
CISCO SYSTEMS

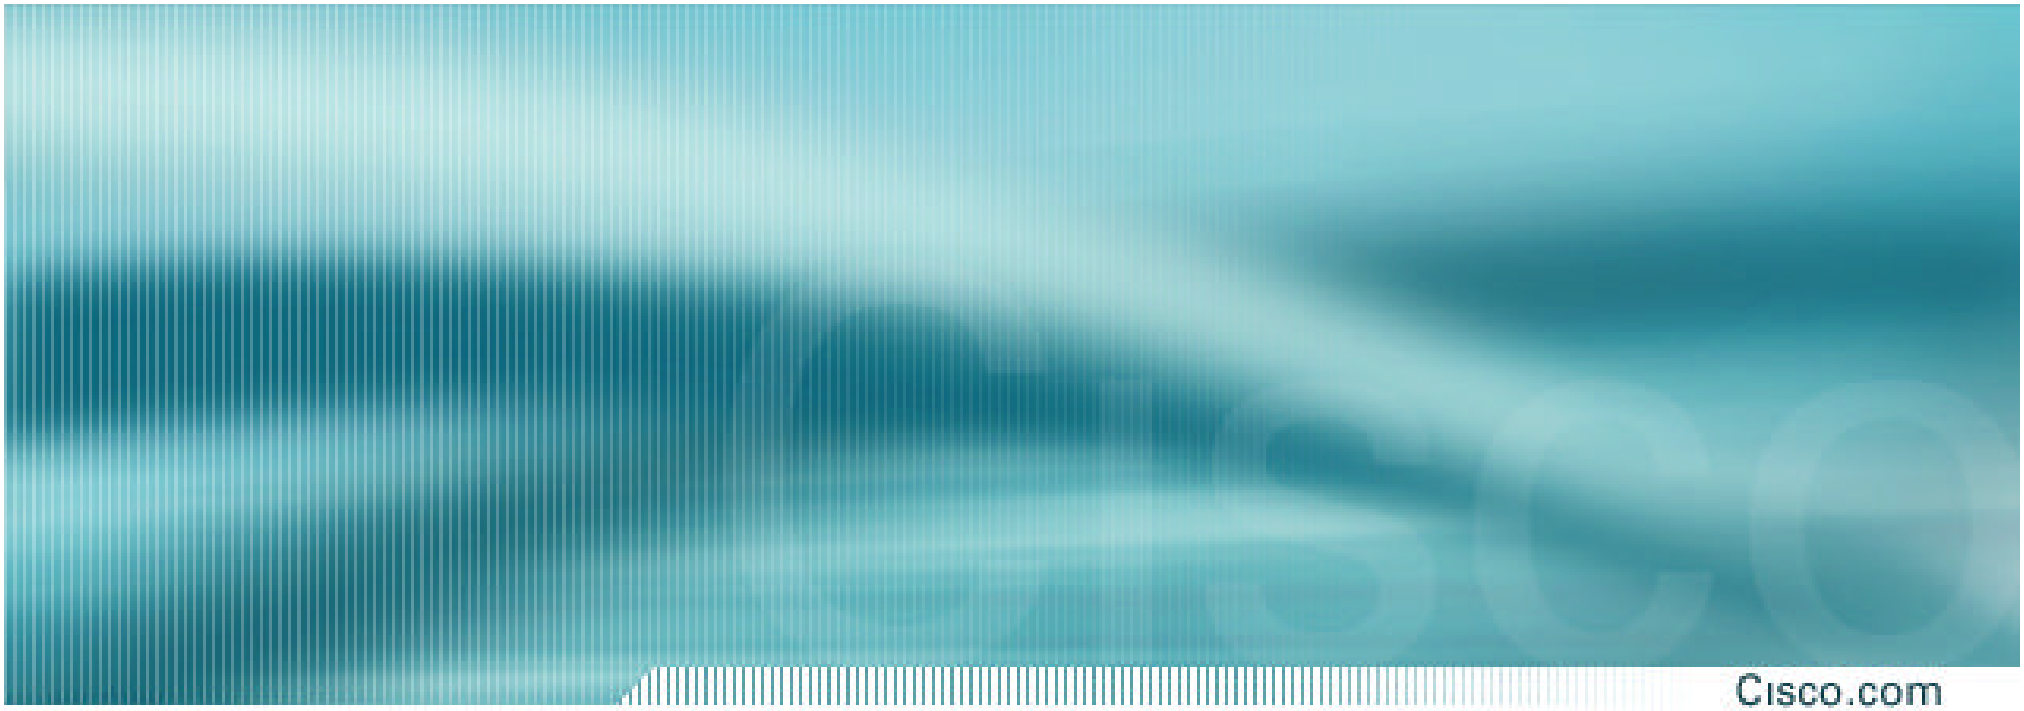

ISP Security Essentials — Best Practice Cisco IOS® and Other Techniques to Help an ISP Survive in Today's Internet

Version Guide

- **Version 1.5 – Pulled from ISP Essentials and Updated**
- **Version 1.6 – Post Code Red and new Backscatter Traceback work by UUNET by Chris Morrow chris@uu.net and Brian Gemberling brian@uu.net**
- **Version 1.7 – Post NANOG23 update with synergy from Michael Behringer's mbehring@cisco.com work.**

The ISP's World Today

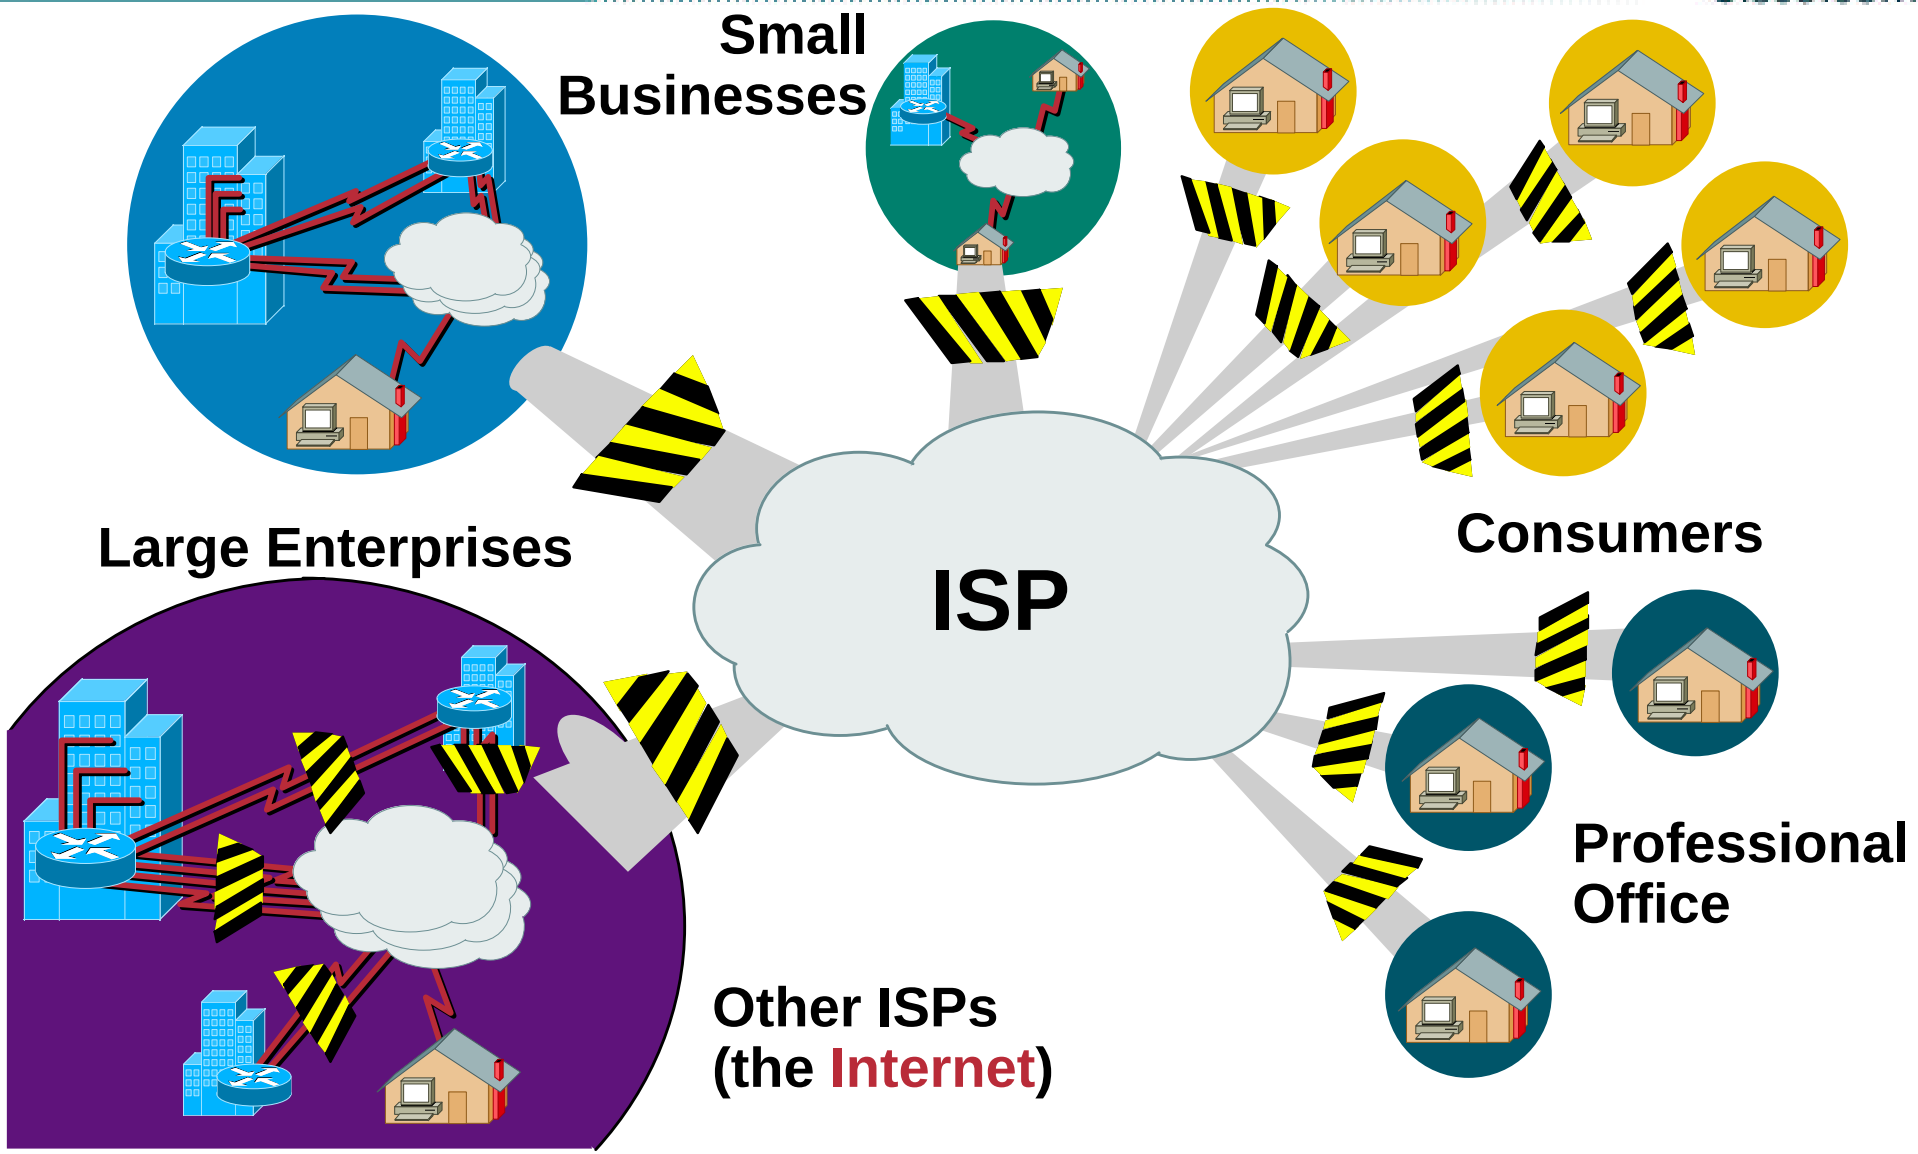

The ISP's World Today

- **Changing threat**
  - **User friendly tools make it easier for the amateur cyberpunks to do more damage**
  - **eCommerce provides a monetary motivation**
  - **Direct attacks on the Internet's core infrastructure means that the NET is not sacred anymore**
  - **Common for ISPs to have several calls per day from their customers to help defend against attacks**

Attack Methods—WinNuke

Attack Methods—Crack Shareware

ISP's Are Today's New Battle Grounds

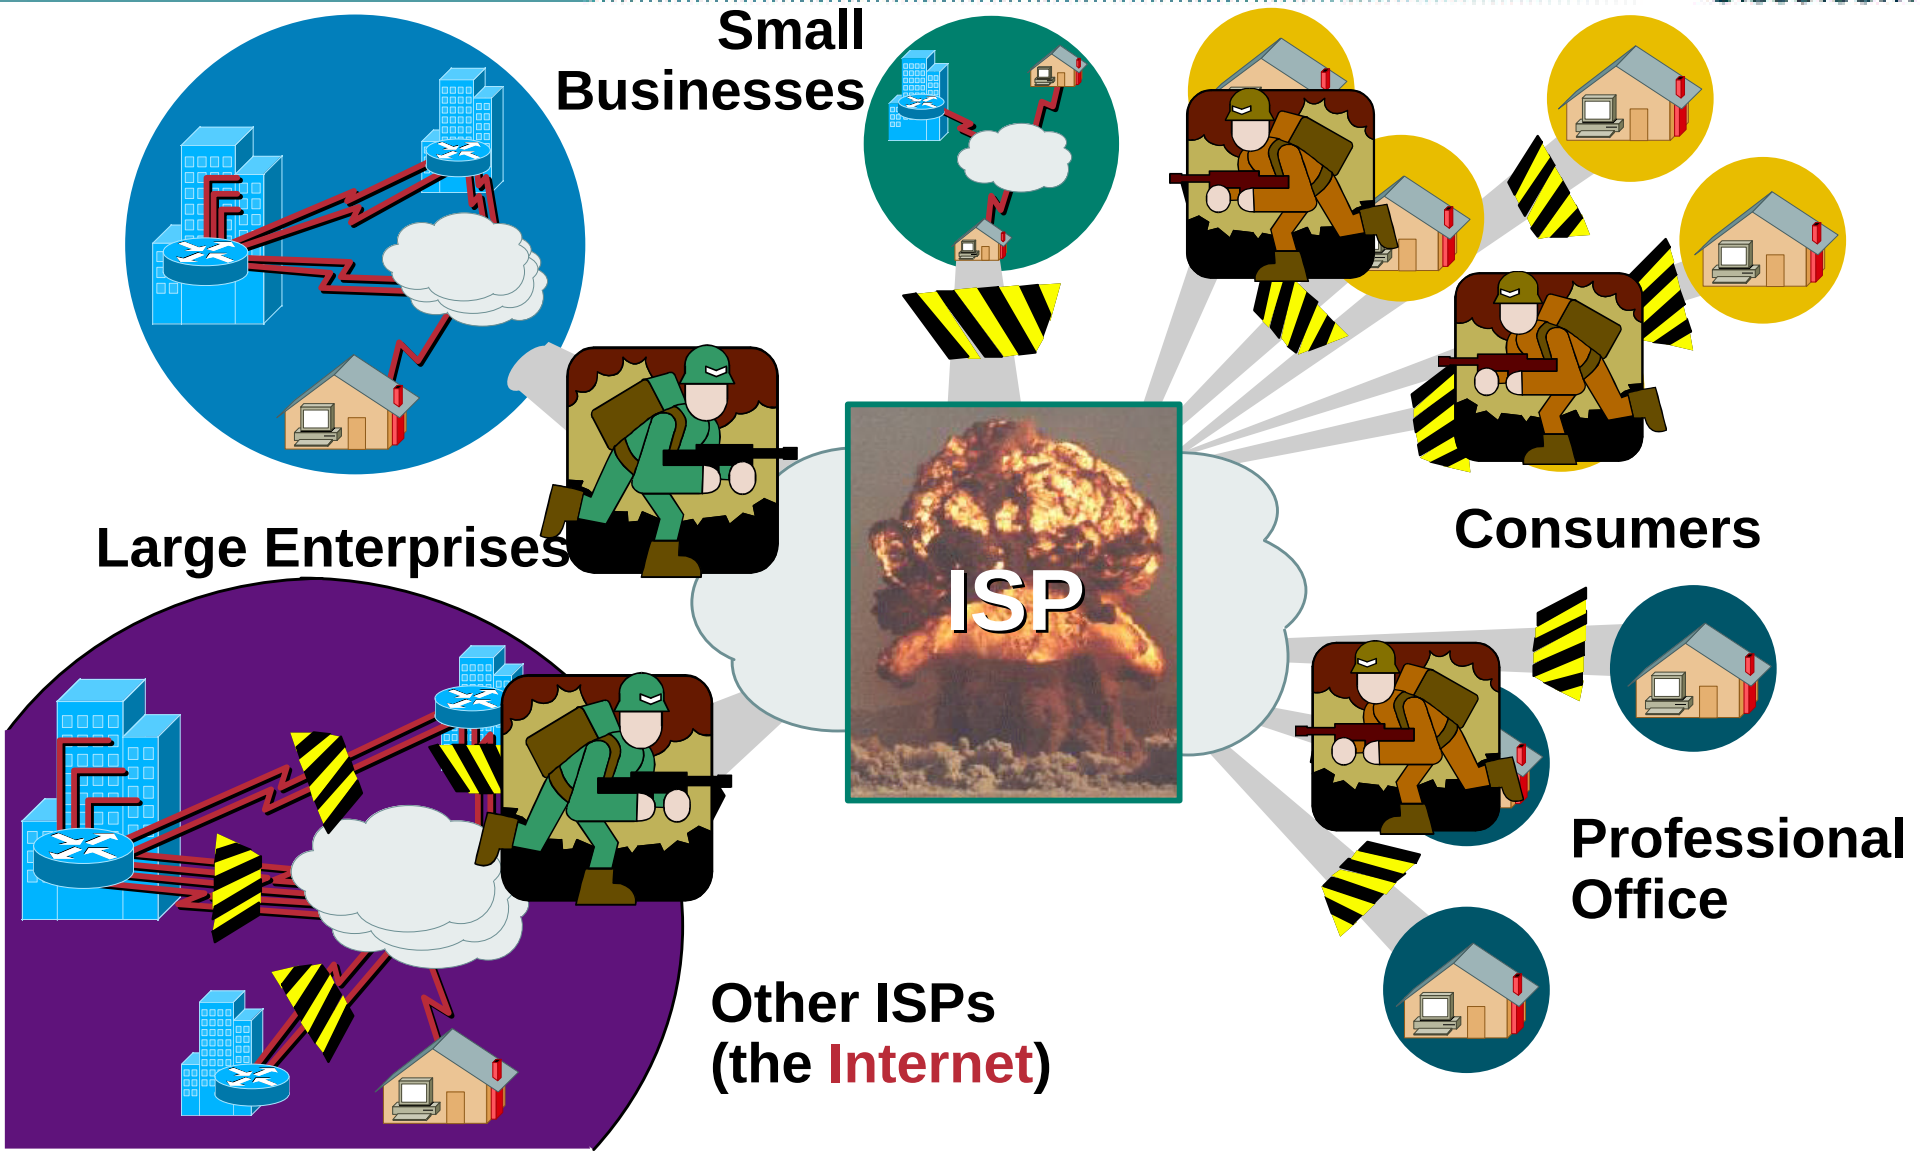

ISP Security

- **ISPs need to:**
  - **Protect themselves**
  - **Help protect their customers from the Internet**
  - **Protect the Internet from their customers**
  - **At any given time there are between 20 to 40 DOS/DDOS attacks on the Net**

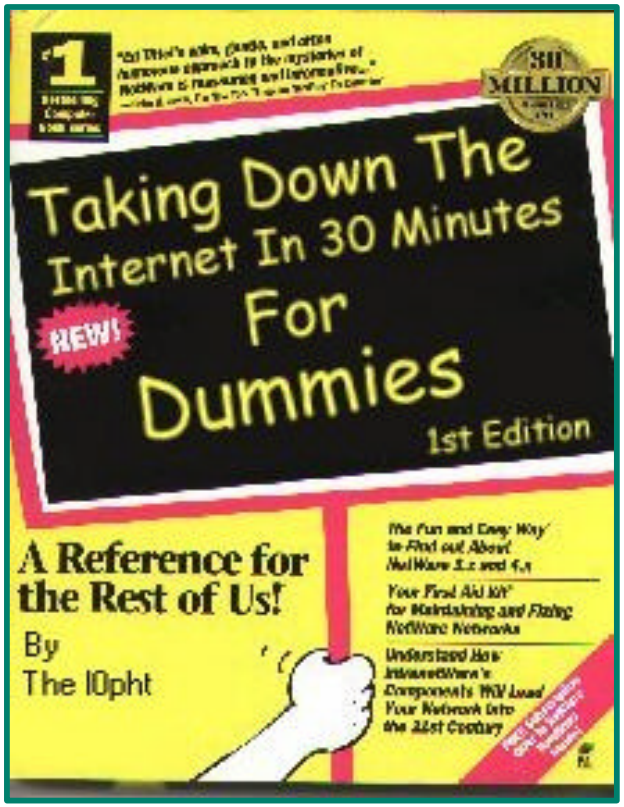

What Do ISPs Need to Do?

Security Is Not Optional!

2) Secure
Firewall, Encryption, Authentication
(PIX, Cisco IOS, FW, IPSEC, TACACS+Radius)

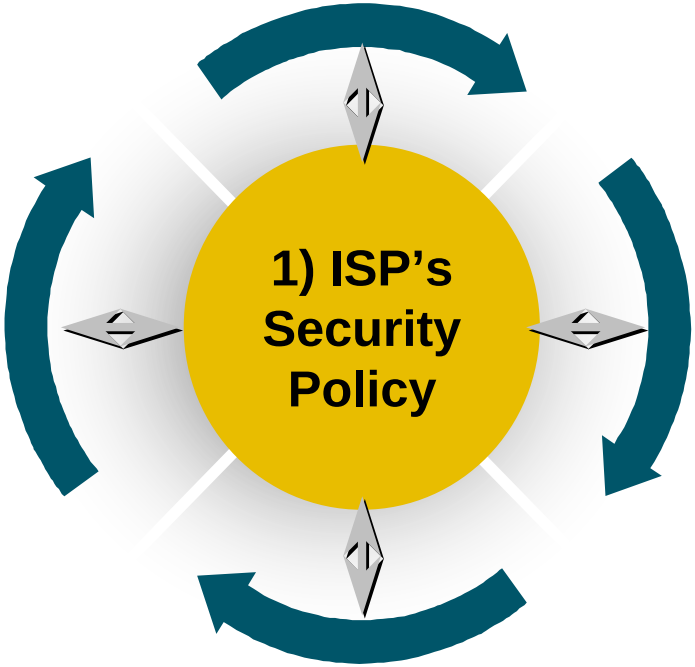

5) Manage and Improve
Network Operations and Security Professionals

3) Monitor and Respond
Intrusion Detection
(i.e. NetRanger)

4) Test
Vulnerability Scanning
(i.e. NetSonar, SPA)

What Do ISPs Need to Do?

- **Implement Best Common Practices (BCPs)**
  - **ISP infrastructure security**
  - **ISP network security**
  - **ISP services security**
- **Work with operations groups, standards organizations, and vendors on new solutions**

Hardware Vendor's Responsibilities

- **The roll of the hardware vendor is to support the network's objectives. Hence, there is a very synergistic relationship between the ISP and the hardware vendor to insure the network is resistant to security compromises**

Hardware Vendor's Responsibilities

- **Cisco System's example:**
  - **Operations people working directly with the ISPs**
  - **Emergency reaction teams (i.e. PSIRT)**
  - **Developers working with customers and IETF on new features**
  - **Security consultants working with customers on attacks, audits, and prosecution**
  - **Individuals tracking the hacker/phracker communities**
  - **Consultants working with governments/law enforcement officials**

ISP Security

- **Where to start…**

  Cisco Internet Security Advisories

  http://www.cisco.com/warp/public/779/largeent/security/advisory.html

  Cisco IOS documentation for 12.0

  http://www.cisco.com/univercd/data/doc/software/11_2/2cbook.html

  RFC2196 (site security handbook)

  Networker's security sessions

Take Note

- There are no magic knobs, grand security solutions, or super vendor features that will solve the ISP Security problem.
- Likewise, there is no rocket science involved. Just hard work that is within all ISP's grasp.
- What follows are tools and techniques that might or might not work for you.

The rest of this module ……

- **ISP Security in a Five Phase Approach:**
  - **Preparation**
  - **Identification**
  - **Classification**
  - **Traceback**
  - **Reaction**
  - **Post Mortem**
- **Examples of How ISPs Work Attacks**

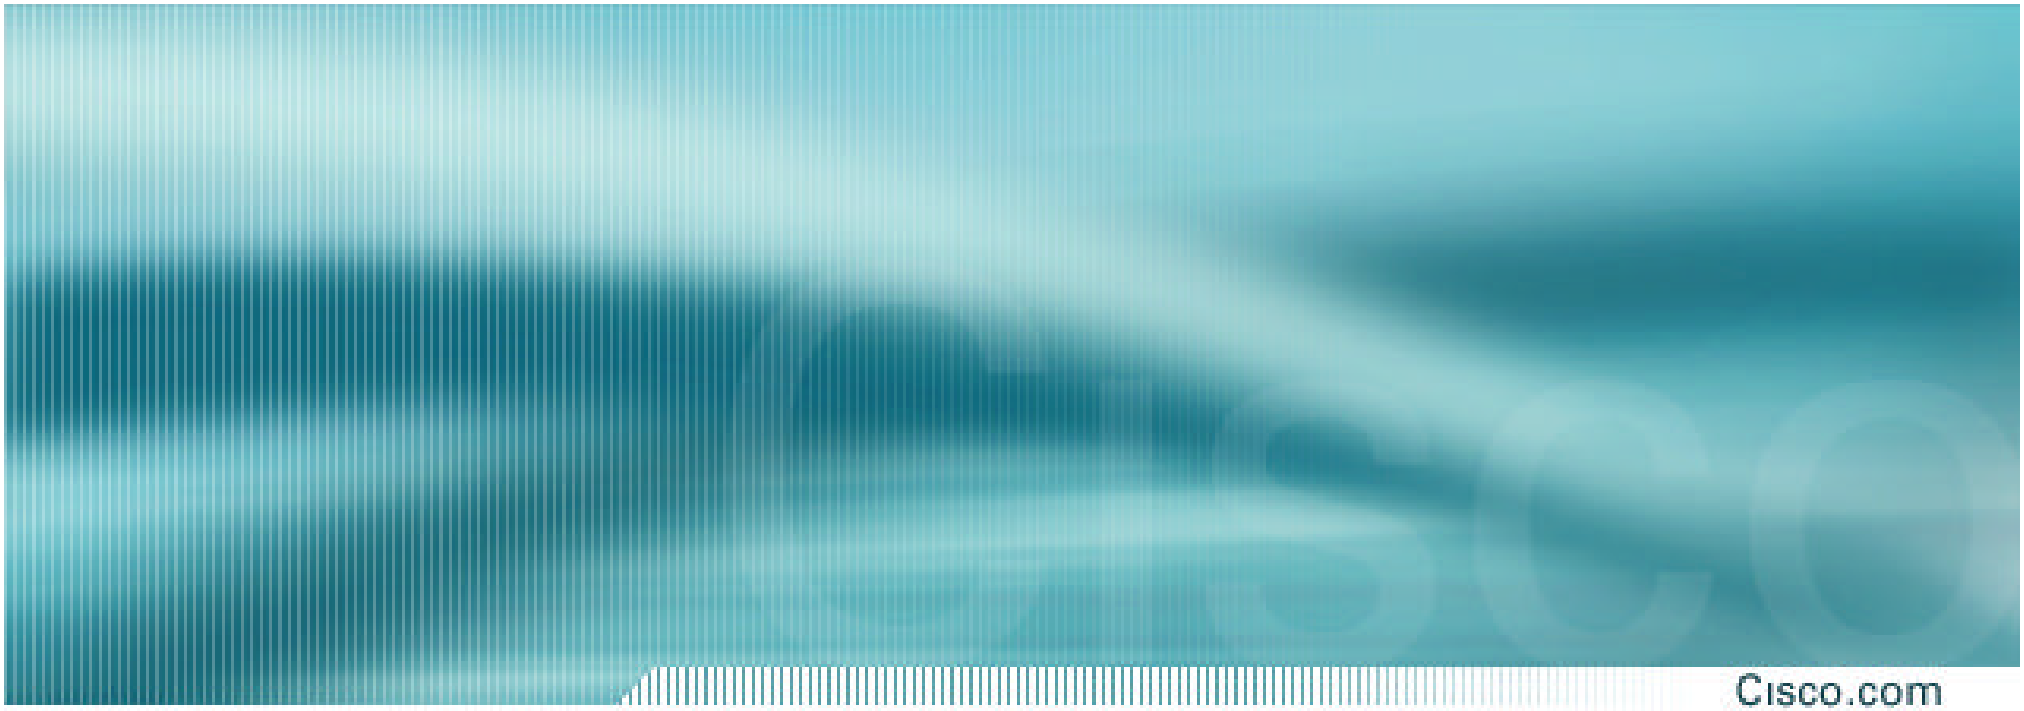

Six Phases of How and ISP Responds to a Security Incident

ISP Security Response

- **ISP's Operations Team response to a security incident can typically be broken down into six phases:**
  - **Preparation**
  - **Identification**
  - **Classification**
  - **Traceback**
  - **Reaction**
  - **Post Mortem**

ISP Security Response

Cisco.com

- **Preparation: All the work the ISP does to prepare he network, create the tools, test the tools, develop the procedures, train the team, and practice.**

  Perhaps the most important phase of how a ISP responds to a security incident.

- **Identification – How do you know you or your customer is under attack?**

ISP Security Response

Cisco.com

- **Classification – Understanding the type of attack and what damage is it causing.**
- **Traceback – From where is the attack originating?**
- **Reaction – Doing something to counter the attack – even if you choose to do nothing.**
- **Post Mortem – Analyzing what just happened. What can be done to build resistance to the attack happening again.**

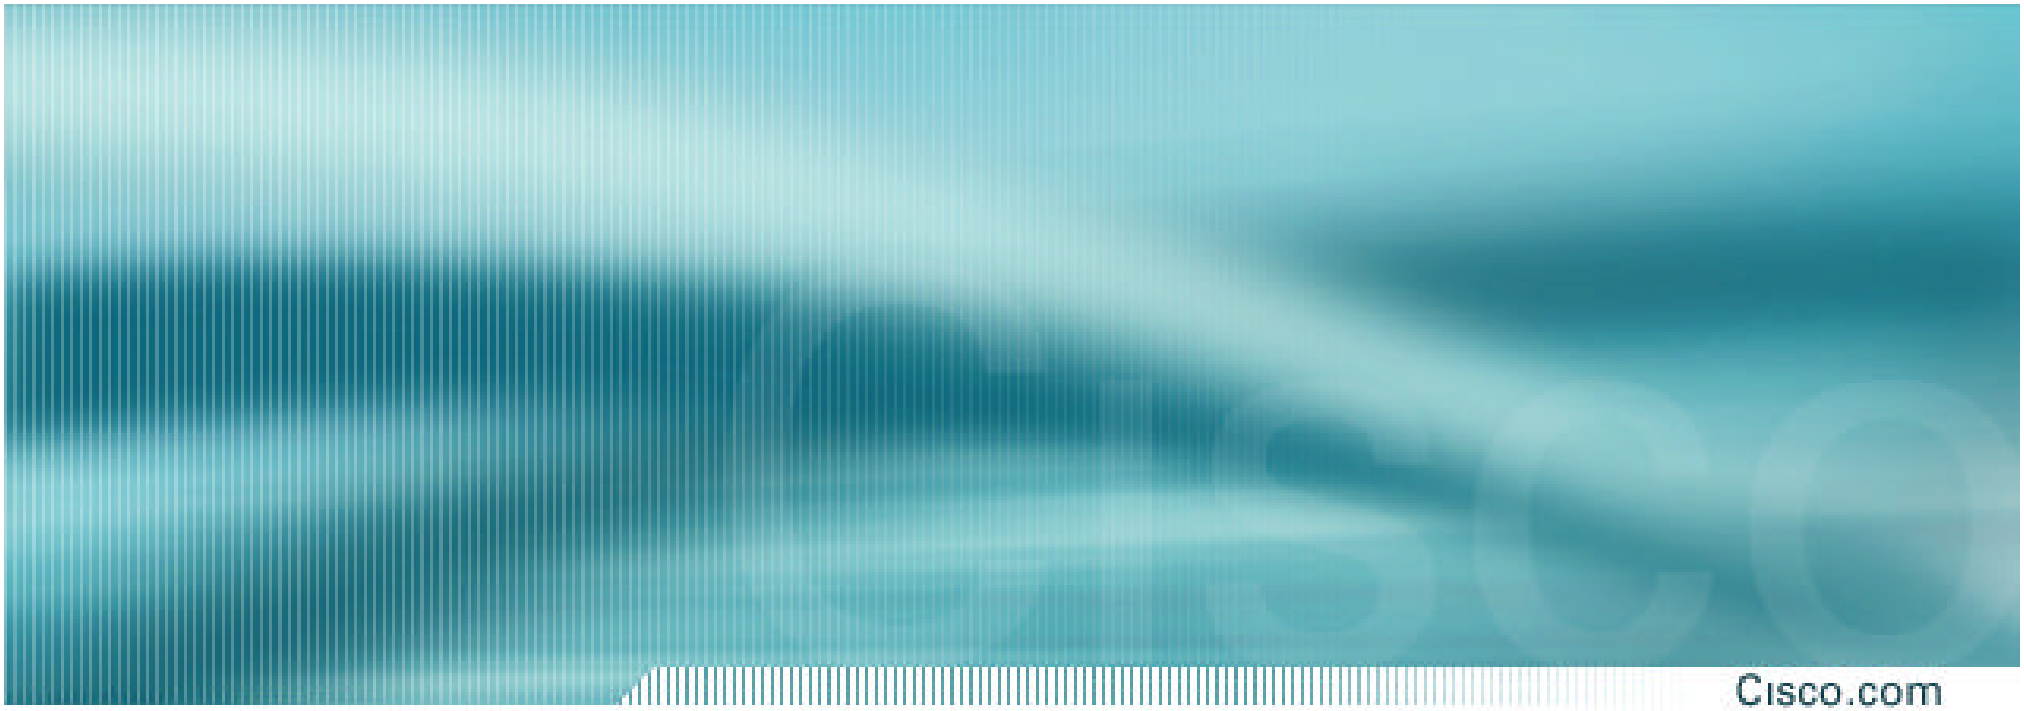

Phase 1 – Preparation for the Attack

Phase 1 - Preparation

- **Preparation is critical!**
  - **You know your *customers* are going to be attacked**
  - **It is not a matter of if but how often and how hard**
  - **The Internet is not a nice place anymore!**
  - **Think battle plans**
- **Militaries know the value of planning, practice, drilling and simulation**
  - **Those that are prepared will be victorious.**

What Do ISPs Need to Do?

Security incidence are a normal part of an ISP's operations!

2) Secure Resources
Firewall, Encryption, Authentication, Audit

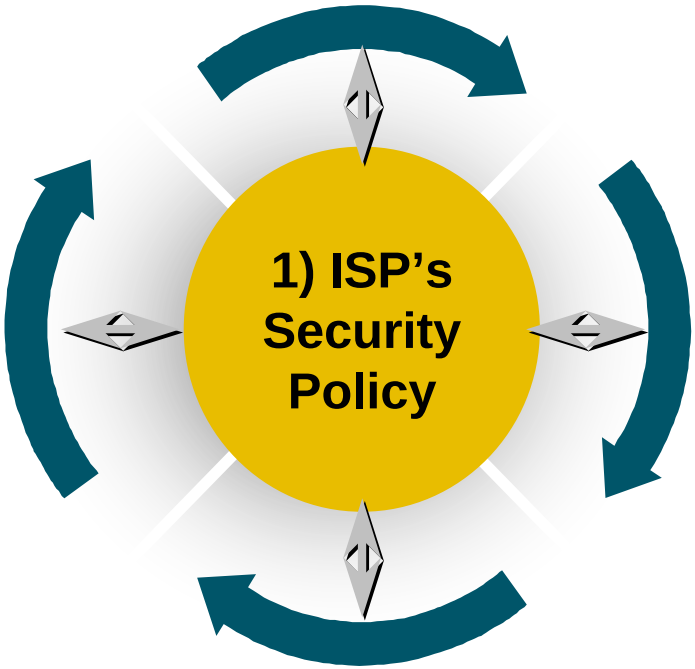

5) Manage and Improve
Post Mortem, Analyze the Incident, modify the plan/procedures

3) Monitor and Respond
Intrusion Detection, work the incidence,

4) Test, Practice, Drill
Vulnerability Scanning

Phase 1 - Preparation

- **The problem - Most ISP NOCs:**
  - **Do not have security plans**
  - **Do not have security procedures**
  - **Do not train in the tools or procedures**
  - **OJT (on the job training)—learn as it happens**

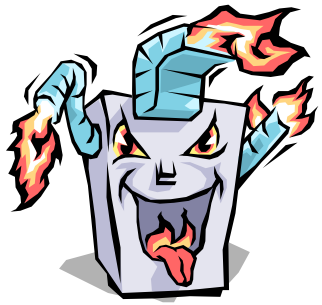

Preparation

Cisco.com

- **It is imperative that an ISP's operations team prepare.**
  - **Contacts for all ISPs who you inter-connect (peers, customers, and upstreams)**
  - **Document your policies. Will you help your customers? Will you classify the attacks? Will you traceback the attacks? Will you drop the attacks on your infrastructure?**

Preparation

- **Prepare you Tools!**

  Do you have your ACLs created?

  Do you have your scripts created?

  Have you built and tested your *Sink Hole* and *Backscatter* tools?

Preparation

- **Test you Tools before you really need to use them!**

  Have you tried putting a classification ACL on various parts of your network?

  Have you tested your scripts to insure they will work?

  Have you simulated attacks?

Phase 1 - Preparation

Cisco.com

- **Red Team/Blue Team exercises**
  - **Divide up into two teams — one defends, one attacks**
  - **Referee assigns the attackers with an objective (get this file, deface the web site, take down the target, etc.)**
  - **Defenders use network/system designs and tools/procedures to defend the target**
  - **One of the most effective ways to get your staff into the depths of TCP/IP, OS, applications, and security**

Preparation

- **Audit your network configs.**
  - **Secure the Router/Switch**
  - **Secure the Routing Protocol**
  - **Secure the Network**

Preparation

- **Know your Equipment and Infrastructure:**

  Know the Performance Envelop of all your equipment (routers, switches, workstation, etc). You need to know what your equipment is really capable of doing. If you cannot do it your self, make is a purchasing requirement.

  Know the capabilities of your network. If possible, test it. Surprises are not kind during a security incident.

What Do ISPs Need to Do?

Security incidence are a normal part of an ISP's operations!

2) Secure Resources
Firewall, Encryption, Authentication, Audit

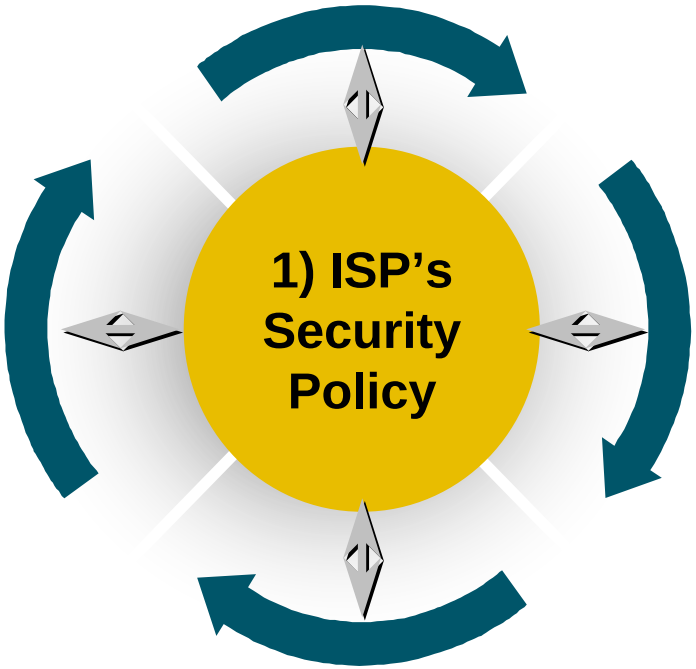

5) Manage and Improve
Post Mortem, Analyze the Incident, modify the plan/procedures

3) Monitor and Respond
Intrusion Detection, work the incidence,

4) Test, Practice, Drill
Vulnerability Scanning

ISP Security

Cisco.com

- **Proactive Step to be Prepared**
- **Securing the Router**
- **Securing the Routing Protocols**
- **Securing the Network**

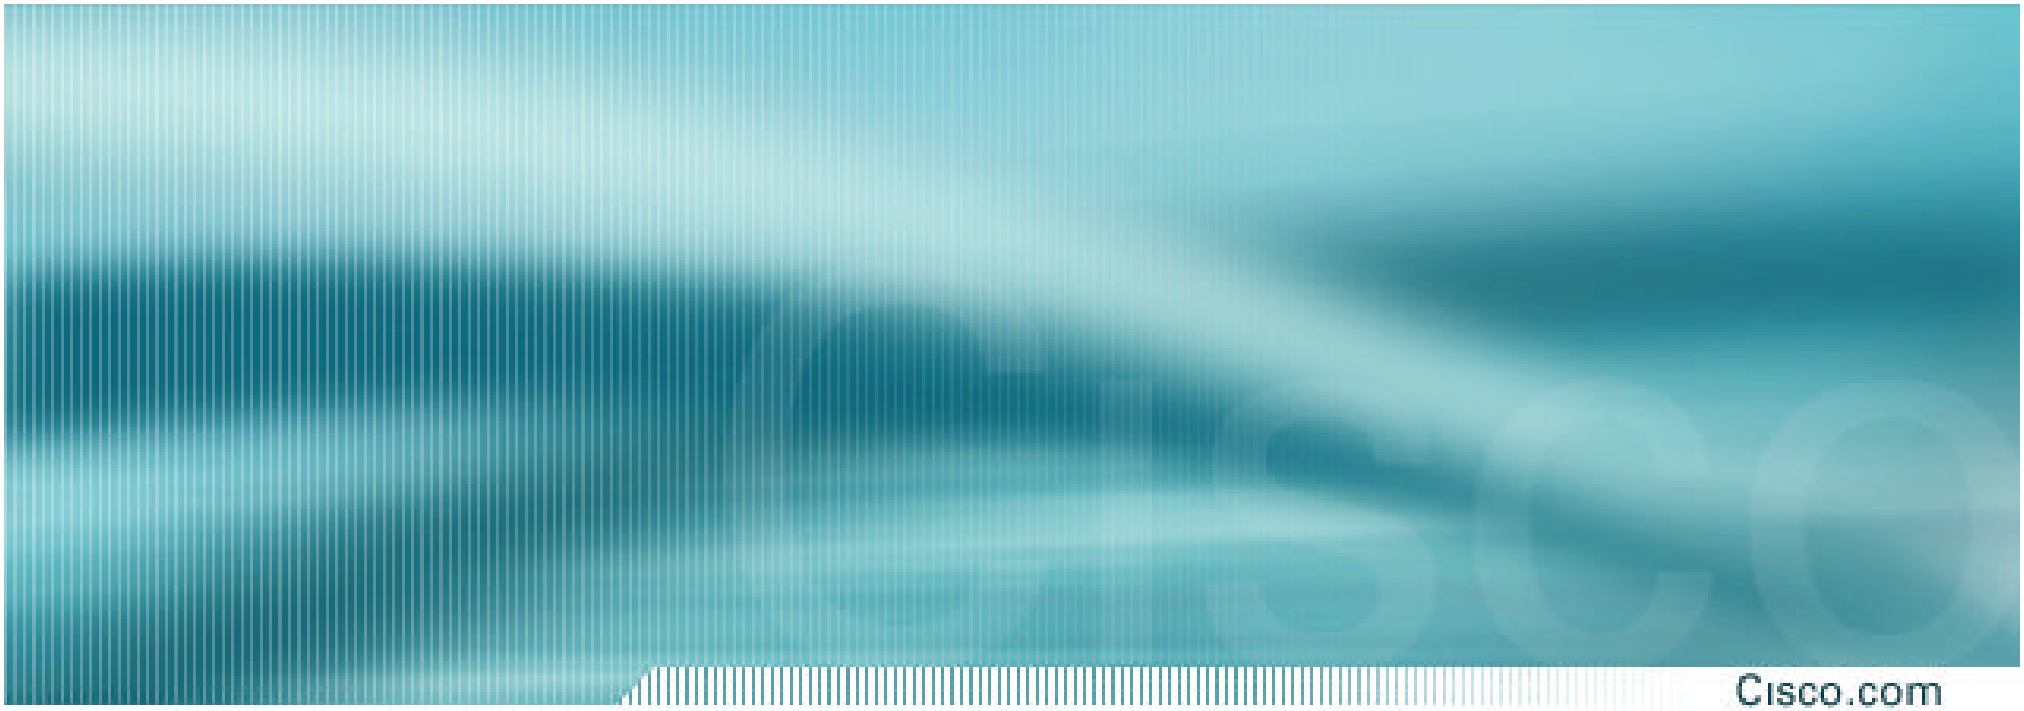

Phase 1 – Preparation for the Attack

Securing the Router

Global Services You Turn OFF

- **Some services turned on by default, should be turned off to save memory and prevent security breaches/attacks**

```
no service finger
no service pad
no service udp-small-servers
no service tcp-small-servers
no ip bootp server
```

Interface Services You Turn OFF

- **Some IP features are great for campus LANs, but do not make sense on a ISP backbone**
- **All interfaces on an ISP's backbone router should have the follow as a default:**

```
no ip redirects
no ip directed-broadcast
no ip proxy-arp
```

Cisco Discovery Protocol

- Lets network administrators discover neighbouring Cisco equipment, model numbers and software versions
- Should not be needed on ISP network

```
no cdp run
```

- Should not be activated on any public facing interface: IXP, customer, upstream ISP – unless part of the peering agreement.
- Disable per interface

```
no cdp enable
```

Cisco Discovery Protocol

Cisco.com

```
Defiant#show cdp neighbors detail
-------------------------
Device ID: Excalabur
Entry address(es):
  IP address: 4.1.2.1
Platform: cisco RSP2,  Capabilities: Router
Interface: FastEthernet1/1,  Port ID (outgoing port): FastEthernet4/1/0
Holdtime : 154 sec

Version :
Cisco Internetwork Operating System Software
IOS (tm) RSP Software (RSP-K3PV-M), Version 12.0(9.5)S, EARLY DEPLOYMENT
   MAINTEN
ANCE INTERIM SOFTWARE
Copyright (c) 1986-2000 by cisco Systems, Inc.
Compiled Fri 03-Mar-00 19:28 by htseng

Defiant#
```

Login Banner

- **Use a good login banner, or nothing at all:**

```
banner login ^
  Authorised access only
  This system is the property of Galactic Internet
  Disconnect IMMEDIATELY if you are not an authorised user!
  Contact noc@net.galaxy +99 876 543210 for help.
^
```

Exec Banner

- **Useful to remind logged in users of local conditions:**

```
banner exec ^
  PLEASE NOTE - THIS ROUTER SHOULD NOT HAVE A DEFAULT ROUTE!
  It is used to connect paying peers. These ‘customers’ should
not be able to default to us.
  The config for this router is NON-STANDARD
  Contact Network Engineering +99 876 543234 for more info.
^
```

Use Enable Secret

- Encryption '7' on a Cisco is reversible
- The “enable secret” password encrypted via a one-way algorithm

```
enable secret <removed>

no enable password

service password-encryption
```

VTY and Console Port Timeouts

- **Default idle timeout on async ports is 10 minutes 0 seconds**

```
exec-timeout 10 0
```

- **Timeout of 0 means permanent connection**
- **TCP keepalives on incoming network connections**

```
service tcp-keepalives-in
```

- **Kills unused connections**

VTY Security

- **Access to VTYs should be controlled, not left open; consoles should be used for last resort admin only:**

```
access-list 3 permit 215.17.1.0 0.0.0.255
access-list 3 deny   any
line vty 0 4
 access-class 3 in
 exec-timeout 5 0
 transport input telnet ssh
 transport output none
 transport preferred none
 password 7 045802150C2E
```

VTY Security

- **Use more robust ACLs with the logging feature to spot the probes on you network**

```
access-list 199 permit tcp 1.2.3.0 0.0.0.255 any
access-list 199 permit tcp 1.2.4.0 0.0.0.255 any
access-list 199 deny   tcp any any range 0 65535
log
access-list 199 deny   ip any any log
```

VTY Access and SSHv1

- Secure shell supported as from IOS 12.0S
- Obtain, load and run appropriate crypto images on router
- Set up SSH on router

```
Beta7200(config)#crypto key generate rsa
```

- Add it as input transport

```
line vty 0 4
  transport input telnet ssh
```

VTY Access and SSHv1

- **SSHv1 client in IOS for router to router SSH (not in docs)**

  ssh [-l <userid>] [-c <des|3des>] [-o numberofpasswdprompts <n>] [-p <portnum>] <ipaddr|hostname> [<IOS command>]

  where

  -l <userid> is the user to login as on the remote machine. Default is the current user id.

  -c <des|3des> specifies the cipher to use for encrypting the session. Triple des is encrypt-decrypt-encrypt with three different keys. The default is 3des if this algorithm is included in the image, else the default is des.

  -o specifies the options which is currently one only numberofpasswdprompts <n> specifies the number of password prompts before ending the attempted session. The server also limits the number of attempts to 5 so it is useless to set this value larger than 5. Therefore the range is set at 1-5 and the default is 3 which is also the IOS server default.

  -p <portnum> Port to connect to on the remote host. Default is 22.

  <ipaddr|hostname> is the remote machine ip address or hostname

  <IOS command> is an IOS exec command enclosed in quotes (ie "). This will be executed on connection and then the connection will be terminated when the command has completed.

VTY Access and SSHv1

- **Example:**

  Insure you have the proper image (post 12.0(10)S with “k3pv”

  ```
  i.e. rsp-k3pv-mz.120-11.S3.bin
  ```

  Set up SSH on the router

  ```
  Beta7200(config)#crypto key generate rsa
  ```

  Use the SSH client:

  ```
  ssh -l myuser myhost "sh users"

  ssh -l myuser -c 3des -o 5 -p 22 myhost
  ```

User Authentication

- **Account per user, with passwords**

```
aaa new-model
aaa authentication login neteng local
username joe password 7 1104181051B1
username jim password 7 0317B21895FE
line vty 0 4
 login neteng
 access-class 3 in
```

- **Username/password is more resistant to attack than a plain password**

User Authentication

- **Use distributed authentication system**
  - **RADIUS—Recommended for user accounting**
  - **TACACS+—Recommended for securing the network**

```
aaa new-model
aaa authentication login default tacacs+ enable
aaa authentication enable default tacacs+ enable
aaa accounting exec start-stop tacacs+
ip tacacs source-interface Loopback0
tacacs-server host 215.17.1.1
tacacs-server key CKr3t#
line vty 0 4
 access-class 3 in
```

User Authentication

Cisco.com

TACACS+ Provides a Detailed Audit Trail of what Is Happening on the Network Devices

User-Name	Group-cmd		priv-lvl	service	NAS-Portname	task_id	NAS-IP-reason
bgreene	NOC	enable <cr>	0	shell	tty0	4	210.210.51.224
bgreene	NOC	exit <cr>	0	shell	tty0	5	210.210.51.224
bgreene	NOC	no aaa accounting exec Workshop <cr>	0	shell	tty0	6	210.210.51.224
bgreene	NOC	exit <cr>	0	shell	tty0	8	210.210.51.224
pfs	NOC	enable <cr>	0	shell	tty0	11	210.210.51.224
pfs	NOC	exit <cr>	0	shell	tty0	12	210.210.51.224
bgreene	NOC	enable <cr>	0	shell	tty0	14	210.210.51.224
bgreene	NOC	show accounting <cr>	15	shell	tty0	16	210.210.51.224
bgreene	NOC	write terminal <cr>	15	shell	tty0	17	210.210.51.224
bgreene	NOC	configure <cr>	15	shell	tty0	18	210.210.51.224
bgreene	NOC	exit <cr>	0	shell	tty0	20	210.210.51.224
bgreene	NOC	write terminal <cr>	15	shell	tty0	21	210.210.51.224
bgreene	NOC	configure <cr>	15	shell	tty0	22	210.210.51.224
bgreene	NOC	aaa new-model <cr>	15	shell	tty0	23	210.210.51.224
bgreene	NOC	aaa authorization commands 0 default tacacs+ none <cr>	15	shell	tty0	24	210.210.51.224
bgreene	NOC	exit <cr>	0	shell	tty0	25	210.210.51.224
bgreene	NOC	ping <cr>	15	shell	tty0	32	210.210.51.224
bgreene	NOC	show running-config <cr>	15	shell	tty66	35	210.210.51.224
bgreene	NOC	router ospf 210 <cr>	15	shell	tty66	45	210.210.51.224
bgreene	NOC	debug ip ospf events <cr>	15	shell	tty66	46	210.210.51.224

User Authentication

- **Ideally, when you have TACACS+ on a router, you do not give out the local username/ password nor enable password**
  - **Lock them in a safe in the NOC in case of total TACACS+ failure**
- **Problem—username/password is a reversible hash**
  - **Some engineer can take a config an reverse the hash**
- **Threat—disgruntled employees can attack TACACS+ then get into the routers**

User Authentication

- **Fix is in CSCds84754**

  Added simple MD5 Encryption mechanism for username password:

  ```
  username barry secret 5 ;2kj45nk5jnt43
  ```

- **Now MD5 Encrypted username/passwords can be used with TACACS+ to keep the system secure from the internal security threat.**

User Authentication

- **So now you can have the following:**

```
aaa new-model
aaa authentication login default tacacs+ local
enable
aaa authentication enable default tacacs+ local
enable
aaa accounting exec start-stop tacacs+
ip tacacs source-interface Loopback0
tacacs-server host 215.17.1.1
tacacs-server key CKr3t#
line vty 0 4
 access-class 3 in
username joe password 6 1104181051B1
username jim password 6 0317B21895FE
```

Source Routing

- **IP has a provision to allow source IP host to specify route through Internet**
- **ISPs should turn this off, unless it is specifically required:**

```
no ip source-route
```

- ***traceroute*-s to investigate network failures—valuable tool; but, it you are not using *traceroute*-s, then turn off the feature!**

ICMP Unreachable Overload

- **Originally, all ICMP Unreachable replies were *punted* from the LC/VIP to the GRP/RP.**
- **The result was that the GRP/RP's CPU resources could be overloaded, just responding to ICMP Unreachables.**
- ***Potential Security Hole* that can be used to overload a router.**
- **Prevented Black Hole Filtering on Router.**

ICMP Unreachable Overload

- **Problem resolved across the the LC/VIP based platforms:**

  CSCds36541 - Traffic received on eng1 LC for null0 punted to RP

  CSCdr46528 - GSR eng0 LC: routes for Null0 have terrible lookup performance

  CSCdt66560 - Engine 2 PSA Punts Null0 Traffic to GRP

  CSCdt68393 - 100% CPU using Null0 to blackhole traffic under DOS

- **All LCs and VIPs now handle the ICMP Unreachables and the *no ip unreachables* command works on all interfaces.**

ICMP Unreachable Overload

Cisco.com

- **All Routers who use any static route to Null0 should put *no ip unreachables (i.e. BGP Advertisements).***

```
interface Null0
 no ip unreachables
!
ip route <dest to drop> <mask> Null0
```

ICMP Unreachable Rate-Limiting

- **New ICMP Unreachable Rate-Limiting Command:**

```
ip icmp rate-limit unreachable [DF] <1-4294967295
  millisecond>

no ip icmp rate-limit unreachable [df]
```

- **Turned on by default and hidden since 12.0(8)S. Default value set to 500 milliseconds.**
- **Peer Review with several top ISP operations engineers are recommending this be set at 1 second for normal and DF.**

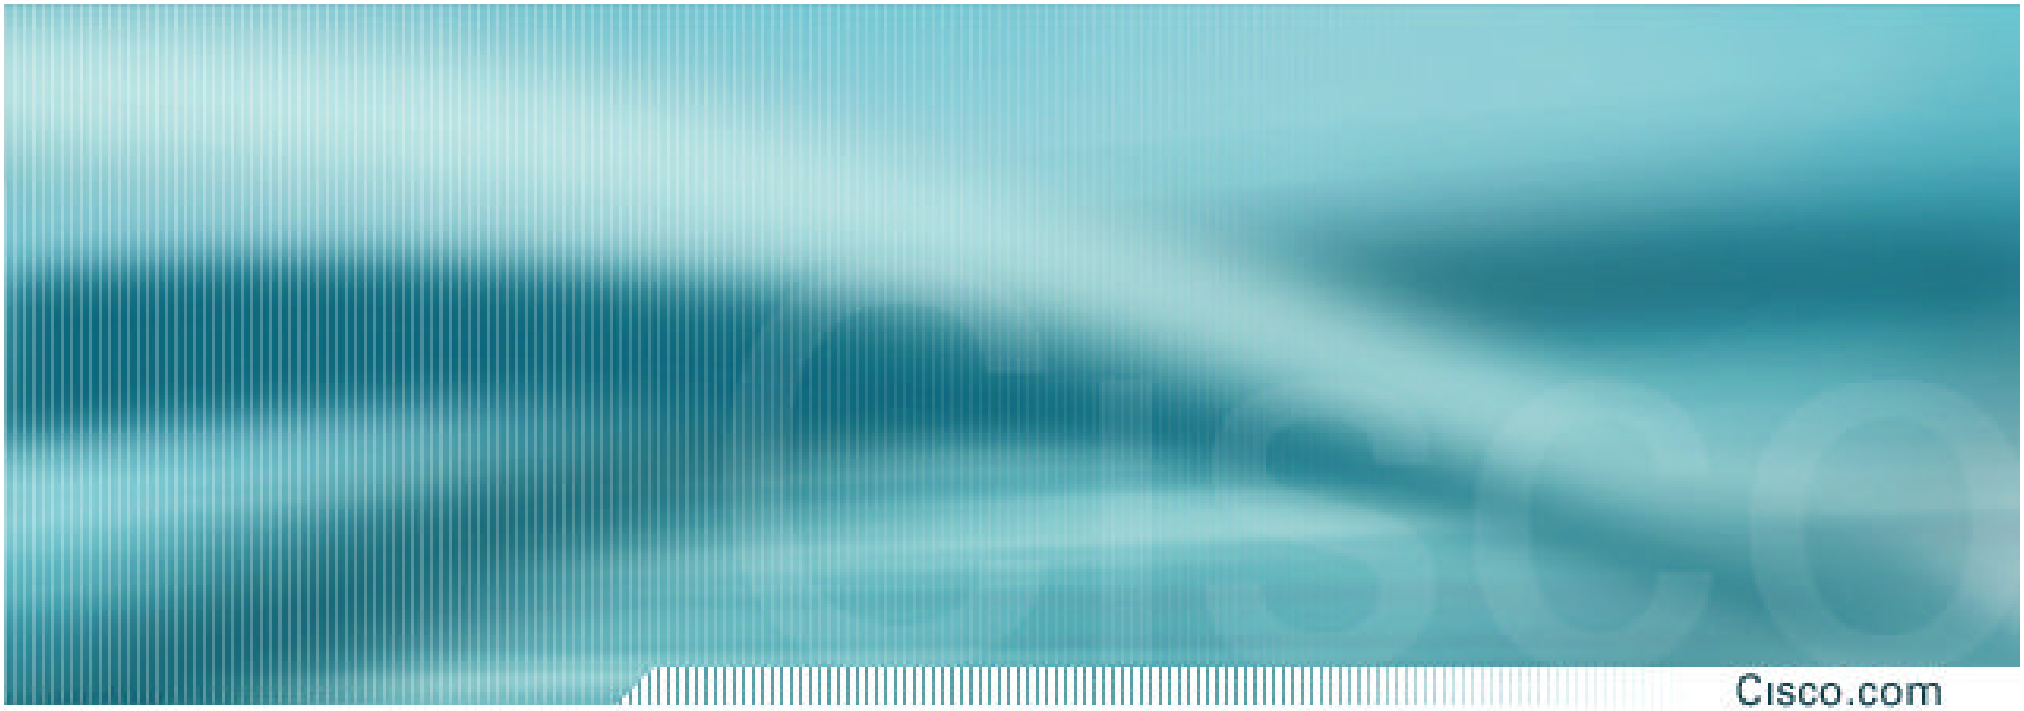

Phase 1 – Preparation for the Attack

Securing the Routing Protocol

Routing Protocol Security

- **Routing protocol can be attacked**
  - **Denial of service**
  - **Smoke screens**
  - **False information**
  - **Reroute packets**

May Be Accidental or Intentional

Secure Routing
Route Authentication

Configure Routing Authentication

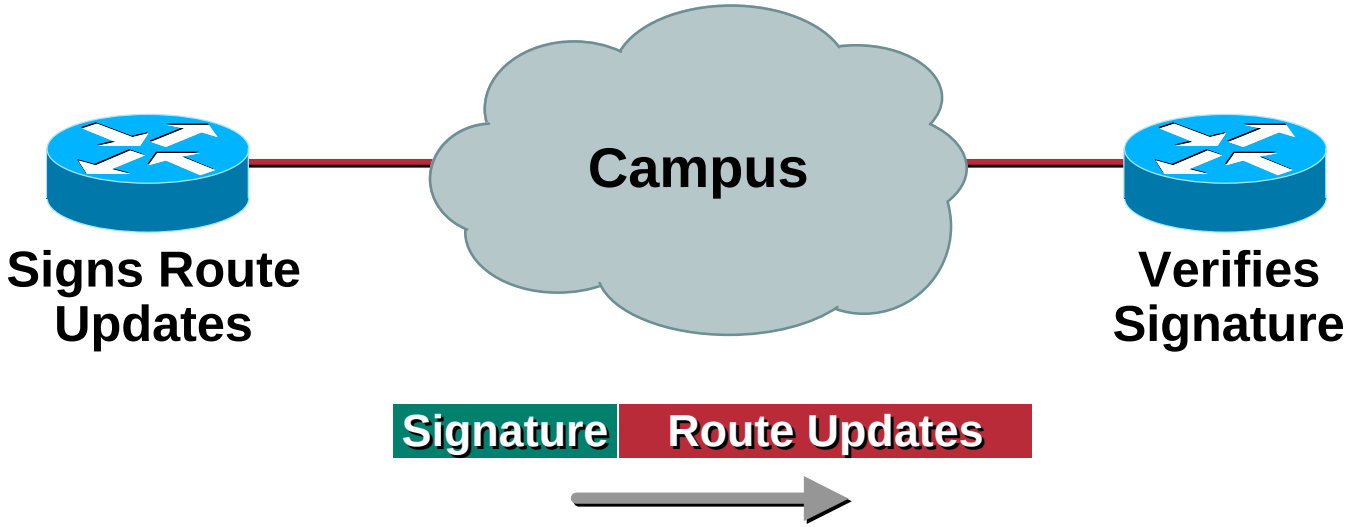

Certifies Authenticity of Neighbor and Integrity of Route Updates

Signature Generation

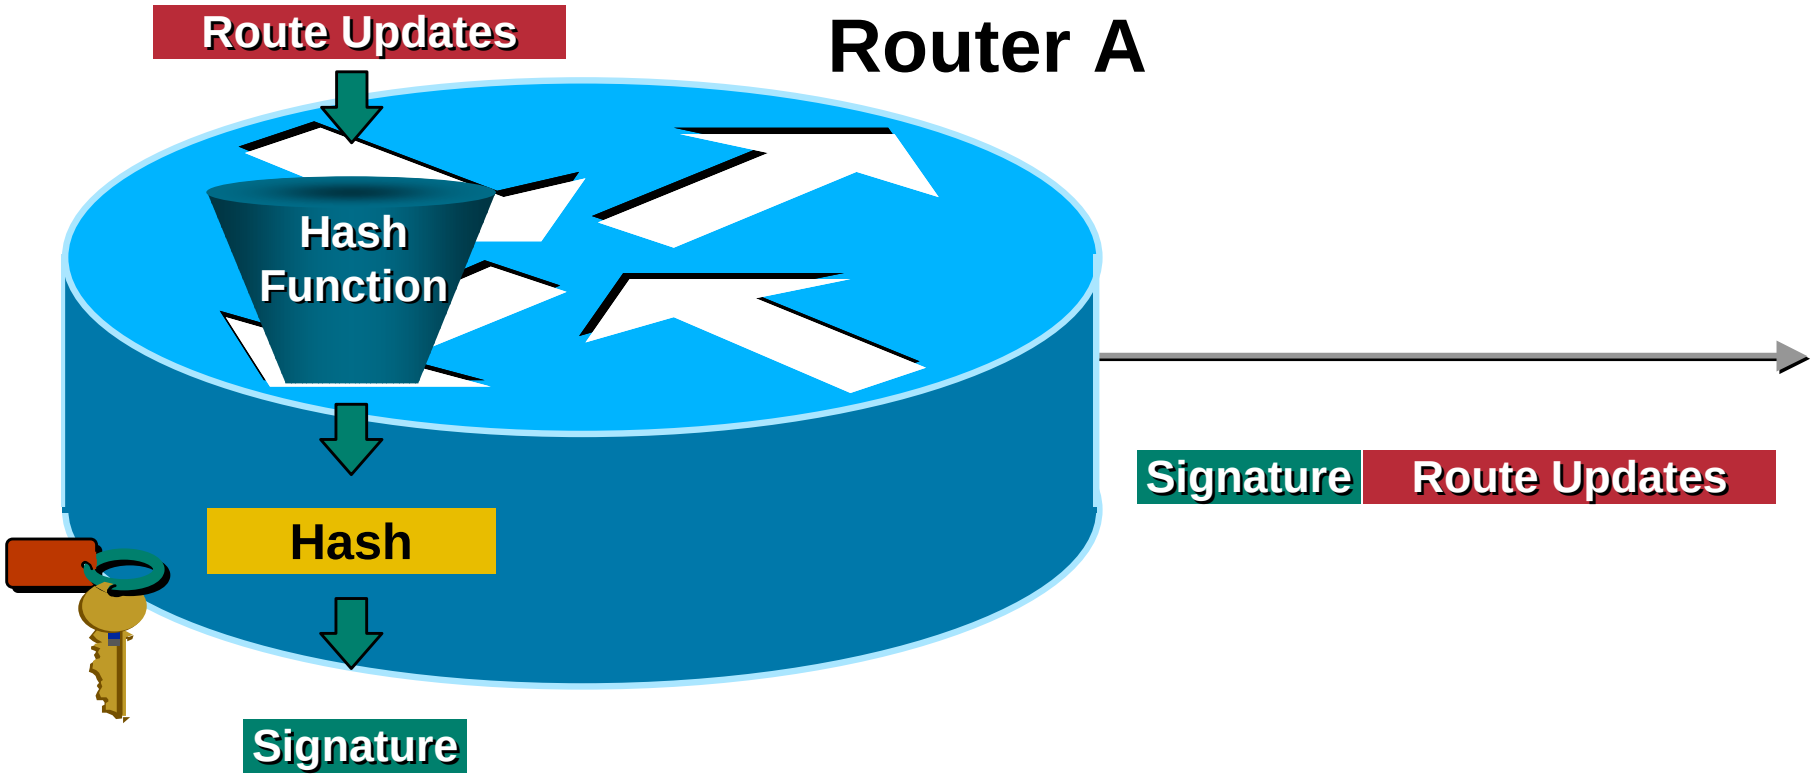

Signature = Encrypted Hash of Routing Update

Signature Verification

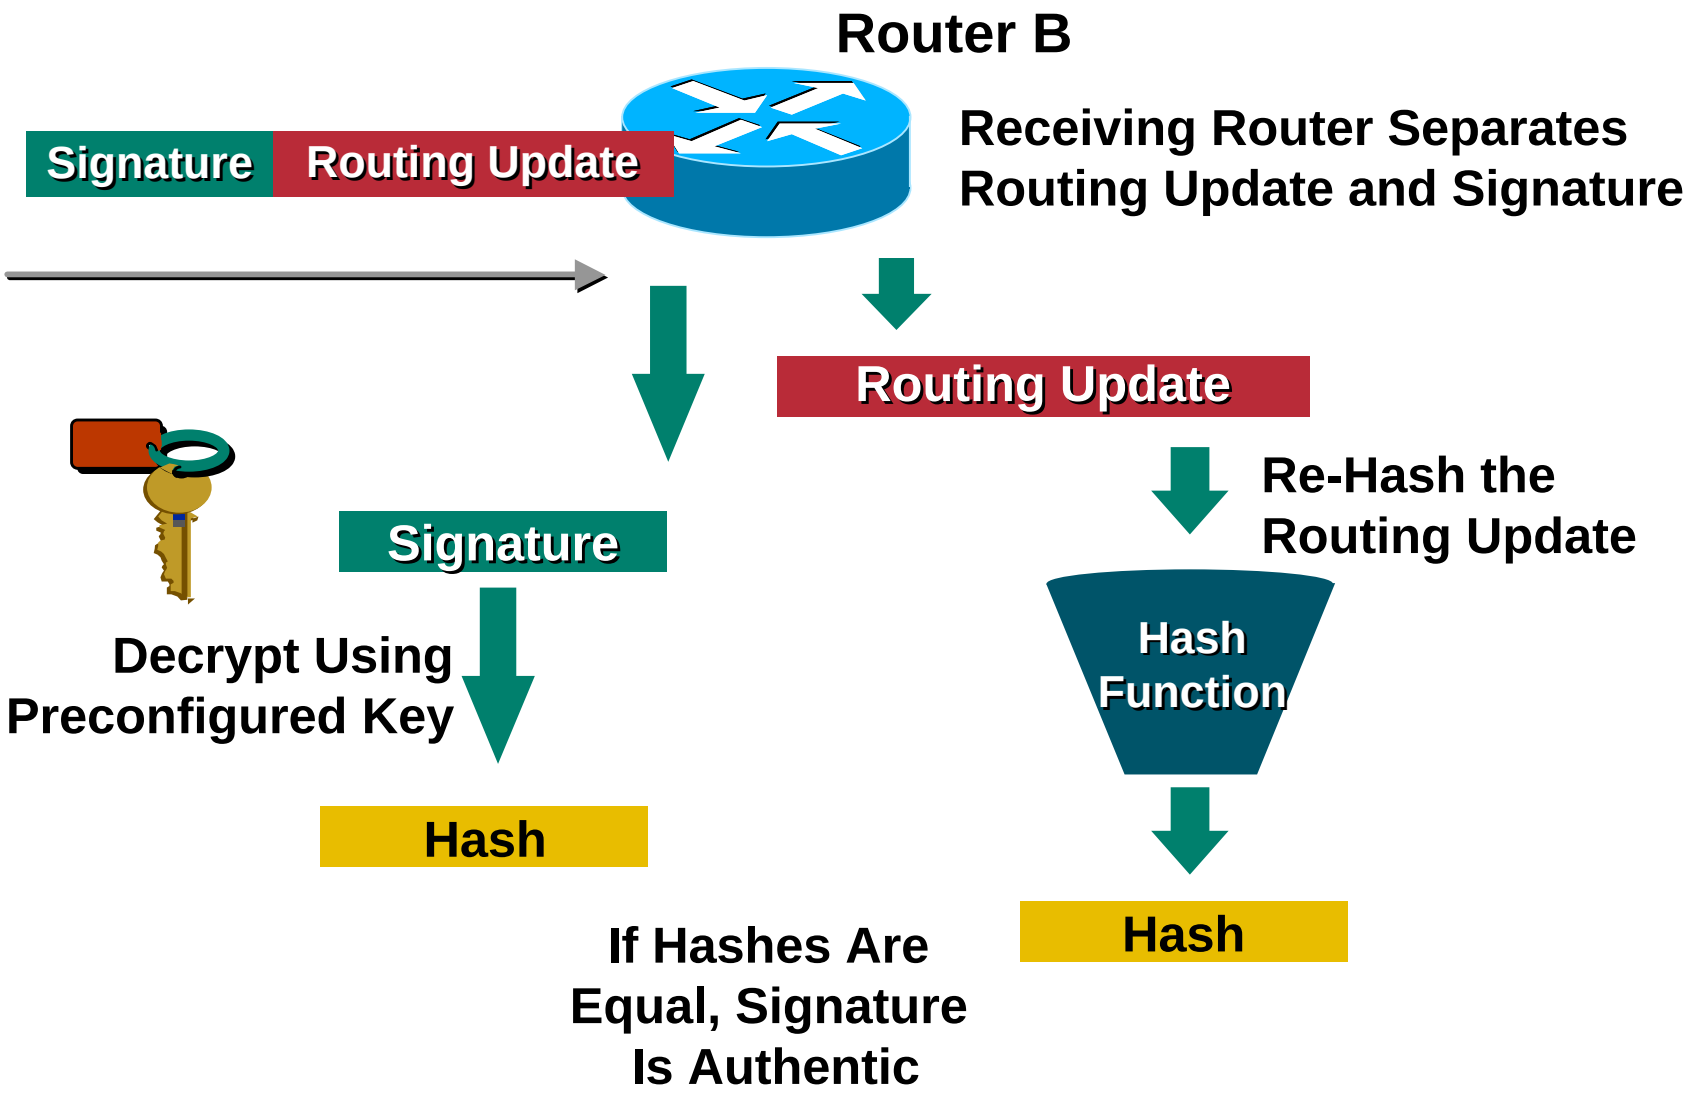

Route Authentication

- **Authenticates routing update packets**
- **Shared key included in routing updates**
  - **Plain text—Protects against accidental problems only**
  - **Message Digest 5 (MD5)—Protects against accidental and intentional problems**

Route Authentication

- **Multiple keys supported**
  - **Key lifetimes based on time of day**
  - **Only first valid key sent with each packet**
- **Supported in: BGP, IS-IS, OSPF, RIPv2, and EIGRP(11.2(4)F)**
- **Syntax differs depending on routing protocol**

OSPF Route Authentication

- **OSPF area authentication**
  - **Two types**
    - **Simple password**
    - **Message Digest (MD5)**

ip ospf authentication-key *key* (this goes under the specific interface)
area *area-id* **authentication** (this goes under "router ospf <process-id>")

ip ospf message-digest-key *keyid* **md5** *key* (used under the interface)
area *area-id* **authentication message-digest** (used under "router ospf <process-id>")

OSPF and ISIS Authentication Example

Cisco.com

- **OSPF**

```
interface ethernet1
ip address 10.1.1.1
255.255.255.0
ip ospf message-digest-
key 100 md5 cisco
!
router ospf 1
network 10.1.1.0
0.0.0.255 area 0
area 0 authentication
message-digest
```

- **ISIS**

```
interface ethernet0
ip address 10.1.1.1
255.255.255.0
ip router isis
isis password cisco
level-2
```

BGP Route Authentication

Cisco.com

```
router bgp 200
 no synchronization
 neighbor 4.1.2.1 remote-as 300
 neighbor 4.1.2.1 description Link to Excalabur
 neighbor 4.1.2.1 send-community
 neighbor 4.1.2.1 version 4
 neighbor 4.1.2.1 soft-reconfiguration inbound
 neighbor 4.1.2.1 route-map Community1 out
 neighbor 4.1.2.1 password 7 cisco
```

BGP Route Authentication

- **Works per neighbor or for an entire peer-group**
- **Two routers with password mis-match:**

  %TCP-6-BADAUTH: Invalid MD5 digest from [peer's IP address]:11004 to [local router's IP address]:179

- **One router has a password and the other does not:**

  %TCP-6-BADAUTH: No MD5 digest from [peer's IP address]:11003 to [local router's IP address]:179

Selective Packet Discard

- When a link goes to a saturated state, you will drop packets; the problem is that you will drop any type of packets—including your routing protocols
- Selective Packet Discard (SPD) will attempt to drop non-routing packets instead of routing packets when the link is overloaded

```
ip spd enable (11.1 CA & CC)
```

Selective Packet Discard

- **Enabled by default from 11.2(5)P and later releases, available option in 11.1CA/CC**
- **12.0 the syntax changes and the default is to enable SPD**

Selective Packet Discard

- **Attack of IP packets with bad TTL are processed switched with ICMP reply—crippling the router**

```
ip spd mode aggressive
```

- **`show ip spd`**

Current mode: normal.

Queue min/max thresholds: 73/74, Headroom: 100

IP normal queue: 0, priority queue: 0

SPD special drop mode: aggressively drop bad packets

What Ports Are open on the Router?

- **It may be useful to see what sockets/ports are open on the router**
- ***Show ip sockets***

```
7206-UUNET-SJ#show ip sockets
Proto    Remote       Port      Local        Port  In Out Stat TTY
OutputIF
 17 192.190.224.195   162 204.178.123.178  2168   0   0    0   0
 17   --listen--          204.178.123.178    67   0   0    9   0
 17 0.0.0.0           123 204.178.123.178   123   0   0    1   0
 17 0.0.0.0             0 204.178.123.178   161   0   0    1   0
```

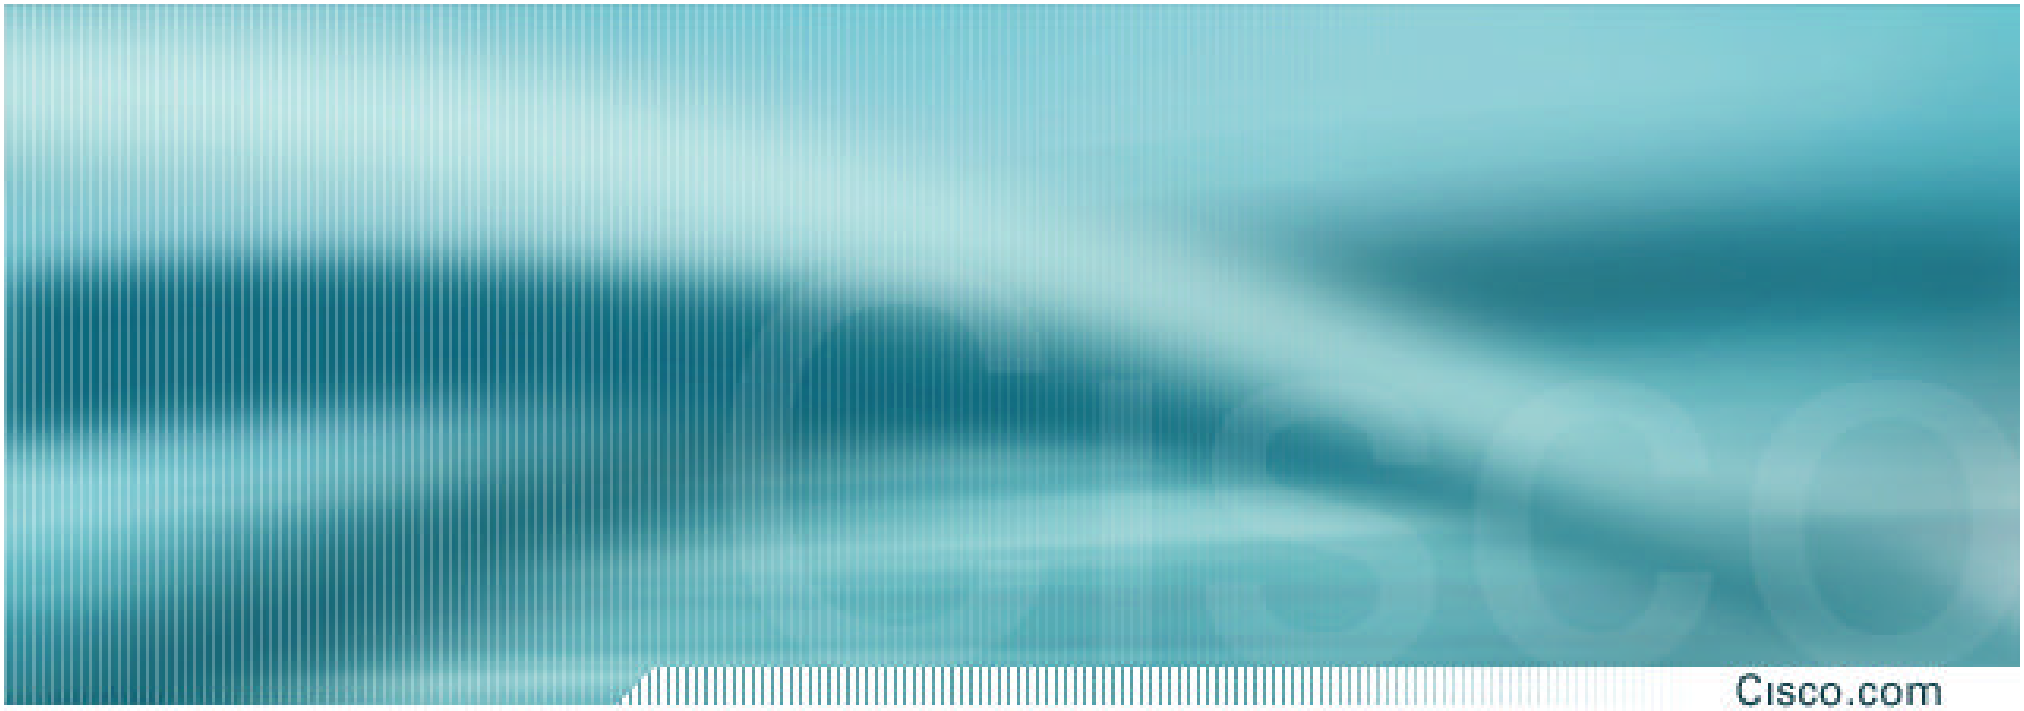

Phase 1 – Preparation for the Attack

Securing the Network

Securing the Network

Cisco.com

- Route filtering
- Packet filtering
- Rate limits

Ingress Filters—Inbound Traffic

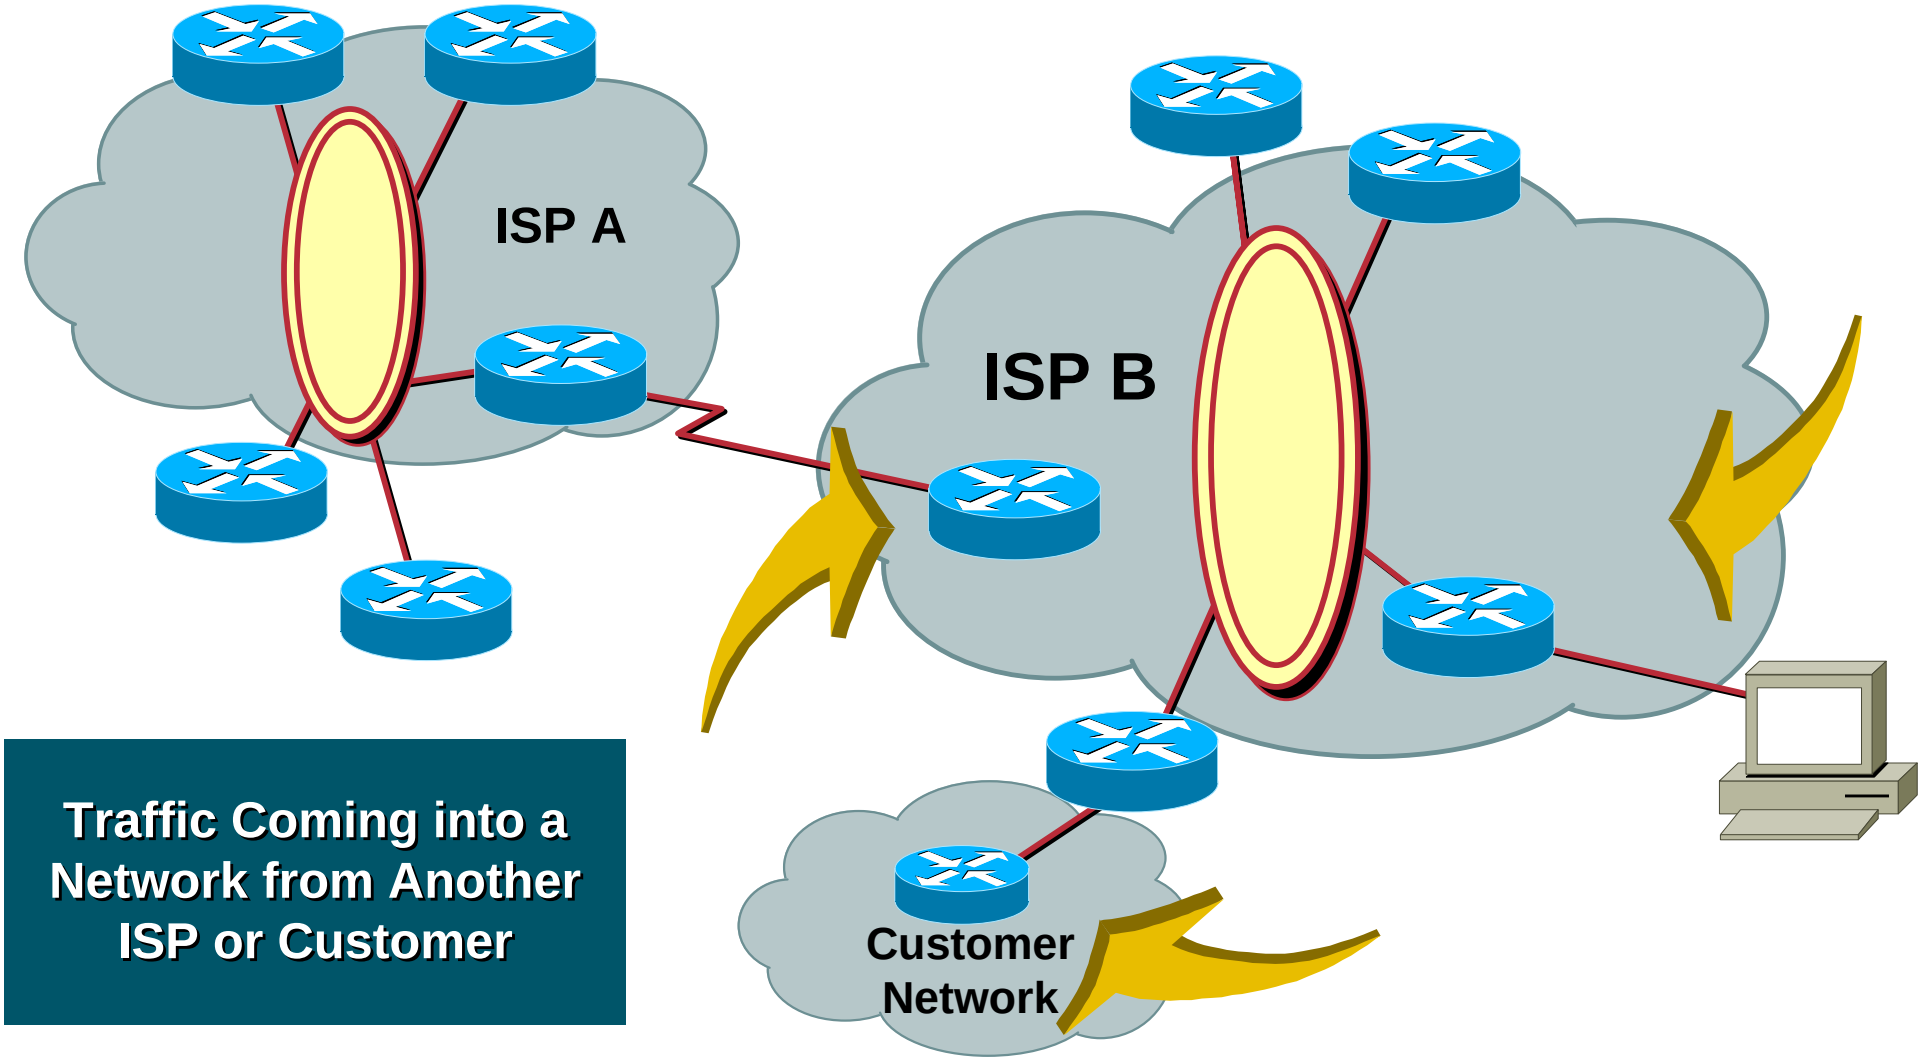

Traffic Coming into a Network from Another ISP or Customer

Egress Filters—Outbound Traffic

Cisco.com

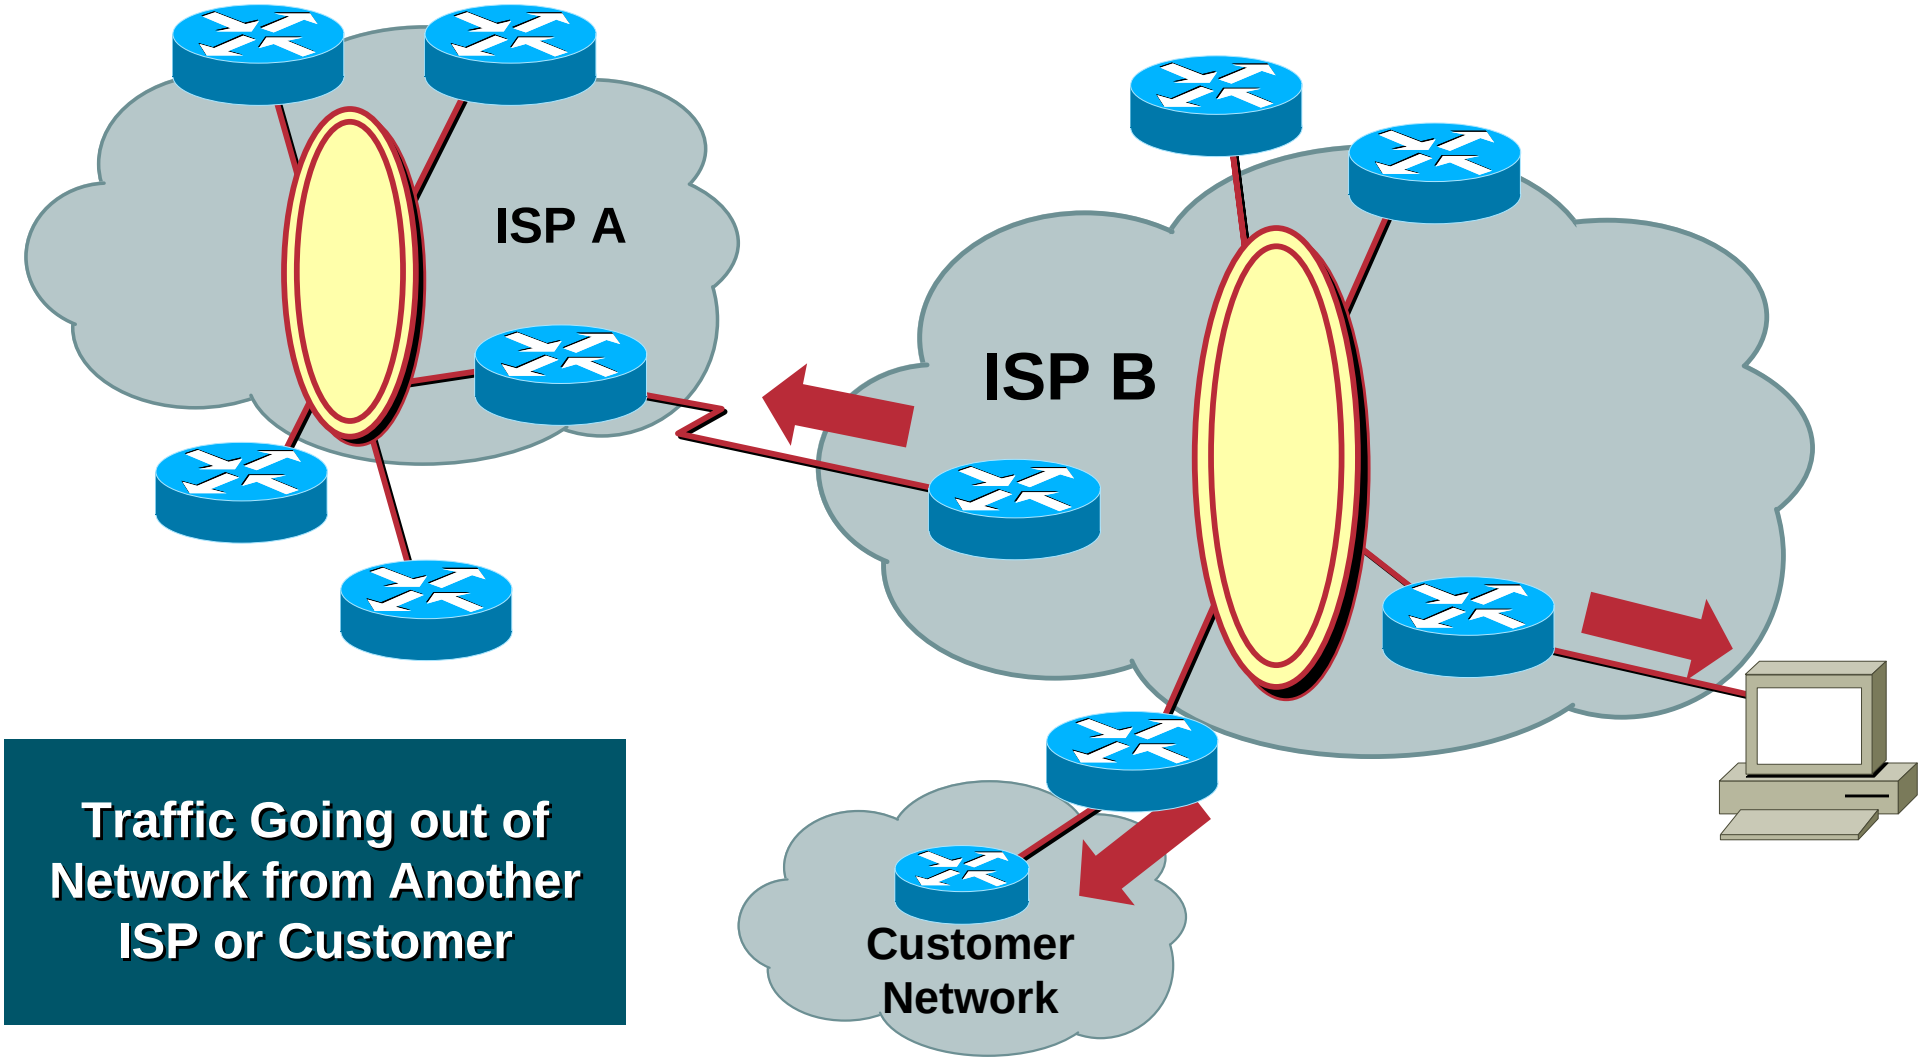

Traffic Going out of Network from Another ISP or Customer

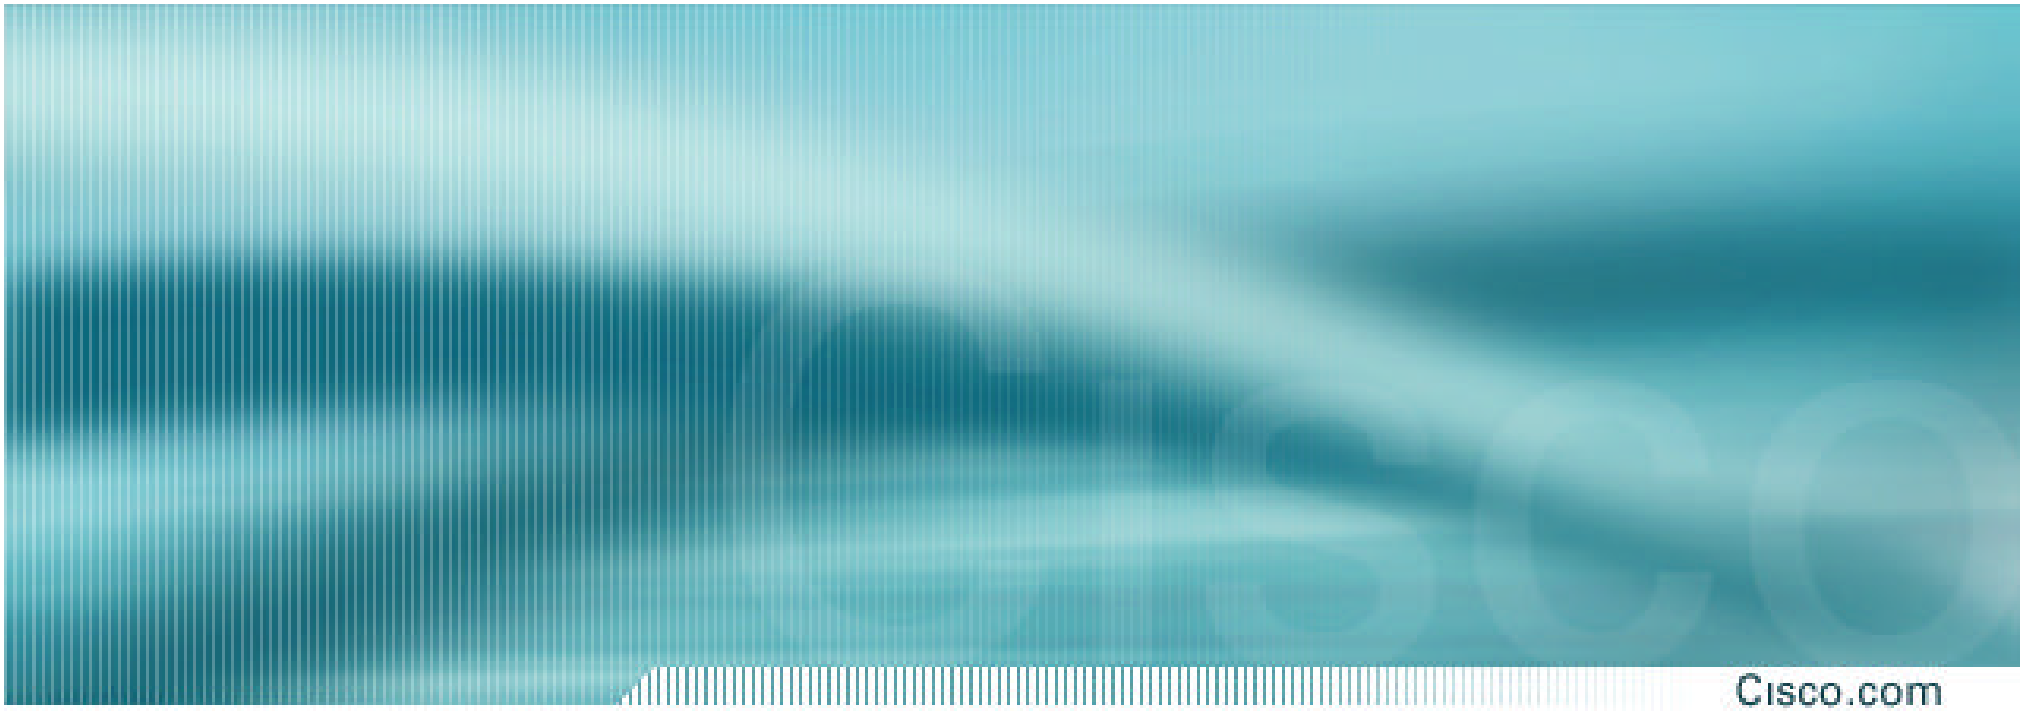

Phase 1 – Preparation for the Attack

Route Filtering

Ingress and Egress Route Filtering

- **There are routes that should NOT be routed on the Internet**
  - **RFC 1918 and “Martian” networks**
  - **127.0.0.0/8 and multicast blocks**
  - **See Bill Manning’s ID for background information:**
    - **ftp://ftp.ietf.org/internet-drafts/draft-manning-dsua-03.txt**
- **BGP should have filters applied so that these routes are not advertised to or propagated through the Internet**

Ingress and Egress Route Filtering

- **Quick review**
  - **0.0.0.0/8 and 0.0.0.0/32—Default and broadcast**
  - **127.0.0.0/8—Host loopback**
  - **192.0.2.0/24—TEST-NET for documentation**
  - **10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16—RFC 1918 private addresses**
  - **169.254.0.0/16—End node auto-config for DHCP**

Ingress and Egress Route Filtering

- Two flavors of route filtering:
  - Distribute list—Widely used
  - Prefix list—Increasingly used
- Both work fine—Engineering preference

Ingress and Egress Route Filtering

Cisco.com

Extended ACL for a BGP Distribute List

```
access-list 150 deny ip host 0.0.0.0 any
access-list 150 deny ip 10.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 150 deny ip 127.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 150 deny ip 169.254.0.0 0.0.255.255 255.255.0.0 0.0.255.255
access-list 150 deny ip 172.16.0.0 0.15.255.255 255.240.0.0 0.15.255.255
access-list 150 deny ip 192.0.2.0 0.0.0.255 255.255.255.0 0.0.0.255
access-list 150 deny ip 192.168.0.0 0.0.255.255 255.255.0.0 0.0.255.255
access-list 150 deny ip 224.0.0.0 31.255.255.255 224.0.0.0 31.255.255.255
access-list 150 permit ip any any
```

Ingress and Egress Route Filtering

Cisco.com

BGP with Distribute List Flavor of Route Filtering

```
router bgp 200
no synchronization
bgp dampening
 neighbor 220.220.4.1 remote-as 210
 neighbor 220.220.4.1 version 4
 neighbor 220.220.4.1 distribute-list 150 in
 neighbor 220.220.4.1 distribute-list 150 out
 neighbor 222.222.8.1 remote-as 220
 neighbor 222.222.8.1 version 4
 neighbor 222.222.8.1 distribute-list 150 in
 neighbor 222.222.8.1 distribute-list 150 out
 no auto-summary
!
```

Ingress and Egress Route Filtering

Prefix-List for a for a BGP Prefix List

```
ip prefix-list rfc1918-dsua deny   0.0.0.0/8 le 32
ip prefix-list rfc1918-dsua deny   10.0.0.0/8 le 32
ip prefix-list rfc1918-dsua deny   127.0.0.0/8 le 32
ip prefix-list rfc1918-dsua deny   169.254.0.0/16 le 32
ip prefix-list rfc1918-dsua deny   172.16.0.0/12 le 32
ip prefix-list rfc1918-dsua deny   192.0.2.0.0/24 le 32
ip prefix-list rfc1918-dsua deny   192.168.0.0/16 le 32
ip prefix-list rfc1918-dsua deny   224.0.0.0/3 le 32
ip prefix-list rfc1918-dsua permit 0.0.0.0/0 le 32
```

Ingress and Egress Route Filtering

Cisco.com

BGP with Prefix-List Flavor of Route Filtering

```
router bgp 200
no synchronization
bgp dampening
 neighbor 220.220.4.1 remote-as 210
 neighbor 220.220.4.1 version 4
 neighbor 220.220.4.1 prefix-list rfc1918-dsua in
 neighbor 220.220.4.1 prefix-list rfc1918-dsua out
 neighbor 222.222.8.1 remote-as 220
 neighbor 222.222.8.1 version 4
 neighbor 222.222.8.1 prefix-list rfc1918-dsua in
 neighbor 222.222.8.1 prefix-list rfc1918-dsua out
 no auto-summary
!
```

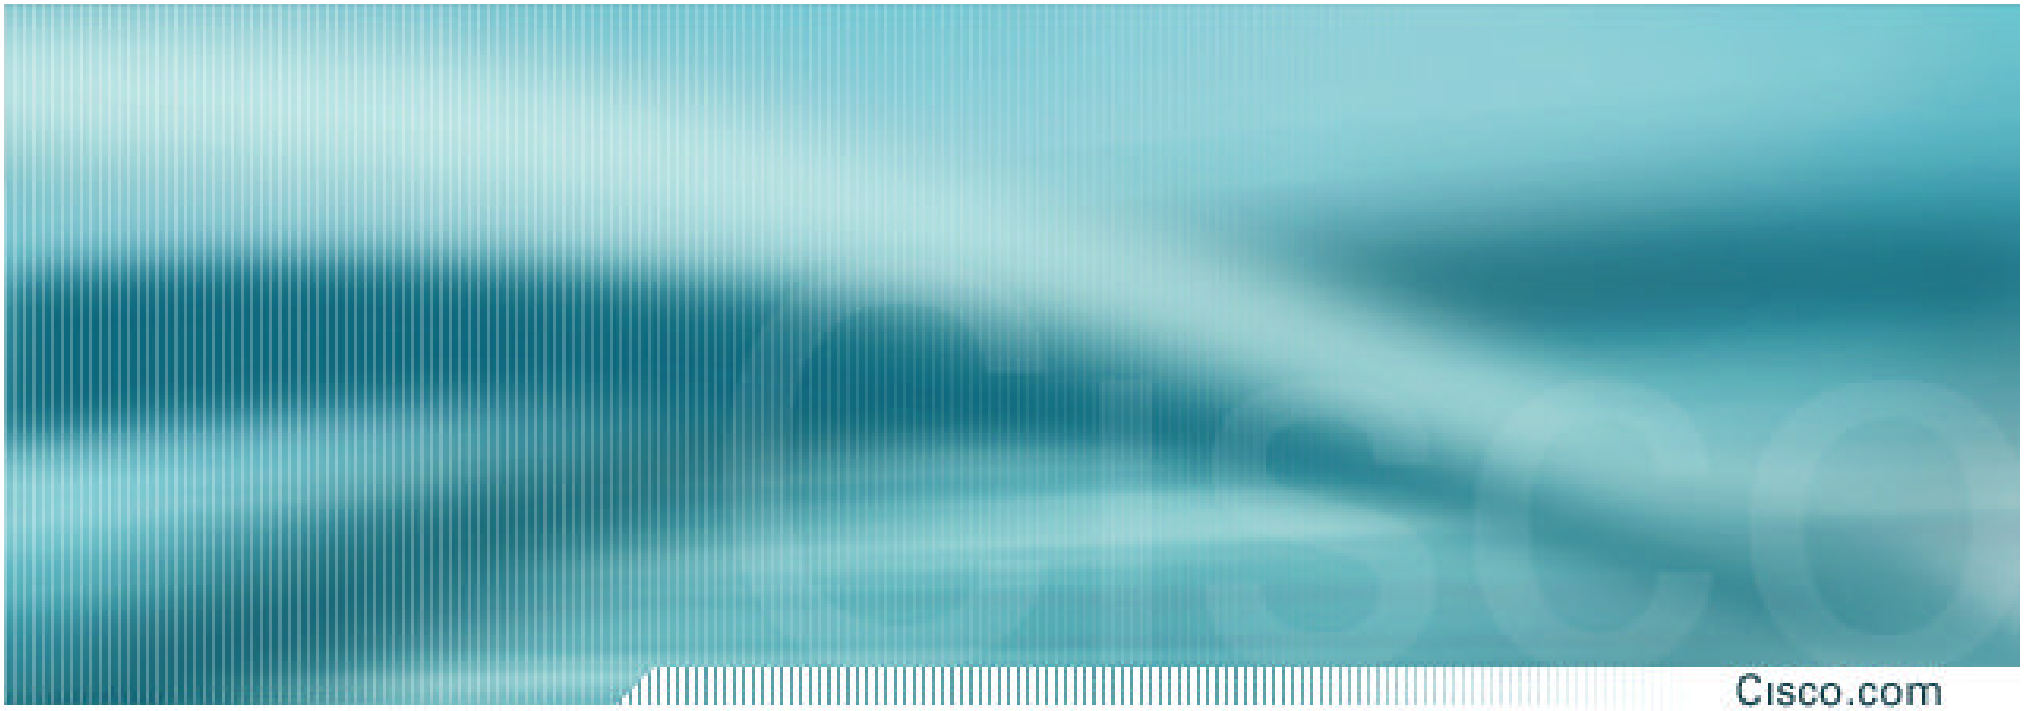

Phase 1 – Preparation for the Attack

Black Hole Filtering

Black Hole Filtering

- ***Black Hole Filtering*** **or** ***Black Hole Routing*** **forwards a packet to a router's** ***bit bucket*****.**

  Also known as "route to Null0"

- **Works only on destination addresses, since it is really part of the forwarding logic.**
- **Forwarding ASICs are designed to work with routes to Null0 – dropping the packet with minimal to no performance impact (depending on the forwarding ASIC).**
- **Used for years as a means to "black hole" un wanted packets.**

Black Hole Filtering

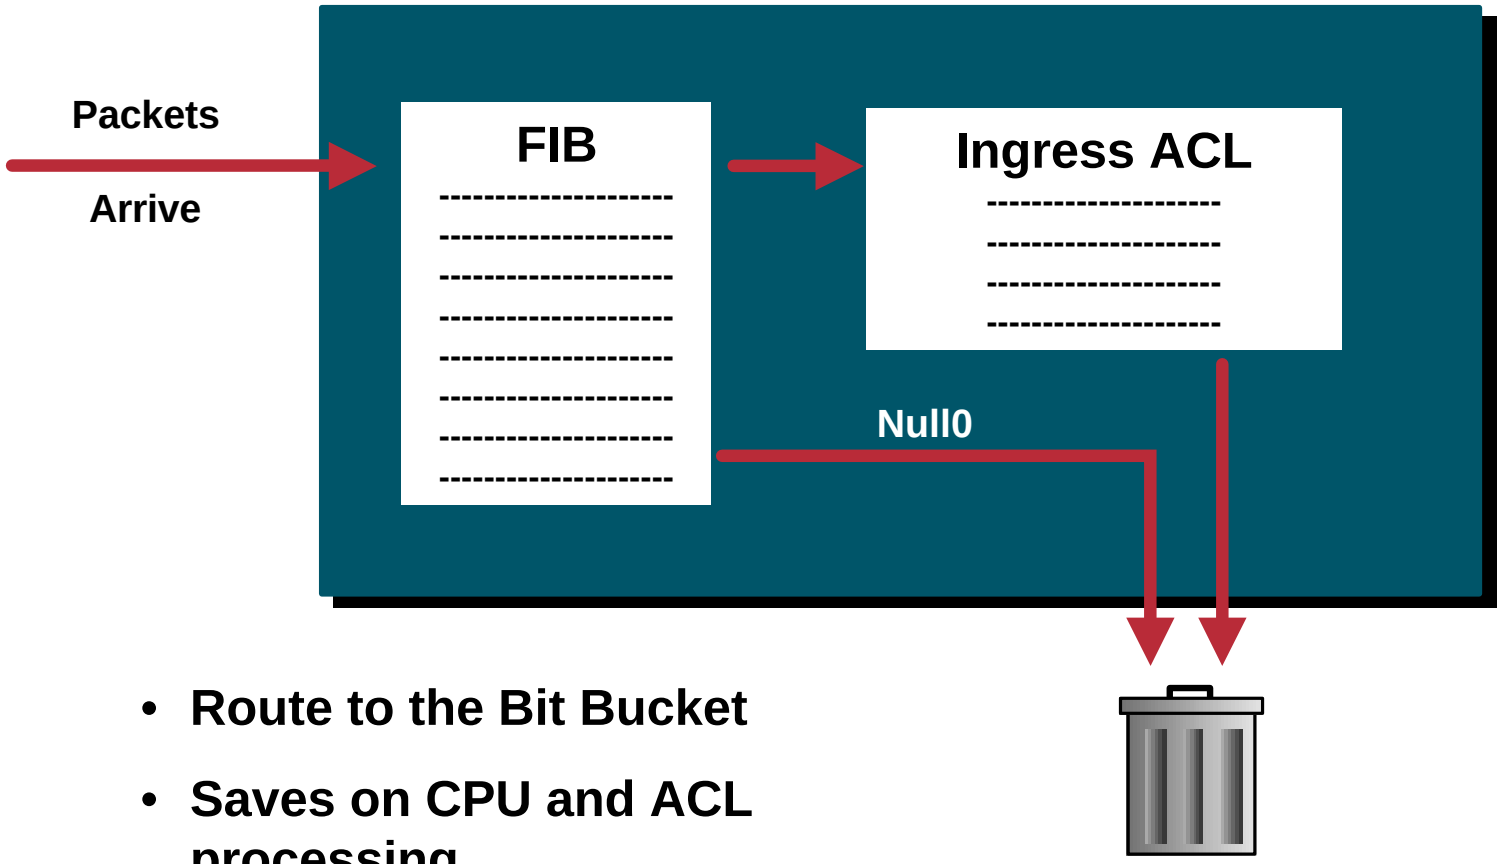

- **Route to the Bit Bucket**
- **Saves on CPU and ACL processing**

Remotely Triggered Black Hole Filtering

- A simple static route and BGP will allow an ISP to trigger network wide black holes as fast as iBGP can update the network.
- This provides ISPs a tool that can be used to respond to security related events or used for DOS/DDOS Backscatter Tracebacks.

Remotely Triggered Black Hole Filtering - Preparation

1. Select a small block that will not be used for anything other than black hole filtering. Test Net (192.0.2.0/24) is optimal since it should not be on the Net and is not really used.
2. Put a static route with Test Net – 192.0.2.0/24 to Null 0 on every router on the network.
3. Prepare a BGP speaking router that will be used to announce the network to be Black Holed (see config example on next slide).

Remotely Triggered Black Hole Filtering - Preparation

Cisco.com

```
router bgp 109
.
redistribute static route-map static-to-bgp
.
!
route-map static-to-bgp permit 10
match tag 66
set ip next-hop 192.0.2.1
set local-preference 50
set community no-export
set origin igp
!
Route-map static-to-bgp permit 20
```

Remotely Triggered Black Hole Filtering - Activation

1. **ISP adds a static route of the destination address they wish to black hole to the advertising router. The static is added with the “tag 66” to keep it separate from other statics on the router.**

```
ip route 171.68.1.1 255.255.255.255 Null0 Tag 66
```

2. **BGP Advertisement goes out to all BGP speaking routers.**
3. **Router hear the announcement, glues it to the existing static on the route, and changes the next-hop for the BGP advertised route to Null0 – triggering black hole routing.**

Remotely Triggered Black Hole Filtering - Activation

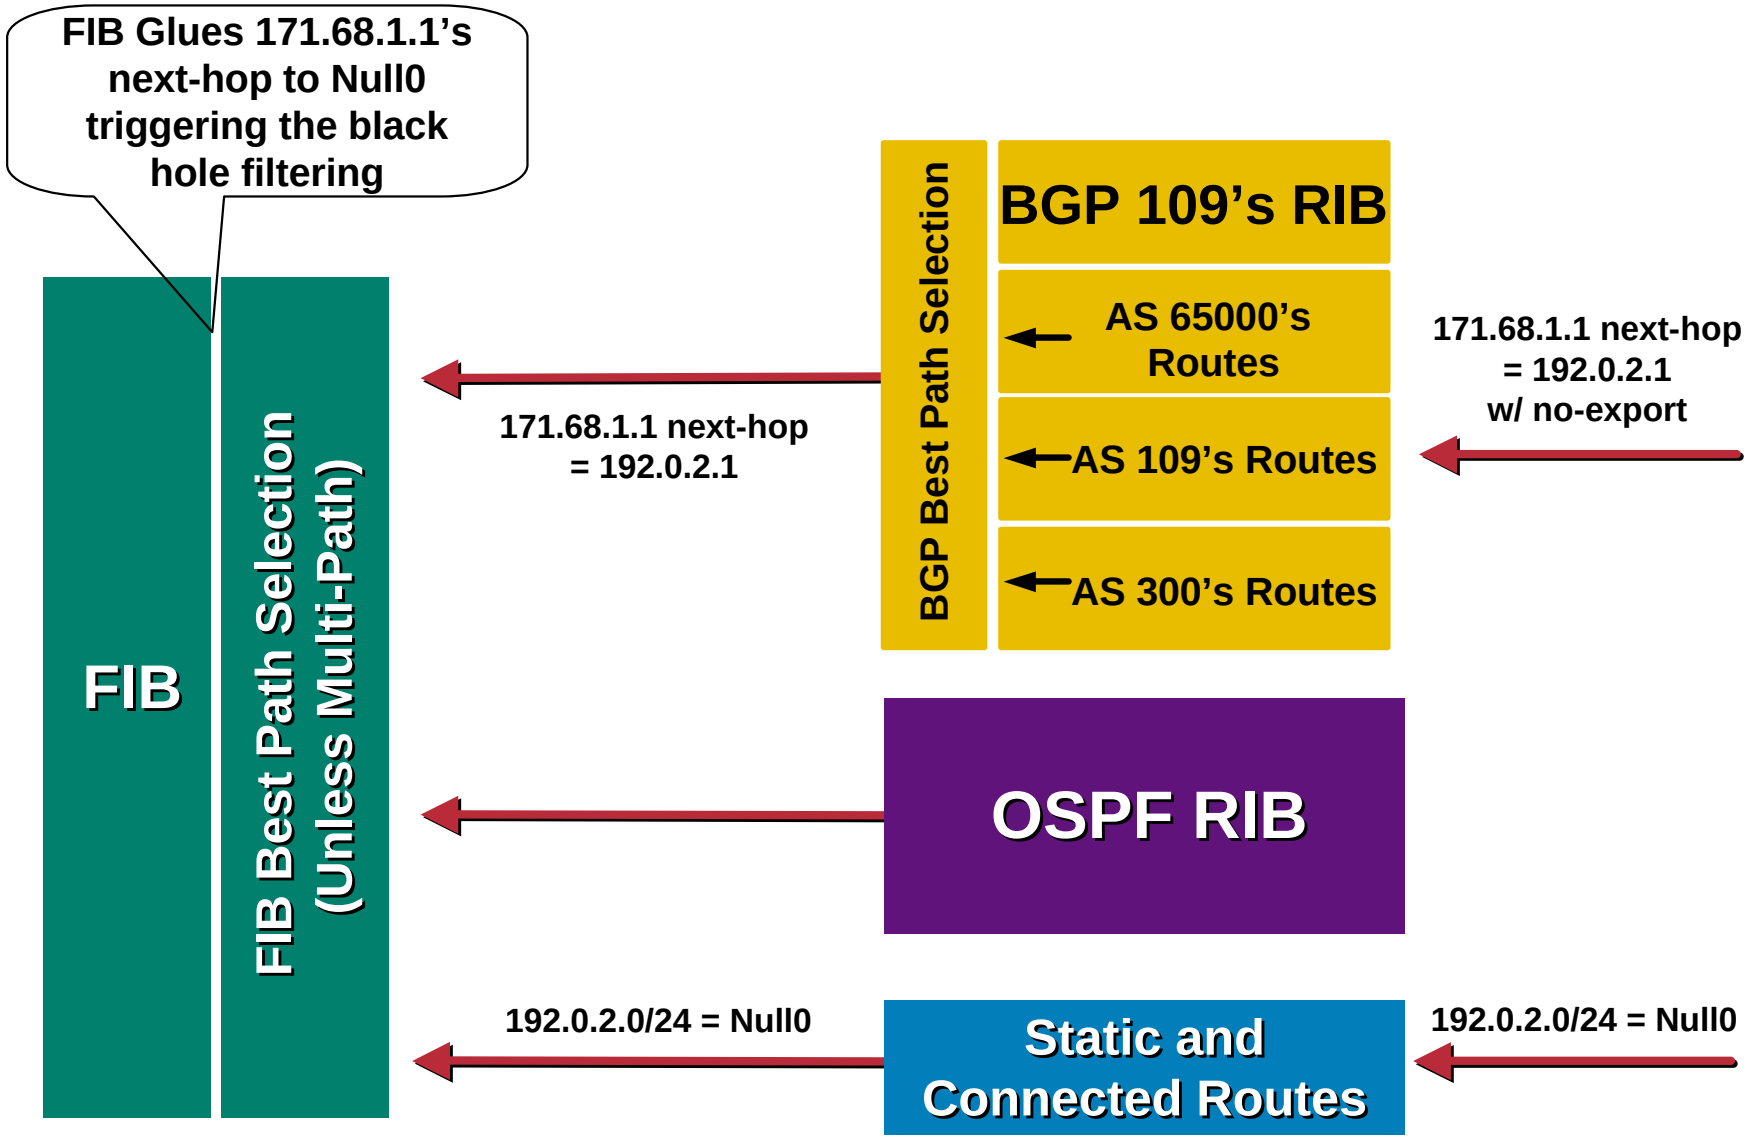

Remotely Triggered Black Hole Filtering - Activation

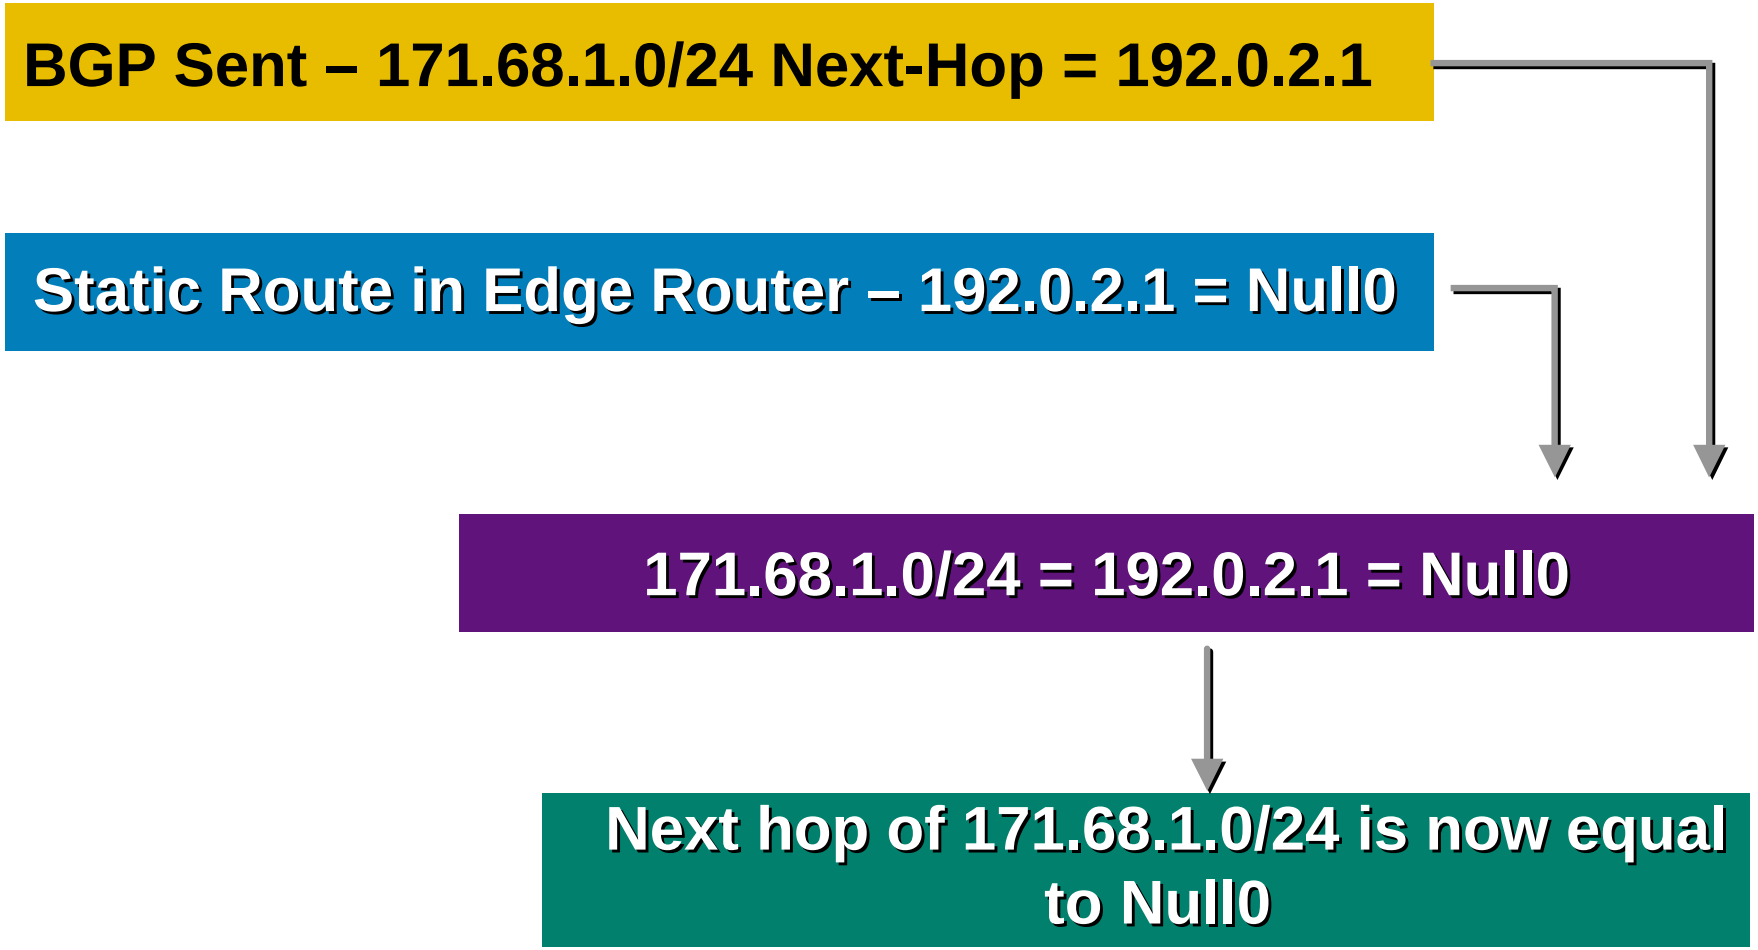

Remotely Triggered Black Hole Filtering - Activation

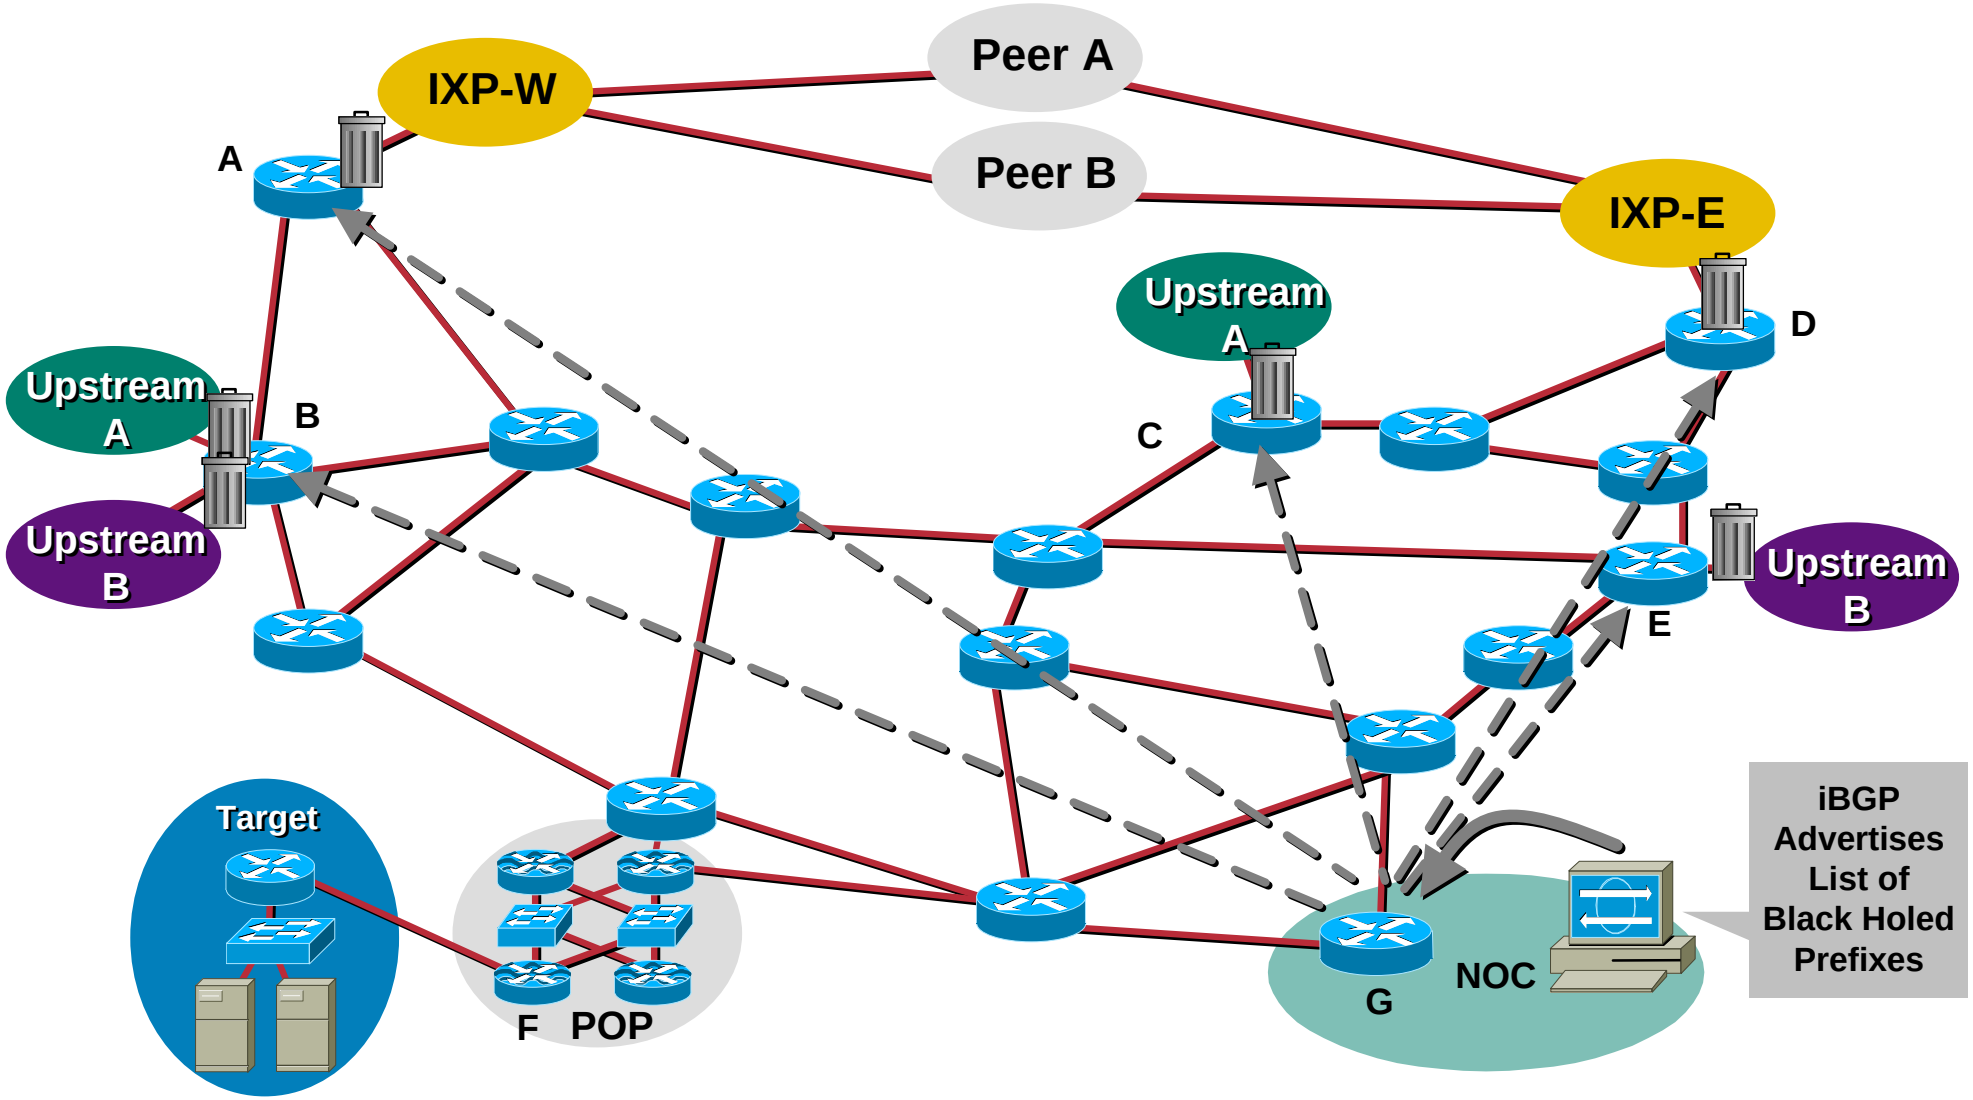

Gotchas with Black Hole Filtering

- Routers were designed to forward traffic, not drop traffic.
- ASIC Based Forwarding can drop traffic at line rate.
- Processor Based Forwarding can have problems dropping large amounts of data.
- Remember the old shunt technique ….

Gotchas with Black Hole Filtering

- **Back in the days when this was in the core of the Internet …..**

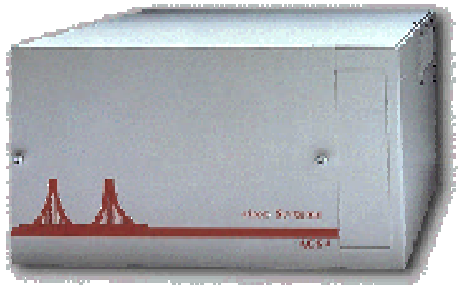

- **All “drops” to Null0 were process switched.**
- **Fast Drops fixed the problem for a while, but traffic loads increased to the to where they could not drop at line rate anymore.**
- **Bottomline – Software based forwarding routers (any vendor) can forward faster then they can drop.**

Black Hole Shunt

- **Black Hole *Shunts* are used to forward traffic out a spoofed interface.**

  Classic Example: AUI/BNC Transceiver with a T connector. A static MAC address is used with a static route.

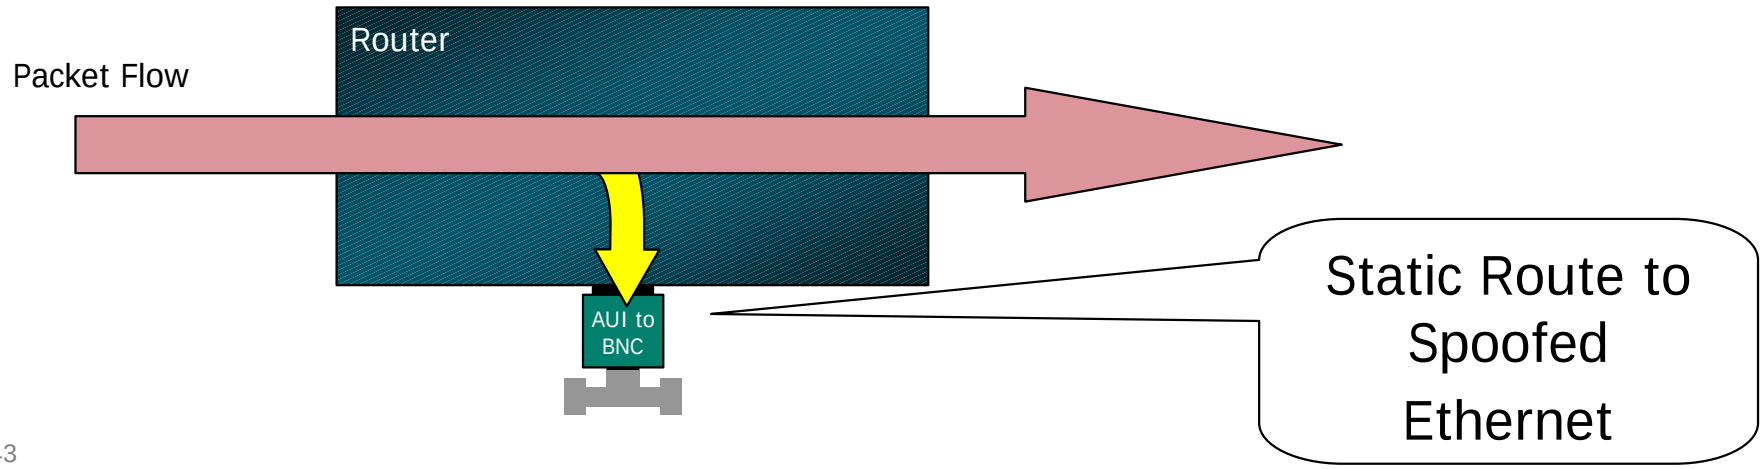

Black Hole Shunt

- **Some ISPs used Black Hole Shunts during Code Red.**

  Routers that injected Default Sucked all traffic to them.

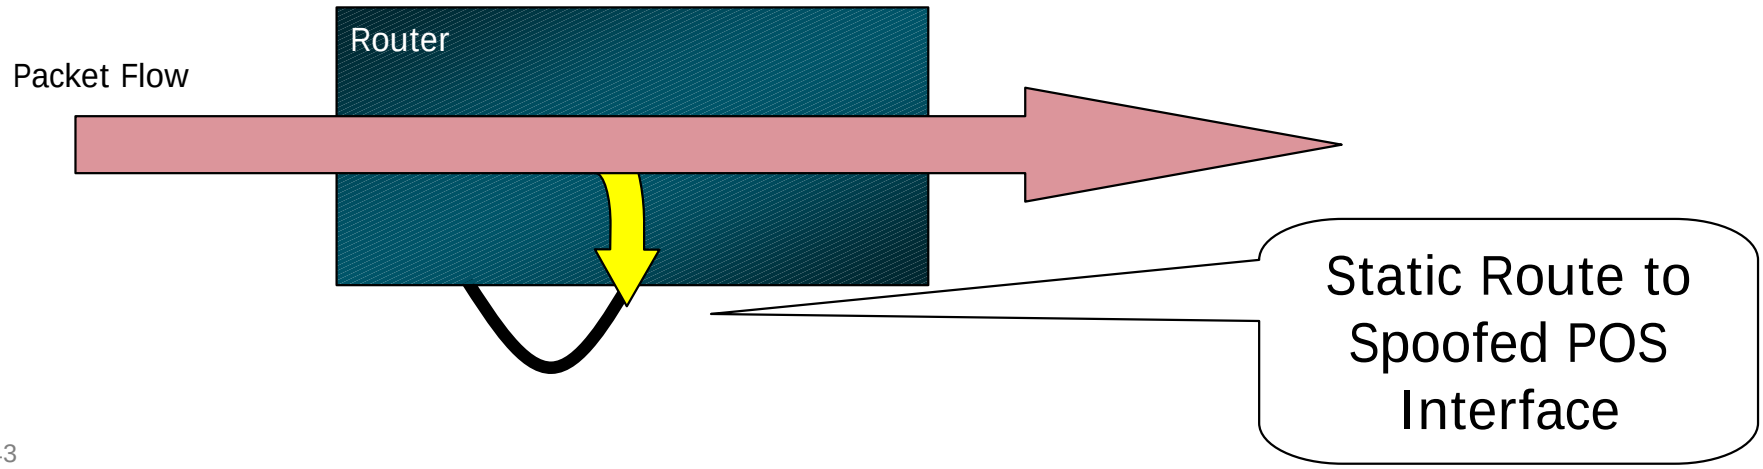

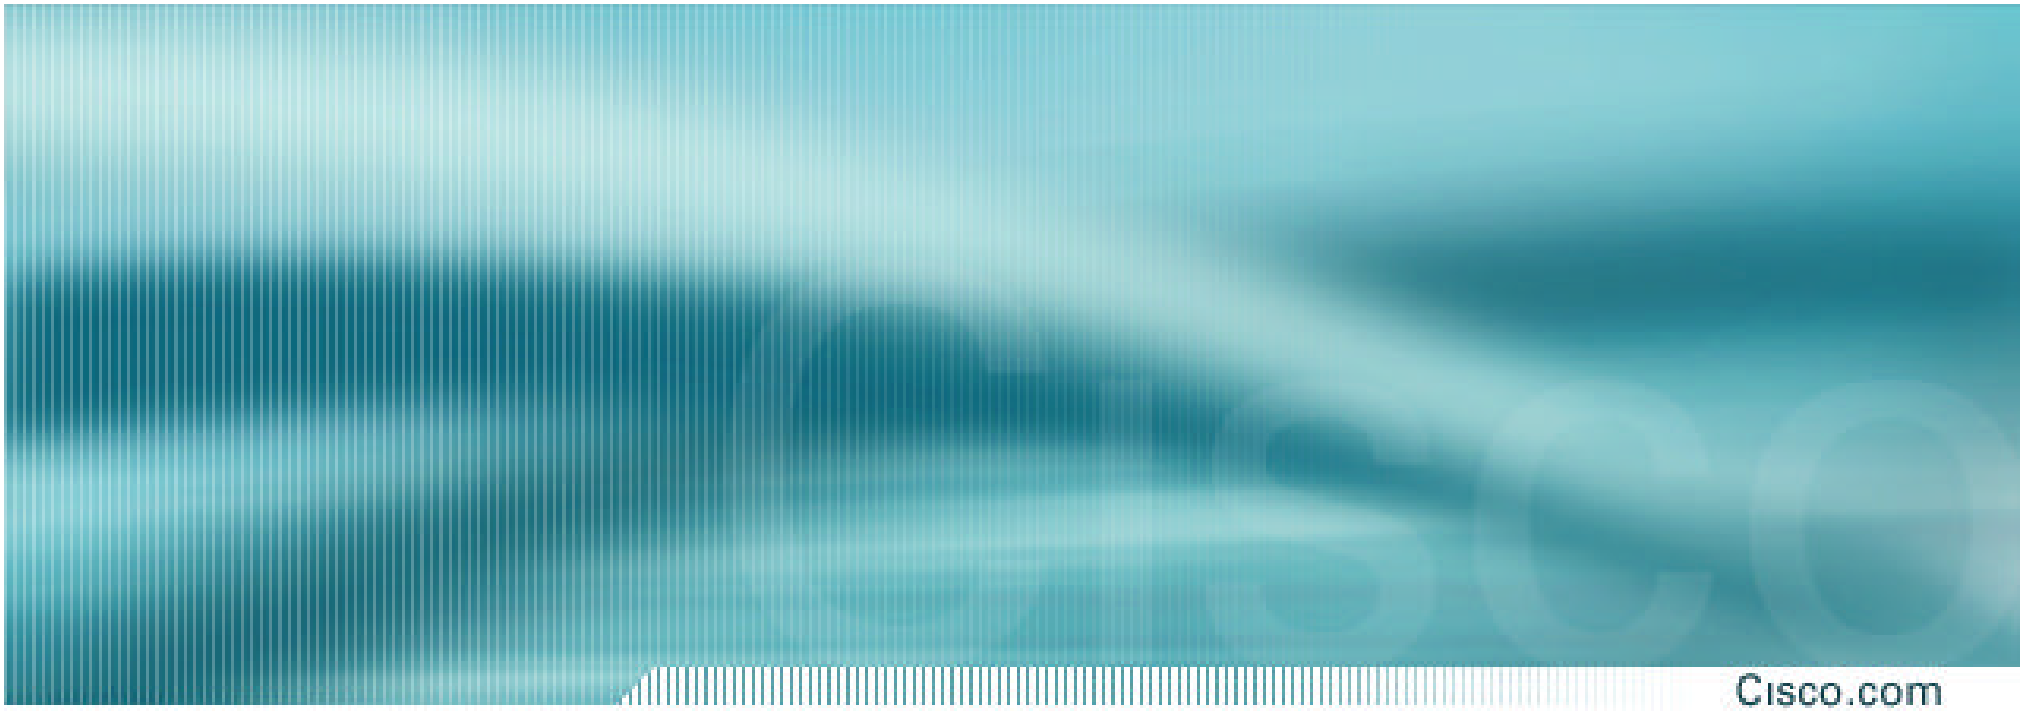

Phase 1 – Preparation for the Attack

Sink Hole Routers/Networks

Sink Hole Routers/Networks

- **Sink Holes are a the network equivalent of a honey pot.**
  - **BGP speaking Router or Workstation that built to *suck in* attacks.**
  - **Used to redirect attacks away from the customer – working the attack on a router built to withstand the attack.**
  - **Used to monitor *attack noise, scans,* and other activity (via the advertisement of default)**

Sink Hole Routers/Networks

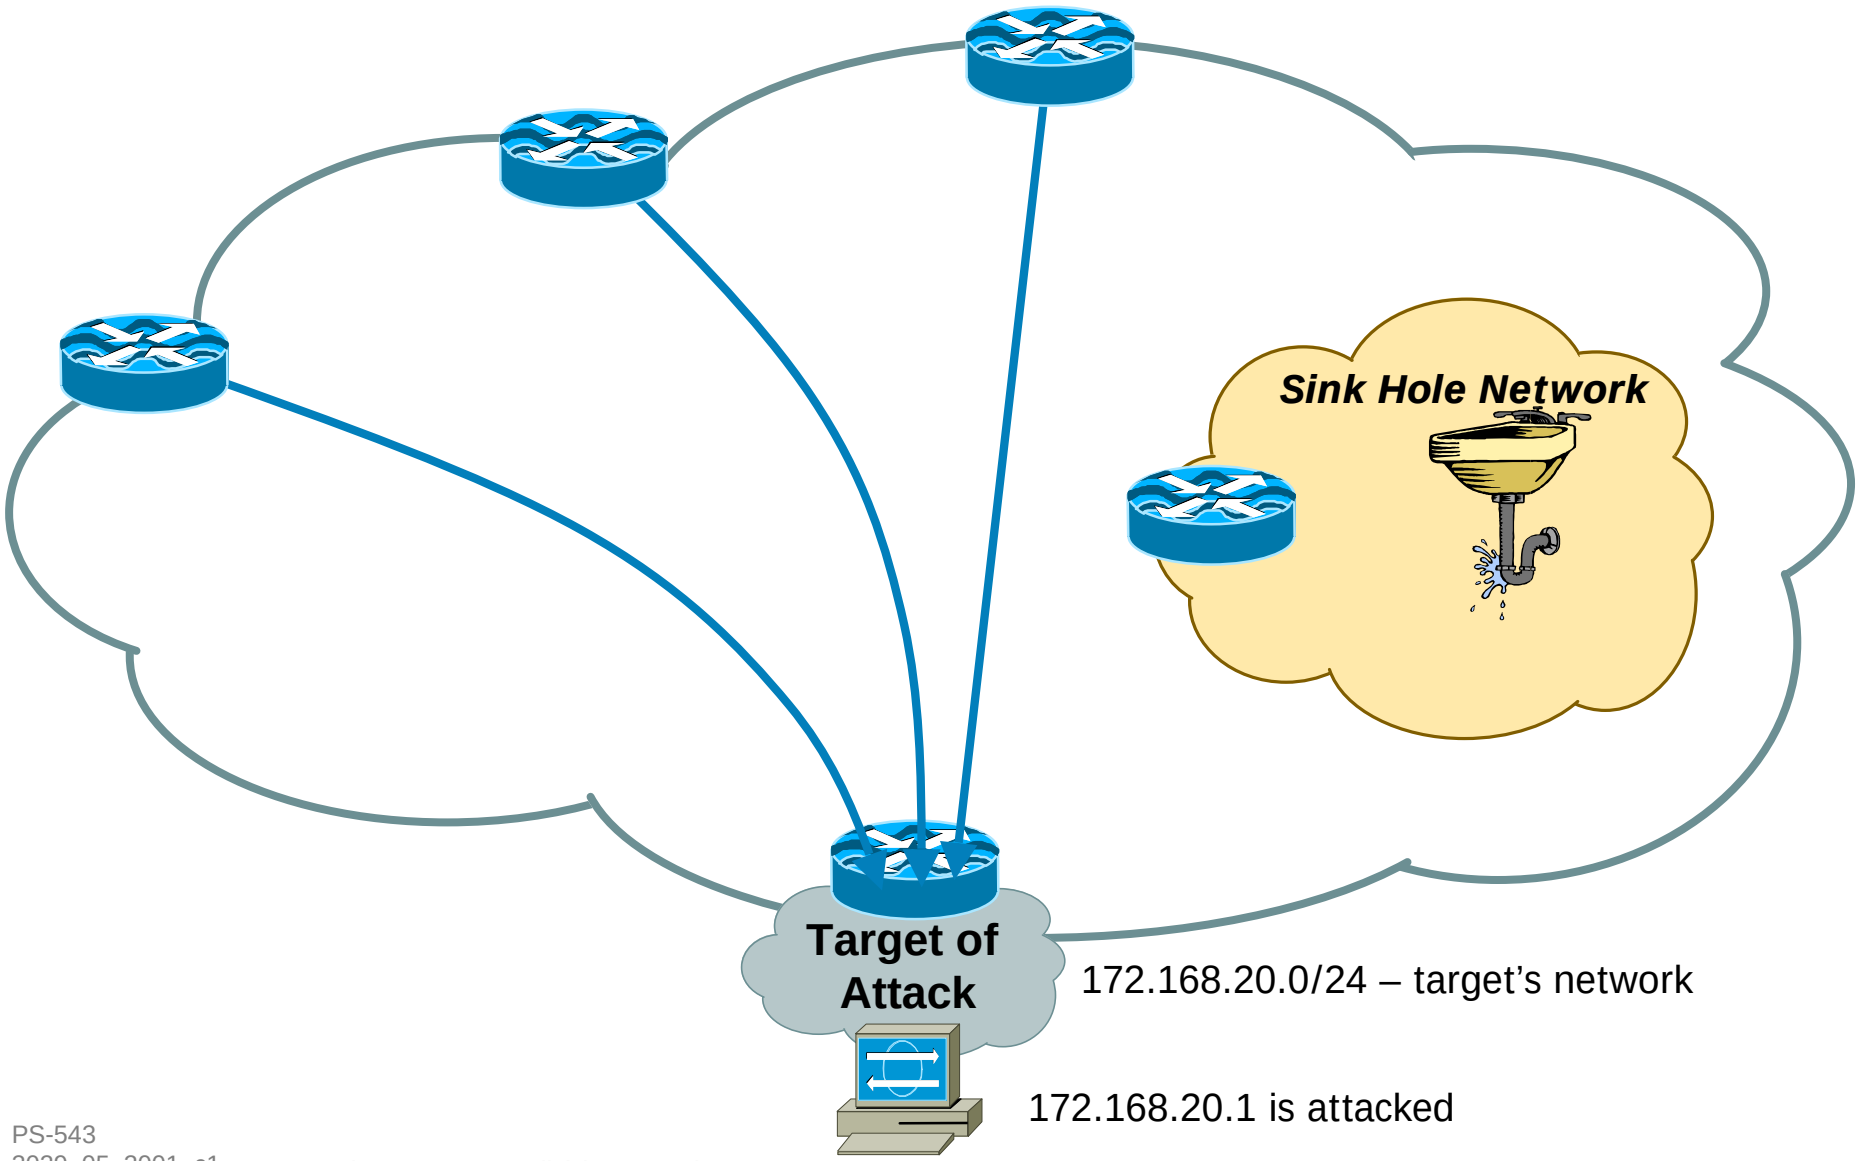

Sink Hole Routers/Networks

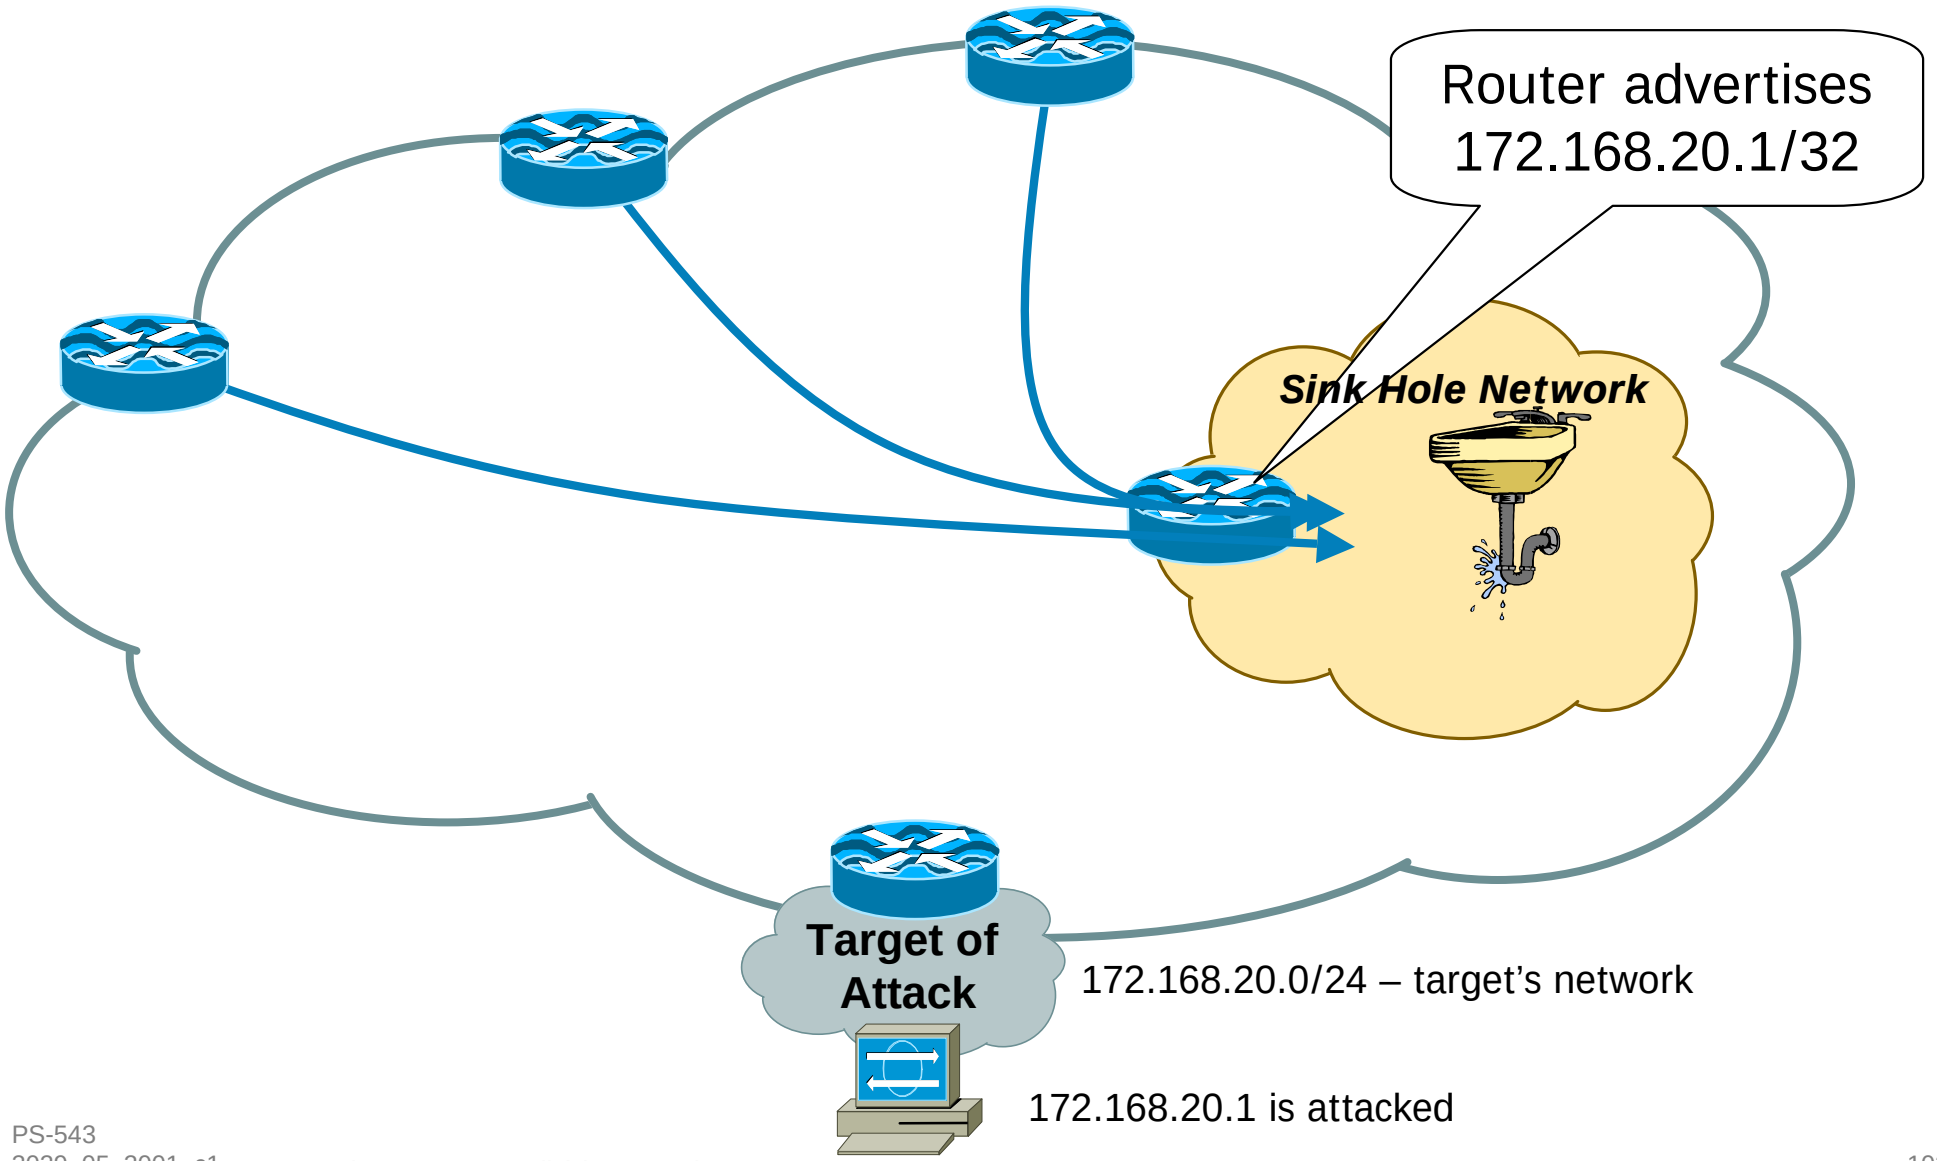

Sink Hole Routers/Networks

- **Attack is pulled off customer and your aggregation router.**
- **Can now do classification ACLs, Flow Analysis, Sniffer Capture, Traceback, etc.**
- **Objective is to minimize the risk to the network while working the attack incident.**

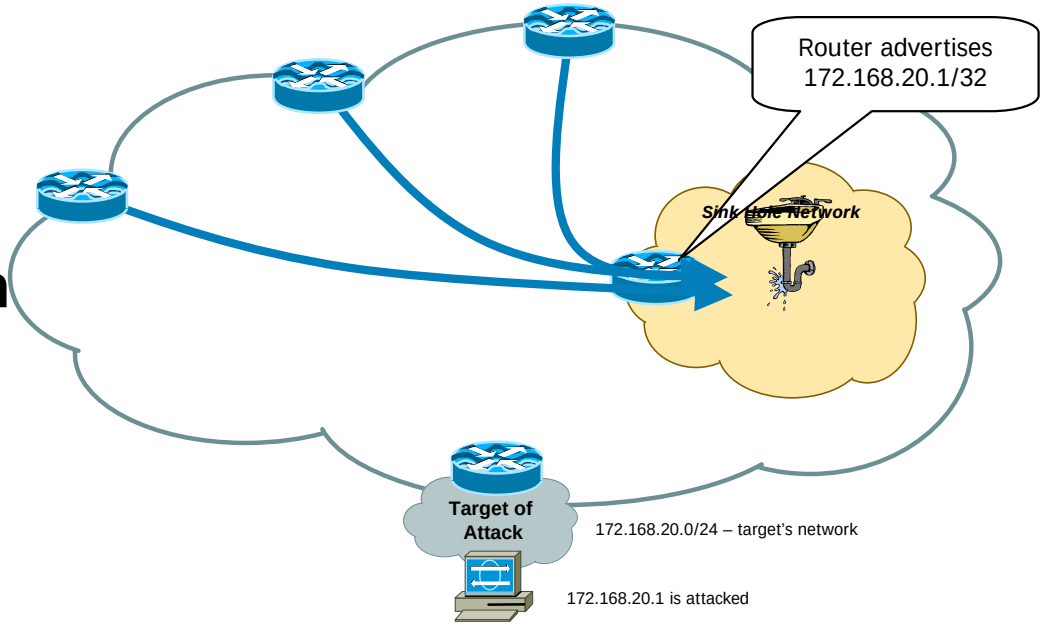

Sink Hole Routers/Networks

- **Advertising Default from the Sink Hole will pull down all sort of *junk* traffic.**
  - **Customer Traffic when circuits flap.**
  - **Network Scans**
  - **Failed Attacks**
  - **Code Red/NIMDA**
  - **Backscatter**
- **Can place tracking tools and IDA in the Sink Hole network to monitor the noise.**

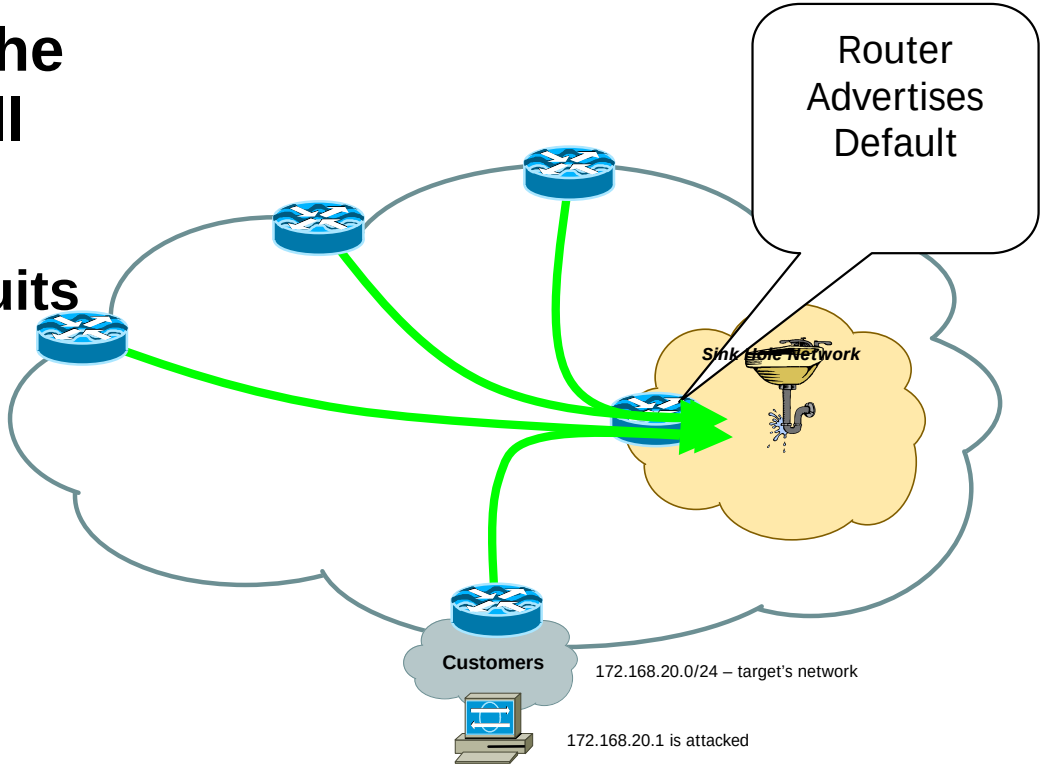

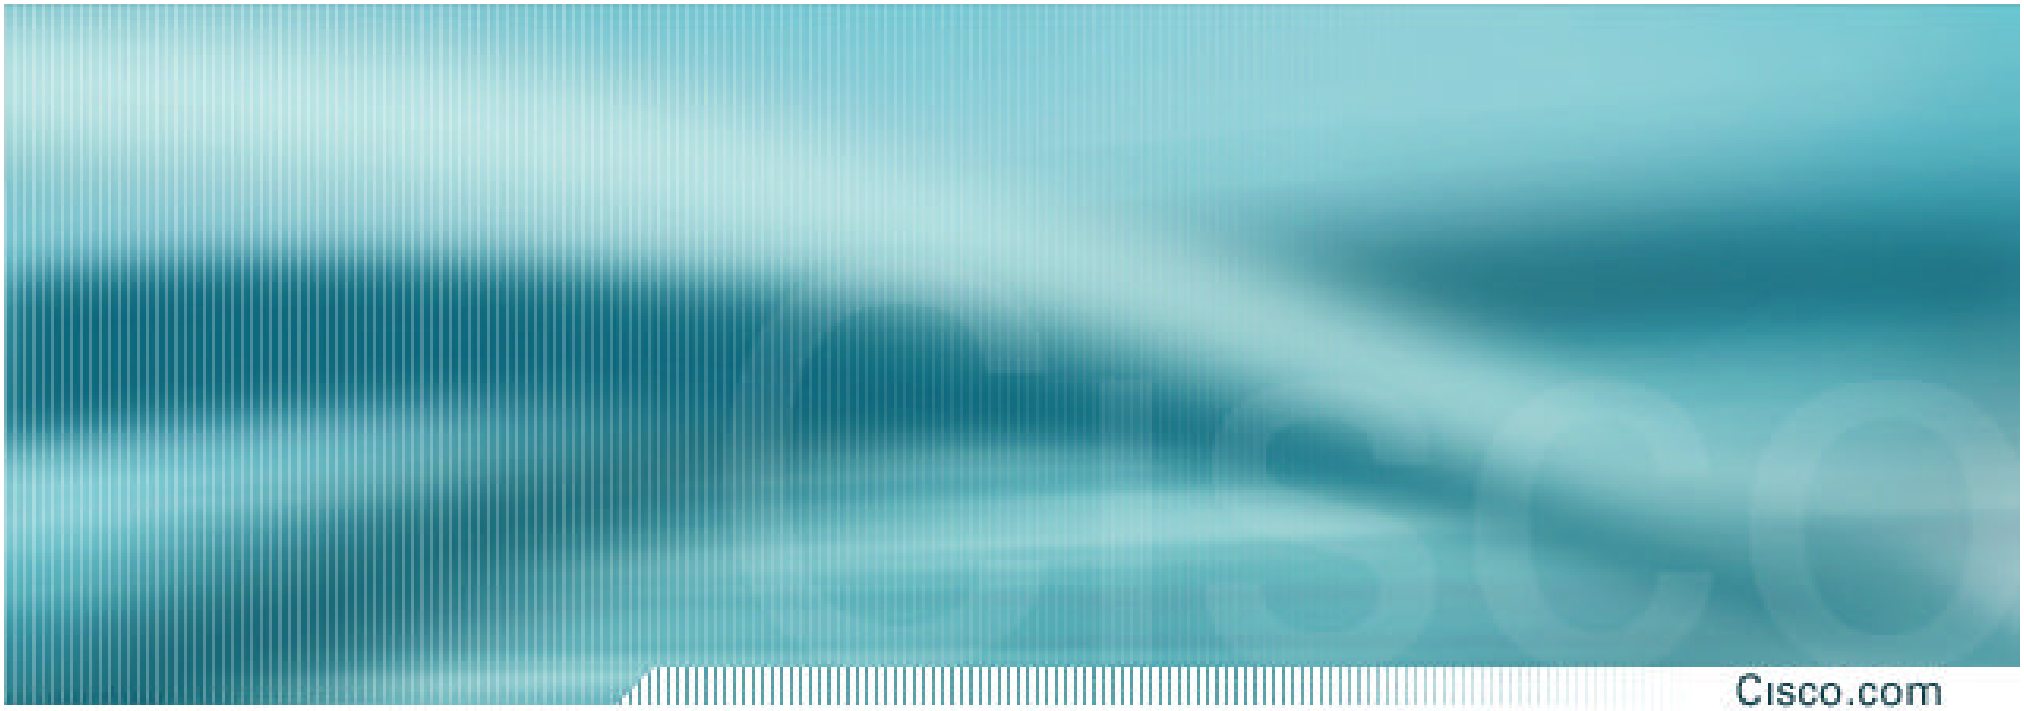

Phase 1 – Preparation for the Attack

Packet Filtering

Ingress and Egress Packet Filtering

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

Ingress and Egress Packet Filtering

- 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

Packet Filtering

Cisco.com

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

Egress Packet Filtering Upstream Border

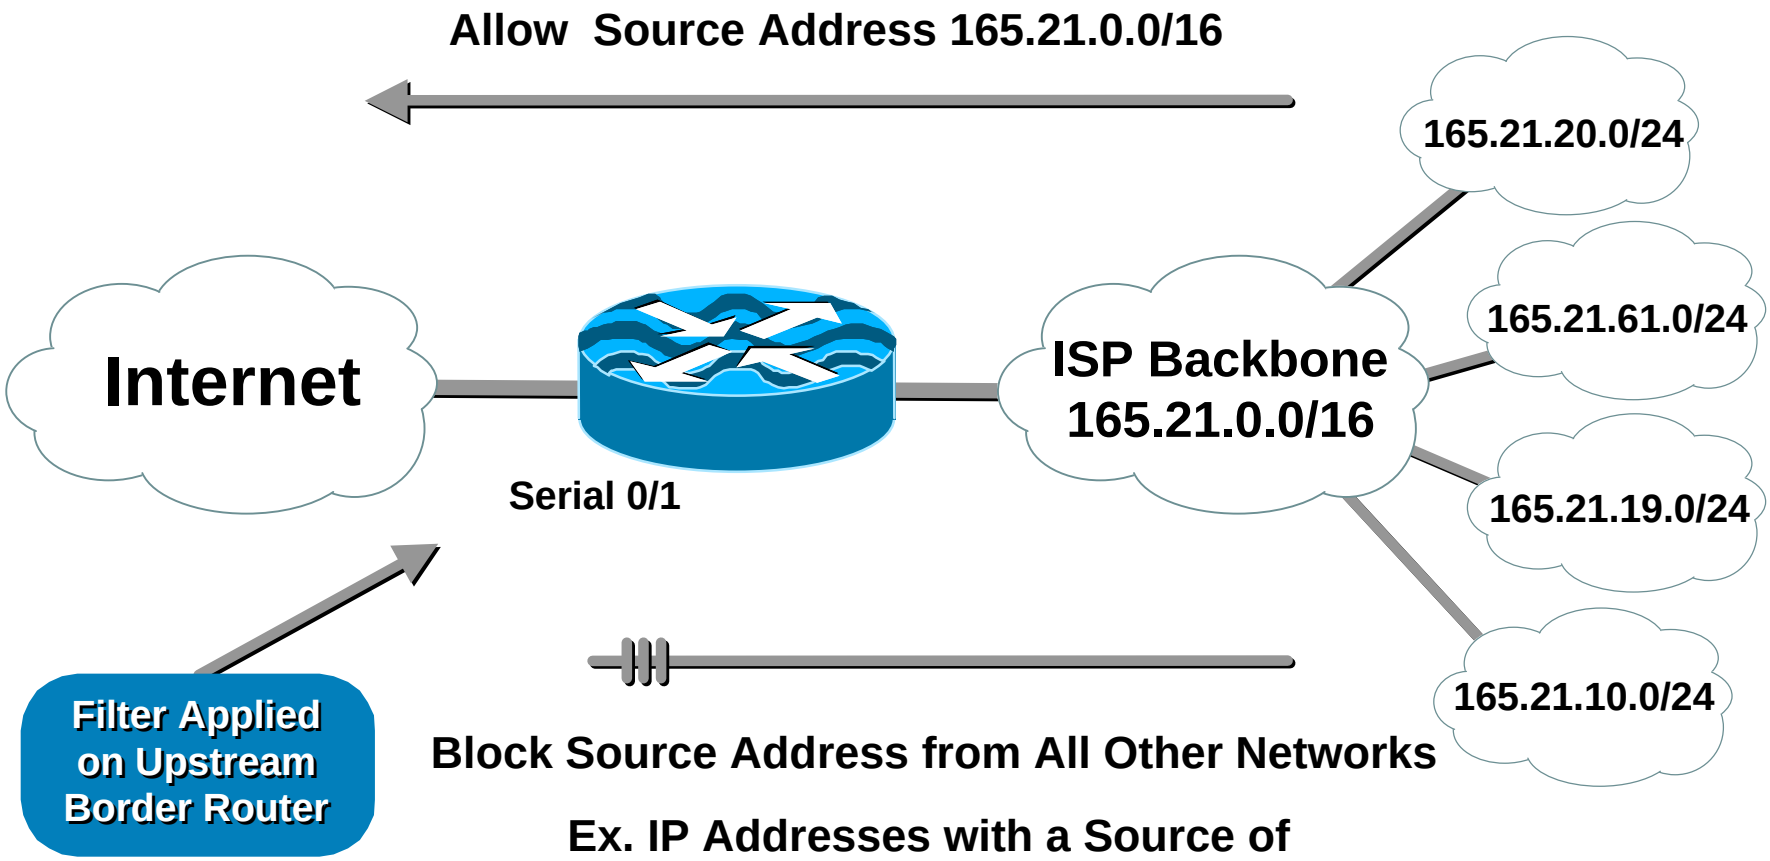

Ingress Packet Filtering Upstream Border

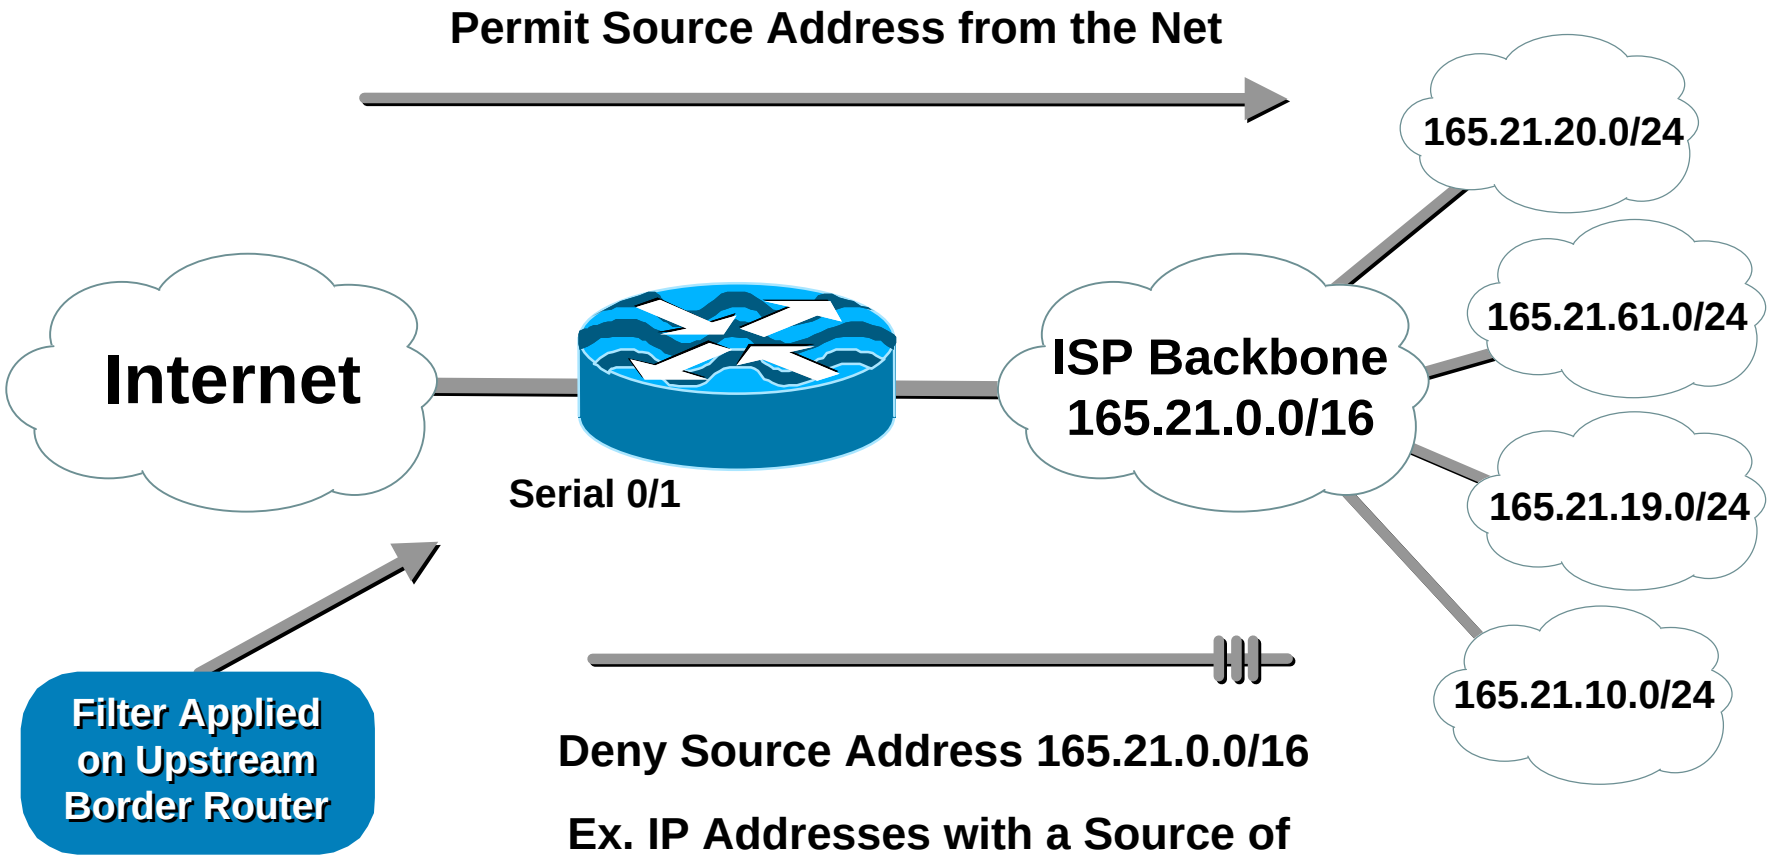

Ingress Packet Filtering Customer Edge

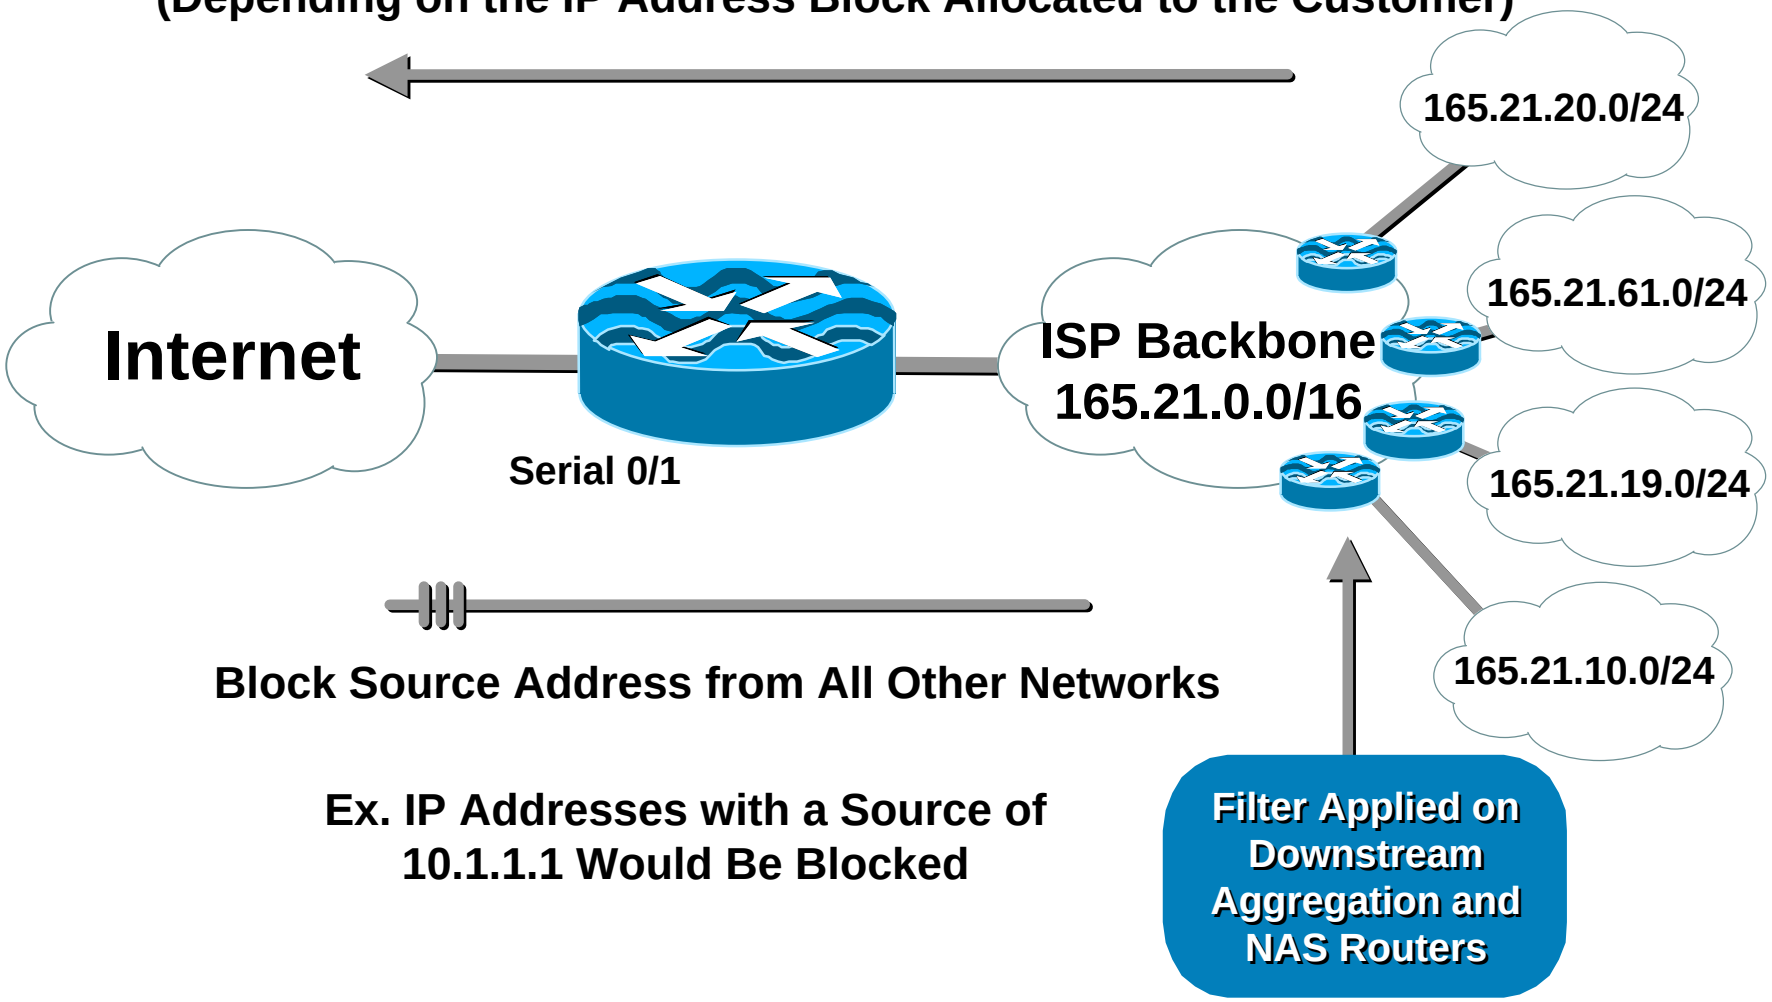

Egress Packet Filtering Customer Edge

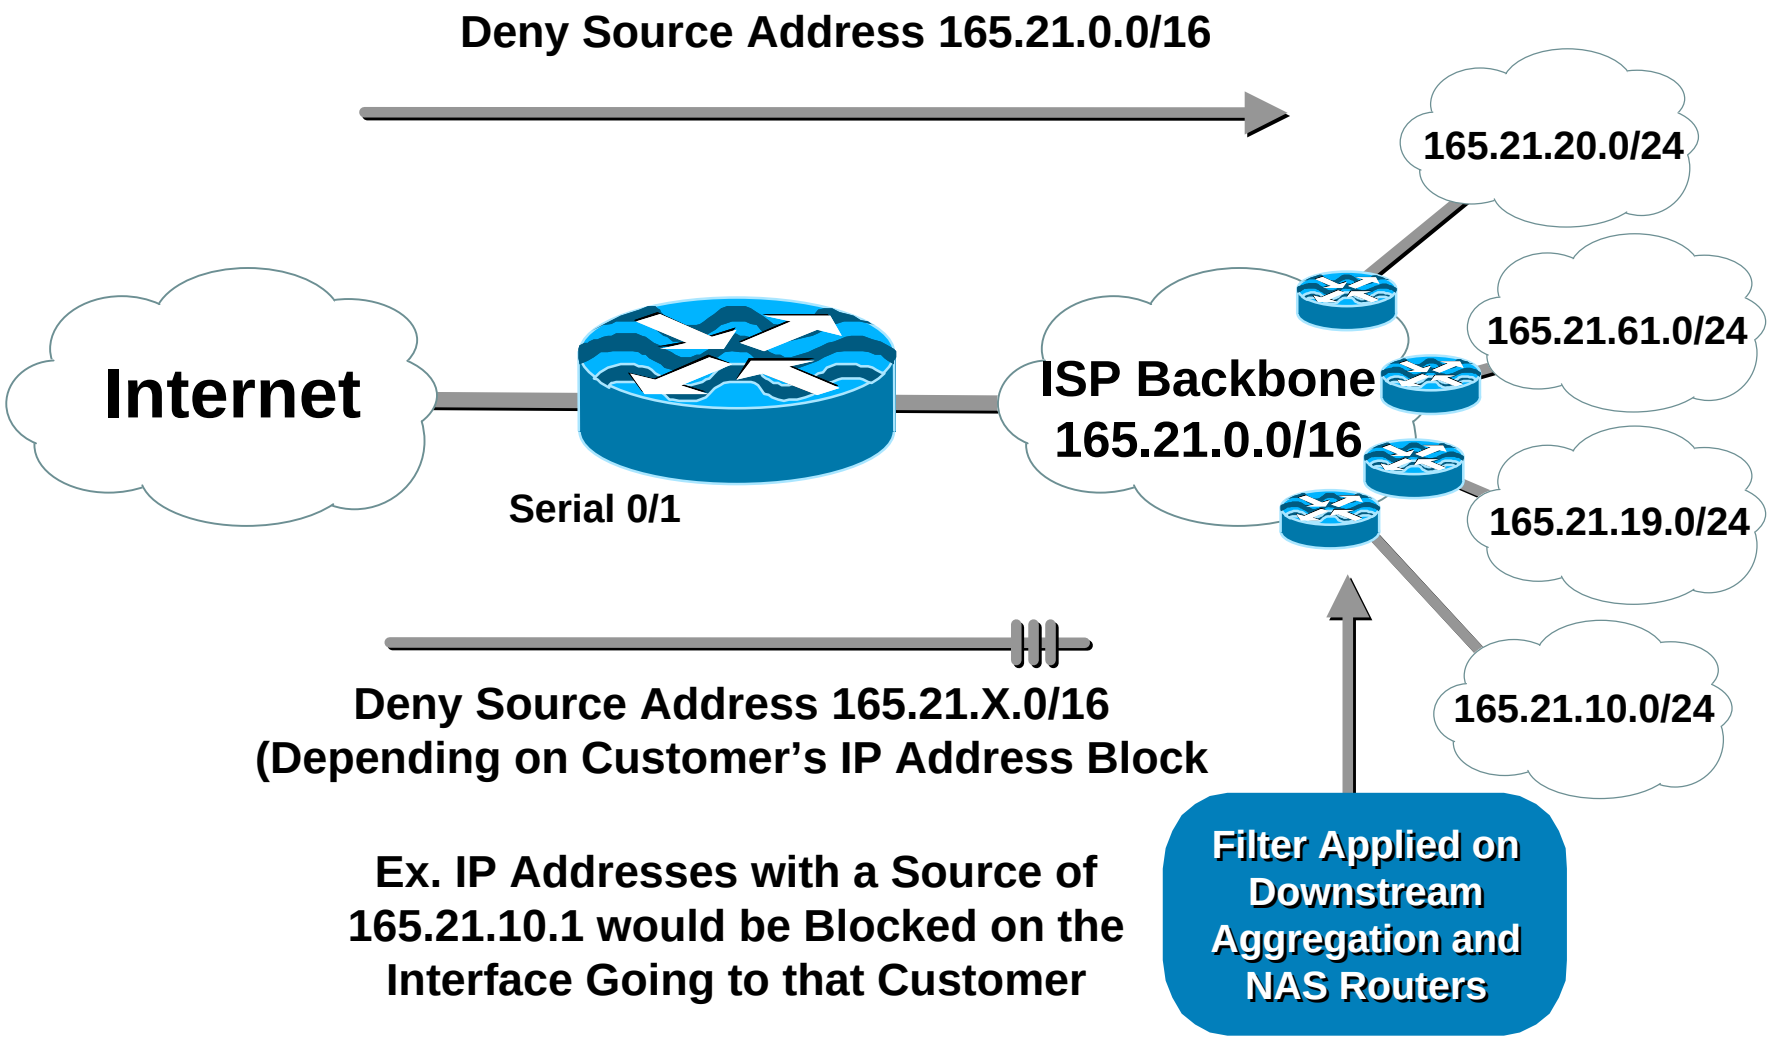

Dynamic ACLs with AAA Virtual Profiles

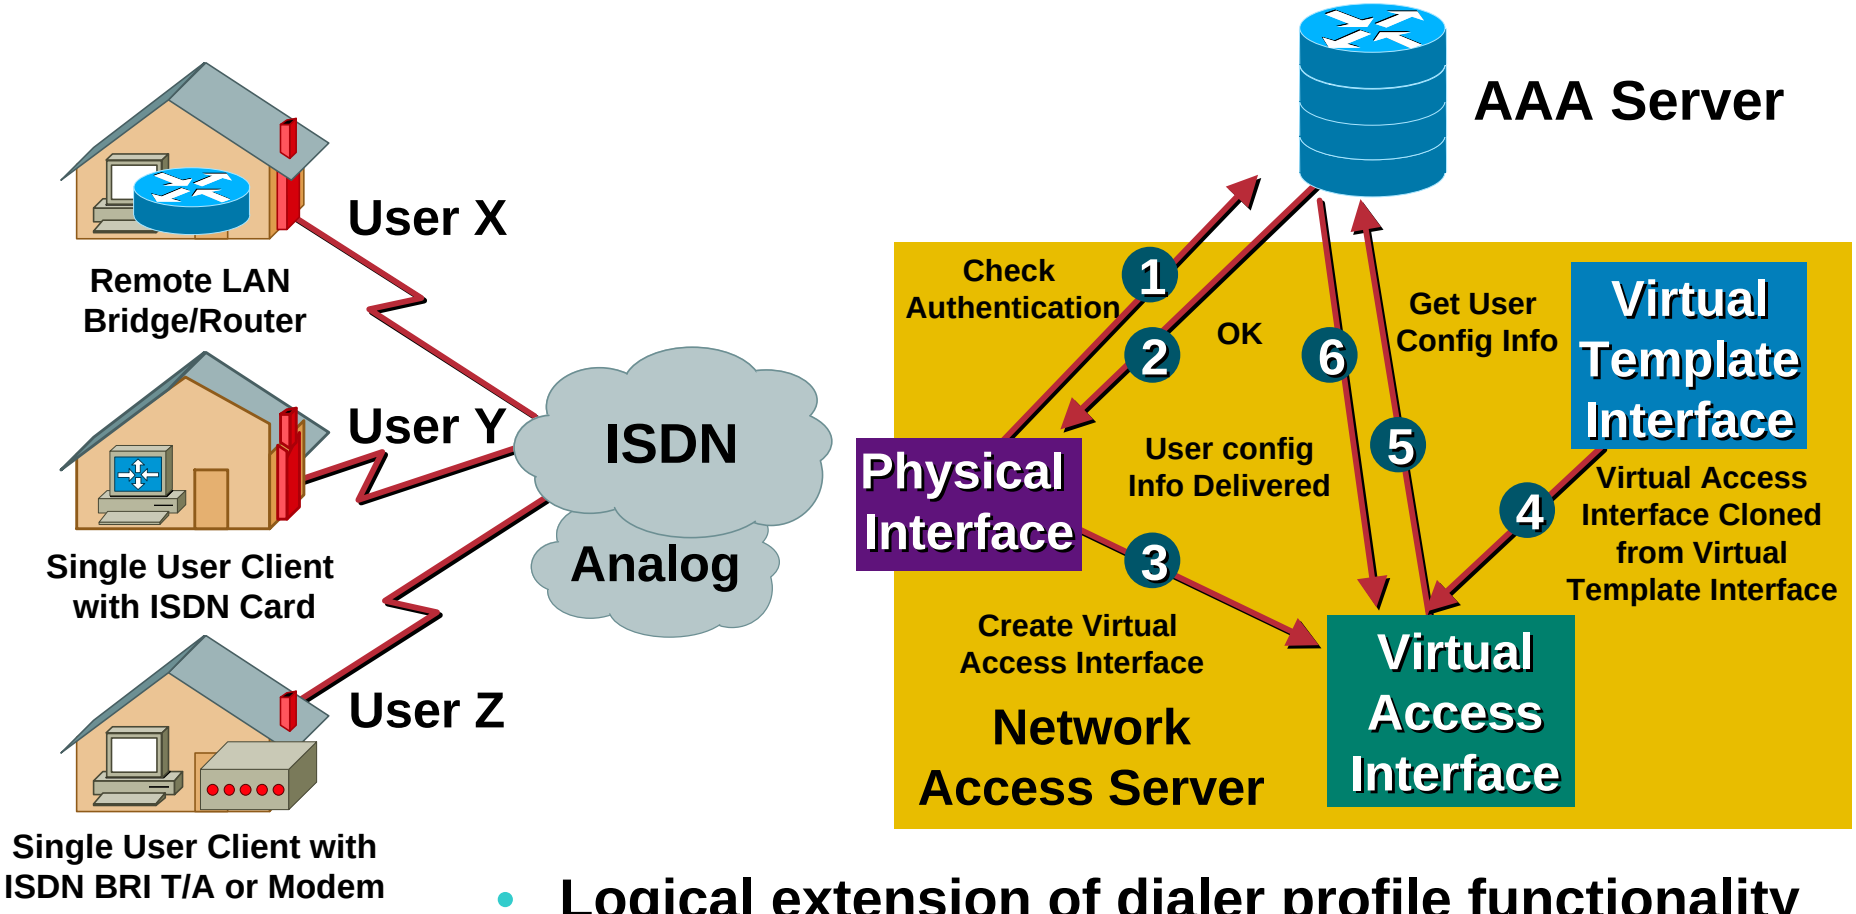

- Logical extension of dialer profile functionality
- ACLs stored in the Central AAA server
- Supports both Radius and Tacacs+

Dynamic ACLs with AAA Virtual Profiles

Cisco.com

- **List of site with information on how to configure radius to download ACLs:**

  Cisco Radius

  http://www.cisco.com/warp/public/480/radius_ACL1.html#secondary

  Ascend/Radius

  http://www.hal-c.org/~ascend/MaxTNT/radius/attrib.htm#216191

  TACACS+

  http://www.cisco.com/warp/public/480/tacacs_ACL1.html

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
- New! Two Flavors of uRPF:

  Strict Mode for BCP 38/ RFC 2827 Filters on Customer Ingress Edge

  Loose Mode for ISP ⇔ ISP Edge

CEF Unicast RPF (Strict Mode)

Cisco.com

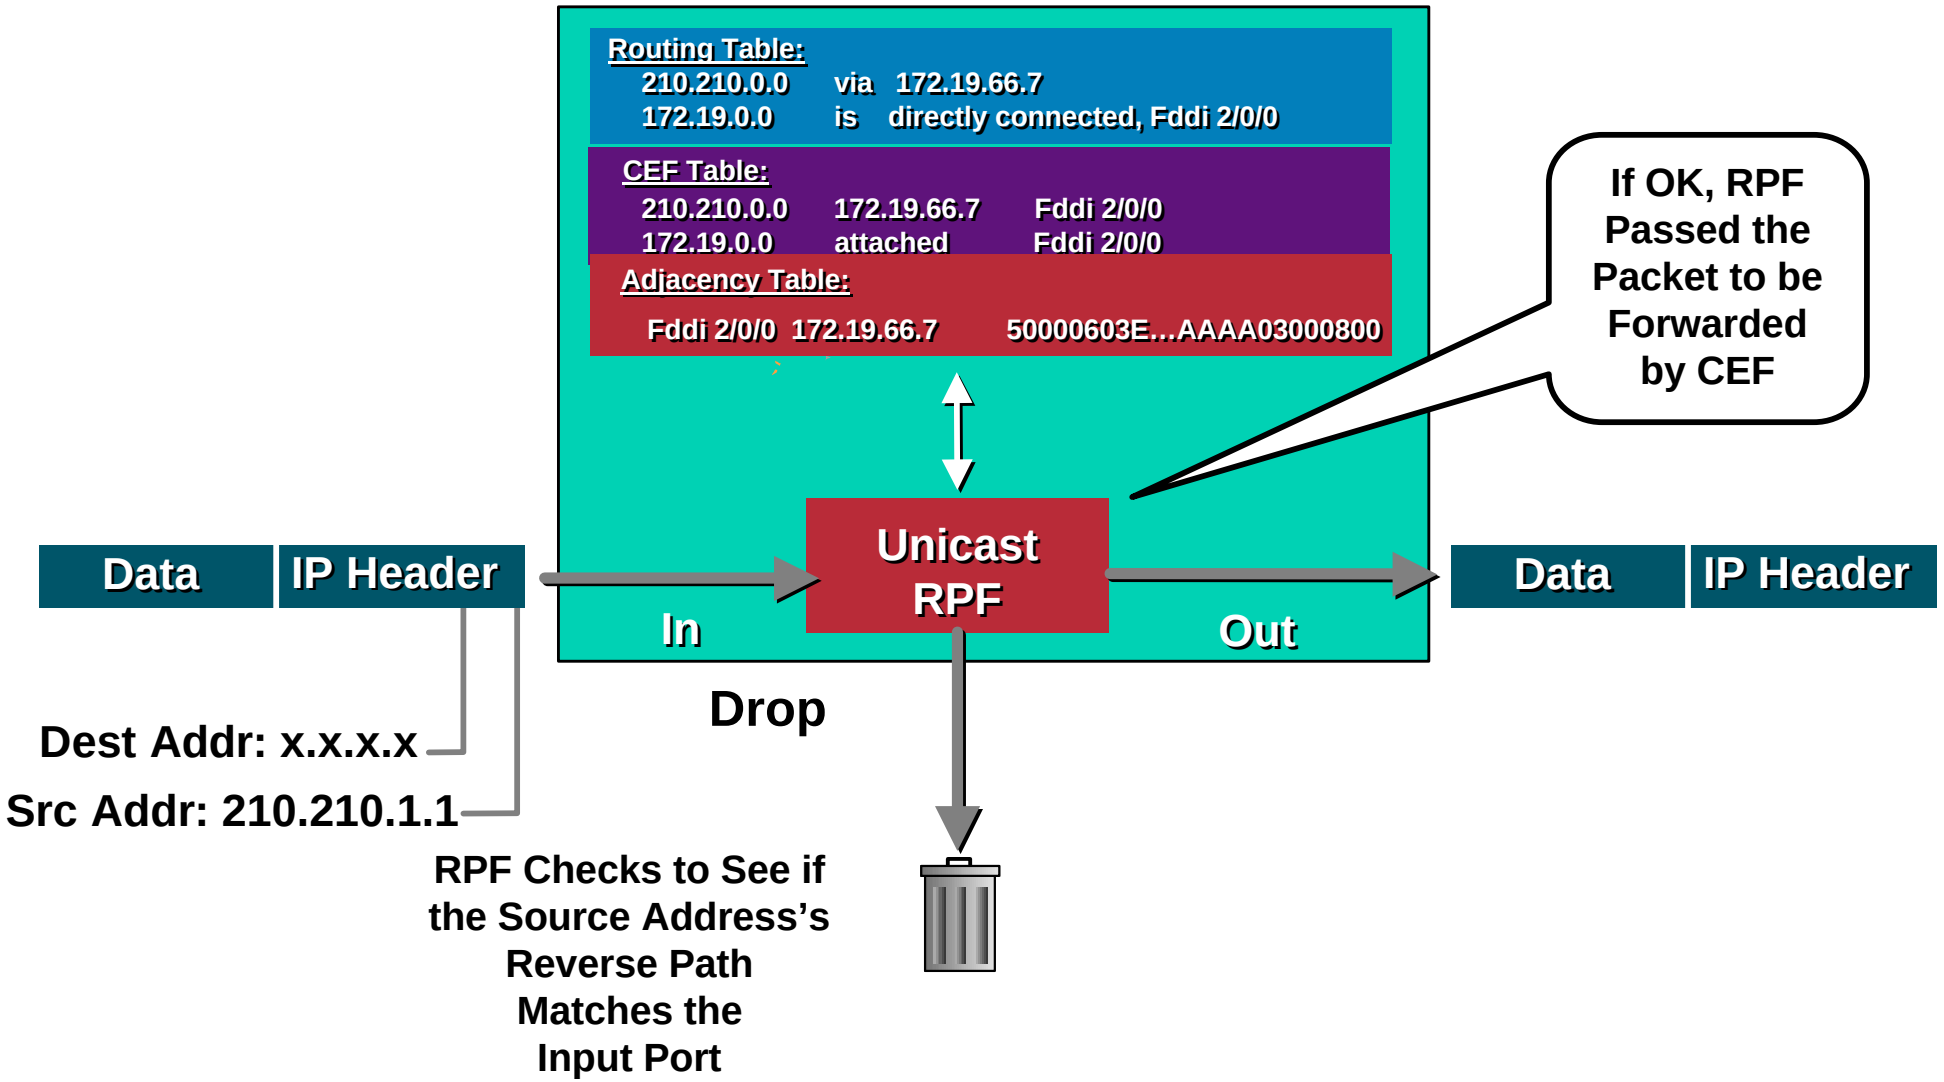

CEF Unicast RPF (Strict Mode)

Cisco.com

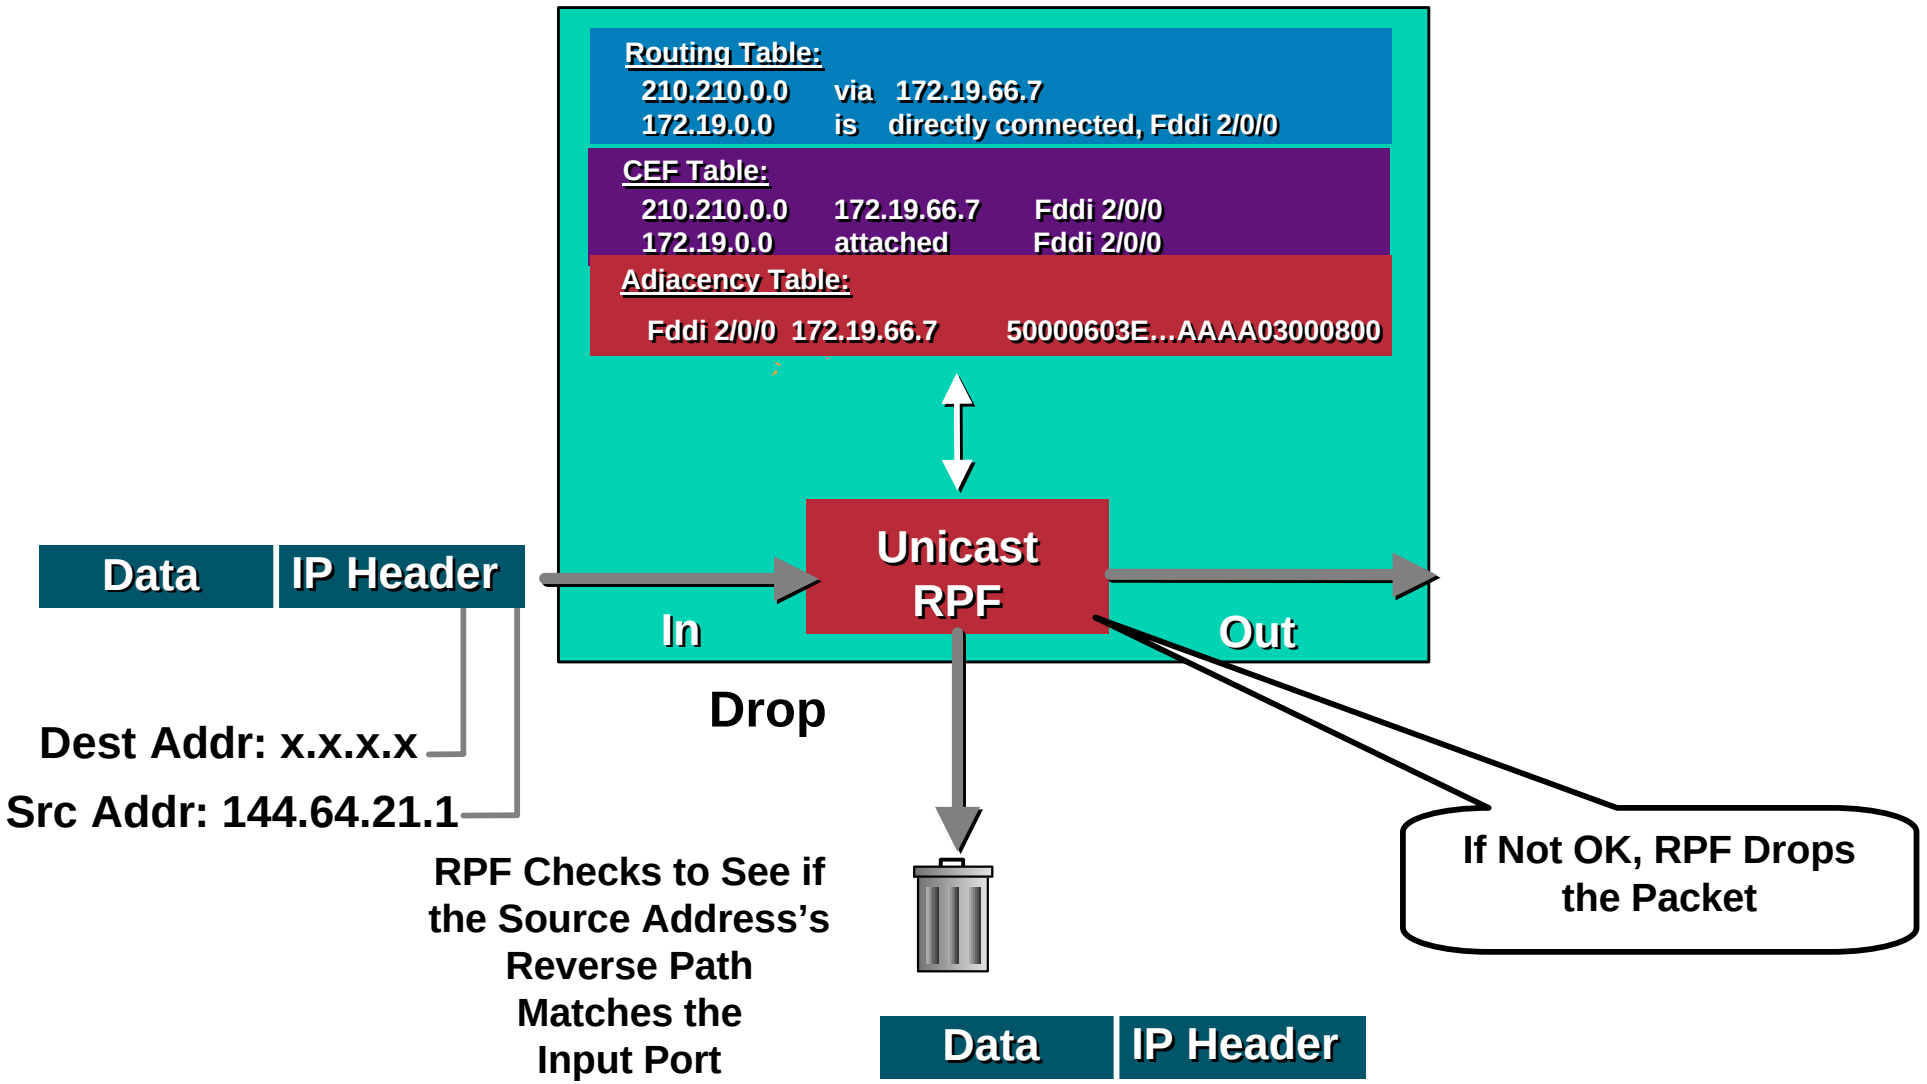

uRPF Originally Designed for the Customer ® ISP Edge

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

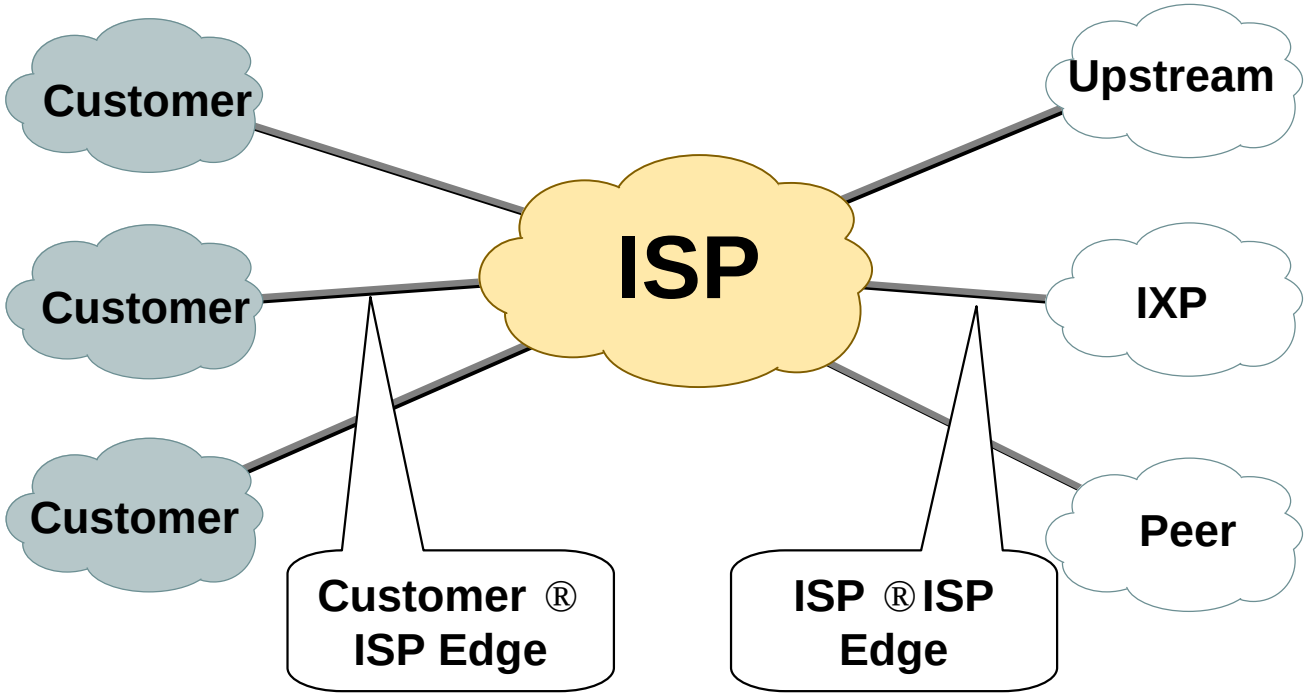

Where to Apply Unicast RPF (Strict Mode)?

Cisco.com

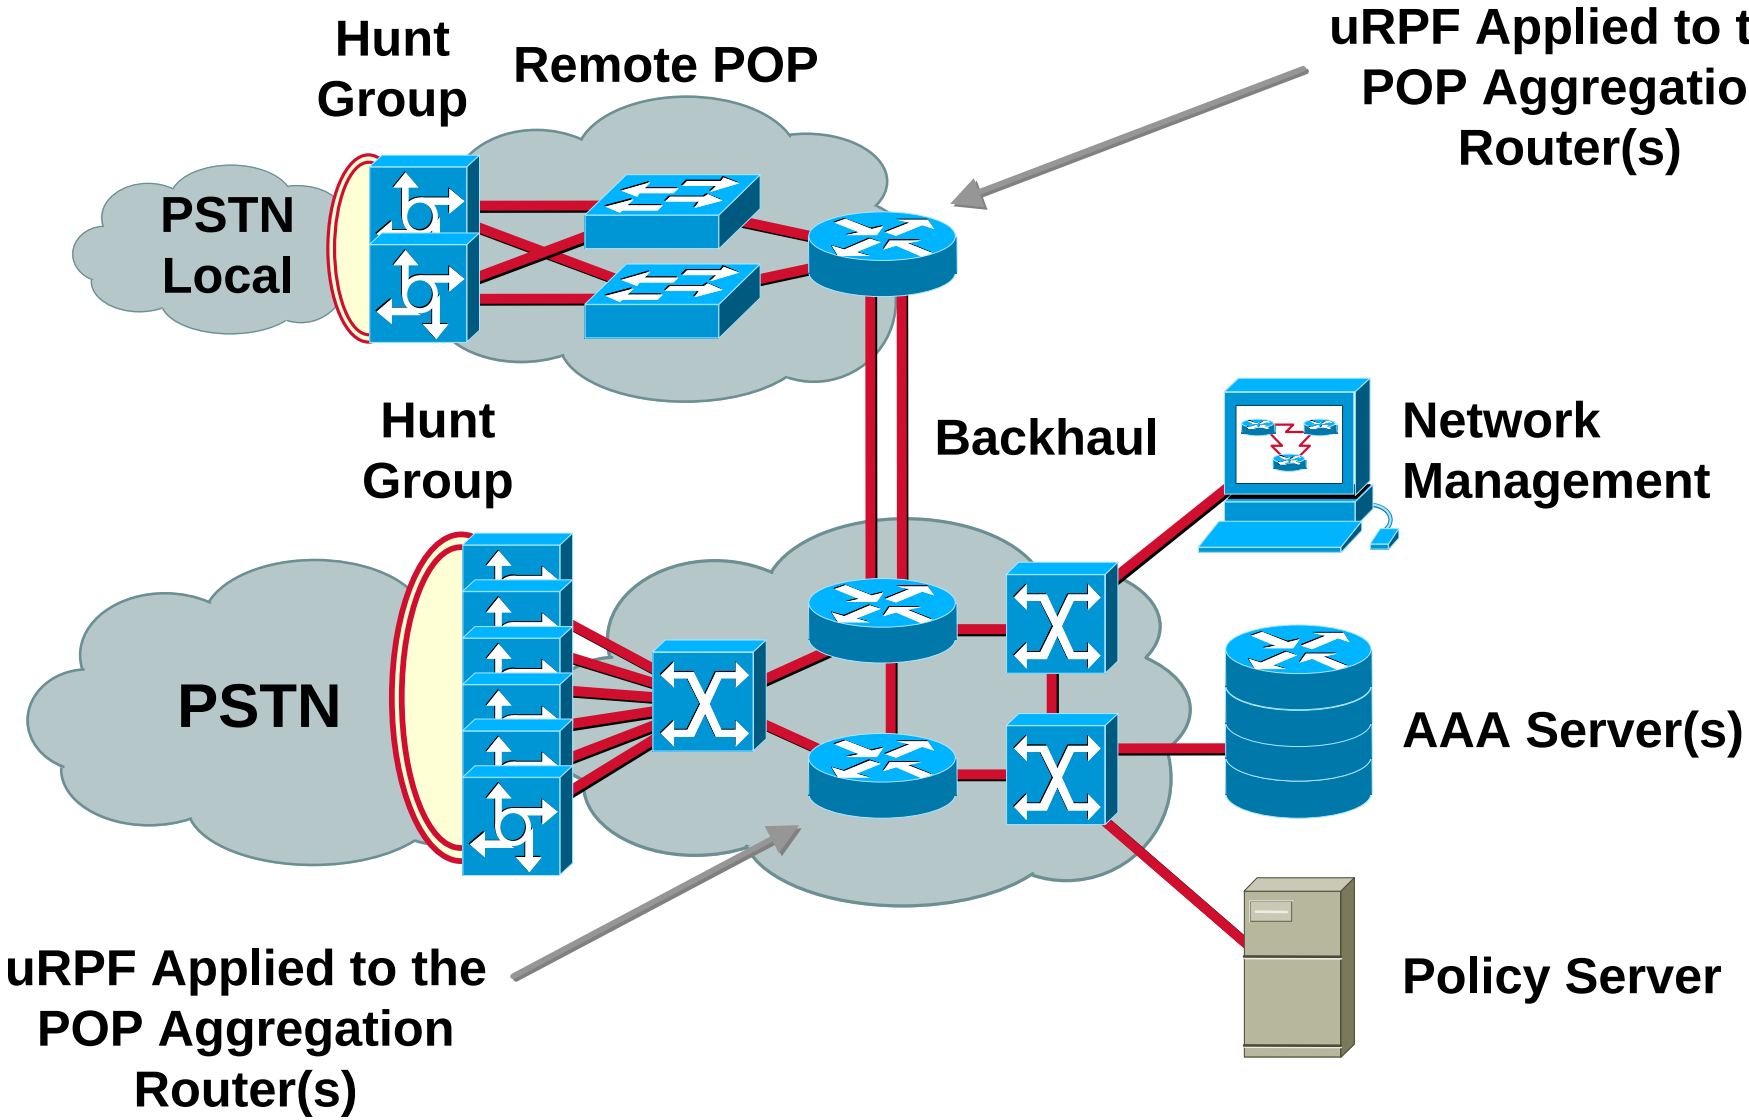

Unicast RPF Commands (Strict Mode)

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)

- **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—Simple Single Homed Customer Example


```
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
```

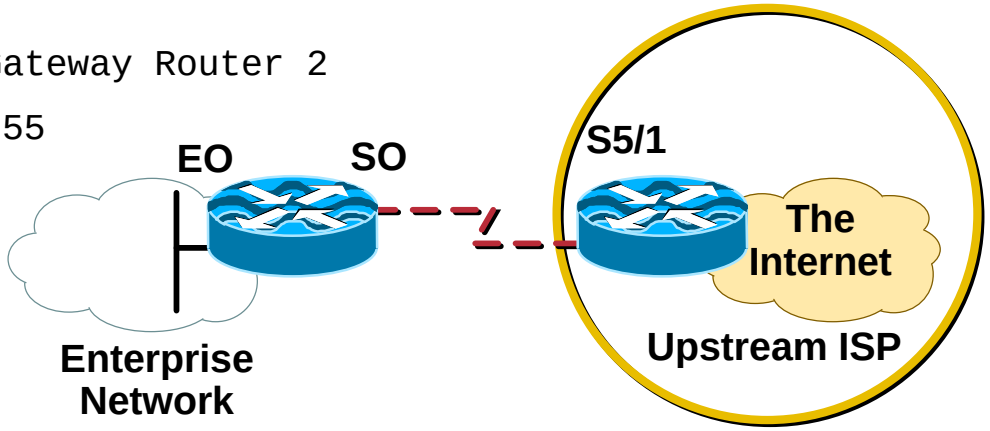

Unicast RPF—Simple Single Homed Customer Example


```
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
```

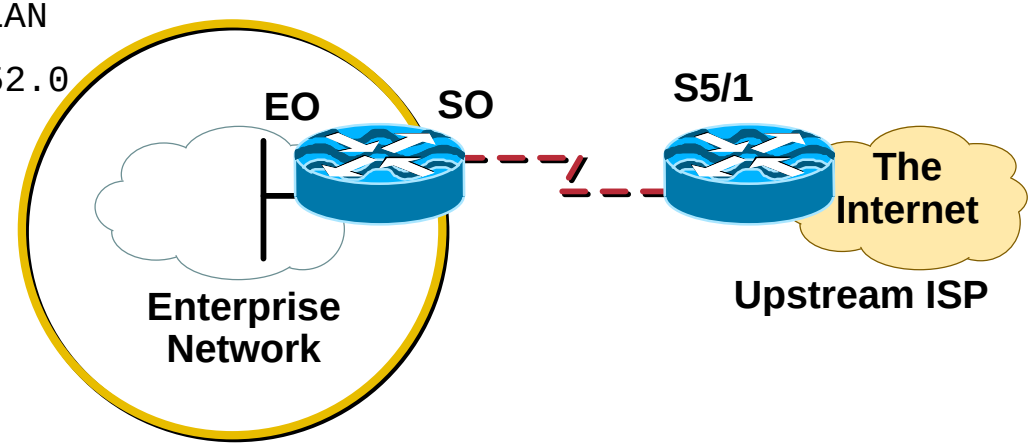

CEF Unicast RPF (Strict Mode)

- **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 Not Widely Deployed?

- 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?

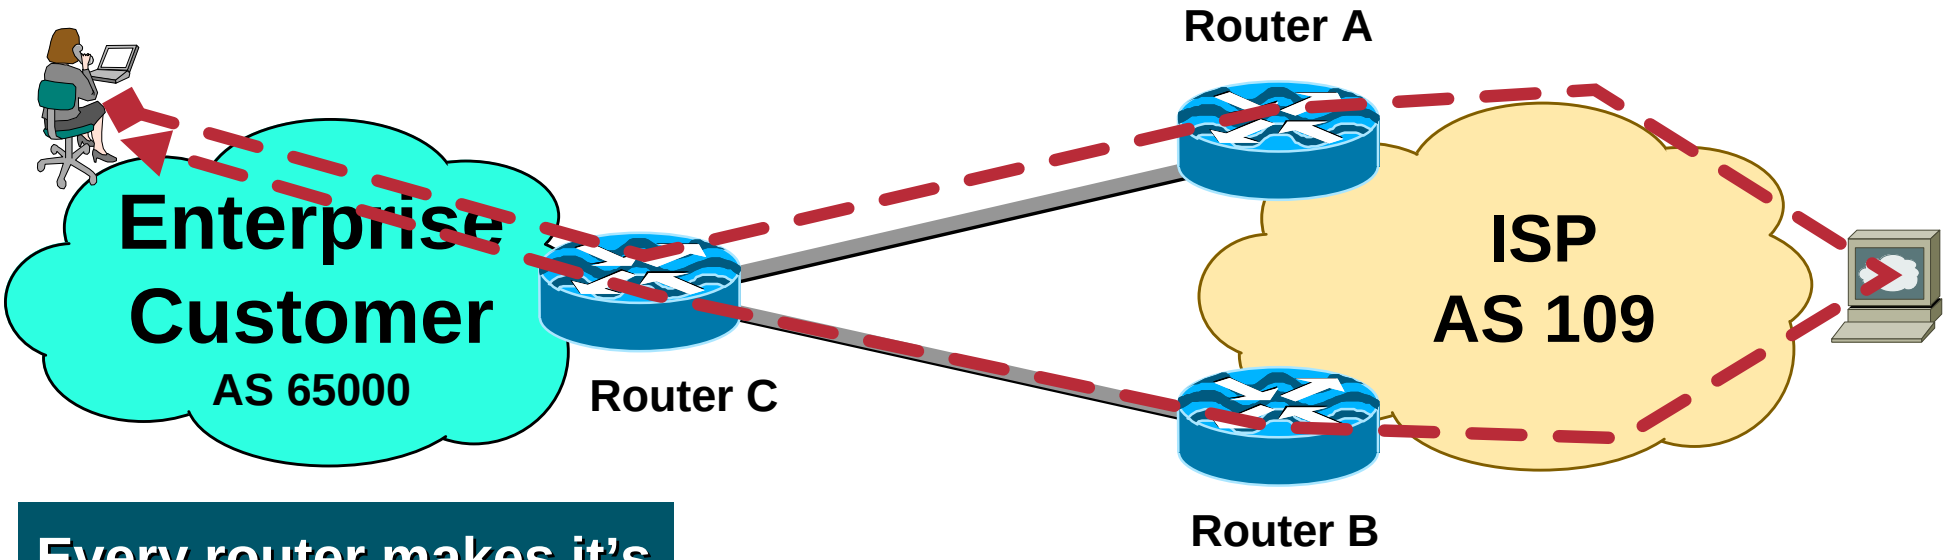

Every router makes it's own best path forwarding decision – resulting in asymmetrical routing

Best Path Routing in the Internet

Cisco.com

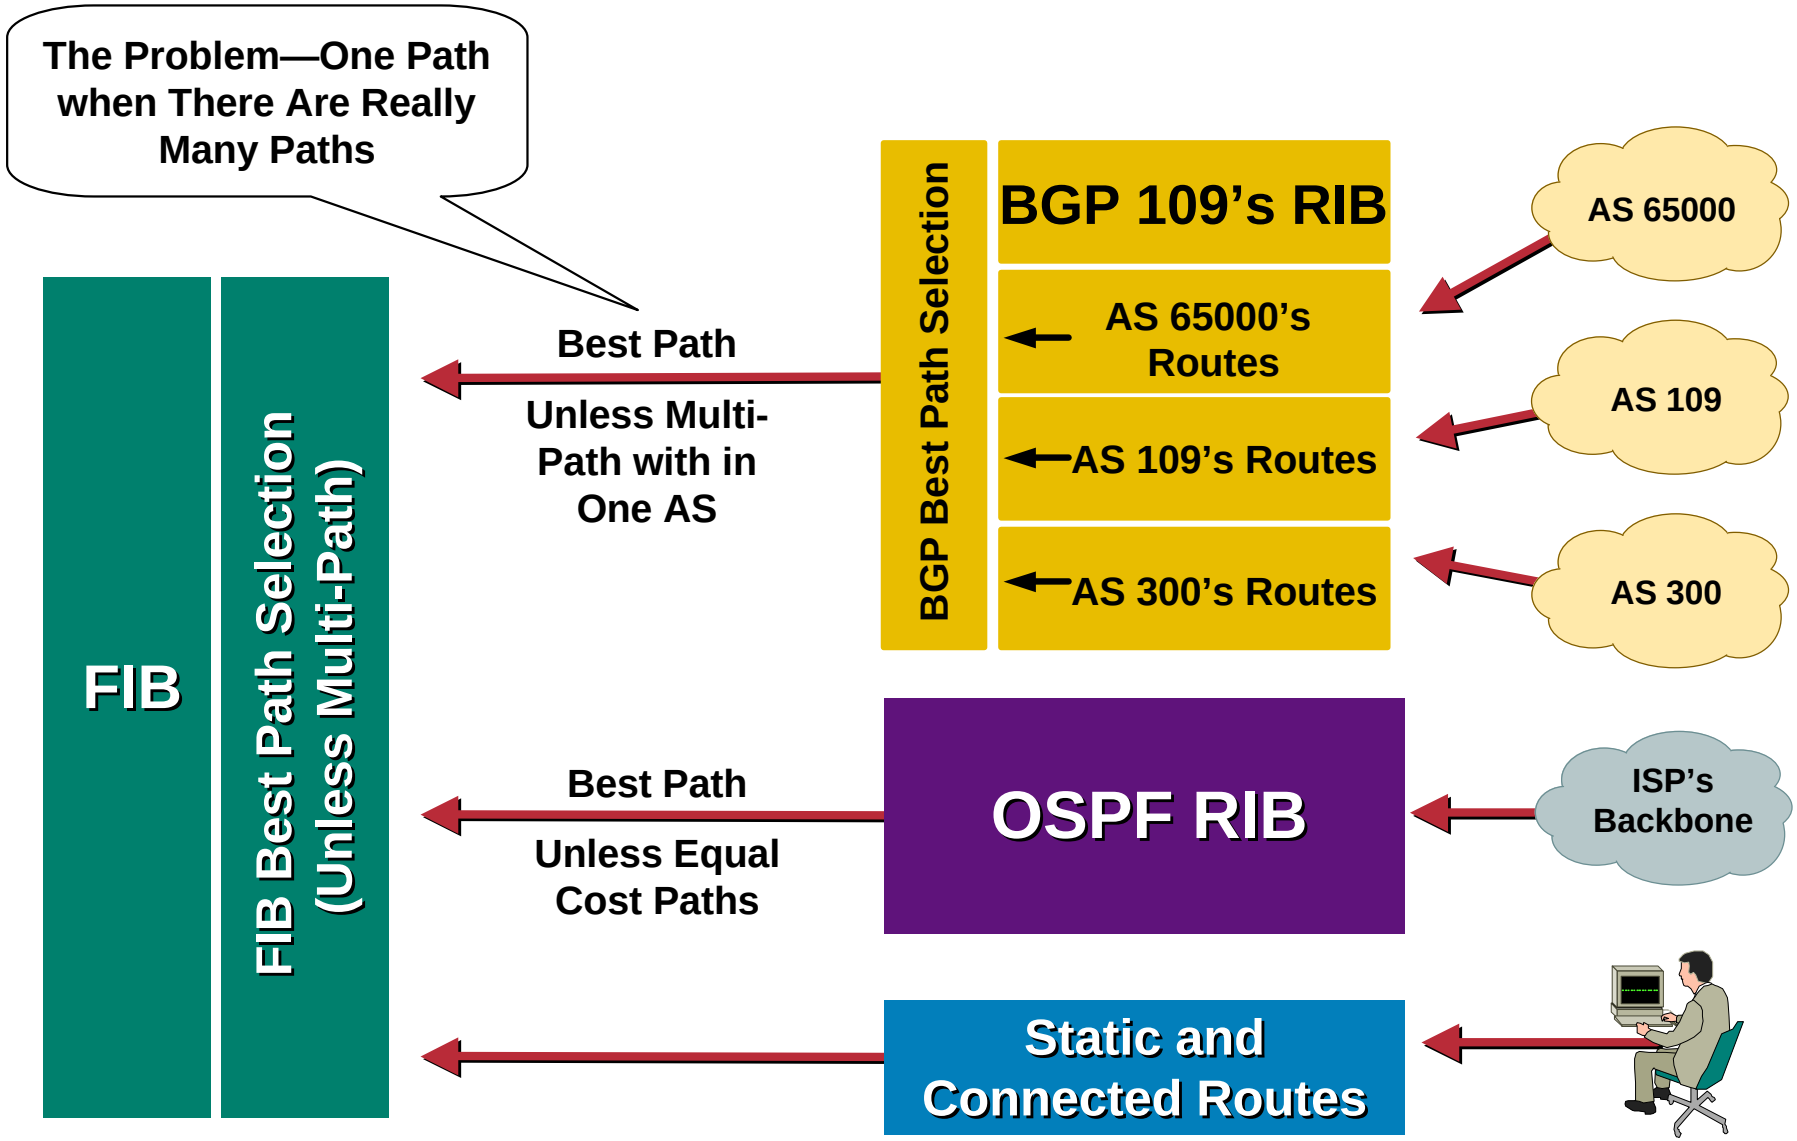

BGP Weight aligns the FIB for uRPF

Cisco.com

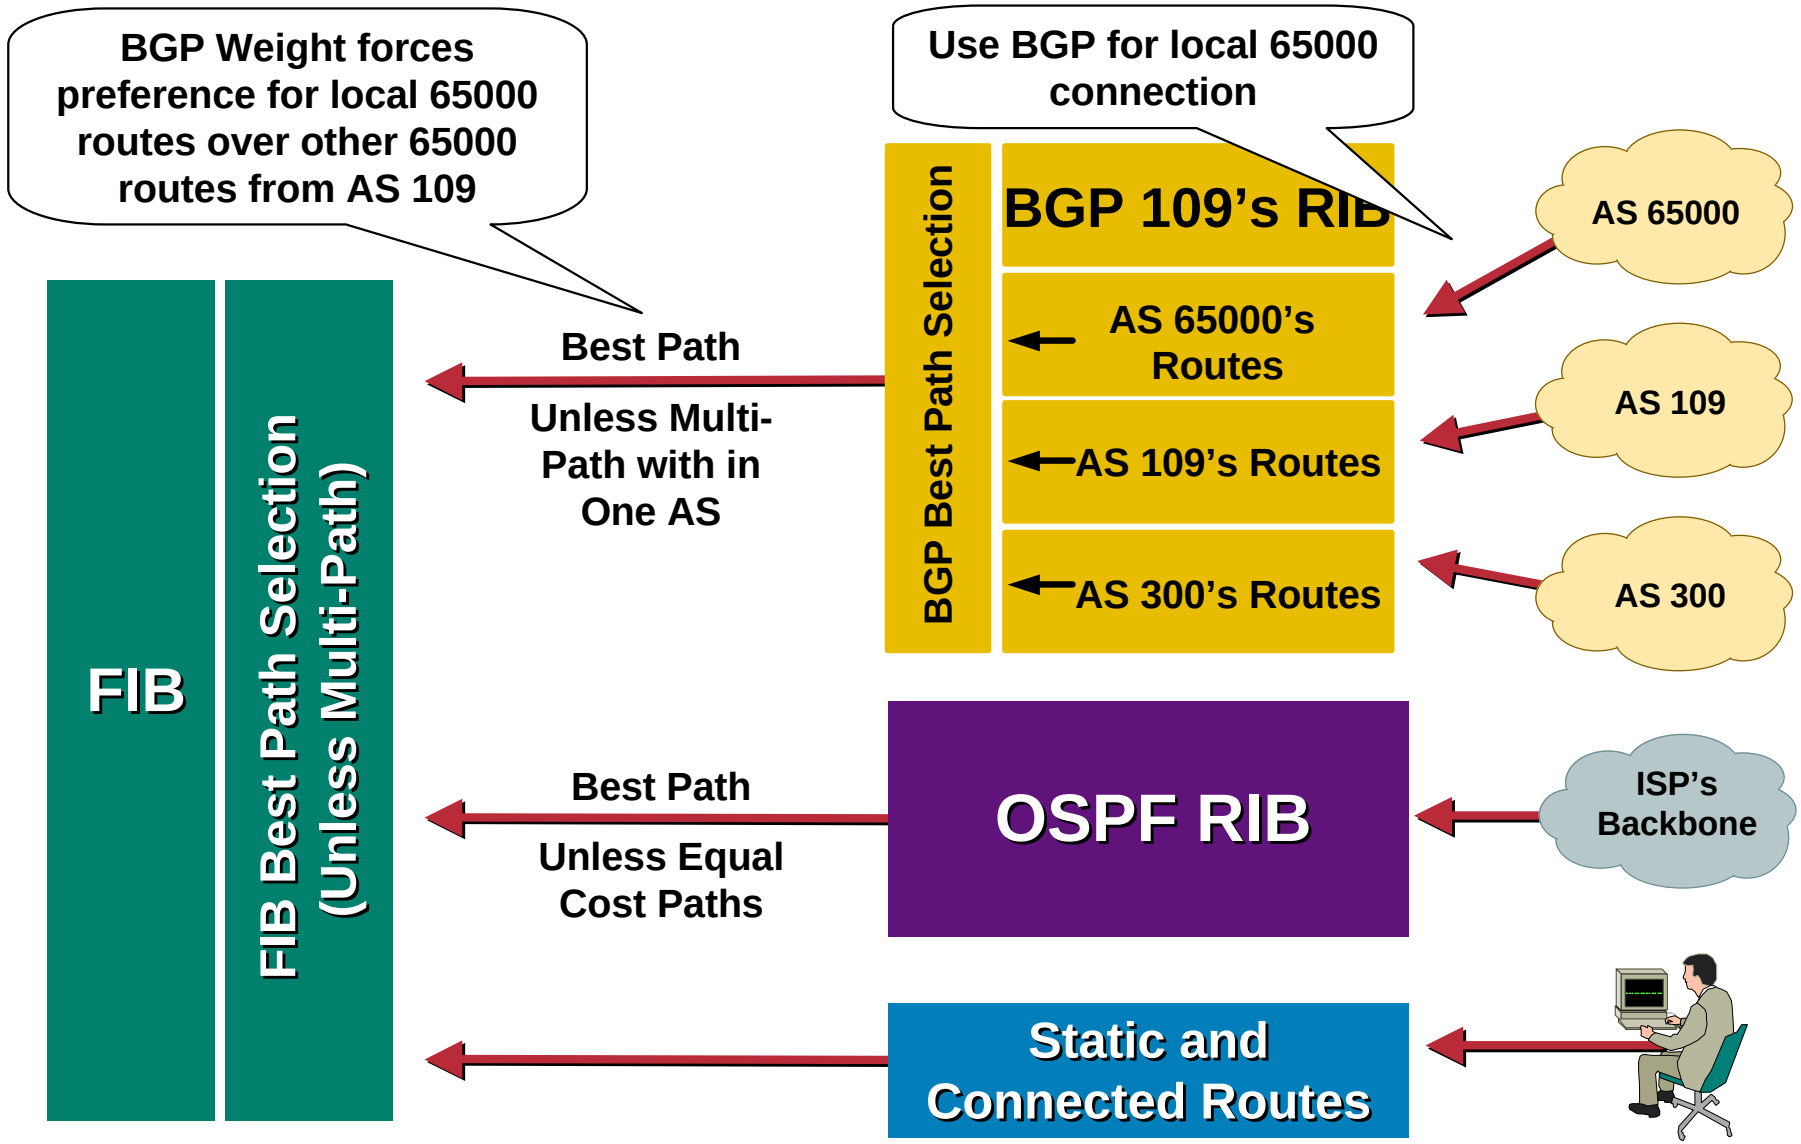

Unicast RPF — Dual Homed Customer

Cisco.com

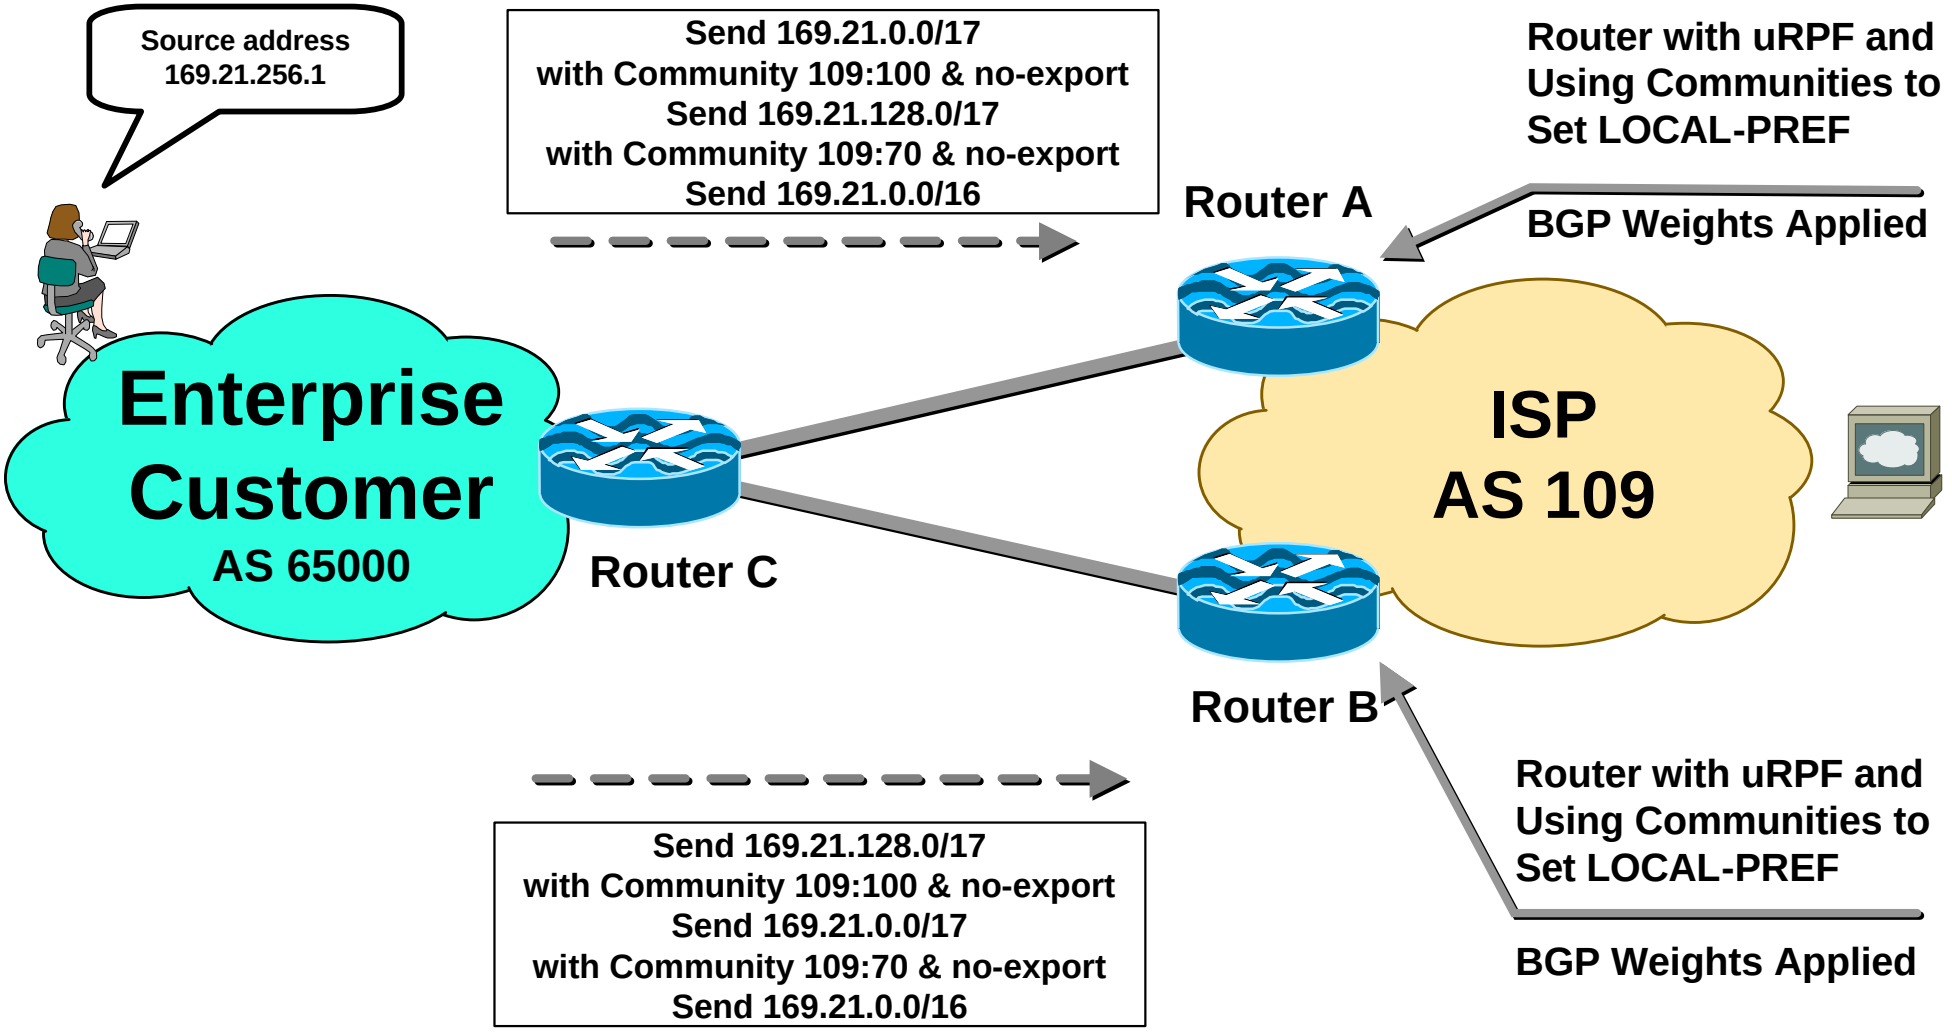

Unicast RPF — Before BGP Weight

Cisco.com

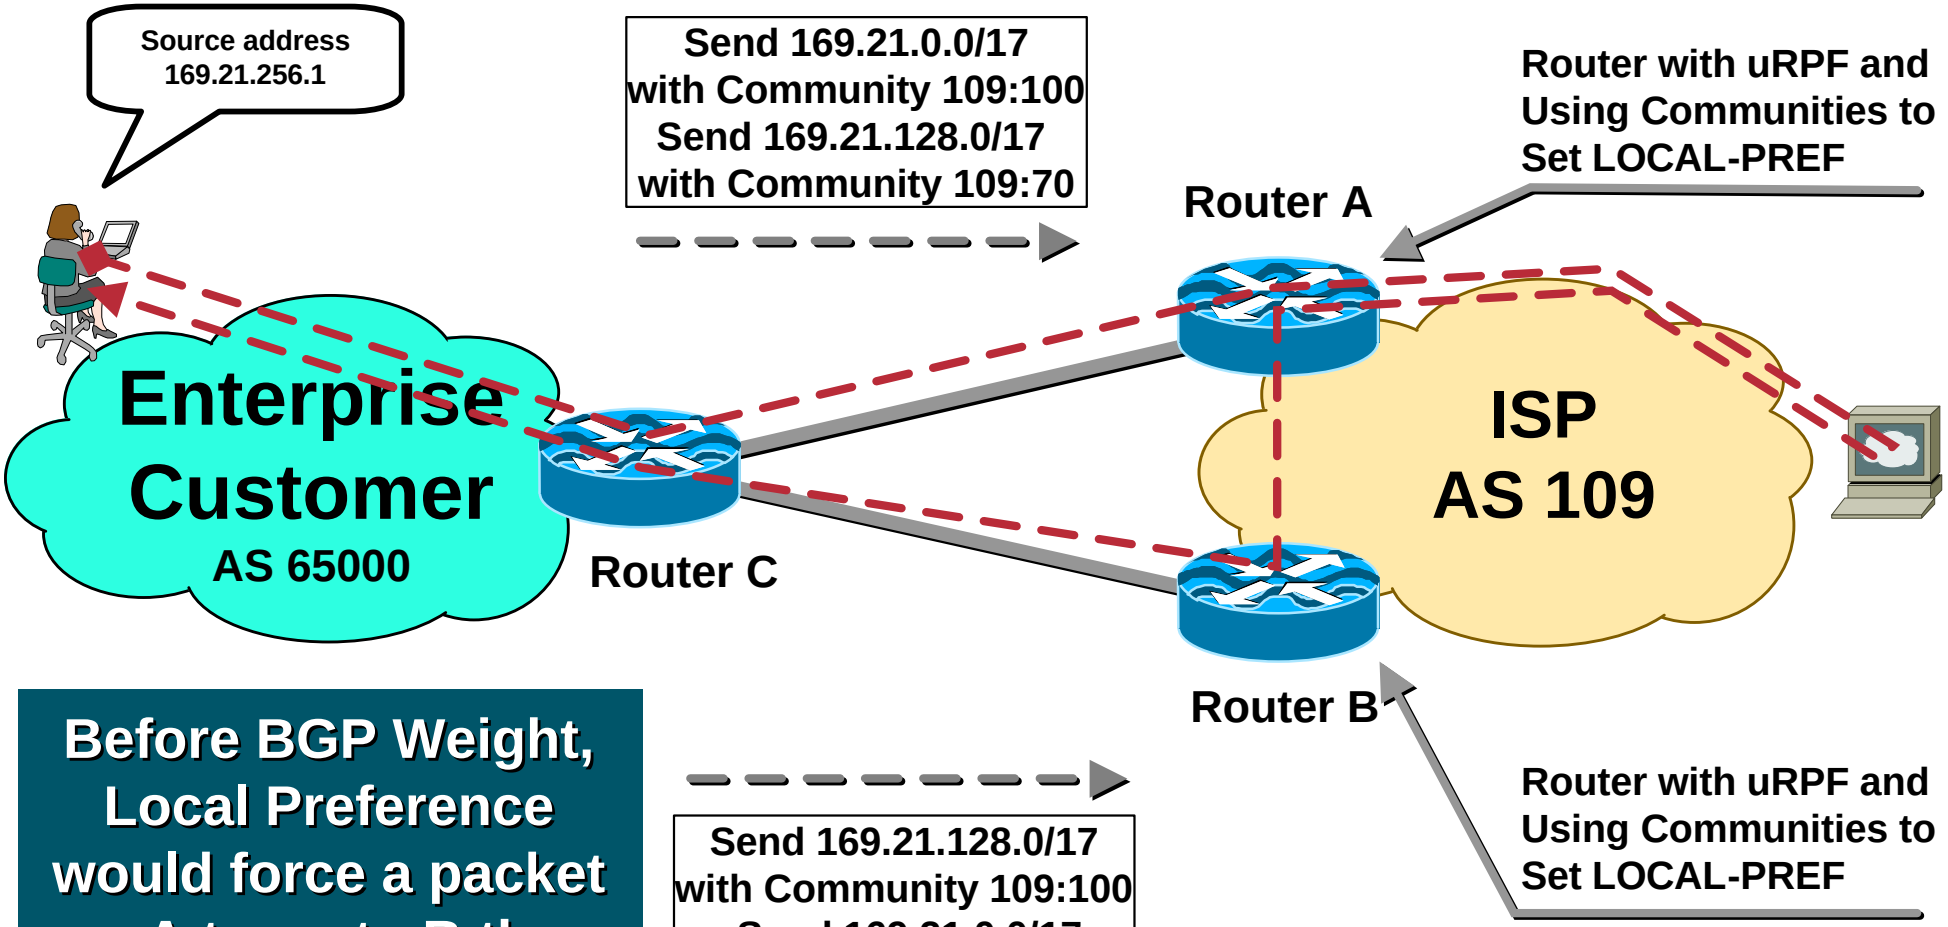

Before BGP Weight, Local Preference would force a packet on A to go to B then C.

Unicast RPF — After BGP Weight

Cisco.com

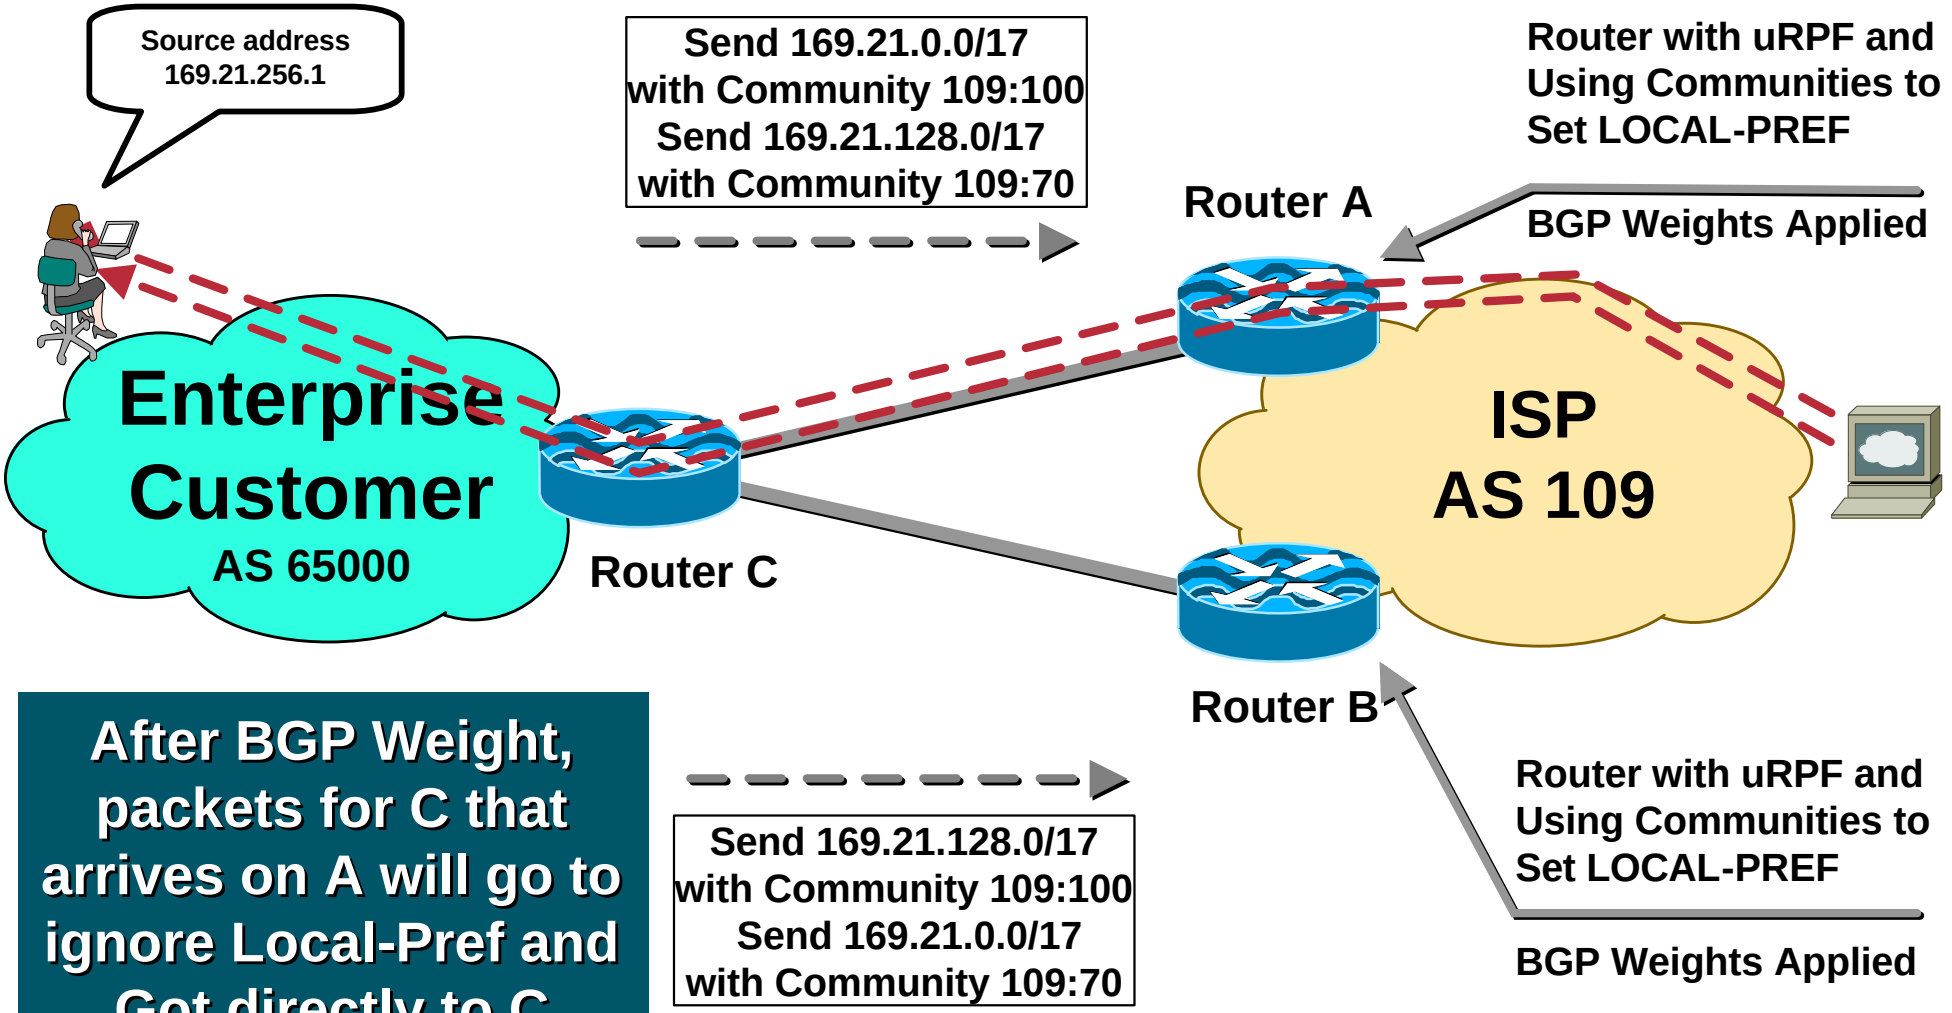

After BGP Weight, packets for C that arrives on A will go to ignore Local-Pref and Got directly to C

Unicast RPF — 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 — 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 — Dual Homed Enterprise to One ISP

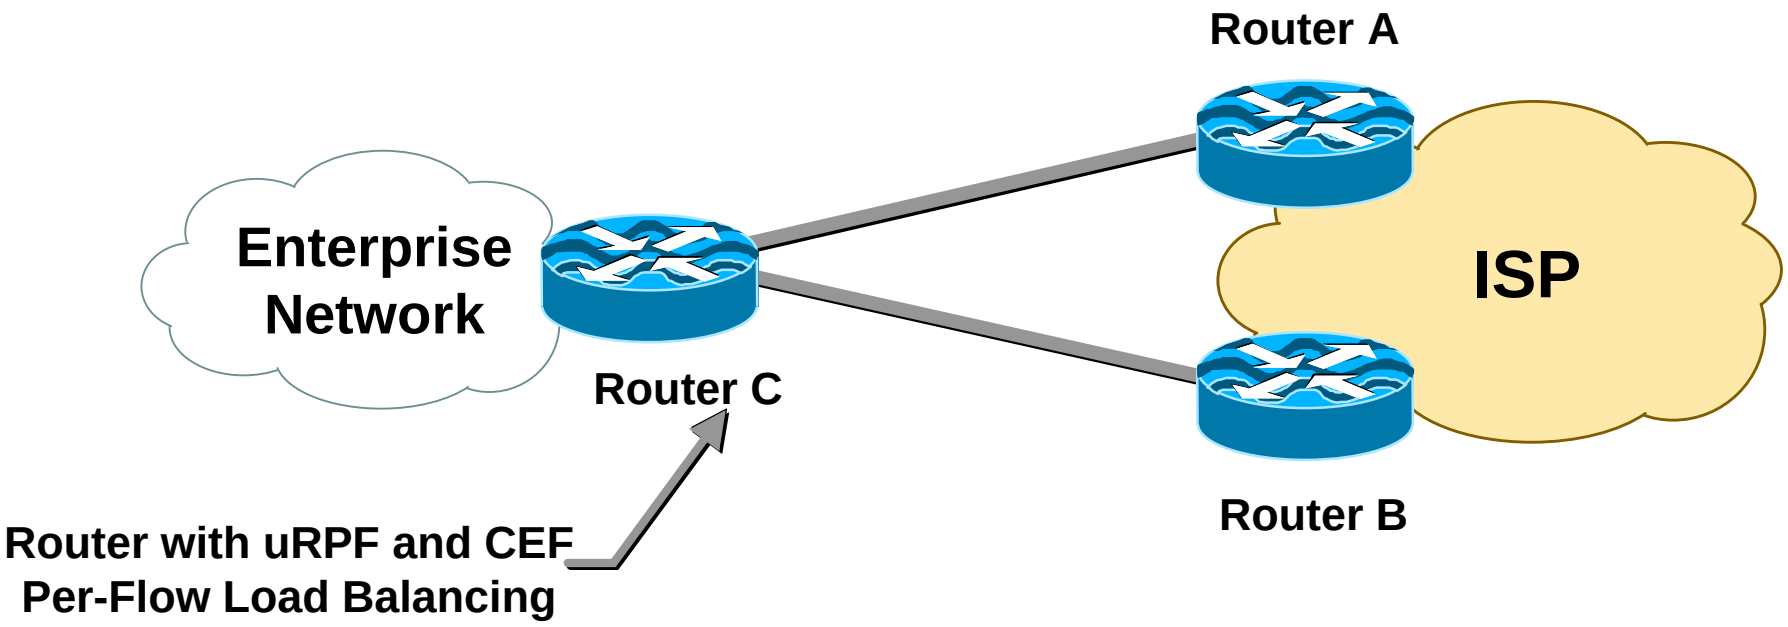

- Used to protect against spoof attacks
- Some attacks get around the RFC1918 filters by using un-allocated IP address space

Unicast RPF — 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 — Dual Homed Enterprise to One ISP


```
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 — Dual Homed Enterprise to One ISP

- **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 — Dual Homed Enterprise to Two ISPs

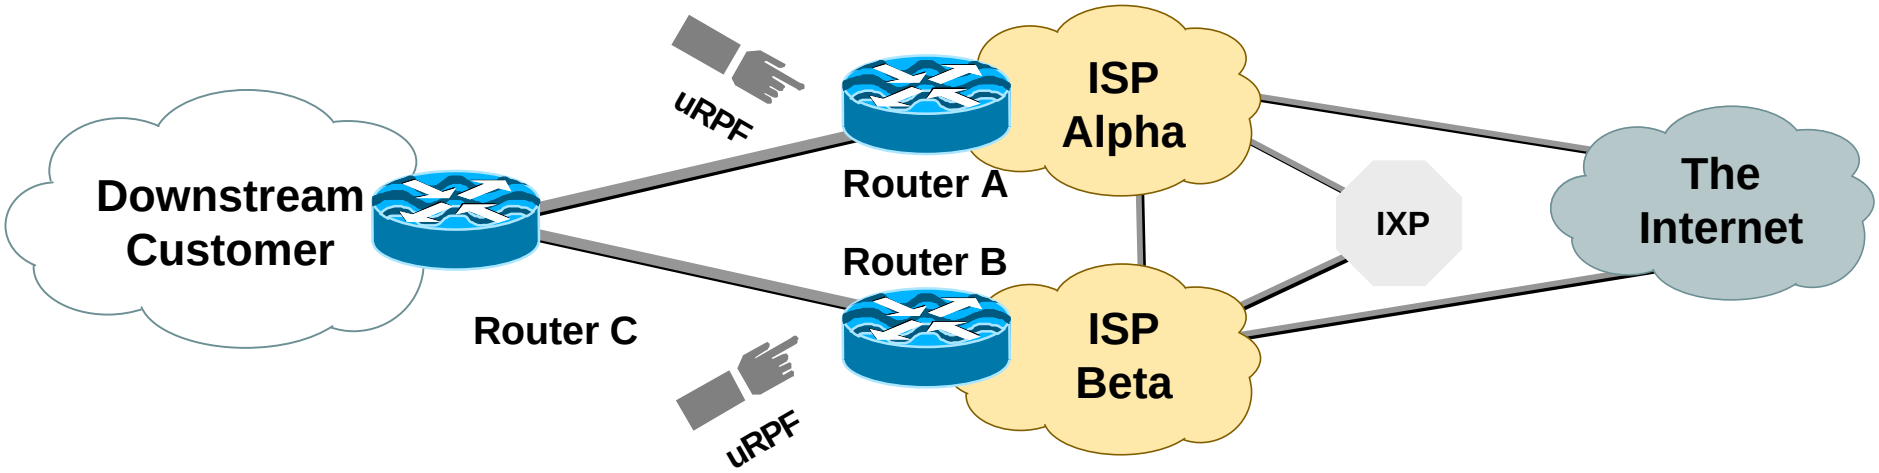

- **ISP Configuration for both ISPs are similar to a dual homed customer.**

  BGP weight is used to over ride AS path prepends

Unicast RPF — Dual Homed Enterprise to Two ISPs

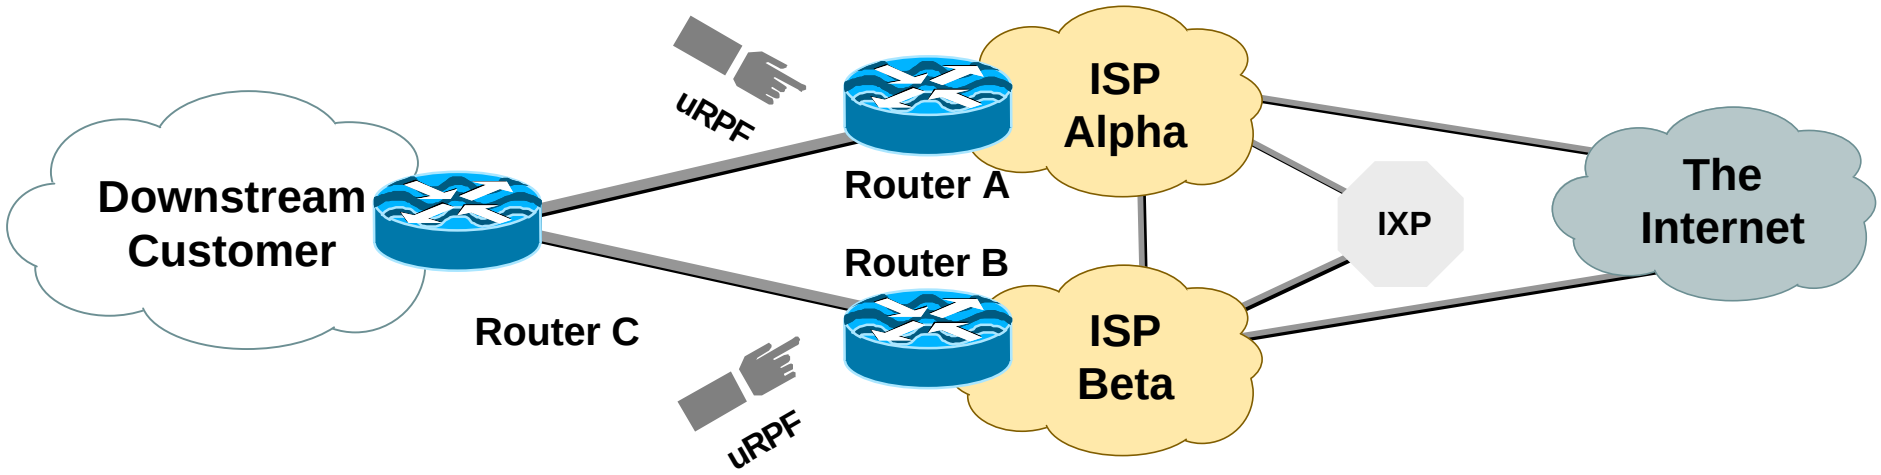

- **BGP weight override an AS path prepend**
  - **BGP weight on Router A will keep the preferred path for packets on that router to be C « A**
  - **BGP weight on Router B will keep the preferred path for packets on that router to be C « B**

Unicast RPF — Dual Homed Enterprise to Two ISPs

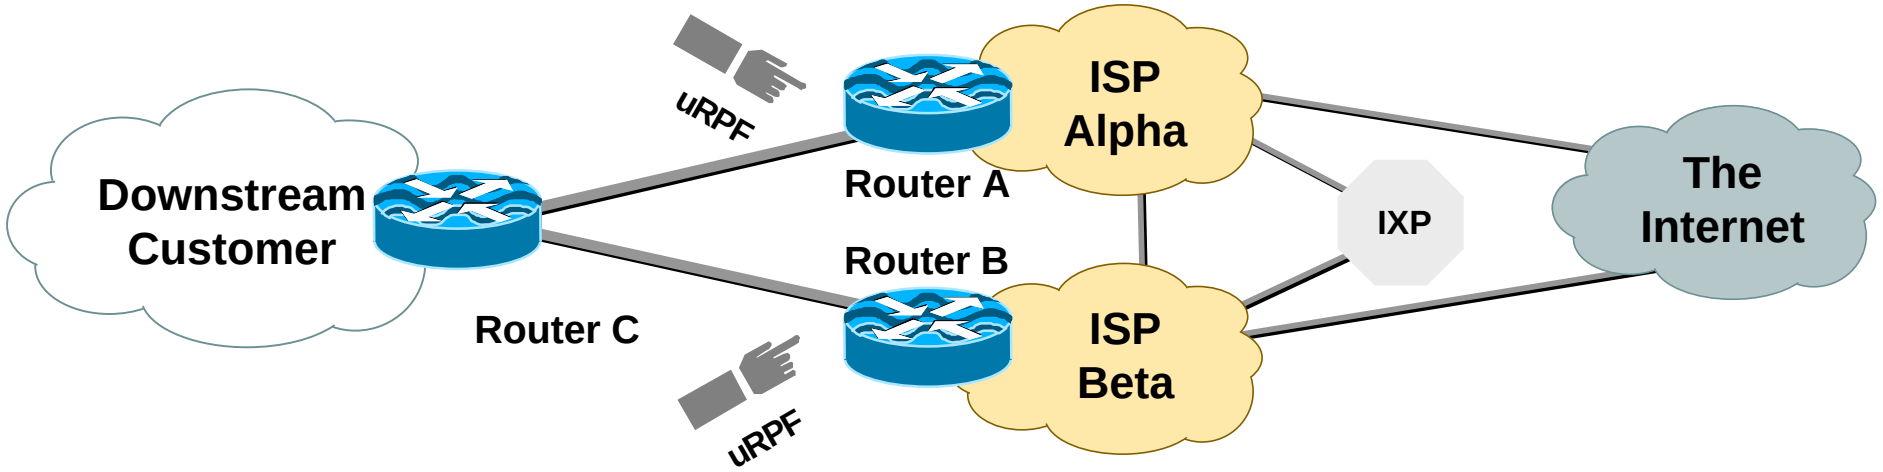

- **Enterprise configuration cannot use maximum-paths**

  Need equal AS paths for maximum-paths to work

Unicast RPF — 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 — 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 — The ACL Bypass Option

Cisco.com

- **Cisco 7500 with a classification filter:**

```
interface ethernet 0/1/1
 ip address 192.168.200.1 255.255.255.0
 ip verify unicast reverse-path 171
!
access-list 171 deny icmp any any echo log-input
access-list 171 deny icmp any any echo-reply log-input
access-list 171 deny udp any any eq echo log-input
access-list 171 deny udp any eq echo any log-input
access-list 171 deny tcp any any established log-input
access-list 171 deny tcp any any log-input
access-list 171 deny ip any any log-input
```

Unicast RPF — 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 — 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 — 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 — Bottom Line

Cisco.com

- Unicast RFP is another tool to help defend the Internet
- Unicast RPF works when it is deployed within it's operational envelop
- Unicast RPF does not work when just thrown into the network; give it some thought

New Unicast RPF Enhancements

- **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

- **Phase 2—Loose check (if exist)**
  - **DDTS CSCdr93424**
  - **12.0(14)S for 7200, 7500, and GSR Engine 0 and 1**
  - **Scheduled 12.0(19)S for GSR Engine 2**
  - **Scheduled 12.1(8)E for CAT6K**

New Unicast RPF Enhancements

Cisco.com

- **Objectives in phase 2:**
  - **Allow for uRPF to work on the ISP ⇔ISP edge of the network**
  - **Create a new tool to drop DOS/DDOS attacks on the edge of an ISP's network**
  - **All for the drop to be activated and controlled by a network protocol**

New Unicast RPF Enhancements

- **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>]
```

uRPF Originally Designed for the Customer ® ISP Edge

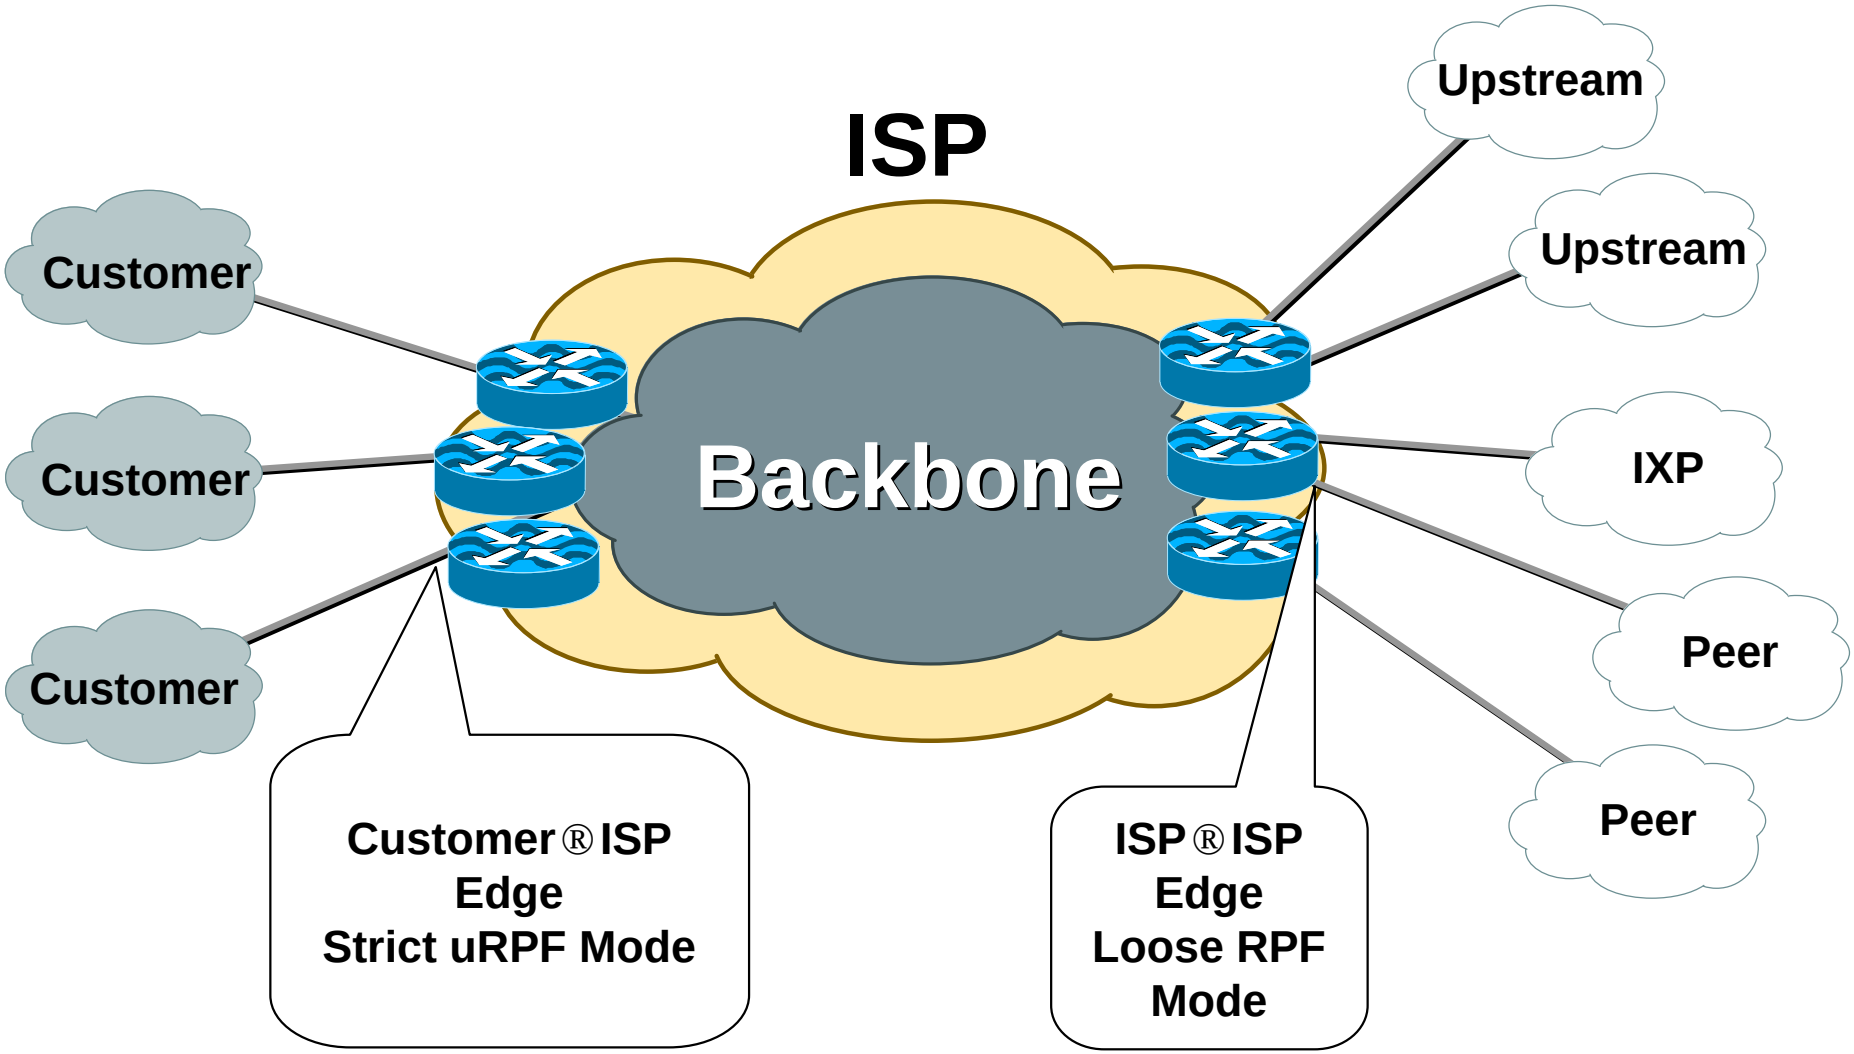

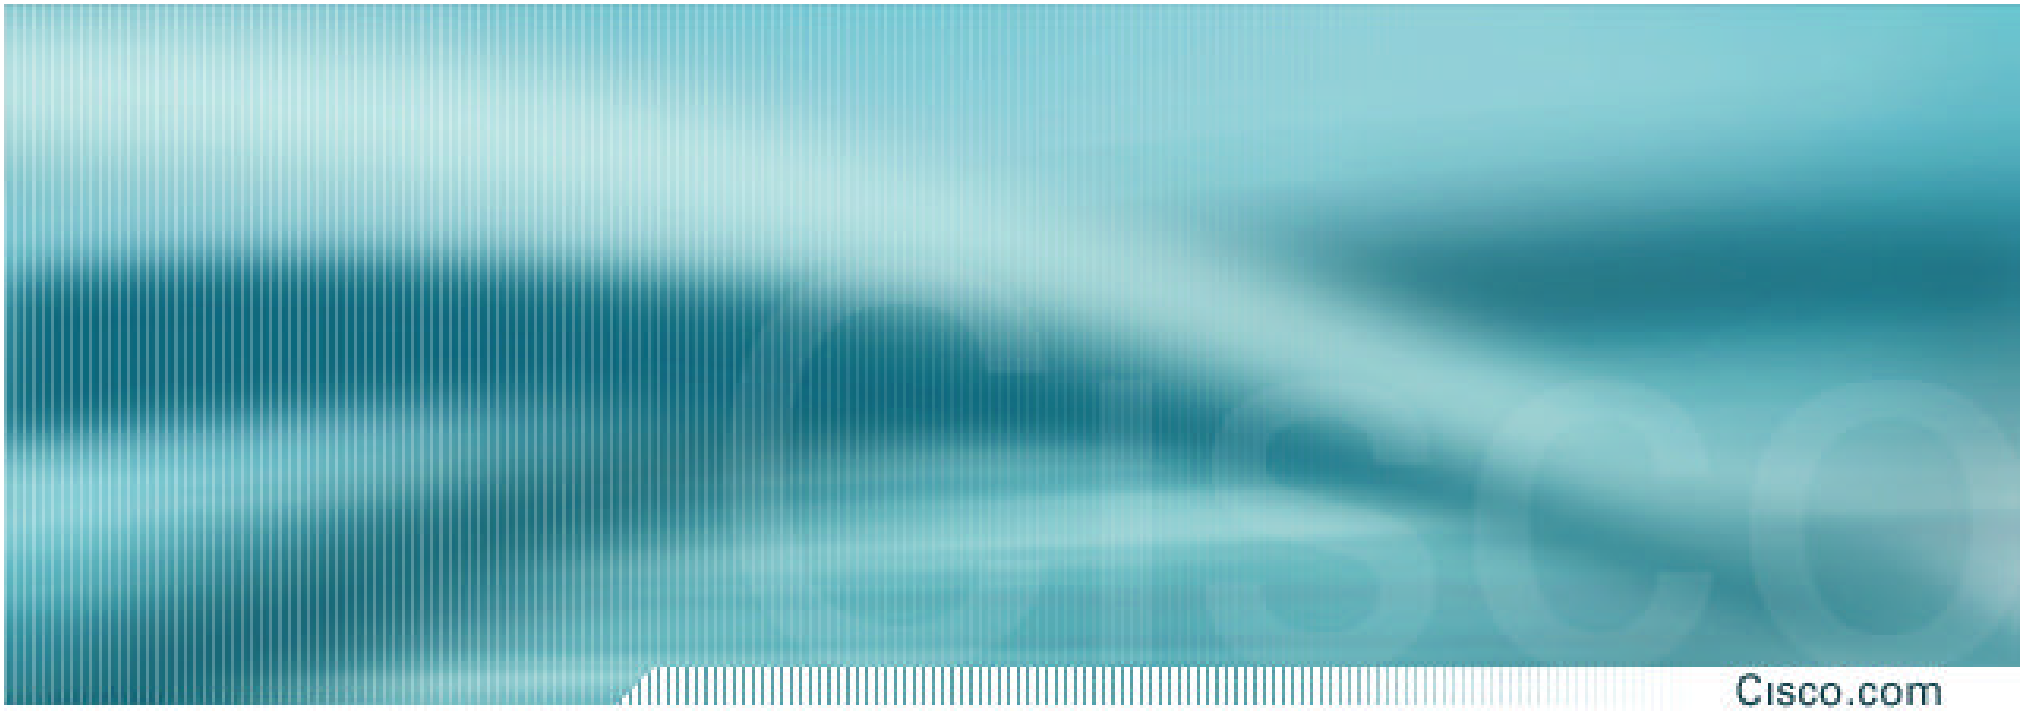

Phase 1 – Preparation for the Attack

Default Routes, ISPs, and Security

Avoid Default Routes

- **ISPs with full BGP feeds should avoid default routes.**
- **DOS/DDOS attack use spoofed addresses from the un-allocated IPV4 space.**

  See http://www.iana.org/assignments/ipv4-address-space for the latest macro allocations.
- **Backscatter traffic from DOS/DDOS targets need to go somewhere. If there is a default, then this traffic will do to this one router and get dropped.**
- **Dropping backscatter traffic might overload the router.**

Network with Default Route – Pointing to Upstream A

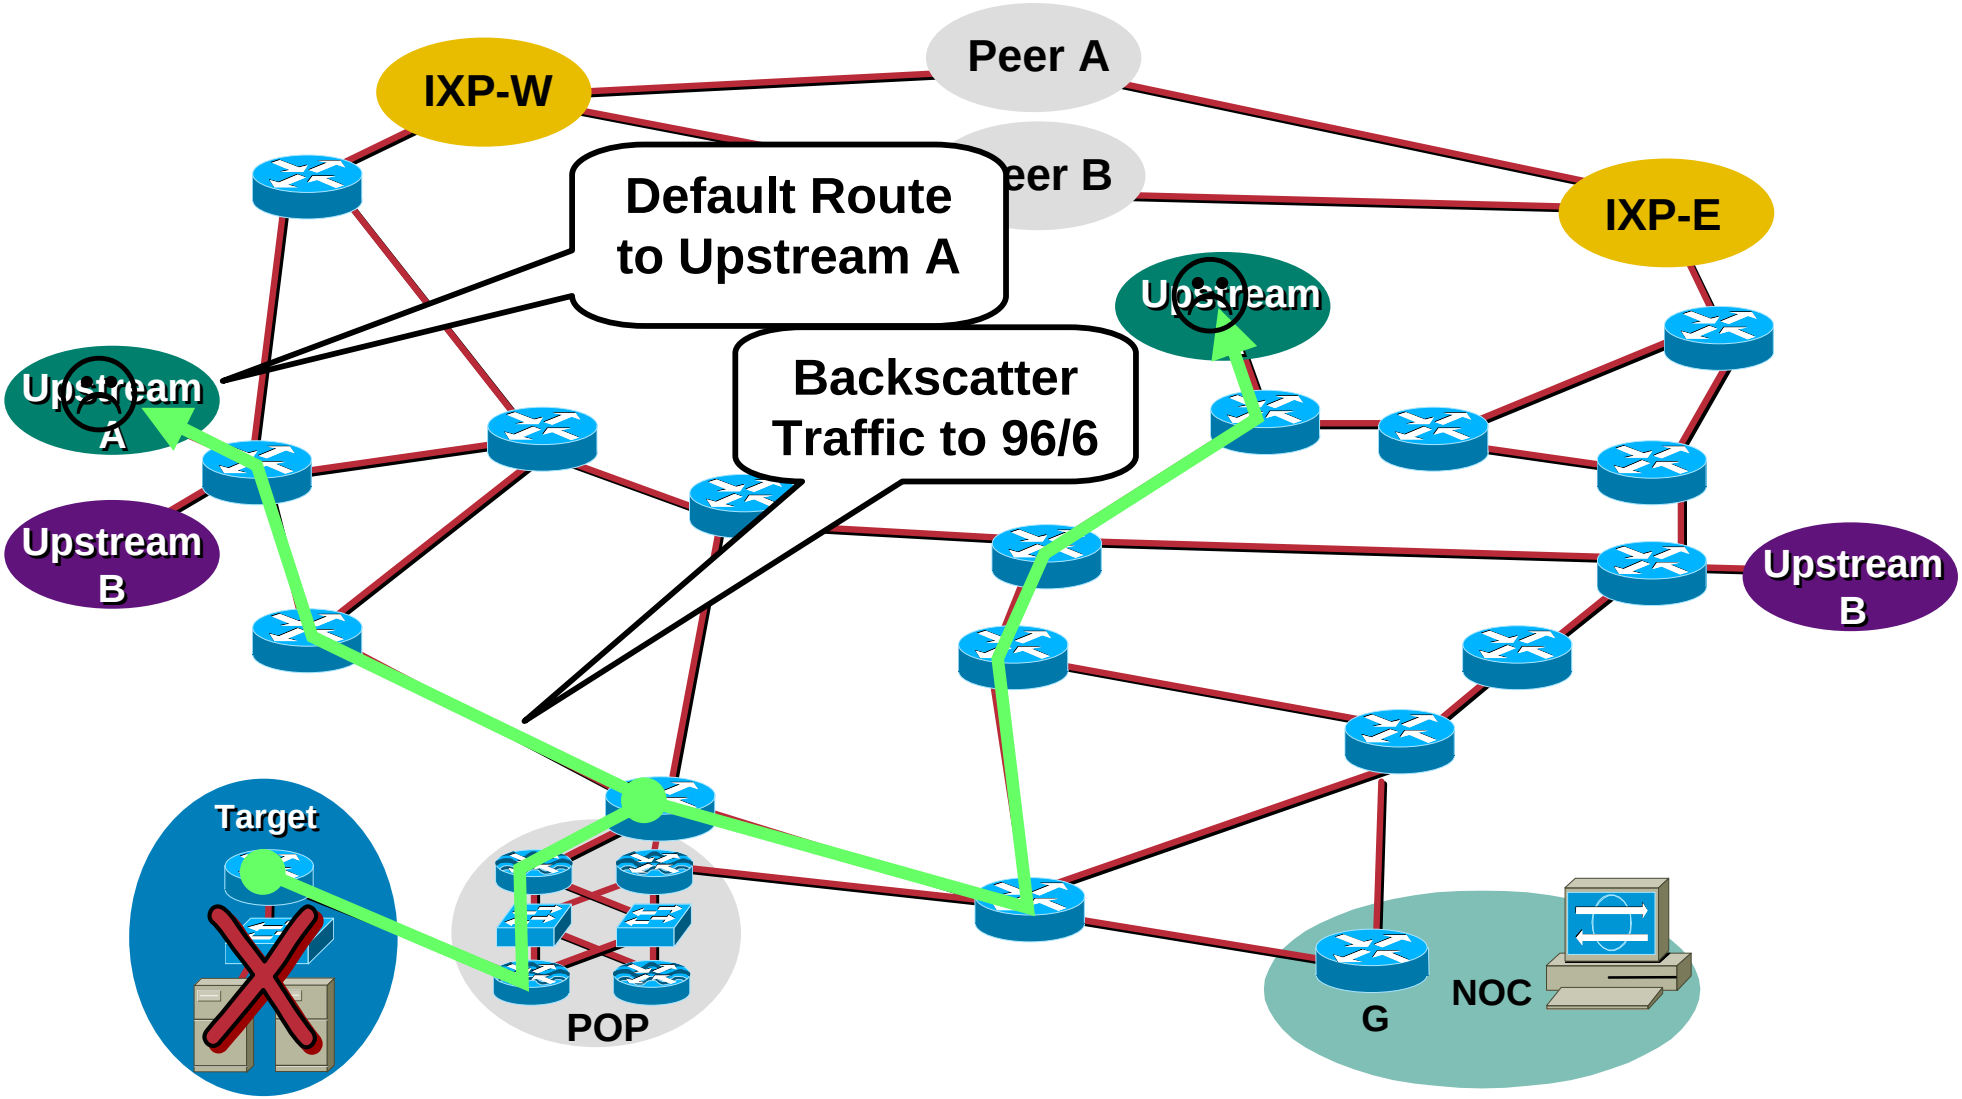

Network with Default Route – But not Pointing to Upstream

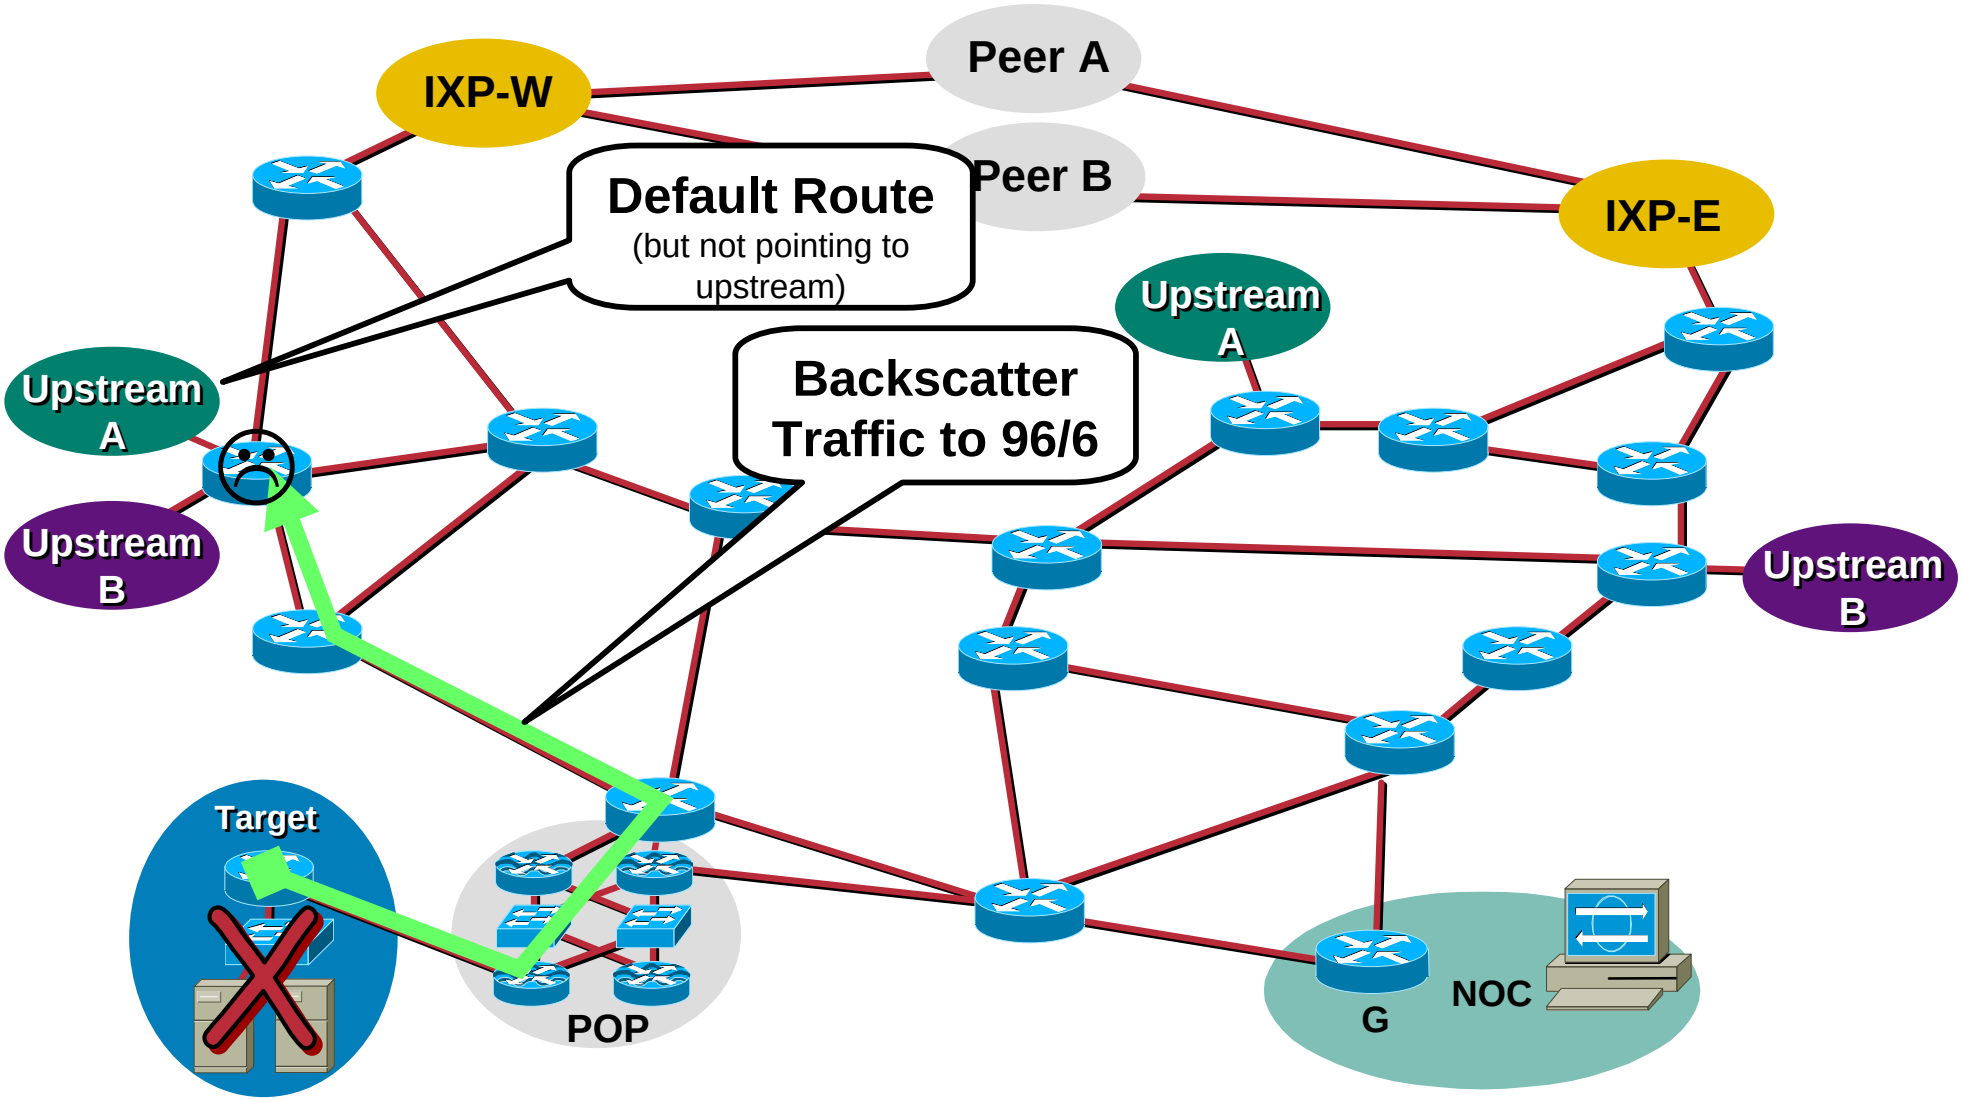

Network with No Default Route

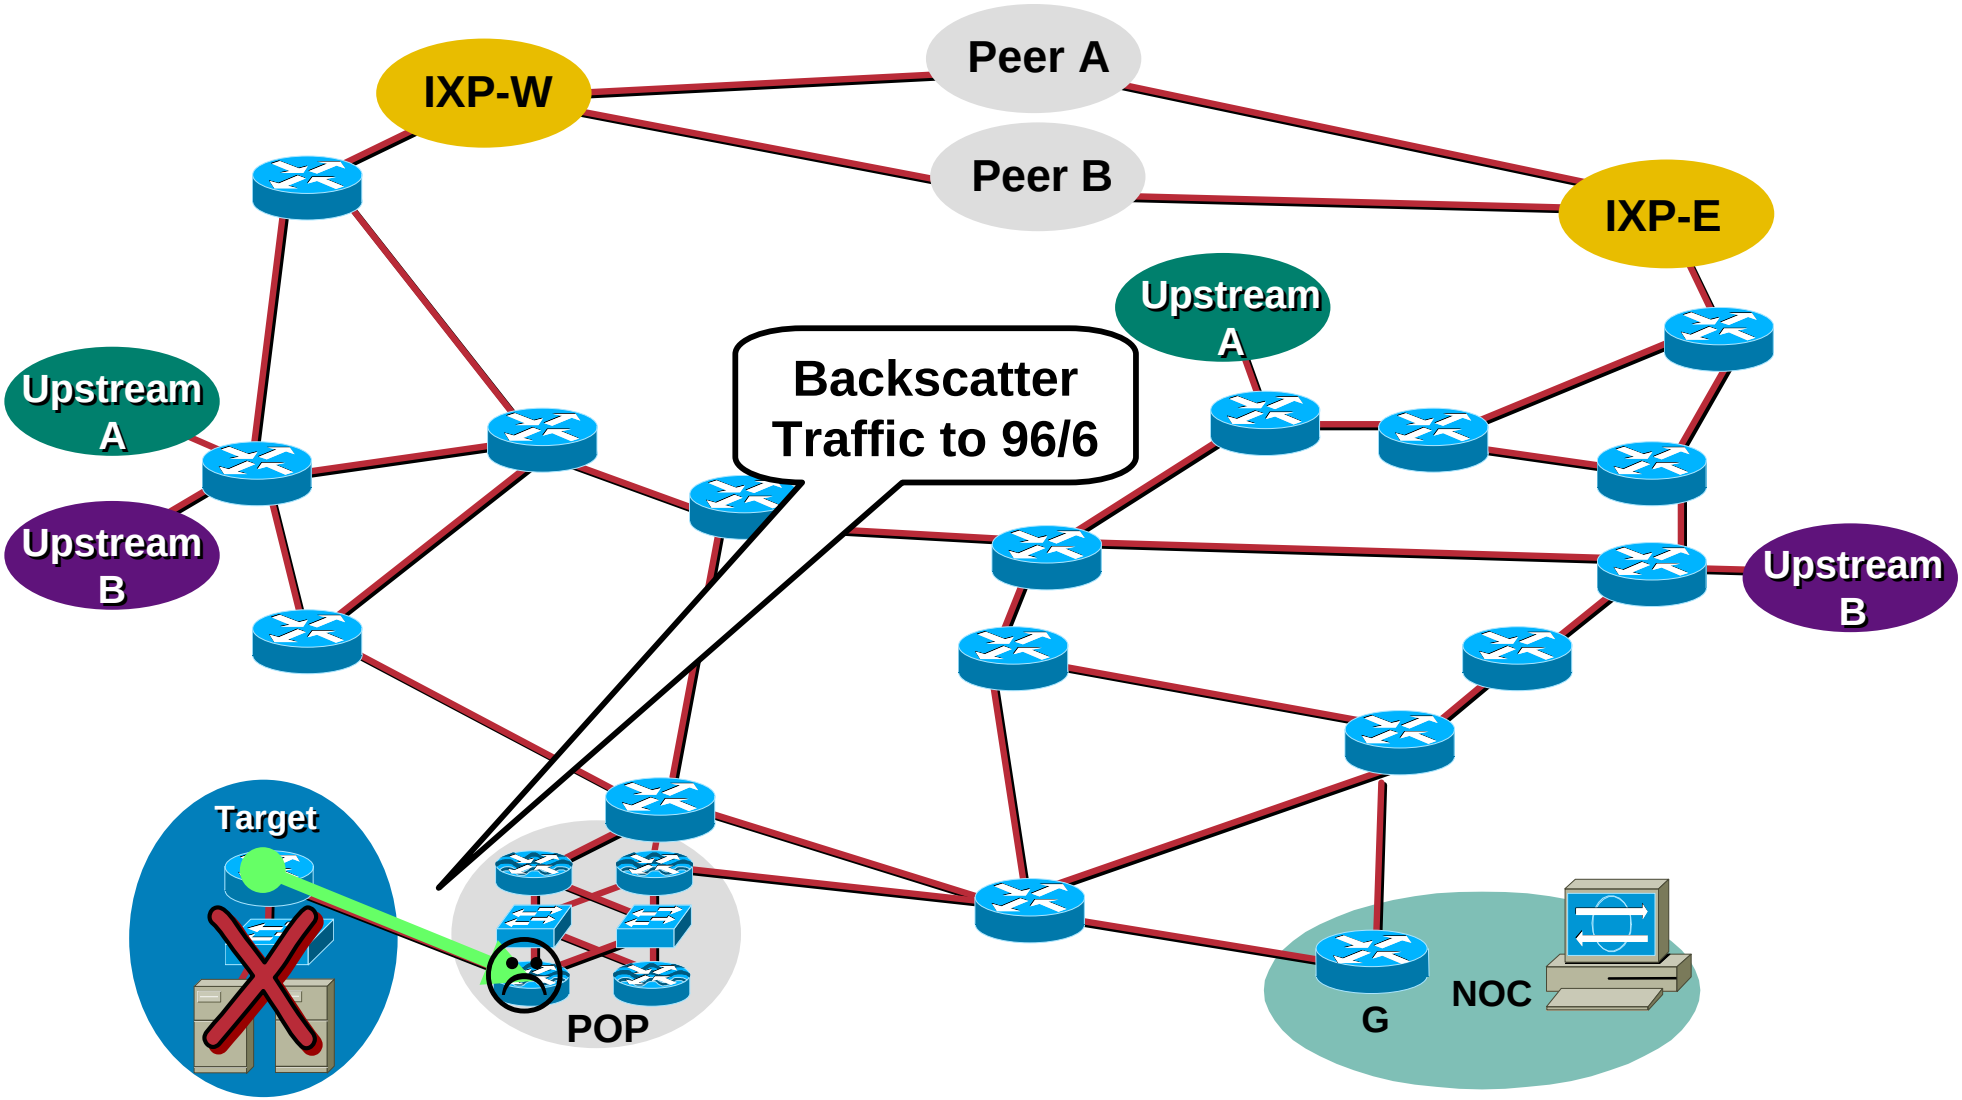

Default Route and ISP Security - Guidance

- **Engineer Default Route with ISP Security as one of the factors.**

  Most just engineer default with routing/forwarding as the only factor

- **If you need to use default, best to forward it upstream or to a Sink-Hole network engineered for packet drops.**

Default to a Sink-Hole Router/Network

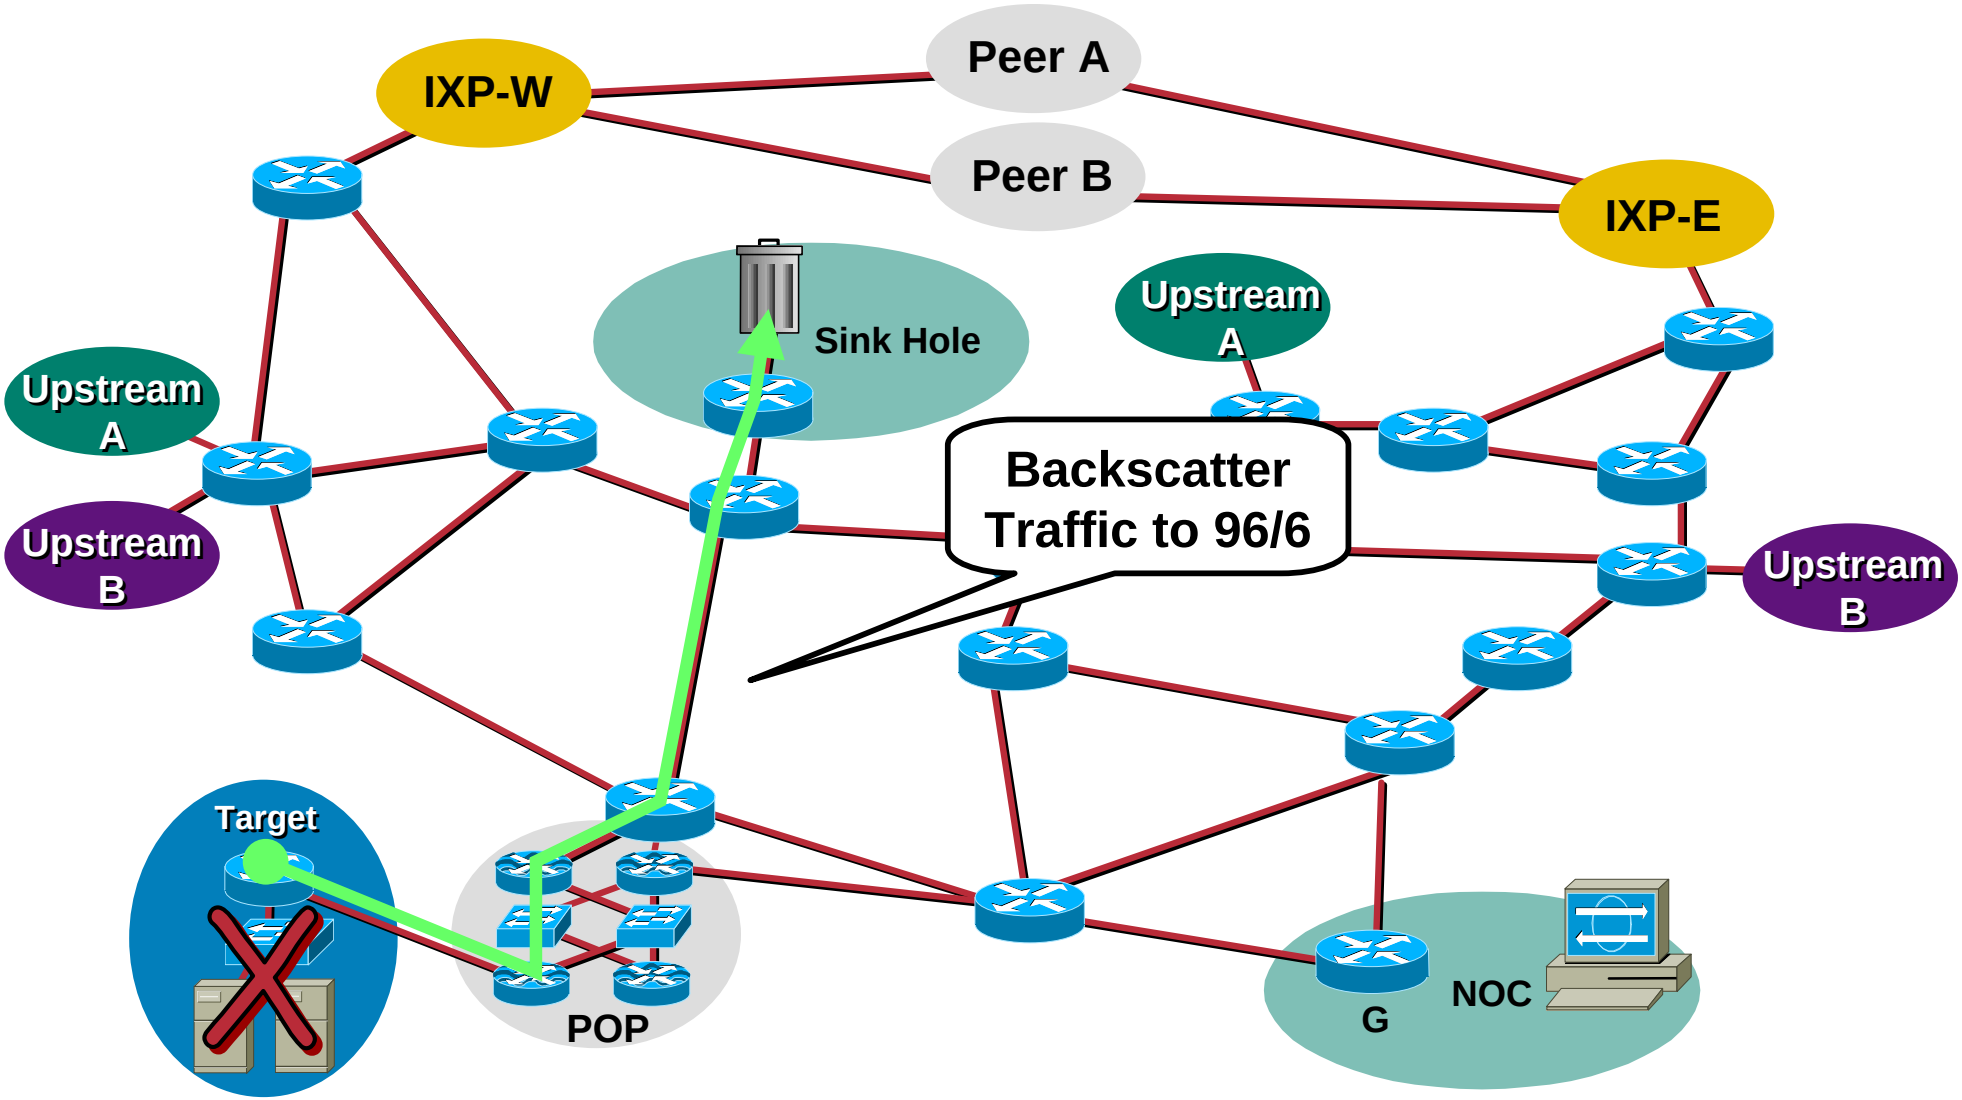

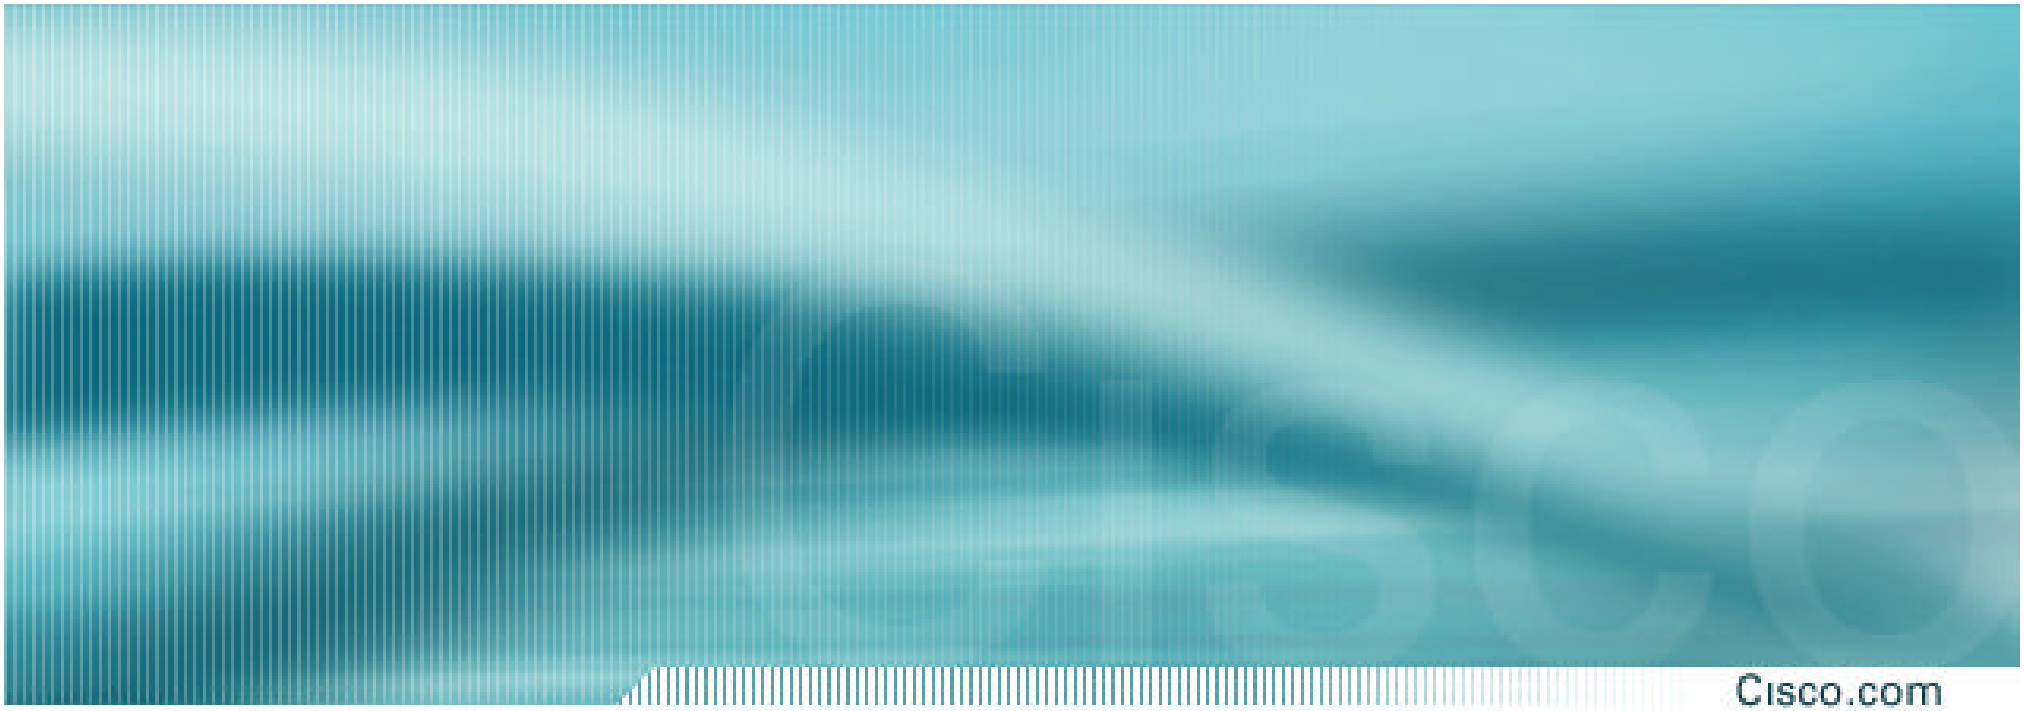

Phase 2 – Identification of the Attack

Identifying an Attack

- **When are we being probed?**

  Probes happen all the time; which ones are important?

  Probes precede an attack; if you can track specific probes, you might get a heads up that an attack is imminent

Identifying an Attack

Cisco.com

- **When are we your customers being attacked?**

  #1 way to identify that there is an attack in progress is when a customer calls the NOC

  New ISP oriented IDS tool are in the works

Identifying an Attack

- **When are you being attack?**

  NOC Alerts – is a problem in the network, a surge in traffic, a killer app, or someone attacking your network?

Identifying an Attack

- **SNMP Data abortion can signal a network problem *or* a security incident.**

CPU spike when nothing else is happening on the network and with no one working on the router.

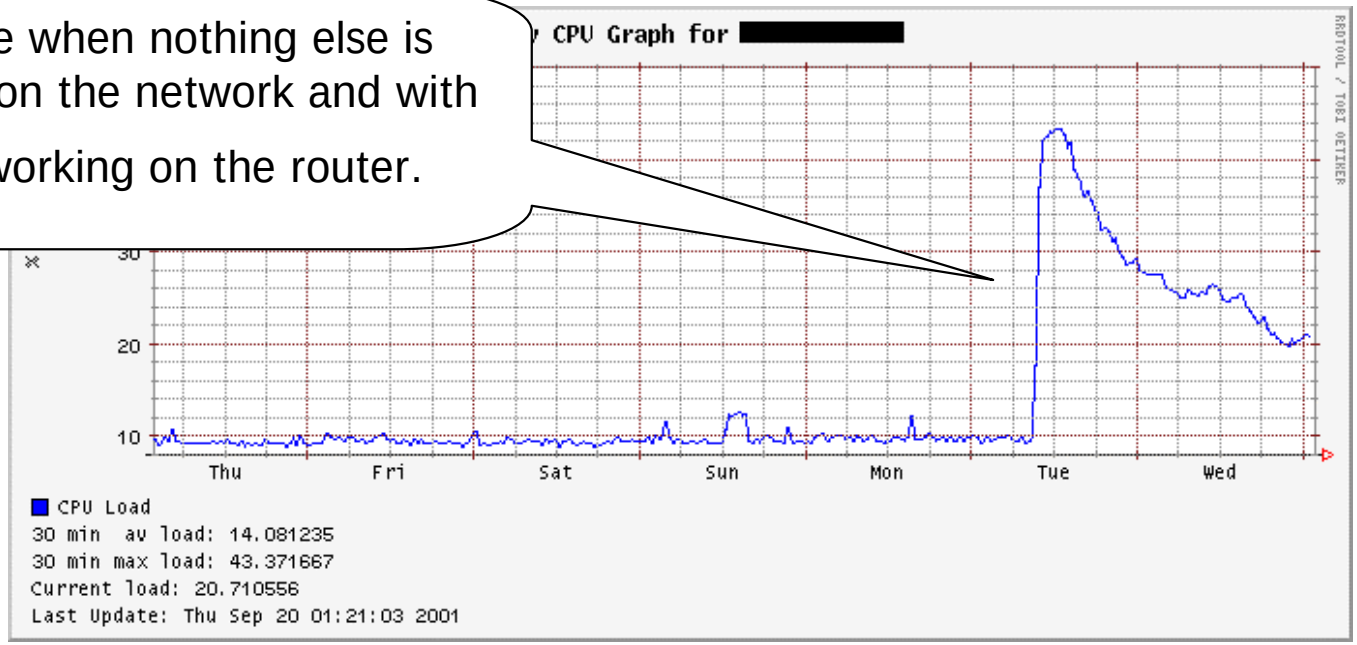

Identifying an Attack through CPU Load

router>sh proc cpu

CPU utilization for five seconds: A% B%, one minute: C%; five minutes: D%

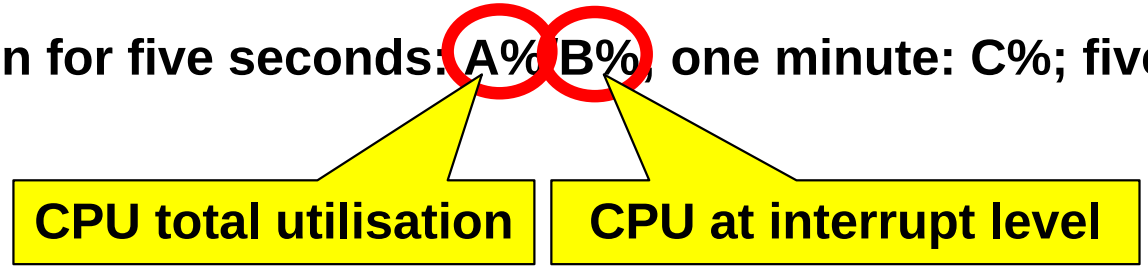

- **A: Total CPU load**
- **B: CPU at Interrupt level (note: B <= A)**
- **A-B: Process switched traffic, CPU processes**

(See: http://www.cisco.com/warp/public/63/highcpu.html)

Identifying an Attack through CPU Load

Cisco.com

router>sh proc cpu

CPU utilization for five seconds: A%/B%; one minute: C%; five minutes: D%

- **If A˜B: "Too much traffic to forward"**

  Interrupts: Packet switching (fast switching)

- **If A >> B: "Too much central processing"**

  Packets to/from the router (eg SNMP, ICMPs, vty and console, IPsec (w/o h/w), routing, …)

  Process switched packets or switching problem

Identifying an Attack through CPU Load

- **If A˜B (Packet Rate Getting too High)**

  Check interfaces to find the source:

  show interface

  Watch load and drops

  show interface switching

  Watch throttles (-> drops due to overload)

  Protocol stats (IP, ARP, …)

Identifying an Attack through CPU Load

- **If A>>B (CPU too busy)**

  Switching problems:

  Cache misses: If flow not in cache, ask CPU! (sh int switching)

  DoS: spoofed addresses -> many cache misses

  Packet from/to router:

  Routing, ARP, ICMP, SNMP, console, telnet, …
  Watch out: Too many ICMP could come from a route null0; use no ip unreachables

  Packet with options (could be DoS)

  …

Identifying Attacks with Netflow

- **Basis: Have Netflow running on the network**

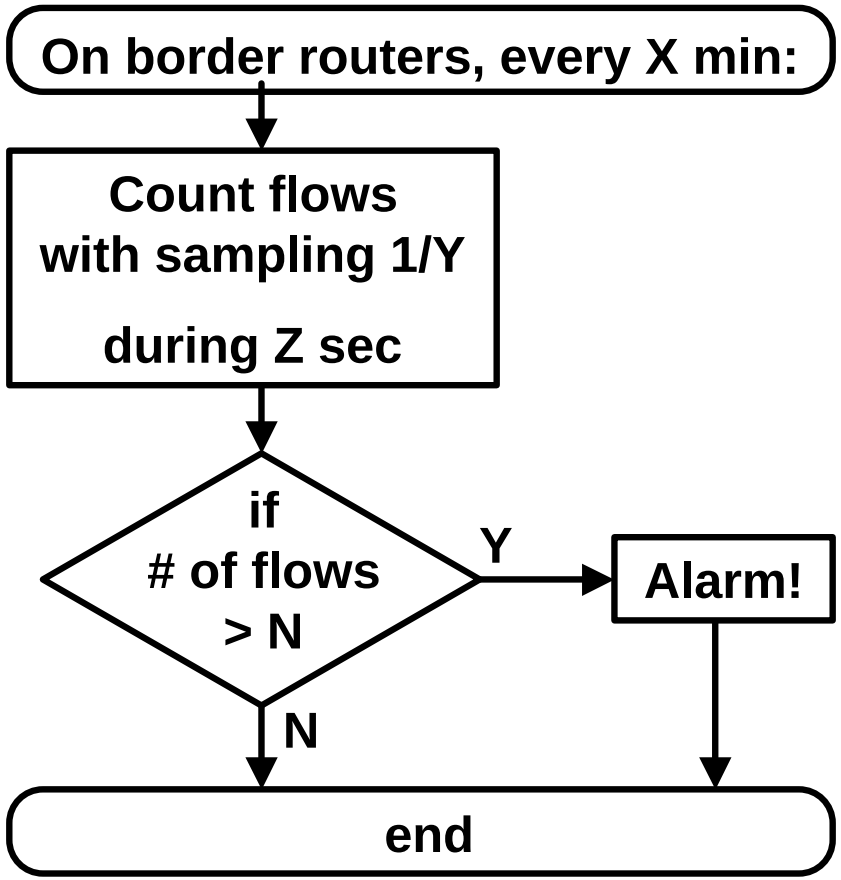

DANTE uses:
X=15 min, Y=200,
Z=10 sec, N=10

Values are empirical

How does a DoS Attack Look Like?

Cisco.com

Potential DoS attack (33 flows) on router1
Estimated: 660 pkt/s 0.2112 Mbps
ASxxx is: …
ASddd is: …

src_ip	dst_ip	in_if	out_if	s_port	d_port	pkts	bytes	prot	src_as	dst_as
192.xx.xxx.69	194.yyy.yyy.2	29	49	1308	77	1	40	6	xxx	ddd
192.xx.xxx.222	194.yyy.yyy.2	29	49	1774	1243	1	40	6	xxx	ddd
192.xx.xxx.108	194.yyy.yyy.2	29	49	1869	1076	1	40	6	xxx	ddd
192.xx.xxx.159	194.yyy.yyy.2	29	49	1050	903	1	40	6	xxx	ddd
192.xx.xxx.54	194.yyy.yyy.2	29	49	2018	730	1	40	6	xxx	ddd
192.xx.xxx.136	194.yyy.yyy.2	29	49	1821	559	1	40	6	xxx	ddd
192.xx.xxx.216	194.yyy.yyy.2	29	49	1516	383	1	40	6	xxx	ddd
192.xx.xxx.111	194.yyy.yyy.2	29	49	1894	45	1	40	6	xxx	ddd
192.xx.xxx.29	194.yyy.yyy.2	29	49	1600	1209	1	40	6	xxx	ddd
192.xx.xxx.24	194.yyy.yyy.2	29	49	1120	1034	1	40	6	xxx	ddd
192.xx.xxx.39	194.yyy.yyy.2	29	49	1459	868	1	40	6	xxx	ddd
192.xx.xxx.249	194.yyy.yyy.2	29	49	1967	692	1	40	6	xxx	ddd
192.xx.xxx.57	194.yyy.yyy.2	29	49	1044	521	1	40	6	xxx	ddd
192.xx.xxx.202	194.yyy.yyy.2	29	49	1840	345	1	40	6	xxx	ddd
192.xx.xxx.90	194.yyy.yyy.2	29	49	1327	176	1	40	6	xxx	ddd
192.xx.xxx.164	194.yyy.yyy.2	29	49	1451	1343	1	40	6	xxx	ddd

….

Observations of DANTE

- **False positive rate estimated at 2% (!)**

  (Biggest false positive targets: DNS root servers)

- **False negative: Attacks < 200 pps**
- **Spoofing: Mostly only host bit spoofed (!)**
- **Most attacks target pps performance with lots of small packets**
- **Most attacks last less than 15 mins**
- **Approx 35 attacks per day, 3-6 concurrent**

Observations of DANTE

- **Tackling Network DoS on Transit Networks**

  http://www.dante.net/pubs/dip/42/42.html

Identifying an Attack

Cisco.com

- **What about those Intrusion Detection Systems (IDS)?**

  Try them.

  Sink Hole Network is a good place to put them (sucks in all the junk and lets the IDS sort it out).

  Always be on the lookout for a new tool, trick, feature, or capability.

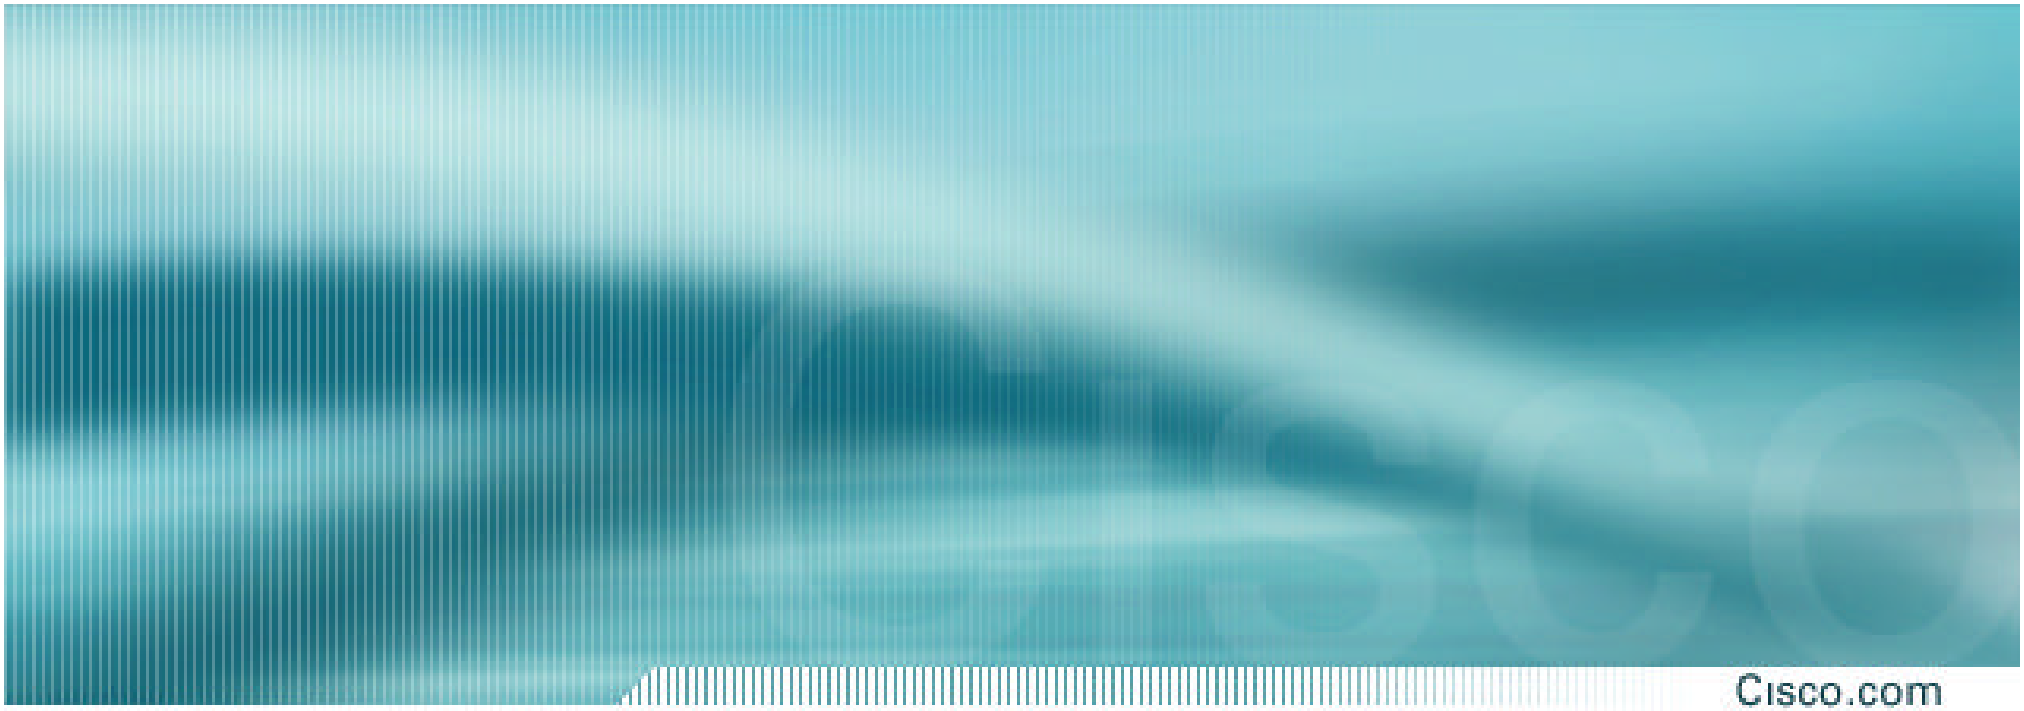

Phase 3 – Classification of the Attack

Phase 3 - Classifying an Attack

Cisco.com

- **How are we being attacked?**
  - **Once the attack starts, how do you find specifics of the attack?**
  - **Customer might provide information**
  - **Tools and procedures needed inside an ISP to specific information on the attack**
  - **Minimum source addresses and protocol type**

Phase 3 - Classifying an Attack

Cisco.com

- **Use ACL with permit for a group of protocols to drill down to the protocol**

```
Extended IP access list 169
        permit icmp any any echo (2 matches)
        permit icmp any any echo-reply (21374 matches)
        permit udp any any eq echo
        permit udp any eq echo any
        permit tcp any any established (150 matches)
        permit tcp any any (15 matches)
        permit ip any any (45 matches)
```

See http://www.cisco.com/warp/public/707/22.html

Sink Hole Classification Technique

- **Is it worth the risk to make config changes while a customer is under attack on a aggregation router with hundreds of customers connected to it?**

  Config changes when the network is under duress can and will cause more problems (it is not an "IOS" think – this applies to any network)

- **What would help is if the attack flow can be shifted from the target (i.e. customer) to some other router where the risk is manageable.**
- **Enter the Sink Hole Router.**

  Similar to a Unix HoneyPot.

Sink Hole Classification Technique

- **Sink Hole Router Preparation:**
  1. **Router with really fast packet dropping capability, software features, and a connection to the network (were traffic to it would not endanger the network). Think 7200 with the fastest NPE you can get.**
  2. **BGP session (Route Reflector Client). The target's more specific address will get advertised from here.**
  3. **Packet Filters, syslog exports, and a way to analyze the logs from the ACL's log-input.**

Sink Hole Classification Technique

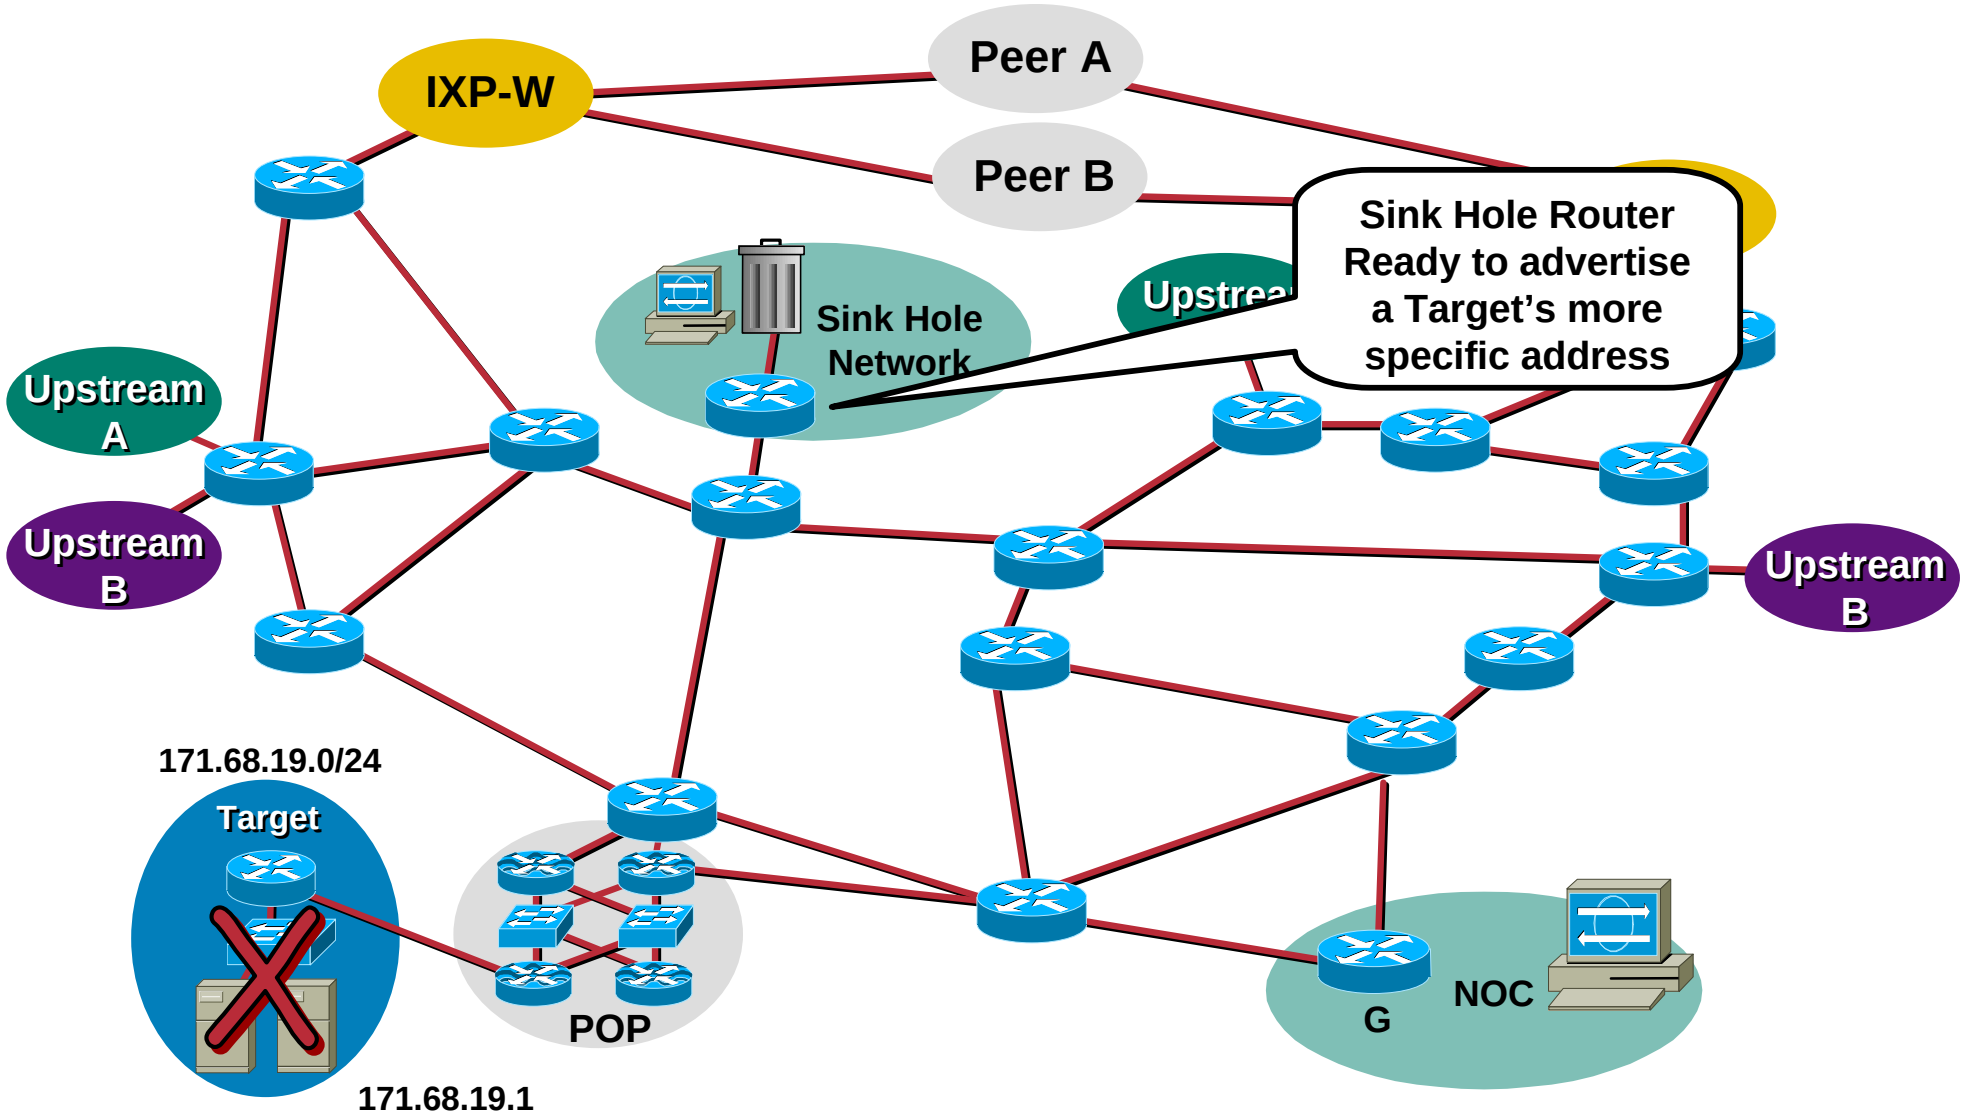

Sink Hole Classification Technique

Cisco.com

- **Sink Hole Classification – Activation**
  1. **Customer notifies ISP that they are under attack and need help. ISP lets the customer know that they will take the targeted host's IP address and redirect it to classify and traceback (see Backscatter Traceback technique).**
  2. **Sink Hole Router advertises the /32 address that is under attack.**
  3. **All traffic for that /32 shifts to the Sink Hole Router. ACL Packet Classification, Netflow Classification, or host based (specialized box) is done on a section of the ISPs network built to be attacked.**
  4. **Massive Aggregation Router is not touched.**

Sink Hole Classification Technique

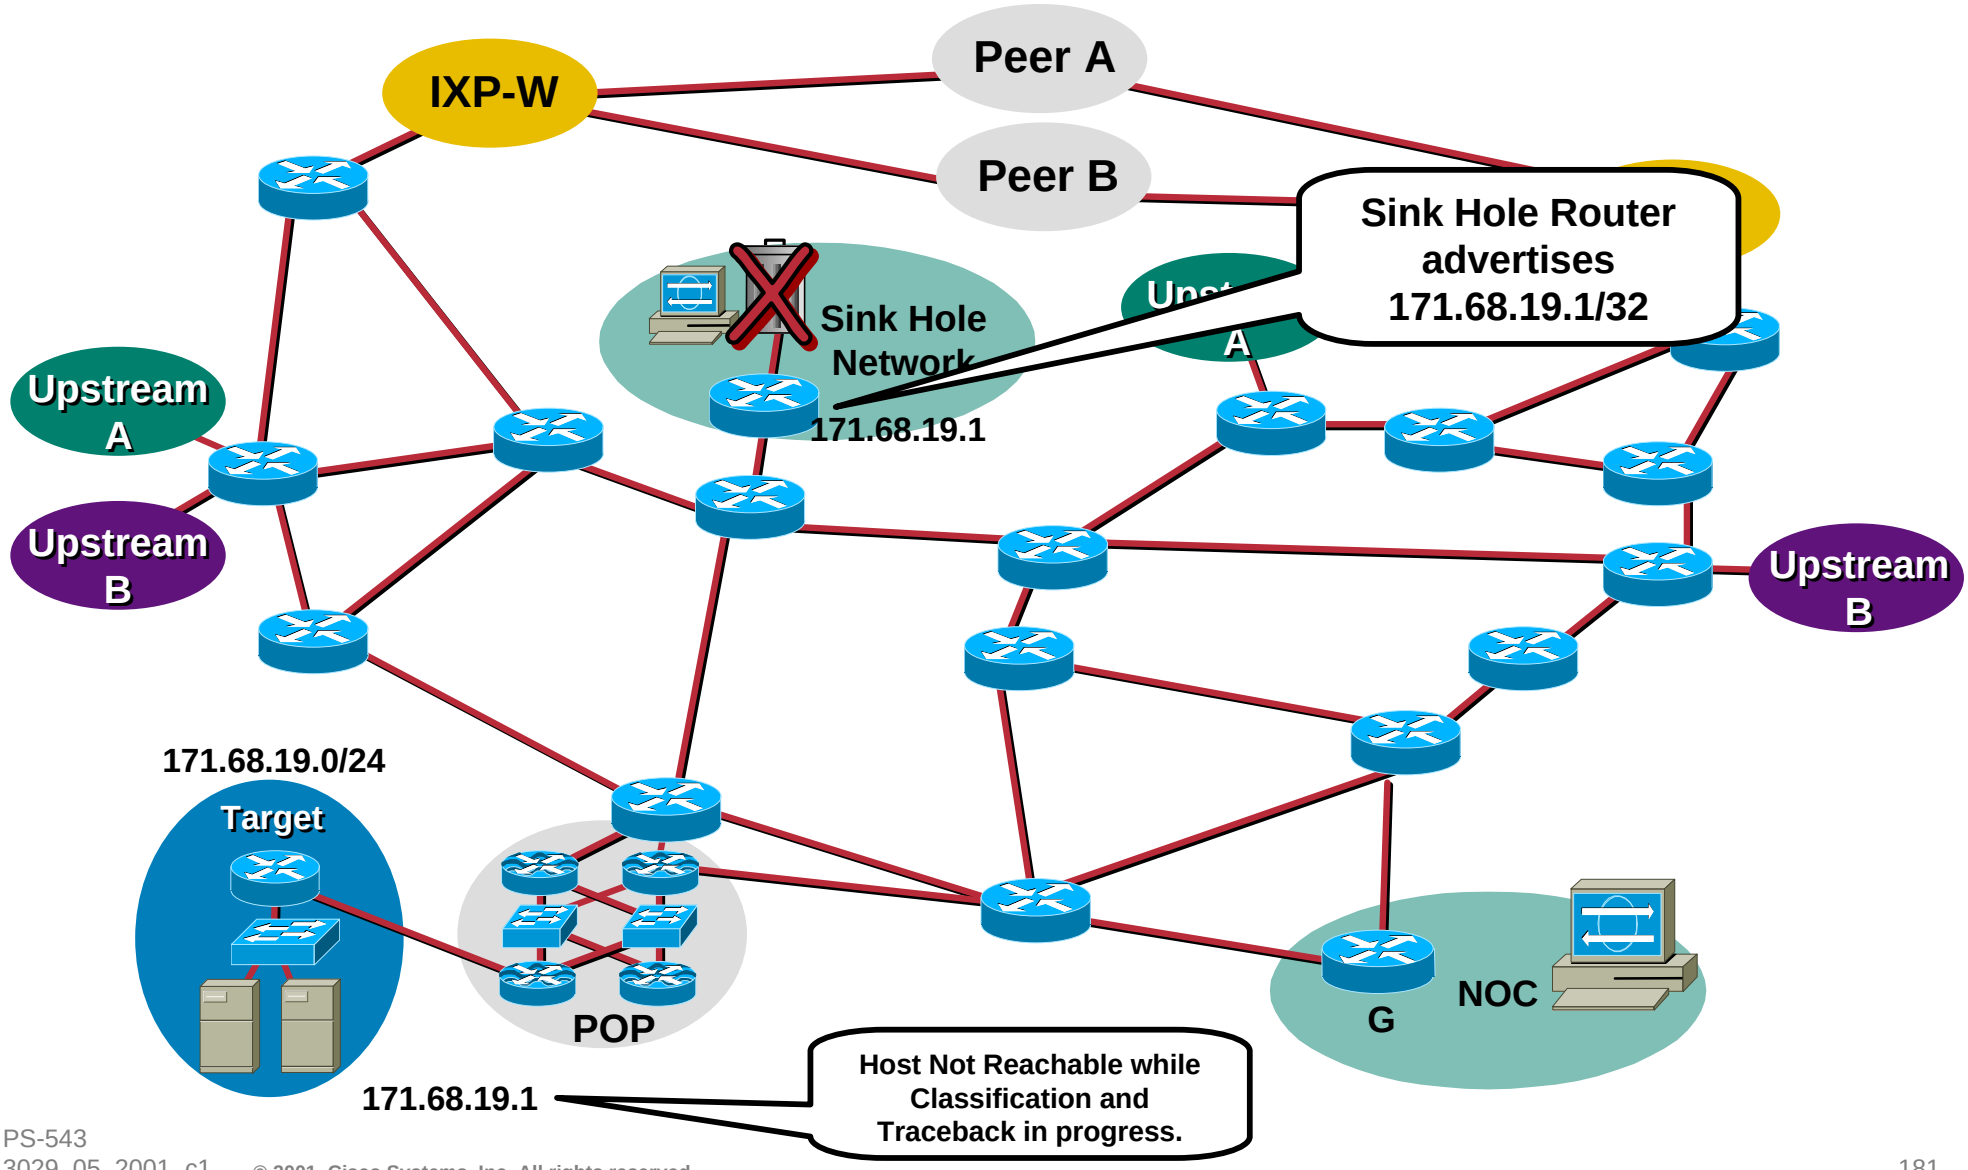

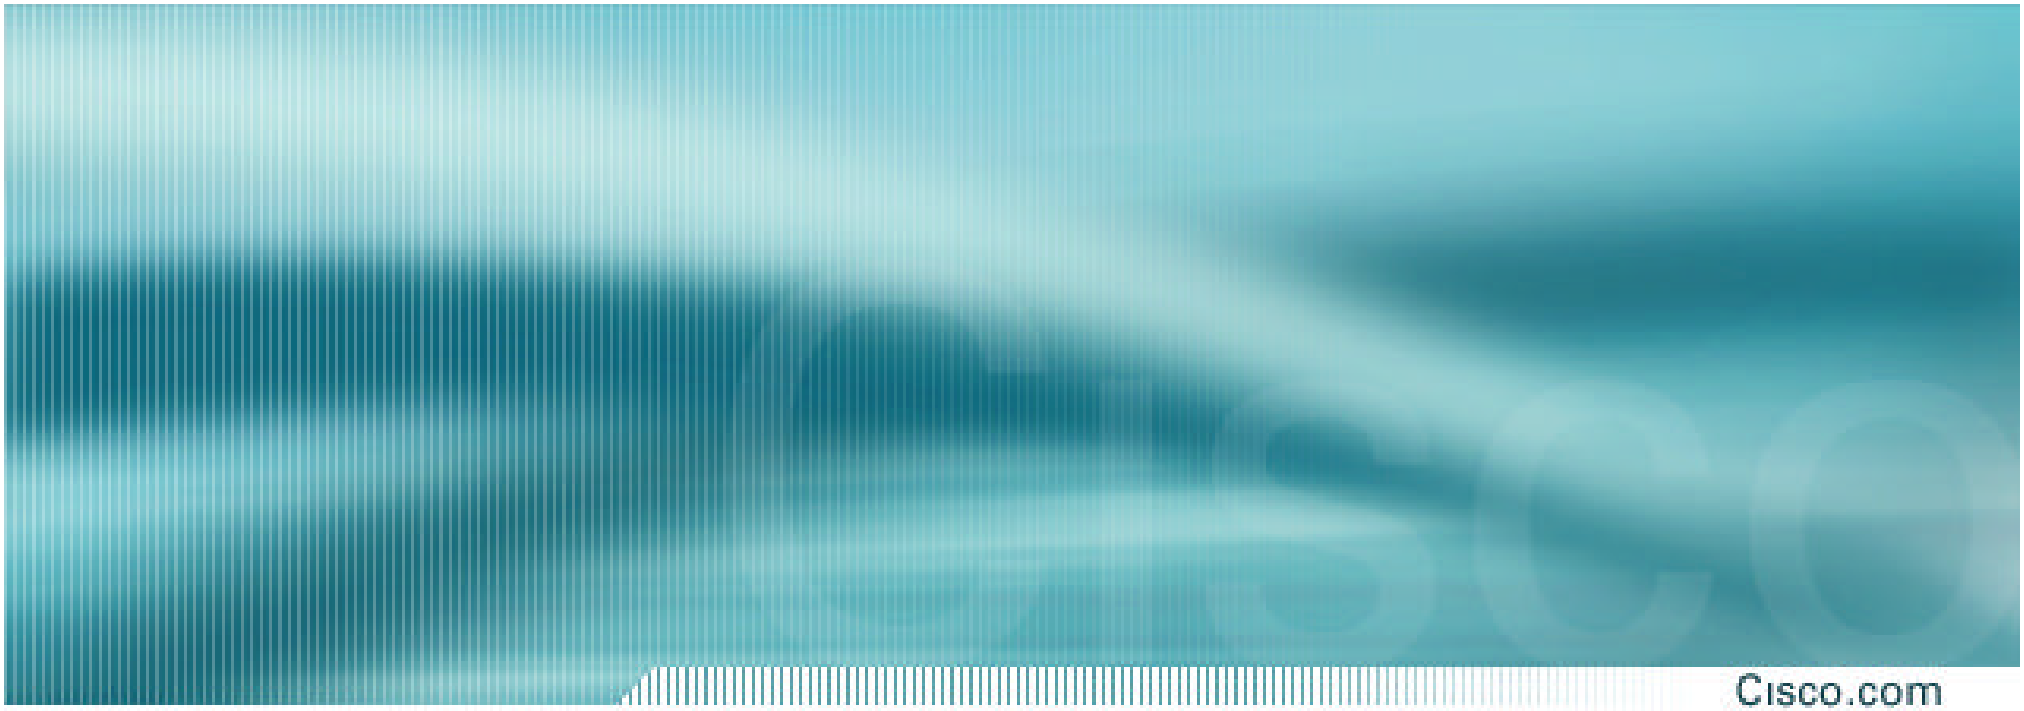

Phase 4 – Traceback the Attack

Traceback Attacks to their Source

- **Valid IPv4 Source Addresses are Easy.**

  Gets harder with DDOS – where there are a multitude of source addresses.

- **Spoofed IPv4 Source Addresses are more challenging.**

  Backscatter Traceback technique makes a difference.

- **Inter-Provider Hand off of the traceback is the big challenge today (end of 2001).**

Traceback Essentials

- **If source prefix is not spoofed:**

  -> Routing table
  -> Internet Routing Registry (IRR)
  -> direct site contact

- **If source prefix is spoofed:**

  -> Trace packet flow through the network
  -> Find upstream ISP
  -> Upstream needs to continue tracing

Traceback Valid IPv4 Source Addresses


```
madrid% whois -h whois.arin.net 64.103.0.0
Cisco Systems, Inc. (NETBLK-CISCO-GEN-6)
   170 West Tasman Drive
   San Jose, CA 95134
   US

   Netname: CISCO-GEN-6
   Netblock: 64.100.0.0 - 64.104.255.255

   Coordinator:
      Huegen, Craig  (CAH5-ARIN)  chuegen@cisco.com
      +1-408-526-8104 (FAX) +1 408 525 2597

   Domain System inverse mapping provided by:

   NS1.CISCO.COM             192.31.7.92
   NS2.CISCO.COM             192.135.250.69
   DNS-SJ6.CISCO.COM           192.31.7.93
   DNS-RTP4.CISCO.COM          192.135.250.70

   Record last updated on 11-Jan-2001.
   Database last updated on 2-Aug-2001 23:12:13 EDT.
```

- **Use Regional Internet Registries (RIRs):**
  - **Europe: whois.ripe.net**
  - **Asia-Pac: whois.apnic.net**
  - **USA and rest: whois.arin.net**

Traceback Valid IPv4 Source Addresses

Cisco.com

```
madrid% whois -h whois.arin.net "as 109"
Cisco Systems, Inc. (ASN-CISCO)
   170 W. Tasman Drive
   San Jose, CA 95134
   US

   Autonomous System Name: CISCOSYSTEMS
   Autonomous System Number: 109

   Coordinator:
      Koblas, Michelle  (MRK4-ARIN)  mkoblas@CISCO.COM
      (408) 526-5269 (FAX) (408) 526-4575

   Record last updated on 20-May-1997.
   Database last updated on 2-Aug-2001 23:12:13 EDT.
```

Also, if domain known: abuse@domain

Traceback Spoofed IPv4 Addresses

Cisco.com

- **From where are we being attacked (inside or outside)?**

  Once you have a fundamental understanding of the type of attack (source address and protocol type), you then need to track back to the ingress point of the network

  Two techniques—hop by hop and jump to ingress

Traceback via Hop by Hop Technique

- **Hop by hop tracebacks takes time**
  - **Starts from the beginning and traces to the source of the problem**
  - **Needs to be done on each router**
  - **Often requires splitting—tracing two separate paths**
  - **Speed is the limitation of the technique**

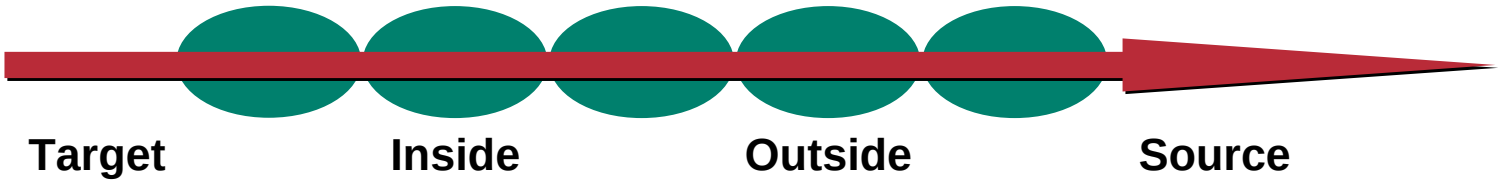

Traceback via Hop by Hop Technique

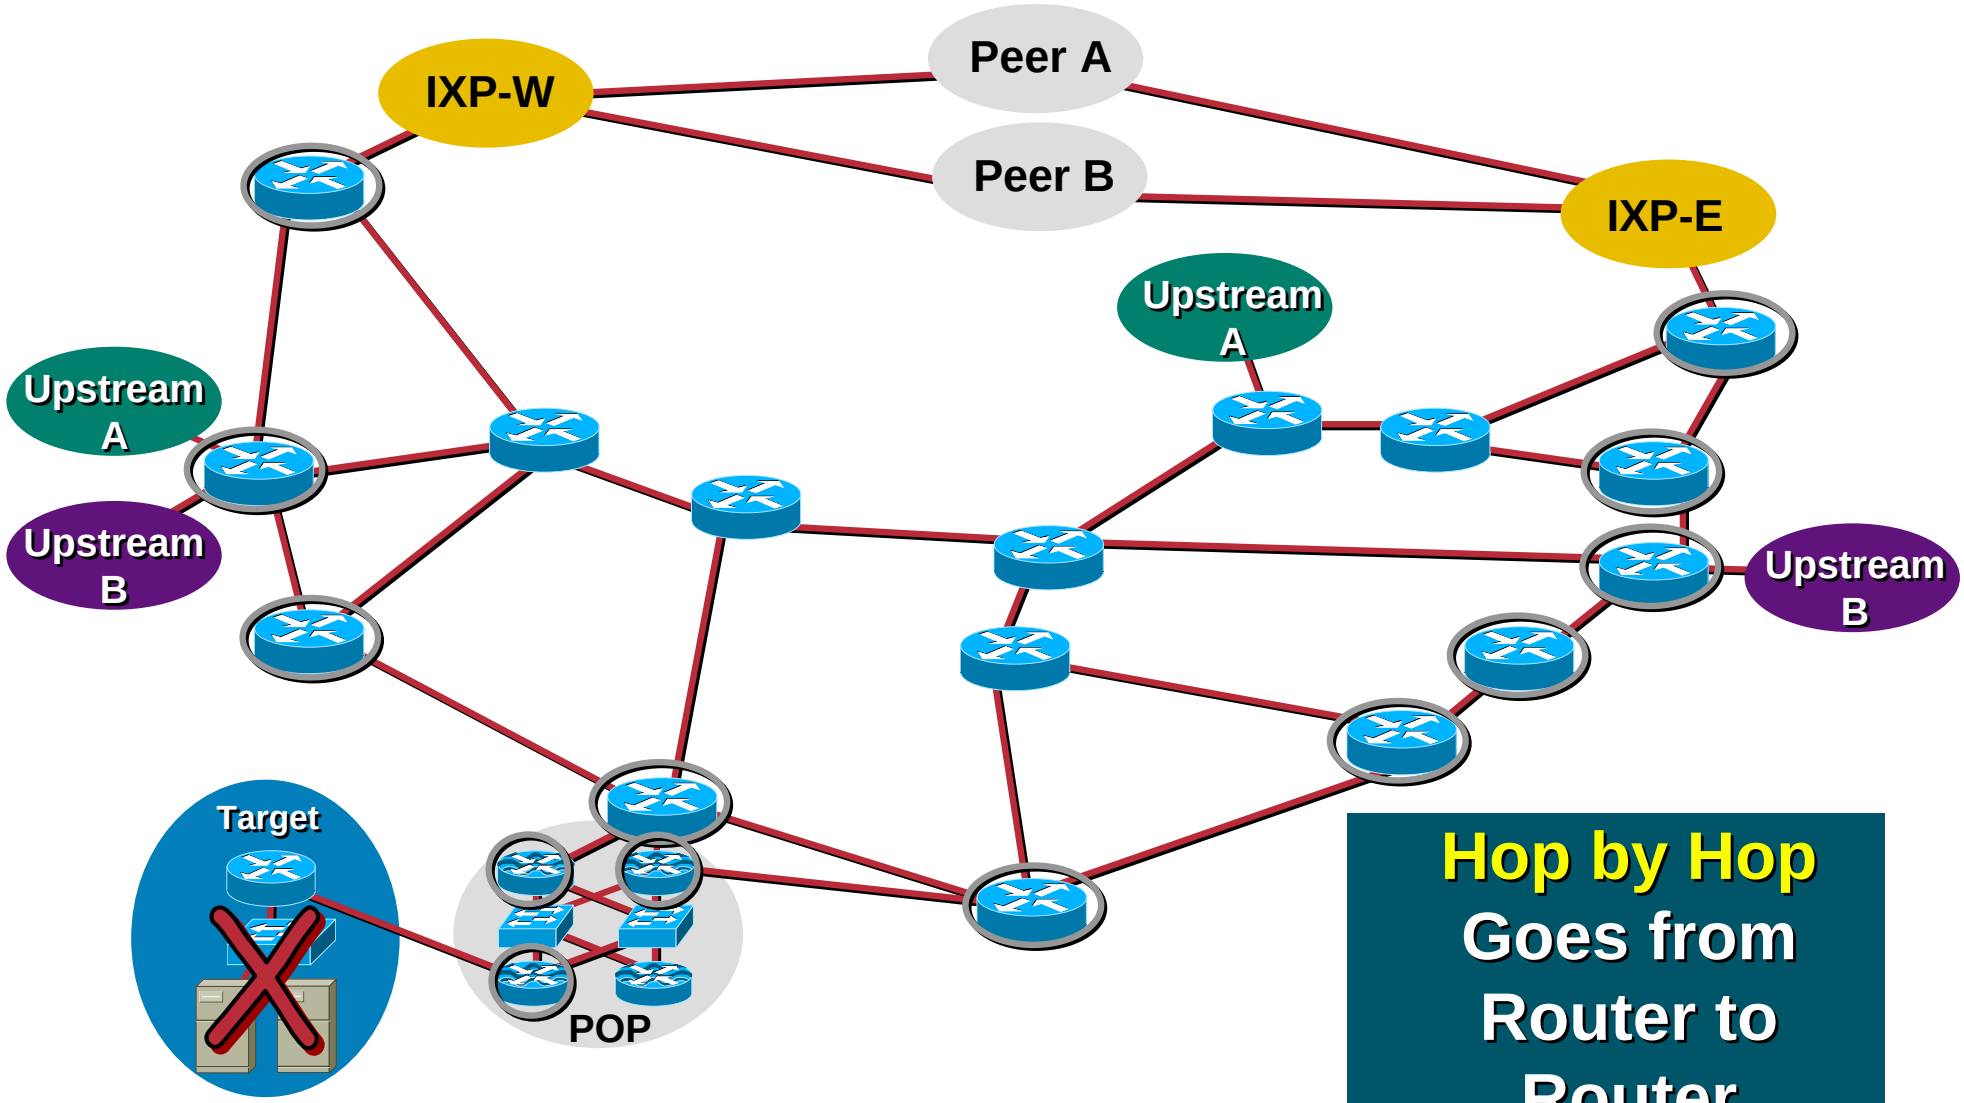

Hop by Hop Goes from Router to Router

Traceback via the Jump to Ingress Technique

- **Jump to ingress tracebacks divides the problem in half**
  - **Is the attack originating from inside the ISP or outside the ISP?**
  - **Jumps to the ISP's ingress border routers to see if the attack is entering the network from the outside**
  - **Advantage of speed—are we the source or someone else the source?**

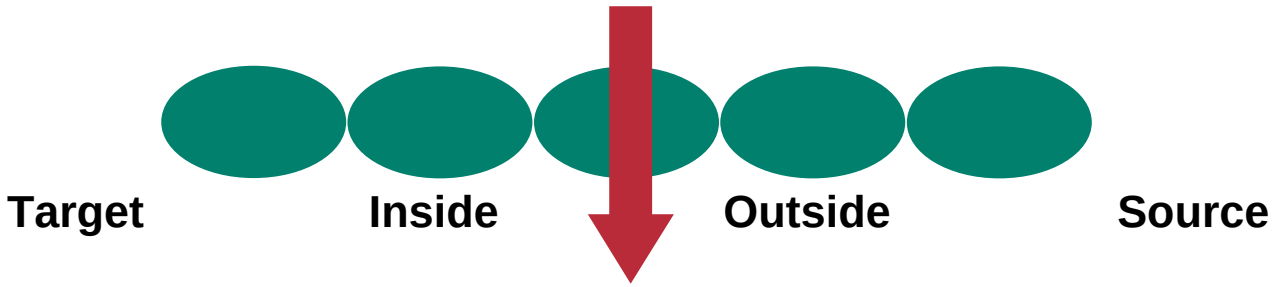

Traceback via the Jump to Ingress Technique

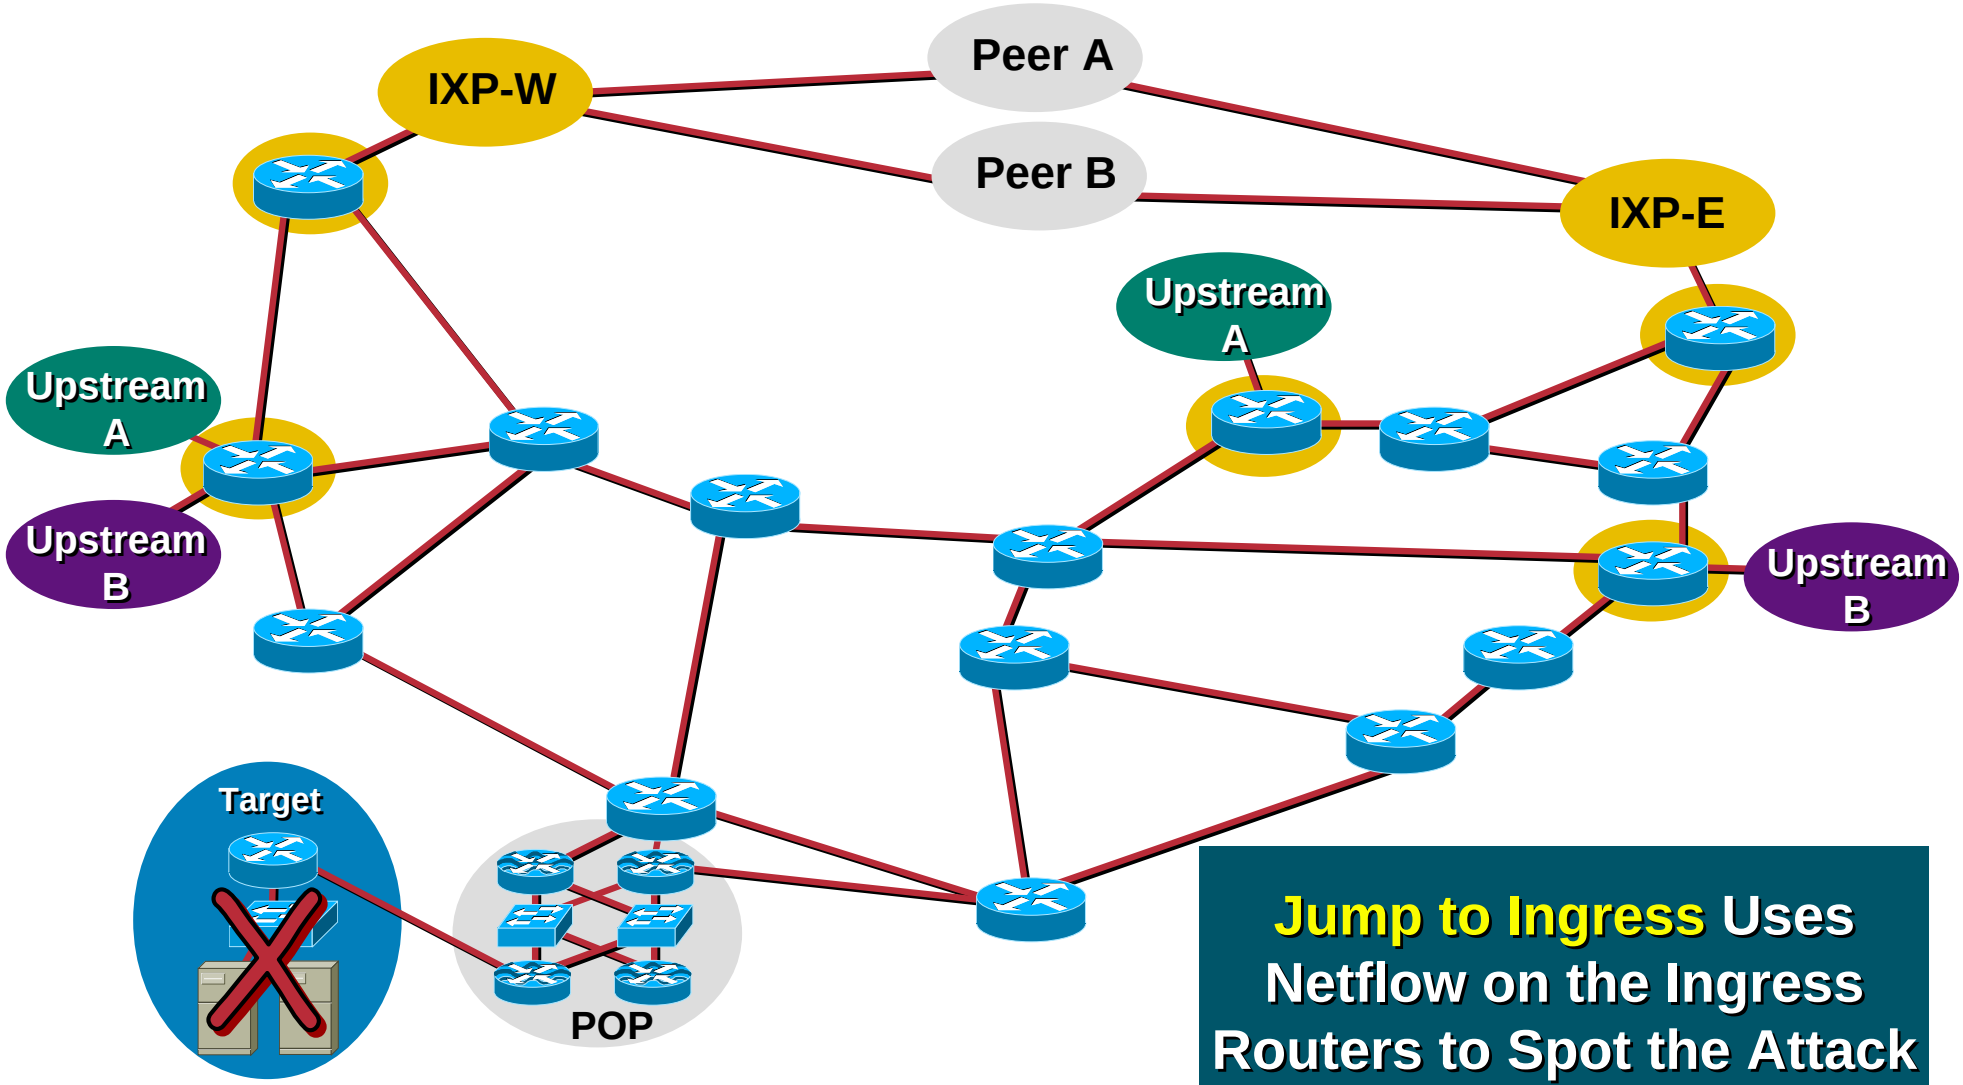

Jump to Ingress Uses Netflow on the Ingress Routers to Spot the Attack

Traceback Spoofed IPv4 Addresses

- **Three techniques**
  - **Apply temporary ACLs with log-input and examine the logs (like step 2)**
  - **Query Netflow's flow table (if show ip cache-flow is turned on)**
  - **Backscatter Traceback Technique**

Traceback with ACLs

Cisco.com

```
access-list 170 permit icmp any any echo
access-list 170 permit icmp any any echo-reply log-input
access-list 170 permit udp any any eq echo
access-list 170 permit udp any eq echo any
access-list 170 permit tcp any any established
access-list 170 permit tcp any any
access-list 170 permit ip any any

interface serial 0
  ip access-group 170 out
! Wait a short time - (i.e 10 seconds)
  no ip access-group 170 out
```

Traceback with ACLs

- Original technique for doing tracebacks
- Hazard—inserting change into a network that is under attack
- Hazard—log-input requires the forwarding ASIC to punt the packet to capture log information
- BCP is to apply the filter, capture just enough information, then remove the filter

Traceback with Netflow

- **Using Netflow for hop-by-hop traceback:**

```
Beta-7200-2>sh ip cache 198.133.219.0 255.255.255.0 verbose flow
IP packet size distribution (17093 total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .000 .735 .088 .054 .000 .000 .008 .046 .054 .000 .000 .000 .000 .000 .000

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
   .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000

IP Flow Switching Cache, 1257536 bytes
  3 active, 15549 inactive, 12992 added
  210043 ager polls, 0 flow alloc failures
  last clearing of statistics never
Protocol         Total    Flows   Packets Bytes  Packets  Active(Sec) Idle(Sec)
--------         Flows     /Sec     /Flow  /Pkt     /Sec      /Flow      /Flow
TCP-Telnet          35      0.0        80    41      0.0       14.5       12.7
UDP-DNS             20      0.0         1    67      0.0        0.0       15.3
UDP-NTP           1223      0.0         1    76      0.0        0.0       15.5
UDP-other        11709      0.0         1    87      0.0        0.1       15.5
ICMP                 2      0.0         1    56      0.0        0.0       15.2
Total:           12989      0.0         1    78      0.0        0.1       15.4

SrcIf         SrcIPaddress    DstIf         DstIPaddress    Pr SrcP DstP  Pkts
Fa1/1         192.168.45.142  POS1/0        198.133.219.25  11 008A 008A     1
Fa1/1         192.168.45.113  POS1/0        198.133.219.25  11 0208 0208     1
Fa1/1         172.16.132.154  POS1/0        198.133.219.25  06 701D 0017    63
```

Spoofed Flows are Tracks in Netflow!

Traceback with Netflow

- **Generic ways to use the Netflow command:**

  show ip cache <addr> <mask> verbose flow

  show ip cache flow | include <addr>

  Proactive approach—create scripts ……

    ssh -x -t -c [des|3des] -l <username> <IPAddr>
    “show ip cache <addr> <mask> verbose flow”

Traceback with Netflow

- **GSR—use the show controllers with sample Netflow (if LC supports SNF)**

```
GSR-2# exec slot 0 sh ip cache <addr> <mask>
verbose flow
```

- **7500 with dCEF—CSCdp91364.**

```
7500# exec slot 0 sh ip cache <addr> <mask>
verbose flow
```

- **Remember! *execute-on all* to get Netflow from all the LC/VIPs.**

Traceback with Netflow

- **Key advantage of Netflow:**
  - **No changes to the router while the network is under attack; passive monitoring**
  - **Scripts can be used to poll and sample throughout the network**
  - **IDS products can plug into Netflow**
  - **Working on a MIB for SNMP access**

Backscatter Traceback Technique

- **Three key advantages:**

  Reduced Operational Risk to the Network while traceback is in progress.

  Speedy Traceback

  Ability to hand off from one ISP to another – potentially tracing back to it's source.

Backscatter Traceback Technique

- **Created by Chris Morrow and Brian Gemberling @ UUNET as a means of finding the entry point of a spoofed DOS/DDOS.**

  http://www.secsup.org/Tracking/

- **Combines the Sink Hole router, Backscatter Effects of Spoofed DOS/DDOS attacks, and remote triggered Black Hole Filtering to create a traceback system that provides a result within 10 minutes.**

Backscatter Traceback Technique

Cisco.com

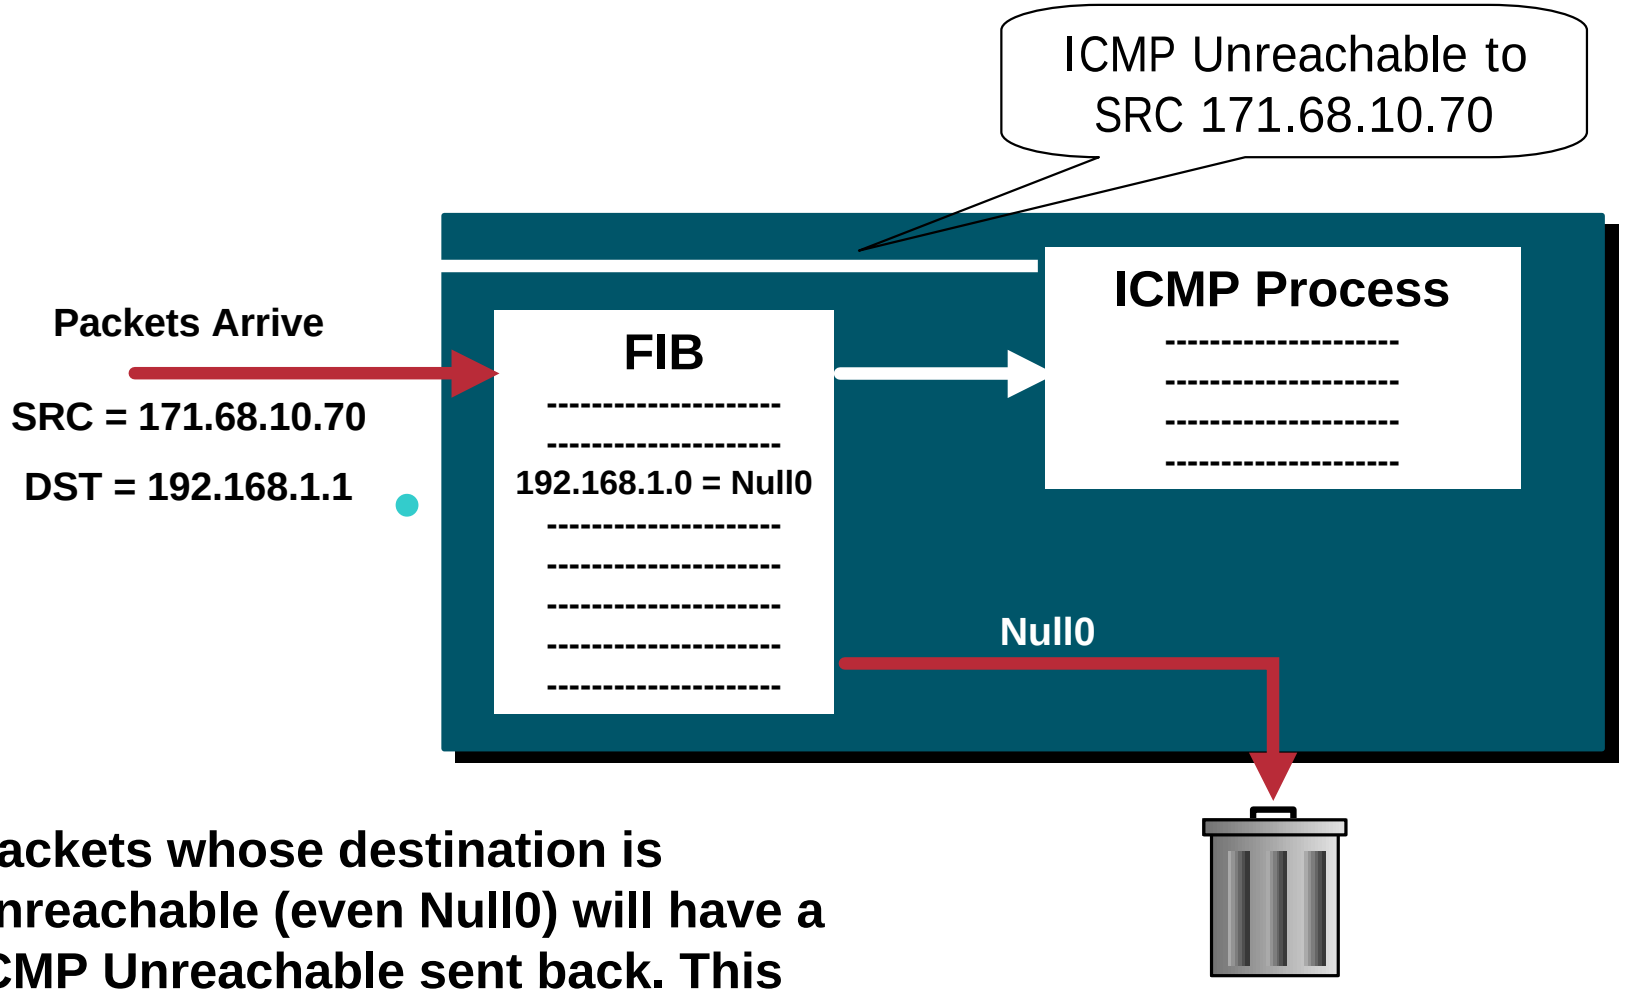

Packets whose destination is unreachable (even Null0) will have a ICMP Unreachable sent back. This "unreachable noise" is backscatter.

Backscatter Traceback *Preparation*

1. **Sink Hole Router/Network connected to the network and ready to classify the traffic. Like before, BGP Route Reflector Client, device to analyze logs, etc.**

   Can use one router to do both the route advertisement and logging OR break them into two separation routers – one for route advertisement and the other to accept/log traffic

   Can be used for other Sink Hole functions while not using the traceback technique.

   Sink Hole Router can be a iBGP Route Reflector into the network.

Backscatter Traceback *Preparation*

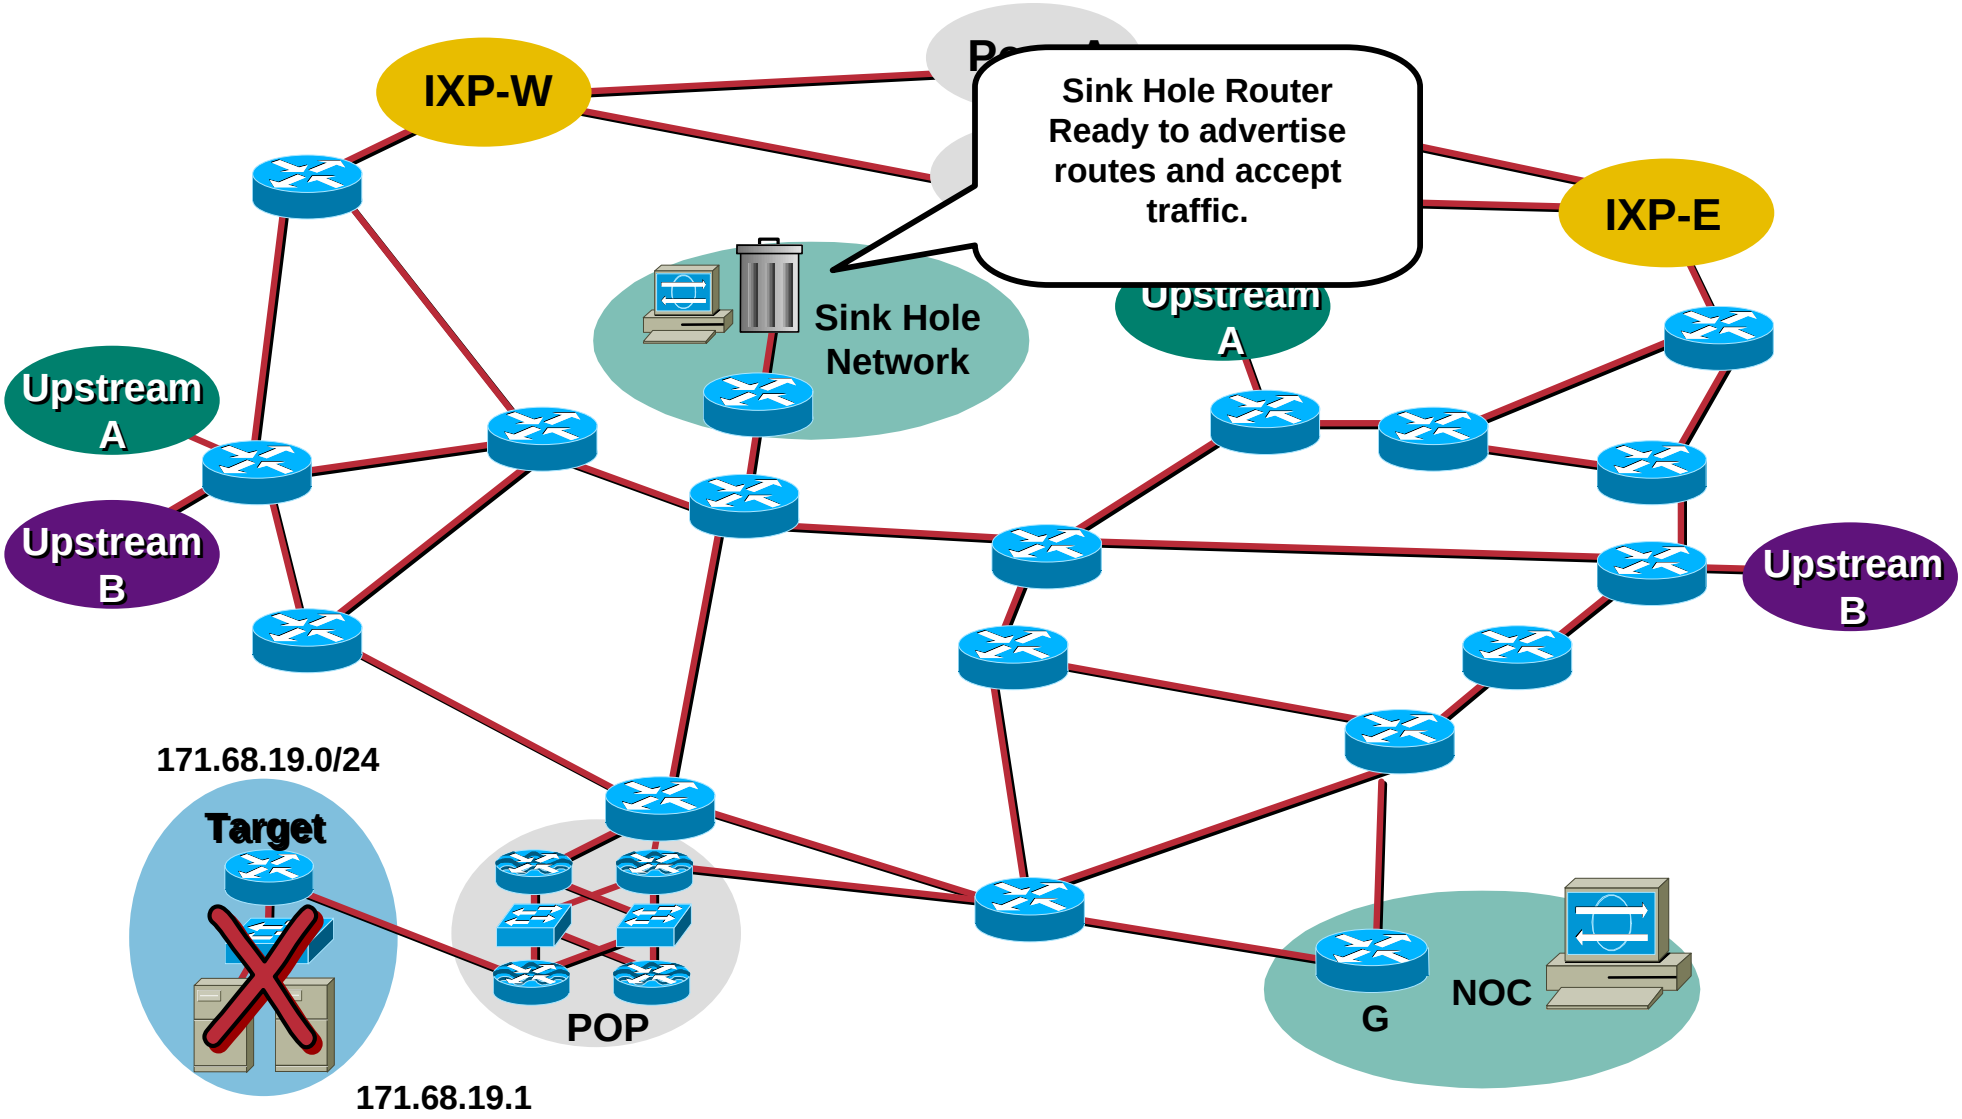

Backscatter Traceback *Preparation*

Cisco.com

```
router bgp 31337
!
! set the static redistribution to include a route-map so we can filter
! the routes somewhat... or at least manipulate them
! redistribute static route-map static-to-bgp
!
! add a stanza to the route-map to set our special next hop
!
 route-map static-to-bgp permit 5
 match tag 666
 set ip next-hop 172.20.20.1
 set local-preference 50
 set origin igp
```

Backscatter Traceback *Preparation*

2. **All edge devices (routers, NAS, IXP Routers, etc) with a static route to Null0. The Test-Net is a safe address to use (192.0.2.0/24) since no one is using it.**

   Cisco: **ip route 172.20.20.1 255.255.255.255 Null0**

   Routers also need to have ICMP Unreachables working. If you have ICMP Unreachables turned off (i.e. *no ip unreachables* on a Cisco), then make sure they are on.

   If ICMP Unreachable Overloads are a concern, use a ICMP Unreachable Rate Limit (i.e. *ip icmp rate-limit unreachable* command on a Cisco).

Backscatter Traceback *Preparation*

Cisco.com

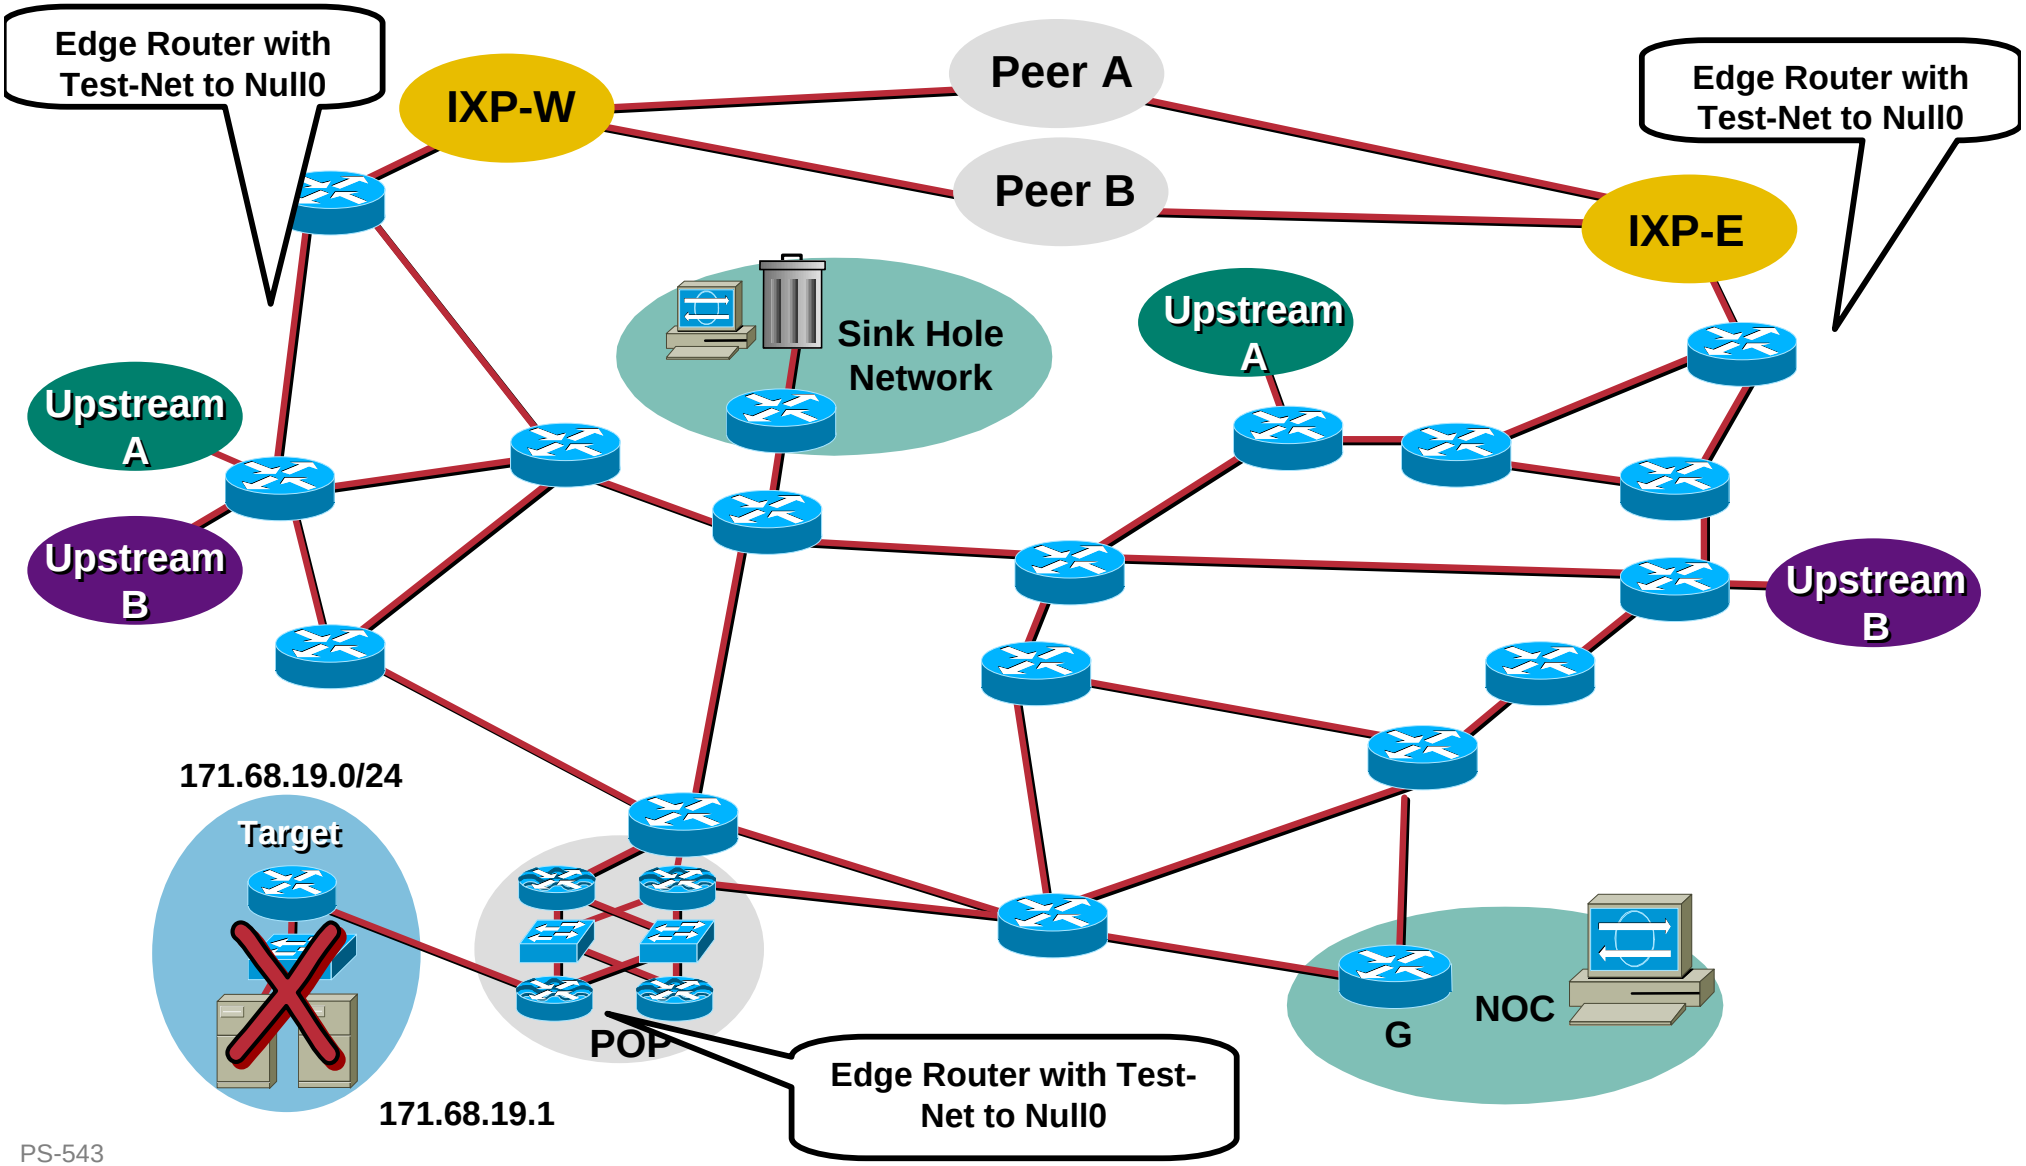

Backscatter Traceback *Preparation*

3. **Sink Hole Router advertising a large block of un-allocated address space with the BGP no-export community and BGP Egress route filters to keep the block inside. 96.0.0.0/3 is an example.**

   Check with IANA for unallocated blocks:
   `www.iana.org/assignments/ipv4-address-space`

   BGP Egress filter should keep this advertisement inside your network.

   Use BGP *no-export* community to insure it stays inside your network.

Backscatter Traceback *Preparation*

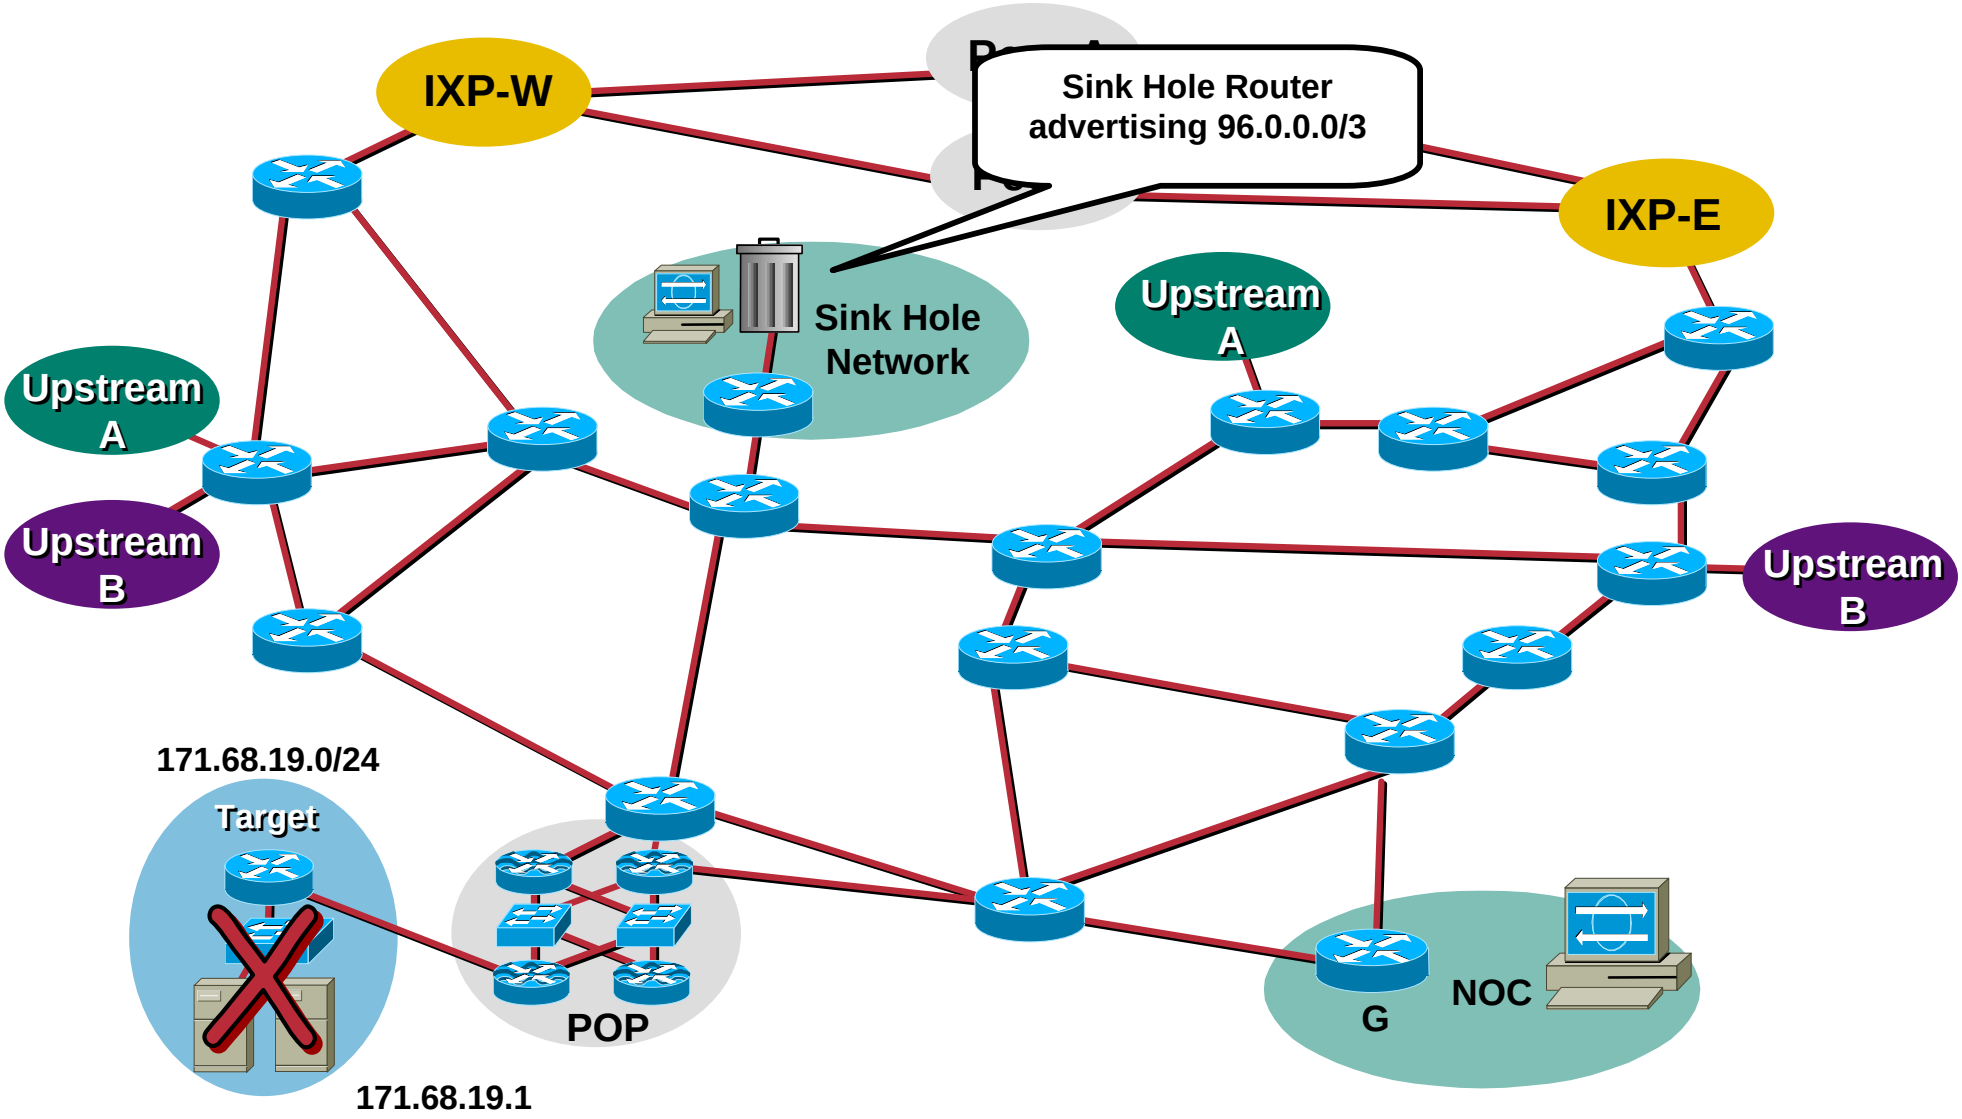

Backscatter Traceback *Activation*

Cisco.com

- **Activation happens when an attack has been identified.**
- **Basic Classification should be done to see if the backscatter traceback will work:**

  May need to adjust the advertised block.

  Statistically, most attacks have been spoofed using the entire Internet block.

Backscatter Traceback *Activation*

Cisco.com

1. **Sink Hole Router Advertises the /32 under attack into iBGP with.**

Advertised with a static route with the “666” tag:

```
ip route victimip 255.255.255.255 Null0 tag 666
```

The static triggers the routers to advertise the customer’s prefix

Backscatter Traceback *Activation*

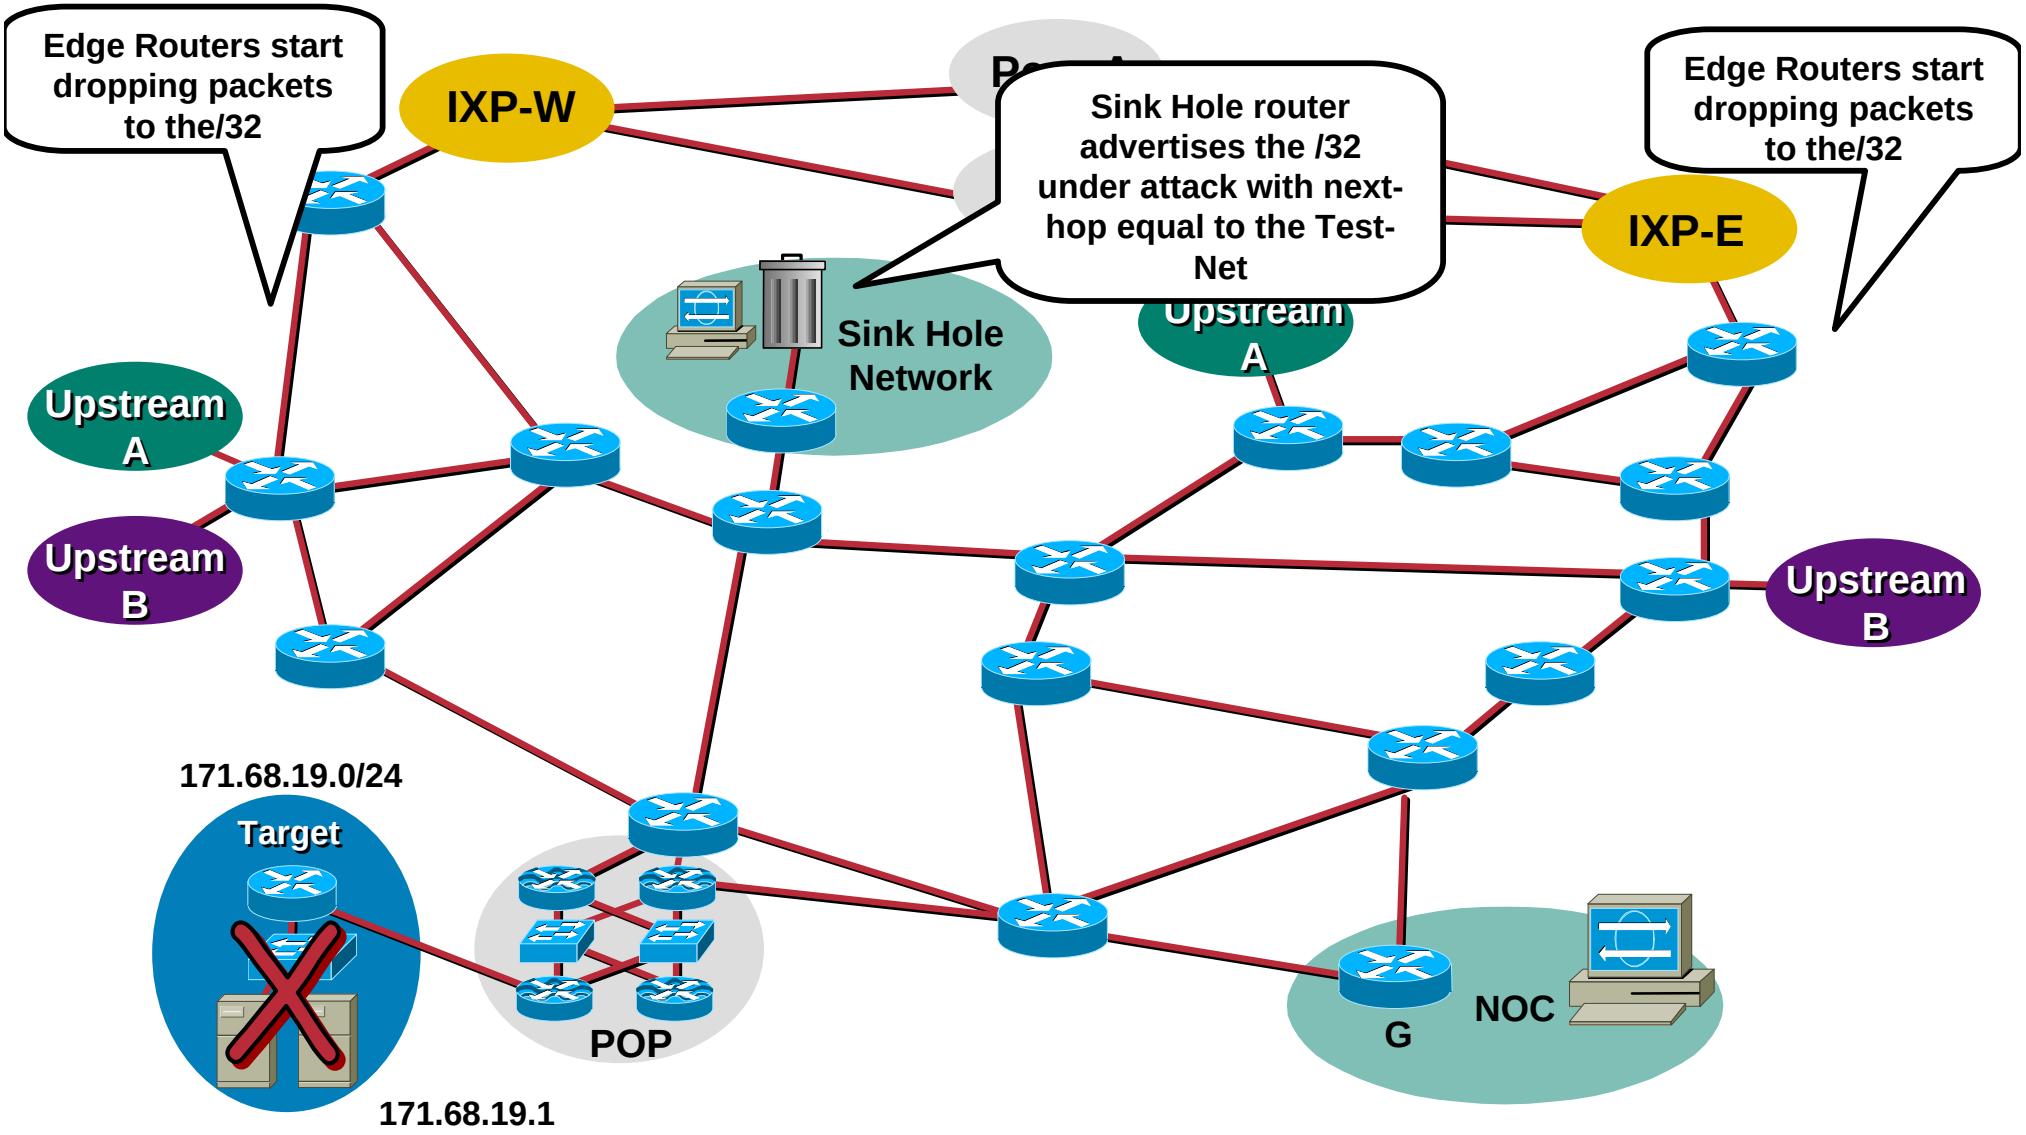

Backscatter Traceback *Activation*

Cisco.com

2. **Black Hole Filtering is triggered by BGP through out the network. Packets to the target get dropped. ICMP Unreachable Backscatter starts heading for 96.0.0.0/3.**

Access list is used on the router to find which routers are dropping packets.

```
access-list 101 permit icmp any any unreachables log

access-list 101 permit ip any any
```

Backscatter Traceback *Activation*

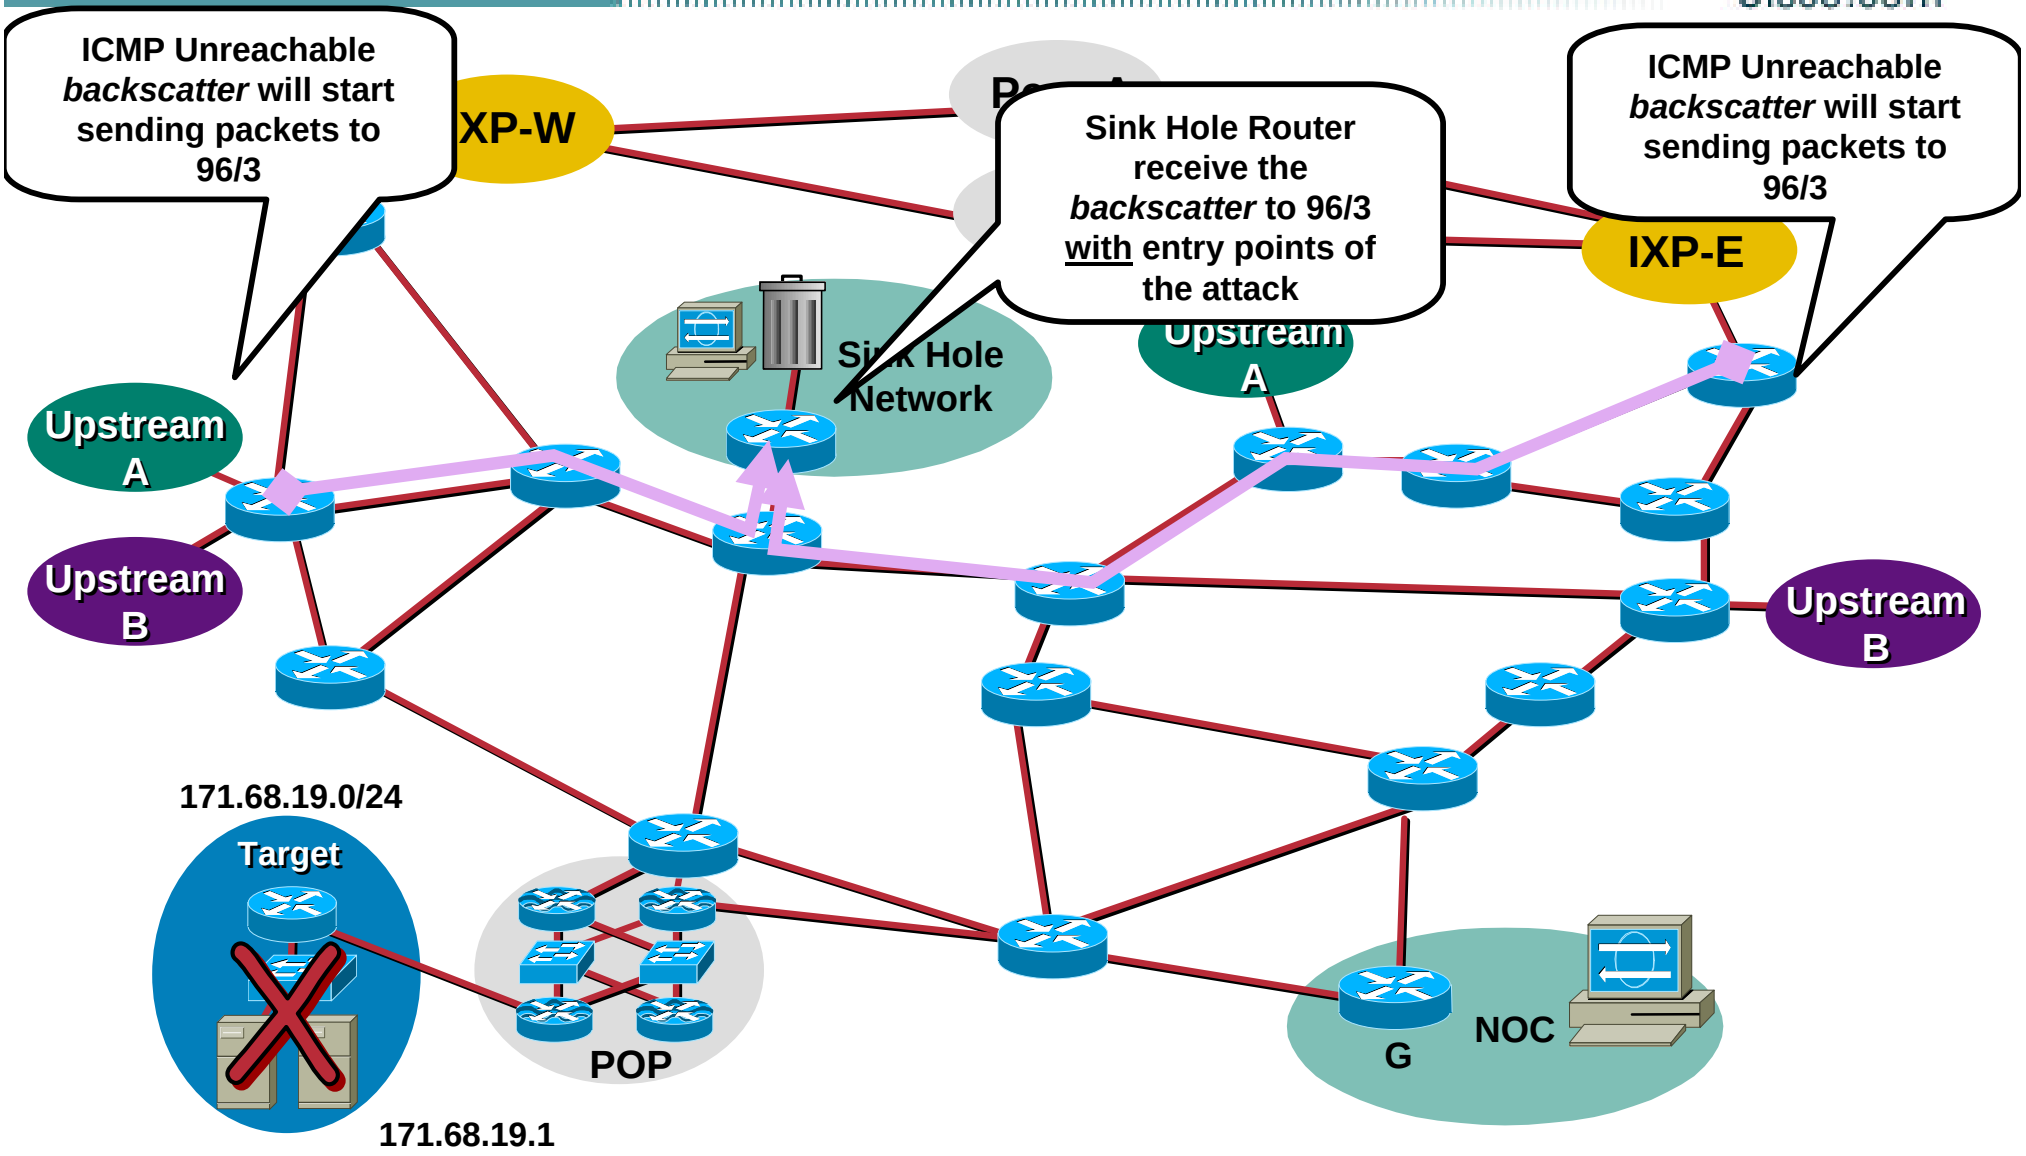

Backscatter Traceback *Activation*

Cisco.com

SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.47.251.104 (3/1), 1 packet

SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.70.92.28 (3/1), 1 packet

SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.222.127.7 (3/1), 1 packet

SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.96.223.54 (3/1), 1 packet

SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.14.21.8 (3/1), 1 packet

SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.105.33.126 (3/1), 1 packet

SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.77.198.85 (3/1), 1 packet

SLOT 5:3w1d: %SEC-6-IPACCESSLOGDP: list 150 permitted icmp 171.68.66.18 -> 96.50.106.45 (3/1), 1 packet

Questions

- **Pulling down all the traffic into a Sink Hole could be very dangerous.**

  Yes. Make sure you've integrated in the network so when it melts down, it will not impact the network.

- **Advertising large chunks of address space (I.e. 64/8) to do the backscatter traceback could be dangerous.**

  Murphy's Law of Networking – Layered checks should be used – Egress BGP filtering + no-export community.

Demo Time

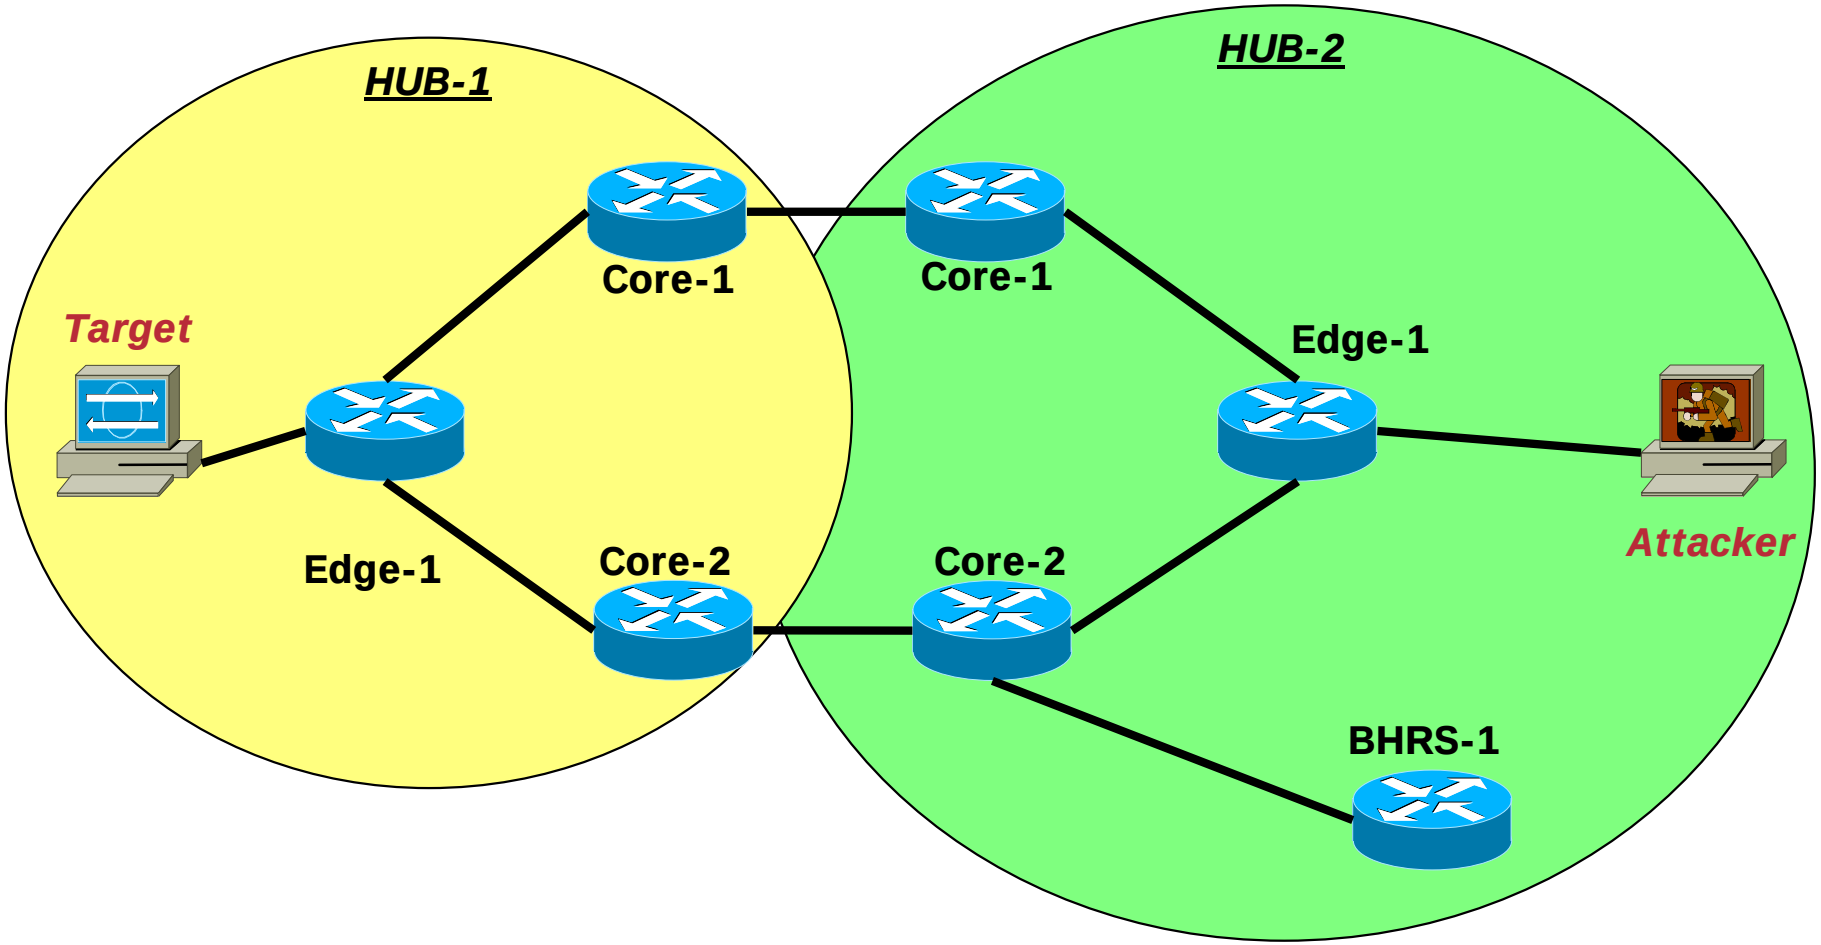

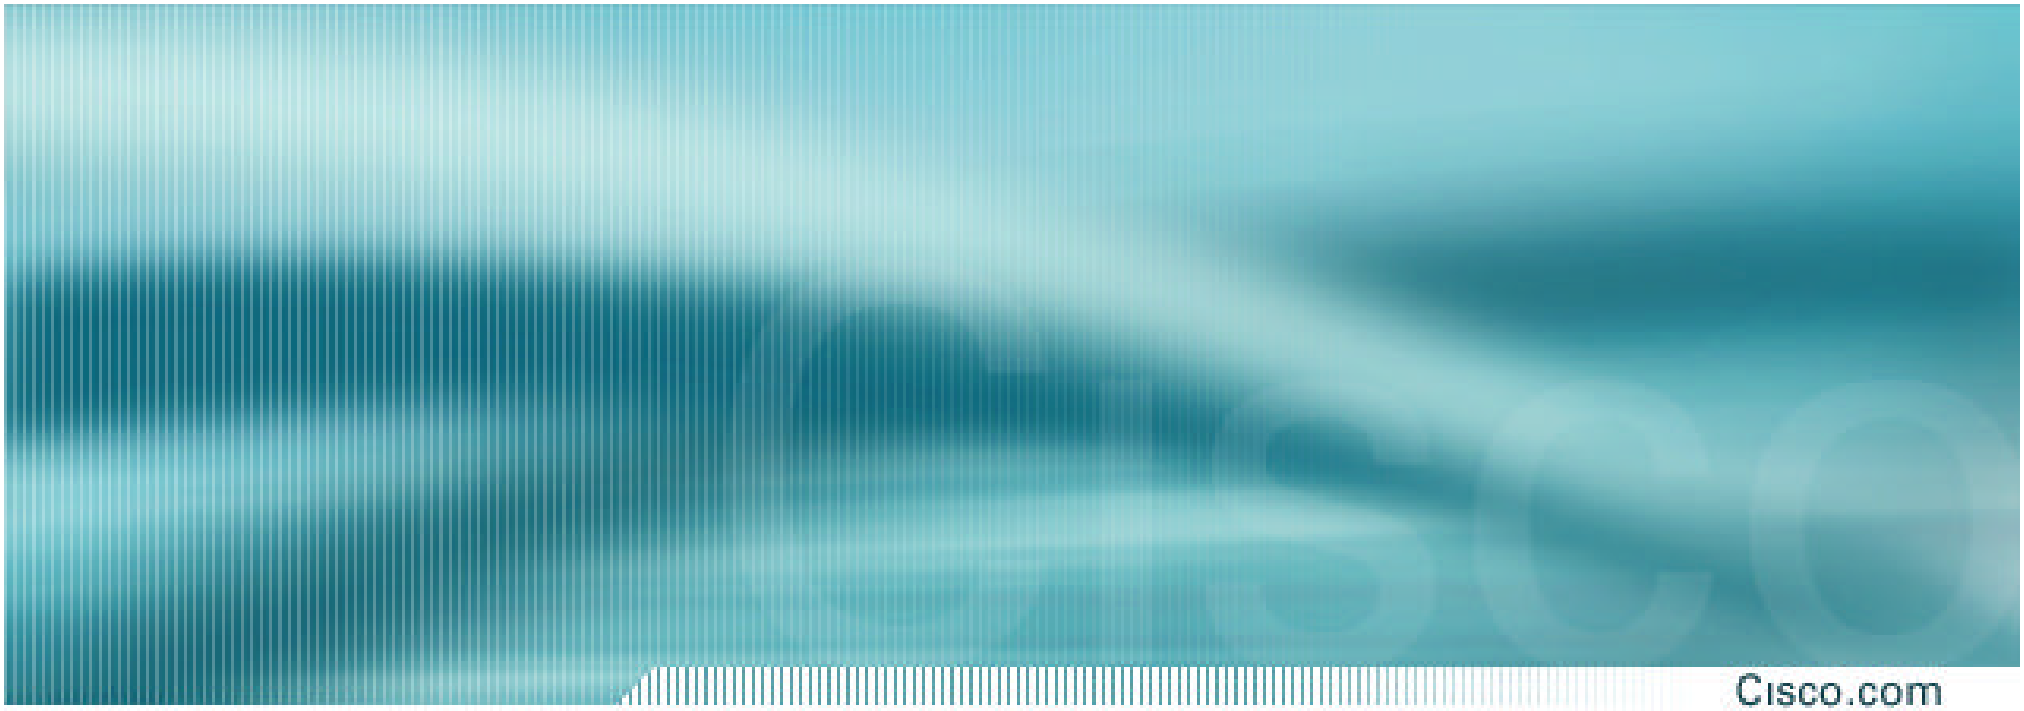

Phase 5 – Reacting to the Attack

Phase 5 - React to the Attack

Cisco.com

- **Do something to mitigate the impact of the attack OR stop the attack**
  - **Options can be everything from do nothing (doing something might cause other problems) to unplug from the source of the attack (another country during a cyberwar attack)**
- **Most ISPs try to help their customers**
  - **Rate-limit the attack**
  - **Drop the packets based on a list of source addresses**
- **Reactions need to be fast and flexible**

Phase 5 - React to the Attack

Cisco.com

- **Three techniques used to drop or rate limit:**
  - **ACLs—Manual upload**
  - **uRPF—Remote trigger via BGP**
  - **CAR—Manual upload or remote trigger via BGP**

Reacting to an Attack with ACL

- Traditional mode of stopping attacks
- Scaling issues encountered:
  - Updates of ACLs on many many routers a pain
  - Additive ACLs when there are multiple attacks on multiple customers are a pain
  - Confusion with the "Line Rate Debate"

Reacting to an Attack with uRPF

Cisco.com

- uRPF loose check mode can be used on the ISP®ISP edge
- Can be used remote trigger drops of a DOS/DDOS flow
- Allows many many routers to be simultaneously updated with a new drop list all via a routing protocol
- Effect L3 filter (source and destination address)

Reacting to an Attack with uRPF

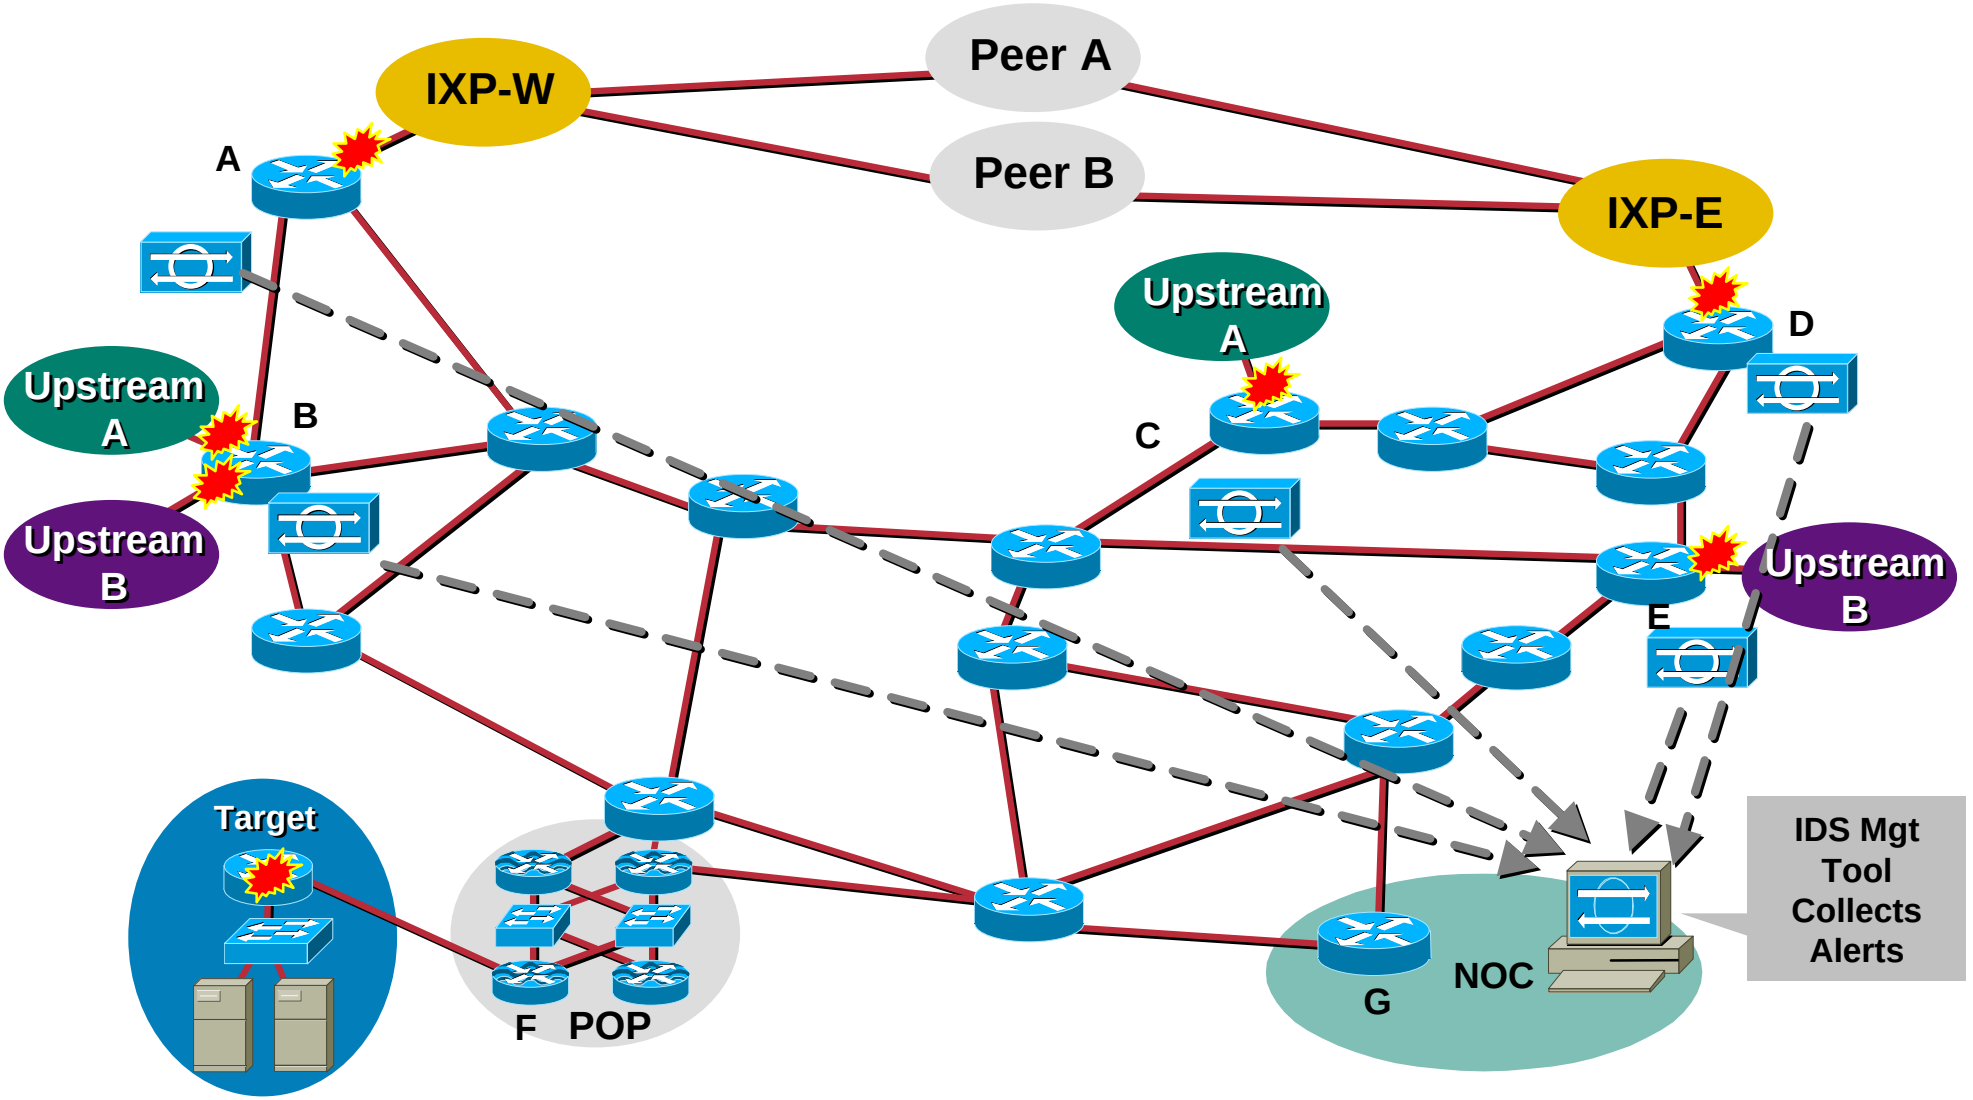

Reacting to an Attack with uRPF

Cisco.com

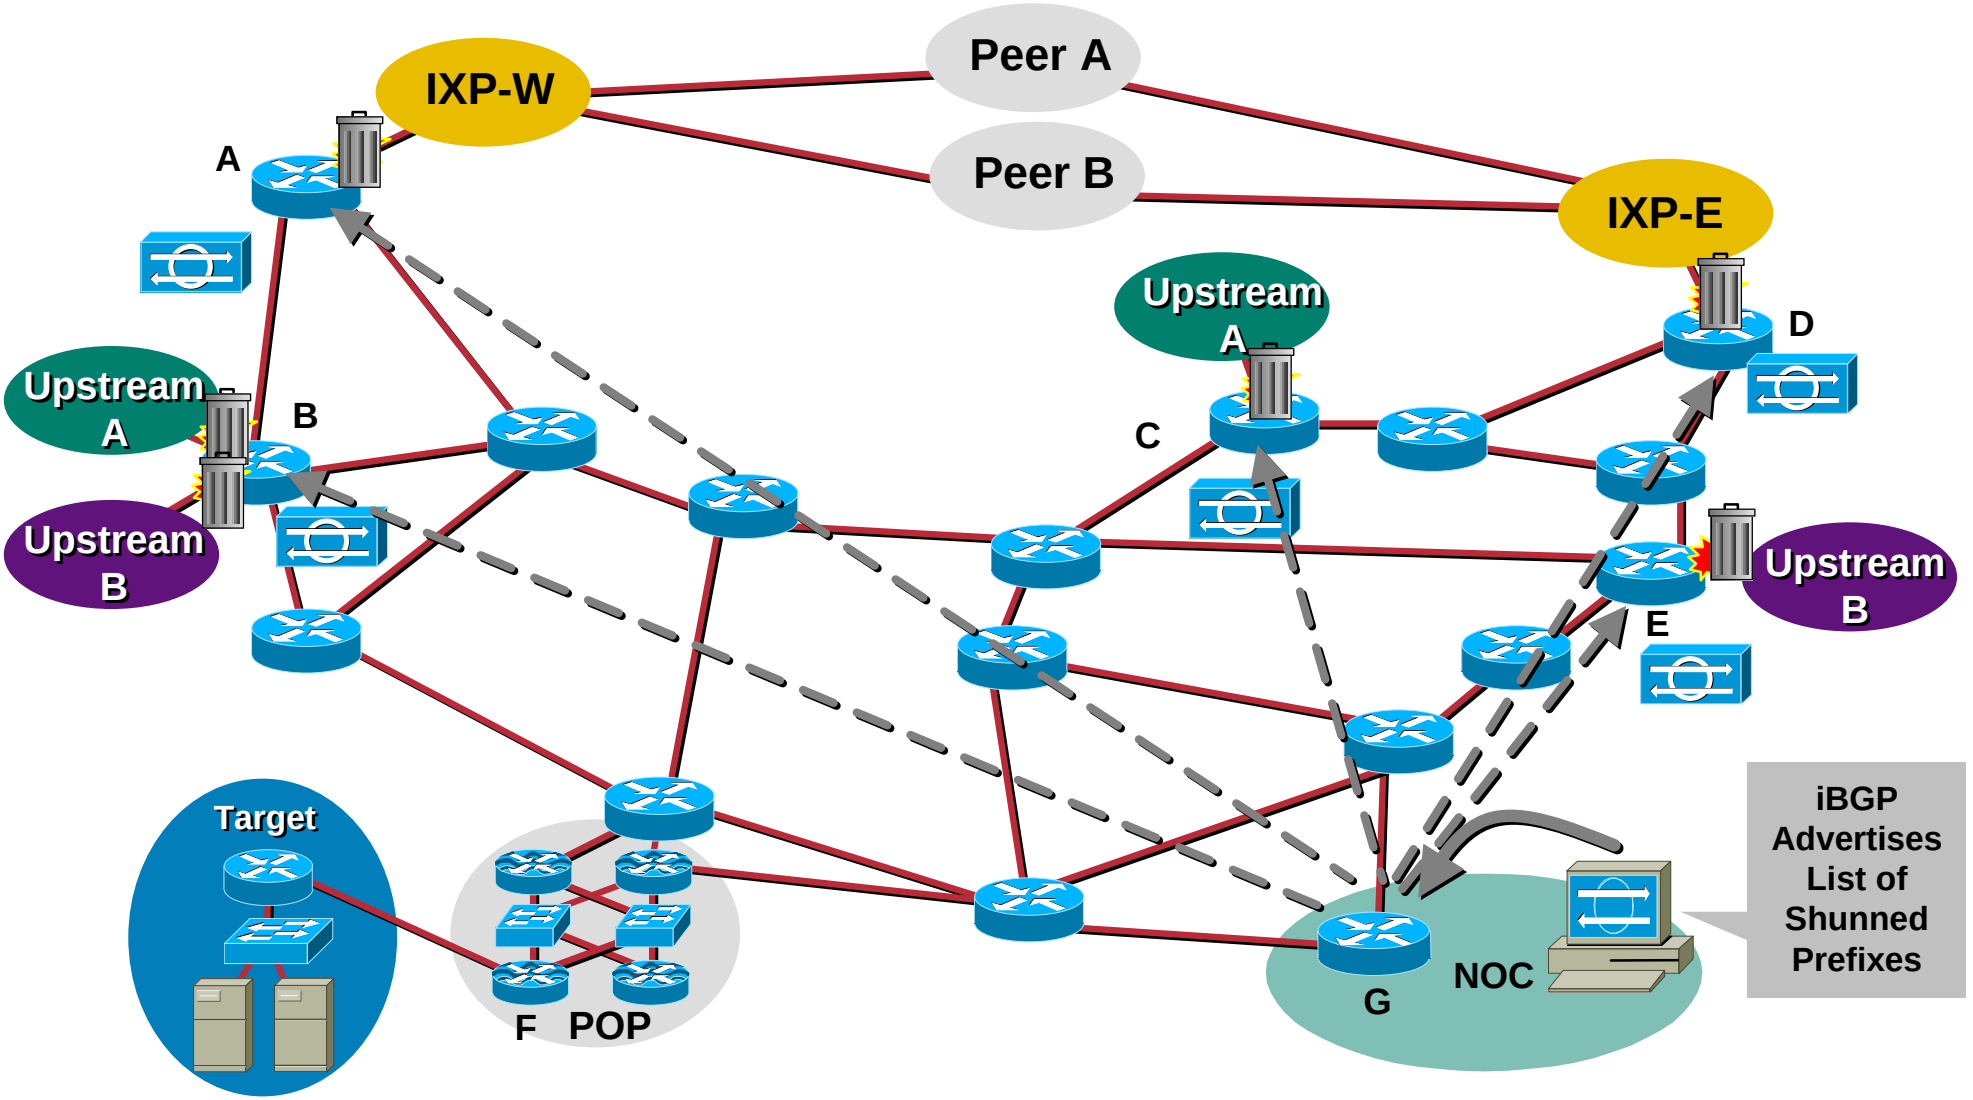

Reacting to an Attack with uRPF

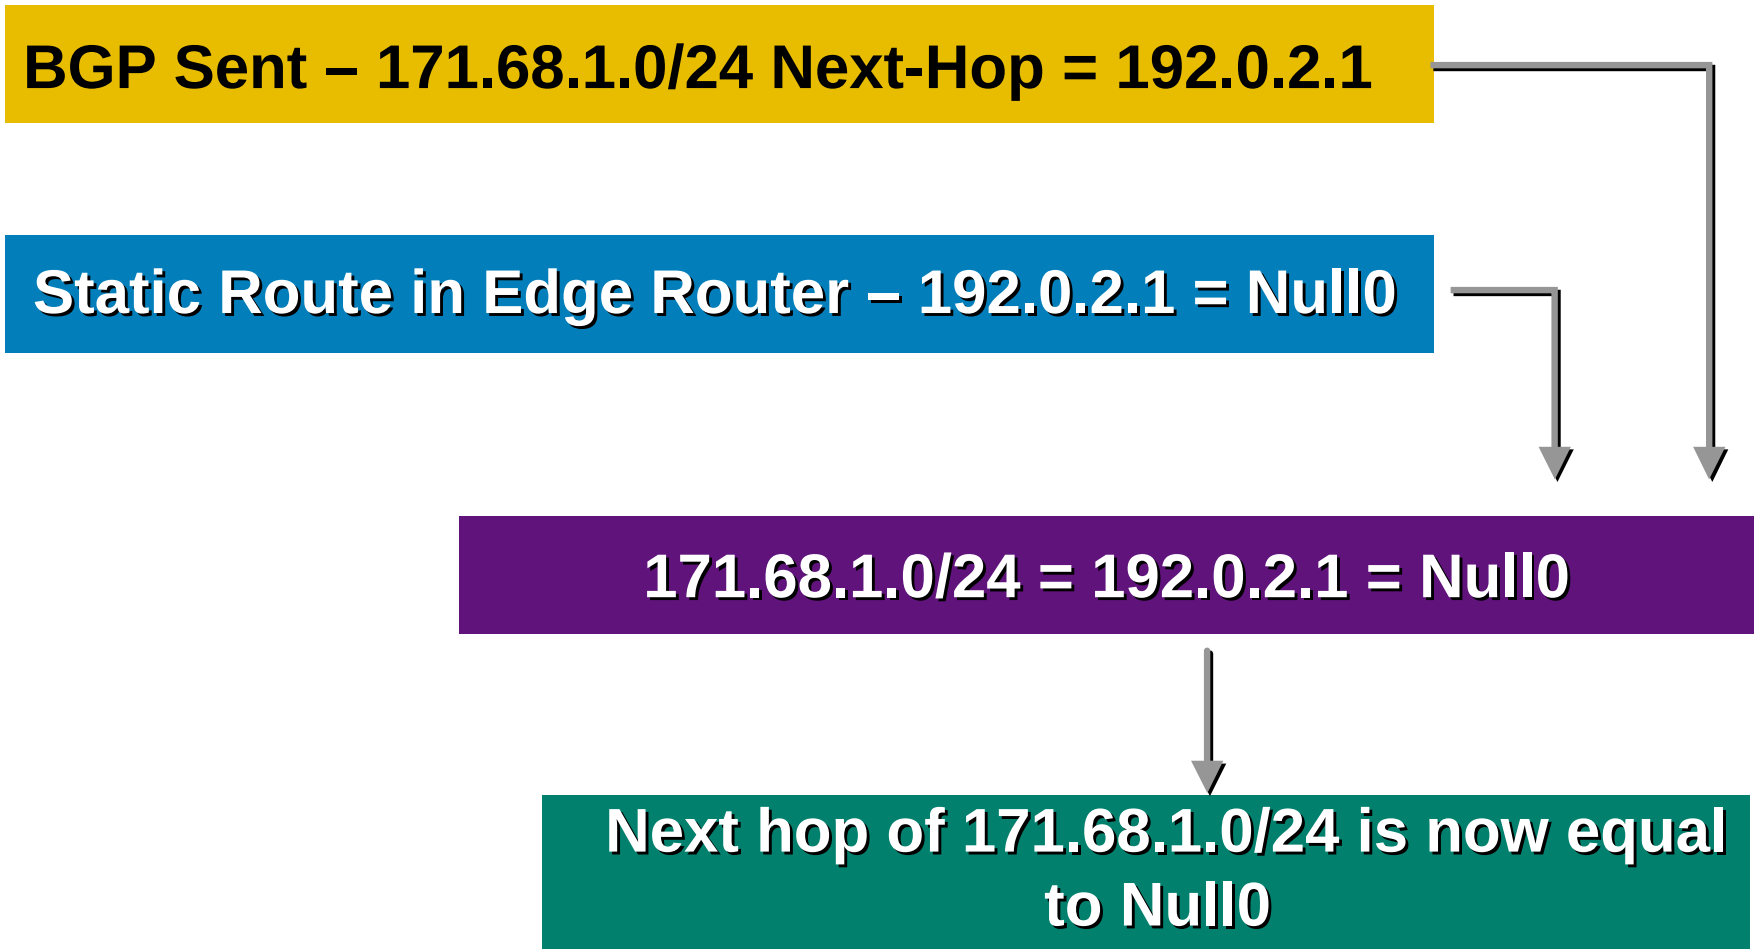

Reacting to an Attack with uRPF

- **What is needed?**

  uRPF loose check on all border routers

  Static to Null0 with an address like the test-net on all border routers

  Way to inject a BGP advertisement into the network with a BGP community that will trigger the drop; (should include the no-export community and have good egress router filters)

Reacting to an Attack with uRPF

- **Key advantages:**
  - **No ACL update**
  - **No change to the router's config**
  - **Drops happen in the forwarding path**
  - **Frequent changes when attacks are dynamic (or multiple attacks on multiple customers)**

Reacting to an Attack with CAR

Cisco.com

- **CAR and other rate-limit features have proven to be an effective reaction to an attack**
  - **Rate limiting attacks allow the attack to be monitored**
  - **Data collection for law enforcement evidence can continue with rate limiting**
  - **QOS group support (QPPB) allows for remote triggering of CAR with out logging into the router**

Reacting to an Attack with CAR

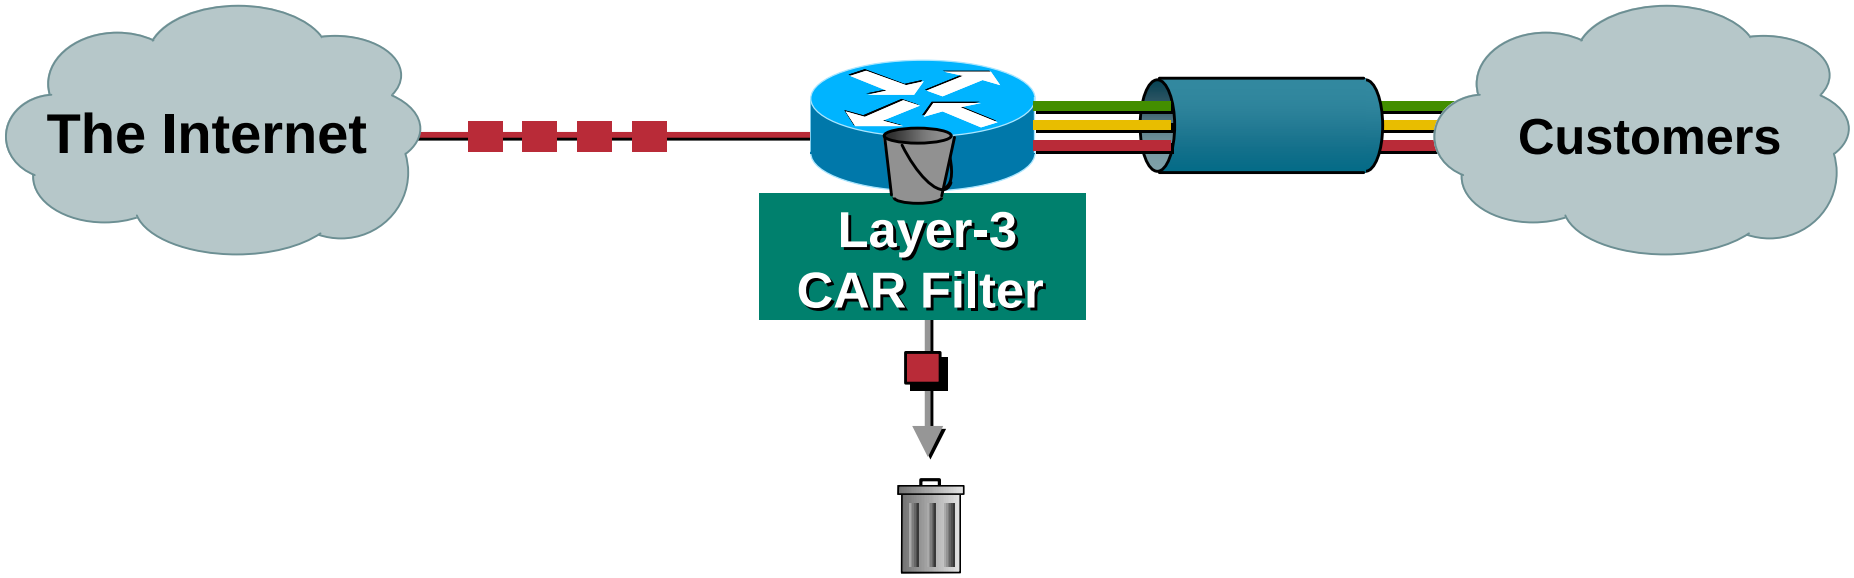

- Layer-3 input and output rate limits ® specifically input rate limits
- Security filters use the input rate limit to drop packets before there are forwarded through the network
- Aggregate and granular limits

  Port, MAC address, IP address, application, precedence, QOS ID
- Excess burst policies

Reacting to an Attack with CAR

- **Limit all ICMP echo and echo-reply traffic received at the borders to 256 Kbps with a small amount of burst:**

```
! traffic we want to limit
access-list 102 permit icmp any any echo
access-list 102 permit icmp any any echo-reply
! interface configurations for borders
interface Serial3/0/0
 rate-limit input access-group 102 256000 8000 8000
conform-action transmit exceed-action drop
```

- **Multiple “rate-limit” commands can be added to an interface in order to control other kinds of traffic as well**

Reacting to an Attack with CAR

- **Use CAR to limit TCP SYN floods to particular hosts—without impeding existing connections; some attackers have started using very high streams of TCP SYN packets in order to harm systems**
- **This example limits TCP SYN packets directed at host 10.0.0.1 to 8 kbps or so:**

```
! We don't want to limit established TCP sessions -- non-SYN
packets

access-list 103 deny tcp any host 10.0.0.1 established

! We do want to limit the rest of TCP (this really only
includes SYNs)

access-list 103 permit tcp any host 10.0.0.1

! interface configurations for network borders

interface Serial3/0/0

 rate-limit input access-group 103 8000 8000 8000 conform-
action transmit exceed-action drop
```

Reacting to an Attack with CAR with Remote Trigger

- CAR's rate limiting has proven to be an effective reaction tool to a DOS/DDOS attack
- The problem is how do quickly update +60 routers on the ingress of a network—especially when the attack character shifts to respond to your countermeasures?
- Answer—CAR is a FIB entry-based feature (CEF feature); so we can use a network protocol to trigger the rate limits on source/destination

Reacting to an Attack with CAR with Remote Trigger

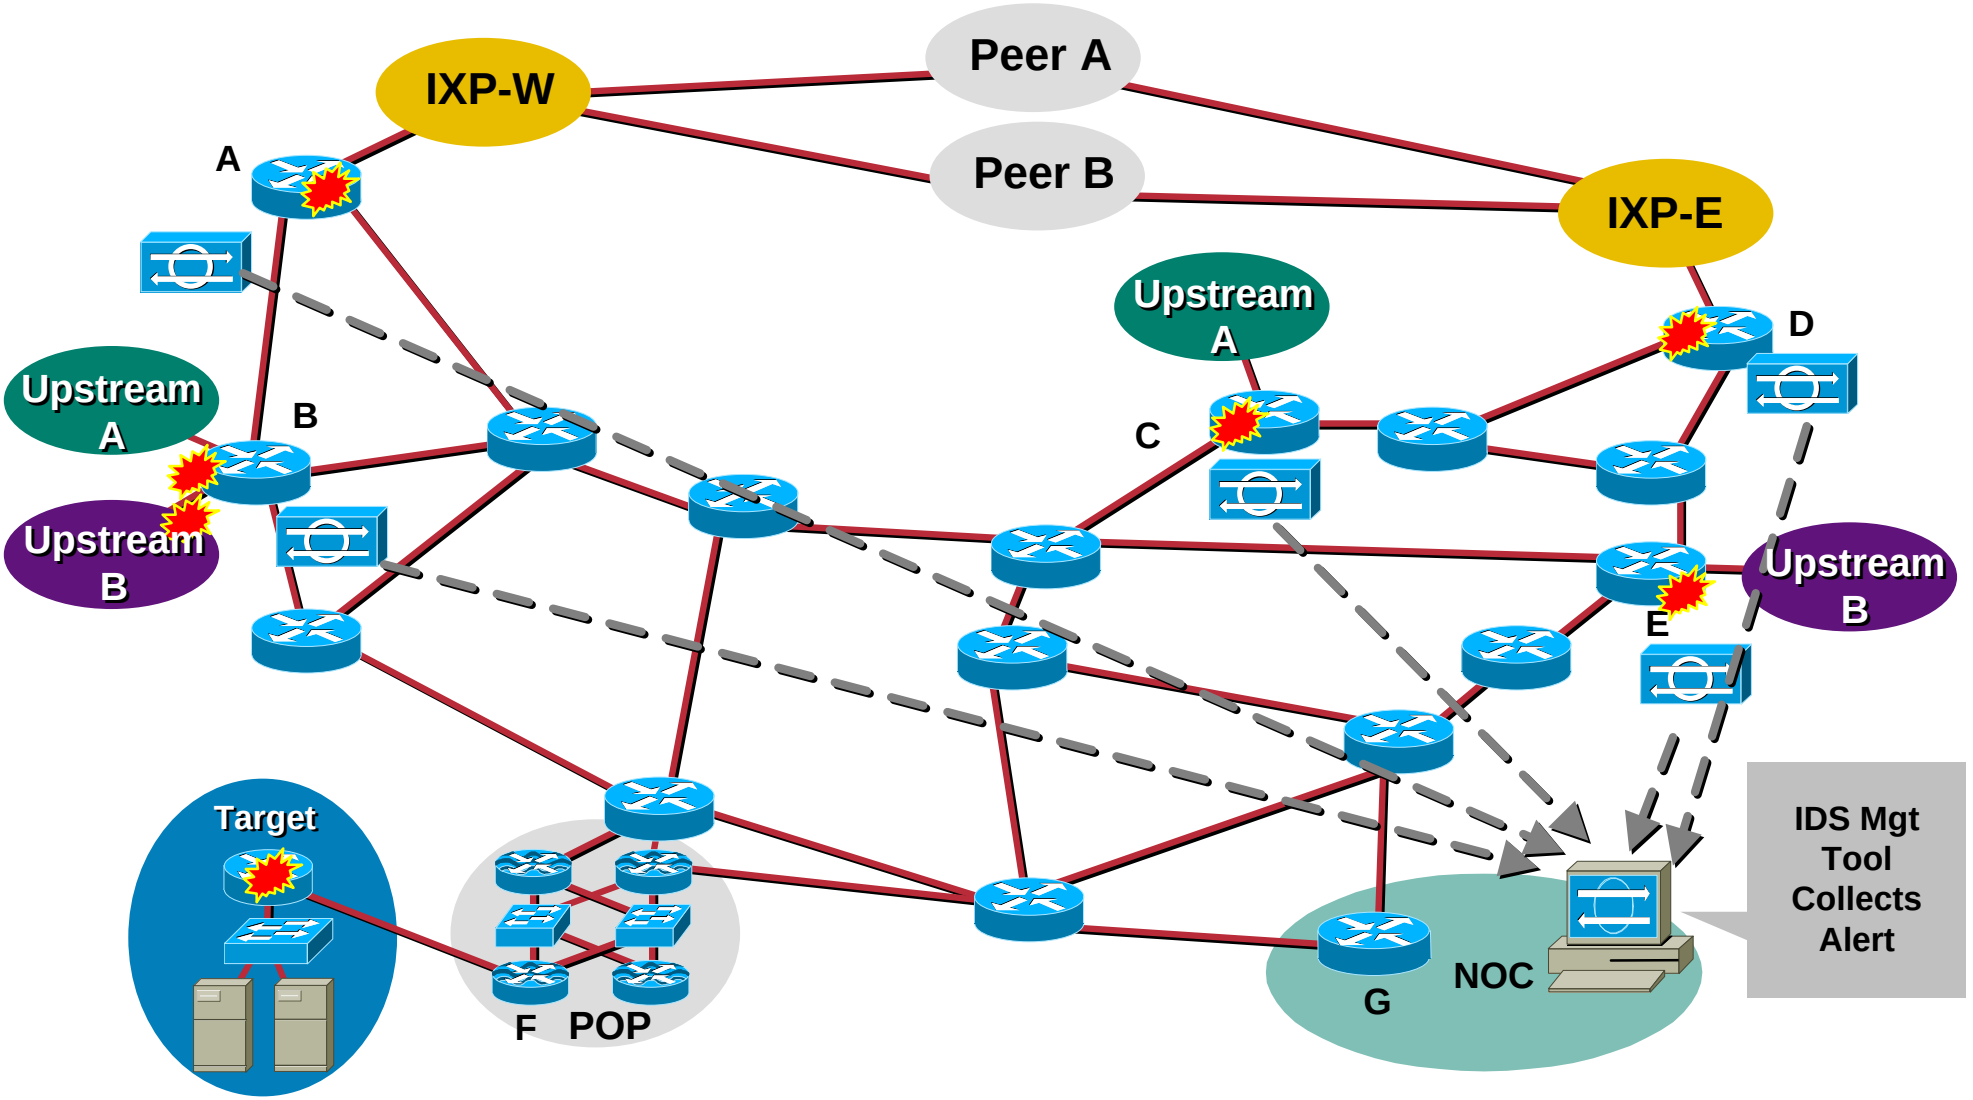

Reacting to an Attack with CAR with Remote Trigger

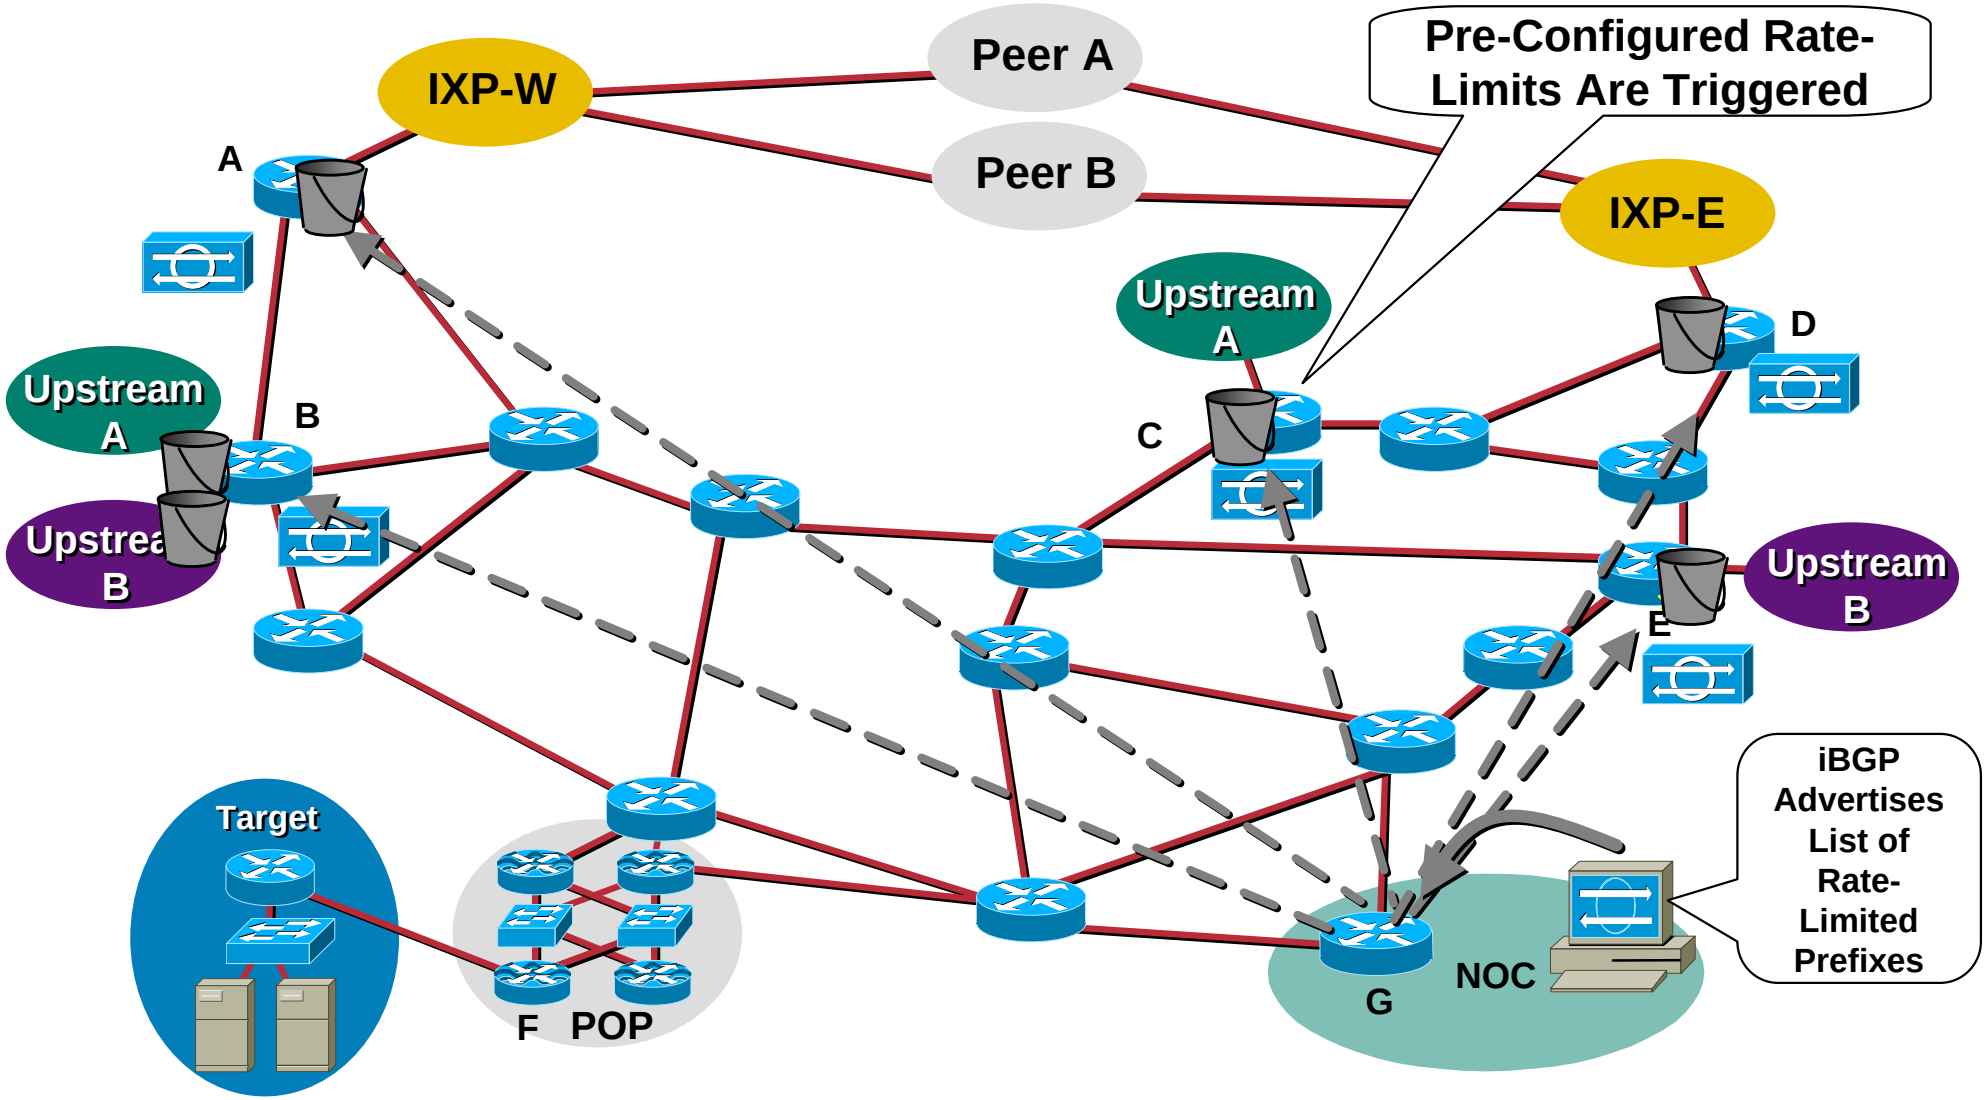

Reacting to an Attack with CAR with Remote Trigger

- Conveys IP precedence to be used in forwarding to specified destination prefix via BGP community tag
- Allows ingress routers to prioritise incoming traffic
- Also allows IP precedence setting based on AS-path attribute or access list
- Inter-ISP Service Level Agreements (SLAs)

Reacting to an Attack with CAR with Remote Trigger

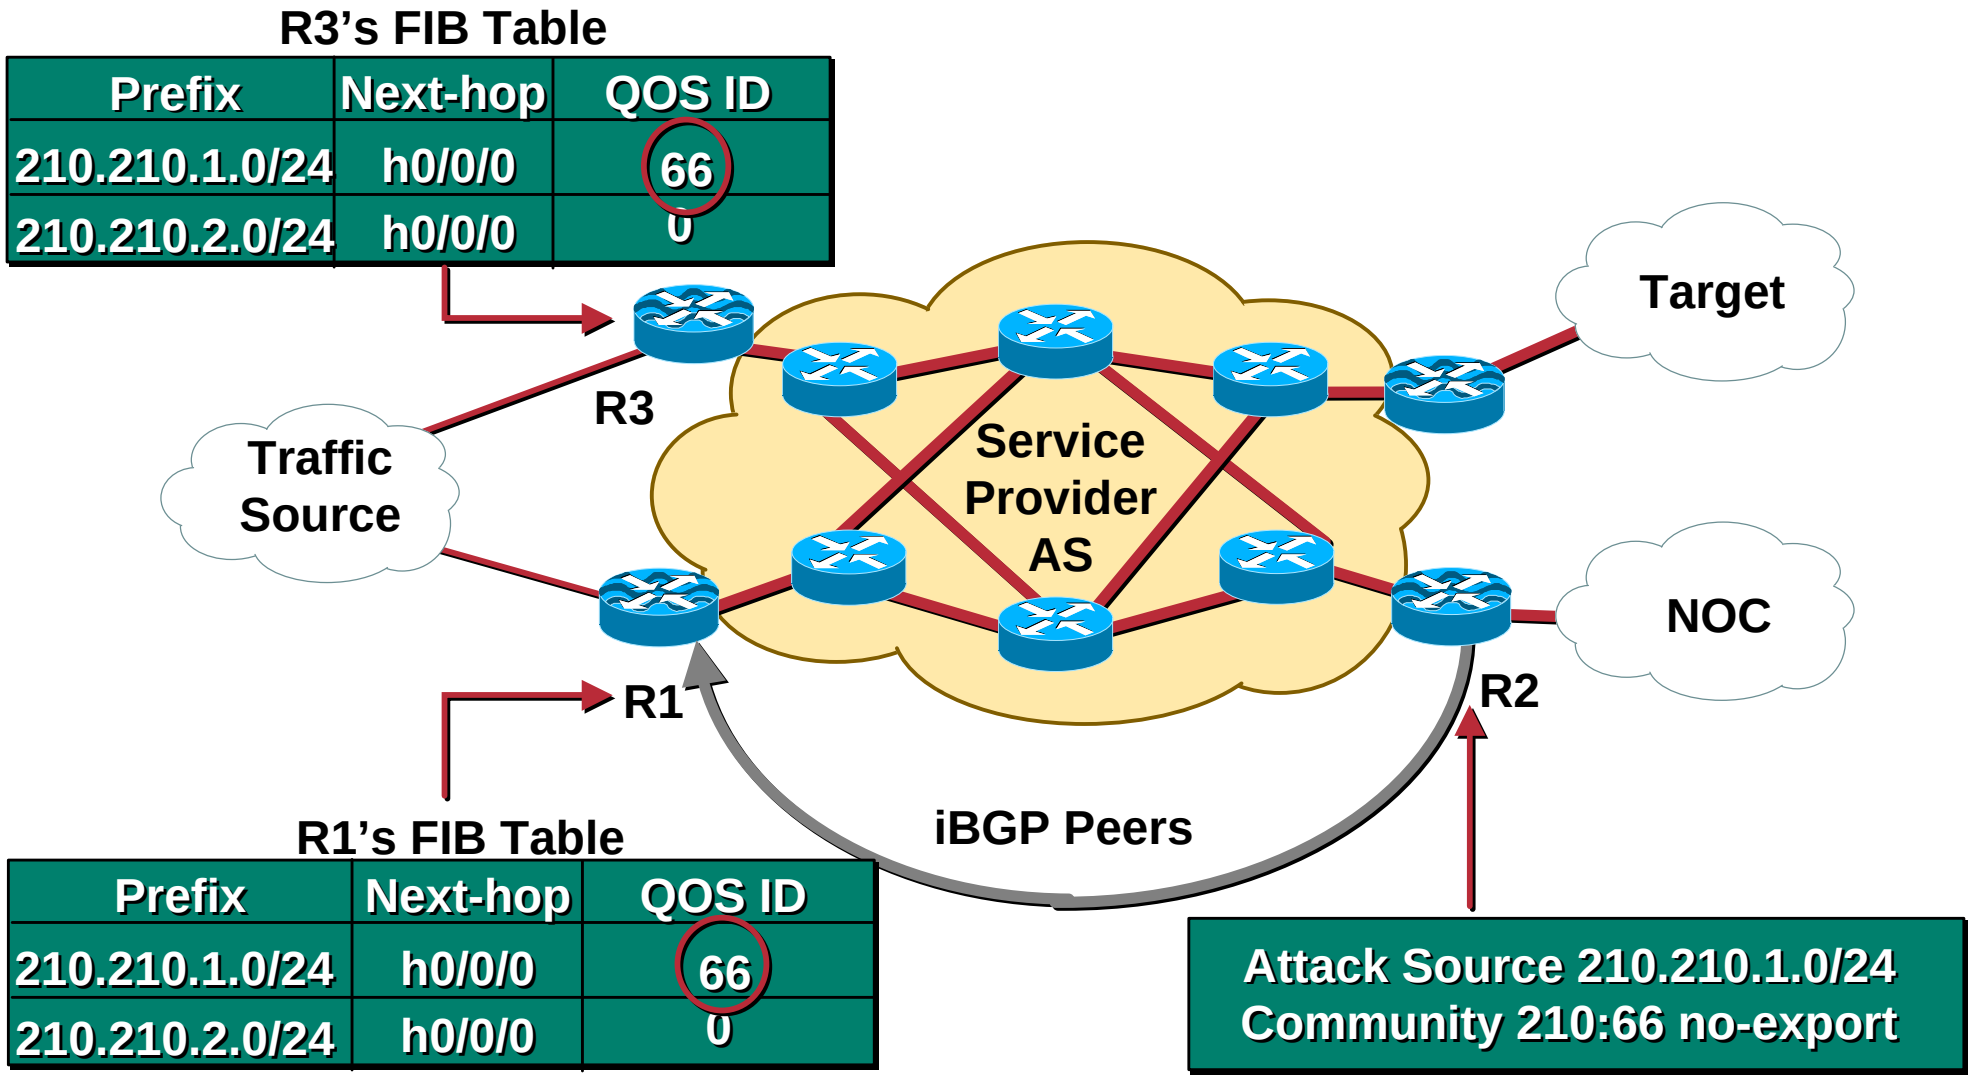

Reacting to an Attack with CAR with Remote Trigger

- **NOC-Router#write term**

```
router bgp 210
 network  210.210.1.0 mask 255.255.255.0
 neighbor 210.210.14.1 remote-as 210
 neighbor 210.210.14.1 route-map DOS-Trigger out
 neighbor 210.210.14.1 send-community
!
ip bgp-community new-format
!
ip route 210.210.1.0 255.255.255.0 Null0 254

access-list 1 permit 210.210.1.0 0.0.0.255
!
route-map DOS-Trigger permit 10
 match ip address 1
 set community 210:66 no-export
!
route-map DOS-Trigger permit 20
```

Note: There Are Other Ways to Originate a Prefix

Reacting to an Attack with CAR with Remote Trigger

- **R1#write term**

```
!
router bgp 210
 table-map DOS-Activate
 neighbor 200.200.14.4 remote-as 210
 neighbor 200.200.14.4 update-source Loopback0
!
ip bgp-community new-format
!
ip community-list 1 permit 210:66
!
route-map DOS-Activate permit 10
 match community 1
 set ip qos-group 66
!
route-map DOS-Activate permit 20
!
```

Directly Updates QOS_ID in the FIB

Matches Community and Sets the QOS Group

Reacting to an Attack with CAR with Remote Trigger

- **Router 1 (cont.):**

```
!
interface HSSI 0/0/0
bgp-policy source ip-qos-map
rate-limit input qos-group 66 256000 8000
8000 conform-action transmit exceed-action
drop
```

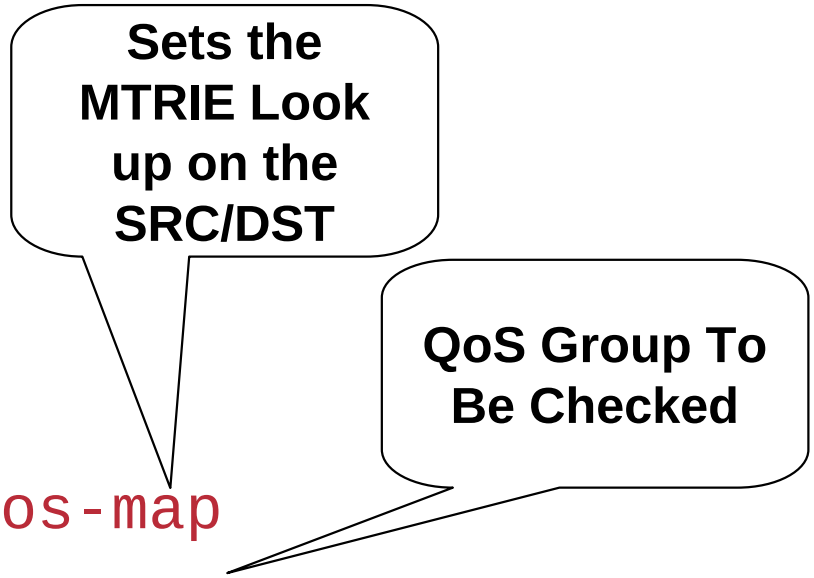

Reacting to an Attack with CAR with Remote Trigger

- **Caveats with CAR:**
  - **Not all platforms support the full version of CAR (I.e. Engine 2)**
  - **Not all platforms support the full version of QoS group (QPPB)**
  - **Some platforms have specialized rate limiting ASICs (7600)**
- **Bottom-line—CAR is not yet cross platform compatible (working on it)**

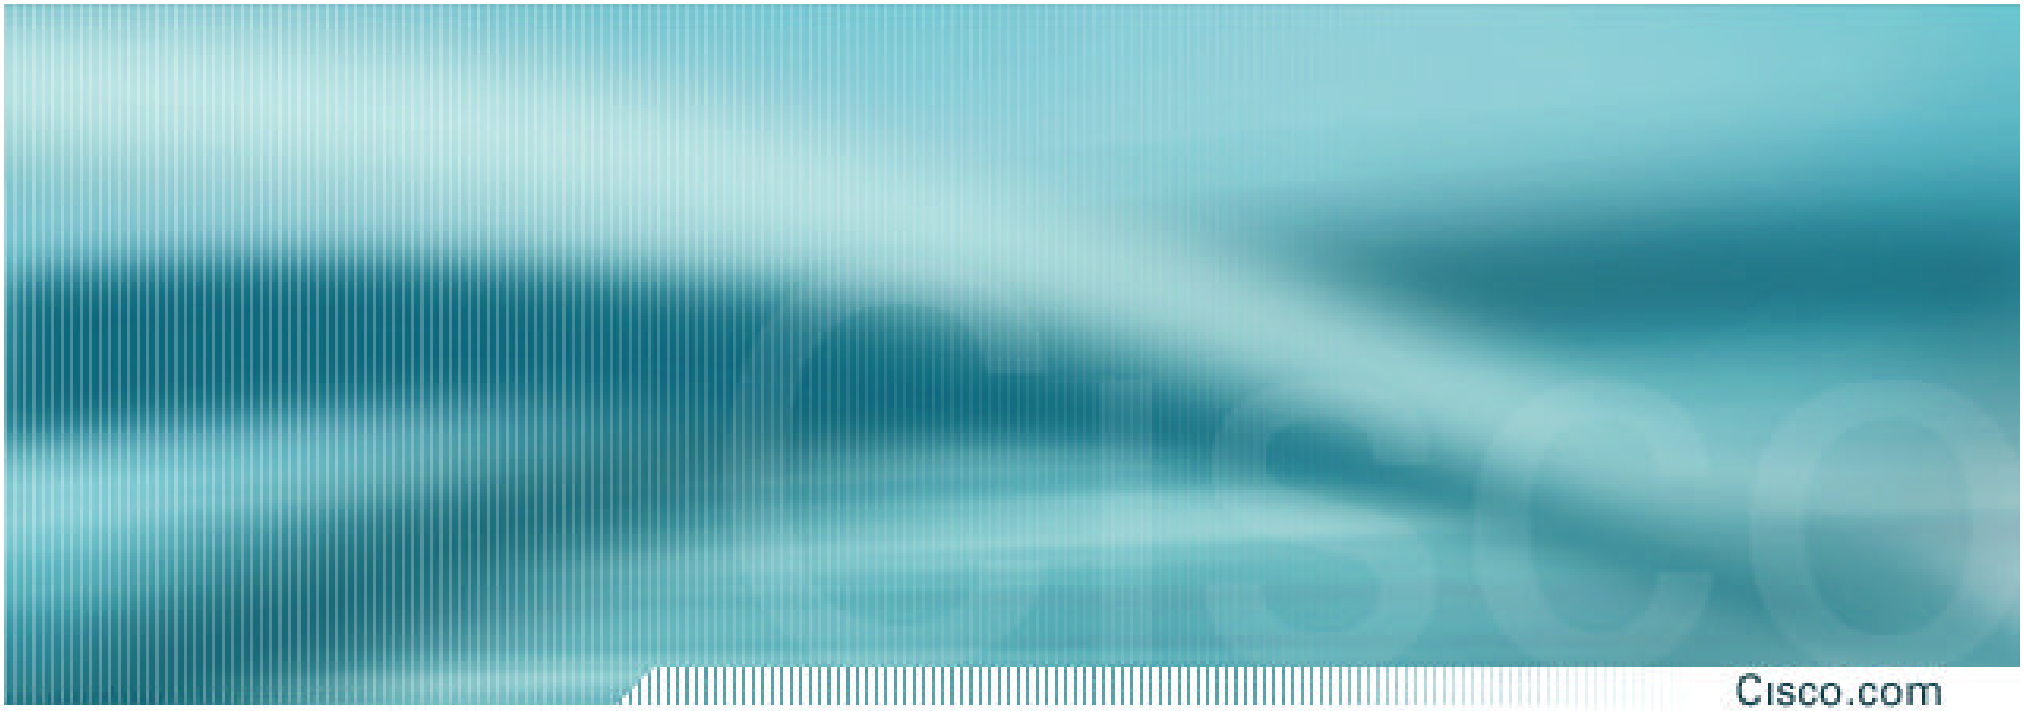

Phase 6 - Post Mortem

Post Mortem

- **Learning from your mistakes is essential.**
- **Do not wait until the next attack to implement the lessons of the last attack.**

  Take time after each incident to see if processes, procedures, tools, techniques, and configurations can be improved.

  It is an arms race. Those who learn from this mistakes excel.

Post Mortem

Cisco.com

- ***Fighting the Last War*** **is the #2 mistake of military planner.**
- **Underestimating the capabilities and commitment of your enemy is the #1 mistake of military planners.**
- **This observation directly applies to ISP Security.**

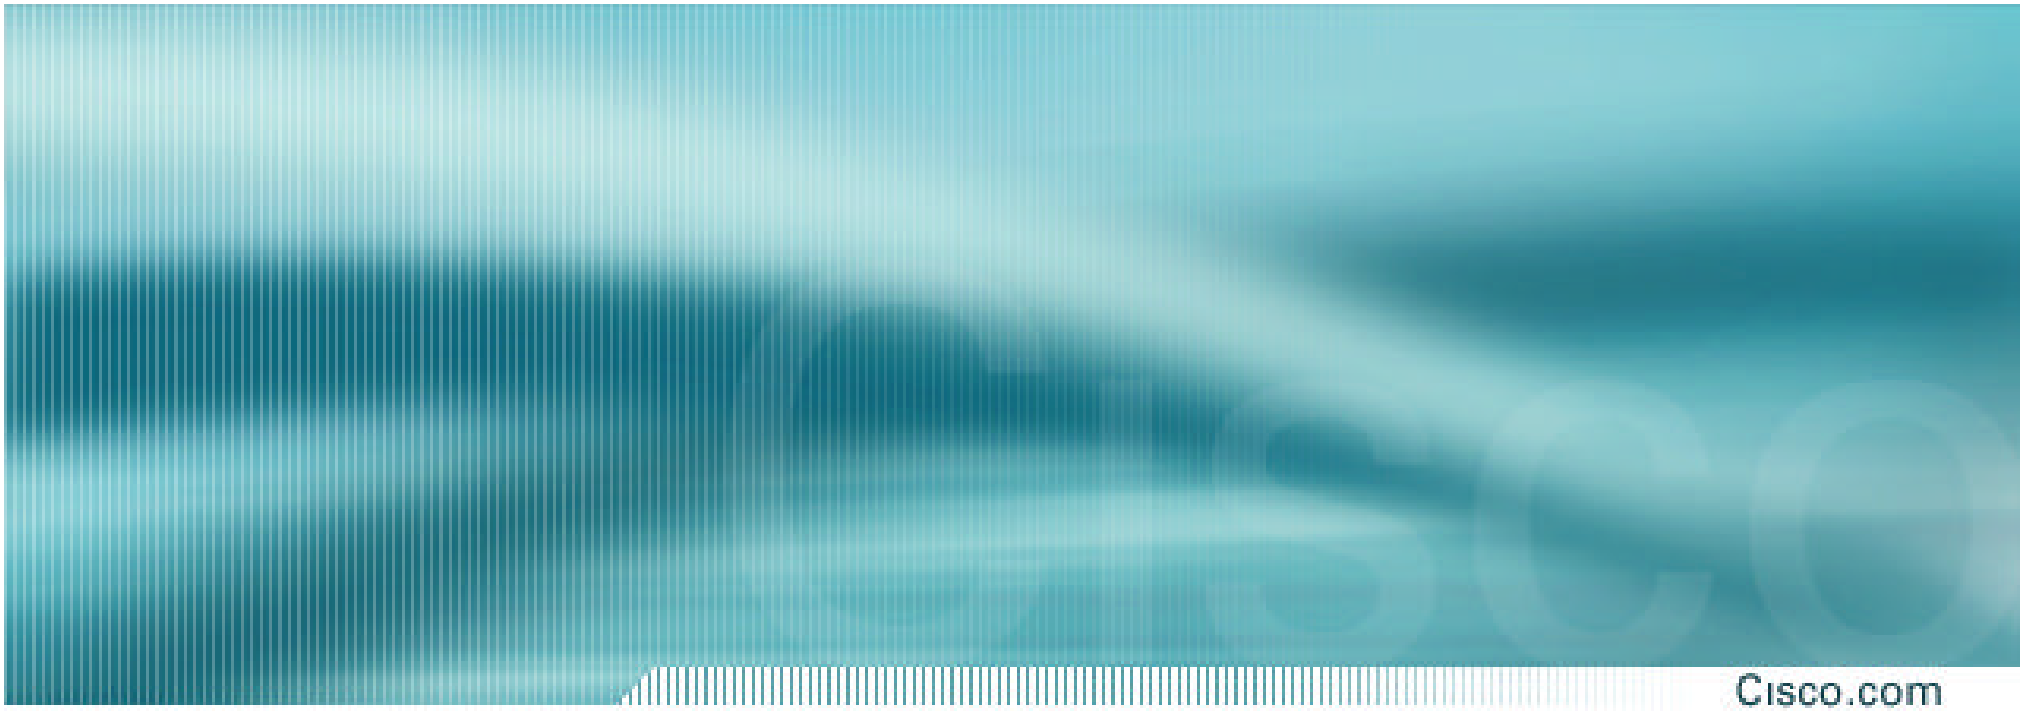

Example of an ISP Tracking DoS/DDoS Attacks through an ISP's Network

Tracking Attacks—ISP POV

Cisco.com

- **Situation in the NOC**
  - **Alarms go off in the NOC—circuits are dropping packets**
  - **Major content customer calls—their site is being hit by a DoS/DDoS attack**
  - **Management calls, they want to know what is going on**
  - **Other customers call, slow network performance**
  - **Reporter calls—not sure how they got the NOC's number, they are looking for a quote**
  - **It's been 5 minutes since the first alarm went off, what do you do?!?!?!?!**

The Network

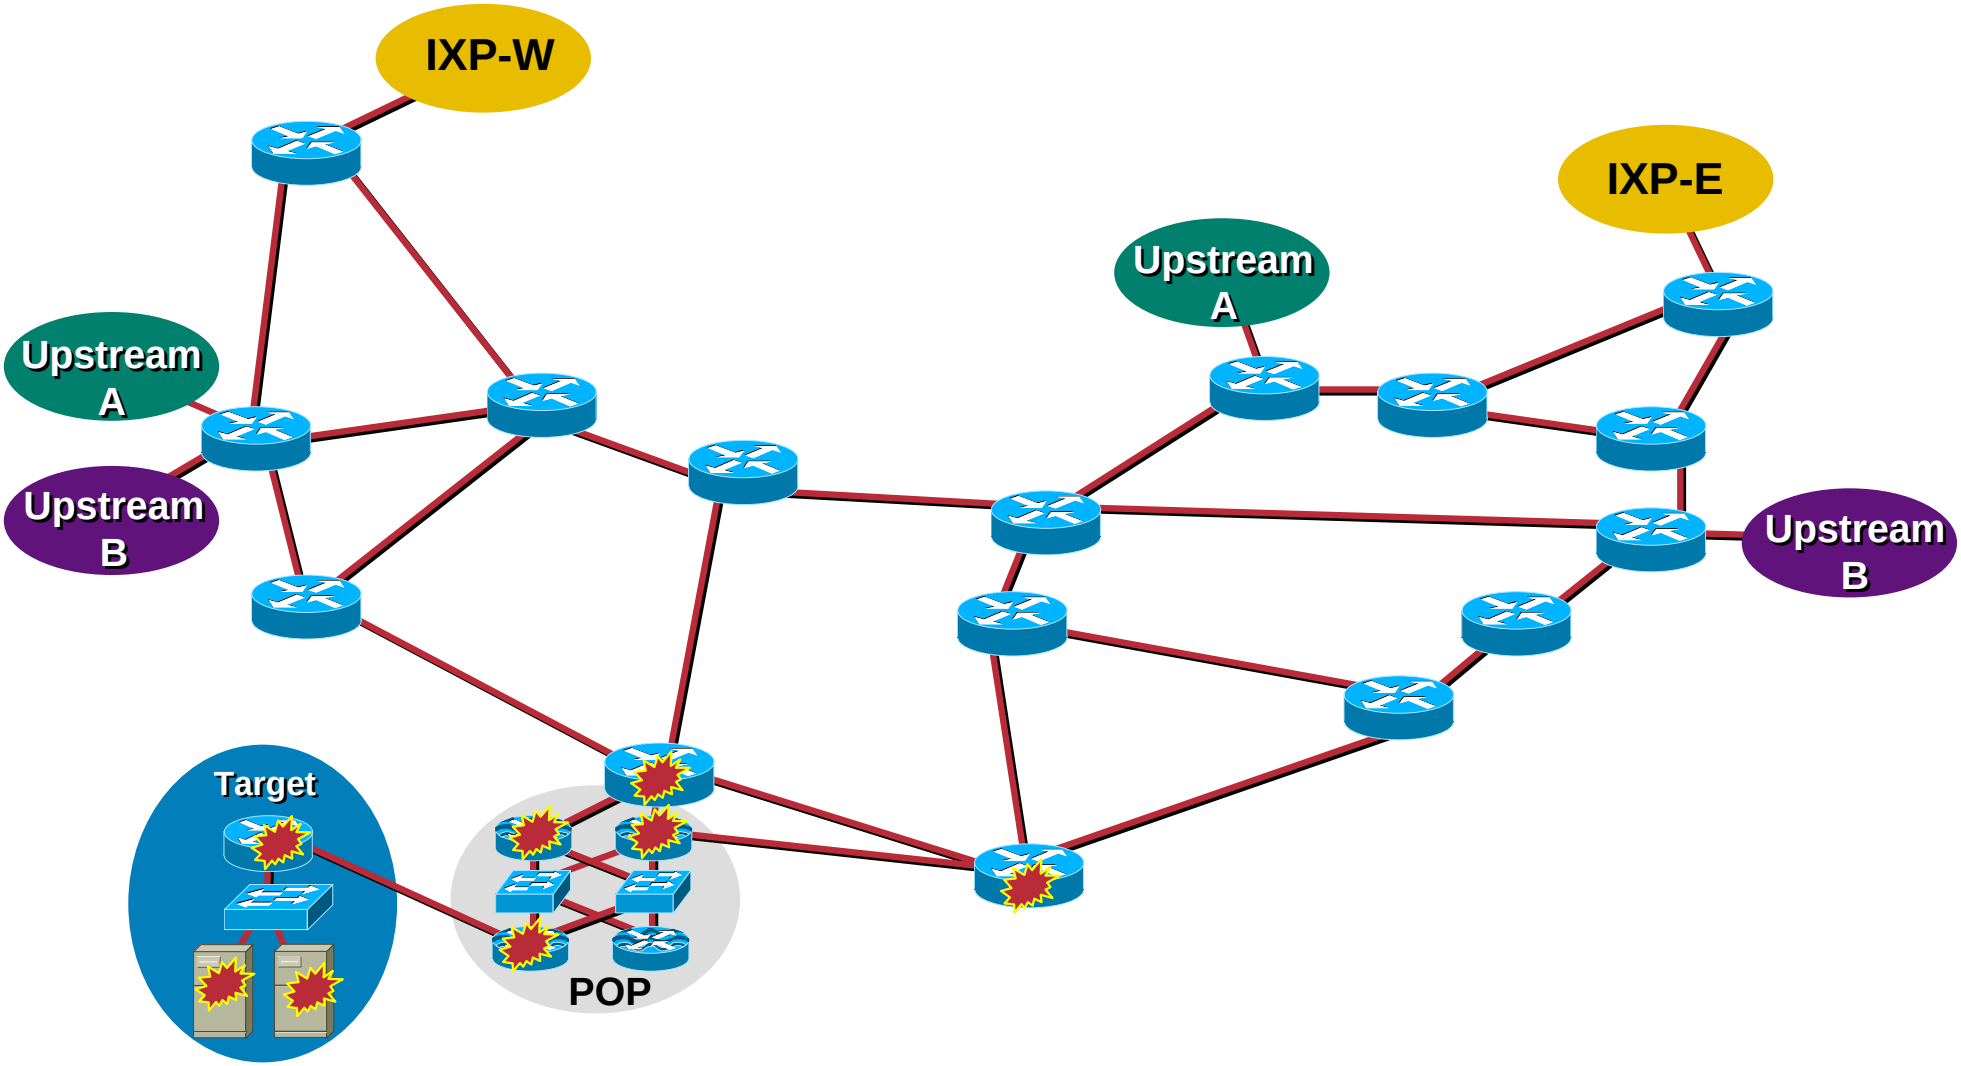

Step 1—Classifying the Attack

- **Use ACL to find out the characteristics of the attack**

```
access-list 169 permit icmp any any echo
access-list 169 permit icmp any any echo-reply
access-list 169 permit udp any any eq echo
access-list 169 permit udp any eq echo any
access-list 169 permit tcp any any established
access-list 169 permit tcp any any range 0 65535
access-list 169 permit ip any any
interface serial 0
ip access-group 169 out
```

Step 1—Classifying the Attack

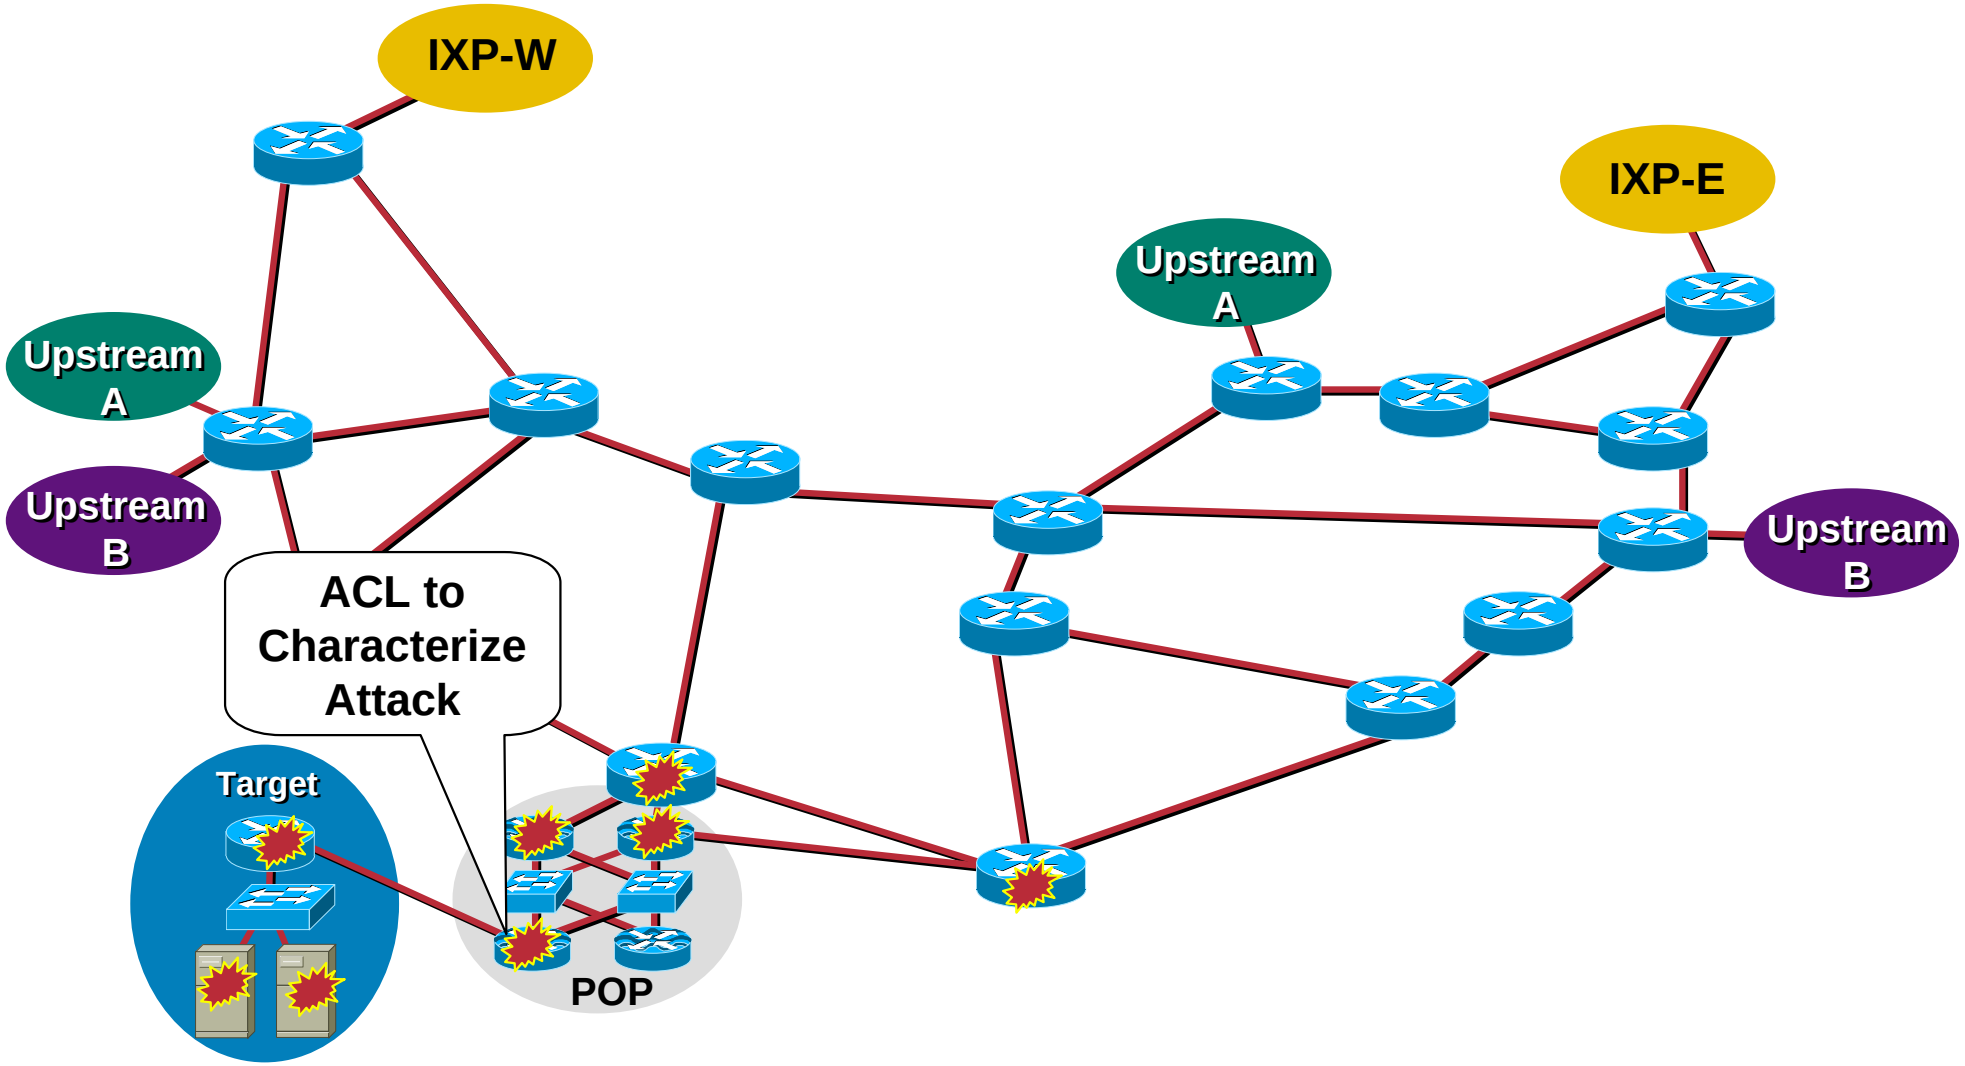

Step 1—Classifying the Attack

- **Use the show access-list 169 to see which protocol is the source of the attack:**

```
Extended IP access list 169
        permit icmp any any echo (2 matches)
        permit icmp any any echo-reply (21374 matches)
        permit udp any any eq echo
        permit udp any eq echo any
        permit tcp any any established (150 matches)
        permit tcp any any (15 matches)
        permit ip any any (45 matches)
```

Step 2—Capture a Source IP

Cisco.com

- Tracing spoofed source IP addresses are a challenge
- Tracing needs to happen hop by hop
- The first step is to use the ACL “log-input” function to grab a few packets
- Quick in and out is needed to keep the router for overloading with logging interrupts to the CPU

Step 2—Capture a Source IP

- **Preparation**

  Make sure your logging buffer on the router is large

  Create the ACL

  Turn off any notices/logging messages to the console or vty (so you can type the command *no access-group 170*

Step 2—Capture a Source IP

Cisco.com

```
access-list 170 permit icmp any any echo
access-list 170 permit icmp any any echo-reply log-input
access-list 170 permit udp any any eq echo
access-list 170 permit udp any eq echo any
access-list 170 permit tcp any any established
access-list 170 permit tcp any any
access-list 170 permit ip any any

interface serial 0
  ip access-group 170 out
! Wait a short time - (i.e 10 seconds)
  no ip access-group 170 out
```

Step 2—Capture a Source IP

- **Validate the capture with *show access-list 170*; make sure it the packets we counted**
- **Check the log with *show logging* for addresses:**

%SEC-6-IPACCESSLOGDP: list 170 permit icmp 192.168.212.72 (Serial0 *HDLC*) -> 198.133.219.25 (0/0), 1 packet

%SEC-6-IPACCESSLOGDP: list 170 permit icmp 172.16.132.154 (Serial0 *HDLC*) -> 198.133.219.25 (0/0), 1 packet

%SEC-6-IPACCESSLOGDP: list 170 permit icmp 192.168.45.15 (Serial0 *HDLC*) -> 198.133.219.25 (0/0), 1 packet

%SEC-6-IPACCESSLOGDP: list 170 permit icmp 192.168.45.142 (Serial0 *HDLC*) -> 198.133.219.25 (0/0), 1 packet

%SEC-6-IPACCESSLOGDP: list 170 permit icmp 172.16.132.47 (Serial0 *HDLC*) -> 198.133.219.25 (0/0), 1 packet

Step 3—Tracing the Source

Cisco.com

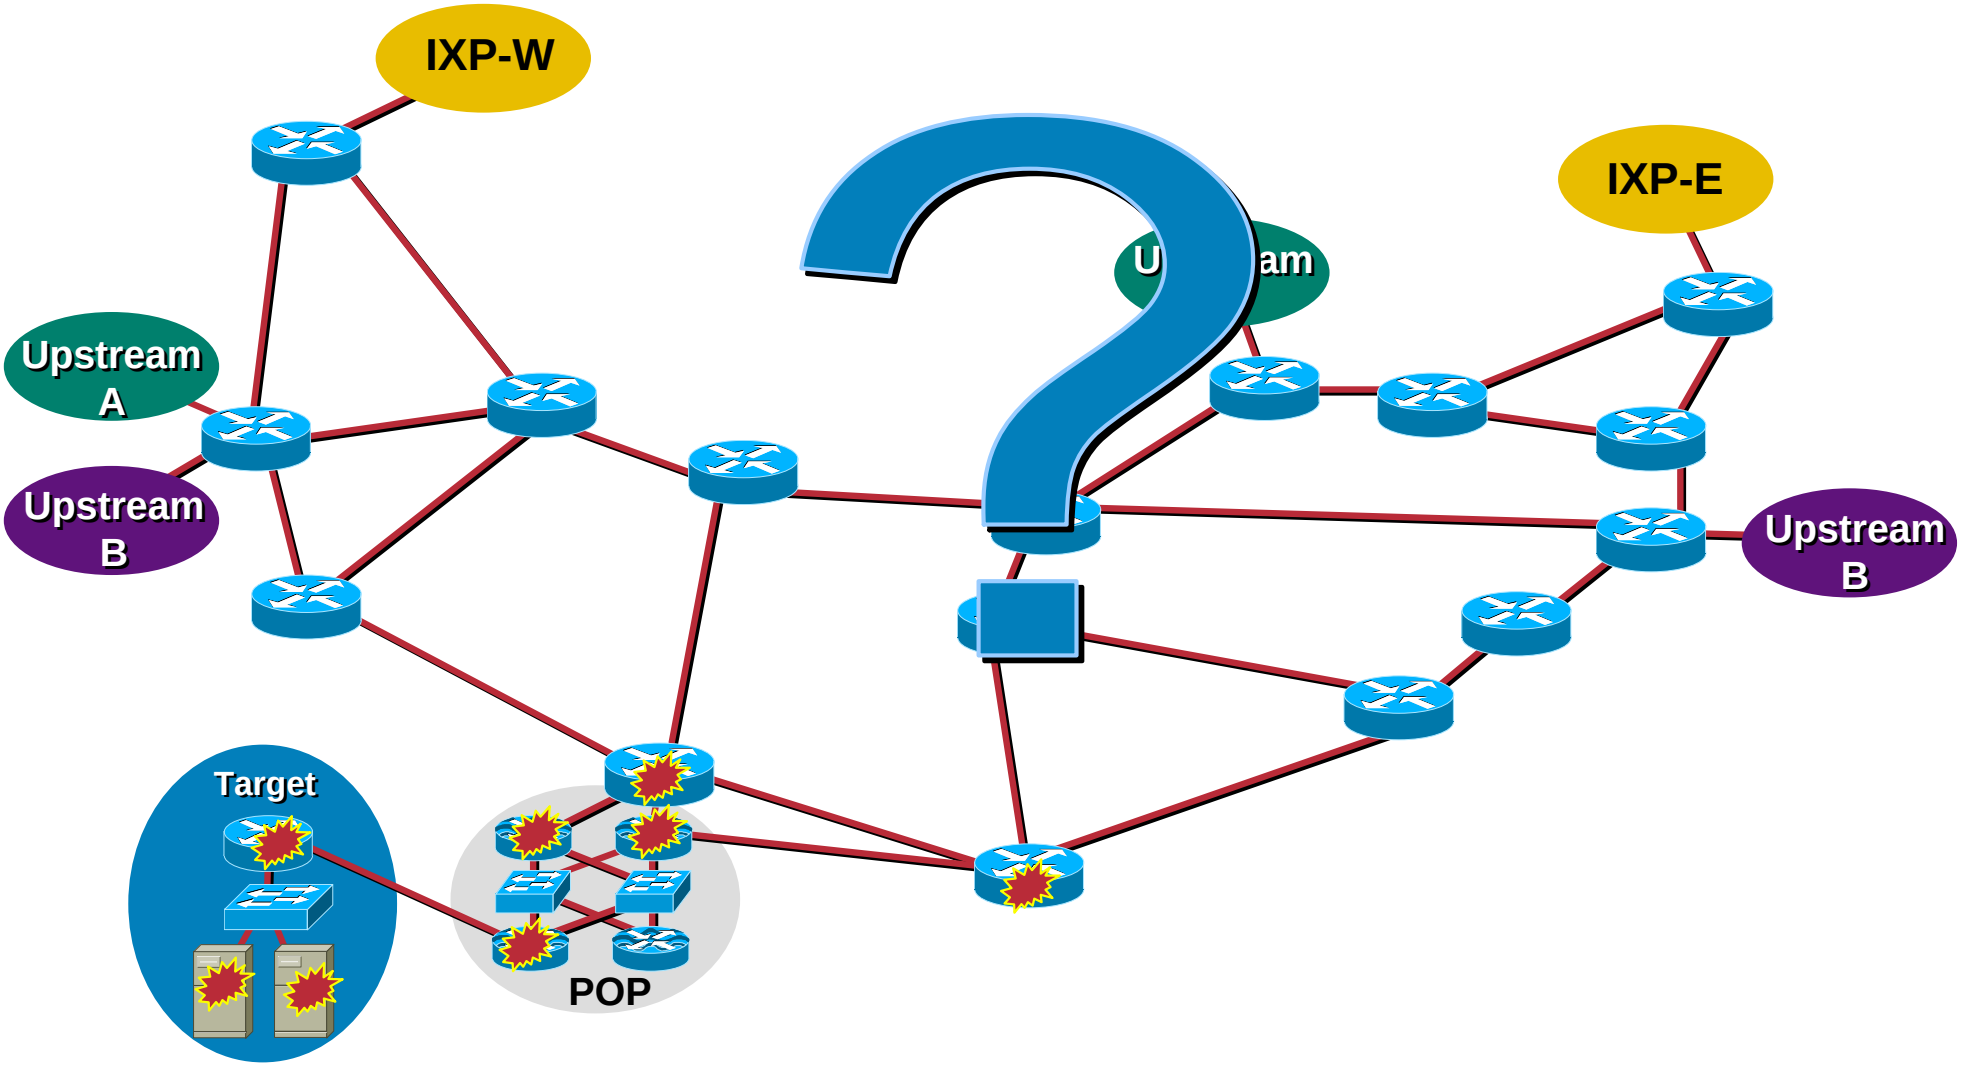

Step 3—Tracing the Source

- **Using Netflow for hop-by-hop traceback:**

```
Beta-7200-2>sh ip cache 198.133.219.0 255.255.255.0 verbose flow
IP packet size distribution (17093 total packets):
   1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
   .000 .735 .088 .054 .000 .000 .008 .046 .054 .000 .009 .000 .000 .000 .000

    512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
   .000 .000 .000 .000 .000 .000 .000 .000 .000 .000 .000

IP Flow Switching Cache, 1257536 bytes
  3 active, 15549 inactive, 12992 added
  210043 ager polls, 0 flow alloc failures
  last clearing of statistics never
Protocol         Total    Flows   Packets Bytes  Packets Active(Sec) Idle(Sec)
--------         Flows     /Sec     /Flow  /Pkt     /Sec      /Flow     /Flow
TCP-Telnet          35      0.0        80    41      0.0       14.5      12.7
UDP-DNS             20      0.0         1    67      0.0        0.0      15.3
UDP-NTP           1223      0.0         1    76      0.0        0.0      15.5
UDP-other        11709      0.0         1    87      0.0        0.1      15.5
ICMP                 2      0.0         1    56      0.0        0.0      15.2
Total:           12989      0.0         1    78      0.0        0.1      15.4

SrcIf         SrcIPaddress    DstIf         DstIPaddress    Pr SrcP DstP  Pkts
Fa1/1         192.168.45.142  POS1/0        198.133.219.25  11 008A 008A     1
Fa1/1         192.168.45.113  POS1/0        198.133.219.25  11 0208 0208     1
Fa1/1         172.16.132.154  POS1/0        198.133.219.25  06 701D 0017    63
```

Step 3—Tracing the Source

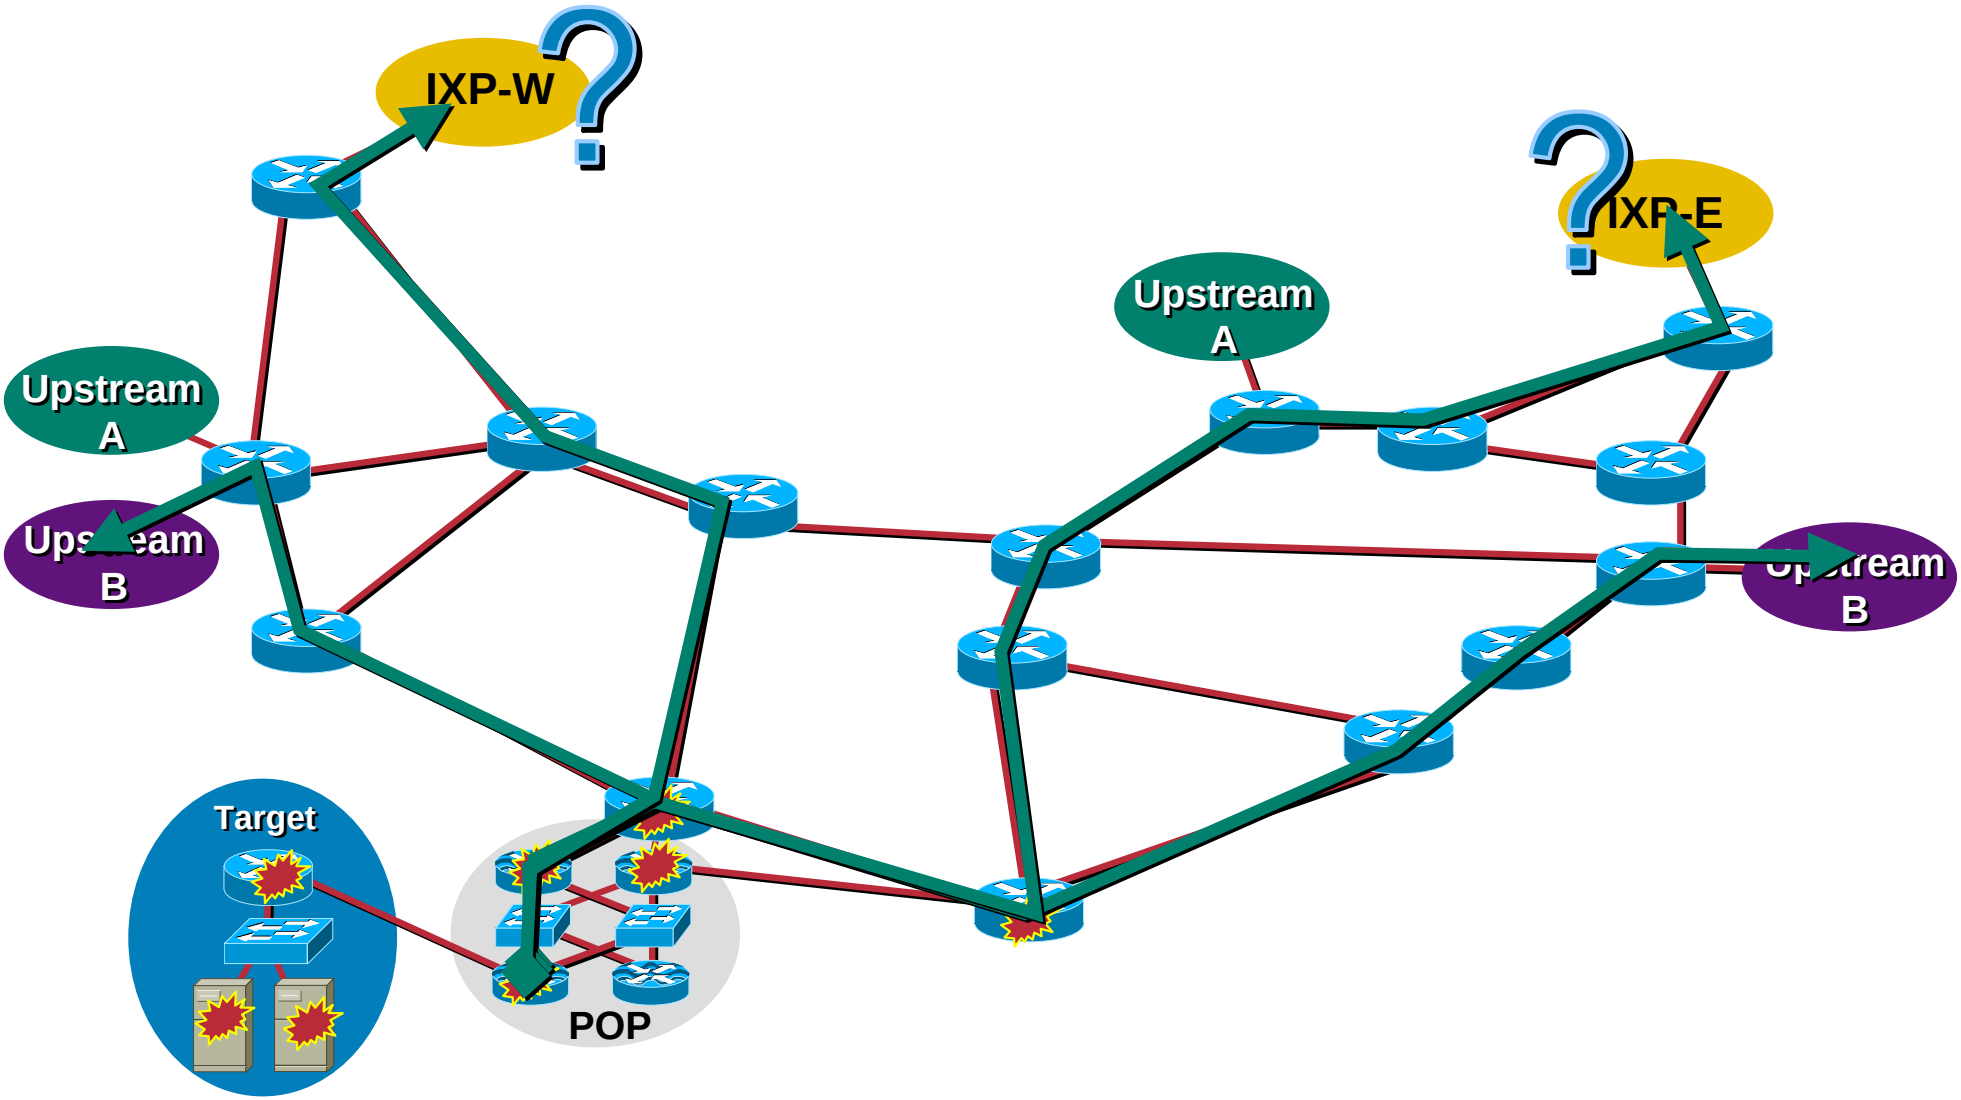

Step 3—Tracing the Source

- **Tracing across a shared access medium (I.e. like IXPs) require that ACL technique**

```
May 23 4:30:04.379: %SEC-6-IPACCESSLOGP: list 170 permitted
   icmp 192.168.45.142(0)(FastEthernet3/0/0 00d0.bc83.58a0)
   -> 198.133.219.25 (0), 1 packet

May 23 4:30:05.379: %SEC-6-IPACCESSLOGP: list 170 permitted
   icmp 192.168.45.142(0)(FastEthernet3/0/0 00d0.bc83.58a0)
   -> 198.133.219.25 (0), 1 packet

May 23 4:30:06.379: %SEC-6-IPACCESSLOGP: list 170 permitted
   icmp 192.168.45.142 (0)(FastEthernet3/0/0 00d0.bc83.58a0)
   -> 198.133.219.25 (0), 1 packet
```

Step 3—Tracing the Source

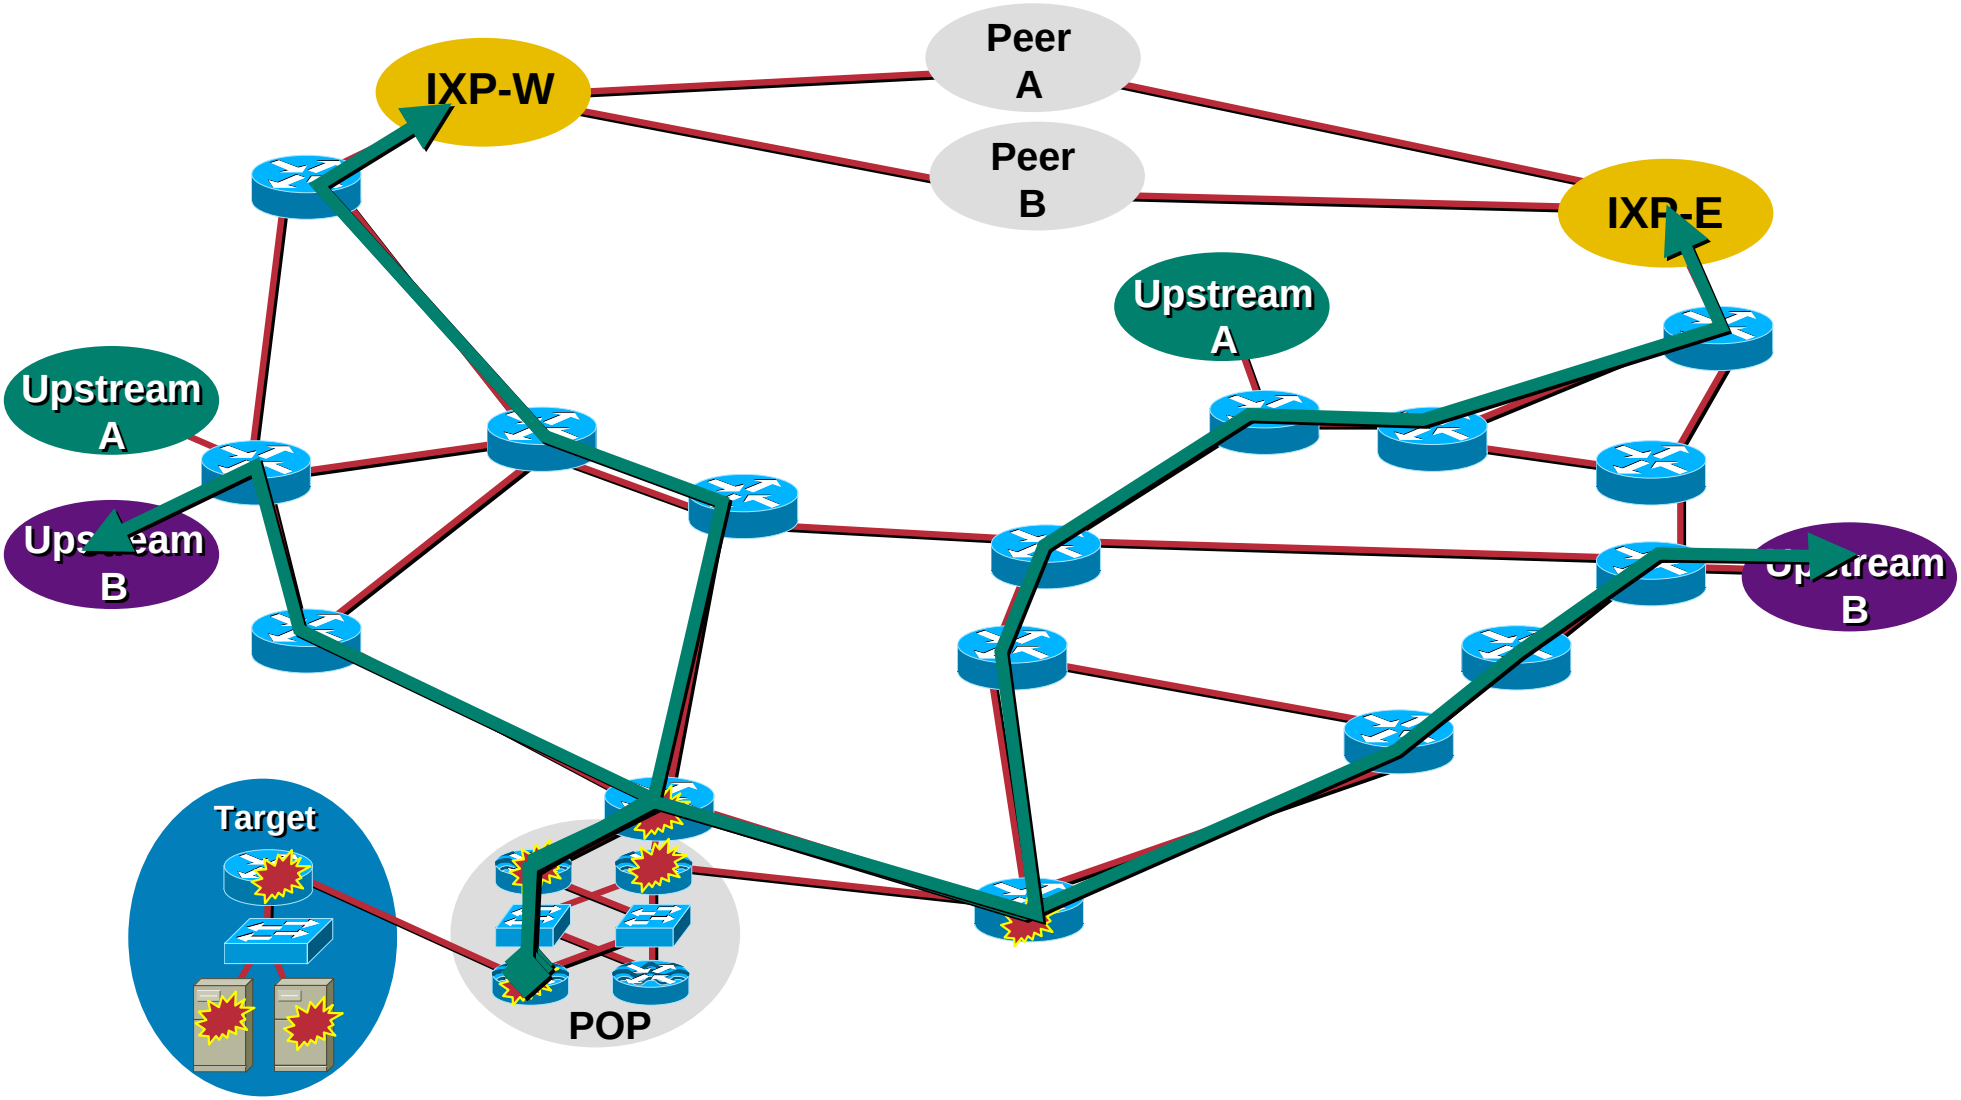

Troubleshooting Split

Cisco.com

- **Split in the security reaction team's flow:**
  - **One team starts calling NOCs**
    - **Upstream 2, Peer A, and Peer B**
  - **Other team drops filters in to push the packet drops to the edge of the network**

Step 4—Pushing the Packet Drops to the Edge

- **Options:**
  - **Rate limit the attack with CAR (input feature)**
  - **ACL to drop the packets**
  - **uRPF (perhaps)**
  - **Drop the connection to the peer/upstream**

Step 4—Pushing the Packet Drops to the Edge

- **Select rate limiting option; limit ICMP echo-reply for everyone and limit the peer's traffic**

```
interface FastEthernet3/0/0

  rate-limit output access-group 2020 256000 16000 24000
conform-action transmit exceed-action drop

  rate-limit input access-group rate-limit 100 8000000 64000
80000 conform-action transmit exceed-action drop

!

access-list 2020 permit icmp any any echo-reply

access-list rate-limit 100 00d0.bc83.58a0
```

Step 4—Pushing the Packet Drops to the Edge

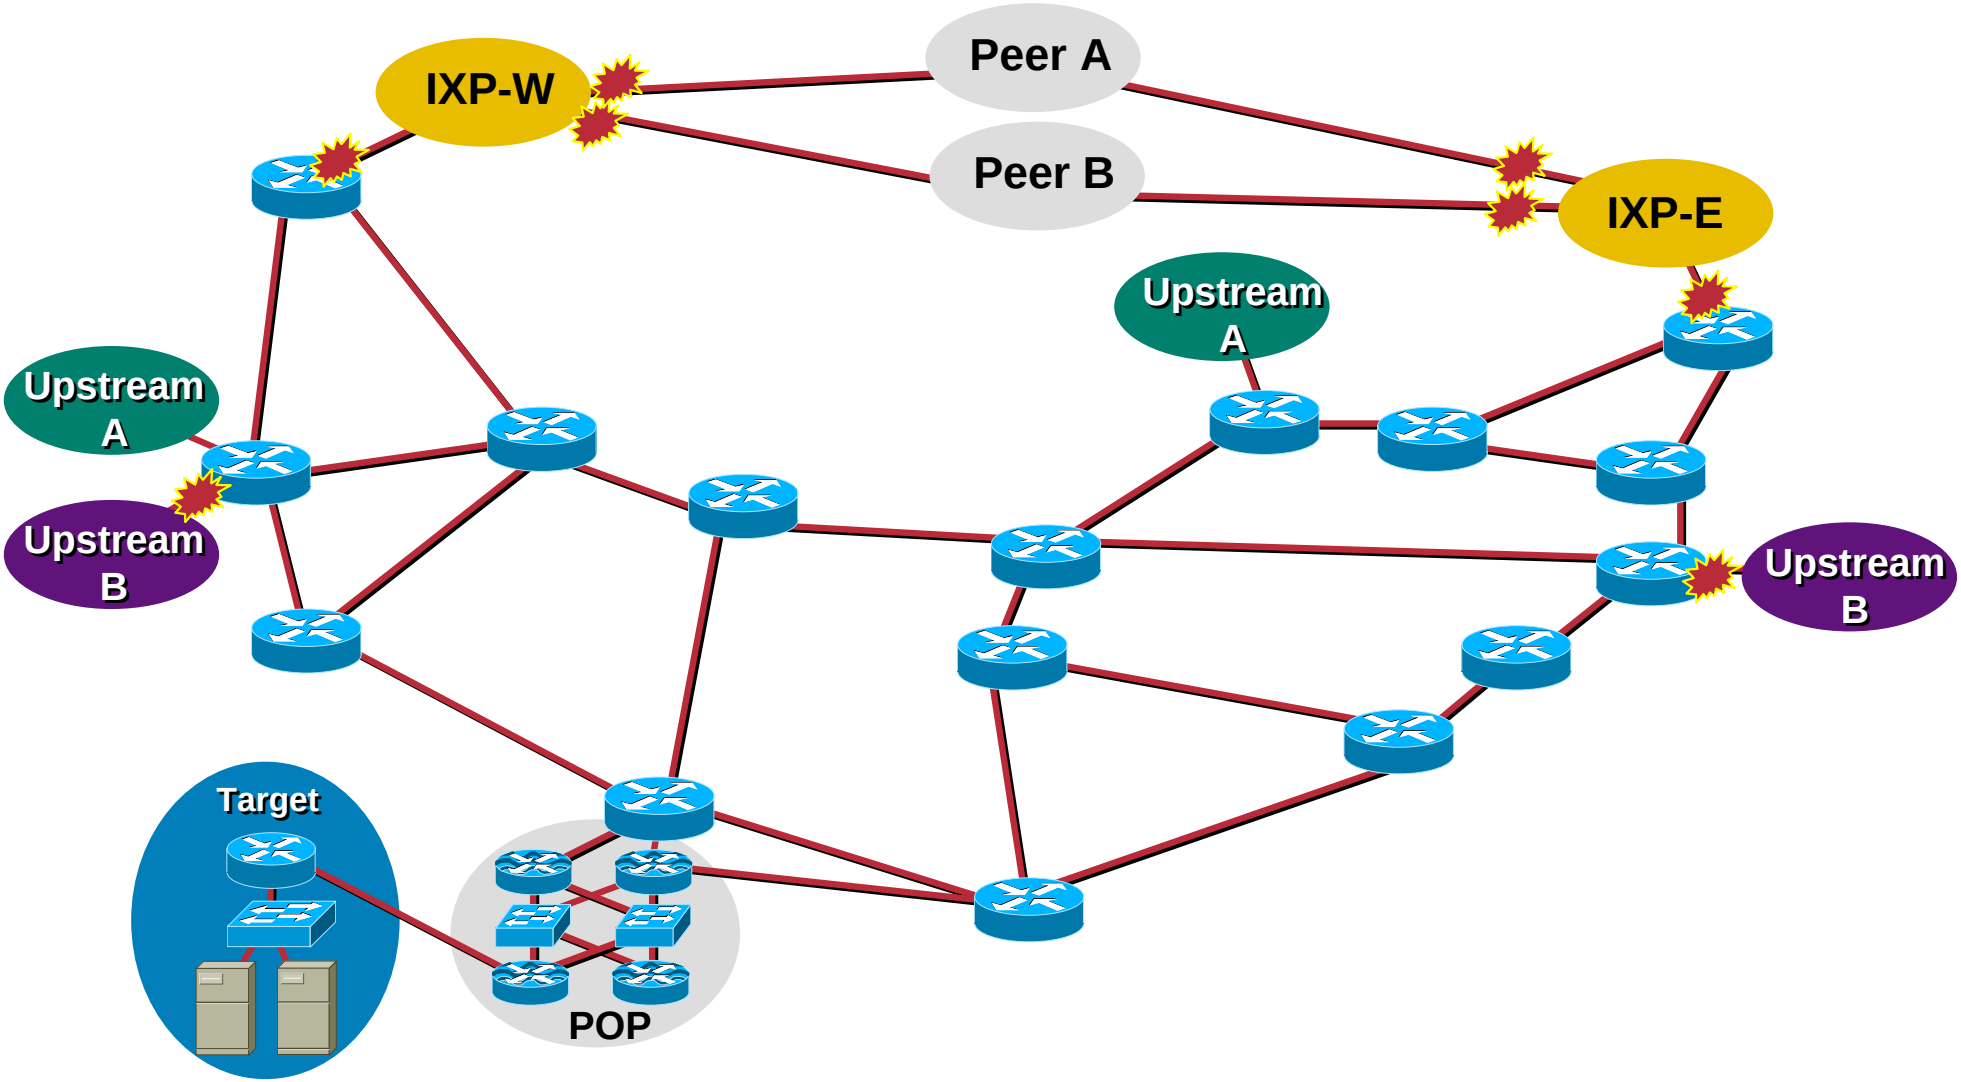

Check Point

- SitRep—attack still in progress—packets being dropped at the ISP edge
- Work with upstream and peer ISP NOCs to continue the trace back to the sources
- Collect evidence—work with customer and call your legal team

Alert!

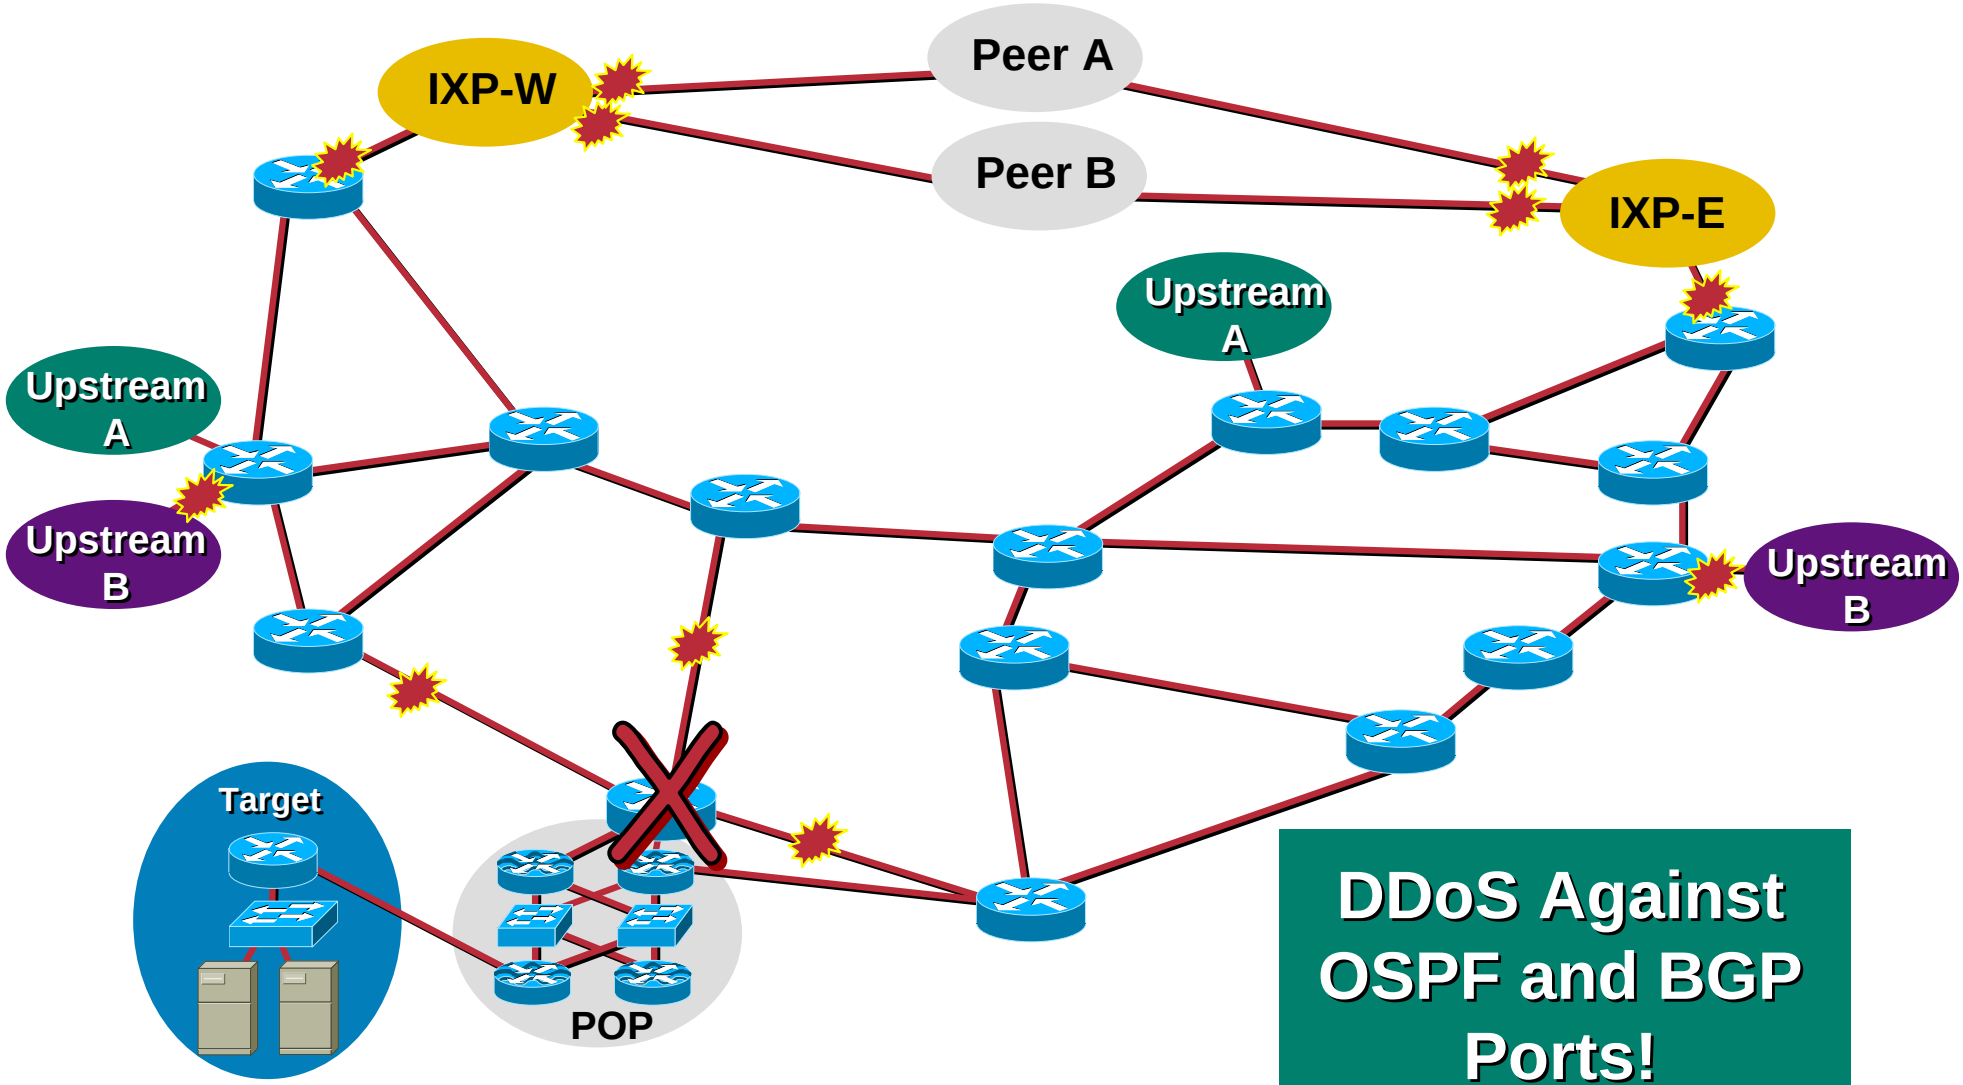

DDoS Against OSPF and BGP Ports!

Next Phase of the Attack

Cisco.com

- **The attackers have shifted the attack to their target's infrastructure**
  - **ISPs and IXPs have and will be directly attacked to get at the target!**
  - **ISP's routers are being directly attacked to take out the target**


The Hacker News Network
www.hackernews.com
HOPE
August 13-14, 1994
Hotel Pennsylvania
New York City
Celebrating 10 years of 2600
2600
Magazine
The HACKER Quarterly
Are
You
phrack
v0.56
http://www.phrack.com
Ready?

In Case You Wondering…

- **How to work a DoS attack against the routing protocol?**
  - **Out of band access to the router!**
  - **Rate limits on traffic to the routing protocol**
  - **ACLs to block outside traffic to the routing protocol ports**

DDoS Links

- http://www.denialinfo.com/
- http://www.staff.washington.edu/dittrich
- http://www.fbi.gov/nipc/trinoo.htm
- http://www.sans.org/y2k/DDoS.htm
- http://www.nanog.org/mtg-9910/robert.html
- http://cve.mitre.org/
- http://packetstorm.securify.com/distributed/