

Spanning-Tree Protocol (STP) - I

A Layer 2 switch, which functions as a transparent bridge, offers no additional links for redundancy purposes. To add redundancy, a second switch must be added. Now two switches offer the transparent bridging function in parallel. LAN designs with redundant links introduce the possibility that frames might loop around the network forever. These looping frames would cause network performance problems. For example, when the switches receive an unknown unicast, both will flood the frame out all their available ports, including the ports that link to the other switch, resulting in what is known as a bridging loop, as the frame is forwarded around and around between two switches. This occurs because parallel switches are unaware of each other. The Spanning Tree Protocol (STP), which allows the redundant LAN links to be used while preventing frames from looping around the LAN indefinitely through those redundant links, was developed to overcome the possibility of bridging loops. It enables switches to become aware of each other so that they can negotiate a loop-free path through the network. Loops are discovered before they are opened for use, and redundant links are shut down to prevent the loops from forming. STP is communicated between all connected switches on a network. Each switch executes the Spanning-Tree Algorithm (STA) based on information received from other neighboring switches. The algorithm chooses a reference point in the network and calculates all the redundant paths to that reference point. When redundant paths are found, STA picks one path to forward frames with and disables or blocks forwarding on the other redundant paths. STP computes a tree structure that spans all switches in a subnet or network. Redundant paths are placed in a blocking or standby state to prevent frame forwarding. The switched network is then in a loop-free condition. However, if a forwarding port fails or becomes disconnected, the STA will run again to recompute the Spanning-Tree topology so that blocked links can be reactivated. By default, STP is enabled on all ports of a switch. STP should remain enabled in a network to prevent bridging loops from forming. However, if STP has been disabled on a CLI-based switch, it can be reenabled with the following command: **Switch (enable) set spantree enable [all | module_number/port_number]** If STP has been disabled on an IOS-based switch, it can be re-enabled with the following command: **Switch (config)# spantree vlan_list** You can use the `show spantree [vlan]` command to view the status of STP on either a CLI- or IOS-based switch. The STA places each bridge/switch port in either a forwarding state or a blocking state. All the ports in forwarding state are considered to be in the current spanning tree. The collective set of forwarding ports creates a single path over which frames are sent between Ethernet segments. Switches can forward frames out ports and receive frames in ports that are in forwarding state; switches do not forward frames out ports and receive frames in ports that are in blocking state. STP uses three criteria to choose whether to put an interface in forwarding state or a blocking state: . STP elects a root bridge and puts all interfaces on the root bridge in forwarding state. . Each nonroot bridge considers one of its ports to have the lowest administrative cost between itself and the root bridge. STP places this lowest-root-cost interface, called that bridge's root port, in forwarding state. . Many bridges can attach to the same Ethernet segment. The bridge with the lowest administrative cost from itself to the root bridge, as compared with the other bridges attached to the same segment, is placed in forwarding state. The lowest-cost bridge on each segment is called the designated bridge, and that bridge's interface, attached to that segment, is called the designated port. All other interfaces are placed in blocking state. Root Bridge Election For all switches in a network to agree on a loop-free topology, a common frame of reference must exist. This reference point is called the Root Bridge. The Root Bridge is chosen by an election process among all connected switches. Each switch has a unique Bridge ID that it uses to identify itself to other switches. The Bridge ID is an 8-byte value. 2 bytes of the Bridge ID is used for a Bridge Priority field, which is the priority or weight of a switch in relation to all other switches. The other 6 bytes of the Bridge ID is used for the MAC Address field, which can come from the Supervisor module, the backplane, or a pool of 1024 addresses that are assigned to every Supervisor or backplane depending on the switch model. This address is hardcoded, unique, and cannot be changed. The election process begins with every switch sending out BPDUs with a Root Bridge ID equal to its own Bridge ID as well as a Sender Bridge ID. The latter is used to identify the source of the BPDU message. Received BPDU messages are analyzed for a lower Root Bridge ID value. If the BPDU message has a Root Bridge ID of the lower value than the switch's own Root Bridge ID, it replaces its own Root Bridge ID with the Root Bridge ID announced in the BPDU. If two Bridge Priority values are equal, then the lower MAC address takes preference. The switch is then nominated the new Root Bridge ID in its own BPDU messages although it will still identify itself as the Sender Bridge ID. Once the process has converged, all switches will agree on the Root Bridge until a new switch is added. The Root Bridge election is based on the idea that one switch is chosen as a common reference point, and all other switches choose ports that are closest to the Root. The Root Bridge election is also based on the idea that the Root Bridge can become a central hub that interconnects other legs of the network. Therefore, the Root Bridge can be faced with heavy switching loads in its central location. If heavy loads of traffic are expected to pass through the Root Bridge, the slowest switch is not the ideal candidate. Furthermore, only one Root Bridge is elected. This is thus not fault tolerant. To overcome these problems, you should set a Root Bridge in a

determined fashion, and set a secondary Root Bridge in case of primary Root Bridge failure. The Root Bridge and the secondary Root Bridge should be placed near the center of the network. To configure a CLI-based Catalyst switch to become the Root Bridge, use the following command to modify the Bridge Priority value so that a switch can be given a lower Bridge ID value to win a Root Bridge election: **Switch (enable) set spantree priority bridge_priority [vlan** Alternatively, you can use the following command: **Switch (enable) set spantree root [secondary [vlan_list [dia diameter [hello hello_time** This command is a macro that executes several other commands. The result is a more direct and automatic way to force one switch to become the Root Bridge. Actual Bridge Priorities are not given in the command. Rather, the switch will modify STP values according to the current values in use within the active network. To configure an IOS-based Catalyst switch to become the Root Bridge, use the following command to modify the Bridge Priority value so that a switch can be given a lower Bridge ID value to win a Root Bridge election: **Switch (config)# spanning-tree [vlan vlan_list priority bridge_priority** Root Ports Election Once a reference point has been nominated and elected for the entire switched network, each non-root switch must find its relation to the Root Bridge. This action can be performed by selecting only one Root Port on each non-root switch. STP uses the Root Path Cost to select a Root Port. The Root Path Cost is the cumulative cost of all the links leading to the Root Bridge. A particular switch link has a cost associated with it called the Port or Path Cost. This cost is inversely proportional to the port's bandwidth. As the Path Cost travels along, other switches can modify its value to make it cumulative. The Path Cost is known only to the local switch where the port or "path" to a neighboring switch resides as it is not contained in the BPDU. Only the Root Path Cost is contained in the BPDU. Path Costs are defined as a one-byte value. The Root Bridge sends out a BPDU with a Root Path Cost value of zero because its ports sit directly on the Root Bridge. When the next closest neighbor receives the BPDU, it adds the Path Cost of its own port where the BPDU arrived. The neighbor then sends out BPDUs with this new cumulative value as the Root Path Cost. This value is incremented by subsequent switch port Path Costs as the BPDU is received by each switch on down the line. After incrementing the Root Path Cost, a switch also records the value in its memory. When a BPDU is received on another port and the new Root Path Cost is lower than the previously recorded value, this lower value becomes the new Root Path Cost. In addition, the lower cost tells the switch that the Root Bridge must be closer to this port than it was on other ports. The switch has now determined which of its ports is the closest to the root-the Root Port. If desired, the cost of a port can be modified from the default value. However, changing one port's cost may influence STP to choose that port as a Root Port. Therefore careful calculation is required to ensure that the desired path will be elected. On a CLI-based switch, the port cost can be modified by using one of the following commands: **Switch (enable) set spantree portcost module_number/port_number cost** or **Switch (enable) set spantree portvlancost module_number/port_number [cost cost [vlan_list** On an IOS-based switch, the port cost for individual VLANs can be modified by using the following command: **Switch (config-if)# spanning-tree [vlan vlan_list cost cost** Designated Ports Election Once the Root Path Cost values have been computed, the Root Ports have been identified; however, all other links are still connected and could be active, leaving bridging loops. To remove the bridging loops, STP makes a final computation to identify one Designated Port on each network segment which would forward traffic to and from that segment. Switches choose a Designated Port based on the lowest cumulative Root Path Cost to the Root Bridge. All ports are still active and bridging loops are still possible. STP has a set of progressive states that each port must go through, regardless of the type or identification. These states will actively prevent loops from forming.