



Building Cisco Multilayer Switched Networks (BCMSN)

Spanning-Tree Protocol (STP)

<http://www.INE.com>

Switching Logic Review

- Layer 2 switches use the CAM table to switch traffic based on destination MAC address
- To populate the CAM table the following logic is used
 - A frame from X going to Y is received on port 1
 - Insert X into the CAM table via port 1
 - Flood the frame out all ports in the VLAN except 1
 - A return frame from Y going to X is received on port 2
 - Insert Y into the CAM table via port 2
 - Subsequent traffic does not require flooding

Copyright © 2009 InternetNetwork Expert, Inc
www.INE.com



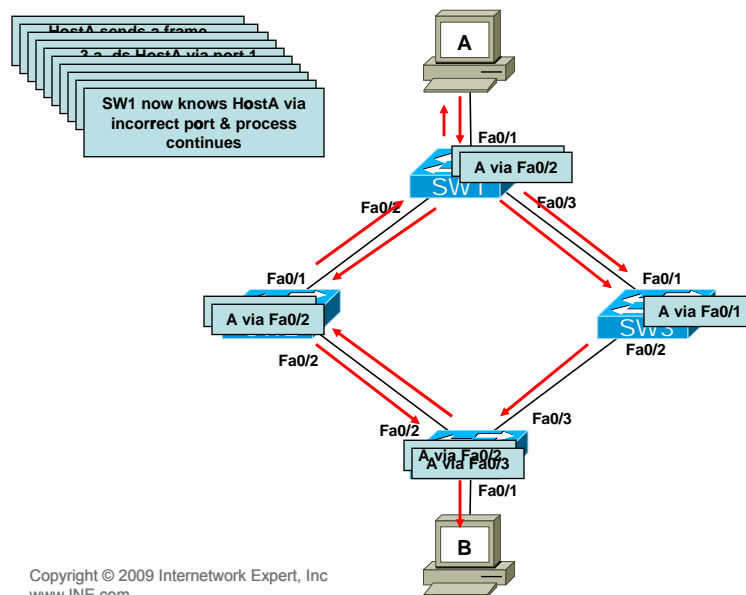
Switching Loop Problems

- When redundant paths exist in the layer 2 network, CAM population logic breaks down and frames are switched out the wrong interfaces
- Looping frames, especially broadcasts, can quickly overwhelm all links with 100% utilization

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



Switching Loop Example



Copyright © 2009 Internetwork Expert, Inc
www.INE.com



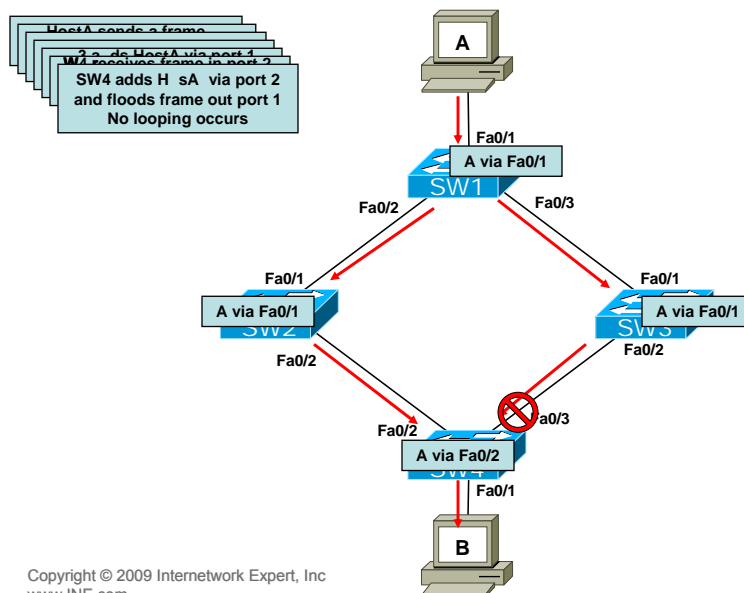
Spanning-Tree Protocol Overview

- STP solves the looping problem by “blocking” redundant paths
 - Blocked links cannot forward traffic or use the CAM table
 - Same effect as removing or shutting down the link
- Since STP is dynamic, layer 2 network can reconverge around network failures
- Standards based per 802.1D

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



Switching Loop Prevention Example



How STP Loop Prevention Works

- All devices agree on a reference point in the network
 - Called the “root bridge”
- Device directly downstream of the root bridge performs the following...
 - Select one upstream facing port to forward traffic towards the root bridge
 - Called the “root port”
 - All other upstream facing ports are disabled
 - Called “blocking” ports
 - All downstream facing ports are called “designated” ports
- Next downstream device performs the same, selecting one upstream facing root port
- Repeat until entire loop-free tree is built

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



How STP Works

- Exchange bridge and link attributes
- Elect one Root Bridge
- Elect one Root Port per bridge
- Elect Designated Ports

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



STP Advertisements

- Uses Bridge Protocol Data Units (BPDUs)
 - Sent as multicast frames between adjacent bridges (0180.C200.0000)
- Used to advertise bridge and link attributes
 - Root ID
 - Root Path Cost
 - Bridge ID
 - Port ID
 - Timers
- Two types of BPDUs
 - Configuration BPDUs
 - Topology Change Notification (TCN) BPDUs

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



Root Bridge Election

- Based on the lowest Bridge ID (BID) in the network
- BID is an 8-byte field that contains
 - Bridge Priority
 - 0-65535
 - Defaults to 32768
 - MAC Address
- New standard splits Bridge Priority into two fields
 - AKA “MAC address reduction” feature
 - Bridge Priority
 - 4 high order bits
 - 0 - 61440 in increments of 4096
 - System ID Extension
 - 12 low order bits
 - 0 - 4095
- Lowest BID in the network becomes everyone’s Root ID (RID) in their BPDUs

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



Root Port Election

- Port closest to the Root Bridge
 - Root is always upstream
- Elected based on lowest Root Path Cost
 - Cumulative cost of all links to get to the root
 - Cost based on inverse bandwidth
 - i.e. higher bandwidth, lower cost
 - Not linear
- If tie in cost...
 - Choose lowest upstream BID
 - Choose lowest upstream Port ID

Copyright © 2009 Internetwork Expert, Inc
www.INE.com




Designated Port Election

- Ports downstream facing away from Root Bridge
- Like Root Port, elected based on...
 - Lowest Root Path Cost
 - Lowest BID
 - Lowest Port ID
- All other ports go into “blocking” mode
 - Receives BPDUs
 - Discards all other traffic
 - Cannot send traffic
- Blocking ports are the key to the loop free topology

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



STP Path Selection Example



Copyright © 2009 Internetwork Expert, Inc
www.INE.com



STP Path Selection Verification (SW1)

```
SW1#show spanning-tree vlan 1
```

```
VLAN0001
```

```
Spanning tree enabled protocol ieee
```

```
Root ID    Priority    32769  
Address    0009.433c.a380  
Cost       38  
Port       18 (FastEthernet0/16)  
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

```
Bridge ID  Priority    32769 (priority 32768 sys-id-ext 1)  
Address    0019.56c8.4e80  
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec  
Aging Time 15
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
Fa0/13	Altn	BLK	19	128.15	P2p
Fa0/16	Root	LRN	19	128.18	P2p

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



STP Path Selection Verification (SW2)

```
SW2#show spanning-tree vlan 1
```

```
VLAN0001
```

```
Spanning tree enabled protocol ieee
```

```
Root ID    Priority    32769  
Address    0009.433c.a380  
Cost       19  
Port       21 (FastEthernet0/19)  
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

```
Bridge ID  Priority    32769 (priority 32768 sys-id-ext 1)  
Address    0019.aa7e.ea00  
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec  
Aging Time 15
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
-----	-----	-----	-----	-----	-----
Fa0/13	Desg	FWD	19	128.15	P2p
Fa0/19	Root	FWD	19	128.21	P2p

Copyright © 2009 Internetnetwork Expert, Inc
www.INE.com



STP Path Selection Verification (SW3)

```
SW3#show spanning-tree vlan 1
```

```
VLAN0001
```

```
Spanning tree enabled protocol ieee
```

```
Root ID    Priority    32769  
Address    0009.433c.a380  
Cost       19  
Port       19 (FastEthernet0/19)  
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec
```

```
Bridge ID  Priority    32769 (priority 32768 sys-id-ext 1)  
Address    000a.f4f3.e780  
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec  
Aging Time 15
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
-----	-----	-----	-----	-----	-----
Fa0/13	Desg	FWD	19	128.13	P2p
Fa0/19	Root	FWD	19	128.19	P2p

Copyright © 2009 Internetnetwork Expert, Inc
www.INE.com



STP Path Selection Verification (SW4)

```
SW4#show spanning-tree vlan 1
```

```
VLAN0001
```

```
Spanning tree enabled protocol ieee
```

```
Root ID    Priority    32769  
Address    0009.433c.a380  
This bridge is the root  
Hello Time 2 sec    Max Age 20 sec    Forward Delay 15 sec
```

```
Bridge ID  Priority    32769 (priority 32768 sys-id-ext 1)  
Address    0009.433c.a380  
Hello Time 2 sec    Max Age 20 sec    Forward Delay 15 sec  
Aging Time 15
```

Interface	Role	Sts	Cost	Prio.Nbr	Type
-----	---	---	---	-----	-----
Fa0/16	Desg	FWD	19	128.16	P2p
Fa0/19	Desg	FWD	19	128.19	P2p

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



STP Verification Detail (SW1)

```
SW1#show spanning-tree detail
```

```
VLAN0001 is executing the ieee compatible Spanning Tree protocol  
Bridge Identifier has priority 32768, sysid 1, address 0019.56c8.4e80  
Configured hello time 2, max age 20, forward delay 15  
Current root has priority 32769, address 0009.433c.a380  
Root port is 18 (FastEthernet0/16), cost of root path is 38  
Topology change flag not set, detected flag not set  
Number of topology changes 1 last change occurred 00:03:34 ago  
  from FastEthernet0/13  
Times: hold 1, topology change 35, notification 2  
       hello 2, max age 20, forward delay 15  
Timers: hello 0, topology change 0, notification 0, aging 300
```

```
Port 15 (FastEthernet0/13) of VLAN0001 is blocking  
Port path cost 19, Port priority 128, Port Identifier 128.15.  
Designated root has priority 32769, address 0009.433c.a380  
Designated bridge has priority 32769, address 0019.aa7e.ea00  
Designated port id is 128.15, designated path cost 19  
Timers: message age 2, forward delay 0, hold 0  
Number of transitions to forwarding state: 0  
Link type is point-to-point by default  
BPDU: sent 6, received 116
```

```
Port 18 (FastEthernet0/16) of VLAN0001 is forwarding  
Port path cost 19, Port priority 128, Port Identifier 128.18.  
Designated root has priority 32769, address 0009.433c.a380  
Designated bridge has priority 32769, address 000a.f4f3.e780  
Designated port id is 128.13, designated path cost 19  
Timers: message age 2, forward delay 0, hold 0  
Number of transitions to forwarding state: 1  
Link type is point-to-point by default  
BPDU: sent 2, received 111
```

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



STP Verification Detail (SW2)

SW2#show spanning-tree detail

VLAN0001 is executing the ieee compatible Spanning Tree protocol
Bridge Identifier has priority 32768, sysid 1, address 0019.aa7e.ea00
Configured hello time 2, max age 20, forward delay 15
Current root has priority 32769, address 0009.433c.a380
Root port is 21 (FastEthernet0/19), cost of root path is 19
Topology change flag not set, detected flag not set
Number of topology changes 2 last change occurred 00:03:19 ago
 from FastEthernet0/19
Times: hold 1, topology change 35, notification 2
 hello 2, max age 20, forward delay 15
Timers: hello 0, topology change 0, notification 0, aging 300

Port 15 (FastEthernet0/13) of VLAN0001 is forwarding
Port path cost 19, Port priority 128, Port Identifier 128.15.
Designated root has priority 32769, address 0009.433c.a380
Designated bridge has priority 32769, address 0019.aa7e.ea00
Designated port id is 128.15, designated path cost 19
Timers: message age 0, forward delay 0, hold 0
Number of transitions to forwarding state: 1
Link type is point-to-point by default
BPDU: sent 117, received 6

Port 21 (FastEthernet0/19) of VLAN0001 is forwarding
Port path cost 19, Port priority 128, Port Identifier 128.21.
Designated root has priority 32769, address 0009.433c.a380
Designated bridge has priority 32769, address 0009.433c.a380
Designated port id is 128.16, designated path cost 0
Timers: message age 1, forward delay 0, hold 0
Number of transitions to forwarding state: 1
Link type is point-to-point by default
BPDU: sent 3, received 118

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



STP Verification Detail (SW3)

SW3#show spanning-tree detail

VLAN0001 is executing the ieee compatible Spanning Tree protocol
Bridge Identifier has priority 32768, sysid 1, address 000a.f4f3.e780
Configured hello time 2, max age 20, forward delay 15
Current root has priority 32769, address 0009.433c.a380
Root port is 19 (FastEthernet0/19), cost of root path is 19
Topology change flag not set, detected flag not set
Number of topology changes 3 last change occurred 00:03:12 ago
 from FastEthernet0/19
Times: hold 1, topology change 35, notification 2
 hello 2, max age 20, forward delay 15
Timers: hello 0, topology change 0, notification 0, aging 300

Port 13 (FastEthernet0/13) of VLAN0001 is forwarding
Port path cost 19, Port priority 128, Port Identifier 128.13.
Designated root has priority 32769, address 0009.433c.a380
Designated bridge has priority 32769, address 000a.f4f3.e780
Designated port id is 128.13, designated path cost 19
Timers: message age 0, forward delay 0, hold 0
Number of transitions to forwarding state: 1
Link type is point-to-point by default
BPDU: sent 114, received 2

Port 19 (FastEthernet0/19) of VLAN0001 is forwarding
Port path cost 19, Port priority 128, Port Identifier 128.19.
Designated root has priority 32769, address 0009.433c.a380
Designated bridge has priority 32769, address 0009.433c.a380
Designated port id is 128.19, designated path cost 0
Timers: message age 1, forward delay 0, hold 0
Number of transitions to forwarding state: 1
Link type is point-to-point by default
BPDU: sent 3, received 114

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



STP Verification Detail (SW4)

SW4#show spanning-tree detail

```
VLAN0001 is executing the ieee compatible Spanning Tree protocol
Bridge Identifier has priority 32768, sysid 1, address 0009.433c.a380
Configured hello time 2, max age 20, forward delay 15
We are the root of the spanning tree
Topology change flag not set, detected flag not set
Number of topology changes 1 last change occurred 00:03:47 ago
  from FastEthernet0/19
Times: hold 1, topology change 35, notification 2
      hello 2, max age 20, forward delay 15
Timers: hello 1, topology change 0, notification 0, aging 300

Port 16 (FastEthernet0/16) of VLAN0001 is forwarding
Port path cost 19, Port priority 128, Port Identifier 128.16.
Designated root has priority 32769, address 0009.433c.a380
Designated bridge has priority 32769, address 0009.433c.a380
Designated port id is 128.16, designated path cost 0
Timers: message age 0, forward delay 0, hold 0
Number of transitions to forwarding state: 1
Link type is point-to-point by default
BPDUs: sent 122, received 3

Port 19 (FastEthernet0/19) of VLAN0001 is forwarding
Port path cost 19, Port priority 128, Port Identifier 128.19.
Designated root has priority 32769, address 0009.433c.a380
Designated bridge has priority 32769, address 0009.433c.a380
Designated port id is 128.19, designated path cost 0
Timers: message age 0, forward delay 0, hold 0
Number of transitions to forwarding state: 1
Link type is point-to-point by default
BPDUs: sent 116, received 3
```

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



STP Port States

- When a bridge boots up, the initial STP convergence time is based on how long the device takes to transition through different port states
- STP port states are...
 - Disabled
 - e.g. shutdown
 - Listening
 - Exchanging BPDUs with adjacent bridges
 - Learning
 - Building the CAM table
 - Forwarding
 - Normal loop-free traffic forwarding
 - Blocking
 - Receiving BPDUs but not forwarding
- Normal progression between states is either...
 - Disabled → Listening → Learning → Forwarding
 - Blocking → Listening → Learning → Forwarding

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



STP Timers

- Timers that affect the transition between port states are...
 - Hello timer
 - How often configuration BPDUs are sent
 - Defaults to 2 seconds
 - MaxAge timer
 - How long to wait in blocking state without hearing a BPDU
 - Defaults to 20 seconds
 - Forward Delay
 - How long to wait in each the listening and learning phases
 - Defaults to 15 seconds

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



STP Reconvergence

- The second BPDU type, TCN, is used to quickly age out the CAM table in case of a port state change
 - e.g. Forwarding→Down, Blocking→Designated
- TCN is sent up to Root Bridge
 - TCN sent out Root Port
 - Upstream switch sends TCACK in Configuration BPDU back
 - Upstream switch sends TCN out Root Port
 - Next upstream switch sends TCACK in Configuration BPDU back
 - Next upstream switch sends TCN out Root Port
 - Process continues until Root Bridge receives TCN
- When Root Bridge receives TCN, it replies with TCN out all ports
- Result is that CAM aging time is reduced to Forward Delay
 - Default of 5 minutes reduced to 15 seconds

Copyright © 2009 Internetwork Expert, Inc
www.INE.com



STP Q&A

Copyright © 2009 Internetwork Expert, Inc
www.INE.com

