Wikipedia

Spanning Tree Protocol

The **Spanning Tree Protocol** (**STP**) is a network protocol that builds a loop-free logical topology for Ethernet networks. The basic function of STP is to prevent bridge loops and the broadcast radiation that results from them. Spanning tree also allows a network design to include backup links providing fault tolerance if an active link fails.

As the name suggests, STP creates a spanning tree that characterizes the relationship of nodes within a network of connected layer-2 bridges, and disables those links that are not part of the spanning tree, leaving a single active path between any two network nodes. STP is based on an algorithm that was invented by Radia Perlman while she was working for Digital Equipment Corporation.[1][2]

In 2001, the IEEE introduced **Rapid Spanning Tree Protocol** (**RSTP**) as 802.1w. RSTP provides significantly faster recovery in response to network changes or failures, introducing new convergence behaviors and bridge port roles to do this. RSTP was designed to be backwards-compatible with standard STP.

STP was originally standardized as IEEE 802.1D but the functionality of spanning tree (802.1D), rapid spanning tree (802.1w), and multiple spanning tree (802.1s) has since been incorporated into IEEE 802.1Q-2014.[3]

Contents

Protocol operation

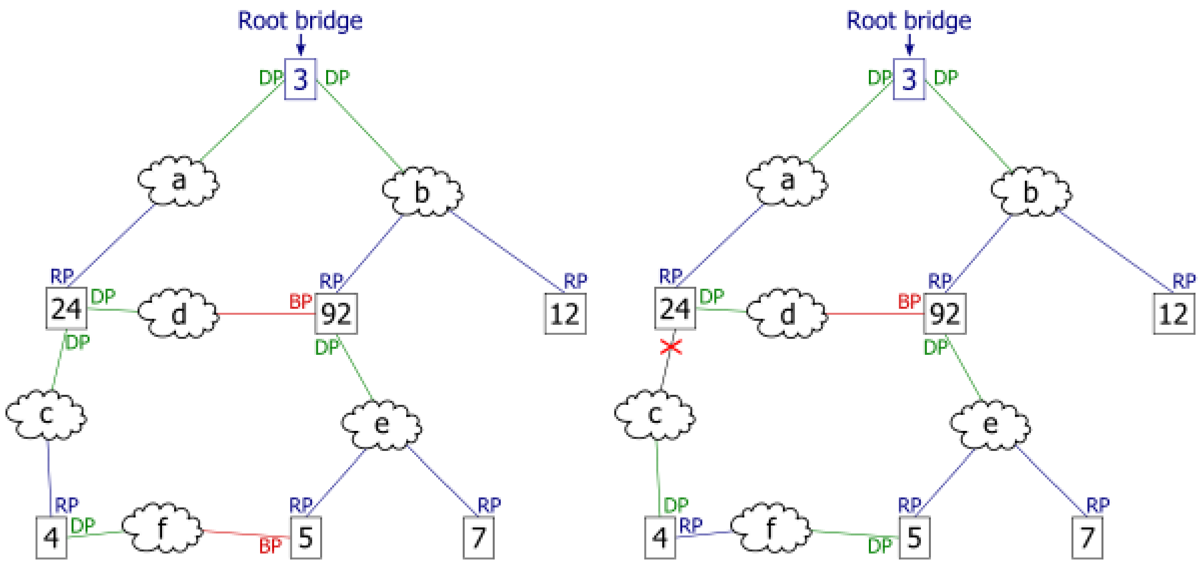

Switches with Spanning Tree Protocol implementation in a local area network (LAN). One switch is the STP *root bridge*. All switch ports that connect a link between two switches are either a *root port* (RP), a *designated port* (DP), or a *blocked port* (BP).

After link failure the spanning tree algorithm computes and spans new least-cost tree.

Switches with Spanning Tree Protocol implementation in a local area network (LAN)

The need for the Spanning Tree Protocol (STP) arose because switches in local area networks (LANs) are often interconnected using redundant links to improve resilience should one connection fail.[4]: 386 However, this connection configuration creates a switching loop resulting in broadcast radiations and MAC table instability.[4]: 388 If redundant links are used to connect switches, then switching loops need to be avoided.[4]: 385

To avoid the problems associated with redundant links in a switched LAN, STP is implemented on switches to monitor the network topology. Every link between switches, and in particular redundant links, are catalogued. The spanning-tree algorithm then blocks forwarding on redundant links by setting up one preferred link between switches in the LAN. This preferred link is used for all Ethernet frames unless it fails, in which case a non-preferred redundant link is enabled. When implemented in a network, STP designates one layer-2 switch as *root bridge*. All switches then select their best connection towards the root bridge for forwarding and block other redundant links. All switches constantly communicate with their neighbors in the LAN using Bridge Protocol Data Units (BPDUs).[4]: 388

Provided there is more than one link between two switches, the STP root bridge calculates the cost of each path based on bandwidth. STP will select the path with the lowest cost, that is the highest bandwidth, as the preferred link. STP will enable this preferred link as the only path to be used for Ethernet frames between the two switches, and disable all other possible links by designating the switch ports that connect the preferred path as *root port*.[4]: 393

After STP enabled switches in a LAN have elected the root bridge, all non-root bridges assign one of their ports as root port. This is either the port that connects the switch to the root bridge, or if there are several paths, the port with the preferred path as calculated by the root bridge. Because not all switches are directly connected to the root bridge they communicate amongst each other using STP Bridge Protocol Data Units (BPDUs). Each switch adds the cost of its own path to the cost received from the neighboring switches to determine the total cost of a given path to the root bridge. Once the cost of all possible paths to the root bridge have been added up, each switch assigns a port as root port which connects to the path with the lowest cost, or highest bandwidth, that will eventually lead to the root bridge.[4]: 394

Path cost

The STP path cost default was originally calculated by the formula $\frac{1\text{ Gbit/s}}{\text{bandwidth}}$. When faster speeds became available, the default values were adjusted as otherwise speeds above 1 Gbit/s would have been indistinguishable by STP. Its successor RSTP uses a similar formula with a larger numerator: $\frac{20\text{ Tbit/s}}{\text{bandwidth}}$. These formulas lead to the sample values in the table.[5]: 154

Path cost for different port speed and STP variation

Data rate (link bandwidth)	Original STP cost (802.1D-1998)	RSTP/MSTP cost (recommended value)[3]: 503
4 Mbit/s	250	5,000,000
10 Mbit/s	100	2,000,000
16 Mbit/s	62	1,250,000
100 Mbit/s	19	200,000
1 Gbit/s	4	20,000
2 Gbit/s	3	10,000
10 Gbit/s	2	2,000
100 Gbit/s	N/A	200
1 Tbit/s	N/A	20

Port states

All switch ports in the LAN where STP is enabled are categorized.[4]: 388

Blocking

A port that would cause a switching loop if it were active. To prevent the use of looped paths, no user data is sent or received over a blocking port. BPDU data is still received in blocking state. A blocked port may go into forwarding mode if the other links in use fail and the spanning tree algorithm determines the port may transition to the forwarding state.

Listening

The switch processes BPDUs and awaits possible new information that would cause it to return to the blocking state. It does not populate the MAC table and it does not forward frames.

Learning

While the port does not yet forward frames, it does learn source addresses from frames received and adds them to the MAC table.

Forwarding

A port in normal operation receiving and forwarding frames. The port monitors incoming BPDUs that would indicate it should return to the blocking state to prevent a loop.

Disabled

A network administrator has manually disabled the switch port.

When a device is first attached to a switch port, it will not immediately start to forward data. It will instead go through a number of states while it processes BPDUs and determines the topology of the network. The port attached to a host such as a computer, printer or server always goes into the forwarding state, albeit after a delay of about 30 seconds while it goes through the listening and learning states. The time spent in the listening and learning states is determined by a value known as the forward delay (default 15 seconds and set by the root bridge). If another switch is connected, the port may remain in blocking mode if it is determined that it would cause a loop in the network. Topology Change Notification (TCN) BPDUs are used to inform other switches of port changes. TCNs are injected into the network by a non-root switch and propagated to the root. Upon receipt of the TCN, the root switch will set the Topology Change flag in its normal BPDUs. This flag is propagated to all other switches and instructs them to rapidly age out their forwarding table entries.

Configuration

Before configuring STP, the network topology should be carefully planned.[6] Basic configuration requires that STP be enabled on all switches in the LAN and the same version of STP chosen on each. The administrator may determine which switch will be the root bridge and configure the switches appropriately. If the root bridge goes down, the protocol will automatically assign a new root bridge based on bridge ID. If all switches have the same bridge ID, such as the default ID, and the root bridge goes down, a tie situation arises and the protocol will assign one switch as root bridge based on the switch MAC addresses. Once the switches have been assigned a bridge ID and the protocol has chosen the root bridge switch, the best path to the root bridge is calculated based on port cost, path cost and port priority.[7] Ultimately STP calculates the path cost on the basis of the bandwidth of a link, however links between switches may have the same bandwidth. Administrators can influence the protocol's choice of the preferred path by configuring the port cost, the lower the port cost the more likely it is that the protocol will choose the connected link as root port for the preferred path.[8] The selection of how other switches in the topology choose their root port, or the least cost path to the root bridge, can be influenced by the port priority. The highest priority will mean the path will ultimately be less preferred. If all ports of a switch have the same priority, the port with the lowest number is chosen to forward frames.[9]

Root bridge and the bridge ID

The *root bridge* of the spanning tree is the bridge with the smallest (lowest) bridge ID. Each bridge has a configurable priority number and a MAC address; the bridge ID is the concatenation of the bridge priority and the MAC address. For example, the ID of a bridge with priority 32768 and MAC *0200.0000.1111* is *32768.0200.0000.1111*. The bridge priority default is 32768 and can only be configured in multiples of 4096.[a] When comparing two bridge IDs, the priority portions are compared first and the MAC addresses are compared only if the priorities are equal. The switch with the lowest priority of all the switches will be the root; if there is a tie, then the switch with the lowest priority and lowest MAC address will be the root. For example, if switches *A* (MAC = *0200.0000.1111*) and *B* (MAC = *0200.0000.2222*) both have a priority of 32768 then switch *A* will be selected as the root bridge.[b] If the network administrators would like switch *B* to become the root bridge, they must set its priority to be less than 32768.[c]

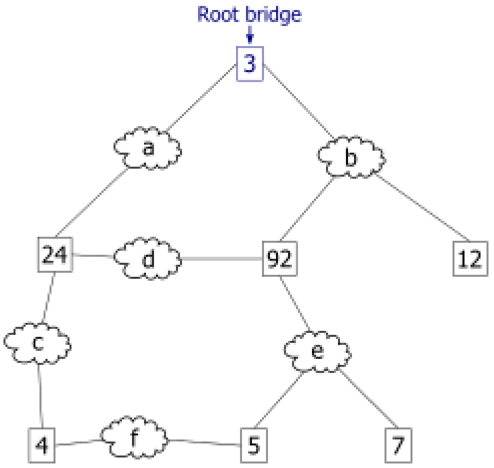

An example network. The numbered boxes represent bridges, that is switches in a LAN. The number is the bridge ID. The lettered clouds represent network segments. The smallest bridge ID is 3. Therefore, bridge 3 is the root bridge.

Path to the root bridge

The sequence of events to determine the best received BPDU (which is the best path to the root) is:

1. Lowest root bridge ID (BID) - Determines the root bridge.
2. Lowest cost to the root bridge - Favors the upstream switch with the least cost to root
3. Lowest sender bridge ID - Serves as a tiebreaker if multiple upstream switches have equal cost to root
4. Lowest sender port ID - Serves as a tiebreaker if a switch has multiple (non-EtherChannel) links to a single upstream switch, where:
   - Bridge ID = priority (4 bits) + locally assigned system ID extension (12 bits) + ID [MAC address] (48 bits); the default bridge priority is 32768, and
   - Port ID = priority (4 bits) + ID (Interface number) (12 bits); the default port priority is 128.

Tiebreakers

Root ports

When multiple paths from a bridge are least-cost paths, the chosen path uses the neighbor bridge with the lower bridge ID. The root port is thus the one connecting to the bridge with the lowest bridge ID. For example, in the figures, if switch 4 were connected to network segment d instead of segment f, there would be two paths of length 2 to the root, one path going through bridge 24 and the other through bridge 92. Because there are two least-cost paths, the lower bridge ID (24) would be used as the tiebreaker in choosing which path to use.

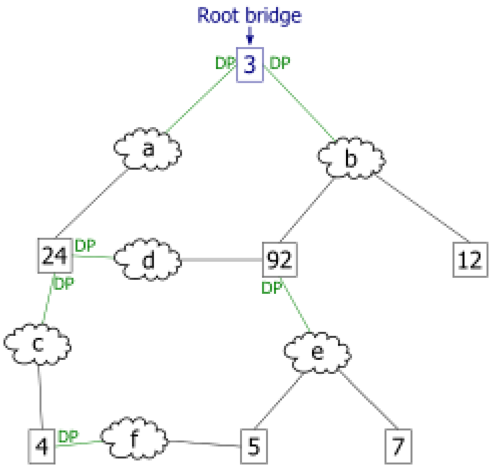

Path tie: The least-cost path to the root from network segment e goes through bridge 92. Therefore, the designated port for network segment e is the port that connects bridge 92 to network segment e.

Paths

When more than one bridge on a segment leads to a least-cost path to the root, the bridge with the lower bridge ID is used to forward messages to the root. The port attaching that bridge to the network segment is the *designated port* for the segment. In the figures, there are two least-cost paths from network segment d to the root, one going through bridge 24 and the other through bridge 92. The lower bridge ID is 24, so the tiebreaker dictates that the designated port is the port through which network segment d is connected to bridge 24. If bridge IDs were equal, then the bridge with the lowest MAC address would have the designated port. In either case, the loser sets the port as being blocked.

Designated ports

When the root bridge has more than one port on a single LAN segment, the bridge ID is effectively tied, as are all root path costs (all equal zero). The port on that LAN segment with the lowest port ID becomes the designated port. It is put into forwarding mode while all other ports on the root bridge on that same LAN segment become non-designated ports and are put into blocking mode.[11] Not all bridge manufacturers follow this rule, instead making all root bridge ports designated ports, and putting them all in forwarding mode.

Final tiebreaker

In some cases, there may still be a tie, as when the root bridge has multiple active ports on the same LAN segment (see above) with equally low root path costs and bridge IDs, or, in other cases, multiple bridges are connected by multiple cables and multiple ports. In each case, a single bridge may have multiple candidates for its root port. In these cases, candidates for the root port have already received BPDUs offering equally-low (i.e. the "best") root path costs and equally-low (i.e. the "best") bridge IDs, and the final tiebreaker goes to the port that received the lowest (i.e. the "best") port priority ID, or port ID.[12]

Bridge Protocol Data Units

The above rules describe one way of determining what spanning tree will be computed by the algorithm, but the rules as written require knowledge of the entire network. The bridges have to determine the root bridge and compute the port roles (root, designated, or blocked) with only the information that they have. To ensure that each bridge has enough information, the bridges use special data frames called Bridge Protocol Data Units (BPDU) to exchange information about bridge IDs and root path costs.

A bridge sends a BPDU frame using the unique MAC address of the port itself as a source address, and a destination address of the STP multicast address *01:80:C2:00:00:00*.

There are two types of BPDUs in the original STP specification[5]:63 (the Rapid Spanning Tree (RSTP) extension uses a specific RSTP BPDU):

- Configuration BPDU (CBPDU), used for spanning tree computation
- Topology Change Notification (TCN) BPDU, used to announce changes in the network topology

BPDUs are exchanged regularly (every 2 seconds by default) and enable switches to keep track of network changes and to start and stop forwarding at ports as required. To prevent the delay when connecting hosts to a switch and during some topology changes, Rapid STP was developed, which allows a switch port to rapidly transition into the forwarding state during these situations.

Bridge Protocol Data Unit fields

IEEE 802.1D and IEEE 802.1aq BPDUs have the following format:

```
 1. Protocol ID:       2 bytes (0x0000 IEEE 802.1D)
 2. Version ID:        1 byte (0x00 Config & TCN / 0x02 RST / 0x03 MST / 0x04 SPT  BPDU)
 3. BPDU Type:         1 byte (0x00 STP Config BPDU, 0x80 TCN BPDU, 0x02 RST/MST Config BPDU)
 4. Flags:             1 byte
   bits  : usage
       1 : 0 or 1 for Topology Change
       2 : 0 (unused) or 1 for Proposal in RST/MST/SPT BPDU
     3-4 : 00 (unused) or
           01 for Port Role Alternate/Backup in RST/MST/SPT BPDU
           10 for Port Role Root in RST/MST/SPT BPDU
           11 for Port Role Designated in RST/MST/SPT BPDU
       5 : 0 (unused) or 1 for Learning in RST/MST/SPT BPDU
       6 : 0 (unused) or 1 for Forwarding in RST/MST/SPT BPDU
       7 : 0 (unused) or 1 for Agreement in RST/MST/SPT BPDU
       8 : 0 or 1 for Topology Change Acknowledgement
 5. Root ID:           8 bytes (CIST Root ID in MST/SPT BPDU)
   bits  : usage
     1-4 : Root Bridge Priority
    5-16 : Root Bridge System ID Extension
   17-64 : Root Bridge MAC Address
 6. Root Path Cost:    4 bytes (CIST External Path Cost in MST/SPT BPDU)
 7. Bridge ID:         8 bytes (CIST Regional Root ID in MST/SPT BPDU)
   bits  : usage
     1-4 : Bridge Priority
    5-16 : Bridge System ID Extension
   17-64 : Bridge MAC Address
  8. Port ID:          2 bytes
  9. Message Age:      2 bytes in 1/256 secs
 10. Max Age:          2 bytes in 1/256 secs
 11. Hello Time:       2 bytes in 1/256 secs
 12. Forward Delay:    2 bytes in 1/256 secs
 13. Version 1 Length: 1 byte (0x00 no ver 1 protocol info present. RST, MST, SPT BPDU only)
```

```
14. Version 3 Length: 2 bytes (MST, SPT BPDU only)

The TCN BPDU includes fields 1-3 only.
```

Spanning Tree Protocol standards

The first spanning tree protocol was invented in 1985 at the Digital Equipment Corporation by Radia Perlman.[1] In 1990, the IEEE published the first standard for the protocol as 802.1D,[13] based on the algorithm designed by Perlman. Subsequent versions were published in 1998[14] and 2004,[15] incorporating various extensions. The original Perlman-inspired Spanning Tree Protocol, called DEC STP, is not a standard and differs from the IEEE version in message format as well as timer settings. Some bridges implement both the IEEE and the DEC versions of the Spanning Tree Protocol, but their interworking can create issues for the network administrator.[16]

Different implementations of a standard are not guaranteed to interoperate, due for example to differences in default timer settings. The IEEE encourages vendors to provide a Protocol Implementation Conformance Statement, declaring which capabilities and options have been implemented,[15] to help users determine whether different implementations will interoperate correctly.

Rapid Spanning Tree Protocol

In 2001, the IEEE introduced Rapid Spanning Tree Protocol (RSTP) as **IEEE 802.1w**. RSTP was then incorporated into IEEE 802.1D-2004 making the original STP standard obsolete.[17] RSTP was designed to be backward-compatible with standard STP.

RSTP provides significantly faster spanning tree convergence after a topology change, introducing new convergence behaviors and bridge port roles to accomplish this. While STP can take 30 to 50 seconds to respond to a topology change, RSTP is typically able to respond to changes within 3 × *hello times* (default: 3 × 2 seconds) or within a few milliseconds of a physical link failure. The hello time is an important and configurable time interval that is used by RSTP for several purposes; its default value is 2 seconds.[18][19]

Rapid Spanning Tree Operation

RSTP adds new bridge port roles in order to speed convergence following a link failure:

- **Root** - A forwarding port that is the best port from non-root bridge to root bridge
- **Designated** - A forwarding port for every LAN segment
- **Alternate** - An alternate path to the root bridge. This path is different from using the root port
- **Backup** - A backup/redundant path to a segment where another bridge port already connects
- **Disabled** - Not strictly part of STP, a network administrator can manually disable a port

The number of switch port states a port can be in has been reduced to three instead of STP's original five:

- **Discarding** - No user data is sent over the port
- **Learning** - The port is not forwarding frames yet, but is populating its MAC-address-table
- **Forwarding** - The port is fully operational

RSTP operational details:

- Detection of root switch failure is done in 3 hello times, which is 6 seconds if the default hello times have not been changed.
- Ports may be configured as edge ports if they are attached to a LAN that has no other bridges attached. These edge ports transition directly to the forwarding state. RSTP still continues to monitor the port for BPDUs in case a bridge is connected. RSTP can also be configured to automatically detect edge ports. As soon as the bridge detects a BPDU coming to an edge port, the port becomes a non-edge port.
- RSTP calls the connection between two or more switches as a "link-type" connection. A port that operates in full-duplex mode is assumed to be point-to-point link, whereas a half-duplex port (through a hub) is considered a shared port by default. This automatic link type setting can be overridden by explicit configuration. RSTP improves convergence on point-to-point links by reducing the Max-Age time to 3 times Hello interval, removing the STP listening state, and exchanging a handshake between two switches to quickly transition the port to forwarding state. RSTP does not do anything differently from STP on shared links.
- Unlike in STP, RSTP will respond to BPDUs sent from the direction of the root bridge. An RSTP bridge will *propose* its spanning tree information to its designated ports. If another RSTP bridge receives this information and determines this is the superior root information, it sets all its other ports to discarding. The bridge may send an *agreement* to the first bridge confirming its superior spanning tree information. The first bridge, upon receiving this agreement, knows it can rapidly transition that port to the forwarding state bypassing the listening/learning state transition. This essentially creates a cascading effect away from the root bridge where each designated bridge proposes to its neighbors to determine if it can make a rapid transition. This is one of the major elements that allows RSTP to achieve faster convergence times than STP.
- As discussed in the port role details above, RSTP maintains backup details regarding the discarding status of ports. This avoids timeouts if the current forwarding ports were to fail or BPDUs were not received on the root port in a certain interval.
- RSTP will revert to legacy STP on an interface if a legacy version of an STP BPDU is detected on that port.

Standards for VLANs

STP and RSTP do not segregate switch ports by VLAN.[20] However, in Ethernet switched environments where multiple Virtual LANs (VLANs) exist, it is often desirable to create multiple spanning trees so that traffic on different VLANs uses different links.

Proprietary standards

Before the IEEE published a Spanning Tree Protocol standard for VLANs, a number of vendors who sold VLAN capable switches developed their own Spanning Tree Protocol versions that were VLAN capable. Cisco developed, implemented and published the **Per-VLAN Spanning Tree** (**PVST**) proprietary protocol using its own proprietary Inter-Switch Link (ISL) for VLAN encapsulation, and PVST+ which uses 802.1Q VLAN encapsulation. Both standards implement a separate spanning tree for every VLAN. Cisco switches now commonly implement PVST+ and can only implement Spanning Trees for VLANs if the other switches in the LAN implement the same VLAN STP protocol. HP provides PVST and PVST+ compatibility in some of its network switches.[21] Some devices from Force10 Networks, Alcatel-Lucent, Extreme Networks, Avaya, Brocade Communications Systems and BLADE Network Technologies support PVST+.[22][23][24] Extreme Networks does so with two limitations: Lack of support on ports where the VLAN is untagged/native, and also on the VLAN with ID 1. PVST+ can tunnel across an MSTP Region.[25]

The switch vendor Juniper Networks in turn developed and implemented its VLAN Spanning Tree Protocol (VSTP) to provide compatibility with Cisco's PVST, so that the switches from both vendors can be included in one LAN.[20] The VSTP protocol is only supported by the EX and MX Series from Juniper Networks. There are two restrictions to the compatibility of VSTP:

1. VSTP supports only 253 different spanning-tree topologies. If there are more than 253 VLANs, it is recommended to configure RSTP in addition to VSTP, and VLANs beyond 253 will be handled by RSTP.
2. MVRP does not support VSTP. If this protocol is in use, VLAN membership for trunk interfaces must be statically configured.[26]

By default, VSTP uses the RSTP protocol as its core spanning-tree protocol, but usage of STP can be forced if the network includes old bridges.[27] More information about configuring VSTP on Juniper Networks switches was published in the official documentation.[28]

Cisco also published a proprietary version of Rapid Spanning Tree Protocol. It creates a spanning tree for each VLAN, just like PVST. Cisco refers to this as **Rapid Per-VLAN Spanning Tree** (**RPVST**).

Multiple Spanning Tree Protocol

The *Multiple Spanning Tree Protocol* (MSTP), originally defined in IEEE 802.1s-2002 and later merged into IEEE 802.1Q-2005, defines an extension to RSTP to further develop the usefulness of virtual LANs (VLANs).

In the standard, a spanning tree that maps one or more VLANs is called *multiple spanning tree* (MST). If MSTP is implemented a spanning tree can be defined for individual VLANs or for groups of VLANs. Furthermore, the administrator can define alternate paths within a spanning tree. VLANs must be assigned to a so-called *multiple spanning tree instance* (MSTI). Switches are first assigned to an MST region, then VLANs are mapped against or assigned to this MST. A *Common Spanning Tree* (CST) is an MST to which several VLANs are mapped, this group of VLANs is called *MST Instance* (MSTI). CSTs are backward compatible with the STP and RSTP standard. A MST that has only one VLAN assigned to it is a *Internal Spanning Tree* (IST).[21]

Unlike some proprietary per-VLAN spanning tree implementations,[29] MSTP includes all of its spanning tree information in a single BPDU format. Not only does this reduce the number of BPDUs required on a LAN to communicate spanning tree information for each VLAN, but it also ensures backward compatibility with RSTP (and in effect, classic STP too). MSTP does this by encoding additional region information after the standard RSTP BPDU as well as a number of MSTI messages (from 0 to 64 instances, although in practice many bridges support fewer). Each of these MSTI configuration messages conveys the spanning tree information for each instance. Each instance can be assigned a number of configured VLANs and frames (packets) assigned to these VLANs operate in this spanning tree instance whenever they are inside the MST region. In order to avoid conveying their entire VLAN to spanning tree mapping in each BPDU, bridges encode an MD5 digest of their VLAN to instance table in the MSTP BPDU. This digest is then used by other MSTP bridges, along with other administratively configured values, to determine if the neighboring bridge is in the same MST region as itself.

MSTP is fully compatible with RSTP bridges, in that an MSTP BPDU can be interpreted by an RSTP bridge as an RSTP BPDU. This not only allows compatibility with RSTP bridges without configuration changes, but also causes any RSTP bridges outside of an MSTP region to see the region as a single RSTP bridge, regardless of the number of MSTP bridges inside the region itself. In order to further facilitate this view of an MST region as a single RSTP bridge, the MSTP protocol uses a variable known as remaining hops as a time to live counter instead of the message age timer used by RSTP. The message age time is only incremented once when spanning tree information enters an MST region, and therefore RSTP bridges will see a region as only one "hop" in the spanning tree. Ports at the edge of an MST region connected to either an RSTP or STP bridge or an endpoint are known as boundary ports. As in RSTP, these ports can be configured as edge ports to facilitate rapid changes to the forwarding state when connected to endpoints.

Shortest path bridging

IEEE 802.1aq also known as Shortest Path Bridging (SPB) allows redundant links between switches to be active through multiple equal cost paths, and provides much larger layer-2 topologies, faster convergence, and improves the use of the mesh topologies through increased bandwidth between all devices by allowing traffic to load share across all paths on a mesh network.[30][31] SPB consolidates multiple existing functionalities, including Spanning Tree Protocol (STP), Multiple Spanning Tree Protocol (MSTP), Rapid Spanning Tree Protocol (RSTP), Link aggregation, and Multiple MAC Registration Protocol (MMRP) into a one link state protocol.[32]

System ID Extension

The bridge ID (BID) is a field inside a BPDU packet. It is eight bytes in length. The first two bytes are the bridge priority, an unsigned integer of 0-65,535. The last six bytes are a MAC address supplied by the bridge. Prior to IEEE 802.1D-2004, the first two bytes gave a 16 bit bridge priority. Since IEEE 802.1D-2004, the first four bits are a configurable priority, and the last twelve bits carry the bridge system ID extension. In the case of MST, the bridge system ID extension carries the MSTP instance number. Some vendors set the bridge system ID extension to carry a VLAN ID allowing a different spanning tree per VLAN, such as Cisco's PVST.

Disadvantages and current practice

Spanning tree is an older protocol with a longer default hold-down time that governs convergence of the protocol state. Improper use or implementation can contribute to network disruptions. The idea of blocking links is something that customers these days do not accept as a proper high availability solution. Modern networks can make use of all connected links by use of protocols that inhibit, control or suppress the natural behavior of logical or physical topology loops.

Newer, more robust protocols include the TRILL (Transparent Interconnection of Lots of Links) protocol, also created by Dr. Perlman.[33]

Switch virtualization techniques like HPE IRF, Aruba VSF and Cisco VSS combine multiple switches into a single logical entity. A multi-chassis link aggregation group works like a normal LACP trunk, only distributed through multiple switches. Conversely partitioning technologies compartmentalize a single physical chassis into multiple logical entities.

On the edge of the network, loop-detection is configured to prevent accidental loops by users.

See also

- Bridge Protocol Data Unit
- Distributed minimum spanning tree
- EtherChannel
- Ethernet Automatic Protection Switching

- Ethernet Ring Protection Switching
- Flex links
- Flooding (computer networking)
- Media Redundancy Protocol
- Minimum spanning tree
- Unidirectional Link Detection
- Virtual Link Trunking

Notes

a. Spanning tree incorporated 802.1t, and per 802.1t, uses the 4 most-significant bits of the 802.1d two-octet priority field as priority, and the least-significant 12 bits of that field as the extended system ID.

b. The original 802.1d envisioned the possibility of the root bridge having more than one port on the same LAN segment, and in that case, the port with the lowest port ID would become the designated port for that LAN segment, and put into forwarding mode, while its other ports on that same LAN segment became non-designated ports put into blocking mode. Not all bridge manufacturers follow that rule, some making all ports designated ports and putting them all into forwarding mode.

c. Alternatively the network administrator can configure the switch as a spanning tree root primary or secondary. When configuring the root primary and root secondary the switch will automatically change the priority accordingly, 24576 and 28672 respectively with the default configuration.[10]

References

1. Perlman, Radia (1985). "An Algorithm for Distributed Computation of a Spanning Tree in an Extended LAN" (https://semanticscholar.org/paper/933d8fc9b5ddc0e5e12116c1eb309ab535e6ae75). *ACM SIGCOMM Computer Communication Review*. **15** (4): 44–53. doi:10.1145/318951.319004 (https://doi.org/10.1145%2F318951.319004). S2CID 61172150 (https://api.semanticscholar.org/CorpusID:61172150).
2. Perlman, Radia (2000). *Interconnections, Second Edition*. USA: Addison-Wesley. ISBN 0-201-63448-1.
3. Bridges and Bridged Networks (http://www.ieee802.org/1/pages/802.1Q-2014.html)
4. Silviu Angelescu (2010). *CCNA Certification All-In-One For Dummies*. John Wiley & Sons. ISBN 9780470635926.
5. "802.1D IEEE Standard for Local and Metropolitan Area Networks. Media Access Control (MAC) Bridges" (http://standards.ieee.org/getieee802/download/802.1D-2004.pdf) (PDF). IEEE. 2004. Retrieved 19 April 2012.
6. Wade Edwards, Terry Jack, Todd Lammle, Toby Skandier, Robert Padjen, Arthur Pfund & Carl Timm (2006). *CCNP Complete Study Guide: Exams 642-801, 642-811, 642-821, 642-831*. John Wiley & Sons. pp. 506 & 511. ISBN 9780782150667.
7. Wade Edwards, Terry Jack, Todd Lammle, Toby Skandier, Robert Padjen, Arthur Pfund & Carl Timm (2006). *CCNP Complete Study Guide: Exams 642-801, 642-811, 642-821, 642-831*. John Wiley & Sons. p. 506. ISBN 9780782150667.
8. Wade Edwards, Terry Jack, Todd Lammle, Toby Skandier, Robert Padjen, Arthur Pfund & Carl Timm (2006). *CCNP Complete Study Guide: Exams 642-801, 642-811, 642-821, 642-831*. John Wiley & Sons. p. 511. ISBN 9780782150667.

9. Wade Edwards, Terry Jack, Todd Lammle, Toby Skandier, Robert Padjen, Arthur Pfund & Carl Timm (2006). *CCNP Complete Study Guide: Exams 642-801, 642-811, 642-821, 642-831*. John Wiley & Sons. p. 513. ISBN 9780782150667.
10. "spanning-tree vlan" (https://www.cisco.com/c/m/en_us/techdoc/dc/reference/cli/nxos/commands/l2/spanning-tree-vlan.html). Cisco Systems. Retrieved 2020-05-04.
11. 802.1d-1998 section 8.3.1: The Designated Port for each LAN is the Bridge Port for which the value of the Root Path Cost is the lowest: if two or more Ports have the same value of Root Path Cost, then first the Bridge Identifier of their Bridges, and their Port Identifiers are used as tie breakers.
12. 802.1d-1998 section 8.3.2 b) A Bridge that receives a Configuration BPDU on what it decides is its Root Port conveying better information (i.e. highest priority Root Identifier, lowest Root Path Cost, highest priority transmitting Bridge and Port), passes that information on to all the LANs for which it believes itself to be the Designated Bridge.
13. LAN/MAN Standards Committee of the IEEE Computer Society, ed. (1990). *ANSI/IEEE Std 802.1D*. IEEE.
14. LAN/MAN Standards Committee of the IEEE Computer Society, ed. (1998). *ANSI/IEEE Std 802.1D, 1998 Edition, Part 3: Media Access Control (MAC) Bridges*. IEEE.
15. LAN/MAN Standards Committee of the IEEE Computer Society, ed. (2004). *ANSI/IEEE Std 802.1D - 2004: IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges*. IEEE.
16. "Understanding Issues Related to Inter-VLAN Bridging" (https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/11072-inter-vlan-11072.pdf) (PDF). Cisco Systems, Inc. 11072.
17. *IEEE 802.1D-2004*, IEEE, 2004-06-04, "Since the original Spanning Tree Protocol (STP) has been removed from the 2004 revision of IEEE Std 802.1D, an implementation of RSTP is required for any claim of conformance for an implementation of IEEE Std 802.1Q-2003 that refers to the current revision of IEEE Std 802.1D"
18. Waldemar Wojdak (March 2003). "Rapid Spanning Tree Protocol: A new solution from an old technology" (http://www.compactpci-systems.com/articles/id/?203). *CompactPCI Systems*. Retrieved 2008-08-04.
19. "Understanding Rapid Spanning Tree Protocol (802.1w)" (http://www.cisco.com/en/US/tech/tk389/tk621/technologies_white_paper09186a0080094cfa.shtml). Retrieved 2008-11-27.
20. Michael G. Solomon, David Kim & Jeffrey L. Carrell (2014). *Fundamentals of Communications and Networking*. Jones & Bartlett Publishers. p. 204. ISBN 9781284060157.
21. Michael G. Solomon, David Kim & Jeffrey L. Carrell (2014). *Fundamentals of Communications and Networking*. Jones & Bartlett Publishers. p. 204. ISBN 9781284060157.
22. "Technical Documentation" (https://www.force10networks.com/CSPortal20/TechTips/0050B_HowDoIConfigureSpanningTree.aspx). Force10. Retrieved 2011-01-25.
23. "ExtremeXOS Operating System, Version 12.5" (http://www.extremenetworks.com/libraries/products/DSExtXOS_1030.pdf) (PDF). Extreme Networks. 2010. Retrieved 2011-01-25.
24. "BLADE PVST+ Interoperability with Cisco" (http://www.bladenetwork.net/userfiles/file/PDFs/WP_PVST_SpanningTree_Cisco.pdf) (PDF). 2006. Retrieved 2011-01-25.
25. "Bridging Between IEEE 802.1Q VLANs" (http://www.cisco.com/en/US/docs/ios/12_1t/12_1t3/feature/guide/dtbridge.html#wp1020686). Cisco Systems. Retrieved 2011-01-25.
26. http://www.juniper.net/techpubs/en_US/junos10.0/topics/concept/bridging-mvrp-ex-series.html
27. https://www.juniper.net/techpubs/en_US/junos9.4/topics/concept/spanning-trees-ex-series-vstp-understanding.html

28. Understanding VSTP (https://www.juniper.net/techpubs/en_US/junos14.1/topics/concept/spanning-trees-ex-series-vstp-understanding.html)
29. "CiscoWorks LAN Management Solution 3.2 Deployment Guide" (https://www.cisco.com/en/US/prod/collateral/netmgtsw/ps6504/ps6528/ps2425/white_paper_c07-552114.html#wp9003215). August 2009. Retrieved 2010-01-25.
30. Peter Ashwood-Smith (24 Feb 2011). "Shortest Path Bridging IEEE 802.1aq Overview" (https://web.archive.org/web/20130515115628/http://meetings.apnic.net/__data/assets/pdf_file/0012/32007/APRICOT_SPB_Overview.pdf) (PDF). Huawei. Archived from the original (http://meetings.apnic.net/__data/assets/pdf_file/0012/32007/APRICOT_SPB_Overview.pdf) (PDF) on 15 May 2013. Retrieved 11 May 2012.
31. Jim Duffy (11 May 2012). "Largest Illinois healthcare system uproots Cisco to build $40M private cloud" (http://www.pcadvisor.co.uk/news/internet/3357242/largest-illinois-healthcare-system-uproots-cisco-build-40m-private-cloud/). PC Advisor. Retrieved 11 May 2012. "Shortest Path Bridging will replace Spanning Tree in the Ethernet fabric."
32. "IEEE Approves New IEEE 802.1aq Shortest Path Bridging Standard" (http://www.techpowerup.com/165594/IEEE-Approves-New-IEEE-802.1aq-Shortest-Path-Bridging-Standard.html). Tech Power Up. 7 May 2012. Retrieved 11 May 2012.
33. "Dr. Radia Perlman: One of the First Female Programmers and Inventor the Internet’s Protocols" (https://www.captechu.edu/blog/dr-radia-perlman-one-of-first-female-programmers-and-inventor-internets-protocols).

External links

- Cisco home page for the Spanning-Tree protocol family (http://www.cisco.com/en/US/tech/tk389/tk621/tsd_technology_support_protocol_home.html) (discusses CST, MISTP, PVST, PVST+, RSTP, STP)
- "Educational explanation of STP" (https://web.archive.org/web/20160304105810/http://www.cisco.com/image/gif/paws/10556/spanning_tree1.swf). Archived from the original (http://www.cisco.com/image/gif/paws/10556/spanning_tree1.swf) on 2016-03-04.
- STP article in the Wireshark wiki (https://wiki.wireshark.org/STP) Includes a sample PCAP-file of captured STP traffic.
- Perlman, Radia. "Algorhyme" (https://web.archive.org/web/20110719212324/http://www.csua.berkeley.edu/~ranga/humor/algorhyme.txt). University of California at Berkeley. Archived from the original (http://www.csua.berkeley.edu/~ranga/humor/algorhyme.txt) on 2011-07-19. Retrieved 2011-09-01.
- IEEE Standards
  - ANSI/IEEE 802.1D-2004 standard (http://standards.ieee.org/getieee802/download/802.1D-2004.pdf), section 17 discusses RSTP (Regular STP is no longer a part of this standard. This is pointed out in section 8.)
  - ANSI/IEEE 802.1Q-2005 standard (http://standards.ieee.org/getieee802/download/802.1Q-2005.pdf), section 13 discusses MSTP
- RFCs
  - RFC 2674–1999, proposed standard, Definitions of Managed Objects for Bridges with Traffic Classes, Multicast Filtering and Virtual LAN Extensions
  - RFC 1525–1993, - SBRIDGEMIB, proposed standard, Definitions of Managed Objects for Source Routing Bridges
  - RFC 1493-1993 - BRIDGEMIB, draft standard, Definitions of Managed Objects for Bridges
- Spanning Tree Direct vs Indirect Link Failures - CCIE Study (https://web.archive.org/web/20110812152242/http://blog.ipexpert.com/2010/03/22/spanning-tree-direct-vs-indirect-link-failu

res/)

- Spanning Tree Protocol Overview (https://web.archive.org/web/20151227040910/http://www.networkel.com/2015/10/spanning-tree-protocol-stp-no-loop.html)

Retrieved from "https://en.wikipedia.org/w/index.php?title=Spanning_Tree_Protocol&oldid=1046975882"

This page was last edited on 28 September 2021, at 11:24 (UTC).

Text is available under the Creative Commons Attribution-ShareAlike License; additional terms may apply. By using this site, you agree to the Terms of Use and Privacy Policy. Wikipedia® is a registered trademark of the Wikimedia Foundation, Inc., a non-profit organization.