



Multiple Spanning Tree Protocol

This chapter introduces you to Multiple Spanning Tree Protocol (MSTP) and Per-VLAN Rapid Spanning Tree (PVRST) and describes how you can configure MSTP and PVRST.

Table 1: Feature History Table

Feature Name	Release Information	Feature Description
MSTP and PVRST	Release 7.6.1	This feature is now supported on routers that have Cisco NC57 line cards installed and operate in native and compatibility modes.

- [Multiple Spanning Tree Protocol, on page 1](#)
- [Per-VLAN Rapid Spanning Tree, on page 7](#)
- [Information About Multiple Spanning Tree Protocol, on page 11](#)

Multiple Spanning Tree Protocol

The Multiple Spanning Tree Protocol (MSTP) is a Spanning Tree Protocols (STPs) variant that allows you to create multiple and independent spanning trees over the same physical network. You can configure the parameters for each spanning tree separately. You can select different network devices as the root bridge or different paths to form the loop-free topology. Therefore, you can block a given physical interface for some of the spanning trees and unblock for others.

After setting up multiple spanning tree instances, you can partition the set of VLANs in use. For example, you can assign VLANs 1–100 to spanning tree instance 1, VLANs 101–200 to spanning tree instance 2, VLANs 201–300 to spanning tree instance 3, and so on. Since each spanning tree has a different active topology with different active links, this has the effect of dividing the data traffic among the available redundant links based on the VLAN—a form of load balancing.

MSTP Supported Features

The routers support MSTP, as defined in IEEE 802.1Q-2005, on physical Ethernet interfaces and Ethernet Bundle interfaces. This includes the Port Fast, Backbone Fast, Uplink Fast and Root Guard features found in Cisco implementations of legacy STP, RSTP and PVST, as these are encompassed by the standard MSTP protocol. The routers can operate in either standard 802.1Q mode, or in Provide Edge (802.1ad) mode. In

provider edge mode, a different MAC address is used for bridge protocol data units (BPDUs), and any BPDUs received with the 802.1Q MAC address are forwarded transparently.

When you have not configured the **allow-bpdu-guard** command on MST default instance, and if one of the bridge ports receives legacy BPDU, the port enters **error-disable** state.



Note MSTP supports interoperation with RSTP as described in the 802.1Q standard. However, these features do not support interoperation with legacy STP.


BPDU Guard

The BPDU Guard feature allows you to protect against misconfiguration of edge ports. It is an enhancement to the MSTP port fast feature. When you configure port fast on an interface, MSTP considers that interface to be an edge port and removes it from consideration when calculating the spanning tree. When you configure BPDU Guard, MSTP additionally shuts down the interface using error-disable when an MSTP BPDU is received.

Flush Containment

The Flush Containment feature allows you to prevent unnecessary MAC flushes due to unrelated topology changes in other areas of a network. The following figure shows a network containing four devices. Two VLANs are in use: VLAN 1 is only used on device D, while VLAN 2 spans devices A, B and C. The two VLANs are in the same spanning tree instance, but do not share any links.

Figure 1: Flush Containment



If the link BC goes down, then in normal operation, as C brings up its blocked port, it sends out a topology change notification on all other interfaces, including towards D. This causes a MAC flush to occur for VLAN 1, even though the topology change which has taken place only affects VLAN 2.

Flush containment helps deal with this problem by preventing topology change notifications from being sent on interfaces on which no VLANs are configured for the MSTI in question. In the example network this would mean no topology change notifications would be sent from C to D, and the MAC flushes which take place would be confined to the right hand side of the network.

Bringup Delay

The Bringup Delay feature allows you to stop MSTP from considering an interface when calculating the spanning tree when the interface is not yet ready to forward traffic. This is useful when a line card first boots up, as the system may declare that the interfaces on that card are *Up* before the dataplane is fully ready to forward traffic. According to the standard, MSTP considers the interfaces as soon as they are declared *Up*, and this may cause it to move other interfaces into the blocking state if the new interfaces are selected instead.

Bringup delay solves this problem by adding a configurable delay period which occurs as interfaces that are configured with MSTP first come into existence. Until this delay period ends, the interfaces remain in blocking state, and are not considered when calculating the spanning tree.

Bringup delay only takes place when interfaces which are already configured with MSTP are created, for example, on a card reload. No delay takes place if an interface which already exists is later configured with MSTP.

Restrictions

These restrictions apply when using MSTP:

- You must enable MSTP must only on interfaces where the interface itself (if it is in L2 mode) or all of the subinterfaces have a simple encapsulation configured. These encapsulation matching criteria are considered simple:
 - Single-tagged 802.1Q frames
 - Double-tagged Q-in-Q frames (only the outermost tag is examined)
 - 802.1ad frames (if MSTP is operating in Provider Bridge mode)
 - Ranges or lists of tags (any of the above)
- If an L2 interface or subinterface is configured with an encapsulation that matches multiple VLANs, then all of those VLANs must be mapped to the same spanning tree instance. There is therefore a single spanning tree instance associated with each L2 interface or subinterface.
- All the interfaces or subinterfaces in a given bridge domain must be associated with the same spanning tree instance.
- Multiple subinterfaces on the same interface must not be associated with the same spanning tree instance, unless those subinterfaces are in the same split horizon group. In other words, hair-pinning is not possible. Across the network, L2 interfaces or subinterfaces must be configured on all redundant paths for all the VLANs mapped to each spanning tree instance. This is to avoid inadvertent loss of connectivity due to STP blocking of a port.



Caution A subinterface with a default or untagged encapsulation leads to an MSTP state machine failure.

- When you have not configured the **allow-bpdu-guard** command on MST default instance, and if one of the bridge ports receives legacy BPDU, the port enters **error-disable** state.

Configure MSTP

By default, STP is disabled on all interfaces. You must enable MSTP on each physical or Ethernet Bundle interface. When you configure MSTP on an interface, all the subinterfaces of that interface are automatically MSTP-enabled.

Perform these tasks to configure MSTP:

- Configure VLAN interfaces

- Configure L2VPN bridge-domains
- Configure MSTP

Configuration Example

```

/* Configure VLAN interfaces */
Router# configure
Router(config)# interface TenGigE0/0/0/2.1001 l2transport
Router(config-subif)# encapsulation dot1q 1001
Router(config)# interface TenGigE0/0/0/3.1001 l2transport
Router(config-subif)# encapsulation dot1q 1001
Router(config)# interface TenGigE0/0/0/14.1001 l2transport
Router(config-subif)# encapsulation dot1q 1001
Router(config)# interface TenGigE0/0/0/2.1021 l2transport
Router(config-subif)# encapsulation dot1q 1021
Router(config)# interface TenGigE0/0/0/3.1021 l2transport
Router(config-subif)# encapsulation dot1q 1021
Router(config)# interface TenGigE0/0/0/14.1021 l2transport
Router(config-subif)# encapsulation dot1q 1021
Router(config-subif)# commit

/* Configure L2VPN bridge-domains */
Router# configure
Router(config)# l2vpn bridge group mstp
Router(config-l2vpn-bg)# bridge-domain mstp1001
Router(config-l2vpn-bg-bd)# int TenGigE 0/0/0/2.1001
Router(config-l2vpn-bg-bd-ac)# exit
Router(config-l2vpn-bg-bd)# int TenGigE 0/0/0/3.1001
Router(config-l2vpn-bg-bd-ac)# exit
Router(config-l2vpn-bg-bd)# int TenGigE 0/0/0/14.1001
Router(config-l2vpn-bg-bd-ac)# exit
Router(config-l2vpn-bg-bd)# exit
Router(config-l2vpn-bg)# exit
Router(config-l2vpn-bg)# bridge-domain mstp1021
Router(config-l2vpn-bg-bd)# int TenGigE 0/0/0/2.1021
Router(config-l2vpn-bg-bd-ac)# exit
Router(config-l2vpn-bg-bd)# int TenGigE 0/0/0/3.1021
Router(config-l2vpn-bg-bd-ac)# exit
Router(config-l2vpn-bg-bd)# int TenGigE 0/0/0/14.1021
Router(config-l2vpn-bg-bd-ac)# commit

/* Configure MSTP */

Router# configure
Router(config)# spanning-tree mst abc
Router(config-mstp)# name mstp1
Router(config-mstp)# instance 1001
Router(config-mstp-inst)# vlan-ids 1001-1020
Router(config-mstp-inst)# exit
Router(config-mstp)# instance 1021
Router(config-mstp-inst)# vlan-ids 1021-1040
Router(config-mstp-inst)# exit
Router(config-mstp)# int tenGigE 0/0/0/2
Router(config-mstp-if)# exit
Router(config-mstp)# int tenGigE 0/0/0/3
Router(config-mstp-if)# exit
Router(config-mstp)# int tenGigE 0/0/0/14
Router(config-mstp-if)# commit

```

```

/* Configure MSTP Parameters */
Router#configure
Router(config)#spanning-tree mst a
Router(config-mstp)#bringup delay for 10 minutes
Router(config-mstp)#flush containment disable
Router(config-mstp)#name m1
Router(config-mstp)#revision 10
Router(config-mstp)#forward-delay 20
Router(config-mstp)#max age 40
Router(config-mstp)#transmit hold-count 8
Router(config-mstp)#provider-bridge
Router(config-mstp)#instance 101
Router(config-mstp-inst)#priority 8192
Router(config-mstp-inst)#vlan-id 2-1005
Router(config-mstp)#interface FastEthernet 0/0/0/1
Router(config-mstp-if)#instance 101 port-priority 160
Router(config-mstp-if)#portfast bpduguard
Router(config-mstp-if)#commit

```

Running Configuration

This section show MSTP running configuration.

```

!
Configure
/* Configure VLAN interfaces */
interface TenGigE0/0/0/2.1001 l2transport
 encapsulation dot1q 1001
!
interface TenGigE0/0/0/3.1001 l2transport
 encapsulation dot1q 1001
!
interface TenGigE0/0/0/14.1001 l2transport
 encapsulation dot1q 1001

interface TenGigE0/0/0/2.1021 l2transport
 encapsulation dot1q 1021
!
interface TenGigE0/0/0/3.1021
 l2transport
 encapsulation dot1q 1021
!
interface TenGigE0/0/0/14.1021 l2transport
 encapsulation dot1q 1021
!
/* Configure L2VPN Bridge-domains */
l2vpn
 bridge group mstp
  bridge-domain mstp1001
   interface TenGigE0/0/0/2.1001
    !
   interface TenGigE0/0/0/3.1001
    !
   interface TenGigE0/0/0/14.1001
    !
  bridge-domain mstp1021
   interface TenGigE0/0/0/2.1021
    !
   interface TenGigE0/0/0/3.1021
    !
   interface TenGigE0/0/0/14.1021
    !
!

```

```

/* Configure MSTP */
spanning-tree mst abc
name mstpl
instance 1001
  vlan-ids 1001-1020
!
instance 1021
  vlan-ids 1021-1040
!
interface TenGigE0/0/0/2
!
interface TenGigE0/0/0/3
!
interface TenGigE0/0/0/14

/* Configure MSTP Parameters */
spanning-tree mst a
bringup delay for 10 minutes
flush containment disable
name m1
revision 10
forward-delay 20
max hops 30
transmit hold-count 8
provider-bridge
instance 101
  priority 8192
  vlan-id 2-1005
interface FastEthernet 0/0/0/1
  instance 101 port-priority 160
  portfast bpduguard
!
!

```

Verification

Verify the MSTP configuration using the **show spanning-tree mst** command.

```

/* Verify the MSTP configuration */
Router# show spanning-tree mst abc instance 121
Mon Jan 23 12:11:48.591 UTC
Role:  ROOT=Root, DSGN=Designated, ALT=Alternate, BKP=Backup, MSTR=Master
State:  FWD=Forwarding, LRN=Learning, BLK=Blocked, DLY=Bringup Delayed

Operating in dot1q mode

MSTI 121:

  VLANS Mapped: 121-130

  Root ID    Priority    32768
            Address    dceb.9456.b9d4
            This bridge is the root
            Int Cost    0
            Max Age 20 sec, Forward Delay 15 sec

  Bridge ID  Priority    32768 (priority 32768 sys-id-ext 0)
            Address    dceb.9456.b9d4
            Max Age 20 sec, Forward Delay 15 sec
            Max Hops 20, Transmit Hold count 6

```

Interface	Port ID Pri.Nbr	Cost	Role	State	Designated Bridge ID	Port ID Pri.Nbr
BE1	128.1	10000	DSGN	FWD	32768 dceb.9456.b9d4	128.1
Te0/0/0/1	128.2	2000	DSGN	FWD	32768 dceb.9456.b9d4	128.2
Te0/0/0/16	128.3	2000	DSGN	FWD	32768 dceb.9456.b9d4	128.3
Te0/0/0/17	128.4	2000	DSGN	FWD	32768 dceb.9456.b9d4	128.4

Related Topics

- [Information About Multiple Spanning Tree Protocol, on page 11](#)
- [Multiple Spanning Tree Protocol, on page 1](#)

Associated Commands

- spanning-tree mst
- show spanning-tree mst

Per-VLAN Rapid Spanning Tree

Per-VLAN Rapid Spanning Tree (PVRST) or Rapid PVST or PVST+ is the IEEE 802.1w (RSTP) standard implemented per VLAN. A single instance of STP runs on each configured VLAN (if you do not manually disable STP). Each Rapid PVST+ instance on a VLAN has a single root switch. When you are running Rapid PVST+ you must configure the feature on all VLANs for a particular port.

PVRST uses point-to-point wiring to provide rapid convergence of the spanning tree. The spanning tree reconfiguration can occur in less than 1 second with PVRST (in contrast to 50 seconds with the default settings in the 802.1D STP).



Note PVRST supports one STP instance for each VLAN.

Using PVRST, STP convergence occurs rapidly. Each designated or root port in the STP sends out a BPDU every 2 seconds by default. On a designated or root port in the topology, if hello messages are missed three consecutive times, or if the maximum age expires, the port immediately flushes all protocol information in the table. A port considers that it loses connectivity to its direct neighbor root or designated port if it misses three BPDUs or if the maximum age expires. This rapid aging of the protocol information allows quick failure detection.

PVRST achieves rapid transition to the forwarding state only on edge ports and point-to-point links. Although the link type is configurable, the system automatically derives the link type information from the duplex setting of the port. Full-duplex ports are assumed to be point-to-point ports, while half-duplex ports are assumed to be shared ports.

PVRST has the following characteristics:

- You can configuration Forward Delay and Max Age timers globally and not per VLAN.

- You can configure Hello timer on per port basis and not per VLAN basis. The Hello timer configured on a port applies to all VLANs on that specific port.
- The cost of a spanning tree bundle port is always 10000. It is not affected by any of the following:
 - Number or speed of the bundle members
 - Logical or administrative operational status of the bundle member ports
 - Addition or deletion of bundle members
- Receiving BPDU on an interface configured with the BPDU Guard error-disables the physical interface as well as any layer-2 or layer-3 sub-interfaces configured on the physical interface.
- Only Ethernet Flow-points (EFPs) that are untagged or have a single VLAN tag can be protected by PVRST.
- If any one EFP in a bridge-domain is protected by PVRST, then all EFPs in that bridge domain must belong to the same VLAN.
- If any one EFP on a port is protected by PVRST, then all EFPs on that port must be protected by PVRST.
- PVRST supports 64 VLANs and 512 EFP's per router.

Configure PVRST

Perform this task to configure PVRST per VLAN.

Prerequisites

- Define L2 transport subinterfaces with VLAN encapsulation. PVRST does not support dot1ad encapsulation.
- Configure L2VPN bridge domains under bridge group for every VLAN running spanning tree.
- Configure corresponding l2transport subinterfaces in the bridge-domain.

Configuration Example

```
/* Configure PVRST */
Router# configure
Router(config)# spanning-tree pvrst stp
Router(config-pvrst)# forward-delay 10
Router(config-pvrst)# maximum age 10
Router(config-pvrst)# transmit hold-count 4
Router(config-pvrst)# vlan 200
Router(config-pvrst-vlan)# exit
Router(config-pvrst)# vlan 300
Router(config-pvrst-vlan)# exit
Router(config-pvrst)# 400
Router(config-pvrst-vlan)# exit
Router(config-pvrst)# interface Bundle-Ether1
Router(config-pvrst-if)# exit
Router(config-pvrst)# interface Bundle-Ether2
Router(config-if)# exit
Router(config-pvrst)# interface TenGigE0/0/0/21
```



```

/* Configure bridge domain */
Router(config)# l2vpn bridge group pvrst
Router(config-l2vpn-bg)# bridge-domain pvrst1001
Router(config-l2vpn-bg-bd)# interface Bundle-Ether1
Router(config-l2vpn-bg-bd-ac)# exit
Router(config-l2vpn-bg-bd)# interface Bundle-Ether2
Router(config-l2vpn-bg-bd-ac)# exit
Router(config-l2vpn-bg-bd)# interface TenGigE0/0/0/21
Router(config-l2vpn-bg-bd-ac)# commit

```

Running Configuration for PVRST

This section shows PVRST running configuration.

```

/* PVRST Configuration */
configure
spanning-tree pvrst stp
  forward-delay 10
  maximum age 10
  transmit hold-count 4
  vlan 200
  !
  vlan 300
  !
  vlan 400
  !
  interface Bundle-Ether1
  !
  interface Bundle-Ether2
  !
  interface TenGigE0/0/0/21
  !
/* Configure bridge-domain */
l2vpn
bridge group pvrst
  bridge-domain pvrst1001
  interface Bundle-Ether1
  !
  !
  interface Bundle-Ether2
  !
  interface TenGigE0/0/0/14.1001
  !
  interface TenGigE0/0/0/21

```

Verification

Verify the PVRST configuration using the **show spanning-tree pvrst** command.

```

Router# show spanning-tree pvrst stp
Mon Jan 20 22:56:16.242 UTC
Role:  ROOT=Root,  DSGN=Designated,  ALT=Alternate,  BKP=Backup
State: FWD=Forwarding,  LRN=Learning,  BLK=Blocked
VLAN 200:
  Root ID      Priority      32768
              Address      008a.9610.08d8
              Max Age 20 sec, Forward Delay 15 sec
  Bridge ID    Priority      32768 (priority 32768 sys-id-ext 0)

```

```

Address      00bc.6021.a4d8
Max Age 10 sec, Forward Delay 10 sec
Transmit Hold count 4
Interface    Port ID          Role State Designated      Port ID
            Pri.Nbr Cost                Bridge ID          Pri.Nbr
-----
Te0/0/0/14  128.2   2000    DSGN FWD  32768 00bc.6021.a4d8  128.2
Te0/0/0/20  128.3   2000    DSGN FWD  32768 00bc.6021.a4d8  128.3
Te0/0/0/26  128.4   2000    DSGN FWD  32768 00bc.6021.a4d8  128.4
Te0/0/0/27  128.5   2000    DSGN FWD  32768 00bc.6021.a4d8  128.5
Te0/0/0/38  128.6   2000    ALT  BLK  32768 008a.9610.08d8  128.4
Te0/0/0/6   128.1   2000    ROOT FWD  32768 008a.9610.08d8  128.2
-----
VLAN 300:
  Root ID    Priority    32768
             Address    008a.9610.08d8
             Max Age 20 sec, Forward Delay 15 sec
  Bridge ID  Priority    32768 (priority 32768 sys-id-ext 0)
             Address    00bc.6021.a4d8
             Max Age 10 sec, Forward Delay 10 sec
             Transmit Hold count 4
Interface    Port ID          Role State Designated      Port ID
            Pri.Nbr Cost                Bridge ID          Pri.Nbr
-----
Te0/0/0/14  128.2   2000    DSGN FWD  32768 00bc.6021.a4d8  128.2
Te0/0/0/20  128.3   2000    DSGN FWD  32768 00bc.6021.a4d8  128.3
Te0/0/0/26  128.4   2000    DSGN FWD  32768 00bc.6021.a4d8  128.4
Te0/0/0/27  128.5   2000    DSGN FWD  32768 00bc.6021.a4d8  128.5
Te0/0/0/38  128.6   2000    ALT  BLK  32768 008a.9610.08d8  128.4
Te0/0/0/6   128.1   2000    ROOT FWD  32768 008a.9610.08d8  128.2
-----
VLAN 400:
  Root ID    Priority    32768
             Address    008a.9610.08d8
             Max Age 20 sec, Forward Delay 15 sec
  Bridge ID  Priority    32768 (priority 32768 sys-id-ext 0)
             Address    00bc.6021.a4d8
             Max Age 10 sec, Forward Delay 10 sec
             Transmit Hold count 4
Interface    Port ID          Role State Designated      Port ID
            Pri.Nbr Cost                Bridge ID          Pri.Nbr
-----
Te0/0/0/14  128.2   2000    DSGN FWD  32768 00bc.6021.a4d8  128.2
Te0/0/0/20  128.3   2000    DSGN FWD  32768 00bc.6021.a4d8  128.3
Te0/0/0/26  128.4   2000    DSGN FWD  32768 00bc.6021.a4d8  128.4
Te0/0/0/27  128.5   2000    DSGN FWD  32768 00bc.6021.a4d8  128.5
Te0/0/0/38  128.6   2000    ALT  BLK  32768 008a.9610.08d8  128.4
Te0/0/0/6   128.1   2000    ROOT FWD  32768 008a.9610.08d8  128.2
-----

```

Related Topics

- [Information About Multiple Spanning Tree Protocol, on page 11](#)
- [Multiple Spanning Tree Protocol, on page 1](#)

Associated Commands

- `spanning-tree pvrst`
- `show spanning-tree pvrst`

Information About Multiple Spanning Tree Protocol

To configure Multiple Spanning Tree Protocol, you must understand these concepts:

Spanning Tree Protocol Overview

Ethernet is no longer just a link-layer technology used to interconnect network vehicles and hosts. Its low cost and wide spectrum of bandwidth capabilities coupled with a simple *plug and play* provisioning philosophy have transformed Ethernet into a legitimate technique for building networks, particularly in the access and aggregation regions of service provider networks.

Ethernet networks lacking a TTL field in the Layer 2 (L2) header and, encouraging or requiring multicast traffic network-wide, are susceptible to broadcast storms if loops are introduced. However, loops are a desirable property as they provide redundant paths. Spanning tree protocols (STP) are used to provide a loop free topology within Ethernet networks, allowing redundancy within the network to deal with link failures.

There are many variants of STP; however, they work on the same basic principle. Within a network that may contain loops, a sufficient number of interfaces are disabled by STP so as to ensure that there is a loop-free spanning tree, that is, there is exactly one path between any two devices in the network. If there is a fault in the network that affects one of the active links, the protocol recalculates the spanning tree so as to ensure that all devices continue to be reachable. STP is transparent to end stations which cannot detect whether they are connected to a single LAN segment or to a switched LAN containing multiple segments and using STP to ensure there are no loops.

STP Protocol Operation

All variants of STP operate in a similar fashion: STP frames (known as bridge protocol data units (BPDUs)) are exchanged at regular intervals over Layer 2 LAN segments, between network devices participating in STP. Such network devices do not forward these frames, but use the information to construct a loop free spanning tree.

The spanning tree is constructed by first selecting a device which is the *root* of the spanning tree (known as the root bridge), and then by determining a loop free path from the *root bridge* to every other device in the network. Redundant paths are disabled by setting the appropriate ports into a blocked state, where STP frames can still be exchanged but data traffic is never forwarded. If a network segment fails and a redundant path exists, the STP protocol recalculates the spanning tree topology and activates the redundant path, by unblocking the appropriate ports.

The selection of the root bridge within a STP network is determined by the lowest Bridge ID which is a combination of configured bridge priority and embedded mac address of each device. The device with the lowest priority, or with equal lowest priority but the lowest MAC address is selected as the root bridge.

Root port: is selected based on lowest root path cost to root bridge. If there is a tie with respect to the root path cost, port on local switch which receives BPDU with lowest sender bridge ID is selected as root port.

Designated port: Least cost port on local switch towards root bridge is selected as designated port. If there is a tie, lowest number port on local switch is selected as designated port.

The selection of the active path among a set of redundant paths is determined primarily by the port path cost. The port path cost represents the cost of transiting between that port and the root bridge - the further the port is from the root bridge, the higher the cost. The cost is incremented for each link in the path, by an amount that is (by default) dependent on the media speed. Where two paths from a given LAN segment have an equal cost, the selection is further determined by the lowest bridge ID of the attached devices, and in the case of

two attachments to the same device, by the configured port priority and port ID of the neighboring attached ports.

Once the active paths have been selected, any ports that do not form part of the active topology are moved to the blocking state.

Topology Changes

Network devices in a switched LAN perform MAC learning; that is, they use received data traffic to associate unicast MAC addresses with the interface out of which frames destined for that MAC address should be sent. If STP is used, then a recalculation of the spanning tree (for example, following a failure in the network) can invalidate this learned information. The protocol therefore includes a mechanism to notify topology changes around the network, so that the stale information can be removed (flushed) and new information can be learned based on the new topology.

A *Topology Change* notification is sent whenever STP moves a port from the blocking state to the forwarding state. When it is received, the receiving device flushes the MAC learning entries for all ports that are not blocked other than the one where the notification was received, and also sends its own topology change notification out of those ports. In this way, it is guaranteed that stale information is removed from all the devices in the network.

Variants of STP

There are many variants of the Spanning Tree Protocol:

- Legacy STP (STP)—The original STP protocol was defined in IEEE 802.1D-1998. This creates a single spanning tree which is used for all VLANs and most of the convergence is timer-based.
- Rapid STP (RSTP)—This is an enhancement defined in IEEE 802.1D-2004 to provide more event-based, and hence faster, convergence. However, it still creates a single spanning tree for all VLANs.
- Multiple STP (MSTP)—A further enhancement was defined in IEEE 802.1Q-2005. This allows multiple spanning tree instances to be created over the same physical topology. By assigning different VLANs to the different spanning tree instances, data traffic can be load-balanced over different physical links. The number of different spanning tree instances that can be created is restricted to a much smaller number than the number of possible VLANs; however, multiple VLANs can be assigned to the same spanning tree instance. The BPDUs used to exchange MSTP information are always sent untagged; the VLAN and spanning tree instance data is encoded inside the BPDU.
- Per-VLAN Rapid Spanning Tree (PVRST)—This feature is the IEEE 802.1w (RSTP) standard implemented per VLAN, and is also known as Rapid PVST or PVST+. A single instance of STP runs on each configured VLAN (if you do not manually disable STP). Each Rapid PVST+ instance on a VLAN has a single root switch. You can enable and disable STP on a per-VLAN basis when you are running Rapid PVST+.

PVRST uses point-to-point wiring to provide rapid convergence of the spanning tree. The spanning tree reconfiguration can occur in less than one second with PVRST (in contrast to 50 seconds with the default settings in the 802.1D STP).


- REP (Cisco-proprietary ring-redundancy protocol)—This is a Cisco-proprietary protocol for providing resiliency in rings. It is included for completeness, as it provides MSTP compatibility mode, using which, it interoperates with an MSTP peer.

MSTP Regions

Along with supporting multiple spanning trees, MSTP also introduces the concept of regions. A region is a group of devices under the same administrative control and have similar configuration. In particular, the configuration for the region name, revision, and the mapping of VLANs to spanning tree instances must be identical on all the network devices in the region. A digest of this information is included in the BPDUs sent by each device, so as to allow other devices to verify whether they are in the same region.


The following figure shows the operation of MST regions when bridges running MSTP are connected to bridges running legacy STP or RSTP. In this example, switches SW1, SW2, SW3, SW4 support MSTP, while switches SW5 and SW6 do not.

Figure 2: MST Interaction with Non-MST Regions



To handle this situation, an Internal Spanning Tree (IST) is used. This is always spanning tree instance 0 (zero). When communicating with non-MSTP-aware devices, the entire MSTP region is represented as a single switch. The logical IST topology in this case is shown in the following figure.

Figure 3: Logical Topology in MST Region Interacting with Non-MST Bridges



The same mechanism is used when communicating with MSTP devices in a different region. For example, SW5 in the above figure could represent a number of MSTP devices, all in a different region compared to SW1, SW2, SW3 and SW4.

MSTP Port Fast

MSTP includes a *Port Fast* feature for handling ports at the edge of the switched Ethernet network. For devices that only have one link to the switched network (typically host devices), there is no need to run MSTP, as there is only one available path. Furthermore, it is undesirable to trigger topology changes (and resultant MAC flushes) when the single link fails or is restored, as there is no alternative path.

By default, MSTP monitors ports where no BPDUs are received, and after a timeout, places them into *edge mode* whereby they do not participate in MSTP. However, this process can be speeded up (and convergence of the whole network thereby improved) by explicitly configuring edge ports as port fast.



Note

- You must disable and re-enable the port for Port Fast configuration to take effect. Use **shutdown** and **no shutdown** command (in interface configuration mode) to disable and re-enable the port.
- Port Fast is implemented as a Cisco-proprietary extension in Cisco implementations of legacy STP. However, it is encompassed in the standards for RSTP and MSTP, where it is known as Edge Port.

MSTP Root Guard

In networks with shared administrative control, it may be desirable for the network administrator to enforce aspects of the network topology and in particular, the location of the root bridge. By default, any device can become the root bridge for a spanning tree, if it has a lower priority or bridge ID. However, a more optimal forwarding topology can be achieved by placing the root bridge at a specific location in the centre of the network.



Note

The administrator can set the root bridge priority to 0 in an effort to secure the root bridge position; however, this is no guarantee against another bridge which also has a priority of 0 and has a lower bridge ID.

The root guard feature provides a mechanism that allows the administrator to enforce the location of the root bridge. When root guard is configured on an interface, it prevents that interface from becoming a root port (that is, a port via which the root can be reached). If superior information is received via BPDUs on the interface that would normally cause it to become a root port, it instead becomes a backup or alternate port. In this case, it is placed in the blocking state and no data traffic is forwarded.

The root bridge itself has no root ports. Thus, by configuring root guard on every interface on a device, the administrator forces the device to become the root, and interfaces receiving conflicting information are blocked.



Note

Root Guard is implemented as a Cisco-proprietary extension in Cisco implementations of legacy STP and RSTP. However, it is encompassed in the standard for MSTP, where it is known as Restricted Role.

MSTP Topology Change Guard

In certain situations, it may be desirable to prevent topology changes originating at or received at a given port from being propagated to the rest of the network. This may be the case, for example, when the network is not

under a single administrative control and it is desirable to prevent devices external to the core of the network from causing MAC address flushing in the core. This behavior can be enabled by configuring Topology Change Guard on the port.



Note Topology Change Guard is known as *Restricted TCN* in the MSTP standard.
