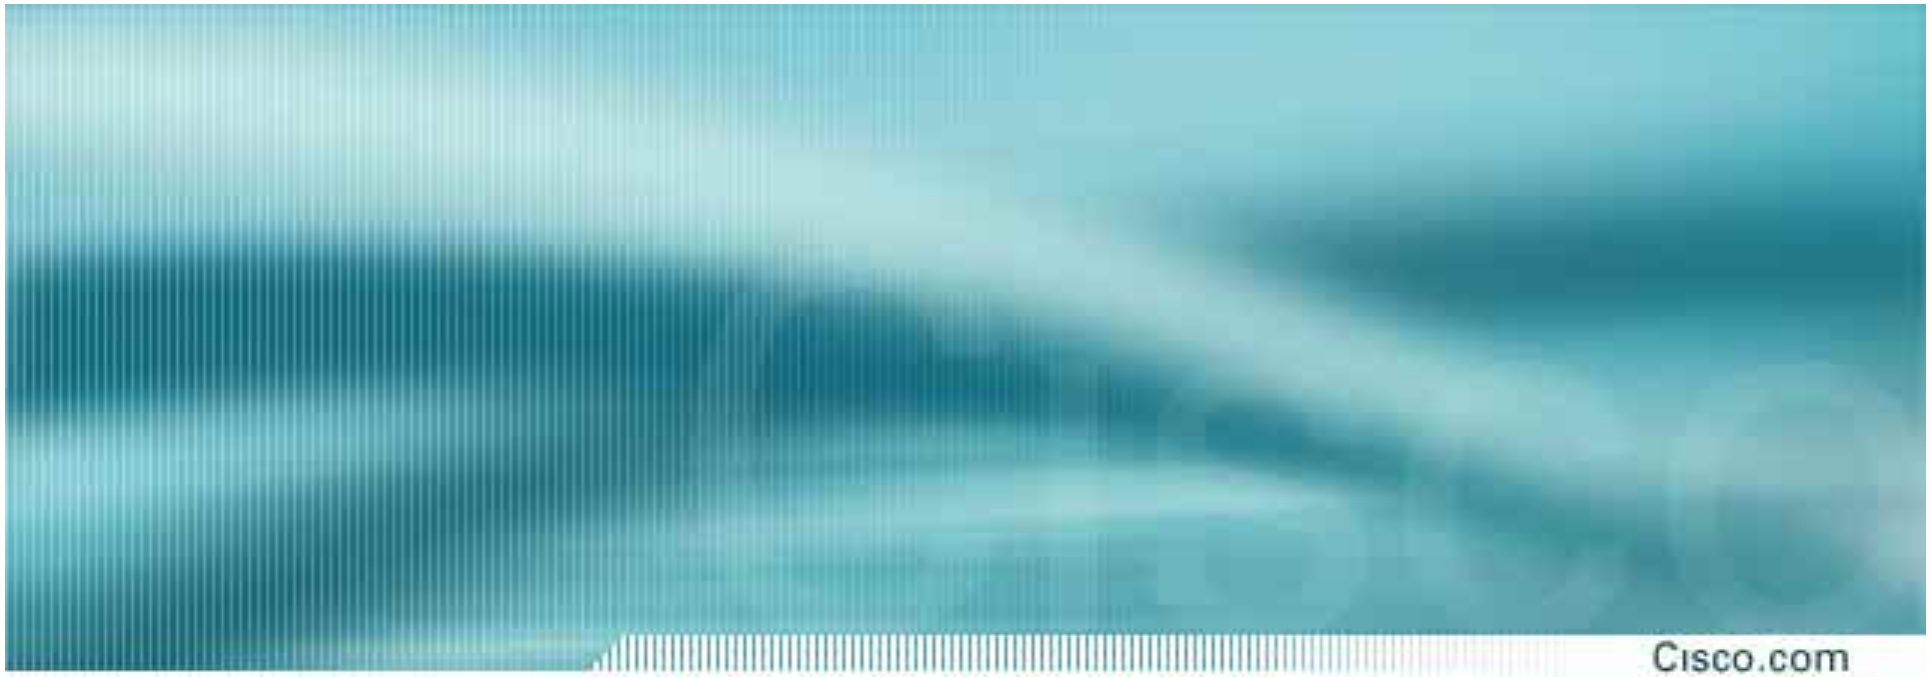

Hacking Layer 2: Fun with Ethernet Switches

Sean Convery, Cisco Systems

sean@cisco.com

Agenda

- Layer 2 Attack Landscape
- Specific Attacks and Countermeasures (Cisco and @Stake Testing)—http://www.atstake.com
  - MAC Attacks
  - VLAN "Hopping" Attacks
  - ARP Attacks
  - Spanning Tree Attacks
  - Layer 2 Port Authentication
  - Other Attacks
- Summary and Case Study

Caveats

- All attacks and mitigation techniques assume a switched Ethernet network running IP
  - If shared Ethernet access is used (WLAN, Hub, etc.) most of these attacks get much easier
  - If you aren't using Ethernet as your L2 protocol, some of these attacks may not work, but you may be vulnerable to different ones ☺
- Attacks in the "theoretical" category can move to the practical in a matter of days
- All testing was done on Cisco equipment, Ethernet switch attack resilience varies widely from vendor to vendor
- This is not a comprehensive talk on configuring Ethernet switches for security; the focus is on L2 attacks and their mitigation

Why Worry about Layer 2 Security?

OSI Was Built to Allow Different Layers to Work without Knowledge of Each Other

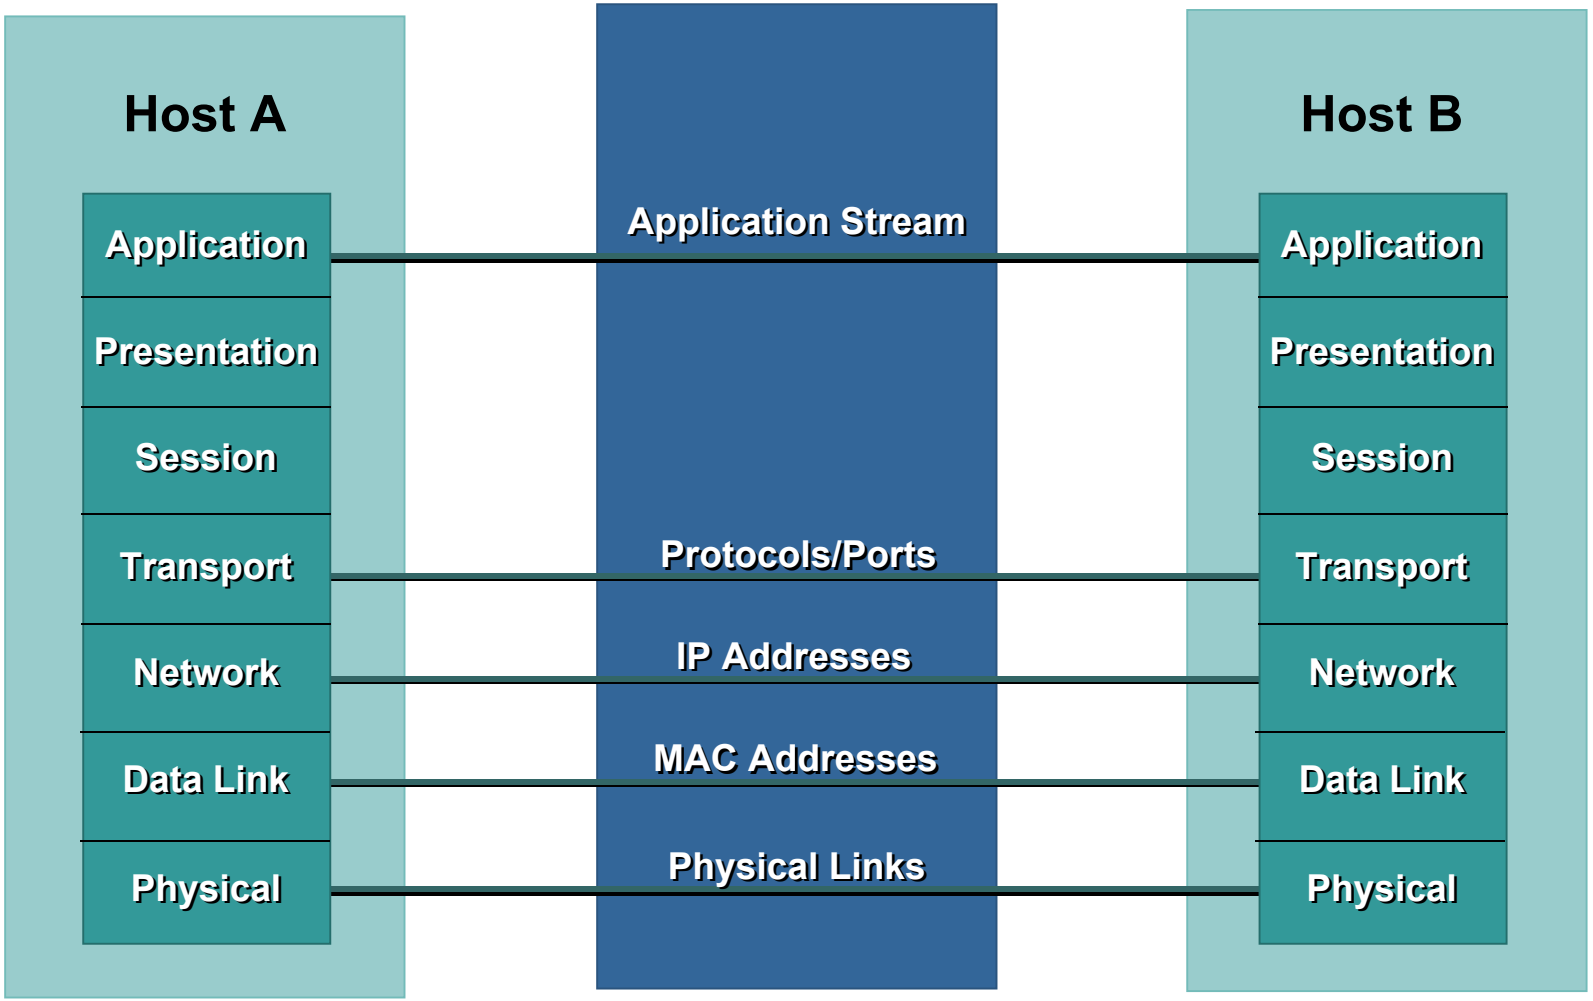

The Domino Effect

- Unfortunately this means if one layer is hacked, communications are compromised without the other layers being aware of the problem
- Security is only as strong as your weakest link
- When it comes to networking, layer 2 can be a VERY weak link

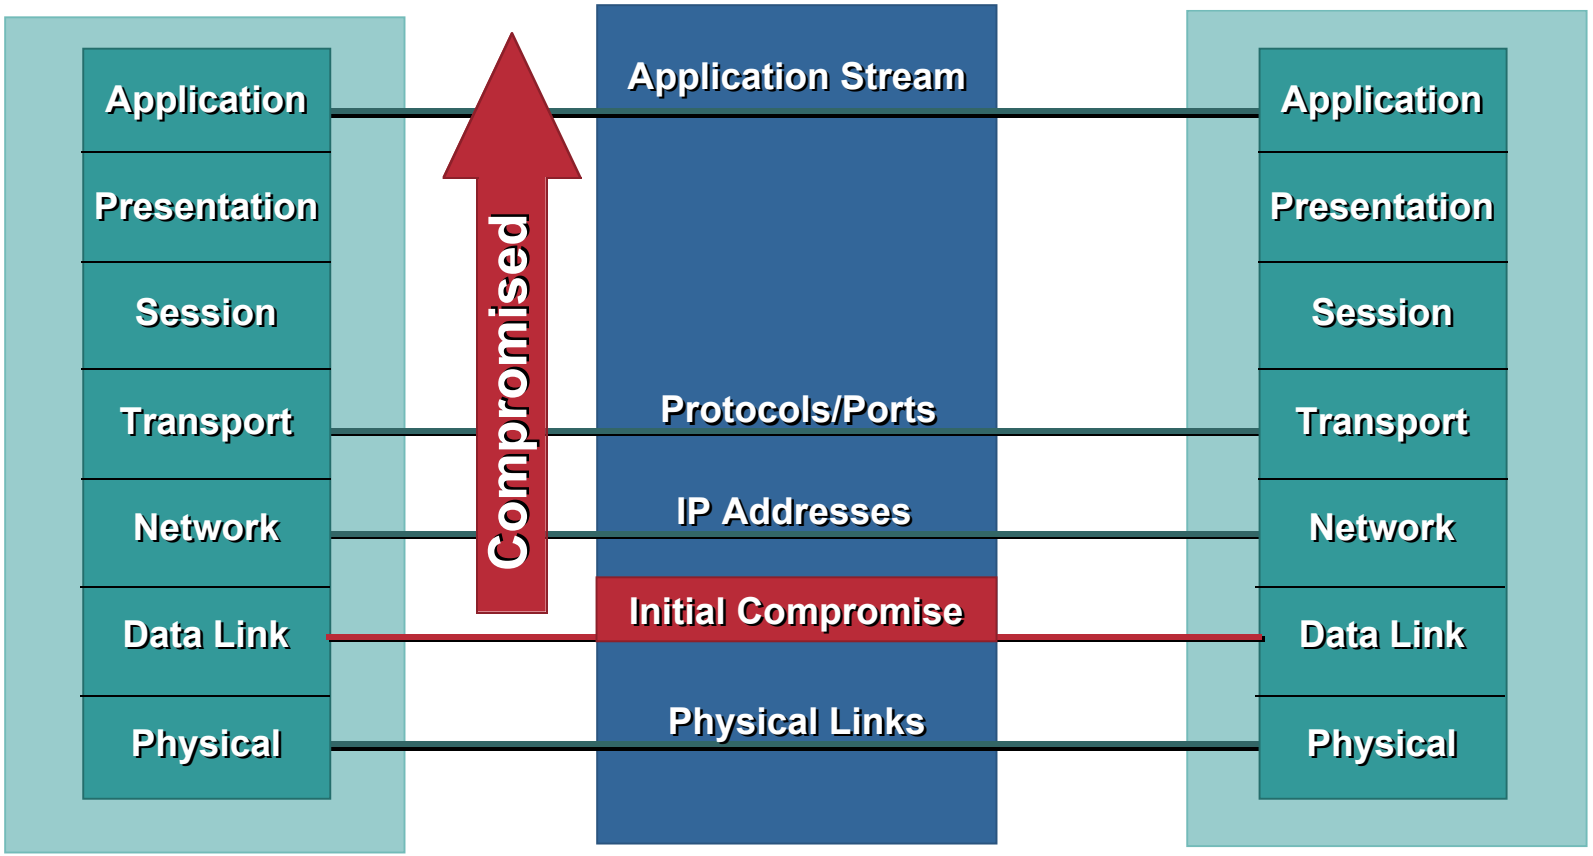

NetOPS/SecOPS, Who's Problem Is It?

Questions:	Most NetOPS	Most SecOPS
• What is your stance on L2 security issues? • Do you use VLANs often? • Do you ever put different security levels on the same switch using VLANs? • What is the process for allocating addresses for segments?	• There are L2 Security issues? • I use VLANs all the time • Routing in and out of the same switch is OK by me! That's what VLANs are for • The security guy asks me for a new segment, I create a VLAN and assign him an address space	• I handle security issues at L3 and above • I have no idea if we are using VLANs • Why would I care what the network guy does with the switch? • I ask Netops for a segment, they give me ports and addresses

The Numbers from CSI/FBI

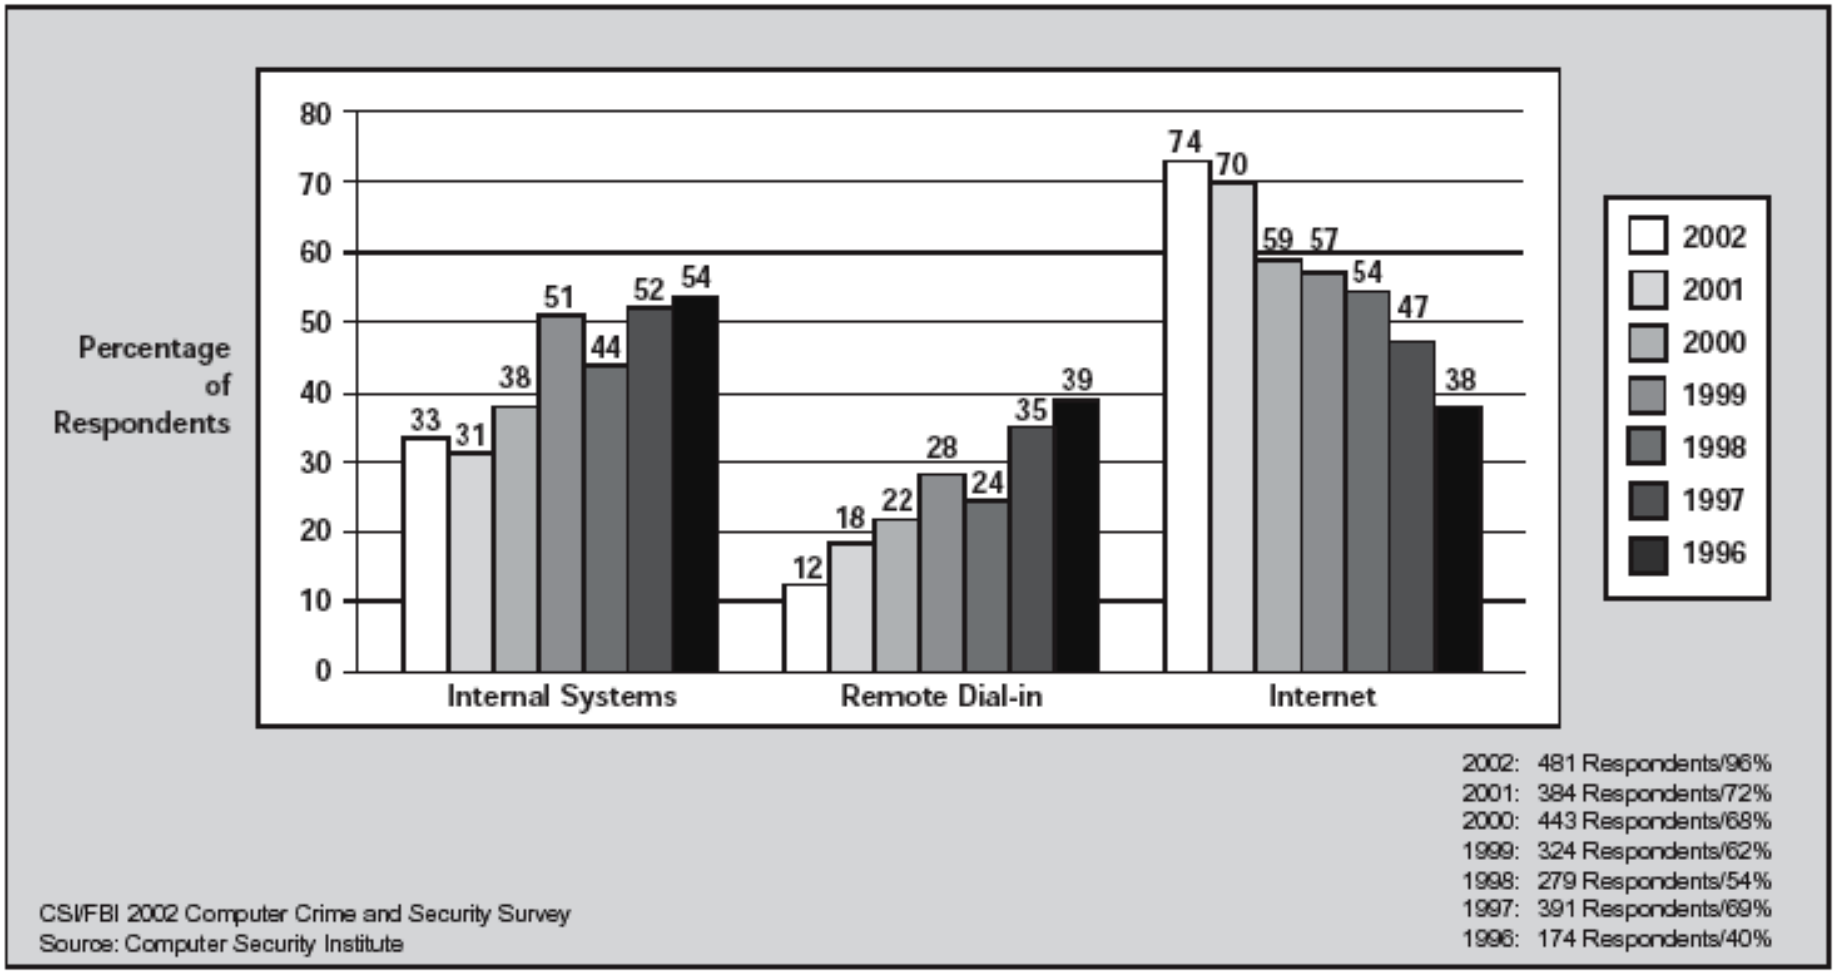

MAC Attacks

MAC Address/CAM Table Review

48 Bit Hexadecimal (Base16) Unique Layer Two Address

1234.5678.9ABC

First 24 bits = Manufacture Code Assigned by IEEE

0000.0cXX.XXXX

Second 24 bits = Specific Interface, Assigned by Manufacture

XXXX.XX00.0001

All F's = Broadcast

FFFF.FFFF.FFFF

- **CAM Table stands for Content Addressable Memory**
- **The CAM Table stores information such as MAC addresses available on physical ports with their associated VLAN parameters**
- **CAM Tables have a fixed size**

Normal CAM Behaviour 1/3

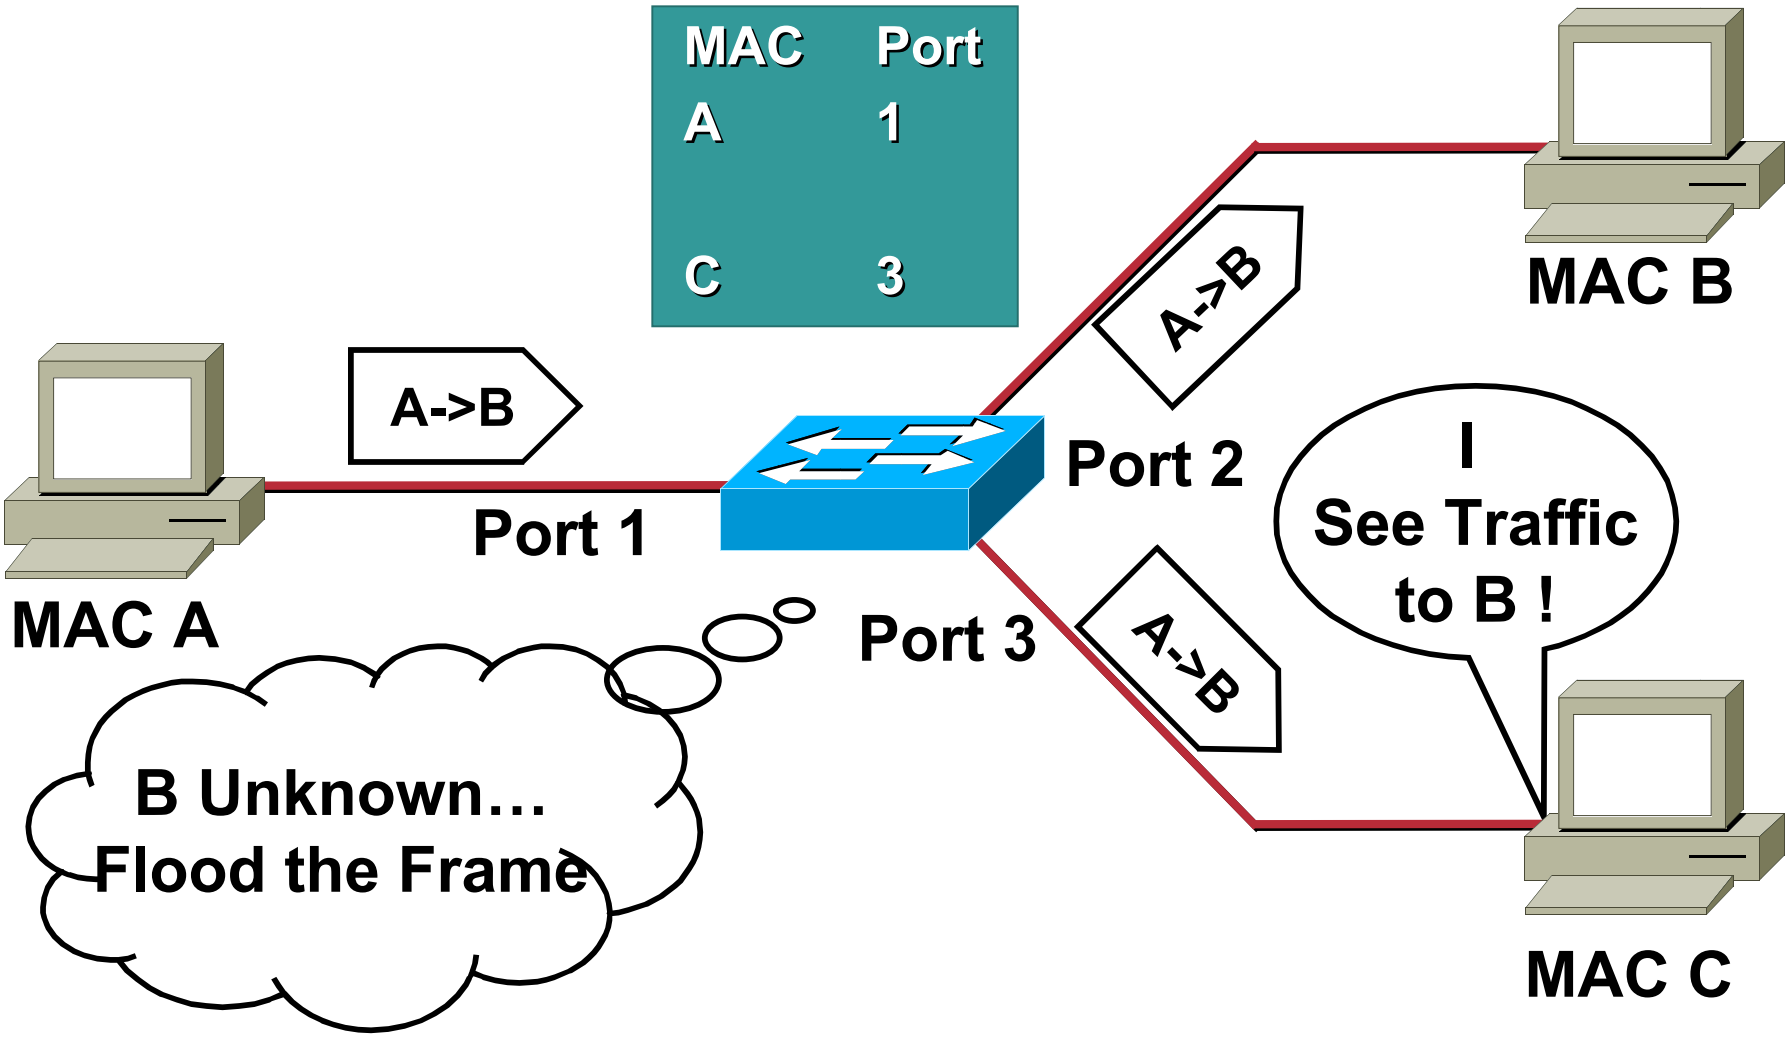

Normal CAM Behaviour 2/3

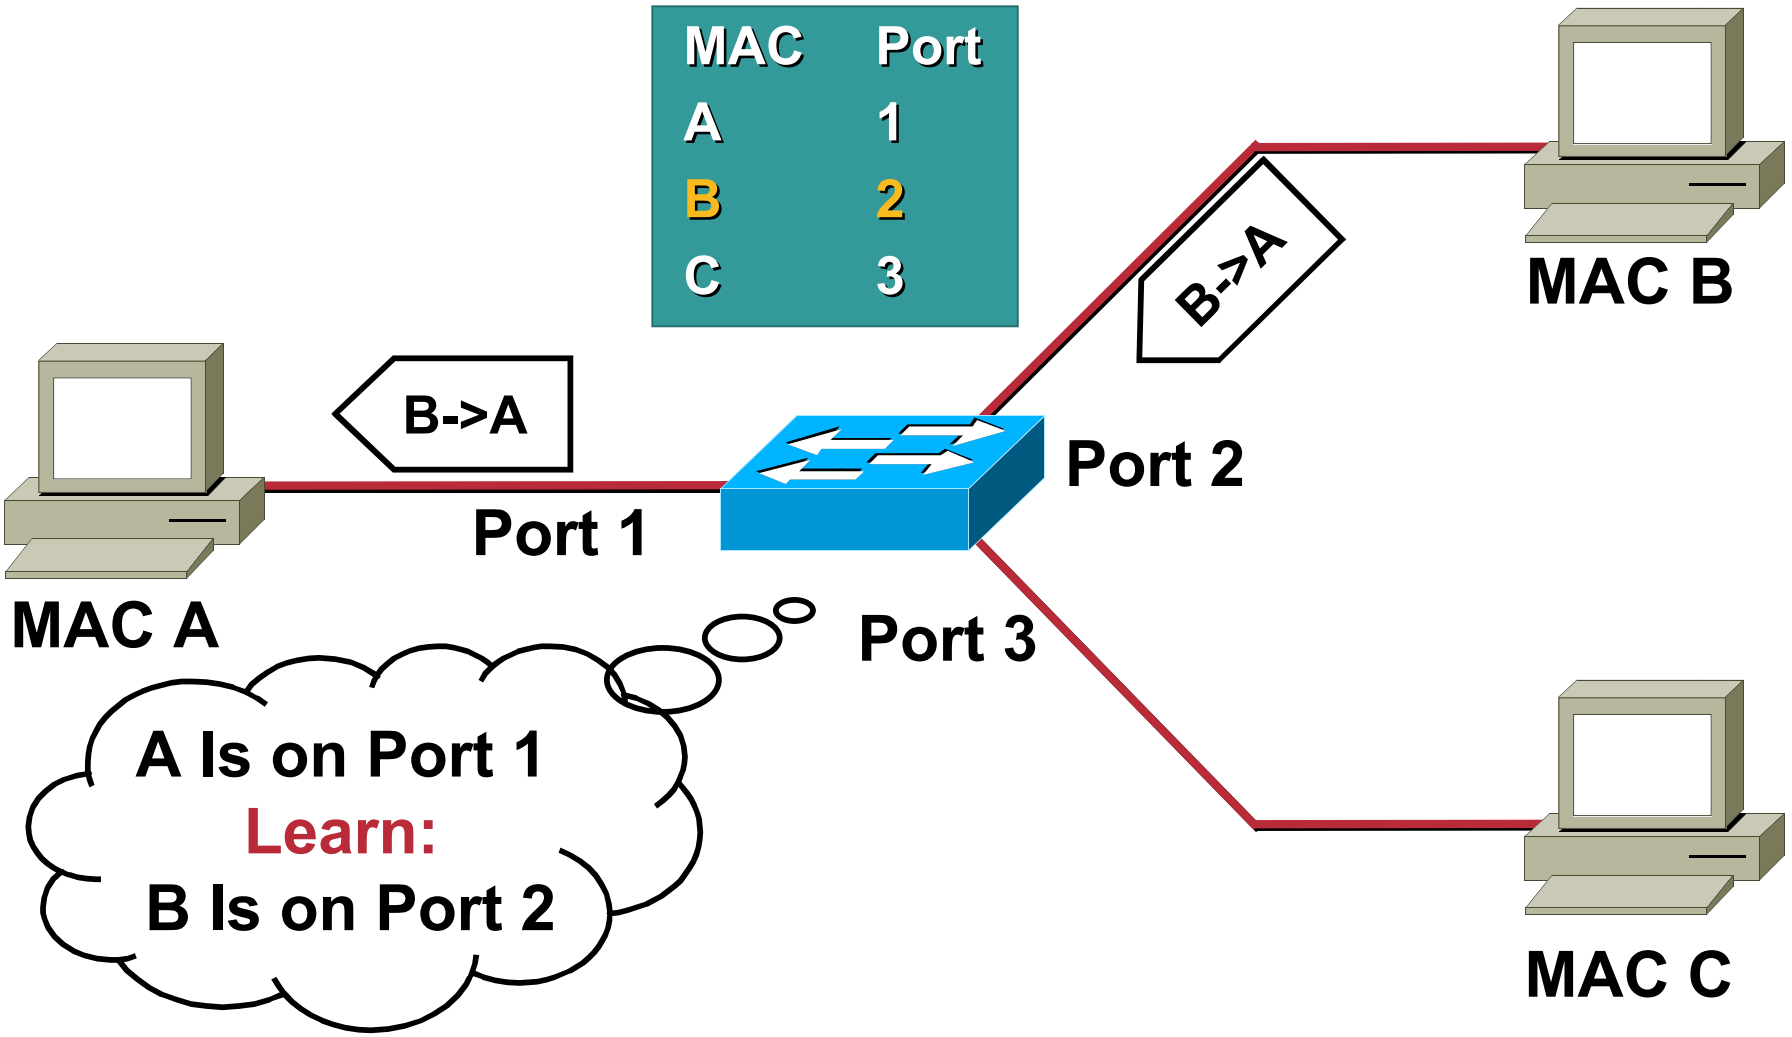

Normal CAM Behaviour 3/3

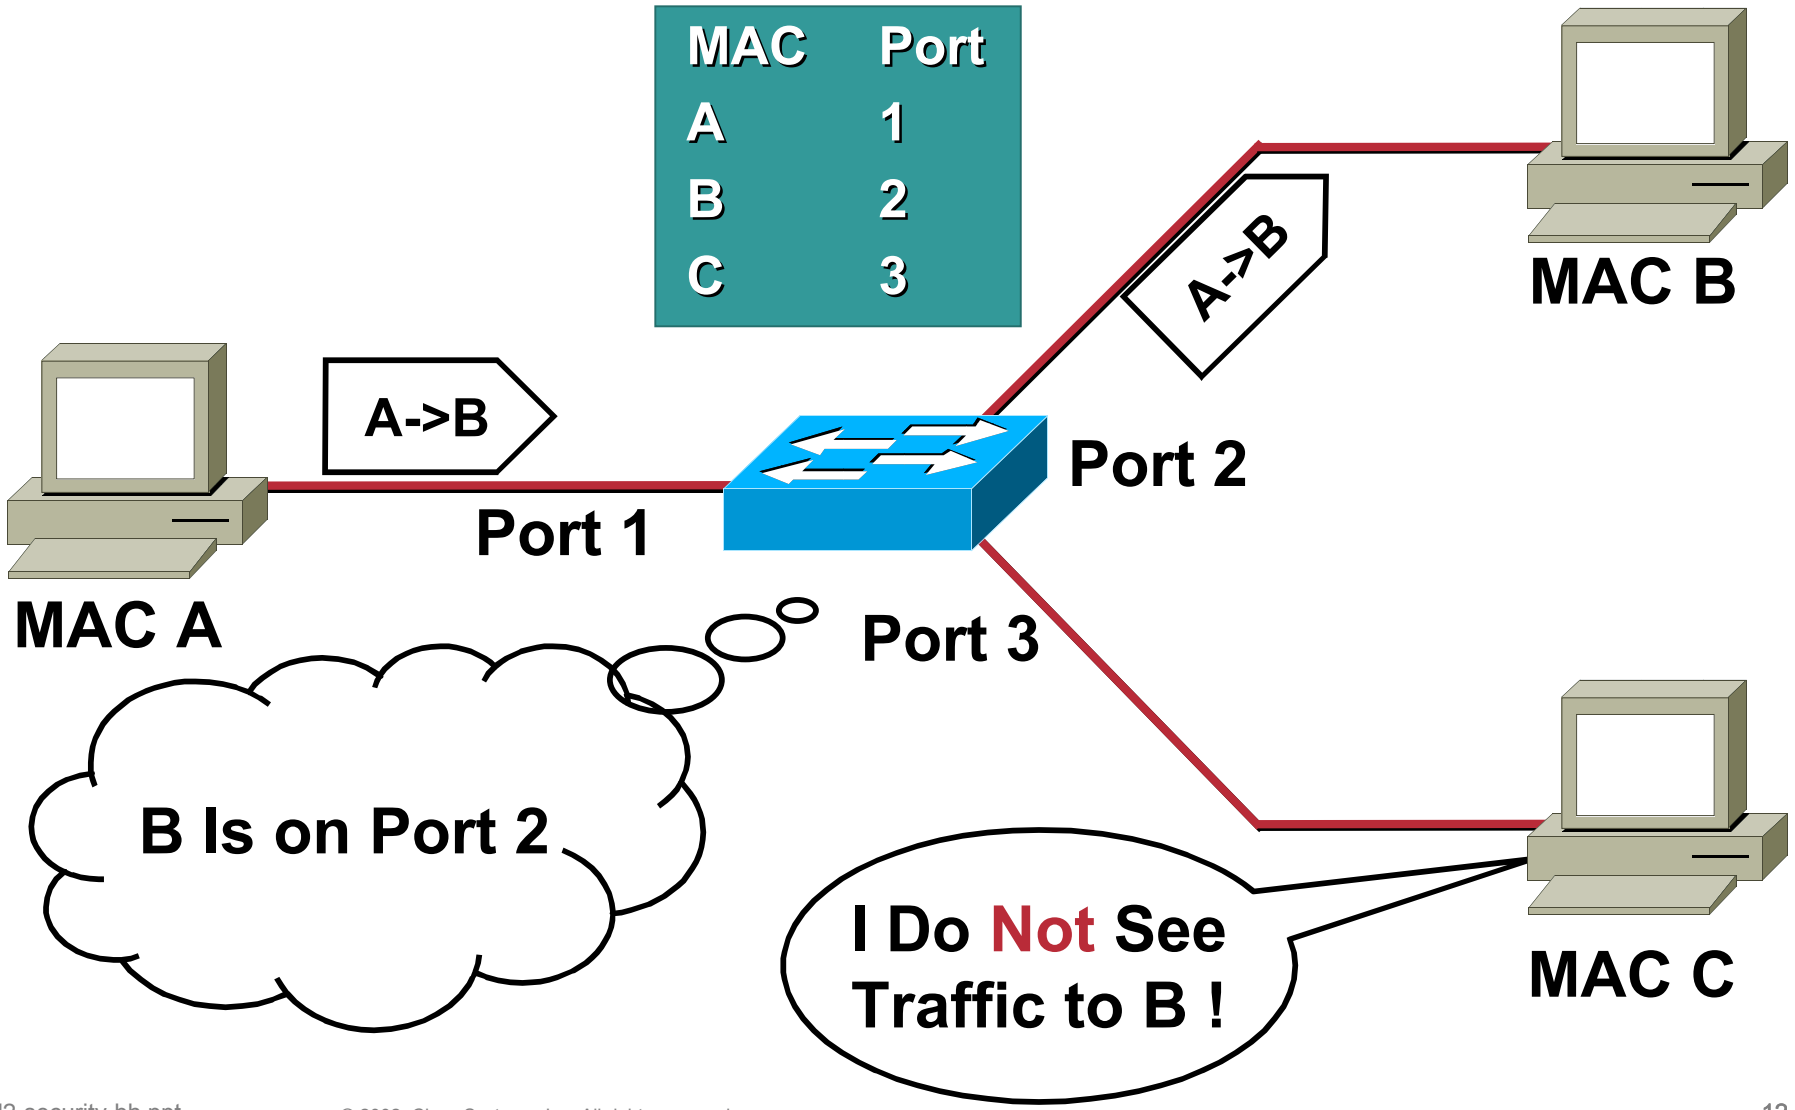

CAM Overflow 1/3

- **Theoretical attack until May 1999**
- ***macof* tool since May 1999**
  - **About 100 lines of perl from Ian Vitek**
  - **Later ported to C by Dug Song for “dsniff”**
- **Based on CAM Table’s limited size**

CAM Overflow 2/3

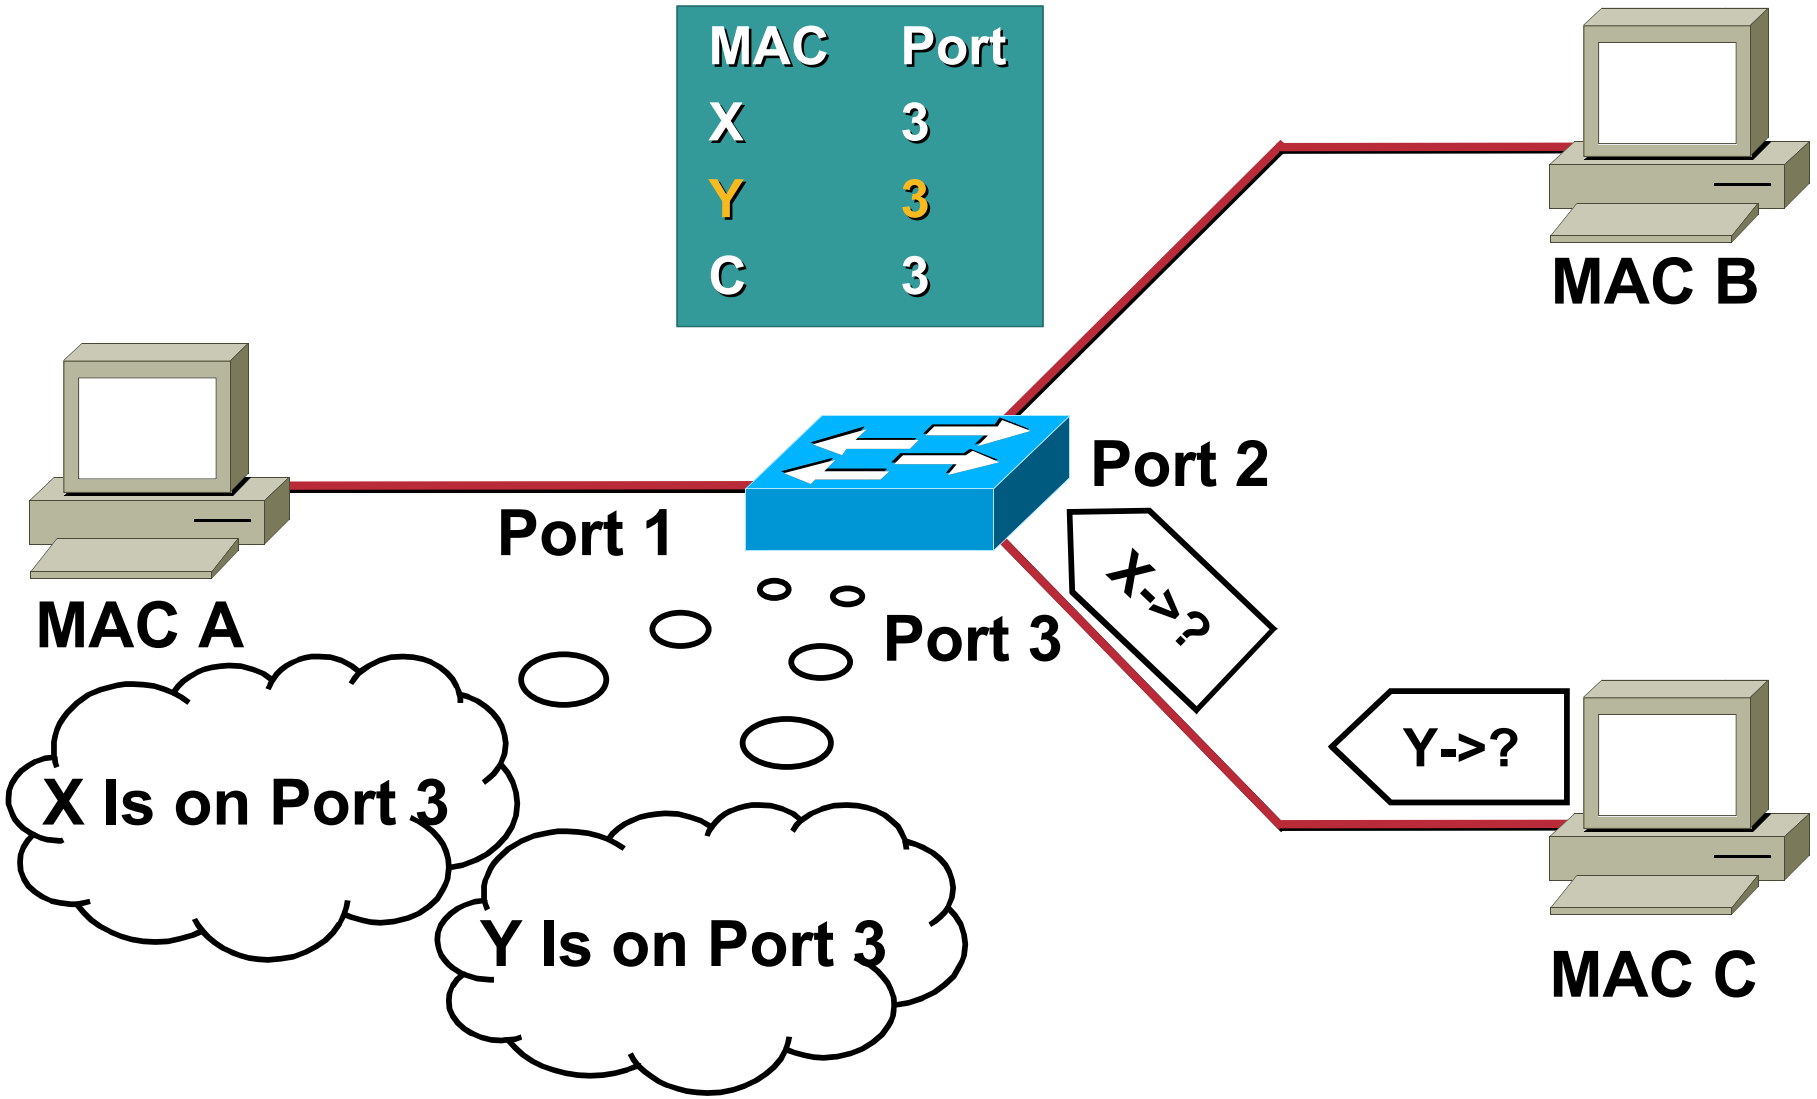

CAM Overflow 3/3

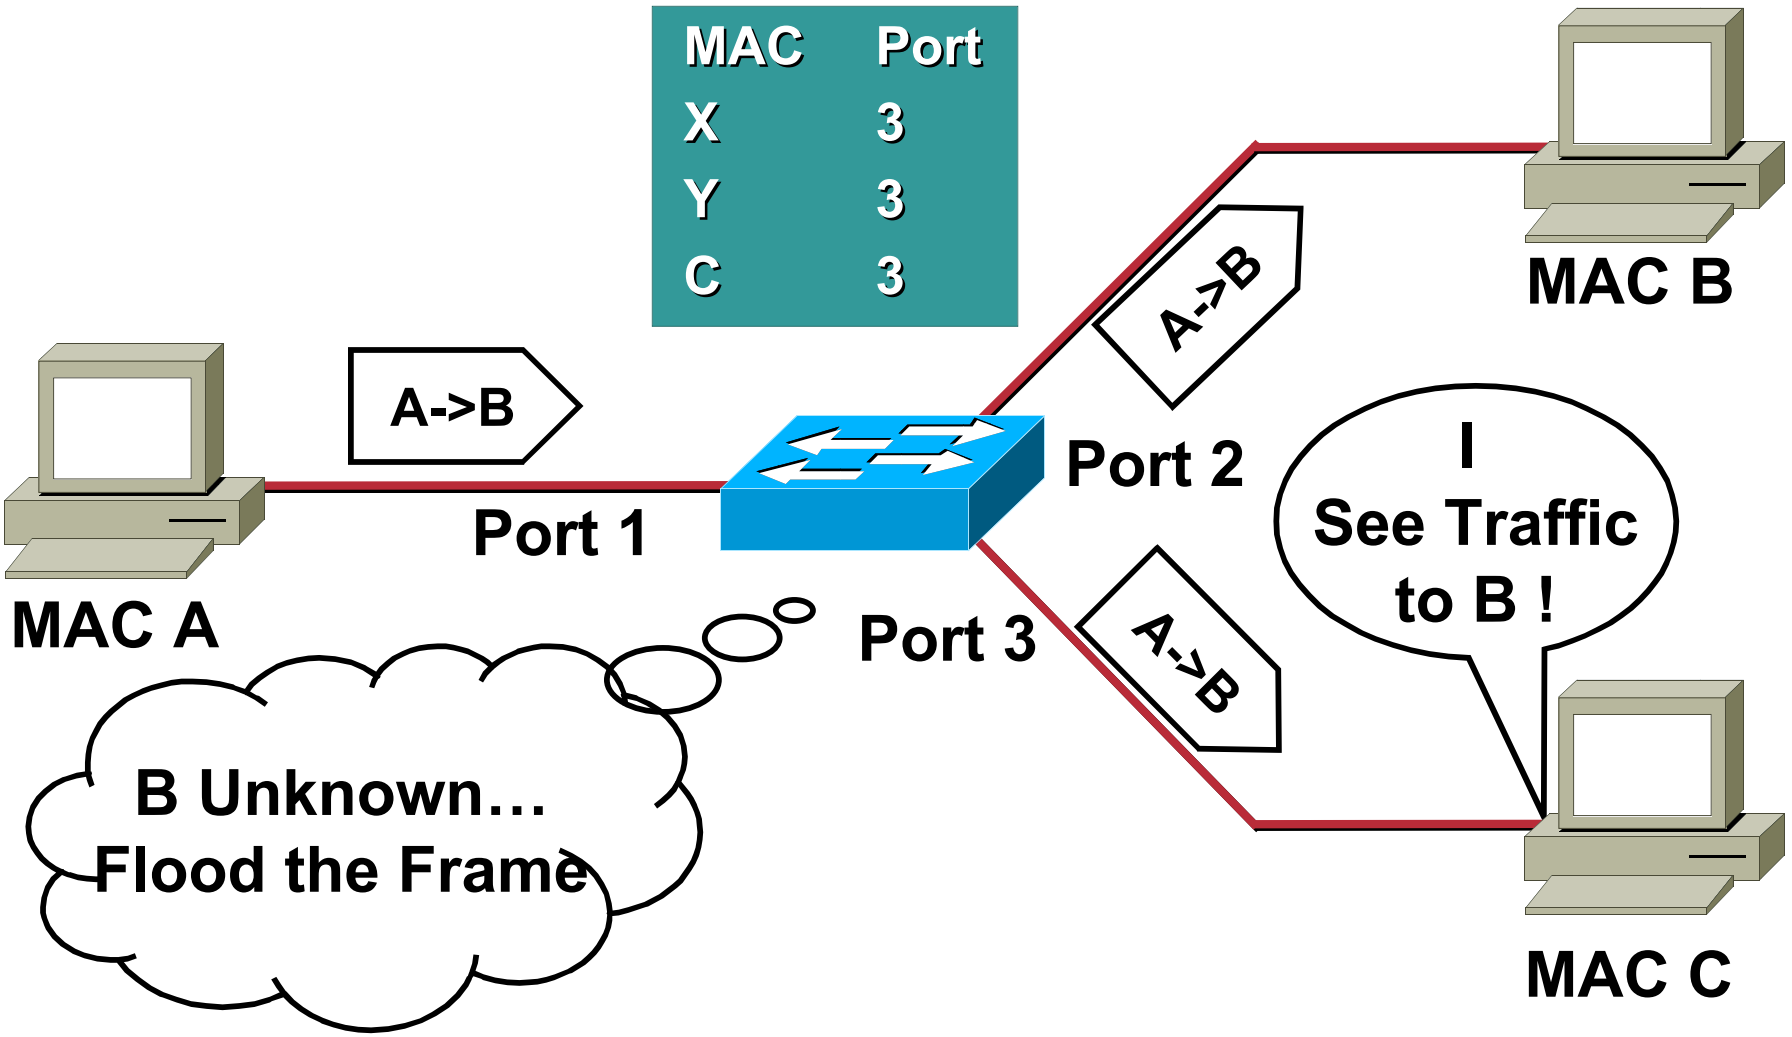

Catalyst CAM Tables

- Catalyst switches use hash to place MAC in CAM table

1	A	B	C						
2	D	E	F	G					
3	H								
.	I								
.	J	K							
16,000	L	M	N	O	P	Q	R	S	T

Flooded!

- 63 bits of source (MAC, VLAN, misc) creates a 17 bit hash value

  If the value is the same there are 8 buckets to place CAM entries, if all 8 are filled the packet is flooded

MAC Flooding Switches with Macof

Cisco.com

```
[root@attack-lnx dsniff-2.3]# ./macof
b5:cf:65:4b:d5:59 2c:01:12:7d:bd:36 0.0.0.0.4707 > 0.0.0.0.28005: S 106321318:106321318(0) win 512
68:2a:55:6c:1c:1c bb:33:bb:4d:c2:db 0.0.0.0.44367 > 0.0.0.0.60982: S 480589777:480589777(0) win 512
1e:95:26:5e:ab:4f d7:80:6f:2e:aa:89 0.0.0.0.42809 > 0.0.0.0.39934: S 1814866876:1814866876(0) win 512
51:b5:4a:7a:03:b3 70:a9:c3:24:db:2d 0.0.0.0.41274 > 0.0.0.0.31780: S 527694740:527694740(0) win 512
51:75:2e:22:c6:31 91:a1:c1:77:f6:18 0.0.0.0.36396 > 0.0.0.0.15064: S 1297621419:1297621419(0) win 512
7b:fc:69:5b:47:e2 e7:65:66:4c:2b:87 0.0.0.0.45053 > 0.0.0.0.4908: S 976491935:976491935(0) win 512
19:14:72:73:6f:ff 8d:ba:5c:40:be:d5 0.0.0.0.867 > 0.0.0.0.20101: S 287657898:287657898(0) win 512
63:c8:58:03:4e:f8 82:b6:ae:19:0f:e5 0.0.0.0.58843 > 0.0.0.0.40817: S 1693135783:1693135783(0) win 512
33:d7:e0:2a:77:70 48:96:df:20:61:b4 0.0.0.0.26678 > 0.0.0.0.42913: S 1128100617:1128100617(0) win 512
f2:7f:96:6f:d1:bd c6:15:b3:21:72:6a 0.0.0.0.53021 > 0.0.0.0.5876: S 570265931:570265931(0) win 512
22:6a:3c:4b:05:7f 1a:78:22:30:90:85 0.0.0.0.58185 > 0.0.0.0.51696: S 1813802199:1813802199(0) win 512
f6:60:da:3d:07:5b 3d:db:16:11:f9:55 0.0.0.0.63763 > 0.0.0.0.63390: S 1108461959:1108461959(0) win 512
bc:fd:c0:17:52:95 8d:c1:76:0d:8f:b5 0.0.0.0.55865 > 0.0.0.0.20361: S 309609994:309609994(0) win 512
bb:c9:48:4c:06:2e 37:12:e8:19:93:4e 0.0.0.0.1618 > 0.0.0.0.9653: S 1580205491:1580205491(0) win 512
e6:23:b5:47:46:e7 78:11:e3:72:05:44 0.0.0.0.18351 > 0.0.0.0.3189: S 217057268:217057268(0) win 512
c9:89:97:4b:62:2a c3:4a:a8:48:64:a4 0.0.0.0.23021 > 0.0.0.0.14891: S 1200820794:1200820794(0) win 512
56:30:ac:0b:d0:ef 1a:11:57:4f:22:68 0.0.0.0.61942 > 0.0.0.0.17591: S 1535090777:1535090777(0) win 512
```

CAM Table Full!

- Dsniff (macof) can generate 155,000 MAC entries on a switch per minute
- Assuming a perfect hash function, the CAM table will be completely filled after 131,052 (approx. 16,000 x 8) entries
  - Since hash isn't perfect it actually takes 70 seconds to fill the CAM table

```
CAT6506 (enable) sho cam count dynamic

Total Matching CAM Entries = 131052
```

- Once table is full, traffic without a CAM entry floods on the local VLAN, but NOT existing traffic with an existing CAM entry
- This attack will also fill CAM tables of adjacent switches

Snoop output on non-SPAN port 10.1.1.50

```
10.1.1.22 -> (broadcast)  ARP C Who is 10.1.1.1, 10.1.1.1 ?
10.1.1.22 -> (broadcast)  ARP C Who is 10.1.1.19, 10.1.1.19 ?
10.1.1.26 -> 10.1.1.25    ICMP Echo request (ID: 256 Sequence number: 7424) ← OOPS
10.1.1.25 -> 10.1.1.26    ICMP Echo reply (ID: 256 Sequence number: 7424) ← OOPS
```

MAC Flooding Attack Mitigation

- **Port Security**
  - **Capabilities are dependant on the platform**
  - **Allows you to specify MAC addresses for each port, or to learn a certain number of MAC addresses per port**
  - **Upon detection of an invalid MAC the switch can be configured to block only the offending MAC or just shut down the port**
  - **Port security prevents macof from flooding the CAM table**

http://cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_5_4/config/sec_port.htm

Port Security Details

- Beware management burden and performance hit
- Lots of platform specific options besides just “ON/OFF”

```
CatOS> (enable) set port security mod/ports... [enable | disable]
[mac_addr] [age {age_time}] [maximum {num_ of_mac}] [shutdown
{shutdown_time}] [violation{shutdown | restrict}]

IOS(config-if)# port security [action {shutdown | trap} | max-mac-
count addresses]
```

- MAC Tables do not have unlimited size (platform dependant)
- “Restrict” option may fail under macof load and disable the port, shutdown option is more appropriate

```
2002 Apr 03 15:40:32 %SECURITY-1-PORTSHUTDOWN:Port 3/21 shutdown due to no space
```

Available in Cat 29XX, 4K, 5K, and 6K in CatOS 5.2; 29/3500XL in 11.2(8)SA; 2950 in 12.0(5.2)WC(1); 3550 in 12.1(4)EA1

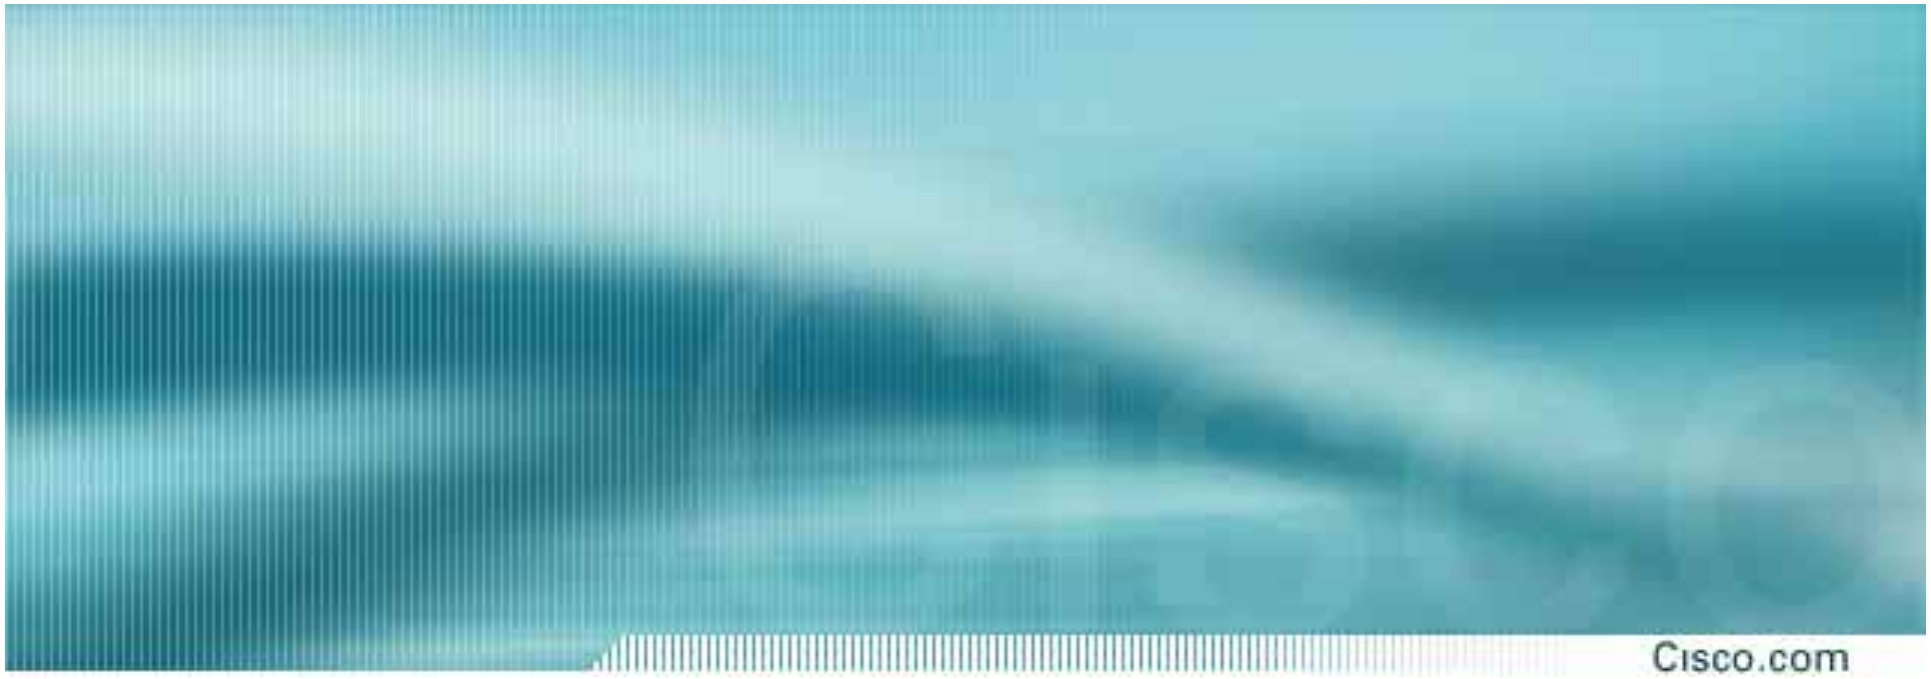

VLAN “Hopping” Attacks

Trunk Port Refresher

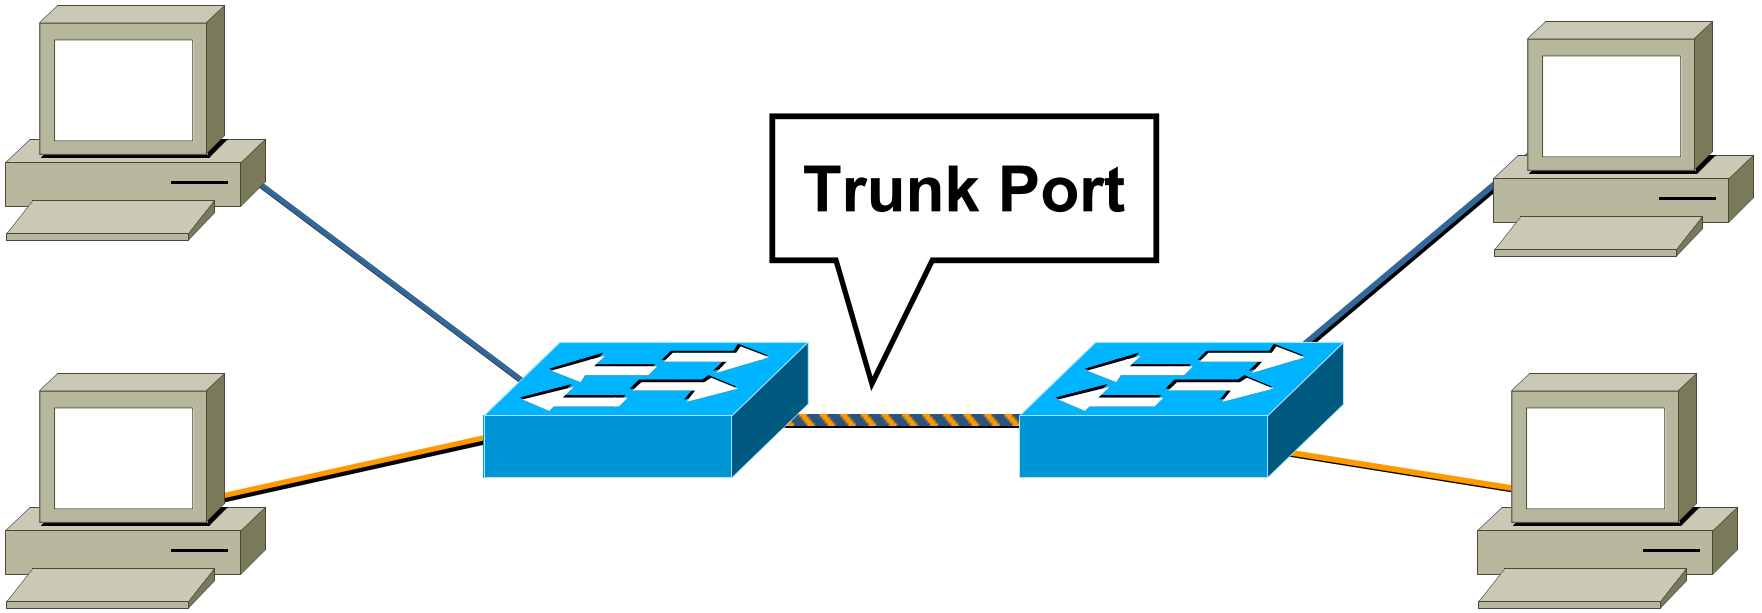

- Trunk ports have access to all VLANs by default
- Used to route traffic for multiple VLANs across the same physical link (generally used between switches)
- Encapsulation can be 802.1Q or ISL

Cisco Switching Control Protocols

- Used to negotiate trunk status, exchange VLAN information, etc.
- The majority use an IEEE 802.3 w/802.2 SNAP encapsulation
  - Includes LLC 0xAAAA03 (SNAP), and the Cisco OUI 0x00000C
  - Most use multicast destination addresses
    - Usually a variation on 0100.0ccc.cccc
  - Source address is derived from a bank of available addresses included in an EPROM on the chassis
  - SNAP Protocol Type varies and will be included through the rest of the talk.
- CDP and VTP (two common Cisco control protocols) are passed over VLAN 1 only. If VLAN 1 is cleared from a trunk, although no user data is transmitted or received, the switch continues to pass some control protocols on VLAN 1.
  - For this reason (and the fact that VLAN 1 can not be deleted) don't use it if you don't need to.

Lots of detail: http://www.cisco.com/warp/public/473/103.html

For the Detail Oriented: 802.3 w/802.2 SNAP

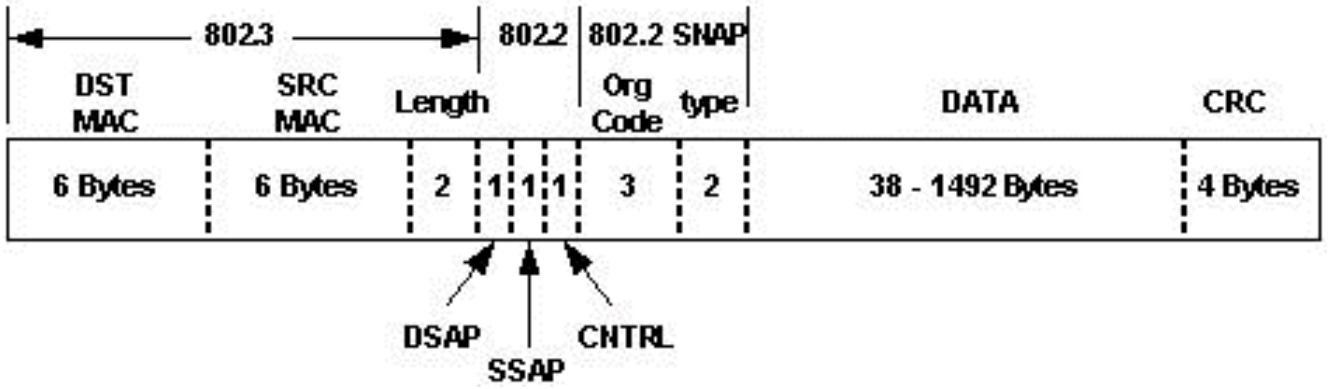

- **DST MAC: Generally a variant of 0100.0ccc.cccc**
- **SRC MAC: Pulled from a pool in the switch EPROM**
- **802.2 LLC Fields**

  DSAP:AA + SSAP:AA + CNTRL:03 = SNAP
- **802.2 SNAP Fields**

  Org Code: 0x00000c (Cisco)

  Protocol Type: Varies

If you like this sort of thing: http://www.cisco.com/warp/public/105/encheat.html

Dynamic Trunk Protocol (DTP)

- **What is DTP?**
  - Automates ISL/802.1Q trunk configuration
  - Operates between switches
  - Does not operate on routers
  - Not supported on 2900XL or 3500XL
- **DTP synchronizes the trunking mode on link ends**
- **DTP state on ISL/1Q trunking port can be set to "Auto", "On", "Off", "Desirable", or "Non-Negotiate"**

DST MAC	0100.0ccc.cccc
SNAP Proto	0x2004

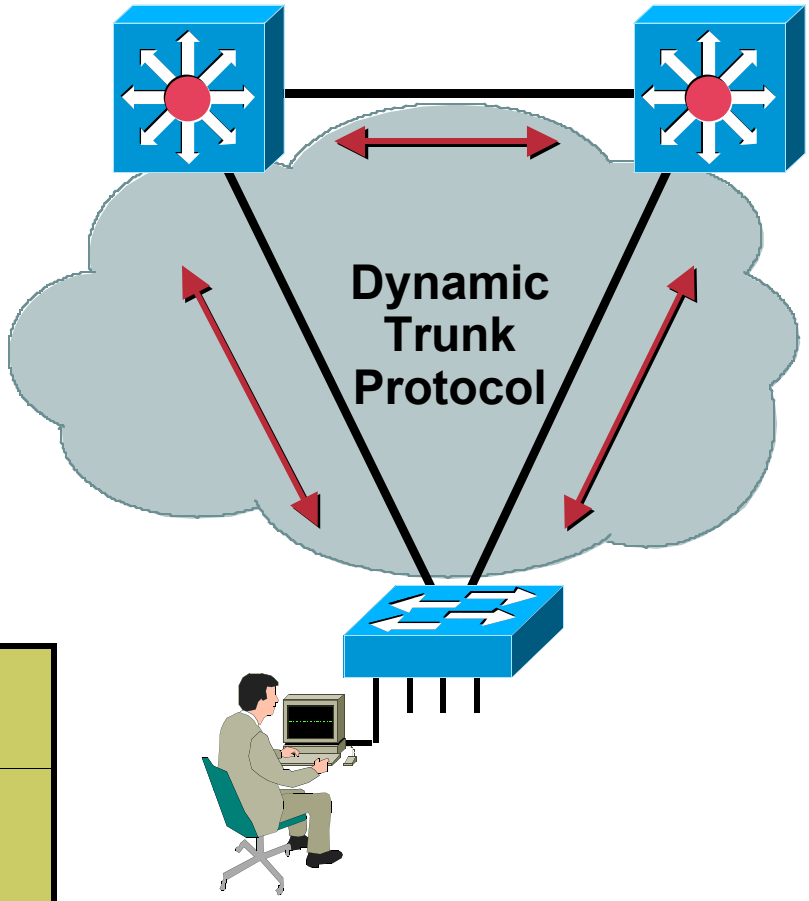

Basic VLAN Hopping Attack

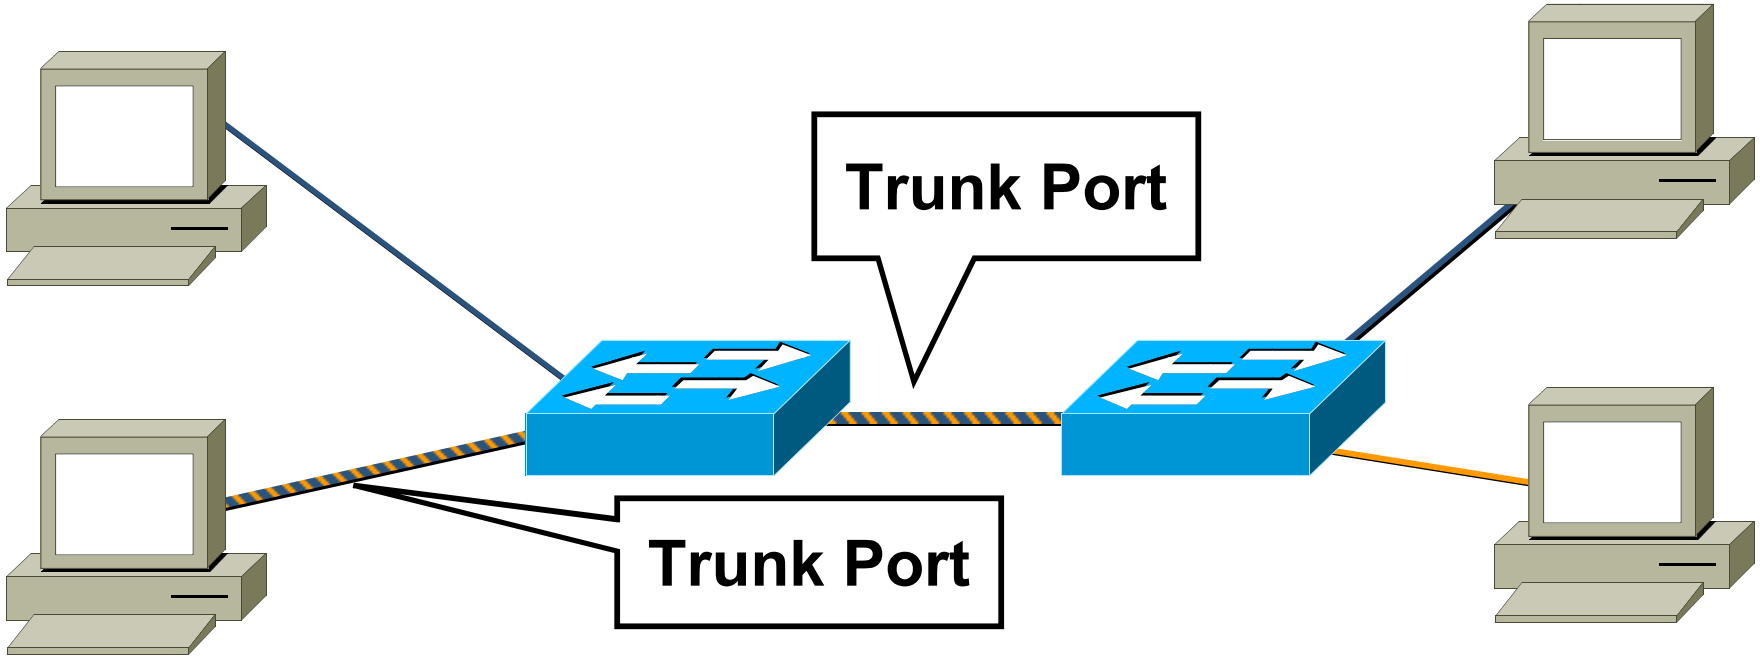

- A station can spoof as a switch with ISL or 802.1Q signaling (DTP signaling is usually required as well, or a rogue DTP speaking switch)
- The station is then member of all VLANs
- Requires a trunking favorable setting on the port (the SANS paper is two years old)

  http://www.sans.org/newlook/resources/IDFAQ/vlan.htm

Double Encapsulated 802.1q VLAN Hopping Attack

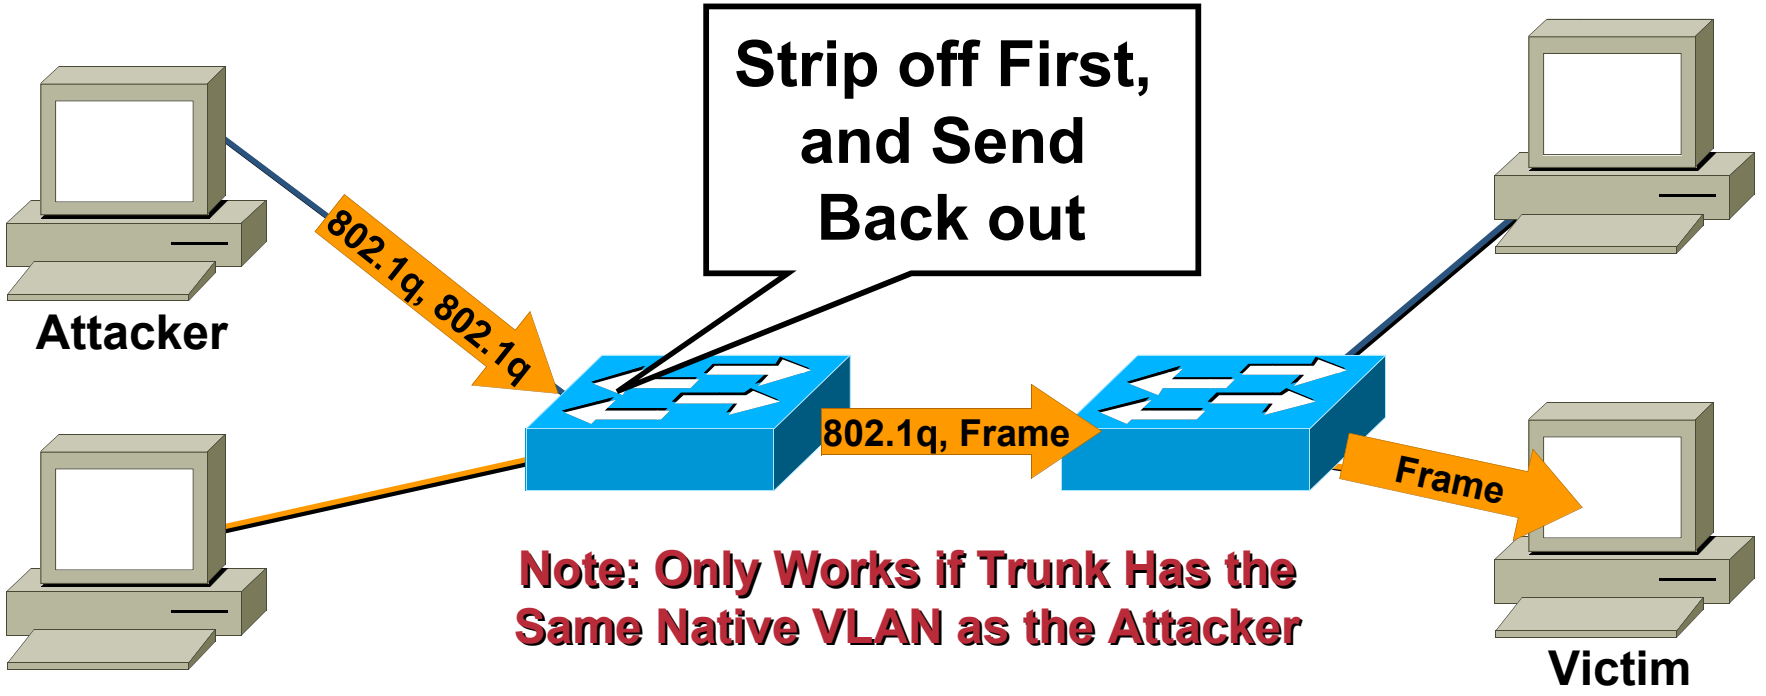

- Send double encapsulated 802.1Q frames
- Switch performs only one level of decapsulation
- Unidirectional traffic only
- Works even if trunk ports are set to off

Double Encap 802.1Q Ethereal Capture

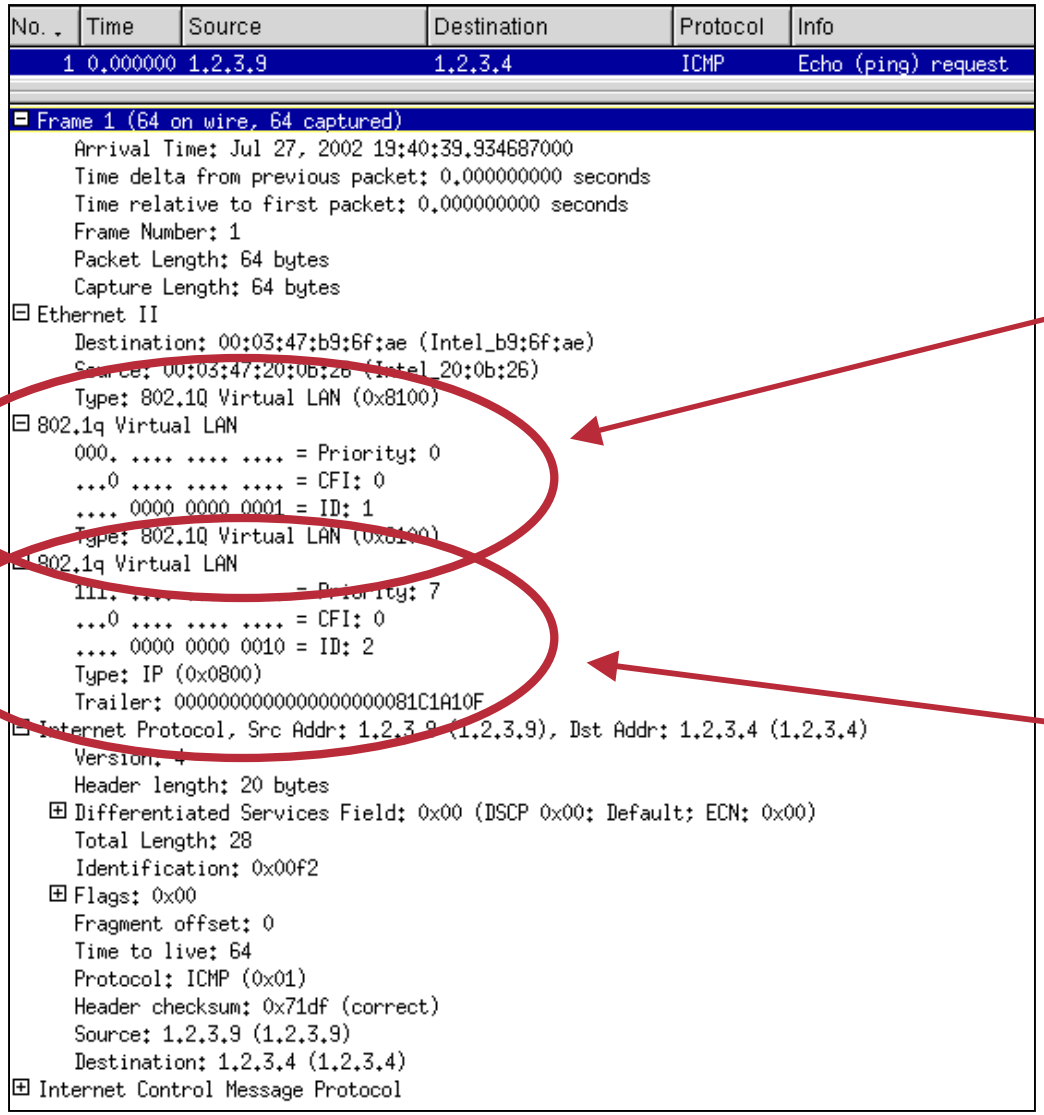

Outer Tag, Attacker VLAN

Inner Tag, Victim VLAN

Disabling Auto-Trunking


```
CatOS> (enable) set trunk <mod/port> off
IOS(config-if)#switchport mode access
```

- **Defaults change depending on switch; always check:**

  From the Cisco docs: "The default mode is dependent on the platform…"

  To check from the CLI:

```
CatOS> (enable) show trunk [mod|mod/port]
IOS# show interface type number switchport
```

Security Best Practices for VLANs and Trunking

- **Always** use a dedicated VLAN ID for all trunk ports
- Disable unused ports and put them in an unused VLAN
- Be paranoid: Do not use VLAN 1 for anything
- Set all user ports to non-trunking (DTP Off)

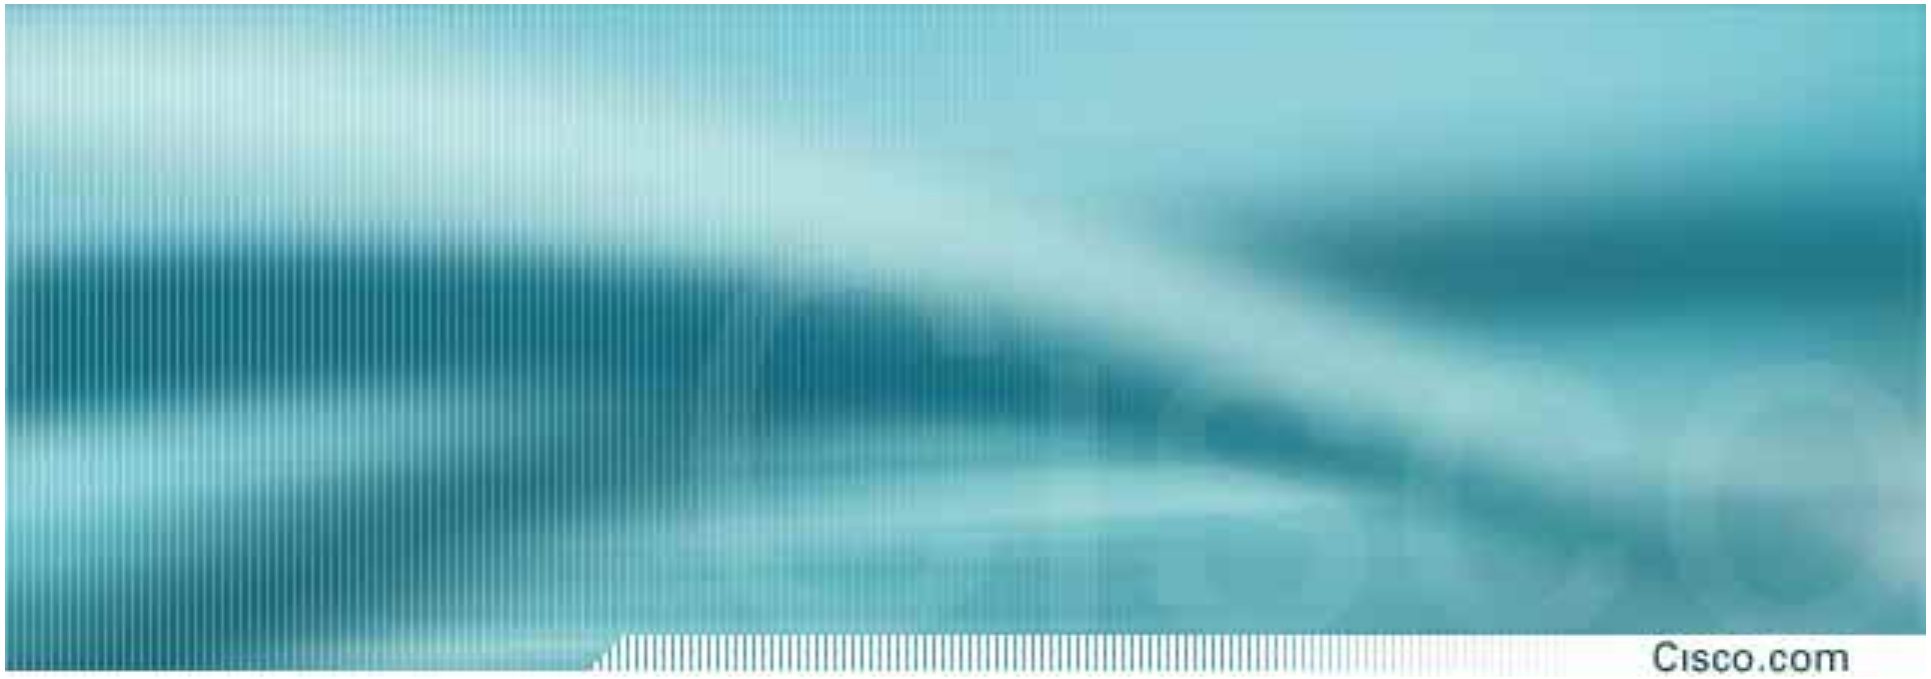

ARP Attacks

ARP Refresher

- An ARP request message should be placed in a frame and broadcast to all computers on the network
- Each computer receives the request and examines the IP address
- The computer mentioned in the request sends a response; all other computers process and discard the request without sending a response

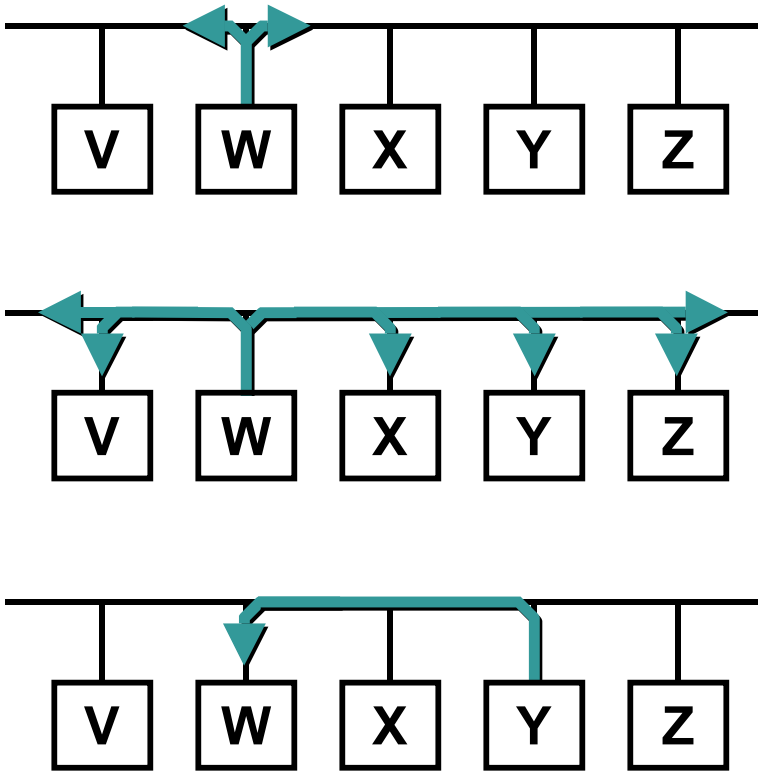

Gratuitous ARP

- Gratuitous ARP is used by hosts to “announce” their IP address to the local network and avoid duplicate IP addresses on the network; routers and other network hardware may use cache information gained from gratuitous ARPs
- Gratuitous ARP is a broadcast packet (like an ARP request)

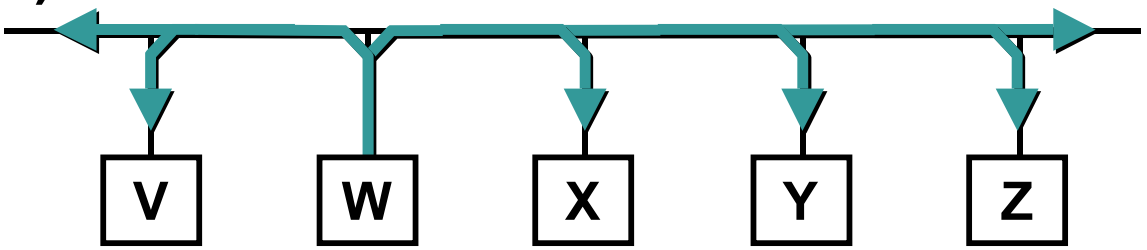

- HOST W: Hey everyone I’m host W and my IP Address is 1.2.3.4 and my MAC address is 12:34:56:78:9A:BC

Misuse of Gratuitous ARP

- ARP has no security or ownership of IP or MAC addresses
- What if we did the following?

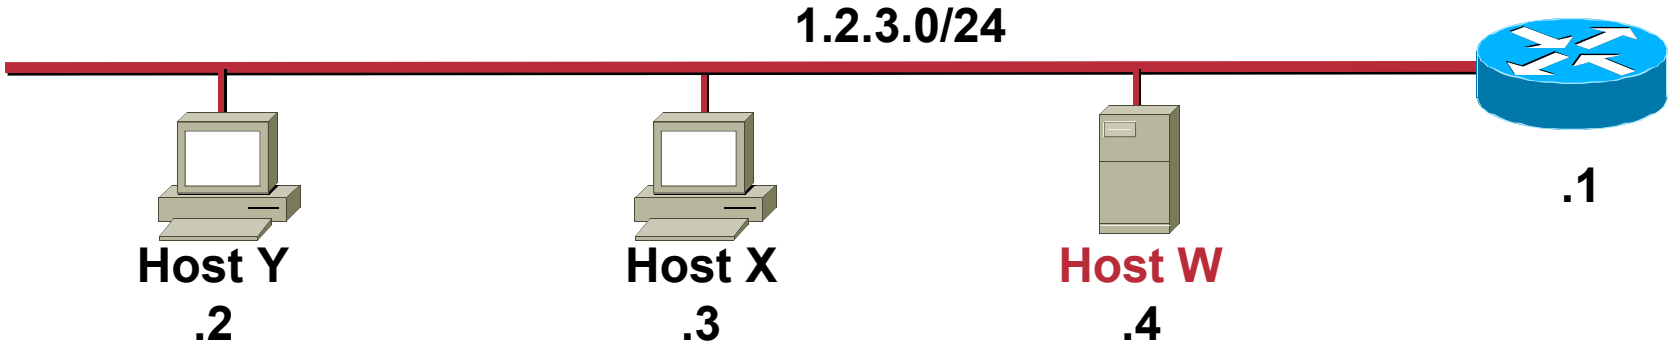

- Host W broadcasts I'm 1.2.3.1 with MAC 12:34:56:78:9A:BC
- (Wait 5 seconds)
- Host W broadcasts I'm 1.2.3.1 with MAC 12:34:56:78:9A:BC

A Test in the Lab

- Host X and Y will likely ignore the message unless they currently have an ARP table entry for 1.2.3.1

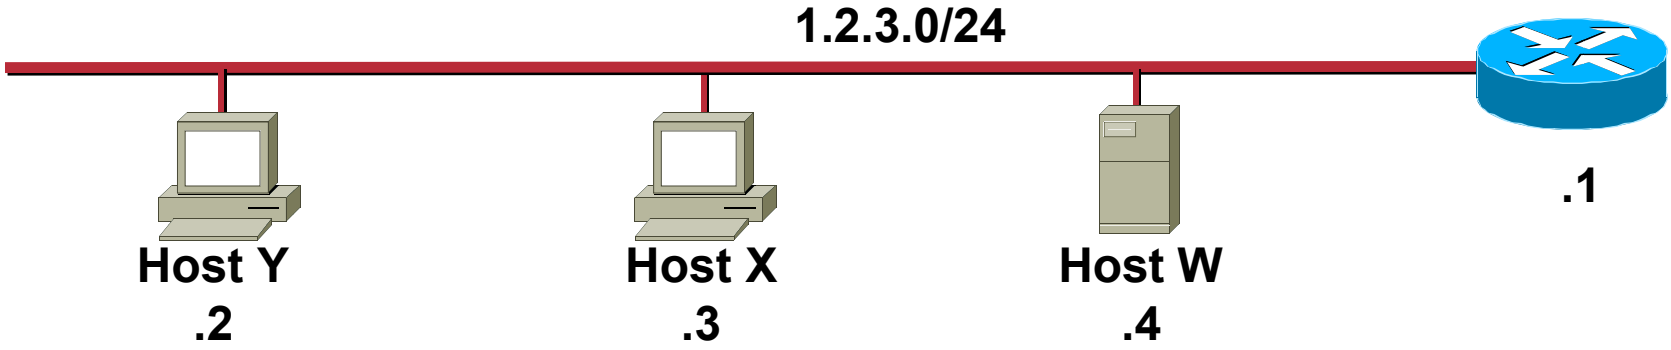

- When host Y requests the MAC of 1.2.3.1 the real router will reply and communications will work until host W sends a gratuitous ARP again
- Even a static ARP entry for 1.2.3.1 on Y will get overwritten by the Gratuitous ARP on some OSs (NT4,WIN2K for sure)

Dsniff—A Collection of Tools to Do:

- ARP spoofing
- MAC flooding
- Selective sniffing
- SSH/SSL interception

Dug Song, Author of dsniff

www.monkey.org/~dugsong/dsniff

Arpspoof in Action


```
[root@attack-lnx dsniff-2.3]# ./arpspoof 10.1.1.1
0:4:43:f2:d8:1 ff:ff:ff:ff:ff:ff 0806 42: arp reply
10.1.1.1 is-at 0:4:4e:f2:d8:1
0:4:43:f2:d8:1 ff:ff:ff:ff:ff:ff 0806 42: arp reply
10.1.1.1 is-at 0:4:4e:f2:d8:1
0:4:43:f2:d8:1 ff:ff:ff:ff:ff:ff 0806 42: arp reply
10.1.1.1 is-at 0:4:4e:f2:d8:1
0:4:43:f2:d8:1 ff:ff:ff:ff:ff:ff 0806 42: arp reply
10.1.1.1 is-at 0:4:4e:f2:d8:1u
```

```
C:\>test

C:\>arp -d 10.1.1.1

C:\>ping -n 1 10.1.1.1

Pinging 10.1.1.1 with 32 bytes

Reply from 10.1.1.1: bytes=32 time<10ms TTL=255

C:\>arp -a

Interface: 10.1.1.26 on Interface 2
  Internet Address      Physical Address      Type
  10.1.1.1              00-04-4e-f2-d8-01     dynamic
  10.1.1.25             00-10-83-34-29-72     dynamic
C:\>arp -a

Interface: 10.1.1.26 on Interface 2
  Internet Address      Physical Address      Type
  10.1.1.1              00-10-83-34-29-72     dynamic
  10.1.1.25             00-10-83-34-29-72     dynamic
```

More on Arpspoof

- All traffic now flows through machine running dsniff in a half-duplex manner
  - Not quite a sniffer but fairly close
- Port security doesn't help
- Note that attack could be generated in the opposite direction by spoofing the destination host when the router sends its ARP request
- Attack could be more selective and just spoof one victim

Selective Sniffing

- **Once the dsniff box has started the arpspoof process, the magic begins:**

```
[root@attack-lnx dsniff-2.3]# ./dsniff -c
dsniff: listening on eth0
-----------------
07/17/01 10:09:48 tcp 10.1.1.26.1126 -> wwwin-abc.cisco.com.80 (http)
GET /SERVICE/Paging/page/ HTTP/1.1
Host: wwwin-abc.cisco.com
Authorization: Basic c2NvdlghV9UNMRH4lejDmaA== [myuser:mypassword]
```

Supports More than 30 Standardized/Proprietary Protocols:

FTP, Telnet, SMTP, HTTP, POP, poppass, NNTP, IMAP, SNMP, LDAP, Rlogin, RIP, OSPF, PPTP MS-CHAP, NFS, YP/NIS, SOCKS, X11, CVS, IRC, AIM, ICQ, Napster, PostgreSQL, Meeting Maker, Citrix ICA, Symantec pcAnywhere, NAI Sniffer, Microsoft SMB, Oracle SQL*Net, Sybase et Microsoft SQL

SSL/SSH Interception

- **Using dnsspoof all web sites can resolve to the dsniff host IP address:**

```
C:\>ping www.amazon.com

Pinging www.amazon.com [10.1.1.25] with 32 bytes of data:

Reply from 10.1.1.25: bytes=32 time<10ms TTL=249
Reply from 10.1.1.25: bytes=32 time<10ms TTL=249
Reply from 10.1.1.25: bytes=32 time<10ms TTL=249
Reply from 10.1.1.25: bytes=32 time<10ms TTL=249
```

- **Once that happens you can proxy all web connections through the dsniff host**

SSL/SSH Interception

- **Using dsniff (webmitm) most SSL sessions can be intercepted and bogus certificate credentials can be presented**

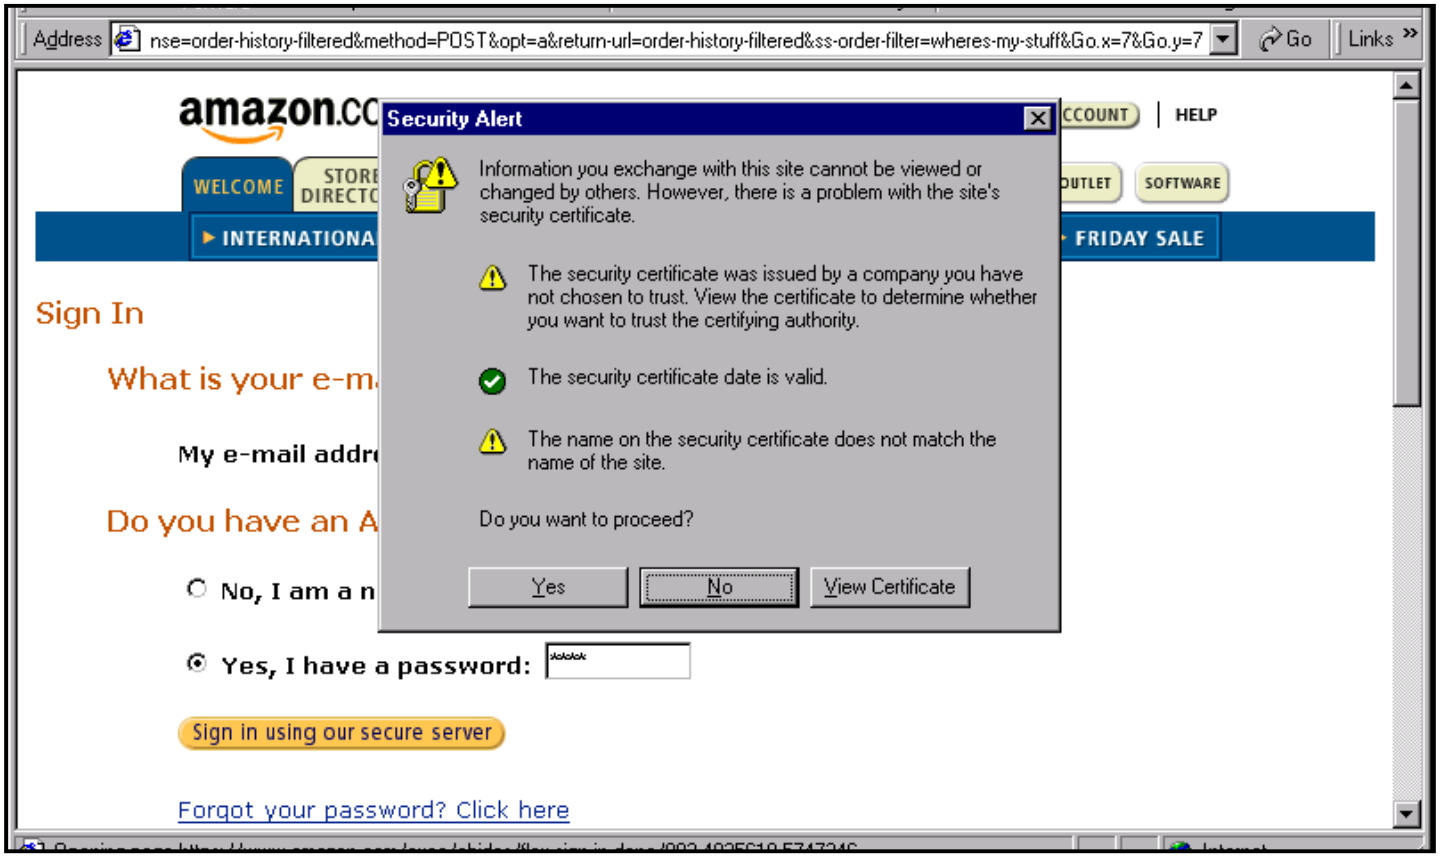

SSL/SSH Interception

- **Upon inspection they will look invalid but they would likely fool most users**

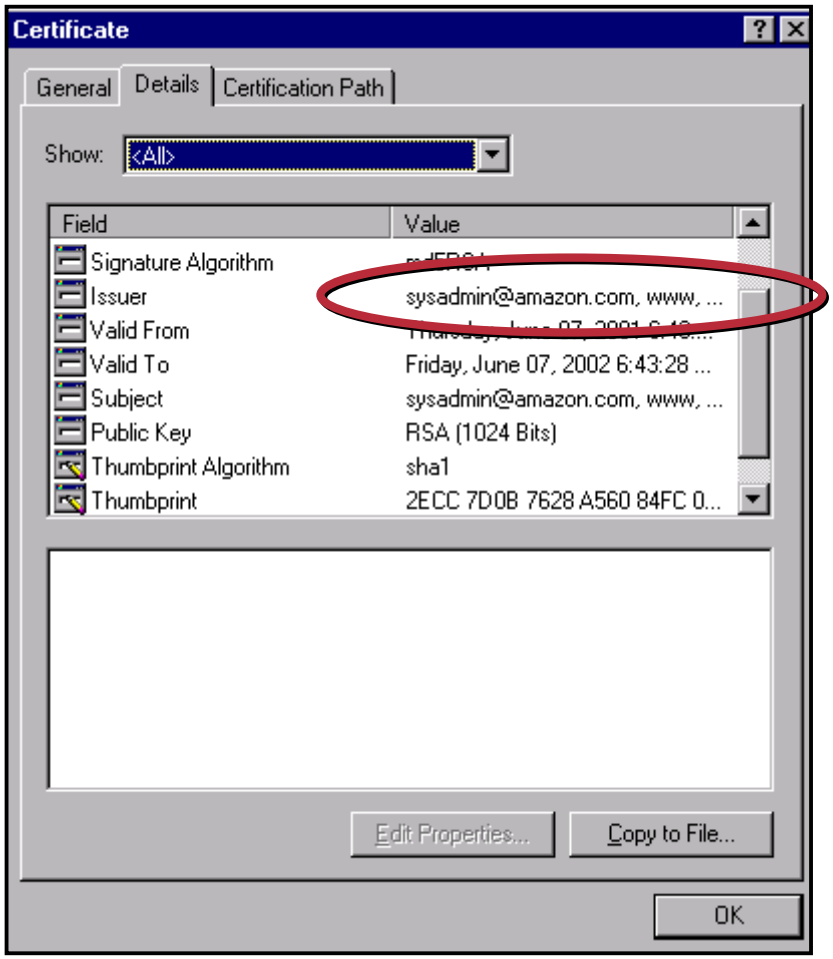

invalid

Dsniff evolves: Ettercap

- Similar to dsniff though not as many protocols supported for sniffing
- Can ARP spoof both sides of a session to achieve full-duplex sniffing
- Allows command insertion into persistent TCP sessions
- Menu driven interface
- http://ettercap.sourceforge.net/

Can It Get Much Easier?

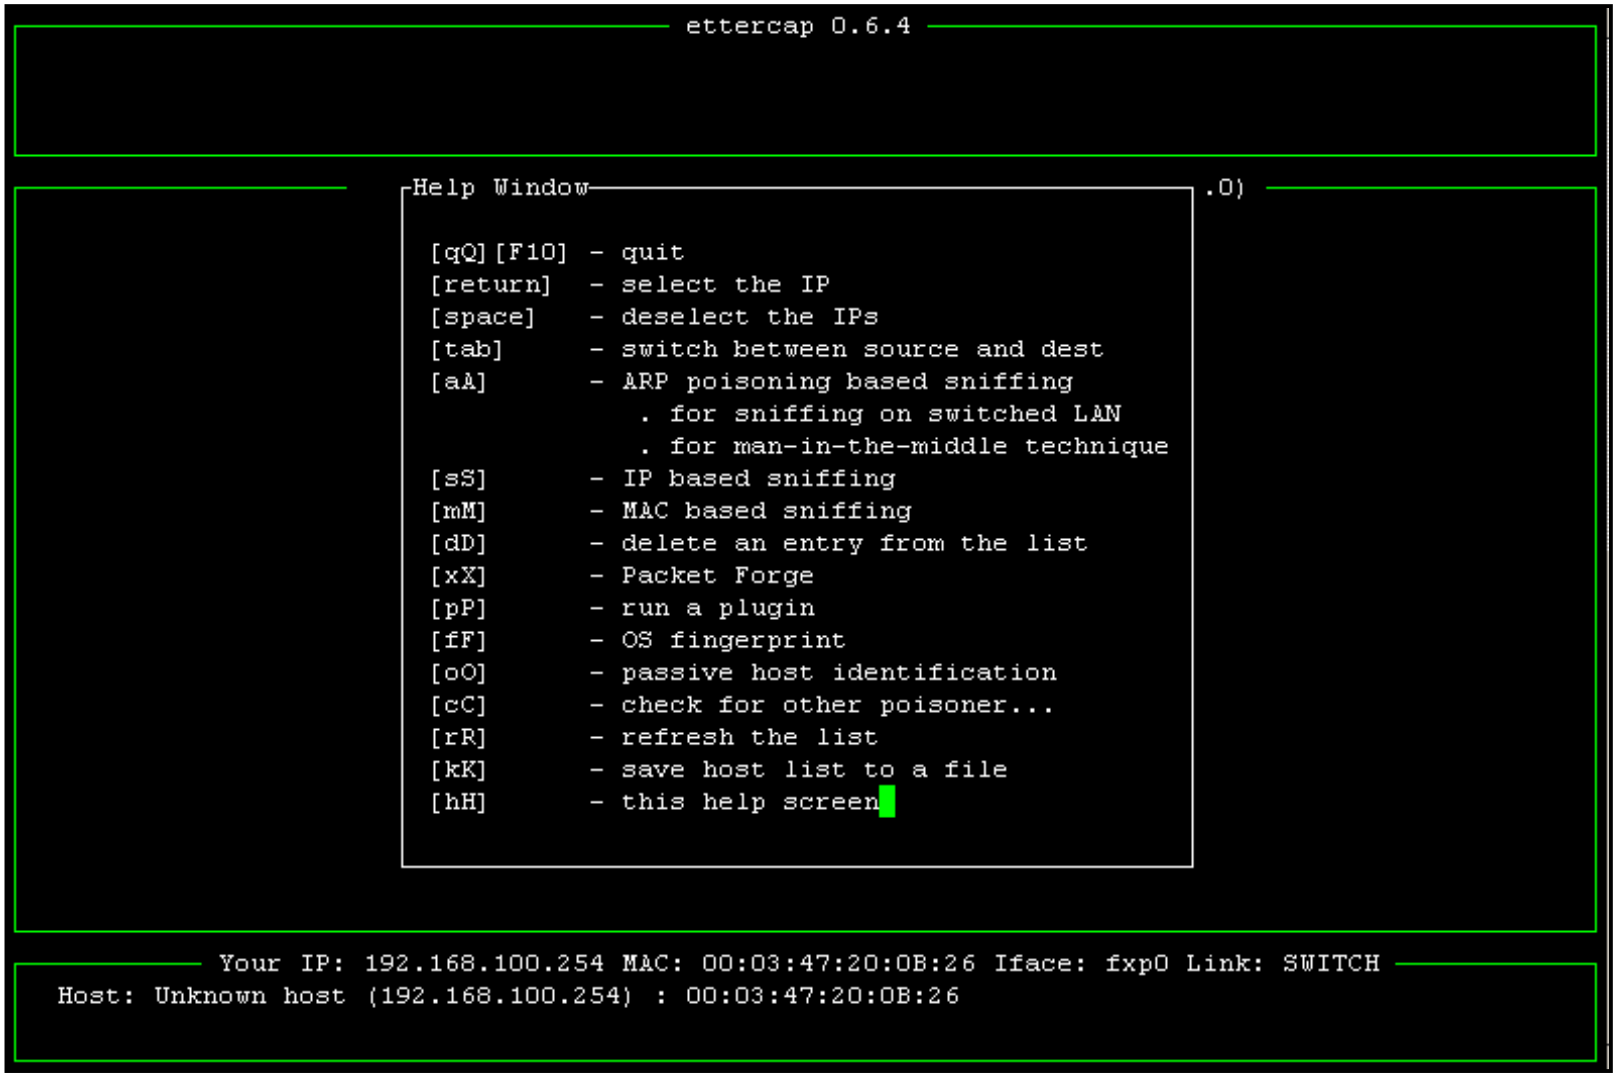

ARP Spoof Mitigation: Private VLANs

Only One Subnet!

- Primary VLAN
- Community VLAN
- Community VLAN
- Isolated VLAN

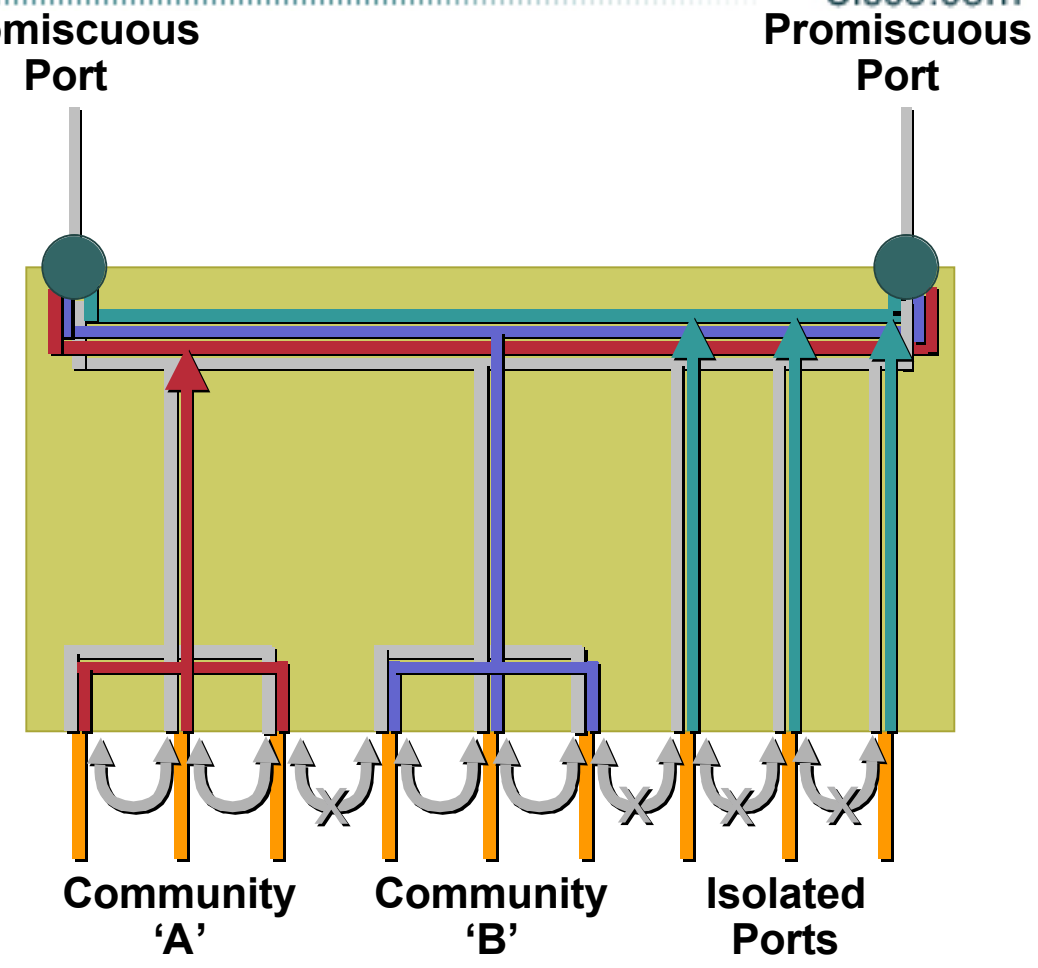

- PVLANs isolate traffic in specific communities to create distinct "networks" within a normal VLAN
- Note: Most inter-host communication is disabled with PVLANs turned on

http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_7_1/conf_gd/vlans.htm#xtocid854519

All PVLANs Are Not Created Equal

- On CAT 4K, 6K they are called Private VLANs
- On CAT 2K, 3K they are called Private VLAN edge or port protected
- CAT 4K,6K PVLANs support the following extra features:
  - Sticky ARP to mitigate default gateway attacks
    - ARP Entries do not age out
    - Changing ARP bindings requires manual intervention
  - PVLANs spanning multiple switches
  - Community Ports
- PVLANs are only compatible with Port Security on Cat 4K and 6K

Private VLAN Configuration

- **Available on: Cat 6K with CatOS 5.4(1); Cat 4K with CatOS 6.2; (no native IOS support); Cat6K IOS with12.1(11b)E and Cat4K IOS with 12.1(8a)EW; config can be a bit trickey (CatOS shown):**

```
CatOS> (enable) set vlan vlan_num pvlan-type primary

CatOS> (enable) set vlan vlan_num pvlan-type {isolated |
community}

CatOS> (enable) set pvlan primary_vlan_num {isolated_vlan_num |
community_vlan_num} mod/port

CatOS> (enable) set pvlan mapping primary_vlan_num
{isolated_vlan_num | community_vlan_num} mod/ports
```

- **Available as private VLAN edge (no community port support) on: 29/3500XL with 12.0(5)XU or later; 2950 with 12.0(5.2)WC(1); 3550 with 12.1(4)EA1**

```
IOS(config-if)#port protected

Any port without this command entered is
promiscuous
```

CatOS PVLAN Configuration Example

Cisco.com

```
bh-2002 (enable) set vlan 41 pvlan primary

VTP advertisements transmitting temporarily stopped,and will resume after the
command finishes. Vlan 41 configuration successful


bh-2002 (enable) show pvlan

Primary Secondary Secondary-Type   Ports

------- --------- ---------------- ------------

41      -         -

bh-2002 (enable) set vlan 42 pvlan isolated

VTP advertisements transmitting temporarily stopped,and will resume after the
command finishes. Vlan 42 configuration successful

bh-2002 (enable) set pvlan 41 42 3/9-10

Successfully set the following ports to Private Vlan 41,42:3/9-10


bh-2002 (enable) set pvlan mapping 41 42 3/35

Successfully set mapping between 41 and 42 on 3/35
```

More ARP Spoof Mitigation

- Some IDS systems will watch for an unusually high amount of ARP traffic
- ARPWatch is a freely available tool that will track IP/MAC address pairings
- Consider static ARP for critical routers and hosts (beware the administrative pain)
- An ARP “Firewall” feature is in development at Cisco for initial deployment on our higher-end switches

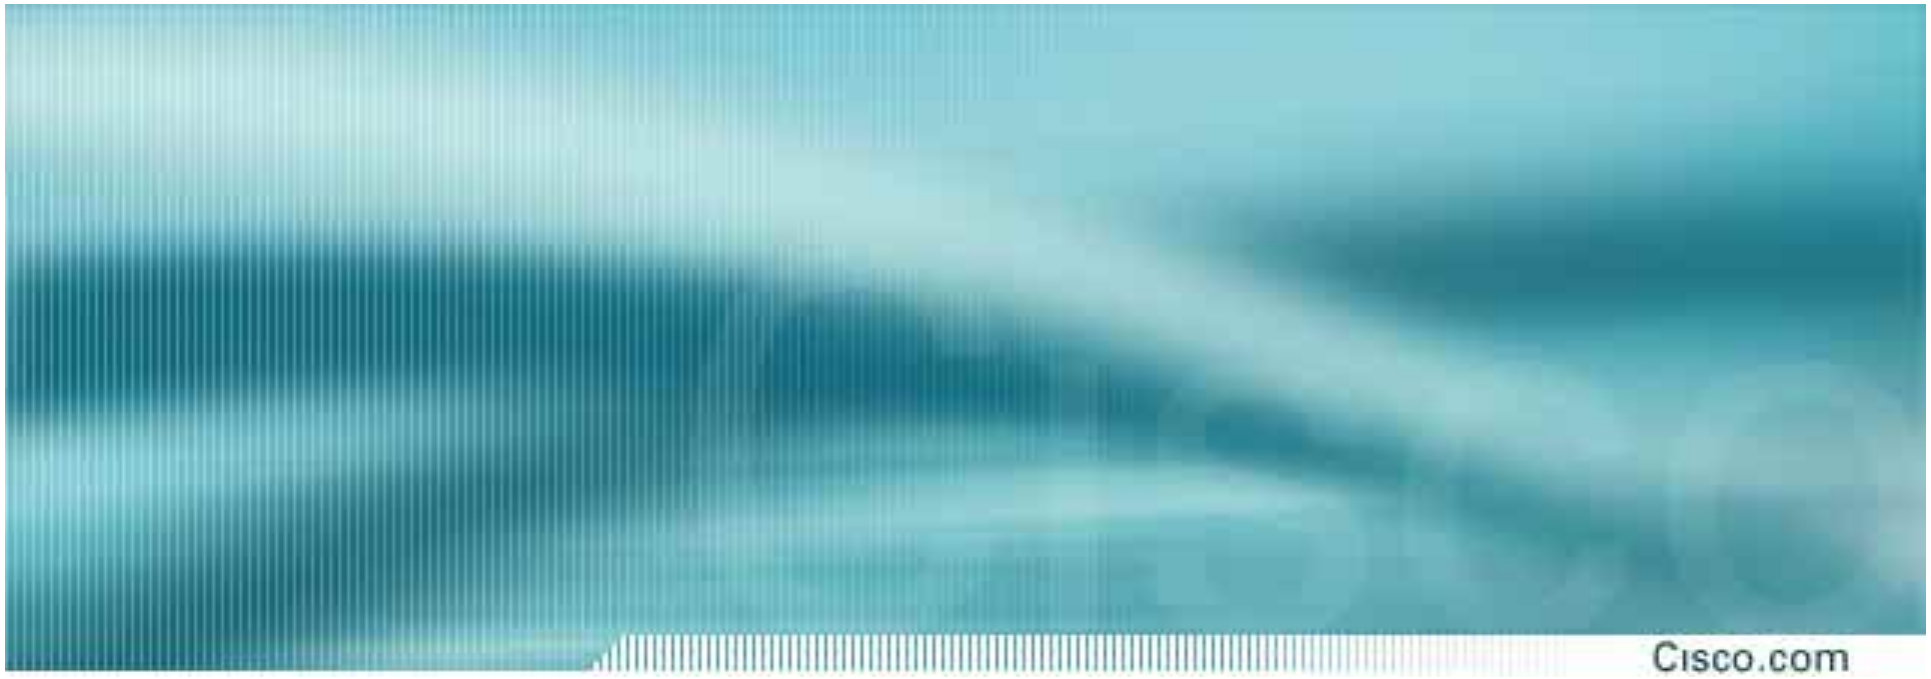

Spanning Tree Attacks

Spanning Tree Basics

STP Purpose: To maintain loop-free topologies in a redundant Layer 2 infrastructure

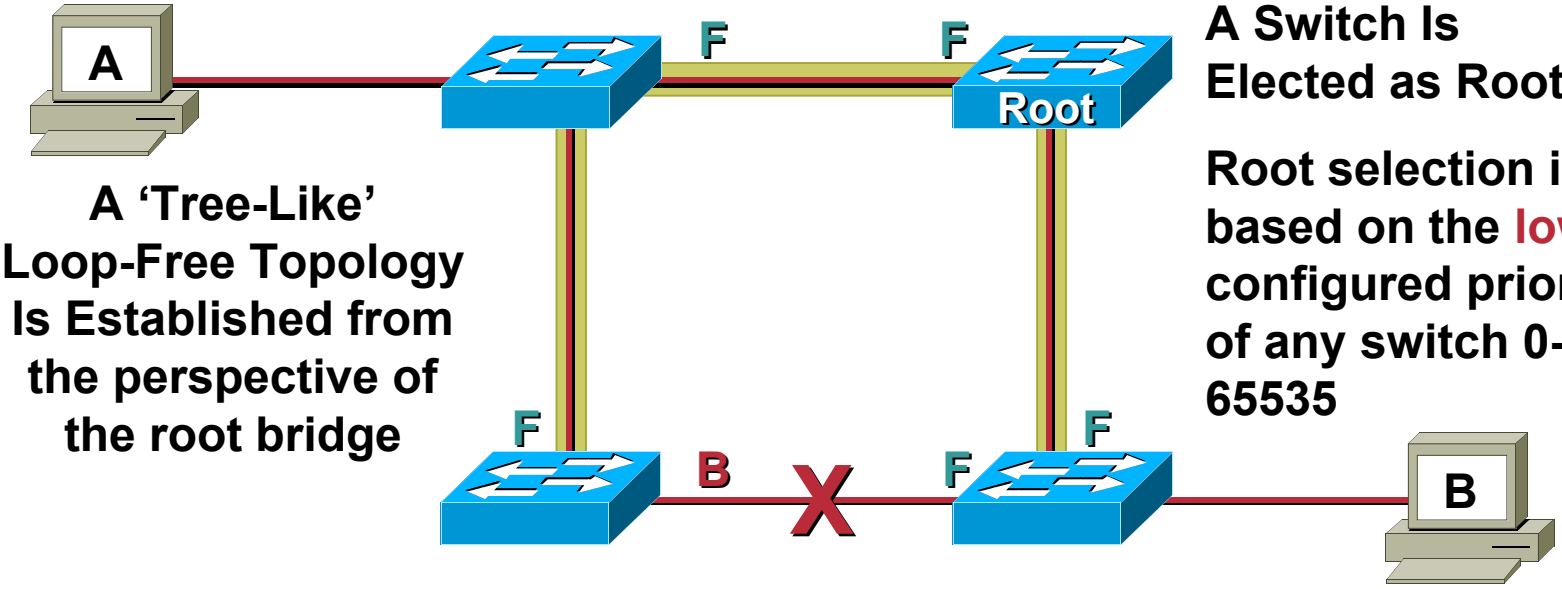

A 'Tree-Like' Loop-Free Topology Is Established from the perspective of the root bridge

A Switch Is Elected as Root

Root selection is based on the lowest configured priority of any switch 0-65535

STP is very simple. Messages are sent using Bridge Protocol Data Units (BPDUs). Basic messages include: configuration, topology change notification/acknowledgment (TCN/TCA); most have no "payload"

Avoiding loops ensures broadcast traffic does not become storms

Spanning Tree Attacks and Methods

- **Standard 802.1d STP takes 30-45 seconds to deal with a failure or Root bridge change (nice DoS)**
  - **Generally only devices affected by the failure notice the issue**
  - **PortFast and UplinkFast can greatly improve this**
- **Sending BPDUs from the attacker can force these changes and create a DoS condition on the network**
- **As a link with macof: the TCN message will result in the CAM table aging all entries in 15 seconds if they do not communicate (the default is 300 seconds)**
- **Testing using brconfig on OpenBSD was easily able to create the DoS condition. Depending on the topology it could also yield more packets available for the attacker**

```
⊟ Frame 25 (64 on wire, 64 captured)
    Arrival Time: Jul 27, 2002 21:02:26.287433000
    Time delta from previous packet: 1.934720000 seconds
    Time relative to first packet: 36.004304000 seconds
    Frame Number: 25
    Packet Length: 64 bytes
    Capture Length: 64 bytes
⊟ IEEE 802.3 Ethernet
    Destination: 01:80:c2:00:00:00 (01:80:c2:00:00:00)
    Source: 00:04:4d:a9:67:c2 (Cisco_a9:67:c2)
    Length: 38
    Trailer: 00000000000000008731E1C5
⊟ Logical-Link Control
    DSAP: Spanning Tree BPDU (0x42)
    IG Bit: Individual
    SSAP: Spanning Tree BPDU (0x42)
    CR Bit: Command
  ⊟ Control field: U, func = UI (0x03)
      000. 00.. = Unnumbered Information
      .... ..11 = Unnumbered frame
⊟ Spanning Tree Protocol
    Protocol Identifier: Spanning Tree Protocol (0x0000)
    Protocol Version Identifier: 0
    BPDU Type: Configuration (0x00)
  ⊟ BPDU flags: 0x00
      0... .... = Topology Change Acknowledgment: No
      .... ...0 = Topology Change: No
    Root Identifier: 32768 / 00:04:4d:a9:67:c0
    Root Path Cost: 0
    Bridge Identifier: 32768 / 00:04:4d:a9:67:c0
    Port identifier: 0x800e
    Message Age: 0
    Max Age: 20
    Hello Time: 2
    Forward Delay: 15
```

Spanning Tree Attack Example 1/2

- **Send BPDU messages to become root bridge**

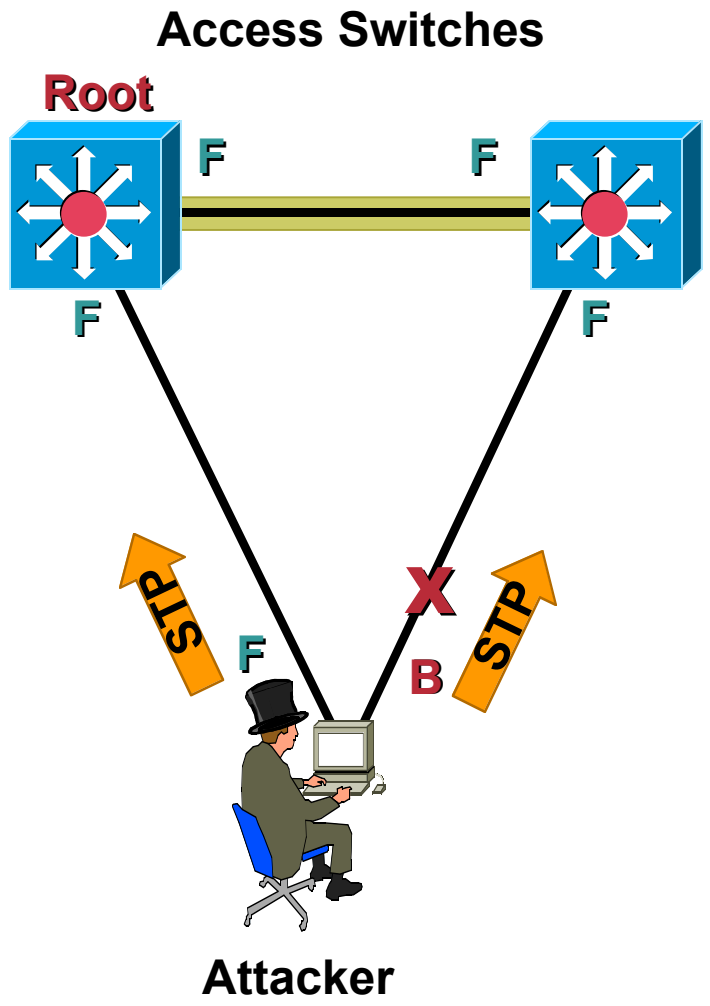

Spanning Tree Attack Example 2/2

- **Send BPDU messages to become root bridge**
  - **The attacker then sees frames he shouldn't**
    - **MITM, DoS, etc. all possible**
    - **Any attack is very sensitive to the original topology, trunking, PVST, etc.**
    - **Although STP takes link speed into consideration, it is always done from the perspective of the root bridge. Taking a Gb backbone to half-duplex 10 Mb was verified**
    - **Requires attacker is dual homed to two different switches (with a hub, it can be done with just one interface on the attacking host)**

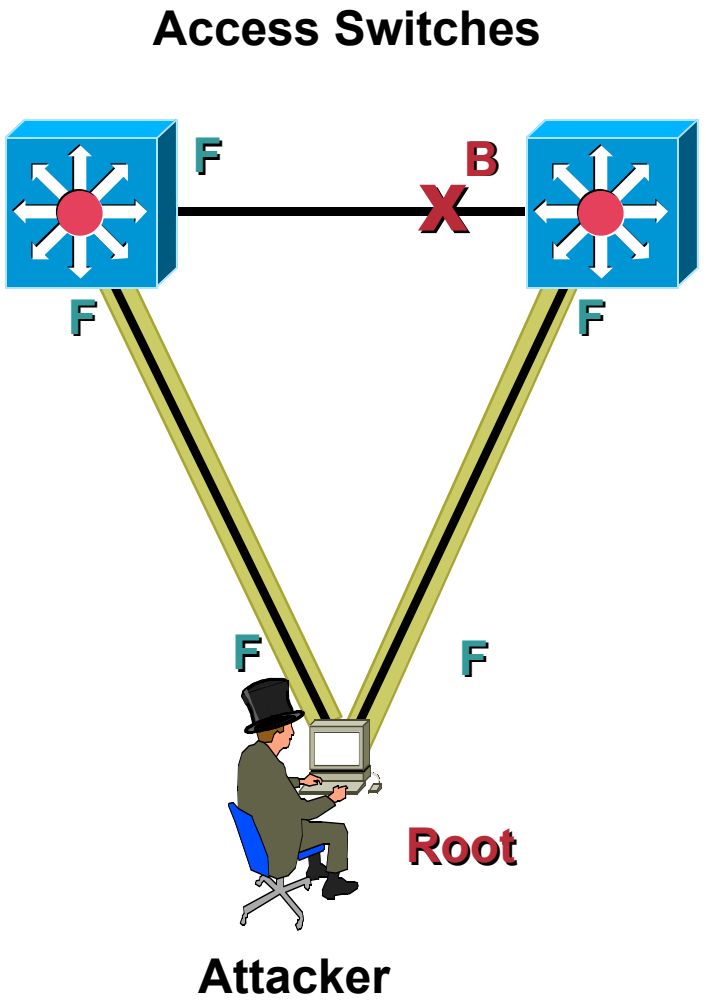

Applied Knowledge: Summary Attack

- Goal: see traffic on the backbone but interesting hosts have static ARP entries and are very chatty (macof will likely never steal their CAM entry)
- Step 1: MAC flood access switch
- Step 2: Run bridging software (brconfig) on attacking host; advertise as a priority zero bridge

  Attacker becomes root bridge

  Spanning Tree recalculates

  GE backbone becomes FE ☹

  CAM table on access switch is full (from macof); there is no room at the inn for the chatty servers. Traffic is flooded

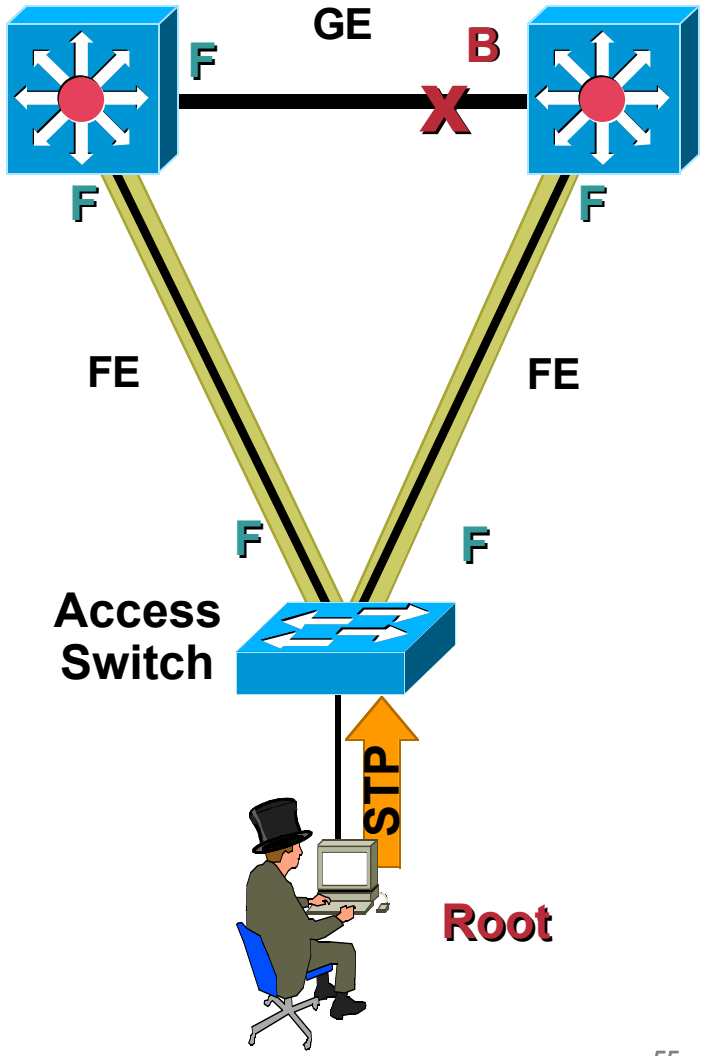

STP Attack Mitigation

- Don't disable STP, introducing a loop would become another attack
- BPDU Guard
  - Disables ports using portfast upon detection of a BPDU message on the port
  - Globally enabled on all ports running portfast
  - Available in CatOS 5.4.1 for Cat 2K, 4K, 5K, and 6K; 12.0XE for native IOS 6K; 12.1(8a)EW for 4K Sup III; 12.1(4)EA1 for 3550; 12.1(6)EA2 for 2950

```
CatOS> (enable)set spantree portfast bpdu-guard enable

IOS(config)#spanning-tree portfast bpduguard
```

- Root Guard
  - Disables ports who would become the root bridge due to their BPDU advertisement
  - Configured on a per port basis
  - Available in CatOS 6.1.1 for Cat 29XX, 4K, 5K, and 6K; 12.0(7) XE for native IOS 6K, 12.1(8a)EW for 4K Sup III; 29/3500XL in 12.0(5)XU; 3550 in 12.1(4)EA1; 2950 in 12.1(6)EA2

```
CatOS> (enable) set spantree guard root 1/1

IOS(config)#spanning-tree guard root (or rootguard)
```

VLAN Trunking Protocol (VTP)

- Used to distribute VLAN configuration among switches
- VTP is used only over trunk ports
- VTP can cause more problems than it solves, consider if it is needed
- If needed, use the VTP MD5 digest:

```
CatOS> (enable) set vtp [domain domain_name] [mode
{client | server | transparent | off}] [passwd
passwd][pruning {enable | disable}] [v2 {enable |
disable}]
IOS(config)#vtp password password-value
```

DST MAC	0100.0ccc.cccc
SNAP Proto	0x2003

Potential VTP Attacks

- After becoming a trunk port, an attacker could send VTP messages as a server with no VLANs configured. All VLANs would be deleted across the entire VTP domain
- Disabling VTP:

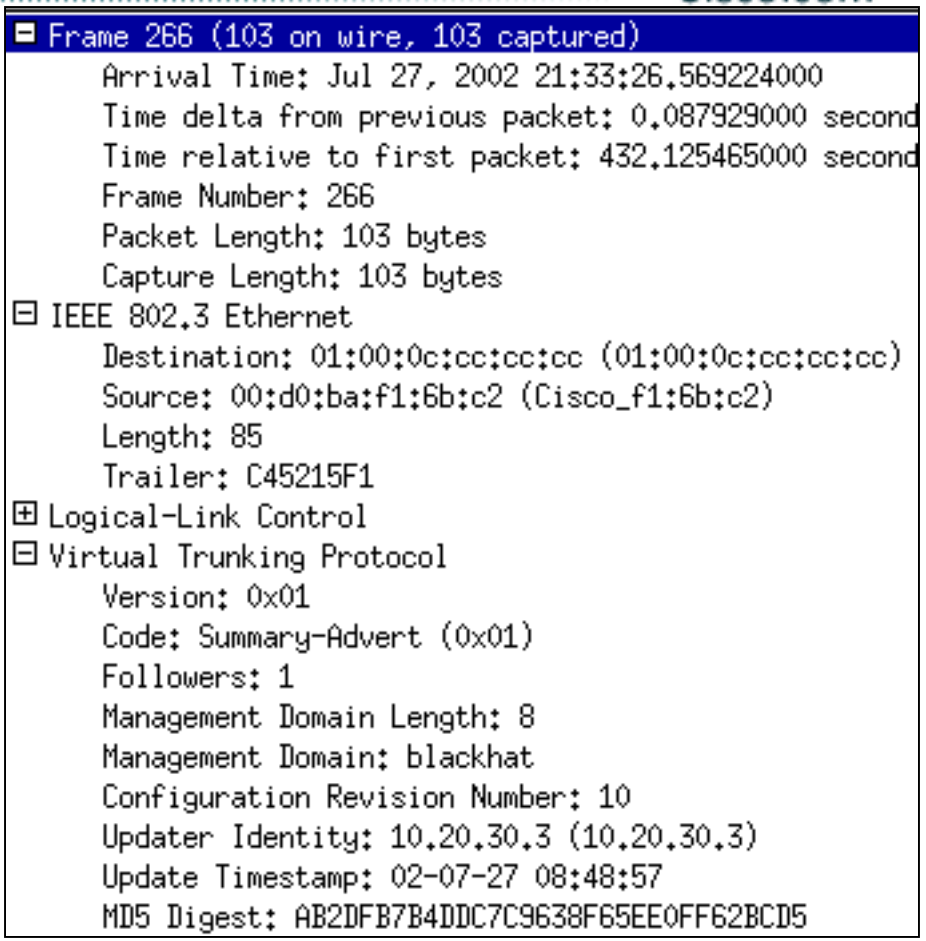


```
CatOS> (enable) set vtp mode transparent | off
IOS(config)#vtp mode transparent
```

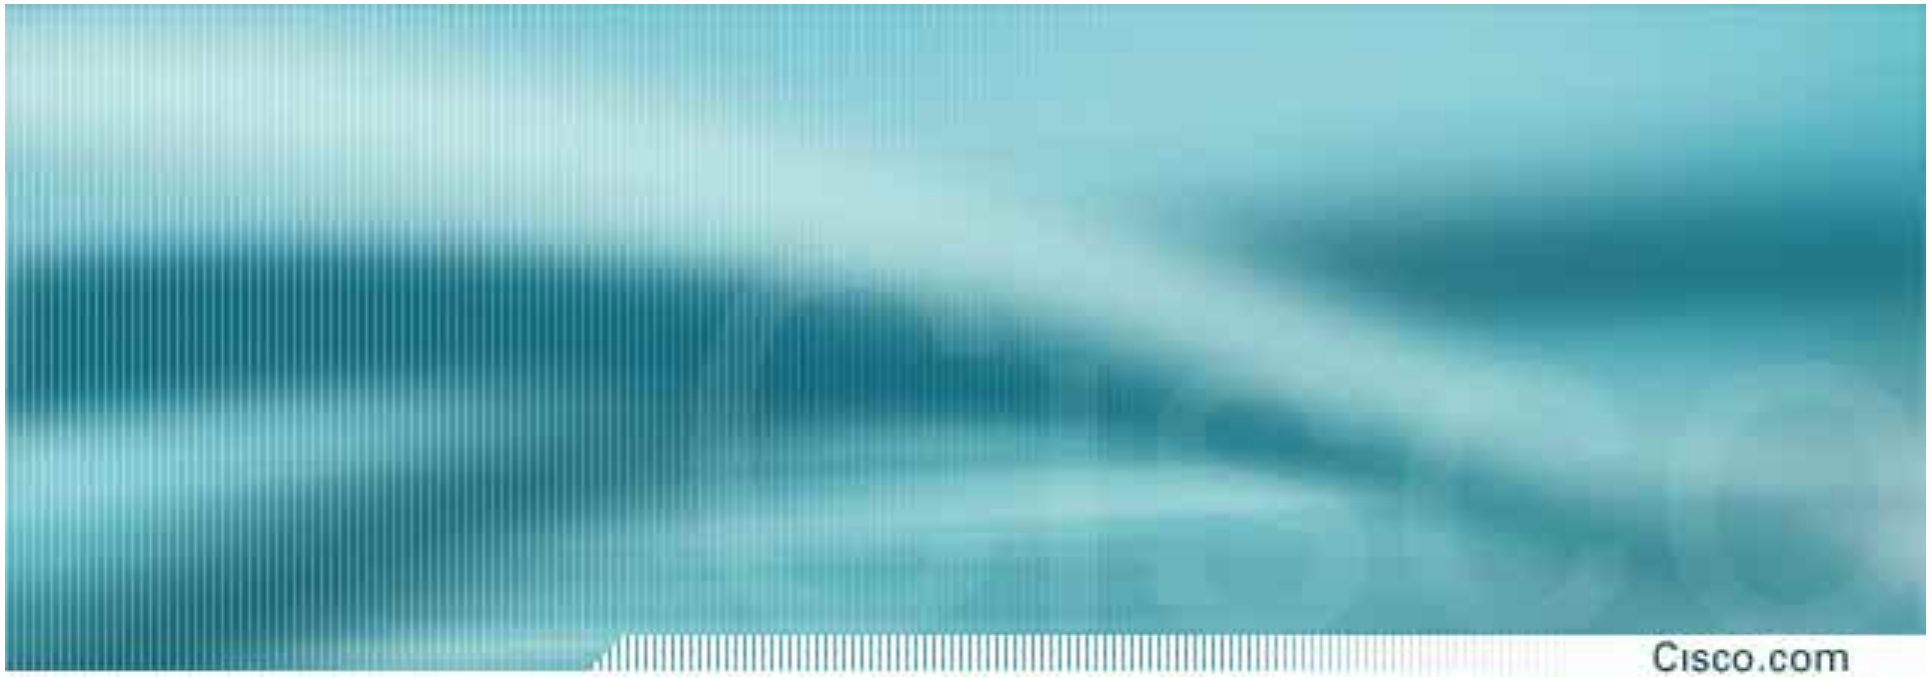

Layer 2 Port Authentication

Dynamic VLAN Access Ports

Cisco.com

- VLAN assignment based on MAC address or HTTP Auth (URT) is possible with a VLAN Management Policy Server (VMPS)
- Requires VLAN to MAC database which is downloaded via TFTP to the VMPS server
- VMPS uses VLAN Query Protocol (VQP) which is unauthenticated and runs over UDP
- Can restrict certain VLANs to certain physical ports
- During access violation, switch can send either an "access denied" response or shutdown the port (depends on configuration)
- Server and client

  Available in Cat 29XX, 4K, 5K, and 6K in CatOS 5.2
- Client only

  Available in 3550 and 2950 in 12.1(4)EA1; 29/3500XL in 11.2(8)SA4

VMPS Architecture

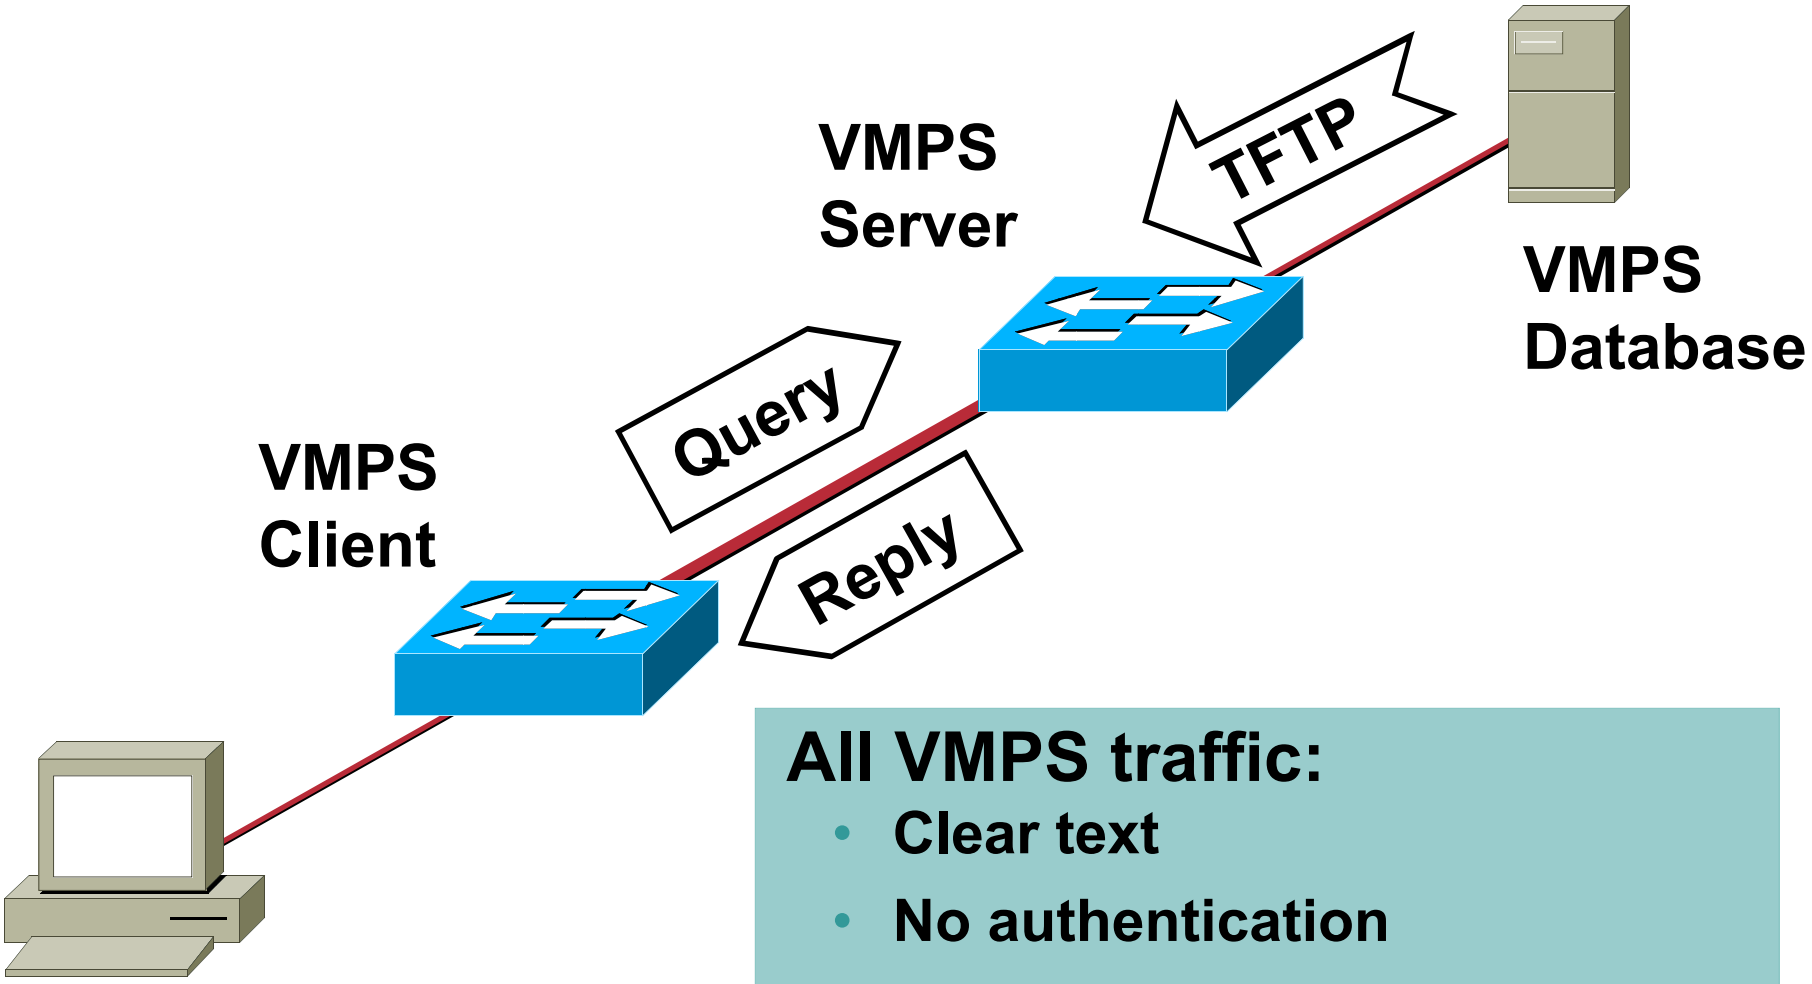

All VMPS traffic:

- Clear text
- No authentication
- UDP based (spoofing trivial)

VMPS/VQP Attacks

- No public domain tools today (Ethereal doesn't even decode)
- VQP/VMPS not frequently used due to administrative burden
- Possible attacks include DoS (prevent login) or Impersonation (Join an unauthorized VLAN)

VQP Query

```
⊞ Ethernet II
⊞ Internet Protocol, Src Addr: 10.20.30.2 (10.20.30.2), Dst Addr: 10.20.30.3 (10.20.30.3)
⊟ User Datagram Protocol, Src Port: 4727 (4727), Dst Port: 1589 (1589)
     Source port: 4727 (4727)
     Destination port: 1589 (1589)
     Length: 142
     Checksum: 0x5e79 (correct)
  Data (134 bytes)
```

```
0000  00 10 7b f7 ae ff 00 04  4d a9 67 c0 08 00 45 00   ..{..... M.g...E.
0010  00 a2 00 00 00 00 ff 11  6b 1e 0a 14 1e 02 0a 14   ........ k.......
0020  1e 03 12 77 06 35 00 8e  5e 79 01 01 00 06 00 00   ...w.5.. ^y......
0030  00 01 00 00 0c 01 00 04  0a 14 1e 02 00 00 0c 02   ........ ........
0040  00 05 46 61 30 2f 38 00  00 0c 03 00 08 2d 2d 4e   ..Fa0/8. .....--N
0050  4f 4e 45 2d 2d 00 00 0c  04 00 08 62 6c 61 63 6b   ONE--... ...black
0060  68 61 74 00 00 0c 07 00  01 00 00 00 0c 05 00 40   hat..... .......@
0070  ff ff ff ff ff ff 00 80  c7 45 5f 9d 08 00 45 00   ........ .E_...E.
0080  00 4e e5 82 00 00 80 11  03 ec 0a 14 1e 0a 0a 14   .N...... ........
0090  1e ff 00 89 00 89 00 3a  06 2b 00 8e 01 10 00 01   .......: .+......
00a0  00 00 00 00 00 00 20 46  48 46 48 46 48 43 4f 45   ...... F HFHFHCOE
00b0  7d a5 03 38                                       }..8
```

VMPS/VQP Attack Mitigation

- **Consider sending VQP messages Out-of-Band (OOB)**
- **If you have the administrative resources to deploy VMPS, you probably have the resources to closely monitor its security**

VQP Response

```
⊞ Ethernet II
⊞ Internet Protocol, Src Addr: 10.20.30.3 (10.20.30.3), Dst Addr: 10.20.30.2 (10.20.30.2)
⊟ User Datagram Protocol, Src Port: 1589 (1589), Dst Port: 4727 (4727)
     Source port: 1589 (1589)
     Destination port: 4727 (4727)
     Length: 42
     Checksum: 0x1dfd (correct)
  Data (34 bytes)
```

```
0000  00 04 4d a9 67 c0 00 10  7b f7 ae ff 08 00 45 00   ..M.g... {.....E.
0010  00 3e 07 07 00 00 1e 11  45 7c 0a 14 1e 03 0a 14   .>...... E|......
0020  1e 02 06 35 12 77 00 2a  1d fd 01 02 00 02 00 00   ...5.w.* ........
0030  00 01 00 00 0c 03 00 08  76 6c 61 6e 30 30 30 38   ........ vlan0008
0040  00 00 0c 08 00 06 00 80  c7 45 5f 9d 52 09 75 e3   ........ .E_.R.u.
```

802.1x/EAP Switch Authentication

Cisco.com

- 802.1x and EAP (Extensible Authentication Protocol) can authenticate a device before allowing access to a switch and can assign a VLAN after authentication
  - EAP allows different authentication types to use the same format (TLS, MD5, OTP)
- Works between the supplicant (client) and the authenticator (network device)
- Maintains backend communication to an authentication (RADIUS) server
- The authenticator (switch) becomes the middleman for relaying EAP received in 802.1x packets to an authentication server by using RADIUS to carry the EAP information
- Available on Cat 2900,4K,6K in CatOS 6.2; Cat 3550 in 12.1(4)EA1; Cat 2950 in 12.1(6)EA2

802.1X Port Authentication

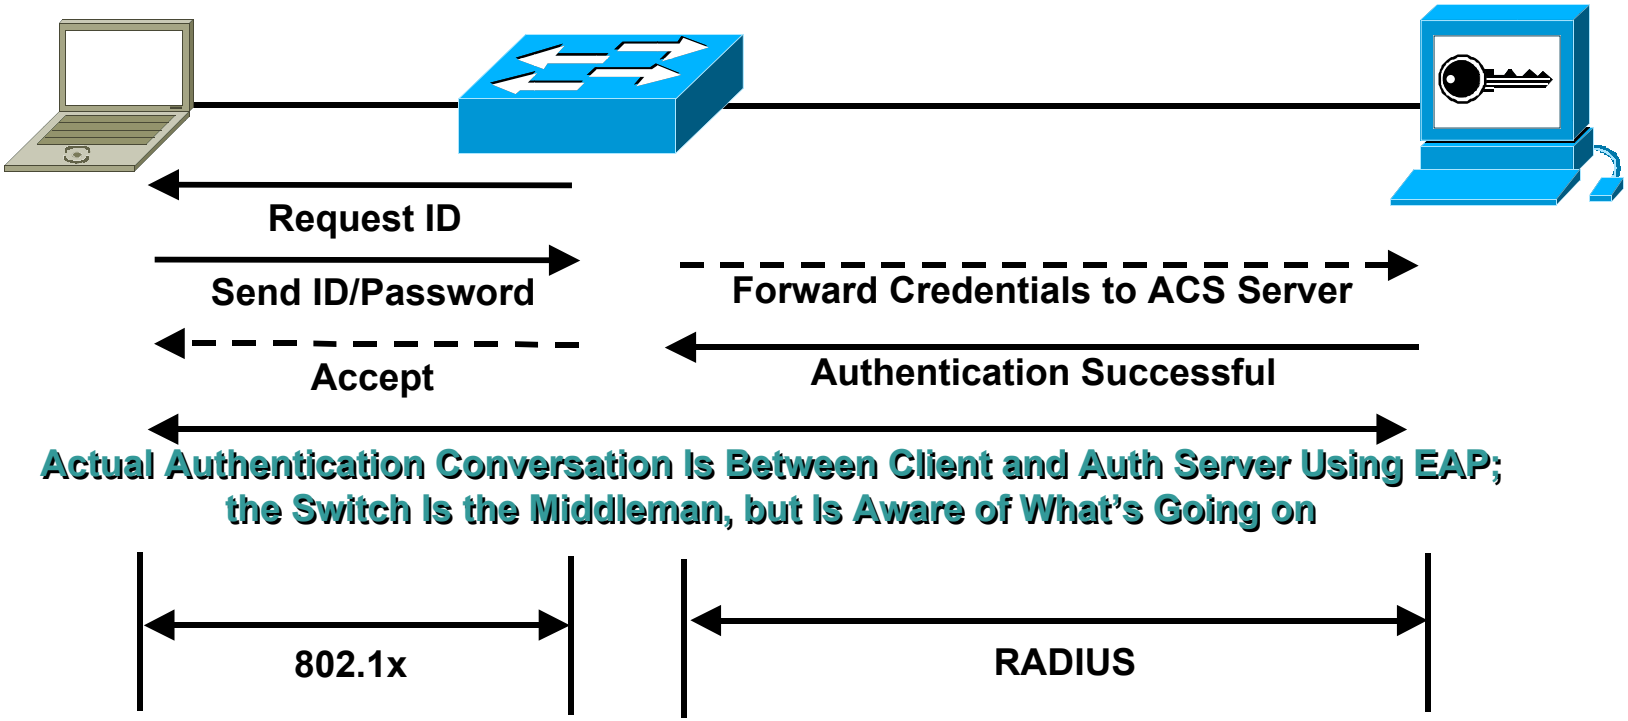

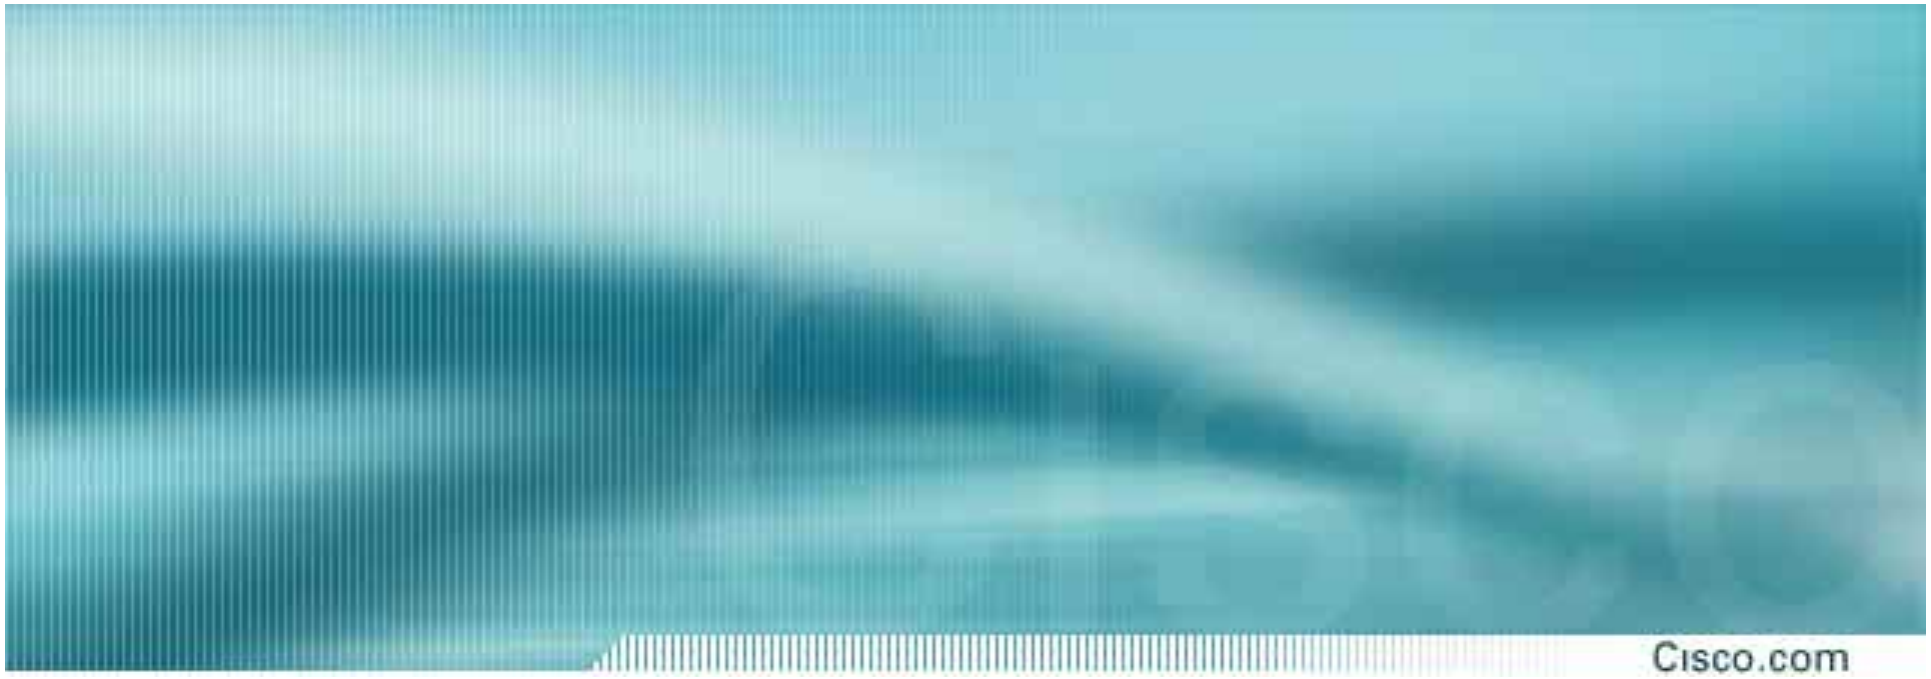

Other Attacks

Cisco Discovery Protocol (CDP)


```
⊟ Capabilities: 0x0000000a
      .... .... .... .... .... .... .... ...0 = Doesn't perform level 3 routing
      .... .... .... .... .... .... .... ..1. = Performs level 2 transparent bridging
      .... .... .... .... .... .... .... .0.. = Doesn't perform level 2 source-route bridging
      .... .... .... .... .... .... .... 1... = Performs level 2 switching
      .... .... .... .... .... .... ...0 .... = Doesn't send or receive packets for network-layer
      .... .... .... .... .... .... ..0. .... = Forwards IGMP Report packets on nonrouter ports
      .... .... .... .... .... .... .0.. .... = Doesn't provide level 1 functionality
⊟ Software Version
      Type: Software version (0x0005)
      Length: 108
      Software Version: WS-C4003 Software, Version McpSW: 7.3(1.0) NmpSW: 7.3(1)
                        Copyright (c) 1995-2002 by Cisco Systems, Inc.
⊟ Platform: WS-C4003
      Type: Platform (0x0006)
      Length: 12
      Platform: WS-C4003
⊞ VTP Management Domain:
⊟ Native VLAN: 1
      Type: Native VLAN (0x000a)
      Length: 6
      Native VLAN: 1
⊞ Duplex: Half
```

DST MAC	0100.0ccc.cccc
SNAP Proto	0x2000

- Runs at Layer 2 and allows Cisco devices to chat with one another
- Can be used to learn sensible information about the CDP sender (IP address, software version, router model …)
- CDP is in the clear and unauthenticated
- Consider disabling CDP, or being very selective in its use in security sensitive environments (backbone vs. user port may be a good distinction)
- Note: there was a reason Cisco developed CDP, some Cisco apps make use of it!

```
CatOS> (enable) set cdp disable
<mod>/<port> | all

IOS(config)#no cdp run

IOS(config-if)#no cdp enable
```

CDP Attacks

- Besides the information gathering benefit CDP offers an attacker, there was a vulnerability in CDP that allowed Cisco devices to run out of memory and potentially crash if you sent it tons of bogus CDP packets
- If you need to run CDP, be sure to use IOS code with minimum version numbers: 12.2(3.6)B, 12.2(4.1)S, 12.2(3.6)PB, 12.2(3.6)T, 12.1(10.1), 12.2(3.6) or CatOS code 6.3, 5.5, or 7.1 and later
- Problem was due to improper memory allocation for the CDP process (basically there was no upper limit)
- Discovered by FX @ Phenolit
- For more information:

  http://www.cisco.com/warp/public/707/cdp_issue.shtml

  http://www.kb.cert.org/vuls/id/139491

DHCP Starvation Attacks

- Anyplace where macof works, you can DoS a network by requesting all of the available DHCP addresses
- With or without the DoS, an attacker could use a rogue DHCP server to provide addresses to clients
- Since DHCP responses include DNS servers and default gateway entries, guess where the attacker would point these unsuspecting users? ☺
- All the MITM attacks are now possible

DHCP Starvation Attack Mitigation

- **Same techniques that mitigate CAM flooding, can mitigate DHCP starvation but not the Rogue DHCP server (from the DHCP RFC 2131):**

  "The client collects DHCPOFFER messages over a period of time, selects one DHCPOFFER message from the (possibly many) incoming DHCPOFFER messages (e.g., the first DHCPOFFER message or the DHCPOFFER message from the previously used server) and extracts the server address from the 'server identifier' option in the DHCPOFFER message. The time over which the client collects messages and the mechanism used to select one DHCPOFFER are implementation dependent."

- **RFC 3118 "Authentication for DHCP Messages" will help, but has yet to be implemented**
- **Consider using multiple DHCP servers for the different security zones of your network**
- **DHCP Option 82 on the 3550 can help: http://www.cisco.com/univercd/cc/td/doc/product/lan/c3550/1219ea1/3550scg/swdhcp82.htm**
- **Cisco is developing a DHCP "firewall" for initial implementation in our higher-end switches**

Private VLAN Attacks 1/2

Cisco.com

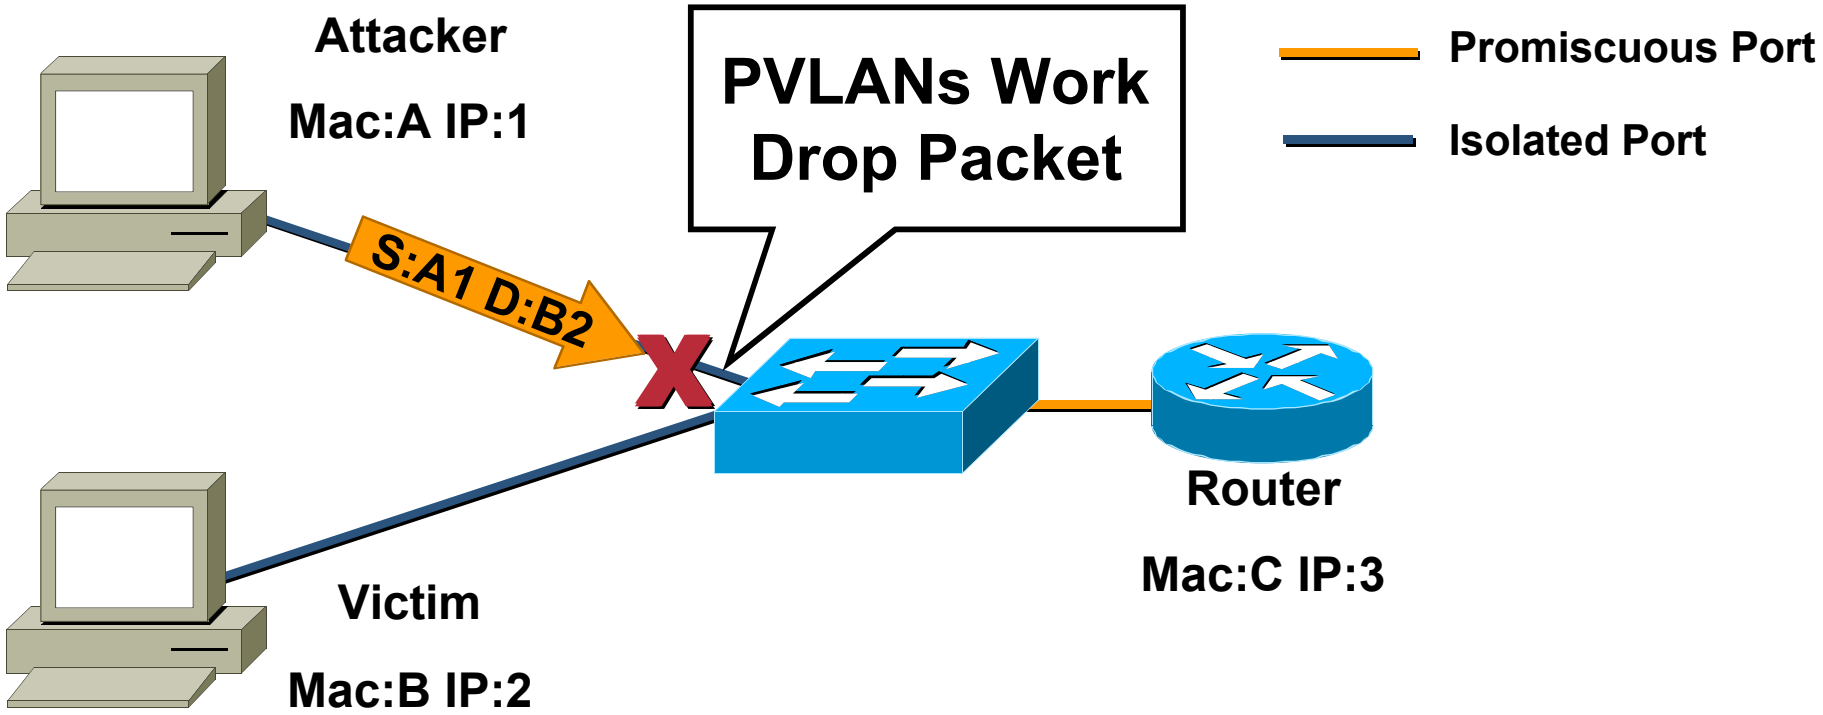

Private VLAN Attacks 2/2

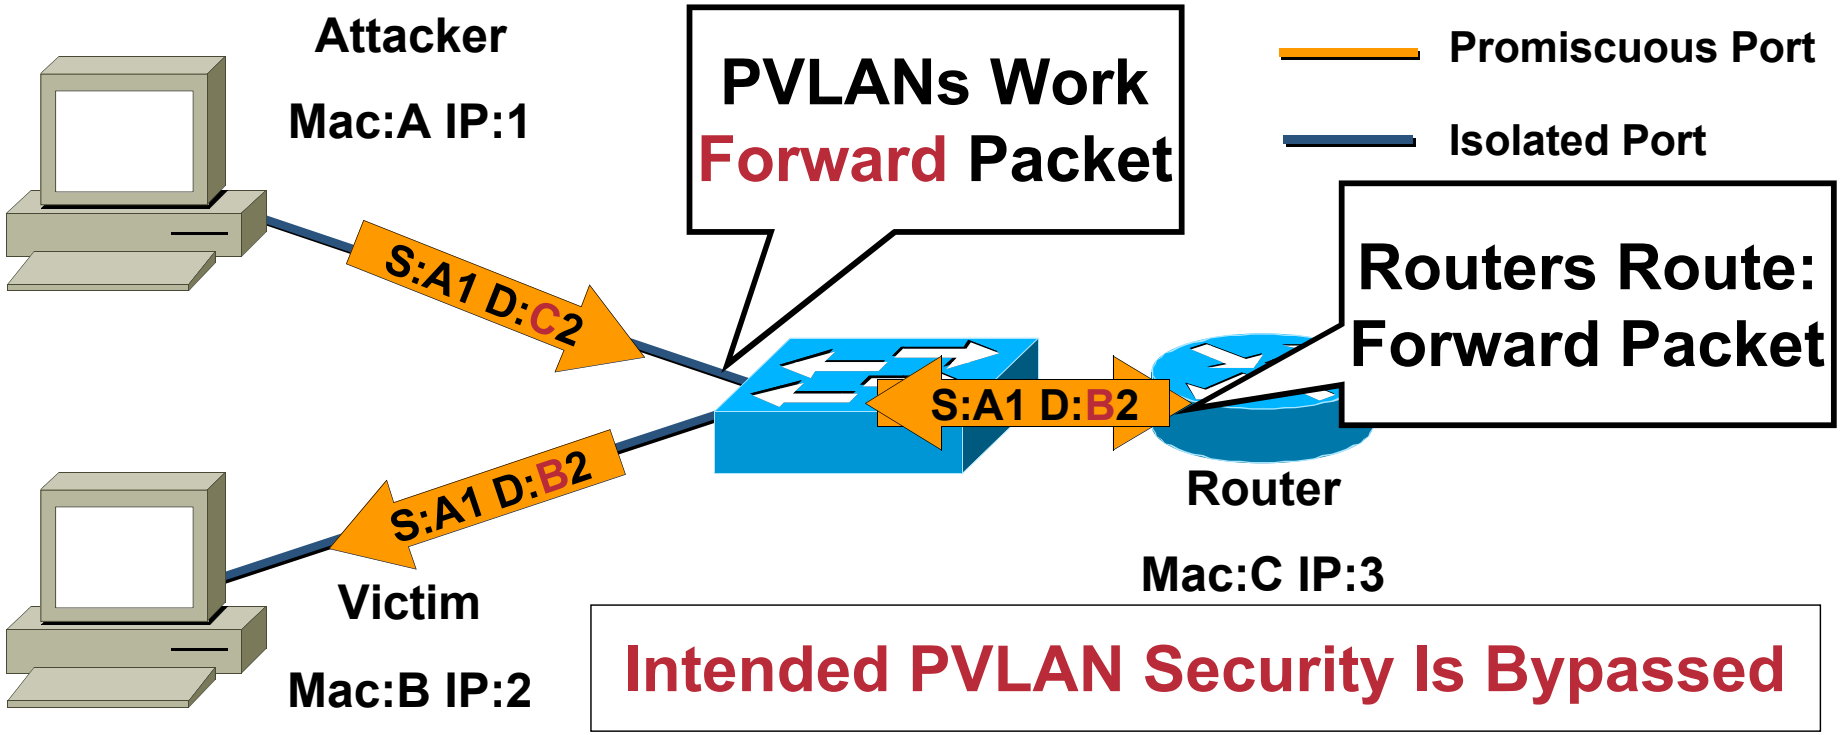

- Only allows unidirectional traffic (Victim will ARP for A and fail)
- If both hosts were compromised, setting static ARP entries for each other via the router will allow bi-directional traffic
- Most firewalls will not forward the packet like a router
- Note: this is not a PVLAN vulnerability as it enforced the rules!

PVLAN Attack Mitigation

- **Setup ACL on ingress router port:**

```
IOS(config)#access-l 101 deny ip
localsubnet lsubmask localsubnet lsubmask
log

IOS(config)#access-l 101 permit ip any any

IOS(config-if)#ip access-group 101 in
```

- **All known PVLAN exploits will now fail**
- **VLAN ACL (VACL) could also be used**

Multicast Brute-Force Failover Analysis

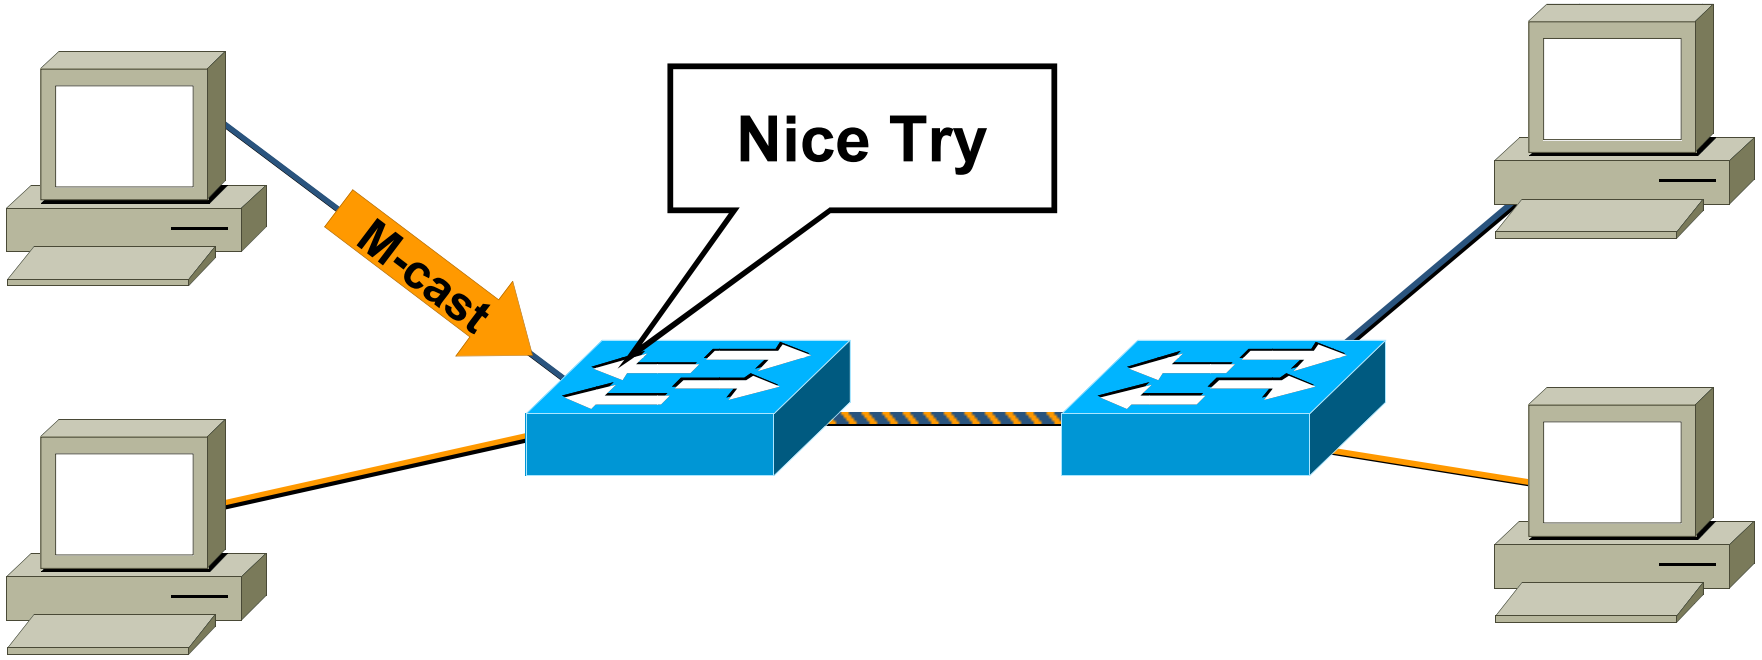

- **Send random Ethernet multicast frames to a switch interface attempting to get frames to another VLAN**

Random Frame Stress Attack

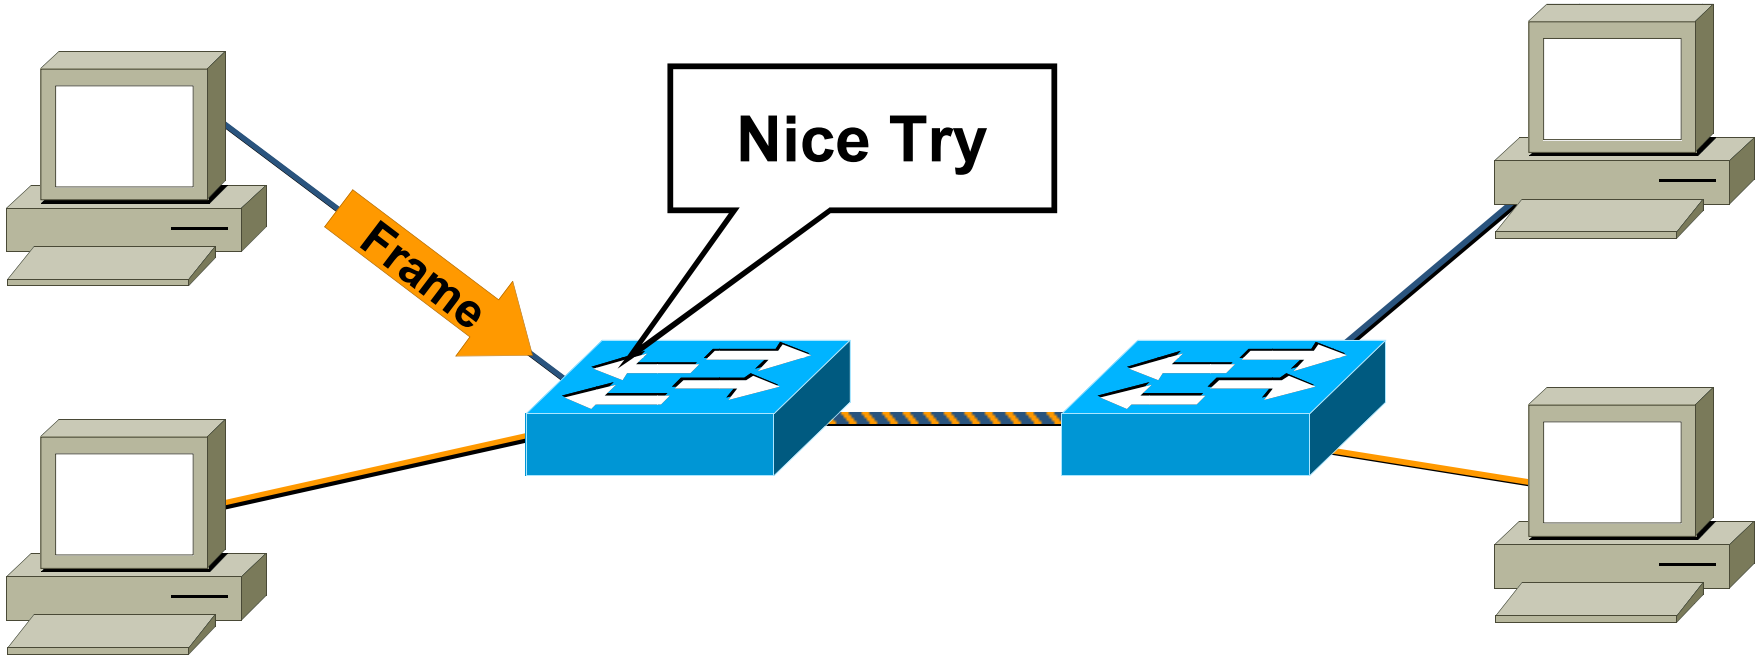

- **Send random frames to a switch interface attempting to get frames to another VLAN**

IP Telephony Considerations

- **Most IP Telephony deployments use a distinct VLAN for voice vs. data traffic**

  Done because of QoS and security considerations

  Voice VLAN is called an "auxiliary" VLAN and is set on the phone via a CDP message (trunking can still be disabled)

Tcpdump Output

```
04:16:06.652765 802.1Q vid 987 pri 0 1:0:c:cc:cc:cd > 0:8:e3:cf:1a:dd sap aa ui/C len=39
04:16:07.095781 0:8:e3:cf:1a:dd > 1:0:c:cc:cc:cd sap aa ui/C len=39
```

All mentioned attack mitigation features work fine except PVLANs and 802.1X which do not yet support aux VLANs

IP Telephony currently does not support confidentiality. Use the techniques discussed in this presentation to mitigate the effects of tools like Vomit. http://vomit.xtdnet.nl

Switch Management

Cisco.com

- **Management can be your weakest link**
  - **All the great mitigation techniques we talked about aren't worth much if the attacker telnets into your switch and disables them**
- **Most of the network management protocols we know and love are insecure (syslog, SNMP, TFTP, Telnet, FTP, etc.)**
- **Consider secure variants of these protocols as they become available (SSH, SCP, SSL, OTP etc.), where impossible, consider out of band (OOB) management**
  - **Put the management VLAN into a dedicated non-standard VLAN where nothing but management traffic resides**
  - **Consider physically back-hauling this interface to your management network**
- **When OOB management is not possible, at least limit access to the management protocols using the "set ip permit" lists on the management protocols**
- **SSH is available on Cat 6K with CatOS 6.1 and Cat 4K/29XXG with CatOS 6.3**

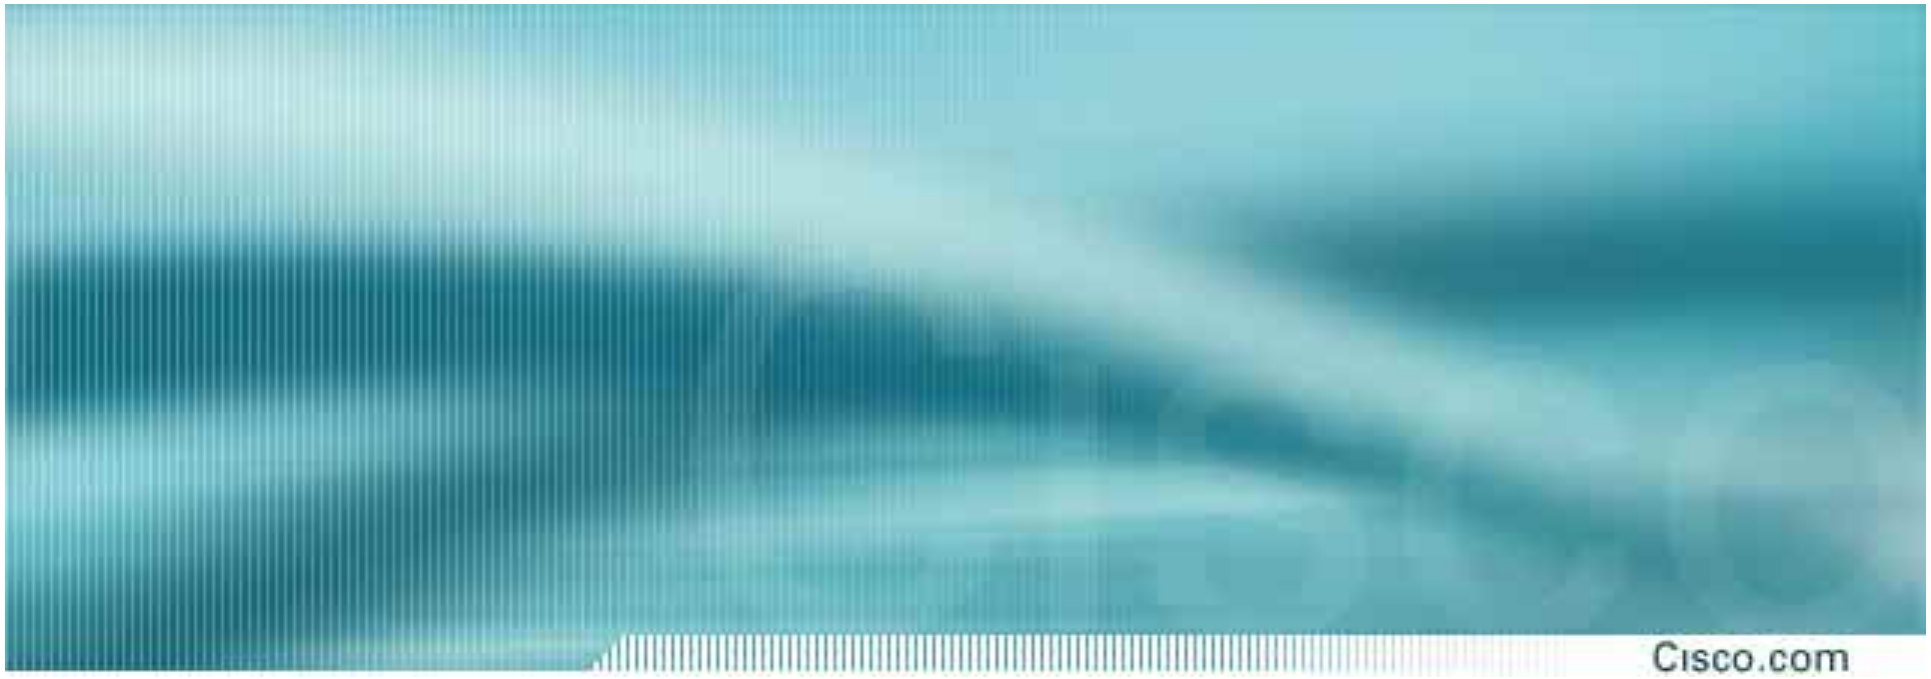

Summary and Case Study

Layer 2 Security Best Practices 1/2

Cisco.com

- Manage switches in as secure a manner as possible (SSH, OOB, permit lists, etc.)
- Always use a dedicated VLAN ID for all trunk ports
- Be paranoid: do not use VLAN 1 for anything
- Set all user ports to non trunking
- Deploy port-security where possible for user ports
- Selectively use SNMP and treat community strings like root passwords
- Have a plan for the ARP security issues in your network

Layer 2 Security Best Practices 2/2

- Enable STP attack mitigation (BPDU Guard, Root Guard)
- Use private VLANs where appropriate to further divide L2 networks
- Use MD5 authentication for VTP
- Use CDP only where necessary
- Disable all unused ports and put them in an unused VLAN
- Consider 802.1X for the future

All of the Preceding Features Are Dependant on Your Own Security Policy

A Relevant Case Study

- Do you have a part of your network that looks like this?

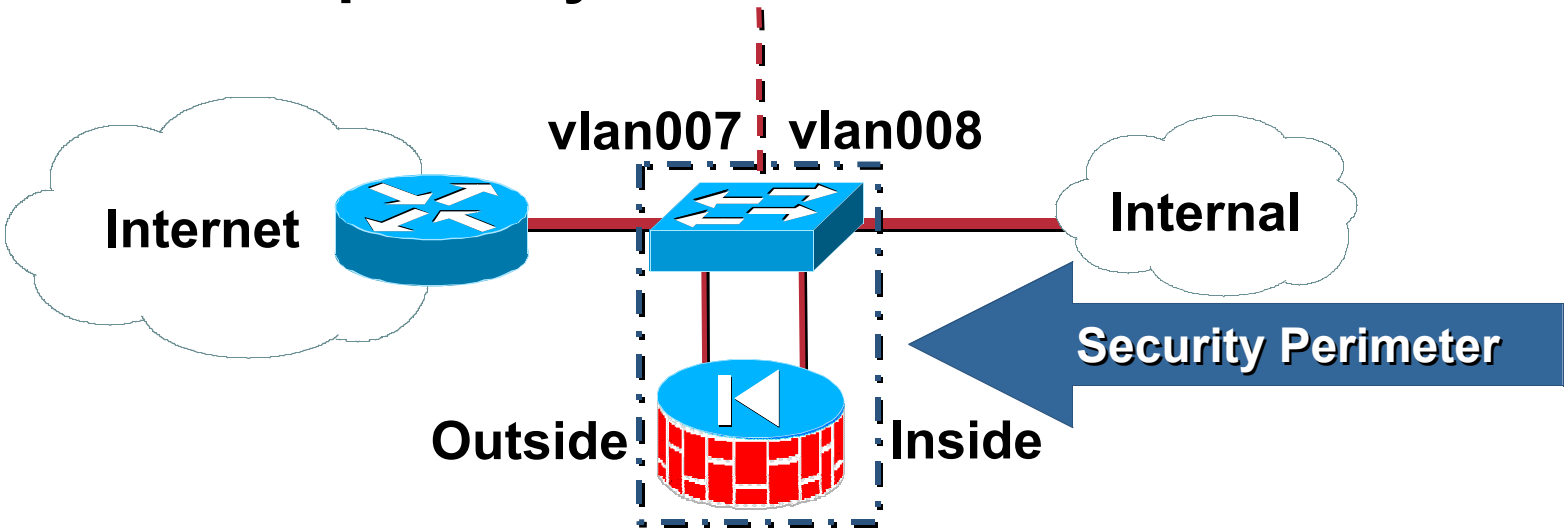

  - While it is technically feasible to make this "secure", consider the ramifications:
    - What happens if the switch is compromised?
    - Does SECOPS control the VLAN settings on the switch? (likely not)
      - This means you now have NETOPS folks taking actions that could adversely affect security
    - Realize your security perimeter now includes the switch

A More Secure Alternative

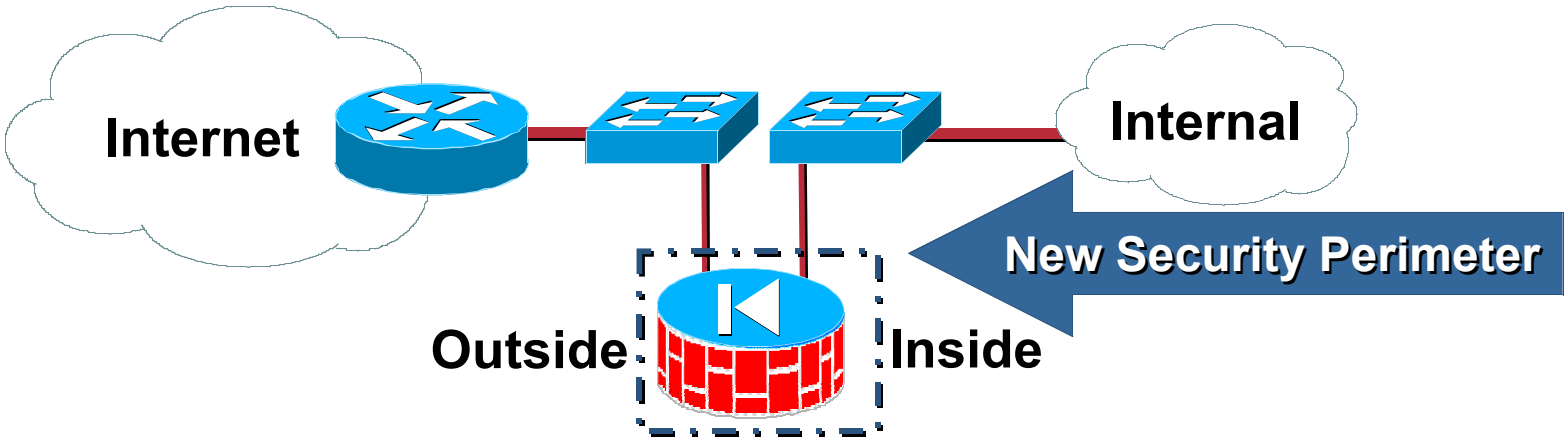

Lessons Learned

- **Carefully consider any time you must count on VLANs to operate in a security role**
  - **If properly configured, our testing did not discover a method of VLAN Hopping using Cisco switches**
  - **Pay close attention to the configuration**
  - **Understand the organizational implications**
- **Evaluate your security policy while considering the other issues raised in this session**
  - **Is there room for improvement?**
  - **What campus risks are acceptable based on your policy?**
- **Deploy, where appropriate, L2 security best practices**

Further Reading

- **SAFE Blueprints**

  http://www.cisco.com/go/safe

- **Improving Security on Cisco Routers**

  http://www.cisco.com/warp/public/707/21.html

- **Links in this presentation:**

  Port security: http://cisco.com/univercd/cc/td/doc/product/lan/cat5000/rel_5_4/config/sec_port.htm

  Switch Control Protocols: http://www.cisco.com/warp/public/473/103.html

  Ethernet Encapsulation Info: http://www.cisco.com/warp/public/105/encheat.html

  SANS VLAN paper (out of date): http://www.sans.org/newlook/resources/IDFAQ/vlan.htm

  Dsniff homepage: http://www.monkey.org/~dugsong/dsniff

  Ettercap homepage: http://ettercap.sourceforge.net/

  PVLAN / VACL Design: http://www.cisco.com/warp/public/473/90.shtml

  PVLAN details: http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/sw_7_1/conf_gd/vlans.htm#xtocid854519

  CDP vulnerability: http://www.cisco.com/warp/public/707/cdp_issue.shtml

  DHCP Option 82: http://www.cisco.com/univercd/cc/td/doc/product/lan/c3550/1219ea1/3550scg/swdhcp82.htm