

Switch Meshing

Contents

Introduction	5-2
Switch Meshing Fundamentals	5-4
Terminology	5-4
Operating Rules	5-5
Using a Heterogeneous Switch Mesh	5-7
Bringing Up a Switch Mesh Domain	5-8
Further Operating Information	5-8
Configuring Switch Meshing	5-9
Preparation	5-9
Menu: To Configure Switch Meshing	5-9
CLI: To View and Configure Switch Meshing	5-12
Viewing Switch Mesh Status	5-12
CLI: Configuring Switch Meshing	5-14
Operating Notes for Switch Meshing	5-15
Flooded Traffic	5-16
Unicast Packets with Unknown Destinations	5-17
Spanning Tree Operation with Switch Meshing	5-17
Filtering/Security in Meshed Switches	5-20
IP Multicast (IGMP) in Meshed Switches	5-20
Static VLANs	5-20
Dynamic VLANs	5-21
Jumbo Packets	5-21
Mesh Design Optimization	5-22
Other Requirements and Restrictions	5-23

Introduction

Switch meshing is a load-balancing technology that enhances reliability and performance in these ways:

- Provides significantly better bandwidth utilization than either Spanning Tree Protocol (MSTP) or standard port trunking.
- Uses redundant links that remain open to carry traffic, removing any single point of failure for disabling the network, and allowing quick responses to individual link failures. This also helps to maximize investments in ports and cabling.
- Unlike trunked ports, the ports in a switch mesh can be of different types and speeds (10 and 100 Mbps, gigabit, and 10 gigabit). For example, a 10Base-FL port and a 1GB port can be included in the same switch mesh.

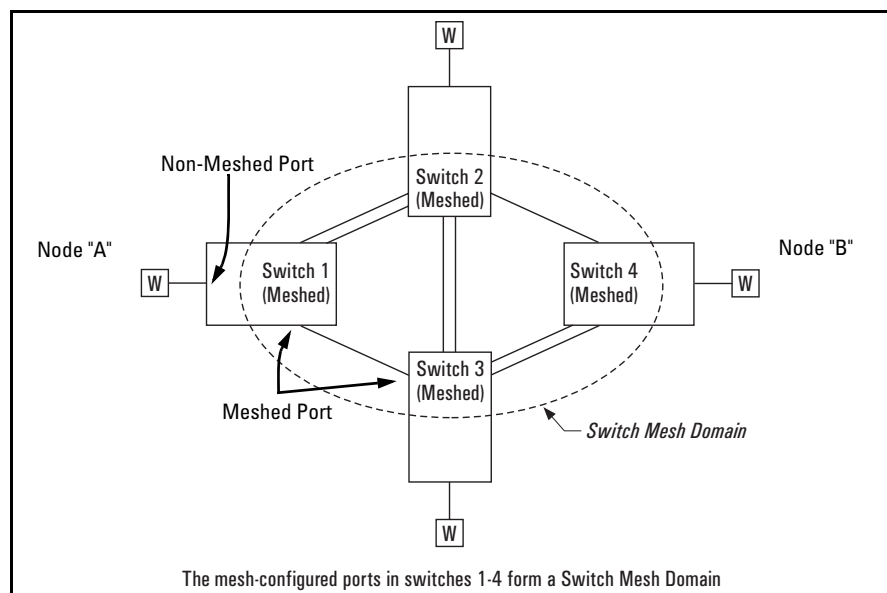


Figure 5-1. Example of Switch Meshing

Finding the Fastest Path. Using multiple switches redundantly linked together to form a *meshed switch domain*, switch meshing dynamically distributes traffic across load-balanced switch paths by seeking the fastest paths for new traffic between nodes. In actual operation, the switch mesh periodically determines the best (lowest latency) paths, then assigns these paths as the need arises. The path assignment remains until the related MAC address entry times out. The mesh sees later traffic between the same nodes as new traffic, and may assign a different path, depending on conditions at the time. For example, at one time the best path from node A to node B is through switch 2. However, if traffic between node A and node B ceases long enough for the path assignment to age out, then the next time node A has traffic for node B, the assigned path between these nodes may be through switch 3 if network conditions have changed significantly.

Note

The **mac-age-time** parameter determines how long an inactive path assignment remains in memory. Refer to “System Information” in the chapter titled “Interface Access and System Information” in the *Management and Configuration Guide* for your switch.

Because Redundant Paths Are Active, Meshing Adjusts Quickly to Link Failures. If a link in the mesh fails, the fast convergence time designed into meshing typically has an alternate route selected in less than a second for traffic that was destined for the failed link.

Meshing Allows Scalable Responses to Increasing Bandwidth Demand. As more bandwidth is needed in a LAN backbone, another switch and another set of links can be added. This means that bandwidth is not limited by the number of trunk ports allowed in a single switch.

Meshing Features

Feature	Default	Menu	CLI	Web
Viewing a mesh configuration	n/a	5-9	5-12	n/a
Configuring a Switch Mesh	n/a	5-9	5-14	n/a

Switch Meshing Fundamentals

Terminology

Switch Mesh Domain. This is a group of meshed switch ports exchanging meshing protocol packets. Paths between these ports can have multiple redundant links without creating broadcast storms.

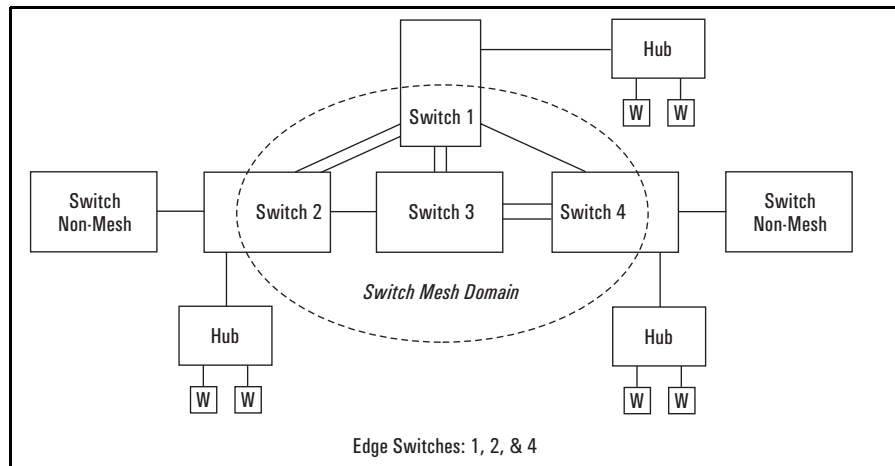


Figure 5-2. Example of a Switch Mesh Domain in a Network

Edge Switch. This is a switch that has some ports in the switch meshing domain and some ports outside of the domain. (See figure 5-2, above.)

Operating Rules

(See also “Mesh Design Optimization” on page 5-22.)

- A meshed switch can have some ports in the meshed domain and other ports outside the meshed domain. That is, ports within the meshed domain must be configured for meshing, while ports outside the meshed domain must not be configured for meshing.
- Meshed links must be point-to-point switch links.
- On any switch, all meshed ports belong to the same mesh domain.
- A switch can have up to 24 meshed ports.
- A mesh domain can include up to 12 switches.
- Up to five inter-switch, meshed hops are allowed in the path connecting two nodes through a switch mesh domain. A path of six or more meshed hops between two nodes is unusable. However, in most mesh topologies, there would normally be a shorter path available, and paths of five hops or fewer through the same mesh will continue to operate.
- Hub links between meshed switch links are not allowed.
- If the switch has multiple static VLANs and you configure a port for meshing, the port becomes a tagged member of all such VLANs. If you remove a port from meshing, it becomes an untagged member of only the default VLAN.
- A port configured as a member of a *static* trunk (LACP or Trunk) cannot also be configured for meshing.
- If a port belongs to a *dynamic* LACP trunk and you impose meshing on the port, it automatically ceases to be a member of the dynamic trunk.
- Meshing is not supported on ports configured with 802.1X access control.
- On a port configured for meshing, if you subsequently remove meshing from the port's configuration and reboot the switch, the port returns to its default configuration. (It *does not* revert to any non-default configuration it had before being configured for meshing).
- In a given mesh domain, switches in the same product family must run the same switch software version. For example, if you update the software version on one 8212zl switch, then you must update the software version on any other 8212zl switch in the mesh. ProCurve recommends that you always use the most recent software version available for the switches in your network.
- If meshing is configured on the switch, the routing features (IP routing, RIP, and OSPF) must be disabled. *That is, the switch's meshing and routing features cannot be enabled at the same time.*

- The spanning-tree configuration must be the same for all switches in the mesh (enabled or disabled). If spanning tree is enabled in the mesh, it must be the same version on all switches in the mesh: 802.1D, 802.1w, or 802.1s.
- If a switch in the mesh has GVRP enabled, then all switches in the mesh must have GVRP enabled. Otherwise, traffic on a dynamic VLAN may not pass through the mesh.
- If a switch in the mesh has a particular static vlan configured, then all switches in the mesh must have that static vlan configured.
- If a switch in the mesh has IGMP enabled, then all switches in the mesh must have IGMP enabled.
- If a switch in the mesh has LLDP enabled, then all switches in the mesh must have LLDP enabled.
- After adding or removing a port from the mesh, you must save the current configuration and reboot the switch in order for the change to take effect.
- Multiple meshed domains require separation by either a non-meshed switch or a non-meshed link. For example:




Figure 5-3. Example of Multiple Meshed Domains Separated by a Non-Mesh Switch or a Non-Mesh Link

- If GVRP is enabled, meshed ports in a switch become members of any dynamic VLANs created in the switch in the same way that they would if meshing was not configured in the switch. (For more on GVRP, refer to chapter 3, “GVRP”.)

Note

- A switch mesh domain (figure 5-1 on page 5-2) cannot include either a switch that is not configured for meshing, or a hub.
- Where a given pair of switches are linked with meshed ports, you must not also link the pair together through non-meshed ports unless you have also enabled STP, RSTP, or MSTP to prevent a loop from forming.




Figure 5-4. Example of an Unsupported Topology

- The switch blocks traffic on a meshed port connected to a non-meshed port on another switch.
- Switch meshing does not allow trunked links (LACP or Trunk) between meshed ports.

Linking a non-mesh device or port into the mesh causes the meshed switch port(s) connected to that device to shut down.

Using a Heterogeneous Switch Mesh

You can use the switches covered in this guide with the ProCurve Series 5300xl switches in normal mode.

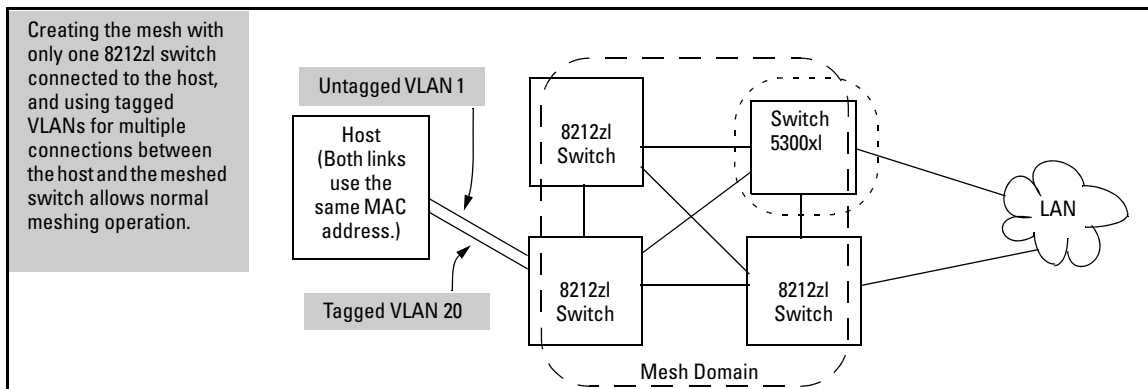


Figure 5-5. Example of a Supported Heterogeneous Topology in Normal Mode

Bringing Up a Switch Mesh Domain

When a meshed port detects a non-meshed port on the opposite end of a point-to-point connection, the link will be blocked. Thus, as you bring up switch meshing on various switches, you may temporarily experience blocked ports where meshed links should be running. These conditions should clear themselves after all switches in the mesh have been configured for meshing and their switches rebooted. To reduce the effect of blocked ports during bring-up, configure meshing and reboot the switches before installing the meshed switches in the network. Also, since adding (or removing) a meshed port requires a switch reboot to implement, you can avoid repeated system disruptions by waiting to implement the mesh until you have finished configuring meshing on all ports in your intended mesh domain.

Further Operating Information

Refer to “Operating Notes for Switch Meshing” on page 5-15.

Configuring Switch Meshing

Preparation

Before configuring switch meshing:

- Review the Operating Rules (page 5-5), and particularly the restrictions and requirements for using switch meshing in environments that include static trunks, multiple static VLANs, GVRP, IGMP, and MSTP.
- To avoid unnecessary system disruption, plan the mesh bring-up to minimize temporary port-blocking. (Refer to “Bringing Up a Switch Mesh Domain” on page 5-8.)
- To view the current switch mesh status on the switch, use the CLI **show mesh** command (page 5-12).

Menu: To Configure Switch Meshing

1. From the Main Menu, select:
 - 2. Switch Configuration**
 - 2. Port/Trunk Settings**
2. Press [E] (for **Edit**) to access the load balancing parameters.

```

----- CONSOLE - MANAGER MODE -----
Switch Configuration - Port/Trunk Settings

Port   Type      Enabled  Mode      Flow Ctrl  Group  Type
-----+-----
A1     1000SX    Yes      Auto      Disable    -----
A2     1000SX    Yes      Auto      Disable
A3     1000LX    Yes      Auto      Disable
A4     1000LX    Yes      Auto      Disable
B1     1000T     Yes      Auto      Disable
B2     1000T     Yes      Auto      Disable
B3     1000T     Yes      Auto      Disable
B4     1000T     Yes      Auto      Disable
C1     10/100TX  Yes      Auto      Disable
C2     10/100TX  Yes      Auto      Disable
C3     10/100TX  Yes      Auto      Disable
C4     10/100TX  Yes      Auto      Disable

Actions->  C Cancel    E Edit     S Save     H Help

Cancel changes and return to previous screen.
Use arrow keys to change action selection and <Enter> to execute action.
    
```

Figure 5-6. Example of the Screen for Configuring Ports for Meshing

Switch Meshing

Configuring Switch Meshing

- In the Group column, move the cursor to the port you want to assign to the switch mesh.
- Press **[M]** to choose **Mesh** for the selected port.
- Use the **up-arrow or down-arrow** key to select the next port you want to include in your mesh domain, then press **[M]** again. For example, if you were adding ports A1 and A2 to your mesh domain, the screen would appear similar to figure 5-7:

```
----- CONSOLE - MANAGER MODE -----
Switch Configuration - Port/Trunk Settings

Port   Type      Enabled  Mode    Flow Ctrl  Group  Type
-----
A1     1000SX    Yes      Auto    Disable    Mesh
A2     1000SX    Yes      Auto    Disable    Mesh
A3     1000LX    Yes      Auto    Disable
A4     1000LX    Yes      Auto    Disable
B1     1000T     Yes      Auto    Disable
B2     1000T     Yes      Auto    Disable
B3     1000T     Yes      Auto    Disable
B4     1000T     Yes      Auto    Disable
C1     10/100TX  Yes      Auto    Disable
C2     10/100TX  Yes      Auto    Disable
C3     10/100TX  Yes      Auto    Disable
C4     10/100TX  Yes      Auto    Disable

Actions->  Cancel   Edit     Save     Help

Select whether the port is part of a trunk or Mesh.
Use arrow keys to change field selection, <Space> to toggle field choices,
and <Enter> to go to Actions.
```

Ports A1 and A2 configured for meshing.

Figure 5-7. Example of Mesh Group Assignments for Several Ports

- Repeat step 5 for all ports you want in the mesh domain.

Notes

For meshed ports, leave the **Type** setting blank. (Meshed ports do not accept a **Type** setting.)

All meshed ports in the switch automatically belong to the same mesh domain. (See figure 5-2 on page 5-4.)

- When you finish assigning ports to the switch mesh, press **[Enter]**, then **[S]** (for **Save**). You will then see the following screen.

The asterisk indicates that you must reboot the switch to cause the **Mesh** configuration change to take effect.

```
----- CONSOLE - MANAGER MODE -----  
Switch Configuration Menu  
  
1. System Information  
*2. Port/Trunk Settings  
3. Network Monitoring Port  
4. Spanning Tree Operation  
5. IP Configuration  
6. SNMP Community Names  
7. IP Authorized Managers  
8. VLAN Menu...  
0. Return to Main Menu...  
  
Configures switch ports: Enabled, Mode, Flow Control, Trunking.  
To select menu item, press item number, or highlight item and pr  
(*Needs reboot to activate changes.)
```

Figure 5-8. After Saving a Mesh Configuration Change, Reboot the Switch

8. Press **[0]** to return to the Main menu.
9. To activate the mesh assignment(s) from the Main menu, reboot the switch by pressing the following keys:
 - a. **[6]** (for **Reboot Switch**)
 - b. Space bar (to select **Yes**).
 - c. **13** (to start the reboot process).

(The switch cannot dynamically reconfigure ports to enable or disable meshing, so it is always necessary to reboot the switch after adding or deleting a port in the switch mesh.)

CLI: To View and Configure Switch Meshing

Port Status and Configuration Features

Feature	Default	Menu	CLI	Web
viewing switch mesh status	n/a	n/a	below	n/a
configuring switch meshing	Disabled	n/a		n/a

Viewing Switch Mesh Status

Syntax: show mesh

*Lists the switch ports configured for meshing, along with the **State** of each mesh-configured connection, the MAC address of the switch on the opposite end of the link (**Adjacent Switch**), and the MAC address of the port on the opposite end of the link (**Peer Port**).*

Reading the Show Mesh Output. For each port configured for meshing, the State column indicates whether the port has an active link to the mesh or is experiencing a problem.

```
ProCurve# show mesh
Status and Counters - Switch Mesh Information

Port  State      | Adjacent Switch Peer Port
-----+-----
C1    Established | 0060b0-880a80  0060b0-880aff
```

Port Configured for Meshing

Operating State of the Link

MAC Address of the Switch to which Port C1 is Connected

MAC Address of the Switch Port to which Port C1 is Connected

Figure 5-9. Example of the Show Mesh Report

Table 5-1. State Descriptions for Show Mesh Output

State	Meaning
Established	The port is linked to a meshed port on another switch and meshing traffic is flowing across the link. The show mesh listing includes the MAC addresses of the adjacent switch and direct connection port on the adjacent switch.
Not Established	The port may be linked to a switch on a port that is not configured for meshing or has gone down.
Initial	The port has just come up as a meshed port and is trying to negotiate meshing.
Disabled	The port is configured for meshing but is not connected to another device.
Error	Indicates a multiple MAC-address error. This occurs when you have two or more mesh ports from the same switch linked together through a hub.
Topology Error	Two meshed switches are connected via a hub, and traffic from other, non-meshed devices, is flowing into the hub. The show mesh listing includes the MAC addresses of the adjacent switch and direct connection port on the adjacent switch.

Topology Example with Show Mesh. Suppose that you have the following topology:




Figure 5-10. Example of a Meshed Topology with Some Mesh Ports Incorrectly Linked

Table 5-2 on page 5-14 describes the meshing operation in the above topology.

Table 5-2. Operating Details for Figure 5-10

Port	Meshing?	Connection
A1	Yes	Connected to a port that may not be configured for meshing
A2	Yes	Connected to a switch port on a device that is not configured for meshing (another switch, or a hub). In this case, the Topology Error message indicates that the switch detects a meshed port on another, non-adjacent device that is also connected to the non-meshed switch or hub. However, meshing will not operate properly through this connection.
B1	Yes	Not connected to another device.
C1	Yes	Connected to a meshed port on the same adjacent switch as D1 with meshing operating properly.
D1	Yes	Connected to a meshed port on the same adjacent switch as C1 with meshing operating properly.

Figure 5-11 lists the show mesh display for the topology and meshing configuration in figure 5-10:

```
ProCurve# show mesh

Status and Counters - Switch Mesh Information

Port  State          | Adjacent Switch Peer Port
-----+-----
A1    Not Established
A2    Topology Error   0060b0-889e00  0060b0-889e7b
B1    Disabled
C1    Established      0060b0-889e00  0060b0-889e7a
D1    Established      0060b0-889e00  0060b0-889e79
```

Figure 5-11. Example of the Show Mesh Listing for the Topology in Figure 5-10

CLI: Configuring Switch Meshing

Syntax: [no] mesh [e] < port-list >

Enables or disables meshing operation on the specified ports.

All meshed ports on a switch belong to the same mesh domain. Thus, to configure multiple meshed ports on a switch, you need to:

1. Specify the ports you want to operate in the mesh domain.
2. Use **write memory** to save the configuration to the startup-config file.
3. Reboot the switch

For example, to configure meshing on ports A1-A4, B3, C1, and D1-D3:

```
ProCurve (config)# mesh e a1-a4,b3,c1,d1-d3
Command will take effect after saving configuration and reboot.
ProCurve (config)# write memory
ProCurve (config)# boot
Device will be rebooted, do you want to continue [y/n]? y
```

Figure 5-12. Example of How To Configure Ports for Meshing

To remove a port from meshing, use the "no" version of **mesh**, followed by **write memory** and rebooting the switch. For example, to remove port C1 from the mesh:

```
ProCurve # config
ProCurve (config)# no mesh c1
Command will take effect after saving configuration and reboot.
ProCurve (config)# write memory
ProCurve (config)# boot
Device will be rebooted, do you want to continue [y/n]? y
```

Figure 5-13. Example of Removing a Port from the Mesh

Operating Notes for Switch Meshing

In a switch mesh domain traffic is distributed across the available paths with an effort to keep latency the same from path to path. The path selected at any time for a connection between a source node and a destination node is based on these latency and throughput cost factors:

- Outbound queue depth, or the current outbound load factor for any given outbound port in a possible path
- Port speed, such as 10Mbps versus 100Mbps; full-duplex or half-duplex
- Inbound queue depth, or how busy a destination switch is in a possible path
- Increased packet drops, indicating an overloaded port or switch

Paths having a lower cost will have more traffic added than those having a higher cost. Alternate paths and cost information is discovered periodically and communicated to the switches in the mesh domain. This information is used to assign traffic paths between devices that are newly active on the mesh.

This means that after an assigned path between two devices has timed out, new traffic between the same two devices may take a different path than previously used.

To display information on the operating states of meshed ports and the identities of adjacent meshed ports and switches, see “Viewing Switch Mesh Status” on page 5-12.

Flooded Traffic

Broadcast and multicast packets will always use the same path between the source and destination edge switches unless link failures create the need to select new paths. (Broadcast and multicast traffic entering the mesh from different edge switches are likely to take different paths.) When an edge switch receives a broadcast from a non-mesh port, it floods the broadcast out all its other non-mesh ports, but sends the broadcast out only those ports in the mesh that represent the path from that edge switch through the mesh domain. (Only one copy of the broadcast packet gets to each edge switch for broadcast out of its nonmeshed ports. This helps to keep the latency for these packets to each switch as low as possible.)




Figure 5-14. Example of a Broadcast Path Through a Switch Mesh Domain

Any mesh switches that are not edge switches will flood the broadcast packets only through ports (paths) that link to separate edge switches in the controlled broadcast tree. The edge switches that receive the broadcast will flood the broadcast out all non-meshed ports. Some variations on broadcast/multicast

traffic patterns, including the situation where multiple VLANs are configured and a broadcast path through the mesh domain leads only to ports that are in the same VLAN as the device originating the broadcast.

Unicast Packets with Unknown Destinations

A meshed switch receiving a unicast packet with an unknown destination does not flood the packet onto the mesh. Instead, the switch sends a query on the mesh to learn the location of the unicast destination. The meshed switches then send 802.2 test packets through their non-meshed ports. After the unicast destination is found and learned by the mesh, subsequent packets having the same destination address will be forwarded. By increasing the **MAC Age Time** you can cause the switch address table to retain device addresses longer. (For more on **MAC Age Time**, refer to “System Information” in the chapter titled “Interface Access and System Information” in the *Management and Configuration Guide* for your switch.) Because the switches in a mesh exchange address information, this will help to decrease the number of unicast packets with unknown destinations, which improves latency within the switch mesh. Also, in an IP environment, ProCurve recommends that you configure IP addresses on meshed switches. This makes the discovery mechanism more robust, which contributes to decreased latency.

Spanning Tree Operation with Switch Meshing

Using MSTP with several switches and no switch meshing configured can result in unnecessarily blocking links and reducing available bandwidth. For example:

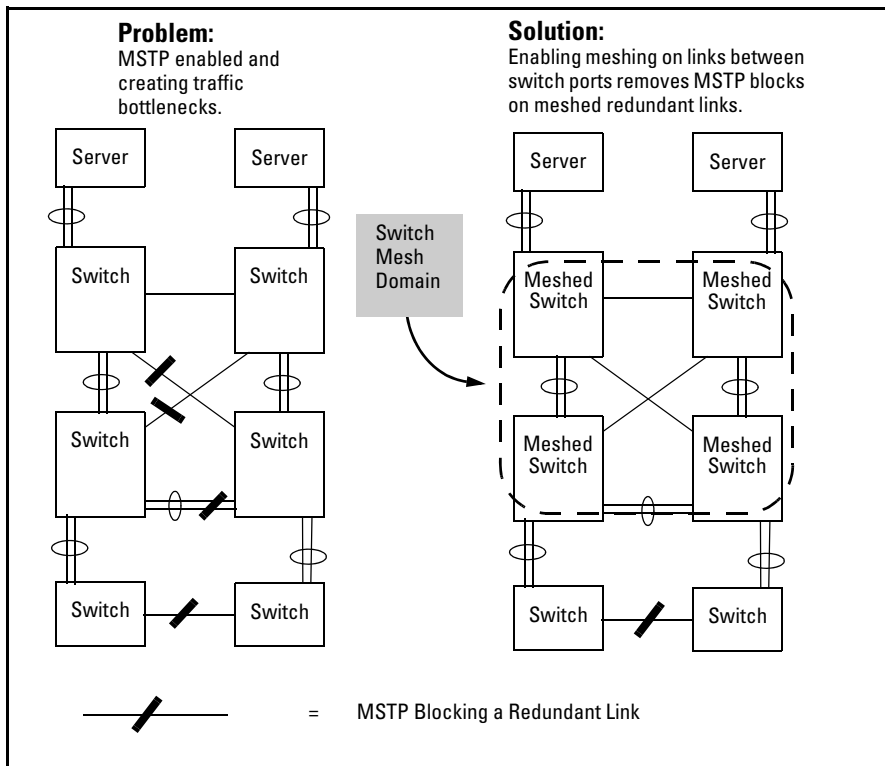


Figure 5-15. Example Using STP Without and With Switch Meshing

If you are going to use spanning-tree in a switch mesh, all switches in the mesh should be configured with the same type of spanning-tree: 802.1d/STP, 802.1w/RSTP, or 802.1s/MSTP. Spanning-Tree interprets a meshed domain as a single link. However, on edge switches in the domain, MSTP will manage non-meshed redundant links from other devices. For example:

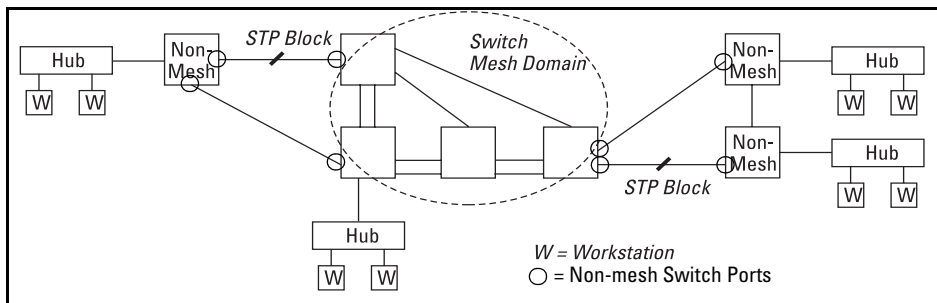


Figure 5-16. Connecting a Switch Mesh Domain to Non-Meshed Devices

Note on the Edge-Port Mode in MSTP

When using MSTP and interconnecting switches covered in this guide in a mesh with switches that are not in the mesh, all the non-mesh switch ports (as indicated in the figure above) should have the **edge-port** parameter disabled.

MSTP should be configured on non-mesh devices that use redundant links to interconnect with other devices or with multiple switch mesh domains. For example:




Figure 5-17. Interconnecting Switch Mesh Domains with Redundant Links

In the above case of multiple switch meshes linked with redundant trunks there is the possibility that spanning-tree will temporarily block a mesh link. This is because it is possible for spanning-tree to interpret the cost on an external trunked link to be less than the cost on a meshed link. However, if this condition occurs, the meshed switch that has a blocked link will automatically increase the cost on the external (non-meshed) link to the point where spanning tree will block the external link and unblock the meshed link. This process typically resolves itself in approximately 30 seconds.

Caution

Spanning tree interprets a switch mesh as a single link. Because the switch automatically gives faster links a higher priority, the default spanning-tree parameter settings are usually adequate for spanning tree operation. Also, because incorrect spanning tree settings can adversely affect network performance, you should not make changes unless you have a strong understanding of how spanning tree operates.

In a mesh environment, the default MSTP timer settings (**Hello Time** and **Forward Delay**) are usually adequate for MSTP operation. Because a packet crossing a mesh may traverse several links within the mesh, using smaller-than-default settings for the MSTP **Hello Time** and **Forward Delay** timers can cause unnecessary topology changes and end-node connectivity problems.

For more on spanning-tree, refer to the chapter titled “Multiple Instance Spanning-Tree Operation” in this guide. Also, you may want to examine the IEEE 802.1d, 802.1w, or 802.1s standards, depending on which version of spanning-tree you are using. The switches covered in this guide use 802.1s.

Filtering/Security in Meshed Switches

Because paths through the mesh can vary with network conditions, configuring filters on meshed ports can create traffic problems that are difficult to predict, and is not recommended. However, configuring filters on nonmeshed ports in an edge switch provides you with control and predictability.

IP Multicast (IGMP) in Meshed Switches

Like trunked ports, the switch mesh domain appears as a single port to IGMP. However, unlike trunked ports, IGMP protocol and multicast traffic may be sent out over several links in the mesh in the same manner as broadcast packets.

Static VLANs

In a network having a switch mesh domain and multiple static VLANs configured, all static VLANs must be configured on each meshed switch, even if no ports on the switch are assigned to any VLAN. (The switch mesh is a member of all static VLANs configured on the switches in the mesh.)

When static VLANs are configured, the mesh is seen as a single entity by each VLAN. All ports in the mesh domain are members of all VLANs and can be used to forward traffic for any VLAN. However, the non-mesh ports on edge switches that allow traffic to move between the mesh and non-meshed devices belong to specific VLANs and do not allow packets originating in a specific VLAN to enter non-meshed devices that do not belong to that same VLAN. (It is necessary to use a router to communicate between VLANs.) For example, in the following illustration, traffic from host A entering the switch mesh can only exit the mesh at the port for hosts B and E. Traffic from host A for any other host (such as C or D) will be dropped because only hosts B and E are in the same VLAN as host A.

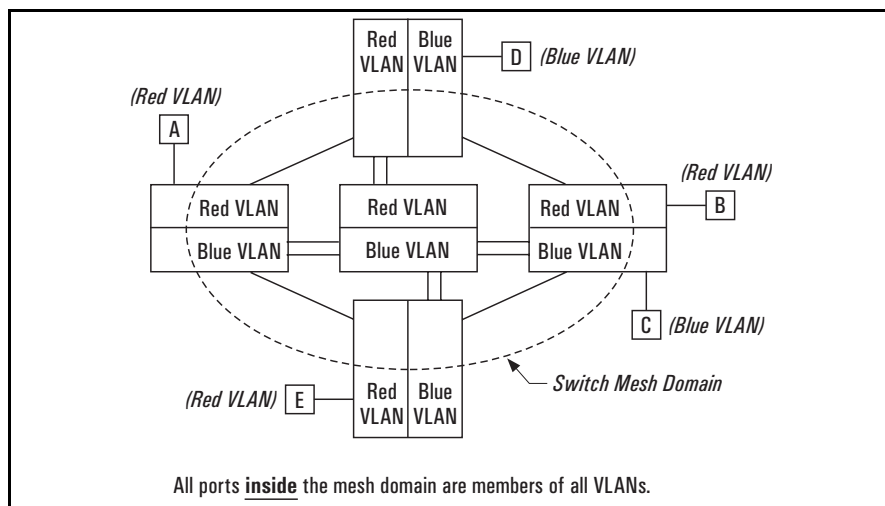


Figure 5-18. VLAN Operation with a Switch Mesh Domain

Dynamic VLANs

If GVRP is enabled, meshed ports in a switch become members of any dynamic VLANs created in the switch in the same way that they would if meshing was not configured in the switch. (For more on GVRP, refer to chapter 3, “GVRP”.)

Jumbo Packets

If you enable jumbo traffic on any VLAN, then all meshed ports on the switch will be enabled to support jumbo traffic. (On a given meshed switch, every meshed port becomes a member of every VLAN configured on the switch.) If a port in a meshed domain does not belong to any VLANs configured to support jumbo traffic, then the port drops any jumbo packets it receives from other devices. In this regard, if a mesh domain includes any ProCurve 8212zl switches, 6200yl switches, Series 5400zl switches, Series 3500yl switches, Series 3400cl or Series 6400cl switches that are configured to support jumbo traffic, only these switches can transmit and receive jumbo packets. Other switch models in the mesh will drop jumbo packets as they are not supported by those switches. For more information on jumbo packets, refer to the chapter titled “Port Traffic Controls” in the *Management and Configuration Guide* for your switch.

Mesh Design Optimization

Mesh performance can be enhanced by using mesh designs that are as small and compact as possible while still meeting the network design requirements. The following are limits on the design of meshes and have not changed:

1. Any switch in the mesh can have up to 24 meshed ports.
2. A mesh domain can contain up to 12 switches.
3. Up to 5 inter-switch meshed hops are allowed in the path connecting two nodes.
4. A fully interconnected mesh domain can contain up to 5 switches.

Mesh performance can be optimized by keeping the number of switches and the number of possible paths between any two nodes as small as possible. As mesh complexity grows, the overhead associated with dynamically calculating and updating the cost of all of the possible paths between nodes grows exponentially. Cost discovery packets are sent out by each switch in the mesh every 30 seconds and are flooded to all mesh ports. Return packets include a cost metric based on inbound and outbound queue depth, port speed, number of dropped packets, etc. Also, as mesh complexity grows, the number of hops over which a downed link has to be reported may increase, thereby increasing the reconvergence time.

The simplest design is the two-tier design because the number of possible paths between any two nodes is kept low and any bad link would have to be communicated only to its neighbor switch.

Other factors affecting the performance of mesh networks include the number of destination addresses that have to be maintained, and the overall traffic levels and patterns. However a conservative approach when designing new mesh implementations is to use the two-tier design and limit the mesh domain to eight switches where possible.

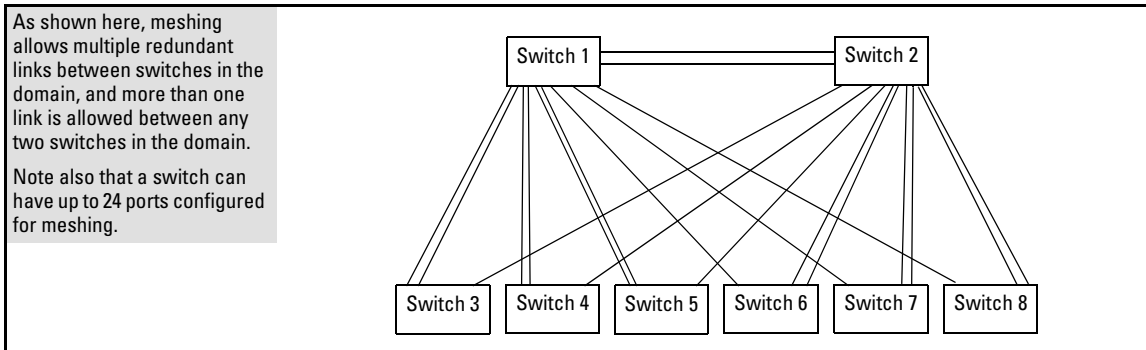


Figure 5-19. Example of a Two-Tier Mesh Design




Figure 5-20. Example of a Fully Interconnected Mesh with the Maximum Switch Count

Other factors affecting the performance of mesh networks include the number of destination addresses that have to be maintained, and the overall traffic levels and patterns. However a conservative approach when designing new mesh implementations is to use the two-tier design and limit the mesh domain to eight switches where possible.

Other Requirements and Restrictions

- **Mesh Support Within the Domain:** All switches in the mesh domain, including edge switches, must support the ProCurve switch meshing protocol.
- **Switch Hop Count in the Mesh Domain:** A maximum of five (meshed) switch hops is allowed in the path connecting two nodes in a switch mesh domain. A path of six meshed hops is unusable. However, this does not interfere with other, shorter paths in the same domain.

- **Connecting Mesh Domains:** To connect two separate switch meshing domains, you must use non-meshed ports. (The non-meshed link can be a port trunk or a single link.) Refer to figure 5-3 on page 5-6.
- **Multiple Links Between Meshed Switches:** Multiple mesh ports can be connected between the same two switches, to provide higher bandwidth. Each port that you want in the mesh domain should be configured as **Mesh** (and not as a trunk—**Trk**). Note that if you configure a port as **Mesh**, there is no “Type” selection for that port.
- **Network Monitor Port:** If a network monitor port is configured, broadcast packets may be duplicated on this port if more than one port is being monitored and switch meshing is enabled.
- **Compatibility with Other Switches:** The switches covered in this guide operate with the Series 5300xl switches in normal mode.
- **Rate-Limiting Not Recommended on Meshed Ports: Rate-Limiting can reduce the efficiency of paths through a mesh domain.**

(See also “Operating Rules” on page 5-5.)

For additional information on troubleshooting meshing problems, refer to “Using a Heterogeneous Switch Mesh” on page 5-7 and “Mesh-Related Problems” in appendix C, “Troubleshooting” of the Management and Configuration Guide for your switch.