

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Conventions](#)

[Purpose of the Topology Change Mechanism](#)

[Principle of Operation](#)

[Notify the Root Bridge](#)

[Broadcast the Event to the Network](#)

[What to Do When There are Many Topology Changes in the Network](#)

[Flooded Traffic](#)

[Problem in ATM LANE Bridged Environments](#)

[Avoid TCN Generation with the portfast Command](#)

[Track the Source of a TCN](#)

[Conclusion](#)

[Related Information](#)

Introduction

When you monitor Spanning-Tree Protocol (STP) operations, you may be concerned when you see topology change counters that increment in the statistics log. Topology changes are normal in STP. But, too many of them can have an impact on network performances. This document explains that the purpose of this topology is to:

- Change the mechanism in per-VLAN spanning tree (PVST) and PVST+ environments.
- Determine what triggers a topology change event.
- Describe issues related to the topology change mechanism.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Conventions

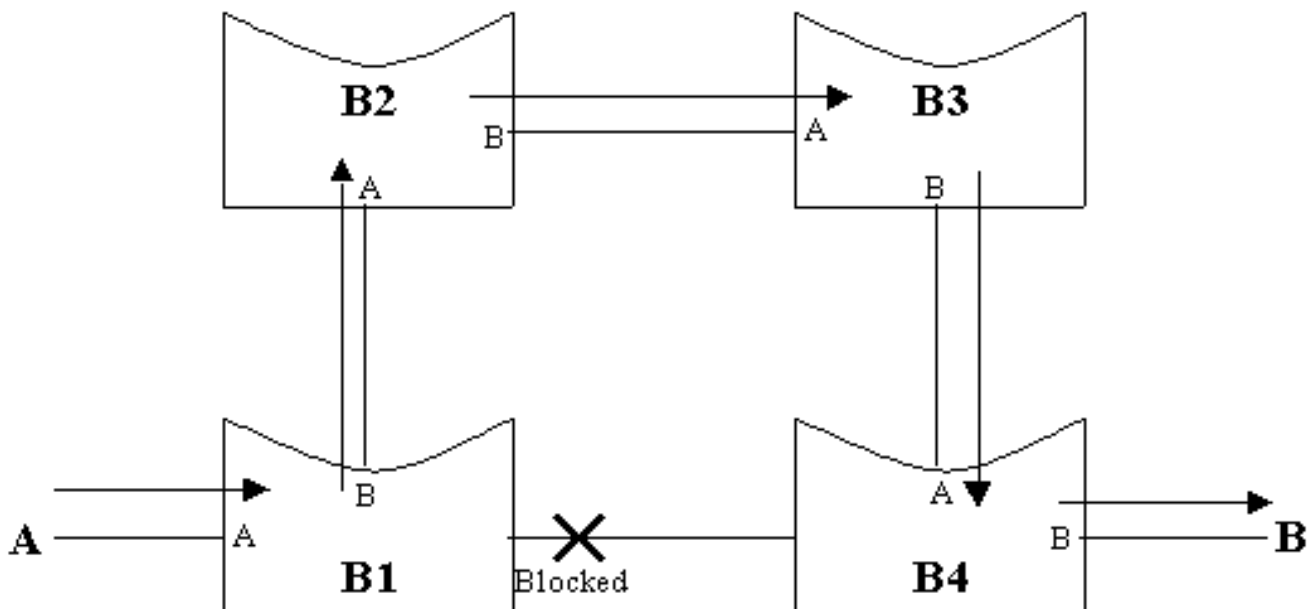
For more information on document conventions, refer to the [Cisco Technical Tips Conventions](#).

Purpose of the Topology Change Mechanism


Learning from the frames it receives, a bridge creates a table that associates to a port the Media Access Control (MAC) addresses of the hosts that can be reached through this port. This table is used to forward frames directly to their destination port. Therefore, flooding is avoided.

Default aging time for this table is 300 seconds (five minutes). Only after a host has been silent for five minutes, its entry disappears from the table of the bridge. Here is an example that shows why you could want this aging to be faster:

In this network, assume that bridge B1 is blocking its link to B4. A and B are two stations that have an established connection. Traffic from A to B goes to B1, B2, B3, and then B4. The scheme shows the MAC addresses table learned by the four bridges in this situation:



Now, assume the link between B2 and B3 fails. Communication between A and B is interrupted at least until B1 puts its port to B4 in forwarding mode (a maximum of 50 seconds with default parameters). However, when A wants to send a frame to B, B1 still has an entry that leads to B2 and the packet is sent to a black hole. The same applies when B wants to reach A. Communication is lost for five minutes, until the entries for A and B MAC addresses age out.



The forwarding databases implemented by bridges are very efficient in a **stable** network. But, there are many situations where the five minute aging time is a problem after the topology of the network has changed. The topology change mechanism is a workaround for that kind of problem. As soon as a bridge detects a change in the topology of the network (a link that goes down or goes to forwarding), it advertises the event to the whole bridged network.

The Principle of Operation section explains how this is practically implemented. Every bridge is then notified and reduces the aging time to `forward_delay` (15 seconds by default) for a certain period of time (`max_age + forward_delay`). It is more beneficial to reduce the aging time instead of clearing the table because currently active hosts, that effectively transmit traffic, are not cleared from the table.

In this example, as soon as bridge B2 or B3 detects the link going down, it sends topology change notifications. All bridges become aware of the event and reduce their aging time to 15 seconds. As B1 does not receive any packet from B on its port leading to B2 in fifteen seconds, it ages out the entry for B on this port. The same happens to the entry for A on the port that leads to B3 on B4. Later when the link between B1 and B4 goes to forwarding, traffic is immediately flooded and re-learned on this link.

Principle of Operation

This section explains how a bridge advertises a topology change at the Bridge Protocol Data Unit (BPDU) level.

It has already been briefly explained when a bridge considers it detected a topology change. The exact definition is:


- When a port that was forwarding is going down (blocking for instance).
- When a port transitions to forwarding and the bridge has a designated port. (This means that the bridge is not standalone.)

The process to send a notification to all bridges in the network involves two steps:

- The bridge notifies the root bridge of the spanning tree.
- The root bridge "broadcasts" the information into the whole network.

Notify the Root Bridge

In normal STP operation, a bridge keeps receiving configuration BPDUs from the root bridge on its root port. But, it never sends out a BPDU toward the root bridge. In order to achieve that, a special BPDU called the topology change notification (TCN) BPDU has been introduced. Therefore, when a bridge needs to signal a topology change, it starts to send TCNs on its root port. The designated bridge receives the TCN, acknowledges it, and generates another one for its own root port. The process continues until the TCN hits the root bridge.




Bridge B notifies a topology change by sending a TCN on its root port. The TCN is acknowledged and forwarded up to the root bridge R.

The TCN is a very simple BPDU that contains absolutely no information that a bridge sends out every hello_time seconds (this is locally configured hello_time, not the hello_time specified in configuration BPDUs). The designated bridge acknowledges the TCN by immediately sending back a normal configuration BPDU with the topology change acknowledgement (TCA) bit set. The bridge that notifies the topology change does not stop sending its TCN until the designated bridge has acknowledged it. Therefore, the designated bridge answers the TCN even though it does not receive configuration BPDU from its root.

Broadcast the Event to the Network

Once the root is aware that there has been a topology change event in the network, it starts to send out its configuration BPDUs with the topology change (TC) bit set. These BPDUs are relayed by every bridge in the network with this bit set. As a result all bridges become aware of the topology change situation and it can reduce its aging time to forward_delay. Bridges receive topology change BPDUs on both forwarding and blocking ports.

The TC bit is set by the root for a period of max_age + forward_delay seconds, which is 20+15=35 seconds by default.



The root R sets the TC bits in its bpdus. This bpdus is relayed to the whole network.

What to Do When There are Many Topology Changes in the Network


Here are some of the problems that can be generated by TCN. It is followed by some information on how to limit topology changes and find from where they come .

If you have the output of a **show-tech support** command from your Cisco device, you can use [Cisco CLI Analyzer](#) (registered customers only) to display potential issues and fixes. To use [Cisco CLI Analyzer](#) (registered customers only) , you must be a registered customer, be logged in, and have JavaScript enabled.


Flooded Traffic

The more hosts are in the network, the higher are the probabilities of getting a topology change.

For instance, a directly attached host triggers a topology change when it is power cycled. In very large (and flat) networks, a point can be reached where the network is perpetually in a topology change status. This is as if the aging time is configured to fifteen seconds, which leads to a high level of flooding. Here is a worst case scenario that happened to a customer who was doing some server backup.



Server S is sending heavy unicast traffic to device D. Due to the nature of the protocol used, D nearly never sends any traffic. If a topology change occurs and B reduces its ageing time, entry for D will be removed from B.



Traffic is then flooded on the whole network and will reduce the bandwidth of every single link until device D sends a frame again.

The aging out of the entry for the device that receives the backup was a disaster because it caused a very heavy traffic to hit all users. See the "Avoid TCN Generation with the portfast Command" section for more information on how to avoid TCN generation.

Problem in ATM LANE Bridged Environments

This case is more critical than the normal flooding of traffic implied by a quick aging. On the receipt of a topology change for a VLAN, a Catalyst switch has its LAN emulation (LANE) blades reconfirming their LE-arp table for the corresponding emulated LAN (ELAN). As every LANE blade in the ELAN issues at the same time the same request, it may put a high stress on the LAN Emulation Server (LES) if there are a lot of entries to reconfirm. Connectivity issues have been seen in this scenario. If the network is sensitive to a topology change, the real problem is not the topology change itself but the design of the network. It is recommended that you limit as much as possible the TCN generation to save the CPU of the LES (at least). See the "Avoid TCN Generation with the portfast Command" section for more information on how to limit TCN

generation.

Avoid TCN Generation with the portfast Command

The **portfast** feature is a Cisco proprietary change in the STP implementation. The command is applied to specific ports and has two effects:

- Ports that come up are put directly in the forwarding STP mode, instead of going through the learning and listening process. The STP still runs on ports with portfast.
- The switch never generates a TCN when a port configured for **portfast** goes up or down.

Enable **portfast** on ports where the connected hosts are very likely to bring their link up and down (typically end stations that users frequently power cycle). This feature should not be necessary for server ports. It should definitely be avoided on ports that lead to hubs or other bridges. A port that directly transitions to forwarding state on a redundant link can cause temporary bridging loops.

Topology changes can be useful, so do not enable **portfast** on a port for which a link that goes up or down is a significant event for the network.

Track the Source of a TCN

In itself, a topology change notification is not a bad thing, but as a good network administrator, it is better to know their origin in order to be sure that they are not related to a real problem. Identifying the bridge that issued the topology change is not an easy task. However, it is not technically complex.

Most bridges only count the number of TCNs they have issued or received. The Catalyst 4500/4000, 5500/5000, and 6500/6000 are able to show the port and the ID of the bridge that sent the last topology change they received. Starting from the root, it is then possible to go downstream to the initiator bridge. Refer to the **show spantree statistics** command for more information.

If you have the output of a **show spantree statistics** command from your Cisco device, use [Cisco CLI Analyzer](#) ([registered](#) customers only) to display potential issues and fixes. To use [Cisco CLI Analyzer](#) ([registered](#) customers only), you must login as a registered customer, and have JavaScript enabled.

Conclusion

An important point to consider here is that a TCN does not start a STP recalculation. This fear comes from the fact that TCNs are often associated with unstable STP environments; TCNs are a consequence of this, not a cause. The TCN only has an impact on the aging time. It does not change the topology nor create a loop.

The number or the rate of topology changes is not an issue in itself. The problem is to know what the topology change means. A healthy network can experience a high rate of topology change. But, a topology change should ideally be related to a significant event in the network like a server that goes up or down or a link that transitions. This is achieved by enabling **portfast** on ports that go up and down as part of their normal operation.

Related Information

- [Understanding and Configuring Spanning Tree Protocol \(STP\) on Catalyst Switches](#)
- [Spanning Tree Protocol Problems and Related Design Considerations](#)
- [Switches Product Support](#)
- [LAN Switching Technology Support](#)
- [Technical Support & Documentation - Cisco Systems](#)