

BLADE 802.1s Spanning Tree and Interoperability with Cisco

Technical Brief

Industry-standard 802.1s Multiple Spanning Tree Protocol with Cisco interoperability for HP GbE2/GbE2c Ethernet Blade Switches

Introduction	1
Multiple Spanning Tree Protocol (IEEE 802.1s).....	1
MSTP region.....	1
Common Internal Spanning Tree (CIST).....	2
Spanning tree groups.....	2
Root Bridge.....	2
Bridge priority.....	2
Port cost.....	2
Port priority	3
Port state changes	3
Edge port	3
Link type.....	3
GbE2/GBE2c Interoperability.....	3
Spanning Tree Protocol (IEEE 802.1s) Configuration Example	4
Configuration Steps	5
About Alteon OS and ISCLI	6
Using BL480C Full height Servers.....	7
Using BL460C Half-Height Servers	7
References.....	8

Introduction

This technical brief addresses the interoperability of the HP GbE2 and GbE2c Ethernet Blade switches within Cisco networks. An introduction to IEEE 802.1s Multiple Spanning Tree Protocol (MSTP), is followed by example configuration commands on the HP GbE2 and GbE2c switches and the corresponding commands for the upstream Cisco switch.

Multiple Spanning Tree Protocol (IEEE 802.1s)

IEEE 802.1s Multiple Spanning Tree extends the IEEE 802.1w Rapid Spanning Tree Protocol through multiple Spanning Tree Groups. MSTP maintains up to 16 spanning-tree instances that correspond to STP Groups 1-16.

In Multiple Spanning Tree Protocol (MSTP), several VLANs can be mapped to each Spanning-Tree instance. Each Spanning-Tree instance is independent of other instances. MSTP allows frames assigned to different VLANs to follow separate paths, each path based on an independent Spanning-Tree instance. This approach provides multiple forwarding paths for data traffic, enabling load balancing, and reducing the number of Spanning-Tree instances required to support a large number of VLANs.

MSTP region

A group of interconnected bridges that share the same attributes is called an MST region. Each bridge within the region must share the following attributes:

- Alphanumeric name
- Version number
- VLAN-to-STG mapping scheme

MSTP provides rapid reconfiguration, scalability, and control due to the support of regions, and multiple Spanning-Tree instances support within each region.

Common Internal Spanning Tree (CIST)

The Common Internal Spanning Tree (CIST) provides a common form of Spanning Tree Protocol, with one Spanning Tree instance that can be used throughout the MSTP region. CIST allows the switch to interoperate with legacy equipment, including devices that run IEEE 802.1d (STP) or IEEE 802.1w (RSTP). CIST allows the MSTP region to act as a virtual bridge to other bridges outside of the region, and provides a single Spanning-Tree instance to interact with them. It is equivalent to Spanning Tree Group 0. CIST port configuration includes Hello time, Edge port status (enable/disable), and Link Type. These parameters do not affect Spanning Tree Groups 1-16. They apply only when the CIST is used.

Spanning tree groups

STP examines the network topology and defines a tree structure spanning all switches in a given layer 2 network domain. These layer 2 network domains are called spanning tree groups (STG). STGs are created by assigning a group of layer 2 switches to be part of a separate layer 2 network domain. When STP examines the network topology it only considers eliminating loops within a single STG. Within a layer 2 domain, there may be multiple STGs each operating its own individual STP algorithm to eliminate layer 2 loops.

Bridging protocol data unit

All network devices that are members of a spanning tree send out packets called bridging protocol data units (BPDU). A BPDU is a 64-byte packet sent by all switches participating in the spanning tree protocol providing information about each other. The BPDU includes information known as switch or bridge priority, port cost, and port priority used to establish a spanning tree root switch and which paths to designate as forwarding and blocking.

Root Bridge

The STP root switch (or root bridge) is the base of the spanning tree topology much like the roots of a tree. All redundant paths to the root bridge within the spanning tree network are placed in the blocked mode. The root bridge is chosen by all the switches based on the results of the BPDU exchange process.

Bridge priority

The bridge priority is used to determine what switch is the root bridge. Bridge priority is a numerical value that may be configured on a switch. The lower a bridge's priority value, the greater the chance it has of becoming the root bridge. If all switches are configured with the same default bridge priority setting, the switch with the lowest MAC address in the STP network becomes the root switch. Bridge priority is automatically assigned by the STP process, or may be manually configured on the GbE2/GbE2c Ethernet Blade Switch.

Port cost

The port cost is a value assigned to each switch port. The port cost information is exchanged within the BPDU to help determine the lowest cost path to the root switch. The port with the lowest cost path is used as the forwarding port between two segments in the STG. All remaining paths within each segment are placed in a blocked state.

The objective is to use the fastest links ensuring the route with the lowest cost is chosen. By default, the GbE2/GbE2c switch assigns fixed costs to all ports regardless of the link speed (default costs 20000 in RSTP modes). If the path cost is set to 0, the cost is set to auto. The spanning tree protocol assigns lower values to high-bandwidth ports, such as Gigabit Ethernet, to encourage their use. In RSTP mode the path

cost is automatically changed to 20000/200000/2000000. The path cost may also be set manually on the GbE2/GbE2c Ethernet Blade Switch.

Port priority

The port priority is yet another STP value assigned to each switch port. In case of identical port costs, the port priority is used as a tie breaker to determine the lowest path cost to the root switch and the resulting forwarding port for each segment. Therefore, in a network topology segment that has multiple paths with the same port cost, the port with the lowest port priority becomes the designated port for the segment. It is also possible for the ports to have identical port priorities. If this is the case, the port number becomes the final decision criteria. Port priority is automatically assigned by the STP process, or manually set on the GbE2/GbE2c Ethernet Blade Switch.

Port state changes

The port state controls the forwarding and learning processes of Spanning Tree.

Table 1: RSTP vs. STP port states

Port operational status	STP port states	RSTP port states
Enabled	Blocking	Discarding
Enabled	Listening	Discarding
Enabled	Learning	Learning
Enabled	Forwarding	Forwarding
Disabled	Disabled	Discarding

Edge port

A port that connects to a server or stub network is called an edge port. Therefore, ports 1-16 should have edge enabled. Edge ports can start forwarding as soon as the link is up.

Edge ports do not take part in Spanning Tree, and should not receive BPDUs. If a port with edge enabled does receive a BPDU, it begins STP processing until it is re-enabled.

Link type

The link type determines how the port behaves in regard to Rapid Spanning Tree. The link type corresponds to the duplex mode of the port. A full-duplex link is point-to-point (p2p), while a half-duplex link should be configured as shared. If you select auto as the link type, the port dynamically configures the link type.

GbE2/GBE2c Interoperability

Spanning tree 802.1D is enabled by default on the GbE2/GbE2c Ethernet Blade Switch to ensure that any existing network layer 2 loops are blocked.

The GbE2/GbE2c Ethernet Blade Switch meets the IEEE 802.1W standard and is compatible with Cisco switches that are 802.1W compliant. The bridge priorities, port costs, port priorities, Link type and edge port may be manually assigned on the GbE2/GbE2c Ethernet Blade Switch, allowing the core or other Catalyst switches to be the root bridge.

Note:

- When Multiple Spanning Tree is turned on, VLAN 1 is moved from Spanning Tree Group 1 to the Common Internal Spanning Tree (CIST).
- When Multiple Spanning Tree is turned off, VLAN 1 is moved back to Spanning Tree Group 1.

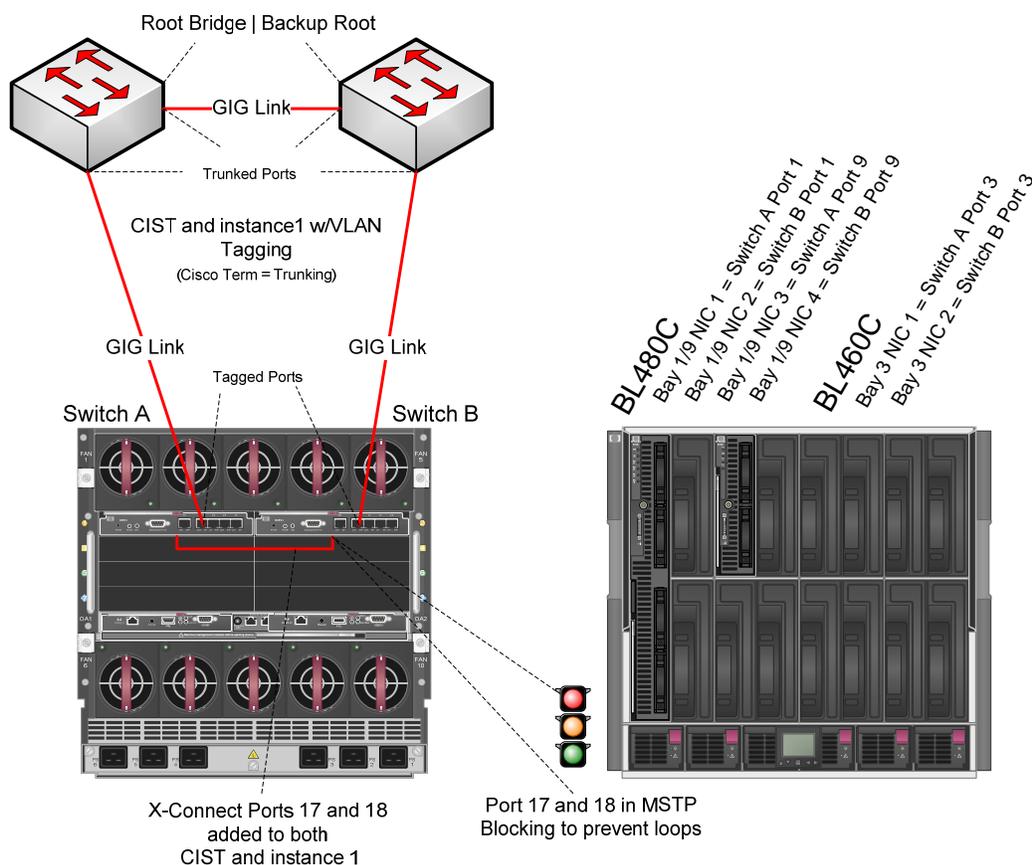
Multiple Spanning Tree Protocol (IEEE 802.1s) Configuration Example

As shown in Figure 1, the example configuration utilizes two instances STP Groups: one for GbE2/GbE2c switch ports (instance 1), and another for ports connecting the GbE2/GbE2c switches to the Cisco 6509 switches (instance 0 = CIST):

- STP 0/all interconnect ports configuration in tagging mode (trunking dot1.q).
- STP 1/VLAN 1 is used for public data traffic within the core.
- STP 2/VLAN 2 will be created for private traffic such as management network and sync traffic within the core
- STP 3/Configuration of MST (Region, version,) on Cisco and GbE2/GbE2c.
- STP 4/Configuration of Root Bridge for MST instance.
- STP 5/Configuration of VLAN 1 in the CIST (instance 0) and VLAN 2 in the instance 1.

NOTE: dot1q must be set on all Cisco Switch ports connecting to the GbE2/GbE2c Switch ports. The cross-connect ports between the two Cisco Switch ports with Spanning Tree mode must be set to MSTP.

Figure 1: Sample Topology



Configuration Steps

Table 1 shows the commands needed for configuring 802.1s Spanning Tree on the Cisco and GbE2/GbE2c switches.

Table 1 Cisco and GbE2/GbE2c Switch Configuration

How to configure an STP Instance for each VLAN	Cisco Commands (Cat OS) ----- Cisco Commands (IOS)	GbE2/GbE2c Commands (Alteon OS)	GbE2/GbE2c Commands (ISCLI)
Set the preferred VLAN trunk protocol mode: on each Cisco 6509 switch	set vtp mode <mode> ----- vtp mode <mode>	Not applicable	Not applicable
Set VLANs (VLAN 1 and 2) within each Cisco 6509 switch and each GbE2/GbE2c switch	set vlan 1-2 ----- vlan 2	/cfg/vlan 2/ena Create and enable VLAN in one command. Note: In order for VLANs to work, they must be enabled. VLAN 1 is created by default on all GbE2/GbE2c switches.	(config)# vlan 2 (config-vlan)# ena
Enable 802.1Q tagging (VLAN trunking) on required ports within each Cisco 6509 switch and each GbE2/GbE2c switch	set trunk <module number>/<port number> non negotiate dot1q 1-2 ----- Interface GigaEthernet <module number>/<port number> switchport mode trunk switchport trunk encapsulation dot1q	/cfg/port <port number>/tag ena Enable tagging (VLAN encapsulation) on the uplink port (21) and the two crosslink ports (17-18) of each GbE2/GbE2c switch /cfg/vlan 2/add <port number> Add uplink port 21 and cross-link ports 17 and 18 to VLAN 2	(config)# interface gigabitethernet <port number> (config-if)# tagging (config)# vlan 2 (config-vlan)# member <port number> Add uplink port 21 and cross-link ports 17 and 18 to VLAN 2
Set MSTP Mode on each switches (Cisco and GbE2/GbE2c)	set spantree mode mst ----- spanning-tree mode mst	/cfg/l2/mrst/mode mstp /cfg/l2/mrst/on	(config)# spanning-tree mrst mode mst spanning-tree mrst enable

Set Configuration MSTP same region and version(Cisco and GbE2/GbE2c)	<pre>set spantree mode mst set spantree mst config name <region name> set spantree mst config revision x ----- spanning-tree mst configuration name <region name> revision <x></pre>	<pre>/cfg/l2/mrst/name <1-32 character region name> /cfg/l2/mrst/version x</pre>	<pre>(config)# spanning-tree mstp name <1-32 character region name> spanning-tree mstp version x</pre>
Set the Core Cisco Switches as root bridge and each GbE2/GbE2c MSTP instance as non root	<pre>set spantree priority 4096 mst 0 set spantree priority 4096 mst 1 ----- spanning-tree mst 0 priority 4096 spanning-tree mst 1 priority 4096</pre>	<pre>/cfg/l2/mrst/cist/brg/prior 65535 /cfg/l2/stp 1/brg/prior 65535</pre>	<pre>spanning-tree mstp cist-bridge priority 65535 spanning-tree stp 1 bridge priority 65535</pre>
Configuration of VLAN in MSTP instance. By default all VLANs are in the CIST, so we only need to put VLAN 2 into instance 1	<pre>set spantree mst 1 vlan 2 set spantree mst config commit ----- spanning-tree mst configuration instance 1 vlan 2</pre>	<pre>/cfg/l2/stp 1/add 2</pre>	<pre>spanning-tree stp 1 vlan 2</pre>

About Alteon OS and ISCLI

The GbE2/GbE2c switch has two CLI entry modes—Alteon OS (AOS) and ISCLI. Alteon OS is a Unix-type menu based command system, while ISCLI has a command structure more familiar to users of Cisco switches.

1. To access the ISCLI, enter the following command from the AOS CLI:

```
Boot Options# mode ISCLI
```

The default command-line interface for the GbE2/GbE2c switch is the AOS CLI.

2. To access the AOS CLI, enter the following command:

```
>> Switch# boot cli-mode aos
```

Users can select the CLI mode upon login, if the `/boot/prompt` command is enabled. Only an administrator connected through the console port can view and enable `/boot/prompt`. When `/boot/prompt` is enabled, the first user to log in can select either CLI mode. Subsequent users must use the selected CLI mode, until all users have logged out.

Using BL480C Full height Servers

1. On GbE2/GbE2c switches A and B, add Public NIC ports to VLAN 1. VLAN 1 member ports are the GbE2/GbE2c Switch A and B ports 1, 3, 5, 7, 9, 11, 13, and 15. Use either the Alteon OS command or the ISCLI command to add the ports to VLAN 1:

Alteon OS:

```
GbE2x>> /cfg/vlan 1/add <port number>
```

ISCLI:

```
GbE2x (ISCLI)>> (config) vlan 1
```

```
GbE2x (ISCLI)>> (config-vlan) member <port number>
```

Note: In ISCLI, you cannot enter a 1 line action item like you can with Alteon OS.

2. On GbE2/GbE2c switches A and B, add Private NIC ports to VLAN 2. VLAN 2 member ports are the GbE2/GbE2c switch A and B ports 2, 4, 6, 8, 10, 12, 14, and 16. Use the following command to add the ports to VLAN 2:

Alteon OS:

```
GbE2x>> /cfg/vlan 2/add <port number>
```

ISCLI:

```
GbE2x (ISCLI)>> (config)#vlan 2
```

```
GbE2x (ISCLI)>> (config-vlan) member <port number>
```

3. On both GbE2/GbE2c switches, add ports 17, 18, and 21 to VLAN 2 using the command identified in step 2.

Using BL460C Half-Height Servers

Half-Height Servers only support 1 NIC per Switch. In order to support multiple VLANs on half-height servers from the same GbE2/GbE2c switch, you will need to setup Tagging on each Server NIC or Virtual NIC if configuring for Network Load Balancing (NLB).

1. On GbE2/GbE2c switches A and B, enable TAGGING on ports 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, and 16. Enabling tagging on these switch ports also requires all servers to have TAGGING enabled, with VLANs 1 and 2 added to both NICs.

Alteon OS:

```
GbE2x>> /cfg/port <port number>/tag ena
```

ISCLI:

```
GbE2x (ISCLI)>> (config)#interface GigabitEthernet <port number>
```

```
GbE2x (ISCLI)>> (config-if)#tagging
```

References

- [Deploying the ProLiant BL p-Class GbE2 Interconnect Switch into a Cisco-based Network](#)
- [ProLiant BL p-Class GbE2 Interconnect Switch Compatibility with Cisco-based Networks](#)

BLADE NETWORK TECHNOLOGIES: The industry's #1 provider of Ethernet, IP, and application switches for blade server systems. Our products represent a strategic control point in a blade server system. They are the gateway for all traffic directed to and from a blade server system. As a strategic control point, our products have a direct influence on a blade server system's performance, high availability, scalability, manageability, security, and total cost-of-ownership.

©2006 Blade Network Technologies, Inc. All rights reserved. Information in this document is subject to change without notice. Blade Network Technologies assumes no responsibility for any errors that may appear in this document.

<http://www.bladenetwork.net>

MKT0610-03 rev0