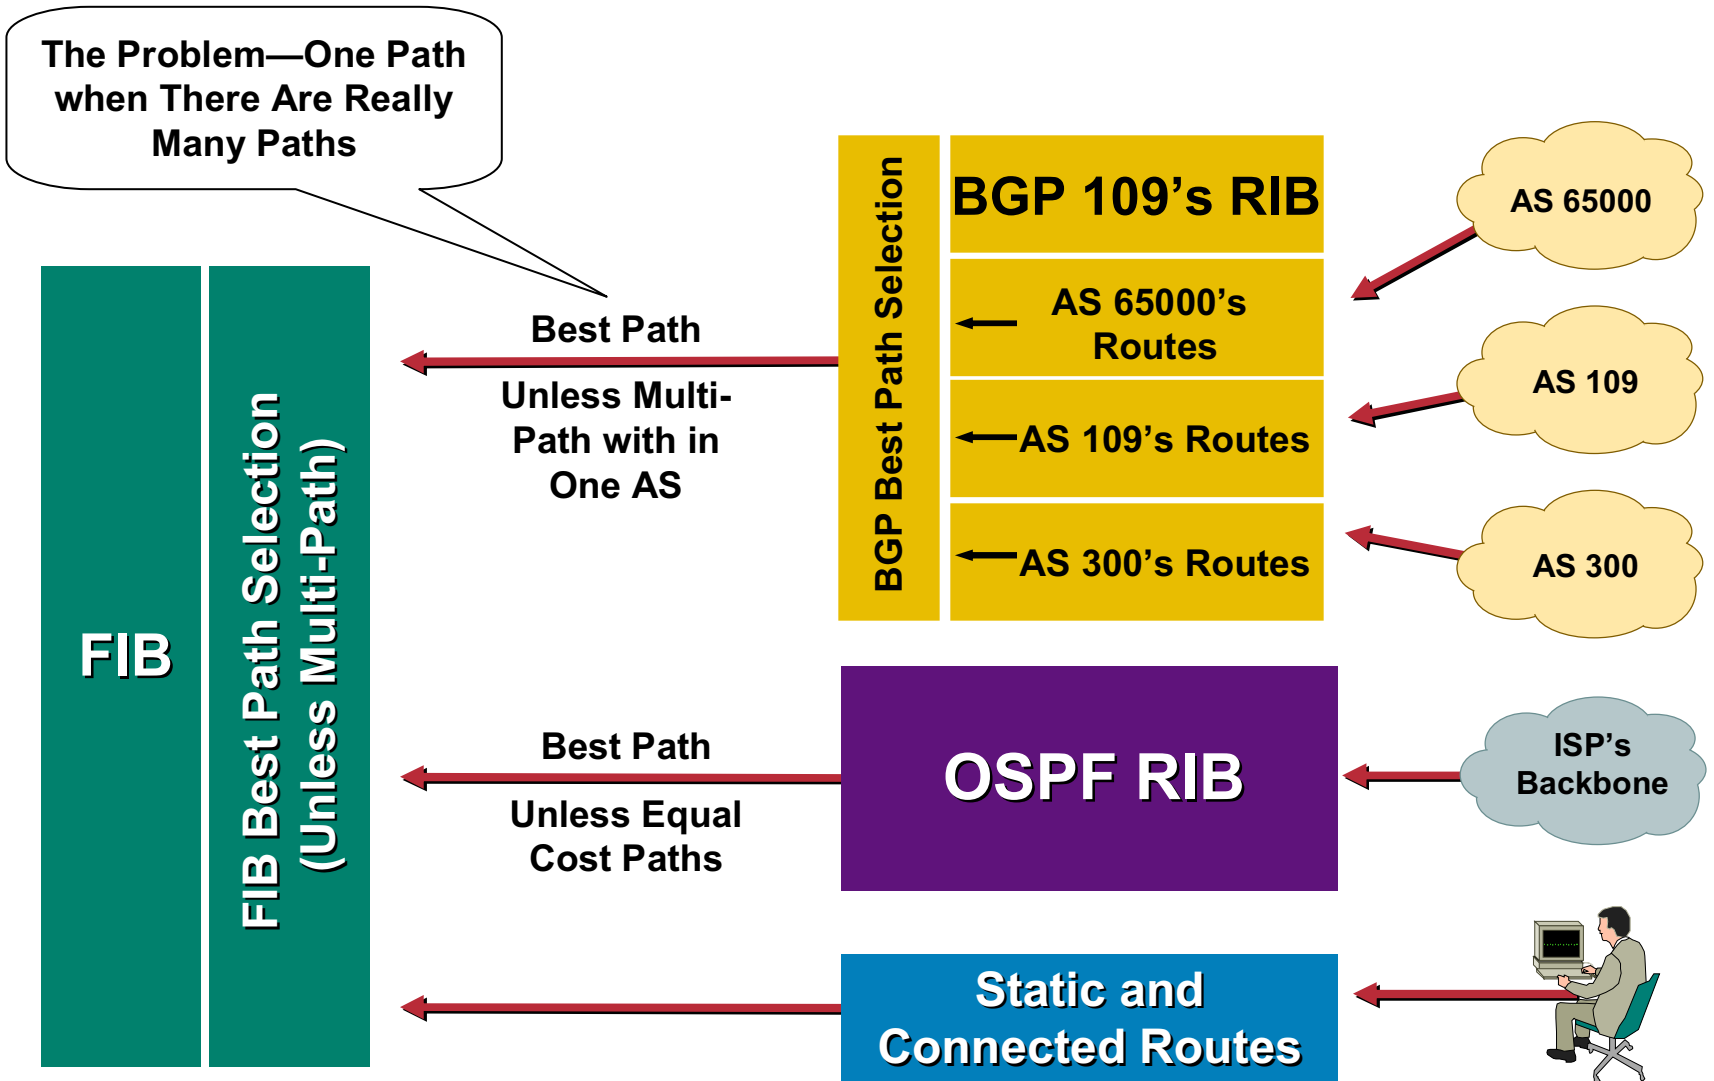
Cisco.com



**Every router makes it's  
own best path  
forwarding decision –  
resulting in  
asymmetrical routing**

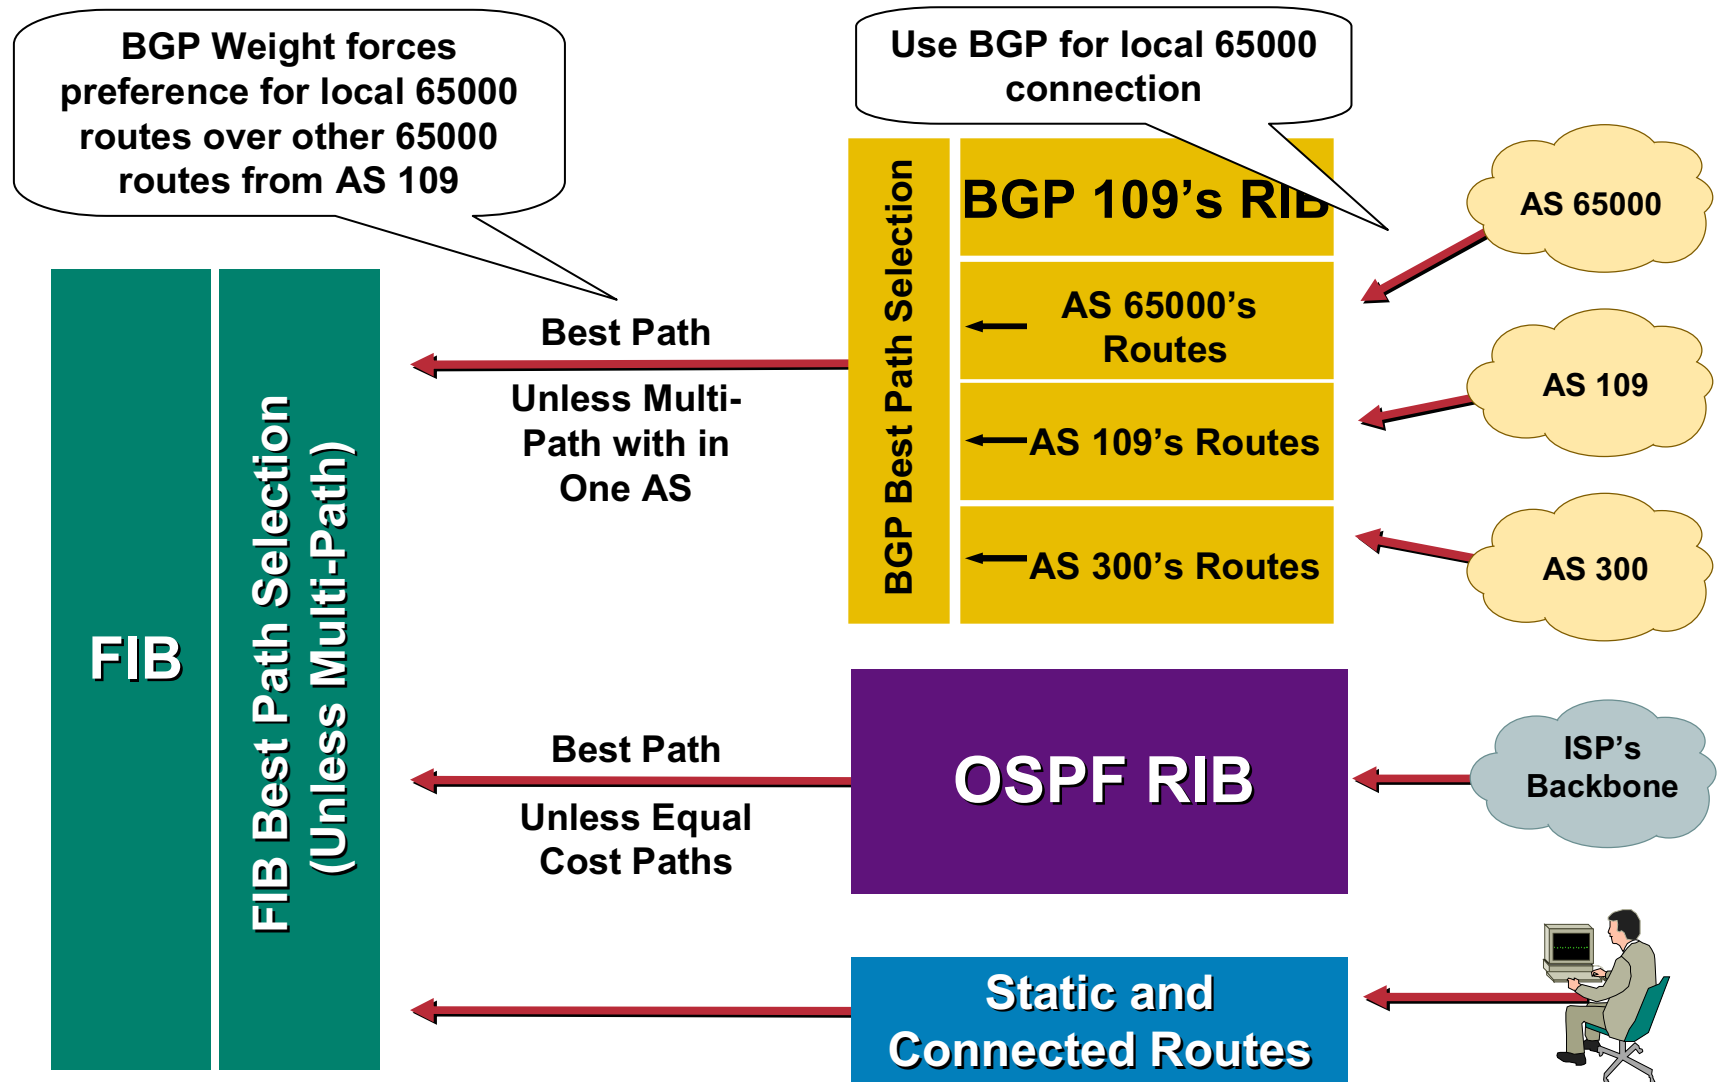
# Best Path Routing in the Internet

Cisco.com



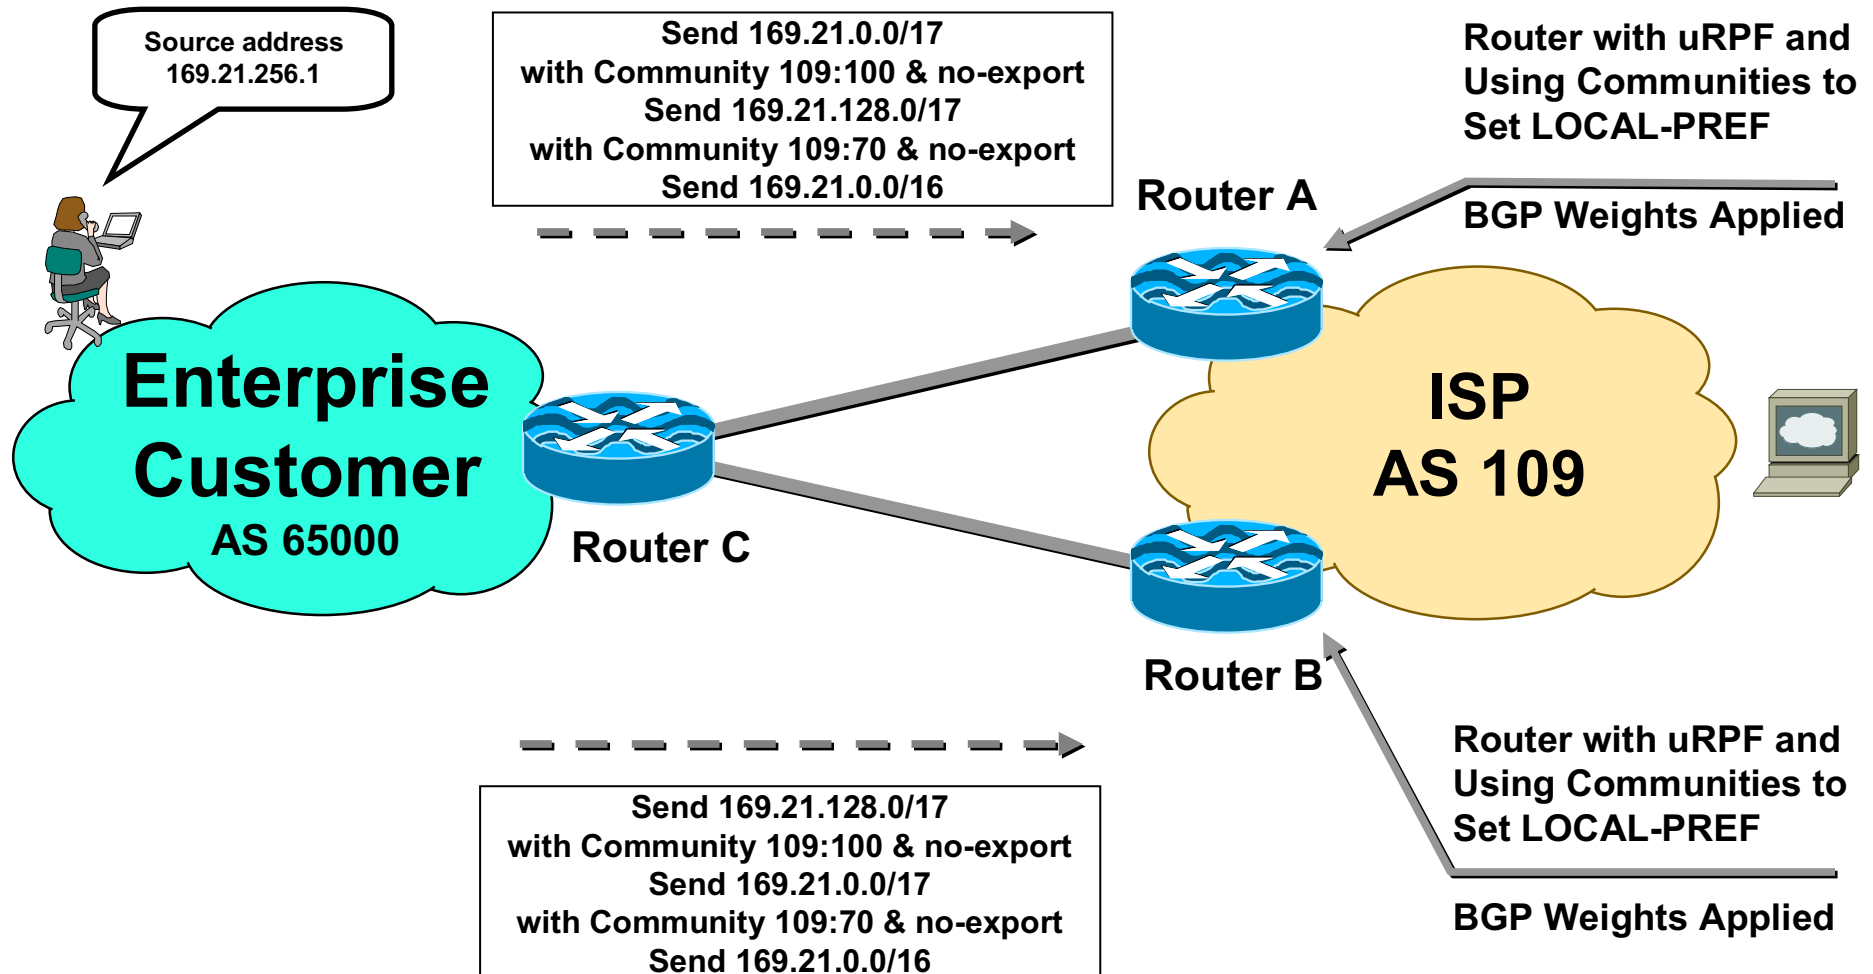
# BGP Weight aligns the FIB for uRPF

Cisco.com



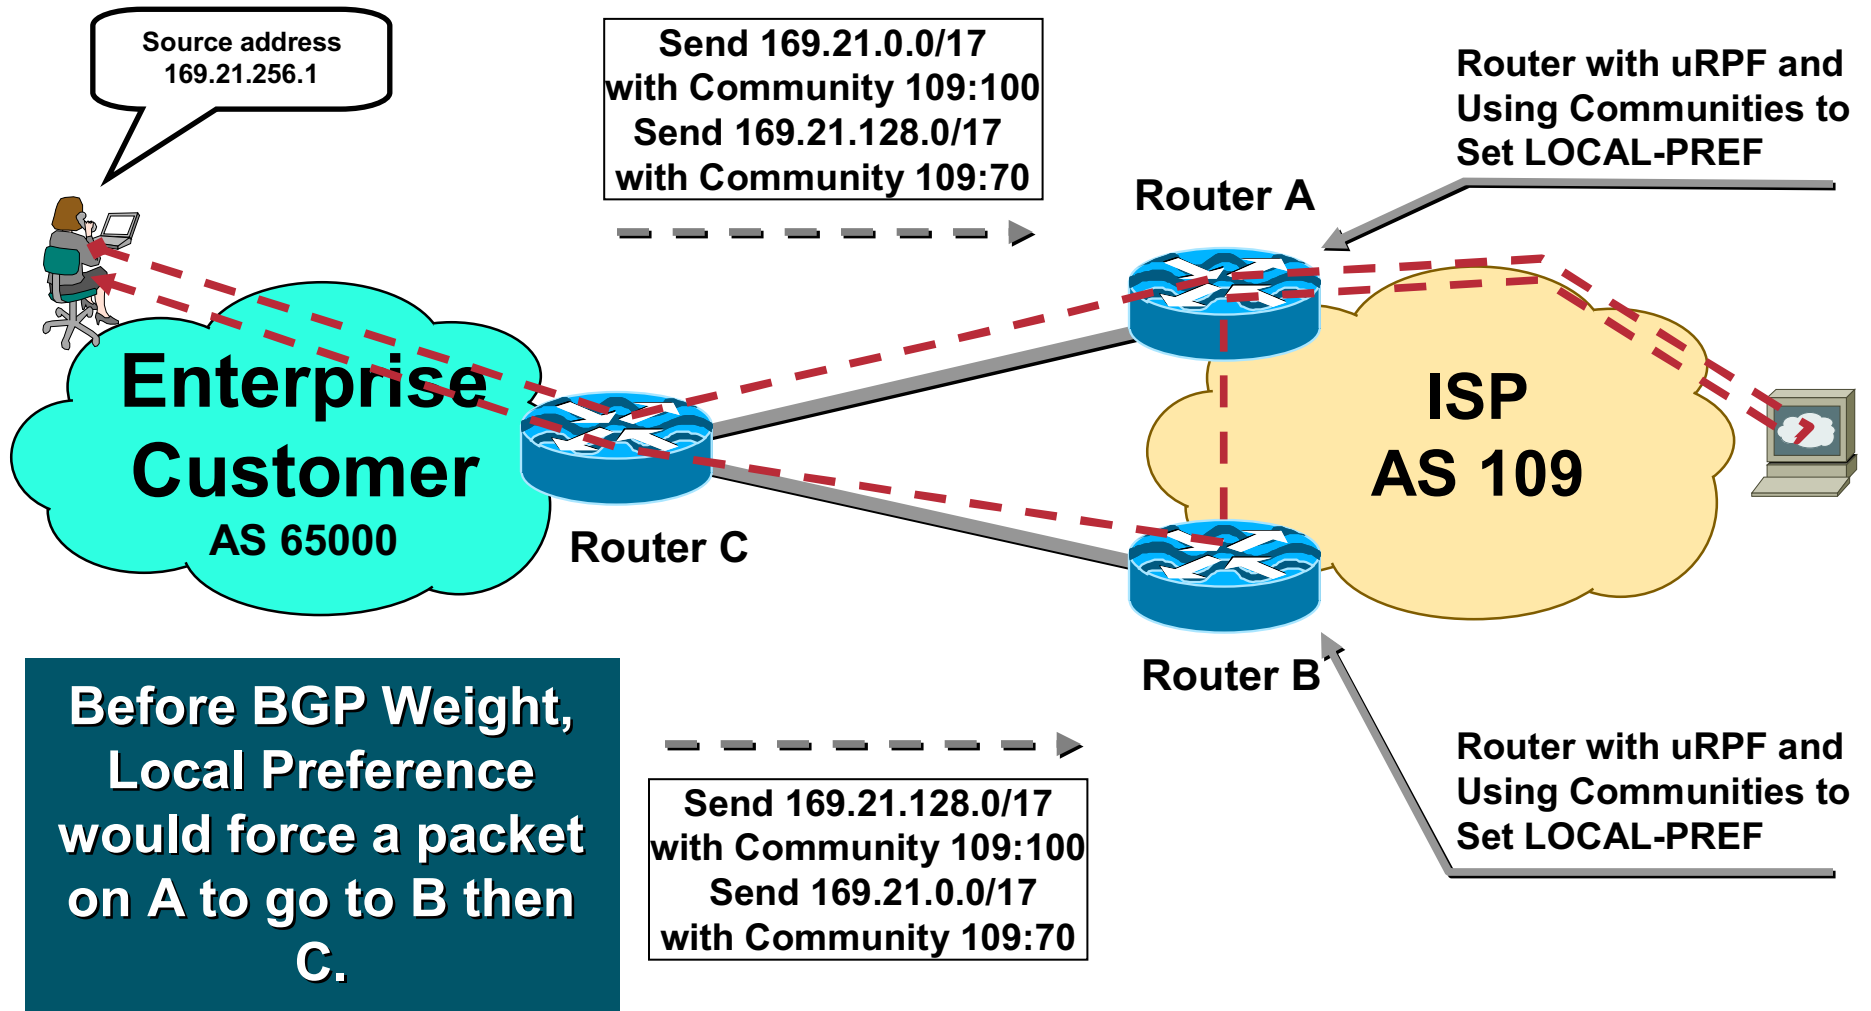
# Unicast RPF (Strict Mode) — Dual Homed Customer

Cisco.com



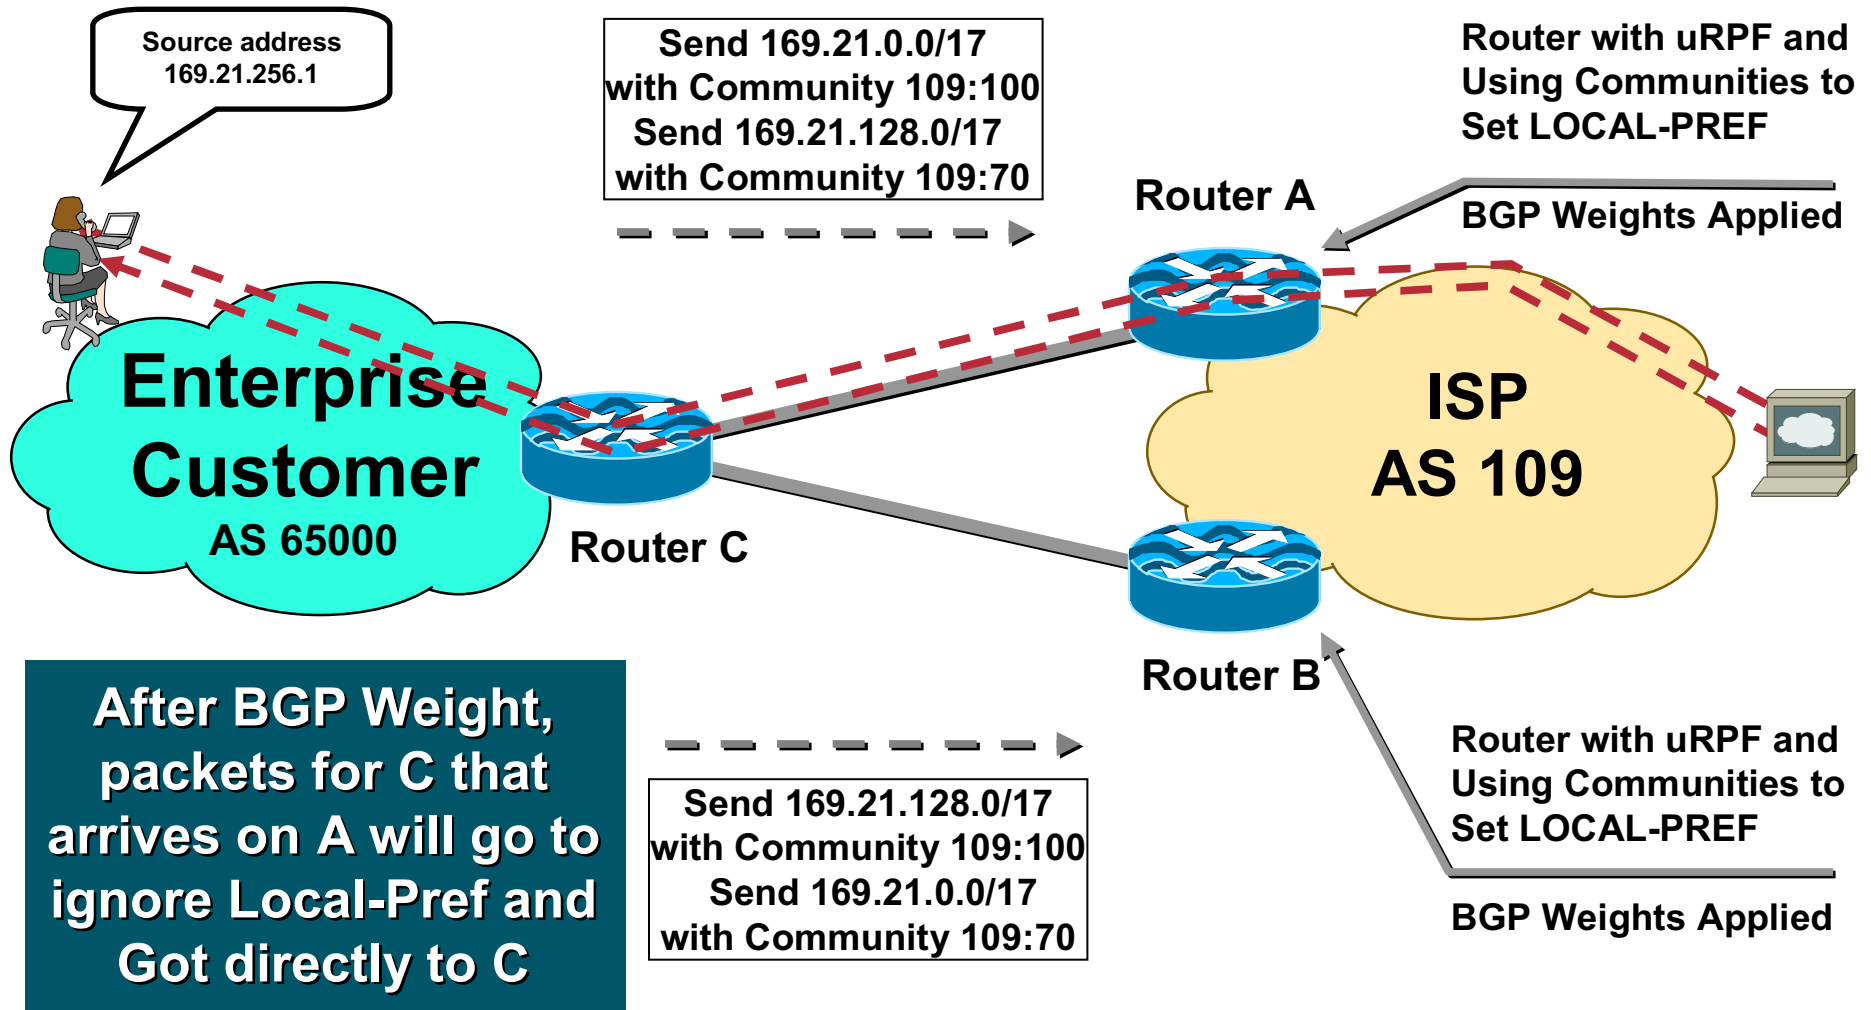
# Unicast RPF (Strict Mode) — Before BGP Weight

Cisco.com



# Unicast RPF (Strict Mode) — After BGP Weight

Cisco.com





# Unicast RPF (Strict Mode) — Dual Homed Customer

Cisco.com

## ISP Router A - Link to Customer Router C

```
interface serial 1/0/1
  description Link to Acme Computer's Router C
  ip address 192.168.3.2 255.255.255.252
  ip verify unicast reverse-path
  no ip redirects
  no ip directed-broadcast
  no ip proxy-arp
  ip route-cache distributed
```

# Unicast RPF (Strict Mode) — Dual Homed Customer


Cisco.com

## ISP Router A - Link to Customer Router C (Cont)

```
router bgp 109
  neighbor 192.168.10.3 remote-as 65000
  neighbor 192.168.10.3 description Multihomed Customer - Acme
  Computers
  neighbor 192.168.10.3 update-source Loopback0
  neighbor 192.168.10.3 send-community
  neighbor 192.168.10.3 soft-reconfiguration inbound
  neighbor 192.168.10.3 route-map set-customer-local-pref in
  neighbor 192.168.10.3 weight 255
.
ip route 192.168.10.3 255.255.255.255 serial 1/0/1
ip bgp-community new-format
```

# Unicast RPF (Strict Mode) — Dual Homed Enterprise to One ISP

Cisco.com



- Used to protect against spoof attacks
- Some attacks get around the RFC1918 filters by using un-allocated IP address space

# Unicast RPF (Strict Mode) — Dual Homed Enterprise to One ISP

Cisco.com

```
router bgp 65000
no synchronization
network 169.21.0.0
network 169.21.0.0 mask 255.255.128.0
network 169.21.128.0 mask 255.255.128.0
neighbor 171.70.18.100 remote-as 109
neighbor 171.70.18.100 description Upstream Connection #1
neighbor 171.70.18.100 update-source Loopback0
neighbor 171.70.10.100 send-community
neighbor 171.70.18.100 soft-reconfiguration inbound
neighbor 171.70.18.100 route-map Router-A-Community out
neighbor 171.70.18.200 remote-as 109
neighbor 171.70.18.200 description Upstream Connection #2
neighbor 171.70.18.200 update-source Loopback0
neighbor 171.70.18.200 send-community
neighbor 171.70.18.200 soft-reconfiguration inbound
neighbor 171.70.18.200 route-map Router-B-Community out
maximum-paths 2
no auto-summary
```

```
route-map Router-A-Community permit 10
match ip address 51
set community 109:70
```

!

```
route-map Router-A-Community permit 20
match ip address 50
set community 109:100
```

!

```
route-map Router-B-Community permit 10
match ip address 50
set community 109:70
```

!

```
route-map Router-B-Community permit 20
match ip address 51
set community 109:100
```

!

```
access-list 50 permit 169.21.0.0 0.0.127.255
```

```
access-list 51 permit 169.21.128.0 0.0.127.255
```

# Unicast RPF (Strict Mode) — Dual Homed Enterprise to One ISP

Cisco.com

```
ip route 169.21.0.0 0.0.255.255 Null 0
ip route 169.21.0.0 0.0.127.255 Null 0
ip route 169.21.128.0 0.0.127.255 Null 0
```

```
ip route 171.70.18.100 255.255.255.255 S 1/0
ip route 171.70.18.200 255.255.255.255 S 1/1
ip bgp-community new-format
!
```

```
interface serial 1/0/
description Link to Upstream Router A
ip address 192.168.3.1 255.255.255.252
ip verify unicast reverse-path
no ip redirects
no ip directed-broadcast
no ip proxy-arp
ip load-sharing per-destination
ip route-cache distributed
!
interface serial 1/0
description Link to Upstream ISP Router B
ip address 192.168.3.5 255.255.255.252
ip verify unicast reverse-path
no ip redirects
no ip directed-broadcast
no ip proxy-arp
ip load-sharing per-destination
ip route-cache distributed
```


# Unicast RPF (Strict Mode) — Dual Homed Enterprise to One ISP

Cisco.com

- **The results:**
  - ✓ The customer has a multihomed connection to the Internet **with** Unicast RPF protecting source spoofing
  - ✓ The ISP provides a multihomed solution with Unicast RPF turned on

# Unicast RPF (Strict Mode) — Dual Homed Enterprise to Two ISPs


Cisco.com



- **ISP Configuration for both ISPs are similar to a dual homed customer.**
  - ✓ **BGP weight** is used to over ride AS path prepends

# Unicast RPF (Strict Mode) — Dual Homed Enterprise to Two ISPs

Cisco.com




- BGP weight override an AS path prepend
  - ✓ BGP weight on Router A will keep the preferred path for packets on that router to be C↔A
  - ✓ BGP weight on Router B will keep the preferred path for packets on that router to be C↔B



# Unicast RPF (Strict Mode) — Dual Homed Enterprise to Two ISPs

Cisco.com



- Enterprise configuration cannot use **maximum-paths**
  - ✓ Need equal AS paths for maximum-paths to work

# Unicast RPF (Strict Mode) — The ACL Bypass Option

Cisco.com

- **ACLs can now be used with Unicast RPF (Strict Mode):**  
`ip verify unicast reverse-path 171`
- **uRPF ACLs are used to:**
  - ✓ Allow exceptions to the Unicast RPF check
  - ✓ Identify characteristics of spoofed packets being dropped by Unicast RPF
- **Software Forwarding Only! Not Supported on uRPF in the Forwarding ASICs (i.e. Engine 2, Engine 4, etc.)**

# Unicast RPF (Strict Mode) — The ACL Bypass Option

Cisco.com

- **Cisco 7206 with bypass ACL**

```
interface ethernet 1/1
```

```
ip address 192.168.200.1 255.255.255.0
```

```
ip verify unicast reverse-path 197
```

```
!
```

```
access-list 197 permit ip 192.168.201.0 0.0.0.255 any log-input
```

```
show ip interface ethernet 1/1 | include RPF
```

```
Unicast RPF ACL 197
```

```
1 unicast RPF drop
```

```
1 unicast RPF suppressed drop
```

# Unicast RPF (Strict Mode) — The ACL Bypass Option

Cisco.com

- **Show the “log-input” results:**

- ✓ **7200—logging done in the RP**

- show logging**

- ✓ **7500—logging done on the VIP**

**Excalabur#sh controllers vip 4 logging**

**show logging from Slot 4:**

▪

**4d00h: %SEC-6-IPACCESSLOGNP: list 171 denied 0 20.1.1.1  
-> 255.255.255.255, 1 packet**

▪

# Unicast RPF (Strict Mode) — The ACL Bypass Option

Cisco.com

- **NOTE – ACL Bypass option will not be in the ASIC implementations of uRPF.**
  - ✓ Cisco 12XXX Engine 2
  - ✓ Cisco 12XXX Engine 3
  - ✓ Cisco 12XXX Engine 4+
  - ✓ Cisco 7600 Sup 2

# Unicast RPF (Strict Mode) — Operations Tools

Cisco.com

```
Excalabur#sh cef inter serial 2/0/0
```

```
Serial2/0/0 is up (if_number 8)
```

```
Internet address is 169.223.10.2/30
```

```
ICMP redirects are never sent
```

```
Per packet loadbalancing is disabled
```

```
IP unicast RPF check is enabled
```

```
Inbound access list is not set
```

# Unicast RPF (Strict Mode) — Operations Tools

Cisco.com

- **Other commands:**
  - ✓ **show ip traffic | include RPF**
  - ✓ **show ip interface ethernet 0/1/1 | include RPF**
  - ✓ **debug ip cef drops rpf <ACL>**

# Unicast RPF (Strict Mode) — Bottom Line

Cisco.com

- **Unicast RFP Strict Mode is designed to help scale BCP 38 filtering.**
- **It is just another tool to help defend the Internet**
- **It is not a perfect tool.**
- **Deploying it will easily cover 80% of your customers (those who are single homed).**
- **Can can be deployed on the last 20% that are multihomed customers.**



- **Module 5 Exercise 5**

# Unicast RPF Enhancements

## *Loose Check*

# New Unicast RPF Enhancements

Cisco.com

- **Objectives—Allow Unicast RPF to work on an ISP-ISP Edge or ISP-Complex multihomed enterprise customer edge**
  - ✓ **Phase 1—Original uRPF (BCP 38/ RFC 2827)**
  - ✓ **Phase 2—Loose check — if exist in FIB**
  - ✓ **Phase 3—Dedicated VRF table per interface**

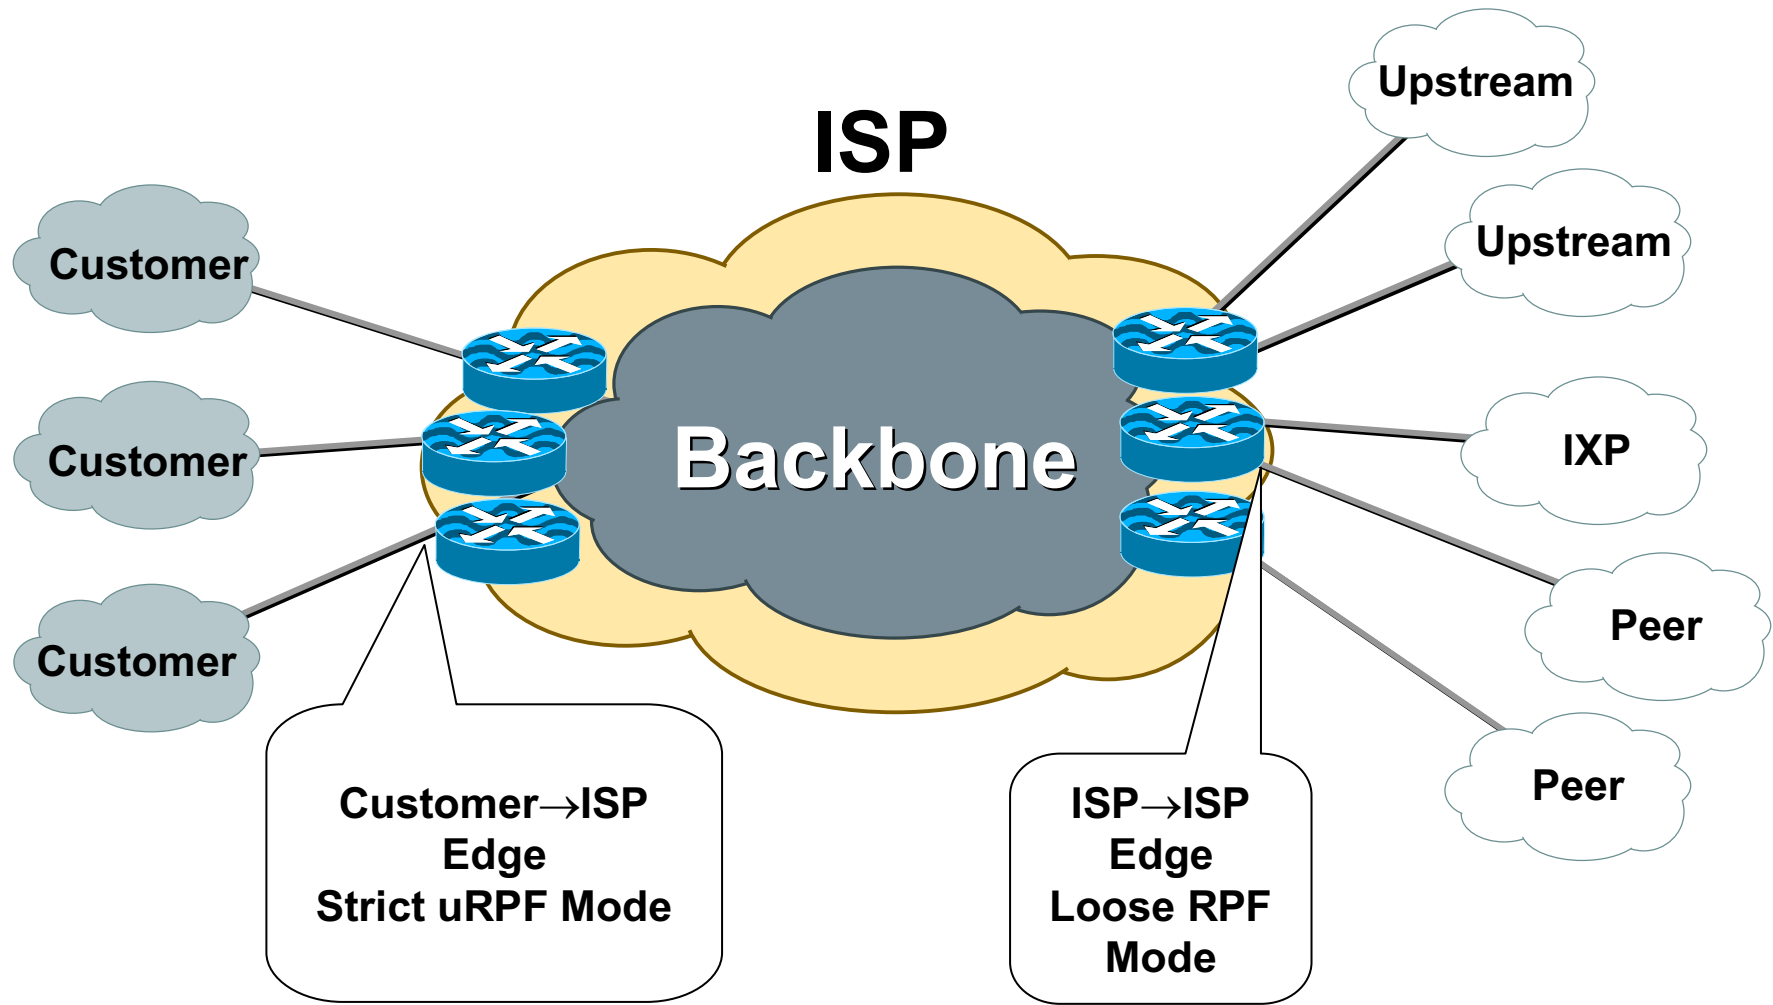
# New Unicast RPF Enhancements

Cisco.com

- **Phase 2—Loose check (if exist)**
  - ✓ **DDTS CSCdr93424**
  - ✓ **12.0(14)S for 7200, 7500, and GSR Engine 0 and 1**
  - ✓ **12.0(19)S for GSR Engine 2**  
**In ASIC with 1.6 Mpps performance**
  - ✓ **Scheduled 12.1(8)E for CAT6K**  
**In ASIC at line rate**
  - ✓ **Cisco 12XXX Engine 3 – 12.0(22)S**  
**In ASIC at line rate**

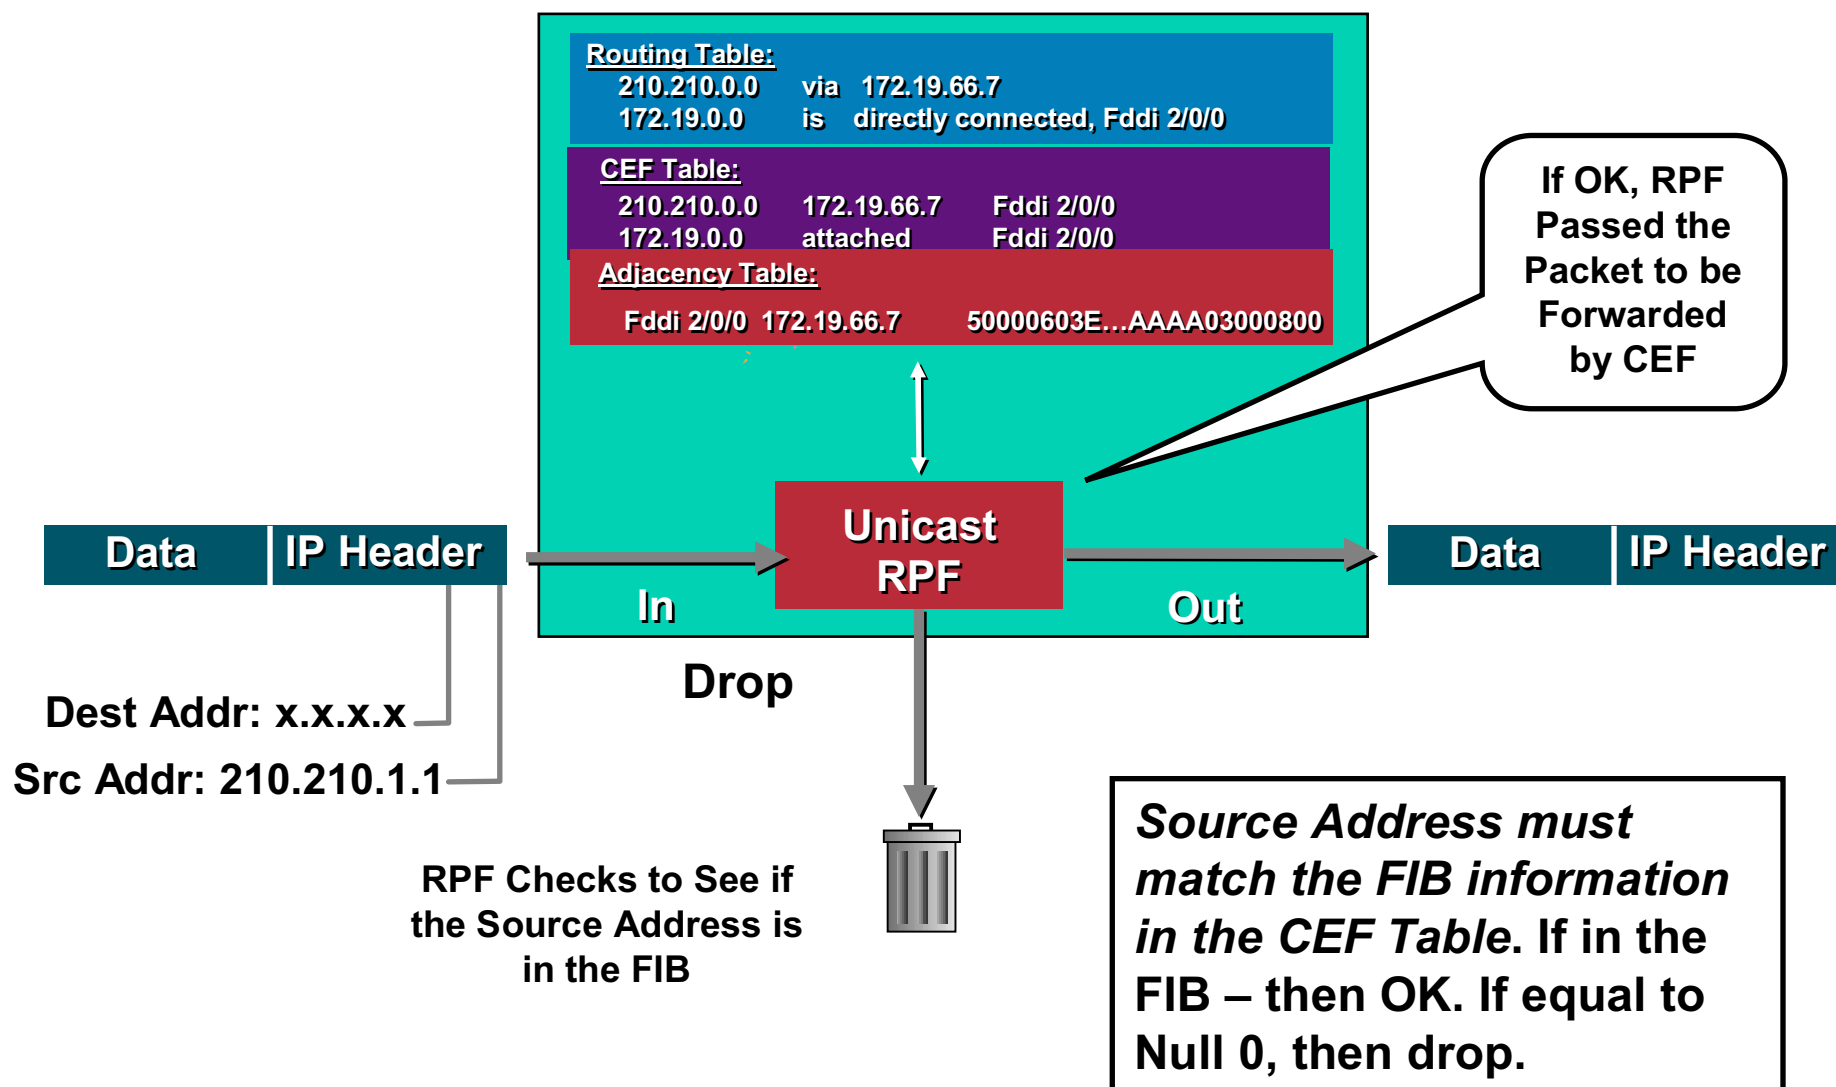
# uRPF Originally Designed for the Customer→ISP Edge

Cisco.com



# CEF Unicast RPF (Loose Check Mode)

Cisco.com



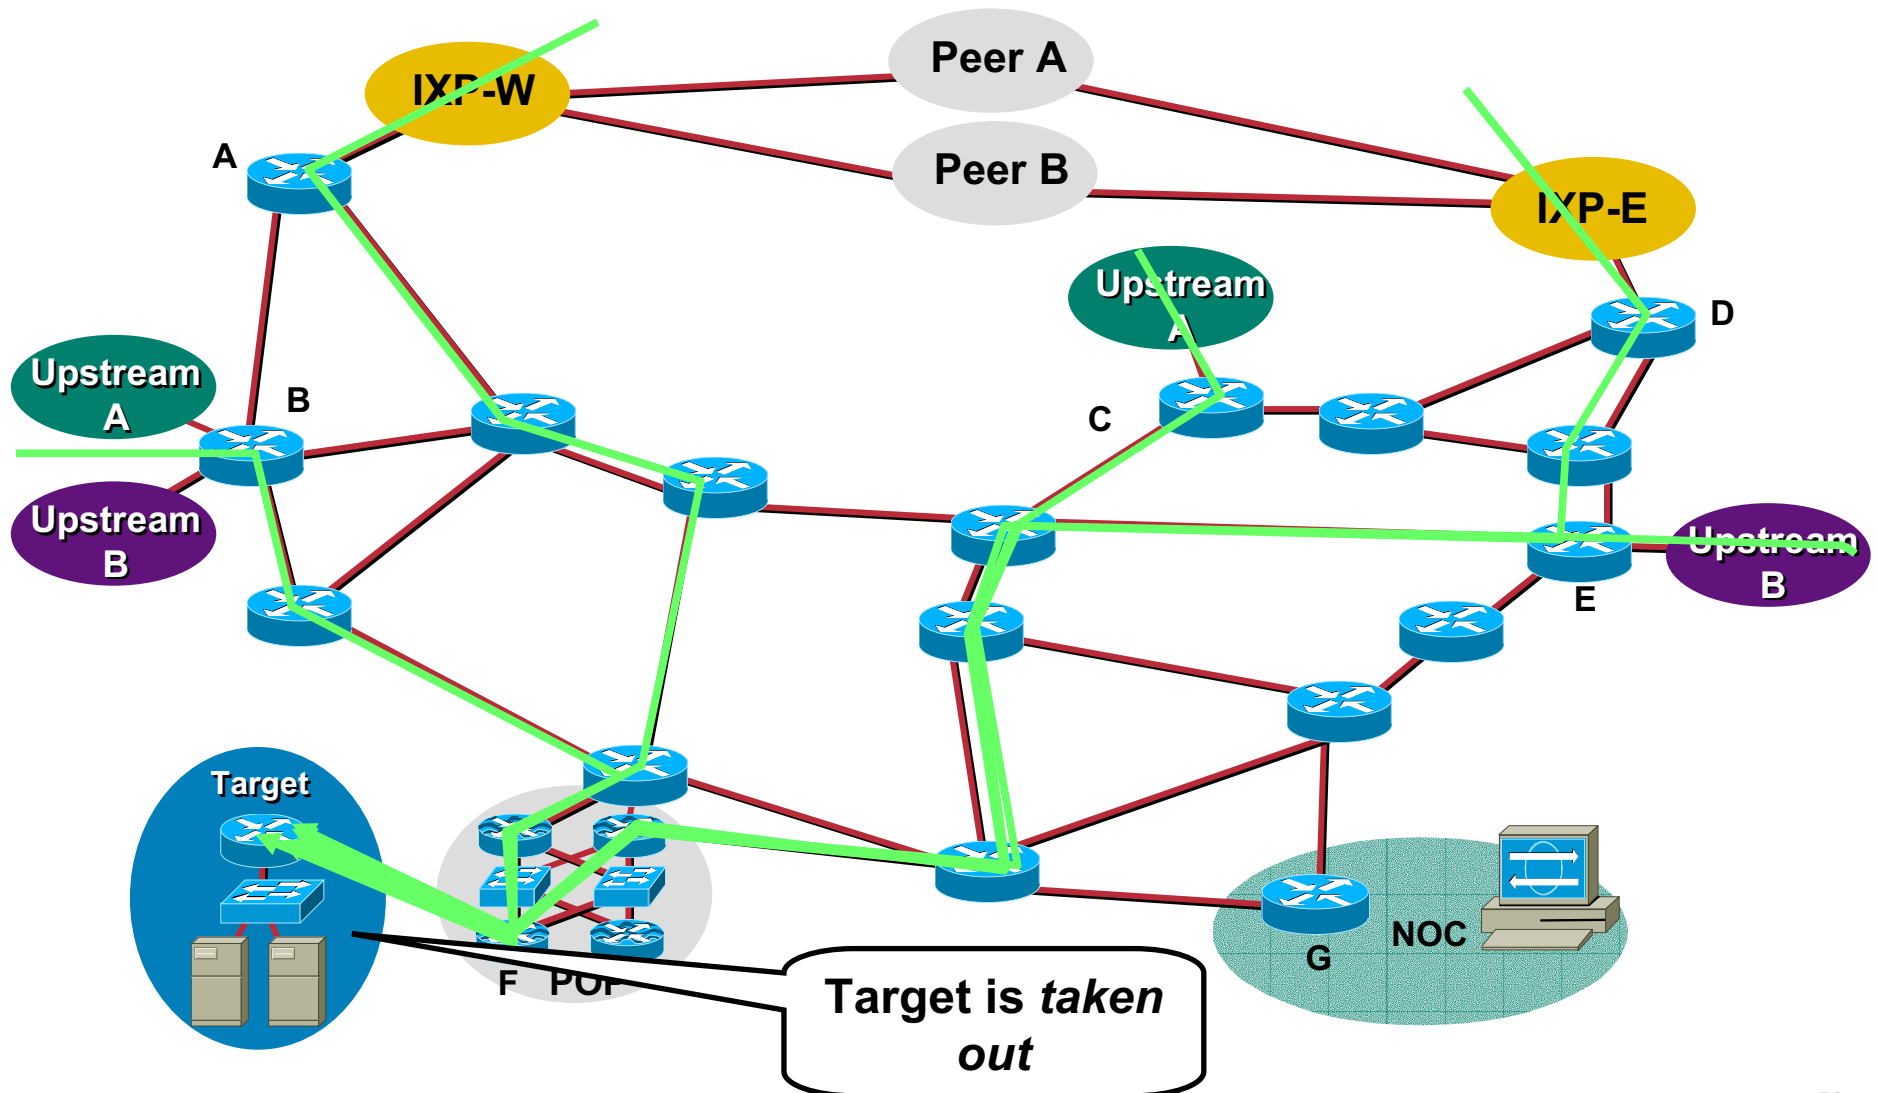
# New Unicast RPF Enhancements

Cisco.com

- **Objectives in phase 2:**
  - ✓ **Allow for uRPF to work on the ISP ↗ ISP edge of the network**
  - ✓ **Source Based – Remote Triggered Black Hole Filtering. Create a new tool to drop DOS/DDOS attacks on the edge of an ISP's network**

# Network Wide Black Hole of an Attack Flow - Before

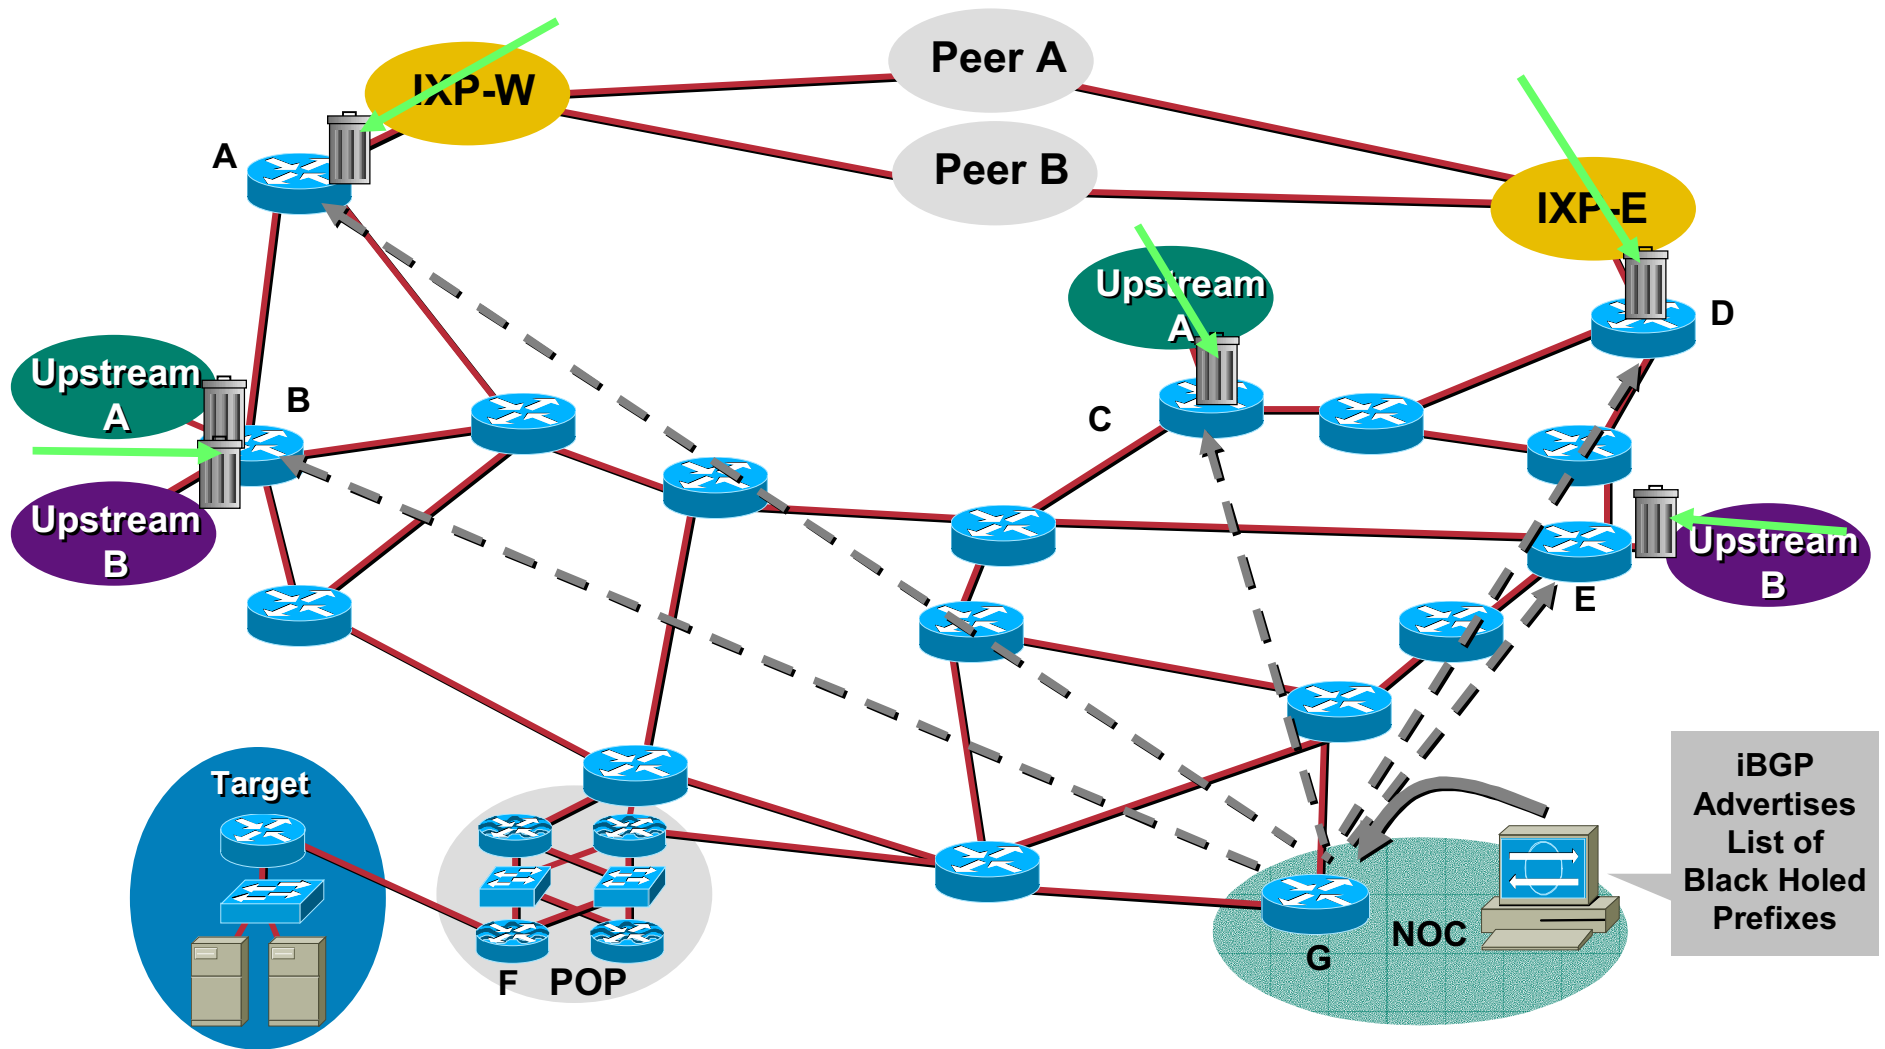
Cisco.com





# Network Wide Black Hole of an Attack Flow - Before

Cisco.com



# New Unicast RPF Enhancements

Cisco.com

- New commands from DDTS CSCdr93424:

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

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

# Data Plane Security

## *Deployment Considerations*

# Deployment Considerations

Cisco.com

- **#1 Consideration ..... will data plane security cause an adverse impact on the service?**
  - ✓ **PPS Rates**
  - ✓ **Stability**
  - ✓ **Complexity**

# Deployment Considerations

Cisco.com

- **Two modes of network gear:**
  - ✓ **CPU Based forwarding/feature processing**
  - ✓ **ASIC Based forwarding/feature processing**
- **Cisco does both.**
- **ASICs – especially today's 2<sup>nd</sup> ASICs – are mission specific.**
  - ✓ **Which means they are designed for a specific niche in the network.**

# Data Plane Security Summary