

Troubleshooting STP

- About Troubleshooting STP, on page 1
- Initial Troubleshooting STP Checklist, on page 1
- Troubleshooting STP Data Loops, on page 2
- Troubleshooting Excessive Packet Flooding, on page 6
- Troubleshooting Convergence Time Issues, on page 7
- Securing the Network Against Forwarding Loops, on page 7

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. For more information on Layer 2, see the *Cisco Nexus 9000 Series Layer 2 Configuration Guide*.

Initial Troubleshooting STP 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	Done
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. For information on these limits, see the *Cisco Nexus 9000 Series NX-OS Layer 2 Switching Configuration Guide*.	
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

When the l2fm logging level is greater than or equal to 4, the switch logs occurrences of host MAC address flapping to help you locate STP data loops. If it detects a MAC address move within less than 1 second and if 10 consecutive moves occur, the switch disables learning on the VLAN for one of the ports between which the MAC address is moving. Learning is disabled for 120 seconds and reenabled automatically. Syslogs are generated while learning is disabled and enabled. You can configure the logging level using the **logging level l2fm** *log-level* command.

SUMMARY STEPS

1. switch# **show interface** *interface-type slot*/*port* **include rate**
2. switch(config)# **interface** *interface-type slot*/*port*
3. switch(config-if)# **shutdown**
4. switch(config-if)# **show spanning-tree vlan** *vlan-id*
5. (Optional) switch(config-if)# **show spanning-tree interface** *interface-type slot*/*port* **detail**
6. (Optional) switch(config-if)# **show system internal pktmgr interface** *interface-type slot*/*port*
7. (Optional) switch(config-if)# **show system internal pktmgr client** *client-id*
8. (Optional) switch(config-if)# **show interface counters errors**

DETAILED STEPS

Procedure

	Command or Action	Purpose

Step 1	switch# **show interface** *interface-type slot*/*port* **include rate**
	Example:
	`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`
		Identifies the ports involved in the loop by looking at the interfaces with high link utilization.

Step 2	switch(config)# **interface** *interface-type slot*/*port* **Example:** `switch(config)# interface ethernet 2/1`	Configures the interface type and location.
Step 3	switch(config-if)# **shutdown** **Example:** `switch(config-if)# shutdown`	Shuts down or disconnects the affected ports. After disconnecting the affected ports, locate every switch in the redundant paths using your network topology diagram.

Step 4	switch(config-if)# **show spanning-tree vlan** *vlan-id*
	Example:
	`switch(config-if)# show spanning-tree vlan 9`
	`VLAN0009`
	`Spanning tree enabled protocol rstp`
	`Root ID Priority 32777''`
	`Address 0018.bad7.db15''`
	`Cost 4`
	`...`
		Verifies that the switch lists the same STP root bridge as the other nonaffected switches.

Step 5	(Optional) switch(config-if)# **show spanning-tree interface** *interface-type slot*/*port* **detail**
	Example:
	`switch(config-if)# 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`
		Verifies that the root port and alternate ports are regularly receiving BPDUs.

	Command or Action	Purpose
	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	

Step 6	(Optional) switch(config-if)# **show system internal pktmgr interface** *interface-type slot/port*
	Example:
	`switch(config-if)# 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`
		Checks if the BPDUs are received by the internal packet manager.

Step 7	(Optional) switch(config-if)# **show system internal pktmgr client** *client-id*
	Example:
	`switch(config-if)# 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`
		Checks if the BPDUs are received by the client.

Step 8	(Optional) switch(config-if)# **show interface counters errors**
	Example:
	`switch(config-if)# 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`
		Checks the hardware packet statistic (error drop) counters.

	Command or Action	Purpose
	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	

Example

This example shows 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
```

This example shows how to 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
```

This example shows how to 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.

SUMMARY STEPS

1. switch# **show spanning-tree vlan** *vlan-id* **detail**
2. switch# **show spanning-tree vlan** *vlan-id* **detail**

DETAILED STEPS

Procedure

	Command or Action	Purpose

Step 1	switch# **show spanning-tree vlan** *vlan-id* **detail**
	Example:
	`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`
	`...`
		Determines the source of the excessive topology change.

	Command or Action	Purpose

Step 2	switch# **show spanning-tree vlan** *vlan-id* **detail**
	Example:
	`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`
	`  ...`
		Determines the interface where the topology change occurred.
		Repeat this step on devices connected to the interface until you can isolate the device that originated the topology change.
		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 of the timers; diameter; Cisco extension features such as bridge assurance, root guard, and BPDU guard; and so on.
- Overloaded switch CPU during convergence that exceeds the recommended logical port (port-vlan) limit.
- 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.

Before you begin

- Enable the Cisco-proprietary Unidirectional Link Detection (UDLD) protocol on all the switch-to-switch links. For information, see the *Cisco Nexus 9000 Series NX-OS Interfaces Configuration Guide*.
- 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. Otherwise, Cisco NX-OS will put the port in the blocked state because of a bridge assurance inconsistency.

- 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 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.

- Enable the Link Aggregation Control Protocol (LACP) for port channels to avoid any port-channel misconfiguration issues. For information, see the *Cisco Nexus 9000 Series 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.

SUMMARY STEPS

1. (Optional) switch(config)# **spanning-tree loopguard default**
2. switch(config)# **spanning-tree bpduguard enable**
3. switch(config)# **vlan** *vlan-range*
4. switch(config)# **spanning-tree vlan** *vlan-range* **root primary**
5. switch(config)# **spanning-tree vlan** *vlan-range* **root secondary**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	(Optional) switch(config)# **spanning-tree loopguard default** **Example:** `switch(config)# spanning-tree loopguard default`	Secures the network STP perimeter with root guard. Root guard and BPDU guard allow you to secure STP against influence from the outside.
Step 2	switch(config)# **spanning-tree bpduguard enable** **Example:** `switch(config)# spanning-tree bpduguard enable`	Enables 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.

	Command or Action	Purpose
		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.
Step 3	switch(config)# **vlan** *vlan-range* **Example:** `switch(config)# vlan 9`	Configures separate VLANs and avoids user traffic on the management VLAN. The management VLAN is contained to a building block, not the entire network.
Step 4	switch(config)# **spanning-tree vlan** *vlan-range* **root primary** **Example:** `switch(config)# spanning-tree vlan 9 root primary`	Configures a predictable STP root.
Step 5	switch(config)# **spanning-tree vlan** *vlan-range* **root secondary** **Example:** `switch(config)# spanning-tree vlan 12 root secondary`	Configures 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.