

This article describes how to identify and resolve problems that might occur when implementing the Spanning Tree Protocol (STP).

Guide Contents
Troubleshooting Overview
Troubleshooting Installs, Upgrades, and Reboots
Troubleshooting Licensing
Troubleshooting VDCs
Troubleshooting CFS
Troubleshooting Ports
Troubleshooting vPCs
Troubleshooting VLANs
<i>Troubleshooting STP (this section)</i>
Troubleshooting Routing
Troubleshooting Unicast Traffic
Troubleshooting WCCP
Troubleshooting Memory
Troubleshooting Packet Flow Issues
Troubleshooting FCoE
Before Contacting Technical Support
Troubleshooting Tools and Methodology

Contents

- [1 Information About Troubleshooting STP](#)
- [2 Initial Troubleshooting Checklist](#)
- [3 Troubleshooting STP Data Loops](#)
- [4 Troubleshooting Excessive Packet Flooding](#)
- [5 Troubleshooting Convergence Time Issues](#)
- [6 Securing the Network Against Forwarding Loops](#)
- [7 See Also](#)
- [8 Further Reading](#)
- [9 External Links](#)

Information About Troubleshooting STP

STP provides a loop-free network at the Layer 2 level. Layer 2 LAN ports send and receive STP frames at regular intervals. Network devices do not forward these frames but use the frames to construct a loop-free path. See the [Cisco NX-OS Layer 2 Switching Configuration Guide](#) for more information on STP.

Follow these guidelines when configuring STP:

- If you are running private VLANs with multiple STP (MST), verify that all secondary VLANs belong to the same MST instance as that of the primary VLANs.
- Disabling spanning tree on the native VLAN of an 802.1Q trunk when you are working in Rapid PVST+ spanning tree mode can cause a spanning tree loop on that VLAN. We recommend that you leave spanning tree enabled on the native VLAN of the 802.1Q trunks. Make sure that your network has no physical loops before you disable spanning tree.
- When you connect two Cisco switches through 802.1Q trunks, the switches exchange spanning tree bridge protocol data units (BPDUs) on each VLAN allowed on the trunks. The BPDUs on the native VLAN of the trunk are sent untagged to the reserved IEEE 802.1D spanning tree multicast MAC address (01-80-C2-00-00-00). The BPDUs on all other VLANs on the trunk are sent tagged to the reserved Cisco Shared Spanning Tree (SSTP) multicast MAC address (01-00-0c-cc-cc-cd).
- In STP, the port-channel bundle is considered as a single port. The port cost is the aggregation of all the configured port costs that are assigned to that channel.
- When a secondary VLAN is associated with the primary VLAN, the STP parameters of the primary VLAN, such as bridge priorities, are propagated to the secondary VLAN. However, STP parameters do not necessarily propagate to other devices. You should manually check the STP configuration to ensure that the spanning tree topologies for the primary, isolated, or community VLANs match exactly so that the VLANs can share the same forwarding database.
- For normal trunk ports, note the following:
 - ◆ There is a separate instance of STP for each VLAN in the private VLAN.
 - ◆ STP parameters for the primary and all secondary VLANs must match.
 - ◆ The primary and all associated secondary VLANs should be in the same MST instance.
 - ◆ The duplex configuration for both sides of the link should be set to full to prevent collisions under heavy traffic conditions.
 - ◆ In MST mode, a misconfiguration cannot be detected if you configure one end of a link in trunk mode and the other end of the link in access mode. This misconfiguration will cause an STP loop.
- For nontrunking ports, note the following:
 - ◆ STP is aware only of the primary VLAN for any private VLAN host port; STP does not run on secondary VLANs on a host port.
- For Rapid PVST+ on private VLANs, note the following:
 - ◆ On a trunk port, the primary and secondary private VLANs are two different logical ports and *must* have the exact same STP topology.
 - ◆ On access ports, STP sees only the primary VLAN.

Note: In some cases, the configuration is accepted with no error messages, but the commands have no effect.

Initial Troubleshooting Checklist

Troubleshooting an STP problem involves gathering information about the configuration and connectivity of individual devices and the entire network.

Begin troubleshooting STP issues by checking the following issues first:

Checklist	Check off
Verify the type of spanning tree configured on your device.	
Verify the network topology including all interconnected ports and switches. Identify all redundant paths on the network and verify that the redundant paths are blocking.	
Use the show spanning-tree summary totals command to verify that the total number of logical interfaces in the Active state are less than the maximum allowed. See the Cisco NX-OS Layer 2 Switching Configuration Guide for information on these limits.	
Verify the primary and secondary root bridge and any configured Cisco extensions.	

Use the following commands to view STP configuration and operational details:

- **show running-config spanning-tree**
- **show spanning-tree summary**
- **show spanning-tree detail**
- **show spanning-tree bridge**
- **show spanning-tree mst**
- **show spanning-tree mst configuration**
- **show spanning-tree interface *interface-type slot/port* [detail]**
- **show tech-support stp**
- **show spanning-tree vlan**

Use the **show spanning-tree blockedports** command to display the ports that are blocked by STP.

Use the **show mac address-table dynamic vlan** command to determine if learning or aging occurs at each node.

Troubleshooting STP Data Loops

Data loops are a common problem in STP networks. Some of the symptoms of a data loop are as follows:

- High link utilization, up to 100 percent
- High CPU and backplane traffic utilization
- Constant MAC address relearning and flapping
- Excessive output drops on an interface

To troubleshoot STP loops, follow these steps:

1. Identify the ports involved in the loop by looking at the interfaces with high link utilization.

```
switch# show interface ethernet 2/1 | include rate
```

```
1 minute input rate 19968 bits/sec, 0 packets/sec
1 minute output rate 3952023552 bits/sec, 957312 packets/sec
```

2. Shut down or disconnect the affected ports.

```
switch(config)# interface ethernet 2/1
```

```
switch(config-if) # shutdown
```

3. Locate every switch in the redundant paths using your network topology diagram.
4. Verify that the switch lists the same STP root bridge as the other nonaffected switches.

```
switch# show spanning-tree vlan 9
```

```
VLAN0009
Spanning tree enabled protocol rstp
  Root ID    Priority    32777''
            Address    0018.bad7.db15''
            Cost      4
...

```

5. Verify that the root port is correctly identified as the port with the lowest cost to the root bridge.

```
switch# show spanning-tree vlan 9
```

```
VLAN0009
Spanning tree enabled protocol rstp
Root ID    Priority    32777
Address    0018.bad7.db15
Cost       4
Port 385 (Ethernet3/1)
Hello Time 2 sec Max Age 20 sec Forward Delay 15 sec

```

6. Verify that the root port and alternate ports are regularly receiving BPDUs.

```
switch# show spanning-tree interface ethernet 3/1 detail
```

```
Port 385 (Ethernet3/1) of VLAN0001 is root forwarding
  Port path cost 4, Port priority 128, Port Identifier 128.385
  Designated root has priority 32769, address 0018.bad7.db15
  Designated bridge has priority 32769, address 0018.bad7.db15
  Designated port id is 128.385, designated path cost 0
  Timers: message age 16, forward delay 0, hold 0
  Number of transitions to forwarding state: 1
  The port type is network by default
  Link type is point-to-point by default
  BPDUs: sent 1265, received 1269

```

7. If the received BPDU counter is not incremented, check if the BPDUs are received by the internal packet manager.

```
switch# show system internal pktmgr interface ethernet 3/1
```

```
Ethernet3/1, ordinal: 36
SUP-traffic statistics: (sent/received)
  Packets: 120210 / 15812
  Bytes: 8166401 / 1083056
  Instant packet rate: 5 pps / 5 pps
  Average packet rates (1min/5min/15min/EWMA):
  Packet statistics:
    Tx: Unicast 0, Multicast 120210
        Broadcast 0
    Rx: Unicast 0, ' Multicast 15812''
        Broadcast 0

```

```
switch# show system internal pktmgr client 303
```

```
Client uuid: 303, 2 filters
  Filter 0: EthType 0x4242, Dmac 0180.c200.0000
  Filter 0: EthType 0x010b, Snap 267, Dmac 0100.0ccc.cccd
Options: TO 0, Flags 0x1, AppId 0, Epid 0
Ctrl SAP: 171, Data SAP 177 (1)
Rx: 28356632, Drop: 0, Tx: 35498365, Drop: 0
```

8. If the BPDUs are not received by the packet manager, check the hardware packet statistic (error drop) counters.

```
switch# show interface counters errors
```

```
-----
Port Align-Err FCS-Err Xmit-Err Rcv-Err UnderSize OutDiscards
-----
mgmt0  --      --      --      --      --      --
Eth1/1  0        0        0        0        0        0
Eth1/2  0        0        0        0        0        0
Eth1/3  0        0        0        0        0        0
Eth1/4  0        0        0        0        0        0
Eth1/5  0        0        0        0        0        0
Eth1/6  0        0        0        0        0        0
Eth1/7  0        0        0        0        0        0
Eth1/8  0        0        0        0        0        0
```

9. Check that the designated port is regularly sending BPDUs.

```
switch# show spanning-tree interface ethernet 3/1 detail
```

```
Port 385 (Ethernet3/1) of VLAN0001 is root forwarding
  Port path cost 4, Port priority 128, Port Identifier 128.385
  Designated root has priority 32769, address 0018.bad7.db15
  Designated bridge has priority 32769, address 0018.bad7.db15
  Designated port id is 128.385, designated path cost 0
  Timers: message age 16, forward delay 0, hold 0
  Number of transitions to forwarding state: 1
  The port type is network by default
  Link type is point-to-point by default
  BPDU: sent 1265, received 1269
```

10. If the BPDU send counter is incrementing, check if BPDUs are transmitted by the packet manager.

```
switch# show system internal pktmgr interface ethernet 3/1
```

```
Ethernet3/1, ordinal: 36
  SUP-traffic statistics: (sent/received)
  Packets: 120210 / 15812
  Bytes: 8166401 / 1083056
  Instant packet rate: 5 pps / 5 pps
  Average packet rates(1min/5min/15min/EWMA):
  Packet statistics:
    Tx: Unicast 0, M'' ulticast 120210''
       Broadcast 0
    Rx: Unicast 0, Multicast 15812
       Broadcast 0
```

```
switch# show system internal pktmgr client 303
```

```
Client uuid: 303, 2 filters
  Filter 0: EthType 0x4242, Dmac 0180.c200.0000
  Filter 0: EthType 0x010b, Snap 267, Dmac 0100.0ccc.cccd
  Options: TO 0, Flags 0x1, AppId 0, Epid 0
  Ctrl SAP: 171, Data SAP 177 (1)
  Rx: 28356632, Drop: 0, Tx: 35498365, Drop: 0
```

11. If the packet manager BPDU sent counters is incrementing, check the hardware packet statistic counters for a possible BPDU error drop.

```
switch# show interface counters errors
```

```
-----
Port  Align-Err  FCS-Err  Xmit-Err  Rcv-Err  UnderSize  OutDiscards
-----
mgmt0  --         --         --         --         --         --
Eth1/1  0          0          0          0          0          0
Eth1/2  0          0          0          0          0          0
Eth1/3  0          0          0          0          0          0
Eth1/4  0          0          0          0          0          0
Eth1/5  0          0          0          0          0          0
Eth1/6  0          0          0          0          0          0
Eth1/7  0          0          0          0          0          0
Eth1/8  0          0          0          0          0          0
```

Troubleshooting Excessive Packet Flooding

Unstable STP topology changes can trigger excessive packet flooding in your STP network. With Rapid STP or Multiple STP (MST), a change of the port's state to forwarding, as well as the role change from designated to root can trigger a topology change. Rapid STP immediately flushes the Layer 2 forwarding table. 802.1D shortens the aging time. The immediate flushing of the forwarding table restores connectivity faster but causes more flooding.

In a stable topology, a topology change should not trigger excessive flooding. Link flaps can cause a topology change, so continuous link flaps can cause repetitive topology changes and flooding. Flooding slows the network performance and can cause packet drops on an interface.

To troubleshoot excessive flooding, follow these steps:

1. Determine the source of the excessive topology change.

```
switch# show spanning-tree vlan 9 detail
```

```
VLAN0009 is executing the rstp compatible Spanning Tree protocol
  Bridge Identifier has priority 32768, sysid 9, address 0018.bad8.27ad
  Configured hello time 2, max age 20, forward delay 15
  Current root has priority 32777, address 0018.bad7.db15
  Root port is 385 (Ethernet3/1), cost of root path is 4
  Topology change flag not set, detected flag not set
  '' Number of topology changes 8 last change occurred 1:32:11 ago''
  '' from Ethernet3/1''
  Times: hold 1, topology change 35, notification 2
...
```

2. Determine the interface where the topology change occurred.

```
switch# show spanning-tree vlan 9 detail
```

```
VLAN0009 is executing the rstp compatible Spanning Tree protocol
  Bridge Identifier has priority 32768, sysid 9, address 0018.bad8.27ad
  Configured hello time 2, max age 20, forward delay 15
  Current root has priority 32777, address 0018.bad7.db15
  Root port is 385 (Ethernet3/1), cost of root path is 4
  Topology change flag not set, detected flag not set
  Number of topology changes 8 last change occurred 1:32:11 ago
    '' from Ethernet3/1''
  Times: hold 1, topology change 35, notification 2
  ...
```

3. Repeat step 2 on devices connected to the interface until you can isolate the device that originated the topology change.

4. Check for link flaps on the interfaces on this device.

Troubleshooting Convergence Time Issues

STP convergence can take longer than expected or result in an unexpected final network topology.

To troubleshoot convergence issues, check the following issues:

- Errors in the documented network topology diagram.
- Misconfiguration-Check that the timers, diameter, Cisco extension features such as Bridge Assurance, Root Guard, BPDU Guard, and so on are not misconfigured.
- Overloaded switch CPU during convergence that exceeds the recommended logical port (port-vlan) limit.

Note: The recommended scalability limits are system wide and not per VDC.

- Software defects that affect STP.

Securing the Network Against Forwarding Loops

To handle the inability of STP to deal correctly with certain failures, Cisco has developed a number of features and enhancements to protect the networks against forwarding loops.

Troubleshooting STP helps to isolate and find the cause for a particular failure, while the implementation of these enhancements is the only way to secure the network against forwarding loops.

To protect your network against forwarding loops, follow these steps:

1. Enable the Cisco-proprietary Unidirectional Link Detection (UDLD) protocol on all the switch-to-switch links. See the UDLD section in the [Cisco NX-OS Interfaces Configuration Guide](#).

2. Set up the Bridge Assurance feature by configuring all the switch-to-switch links as the spanning tree network port type.

Note: You should enable the Bridge Assurance feature on both sides of the links or Cisco NX-OS will put the port in the blocked state because of a Bridge Assurance inconsistency.

3. Set up all the end-station ports as a spanning-tree edge port type.

You must set up the STP edge port to limit the amount of topology change (TC) notices and subsequent flooding that can affect the performance of the network. Use this command only with ports that connect to end stations. Otherwise, an accidental topology loop can cause a data-packet loop and disrupt the device and network operation.

4. Enable the Link Aggregation Control Protocol (LACP) for port channels to avoid any port-channel misconfiguration issues. See the LACP section in the [Cisco NX-OS Interfaces Configuration Guide](#).

Do not disable autonegotiation on the switch-to-switch links. Autonegotiation mechanisms can convey remote fault information, which is the quickest way to detect failures at the remote side. If failures are detected at the remote side, the local side brings down the link even if the link is still receiving pulses.

Caution! Be careful when you change STP timers. STP timers are dependent on each other and changes can impact the entire network.

5. (Optional) To prevent denial-of-service attacks, use the **spanning-tree loopguard default** command to secure the network STP perimeter with Root Guard. Root Guard and BPDU Guard allow you to secure STP against influence from the outside.

6. Use the **spanning-tree bpduguard enable** command to enable BPDU Guard on STP edge ports to prevent STP from being affected by unauthorized network devices (such as hubs, switches, and bridging routers) that are connected to the ports.

Root Guard prevents STP from outside influences. BPDU Guard shuts down the ports that are receiving any BPDUs (not only superior BPDUs).

Note: Short-living loops are not prevented by Root Guard or BPDU Guard if two STP edge ports are connected directly or through the hub.

7. Use the **vlan** command to configure separate VLANs and avoid user traffic on the management VLAN. The management VLAN is contained to a building block, not the entire network.

8. Use the **spanning-tree vlan *vlan-range* root primary** command to configure a predictable STP root.

9. Use the **spanning-tree vlan *vlan-range* root secondary** command to configure a predictable backup STP root placement.

You must configure the STP root and backup STP root so that convergence occurs in a predictable way and builds optimal topology in every scenario. Do not leave the STP priority at the default value.

See Also

[Cisco NX-OS/IOS STP Comparison](#)

Further Reading

The following links contain further information on this topic from Cisco.com:

[Cisco Nexus 7000 Series NX-OS Layer 2 Switching Configuration Guide](#)

External Links

External links contain content developed by external authors. Cisco does not review this content for accuracy.