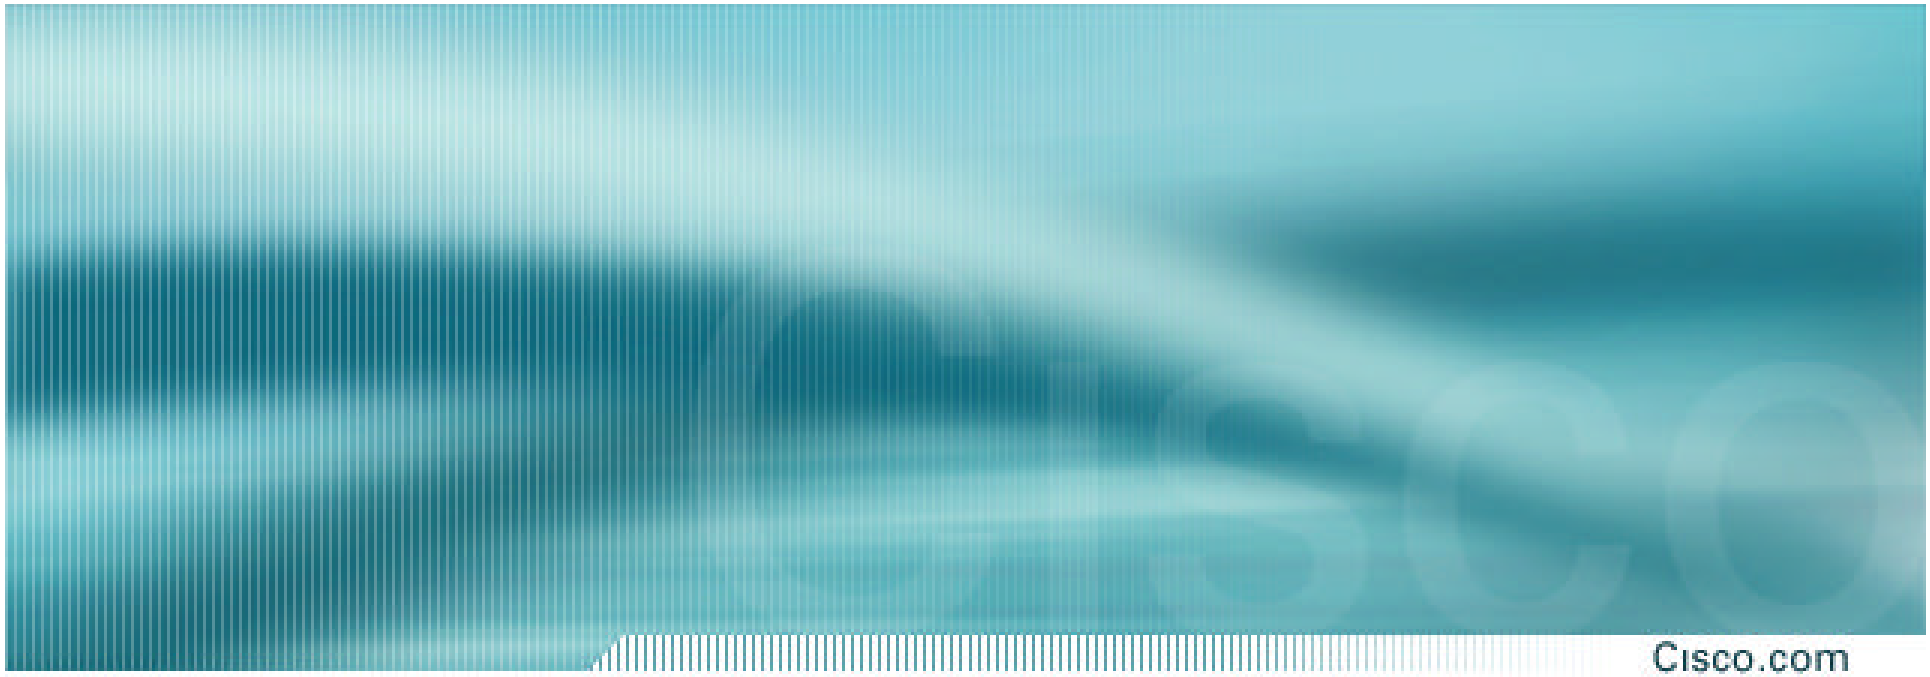

ISP Essentials

Essential IOS Features every ISP should Consider

Version 3.0

SANOG 2, Sri Lanka, July 2003

Philip Smith <pfs@cisco.com>

Presentation Slides

Cisco.com

- **Will be available on**

  ftp://ftp-eng.cisco.com/pfs/seminars

- **Feel free to ask questions any time**

Background

- **This presentation is based on content from the Cisco ISP Essentials book**

  Cisco Press ISBN 1-58705-041-2

  www.ciscopress.com to buy it ☺

  www.ispbook.com for updates

Cisco ISP Essentials

Cisco.com

- **IOS Software and Router Management**
- **General Features**
- **Routing Configuration Guidelines**
- **Securing the Router**
- **Securing the Network**

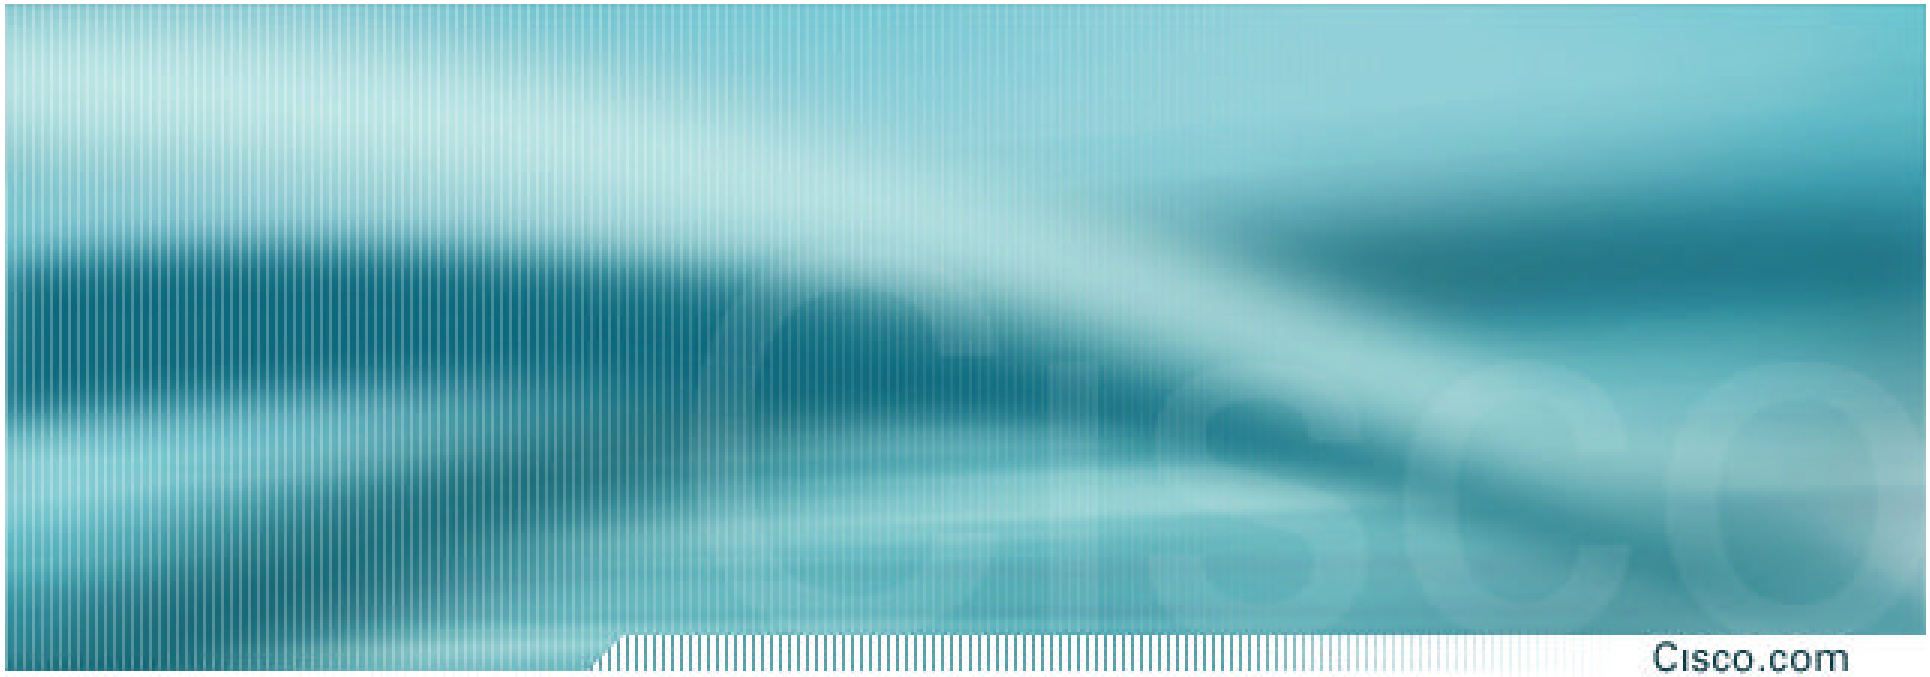

IOS Software and Router Management

Which IOS?

- **IOS is a feature rich and highly complex router control system**
- **ISPs should choose the IOS variant which is suitable for their needs**

  As with any router hardware, one size does not fit all!

Which IOS?

- **There is an exclusive service provider train in IOS**
  - **This is 12.0S, for 7200, 7500, 10000 and 12000**
  - **Images also available for 2500, 2600, 3600 and 4500, but are completely unsupported**
- **There is a service provider image in most IOS releases**
  - **This is the image with –p– in its name, for example:**
    - **c7200-p-mz.122-8.T1 and c2600-p-mz.121-14**
  - **The –p– image is IP-only plus ISIS/CLNS**

Which IOS?

- **12.*n* – for example 12.2**
  - **This means the IOS is a mainline image**
  - **NO new features**
  - **ONLY bug fixes**
  - **The aim is stability!**
- **12.*n*T – for example 12.2T**
  - **This means the IOS is the technology release**
  - **NEW features**
  - **Bug fixes**
  - **Avoid unless you need the feature!**

12.0 IOS release images

- **12.0S is the release for all ISPs**
  - **for 7200, 7500, 10000 and GSR/12000**
  - **currently at 12.0(25)S1**
  - **Non-MPLS images at 12.0(21)S7**
- **12.0 is the “mainline” train**
  - **for the older platforms not able to support 12.1+**
  - **currently at 12.0(27)**
- **Available on CCO, supported by TAC**

12.1 IOS release images

- **12.1 is the more recent “mainline” train**
  - **Comes from 12.0T, currently at 12.1(20)**
  - **Supports more platforms and has more features than 12.0**
  - **For older platforms not able to support 12.2+**
- **12.1E is the enterprise train**
  - **Started off as the 7600/Cat6500 train**
  - **Has many of the features from 12.0S**
  - **Last release was 12.1(14)E4, now part of 12.2S/12.3**
- **Available on CCO, supported by TAC**

12.2 IOS release images

- **12.2 is the previous “mainline” train**
  - **Originated from 12.1T, currently at 12.2(17a)**
  - **Supports more platforms and has more features than 12.1**
  - **Recommended for non-”S-image” platforms**
- **12.2S merges 12.0S and 12.1E**
  - **for 7100, 7200, 7400, 7500 and 7600/Cat6K**
  - **currently at 12.2(14)S3**
- **Available on CCO, supported by TAC**

12.3 IOS release images

- **12.3 is the new “mainline” train**
  - **Originated from 12.2T, currently at 12.3(1a)**
  - **Includes IPv6 in IP+ images**
  - **Still early in development cycle**
- **12.3T is the “technology train”**
  - **New features introduced for IOS 12.3**
  - **First release due end July 2003**
- **Available on CCO, supported by TAC**

Cisco IOS Roadmap

Cisco.com

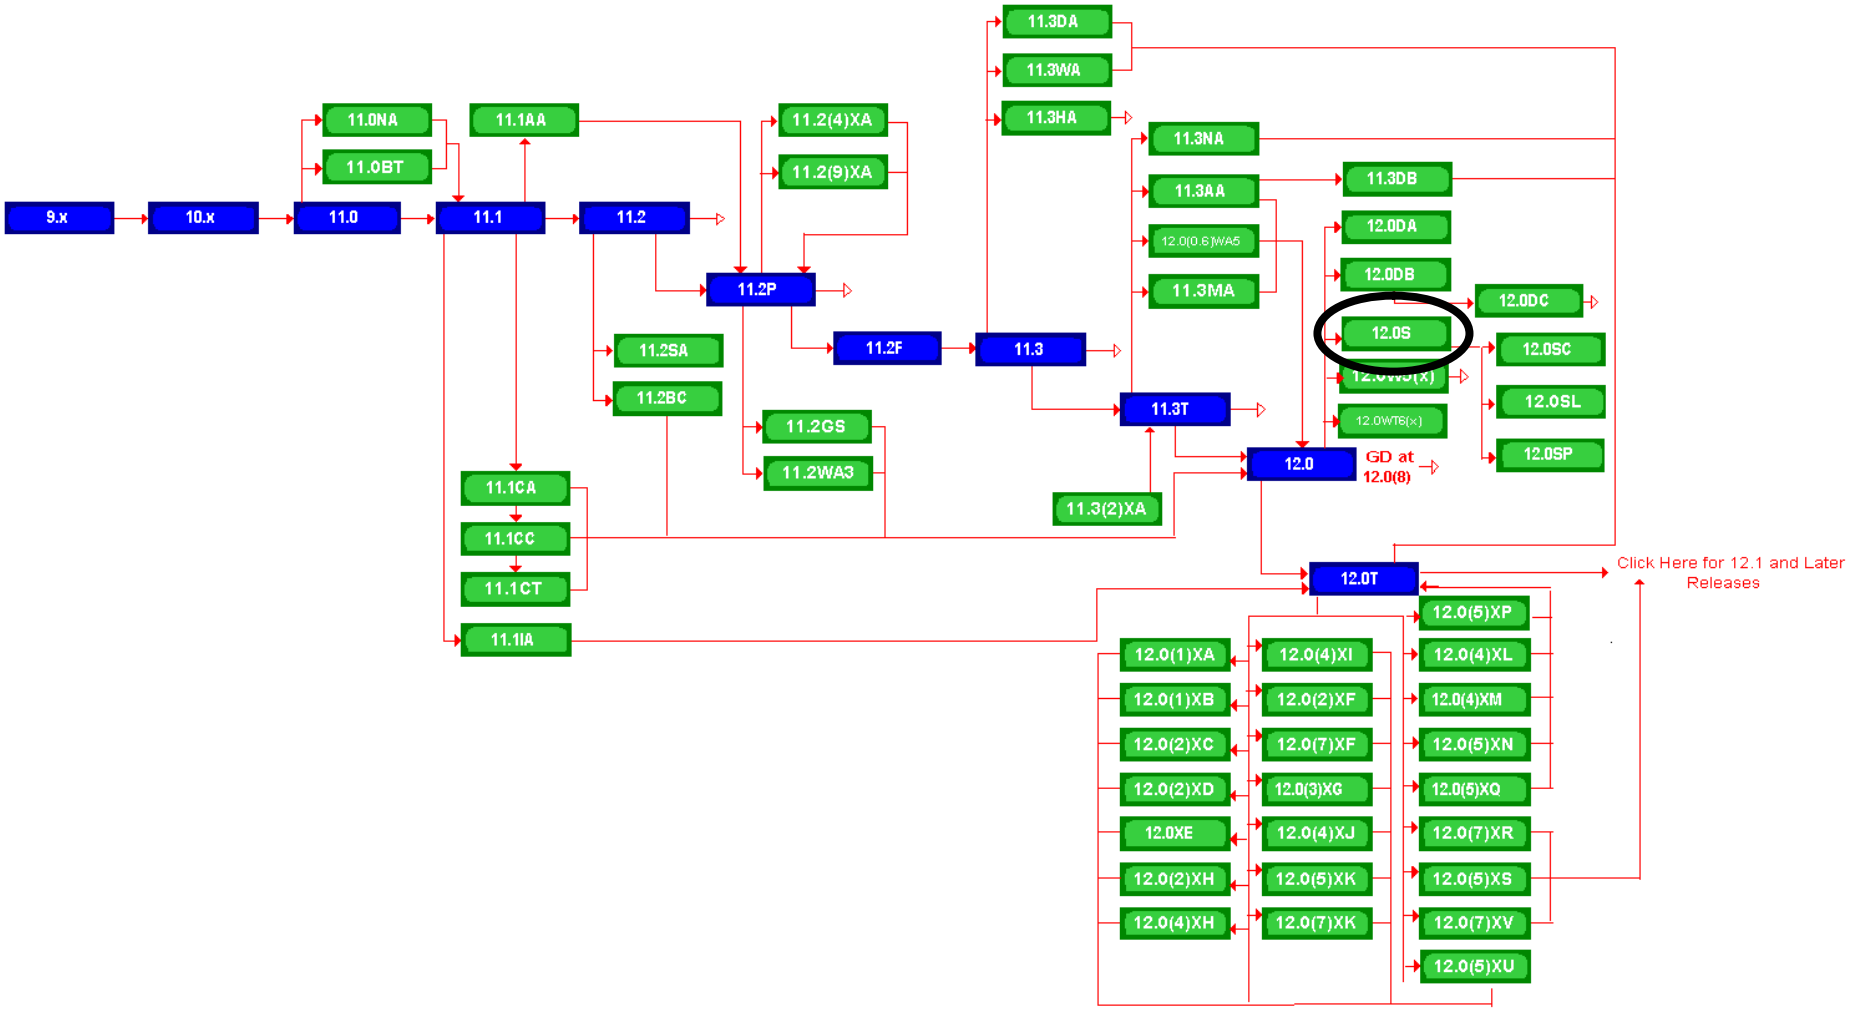

http://www.cisco.com/warp/public/620/roadmap.shtml

Cisco IOS Roadmap

Cisco.com

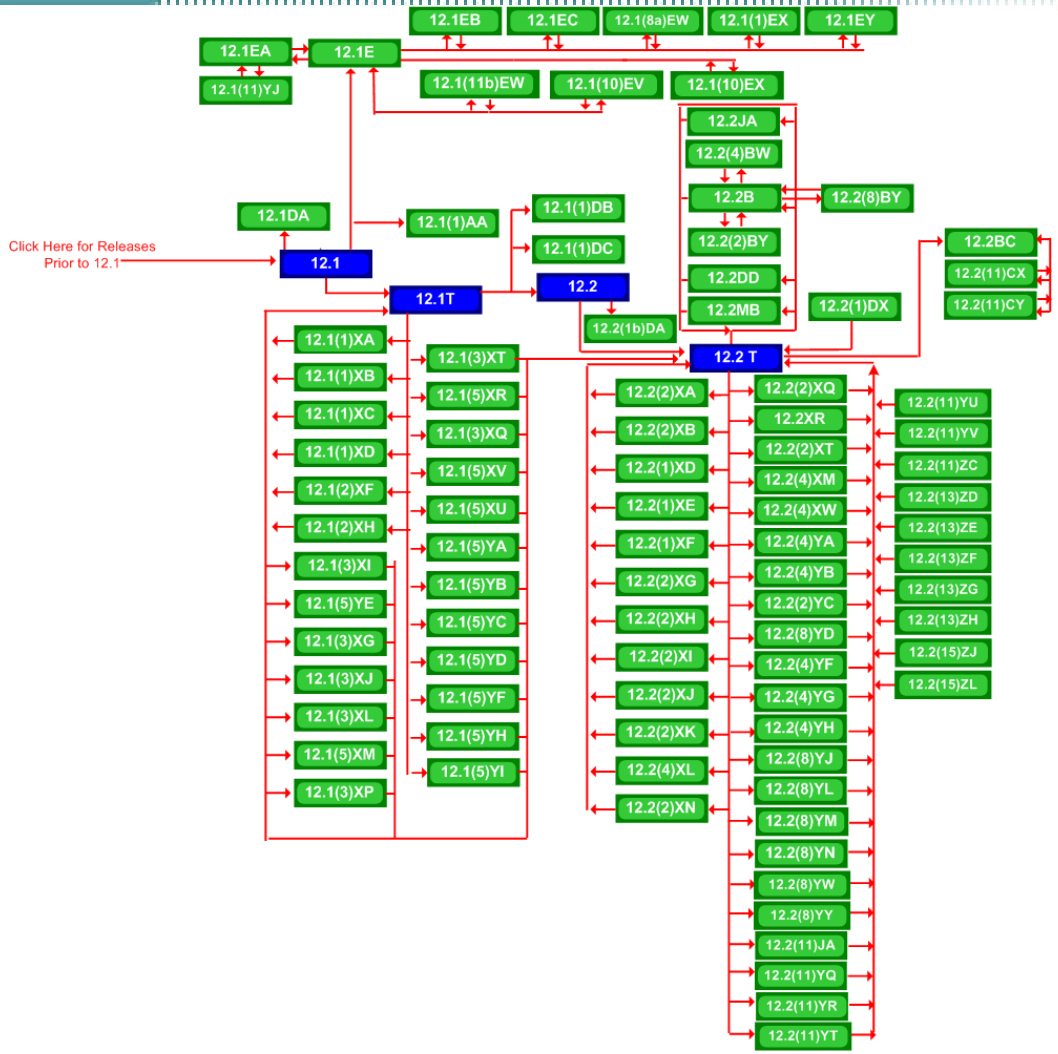

http://www.cisco.com/warp/public/620/roadmap_b.shtml

Which IOS?

Cisco.com

- **IOS Choice Matrix:**
  - **Is there a 12.0S for my router?**
    - **If not:**
      - **Is there a 12.2 for my router?**
        - **If not:**
          - **Use 12.1, and failing that, 12.0**
  - **Only use 12.*n*T image if you need the feature in the Technology Train**
- **Pick the –p– image if it exists, otherwise –i– is usually all that is necessary, e.g.:**

c7200-k4p-mz.120-21.S7	**ISP SSH/3DES image for 7200**
c3620-p-mz.122-12	**ISP ServiceProvider image for 3620**
c2500-i-l.121-20	**IP image for 2500**

IOS Software Management
Flash Memory

- **Good practice is to have at least two distinct flash memory volumes**
  - **allows backup image(s)**
  - **back out path in case of upgrade problems**
- **Partition the built-in flash**

```
partition flash 2 16 16
```

- **Install a PCMCIA flash card in external slot(s)**

IOS Software Management
Flash Memory

- **Ensure there is a configured back for the selected IOS image**
  - **Backup image is previous good image**

```
boot system flash slot0:rsp-k4pv-mz.120-23.S1
boot system flash slot1:rsp-k4pv-mz.120-21.S7
boot system flash
```

  - **Which means:**
    - **Boot quoted image from slot0:. If it isn't there, boot the quoted image in slot1:. If that isn't there, try the first image available in flash**

IOS Software Management
System Memory

- **Good practice is to maximise router memory**
  - **allows for the rapidly growing Internet**
- **At least 128Mbytes RAM needed for full Internet routing table**
- **Recognised that equipment works best when “left alone”**

IOS Software Management
When to Upgrade

- **Upgrades needed when:**
  - **bug fixes released**
  - **new hardware support**
  - **new software features required**
- **Otherwise:**

If it isn’t broken, don’t fix it!

(Digression)
Loopback Interface

- **Most ISPs make use of the router loopback interface**
  - **OSPF router id**
  - **iBGP peering mesh**
- **IP address configured is a host address**
- **Configuration example:**

```
interface loopback 0
 description Loopback Interface of CORE-GW3
 ip address 215.18.3.34 255.255.255.255
```

(Digression)
Loopback Interface

- **Loopback interfaces on ISP backbone usually numbered:**
  - **out of one contiguous block, or**
  - **using a geographical scheme, or**
  - **using a per PoP scheme**
- **Aim is to increase stability, aid administration, and improve security**

(Digression)
Loopback Interface

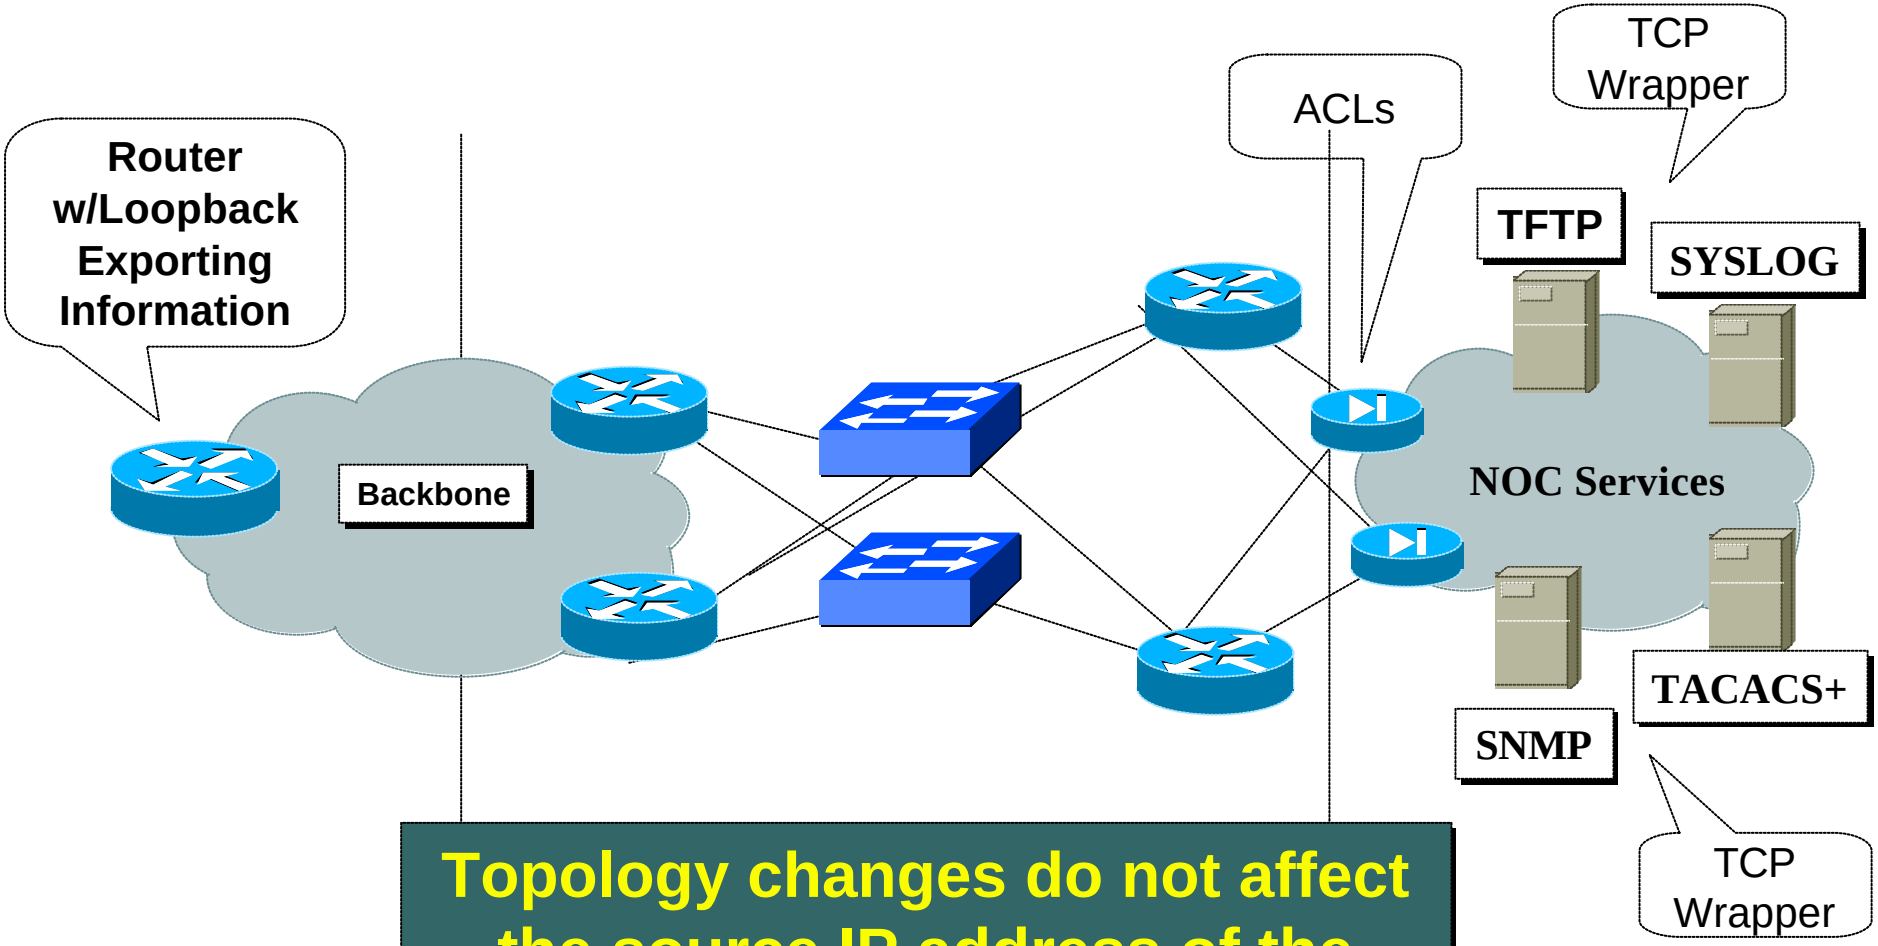

Topology changes do not affect the source IP address of the packets coming from the Router

(Digression)
Loopback Interface

- **Loopback interface is not “redundant” or “superfluous”**
- **Multitude of uses to ease security, access, management, information and scalability of router and network**
- **Protects the ISP’s Management Systems**
- **Use the loopback!**

Configuration Management

Cisco.com

- **Backup NVRAM configuration off the router:**
  - **write configuration to TFTP server**
  - **TFTP server files kept under revision control**
  - **router configuration built from master database**
- **Allows rapid recovery in case of emergency**

Configuration Management

- **Secure the TFTP server**
  - **TFTP source interface Loopback 0 on Router**
  - **Firewall/ACL**
  - **Wrapper on TFTP Server which only allows the router's loopback address**

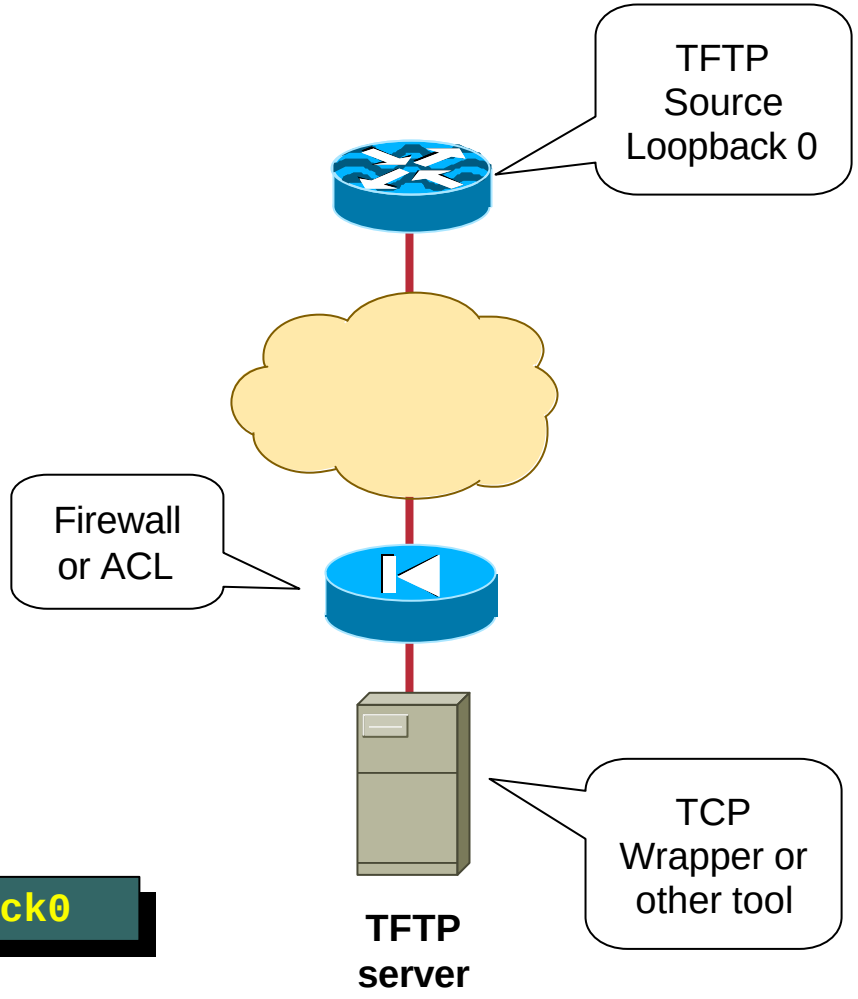


```
ip tftp source-interface Loopback0
```

FTP Client Support

- **TFTP has security and file size limitations**
- **FTP Client support is added in 12.0; this allows for FTP upload/downloads.**
- **Remember to use the same security/redundancy options with loopback 0:**

```
ip ftp source-interface loopback 0
```

FTP Client Support


```
7206-pfs-bne#copy ftp://pfs:XXX@ftp.cisco.com slot0:

Source filename []? /cisco/ios/12.0/12.0.21S7/7200/c7200-k4p-
mz.120-21.S7.bin

Destination filename [c7200-k4p-mz.120-21.S7.bin]?

Accessing ftp://pfs:XXX@ftp.cisco.com
//cisco/ios/12.0/12.0.21S7/7200/c7200-k4p-mz.120-
21.S7.bin...Translating "ftp.cisco.com"...domain server
(207.126.96.162) [OK]


Loading /cisco/ios/12.0/12.0.21S7/7200/c7200-k4p-mz.120-
21.S7.bin
```

Larger Configurations

Cisco.com

- **Compress Configuration**

  Used when configuration required is larger than configuration memory (NVRAM) available.

  ```
  service compress-config
  ```

- **FLASH or remote server**

  Used when NVRAM compression is not enough

Command Line Interface Features

Cisco.com

- **Some Convenient Editing Keys**

TAB	**command completion**
arrow keys	**scroll history buffer**
ctrl A	**beginning of line**
ctrl E	**end of line**
ctrl K	**delete all chars to end of line**
ctrl X	**delete all chars to beginning of line**
ctrl W	**delete word to left of cursor**
esc B	**back one word**
esc F	**forward one word**

Command Line Interface Features

- **CLI now has string searches**

  ```
  show configuration | [begin|include|exclude] <regexp>
  ```

- **Pager "--more--" now has string searches**

  ```
  /<regexp>, -<regexp>, +<regexp>
  ```

- **"More" command has string searches**

  ```
  more <filename> | [begin|include|exclude] <regexp>
  ```

Command Line Interface Features

- **Example:**

  Show running configuration from the point where BGP is configured

```
Defiant#show running-config | begin ^router bgp
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 route-map Community1 out
 neighbor 4.1.2.1 route-map Community2 in
!
```

System Logs

- Off load router syslog information to a syslog 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
!
no logging console
logging buffered 16384
logging trap debugging
logging facility local7
logging 169.223.32.1
logging 169.223.35.8
logging source-interface loopback0
```

System Logs: Topologies Used

- **Centralised Syslog Servers in Operations Centres**

  All logs in one place – easy to view, but could be single point of failure

  Network congestion may cause loss of messages

- **Syslog Servers in Major POPs**

  Distributed system, may be hard to view, and need collation

  Solves network congestion problem

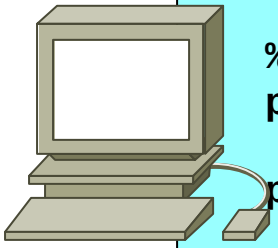

```
pfs-pc% tail -1 cisco.log
  Feb 17 21:48:26 [10.1.1.101.9.132] 31: Feb 17 11:51:55 AEST:
%SYS-5-CONFIG_I: Configured from console by vty0 (10.1.1.2)
pfs-pc% date
  Tue Feb 17 21:49:53 AEST 1998
pfs-pc%
```

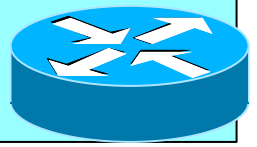

Network Time Protocol

- **If you want to cross compare logs, you need to synchronize the time on all the devices**
- **Use NTP**

  from external time source

  Upstream ISP, Internet, GPS, atomic clock

  from internal time source

  router can act as stratum 1 time source

Network Time Protocol

- **Set timezone**

```
clock timezone <name> [+/-hours [mins]]
```

- **Router as source**

```
ntp master 1
```

- **External time source (higher stratum)**

```
ntp server a.b.c.d
```

- **External time source (equivalent stratum)**

```
ntp peer e.f.g.h
```

Network Time Protocol

- **Example Configuration:**

```
clock timezone SST 8
ntp update-calendar
ntp source loopback0
ntp server <other time source>
ntp peer <other time source>
ntp peer <other time source>
```

Network Time Protocol

- **Where to get NTP Reference Sources?**

  http://www.eecis.udel.edu/~ntp/hardware.html

- **Attaching a Telecom Solutions GPS Clock to the Router's AUX port:**

```
Excalabur(config)#line aux 0
Excalabur(config-line)#ntp refclock telecom-solutions pps ?
  cts   PPS on CTS
  none  No PPS signal available
  ri    PPS on RI
```

SNMP

- **Remove any SNMP commands if SNMP is not going to be used.**
- **If SNMP is going to be used:**

```
access-list 98 permit 169.223.1.1
access-list 98 deny   any
snmp-server community 5nmc02m RO 98
snmp-server trap-source Loopback0
snmp-server trap-authentication
snmp-server host 169.223.1.1 5nmc02m
```

HTTP Server

- **HTTP Server in IOS from 11.1CC and 12.0S**

  router configuration via web interface

- **Disable if not going to be used:**

```
no ip http server
```

- **Configure securely if going to be used:**

```
ip http server
ip http port 8765
ip http authentication aaa
ip http access-class <1-99>
```

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 network operations staff access only

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
```

Cisco ISP Essentials

Cisco.com

- IOS Software and Router Management
- General Features
- Routing Configuration Guidelines
- Securing the Router
- Securing the Network

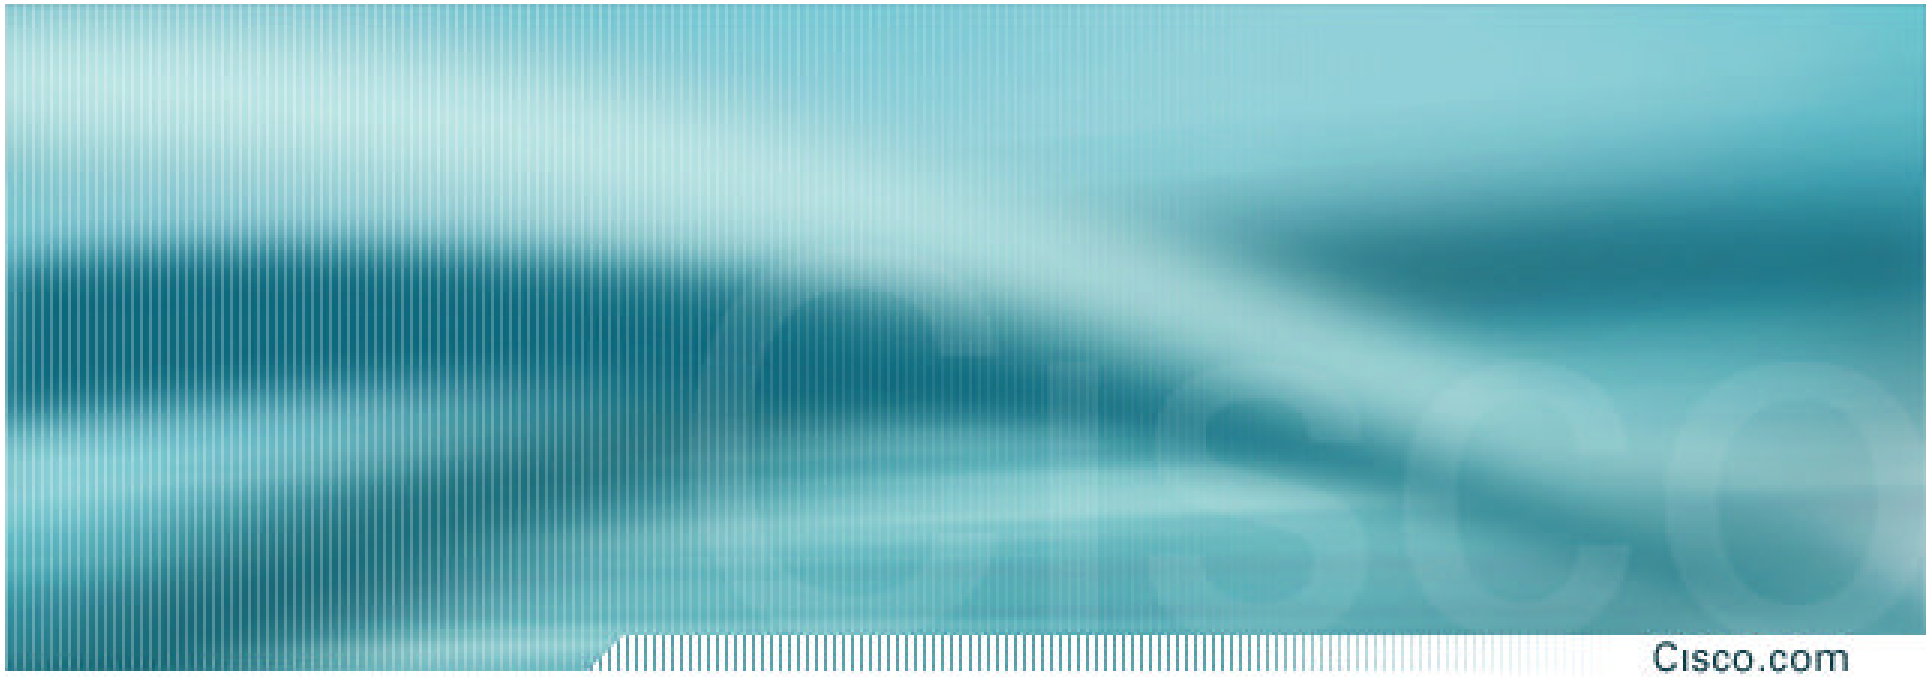

General Features

Interface Configuration

- **“ip unnumbered”**
  - **no need for an IP address on point-to-point links**
  - **keeps IGP small**
- **“description”**
  - **customer name, circuit id, cable number, etc**
  - **on-line documentation!**
- **“bandwidth”**
  - **used by IGP**
  - **documentation!**

Interface Configuration – Example

Cisco.com

- **ISP router**

```
!
interface loopback 0
description Loopback interface on GW2 Router
ip address 215.17.3.1 255.255.255.255
!
interface Serial 5/0
description 128K HDLC link to Galaxy
Publications Ltd [galpub1] WT50314E  R5-0
bandwidth 128
ip unnumbered loopback 0
!
ip route 215.34.10.0 255.255.252.0 Serial 5/0
```

- **Customer router**

```
!
interface Ethernet 0
description Galaxy Publications LAN
ip address 215.34.10.1 255.255.252.0
!
interface Serial 0
description 128K HDLC link to Galaxy
Internet Inc WT50314E  C0
bandwidth 128
ip unnumbered ethernet 0
!
ip route 0.0.0.0 0.0.0.0 Serial 0
```

Interface Status Checking

Cisco.com

- **`show interface switching`**

  Hidden command which provides information about the switching status of the router interfaces

- **`show interface stats`**

  Hidden command which provides inbound and outbound packet information on the router interfaces

- **`show idb` (interface descriptor blocks)**

  Shows how many IDBs are configured on the router

  Early routers (such as AGS+) could only support 300 IDBs

More Interface Features

- **By default, the load on the interface is calculated as an average over 5 minutes**

  ISPs tend to want higher resolution, for example, averaged over 30 seconds:

```
interface serial 0/0
 load-interval 30
```

- **Inbound hold-queue is only 75 spots**

  Should be increased to something more reasonable, especially with routers with large numbers of peers

```
interface serial 0/0
 hold-queue 4096 in
```

NetFlow

- **Provides network administrators with “packet flow” information**
- **Allows:**
  - **security monitoring**
  - **network management and planning**
  - **customer billing**
  - **traffic flow analysis**
- **Available from 11.1CC for 7x00 and 12.0 for remaining router platforms**

NetFlow – Capacity Planning

Cisco.com

Public Routers 1 , 2, 3 Month of September Outbound Traffic

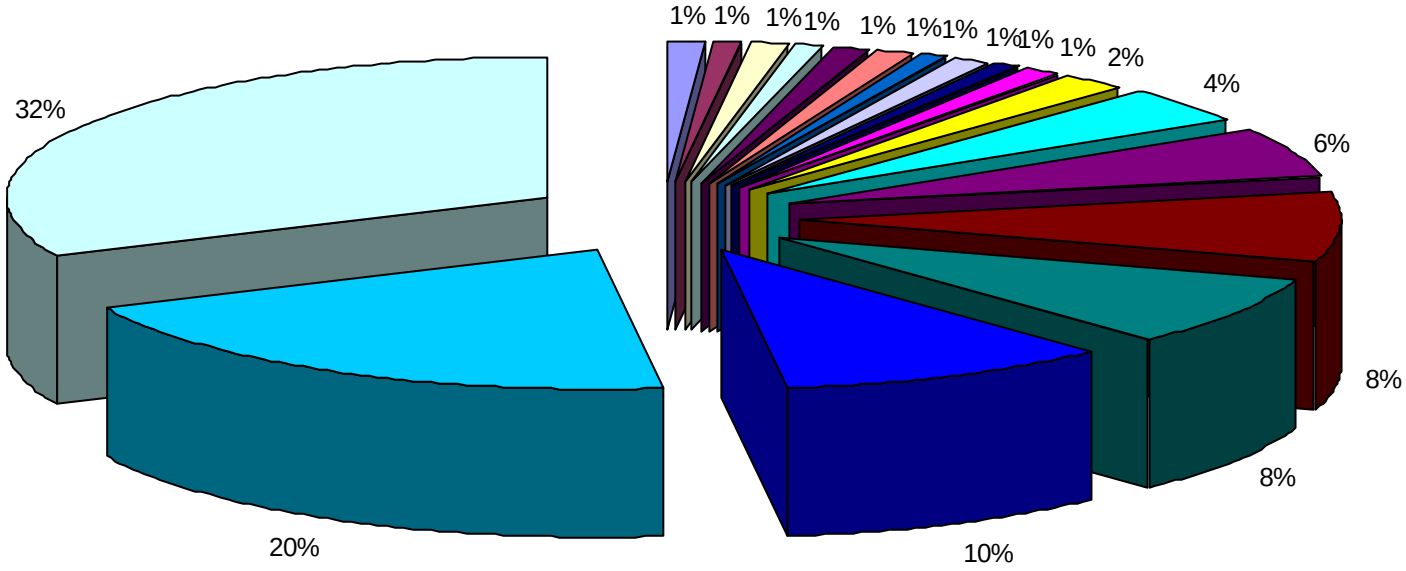

WEC	WebTV	ABSNET	AOL	Compuserve
SURAnet	IBM	OARNet	NIH	PacBell Internet Service
JHU	C&W	UMD	AT&T	BBN
Erols	Digex	Other		

NetFlow

- **Configuration example:**

```
interface serial 5/0
 ip route-cache flow
```

- **If CEF not configured, NetFlow enhances existing switching path**
- **If CEF configured, NetFlow becomes a flow information gatherer**

NetFlow

- **Information export:**

  router to collector system

  ```
  ip flow-export version 5 [origin-as|peer-as]
  ip flow-export destination x.x.x.x <udp-port>
  ```

- **Flow aggregation (new in 12.0S):**

  router sends aggregate records to collector system

  ```
  ip flow-aggregation cache as|prefix|dest|source|proto
    enabled
    export destination x.x.x.x <udp-port>
  ```

NetFlow

Cisco.com

- **Sample Output on router:**

```
Beta-7200-2>sh ip cache 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
Et1/1         144.254.153.10  Null          144.254.153.127 11 008A 008A     1
Et1/1         144.254.153.112 Null          255.255.255.255 11 0208 0208     1
Et1/1         144.254.153.50  Local         144.254.153.51  06 701D 0017    63
```

Using DNS

- **Map names to addresses**
- **Descriptive names**
  - `ip domain-name`
  - `ip name-server`
- **Sample trace through network:**

```
4:Received echo from sj-wall-2.cisco.com [198.92.1.138] in 440 msec
5:Received echo from barrnet-gw.cisco.com [192.31.7.37] in 335 msec
6:Received echo from paloalto-cr1.bbnplanet.net [131.119.26.9] in 335 msec
7:Received echo from paloalto-br2.bbnplanet.net [131.119.0.194] in 327 msec
8:Received echo from core6-hssi6-0.SanFrancisco.mci.net [206.157.77.21] in 468 msec
9:Received echo from bordercore1-loopback.Washington.mci.net[166.48.36.1] in 454 msec
10:Received 48 bytes from www.getit.org [199.233.200.55] in 466 msec
```

Cisco ISP Essentials

Cisco.com

- IOS Software and Router Management
- General Features
- Routing Configuration Guidelines
- Securing the Router
- Securing the Network

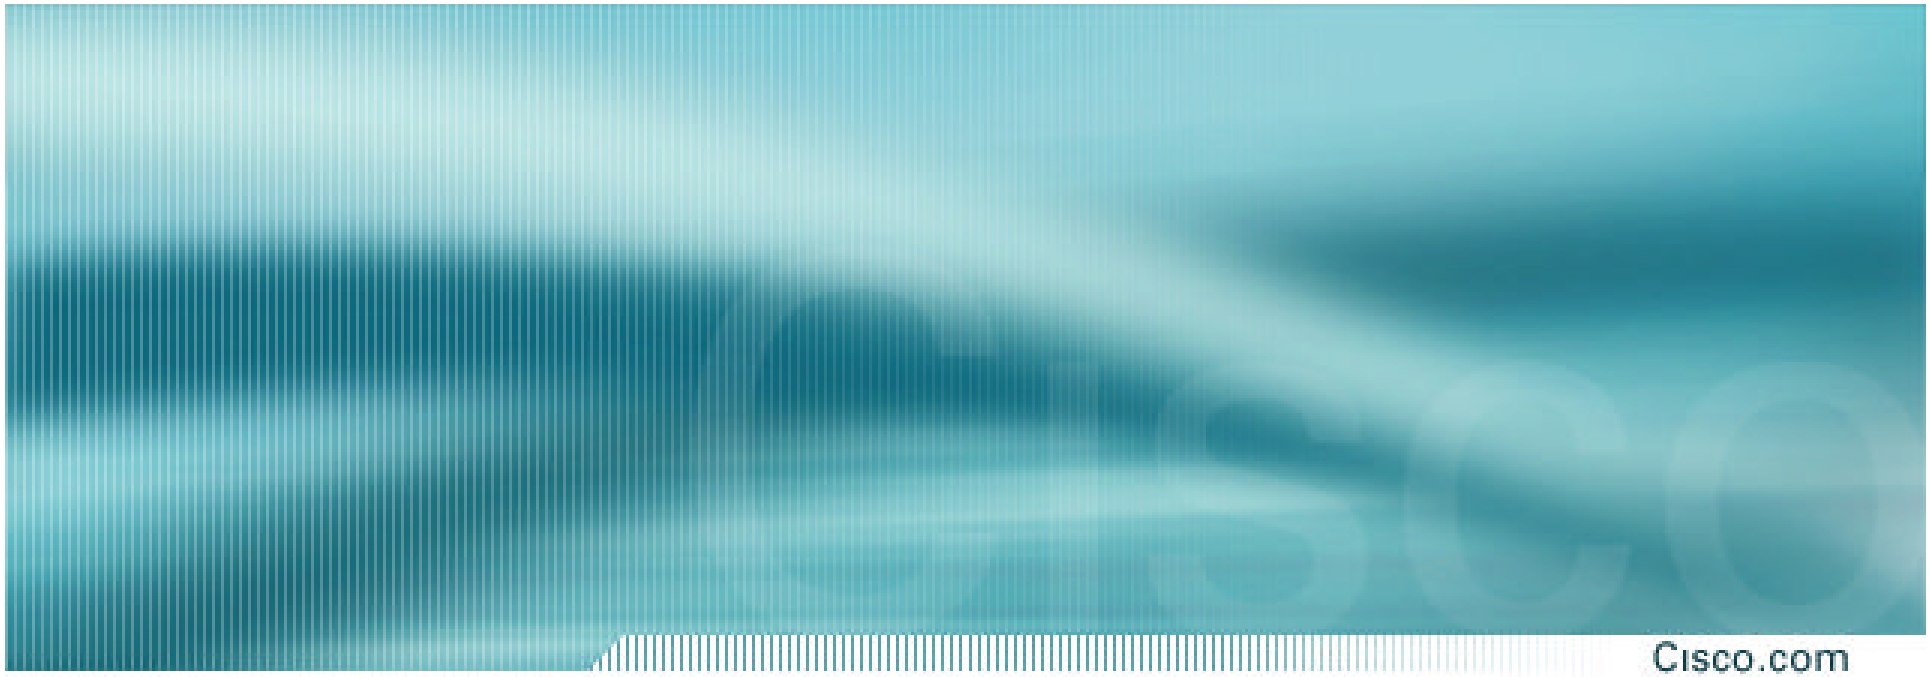

Routing

Routing Tables Feed the Forwarding Table

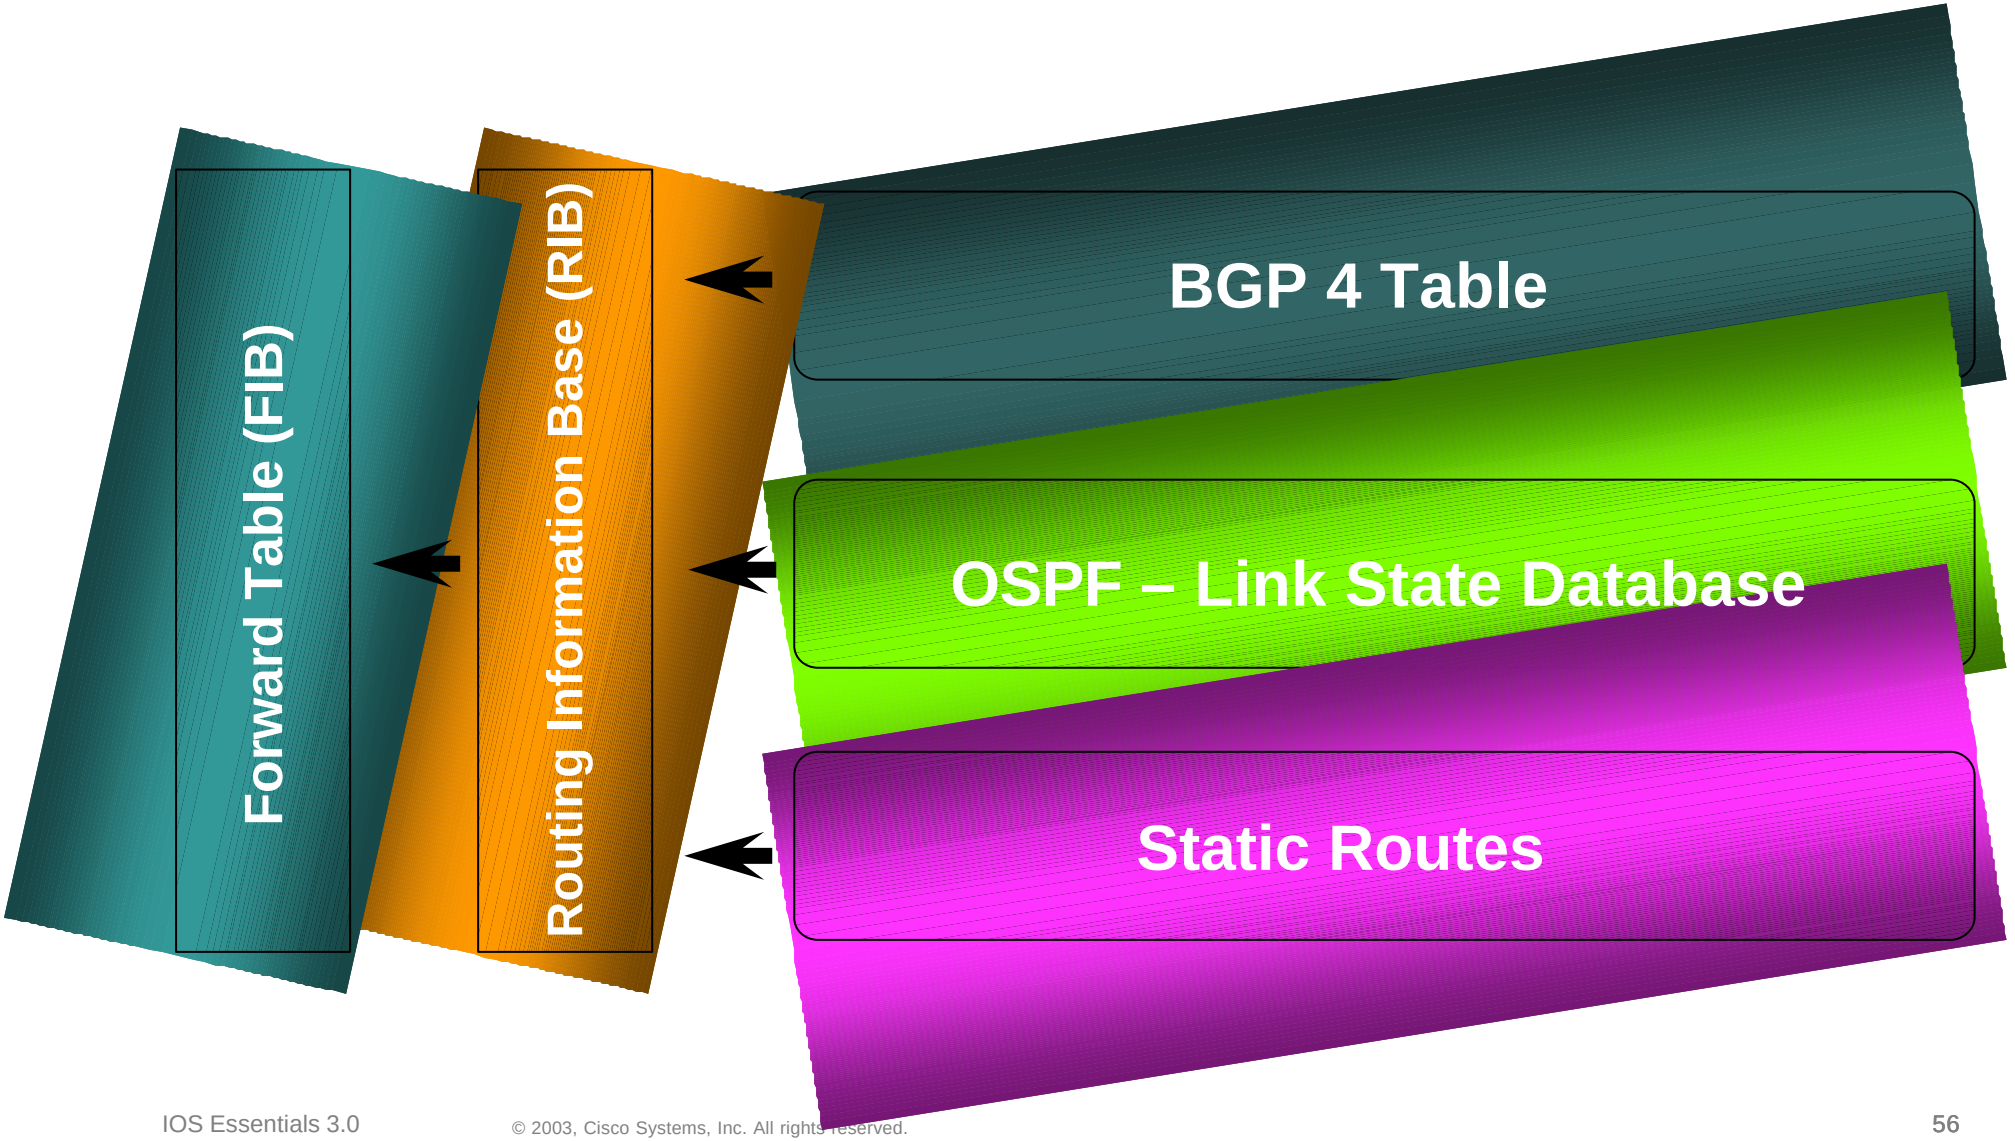

HSRP

- **Hot Standby Routing Protocol**

  virtual default gateway for dumb system LAN

  transparent cut-over in case of failure

```
Router1:
  interface ethernet 0/0
  description Service LAN
  ip address 169.223.10.1 255.255.255.0
  standby 10 ip 169.223.10.254
Router2:
  interface ethernet 0/0
  description Service LAN
  ip address 169.223.10.2 255.255.255.0
  standby 10 priority 150
  standby 10 preempt
  standby 10 ip 169.223.10.254
```

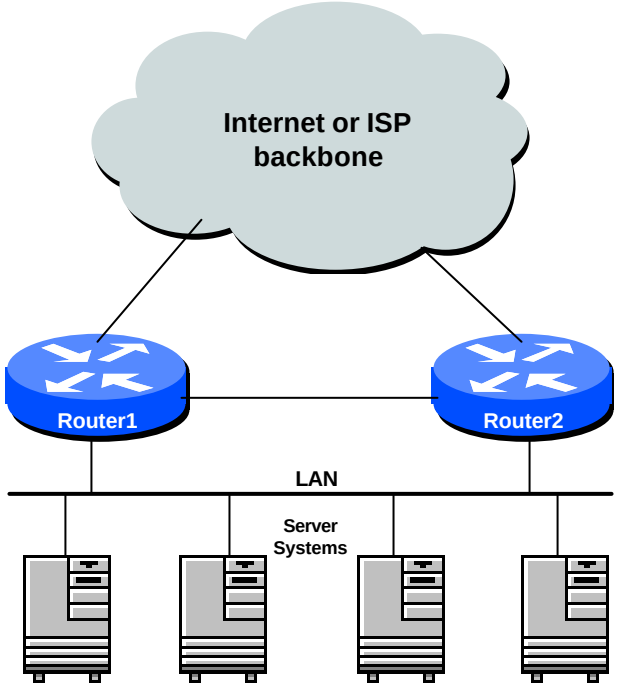

CIDR Features

- The Internet is a classless world. All routers connect to the Internet must be CIDR compliant, else there will be problems with the network connection to the Internet.
- All Cisco routers should have the following commands configured for CIDR:

```
ip subnet-zero
ip classless
```

- These are default from IOS 12.0 onwards

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
```

- Enabled by default from 11.2(5)P and later releases, available option in 11.1CA/CC.

Source Routing

- IP has 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
```

Path MTU Discovery

- **Path MTU discovery should be enabled**
  - **Allows communications from router to other devices to find optimum MTU for the path used**
  - **By default, MTU is fixed at 576 bytes – very inefficient for BGP, especially with large number of peers and prefixes**

```
ip tcp path-mtu-discovery
```

OSPF – configuration hot tips

- There are key OSPF features important for ISPs:
  - Configure one loopback per router – OSPF router id
  - Adding networks
  - `passive-interface default`
  - `log-adjacency-changes`
  - Cost and reference bandwidth
  - New `clear` commands

OSPF – Router ID

- If the loopback interface exists and has an IP address, that address is used as the router ID in routing protocols – stability!
- If the loopback interface does not exist, or has no IP address, the router ID is the highest IP address configured – danger!
- New sub command to manually set the OSPF Router ID:

```
router ospf 100
 router-id <ip address>
```

OSPF – Adding Networks

Cisco.com

- **Use specific network statements**

  Every active interface with an IP address needs a OSPF network statement

  Interface that should not be broadcasting OSPF Hello packets needs *passive-interface*

```
router ospf 100
 network 192.168.1.1 0.0.0.3 area 51
 network 192.168.1.5 0.0.0.3 area 51
 passive interface Serial 1/0
```

OSPF – Adding Networks

Cisco.com

- **Large numbers of interfaces**

  Every interface covered by wildcard mask used in OSPF network statement

  Interfaces that should not be broadcasting OSPF Hello packets need individual *passive-interface* statements or *passive-interface default*

```
router ospf 100
 network 192.168.1.0 0.0.0.255 area 51
 passive-interface default
 no passive interface POS 4/0
```

OSPF – Logging Neighbour Changes

Cisco.com

- **The router will generate a log message whenever an OSPF neighbour changes state**
- **Syntax:**

```
router ospf 100
 [no] log-adjacency-changes
```

- **Example of a typical log message:**

```
%OSPF-5-ADJCHG: Process 1, Nbr
223.127.255.223 on Ethernet0 from LOADING to
FULL, Loading Done
```

OSPF – Cost & Reference Bandwidth

- **Bandwidth used in Metric calculation**

  Cost = 10^8/BW

  Not useful for BW > 100 Mbps

- **Syntax:**

  `ospf auto-cost reference-bandwidth <reference-bw>`

- **Default reference bandwidth still 100 Mbps for backward compatibility**

- **Most ISPs simply choose to develop their own cost strategy and apply to each interface type**

OSPF – Cost: Example Strategy

Cisco.com

10GE/OC192	10Gbps	cost = 10
OC48	2.5Gbps	cost = 50
GigEthernet	1Gbps	cost = 100
OC12	622Mbps	cost = 200
OC3	155Mbps	cost = 500
FastEthernet	100Mbps	cost = 1000
Ethernet	10Mbps	cost = 5000
E1	2Mbps	cost = 10000

OSPF – Clear/Restart

Cisco.com

- **New OSPF `clear` commands**

  If no pid is given, all OSPF processes on the router are assumed

- **`clear ip ospf [pid] redistribution`**

  This command clears redistribution based on OSPF routing process ID

- **`clear ip ospf [pid] counters`**

  This command clears counters based on OSPF routing process ID

- **`clear ip ospf [pid] process`**

  This command will restart the specified OSPF process. It attempts to keep the old router-id, except in cases, where a new router-id was configured, or an old user configured router-id was removed. Since this command can potentially cause a network churn, a user confirmation is required before performing any action.

BGP – configuration hot tips

Cisco.com

- There are many features within BGP in Cisco IOS
- Designed to make life easier for ISPs
- Designed to make the Internet safer and more secure
- Each should be considered for applicability to the network

BGP – useful features

Cisco.com

```
no synchronization
no auto-summary
update-source loopback 0 (for iBGP)
ip bgp-community new-format
bgp neighbor shutdown
BGP Route Refresh Capability
bgp dampening
bgp deterministic-med
bgp neighbor next-hop-self
bgp neighbor remove-private-AS
bgp neighbor local-as
bgp neighbor authentication
```

```
bgp neighbor maximum-prefix
bgp neighbor maxas-limit
bgp log-neighbor-changes
no bgp fast-external-fallover
bgp peer-groups
ip prefix-lists
route-maps
policy-lists
route-map continue
peer-templates
Dynamic peer-groups
ibgp multi-path
```

BGP Synchronization

- **Archaic Default Number One**
- **By default BGP does not advertise a route before all routers in the AS have learned it via an IGP**

  i.e., if the prefix isn't in the IGP, BGP won't announce it
- **Synchronization must be disabled in every ISP network**

  ISPs use iBGP across backbone, IGP simply provides internal reachability

```
no synchronization
```

BGP Auto Summarisation

- **Archaic Default Number Two**
- **Automatically summarises subprefixes to the classful network when redistributed to BGP from another routing protocol**
- **Must be turned off for any Internet connected site using BGP.**
- **Internet is classless – class A, class B and class C are no more.**

```
no auto-summary
```

iBGP configuration

- **Use loopback interface**
  - **it never goes away**
  - **routers have multiple external paths**
  - **has multiple uses**

```
interface loopback 0
 ip address 215.17.1.34 255.255.255.255
router bgp 200
 neighbor 215.17.1.35 remote-as 200
 neighbor 215.17.1.35 update-source loopback 0
 neighbor 215.17.1.36 remote-as 200
 neighbor 215.17.1.36 update-source loopback 0
```

BGP Community Format

- **Communities are used extensively**
- **Cisco IOS supports two formats**
  - **One 32 bit integer** **e.g. 13107210**
  - **Two 16 bit integers** **e.g. 200:10**
- **RFC1998 recommends 16:16 format**
  - **Format AS:xxxx**

```
ip bgp-community new-format
```

Route Refresh Capability

- Facilitates non-disruptive policy changes
- No configuration is needed
- No additional memory is used
- Requires peering routers to support “route refresh capability” – RFC2918
- clear ip bgp x.x.x.x in tells peer to resend full BGP announcement
- clear ip bgp x.x.x.x out resends full BGP announcement to peer

Dynamic Reconfiguration

- **Use Route Refresh capability if supported**

  find out from “show ip bgp neighbor”

  Non-disruptive, “Good For the Internet”

- **Otherwise use Soft Reconfiguration IOS feature**

  `neighbor x.x.x.x soft-reconfiguration in`

- **Only hard-reset a BGP peering as a last resort**

Consider the impact to be equivalent to a router reboot

Soft Reconfiguration

- Router normally stores prefixes which have been received from peer after policy application
  - Enabling soft-reconfiguration means router also stores prefixes/attributes received prior to any policy application
- New policies can be activated without tearing down and restarting the peering session
- Configured on a per-neighbour basis
- Uses more memory to keep prefixes whose attributes have been changed or have not been accepted
- Also advantageous when operator requires to know which prefixes have been sent to a router prior to the application of any inbound policy

Managing Policy Changes

Cisco.com

- **Ability to clear the BGP sessions of groups of neighbours configured according to several criteria**
- **`clear ip bgp <addr> [soft] [in|out]`**

  <addr> may be any of the following

x.x.x.x	**IP address of a peer**
*****	**all peers**
ASN	**all peers in an AS**
external	**all external peers**
peer-group <name>	**all peers in a peer-group**

BGP Neighbour Shutdown

- **Shutdown BGP peering**
  - **previously required operator to delete configuration**
  - **now can simply “shutdown” the peering**
- **Configuration example:**

```
router bgp 200
 neighbor 215.7.1.1 remote-as 210
 neighbor 215.7.1.1 shutdown
```

- **Can be reactivated with**

```
no neighbor 215.7.1.1 shutdown
```

BGP Damping

- **Route flap damping to minimise instability in local network and Internet**

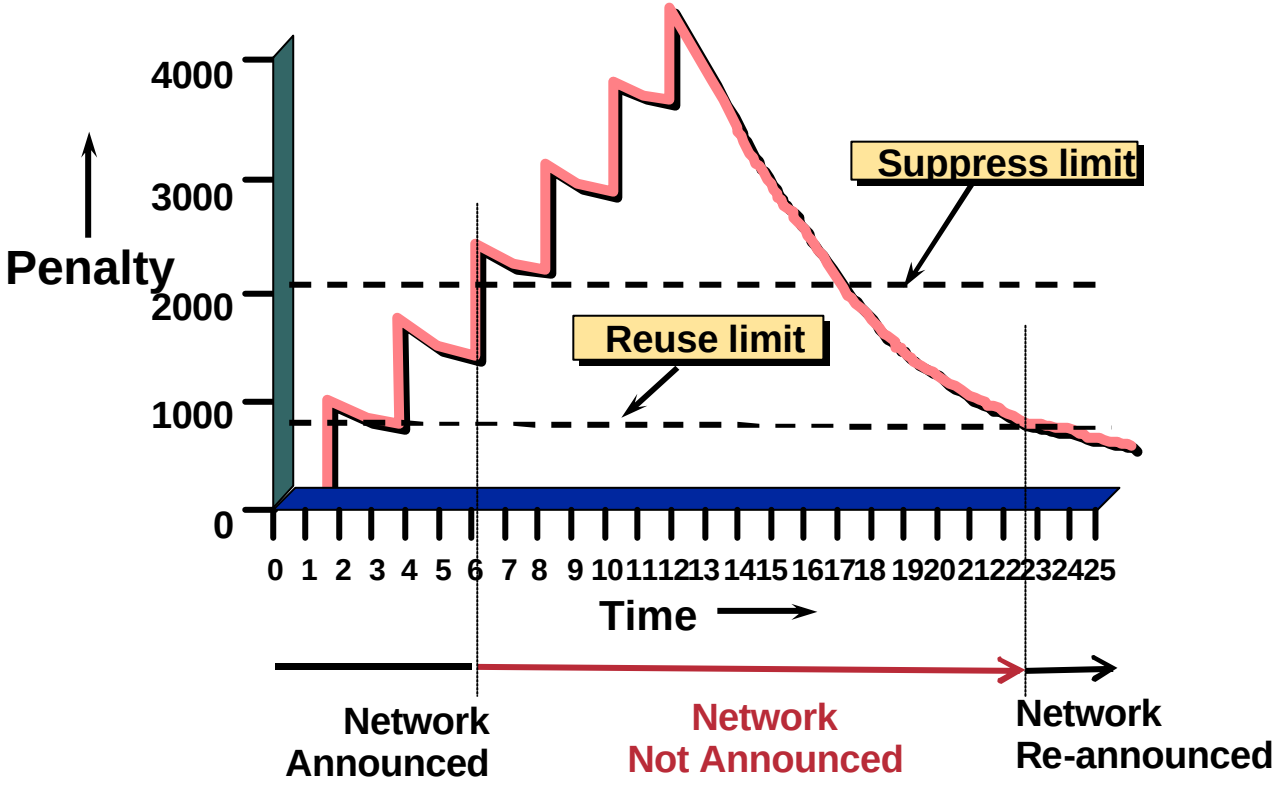

BGP Damping

Cisco.com

- **Recommended values and sample configurations for ISPs at:**

  http://www.ripe.net/docs/ripe-229.html

- **Example techniques:**

  Internet Routing Architectures 2nd Edition – Sam Halabi & Danny McPherson

  ```
  bgp dampening
  ```

Deterministic MED

- RFC1771 says that MED is not always compared
- As a result, the ordering of the paths can effect the decision process
- By default in Cisco IOS, the prefixes are compared in order of arrival (most recent to oldest)

  Use bgp deterministic-med to order paths consistently

  The bestpath is recalculated as soon as the command is entered

  Enable in all the routers in the AS

Deterministic MED—Operation

Cisco.com

- The paths are ordered by Neighbour AS
- The bestpath for each Neighbour AS group is selected
- The overall bestpath results from comparing the winners from each group
- The bestpath will be consistent because paths will be placed in a deterministic order

Next-hop-self
iBGP versus IGP

- **Make sure loopback is configured on router**
  - **iBGP between loopbacks, NOT real interfaces**
- **Make sure IGP carries loopback /32 address**
- **Make sure IGP carries DMZ nets**
  - **Use ip-unnumbered where possible**
  - **Or use next-hop-self on iBGP neighbours**
  - **neighbor x.x.x.x next-hop-self**

Next-hop-self
“Scaling IGP”

- **Used by many ISPs on edge routers**
  - **Preferable to carrying DMZ /30 addresses in the IGP**
  - **Reduces size of IGP to just core infrastructure**
  - **Alternative to using ip unnumbered**
  - **Helps scale network**
  - **BGP speaker announces external network using local address (loopback) as next-hop**

Default Administrative Distances

Route Source	Default Distance
Connected Interface	0
Static Route	1
Enhanced IGRP Summary Route	5
External BGP	20
Internal Enhanced IGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
EGP	140
External Enhanced IGRP	170
Internal BGP	200
Unknown	255

BGP Distance

- **Set BGP distance to be longer than any other routing protocol**

  OSPF distance = 110

  eBGP default = 20, iBGP default = 200

  By default prefixes learned by eBGP which have identical match in iBGP or OSPF will override the iBGP or OSPF entries

  - **Disaster for internal network**

- **Change to 200 for both eBGP and iBGP**

  ```
  distance bgp 200 200 200
  ```

  eBGP can never override internal routing protocols

Private-AS Removal

- **Private ASes range from 64512 to 65534**

  Used for internal policy – must not appear on Internet

- **neighbor x.x.x.x remove-private-AS**
- **Rules:**

  available for eBGP neighbors only

  if the update has AS_PATH made up of private-AS numbers, the private-AS will be dropped

  if the AS_PATH includes private and public AS numbers, private AS number will not be removed…it is a configuration error!

  if AS_PATH contains the AS number of the eBGP neighbor, the private-AS numbers will not be removed

  if used with confederations, it will work as long as the private AS numbers are after the confederation portion of the AS_PATH

local-AS

- Allows masquerading as a different AS
  - Especially useful during mergers and acquisitions of ISP networks
- Migrating internal network can be done during ISP's maintenance periods
- During this work, the eBGP sessions need to be migrated to the new AS
  - But peers or customers or upstreams may not be available during ISP maintenance period
  - local-AS comes to the rescue
- Local-AS configured on specific eBGP peerings so that router in new AS appears as though it is still in its original AS

local-AS – Example

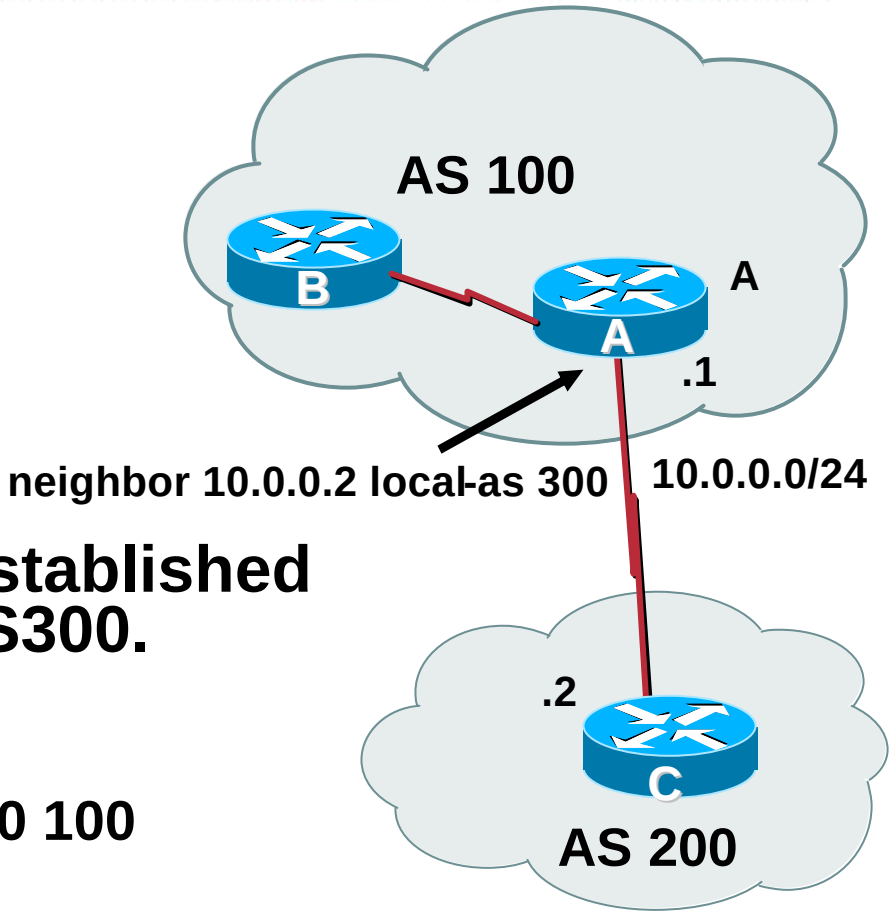

- **Router A is in AS100**
- **The peering with AS200 is established as if router A belonged to AS300.**
- **On Router C**

  routes originated in AS100 = 300 100
- **On Router A**

  routes received from AS200 = 300 200

local-AS – Example

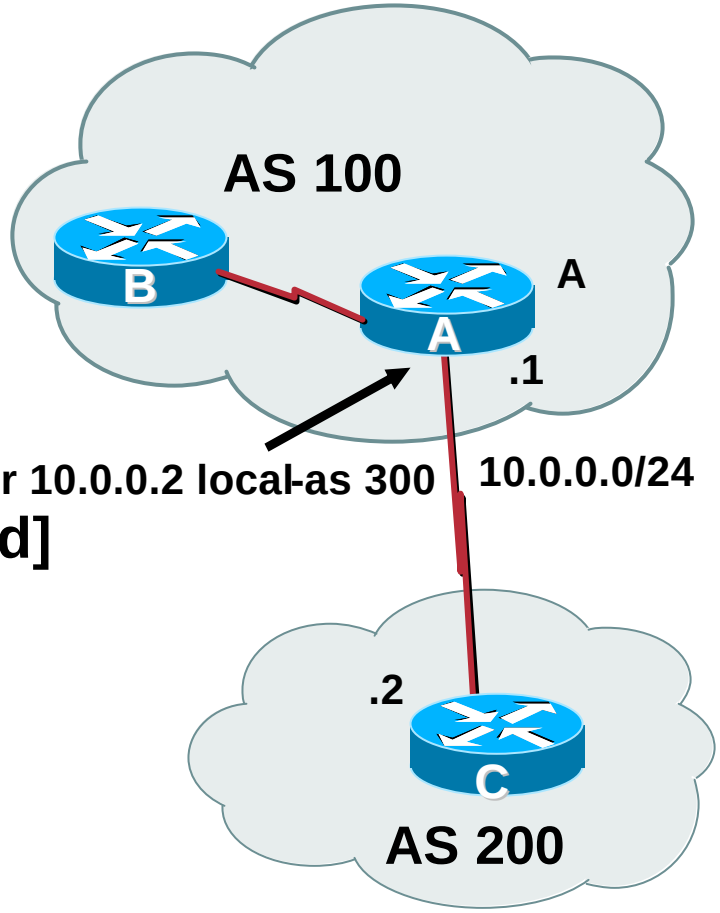

- **Router A sees the old AS300 in the path**

  If this is not desired, the [no-prepend] option can be used

```
router bgp 100
 neigh 10.0.0.2 local-as 300 no-prepend
```

- **routes received now appear as though they come directly from AS200 and not through AS300**

BGP Neighbour Authentication

- **MD5 authentication between two peers**
  - **password must be known to both peers**
- **peer-group can be used to apply to multiple peerings**

```
neighbor 169.222.10.1 password v61ne0qkel33&
```

BGP Maximum Prefix Tracking

- Allow configuration of the maximum number of prefixes a BGP router will receive from a peer
- Three level control

  Warning threshold: log warning message

  ```
  Mar 21 21:58:47.798 AEST: %BGP-4-MAXPFX: No. of
  unicast prefix received from 1.2.3.4 reaches
  122858, max 150000
  ```

  Maximum: tear down the BGP peering, manual intervention required to restart

  ```
  Mar 21 21:58:47.798 AEST: %BGP-3-MAXPFXEXCEED:
  No. of unicast prefix received from 1.2.3.4:
  150313 exceed limit 150000
  ```

  Restart interval: automatically restarts the BGP session after specified interval

BGP Maximum Prefix Tracking


```
neighbor <x.x.x.x> maximum-prefix <max> [<threshold>]
  [warning-only] [restart <restart-interval>]
```

- **threshold is an optional parameter between 1 to 100 percent**

  Specify the percentage of <max> that a warning message will be generated. Default is 75%.

- **warning-only is an optional keyword which allows log messages to be generated but peering session will not be torn down**
- **restart-interval specifies how long in minutes the router will wait before attempting to restart the BGP peering**

Limiting AS Path Length in BGP

- **Some BGP implementations have problems with long AS_PATHS**
  - **Memory corruption**
  - **Memory fragmentation**
- **Even using AS_PATH prepends, it is not normal to see more than 20 ASes in a typical AS_PATH in the Internet today**
  - **The Internet is around 5 ASes deep on average**
  - **Largest AS_PATH is usually 16-20 ASNs**

Limiting AS Path Length in BGP

Cisco.com

- **Some announcements have ridiculous lengths of AS-paths:**

```
*> 3FFE:1600::/24   3FFE:C00:8023:5::2   22 11537 145 12199 10318
10566 13193 1930 2200 3425 293 5609 5430 13285 6939 14277 1849 33
15589 25336 6830 8002 2042 7610 i
```

This example is an error in one IPv6 implementation

Use bgp maxas-limit to ignore this bogus announcement

```
router bgp 100
 bgp maxas-limit 15
```

Limits the AS-path length to 15 ASNs only

BGP log-neighbor-changes

- **Log neighbour up/down events, and the reason for the last neighbour peering reset**
- **Available from 11.1 CC and 12.0 releases**
- **Syntax (router subcommand):**

  [no] log-neighbor-changes

- **Typical log messages:**

  %BGP-6-ADJCHANGE: neighbor x.x.x.x Up

  %BGP-6-RESET: neighbor x.x.x.x reset
  (User reset request)

Reason for Last Peer Reset

- Router keeps reason for the last BGP peer reset for each of its peers. Useful for analysing BGP session resets
- Available as part of the **show ip bgp neighbor** command output
- Accessible through SNMP
- Has been available since 11.1CC, 11.2(12) and 11.3(2)

BGP Peering

Cisco.com

- **By default, peerings are reset immediately the line protocol to an external neighbour goes down**
  - **bad for high latency, unreliable, long distance, or congested links**
- **IOS option to disable this**
  - **recommended in RIPE-229**
  - **uses standard keepalive/hold timers (60s/180s)**

```
no bgp fast-external-fallover
```

BGP peer groups

- **Reduces CPU load and memory**
  - **update generation processed once**
  - **BGP configuration simplified**

```
router bgp 109
 neighbor internal peer-group
 neighbor internal remote-as 109
 neighbor internal update-source loopback 0
 neighbor 131.108.10.1 peer-group internal
 neighbor 131.108.20.1 peer-group internal
```

Prefix Lists

- High performing access-list
- Faster loading of large lists
- Incremental configuration
  - sequence numbers optional
  - `no ip prefix-list sequence-number`
- Available from 11.1(17)CC and 12.0
- Configured by:
  - `ip prefix-list <list-name>`

Prefix-list Command

Cisco.com

[no] ip prefix-list <*list-name*> [seq <*seq-value*>] deny | permit <*network*>/<*len*> [ge <*ge-value*>] [le <*le-value*>]

<*network*>/<*len*>: The prefix and its length

ge <*ge-value*>: "greater than or equal to"

le <*le-value*>: "less than or equal to"

Both "ge" and "le" are optional. Used to specify the range of the prefix length to be matched for prefixes that are more specific than <*network*>/<*len*>

Prefix Lists – Examples

Cisco.com

- Deny default route

```
ip prefix-list EG deny 0.0.0.0/0
```

- Permit the prefix 35.0.0.0/8

```
ip prefix-list EG permit 35.0.0.0/8
```

- In 192/8 allow up to /24

```
ip prefix-list EG permit 192.0.0.0/8 le 24
```

- In 192/8 deny /25 and above

```
ip prefix-list EG deny 192.0.0.0/8 ge 25
```

- Permit all

```
ip prefix-list EG permit 0.0.0.0/0 le 32
```

Prefix Lists in BGP

- **Prefix-list should be used instead of distribute-list**

  “distribute-list”, i.e. using access-lists for filtering prefixes, should be considered obsolete

```
router bgp 200
 neighbor 169.222.1.1 remote-as 200
 neighbor 169.222.1.1 prefix-list FILTER-IN in
 neighbor 169.222.1.1 prefix-list FILTER-OUT out
```

- **Prefix-lists and access-lists are mutually exclusive**

Prefix-list route-map command


```
route-map <name> permit|deny <seq-num>
  match ip address prefix-list <name> [<name> ...]
```

- **Used for route filtering, originating default, and redistribution in other routing protocols as well**
- **Not for packet filtering**

Prefix-List ORF

- **Outbound Route Filter Capability when using prefix-lists**

  new from 12.0(5)S release

- **If remote BGP peer supports ORF capability, local BGP router can send inbound prefix-list to remote router**
- **Remote router installs received prefix-list in addition to its own outbound filters**
- **Reduces unwanted routing updates from peers**

BGP Policy Configuration and Maintenance

- **The main vehicle for policy configuration in BGP are route-maps**
  - **Allow for the application of conditions and specific actions in case of a match**
  - **Older IOS versions have no provisions for complex (or multiple) condition/action pairs**
- **Peer-groups are used to group peers with common outgoing policy**
  - **Older IOS versions do not allow exceptions in the outgoing policy**

BGP Policy
Route-map Features

- **A route-map is like a “programme” for IOS**
- **Has “line” numbers, like programmes**
- **Each line is a separate condition/action**
- **Concept is basically:**

  if *match* then do *expression* and *exit*

  else

  if *match* then do *expression* and *exit*

  else *etc*

BGP Policy
Route-map Features

- **Multiple matches on the same line mean they are ANDed together**

```
route-map infilter permit 10
 match community 1 2 3
 set local-preference 120
!
```

Community-list 1 AND 2 AND 3 must match before condition is TRUE

- **Multiple matches on different lines mean they are ORed**

```
route-map infilter permit 10
 match community 1
 match community 2
 match community 3
 set local-preference 120
!
```

Community-list 1 OR 2 OR 3 must match before condition is TRUE

BGP Policy Configuration and Maintenance

policy-list

- **In short, it is a ‘macro’ for route-maps**

  Conditions can be grouped and then applied to a route-map

```
ip policy-list foo
 match as-path 10
 match ip address 100
!
route-map bar permit 10
 match ip policy-list foo
 set community 100:200
```

route-map continue

- Currently, once a match is found in a route-map, any applicable action is applied and the route-map exits
  - This behavior doesn't allow for multiple conditional actions

```
continue [route-map name|current route-map clause]
```

BGP Policy Configuration and Maintenance

route-map continue

- Provides the ability to jump to a specific step within the current route-map or to jump to the beginning of a different route-map
  - All the ‘match’ statements are evaluated against the original set of attributes

```
route-map local-policy-map
 set community 10:10
!
route-map foo-out permit 10
 match ip address 1
 match metric 10
 continue 30
!
route-map foo-out permit 20
 match ip address 2
 match metric 20
 set as-path prepend 10 10
!
route-map foo-out permit 30
 match community 10:1
 set local-preference 104
 continue local-policy-map
```

BGP Policy Configuration and Maintenance

- **The main benefits of peer-groups are:**

  UPDATE replication: only one UPDATE message is created per peer-group—It is then sent to each individual member

  Configuration grouping: All the members of a peer-group MUST have the same outgoing policy

- **Any deviation from the peer-group's outgoing policy causes the peer not to be able to be a part of the peer-group**

  Results in longer configuration files

BGP Policy Configuration and Maintenance

peer-templates

- Used to group common configurations
  - Uses peer-group-like syntax
  - No associated UPDATE replication assistance
- Hierarchical policy configuration mechanism
  - A peer-template may be used to provide policy configurations to an individual neighbor, a peer-group or another peer-template
  - The more specific user takes precedence if policy overlaps
    - individual neighbor > peer-group > peer-template

BGP Policy Configuration and Maintenance

Cisco.com

peer-templates Example

```
router bgp 100
 neighbor customer peer-group
 neighbor customer route-map martian-filter in
 neighbor customer route-map out-filter out
 neighbor customer send-community
 neighbor 1.1.1.1 remote-as 1
 neighbor 1.1.1.1 peer-group customer
 neighbor 2.2.2.2 remote-as 2
 neighbor 2.2.2.2 peer-group customer
 …..
 neighbor 10.10.10.10 remote-as 10
 neighbor 10.10.10.10 route-map martian-filter in
 neighbor 10.10.10.10 route-map out-filter out
 neighbor 10.10.10.10 send-community
 neighbor 10.10.10.10 default-information
```

```
router bgp 100
 neighbor customer peer-template
 neighbor customer route-map martian-filter in
 neighbor customer route-map out-filter out
 neighbor customer send-community
 neighbor 1.1.1.1 remote-as 1
 neighbor 1.1.1.1 peer-template customer
 neighbor 2.2.2.2 remote-as 2
 neighbor 2.2.2.2 peer-template customer
 …..
 neighbor 10.10.10.10 remote-as 10
 neighbor 10.10.10.10 peer-template customer
 neighbor 10.10.10.10 default-information
```

The Common Part of the Configuration Doesn’t Have to Be Duplicated

BGP Policy Configuration and Maintenance

peer-templates Example 2

```
router bgp 100
 neighbor customer peer-template
 neighbor customer route-map martian-filter in
 neighbor customer route-map out-filter out
 neighbor customer send-community
 neighbor customer2 peer-template
 neighbor customer2 peer-template customer
 neighbor customer2 route-map max-filter out
 neighbor 1.1.1.1 remote-as 1
 neighbor 1.1.1.1 peer-template customer
 neighbor 2.2.2.2 remote-as 2
 neighbor 2.2.2.2 peer-template customer2
 …..
 neighbor 10.10.10.10 remote-as 10
 neighbor 10.10.10.10 peer-template customer2
 neighbor 10.10.10.10 route-map peer10 out
 neighbor 10.10.10.10 default-information
```

peer-template *customer2* Inherits the *customer* Configuration

route-map *max-filter* Is Applied to Neighbor 2.2.2.2

route-map *peer10* Is Applied to Neighbor 10.10.10.10

BGP Policy Configuration and Maintenance

Dynamic peer-groups

- The use of *policy-lists*, *route-map continue* and *peer-templates* permit complex policy configurations, BUT…
- peer-group members MUST have the same outgoing policy
- Dynamic peer-groups eases the configuration by internally (no configuration needed) determining which peers have the same outgoing policy and then generating only one UPDATE for such peers

iBGP Multipath

- Allows BGP to install more than one internal path to a destination
  - Useful for load sharing
- The paths MUST be equivalent: all the absolute attributes MUST tie during the best path selection process
  - router-id, peer-address are not absolute attributes
- The best path (as determined by the selection process) is advertised
  - All eligible paths are installed in the RIB/FIB
  - Each path has a unique NEXT_HOP

iBGP Multipath—Example

- R1 has two paths for 10.0.0.0/8
- Both paths are flagged as "multipath"

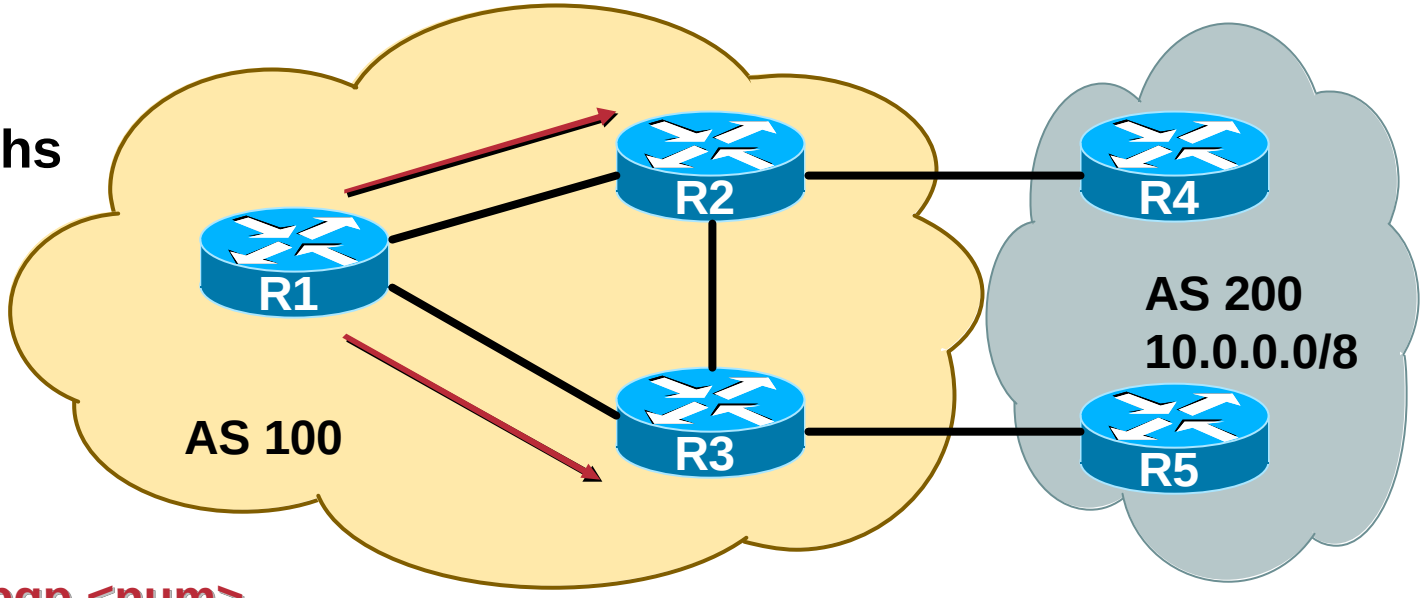

maximum-paths ibgp <num>

```
R1#sh ip bgp 10.0.0.0
  200
    20.20.20.3 from 20.20.20.3 (3.3.3.3)
    Origin IGP, metric 0, localpref 100, valid, internal, multipath
  200
    20.20.20.2 from 20.20.20.2 (2.2.2.2)
    Origin IGP, metric 0, localpref 100, valid, internal, multipath, best
```

BGP Templates

- **Good practice to configure templates for everything**
  - **Cisco defaults tend not to be optimal or even very useful for ISPs, not even in service provider images**
  - **ISPs create their own defaults by using configuration templates**
  - **Sample iBGP and eBGP templates follow**

BGP Template – iBGP peers


```
router bgp 100
neighbor internal peer-group
neighbor internal description ibgp peers
neighbor internal remote-as 100
neighbor internal update-source Loopback0
neighbor internal next-hop-self
neighbor internal send-community
neighbor internal version 4
neighbor internal password 7 03085A09
neighbor 1.0.0.1 peer-group internal
neighbor 1.0.0.2 peer-group internal
```

BGP Template – iBGP peers

Cisco.com

- Use peer-groups
- iBGP between loopbacks!
- Next-hop-self
  - Keep DMZ and point-to-point out of IGP
- Always send communities in iBGP
  - Otherwise accidents will happen
- Hardwire BGP to version 4
  - Yes, this is being paranoid!
- Use passwords on iBGP session
  - Not being paranoid, VERY necessary

BGP Template – eBGP peers

```
Router B:
router bgp 100
bgp dampening route-map RIPE229-flap
bgp deterministic-med
network 10.60.0.0 mask 255.255.0.0
neighbor external peer-group
neighbor external remote-as 200
neighbor external description ISP connection
neighbor external remove-private-AS
neighbor external version 4
neighbor external prefix-list ispout out ! “real” filter
neighbor external filter-list 1 out      ! “accident” filter
neighbor external route-map ispout out
neighbor external prefix-list ispin in
neighbor external filter-list 2 in
neighbor external route-map ispin in
neighbor external password 7 020A0559
neighbor external maximum-prefix 150000 [warning-only]
neighbor 10.200.0.1 peer-group external
!
ip route 10.60.0.0 255.255.0.0 null0 250
```

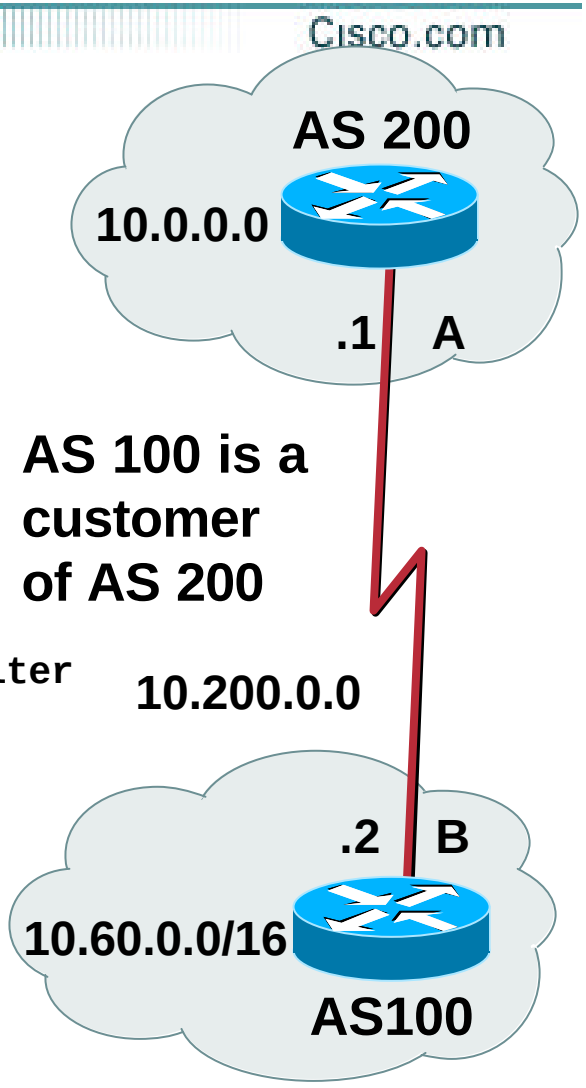

BGP Template – eBGP peers

Cisco.com

- BGP damping – use RIPE-229 parameters
- Remove private ASes from announcements
  - Common omission today
- Use extensive filters, with “backup”
  - Use as-path filters to backup prefix-lists
  - Use route-maps for policy
- Use password agreed between you and peer on eBGP session
- Use maximum-prefix tracking
  - Router will warn you if there are sudden increases in BGP table size, bringing down eBGP if desired

Cisco ISP Essentials

Cisco.com

- IOS Software and Router Management
- General Features
- Routing Configuration Guidelines
- Securing the Router
- Securing the Network

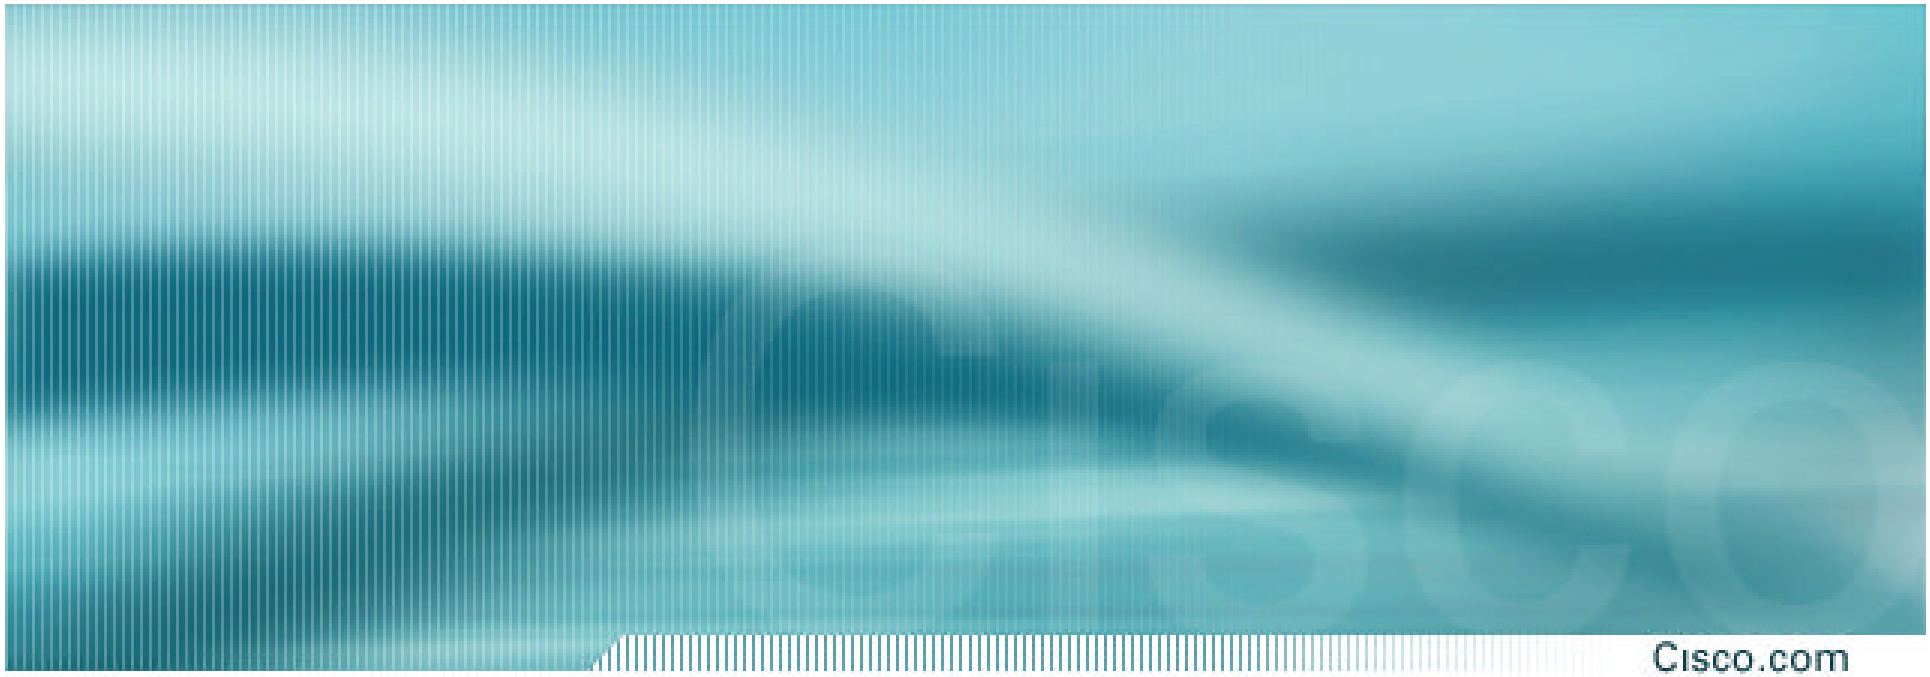

Securing the Router

ISP Security

- **ISPs need to:**
  - **Protect themselves**
  - **Help protect their customers from the Internet**
  - **Protect the Internet from their customers**

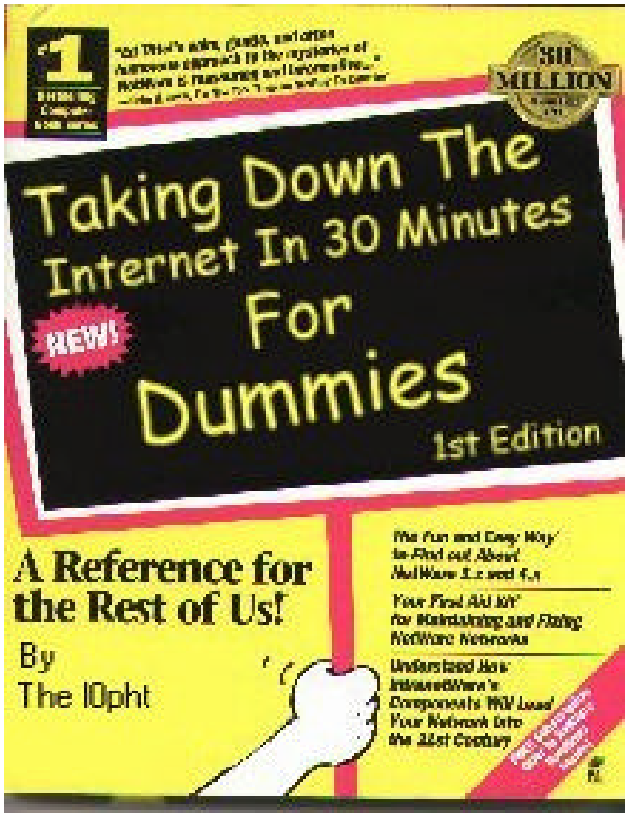

ISP Security

- **Where to start …..**
  - **Cisco Internet Security Advisories**
    - **www.cisco.com/warp/public/707/advisory.html**
  - **Cisco IOS documentation**
    - **www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_c/**
  - **RFC2196 (Site Security Handbook)**
  - **Networker's Security Sessions**

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 (before 12.0)
no ip finger (from 12.0)
no service pad
no service udp-small-servers
no service tcp-small-servers
no ip bootp server
```

- **Small tcp/udp port servers disabled by default as from 12.0S and 12.0**

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 following as a *default:***

```
no ip redirects
no ip directed-broadcast  (default from 12.0)
no ip proxy-arp
```

Cisco Discovery Protocol

- **Lets network administrators discover neighbouring Cisco equipment, model numbers and software versions**
- **Not needed on ISP network**
  - **Operators should know their equipment!**
  - `no cdp run`
- **Should not be activated on any public facing interface: IXP, customer, upstream ISP**
- **Disable per interface**
  - `no cdp enable`

Cisco Discovery Protocol Example


```
Router2#sh cdp neighbors detail
Device ID: router4
Entry address(es):
  IP address: 200.200.9.2
Platform: cisco 2611,  Capabilities: Router
Interface: Serial0/0,  Port ID (outgoing port): Serial0/1
Holdtime : 168 sec

Version :
Cisco Internetwork Operating System Software
IOS (tm) C2600 Software (C2600-I-M), Version 12.1(5)T9, RELEASE
  SOFTWARE (fc1)
TAC Support: http://www.cisco.com/tac
Copyright (c) 1986-2001 by cisco Systems, Inc.
Compiled Sat 23-Jun-01 20:13 by cmong
```

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**

  No need for a specific `enable password` – superceded by `enable secret`

```
enable secret <removed>
no enable password
service password-encryption
```

Turn on Nagle

- Telnet was designed to do one character, one packet dialog.
- John Nagle's algorithm (RFC 896) helps alleviate the small-packet problem in TCP.

```
service nagle
```

- Lessens the load on the CPU when using “show XXXX” commands

ident Feature

- **Identification (ident) support allows you to query a Transmission Control Protocol (TCP) port for identification.**
- **This feature enables an insecure protocol, described in RFC 1413, to report the identity of a client initiating a TCP connection and a host responding to the connection. No attempt is made to protect against unauthorized queries.**

```
ip ident
```

- **ISPs are very unlikely to need ident capability on any router**

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
```

Ratelimiting connections to router ICMP echo/echo-reply

- **Rate limit ICMP echo and echo-replies entering network**
  - **stops ICMP flood attacks**
  - **Example: rate-limit ICMP to 8kbps**

```
interface serial 2/0
rate-limit input access-group 190 8000 8000 8000
conform-action transmit exceed-action drop
!
access-list 190 permit icmp any any echo
access-list 190 permit icmp any any echo-reply
```

Ratelimiting connections to router TCP connections

- **Rate limit new TCP connection attempts**

  stops TCP flood attacks

  Example: rate-limit new TCP connections to 32kbps

```
interface serial 2/0
rate-limit input access-group 191 32000 8000 8000
conform-action transmit exceed-action drop
!
access-list 191 deny   tcp any any established
access-list 191 permit tcp any any
```

Compiled Access-Lists

Cisco.com

- **Traditional access-lists are processed sequentially by router CPU**
  - **shows degrading performance with increasing length of the list**
- **Compiled access-lists introduced on 7200 and higher platforms from 12.0(6)S**
  - **Uniform CPU performance, regardless of length of list**

```
access-list compiled
```

ASIC Access-lists

Cisco.com

- **GSR/12000 and 7600 series introduces ASIC based access-list processing**
  - **Has no impact on router CPU – packet filtering operates at "line-rate"**
  - **Linecard dependent though e.g.:**
    - **GSR Engine 1 has no maximum – dependent on line card memory**

      `access-list hardware salsa`

    - **GSR Engine 2 has maximum of 448 entries**

      `access-list hardware psa` **(default)**

    - **7600 TCAM has maximum of 15000 entries**

Black Hole Routing
Forwarding to Null0

- **Null0 is often used as a black hole**

  And for generating Aggregate in BGP

  ```
  ip route 213.13.0.0 255.255.192.0 null 0
  ```

- **Packets without a specific destination are dumped in the null0 interface**

  (as part of CEF – not process switched)

  Disable icmp unreachables for these packets

  ```
  interface null 0
   no icmp unreachables
  ```

Black Hole Routing
Ratelimiting ICMP unreachables

- **ICMP unreachables also rate-limited**

  Want to avoid the router CPU being swamped sending responses to dumped packets

  ```
  ip icmp rate-limit unreachable DF 2000
  ```

  Community consensus is to set ICMP unreachable response to one every 2 seconds with the DF bit set

  (IOS default is one response every 500ms)

VTY and Console port timeouts

Cisco.com

- **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`

  Disconnects unused connections

VTY Security

- Consoles should be used for last resort admin only
- Access to VTYs should be controlled, not left open
  - Use the ACL log function to spot the probes on your network

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

VTY Access and SSHv1
SSHv1 Server

- **Secure Shell v1 supported as from IOS 12.0S**
  - **Also in 12.1+ 3DES images**
- **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

Cisco.com

ssh [-l <userid>] [-c <des|3des>] [-o num-attempts <n>] [-p <port>] <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 num-attempts <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 <port> 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.

User Authentication – take 1

- **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
```

User Authentication – take 2

- **More recent versions of IOS add MD5 encryption for user passwords**

```
aaa new-model
aaa authentication login neteng local
username joe secret 5 $1$j6Ac$3KarJszBV3VMaL/2Nio3E.
username jim secret 5 $1$LPV2$Q04NwAudy0/4AHHHQHvWj0
line vty 0 4
 login neteng
 access-class 3 in
```

User Authentication

- **Use centralised authentication system**
  - **RADIUS recommended for dial access AAA**
  - **TACACS+ recommended for system security**

```
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 host 215.17.5.35
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-N	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 Worksho	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 de	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

Cisco ISP Essentials

Cisco.com

- IOS Software and Router Management
- General Features
- Routing Configuration Guidelines
- Securing the Router
- Securing the Network

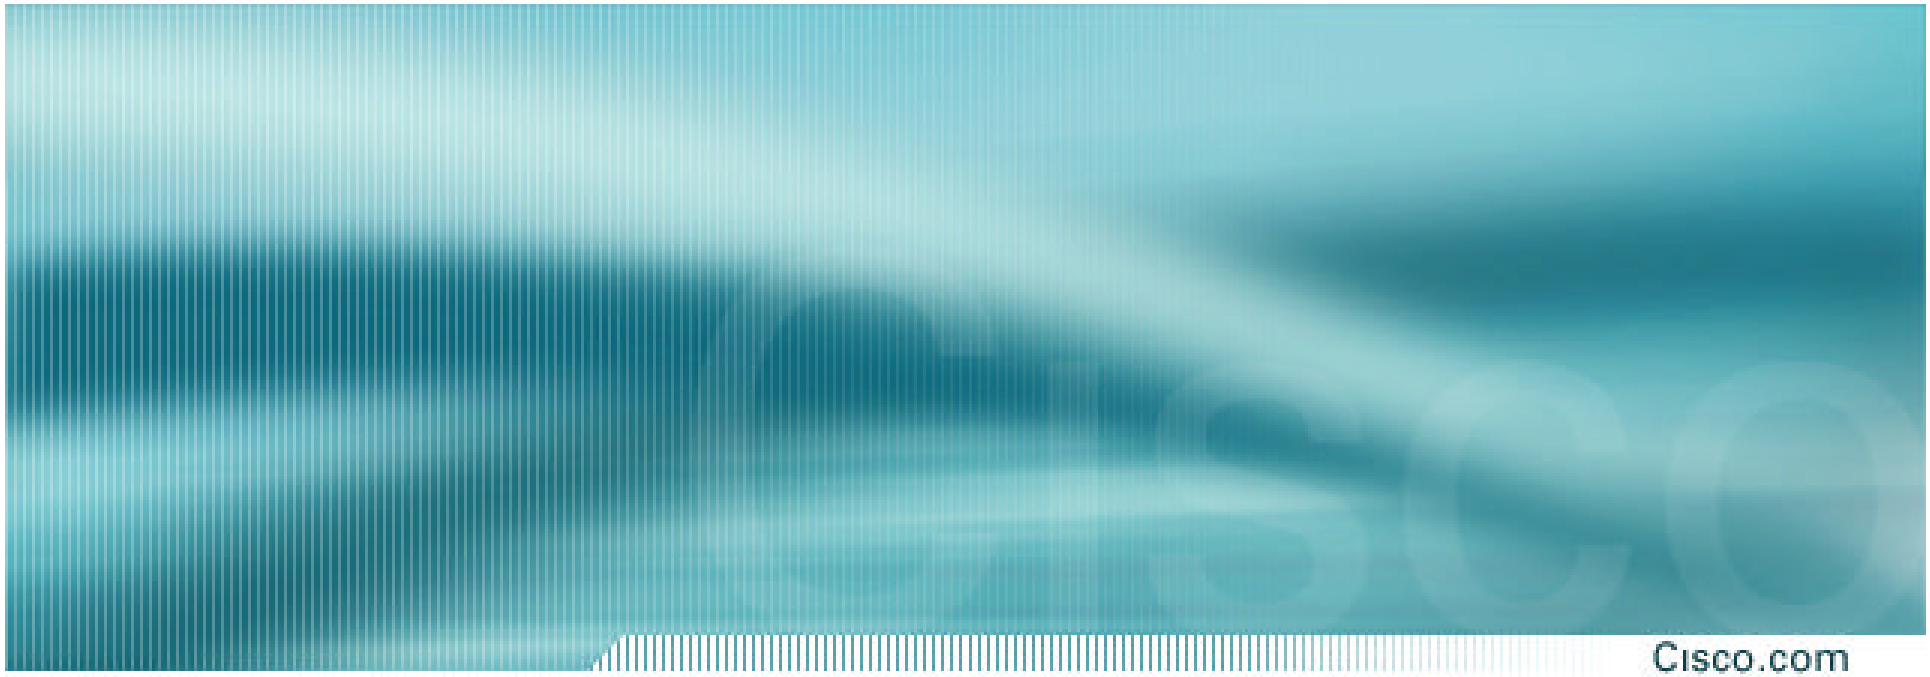

Securing the Network

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 RFC3330 for background information on special networks
- Check Rob Thomas’ list of “bogons”
  - http://www.cymru.org/Documents/bogon-list.html
- BGP should have filters applied so that these routes are not advertised to or propagated through the Internet

Ingress and Egress Route Filtering

BGP Configuration

```
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 bogons in
 neighbor 220.220.4.1 prefix-list bogons out
 neighbor 222.222.8.1 remote-as 220
 neighbor 222.222.8.1 version 4
 neighbor 222.222.8.1 prefix-list bogons in
 neighbor 222.222.8.1 prefix-list bogons out
 no auto-summary
!
```

Ingress and Egress Route Filtering

Cisco.com

Prefix List

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

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**

Outbound Packet Filtering

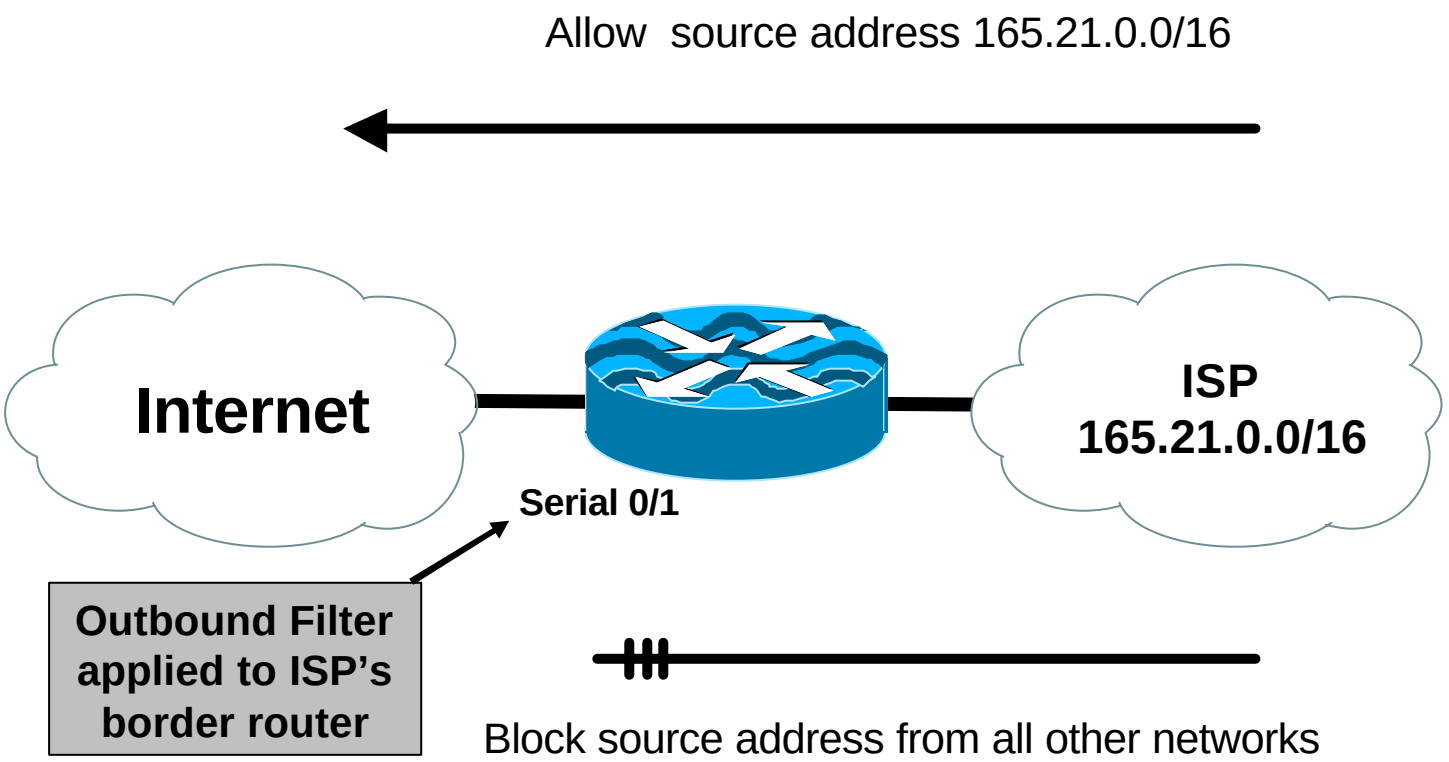

Ex. IP addresses with a source of 10.1.1.1 would be blocked

Inbound Packet Filtering

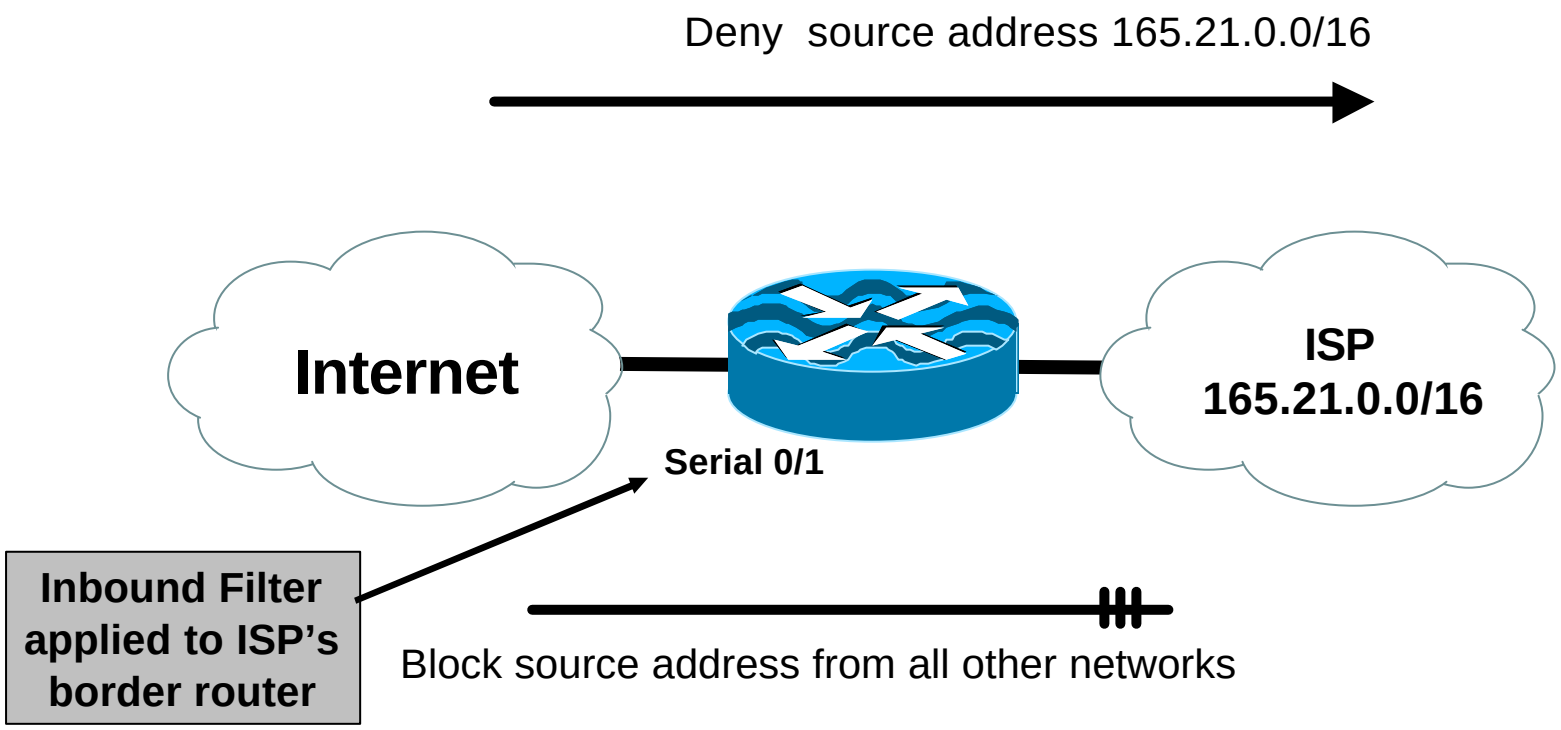

Ex. IP addresses with a source of 10.1.1.1 would be blocked

Dynamic ACLs with AAA Virtual Profiles

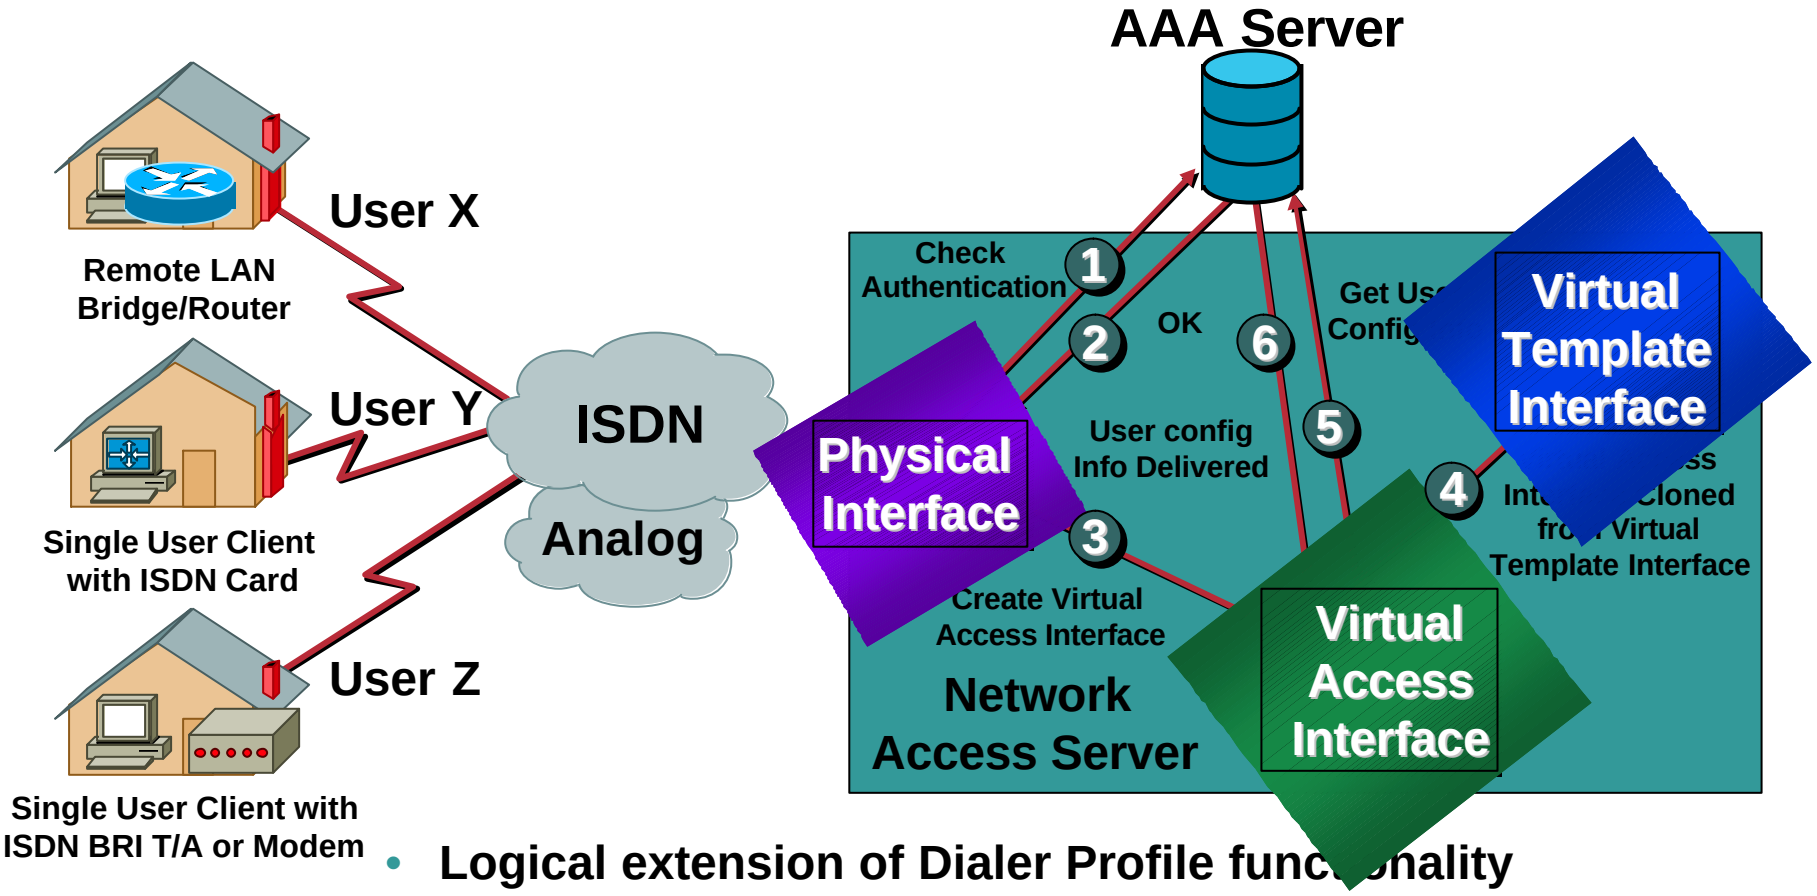

- Logical extension of Dialer Profile functionality
- ACLs stored in the Central AAA Server
- Supports both Radius and Tacacs+

Reverse Path Forward Check

- 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
- Thought/planning required in multihoming situations

Reverse Path Forward Check

- **IOS Command**

```
interface serial 1/0
 ip verify unicast reverse-path <acl>
```

- **Access-list has two uses**

  To allow prefixes which have failed the uRPF test (access-list permit statement)

  To log uRPF failures (access-list deny log statement)

CEF Unicast RPF

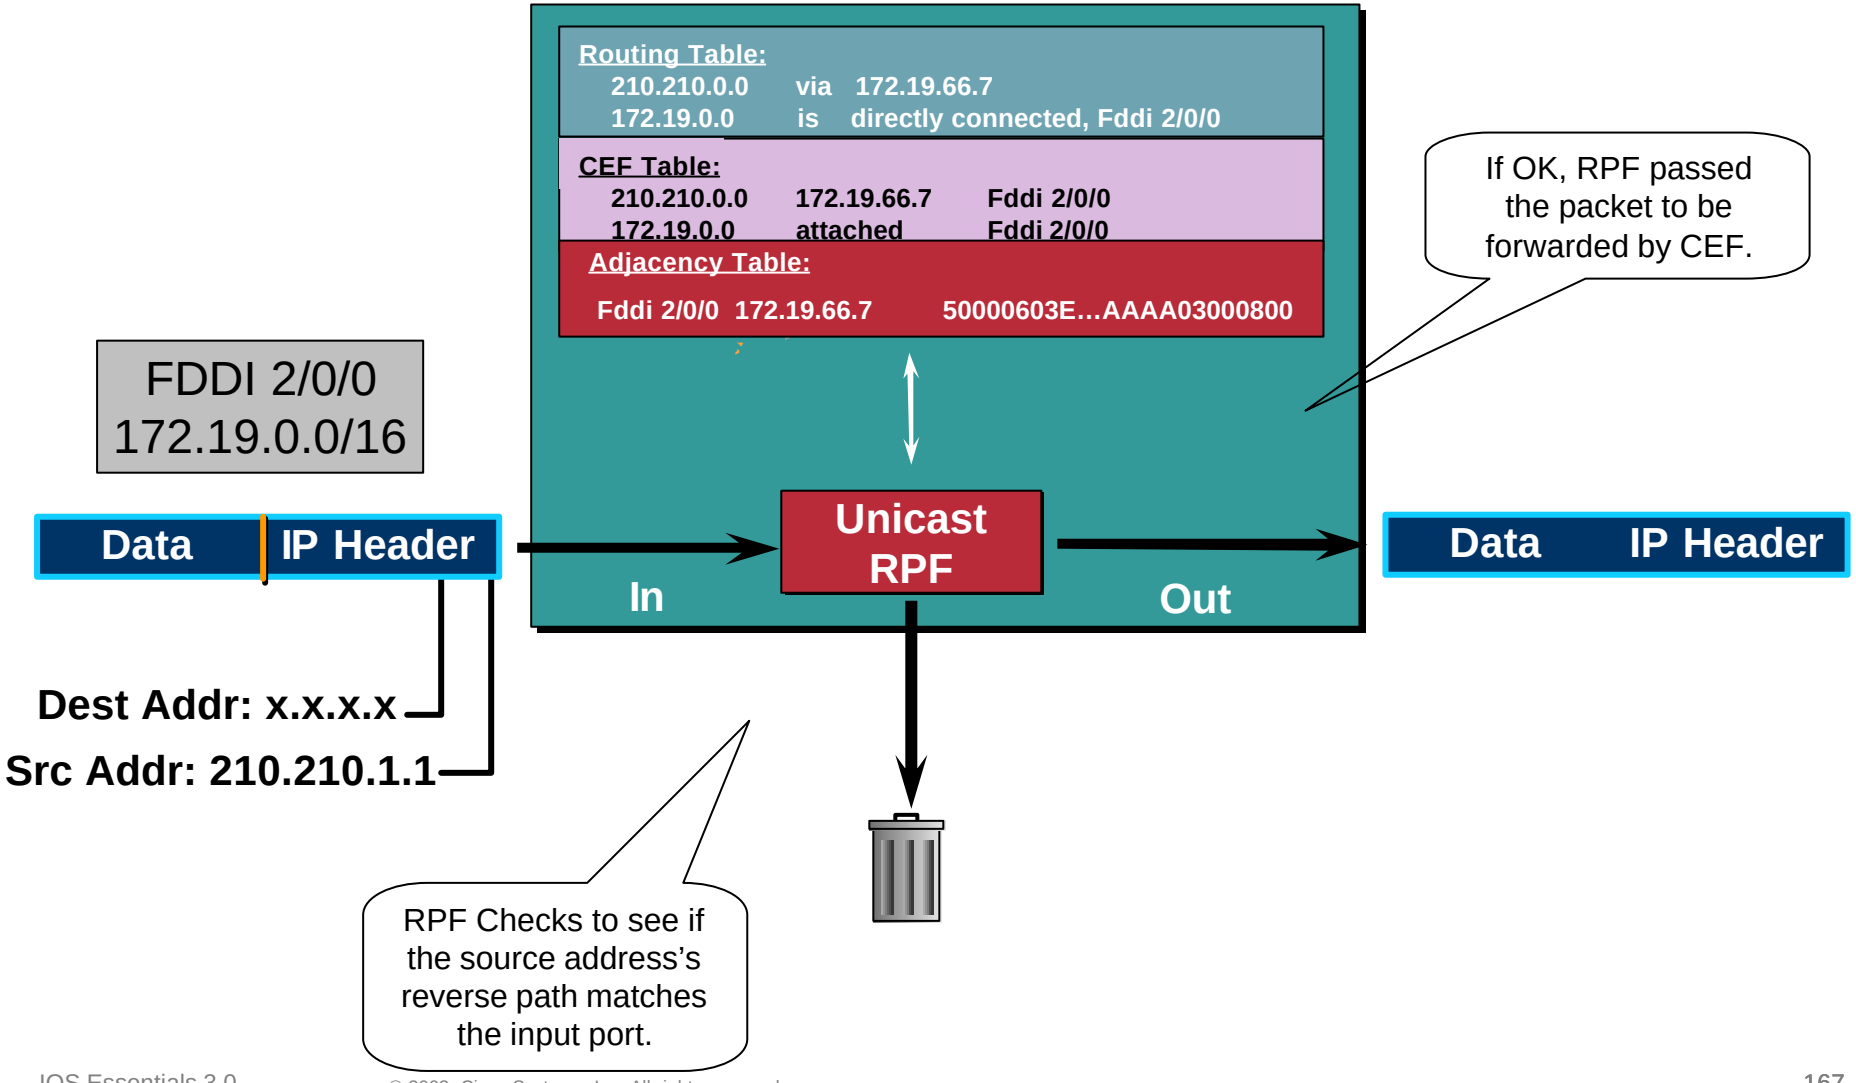

CEF Unicast RPF

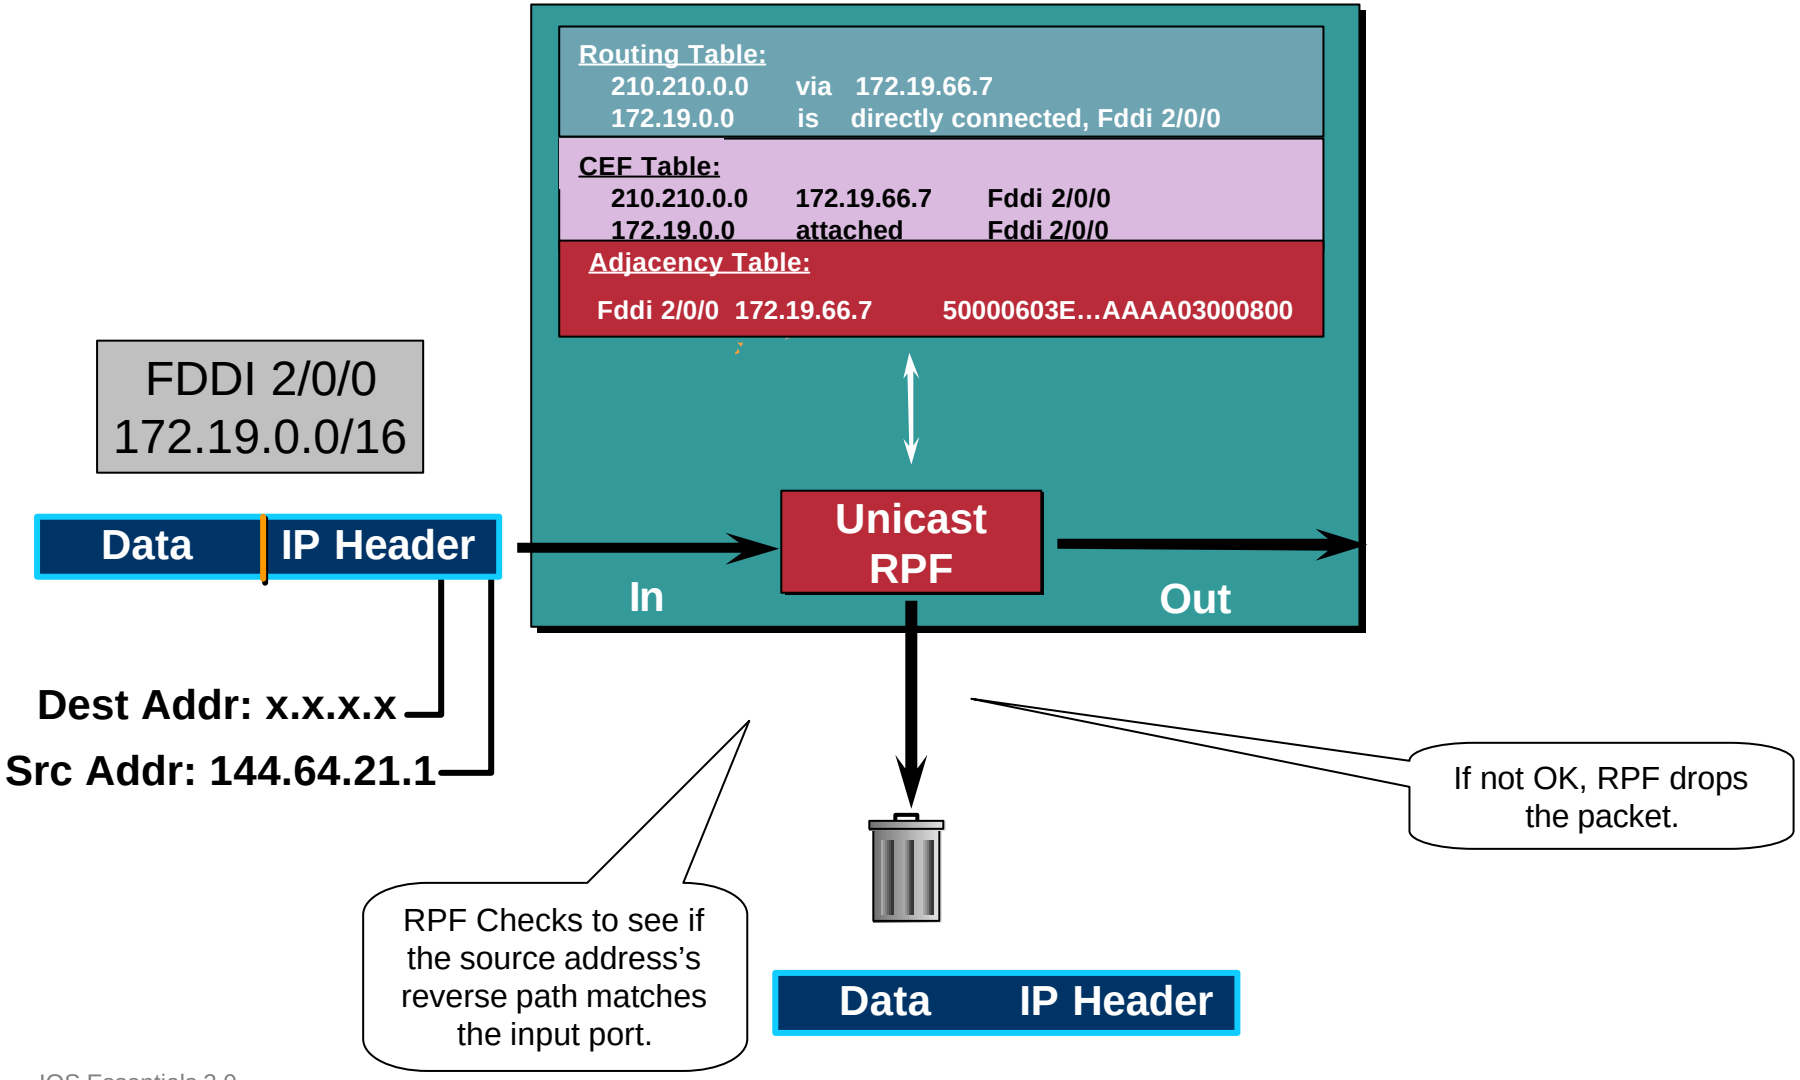

Unicast RPF – Simple Single Homed Customer Example


```
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
```

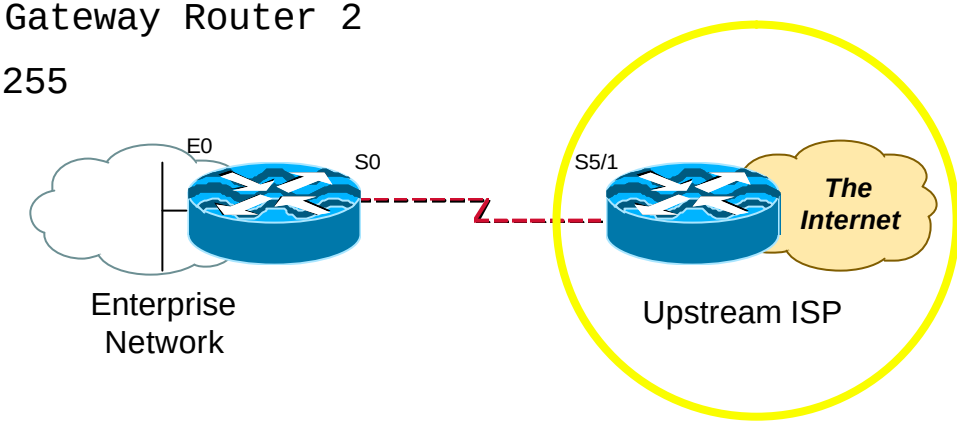

Unicast RPF Check

- **Should be mandatory command on all ISP's edge routers connecting customers to the Internet**
  - **Part of IOS Essentials ISP router template**
- **Multihomed customers require a little more thought and planning**
  - **Use BGP weight**
  - **Use uRPF enhancements (ACL and FIB comparison) in 12.0(14)S**

```
ip verify unicast reverse-path <acl>
ip verify unicast source reachable-via [any|rx]
[allow-default] [allow-self-ping] [<acl>]
```

Unicast RPF – ACL

- **ACLs can now be used with Unicast RPF:**

  ip verify unicast reverse-path 171

- **ACLs are used to:**

  Allow exceptions to the Unicast RPF check

  Identify characteristics of spoofed packets being dropped by Unicast RPF

Unicast RPF – ACL

- **Cisco 7206 with Bypass ACL**

```
interface fastethernet 1/0
  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

beta7200# show ip interface ethernet 1/1 | include RPF
  Unicast RPF ACL 197
  1 unicast RPF drop
  1 unicast RPF suppressed drop
```

Unicast RPF – ACL

- **Cisco 7206 with a classification filter:**

```
interface fastethernet 1/0
 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
```

Description of “Smurfing”

- **Smurf is a Denial of Service attack**
  - **Network-based, fills access pipes**
  - **Uses ICMP echo/reply packets with broadcast networks to multiply traffic**
  - **Requires the ability to send spoofed packets**
  - **Would hardly exist if ISPs used uRPF checks and disabled directed-broadcast on LANs**
- **Abuses “bounce-sites” to attack victims**
  - **Traffic multiplied by a factor of 50 to 200**

Description of “Smurfing”

ICMP echo (spoofed source address of victim)
Sent to IP broadcast address

ICMP echo reply

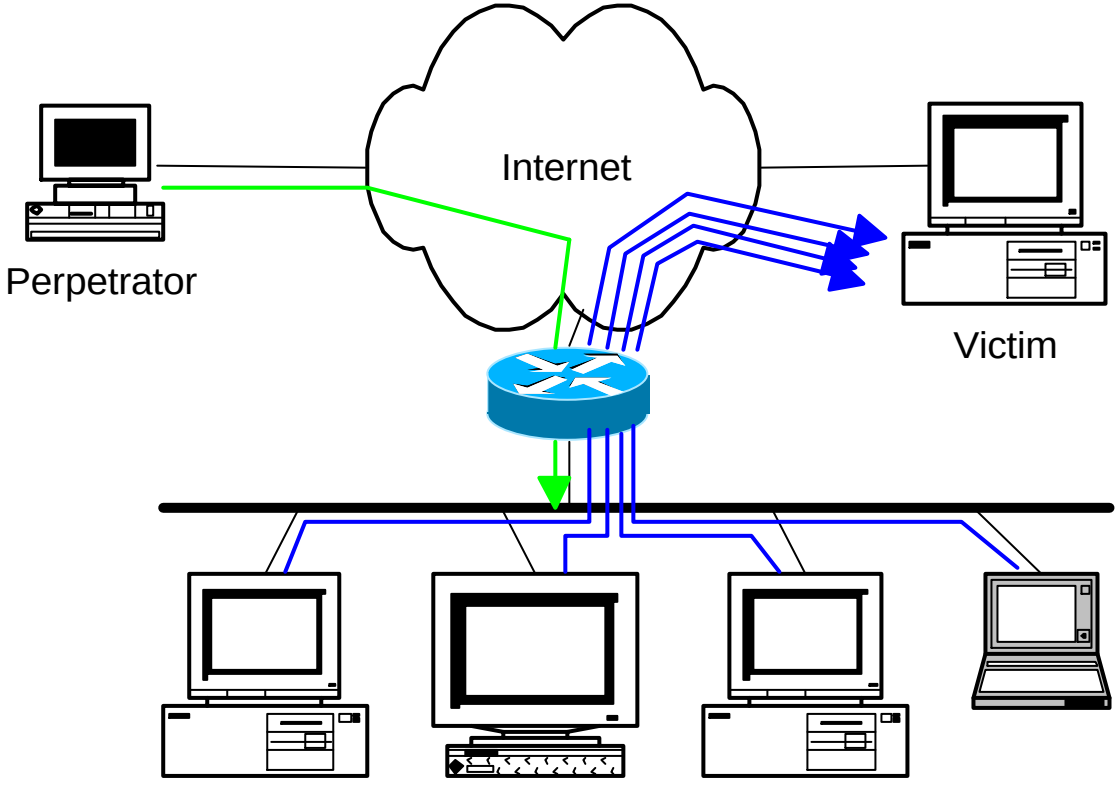

Multiplied Bandwidth – Example

- **Perpetrator has T1 bandwidth available (typically a cracked account), and uses half of it (768 Kbps) to send spoofed packets, half to bounce site 1, half to bounce site 2**
- **Bounce site 1 has a switched co-location network of 80 hosts and T3 connection to net**
- **Bounce site 2 has a switched co-location network of 100 hosts and T3 connection to net**

Multiplied Bandwidth – Consequences

Cisco.com

- (384 Kbps * 80 hosts) = 30 Mbps outbound traffic for bounce site 1
- (384 Kbps * 100 hosts) = 37.5 Mbps outbound traffic for bounce site 2
- Victim is pounded with 67.5 Mbps (!) from half a T1!

Profiles of Participants

Cisco.com

- **Typical Perpetrators**
  - Cracked superuser account on well-connected enterprise network
  - Superuser account on university residence hall network (Ethernet)
  - Typical PPP dial-up account (for smaller targets)
- **Typical Bounce Sites**
  - Large co-location subnets
  - Large switched enterprise subnets
  - Typically scanned for large numbers of responding hosts
- **Typical Victims**
  - IRC Users, Operators, and Servers
  - Providers who eliminate troublesome users' accounts

Prevention Techniques

Cisco.com

- **How to prevent your network from being the source of the attack:**
  - **Apply filters to each customer network**
    - **Ingress: Allow only those packets with source addresses within the customer's assigned netblocks**
  - **Apply filters to your upstreams**
    - **Egress: Allow only those packets with source addresses within your netblocks to protect others**
    - **Ingress: Deny those packets with source addresses within your netblocks to protect yourself**

Prevention Techniques

Cisco.com

- **How to suppress an attack if you're the victim:**
  - **Implement ACL's at network edges to block ICMP echo responses to your high-visibility hosts, such as IRC servers**
    - **Will impair troubleshooting – "ping" breaks**
    - **Will still allow your access pipes to fill**
  - **Work with upstream providers to determine the help they can provide to you**
    - **Blocking ICMP echoes for high-visibility hosts from coming through your access pipes**
    - **Tracing attacks**

DDoS versus DoS

- Same methods and tools as DoS
- Much larger scale attacks
  - Elephant hunting
- Uses hundreds or even thousands of attacking points to overwhelm targets
- Very difficult to determine difference between DDoS and network outage

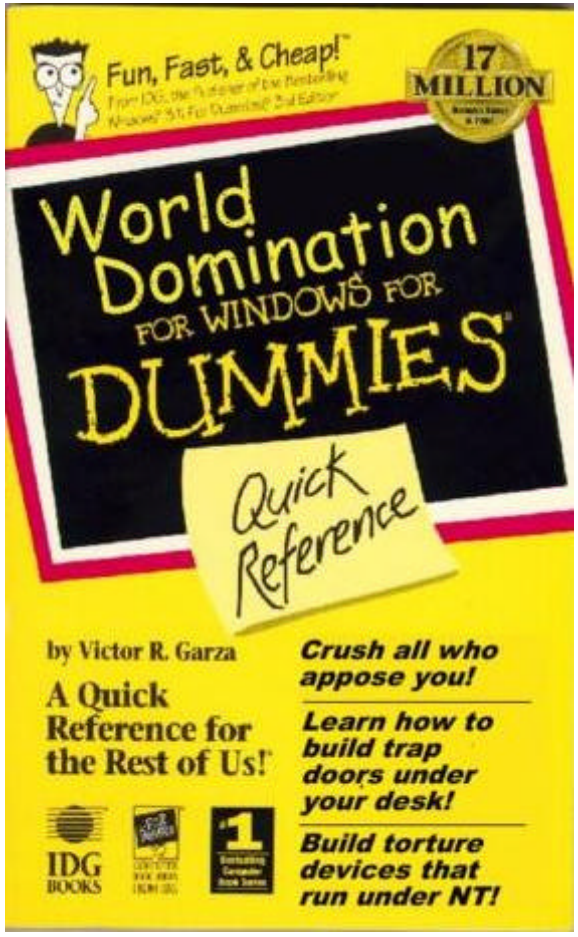

DDoS Links

Cisco.com

- http://www.denialinfo.com/
- http://www.staff.washington.edu/dittrich
- http://www.sans.org/y2k/DDoS.htm
- http://www.nanog.org/mtg-9910/robert.html
- http://cve.mitre.org/

Cisco ISP Essentials

Cisco.com

- IOS Software and Router Management
- General Features
- Routing Configuration Guidelines
- Securing the Router
- Securing the Network

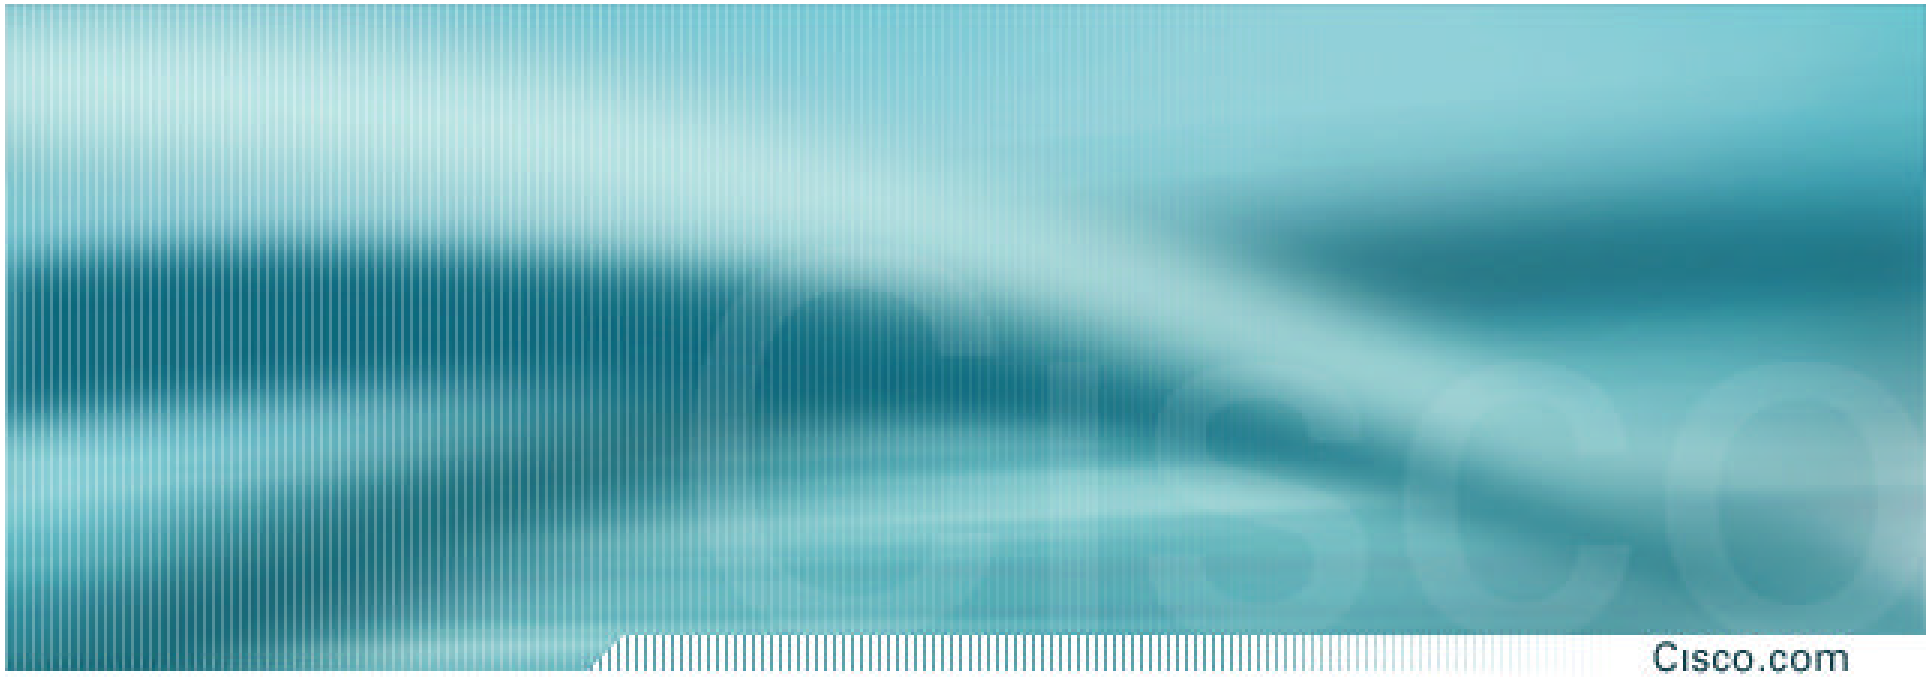

More Information?

Where to get more information

Cisco.com

- **Supporting *Cisco ISP Essentials* Book**

  `http://www.ispbook.com`

- **Check the CTO Consulting Engineering ISP Resources page:**

  `ftp://ftp-eng.cisco.com/cons/`

- **Join the cisco-nsp mailing list – set up by ISPs for ISPs**

  `send e-mail to cisco-nsp-request@puck.nether.net with subject of “subscribe”`

  `http://puck.nether.net/mailman/listinfo/cisco-nsp`

For Further Reference…

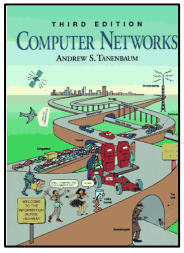

- **Computer Networks, Third Edition**
  by Andrew Tanenbaum (ISBN: 0-13349-945-6)

- **Interconnections : Bridges and Routers (second Ed)**
  by Radia Perlman (ISBN: 0-20163-448-1)

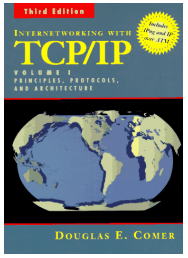

- **Internetworking with TCP / IP, Volume 1: Principles, Protocols, and Architecture**
  by Douglas Comer (ISBN: 0-13216-987-8)

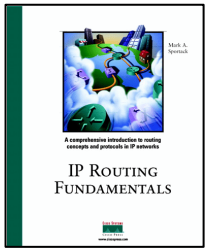

- **IP Routing Fundamentals**
  by Mark Sportack (ISBN: 1-57870-071-x)

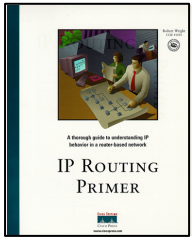

- **IP Routing Primer**
  by Robert Wright (ISBN: 1-57870-108-2)

For Further Reference…

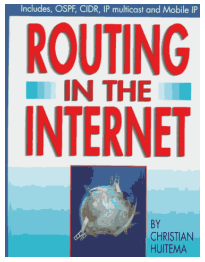

- **Routing in the Internet**
  by Christian Huitema (ISBN: 0-13132-192-7)

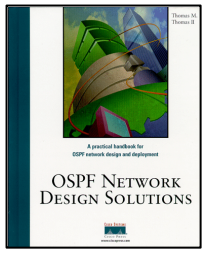

- **OSPF Network Design Solutions**
  by Thomas, Thomas M. (ISBN: 1-57870-046-9)

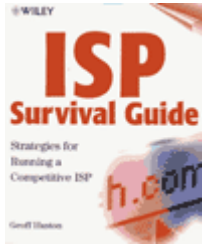

- **ISP Survival Guide : Strategies for Running a Competitive ISP**
  by Geoff Huston (ISBN:0-47131-499-4)

- **Internet Routing Architectures: 2nd Edition**
  by Sam Halabi & Danny Mcpherson

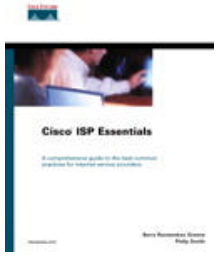

- **Cisco ISP Essentials**
  by Barry Greene & Philip Smith

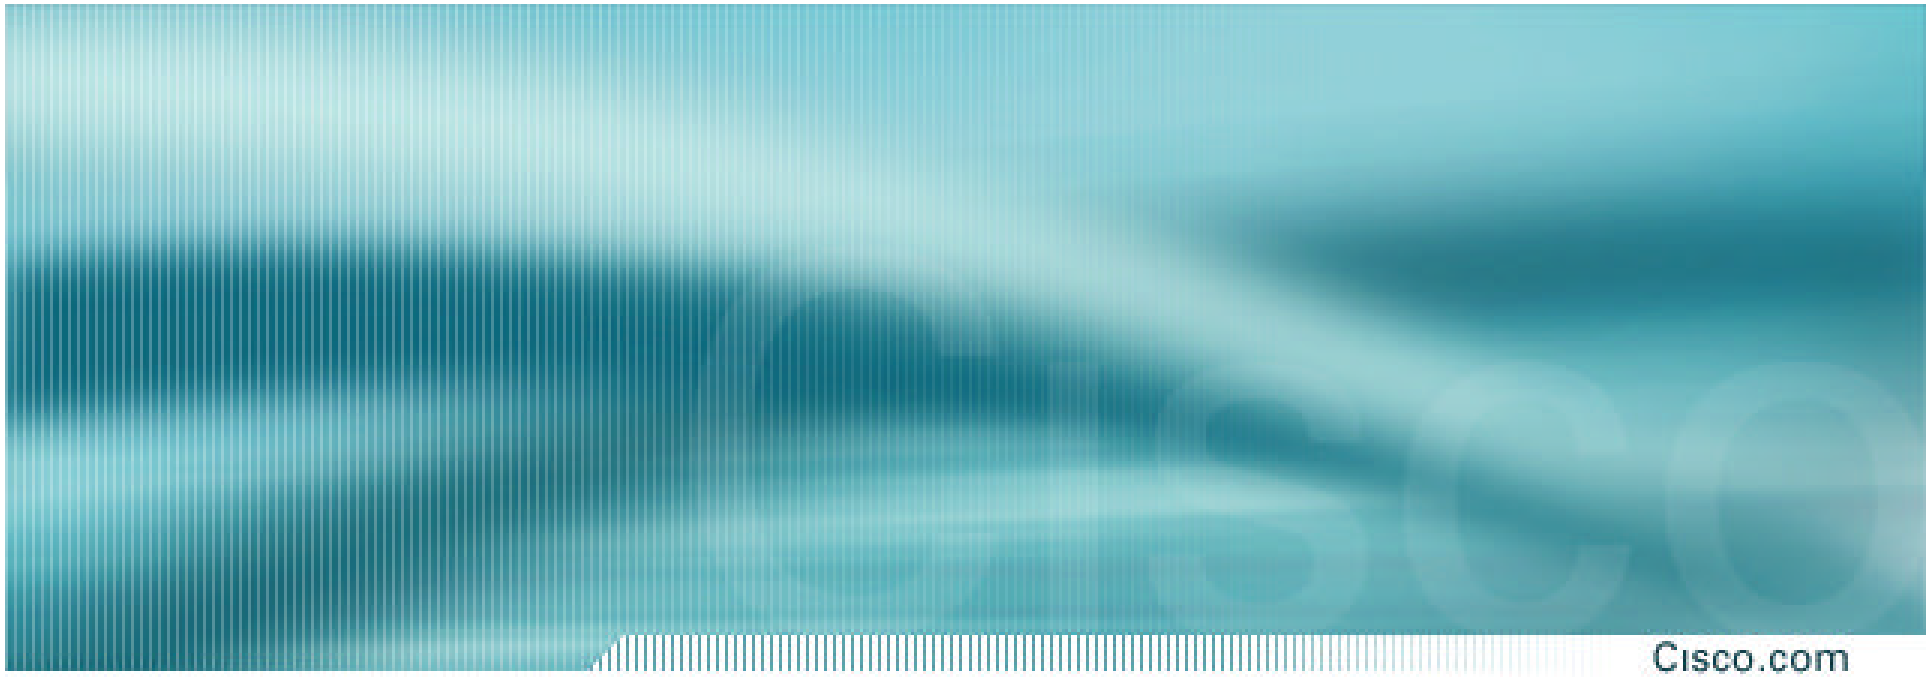

ISP Essentials

Essential IOS Features every ISP should Consider

End of Tutorial