



**NETWORKERS 2004**

**TROUBLESHOOTING  
CATALYST 6000/6500 SWITCHES  
SESSION RST-3509**

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

1

## Agenda

Cisco.com

- **Hybrid/Native**
- **Redundancy**
- **Unicast Forwarding**
- **Multicast Forwarding**
- **Pot Luck Troubleshooting**

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

2

## Hybrid/Native Hybrid vs. Native

Cisco.com

Two of these switches are in Native...

Two are in Hybrid...



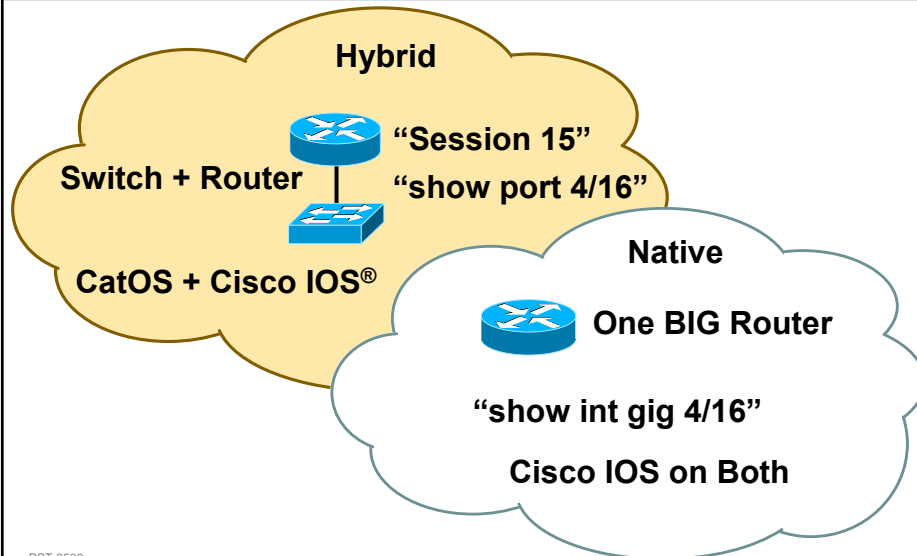
Can you tell me which is which?

RST-3509  
9806\_05\_2004\_c1

3

## Hybrid/Native Hybrid vs. Native

Cisco.com



RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

4

## Hybrid/Native Hybrid: Software Conventions

Cisco.com

### MSFC Type

- MSFC1/2/3—match your hardware
  - c6msfc-dsv-mz.121-20.E2
  - c6msfc2-jsv-mz.121-13.E13
  - c6msfc2-jo3sv-mz.121-13.E13
  - c6msfc3-jsv-mz.122-14.SX2

### Memory (DRAM) Required?

- Depends...
- MSFC min/max 64/128 MB recommend 128MB
- MSFC2 min/max 128/512 MB recommend 256MB+

### What about the MSFC3?

- Good to go w/ 512MB DRAM and 64MB flash

### Boot Image

- MSFC1—use boot image  
c6msfc-boot-mz.121-20.E2
- MSFC2—not needed, but nice to have
- MSFC3—not needed, not available

### Memory (Bootflash) Required?

- File Sizes vary, but generally < 15MB, unless crypto
- MSFC has 16MB, can be tight w/ boot image
- MSFC2 upgrades to 32MB if needed

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

5

## Hybrid/Native Hybrid: Software Conventions

Cisco.com

### Supervisor Type

- Sup1/2/720—match hardware
  - cat6000-supcv.6-3-9.bin
  - cat6000-sup2k8.7-4-3.bin
  - cat6000-sup720k9.8-2-1.bin

### Memory Required (DRAM)

- 5.x—Sup1—64MB
- 6.x—Sup1—64MB  
Sup2—128MB
- 7.x—Sup 1—pre 7.5(1) 64MB  
7.5(1)+ 128M **suggested**  
7.6(4)+ 128MB **required**  
Sup2—128MB
- 8.x—Sup1/2—128MB  
Sup720—512MB

### Features

- cv—CiscoView Web Interface
- k9—Crypto SSH Software

### Memory Required (Bootflash)

- File Sizes vary, but often > 16MB
- Sup1 has 16MB
- Sup2 upgrades to 32MB
- Otherwise boot from slot0:

### What about the Sup720?

- non-issue w/ 512MB DRAM and 64MB flash

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

6

## Hybrid/Native Native: Software Conventions

Cisco.com

### Hardware Type

- Match your hardware

c6sup11-dsv-mz.121-20.E2

c6sup11 = Sup1/MSFC1

c6sup12-is-mz.121-13.E5

c6sup12 = Sup1/MSFC2

s72033-psv-mz.122-17a.SX1

s72033=Sup720/MSFC3/PFC3

### Memory (DRAM) Required?

- Still depends...
- MSFC1/2 on Sup1 min 128 MB
- MSFC2 on Sup2... recommend 256MB+
- Should have same memory in Sup2/DFC as in the MSFC

### Memory (Bootflash) Required?

- MSFC1—still use boot image  
c6msfc-boot-mz.121-20.E2
- Native sup11 images are often > 16MB boot from slot0: or sup-bootflash:

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

7

## Checking Memory in RP/SP/DFC Native

Cisco.com

- How much memory in my RP?

```
Router# show ver  
cisco WS-C6513 (R7000) processor with 458752K/65536K bytes of memory.
```

- How much memory in my SP?

```
Router# remote command switch show ver  
cisco WS-C6513 (R7000) processor with 112640K/18432K bytes of memory.
```

- How much memory in my DFC?

```
Router# remote command module 5 show ver  
cisco WS-C6513 (R7000) processor with 112640K/18432K bytes of memory.
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

8

# Agenda

Cisco.com

- Hybrid/Native
- **Redundancy**
- Unicast Forwarding
- Multicast Forwarding
- Pot Luck Troubleshooting

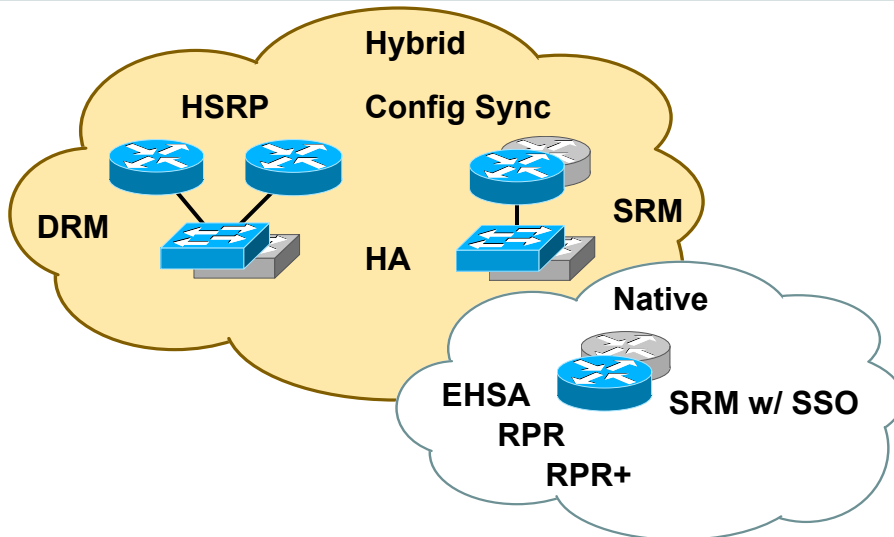
RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

9

# Redundancy Hybrid/Native Options

Cisco.com



RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

10

## Redundancy Hybrid with Dual Supervisors

Cisco.com



- For all Hybrid configurations make sure the supervisor has high-availability **enabled** AND that it is **ON**
- First introduced in 5.4(1)
- If high availability **versioning** is enabled, Sup's code will **not** synchronize

```
CatOS> (enable) show system highavailability
Highavailability: enabled
Highavailability versioning: disabled
Highavailability Operational-status: ON
```


RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

11

## Redundancy Hybrid with Dual MSFCs-DRM

Cisco.com



- Make sure that HSRP is configured and active on both MSFCs

```
MSFC# show standby vlan 2
Vlan2 - Group 0
Local state is Active, priority 110, may preempt
Hellotime 3 holdtime 10
Next hello sent in 00:00:00.628
Hot standby IP address is 172.10.1.254 configured
Active router is local
Standby router is 172.10.1.3 expires in 00:00:09
Standby virtual mac address is 0000.0c07.ac00
2 state changes, last state change 00:00:13
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

12

## Redundancy Hybrid with Dual MSFCs-DRM

Cisco.com

- Know which MSFC is **designated**

```
MSFC# show redundancy
Designated Router: 1 Non-designated Router: 2

Redundancy Status: designated

Config Sync AdminStatus : disabled
Config Sync RuntimeStatus: disabled

Single Router Mode AdminStatus : disabled
Single Router Mode RuntimeStatus: disabled
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

13

## Redundancy Hybrid with Config-Sync-DRM

Cisco.com

- Config-Sync makes the designated MSFC sync its configuration to the non-designated MSFC
- Available in 12.1(3a)E1

```
redundancy
high-availability
config-sync
!
interface Serial4/0/0
ip address 10.1.1.2 255.255.255.0
dsu bandwidth 44210
framing c-bit
!
interface Vlan2
ip address 172.10.1.2 255.255.255.0 alt ip address 172.10.1.3 255.255.255.0
ip pim dense-mode
standby priority 110 preempt alt standby priority 90 preempt
standby ip 172.10.1.1 alt standby ip 172.10.1.1
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

14

## Redundancy Hybrid with Config-Sync-DRM

Cisco.com

- In a dual supervisor chassis utilizing config-sync here is what steady-state should look like

```
CatOS> (enable) show system highavailability
Highavailability: enabled
Highavailability versioning: disabled
Highavailability Operational-status: ON

MSFC# show redundancy
Designated Router: 1 Non-designated Router: 2
Redundancy Status: designated
Config Sync AdminStatus : enabled
Config Sync RuntimeStatus: enabled
Single Router Mode AdminStatus : disabled
Single Router Mode RuntimeStatus: disabled
```

RST-3509  
9806\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

15

## Redundancy Hybrid with Single-Router-Mode-SRM

Cisco.com

- **Single-Router-Mode** makes the **designated** MSFC sync its configuration to the non-designated MSFC and then puts it into standby mode
- Available in 12.1(8a)E4 / 6.3(1)



```
redundancy
high-availability
single-router-mode
!
interface Serial4/0/0
ip address 10.1.1.2 255.255.255.0
dsu bandwidth 44210
framing c-bit
!
interface Vlan2
ip address 172.10.1.2 255.255.255.0
ip pim dense-mode
standby priority 110 preempt
standby ip 172.10.1.1
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

16

## Redundancy Hybrid with SRM

Cisco.com

- Utilizing single-router-mode, here is how steady-state should look

```
CatOS> (enable) show module
Mod Slot Ports  Module-Type           Model                Sub    Status
-----
1   1   2   1000BaseX Supervisor  WS-X6K-SUP1A-2GE  yes   ok
15  1   1   Multilayer Switch Feature WS-F6K-MSFC       no    ok
2   2   2   1000BaseX Supervisor  WS-X6K-SUP1A-2GE  yes   standby
16  2   1   Multilayer Switch Feature WS-F6K-MSFC       no    standby

CatOS> (enable) show system highavailability
Highavailability: enabled
Highavailability versioning: disabled
Highavailability Operational-status: ON

MSFC# show redundancy
Designated Router: 1 Non-designated Router: 2
Redundancy Status: designated
Config Sync AdminStatus : enabled
Config Sync RuntimeStatus: enabled
Single Router Mode AdminStatus : enabled
Single Router Mode RuntimeStatus: enabled
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

17

## Redundancy Hybrid with SRM

Cisco.com

- **Show redundancy** may indicate **enabled** but...
  - 1) Config doesn't get synchronized
  - 2) Failover times are much longer than expected
- Take care to follow the **configuring SRM** procedure on CCO
- Steps often overlooked...
  - 1) **Set system high availability** on the Sups (step 2)
  - 2) Write mem on the DR (step 9)
  - 3) Reload of the Non-DR (step 12)

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

18

## Redundancy Native with RPR+

Cisco.com

- **Route processor redundancy plus** provides for much faster failover as cards **are no longer** reset on failover
- **30 sec or more failover times independent of what cards are in the chassis**
- **Available from 12.1(13)E**

```
redundancy
mode rpr-plus
```



RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

19

## Redundancy Native with RPR+

Cisco.com

- **Utilizing RPR+, here is what steady-state should look like**

```
Native(6k)# show redundancy states
my state = 13 -ACTIVE
peer state = 8 -STANDBY HOT
Mode = Duplex
Unit = Primary
Unit ID = 1

Redundancy Mode (Operational) = Route Processor Redundancy Plus
Redundancy Mode (Configured) = Route Processor Redundancy Plus
Split Mode = Disabled
Manual Swact = Enabled
Communications = Up
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

20

## Redundancy Native SRM with SSO

Cisco.com

- First supported 12.2(17b)SXA
- 1 or more sec failover
- Sup 720 only
- Unicast traffic only
- Layer 2 protocols—link negotiation, VLANs, VTP, DTP, STP, PAgP/LACP, CDP, UDLD, SPAN, Voice VLAN, inline power, IGMP snooping
- Notice the absence of layer 3 routing protocols

```
redundancy
mode sso

Native(6k)# show redundancy states
my state = 13 -ACTIVE
peer state = 8 -STANDBY HOT
Mode = Duplex
Unit = Primary
Unit ID = 5
Redundancy Mode (Operational) = Stateful Switchover
Redundancy Mode (Configured) = Stateful Switchover
```



RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

21

## Agenda

Cisco.com

- Hybrid/Native
- Redundancy
- **Unicast Forwarding**
- Multicast Forwarding
- Pot Luck Troubleshooting


RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

22

## Unicast Forwarding Connectivity Loss: MLS

Cisco.com



- Connectivity problems b/w these 2 routers
- What can you check on the 6k?

RST-3509  
9806\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

23

## Unicast Forwarding Hybrid: Sup1: Checking the RP

Cisco.com

- Check the software **routing tables** and **ARP caches**
- Do this for both **source** and **destination**



```
MSFC# show ip route 10.1.1.5
Routing entry for 10.1.1.0/24
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Routing Descriptor Blocks:
    * directly connected, via Vlan1
      Route metric is 0, traffic share count is 1

MSFC# show ip arp 10.1.1.5
Protocol Address      Age (min)  Hardware Addr  Type   Interface
Internet 10.1.1.5          18        0000.0c5d.143c  ARPA   Vlan1
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

24

## Unicast Forwarding Hybrid: Sup1: Checking the RP

Cisco.com

- Still in the MSFC...
- Verify that the **CEF** and **adjacency** table match the **routing** table and **ARP** cache; Do this for both source and destination
- Use **no ip domain-lookup** and **| begin <ip address>** to save time

```
Router# show ip cef 10.1.1.5
10.1.1.5/32, version 78, epoch 0, connected, cached adjacency 10.1.1.5
0 packets, 0 bytes
via 10.1.1.5, Vlan1, 0 dependencies
next hop 10.1.1.5, Vlan1
valid cached adjacency

Router# sho adj vlan 1 detail | begin 10.1.1.5
Protocol Interface      Address
IP          Vlan1              10.1.1.5(5)
                2 packets, 228 bytes
                00000C5D143C ← destination mac
                0006526141020800 ← source mac
                ARP          03:55:06
                Epoch: 0
```

RST-3509  
9806\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

25

## Unicast Forwarding Hybrid: Sup1: Checking the SP

Cisco.com

- Verify the supervisor hardware has the vlan configured and that the hardware tables are not overflowing



```
CatOS(enable) sho mls
Total packets switched = 7036256
Total Active MLS entries = 2 ← 32000 max
Long-duration flows aging time = 1920 seconds
<some output removed>

IP MSFC ID      Module XTAG MAC          Vlans
-----
10.1.1.1       15     1     00-06-52-61-41-02     1,2
```

RST-3509  
9806\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

26

## Unicast Forwarding Hybrid: Sup1: Checking the SP

Cisco.com

- Check if the **flow** has been installed into the **MLS cache** and matches the **CEF** information
- Check to see the flows are set to the correct **physical ports**



```
CatOS(enable) show mls entry ip
Destination-IP  Source-IP      Prot  DstPrt  SrcPrt  Destination-Mac  Vlan  EDst  ESrc
-----
MSFC 10.1.1.1 (Module 15):
10.1.1.5      10.1.2.5      ICMP  0        0       00-00-0c-5d-14-3c  1     ARPA ARPA
10.1.2.5      10.1.1.5      ICMP  0        0       00-e0-b0-64-23-fa  2     ARPA ARPA
Total entries displayed: 2
```

DPort	SPort	Stat-Pkts	Stat-Bytes	Uptime	Age
2/2	2/1	829	82900	00:00:04	00:00:00
2/1	2/2	861	86100	00:00:05	00:00:00

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

27

## Unicast Forwarding In Hardware or In Software?

Cisco.com

- Not everything can be forwarded in hardware.
- Hardware forwarding occurs b/c we store MLS flow information in memory called **TCAM**
- TCAM on a Sup1 can hold **32K MLS entries**, but QoS microflow policers also use this same TCAM
- The following IP packets cannot be forwarded in hardware
  - Packets with **IP options set**
  - Packets with **TTL=1**
  - Packets that are fragments or require fragmentation
- **Use this to your advantage!**

RST-3509  
9806\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

28

## Unicast Forwarding Testing Software Forwarding Path

Cisco.com

- Source an extended ping with the **record route** option set
- If these pings get through, the MSFC is correct, but something is amiss with the Supervisor



```
Router#ping
Protocol [ip]:
Target IP address: 10.1.2.5
...
Extended commands [n]: y
Source address or interface: 10.1.1.5
...
Loose, Strict, Record, Timestamp, Verbose[none]: r
```

- From a PC...  
`ping -r 9 10.1.2.5`
- From a Unix box...  
`ping -R 10.1.2.5`

RST-3509  
9806\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

29

## Unicast Forwarding Native: Sup1: Checking the RP

Cisco.com

- With Native, there is still a MSFC and a Supervisor
- However, telnet and default console access is to the MSFC (RP)



- Verify consistent information using all the same commands as you did in the Hybrid MSFC...
- `show ip route`
- `show arp`
- `show ip cef`
- `show adjacency`

RST-3509  
9806\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

30

## Unicast Forwarding Native: Sup1: Checking the SP

Cisco.com

- Verify there is a **flow**
- Verify it appears correctly



```
Router# sho mls ip destination 192.168.50.1
DstIP          SrcIP          Dst i/f:DstMAC      Pkts
-----
192.168.50.1  0.0.0.0       10 :0000.0c14.9d08  16680
```

```
Bytes  SrcDstPorts  SrcDstEncap  Age  LastSeen
-----
1668000 Fa4/1,Fa4/48  ARPA,ARPA   27  15:59:53
```

- If traffic is being switched using this flow, these numbers will increment


RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

31

## Unicast Forwarding Connectivity Loss: CEF Sup2

Cisco.com



- Connectivity problems b/w these 2 routers
- What can you check on the Catalyst 6000?

RST-3509  
9806\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

32

## Unicast Forwarding Native: CEF Sup2: Checking the RP

Cisco.com

- Check the software **routing tables** and **ARP caches**
- Do this for both **source and destination**



```
Native# sho ip route 10.1.1.5
Routing entry for 10.1.1.0/24
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Routing Descriptor Blocks:
    * directly connected, via Vlan1
      Route metric is 0, traffic share count is 1

Native# sho ip arp 10.1.1.5
Protocol Address      Age (min)  Hardware Addr  Type   Interface
Internet 10.1.1.5          0          0000.0c5d.143c ARPA   Vlan1
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

33

## Unicast Forwarding Native: CEF Sup 2: Checking the RP

Cisco.com

- Still in the RP...
- Verify that the **CEF** and **adjacency** table match the **routing table** and **ARP cache**; do this for both source and destination
- Use **no ip domain-lookup** and **| begin <ip address>** to save time

```
Native# sho ip cef 10.1.1.5
10.1.1.5/32, version 33, epoch 0, connected, cached adjacency 10.1.1.5
0 packets, 0 bytes
  via 10.1.1.5, Vlan1, 0 dependencies
  next hop 10.1.1.5, Vlan1
  valid cached adjacency

Native# sho adj vlan 1 det
Protocol Interface      Address
IP        Vlan1                10.1.1.5(5)
                        4 packets, 400 bytes
                        0000C5D143C ← destination mac
source mac → 00D079550C0A0800
                        ARP      03:57:16
                        Epoch: 0
```

RST-3509  
9806\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

34

## Unicast Forwarding Native: CEF Sup 2: Checking the SP

Cisco.com

- Check if there is a **CEF entry and adjacency entry**
- Ensure they match the info in the RP



```
Native# sho mls cef ip 10.1.1.5

Native-sp#
Index      Prefix      Mask      Adjacency
12         10.1.1.5   255.255.255.255  0000.0c5d.143c

Native# sho mls cef adj mac-address 0000.0c5d.143c

Native-sp#
Index 17418 : mac-sa: 00d0.7955.0c0a, mac-da: 0000.0c5d.143c
        interface: V11, mtu: 1514
        packets: 0000000000000000, bytes: 0000000000000000
```

RST-3509  
9806\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

35

## Unicast Forwarding Hybrid: CEF Sup 2/720: Checking the RP

Cisco.com

- Verify consistent information in the RP using all the same commands as you did in the Native example
- **show ip route**
- **show arp**
- **show ip cef**
- **show adjacency**



RST-3509  
9806\_05\_2004\_c1

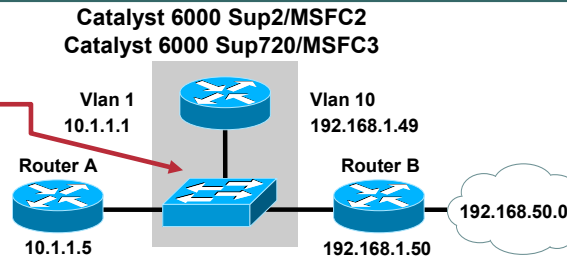
© 2004 Cisco Systems, Inc. All rights reserved.

36

## Unicast Forwarding Hybrid: CEF Sup 2/720: Checking the SP

Cisco.com

- Check if there is a **CEF entry and adjacency entry**
- Ensure they match the info in the RP



```
CatOS> (enable) sh mls entry cef ip 192.168.50.0/24
Mod  FIB-Type  Destination-IP  Destination-Mask  NextHop-IP  Weight
-----
15  resolved  192.168.50.0  255.255.255.0  192.168.1.50  1

CatOS> (enable) sh mls entry cef ip 192.168.1.50/32 adjacency
Mod: 15
Destination-IP: 192.168.1.50  Destination-Mask: 255.255.255.255
FIB-Type: resolved
AdjType  NextHop-IP  NextHop-Mac  Vlan  Encp  Tx-Packets  Tx-Octets
-----
Connect  192.168.1.50  00-07-0e-8f-08-8a  10  ARPA  0  0
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

37

## Unicast Forwarding In Hardware or in Software?

Cisco.com

- Not everything can be forwarded in hardware
- Hardware forwarding occurs b/c we store CEF and Adjacency information in memory called **TCAM**
- Fib TCAM on a Sup2 can hold **256,000** route entries
- Adjacency size is **256,000** entries; Why important?
- The following IP packets cannot be forwarded in hardware
  - Packets with **IP options set**
  - Packets with **TTL=1**
  - Packets that require fragmentation
- **Again, use this to your advantage! Remember the ping test**

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

38

## Unicast Forwarding Hybrid: Fib TCAM Full Check

Cisco.com

- Fib TCAM on a Sup2 can hold **256,000** unicast route entries

```
Console> (enable) show mls cef
Total L3 packets switched:      599997
Total L3 octets switched:      59999700
Total route entries:           14
  IP route entries:            14
```

- Cut in ½ to **128,000** entries if uRPF is enabled on **ANY** interface in RP

```
Console> (enable) sho mls cef interface
Module 15: vlan 10, IP Address 10.1.1.1, Netmask 255.255.255.0
MTU = 1500, State = down, ICMP-Unreach = disabled, ICMP-Redirect = disabled
Unicast RPF = disabled
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

39

## Unicast Forwarding Native: Fib TCAM Full Check

Cisco.com

```
%MLSCEF-SP-7-FIB_EXCEPTION: FIB TCAM exception, Some
entries will be software switched
```

- Fib TCAM on a Sup2 can hold **256,000** unicast route entries

```
NATIVE# sho mls cef summary
NATIVE-sp#
Total CEF switched packets:  0000000000000007
Total CEF switched bytes:    0000000000000322
Total routes:                11
  IP unicast routes:         11
```

- Cut in ½ to **128,000** entries if uRPF is enabled on **ANY** interface in RP

```
NATIVE# sho mls cef hardware
...
rpf mode: off
...
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

40

## Unicast Forwarding Adjacency Usage Check

Cisco.com

- Why is adjacency usage important?

```
Switch> (enable) sh proc cpu
CPU utilization for five seconds: 99.66%
                             one minute: 100.00%
                             five minutes: 100.00%
PID Runtime(ms) Invoked uSecs 5Sec 1Min 5Min TTY Process
-----
29 931842497 192494691 5317000 98.74% 98.00% 98.00% 0 Fib
```

- How to check—hybrid

```
Console> (enable) show polaris fibmgr usage
Total adjacencies: 262144
Allocated adjacencies: 1042
Free adjacencies: 261102
```

- How to check—native

```
NATIVE# sho mls cef adj count
Total adjacencies: 0
```

RST-3509  
9806\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

41

## Unicast Forwarding Adjacency Usage Check

Cisco.com

- Why is adjacency usage important in a topology like this?



- In hybrid 2 adjacencies for **each** prefix (route) w/ this topology
- Possible to run out of adjacencies
- Can **share** adjacencies (starting w/ 7.3(1))

```
Hybrid>(enable) set mls cef per-prefix-stats disable
```

- In native 2 adjacencies for **all** prefixes
- It shares adjacencies by default


RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

42

## Unicast Forwarding Connectivity Loss: CEF Sup 720

Cisco.com



- Connectivity problems to 91.91.91.0
- What can you check on the Catalyst 6000?


RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

43

## Unicast Forwarding Native: CEF Sup 720: Checking the RP

Cisco.com



```
Native720# show ip route 91.91.91.0
Routing entry for 91.91.91.0/24
  Known via "ospf 10", distance 110, metric 2, type intra area
  Last update from 11.11.11.2 on Port-channel1, 00:08:27 ago
  Routing Descriptor Blocks:
    * 11.11.11.2, from 10.10.10.2, 00:08:27 ago, via Port-channel1
      Route metric is 2, traffic share count is 1

Native720# show ip arp 11.11.11.2
Protocol Address          Age (min)  Hardware Addr  Type   Interface
Internet 11.11.11.2         12        0005.7419.5980 ARPA   Port-channel1
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

44

## Unicast Forwarding Native: CEF Sup 720: Checking the RP

Cisco.com

- Still in the RP...
- Verify that the **CEF** and **adjacency** table match the **routing** table and **ARP** cache; do this for both source and destination
- Use **no ip domain-lookup** and **| begin <ip address>** to save time

```
Native720# show ip cef 91.91.91.0 det
91.91.91.0/24, version 4883, epoch 0, cached adjacency 11.11.11.2
0 packets, 0 bytes
via 11.11.11.2, Port-channell, 0 dependencies
next hop 11.11.11.2, Port-channell
valid cached adjacency

Native720# show adj det | begin 11.11.11.2
IP          Port-channell          11.11.11.2(33)
              6019 packets, 9004424 bytes
              000574195980 ← destination mac
              00D000757C000800 ← source mac
              ARP          03:49:07
              Epoch: 0
```

RST-3509  
9806\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

45

## Unicast Forwarding Native: CEF Sup 720: Checking the SP

Cisco.com

- Check if there is a **CEF entry** and **adjacency entry**
- Ensure they match the info in the RP



```
Native720# remote command switch show ip cef 91.91.91.0 det
91.91.91.0/24, version 4883, epoch 0, cached adjacency 11.11.11.2
0 packets, 0 bytes
via 11.11.11.2, Port-channell, 0 dependencies
next hop 11.11.11.2, Port-channell
valid cached adjacency

Native720# remote command switch show adj det | begin 11.11.11.2
IP          Port-channell          11.11.11.2(33)
              0 packets, 0 bytes
              00057419598000D000757C000800
              FIB LC      00:00:00
              Epoch: 0
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

46

## Unicast Forwarding In Hardware or in Software?

Cisco.com

- Not everything can be forwarded in hardware
- Hardware forwarding occurs b/c we store CEF and Adjacency information in memory called **TCAM**
- Fib TCAM on a Sup720 can hold **256,000** route entries and scale up to **1M** routes with **PFC3BXL**
- Adjacency table size goes up to 1M w/ sup720 PFC3a
- The following IP packets cannot be forwarded in hardware
  - Packets with **IP options set**
  - Packets with **TTL=1**
  - Packets that require fragmentation
- **Again, use this to your advantage! Remember the ping test**

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

47

## Unicast Forwarding Hybrid/Native: Sup 720 Fib TCAM Full Check

Cisco.com

- Fib TCAM on a Sup720 can hold **256,000** unicast route entries

```
Hybrid> (enable) show mls cef
Total L3 packets switched:          599997
Total L3 octets switched:          59999700
Total route entries:                14
IP route entries:                  14
```

```
Native# sho mls cef summary
Total routes:                       24
IPv4 unicast routes:                18
```

- **No fib size penalty when uRPF is enabled**

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

48

# Unicast Forwarding Policy-Based Routing

Cisco.com

- **Sup 1**
  - Requires **mls ip pbr** command
  - Behavior is non-deterministic
    - Based on traffic flow and order in which flows are established and removed
- **Sup 2/ Sup720**
  - Only **match ip address <acl>** and **set ip next-hop** are supported in hardware
  - Other match and set operations processed in software
  - Policy routing ACL programmed in hardware—ACE results point to next-hop adjacency information

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

49

# Agenda

Cisco.com

- Hybrid/Native
- Redundancy
- Unicast Forwarding
- **Multicast Forwarding**
- Pot Luck Troubleshooting

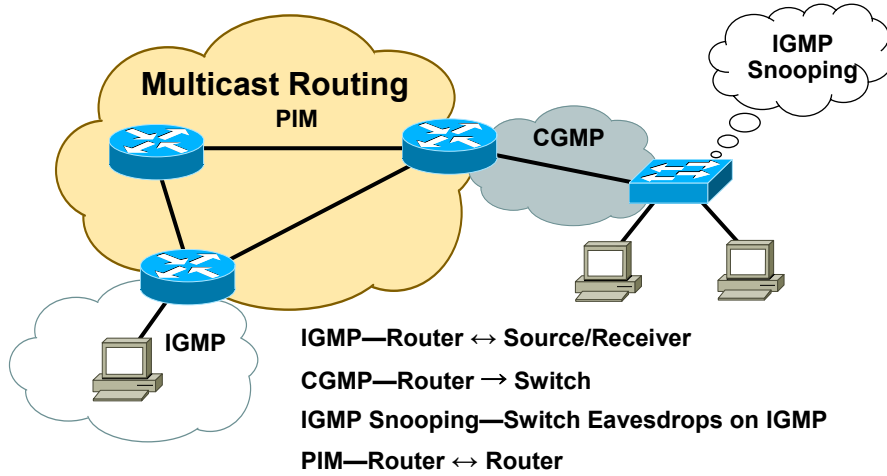
RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

50

# Multicast World of Multicast

Cisco.com



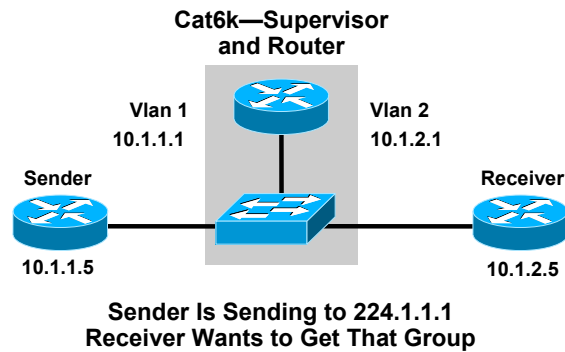
RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

51

# Multicast Path Verification

Cisco.com



RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

52

## Multicast Hybrid: Path Verification

Cisco.com

- Check routing in MSFC (software)
  - show ip mroute
  - show ip mroute count
  - show ip igmp group
- Check routing in MSFC (for hardware entry)
  - show mls ip multicast
- Check hardware entry in supervisor
  - show mls multicast entry group
- Check any multicast constraint table
  - show multicast group
  - show multicast router

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

53

## Multicast Hybrid: Path Verification

Cisco.com

- Check routing in MSFC (software)
- Is the source there?
- Is the **incoming interface** and **oilist** correct?

```
MSFC# show ip mroute
IP Multicast Routing Table
Outgoing interface flags: H - Hardware switched
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.1.1.1), 00:00:17/stopped, RP 0.0.0.0, flags: DJC
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
  Vlan2, Forward/Sparse-Dense, 00:00:02/00:00:00

(10.1.1.5, 224.1.1.1), 00:00:17/00:02:42, flags: T
Incoming interface: Vlan1, RPF nbr 0.0.0.0, RPF-MFD
Outgoing interface list:
  Vlan2, Forward/Sparse-Dense, 00:00:02/00:00:00, H
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

54

# Multicast Hybrid: Path Verification

Cisco.com

- How much am I forwarding?

```
Core_2_6513# show ip mroute 234.92.4.1 count
IP Multicast Statistics
122 routes using 64882 bytes of memory
115 groups, 0.06 average sources per group
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per sec
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 234.92.4.1, Source count: 2, Packets forwarded: 78307
RP-tree: Forwarding: 0/0/0/0, Other: 0/0/0
Source: 10.92.4.1/32, Forwarding: 78008/249/550/1101, Other: 78008/0/0
Source: 10.92.8.1/32, Forwarding: 299/1/68/0, Other: 299/0/0
```

- Is there a host on that interface that wants this group?

```
MSFC# show ip igmp group
IGMP Connected Group Membership
Group Address      Interface      Uptime        Expires       Last Reporter
224.0.1.40         Vlan2         00:30:40     00:02:42    10.1.2.1
224.1.1.1          Vlan2         00:00:10     00:02:49    10.1.2.5
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

55

# Multicast Hybrid: Path Verification

Cisco.com

- Check the routing in the MSFC (for hardware entry)
- Does the MSFC think there is a hardware entry down in the sup for this group?

```
MSFC# show mls ip multicast
Multicast hardware switched flows:
(10.1.1.5, 224.1.1.1) Incoming interface: Vlan1, Packets switched: 30
Hardware switched outgoing interfaces: Vlan2
RPF-MPD installed
Total hardware switched flows : 1
```

- This number is periodically updated from the Sup
- Check hardware entry in Supervisor

```
Switch (enable) show mls multicast entry group 224.1.1.1
Router IP  Dest IP      Source IP      Pkts  Bytes  InVlans  OutVlans
-----
10.1.1.1   224.1.1.1    10.1.1.5      135   13500  1         2
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

56

## Multicast Hybrid: Path Verification

Cisco.com

- Check any multicast constraint tables
- Are the ports that are supposed to get the traffic set up to receive it?

```
Switch (enable) show multicast group 01-00-5e-01-01-01
VLAN  Dest MAC/Route Des      [CoS]  Destination Ports or VCs
-----
1      01-00-5e-01-01-01          15/1
2      01-00-5e-01-01-01          1/1,15/1
```

- Does the switch know the ports where multicast routers are located?

```
Switch (enable) show multicast router
Port      Vlan
-----
15/1      1,2
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

57

## Multicast Native: Path Verification

Cisco.com

- Check routing in RP (software)
- Is the source there?
- Is the **incoming interface** and **oilist** correct?

```
Native# show ip mroute 234.92.4.1
IP Multicast Routing Table

(*, 234.92.4.1), 00:01:04/stopped, RP 10.92.8.253, flags: SJC
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Vlan902, Forward/Sparse-Dense, 00:01:04/00:02:23
    Vlan901, Forward/Sparse-Dense, 00:01:04/00:02:27

(10.92.4.1, 234.92.4.1), 00:00:19/00:02:59, flags: T
  Incoming interface: Vlan901, RPF nbr 0.0.0.0, RPF-MFD
  Outgoing interface list:
    Vlan902, Forward/Sparse-Dense, 00:00:19/00:02:40, H

(10.92.8.1, 234.92.4.1), 00:00:12/00:02:47, flags: T
  Incoming interface: Vlan902, RPF nbr 0.0.0.0, RPF-MFD
  Outgoing interface list:
    Vlan901, Forward/Sparse-Dense, 00:00:15/00:02:44, H
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

58

## Multicast Native: Path Verification

Cisco.com

- Check the routing in the MSFC (for hardware entry)
- Does the MSFC think there is a hardware entry down in the sup for this group?

```
Native# sho mls ip multicast group 234.92.4.1
Multicast hardware switched flows:
(10.92.8.1, 234.92.4.1) Incoming interface: Vlan902, Packets switched: 592
Hardware switched outgoing interfaces: Vlan901
RPF-MFD installed

(10.92.4.1, 234.92.4.1) Incoming interface: Vlan901, Packets switched: 152709
Hardware switched outgoing interfaces: Vlan902
RPF-MFD installed

Total hardware switched flows : 2
```

```
Native# show mls ip multicast summary
1 MMLS entries using 140 bytes of memory
Number of partial hardware-switched flows: 0
Number of complete hardware-switched flows: 2
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

59

## Multicast Native: Path Verification

Cisco.com

- Check hardware entry in Supervisor

```
Native-sp# sho mls cef ip multicast source 10.92.8.1 group 234.92.4.1
Source Destination RPF #packets #bytes Type Output Vlans/Info
+-----+-----+-----+-----+-----+-----+-----+
10.92.8.1 234.92.4.1 V1902 817 55556 MFD V1901 [1 oifs]

Native-sp# sho mls cef ip multicast source 10.92.8.1 group 234.92.4.1
Source Destination RPF #packets #bytes Type Output Vlans/Info
+-----+-----+-----+-----+-----+-----+-----+
10.92.8.1 234.92.4.1 V1902 820 55760 MFD V1901 [1 oifs]
```

- These numbers are updated nearly real-time

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

60

## Multicast Native: Path Verification

Cisco.com

- Check any multicast constraint tables
- Are the ports that are supposed to get the traffic set up to receive it?

```
Native# show mac-address-table multicast vlan 2
-----+-----+-----+-----+-----+-----+-----+-----
vlan  mac address      type  learn qos      ports
-----+-----+-----+-----+-----+-----+-----+-----
  2    0100.5e00.0128    static Yes  -- Router
  2    0100.5e01.0101    static Yes  -- Gi1/1,Router
```

- Does the switch know the ports where multicast routers are located?

```
Native> show ip igmp snooping mrouter
vlan      ports
-----+-----
431 Fa1/13
432 Fa1/13,Router
444 Fa1/13
```


RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

61

## Multicast High CPU on Non-DR Router

Cisco.com



- In a redundant routed multicast environment, the multicast non-designated router will see high CPU utilization due to non-reverse path forwarding

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

62

## Multicast High CPU on Non-DR Symptoms

Cisco.com

- IP input in the process path

```
Core_2_6513# sho proc cpu | exc 0.00
CPU utilization for five seconds: 88%/71%; one minute: 71%; five minutes: 27%
PID Runtime(ms)   Invoked    uSecs   5Sec   1Min   5Min  TTY Process
 35      5212    251082      20  0.08%  0.06%  0.28%  1 Virtual Exec
 69     71888   504110     142 15.30% 12.30%  4.00%  0 IP Input
```

- Quickly incrementing input queue drops

```
Core_2_6513# sho int vlan 901 | inc drop
Input queue: 37/75/4595/4595 (size/max/drops/flushes); Total output drops: 0
Core_2_6513# sho int vlan 901 | inc drop
Input queue: 44/75/4631/4631 (size/max/drops/flushes); Total output drops: 0
```

- RPF counter increments quickly

```
Core_2_6513# sho ip mroute 234.92.15.1 count
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kilobits per second
Other counts: Total/RPF failed/Other drops(OIF-null, rate-limit etc)

Group: 234.92.15.1, Source count: 2, Packets forwarded: 0
Source: 10.92.10.135/32, Forwarding: 0/0/0/0, Other: 0/286/0
Source: 10.92.15.1/32, Forwarding: 0/0/0/0, Other: 0/59170/0
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

63

## Multicast High CPU on Non-DR Router

Cisco.com

- How does Cisco deal with non-RPF traffic with dual MSFC's in a single chassis?

```
MSFC# show ip mroute
(10.1.1.1.5, 224.1.1.1), 00:00:17/00:02:42, flags: T
Incoming interface: Vlan1, RPF nbr 0.0.0.0, RPF-MFD
Outgoing interface list:
Vlan2, Forward/Sparse-Dense, 00:00:02/00:00:00, H
```

- RPF-MFD = **R**everse **P**ath **F**orward **M**ulti **F**ast **D**rop
- Sup 2—On by default after 6.2(1)
- Sup 1—**RPF-MFD** not available internally for dual Sup chassis, must use an ACL on Non-DR, and then should only use if SM

```
MSFC(config)# mls ip multicast stub
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

64

## Multicast Input Queue Drops: Packets w/ ttl = 1

Cisco.com


- Input queue drops are caused by too much **process switched** traffic

```
Native# sho int vlan 901 | inc drop
Input queue: 37/75/4595/4595 (size/max/drops/flushes); Total output drops: 0
```

- Packets w/ **ttl = 1** must be process switched

```
Native# show buffer input vlan901 packet

Buffer information for Middle buffer at 0x423D8BC0
...
source: 10.92.10.5, destination: 234.92.12.5, id: 0x32F7, ttl: 1,
TOS: 0 prot: 17, source port 49450, destination port 49450
```



- Look for applications that allow the setting of **Maximum TTL AND Minimum TTL**

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

65


## Multicast Unexpected Process Switched Traffic

Cisco.com

- Why is this multicast packet showing up in the buffer?

```
Native# show buffer input vlan902 packet

Buffer information for Middle buffer at 0x423D8BC0
...
source: 10.92.12.5, destination: 234.92.12.5, id: 0x81F8, ttl: 58,
TOS: 0 prot: 17, source port 49440, destination port 49440
```



RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

66

## Multicast Packets Handled in Software

Cisco.com

- Packets requiring PIM register encapsulation
- Packets with **IP options** in the header
- Packets with **TTL = 1**
- Packets requiring fragmentation
- Partial switched flows

A single multicast flow, with some OIFs handled in hardware, some in software

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

67

## Multicast IGMP Rate-Limiting

Cisco.com

- Feature introduction in **Hybrid 6.3(8) 7.4(2) 7.5(1) 8.1(1)**
- **Protects NMP from excessive** multicast control packets when IGMP snooping is enabled by limiting # of packets in 30 second period
- Unfortunate side effect is **multicast router port** may disappear and no syslog message given

```
Switch (enable) show multicast router
Port          Vlan
-----
15/1          1,2
```

- Suitable workaround in most networks is to raise the igmp/pim rate-limit thresholds

```
Console> (enable) set igmp ratelimit pimv2 1000
Console> (enable) set igmp ratelimit general-query 1000
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

68

# Multicast IGMP Rate-Limiting

Cisco.com

- Improvements to feature **6.4(9) 7.6(6) 8.2(1)**
  - Syslog
  - Disabled by default
  - CLI configurable

```
Console> (enable) sho igmp ratelimit-info
IGMP Ratelimiting is disabled
IGMP Ratelimiting: No of messages allowed in 30 seconds
-----
Dvmrp Probes           : 100
Mospf1 Hellos          : 100
Mospf2 Hellos          : 100
PimV2 Hellos           : 100
Console> (enable) set igmp ratelimit ?
  disable              Disable IGMP ratelimit
  dvmrp                Set IGMP rate limit for dvmrp
  enable               Enable IGMP ratelimit
  mospf1               Set IGMP rate limit for mospf v1
  mospf2               Set IGMP rate limit for mospf v2
  pimv2                Set IGMP rate limit for pim v2
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

69

# Agenda

Cisco.com

- Hybrid/Native
- Redundancy
- Unicast Forwarding
- Multicast Forwarding
- **Pot Luck Troubleshooting**

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

70

# Pot Luck

## QOS Hybrid: Verifying Port Settings

Cisco.com

- **Show qos info runtime** gives the current qos settings per port

```
6506 (enable) show qos info runtime 4/1
Run time setting of QoS:
QoS is enabled ← Is QOS Enabled?
Policy Source of port 4/1: Local
Tx port type of port 4/1 : 2q2t ← How Many Queues/Thresholds?
Rx port type of port 4/1 : 1q4t
Interface type: port-based
ACL attached:
The qos trust type is set to untrusted. ← Trust Status
Default CoS = 0
Queue and Threshold Mapping for 2q2t (tx): ← COS to Queue/Threshold Mapping
Queue Threshold CoS
-----
1      1      0 1
1      2      2 3
2      1      4 5
2      2      6 7
Queue and Threshold Mapping for 1q4t (rx):
All packets are mapped to a single queue.
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

71

# Pot Luck

## QOS Hybrid: Verifying Port Settings (Cont.)

Cisco.com

```
Rx drop thresholds:
Rx drop thresholds are disabled. ← Drop Thresholds
Tx drop thresholds:
Queue #  Thresholds - percentage (* abs values)
-----
1      80% (288384 bytes) 100% (360192 bytes) ← Queue Threshold Ratios
2      80% (65536 bytes) 100% (77824 bytes)
Rx WRED thresholds:
Rx WRED feature is not supported for this port type. ← WRED Thresholds
Tx WRED thresholds:
WRED feature is not supported for this port type.
Tx queue size ratio:
Queue #  Sizes - percentage (* abs values) ← Queue Size Ratios
-----
1      80% (360448 bytes)
2      20% (81920 bytes)
Rx queue size ratio:
Rx queue size-ratio feature is not supported for this port type.
Tx WRR Configuration of ports with speed 1000Mbps:
Queue #  Ratios (* abs values)
-----
1      100 (25600 bytes)
2      255 (65280 bytes)
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

72

# Pot Luck QoS Native: Verifying Port Settings

Cisco.com

- **Show queuing interface** gives the current qos settings per port

```
Router# sho queuing int gig 1/1
Interface GigabitEthernet1/1 queuing strategy: Weighted Round-Robin
Port QoS is enabled ← Is QoS Enabled?
Port is untrusted ← Trust Status
Extend trust state: not trusted [COS = 0]
Default COS is 0
Transmit queues [type = lp2q2t]: ← Tx Queues/Thresholds
-----
Queue Id      Scheduling  Num of thresholds
-----
1             WRR low     2
2             WRR high   2
3             Priority    1
WRR bandwidth ratios: 100[queue 1] 255[queue 2] ← WRED Ratios
queue-limit ratios:  70[queue 1] 15[queue 2] ← WRED Thresholds
queue random-detect-min-thresholds
-----
1      40[1] 70[2]
2      40[1] 70[2]
queue random-detect-max-thresholds
-----
1      70[1] 100[2]
2      70[1] 100[2]
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

73

# Pot Luck QoS Native: Verifying Port Settings (Cont.)

Cisco.com

```
queue thresh cos-map ← Tx Cos to Queue/
-----
1      1                               Threshold Map
1      2
2      1
2      2
3      1
Receive queues [type = lp1q4t]: ← Rx Queues/Thresholds
Queue Id      Scheduling  Num of thresholds
-----
1             Standard    4
2             Priority    1
queue tail-drop-thresholds ← Rx Drop Thresholds
-----
1      100[1] 100[2] 100[3] 100[4]
queue thresh cos-map ← Rx Cos to Queue/
-----
1      1                               Threshold Map
1      2
1      3
1      4
2      1
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

74

## Pot Luck QoS Marking Map Verification

Cisco.com

- **Show qos map runtime cos-dscp** gives cos-to-dscp map

```
Console>(enable) show qos map runtime cos-dscp
CoS - DSCP map:
CoS   DSCP
---   ----
 0     0
 1     8
 2    16
 3    24
 4    32
 5    40
 6    48
 7    56
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

75

## Pot Luck QoS Classification and Policing

Cisco.com

### Verification—Single Rate Policer

```
Native# sho mls qos ip fast 2/1
[In] Policy map is myPOLICER [Out] Default.
QoS Summary [IP]: (* - shared aggregates, Mod - switch module, F - install error)

Int          Mod Dir Cl-map DSCP AgId Trust FlId  AgForward-Pk  AgPoliced-Pk
-----
Fa2/1        1 I  HOSTCHRIS  0   1  dscp  0           45652          0
<run command again>
Fa2/1        1 I  HOSTCHRIS  0   1  dscp  0           45835          0
```

**AgForward-Pk Will Increment if Traffic Is Matching the ACL**

```
Native# sho mls qos ip fast 2/1
[In] Policy map is myPOLICER [Out] Default.
QoS Summary [IP]: (* - shared aggregates, Mod - switch module, F - install error)
Int          Mod Dir Cl-map DSCP AgId Trust FlId  AgForward-Pk  AgPoliced-Pk
-----
Fa2/1        1 I  HOSTCHRIS  0   1  dscp  0           188631         4625
<run command again>
Fa2/1        1 I  HOSTCHRIS  0   1  dscp  0           198852         5975
```

**AgPoliced-Pk Will Increment if Traffic Is Being Policed**

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

76

# Pot Luck QOS Classification and Policing

Cisco.com

## Verification—Dual Rate Policer (Native)

```
Native# sho mls qos ip gig 1/1
[In] Policy map is myPOLICER [Out] Default.
QoS Summary [IP]: (* - shared aggregates, Mod - switch module, F - install error)
Int      Mod Dir Cl-map DSCP AgId Trust FlId  AgForward-Pk  AgPoliced-Pk
-----
Gig1/1   1 I  HOSTCHRIS  0  1  dscp  0          14478         0
<run command again>
Gig1/1   1 I  HOSTCHRIS  0  1  dscp  0          27838         0
```

**AgForward-Pk Will Increment if Traffic Is Matching the ACL**

```
Native# sho mls qos ip gig 1/1
[In] Policy map is myPOLICER [Out] Default.
QoS Summary [IP]: (* - shared aggregates, Mod - switch module, F - install error)
Int      Mod Dir Cl-map DSCP AgId Trust FlId  AgForward-Pk  AgPoliced-Pk
-----
Gig1/1   1 I  HOSTCHRIS  0  1  dscp  0          197349        2899
<run command again>
Gig1/1   1 I  HOSTCHRIS  0  1  dscp  0          218985        7863
```

**AgPoliced-Pk Will Increment if Traffic Is Being Policed**

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

77

# Pot Luck QOS Classification and Policing

Cisco.com

## Verification—Dual Rate Policer (Native)

```
Native# sho mls qos
QoS global counters:
Total packets: 549639
IP shortcut packets: 0
Packets dropped by policing: 0
IP packets with TOS changed by policing: 39183
IP packets with COS changed by policing: 42587
```

```
Native# sho mls qos last
Packet was transmitted
Packet L3 Prot: 0, packet length: 46, dont_plc: No
Input COS: 0, TOS/DSCP: 0x80/32
Output TOS/DSCP: 0x0/0[rewritten] Output COS: 0[unchanged]
NT&NS: l3_prot: 0(1), 10.1.1.5.0x0000 ==> 10.1.2.5.0x0000
<run command again>
Packet was transmitted
Packet L3 Prot: 0, packet length: 46, dont_plc: No
Input COS: 0, TOS/DSCP: 0x80/32
Output TOS/DSCP: 0x80/32[unchanged] Output COS: 4[rewritten]
<sod>
NT&NS: l3_prot: 0(1), 10.1.1.5.0x0000 ==> 10.1.2.5.0x0000
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

78

## Pot Luck QOS Classification and Policing

Cisco.com

### To Verify Policing Action on Sup1 (Hybrid)

```
Console> (enable) show qos statistics l3stats
Packets dropped due to policing:          0
IP packets with ToS changed:             377218
IP packets with CoS changed:             22405
Non-IP packets with CoS changed:         0
```

### To Verify Policing Action on Sup2 (Hybrid)

```
Console> (enable) show qos statistics aggregate-policer user1Mbps
QoS aggregate-policer statistics:
Aggregate policer      Allowed packet  Packets exceed  Packets exceed
                       count           normal rate    excess rate
-----
user1Mbps              115728         884731         884731
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

79

## Pot Luck QOS Classification and Policing

Cisco.com

### To Verify Policing Action on Sup720 (Hybrid)

```
W2S-5> (enable) show qos statistics l3stats
Packets dropped due to policing:          0
IP packets with ToS changed:             0
IP packets with CoS changed:             0
Non-IP packets with CoS changed:         0

W2S-5> (enable) show qos statistics aggregate-policer travis
QoS aggregate-policer statistics:
Aggregate policer      Allowed byte    Bytes exceed
                       count           excess rate
-----
travis                 986            0
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

80

## Pot Luck QOS Classification and Policing

Cisco.com

- **Sample policer on Sup720 (Native)**

```
class-map match-all ip-traffic
  match access-group 10
  !
policy-map police-ip-300mbps
  class ip-traffic
    police 295000000 9218750 9218750 conform-action transmit exceed-action drop
  !
interface GigabitEthernet1/8
  service-policy output police-ip-300mbps
  !
access-list 10 permit any
```

- **Show interface output may be used as an indicator, but shouldn't expect real-time accuracy**

```
Native# show int gig 1/8 | include rate
Queuing strategy: fifo
30 second input rate 0 bits/sec, 0 packets/sec
30 second output rate 299825000 bits/sec, 585595 packets/sec
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

81

## Pot Luck QOS Classification and Policing

Cisco.com

- **Check policer on Sup720 (Native) with show policy**

```
Native# show policy int gig 1/8
GigabitEthernet1/8
Service-policy output: police-ip-300mbps
  class-map: ip-traffic (match-all)
    Match: access-group 10
      police :
        295000000 bps 9218000 limit 9218000 extended limit
    Earl in slot 5 :
      40328788032 bytes
      30 second offered rate 381034704 bps
      aggregate-forwarded 31238803328 bytes action: transmit
      exceeded 9089984704 bytes action: drop
      aggregate-forward 294919728 bps exceed 0 bps
```

- **Check policer on Sup720 with show mls qos ip egress**

```
Native# show mls qos ip egress
Int Mod Dir Class-map DSCP Agg Trust Fl AgForward-By AgPoliced-By
-----
-
Gil/8 5 Out imp-traffic 0 1 -- 0 724688256 199581376
```

RST-3509  
9806\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

82

## Pot Luck Output Port Oversubscription

Cisco.com

- Does anyone see a problem with this picture?



- Does a switch make this problem disappear?



- **Answer:** Only for short periods of time!

RST-3509  
9806\_05\_2004\_c1


© 2004 Cisco Systems, Inc. All rights reserved.

83

## Pot Luck Output Port Oversubscription

Cisco.com

- Switch makes one-to-many connectivity possible



- Performing data transfers across high-to-slow speed media transitions for extended periods of time results in... **Out Discards**

```
Hybrid> sho counters 4/2
 9  ifInUnknownProtos          =          0
10  ifOutDiscards               =         2309
11  txDelayExceededDiscards    =          0
```

- This counter rapidly increased during file transfer

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

84

# Pot Luck

## Understanding Switching Paths

Cisco.com

- Switching in hardware (MLS/CEF) has **little effect** on router CPU
- **show process cpu | exclude 0.00**

```
Native# show process cpu | exc 0.00
CPU utilization for five seconds: 2%/0%; one minute: 2%; five minutes: 1%
PID Runtime(ms)   Invoked    uSecs   5Sec   1Min   5Min  TTY Process
 35    5212    251082      20  0.08%  0.06%  0.28%  1 Virtual Exec
```

- If not hardware switched, then software switched in **CEF/FAST** path
- Referred to as **interrupt** traffic

```
Native# show process cpu
CPU utilization for five seconds: 18%/18% ; one minute: 19%; five minutes: 18%
PID Runtime(ms)   Invoked    uSecs   5Sec   1Min   5Min  TTY Process
  1         0         4         0  0.00%  0.00%  0.00%  0 Chunk Manager
...
```

- Without CEF/FAST path, then software switched in **process** path

```
Native# show process cpu | exc 0.00
CPU utilization for five seconds: 88%/71% ; one minute: 71%; five minutes: 27%
PID Runtime(ms)   Invoked    uSecs   5Sec   1Min   5Min  TTY Process
 69   71888    504110    142 15.30% 12.30%  4.00%  0 IP Input
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

85

# Pot Luck

## Card Unsupported

Cisco.com

- New linecard inserted, doesn't power up, message generated

```
%C6KPWR-SP-4-UNSUPPORTED: unsupported module in slot 11, power not allowed: Unknown Card Type
```

- Go to **release notes** on CCO, Check **minimum software required** for particular part being installed

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/relnotes/index.htm>

Product Number Append with "=" for Spares	Product Description	Minimum Supervisor Engine Software Version	Recommended Supervisor Engine Software Version
Gigabit Ethernet Switching Modules			
WS-X6148-GE-TX WS-X6148V-GE-TX	48-port 10/100/1000BASE-TX switching module (WS-X6148V-GE-TX provides inline power to IP telephones) QoS port architecture (Rx/Tx): 1q2t/1p2q2t	7.6(1)	7.6(1)

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

86

## Pot Luck High CPU on Switch: IP Phones

Cisco.com

- **SptBpduTx** causing high CPU

```
Hybrid> show proc cpu
CPU utilization for five seconds: 68.70%
                             one minute: 67.16%
                             five minutes: 67.04%

PID Runtime(ms) Invoked  uSecs  5Sec   1Min   5Min   TTY Process
-----
37 309394136 82568023 164000 63.62% 68.43% 68.06% -2 SptBpduTx
```

- IP phones not set up with **auxiliary vlan** but trunking **all vlans**
- Effect increased with more phones

```
Hybrid> show trunk

Port      Vlans allowed on trunk
-----
3/1      1-1005,1025-4094
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

87

## Pot Luck Not Enough Power: Commands

Cisco.com

```
%C6KPWR-SP-4-POWERDENIED: insufficient power, module in slot 9 power denied.
```

```
Hybrid> (enable) show environment power
PS1 Capacity: 2331.00 Watts (55.50 Amps @42V)
PS2 Capacity: none
PS Configuration : PS1 and PS2 in Redundant Configuration.
Total Power Available: 2331.00 Watts (55.50 Amps @42V)
Total Power Available for Line Card Usage: 2331.00 Watts (55.50 Amps @42V)
Total Power Drawn From the System: 730.38 Watts (17.39 Amps @42V)
Total Inline Power Drawn From the System: 0.000 ( 0.000 Amps @42V)
Remaining Power in the System: 1600.62 Watts (38.11 Amps @42V)
```

```
Router# show power
system power redundancy mode = redundant
system power total = 3830.40 Watts (91.20 Amps @ 42V)
system power used = 813.54 Watts (19.37 Amps @ 42V)
system power available = 3016.86 Watts (71.83 Amps @ 42V)

Power-Capacity PS-Fan Output Oper
Watts A @42V Status Status State
-----
1 WS-CAC-4000W-US 3830.40 91.20 OK OK on
2 DS-CAC-4000W-US 3830.40 91.20 OK OK on
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

88

## Pot Luck Not Enough Power: Gotchas

Cisco.com

- System always reserves power for 2nd sup, even if not present

Slot	Card Type	PowerRequested		PowerAllocated		CardStatus
		Watts	A @42V	Watts	A @42V	
3	WS-X6348-RJ-45	100.38	2.39	100.38	2.39	ok
5	WS-SUP720-BASE	315.00	7.50	315.00	7.50	ok
6		0.00	0.00	315.00	7.50	none

- Sup 720 reserves 315W, Sup2 reserves 145W, Sup1 reserves 138W
- Go to **installation guide** on CCO, check **power requirements**  
[www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/6000hw/6500\\_ins/02prep.htm](http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/6000hw/6500_ins/02prep.htm)

Model Number/ Module Type	Card Current (A)	Card Power (Watts)	AC Power (Watts)	Heat Diss (BTU/HR)
WS-X6K-SUP1A-MSFC Sup Engine 1A with PFC and MSFC daughter card	3.30	138.60	173.25	519.68

- **2500W PS @ 120VAC provide same power as 1300W PS**


RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

89

## Pot Luck ACL Is Being Software Switched

Cisco.com



```
interface Vlan2
  no ip unreachable
  ip access-group 105 in
  !
  access-list 105 deny ip any 10.1.1.1 255.255.255.255
  access-list 105 permit ip any 10.1.1.2 255.255.255.255 log
  access-list 105 permit ip any any
```

- Sup 1—**log** keyword, hits on the ACE go to the MSFC
- Sup 1—**ip unreachable**, hits on the deny ACE, ICMP unreachable sent by the MSFC
- Sup 2/Sup720—rate-limits both to the MSFC

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

90

## Pot Luck ACL Errors Messages

Cisco.com

- Applying an access-list results in a some error message

```
%ACL-3-NOLOU:Acl engine is out of logical operation unit
%ACL-3-RACLMAPCOMMITFAIL:Failed to map Router ACL to VLAN 1
%FM-4-TCAM_LOU: Hardware TCAM LOU capacity exceeded
%FM-4-RACL_REDUCED: Interface Vlan1 routed traffic will be software switched
in ingress direction(s)
```

- TCAM is unable to store all the ACL in hardware
- Two major **merge algorithms** used
  - Order Independent Merge—original, aka **(Binary Decision Diagram)**
  - Order Dependent Merge—newer method **(ODM)**

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

91

## Access Control Lists Merge Algorithms

Cisco.com

### Options/Defaults

#### Native

- <12.1(11b)E4 **BDD only**
- ≥12.1(11b)E4 BDD default, **ODM configurable**
- 12.2S **ODM only**

#### Hybrid

- <7.1 **BDD only**
- ≥ 7.1 BDD default, **ODM configurable**
- ≥ 8.x **ODM only**

```
CatOS(6k)> (enable) set aclmerge algo odm
Acl merge algorithm set to odm.
```

```
MSFC(config)# mls acl algorithm odm
```

```
Native(6k)(config)#mls aclmerge algorithm odm
The algorithm chosen will take effect for new ACLs which are being
applied, not for already applied ACLs
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

92

## Pot Luck Merge Algorithms

Cisco.com

- Make sure that the ACL was programmed into hardware by the MSFC and is ACTIVE

```
MSFC# show fm summary
Current global ACL merge algorithm: ODM
ODM optimizations disabled
Interface: Vlan20 is up
ACL merge algorithm used:
  inbound direction: ODM
  outbound direction: ODM
TCAM screening for features is ACTIVE inbound
```

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

93

## Pot Luck Software Recommendations

Cisco.com

- Some quotes...
  - “What code should I run?”
  - “I want the most stable code, what is it?”
  - “I want the latest and greatest, will it work in my network?”
- TAC engineers are trained...
  - Not to recommend code**
    - Too many variables
    - Misperception that code recommended is **bug-free**
    - Some accounts have ANS support
    - Point customers to **release notes** for general guidance
- TAC engineers **DO** have resources to...
  - Determine if a version is exposed to a particular bug
  - Tell you if they **personally** know of any major issues

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

94

## Pot Luck CCO Top Ten

Cisco.com

- **Release Notes**

<http://www.cisco.com/univercd/cc/td/doc/product/lan/cat6000/relnotes/index.htm>

- **Tech Tips**

<http://www.cisco.com/cgi-bin/Support/browse/index.pl?i=Technologies>

- **Power Calculator**

<http://tools.cisco.com/cpc/launch.jsp>

- **Best Practices (Hybrid and Native)**

[http://www.cisco.com/en/US/products/hw/switches/ps663/products\\_tech\\_note09186a0080094713.shtml](http://www.cisco.com/en/US/products/hw/switches/ps663/products_tech_note09186a0080094713.shtml)

[http://www.cisco.com/en/US/products/hw/switches/ps700/products\\_white\\_paper09186a00801b49a4.shtml](http://www.cisco.com/en/US/products/hw/switches/ps700/products_white_paper09186a00801b49a4.shtml)

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

95

## Pot Luck CCO Top Ten (Cont.)

Cisco.com

- **Bug Toolkit**

[http://www.cisco.com/cgi-bin/Support/Bugtool/launch\\_bugtool.pl](http://www.cisco.com/cgi-bin/Support/Bugtool/launch_bugtool.pl)

- **Field Notices**

[http://www.cisco.com/en/US/products/hw/switches/ps708/prod\\_field\\_notices\\_list.html](http://www.cisco.com/en/US/products/hw/switches/ps708/prod_field_notices_list.html)

- **Multicast in the Campus (Tech Tip)**

[http://www.cisco.com/en/US/products/hw/switches/ps700/products\\_tech\\_note09186a00800b0871.shtml](http://www.cisco.com/en/US/products/hw/switches/ps700/products_tech_note09186a00800b0871.shtml)

- **Software Advisor**

<http://tools.cisco.com/Support/Fusion/FusionHome.do>

- **Troubleshooting CEF on a Sup2 (Tech Tip)**

[http://www.cisco.com/en/US/tech/tk827/tk831/technologies\\_tech\\_note09186a0080094b27.shtml](http://www.cisco.com/en/US/tech/tk827/tk831/technologies_tech_note09186a0080094b27.shtml)

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

96

## Complete Your Online Session Evaluation!

Cisco.com

- WHAT:** Complete an online session evaluation and your name will be entered into a daily drawing
- WHY:** Win fabulous prizes! Give us your feedback!
- WHERE:** Go to the Internet stations located throughout the Convention Center
- HOW:** Winners will be posted on the onsite Networkers Website; four winners per day

RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

97

# CISCO SYSTEMS



RST-3509  
9806\_05\_2004\_c1

© 2004 Cisco Systems, Inc. All rights reserved.

98