



CHAPTER 5

Spanning Tree Scalability

This chapter provides details about spanning tree design in the data center.

Extending VLANs in the Data Center

The ability to extend VLANs across the data center is not only necessary to meet application requirements such as Layer 2 adjacency, but to permit a high level of flexibility in administering the servers. Many customers require the ability to group and maintain departmental servers together in a common VLAN or IP subnet address space. This makes management of the data center environment easier with respect to additions, moves, and changes.

For example, consider a medium-to-large data center environment where rows of servers cover a fairly large area. The HR department might have a small cluster of web and application servers physically located in rack 4 of row A in the main data center. They need to add another server that requires Layer 2 adjacency with the existing application cluster. All available space in row A is currently filled. The new server is physically located in row F, which is currently served by a different access layer switch. Because the data center design uses a Layer 2 loop-based topology, it is quite simple to add the new server port in row F by extending the VLAN to the associated access layer switch. Without this type of data center topology, the new server would have to be physically moved or located closer to the existing access layer switch, which can be disruptive and does not provide a very flexible or scalable data center design. This example demonstrates the type of flexibility that is becoming increasingly important in the data center.

Efforts to consolidate many data centers into a few is also driving the need for larger Layer 2 domains. As data centers become larger, the geographical distance between servers is increasing, creating more opportunities for the need to extend VLANs. Customers want to meet server installation requirements without the need to place them in the same physical proximity. There are many factors associated with having to physically rearrange servers in the data center, including such things as labor and administrative costs, and possible server downtime, driving up the overall total cost of ownership.

1RU access layer designs places a smaller number of servers on an access switch. This in itself creates the likelihood that the use of VLAN extension will also increase. 1RU switches also increase the number of uplinks required on the aggregation layer switch, which increases the number of STP logical and virtual ports on the aggregation layer switches.

[Table 5-1](#) provides an example of a 4000 server data center design.

Table 5-1 4000 Server Data Center Design

Modular Access w/4000 servers	1RU Access w/4000 servers
200 servers per 6509=20 access layer 6509s	48 servers per 4948=84 access layer 4948s
Uplink Ports Required (port-channel or 10GE)	Uplink Ports Required (port-channel or 10GE)
20x2 with triangle loop access	84x2 with triangle loop access
10x2 with square loop access	42x2 with square loop access

With a triangle loop access topology, the modular access requires a total of 40 uplinks (20 per agg switch) while the 1RU requires 168 for the same number of servers. This can have a large impact on the spanning tree processing at the aggregation layer. A subsequent section of this chapter covers how uplinks affect the total spanning tree logical and virtual port counts, which can impact performance and stability of a Layer 2 looped design.

When using a Layer 2 looped topology, a loop protection mechanism such as a Spanning Tree Protocol (STP) type is required. STPs automatically break loops, preventing broadcast packets from continuously circulating and melting down the network. The STPs recommended in the data center design consist of 802.1w-Rapid PVST+ and 802.1s-MST. Both 802.1w and 802.1s have the same quick convergence characteristics but differ in flexibility and operation.

STP Active Logical Ports and Virtual Ports per Line Card

In a Layer 2 looped topology design, spanning tree processing instances are created on each interface for each active VLAN. These logical instances are used by the spanning tree process in processing the spanning tree-related packets for each VLAN. These instances are referred to as active logical ports and virtual ports. Both active logical ports and virtual ports are important values to consider in spanning tree designs because they affect STP convergence time and stability. These values are usually only of concern on the aggregation layer switches because they typically have a larger number of trunks and VLANs configured than other layers in the data center topology. The rest of this section is focused on these values in the aggregation layer for this reason.

Active logical ports are a system-wide value that reflects the total number of spanning tree processing instances used in the whole system. This value can be determined by entering a **show spantree summary total** command on the console, as shown in [Figure 5-1](#).

Figure 5-1 Total Active Logical Ports Used

```

172.28.214.20 - AGG1 - SSH Secure Shell
File Edit View Window Help
Quick Connect Profiles
Aggregation-1#
Aggregation-1#
Aggregation-1#
Aggregation-1#sh spanning-tree summary total
Switch is in rapid-pvst mode
Root bridge for: VLAN0001, VLAN0003, VLAN0005-VLAN0006, VLAN0010, VLAN0020
  VLAN0041-VLAN0042, VLAN0044-VLAN0046, VLAN0100-VLAN0102, VLAN0110
  VLAN0200-VLAN0203, VLAN0300-VLAN0600
EtherChannel misconfig guard is enabled
Extended system ID is enabled
Portfast Default is disabled
PortFast BPDU Guard Default is disabled
Portfast BPDU Filter Default is disabled
Loopguard Default is enabled
UplinkFast is disabled
BackboneFast is disabled
Pathcost method used is long

Name          Blocking Listening Learning Forwarding STP Active
-----
321 vlans          0          0          0          1046          1046
Aggregation-1#

```

Connected to 172.28.214.20 | SSH2 - aes128-cbc - hmac-md5 - none | 92x22 | 141597

Virtual ports are a per-line card value that reflects the total number of spanning tree processing instances used on a particular line card. This value can be determined by entering a `show vlan virtual-port slot X` command on the console, as shown in Figure 5-2.

Figure 5-2 Total Virtual Ports used Per Line Card

```

172.28.214.20 - AGG1 - SSH Secure Shell
File Edit View Window Help
Quick Connect Profiles
Aggregation-1#
Aggregation-1#
Aggregation-1#
Aggregation-1#
Aggregation-1#
Aggregation-1#
Aggregation-1#show vlan virtual-port slot 3
Slot 3
Port          Virtual-ports
-----
G13/1          321
G13/2          321
G13/3          321
G13/4          321
Total virtual ports:1284
Aggregation-1#

```

Connected to 172.28.214.20 | SSH2 - aes128-cbc - hmac-md5 - none | 89x16 | 141598

Table 5-2 shows the current maximum scalability for both 802.1w and 802.1s spanning tree protocols on the Cisco Catalyst 6500 Series switch when using a Sup720 supervisor module.

Table 5-2 Total Virtual Ports used Per Line Card

	802.1s /MST	802.1w /R-PVST+
Total active STP logical ports	50,000 total 30,000 total with Release 12.2(17b)SXA ¹	10,000 total
Total virtual ports per line card	6000 per switching module ² [2]	1800 per switching module [2]

1. CSCed33864 is resolved in Release 12.2(17d)SXB and later releases


2. 10 Mbps, 10/100 Mbps, and 100 Mbps switching modules support a maximum of 1200 logical interfaces per module

Calculating the Active Logical Ports

When designing a large data center using extended Layer 2 VLAN topologies, it is necessary to calculate the spanning tree logical and virtual ports in advance to ensure that spanning tree operates with optimal convergence and stability characteristics.

Figure 5-3 shows an example of calculating the active logical ports in an aggregation layer switch.

Figure 5-3 Calculating STP Logical Ports



141599

Figure 5-3 shows a pair of aggregation switches (single aggregation module) connecting to 45 access layer switches in a looped access topology using four Gig-EtherChannel 802.1Q trunks. This might seem like a large number of access switches to have on a single aggregation pair, but with 1RU switch implementations, this amount or more is not uncommon. The following formula is used to determine the total active logical interfaces on the system:

trunks on the switch * active VLANs on trunks + number of non-trunking interfaces on the switch

Using Figure 5-3 for example, the calculation for aggregation 1 and 2 is as follows:

46 trunks * 120 VLANs + 2 = 5,522 active logical interfaces

This value is below the maximum values of both 802.1s MST and 802.1w Rapid-PVST+.

**Note**

An STP instance for all 120 VLANs defined in the system configuration is present on each trunk unless manual VLAN pruning is performed. For example, on each trunk configuration the **switchport trunk allowed vlan X,Y** command must be performed to reduce the number of spanning tree logical interfaces being used on that port. The VTP Pruning feature does not remove STP logical instances from the port.

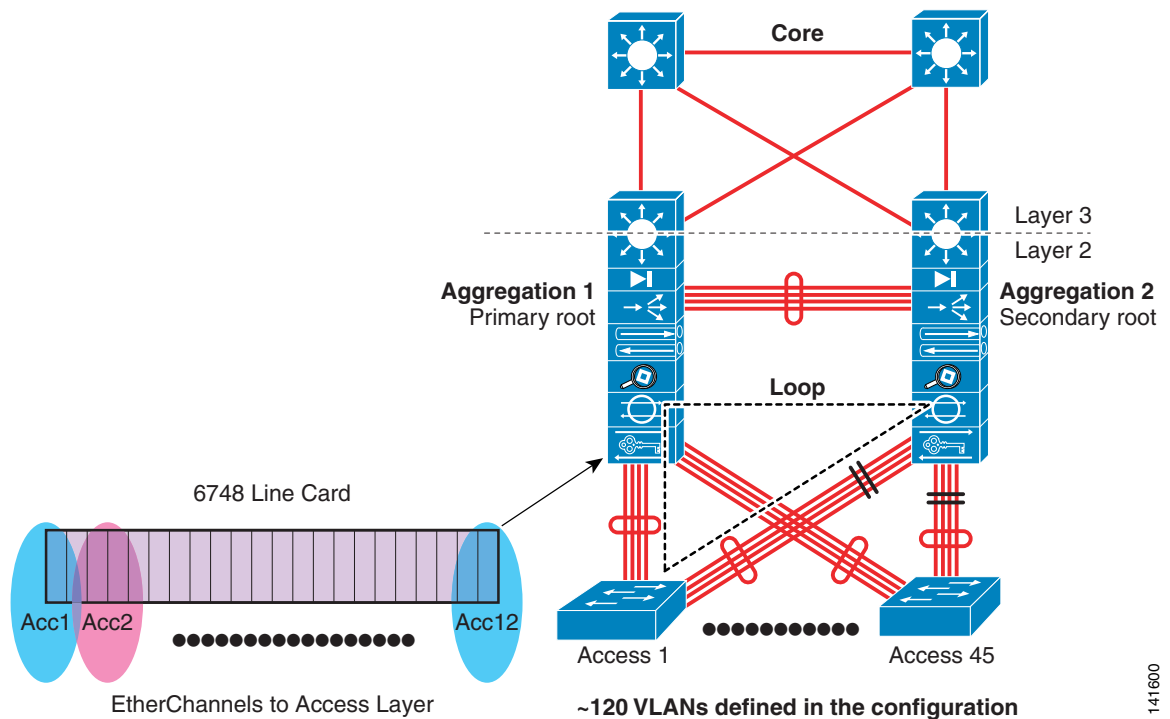
Calculating Virtual Ports per Line Card

Virtual ports are instances allocated to each trunk port on a line card. These ports are used to communicate the spanning tree-related state to the switch processor on the Sup720. A maximum number can be supported on each particular line card, as shown in Table 5-2. The following formula is used to determine the number of spanning tree virtual instances used per line card:

sum of all ports used as trunks or part of a port-channel in a trunk * active VLANs on trunks

Figure 5-4 shows a single line card on Aggregation 1 switch connecting to 12 access layer switches in a looped access topology using four Gig-EtherChannel 802.1Q trunks.

Figure 5-4 Calculating STP Virtual Ports per Line Card



A similar example could show two 6748 line cards connecting to 24 access switches using distributed EtherChannel. The formula below applies equally to both scenarios.

Using Figure 5-4 for example, the calculation for the 6748 line card is as follows:

12 trunks with 4 ports each=48 ports * 120 vlans = 5,760 virtual ports


Note

Virtual ports are allocated for each VLAN for every port participating in a trunk.

This value is well over the maximum values of 802.1w Rapid-PVST+ and very close to the maximum for 802.1s MST. This scenario experiences various issues such as long convergence times and possibly degraded system level stability.

Steps to Resolve Logical Port Count Implications

The following steps can be taken to reduce the total number of logical ports or to resolve the issues related to a large amount of logical ports being used in a system:

- Implementing multiple aggregation modules—As covered in Chapter 2, “Data Center Multi-Tier Model Design,” using multiple aggregation modules permits the spanning tree domain to be distributed, thus reducing total port count implications.
- Performing manual pruning on switchport trunk configurations—Although this can be somewhat cumbersome, it dramatically reduces the total number of both active logical and virtual port instances used.

- Using 801.1s MST—MST supports a very large number of logical port instances and is used in some of the largest data centers in the world. The drawbacks of using MST are that it does not have as much flexibility as other STP protocols, such as Rapid-PVST+, and it might not be supported in certain service module configurations.



Note Refer to the Release Notes for information on MST interoperability with service modules.

- Distributing trunks and non-trunk ports across line cards—This can reduce the number of virtual ports used on a particular line card.
- Remove unused VLANs going to Content Switching Modules (CSMs)—The CSM automatically has all available VLANs defined in the system configuration extended to it via the internal 4GEC bus connection. This is essentially the same as any other trunk configured in the system. Although there is no officially documented method of removing unnecessary VLANs, it can be performed. [Figure 5-5](#) provides an example of the steps to remove VLANs from the CSM configuration. Note that the “command rejected” output is not valid.

Figure 5-5 Removing VLANs from CSM Configuration

```

172.28.214.5-port14_AGG1-sup6.r2w - WRQ Reflection for UNIX and Digital
File Edit Connection Setup Macro Window Help
Aggregation-1#
Aggregation-1#
Aggregation-1#
Aggregation-1#
Aggregation-1#
Aggregation-1#show spannin
Aggregation-1#show spanning-tree int port-ch 259
-----
Vlan          Role Sts Cost      Prio.Nbr Type
-----
VLAN0005     Desg FWD 5000      128.1674 Edge P2p
VLAN0044     Desg FWD 5000      128.1674 Edge P2p
VLAN0102     Desg FWD 5000      128.1674 Edge P2p
Aggregation-1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Aggregation-1(config)#int range port-ch 255 - 259
Aggregation-1(config-if-range)#sw tr all vl rem 5
Command rejected: Po255 not a switching port.
Aggregation-1(config-if-range)#exit
Aggregation-1(config)#exit
Aggregation-1#
Oct 20 13:40:09.156: %SYS-5-CONFIG_I: Configured from console by mnoe onconsole
Aggregation-1#show spanning-tree int port-ch 259
-----
Vlan          Role Sts Cost      Prio.Nbr Type
-----
VLAN0044     Desg FWD 5000      128.1674 Edge P2p
VLAN0102     Desg FWD 5000      128.1674 Edge P2p
Aggregation-1#
Aggregation-1#

```



Note There is no way to view which VLANs are attached to the CSM module via the configuration on the CLI. The only way to determine which VLANs are present is with the **show spanning-tree interface port-channel 259** command.

