Spanning Tree

Saturday, March 29, 2008
7:45 PM

Spanning Tree - 802.1d

Four-Step STP Decision Sequence
- Lowest Root Bridge ID
- Lowest Path Cost to Root Bridge
- Lowest Sender Bridge ID
- Lowest Port ID

Three Steps of Initial STP Convergence
- Elect one Root Bridge
- Elect Root Ports
- Elect Designated Ports

Five STP States
- Blocking
- Listening
- Learning
- Forwarding
- Disabled

STP State and Port Roles
- Blocking
- Forwarding
- Designated Port
- Root Port
- Non-Designated Port

STP Timers
- Hello Time 2 seconds
  Controls the Time Interval between the sending of config bpdus
- Forward Delay 15 seconds
  Time taken in the listening/learning state
- Max Age 20 seconds
  Time bpdu is stored before discarding

Portfast
Forces port into the "Forwarding" state immediately

Uplinkfast
Allows an access switch to transition a blocking port towards the root bridge to the forwarding state in a couple of seconds.

BackboneFast
- Whereas UplinkFast is designed to quickly respond to failures on links directly connected to leaf-node switches, it does not help in the case of indirect failures in the core of the backbone.
- BackboneFast is a Max Age optimization. It allows the default convergence time for

indirect failures to be reduced from 50 seconds to 30 seconds.

- Enabled on all switches in the network, BackboneFast allows all the switches to propigate information about link failures throughout the network.

Per-VLAN Spanning - PVST+ - 802.1s

PVST+ created an STP instance for each VLAN.

Rapid Spanning Tree - 802.1w

Key Difference

- Waiting for only three missed Hellos on the RP before reacting (versus ten missed Hellos via Maxage timer with 802.1d)
- New processes that allow transition from the disabled state (replaces the blocking state in 802.1d) to learning state, bypassing the concept of the 802.1d listening state.
- Standardization of PortFast, UplinkFast, BackboneFast
  - Edge Port = PortFast
  - Alternate Port = UplinkFast
- Allows a backup DP when a switch has multiple ports connected to the same shared LAN segment

RSTP State

- Discarding
- Learning
- Forwarding

802.1d	802.1w
Disabled	Discarding
Blocking	Discarding
Listening	Discarding
Learning	Learning
Forwarding	Forwarding

RSTP Port Roles

- Root Port
- Designated Port
- Alternate Port
- Backup Port
- Disabled Port
- Edge Port

Rapid Per VLAN Spanning Tree Plus - RPVST+

Combination of PVST+ and 802.1w.

Multiple Spanning Tree - 802.1s

Benefits

- Like PVST+ it allows the tuning of STP so that while some ports block for one VLAN they can forward for another
- Always uses 802.1w convergence
- Does not require STP on each VLAN. Best designs use one STP instance per redundant path

Basically MST creates an MST region.

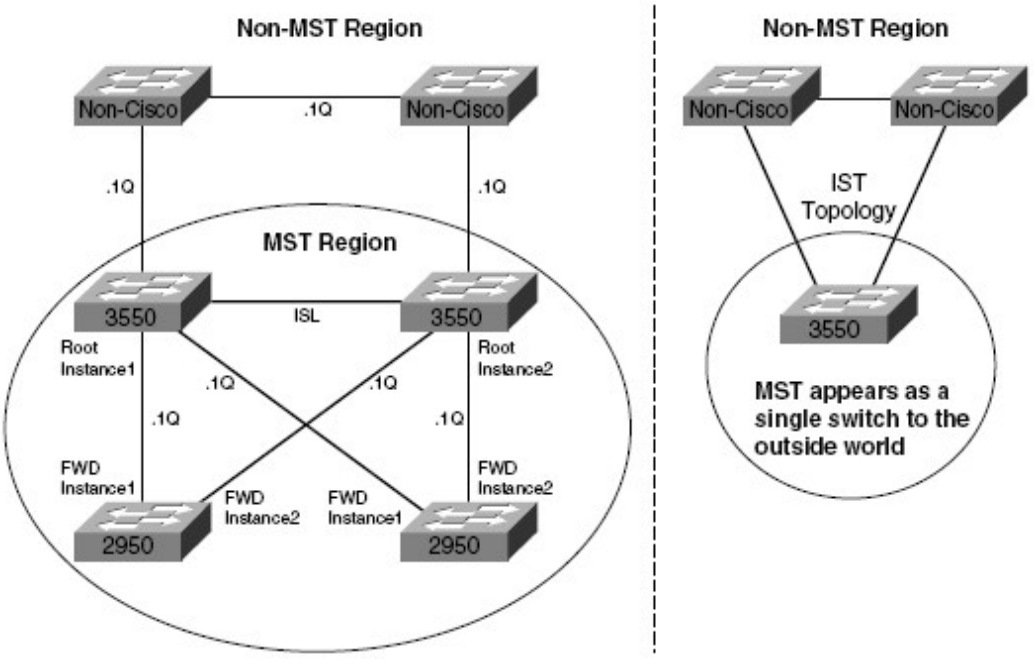

Spanning Tree Technologies

BPDU Guard

Enabled per port; error disables port on receiving of BPDU

Root Guard

Enables per port; ignores any received superior BPDUs to prevent a switch connected to this port from becoming root. Upon receipt of superior BPDUs this switch puts the port in a loop-inconsistent state, ceasing forwarding and receiving frames until the superior BPDUs cease.

With BPDU Guard, the port does not recover from the err-disabled state unless additional configuration is added. With Root Guard, the port recovers when the undesired superior BPDUs are no longer received.

UDLD and Loop Guard protect a switch trunk port from causing loops. Both features prevent switch ports from errantly moving from a blocking to a fowarding state when a unidirectional link exists in the network

Unidirectional links are simply links for which one of the two trasmission paths on the link has failed, but not both. This can happen as a result of miscabling, butting

one fiber cavle, unplugging on fiber, GBIC problems, etc.

Unidirectional Link Detection - UDLD

Uses Layer 2 messaging to decide when a switch can no longer receive frames from a neighbor. The switch whose transmit interface did not fail is placed into an err-disabled state.

UDLD Aggressive Mode

Attempts to reconnect with the other switch (eight times) after realizing no messages have been received. If the other switch does not reply to the repeated additional messages, both sides become err-disabled

Loop Guard

When normal BPDUs are no longer received, the port does not go through normal STP convergence, but rather falls into an STP loop-inconsistent state.

Storm Control

The Storm Control feature protects a LAN from being affected by unicast, broadcast, or multicast storms that might develop. The switch implements storm control by counting the number of packets of a specified type received within the one-second time interval and compares the measurement with a predefined suppression-level threshold. Storm Control can typically enable the administrator to control traffic by a percentage of total bandwidth or the traffic rate at which packets are received. It is important to note that when the rate of multicast traffic exceeds a set threshold, all incoming traffic (broadcast, multicast, and unicast) is dropped until the level drops below the specified threshold level. Only spanning-tree packets are forwarded in this situation. When broadcast and unicast thresholds are exceeded, traffic is blocked for only the type of traffic that exceeded the threshold.

EIGRP

Monday, March 31, 2008
9:08 PM

For neighbors to join
- Must pass authentication
- Must use the same AS number
- Must believe that the source IP address of a received hello is in that routers primary connected subnet on that interface
- K values must match

Note: Hello and Hold timers do NOT need to match.

EIGRP Updates
- Initially, full updates are sent, including all routes except those omitted due to split horizon
- Once all routes have been exchanged, the updates cease
- Future partial updates occur when one or more routes change
- If neighbors fail and recover, or new neighbors join, full updates are sent.

EIGRP Tables
- Neighbor
  - State information re neighbors (sh ip eigrp neighbors)
- Topology
  - EIGRP update messages
- Routing
  - Based on Topology table, chooses the best routes for routing table

EIGRP Packet Types
- Hello
  - Indentifies neighbors and services as a keepalive
- Update
  - Reliably sends route information
- Query
  - Reliably requests specific route information
- Reply
  - Reliably responds to a query
- ACK

EIGRP Metric Calculation (shortened - ON TEST)

Metric = 256(10,000,000/bandwidth) + 256(delay)

Lowest metric wins. Bandwidth is the lowest bandwidth along the path. Delay is the sum of all the delay on the path.

K1 - Bandwidth
K2 - Load
K3 - Delay
K4 - Reliability
K5 - MTU

EIGRP Convergence Function	Description
Reported distance (RD)	The metric (distance) of a route as reported by a neighboring router
Feasible distance (FD)	The metric value for the lowest-metric path to reach a particular subnet
Feasibility condition	When multiple routes to reach one subnet exist, the case in which one route's RD is lower than the FD
Successor route	The route to each destination prefix for which the metric is the lowest metric
Feasible successor (FS)	A route that is not a successor route but meets the feasibility condition; can be used when the successor route fails, without causing loops
Input event	Any occurrence that could change a router's EIGRP topology table
Local computation	An EIGRP router's reaction to an input event, leading to the use of a feasible successor or going active on a route

Local Computation

- If FS routes exist, install the lowest-metric FS route into the routing table and send updates to neighbors to notify them of the new route
- If no FS route exists, actively query neighbors for a new route

Going Active

- If the router does not have an entry in its topology table for that subnet it sends an EIGRP reply packet stating that it has no route
- If the routers successor for that subnet is unchanged, or an FS is found, the neighbor sends back an EIGRP reply message with the details of the route
- If 1 and 2 don't exist, the router goes into active and withholds its RIGRP response to the original query until all of the neighbors respond

Stuck in Active

- When the active timer expires before a router receives all the reply message

EIGRP Stub

- Stub routers are not sent query messages.

Option	This Router Is Allowed To...
connected	Advertise connected routes, but only for interfaces matched with a **network** command.
summary	Advertise auto-summarized or statically configured summary routes.
static	Advertise static routes, assuming the **redistribute static** command is configured.
redistributed	Advertise redistributed routes, assuming redistribution is configured.
receive-only	Not advertise any routes. This option cannot be used with any other option.

EIGRP Variance

- Allows EIGRP to load-balance over unequal paths.
- Calculated with a multiplier * successor route.

OSPF

Tuesday, April 01, 2008
9:42 PM

Router IDs

- Router-id command
- Highest up/up loopback
- Highest up/up non-loopback

RID Notes

- RID doesn't have to match any OSPF network commands
- Does not have to advertise a route to reach the RIDs subnet
- Does not have to be reachable
- Steps 2 and 3 look at the then-current interface state for RID choice
- Router consider the RID when the process is started or RID changed via config
- If RID changes the rest of the area routers have to perform SPF calculation
- When RID configured with router-id command the RID will never change

OSPF Messages

Message	Description
Hello	Used to discover neighbors, bring a neighbor relationship to a 2-way state, and monitor a neighbor's responsiveness in case it fails
Database Description (DD or DBD)	Used to exchange brief versions of each LSA, typically on initial topology exchange, so that a router knows a list of that neighbor's LSAs
Link-State Request (LSR)	A packet that identifies one or more LSAs about which the sending router would like the neighbor to supply full details about the LSAs
Link-State Update (LSU)	A packet that contains fully detailed LSAs, typically sent in response to an LSR message
Link-State Acknowledgement (LSAck)	Sent to confirm receipt of an LSU message

OSPF Adjacency rules

- Must pass authentication
- Must be in the same primary subnet, including mask
- Same area
- Same type (stub, nssa, etc)
- Non duplicate RIDS
- Hello and Dead timers must match

Note: These do not have to match

- OSPF Process ID

Adjacency Progress

- Down
- Init

- Two-way
- Exstart
- Exchange
- Loading
- Full

DR election process

- Any router with OSPF priority 1-255 can try to become DR by putting its RID into the DR field of its sent hellos
- Routers examine the Hellos
- If a router sees a "better" hello, it starts advertising the better router as DR
- Better = higher priority
- If priorities tie then higher RID wins
- The router not claiming to be the DR with the higher priority (or higher RID, in case of a tie) becomes the BDR
- If a new router arrives after the election, or an existing router improves the priority, it cannot preempt the existing DR and take over as DR, or BDR
- When a DR is selected, and the DR fails, the BDR becomes DR, a new election is held for a new BDR

OSPF Network Types

Interface Type	Uses DR/ BDR?	Default Hello Interval	Requires a neighbor Command?	More than Two Hosts Allowed in the Subnet?
Broadcast	Yes	10	No	Yes
Point-to-point[1]	No	10	No	No
Nonbroadcast[2] (NBMA)	Yes	30	Yes	Yes
Point-to-multipoint	No	30	No	Yes
Point-to-multipoint nonbroadcast	No	30	Yes	Yes
Loopback	No	—	—	No

[1] Default on Frame Relay point-to-point subinterfaces.

[2] Default on Frame Relay physical and multipoint subinterfaces.

OSPF Area Types

- Stub
  - Type 4 and 5 LSAs blocked
- Totally Stubby Area
  - Block 3,4,5 LSAs except type 3 default LSA
- Not so stubby area
  - Stub area that creates type 7 LSAs for external routes
- Totally not so stubby area
  - Totally stubby area that creates type 7 LSAs for external routes

LSA Types

Type	Description
1	Router Link Entry
2	Network Entry
3/4	Summary Entry
5	External Entry
6	Multicast OSPF
7	NSSA

Here is a summary of the LSA types permitted in each area.

Area	LSA 1	LSA 2	LSA 3	LSA 4	LSA 5	LSA 7
Backbone	Yes	Yes	Yes	Yes	Yes	No
Nonbackbone	Yes	Yes	Yes	Yes	Yes	No
Stub	Yes	Yes	Yes	Yes	No	No
Totally stubby	Yes	Yes	No	No	No	No
NSSA	Yes	Yes	Yes	Yes	No	Yes