Securing the network infrastructure of an ISP

Rehan Nedaria

CCIE:10971 (R&S/SP)

July 15 2009

Abstract

- As more services are transitioned to an all-IP infrastructure, the availability of the core IP network becomes a critical concern. This session reviews Cisco IOS® features, general security techniques, and best practices used by service providers to increase the availability of the IP network. This session is designed for service provider network engineers, and many examples are based on service provider platforms . However, the session also provides insight into service provider security for network and security professionals in enterprise environments.

Agenda

- Understanding Routers and Planes
- Securing your router - IOS Security

Routers and Planes

- **There are nuances to the definition of these planes**

  IETF RFC3654 defines two planes: control and forwarding

  ITU X805 defines three planes: control, management, and end-user

  Cisco defines three planes: control, management, and data

- **Three Plane Conceptual Model:**

  Control Plane – The routing protocols gluing the network together.

  Management Plane – The tools and protocols used to manage

  the device.

  Data Plane – Packets going through the router.– Packets going through the router.

ASIC Based Platform—Main Components

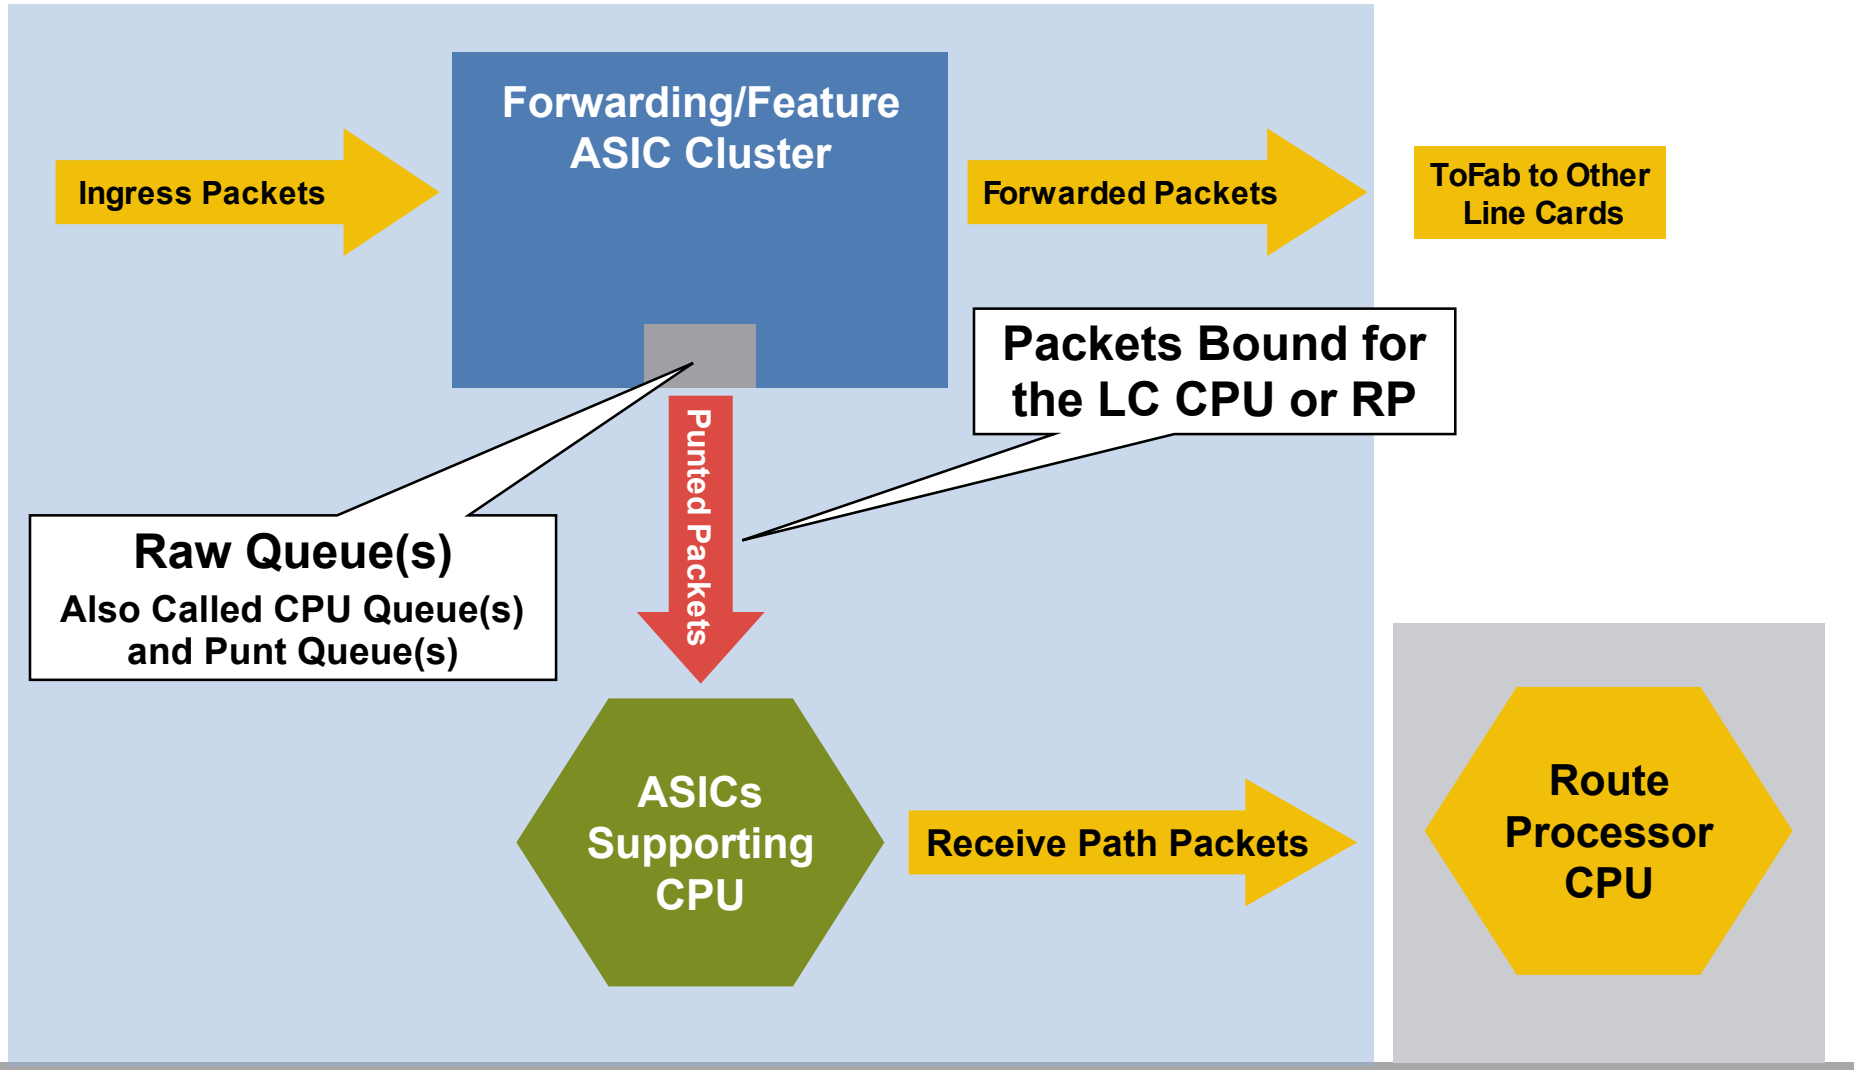

Data Plane

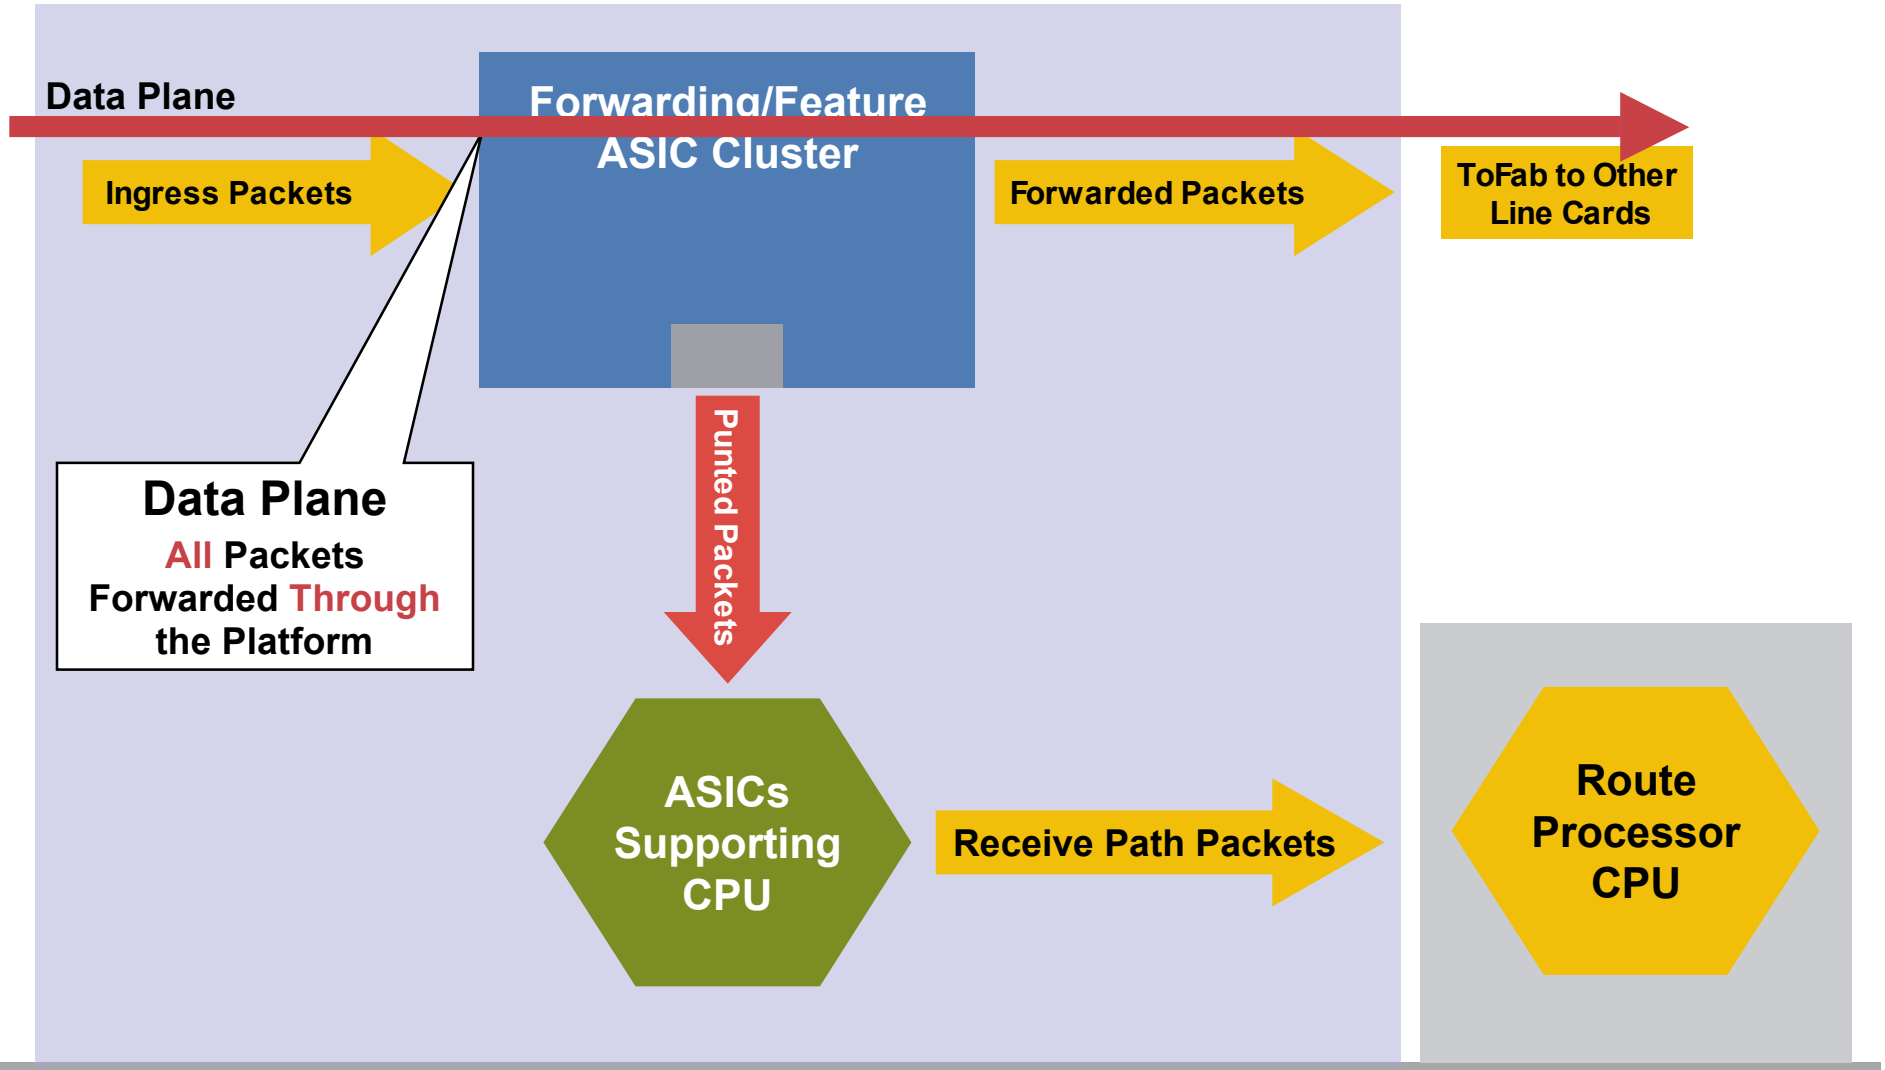

Control Plane

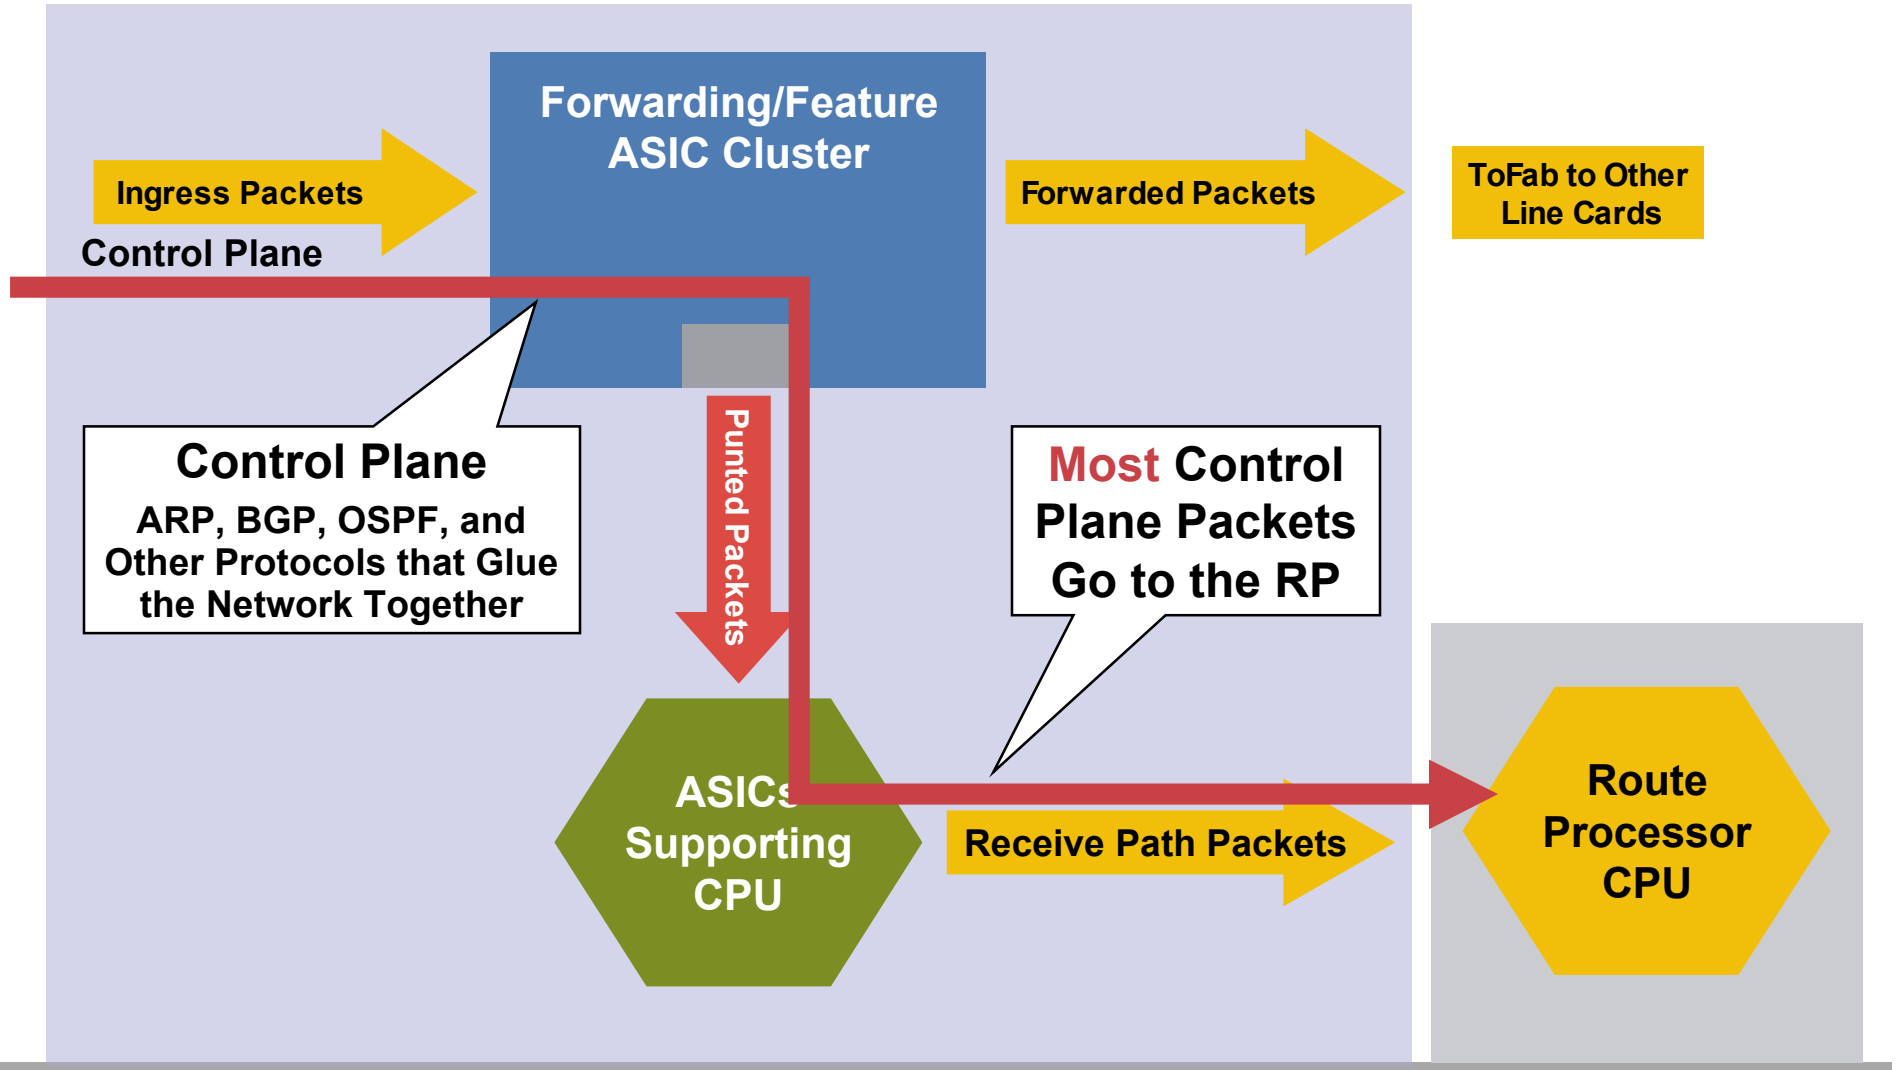

Management Plane

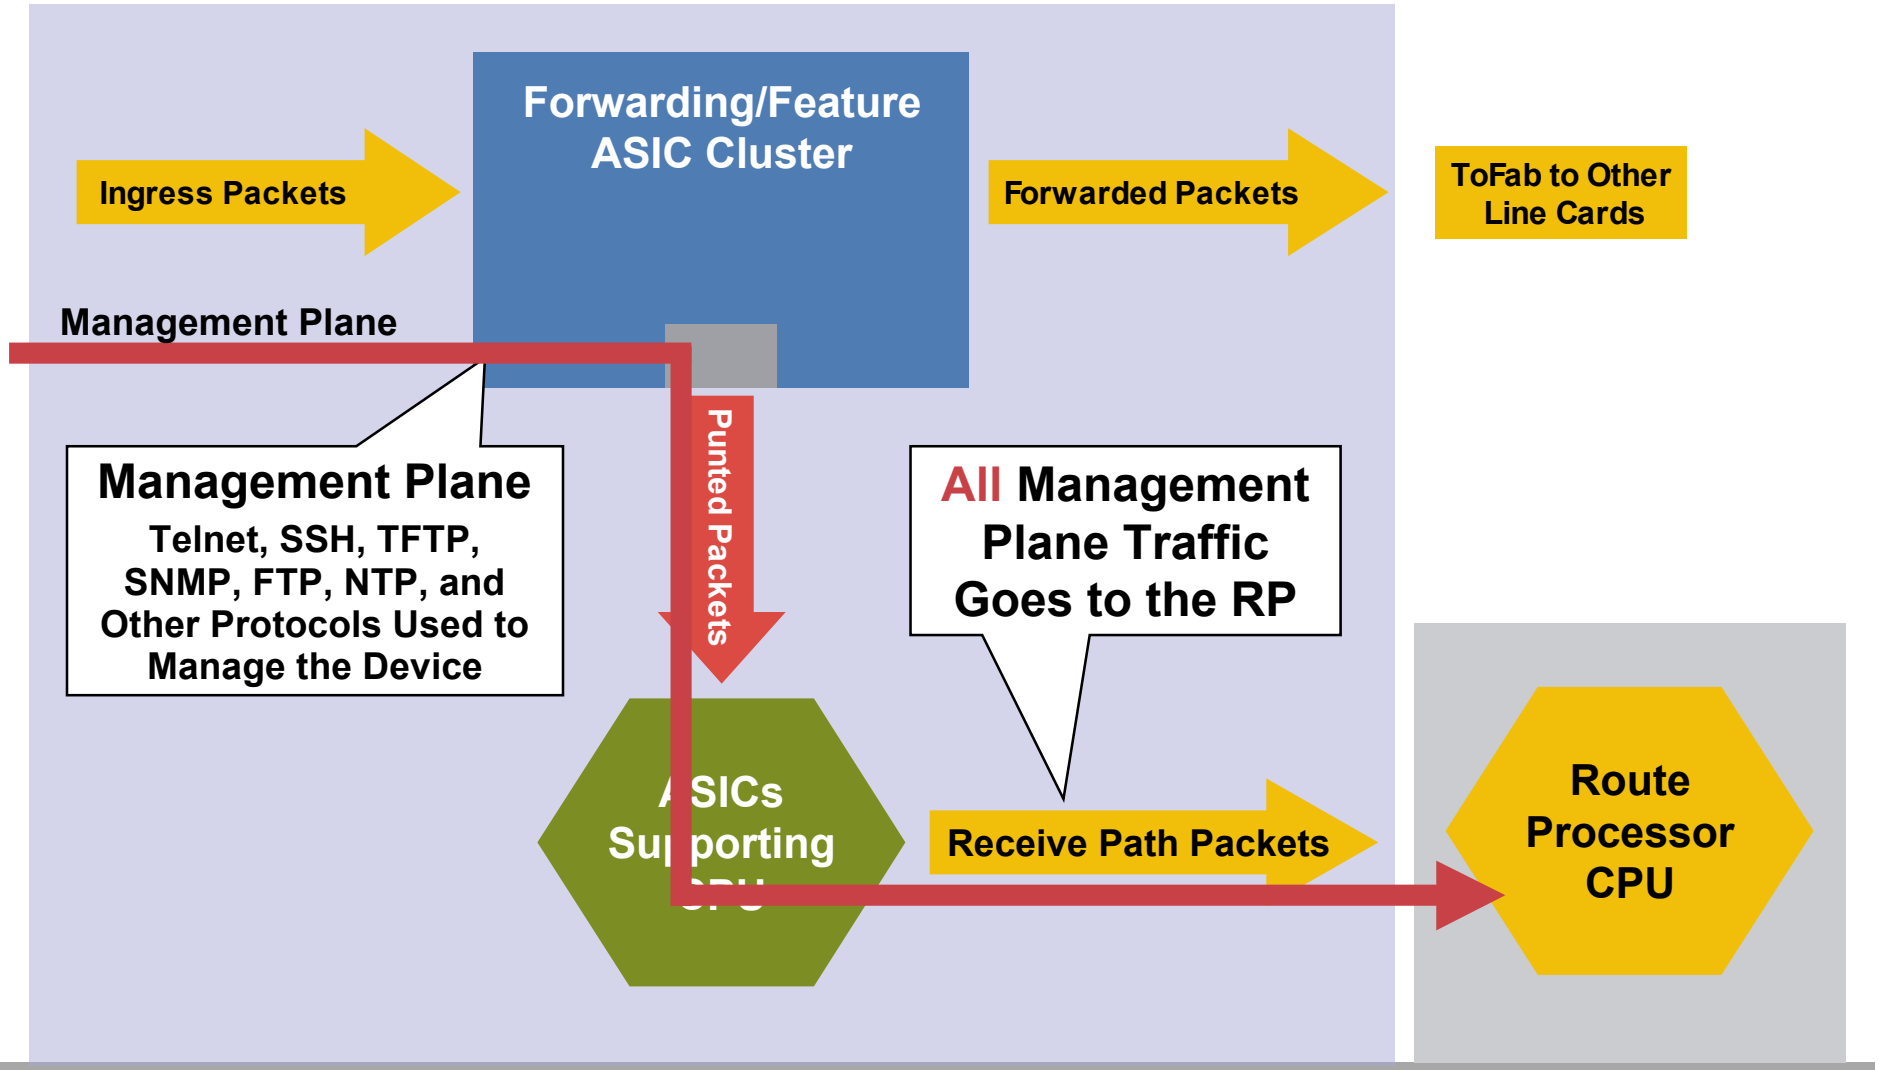

Attack Vectors

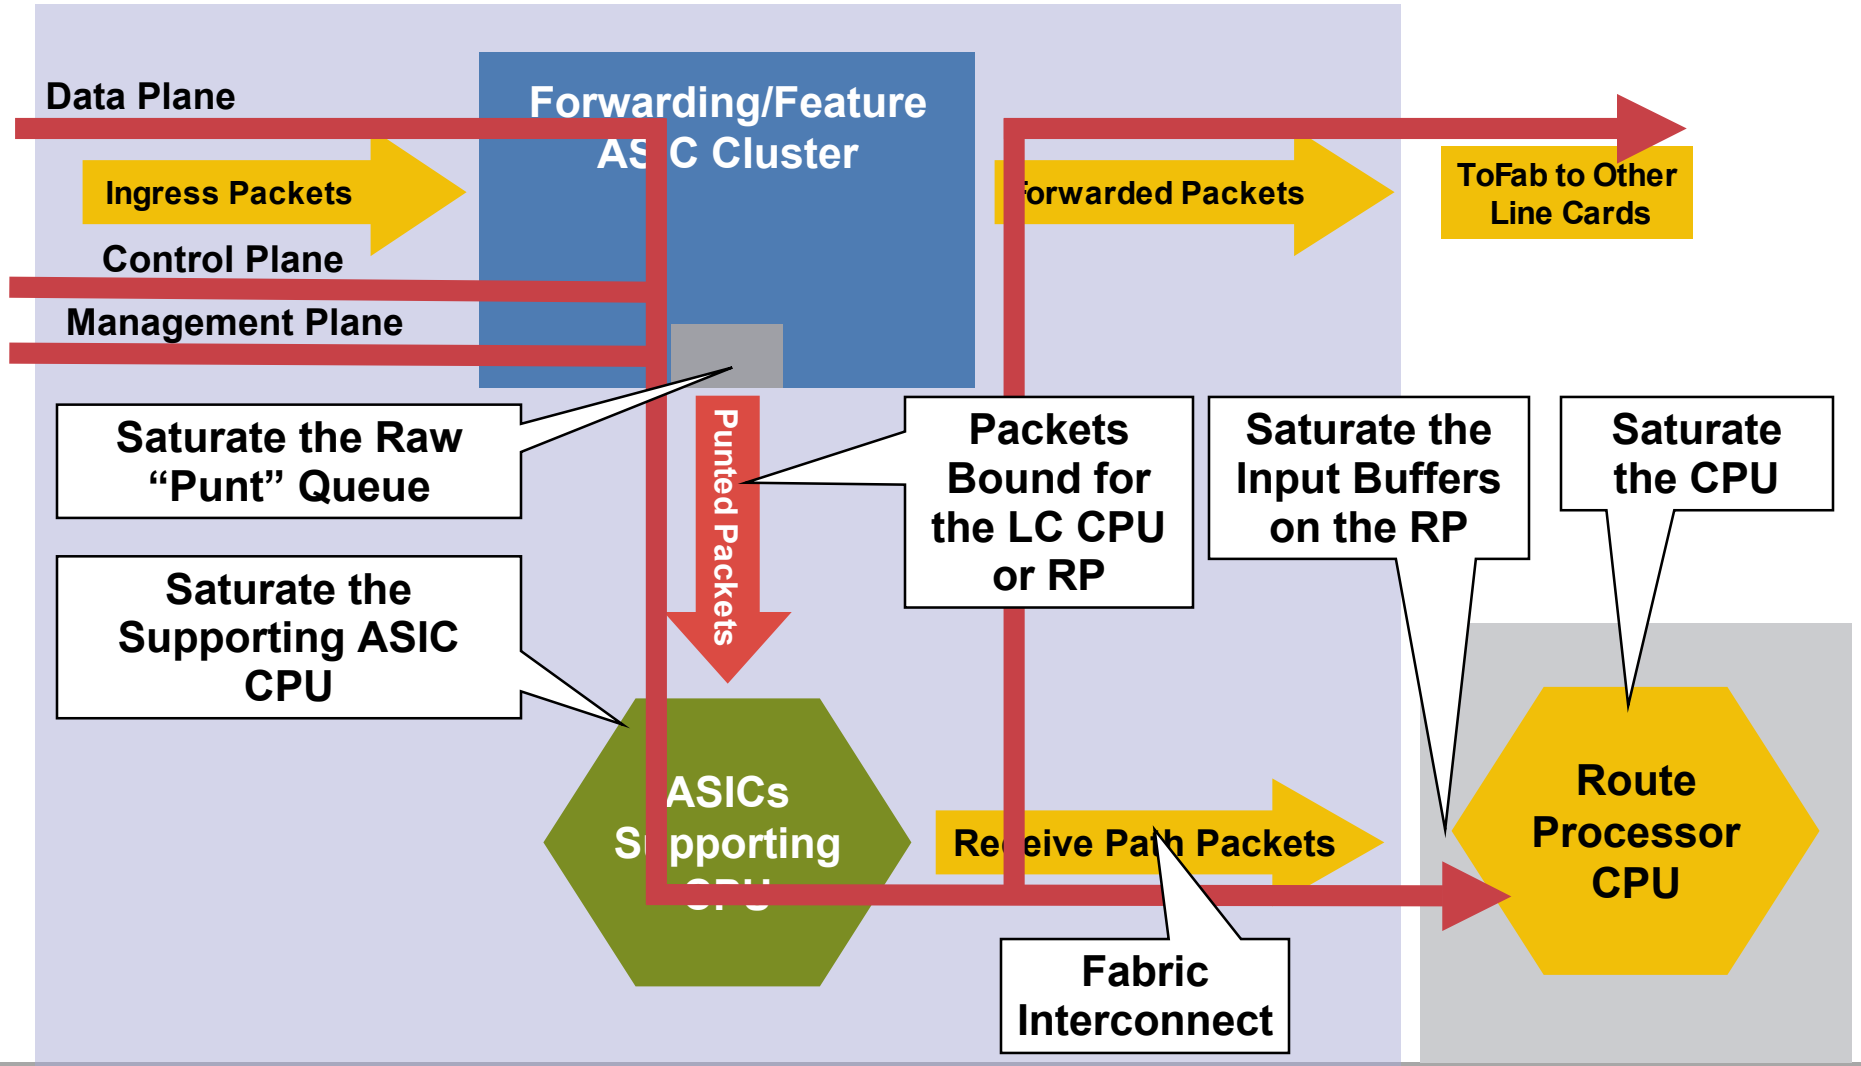

Agenda

- Understanding Routers and Planes
- Securing your router - IOS Security

Securing the Router

Essential Lock Downs

Global Services You Turn OFF

- **Finger**

  Find out who is logged in, from where, how long for

- **PAD**

  Historical – from the days of X.25

- **Small servers**

  Tcp and udp ports < 20 are for developing IP stacks and not needed in day to day operations

- **Bootp**

  Used by systems to bootstrap themselves onto the network – e.g. X-terminals

Global Services You Turn OFF

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

no ip finger

no service pad

no service udp-small-servers

no service tcp-small-servers

no ip bootp server

Interface Services You Turn OFF

- **IP redirects**

  Router will send redirect message if it has to resend a packet through the same interface it was received on.

- **Direct-broadcast**

  If packet intended for network broadcast address, router will physically broadcast it onto the attached network The cause of all SMURF attacks on the Internet.

- **Proxy-arp**

  Dumb host sends arp request for destination – documented in RFC1027

  If router knows how to get to that destination, it will install an entry in the arp table for that destination.

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

```
switch#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 MAINTENANCE INTERIM SOFTWARE
Copyright (c) 1986-2000 by cisco Systems, Inc.
Compiled Fri 03-Mar-00 19:28 by htseng
```

Login Banner

- **Use a good login banner.**

banner login ^

Authorised access only

This system is the property of MattNet Internet

Disconnect IMMEDIATELY if you are not an authorised user!

Contact noc@mattnet.net +99 999 999999 for help.^

Exec Banner

- **Useful to remind logged in staff 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 +99 999 999999 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

Securing Access to the Router

ISP Tools to Secure Access to the Router

- **VTY**
- **SSH Configuration**
- **Username based on the router**
- **AAA**

VTY Security

Access to VTYs should always be controlled and authenticated

- **Network Access is controlled via access lists**
- **Authentication: Local password, TACACS+, RADIUS, Kerberos**
- **Transport mechanism should be encrypted**
  - SSH
  - Telnet over IPSEC

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 (timers for exec Sessions)

transport input telnet ssh (SSH telnet sessions)

password 7 045802150C2E (encrypted passwords)

VTY Security (Logging)

- **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
```

SSH Support in ISP Code

- **Cisco IOS Software supports SSHv1 and SSHv2**

- **SSH Server**

  Server for remote connections

- **SSH Client**

  SSH connections can be generated from router to another router or to the SSH server

- **SCP**

  Enables the copying of config and image files to and from the router via SSH

- **SSH Terminal-Line Access (Reverse-SSH)**

  Accessing terminal lines of the router via SSH

Configuring SSH: Cisco IOS

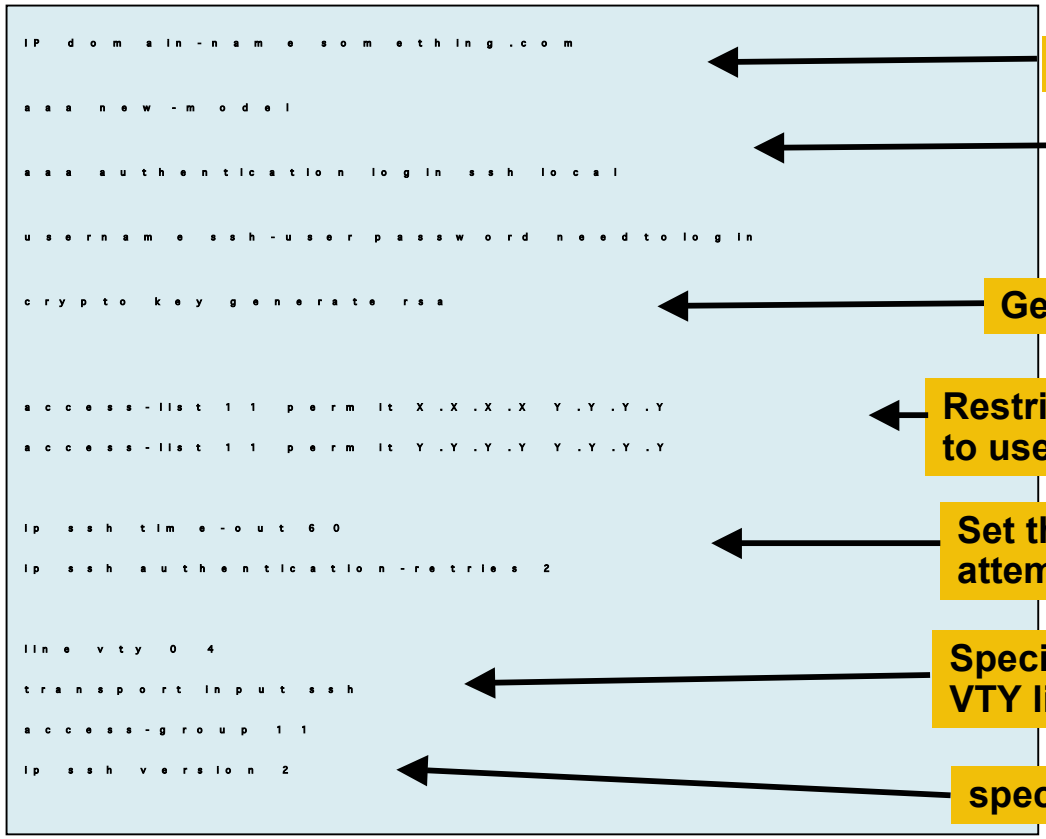

Must have set a domain name

Must have AAA configured

Generate rsa key

Restrict which host/subnets are allowed to use SSH to the router

Set the timeout and retry attempts

Specify SSH on the router VTY lines

specific version of SSH

SSH Terminal-Line Access

- How can security be increased on Telnet sessions, when using terminal servers for access to console ports in data center
- Included a feature to allow SSH access to lines in async modules, allowing encrypted access to TTY connections

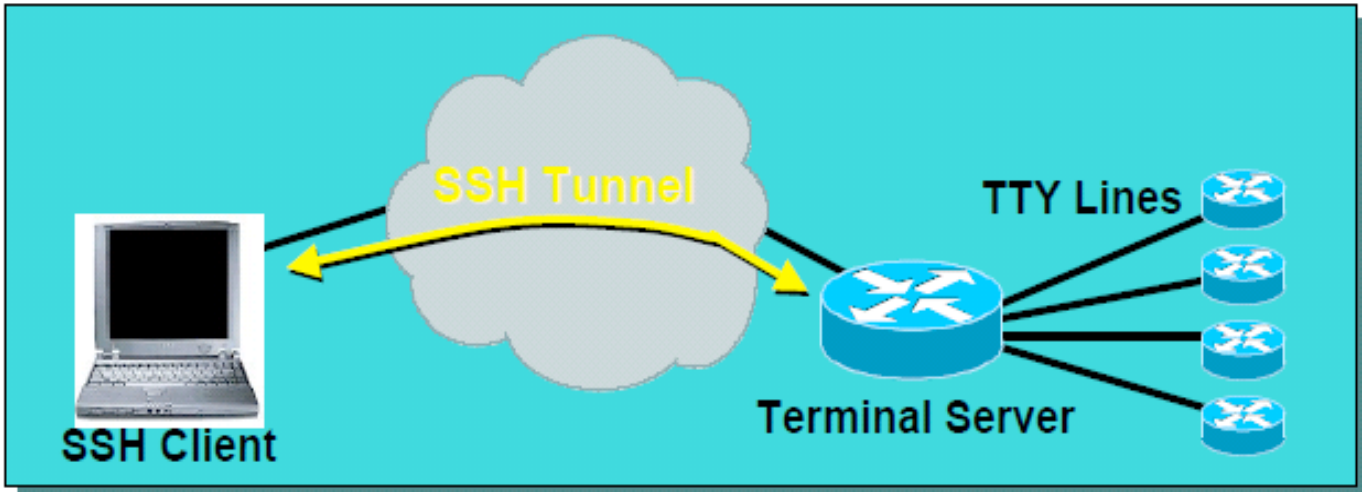

SCP: Secure Copy

- **SCP supported in:**
  - **12.1.(1)T—server**
  - **12.1.(3)T—client**

- **After configuring SSH enable SCP**

```
ip scp server enable
```

- **SCP client examples:**

```
rtr-us-1#   copy running-config scp://tiger@10.1.1.2/
Address or name of remote host [10.1.1.2]? <ret>
  Destination username [tiger]? <ret>
              Destination filename [rtr-us-1-confg]? <ret>
              Writing rtr-us-1-confg
              Password: f00bar
rtr-us-1#  exit
```

Authentication, Authorization, Accounting

- It is good practice to register each individual user with a separate user-id
- Generic account setup is easier to fall into the wrong hands
- Default password only login is easier to use a brute force crack
- A username/password pair makes brute force techniques harder, but not impossible

Configuring:
username joe password 7 045802150C2E
username jim password 7 0317B21895FE
!
line vty 0 4
login local

Authentication, Authorization, Accounting

- **TACACS+**
- **RADIUS**
- **Kerberos**

Advantages for using a central database for authentication

- Scalability
- Logging
- Secured

RADIUS vs. TACACS+ vs Kerberos

	RADIUS	TACACS+	KERBEROS
Uses UDP	X		
Uses TCP		X	X
Encryption	Password Only	All But Header	All But Header
Multiprotocol Support		X	
Router Mgt Acct Control		X	X
Router Mgt Auth Control		X	X
LEAP Support	X		
XAUTH Support	X	X	X

What to Configure?

- **Local Authentication to provide a failsafe**
- **Authentication with AAA Server with (TACACS+ used in this example).**
- **Accounting/Audit with AAA Server**
- **Authorization with AAA Server**

Simple Staff Authentication and Failsafe

- **Username/Passwords on the Router are used as a back-up when the AAA system goes down.**

```
aaa new-model
aaa authentication login default tacacs+ enable
aaa authentication enable default tacacs+ enable
username joe password 7 1104181051B1
username jim password 7 0317B21895FE

ip tacacs source-interface Loopback0
tacacs-server host 215.17.1.2 (Ip address of tacacs server)
tacacs-server host 215.17.34.10 (Redundancy of the tacacs server )
```

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

Distribute AAA Servers and Config Backup

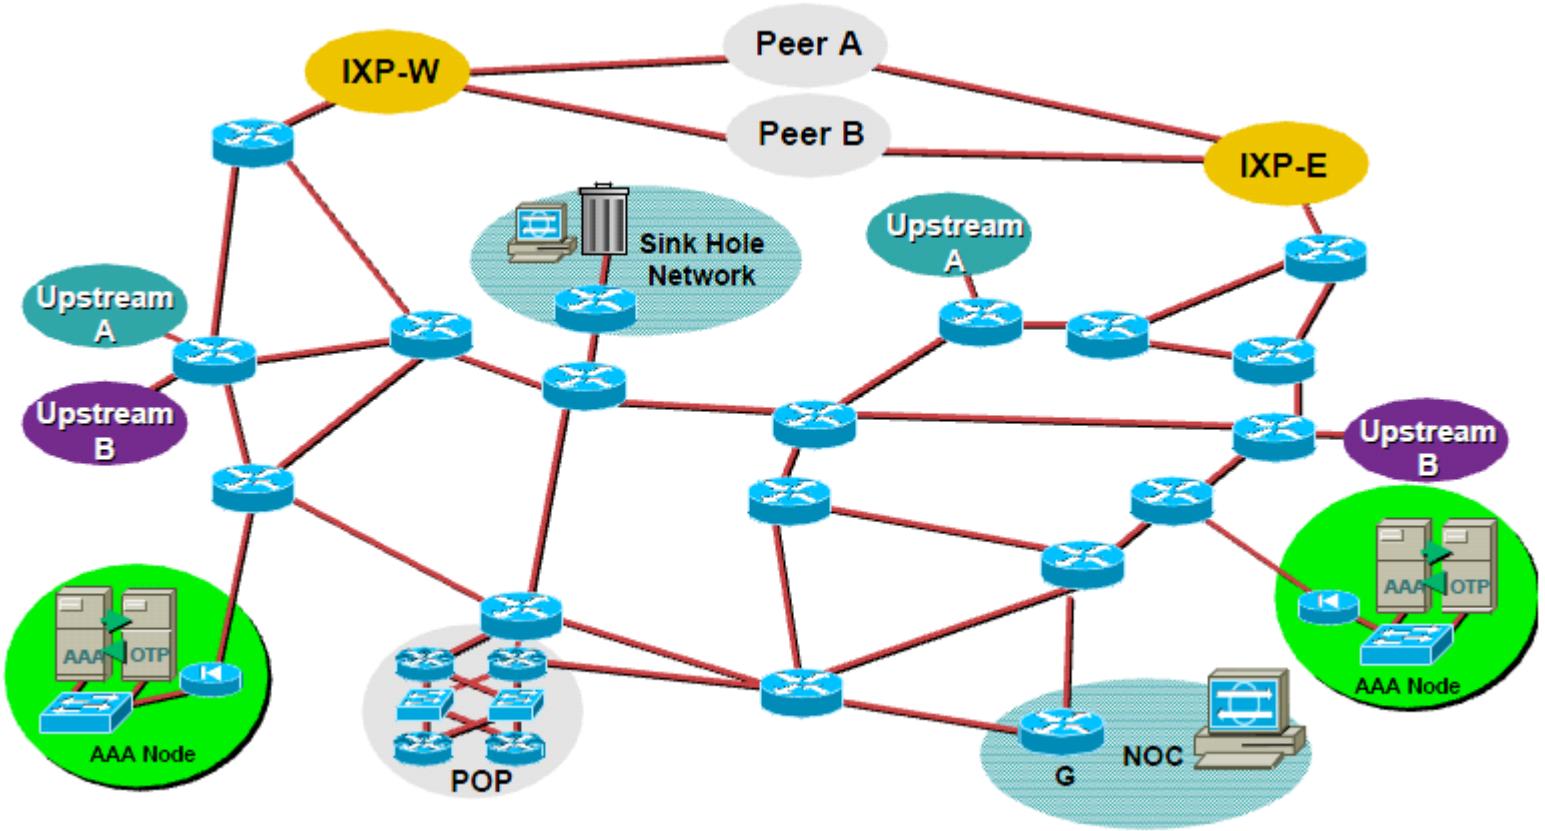

Source Routing

- **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

Source Routing

Example of Source routing disabled on the 9th router

- Unix% /usr/local/bin/traceroute -g 192.121.154.170 www.sprint.net
- traceroute to www.sprint.net (208.27.196.10), 30 hops max, 40 byte packets
- 1 fe0-1-0.hr1.cbg1.gbb.uk.uu.net (158.43.128.192) 1.294 ms 0.703 ms 0.539 ms
- 2 pos0-0.cr1.cbg1.gbb.uk.uu.net (158.43.129.129) 0.748 ms 0.993 ms 0.747 ms
- 3 pos0-0.cr2.cbg1.gbb.uk.uu.net (158.43.129.133) 1.586 ms 1.146 ms 2.145 ms
- 4 pos0-2.cr2.lnd6.gbb.uk.uu.net (158.43.254.2) 4.43 ms 4.143 ms 3.731 ms
- 5 ge2-0.cr1.lnd6.gbb.uk.uu.net (158.43.254.65) 4.395 ms 4.044 ms 4.148 ms
- 6 POS11-0-0.GW2.LND1.Alter.Net (146.188.5.41) 4.898 ms 10.705 ms 5.082 ms
- 7 122.at-2-0-0.XR2.LND2.Alter.Net (146.188.15.170) 5.995 ms 6.179 ms 6.039 ms
- 8 194.ATM1-0-0.HR2.LND1.Alter.Net (146.188.15.129) 12.422 ms 7.229 ms 6.018 ms
- 9 sl-bb5-dc-4-0-0.sprintlink.net (144.232.7.166) 86.662 ms !S 88.132 ms !S *

ICMP Unreachable Overload

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

ICMP Unreachable Overload

- **All Routers who use any static route to Null0 should put *no ip unreachables***

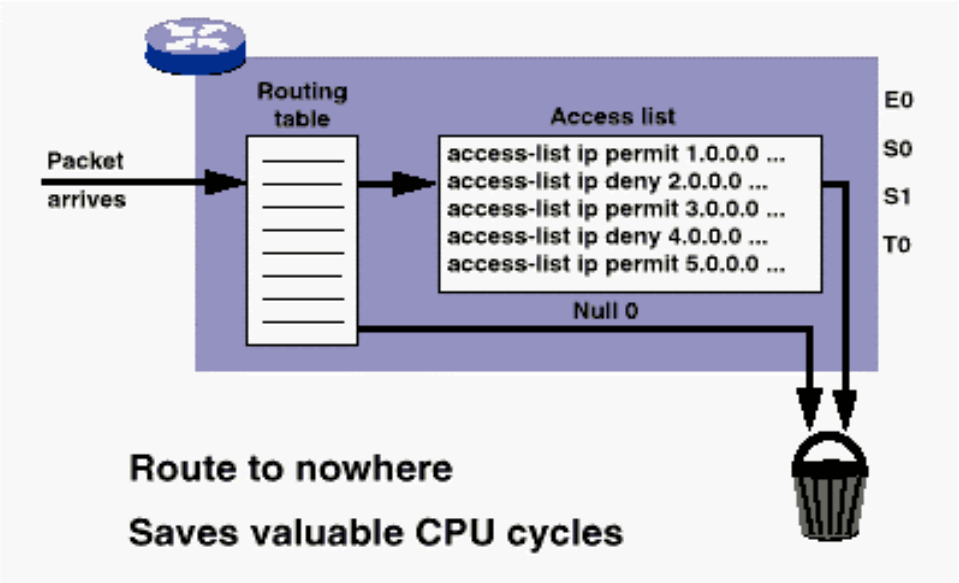

interface Null0

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 2 seconds.

Selective Packet Discard (SPD)

- 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 (SPD)

- **Recommended Settings:**

  ip spd headroom 1000 **Default is 100.**

  Specifies how many high-precedence packets we will enqueue over the normal input hold queue limit. This is to reserve room for incoming high precedence packets.

Experience shows that the higher settings help.

Receive ACL (rACL)

Excessive traffic destined to GRP can lead to high CPU ! DoS

- Receive ACLs filter traffic destined to the GRP
- rACLs explicitly permit or deny traffic destined to the GRP
- rACL do NOT affect transit traffic
- Traffic is filtering on the ingress LC, prior to GRP processing
- rACLs enforce security policy by filtering who/what can access the router

Receive ACL (rACL)

Introduced in 12.0(21)S2/12.0(22)S

- ip receive access-list [number]

  Standard or Extended
- As with other ACL types, show access-list provide ACE hit counts
- Log keyword can be used for more detail

Administrative and Operational Practices

Configuration Management

- **Backup NVRAM configuration off the router:**

  Write configuration to TFTP server

  TFTP server files kept under revision control

- **Allows rapid recovery in case of emergency**

Configuration Management

- Secure the TFTP server
- TFTP loopback 0 on router
- Wrapper on TFTP server which only allows the router's loopback address

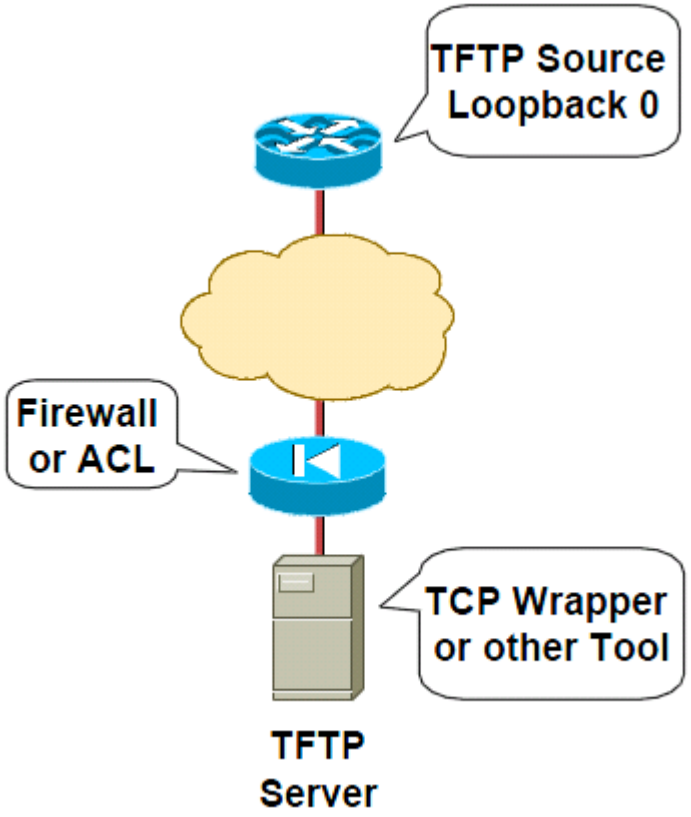

ip tftp source-interface Loopback0

Use Detailed Logging

- **Off load logging information to a logging server**
- **Use the full detailed logging features to keep exact details of the activities**

service timestamps debug **datetime** msec **localtime** show-timezone

service timestamps log **datetime** msec **localtime** show-timezone

logging buffered 16384

logging trap debugging

logging facility local7

logging 169.223.32.1

logging 169.223.55.37

logging source-interface loopback0 (Default outgoing interface)

no logging console ! Recommended - keeps the console port free

Jul 27 15:53:23.235 AEST: %SYS-5-CONFIG_I: Configured from console by philip on console

NTP: Network Time Protocol

- NTP is an open standard referred to in RFC 1305
- NTP keeps clocks on network/systems gear in constant synchronization with one another
- NTP is supported in all Cisco gear, as well as in many if not all operating systems
- NTP is crucial for:

  Accurate logging

  Validating certificates

NTP Authentication

- Authenticating the source for clock adjustments is important to ensure that a remote-based attack doesn't alter the clocks on machines
- Stops illegal time sync updates from illegal sources
- As a result of logging, you can see which sites are trying to change the clocks

NTP Authentication: Cisco IOS

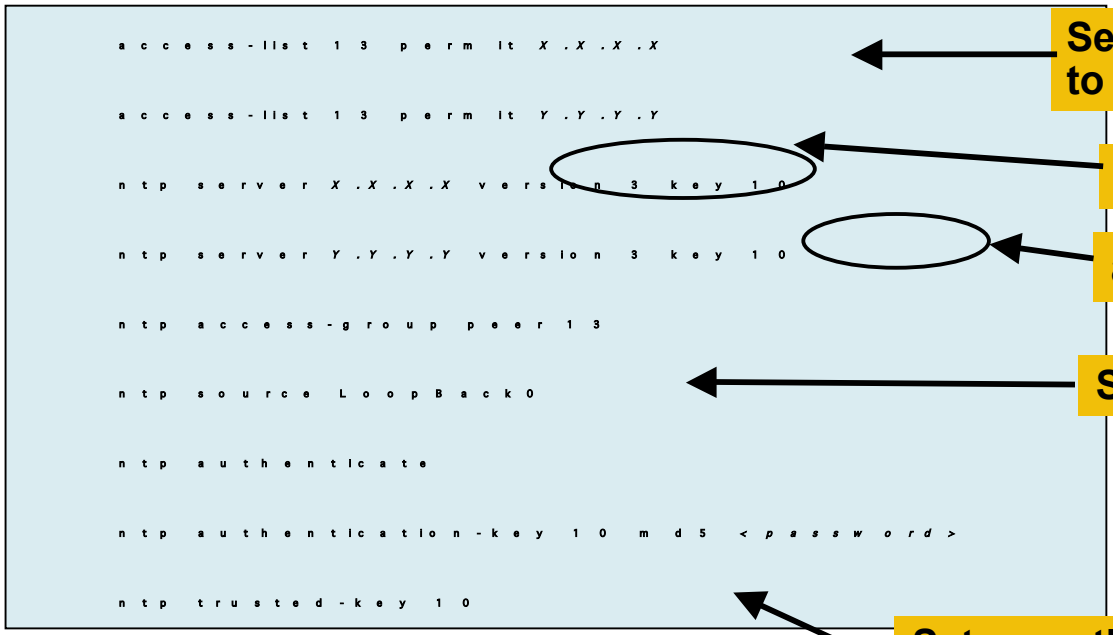

Set time synchronization to X.X.X.X and Y.Y.Y.Y

Specify version 3

authentication-key 10

Specify source interface

Set up authentication for NTP, configure password and define trusted key

IP HTTP Server

- Without secure-http or authentication embedded in the HTTP server with an associated ACL, the HTTP server is at risk of attack and exploitation

- Where you may require the HTTP server:
  - QoS policy manager
  - Secure device manager
  - PIX device manager

IP HTTP Server

- If you need to enable HTTP server, secure access with HTTPS and use a local username/password, TACACS+ or RADIUS

```
aaa new-model
aaa authentication login default radius
aaa authorization exec radius
crypto key generate rsa usage 1024
ip http server 8080 (use a non-standard port)
ip http authentication aaa
ip http access-class <1-99>
```

- Disable

```
no ip http server
```

Core Dumps

- **Cisco routers have a core dump feature that will allow ISPs to transfer a copy of the core dump to a specific FTP server**
- **Set up a FTP account on the server the router will send the core dump to**
- **The server should NOT be a public server**

  Use filters and secure accounts

  Locate in NOC with NOC staff access only

  Enough disk space to handle the dumps

Core Dumps

- **Example configuration:**

```
ip ftp username cisco
ip ftp password 7 045802150C2E
ip ftp source-interface loopback 0
exception protocol ftp
exception dump 169.223.32.1
```

Routing Protocol Security

- **Routing protocol can be attacked**

  Denial of service

  Smoke screens

  False information

  Reroute packets

May Be Accidental or Intentional

Routing Protocol Authentication

- **Certifies Authenticity of Neighbor and Integrity of Route Updates**

- **Shared key included in routing updates**

  Plain text—protects against accidental problems only

  Message Digest 5 (MD5)—protects against accidental and intentional problems

Plain Text Authentication

- A router sends a routing update with a key and the corresponding key number to the neighbour router. For protocols that can have only one key, the key number is always zero.
- The receiving (neighbour) router checks the received key against the same key stored in its own memory.
- If the two keys match, the receiving router accepts the routing update packet. If the two keys did not match, the routing update packet is rejected.

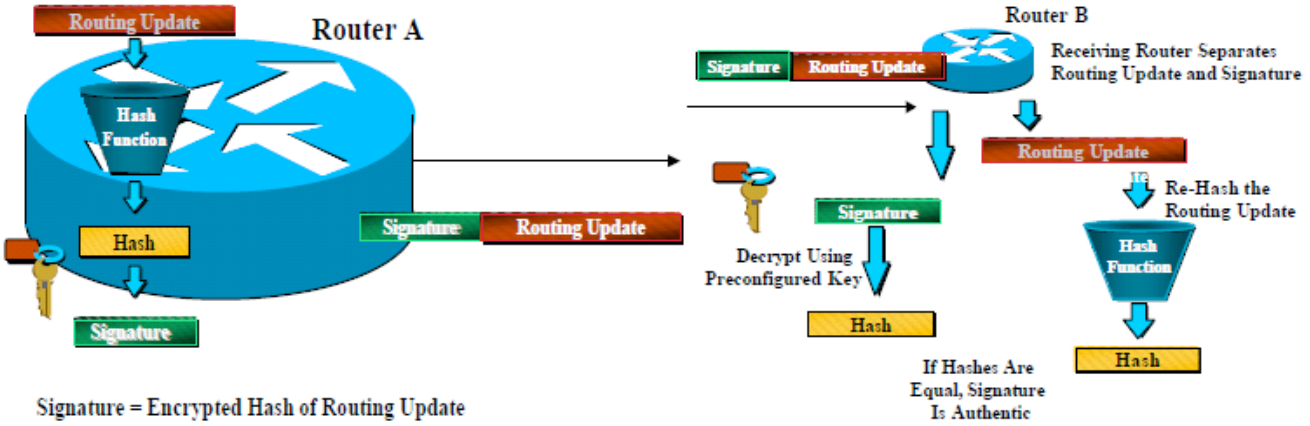

MD5 Authentication

- MD5 authentication works similarly to plain text authentication, except that the key is never sent over the wire.
- Instead, the router uses the MD5 algorithm to produce a “message digest” of the key (also called a “hash”).
- The message digest is then sent instead of the key itself. This ensures that nobody can eavesdrop on the line and learn keys during transmission. These protocols use MD5 authentication:

Routing Protocol Authentication

- **Interior Routing Protocols**
  - RIP v2 route authentication
  - EIGRP route authentication
  - OSPF route authentication
  - ISIS route authentication
- **Exterior Routing Protocols**
  - BGP route authentication

Route Authentication RIP v2: Cisco IOS

```
int Serial AA/BB

  ip rip authentication mode md5

  ip rip authentication key-chain rip-keys

key chain rip-keys

      key 1

      key-string <some password>



  router rip

  version 2

  passive interface default

  no passive Serial AA/BB

  redistribute static

  network X.X.X.X

  no default-information out

  no auto-summary
```

Route Authentication EIGRP

```
int Serial AA/BB
  ip authentication mode eigrp 16799 md5
  ip authentication key-chain eigrp 16799 eigrp-keys
key chain eigrp-keys
      key 16799
      key-string <some password>

  router eigrp 6799
  eigrp log-neighbor-changes
  eigrp log-neighbor-warnings 60
  passive interface default
  no passive Serial AA/BB
  redistribute connected
  redistribute static
  network X.X.X.X
  no auto-summary
```

Route Authentication OSPF

```
! OSPF Route Authentication to our Europe ISP



int Serial A A / B B

   ip ospf network non-broadcast

   ip ospf message-digest-key 1 md5 < password >

router ospf 1

   log-adjacency-changes

   passive-interface default

   no passive interface SerialA A / B B

   neighbor X.X.X.X

   network X.X.X.X Y.Y.Y.Y area 0

   area 0 authentication message-digest
```

BGP Route Authentication

```
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 C2Ebgp
```

BGP Route-Filtering

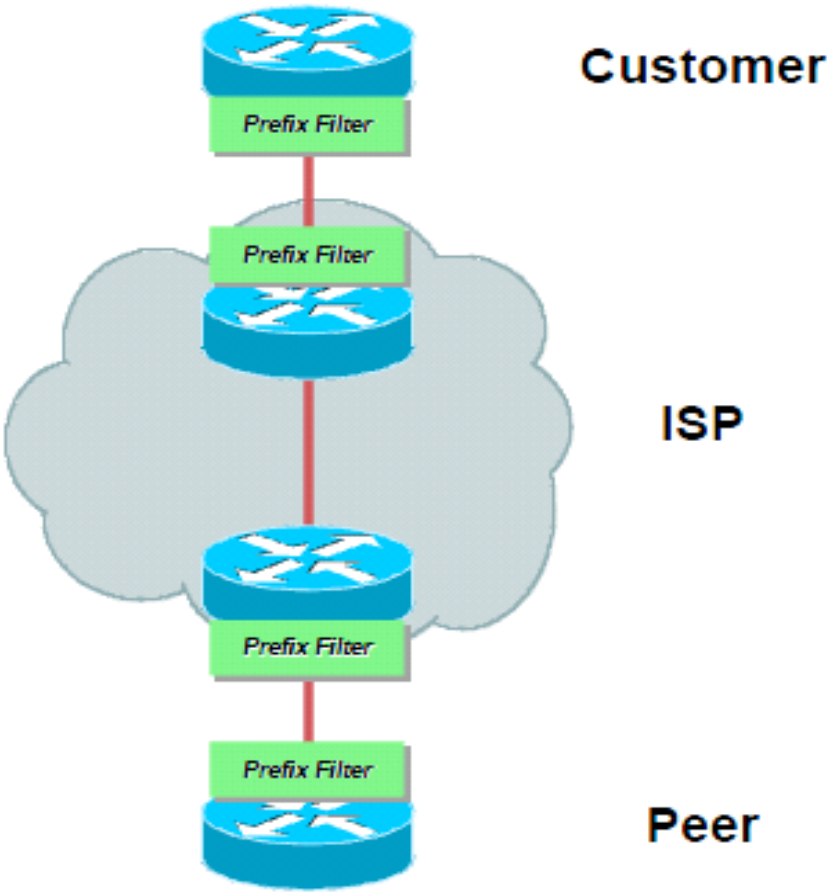

Ingress and Egress Route Filtering

- **Two flavours of prefix filtering**

  Distribute list—Now obsolete

  Prefix list—Widely used, higher performance

- **Two filtering techniques:**

  Explicit Permit (permit then deny any)

  Explicit Deny (deny then permit any)

Ingress and Egress Route Filtering

ip prefix-list deny-sua permit 192.168.0.0/16 **(Exact Match)**

ip prefix-list deny-sua permit 192.168.0.0/16 le 20

192.168.0.0/16 192.168.0.0/18 192.168.0.0/24

ip prefix-list deny-sua permit 192.168.0.0/16 ge 18

192.168.0.0/16 192.168.17.0/20 192.168.2.0/24

BGP with Prefix-List Flavour of Route Filtering

```
router bgp 200
no synchronization
neighbor 220.220.4.1 remote-as 210
neighbor 220.220.4.1 version 4
neighbor 220.220.4.1 prefix-list deny-sua in
neighbor 220.220.4.1 prefix-list deny-sua out
neighbor 222.222.8.1 remote-as 220
neighbor 222.222.8.1 version 4
neighbor 222.222.8.1 prefix-list deny-sua in
neighbor 222.222.8.1 prefix-list deny-sua out
no auto-summary
```

Where to use Prefix List

- **Customer's Ingress/Egress**
- **ISP Ingress on Customer** (may Egress to Customer)
- **ISP Egress to Peer**

  and Ingress from Peer
- **Peer Ingress from**

  ISP and Egress to ISP

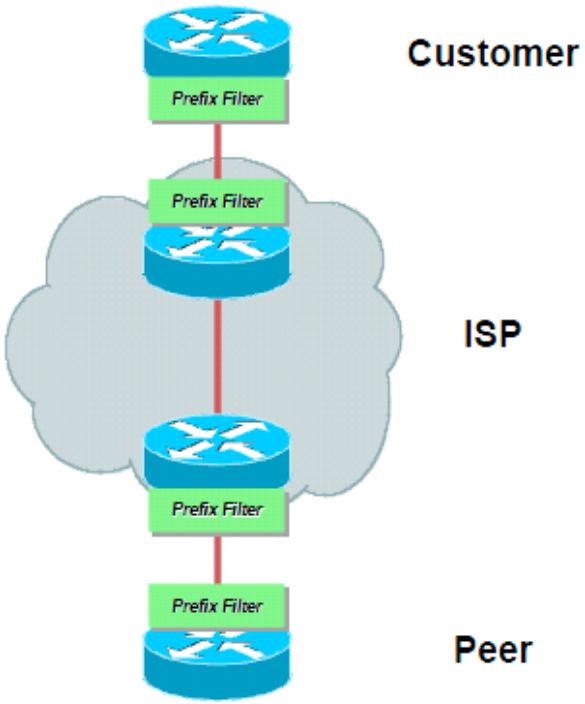

What to Prefix List

- **There are routes that should NOT be routed on the Internet**

RFC 1918

127.0.0.0/8 and multicast blocks

Certain RFC3330 addresses:

http://www.rfc-editor.org/rfc/rfc3330.txt

- **BGP should have filters applied so that these routes are not advertised to or propagated through the Internet**

Prefix Filters on Customers

- **Prefix filter all routes from your customers!**

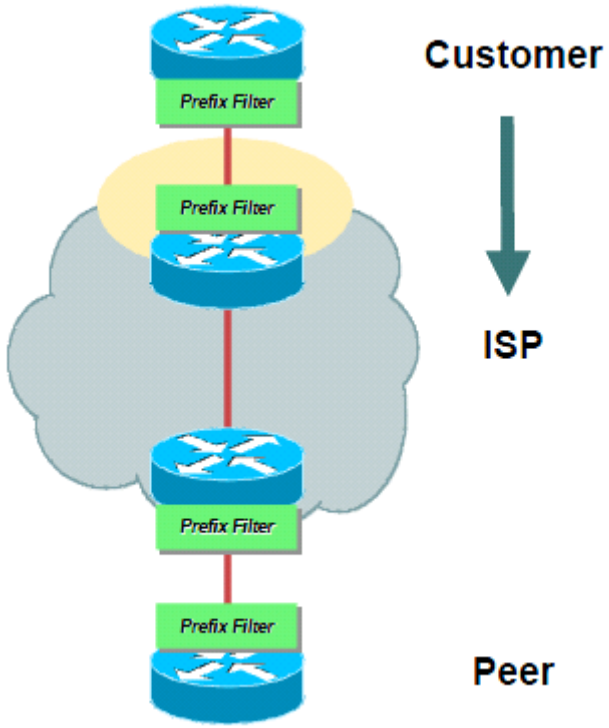

Prefix Filters on Customers

- **ISPs should only accept prefixes which have been assigned or allocated to their downstream peer/customer.**
- **For example**

  Downstream has 220.50.0.0/20 block Should only announce this to peers

  Peers should only accept this from them

Prefix Filters on Customers

- **Configuration example on upstream:**

  router bgp 100

  neighbor 222.222.10.1 remote-as 101

  neighbor 222.222.10.1 prefix-list customer in!

  ip prefix-list customer permit 220.50.0.0/2

  ip prefix-list customer deny 0.0.0.0/0 le 32

Prefixes to Peers

- **Prefix filter all routes to your peers!**

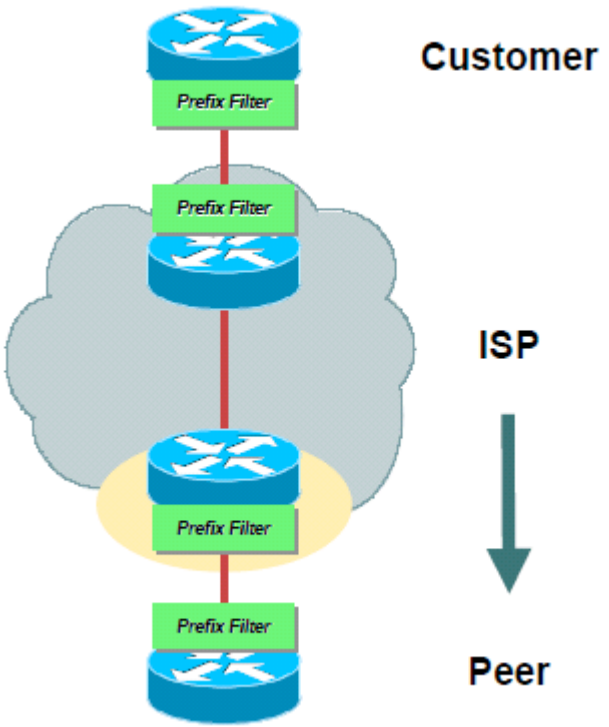

Prefixes to Peers

- **What do you send to the Internet?**

  Your prefixes.

  More specific customers prefixes (customers who are multihoming)

- **What do you *not* send to the Internet?**

  RFC3330 Prefixes – assume junk will leak into your iBGP.

  Bogons – assume garbage will leak into your iBGP.

  Lower Prefix Boundary – Unless absolutely necessary, do no allow anything in the /25 - /32 range.

Prefixes from Peers

- Prefix filter all routes from your peers!

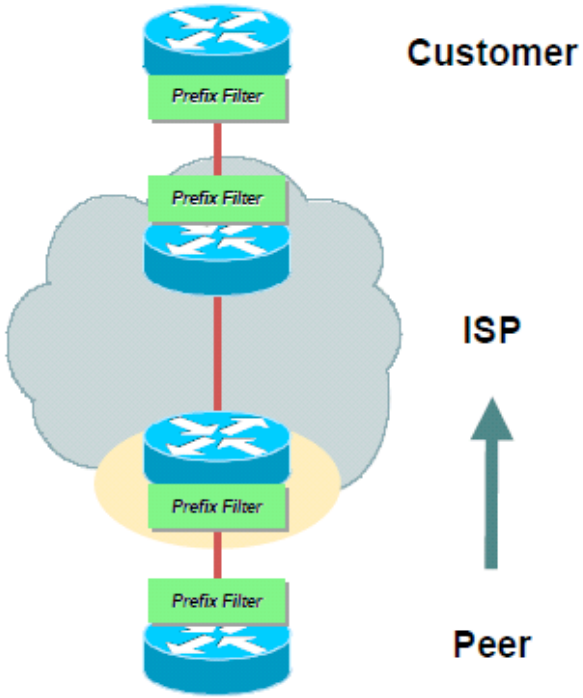

Ingress Routes from Peers or Upstream

- **Ingress Routes from Peers and/or the Upstream ISP are the nets of the Internet.**
- **Ideally, the peering policy should be specific so that exact filters can be put in place.**

  Dynamic nature of the peering makes it hard to maintain specific route filters.

Receiving Prefixes from Upstream & Peers (ideal case)

Don’t accept RFC1918 etc prefixes

Don’t accept your own prefix

Don’t accept default (unless you need it)

Don’t accept prefixes longer than /24

Consider *Net Police* Filtering

Receiving Prefixes — Cisco IOS

- **router bgp 100**
- **network 221.10.0.0 mask 255.255.224.0**
- **neighbor 221.5.7.1 remote-as 101**
- **neighbor 221.5.7.1 prefix-list in-filter in**
- **!**
- **ip prefix-list in-filter deny 0.0.0.0/0** ! Block default
- **ip prefix-list in-filter deny 0.0.0.0/8 le 32**
- **ip prefix-list in-filter deny 10.0.0.0/8 le 32**
- **ip prefix-list in-filter deny 127.0.0.0/8 le 32**
- **ip prefix-list in-filter deny 169.254.0.0/16 le 32**
- **ip prefix-list in-filter deny 172.16.0.0/12 le 32**
- **ip prefix-list in-filter deny 192.0.2.0/24 le 32**
- **ip prefix-list in-filter deny 192.168.0.0/16 le 32**
- **ip prefix-list in-filter deny 221.10.0.0/19 le 32 ! Block local prefix**
- **ip prefix-list in-filter deny 224.0.0.0/3 le 32**
- **ip prefix-list in-filter deny 0.0.0.0/0 ge 25 ! Block prefixes >/24**
- **ip prefix-list in-filter permit 0.0.0.0/0 le 32**

AS-Path Filters

AS Path Regular Expressions

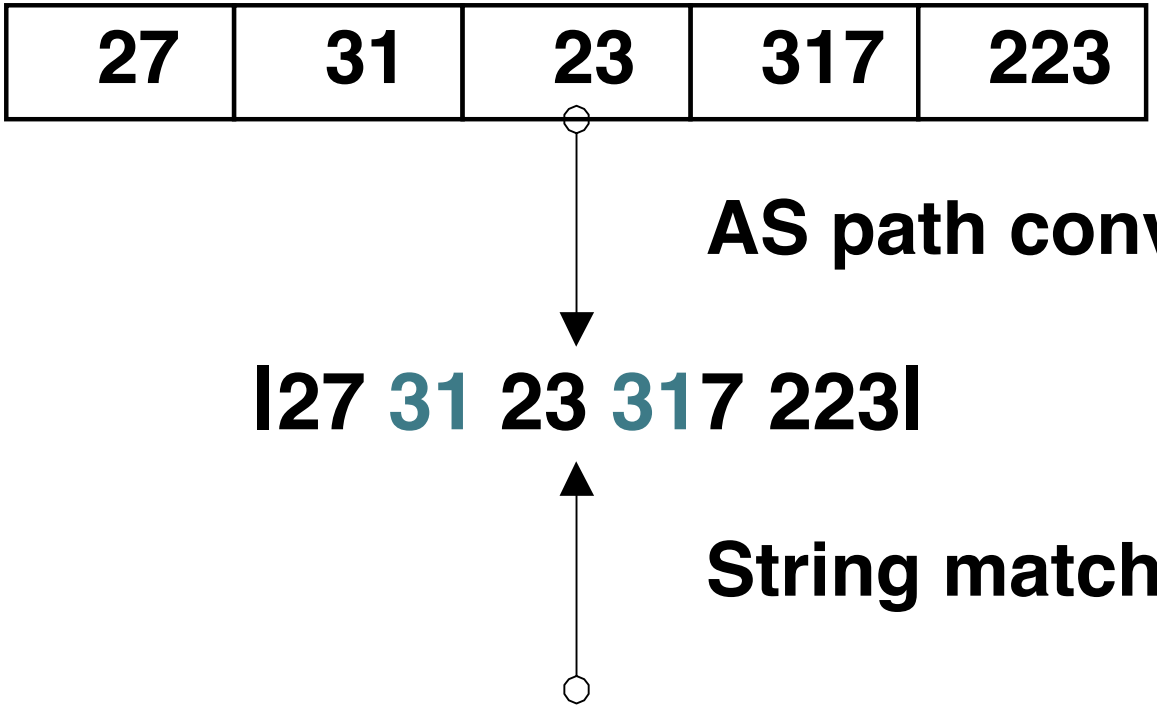

I.e. When to use AS Path Filtering

When you require broad level of filtering

Regular Expressions Alternatives

- **Expression *expr1*|*expr2***

Pipe means OR

how many times does 21|31 match

|2 1 3 3 1 7 2 3 1 6 3 1 |

answer:

|2 1 3 3 1 7 2 3 1 6 3 1 |

Regular Expressions
Ranges and Wildcard Characters

- **Bracket indicates type of range in this case you have two type of range**
  examples:[1234] or [1-4]
  Can 1 or 2 or 3 or 4. put a comma
  Any number in this range
- **dot (.) matches any single character that includes a space**

how many times does [1-3].[34] match

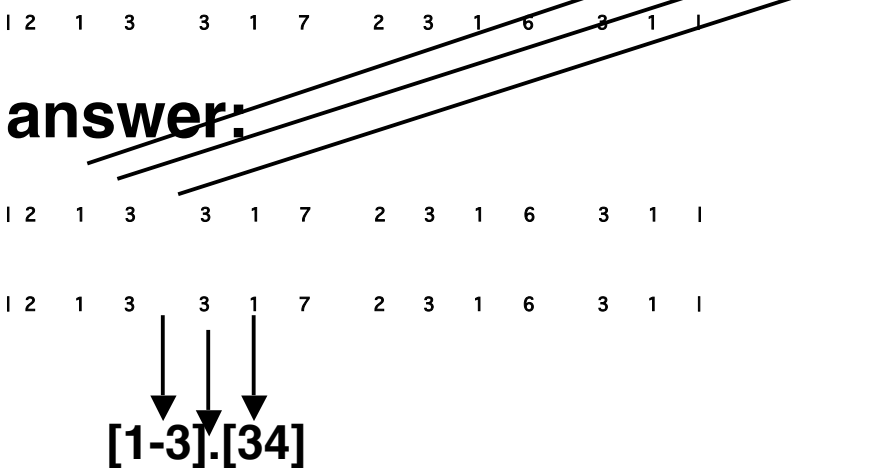

answer:

Regular Expressions Matching Delimiters

^ matches beginning of string

$ matches end of string

_ matches any delimiter (beginning, end, whitespace, tab, comma but no number)

how many times does ^21, 31$, _31_ match

| 2 | 1 | 3 | 3 | 1 | 7 | 2 | 1 | 8 | 3 | 1 | 7 | 3 | 1 |

answer:

| 2 | 1 | 3 | 3 | 1 | 7 | 2 | 1 | 8 | 3 | 1 | 7 | 3 | 1 |

Regular Expressions Grouping

Parenthesis can be used to group smaller regular expressions into larger expressions

how many times does (213|218)_31 match

| 2 1 3 3 1 7 1 2 1 8 3 1 6 3 1 |

answer:

| 2 1 3 3 1 7 1 2 1 8 3 1 6 3 1 |

Regular Expressions
Repeating Operators

*	matches zero or more atoms
?	matches zero or one atom
+	matches one or more atoms

Atom is a single character or a grouping

how do you match AS sequences “23 45” and “23 78 45” in single regular expression

answer:

```
_ 2 3 ( _ 7 8 ) ? _ 4 5 _
```

Sample Regular Expressions

_100_	Going through AS 100
^100$	Directly connected to AS 100
_100$	Originated in AS 100
^100_.	networks behind AS 100
^$	networks originated in local AS
.*	matches everything

Configuring BGP AS-path Filters

router(config)#

ip as-path access-list *number* permit|deny *regexp*

- **Configures AS-path access list**

router(config-router)#

neighbor *ip-address* filter-list *as-path-filter* in|out

- **Configures inbound or outbound AS-path filter for specified BGP neighbor**

Testing your Regular Expressions

router#

show ip bgp regexp *expression*

- **Displays all routes in BGP table matching regular expression**

router#

show ip bgp filter *filter-list*

- **Displays all routes in BGP table permitted by the specified AS-path access list**

router#

show ip as-path-access-list [*filter-list*]

- **Displays one or all filter lists**

BGP MAXIMUM PREFIX

BGP Maximum Prefix Tracking

- **Allow configuration of the maximum number of prefixes a BGP router will receive from a peer**

- **Two level control**

  Warning threshold: Log warning message

  Maximum: Tear down the BGP peering

BGP Maximum-Prefix Tracking

router(config-router)#

neighbor *ip-address* maximum-prefix *maximum* [*threshold*] [warning-only]

- **Controls how many prefixes can be received from a neighbor**
- **Optional threshold parameter specifies the percentage where a warning message is logged (default is 75%)**
- **Optional warning-only keyword specifies the action on exceeding the maximum number (default is to drop neighborship)**

Monitoring Maximum-Prefix Operation

router#

show ip bgp neighbor [*address*]

- **For neighbors with maximum-prefix configured displays the maximum number of prefixes and the warning threshold**
- **For neighbors exceeding the maximum number of prefixes displays the reason the BGP session is idle**

THANKS