

Understanding and Tuning Spanning Tree Protocol Timers

Document ID: 19120

Introduction

Prerequisites

Requirements

Components Used

Conventions

Spanning Tree Protocol Timers

Other Parameters of the Spanning Tree Protocol

Default Values of Spanning Tree Protocol Timers

Tune max age and forward delay Timers

Decrease of the hello Time to 1 Second

Calculate the Diameter

Change the Spanning Tree Protocol Timers

Related Information

Introduction

This document describes the Spanning Tree Protocol (STP) timers and the rules to follow in order to tune the timers.

Note: This document only discusses how to tune STP timers for regular 802.1D spanning tree. This document does not discuss Rapid STP (RSTP) (IEEE 802.1w) or Multiple Spanning Tree (MST) protocol (IEEE 802.1s). For more information on RSTP and MST, refer to these documents:

- Understanding Multiple Spanning Tree Protocol (802.1s)
- Understanding Rapid Spanning Tree Protocol (802.1w)

Prerequisites

Requirements

This document assumes a good understanding of STP. For more information on the operation of STP, refer to Understanding and Configuring Spanning Tree Protocol (STP) on Catalyst Switches .



Caution: You can use this document to help you solve your networking issues, but only if you are familiar with the process or if someone who is familiar with the process has directed you. If you are unfamiliar with STP, changes that you make can cause any of these occurrences:

- Instabilities
- Application slowups
- CPU spikes
- LAN meltdown

Refer to 802.1D – IEEE Standards for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges (Clause 8) for additional details and references on all the parameters that this document discusses.

Components Used

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

Conventions

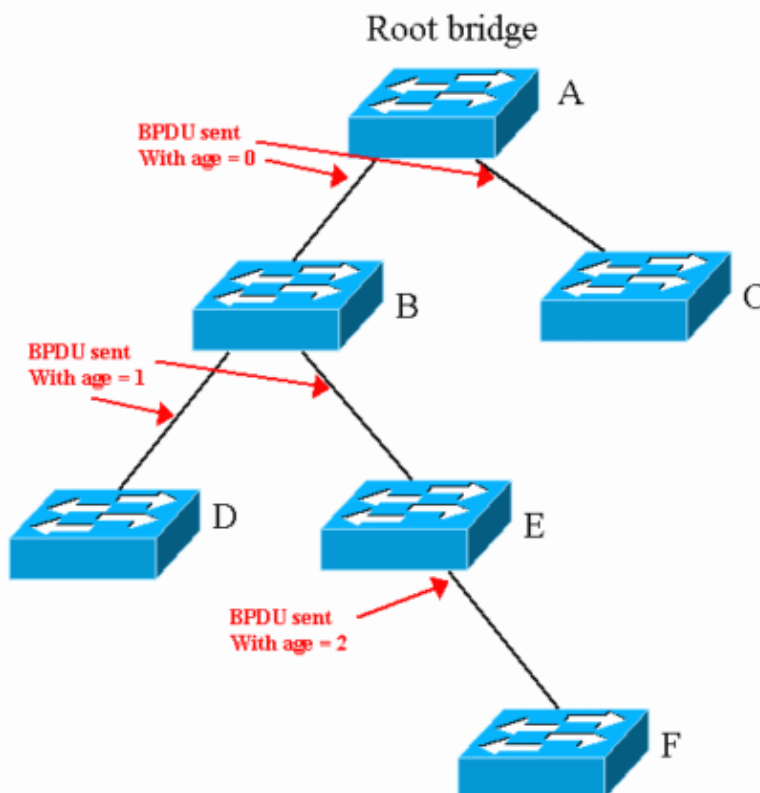
Refer to Cisco Technical Tips Conventions for more information on document conventions.

Spanning Tree Protocol Timers

There are several STP timers, as this list shows:

- **hello** The hello time is the time between each bridge protocol data unit (BPDU) that is sent on a port. This time is equal to 2 seconds (sec) by default, but you can tune the time to be between 1 and 10 sec.
- **forward delay** The forward delay is the time that is spent in the listening and learning state. This time is equal to 15 sec by default, but you can tune the time to be between 4 and 30 sec.
- **max age** The max age timer controls the maximum length of time that passes before a bridge port saves its configuration BPDU information. This time is 20 sec by default, but you can tune the time to be between 6 and 40 sec.

Each configuration BPDU contains these three parameters. In addition, each BPDU configuration contains another time-related parameter that is known as the message age. The message age is not a fixed value. The message age contains the length of time that has passed since the root bridge initially originated the BPDU. The root bridge sends all its BPDUs with a message age value of 0, and all subsequent switches add 1 to this value. Effectively, this value contains the information on how far you are from the root bridge when you receive a BPDU. This diagram illustrates the concept:



When a new configuration BPDU is received that is equal to or better than the recorded information on the port, all the BPDU information is stored. The age timer begins to run. The age timer starts at the message age

that is received in that configuration BPDU. If this age timer reaches max age before another BPDU is received that refreshes the timer, the information is aged out for that port.

Here is an example that applies to the diagram in this section:

- Switches B and C receive a configuration BPDU from switch A with a message age of 0. On the port that goes to A, the information ages out in (max age - 0) sec. This time is 20 sec by default.
- Switches D and E receive the BPDU from switch B with a message age of 1. On the port that goes to A, the information ages out in (max age - 1) sec. This time is 19 sec by default.
- Switch F receives the BPDU from switch E with a message age of 2. On the port that goes to E, the information ages out in (max age - 2) sec. This time is 18 sec by default.

Other Parameters of the Spanning Tree Protocol

IEEE 802.1D defines STP. In addition to the timers that the Spanning Tree Protocol Timers section describes, IEEE also defines these parameters that relate to STP:

- **diameter of the STP domain (dia)** This value is the maximum number of bridges between any two points of attachment of end stations. The IEEE recommendation is to consider a maximum diameter of seven bridges for the default STP timers.
- **bridge transit delay (transit delay)** This value is the time that elapsed between the reception and the transmission of the same frame by the bridge. This is logically the latency through the bridge. The IEEE recommendation is to consider 1 sec as the maximum bridge transit delay.
- **BPDU transmission delay (bpdu_delay)** This value is the delay between the time that a BPDU is received on a port and the time that the configuration BPDU is effectively transmitted to another port. The IEEE recommends 1 sec as the maximum BPDU transmission delay.
- **message age increment overestimate (msg_overestimate)** This value is the increment that each bridge adds to the message age before forwarding a BPDU. As the Spanning Tree Protocol Timers section states, Cisco switches (and probably all switches) add 1 sec to the message age before the switches forward a BPDU.
- **lost message (lost_msg)** This value is the number of BPDUs that can be lost as a BPDU moves from one end of the bridged network to the other end. The IEEE recommendation is to use three as the number of BPDUs that can be lost.
- **transmit halt delay (Tx_halt_delay)** This value is the maximum amount of time that is necessary for a bridge to effectively move a port into the blocking state after the determination that the port needs to be blocked. The IEEE recommendation is to use 1 sec for this parameter.
- **medium access delay (med_access_delay)** This value is the time that is necessary for a device to gain access to the media for initial transmission. It is the time between the CPU decision to send a frame and the moment when the frame effectively begins to leave the bridge. The IEEE recommendation is to use 0.5 sec as the maximum time.

From these parameters, you can calculate other values. This list provides the additional parameters and the calculations. The calculations assume that you use the default recommended IEEE values for all parameters.

- **End-to-end BPDU propagation delay** This value is the amount of time that is necessary for a BPDU to travel from one end of the network to the other end. Assume a diameter of seven hops, three BPDUs that can be lost, and a hello time of 2 sec. In this case, the formula is:

$$\begin{aligned} \text{End-to-end_BPDU_propa_delay} &= ((\text{lost_msg} + 1) \times \text{hello}) + ((\text{BPDU_Delay} \times (\text{dia} - 1)) \\ &= ((3 + 1) \times \text{hello}) + ((1 \times (\text{dia} - 1)) \\ &= 4 \times \text{hello} + \text{dia} - 1 \\ &= 4 \times 2 + 6 \\ &= 14 \text{ sec} \end{aligned}$$

- **Message age overestimate** The purpose of this parameter is to account for the age of the BPDU since origination. Assume that each bridge increases the BPDU message age by 1 sec. The formula is:

$$\begin{aligned}
 &\text{Message_age_overestimate} \\
 &= (\text{dia} - 1) \times \text{overestimate_per_bridge} \\
 &= \text{dia} - 1 \\
 &= 6
 \end{aligned}$$

- **Maximum frame lifetime** This value is the maximum time that a frame that previously was sent to the bridge network remains in the network before the frame reaches that destination. The formula is:

$$\begin{aligned}
 &\text{Maximum_frame_lifetime} \\
 &= \text{dia} \times \text{transit_delay} + \text{med_access_delay} \\
 &= \text{dia} + 0.5 \\
 &= 7.5 \\
 &= 8 \text{ (rounded)}
 \end{aligned}$$

- **Maximum transmission halt delay** This value is the time that is necessary in order to effectively block a port, after the decision to block is made. The IEEE counts 1 sec as the maximum for this event. The formula is:

$$\begin{aligned}
 &\text{Maximum_transmission_halt_delay} \\
 &= 1
 \end{aligned}$$

Default Values of Spanning Tree Protocol Timers

This section details how to reach the default value for max age and forward delay if you use the recommended value for each parameter. The recommended values are a diameter of seven and a hello time of 2 sec.

max age

The max age takes into account the fact that the switch that is at the periphery of the network does not time out the root information under stable conditions (that is, if the root is still alive). The max age value needs to take into account the total BPDU propagation delay and the message age overestimate. Therefore, the formula for max age is:

$$\begin{aligned}
 &\text{max_age} \\
 &= \text{End-to-end_BPDU_propa_delay} + \text{Message_age_overestimate} \\
 &= 14 + 6 \\
 &= 20 \text{ sec}
 \end{aligned}$$

This calculation shows how the IEEE reaches the default recommended value for max age.

forward delay

The movement of a port into the listening state indicates that there is a change in the active STP topology and that a port will go from blocking to forwarding. So the listening and learning periods during which the forward delay runs must cover this consecutive period:

- Time from when the first bridge port enters the listening state (and stays there through the subsequent reconfiguration) to when the last bridge in the bridged LAN hears of the change in the active topology

In addition, you need to count the same delay that you use to calculate max age (message age overestimate and BPDU propagation delay).

- Time for the last bridge to stop forwarding frames that are received on the previous topology (maximum transmission halt delay), until the last frame that is forwarded on the previous topology disappears (maximum frame lifetime)

This amount of time is necessary in order to be sure that you do not get duplicated frames.

Therefore, twice the time of the forward delay (listening time + learning time) contains all these parameters. The formula is:

```
2 x forward delay
= end-to-end_BPDU_propagation_delay + Message_age_overestimate +
  Maximum_frame_lifetime + Maximum_transmission_halt_delay
= 14 + 6 + 7.5 + 1 = 28.5

forward_delay
= 28.5 / 2
= 15 (rounded)
```

Tune max age and forward delay Timers

Among all these parameters, the only ones which you can tune are:

Note: Your ability to tune these parameters depends on the network.

- hello From 1 to 6
- max age
- forward delay
- diameter This depends on the network.

Do not modify any of the values in this list. Leave these values at the recommended IEEE value:

- lost_msg = 3
- transit_delay = 1
- bpdu_delay = 1
- msg_overestimate = 1
- Tx_halt_delay = 1
- med_access_delay = 0.5
- maximum_transmission_halt_delay = 1

These values can seem quite conservative in a modern network, in which you are not likely to lose three BPDUs or to have 1 sec of latency for a frame through a switch. However, remember that these values exist in order to prevent STP loops that can occur in stress conditions, such as:

- Very high CPU utilization
- An overloaded port

Therefore, you must consider these parameters as fixed values. If you use the formulas that the Default Values of Spanning Tree Protocol Timers section shows, you then have:

```
max_age
= End-to-end_BPDU_propa_delay + Message_age_overestimate
= ((lost_msg + 1) x hello) + ((BPDU_Delay x (dia 1)) + (dia 1) x overestimate_per_
  bridge
= (4 x hello) + dia 1 + dia 1
= (4 x hello) + (2 x dia) 2

forward_delay
= (End-to-end_BPDU_propa_delay + Message_age_overestimate +
  Maximum_frame_lifetime + Maximum_transmission_halt_delay ) / 2
= ((lost_msg + 1) x hello) + ((BPDU_Delay x (dia 1)) + ((dia 1)
```

$$\begin{aligned}
& \text{x overestimate_per_bridge) + (dia x transit_delay) + med_access_delay} \\
& + \text{Maximum_transmission_halt_delay) / 2} \\
= & ((4 \text{ x hello}) + \text{dia} \quad 1 + \text{dia} \quad 1 + \text{dia} + 0.5 + 1) / 2 \\
= & ((4 \text{ x hello}) + (3 \text{ x dia}) \quad 0.5) / 2
\end{aligned}$$

These calculations leave you with these two final formulas (if you round the 0.5 value):

$$\begin{aligned}
\text{max_age} &= (4 \text{ x hello}) + (2 \text{ x dia}) \quad 2 \\
\text{forward_delay} &= ((4 \text{ x hello}) + (3 \text{ x dia})) / 2
\end{aligned}$$

If you want to tune the STP timers in order to achieve a better convergence time, you need to strictly follow these two formulas.

Here is an example. If you have a diameter of four for a bridged network, you need to use these parameters:

```

hello = 2 (default) then
max_age = 14 sec
forward_delay = 10 sec
If hello = 1 then
max_age = 10 sec
forward_delay = 8 sec

```

Note: hello = 1 is the lowest value. There is no way that you can tune this parameter below 10 sec for max age and 8 sec for forward delay if your diameter is equal to four.

Decrease of the hello Time to 1 Second


A decrease of the hello time to 1 sec is the easiest and surest way to decrease the STP parameters. However, remember that if you drop the hello time from 2 sec to 1 sec, you double the number of BPDUs that are sent/received by each bridge. This increase causes an additional load on the CPU, which needs to process twice as many BPDUs. This load can be an issue if you have several VLANs and trunks.

Calculate the Diameter

The diameter is completely dependent on the network design. The diameter is the maximum number of switches that you cross in order to link any two switches in the bridged network (which includes source and destination), if you assume the worst cases. You do not cross the same switch twice when you determine the diameter. In the diagram in the Spanning Tree Protocol Timers section of this document, you can see that you have a diameter of 5 (path F–E–B–A–C).

Now, look at the diagram in this section. The diagram contains some access switches (switches C, D, and E) that connect to two distribution switches (switches A and B). There is a Layer 3 (L3) boundary between the distribution switches and the core. The bridged domain is stopped at the distribution switches. The STP diameter is 5:

- C–A–D–B–E
- D–A–C–B–E



You can see from the diagram that there is no pair of switches that gives a diameter that is greater than 5.

Change the Spanning Tree Protocol Timers

As the Spanning Tree Protocol Timers section mentions, each BPDU includes the hello, forward delay, and max age STP timers. An IEEE bridge is not concerned about the local configuration of the timers value. The IEEE bridge considers the value of the timers in the BPDU that the bridge receives. Effectively, only a timer that is configured on the root bridge of the STP is important. If you lose the root, the new root starts to impose its local timer value on the entire network. So, even if you do not need to configure the same timer value in the entire network, you must at least configure any timer changes on the root bridge and on the backup root bridge.

If you use a Cisco switch that runs Catalyst OS (CatOS) software, there are some macros that enable you to set up the root and tune the parameters in accordance with the formulas. Issue the **set spantree root vlan dia diameter hello hello_time** command in order to set the diameter and hello time. Here is an example:

```
Taras> (enable) set spantree root 8 dia 4 hello 2
VLAN 8 bridge priority set to 8192.
VLAN 8 bridge max aging time set to 14.
VLAN 8 bridge hello time set to 2.
VLAN 8 bridge forward delay set to 10.
Switch is now the root switch for active VLAN 8.
```

If you have the STP network diameter configured, the configured diameter value is not displayed in either the configuration or in the output of any **show** command.

Related Information

- [LAN Product Support Pages](#)
 - [LAN Switching Support Page](#)
 - [Technical Support & Documentation – Cisco Systems](#)
-

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2008 – 2009 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Jan 30, 2006

Document ID: 19120
