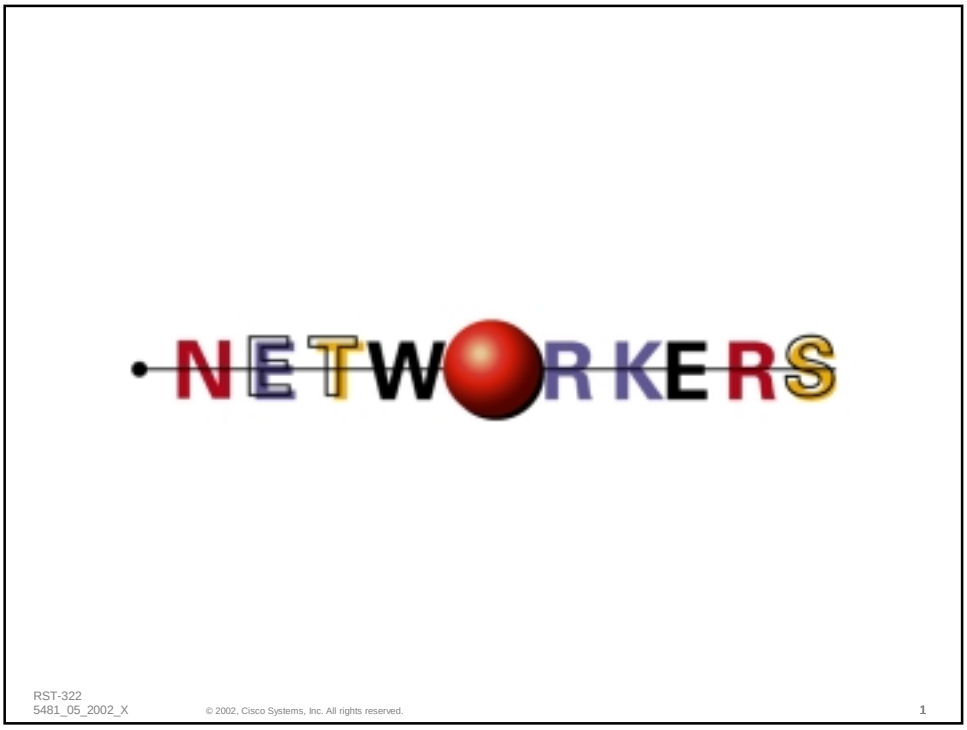

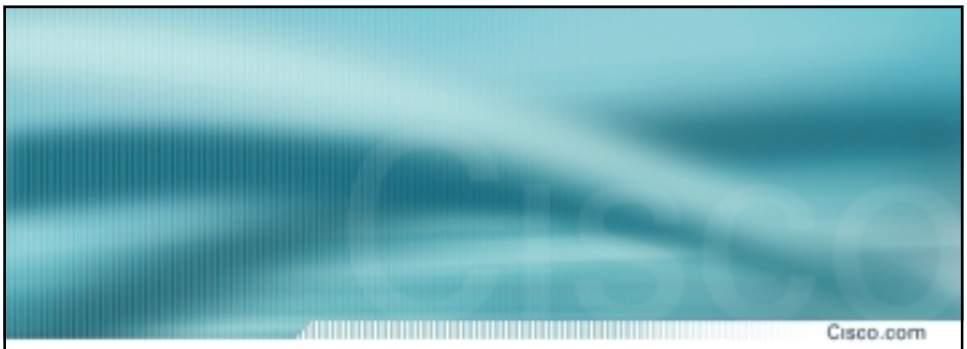

Troubleshooting Catalyst® Switches

Session RST-322

Agenda

- Terminology
- Spanning Tree
- Access Control Lists
- Multicast
- Multi-Layer Switching
- Redundancy

Catalyst Operating Systems–Terminology

Cisco.com

- **CatOS**
  - Used on the Catalyst 4000, 5000, and 6000 series switches
  - CatOS> prompt used for CatOS commands
- **Native IOS**
  - Used on Catalyst 2950, Catalyst 3550, Catalyst 4000 SupIII, and Catalyst 6000 MSFC systems
  - Native# prompt used for Native IOS specific commands
- **IOS**
  - Used on all Cisco Routers, including the MSFC on a Supervisor running CatOS
  - IOS# prompt used for general IOS commands

Network Layers–Picture

Cisco.com

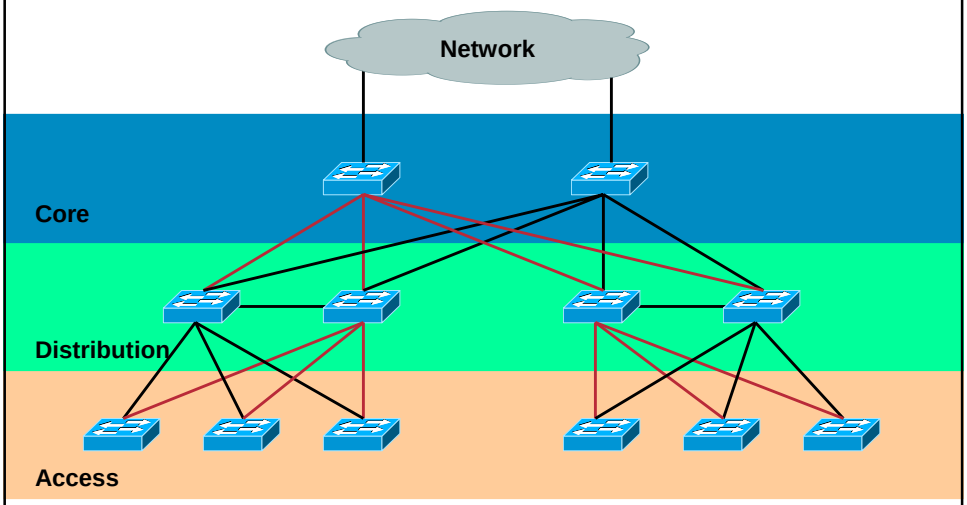

Agenda

Cisco.com

- Terminology
- Spanning Tree
- Access Control Lists
- Multicast
- Multi-Layer Switching
- Redundancy

Spanning Tree—Problem Bridge Loop

Cisco.com

- Bridge loops in the layer 2 redundant network causes attached routers to have high CPU, and other links in the network to be filled with traffic
- Effectively makes the network unusable

Spanning Tree —How It Works

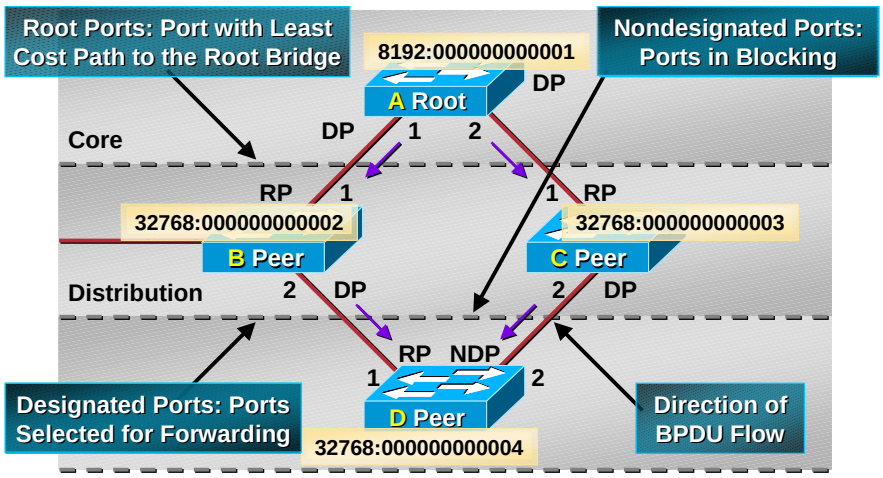

Spanning Tree — How It Works Parameters

- **Network parameters**
  - **Hello interval**
  - **Forward delay**
  - **Max age**
  - **Bridge priority (per bridge)**
- **Port-specific parameters**
  - **Port cost**
  - **Port priority**

Spanning Tree—Commands
Show Spantree

Cisco.com

```
CatOS> (enable) show spantree 1
VLAN 1
Spanning tree mode          PVST+
Spanning tree type          ieee
Spanning tree enabled

Designated Root             00-60-83-55-7b-00
Designated Root Priority    100
Designated Root Cost        100
Designated Root Port        3/1
Root Max Age  20 sec        Hello Time  2  sec  Forward Delay  15 sec

Bridge ID MAC ADDR          00-d0-06-24-6c-00
Bridge ID Priority          32768
Bridge Max Age  20 sec      Hello Time  2  sec  Forward Delay  15 sec

Port      Vlan      Port-State      Cost      Prio      Portfast   Channel_id
------    ------    ----------      --------  ------    ---------- ----------
1/1       1         not-connected   4         32        disabled   0
3/1       1         forwarding      100       32        enabled    0
```


Spanning Tree—Commands
Show Spanning-tree

Cisco.com

```
Native#show spanning-tree vlan 1 brief

VLAN0001
  Spanning tree enabled protocol ieee
  Root ID          Priority   1
                   Address    0060.8355.7b00
                   Cost       23
                   Port       1 (GigabitEthernet1/1)
                   Hello Time  2 sec   Max Age  20 sec    Forward Delay  15 sec

  Bridge ID        Priority   32769  (priority 32768 sys-id-ext 1)
                   Address    0007.0e8f.0880
                   Hello Time  2 sec   Max Age  20 sec    Forward Delay  15 sec
                   Aging Time 300

Interface                                                  Designated
Name                 Port ID  Prio  Cost  Sts  Cost  Bridge ID                 Port ID
-------------------- -------- ----- ----- ---- ----- ------------------------- -------
GigabitEthernet1/1 128.1     128       4 FWD     67  32768 0005.5f33.dc01    128.1
FastEthernet3/48   128.176   128      19 FWD     48  32768 0030.7bdd.5080    128.16
```

Spanning Tree—Solution
Where Not to Put the Root Bridge

Cisco.com

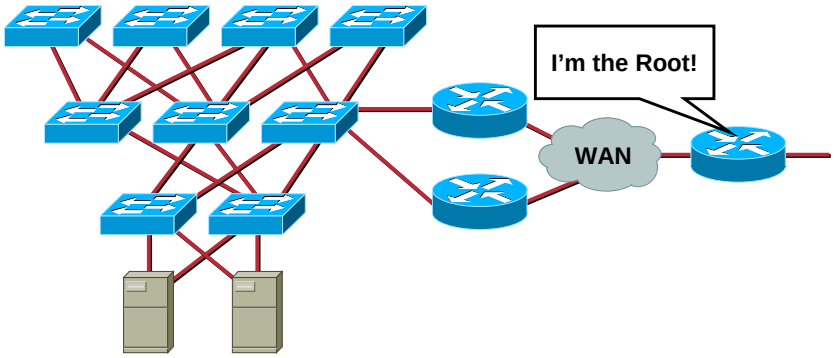

Spanning Tree—Solution
Root Bridge

Cisco.com

- The root of the tree should be in the very core of the network; typically server farm
- Use the set spantree root or spanning-tree vlan <vlan> root macro on the root switch

Spanning Tree—Commands
Set Spantree Root

Cisco.com

- A macro that sets the bridge priority parameter to 8192 for the specified VLANs, 16384 for secondary roots
- Root bridge dictates hello interval for entire network

```
Switch> (enable) set spantree root 1 dia 2 hello 2
VLAN 1 bridge priority set to 8192.
VLAN 1 bridge max aging time set to 10.
VLAN 1 bridge hello time set to 2.
VLAN 1 bridge forward delay set to 7.
Switch is now the root switch for active VLAN 1.
```


Spanning Tree—Commands
Spanning-tree VLAN Root

Cisco.com

- The same macro applied to a Native switch.

```
Native(config)#Spanning-tree vlan 1 root primary
diameter 2 hello-time 2
 VLAN 1 bridge priority set to 8192
 VLAN 1 bridge max aging time unchanged at 10
 VLAN 1 bridge hello time unchanged at 2
 VLAN 1 bridge forward delay unchanged at 7
```

Spanning Tree—Solution UDLD

Cisco.com

- Fiber links can become Uni-Directional for many reasons, and UDLD will provide protection for most of those reasons
- For the best protection use UDLD Aggressive Mode
- Completely available in CatOS 5.4(3)

Spanning Tree—Commands Set/Show UDLD

Cisco.com

- UDLD is a global configuration, but Aggressive Mode is configured per port

```
CatOS> (enable) set udld enable
UDLD enabled globally
CatOS > (enable) set udld aggressive-mode enable 1/1
Aggressive UDLD enabled on port 1/1.
CatOS > (enable) sh udld port
UDLD                    : enabled
Message Interval        : 15 seconds
Port    Admin Status    Aggressive Mode     Link State
--------  ------------  ---------------     ----------------
 1/1      enabled       enabled             undetermined
 1/2      enabled       disabled            undetermined
```

Spanning Tree—Commands UDLD Enable/Agressive

Cisco.com

- Native can have standard or aggressive configured globally and per port exceptions.

```
Native(config)#udld enable
Native(config)#interface gigabitEthernet 1/1
Native(config-if)#udld aggressive
```


Spanning Tree—Commands Show UDLD

Cisco.com

```
Native#show udld gigabitEthernet 1/1
Interface Gi1/1
---
Port enable administrative configuration setting: Enabled / in aggressive mode
Port enable operational state: Enabled / in aggressive mode
Current bidirectional state: Bidirectional
Current operational state: Advertisement - Single neighbor detected
Message interval: 15
Time out interval: 5

   Entry 1
   ---
   Expiration time: 35
   Device ID: 1
   Current neighbor state: Bidirectional
   Device name: SAL06090FCJ
   Port ID: Gi1/1
   Neighbor echo 1 device: SAD044204Y8
   Neighbor echo 1 port: Gi1/1

   Message interval: 5
   CDP Device name: ls-7603-16a
```

Spanning Tree—Solution Root Guard

Cisco.com

- Root Guard forces a Layer 2 LAN interface to be a designated port, and if any device accessible through the interface becomes the root bridge, root guard puts the interface into the root-inconsistent (blocked) state.
- Available in CatOS 6.1(1)

Spanning Tree—Solution BPDU Guard

Cisco.com

- PortFast BPDU guard can prevent loops by moving PortFast-configured interfaces that receive BPDUs to errdisable, rather than running spanning tree across that port
- This keeps ports configured with portfast from being incorrectly connected to another switch
- Available in CatOS 5.4(1)

Spanning Tree—Feature Backbonefast

Cisco.com

- A feature which can bypass waiting for Max Age timer to expire
- Uses a Remote Link Query (RLQ) PDU to test path to the root upon receipt of an inferior BPDU
- Available in CatOS 4.1(1)

Spanning Tree—Feature Uplinkfast

Cisco.com

- A feature that when enabled ONLY on an "Access Layer" switch tunes Spanning Tree such that the blocking redundant uplink is activated immediately upon failure of the primary root port
- Available in CatOS 3.1(1)

Spanning Tree—Feature Uplinkfast

Cisco.com

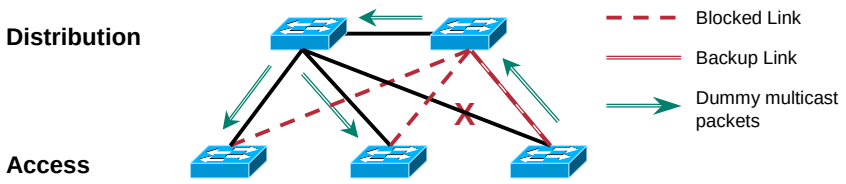

Spanning Tree—Commands Guards and Fasts

Cisco.com

- **Notice how uplinkfast changes STP parameters to try and keep the switch from being root**

```
CatOS> (enable) set spantree guard root 1/1
Do you want to continue (y/n) [n]? y
Rootguard on port 1/1 is enabled.
Warning!! Enabling rootguard may result in a topology change.

CatOS> (enable) set spantree portfast bpdu-guard enable
Spantree portfast bpdu-guard enabled on this switch.

CatOS> (enable) set spantree backbonefast enable
Backbonefast enabled for all VLANs.

CatOS> (enable) set spantree uplinkfast enable
VLANs 1-4094 bridge priority set to 49152.
The port cost and portvlancost of all ports set to above 3000.
Station update rate set to 15 packets/100ms.
uplinkfast all-protocols field set to off.
uplinkfast enabled for bridge.
```

Spanning Tree—Commands Guards and Fasts

Cisco.com

- **And in Native…**

```
Native(config)#interface gigabitEthernet 1/1
Native(config-if)#spanning-tree guard root
Native(config-if)#exit
Native(config)#spanning-tree port bpduguard
Native(config)#spanning-tree backbonefast
Native(config)#spanning-tree uplinkfast
```


Spanning Tree—Commands Show Spantree Summary

Cisco.com

```
CatOS> (enable) show spantree summary
MAC address reduction: disabled
Root switch for vlans: 1.
Global loopguard is disabled on the switch.
Global portfast is disabled on the switch.
BPDU skewing detection disabled for the bridge.
BPDU skewed for vlans:  none.
Portfast bpdu-guard enabled for bridge.
Portfast bpdu-filter disabled for bridge.
Uplinkfast disabled for bridge.
Backbonefast enabled for bridge.

Summary of connected spanning tree ports by vlan

VLAN  Blocking  Listening  Learning  Forwarding  STP Active
------ ---------- ---------- ---------- ---------- ----------
   1          0          0          0          4          4

        Blocking  Listening  Learning  Forwarding  STP Active
------ ---------- ---------- ---------- ---------- ----------
Total          0          0          0          4          4
```

Spanning Tree—Commands Show Spantree Summary

Cisco.com

```
Native#show spanning-tree summary
Root bridge for: VLAN0010.
Extended system ID is enabled
PortFast BPDU Guard is enabled
EtherChannel misconfiguration guard is disabled
UplinkFast is disabled
BackboneFast is enabled
Default pathcost method used is short

Name                 Blocking  Listening  Learning  Forwarding  STP Active
-------------------- --------- ---------- --------- ----------- ----------
VLAN0001                    0          0         0           2          2
VLAN0010                    0          0         0           1          1
VLAN1002                    0          0         0           1          1
VLAN1003                    0          0         0           1          1
VLAN1004                    0          0         0           1          1
VLAN1005                    0          0         0           1          1
-------------------- --------- ---------- --------- ----------- ----------
6 VLANs                     0          0         0           7          7
```


Spanning Tree—Be Proactive

Cisco.com

- Know your network before trouble begins
- Know where the root of each tree is
- Know where all the blocked ports are
- Don’t do any tuning you don’t understand

Spanning Tree—Recovery

Cisco.com

- DO NOT POWER OFF SWITCHES
- Check and physically remove the connections to the ports that should be blocking—this will get your network functional again
- Set up remote access to your network and call the TAC

Spanning Tree—Be Careful

Cisco.com

- Spanning Tree is rather complex
- Do not “tune” Spanning Tree parameters manually unless you fully understand, have a plan, and have a reason
- Abide by recommended guidelines for number of Spanning Tree instances

Spanning Tree—Instances

Cisco.com

= (# of non-ATM trunks * # of vlans on trunk)

+ (# of ATM trunks * vlans on trunk *2)

+ # of non-trunking ports

	Max Recommended Instances
2950	64 VLANs
3550	128 VLANs
4000 Sup I or II	1,500
4000 Sup III	2,000
5000 Sup I	400
5000 Sup II	1,500
5000 Sup IIG or IIIG	1,800
5000 Sup III	4,000
6000 Sup I	4,000
6000 Sup II	20,000

Agenda

Cisco.com

- Terminology
- Spanning Tree
- Access Control Lists
- Multicast
- Multi-Layer Switching
- Redundancy

Layer 3 ACLs—Problem
RACL Is Being Software Switched

Cisco.com

- A new IP access-group is configured on a VLAN interface, and now the MSFC is having to route many more packets on that interface then it did before the access-group was applied

Layer 3 ACLs—Solution
RACL Is Being Software Switched

Cisco.com

- If an ACL is configured with the log key word, any hits on that ACE will be passed to the MSFC to create the log entry
- If ip unreachables in enabled for a vlan interface, then any ACL denys will be passed to the MSFC to send the ICMP unreachable message

Layer 3 ACLs—Solution
RACL Is Being Software Switched

Cisco.com

- **Also check to make sure that the ACL was programmed into hardware by the MSFC**

```
IOS(6k)#show fm summary
Interface: Vlan1
  TCAM screening for features is INACTIVE inbound
Interface: Vlan2
  TCAM screening for features is ACTIVE outbound
  TCAM screening for features is ACTIVE inbound
```


Layer 3 ACLs—Problem
Errors When Applying RACL

Cisco.com

- **When you apply an access-list to a VLAN interface one of the following messages appear:**

```
IOS(6k)(config)#interface vlan 1
IOS(6k)(config-if)#ip access-group 100 in
<router>%ACL-3-NOLOU:Acl engine is out of logical operation unit
<router>%ACL-3-RACLMAPCOMMITFAIL:Failed to map Router ACL to VLAN 1
<switch>%FM-4-TCAM_LOU: Hardware TCAM LOU capacity exceeded
<switch>%FM-4-RACL_REDUCED: Interface Vlan1 routed traffic will be
software switched in ingress direction(s)
…OR…
<router>%ACL-3-TCAMFULL:Acl engine TCAM table is full
<router>%ACL-3-RACLMAPCOMMITFAIL:Failed to map Router ACL to VLAN 1
<switch>%FM-4-TCAM-ENTRY: Hardware TCAM entry capacity exceeded
<switch>%FM-4-RACL_REDUCED: Interface Vlan1 routed traffic will be
software switched in ingress direction(s)
```

Layer 3 ACLs—Solution
Errors When Applying RACL

Cisco.com

- The “Acl engine is out of logical operation unit” message is caused when too many Layer 4 operators are included in a ACL
  - L4Ops are gt (1/2 LOU), lt (1/2 LOU), neq (1/2 LOU), range (1 LOU)
- A maximum of 9 L4Ops per ACL and 32 LOUs per system (6500)
- A maximum of 6 L4Ops per ACL (4000 SupIII)

Layer 3 ACLs—Solution
Errors When Applying RACL

Cisco.com

- The “Acl engine TCAM table is full” message is caused when the total ACEs for all ACLs on the 6500 are more than what can be programmed into the TCAM
  - The 3550 has a maximum of 4K ACEs
  - The 4000 SupIII has a maximum of 16K ACEs input and 16K ACEs output
  - The PFC has a maximum of 16K ACEs, shared by RACLs, VACLs, and QoS
  - The PFC2 has a maximum of 32K ACEs, shared by RACLs and VACLs

Layer 3 ACLs—Solution
Errors When Applying RACL

Cisco.com

- **Change the algorithm used to merge all the programmed ACLs together into hardware TCAM.**
  - **Order Independent Merge – the original merge algorithm based on Binary Decision Diagrams (BDD)**
  - **Order Dependent Merge – a new merge algorithm available in CatOS 7.1(1) and Native IOS 12.1(8a)EX based on "Value, Mask, Result" (VMR)**

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


Layer 3 ACLs—Example
Errors When Applying RACL

Cisco.com

```
Native(config)#interface gigabitEthernet 1/1
Native(config-if)#ip access-group security_RACL in
Native(config-if)#ip nat outside
Native(config-if)#exit
Native(config)#ip nat outside source nat_RACL pool natpool
```

Order Independent Merge

```
Native#show tcam counts
```

	Used	Free
Labels:	3	509
ACL_TCAM		
Masks:	136	3960
Entries:	187	32581

Order Dependent Merge

```
Native#show tcam counts
```

	Used	Free
Labels:	3	509
ACL_TCAM		
Masks:	21	4075
Entries:	104	32664

```
CatOS(6k)> show security acl resource-usage
```

Access Lists–Additional Info

Cisco.com

- **Understanding RACL/VACL/QoS ACL Hardware Resources in Catalyst 6000 Family Switches**

  http://www.cisco.com/warp/public/473/79.html

Agenda

Cisco.com

- **Terminology**
- **Spanning Tree**
- **Access Control Lists**
- **Multicast**
- **Multi-Layer Switching**
- **Redundancy**

Multicast–Problem
IGMP Snooping on a Routerless Segment

Cisco.com

- There are two issues involved with using IGMP Snooping on a routerless segment
  - IGMP normally only forwards IGMP reports to the multicast router port
  - Applications that send a join only on startup, but not after link-up, will never be re-added to the multicast group after a network disruption

Multicast–Solution
IGMP Snooping on a Routerless Segment

Cisco.com

- To get around the first issue, static multicast router ports will need to be configured toward the multicast sources

```
CatOS> (enable) set multicast router 3/45
```

```
Native(6k)(config)#interface vlan 1
Native(6k)(config-if)#ip igmp snooping mrouter interface gigabitEthernet 1/2
```

```
Native(4k/3550)(config)#ip igmp snooping vlan 1 mrouter interface gigabitEthernet 1/2
```

Multicast–Solution
IGMP Snooping on a Routerless Segment

Cisco.com

- The second problem is due to the lack of a IGMP Querier on the network
- Catalyst switches can be configured as IGMP Queriers
- Available in a Layer 2 device only in CatOS 7.1(1)

```
CatOS> (enable) set igmp querier enable 1
IGMP switch querier enabled for VLAN 1

CatOS> (enable) show igmp querier information
VLAN  Querier State        Query Tx Count  QI (seconds)  OQI (seconds)
----  -------------------- --------------  ------------  -------------
1     QUERIER              0               125           300
```


Multicast–Problem
Stream Loss to Groups of Clients

Cisco.com

- IGMP Fast Leave processing is enabled on a core Catalyst switch and a group of clients are attached via a non-IGMP/CGMP enabled access switch
- As soon as one of the clients leave the multicast group, all the clients on that access switch lose that multicast stream

Multicast–Solution
Stream Loss to Groups of Clients

Cisco.com

- With IGMP Fast Leave processing, as soon as we receive a IGMP leave from a port, we remove that port from that group, without sending a group specific query to determine if any other clients need that group on that port
- Do not enable IGMP Fast Leave processing on a switch unless all downstream ports connect to only a single device or a IGMP/CGMP enabled switch

Multicast–Commands
Show Multicast Groups

Cisco.com

```
CatOS> (enable) show multicast group
VLAN     Dest MAC/Route Des              [CoS]    Destination Ports or VCs /
                                                  [Protocol Type]
--------  ------------------------------  --------  ------------------------------------
2         01-00-5e-00-01-28                         15/1

Total Number of Entries = 1
```

```
Native#show ip igmp groups
IGMP Connected Group Membership
Group Address    Interface       Uptime    Expires   Last Reporter
224.0.1.40       Vlan10          02:26:48  00:02:25  192.168.1.13
```

Multicast–Commands Show Multicast Router

Cisco.com

```
CatOS> (enable) show multicast router
Port      Vlan
--------- ----------------
 3/45     1,10
15/1      2

Total Number of Entries = 2
'*' - Configured
'+' - RGMP-capable
```

```
Native#show ip igmp snooping mrouter
Vlan      ports
------    -----
  1       Gi1/1(dynamic)
 13       Gi1/1(static)
```


Multicast MLS—Problem 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

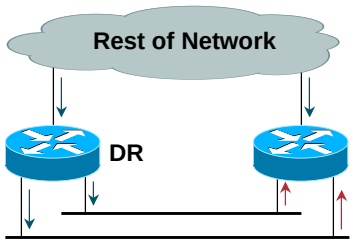

Multicast MLS—Solution
High CPU on non-DR Router

Cisco.com

- In a dual chassis with single supervisor configuration, a router ACL can be applied to kill the multicast non-RPF traffic for pim sparse mode

```
interface Vlan2
 ip address 172.10.1.2 255.255.255.0
 ip access-group 105 in
!
access-list 105 permit ip 172.10.1.0 0.0.0.255 any
access-list 105 permit ip any 224.0.0.0 0.0.0.255
access-list 105 permit ip any 224.0.1.0 0.0.0.255
access-list 105 deny   ip any 224.0.0.0 15.255.255.255
```


Multicast MLS—Solution
High CPU on non-DR Router

Cisco.com

- In Catalyst 6000 Sup2 and 4000 Sup3 environments, the solution is to use the non-RPF Multicast Fast Drop (MFD) feature
- This feature is not supported on a Catalyst 6000 Sup1
- Available in CatOS 6.2(1)

Multicast MLS—Solution
High CPU on non-DR Router

Cisco.com

- Multicast Fast Drop is enabled by default in CatOS
- Enabling Multicast Fast Drop in Native

```
Native(6k)(config)#mls ip multi non-rpf cef
```

```
Native(4k)(config)#ip mfib fastdrop
```


Agenda

Cisco.com

- Terminology
- Spanning Tree
- Access Control Lists
- Multicast
- Multi-Layer Switching
- Redundancy

MLS––Demand-Based Theory of Operation

Cisco.com

- Used on Catalyst 5000 and 6000 Supervisor I
- Both the Candidate Packet and the Enabler Packet must be present for a flow to be built

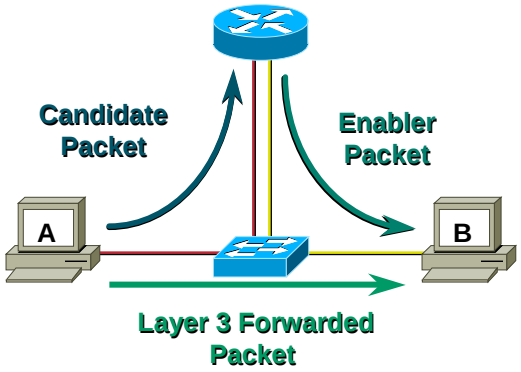

MLS––CEF-Based Theory of Operation

Cisco.com

- Used on Catalyst 3550, 4000 Supervisor III, and 6000 Supervisor II
- The hardware forwarding engine on the Switch Processor is programmed by the software on the Route Processor with a copy of the software FIB and adjacency table.
- The SP depends on the RP for routing and ARP information corresponding to FIB and adjacency updates.

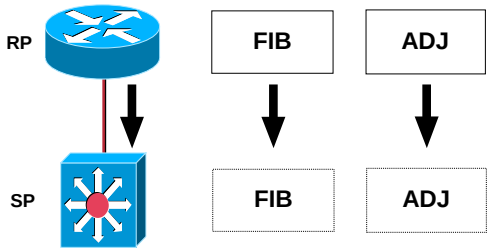

Multi-Layer Switching—Problem Connectivity Loss

Cisco.com

- Communication between two network stations is not functioning properly
- This loss of connectivity can be traced back to a Catalyst switch performing Multi-Layer Switching

Multi-Layer Switching—Step 1 Connectivity Loss

Cisco.com

- First, check the software routing tables and arp caches to verify full ip connectivity

```
IOS#show ip route 192.168.50.0
Routing entry for 192.168.50.0/24
  Known via "eigrp 192", distance 90, metric 3072, type internal
  Redistributing via eigrp 192
  Last update from 192.168.1.50 on Vlan10, 01:23:08 ago
  Routing Descriptor Blocks:
  * 192.168.1.50, from 192.168.1.50, 01:23:08 ago, via Vlan10
      Route metric is 3072, traffic share count is 1
      Total delay is 20 microseconds, minimum bandwidth is 1000000 Kbit
      Reliability 255/255, minimum MTU 1500 bytes
      Loading 1/255, Hops 1

IOS#show ip arp 192.168.1.50
Protocol  Address       Age (min) Hardware Addr   Type    Interface
Internet  192.168.1.50     84     0007.0e8f.088a  ARPA    Vlan10
```

Multi-Layer Switching—Step 2 Connectivity Loss

Cisco.com

- Second, verify that the CEF and Adjacency table match the Routing Table and ARP cache

```
IOS#show ip cef 192.168.50.0
192.168.50.0/24, version 43, epoch 0, cached adjacency 192.168.1.50
0 packets, 0 bytes
  via 192.168.1.50, Vlan10, 0 dependencies
    next hop 192.168.1.50, Vlan10
    valid cached adjacency

IOS#show adjacency vlan 10 detail
Protocol        Interface       Address
IP              Vlan10          192.168.1.50(7)
                                30 packets, 3000 bytes
                                00070E8F088A
                                00055F33DC0A0800
                                ARP       02:34:56
                                Epoch: 0
```


MLS—Step 3a (Cat6 Sup1 Hybrid) Connectivity Loss

Cisco.com

- Next, verify the hardware has the vlan configured and that the hardware tables are not overflowing

```
CatOS> (enable) show mls
Total packets switched = 50
Total Active MLS entries = 2
  MSFC 14.18.3.184 (Module 15) entries = 2
  MSFC 14.18.3.185 (Module 16) entries = 0
Long-duration flows aging time = 1920 seconds
IP Multilayer switching aging time = 256 seconds
IP Multilayer switching fast aging time = 0 seconds, packet threshold = 0
IP Current flow mask is Destination flow
Active IP MLS entries = 2
Netflow Data Export version: 7
Netflow Data Export disabled
Netflow Data Export port/host is not configured.
Total packets exported = 0

IP MSFC ID          Module  XTAG  MAC                 Vlans
------------------  ------  ----  ------------------  ----------
14.18.3.184         15      1     00-02-7e-26-df-44   1,12,10
14.18.3.185         16      2     00-d0-bc-ee-48-e4   1,12,10
```

MLS—Step 3b (Cat6 Sup1 Hybrid) Connectivity Loss

Cisco.com

- Then check to see if the flow has been installed into the MLS cache and matches the CEF information

```
CatOS> (enable) show mls entry ip destination 192.168.50.1
Destination-IP       Source-IP  Prot       DstPrt     SrcPrt     Destination-Mac    Vlan
Edst        Esrc     Dport      Sport      Stat-Pkts  Stat-Bytes Uptime     Age
---------------      ---------- -----      ------     ------     -----------------  ----
----        ----     ---------  ---------  ---------- ---------- --------   --------
MSFC 14.18.3.184 (Module 15):
192.168.50.1         -          -          0          0          00-07-0e-8f-08-8a  10
ARPA        ARPA     3/45       1/3        4          336        00:00:33   00:00:33
MSFC 14.18.3.185 (Module 16):

Total entries displayed: 1
```


MLS—Step 3 (Cat6 Sup1 Native) Connectivity Loss

Cisco.com

- Next, verify that the hardware tables are not overflowing and that a flow has been installed into the MLS cache and matches the CEF information

```
Native(6k)#sh mls ip count
Displaying Netflow entries in Supervisor Earl

 Number of shortcuts = 1

Native(6k)#show mls ip destination 192.168.50.1
Displaying Netflow entries in Supervisor Earl
DstIP                SrcIP                Prot:SrcPort:DstPort          Src i/f:AdjPtr
-------------------------------------------------------------------------------------------------------------------
Pkts       Bytes      Age        LastSeen            Attributes
-------------------------------------------------------------------------
192.168.50.1         0.0.0.0              0     :0       :0                 0       : 0
5          500        18         14:17:13            L3 - Dynamic
```

MLS—Step 3 (Cat6 Sup2 Hybrid) Connectivity Loss

Cisco.com

- Next, verify that the entries in the hardware tables have been installed correctly and match the CEF information

```
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
```


MLS—Step 3 (Cat6 Sup2 Native) Connectivity Loss

Cisco.com

- Next, verify that the entries in the hardware tables have been installed correctly and match the CEF information

```
Native(6k)#show mls cef 192.168.50.0

Native(6k)-sp#
Index      Prefix          Mask            Adjacency
6405       192.168.50.0    255.255.255.0   0007.0e8f.088a

Native(6k)#show mls cef adjacency mac-address 0007.0e8f.088a

Native(6k)-sp#
Index 17424 :  mac-sa: 0005.5f33.dc0a, mac-da: 0007.0e8f.088a
               interface: Vl10, mtu: 1514
               packets: 0000000000000000, bytes: 0000000000000000
```

MLS—Step 3a (Cat4 Sup3 Native) Connectivity Loss

Cisco.com

- Verify that the entries in the software tables used by hardware have been installed correctly and match the CEF information

```
Native(4k)#show platform software ip route network 192.168.50.0 255.255.255.0
---------- Unicast /24 routes
075748: u (10221) 192.168.50.0/24 -> (4083) v: 10 192.168.1.50
                  00:07:0E:8F:08:8A (IrmFibAdjTypeNormal)

Native(4k)#show platform software ip route host 192.168.1.50
---------- Unicast /32 routes
016426: u (10222) 192.168.1.50/32 -> (4083) v: 10 192.168.1.50
                  00:07:0E:8F:08:8A (IrmFibAdjTypeNormal)
```


MLS—Step 3b (Cat4 Sup3 Native) Connectivity Loss

Cisco.com

- Then verify that the hardware was programmed correctly from the software tables and match the CEF information

```
Native(4k)#show platform hardware ip route network 192.168.50.0 255.255.255.0
---------- Unicast /24 routes
075748: 192.168.50.0/24 -> rpf: 65535 adj: 32759 v: 10 p: 248 (Gi1/1)
                  sa: 00:05:9B:F6:70:FF da: 00:07:0E:8F:08:8A

Native(4k)#show platform hardware ip route host 192.168.1.50
---------- Unicast /32 routes
016426: 192.168.1.50/32 -> rpf: 65535 adj: 32759 v: 10 p: 248 (Gi1/1)
                  sa: 00:05:9B:F6:70:FF da: 00:07:0E:8F:08:8A
```

Multi-Layer Switching—Solution Connectivity Loss

Cisco.com

- If all the output is consistent, but packets are not flowing in the proper direction then there is probably an inconsistency in the Layer 3 TCAMs
- Recovery can usually be achieved by clearing the IP route table.
- On a Catalyst 6000 Supervisor II be sure to run 6.3(3) or 12.1(8b)E6 and later code. This enables the new TCAM consistency checker.

Multi-Layer Switching—Problem High CPU on Route Processor

Cisco.com

- The switch is configure correctly
  - route-caches are enabled
  - hardware switching has not been disabled
- Still a lot of traffic is being switched by the Route Processor, causing higher then expected CPU utilization

Multi-Layer Switching—Solution High CPU on Route Processor

Cisco.com

- High CPU on the RP is usually caused by one of the following
  - Large number of broadcasts on one of more interfaces
  - Multicast non-RPF
  - Instability of the routing table
  - SNMP management queries
  - Packet switching

Multi-Layer Switching—Solution High CPU on Route Processor

Cisco.com

- Looking at packet switching. Why isn't it being hardware switched?
  - Overflowing hardware tables
  - Using a feature that isn't supported in hardware

MLS—Restrictions (Cat6 Sup1) High CPU on Route Processor

Cisco.com

- There are 32K MLS cache entries, these are shared with the microflow policer.
- 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
- The following features are not supported in hardware
  - Unicast RPF check
  - Per-packet load balancing
  - NBAR, WCCP, IP accounting

MLS—Restrictions (Cat6 Sup2) High CPU on Route Processor

Cisco.com

- There are 256K route entries, which are cut in half, if Unicast RPF is enabled
- The following IP packets cannot be forwarded in hardware
  - Packets with IP options set
  - Packets with TTL<=1
  - Packets that require fragmentation
- The following features are not supported in hardware
  - Per-packet load balancing
  - Different MAC BIAs per interface

MLS—Restrictions (Cat4 Sup3) High CPU on Route Processor

Cisco.com

- 128K route entries
- The following IP packets cannot be forwarded in hardware
  - Packets with IP options set
  - Packets with TTL<=1
  - Packets that require fragmentation
- The following features are not supported in hardware
  - Policy-based routing
  - Unicast RPF check
  - Manual MAC address assignment
  - Per-packet load balancing
  - NBAR, WCCP, IP accounting

MLS—Restrictions (Cat3550) High CPU on Route Processor

Cisco.com

- By default between 16K and 24K route entries
- The following IP packets cannot be forwarded in hardware
  - Packets with IP options set
  - Packets with TTL<=1
  - Packets that require fragmentation
- The following features are not supported in hardware
  - Policy-based routing
  - Unicast RPF check
  - Per-packet load balancing
  - NBAR, WCCP, IP accounting

Agenda

Cisco.com

- Terminology
- Spanning Tree
- Access Control Lists
- Multicast
- Multi-Layer Switching
- Redundancy

Redundancy—Problem Long Fail-Over Times

Cisco.com

- When a supervisor fails over for any reason, the network should regain full availability as quickly as possible
- There are multiple options available on the Catalyst 6000 products to provide this ability

Redundancy—Options
Long Fail-Over Times

Cisco.com

- For single chassis dual supervisor

  Dual Routers (HSRP) - Hybrid
  Single Router Mode - Hybrid

  EHSA - Native
  RTR+ - Native

- For dual chassis single supervisors

  Dual Routers (HSRP)

- For dual chassis dual supervisors

  Quad Routers (HSRP) - Hybrid
  Single Router Mode (HSRP) - Hybrid

  EHSA (HSRP) - Native
  RTR+ (HSRP) - Native

Redundancy—Solution
Hybrid with Dual Supervisors

Cisco.com

- For all Hybrid configurations make sure the supervisor has High-Availability enabled AND that it is operational

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

Redundancy—Solution Hybrid with Dual Routers

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
```


Redundancy—Problem Hybrid with Dual Routers

Cisco.com

- In dual MSFC systems one MSFC is selected as the designated router (MSFC)

```
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
```

Redundancy—Solution Hybrid with Dual Routers

Cisco.com

- The Designated MSFC does all programming of the Layer 3 switching hardware on the supervisor
  - VLANs that are only configured on the non-designated MSFC will not be routed in hardware and might have other connectivity issues
  - ACLs that are only configured on the non-designated MSFC will not applied
  - On Supervisor 2s routes that appear on only on the non-designated MSFC do not get routed
  - Flexwan and Optical Service Modules will have their interfaces seen only by the designated MSFC

Redundancy—Solution Hybrid with Config-Sync

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
```

Redundancy—Solution Hybrid with Config-Sync

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
```


Redundancy—Problem Hybrid with Single-Router-Mode

Cisco.com

- In a dual chassis with dual supervisor redundancy configuration, by default there are four active routers.
  - This is a management hassle
  - If these routers do redistribution, then there will be 4 redistributed routes for each real route, instead of the “normal” 2 routes
  - HSRP gets more complex

Redundancy—Solution Hybrid with Single-Router-Mode

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

```
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
```


Redundancy—Solution Hybrid with Single-Router-Mode

Cisco.com

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

```
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
```

Redundancy—Solution Hybrid with Single-Router-Mode

Cisco.com

- Single-Router-Mode provides a stateful failover for the Layer 3 switching hardware on the Supervisor, so there is minimal outage for any packet flowing through the switch
- Multicast Layer 3 switching was not stateful until CatOS 7.1(1), and then still not for a Supervisor 1/MSFC1.

Redundancy—Solution Native with EHSA

Cisco.com

- Native redundancy defaults to Enhanced High System Availability
- Provides for 1 to 2 minute failover times, which is dependent on what cards are in the chassis, as all cards are reset on failover

Redundancy—Solution Native with RPR+

Cisco.com

- Route Processor Redundancy Plus provides for much faster failover as cards are no longer reset on failover
- Provides for approximately 30 second failover times independent of what cards are in the chassis
- Available in 12.1(11b)EX

Redundancy—Solution 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

<...some output deleted...>
```

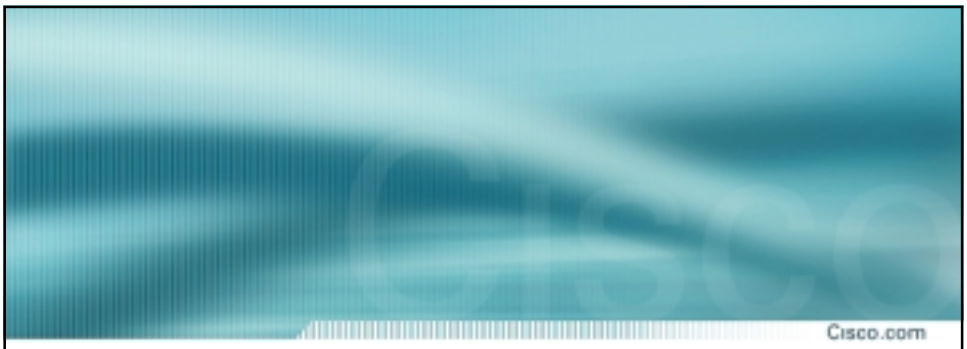

Troubleshooting Catalyst® Switches

Session RST-322

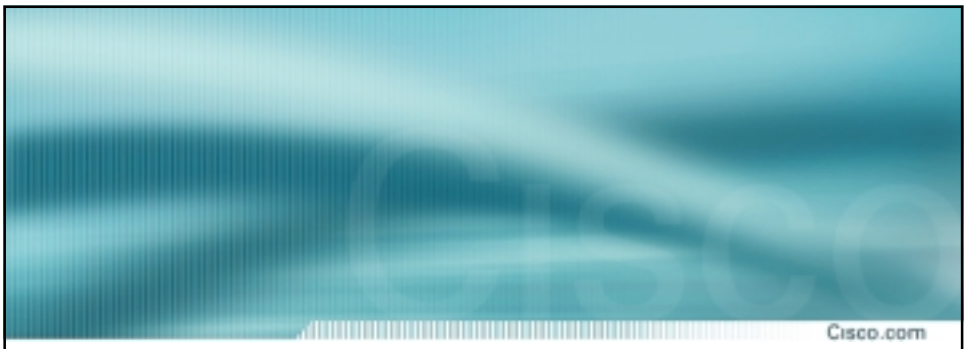

Please Complete Your Evaluation Form

Session RST-322