Technical Guide

Spanning Tree Protocols: STP, RSTP, and MSTP

FEATURE OVERVIEW AND CONFIGURATION GUIDE

Introduction

This guide describes and provides configuration procedures for:

- Spanning Tree Protocol (STP)
- Rapid Spanning Tree Protocol (RSTP)
- Multiple Spanning Tree Protocol (MSTP)

For detailed information about the commands used to configure spanning trees, see the switch's Command Reference on our website at alliedtelesis.com.

Products and software version that apply to this guide

This guide applies to AlliedWare Plus™ products that support STP, RSTP and/or MSTP, running version **5.4.4** or later.

However, support varies between products. To see whether a product supports a particular feature or command, see the following documents:

- The product's Datasheet
- The AlliedWare Plus Datasheet
- The product's Command Reference

These documents are available from the above links on our website at alliedtelesis.com.

Feature support may change in later software versions. For the latest information, see the above documents.

Content

Overview of Spanning Trees

The concept of the spanning tree protocol was devised to address broadcast storming. The spanning tree algorithm itself is defined by the IEEE standard 802.1D and its later revisions.

The IEEE Standard 802.1 uses the term **bridge** to define the spanning tree operation, and uses terms such as Bridge Protocol Data Units and Root Bridge when defining spanning tree protocol functions.

A bridge effectively means "a Layer 2 Ethernet forwarding device that forwards packets based on MAC address". So, it is a term that encompasses hardware switches as well as software-based Layer 2 forwarding devices.

For consistency, the term **bridge** rather than 'switch' will be used in this document.

When a bridge receives a frame, it reads the source and destination address fields. The bridge then enters the frame's source address in its forwarding database. In doing this the bridge associates the frame's source address with the network attached to the port on which the frame was received. The bridge also reads the destination address and if it can find this address in its forwarding database, it forwards the frame to the appropriate port. If the bridge does not recognize the destination address, it forwards the frame out from all its ports except for the one on which the frame was received, and then waits for a reply. This process is known as "flooding". Similarly, packets with broadcast or multicast destination MAC addresses will be flooded by a bridge.

A significant problem arises where bridges connect via multiple paths. A frame that arrives with an unknown or broadcast/multicast destination address is flooded over all available paths. The arrival of these frames at another network via different paths and bridges produces major problems. The bridges find the same source MAC address arriving on multiple different ports, making it impossible to maintain a reliable forwarding database. As a result, increasing numbers of packets will be forwarded to multiple paths. This process is self-perpetuating and produces a condition known as a **packet storm**, where the increase of circulating frames can eventually overload the network.

Spanning Tree operation

Where a LAN's topology results in more than one path existing between bridges, there is always a risk of the packet storm scenario described above. However, multiple paths through the extended LAN are often required in order to provide redundancy and backup in the event of a bridge or link failure.

Therefore, network designers face a problem - multiple paths are desired for resiliency purposes, but multiple paths can lead to broadcast storms. A solution to this problem is to eliminate some physical paths from the active forwarding topology, so that the active forwarding topology has only one path between any two locations. Then, if a link in the active forwarding topology becomes unavailable, one or more of the previously eliminated paths can be brought into the active forwarding topology, to restore full connectivity through the network.

The loop-free active forwarding topology is referred to as a Spanning Tree, as it is a tree topology that spans the whole network.

The spanning tree is created through the exchange of Bridge Protocol Data Units (BPDUs) between the bridges in the LAN. The spanning tree algorithm operates by:

- Automatically computing a loop-free portion of the topology, called a spanning tree. The topology is dynamically pruned to the spanning tree by declaring certain ports on a switch to be redundant, and placing them into a 'blocking' state.
- Automatically recovering from a switch failure that would partition the extended LAN by reconfiguring the spanning tree to use redundant paths, if available.

The logical tree computed by the spanning tree algorithm has the following properties:

- A single bridge is selected to become the spanning tree's unique root bridge. This is the device that advertises the lowest Bridge ID. Each bridge is uniquely identified by its Bridge ID, which comprises the bridge's root priority (a spanning tree parameter) followed by its MAC address.
- Each bridge or LAN segment in the tree, except the root bridge, has a unique parent, known as the designated bridge. The designated bridge, connects a LAN segment to the next segment on the path towards the root bridge.
- Each port connecting a bridge to a LAN segment has an associated cost, called the root path cost. This is the sum of the costs for each link in the path between the particular bridge port and the root bridge. The designated bridge for a LAN segment is the one that advertises the lowest root path cost. If two bridges on the same LAN segment have the same lowest root path cost, then the switch with the lowest bridge ID becomes the designated bridge.

The spanning tree computation is a continuous, distributed process to establish and maintain a spanning tree (Table 1). The basic algorithm is similar for STP, RSTP and MSTP modes.

Table 1: Spanning tree process

THE SPANNING TREE ALGORITHM ...	BY ...
Selects a root bridge	It selects as the root bridge for the spanning tree the device with the (numerically) lowest bridge identifier (that is, the device with lowest root bridge priority value, or if multiple bridges have the same priority, the bridge with the lowest MAC address).
Selects root ports	On each device, it selects the root port according to: ■ the port with the lowest path cost to the root bridge ■ the port connected to the bridge with the lowest root identifier ■ MSTP and RSTP only: the port with the lowest port priority value ■ the port with the lowest port number[1]
Blocks alternate ports	In order to prevent loops, it blocks alternate ports (discarding state) that provide higher cost paths to the root bridge.
Blocks backup ports	Where a second port connects one switch back to itself, it blocks the backup port that has the highest path cost or port number.

Table 1: Spanning tree process

THE SPANNING TREE ALGORITHM ...	BY ...
Selects designated ports	All other ports that are not disabled are selected as designated ports and are eventually made active (Forwarding state).
Maintains the spanning tree	If a switch or port fails, the spanning tree configures a new active topology, changing some port states, to re-establish connectivity and block loops. Depending on where the failure occurs, the changes may be widespread (e.g. if the root bridge fails), or local (e.g. if a designated port fails).

1. The whole three part port number (x.y.z) is used to find the lowest port number, where x is the device number within a stack (1 for a non stacked device), y is the module number (for example, the card or XEM number) within the device (note that 0 is used for all base-board connected ports), and z is the number of the port within the module or base-board.

The logical spanning tree, sometimes called the active topology, includes all root ports and all designated ports. These ports are in the forwarding state. Ports removed from the logical spanning tree are not in the forwarding state. To implement the spanning tree algorithm, devices communicate with one another using the Spanning Tree Protocol.

Spanning Tree modes

STP can run in one of three modes: STP, RSTP or MSTP. A device running RSTP is compatible with other devices running STP; a device running MSTP is compatible with other devices running RSTP or STP. By default, on a device in MSTP mode each port automatically detects the mode of the device connected to it (MSTP, RSTP or STP), and responds in the appropriate mode by sending messages (BPDUs) in the corresponding format. Ports on a device in RSTP mode can automatically detect and respond to connected devices in RSTP and STP mode. Particular ports can also be forced to only operate in a particular mode, by using the **spanning-tree force-version** command.

STP The Spanning Tree Protocol (STP) is the original protocol defined by IEEE standard 802.1D-1988. It creates a single spanning tree over a network. STP mode may be useful for supporting applications and protocols whose frames may arrive out of sequence or duplicated, for example NetBeui.

RSTP Rapid Spanning Tree Protocol (RSTP) also creates a single spanning tree over a network. Compared with STP, RSTP provides for more rapid convergence to an active spanning tree topology. RSTP is defined in IEEE standard 802.1D-2004. By default, the device operates in RSTP mode.

MSTP The Multiple Spanning Tree Protocol (MSTP) addresses the limitations in the previous spanning tree protocols, STP and RSTP, within networks that use multiple VLANs with topologies that employ alternative physical links. It supports multiple spanning tree instances on any given link within a network, and supports large networks by grouping bridges into regions that appear as a single bridge to other devices.

MSTP is defined in IEEE standard 802.1Q-2005. The protocol builds on, and remains compatible with, the previous IEEE standards defining STP and RSTP.

Spanning Tree Protocol (STP)

STP uses the process described in Table 1 on page 4, to avoid loops.

STP port states

In STP mode, each switch port can be in one of five spanning tree states, and one of two switch states. The state of a switch port is taken into account by STP. The STP port states (shown in Table 2) affect the behavior of ports whose switch state is enabled.

Table 2: STP port states

STATE	MEANING
DISABLED	STP operations are disabled on the port. The port does not participate in the operation of the Spanning Tree Algorithm and Protocol. The port can still switch if its switch state is enabled.
BLOCKING	The forwarding process discards received frames and does not submit forwarded frames for transmission. This is the "standby" mode.
LISTENING	The port is enabled for receiving frames only. The port is preparing to participate in frame forwarding. The forwarding process discards received frames and does not submit forwarded frames for transmission.
LEARNING	The port is enabled for receiving frames only, and the Learning Process can add new source address information to the Forwarding Database.
FORWARDING	The normal state for a switch port. The forwarding process and the Spanning Tree entity are enabled for transmit and receive operations on the port.

Interoperation with link aggregation

If multiple ports are aggregated together into a dynamic (LACP) or static channel group, then the spanning-tree process is aware of the link aggregation and treats the aggregated ports as a single logical path.

Configuring STP

By default, RSTP is enabled on all switch ports. This section provides a procedure for configuring STP (Table 3). To configure other modes, see "Configuring RSTP" on page 8 or "Configuring MSTP" on page 18.

Table 3: Configuration procedure for STP

COMMAND	DESCRIPTION
Step 1. **Configure STP** RSTP is enabled by default with default settings on all switch ports to prevent Layer 2 loops in your network.	
`awplus#` `configure terminal`	Enter Global Configuration mode.
`awplus(config)#` `spanning-tree mode stp`	By default, the device is in RSTP mode. Change to STP mode.
`awplus(config)#` `spanning-tree enable`	By default, spanning tree is enabled on all switch ports. If it has been disabled, enable it for STP.

Table 3: Configuration procedure for STP (Continued)

COMMAND	DESCRIPTION
`awplus(config)#` `spanning-tree priority <priority>`	By default, all devices have the same root bridge priority, 32768 (8000 in hexadecimal), so the device with the lowest MAC address becomes the root bridge. If you want the device to be the root bridge, set the root bridge priority to a value lower than 32768. Enter a value in the range 0 to 61440. If you enter a number that is not a multiple of 4096, the switch rounds the number down.
Step 2. **Configure Root Guard** The Root Guard feature makes sure that the port on which it is enabled is a designated port. If the Root Guard enabled port receives a superior BPDU, it goes to a Listening state (for STP) or discarding state (for RSTP and MSTP).	
`awplus(config)#` `interface <port-list>`	Enter Interface Configuration mode for the switch ports you want to enable Root Guard for.
`awplus(config-if)#` `spanning-tree guard root`	Enable the Guard Root feature for these ports.
`awplus(config-if)#` `exit`	Return to Global Configuration mode.
`awplus(config)#` `exit`	Return to Privileged Exec mode.
Step 3. **Check STP configuration**	
`awplus#` `show spanning-tree [interface <port-list>]`	Display the spanning tree configuration for the device, and confirm the new root bridge priority (Bridge Priority). Note that the Bridge ID is in a form like this: 80000000cd240331, and that other IDs follow the same pattern. This is made up of: 8000—the devices' root bridge priority in hexadecimal 0000cd240331—the devices' MAC address.

Advanced configuration

For most networks the default settings for path costs will be suitable, however, you can configure them if required (**spanning-tree path-cost** command).

Rapid Spanning Tree Protocol (RSTP)

RSTP uses the process described in Table 1 on page 4, to avoid loops.

A spanning tree running in STP mode can take up to one minute to rebuild after a topology or configuration change. The **RSTP** algorithm provides for a faster recovery of connectivity following the failure of a bridge, bridge port, or a link. RSTP provides rapid recovery by including port roles in the computation of port states, and by allowing neighboring bridges to explicitly acknowledge signals on a point-to-point link that indicate that a port wants to enter the forwarding mode.

In rapid mode, the rapid transition of a port to the forwarding state is possible when the port is considered to be part of a point-to-point link, or when the port is considered to be an **edge** port. An edge port is one that attaches to a LAN that has no other bridges attached, e.g. a port that is connected to a workstation, a printer, a VoIP phone, or other end-point device.

Table 4: RSTP port states

STATE	MEANING
DISABLED	STP operations are disabled on the port.
DISCARDING	The port does not participate in frame forwarding. The forwarding process discards received frames and does not submit forwarded frames for transmission.
LEARNING	The port is enabled for receiving frames only, and the learning process can add new source address information to the forwarding database. The port does not forward any frames.
FORWARDING	The normal state for a switch port. The forwarding process and the Spanning Tree entity are enabled for transmit and receive operations on the port.

Configuring RSTP

RSTP is enabled by default with default settings on all switch ports to prevent Layer 2 loops in your network. No further configuration is required if you want to use RSTP with these default settings. For further RSTP configuration, see Table 5 below.

To configure other modes, see "Configuring MSTP" on page 18 or "Configuring STP" on page 6.

For detailed configuration examples, see the How To Note How To Configure Basic Switching Functionality, available from alliedtelesis.com.

Table 5: Configuration procedure for RSTP

COMMAND	DESCRIPTION
Step 1. Configure RSTP RSTP is enabled by default with default settings on all switch ports to prevent Layer 2 loops in your network. No further configuration is required if you want to use RSTP with these default settings. If you need to restore the device to RSTP after it has been set to another mode, or modify the default RSTP settings, follow the procedure below.	
`awplus#` `configure terminal`	Enter Global Configuration mode.
`awplus(config)#` `spanning-tree mode rstp`	By default, the device is in RSTP mode. If it has been changed to STP or MSTP mode, change it back to RSTP.
`awplus(config)#` `spanning-tree enable`	By default, spanning tree is enabled on all switch ports. If it has been disabled, enable it for RSTP.
`awplus(config)#` `spanning-tree priority <priority>`	By default, all devices have the same root bridge priority, 32768 (8000 in hexadecimal), so the device with the lowest MAC address becomes the root bridge. If you want the device to be the root bridge, set the root bridge priority to a value lower than 32768. Enter a value in the range 0 to 61440. If you enter a number that is not a multiple of 4096, the switch rounds the number down.
Step 2. Configure edge ports If some switch ports are connected to devices that cannot generate BPDUs (such as workstations), you can set particular switch ports as edge ports, or set them to automatically detect whether they are edge ports.	
`awplus(config)#` `interface <port-list>`	Enter Interface Configuration mode for these switch ports.

`awplus(config-if)#`
`spanning-tree edgeport`
or
`awplus(config-if)#`
`spanning-tree autoedge`
	Set these ports to be edge ports,
	or
	set these ports to automatically detect whether they are edge ports.

Step 3. Configure Root Guard	
`awplus(config-if)#` `exit`	Return to Global Configuration mode.
`awplus(config)#` `interface <port-list>`	Enter Interface Configuration mode for the switch ports you want to enable Root Guard for.
`awplus(config-if)#` `spanning-tree guard root`	The Root Guard feature makes sure that the port on which it is enabled is a designated port. If the Root Guard enabled port receives a superior BPDU, it goes to a Listening state (for STP) or discarding state (for RSTP and MSTP). Enable the Guard Root feature if required.

Table 5: Configuration procedure for RSTP (Continued)

COMMAND	DESCRIPTION
Step 4. **Configure BPDU Guard**	
`awplus(config-if)#` `exit`	Return to Global Configuration mode.
`awplus(config)#` `spanning-tree portfast bpdu-guard`	If required, enable the BPDU Guard feature.
`awplus(config)#` `spanning-tree errdisable-timeout enable`	Set a timeout for ports that are disabled due to the BPDU guard feature.
`awplus(config)#` `spanning-tree errdisable-timeout interval`	Specify the time interval after which a port is brought back up when it has been disabled by the BPDU guard feature.
Step 5. **Check RSTP configuration**	
`awplus(config)#` `exit`	Return to Privileged Exec mode.
`awplus#` `show spanning-tree [interface <port-list>]`	Display the spanning tree configuration for the device, and confirm the new root bridge priority (Bridge Priority). Note that the Bridge ID is in a form like this: 80000000cd240331, and that other IDs follow the same pattern. This is made up of: 8000—the devices' root bridge priority in hexadecimal 0000cd240331—the devices' MAC address.

Advanced configuration

For most networks the default settings for path costs will be suitable, however, you can configure them if required (**spanning-tree path-cost** command).

Multiple Spanning Tree Protocol (MSTP)

Conceptually, MSTP views the total bridged network as one that comprises a number of Multiple Spanning Tree Regions (MSTRs), where each region can contain up to 64 spanning trees, which operate locally, called Multiple Spanning Tree Instances (MSTIs). AlliedWare supports up to 15 MSTIs. The regions are linked by the Common Internal Spanning Tree (CIST).

MSTP uses BPDUs to exchange information between spanning-tree compatible devices, to prevent loops in each MSTI and also in the CIST, by selecting active and blocked paths. This process is described in Table 1 on page 4.

Advantage of MSTP over RSTP

MSTP is similar to RSTP, in that it provides loop resolution and rapid convergence. However, RSTP can keep track of only one spanning-tree. MSTP can track many spanning-trees, referred to as instances. MSTP makes it possible to have different forwarding paths for different MST instances. This enables load balancing of network traffic across redundant links, so that all the links in a network can be used by at least one MSTI, and no link is left completely idle. That is to say that no link is unnecessarily shut down by spanning-tree.

Essentially, MSTP is VLAN aware and RSTP is not VLAN aware. MSTP BPDUs and RSTP BPDUs are compatible, so a network can have a mixture of MSTP and RSTP areas.

Multiple Spanning Tree Instances (MSTI)

MSTP enables the grouping and mapping of VLANs to different spanning tree instances. So, an MST Instance (MSTI) is a particular set of VLANs that are all using the same spanning tree.

In a network where all VLANs span all links of the network, judicious choice of bridge priorities for different MSTIs can result in different switches becoming root bridges for different MSTIs. That will result in the different MSTIs choosing different active topologies on the network. An example of how different MSTIs can choose different active topologies on the same physical set of links is illustrated in Figure 1 on page 12.

MSTP is compatible with RSTP and STP—see "Common and Internal Spanning Tree (CIST)" on page 14.

Figure 1: Different spanning trees created by different MSTIs on the same physical layout

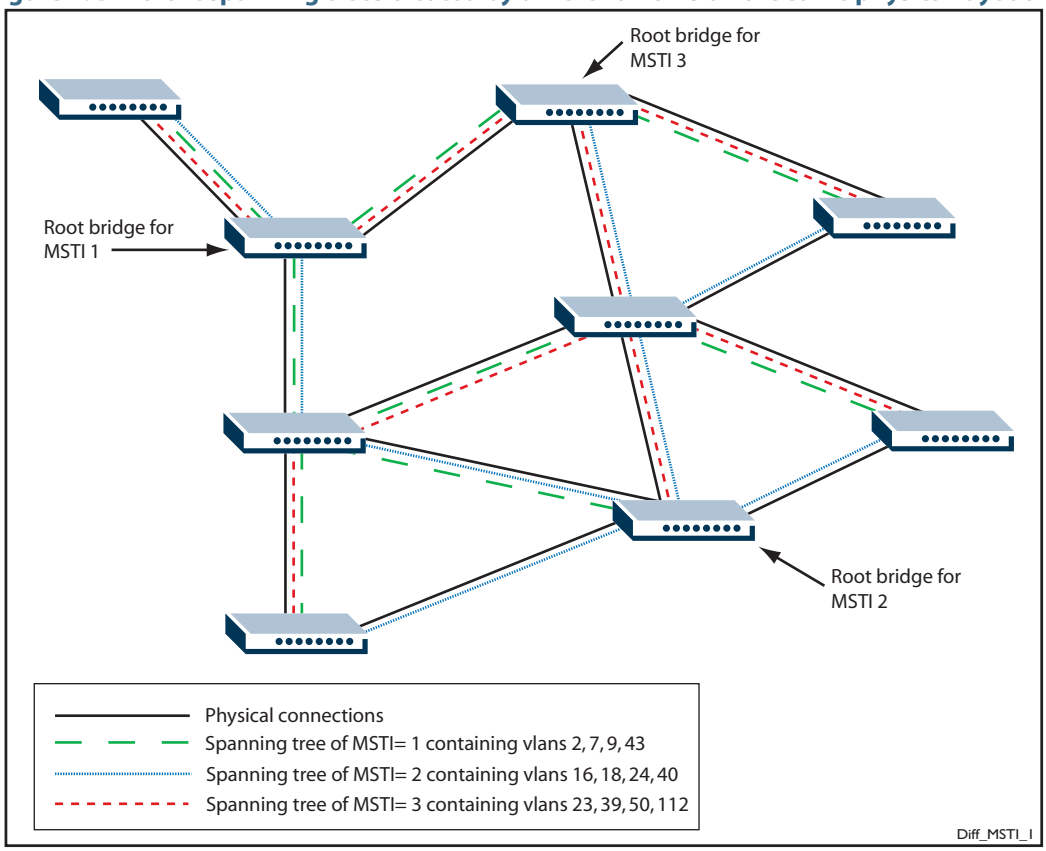

MSTP regions

An MST region is a set of interconnected switches that all have the same values for the following MST configuration identification elements:

- MST configuration name - the name of the MST region
- Revision level - the revision number of configuration
- Configuration digest - the mapping of which VLANs are mapped to which MST instances

Each of the MST instances created are identified by an MSTI number. This number is locally significant within the MST region. Therefore, an MSTI will not span across MST regions.

Figure 2: MSTIs in different regions

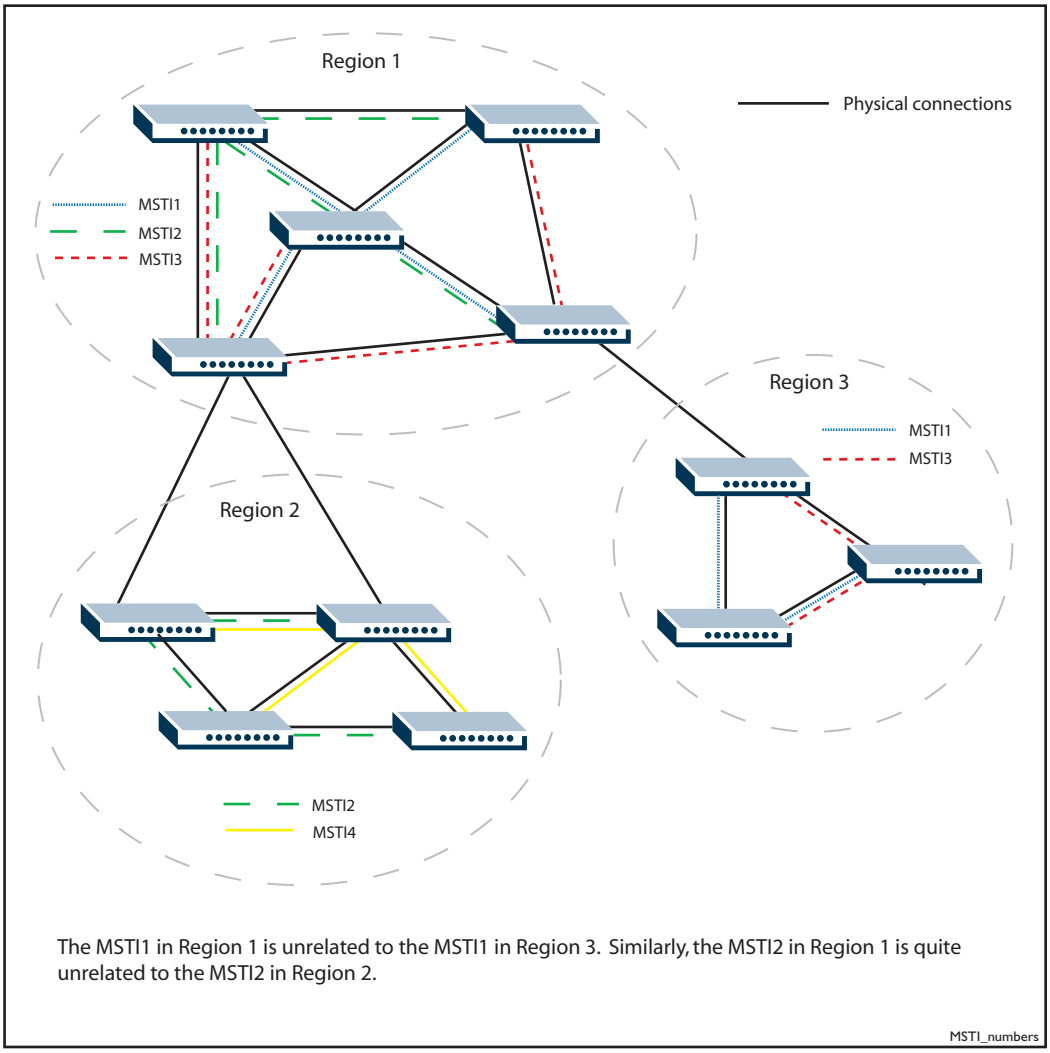

The task of assigning each bridge to a particular region is achieved by the member bridges each comparing their MST Configuration Identifiers. More information on configuration identifiers is provided in Table 6, but for the moment an MST Configuration Identifier can simply be thought of as an identifier that represents the mapping of VLANs to MSTIs within each bridge. Therefore, bridges with identical MST Configuration Identifiers, must have identical MSTI mapping tables.

While each MSTI can have multiple VLANs, each VLAN can be associated with only one MSTI. Once these associations have been made, the bridges in each region can transmit their spanning tree BPDUs and advertise their MSTIs. This in turn establishes the active data paths between the bridges for each group of VLANs (that is, for each MSTI) and block any duplicate paths within each instance. A particular advantage of this enhancement applies where a large number of VLANs share a few internetwork paths. In this situation there need only be as many Multiple Spanning Tree Instances (MSTIs) as there are source and destination bridge pairs, remembering that a pair of bridges probably has multiple paths between them.

In order to ensure that each bridge within a region maintains the same configuration information (particularly their VID to MSTI mappings) and to ensure each bridge's membership of a particular region, the bridges exchange configuration information in the

form of MST Configuration Identifiers. Table 6 below, provides a breakdown of an MST Configuration Identifier. A detailed explanation of bridge configuration identifiers can be found in Section 13.7 of the IEEE 802.1Q-2003 standard.

Table 6: MST Configuration Identifier

FIELD NAME	DESCRIPTION
Format Selector	A single octet field whose value of 0 indicates MSTP operation
Region Name	A name (up to 32 characters long) that identifies a particular MST region, defined using the **region** command.
Revision Level	A number representing the region's revision level, defined using the **revision** command.
Configuration Digest	A 16 octet (HMAC-MD5 based) signature created from the MST configuration table

Common and Internal Spanning Tree (CIST)

The CIST is the default spanning tree instance of MSTP, i.e. all VLANs that are not members of particular MSTIs are members of the CIST. Also, an individual MST region can be regarded as a single virtual bridge by other MST regions. The spanning tree that runs between regions is the CIST. The CIST is also the spanning tree that runs between MST regions and Single Spanning Tree (SST) entities. In Figure 3, the STP that is running between the regions, and to the SST bridges, is the CIST.

Figure 3: The CIST operates on links between regions and to SST devices

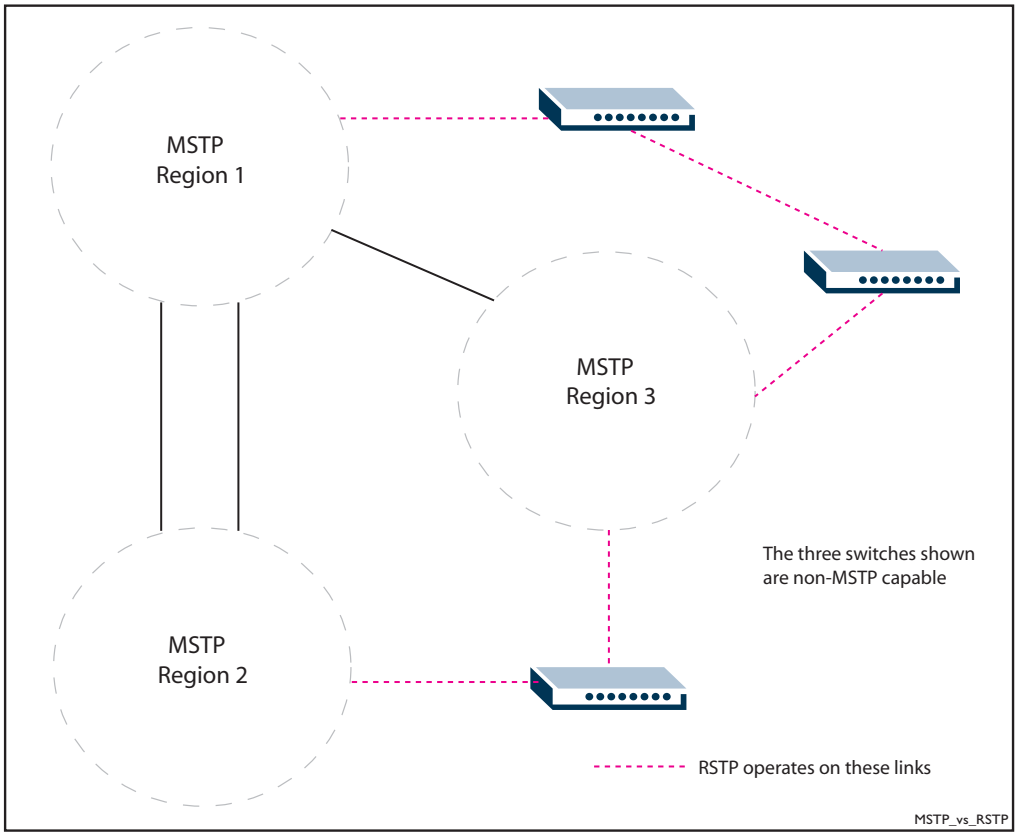

In common with legacy spanning tree systems, the CIST protocol first determines its root bridge from all the bridges on the network. This is the bridge that contains the lowest bridge identifier. The protocol then selects a regional root bridge for each MSTR. This is the bridge that provides the best path to the CIST root. After the MSTR root bridges have been chosen, they then act on the region's behalf in such a way that the region appears to the Common Spanning Tree (CST) as a virtual bridge. So in addition to having multiple MSTIs, each region operates as a bridge in a CST.

CIST

In addition to the individual MSTIs within each MSTP region, the MSTP region is a member of a network-wide spanning tree called the Common and Internal Spanning Tree (CIST). Conceptually, each region represents a virtual bridge. Internal and external bridge connectivity are two independent functions.

Frames with VIDs allocated to the CIST are subject to the rules and path costs of the complete bridged LAN as determined by the CIST's vectors. Frames other than these are subject to the CIST when traveling outside their region, and subject to its particular MSTI inside the region.

The following operational rules apply:

- Each bridge can be a member of only one region.
- A data frame is associated with a single VID.
- Data frames with a given VID are associated with either the CIST or their particular MSTI, but not both.

The role of the Common Spanning Tree (CST) in a network, and the Common and Internal Spanning Tree (CIST) configured on each device, is to prevent loops within a wider network that may span more than one MSTP region and parts of the network running in legacy STP or RSTP mode.

CIST first allocates root and designated bridges by selecting the bridge with the lowest identifier as the root. MSTP then deals with any loops between the regions in the CST. It does this by considering the CIST "vectors" in the following order:

1. CIST External Root Path Cost
2. CIST Regional Root Identifier
3. CIST Internal Root Path Cost
4. CIST Designated Bridge Identifier
5. CIST Designated Port Identifier
6. CIST Receiving Port Identifier

MSTP Bridge Protocol Data Units (BPDUs)

The main function of bridge protocol data units is to enable MSTP to select its root bridges for the CIST ("Common and Internal Spanning Tree (CIST)" on page 14) and each MSTI. MSTP is compatible with earlier spanning tree versions; its Bridge Protocol Data Unit (BPDU) formats build on earlier versions.

Table 7 shows the standardized format for MSTP BPDU messages. The general format of the BPDUs comprise a common generic portion—octets 1 to 36—that are based on those defined in IEEE Standard 802.1D, 1998, followed by components that are specific to CIST—octets 37 to 102. Components specific to each MSTI are added to this BPDU data block.

Table 7: MSTP Bridge Protocol Data Units (BPDUs)

FIELD NAME	OCTETS	DESCRIPTION
Protocol Identifier	1–2	Protocol being used. The value 0000 0000 0000 0000 identifies the spanning tree algorithm and protocol.
Protocol Version Identifier	3	Identifies the protocol version used.
BPDU Type	4	Value 0000 0000 specifies a configuration BPDU.
CIST Flags	5	Bit 1 is the topology change flag. Bit 2 conveys the CIST proposal flag in RST and MST BPDUs - unused in STP. Bits 3 & 4 convey the CIST port role in RST, and MST BPDUs - unused in STP. Bit 5 conveys the CIST learning flag in RST and MST BPDUs - unused in STP. Bit 6 conveys the CIST forwarding flag in RST and MST BPDUs - unused in STP. Bit 7 conveys the CIST agreement flag in RST and MST BPDUs - unused in STP. Bit 8 conveys the topology change acknowledge flag in STP configuration BPDUs - unused in RSTP and MSTP BPDUs.
CIST Root Identifier	6–13	The Bridge identifier of the CIST Root
CIST External Path Cost	14–17	The path cost between MST regions from the transmitting bridge to the CIST root.
CIST Regional Root Identifier	18–25	ID of the current CIST regional root bridge.
CIST Port Identifier	26–27	CIST port identifier of the transmitting bridge port.
Message Age	28–29	Message age timer value.
Max Age	30–31	Timeout value to be used by all bridges in the bridged network. This value is set by the root. Some implementations of MSTP may choose not to use this value.
Hello Time	32–33	Time interval between the generation of configuration BPDUs by the root bridge.
Forward Delay	34–35	A timeout value used to ensure forward delay timer consistency when transferring a port to the forwarding state. It is also used for ageing filtering database dynamic entries following changes in the active topology.
Version 1 Length	36	Used to convey the Version 1 length. It is always transmitted as 0.

Table 7: MSTP Bridge Protocol Data Units (BPDUs) (Continued)

FIELD NAME	OCTETS	DESCRIPTION
Version 3 Length	37–38	Used to convey the Version 3 length. It is the number of octets taken by the parameters that follow in the BPDU.
MST Configuration Identifier	39–89	An identifier comprising elements of the following: ■ Format Selector ■ Configuration Name ■ Revision Level ■ Configuration Digest.
CIST Internal Root Path Cost	90–93	Path cost to the CIST regional root.
CIST Bridge Identifier	94–101	CIST bridge identifier of the transmitting bridge.
CIST Remaining Hops	102	Remaining hops which limits the propagation and longevity of received spanning tree information for the CIST.
MSTI Configuration Messages (may be absent)	103–39 plus Version 3 Length	See Table 8.

Table 8: MSTI configuration messages

FIELD NAME	OCTETS	DESCRIPTION
MSTI Flags	1	Bits 1 through 8, convey the topology change flag, proposal flag, port role (two bits), Learning flag, forwarding flag, agreement flag, and master flag for this MSTI.
MSTI Regional Root Identifier	2–9	This includes the value of the MSTID for this configuration message encoded in bits 4 through 1 of octet 1, and bits 8 through 1 of octet 2.
MSTI Internal Root Path Cost	10-13	Internal Root Path Cost.
MSTI Bridge Priority	14	Bits 5 through 8 convey the value of the bridge identifier priority for this MSTI. Bits 1 through 4 of Octet 14 are transmitted as 0, and ignored on receipt.
MSTI Port Priority	15	Bits 5 through 8 are used to convey the value of the port identifier priority for this MSTI. Bits 1 through 4 are transmitted as 0, and ignored on receipt.
MSTI Remaining Hops	16	Value of remaining hops for this MSTI.

Configuring MSTP

By default, RSTP is enabled with default settings on all switch ports. To configure MSTP, see the configuration procedure in Table 9.

To configure other modes, see "Configuring RSTP" on page 8 or "Configuring STP" on page 6.

For detailed configuration examples, see the How To Note How To Configure Basic Switching Functionality, available from website at alliedtelesis.com.

Configuration guidelines for MSTP

- Switches must have the same MST configuration identification elements (region name, revision level and VLAN to MSTI mapping) to be in the same MST region. When configuring multiple MST regions for MSTP, MSTIs are locally significant within an MST region. MSTIs will not span from one region to another region.
- Common and Internal Spanning Tree (CIST) is the default spanning tree instance for MSTP. This means that all VLANs that are not explicitly configured into another MSTI are members of the CIST.
- The software supports a single instance of the MSTP Algorithm consisting of the CIST and up to 15 MSTIs.
- A VLAN can only be mapped to one MSTI or to the CIST. One VLAN mapped to multiple spanning trees is not allowed. All the VLANs are mapped to the CIST by default. Once a VLAN is mapped to a specified MSTI, it is removed from the CIST.
- To avoid unnecessary STP processing, a port that attaches to a LAN that is known to have no other bridges/switches attached can be configured as an edge port.

Before configuring MSTP

Before configuring MSTP, configure VLANs and associate them with switch ports (see the VLAN Feature Overview and Configuration Guide), and determine for your network:

- which MSTP regions, revision level and instances are required
- which VLANs and switch ports will belong to which MSTIs,
- which devices you want to be root bridges for each MSTI

Table 9: Configuration procedure for MSTP

COMMAND	DESCRIPTION
`awplus#` `configure terminal`	Enter Global Configuration mode.
`awplus(config)#` `spanning-tree mode mstp`	By default, the device is in RSTP mode. Change to MSTP mode.
`awplus(config)#` `spanning-tree enable`	By default, spanning tree is enabled on all switch ports. If it has been disabled, enable it for MSTP.

Table 9: Configuration procedure for MSTP (Continued)

COMMAND	DESCRIPTION
Step 1. Configure MSTP region, revision, and instances All MSTP devices in this region of the network must have the same region name, revision number, and VLAN to MSTI mappings.	
`awplus(config)#` `spanning-tree mst configuration`	Enter MST Configuration mode.
`awplus(config-mst)#` `region <region-name>`	Specify the MSTP region. The **region-name** parameter is an arbitrary string that specifies the name you want to assign to the MST region for identification.
`awplus(config-mst)#` `revision <revision-number>`	The **revision-number** parameter specifies the revision of the current MST configuration. The revision is an arbitrary number that you assign to an MST region. It can be used to keep track of the number of times that MST configuration has been updated for the network. Specify the MST revision number in the range 0 to 255.
`awplus(config-mst)#` `instance <msti-id> vlan {<vid>|<vid-list>}`	To allow MSTP to block traffic for different VLANs in different places in a loop, create multiple MSTP instances and associate VLANs with them. Each VLAN can only be in one instance. Specify the MST instance ID in the range 1 to 15.
Step 2. Advanced configuration The commands above are the minimum required to configure MSTP. The following commands allow more advanced configuration.	
Step 3. Assign root bridge priorities MSTP lets you distribute traffic more efficiently across a network by blocking different links for different VLANs. You do this by configuring different devices to be the root bridge for each MSTP instance, and for the CIST, so that each instance blocks a different link. By default, all devices have the same root bridge priority, 32768 (8000 in hexadecimal), so the device with the lowest MAC address becomes the root bridge. If you want the device to be the root bridge for an instance or for the CIST, set the priority to a lower value (a higher priority) than other devices for this instance. (If you enter a number that is not a multiple of 4096, the device rounds the number down.)	
`awplus(config)#` `spanning-tree mst configuration`	Enter MST Configuration mode.
`awplus(config-mst)#` `instance <msti-id> priority <priority>`	Set the priority for the device to become the root bridge for each instance. Specify the MST instance ID in the range 1 to 15. Specify the root bridge priority in the range 0 to 61440. If you enter a number that is not a multiple of 4096, the switch rounds the number down.
`awplus(config-mst)#` `exit`	Return to Global Configuration mode.

Table 9: Configuration procedure for MSTP (Continued)

COMMAND	DESCRIPTION
`awplus(config)#` `spanning-tree priority <priority>`	Set the priority for the device to become the root bridge for the CIST. Specify the bridge priority in the range 0 to 61440. If you enter a number that is not a multiple of 4096, the switch rounds the number down.
Step 4. Configure edge ports If some switch ports are connected to devices that cannot generate BPDUs (such as workstations), you can set particular switch ports as edge ports, or set them to automatically detect whether they are edge ports.	
`awplus(config)#` `interface <port-list>`	Enter Interface Configuration mode for these switch ports.

`awplus(config-if)#`
`spanning-tree edgeport`
or
`awplus(config-if)#`
`spanning-tree autoedge`
	Set these ports to be edge ports,
	or
	set these ports to automatically detect whether they are edge ports.

Step 5. Configure Root Guard	
`awplus(config-if)#` `spanning-tree guard root`	The Root Guard feature makes sure that the port on which it is enabled is a designated port. If the Root Guard enabled port receives a superior BPDU, it goes to a Listening state (for STP) or discarding state (for RSTP and MSTP). Enable the Guard Root feature if required.
`awplus(config-if)#` `exit`	Return to Global Configuration mode.
Step 6. Configure BPDU Guard	
`awplus(config)#` `spanning-tree portfast bpdu-guard`	If required, enable the BPDU Guard feature.
`awplus(config)#` `spanning-tree errdisable-timeout enable`	Set a timeout for ports that are disabled due to the BPDU guard feature.
`awplus(config)#` `spanning-tree errdisable-timeout interval <10-1000000>`	Specify the time interval after which a port is brought back up when it has been disabled by the BPDU guard feature.

Table 9: Configuration procedure for MSTP (Continued)

COMMAND	DESCRIPTION
Step 7. **Check MSTP configuration**	
`awplus(config)#` `exit`	Return to Privileged Exec mode.
`awplus#` `show spanning-tree mst config`	Check that the digest is the same on this device as for all other devices in the same region.
`awplus#` `show spanning-tree mst`	Check the MST to VLAN and port mapping.
`awplus#` `show spanning-tree mst instance <instance>`	Check the detailed information for a particular instance, and all switch ports associated with that instance. Specify the MST instance ID in the range 1 to 15.
`awplus#` `show spanning-tree mst interface <port>`	Check general information about MSTP, and the CIST settings.

Advanced configuration

For most networks, the default settings of the following will be suitable. However, you can also configure them.

- path costs for ports in an MSTI (**spanning-tree mst instance path-cost** command) or for the CIST (**spanning-tree path-cost** command)
- port priority for ports in an MSTI (**spanning-tree mst instance priority** command) or for the CIST (**spanning-tree priority** command)

C613-22026-00 REV A

Allied Telesis the solution : the network

North America Headquarters | 19800 North Creek Parkway | Suite 100 | Bothell | WA 98011 | USA | T: +1 800 424 4284 | F: +1 425 481 3895
Asia-Pacific Headquarters | 11 Tai Seng Link | Singapore | 534182 | T: +65 6383 3832 | F: +65 6383 3830
EMEA & CSA Operations | Incheonweg 7 | 1437 EK Rozenburg | The Netherlands | T: +31 20 7950020 | F: +31 20 7950021

alliedtelesis.com