Chapter 7

Configuring Trunk Groups and Dynamic Link Aggregation

This chapter describes how to configure trunk groups and 802.3ad link aggregation.

- Trunk groups are manually-configured aggregate links containing multiple ports.
- 802.3ad link aggregation is a protocol that dynamically creates and manages trunk groups.

NOTE: You can use both types of trunking on the same device. However, you can use only one type of trunking for a given port. For example, you can configure port 1/1 as a member of a static trunk group or you can enable 802.3ad link aggregation on the port, but you cannot do both.

Configuring Trunk Groups

The Trunk Group feature allows you to manually configure multiple high-speed load-sharing links between two ProCurve Routing Switches or between a ProCurve Routing Switch and a server. You can configure up to 8 ports as a trunk group, supporting transfer rates of up to 8 Gbps of bi-directional traffic.

In addition to enabling load sharing of traffic, trunk groups provide redundant, alternate paths for traffic if any of the segments fail.

Figure 7.1 shows an example of a configuration that uses trunk groups.

Figure 7.1 Trunk Group application within a ProCurve Routing Switch network

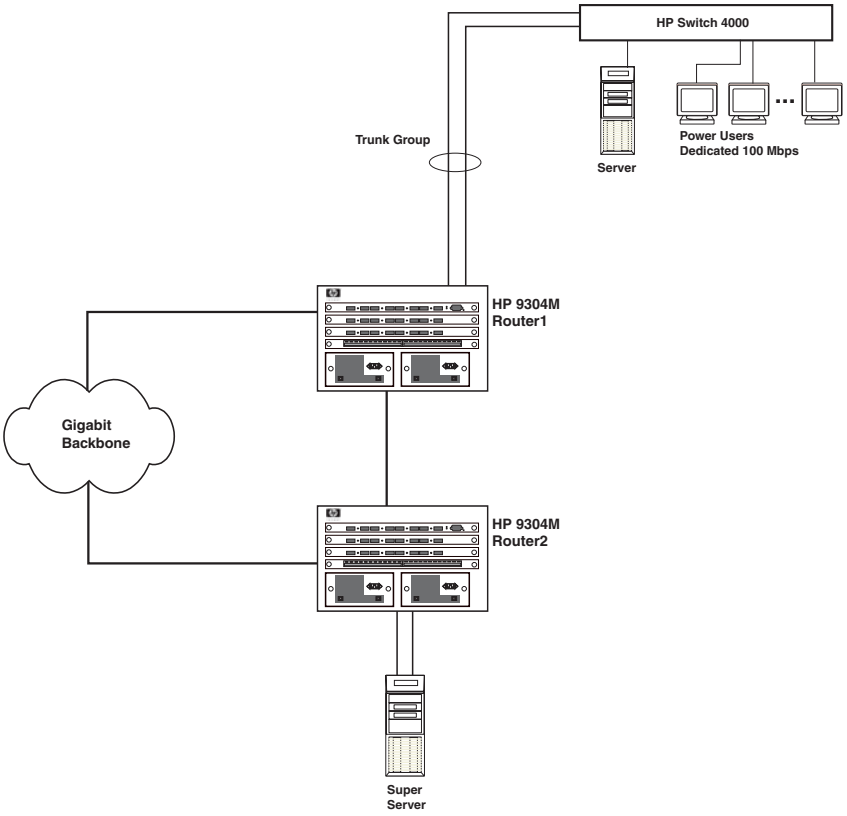

NOTE: The ports in a trunk group make a single logical link. Therefore, all the ports in a trunk group must be connected to the same device at the other end.

Trunk Group Connectivity to a Server

To support termination of a trunk group, the server must have either multiple network interface cards (NICs) or either a dual or quad interface card installed. The trunk server is designated as a server with multiple adapters or a single adapter with multiple ports that share the same MAC and IP address. Figure 7.2 shows an example of a trunk group between a server and an HP device.

Figure 7.2 Trunk group between a server and an HP device

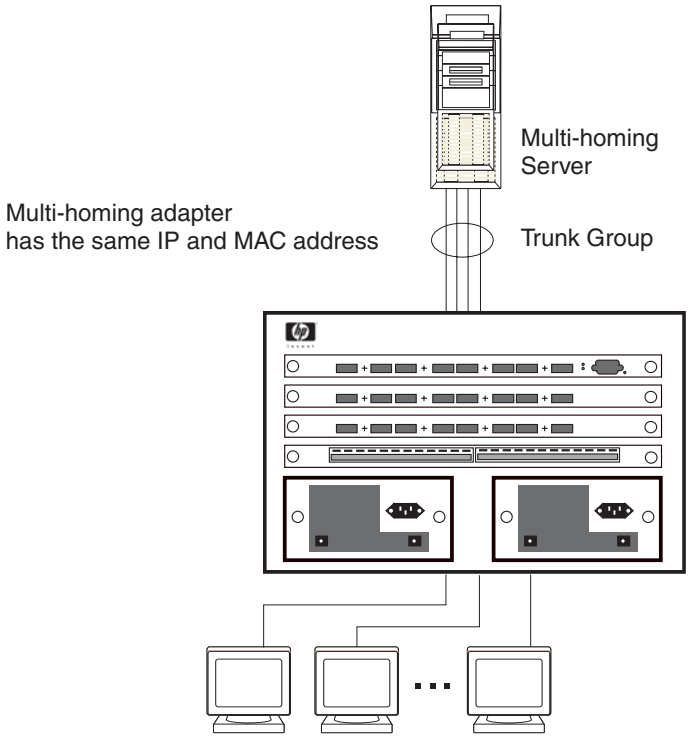

Trunk Group Rules

- You cannot configure a port as a member of a trunk group if 802.3ad link aggregation is enabled on the port.
- The following table lists the maximum number of trunk groups you can configure on a Routing Switch, the valid number of ports in a trunk group, and the port ranges on a device.

Table 7.1: Trunk Group Support

Maximum Number of Trunk Groups		Valid Number of Ports in a Group		Port Ranges and Primary Ports	
10/100	**Gigabit**	**10/100**	**Gigabit**	**10/100**	**Gigabit**
64	8-port slot: 4 16-port slot: 8	2, 4, or 8	2 or 4	1 – 8, 9 – 16, 17 – 24	2 ports: multiples of 2, for example: 1-2, 3-4, 5-6, etc. 4 ports: multiples of 4, for example: 1-4, 5-8, etc.

- You cannot combine 10/100 ports and Gigabit ports in the same trunk group. Each trunk group must start with a primary port. The primary port is always the lowest number in the port range. For example, on the J4140A 10/100 module:
  - Ranges for four-port trunk groups: 1 – 4, 5 – 8, 9 – 15, 16 – 20, 21 – 24
  - Ranges for two-port trunk groups: 1 – 2, 3 – 4, 5 – 6, 7 – 8, 9 – 10, 11 – 12, 13 – 14, 15 – 16, 17 – 18, 19 – 20, 21 – 22, 23 – 24

NOTE: You can configure up to 12 trunk groups on an 9300 series 24-port 10/100 module. The 24-port 10/100 modules have the following primary ports: 1, 3, 5, 7, 9, 11, 13, 15, 17, 19, 21, and 23. See Figure 7.5.

- Port assignment on a module must be contiguous. The port range on the module cannot contain gaps. For example, you can configure ports 1, 2, 3, and 4 on a module together as a trunk group but not ports 1, 3, and 4 (excluding 2).
- Port assignment cannot be across multiple trunk group boundaries.All the ports must be connected to the same device at the other end.
- All trunk group member properties must match the lead port of the trunk group with respect to the following parameters:
  - Port tag type (untagged or tagged port)
  - Port speed and duplex
  - QoS priority

  To change port parameters, you must change them on the primary port. The software automatically applies the changes to the other ports in the trunk group.
- Make sure the device on the other end of the trunk link can support the same number of ports in the link.
- You can trunk two 10 Gigabit Ethernet ports together. The first port must be in an odd-numbered chassis slot and the second port must be in the following even-numbered slot. Trunking of 10-Gigabit Ethernet ports requires software release 07.6.01b or later. See “Configuring a Trunk Group of 10-Gigabit Ethernet Ports” on page 7-14.

Figure 7.3 shows an example of a valid 2-port trunk group link between devices. The trunk groups in this example are switch trunk groups, between two HP devices. Ports in a valid 2-port trunk group on one device are connected to two ports in a valid 2-port trunk group on another device. The same rules apply to 4-port trunk groups.

Figure 7.3 Example of 2-port trunk groups

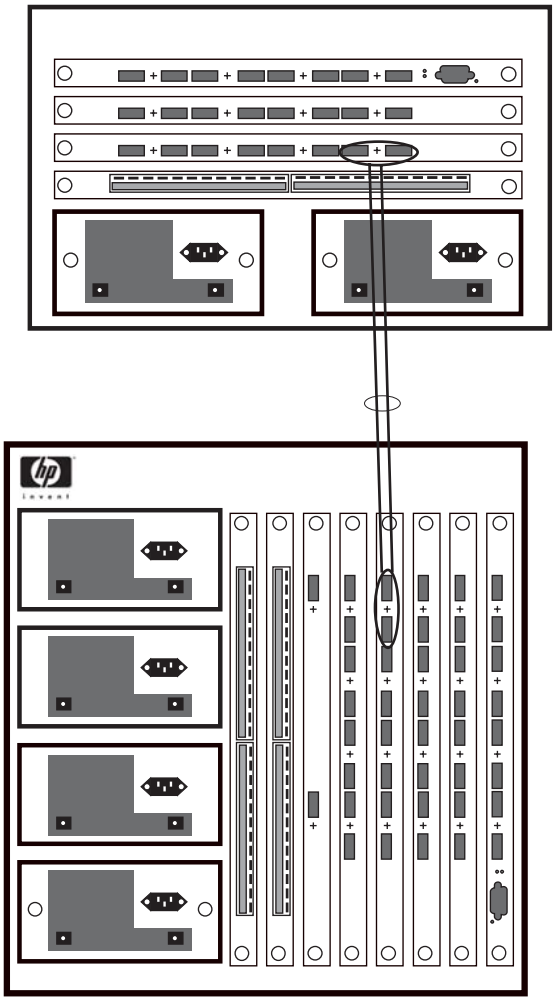

Figure 7.4 shows examples of two Routing Switches connected by multi-slot trunk groups.

Figure 7.4 Examples of multi-slot trunk groups

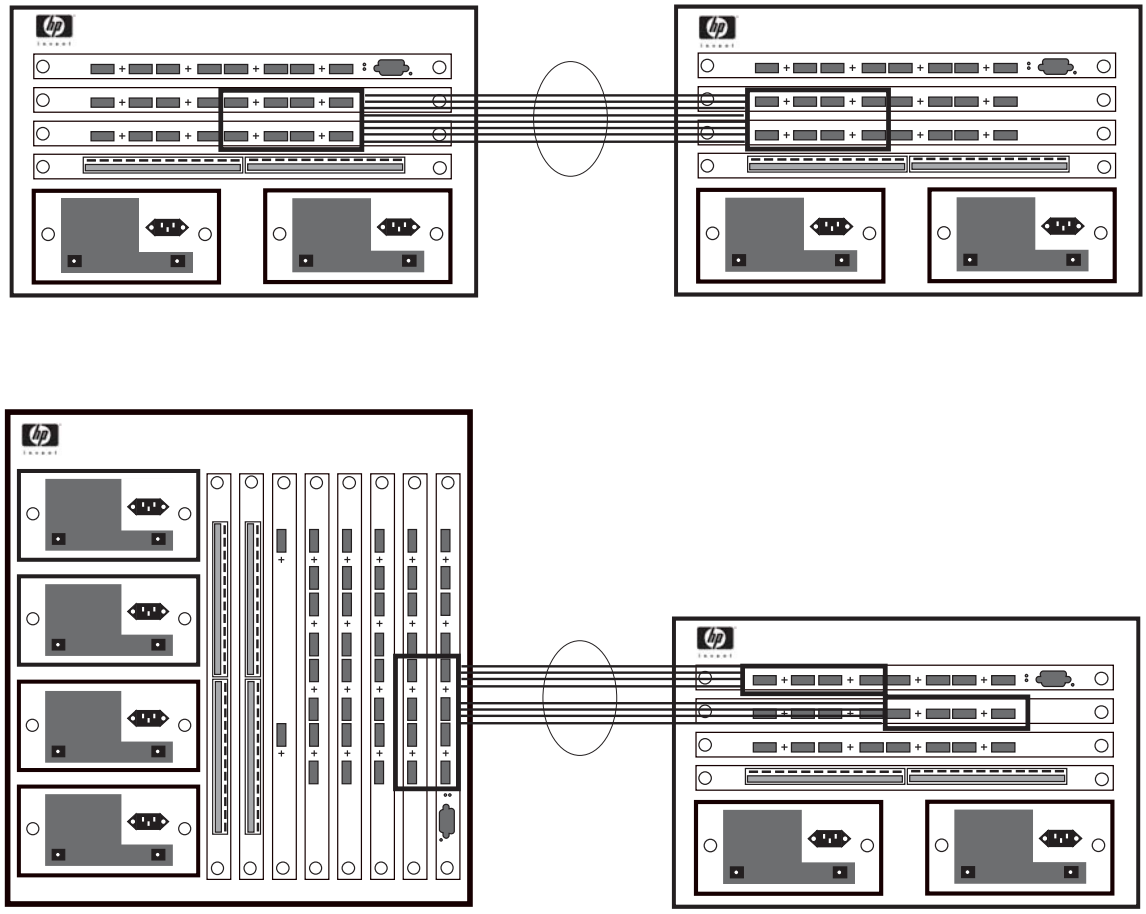

Figure 7.5 shows the valid 2-port and 4-port trunk groups on chassis 10/100 modules.

Figure 7.5 Valid 2-port and 4-port trunk groups on chassis 10/100 modules

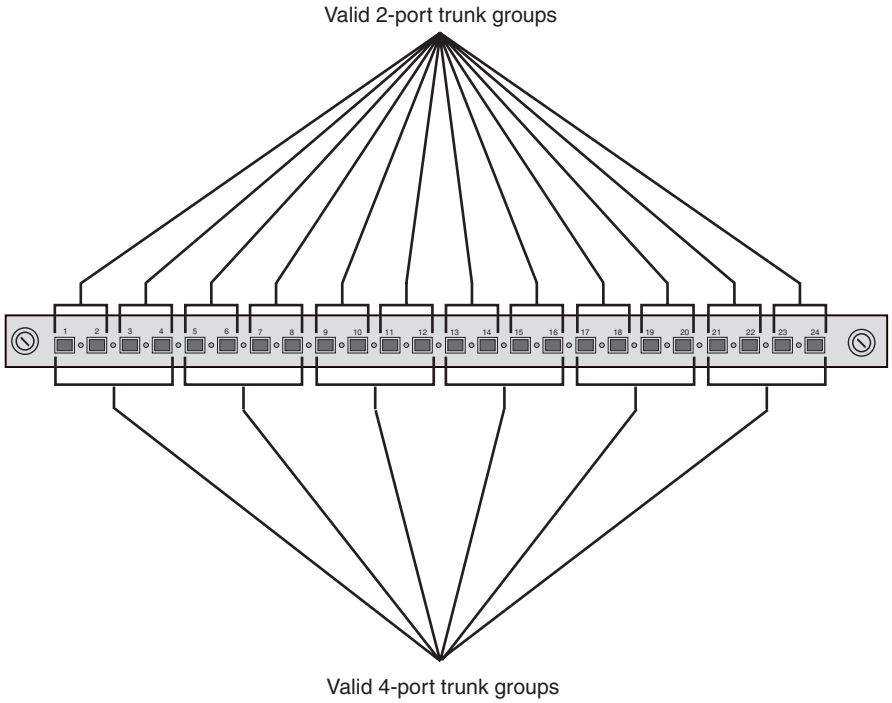

Additional Trunk Group Rules for Multi-Slot Trunk Groups

- You can configure a multi-slot trunk group on two Gigabit Ethernet modules.
- You can configure a maximum of eight ports in the trunk group.
- You can configure up to two groups of ports to make the trunk group and the groups must be alike. For example, you can group two sets of two ports together or two sets of four ports together but you cannot group a set of two ports with a set of four ports. Each group of ports can contain two or four ports.
- Each group of ports must begin with a primary port. On Gigabit Ethernet modules, the primary ports are 1, 3, 5, and 7.
- When you specify the ports in the trunk group, you must specify them in ascending numerical order, beginning with the primary port. For example, to specify a group containing ports 1/1 – 1/4 and 3/1 – 3/4, you must specify them in the order shown. You cannot specify 3/1 – 3/4 first.
- Port configuration for each trunk group is based on the configuration of the primary port. To change port parameters, you must change them on the primary port. The software automatically applies the changes to the other ports in the trunk group.
- If you plan to configure ports on a module into a server trunk group, use the following additional guidelines:
  - In software releases prior to 07.5.04, the management module(s) and the module that has the server trunk group's ports must be in the same set of slots (slots 1 – 7 or 9 – 15). Do not place the management module(s) and the module containing the trunk ports in separate sets of slots.

    This restriction was removed in software release 07.5.04 and later. Software releases 07.5.04 and later support placing the management module and forwarding module in different sets of slots. For single-slot server trunk group, the management module and forwarding module can be on any slot of the chassis.
  - Do not place the management module(s) or the module that has the server trunk group's ports in slot 8. The modules must both be in the same set or slots (slots 1–7 or 9–15).

These guidelines apply to a server trunk group that is configured on a single module or on a pair of modules (multi-slot trunk group). You do not need to follow these guidelines for a switch trunk group.

Trunk Group Load Sharing

When you configure a trunk group, you specify whether the trunk group is a “switch” trunk group or a “server” trunk group:

- Switch trunk group – Use this type of trunk group to connect one ProCurve Routing Switch to another ProCurve Routing Switch.
- Server trunk group – Use this type of trunk group to connect a ProCurve Routing Switch to a file server or single host device.

The HP device load shares across the ports in the trunk group. The method used for the load sharing depends on the following:

- Traffic type – Layer 2 or Layer 3
- Trunk type – Switch or server
- For certain traffic, port type on which the traffic enters the HP device (Gigabit or 10/100)

NOTE: The port type applies only to Layer 2 traffic on a server trunk group configured on a Routing Switch.

NOTE: On a device managed by a T-Flow, you can optimize server trunk load sharing on individual ports. See “Enabling Optimized Server Trunk Load Balancing (T-Flow only)” on page 7-19.

NOTE: Starting in software release 07.7.00, EP devices load balance IP traffic on server trunks based on source and destination TCP and UDP application ports (Layer 4 information), as well as on source and destination IP addresses (Layer 3 information). In addition, software release 07.7.00 enables you to configure server trunk load balancing per packet and to specify the maximum number of hash buckets per server trunk. See “Server Trunk Group Load Sharing Enhancements and Options (Release 07.7.00 and Higher)” on page 7-17.

For switch trunk groups, load sharing occurs as follows:

- For switched traffic, load sharing is based on the destination MAC address
- For routed traffic, load sharing is based on the destination IP address

To select a port, the device determines the port in the trunk group that has the fewest number of flows. Traffic for the destination address is sent over the selected port.

For server trunk groups, which connect a ProCurve Routing Switch to a file server or single host device, load sharing occurs as follows:

- For switched traffic, load sharing is based on source MAC address
- For routed traffic, load sharing is based on the combination of source and destination IP addresses

To select a port, the device calculates a hash value derived from source and destination IP addresses, then uses the hash value to select a port from the available ports in the trunk group.

Trunk Load Sharing with EP Modules

Table 7.2 lists how Enhanced Performance devices load balance traffic .

NOTE: The load sharing methods for server trunk groups also apply to trunks dynamically configured by 802.3ad link aggregation.

Table 7.2: HP Trunk Group Load Sharing – EP devices

Traffic Type	Trunk Type	Input Port Type	Load Balancing Method
Layer 2	Switch	10/100 Ethernet	Destination MAC address
		Gigabit Ethernet	Destination MAC address
		10 Gigabit Ethernet	Destination MAC address
	Server	10/100 Ethernet	Source MAC address
		Gigabit Ethernet	Source MAC address
		10 Gigabit Ethernet	Source MAC address
Layer 2 IP	Switch	10/100 Ethernet	Destination MAC address
		Gigabit Ethernet	Destination MAC address
		10 Gigabit Ethernet	Destination MAC address
	Server	10/100 Ethernet	Source and destination IP addresses
		Non-fragmented traffic	Source and destination IP addresses and source and destination TCP and UDP application port numbers
		Gigabit Ethernet	Source and destination IP addresses
		10 Gigabit Ethernet	Source and destination IP addresses

Table 7.2: HP Trunk Group Load Sharing – EP devices (Continued)

Traffic Type	Trunk Type	Input Port Type	Load Balancing Method
Layer 3 IP	Switch[a]	10/100 Ethernet	Destination IP address
		Gigabit Ethernet	Destination IP address
		10 Gigabit Ethernet	Destination IP address
	Server	10/100 Ethernet	Source and destination IP addresses
		Non-fragmented traffic[b]	Source and destination IP addresses and source and destination TCP and UDP application port number[b]
		Gigabit Ethernet	Source and destination IP addresses
		10 Gigabit Ethernet	Source and destination IP addresses

a.By default, Layer 3 IP traffic uses ip load-sharing by-host. Refer to the *Command Line Interface Reference for ProCurve 9300/9400 Series Routing Switches* for details on this command.
b.New in 07.6.01

Trunk Load Sharing with Standard (non-EP) Modules

Table 7.3 lists how Standard (non-EP) Routing Switches load balance traffic .

Table 7.3: HP Trunk Group Load Sharing

Traffic Layer	Trunk Group Type	Traffic Type	Load-Sharing Basis
Layer 3	Switch	IP	Destination IP address
		IPX	Destination IPX address
		AppleTalk	Destination AppleTalk address
		All other traffic types	Destination MAC address
	Server	IP	Destination IP address
		IPX	Destination IPX address
		AppleTalk	Destination AppleTalk address
		All other traffic types	Destination MAC address

Configuring a Trunk Group

1. Disconnect the cables from those ports on both systems that will be connected by the trunk group. Do not configure the trunk groups with the cables connected.

   NOTE: If you connect the cables before configuring the trunk groups and then rebooting, the traffic on the ports can create a spanning tree loop.

2. Configure the trunk group on one of the two Routing Switches involved in the configuration.
3. Save the configuration changes to the startup-config file.
4. Dynamically place the new trunk configuration into effect by entering the **trunk deploy** command at the global CONFIG level of the CLI.

   NOTE: If you are running a software release earlier than 07.5.04, you must reload the software to place a trunk configuration change into effect.

5. If the device at the other end of the trunk group is another Routing Switch, repeat Steps 2 – 4 for the other device.
6. When the trunk groups on both devices are operational, reconnect the cables to those ports that are now configured as trunk groups, starting with the first port (lead port) of each trunk group.
7. To verify the link is operational, use the **show trunk** command.

Example 1: Configuring the Trunk Groups Shown in Figure 7.1

To configure the trunk groups shown in Figure 7.1, enter the following commands. Notice that the commands are entered on multiple devices.

USING THE CLI

To configure the trunk group link between Router1 and Router2:

NOTE: The text shown in italics in the CLI example below shows messages echoed to the screen in answer to the CLI commands entered.

```
Router1(config)# trunk switch e 1/5 to 1/8
Trunk 2 is created for next power cycle.
Please save configuration to flash and reboot.
Router1(config)# write memory
Write startup-config in progress.
.Write startup-config done.
Router1(config)# exit
Router1# reload
```

NOTE: This example uses devices that are not running software release 07.5.004 or later. Devices running software earlier than 07.5.004 must be reloaded in order to place trunk configuration changes into effect. On devices running 07.5.004 or later, you can dynamically place trunk configuration changes into effect by entering the **trunk deploy** command at the global CONFIG level of the CLI.

To configure the trunk group link between Router2 and the server:

```
Router2(config)# trunk server e 1/3 to 1/4
Trunk 0 is created for next power cycle.
Please save configuration to flash and reboot.
Router2(config)# write memory
Write startup-config in progress.
.Write startup-config done.
Router2(config)# exit
Router2# reload
```

You then configure the trunk group on the HP 4000 Switch.

```
HP4000(config)# trunk e 17/18 trk1 trunk
HP4000(config)# write memory
Write startup-config in progress.
.Write startup-config done.
HP4000(config)# exit
HP4000# reload
```

USING THE WEB MANAGEMENT INTERFACE

To configure ports 5 – 8 as a trunk group between two Routing Switches or a Routing Switch and a server:

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Select the Trunk link.
   - If the device does not have any trunk groups configured, the Trunk configuration panel is displayed, as shown in the following example.
   - If a trunk group is already configured and you are adding a new one, click on the Add Trunk Group link to display the Trunk configuration panel, as shown in the following example.
   - If you are modifying an existing trunk group, click on the Modify button to the right of the row describing the trunk group to display the Trunk configuration panel, as shown in the following example.

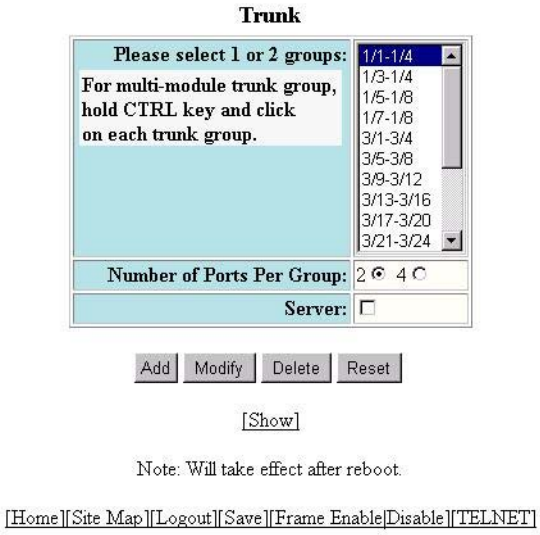

NOTE: This panel lists port ranges only for the slots that contain an active module. In addition, only the ranges that are valid for the module are listed.

The port ranges listed by the panel contain four ports, but the default number of ports in a group is two. If you select a group and leave the number of ports in a group at two, the software assigns the first two ports in the group you select to the trunk group. The last two ports do not become members of the trunk group.

---

4. Select a port range. For example, you can select 1/5 – 1/8.
5. Select the number of ports you want to use in the trunk group. You can select 2 or 4.

6. Click in the checkbox next to Server to place a checkmark in the box if the other end of the trunk group is a server. If the other end of the connection is a ProCurve Routing Switch, do not click this checkbox.
7. Click Apply to save the changes to the device's running-config file.
8. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.
9. Click on the plus sign next to Command in the tree view to list the command options.
10. Select the Reload link and select Yes when the Web management interface asks you whether you really want to reload the software.
11. 4If the other end of the trunk group is a Routing Switch, log in to the other device and follow the steps above.

Example 2: Configuring a Trunk Group That Spans Multiple Gigabit Ethernet Modules in a Routing Switch

To configure a trunk group that spans two modules in a 9300 series Routing Switch, use one of the following methods.

USING THE CLI

To configure a trunk group consisting of two groups of ports, 1/1 – 1/4 on module 1 and 4/5 – 4/8 on module 4, enter the following commands:

```
ProCurveRS(config)# trunk ethernet 1/1 to 1/4 ethernet 4/5 to 4/8
ProCurveRS(config-trunk-1/1-4/8)# write memory
ProCurveRS(config-trunk-1/1-4/8)# exit
ProCurveRS(config)# trunk deploy
```

NOTE: The **trunk deploy** command dynamically places trunk configuration changes into effect, without a software reload. This command is supported only in software release 07.5.04 and later. If you are running a release earlier than 07.5.04, you must reload the software to place trunk configuration changes into effect.

CLI Syntax

Syntax: [no] trunk [server | switch] ethernet <primary-portnum> to <portnum>
ethernet <primary-portnum> to <portnum>

Syntax: trunk deploy

The **server** | **switch** parameter specifies whether the trunk ports will be connected to a server or to another Routing Switch. This parameter affects the type of load balancing performed by the HP device. See "Trunk Group Load Sharing" on page 7-8. The default is **switch**.

Each **ethernet** parameter introduces a port group.

The <primary-portnum> **to** <portnum> parameters specify a port group. Notice that each port group must begin with a primary port. After you enter this command, the primary port of the first port group specified (which must be the group with the lower port numbers) becomes the primary port for the entire trunk group. For Gigabit Ethernet modules, the primary ports are 1, 3, 5, and 7.

To configure a trunk group consisting of two groups of two ports each, enter commands such as the following:

```
ProCurveRS(config)# trunk ethernet 1/1 to 1/2 ethernet 3/3 to 3/4
ProCurveRS(config)# write memory
ProCurveRS(config)# trunk deploy
```

Notice that the groups of ports meet the criteria for a multi-slot trunk group. Each group contains the same number of ports (two) and begins on a primary port (1/1 and 3/3).

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.

3. Select the Trunk link.
   - If the device does not have any trunk groups configured, the Trunk configuration panel is displayed, as shown in the following example.
   - If a trunk group is already configured and you are adding a new one, click on the Add Trunk Group link to display the Trunk configuration panel, as shown in the following example.
   - If you are modifying an existing trunk group, click on the Modify button to the right of the row describing the trunk group to display the Trunk configuration panel, as shown in the following example.

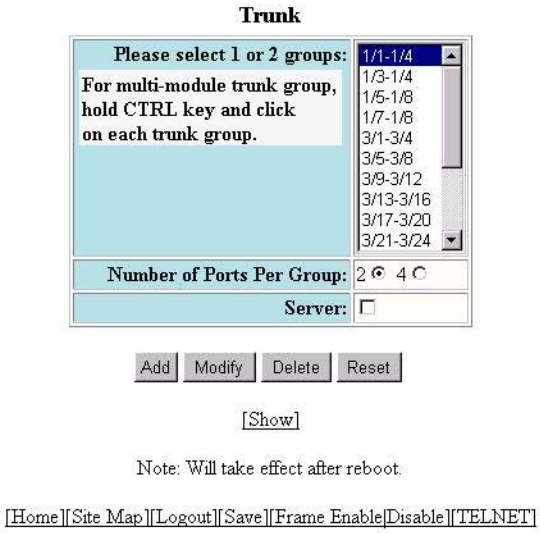

4. Select a port range. For example, you can select 1/5 – 1/8.
5. Select 2 or 4 to indicate the number of ports in each group. Each group must have the same number of ports.
6. Select the port groups. Each group begins with the primary port number for that group. To select two groups, click on the first group, then hold down the CTRL key and click on the second group. Do not select more than two groups.
7. Select Server if you are connecting the trunk group ports to a server. Otherwise, the software assumes you are connecting the trunk group ports to another Routing Switch and uses the default value Switch.
8. Click Apply to save the changes to the device's running-config file.
9. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.
10. Click on the plus sign next to Command in the tree view to list the command options.
11. Select the Reload link and select Yes when the Web management interface asks you whether you really want to reload the software.
12. If the other end of the trunk group is a Routing Switch, log in to the other device and follow the steps above.

NOTE: Hewlett-Packard recommends that you reload the software immediately after saving a trunk group configuration to flash memory, before making further configuration changes.

Configuring a Trunk Group of 10-Gigabit Ethernet Ports

Software release 07.6.04 enables you to configure 10 Gigabit Ethernet ports together in a trunk group (aggregate link).

To configure a trunk group containing two 10 Gigabit Ethernet ports, enter commands such as the following:

```
ProCurveRS(config)# trunk ethernet 1/1 to 2/1
ProCurveRS(config-trunk-1/1-2/1)# write memory
ProCurveRS(config-trunk-1/1-2/1)# exit
ProCurveRS(config)# trunk deploy
```

These commands configure a trunk group consisting of 10 Gigabit Ethernet ports 1/1 and 2/1, then deploy the trunk group. The trunk configuration does not take effect until you deploy it.

Syntax: [no] trunk [server | switch] ethernet <primary-portnum> to <secondary-portnum>

Syntax: trunk deploy

The **server** | **switch** parameter specifies whether the trunk ports will be connected to a server or to another Routing Switch. This parameter affects the type of load balancing performed by the HP device. See "Trunk Group Load Sharing" on page 7-8. The default is **switch**.

The <primary-portnum> parameter specifies the trunk group's primary port. You must specify an odd-numbered slot. See the table "Trunk Group Support" on page 7-3 for valid primary ports.

The <secondary-portnum> parameter specifies the secondary port in the trunk group. You must specify a port that is in the next slot number up from the primary port. For example, if the primary port is 1/1, specify 2/1 as the secondary port.

NOTE: Two-port trunk groups are supported for 10 Gigabit Ethernet. You cannot specify more than two ports.

To display configuration information and load-sharing statistics for the trunk group, enter the **show trunk** command. See "Displaying Trunk Group Configuration Information" on page 7-22.

Additional Trunking Options

The CLI contains commands for doing the following:

- Naming a trunk port
- Disabling or re-enabling a trunk port
- Deleting a trunk group

NOTE: To monitor the traffic on a trunk port, see "Monitoring an Individual Trunk Port" on page 6-52.

Naming a Trunk Port

To name an individual port in a trunk group, enter a command such as the following at the trunk group configuration level:

```
ProCurveRS(config-trunk-4/1-4/4)# port-name customer1 ethernet 4/2
```

Syntax: [no] port-name <text> ethernet <portnum>

The <text> parameter specifies the port name. The name can be up to 50 characters long.

This command assigns the name "customer1" to port 4/2 in the trunk group consisting of ports 4/1 – 4/4.

Disabling or Re-Enabling a Trunk Port

You can disable or re-enable individual ports in a trunk group. To disable an individual port in a trunk group, enter commands such as the following at the trunk group configuration level:

```
ProCurveRS(config-trunk-4/1-4/4)# config-trunk-ind
ProCurveRS(config-trunk-4/1-4/4)# disable ethernet 4/2
```

Syntax: [no] config-trunk-ind

Syntax: [no] disable ethernet <portnum>

The **config-trunk-ind** command enables configuration of individual ports in the trunk group. If you do not use this command, the **disable** command will be valid only for the primary port in the trunk group and will disable all ports

in the trunk group. You need to enter the **config-trunk-ind** command only once in a trunk group. After you enter the command, all applicable port configuration commands apply to individual ports only.

NOTE: If you enter **no config-trunk-ind**, all port configuration commands are removed from the individual ports and the configuration of the primary port is applied to all the ports. Also, once you enter the **no config-trunk-ind** command, the **enable**, **disable**, and **monitor** commands are valid only on the primary port and apply to the entire trunk group.

The **disable** command disables the port. The states of other ports in the trunk group are not affected.

If you have configured a name for the trunk port, you can specify the port name, as shown in the following example:

```
ProCurveRS(config-trunk-4/1-4/4)# config-trunk-ind
ProCurveRS(config-trunk-4/1-4/4)# disable customer1
```

Syntax: disable <portname>

To enable an individual port in a trunk group, enter commands such as the following at the trunk group configuration level:

```
ProCurveRS(config-trunk-4/1-4/4)# config-trunk-ind
ProCurveRS(config-trunk-4/1-4/4)# enable ethernet 4/2
```

Syntax: enable ethernet <portnum>

Syntax: enable <portname>

Disabling or Re-Enabling a Range or List of Trunk Ports

To disable a range of ports in a trunk group, enter commands such as the following:

```
ProCurveRS(config)# trunk switch ethernet 2/1 to 2/8
ProCurveRS(config-trunk-2/1-2/8)# config-trunk-ind
ProCurveRS(config-trunk-2/1-2/8)# disable ethernet 2/2 to 2/5
```

This command disables ports 2/2 – 2/5 in trunk group 2/1 – 2/8.

To disable a list of ports, enter a command such as the following:

```
ProCurveRS(config-trunk-2/1-2/8)# disable ethernet 2/2 ethernet 2/4 ethernet 2/7
```

This command disables ports 2/2, 2/4, and 2/7 in the trunk group.

You can specify a range and a list on the same command line. For example, to re-enable some trunk ports, enter a command such as the following:

```
ProCurveRS(config-trunk-2/1-2/8)# enable ethernet 2/2 to 2/5 ethernet 2/7
```

Syntax: [no] disable ethernet <portnum> [to <portnum> | ethernet <portnum>]

Syntax: [no] enable ethernet <portnum> [to <portnum> | ethernet <portnum>]

The **to** <portnum> parameter indicates that you are specifying a range. Specify the lower port number in the range first, then **to**, then the higher port number in the range.

The **ethernet** <portnum> parameter specifies an individual port. You can enter this parameter multiple times to specify a list, as shown in the examples above.

Deleting a Trunk Group

To delete a trunk group, use either of the following methods.

USING THE CLI

To delete a trunk group, use "**no**" in front of the command you used to create the trunk group. For example, to remove one of the trunk groups configured in the examples above, enter the following command:

```
ProCurveRS(config)# no trunk ethernet 1/1 to 1/2 ethernet 3/3 to 3/4
```

Syntax: no trunk ethernet <portnum> to <portnum>

USING THE WEB MANAGEMENT INTERFACE

To delete a trunk group:

1. Disconnect the ports to the server or Routing Switch at the other end of the trunk.
2. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
3. Click on the plus sign next to Configure in the tree view to display the configuration options.
4. Select the Trunk link to display a table listing the configured trunk groups.
5. Click the Delete button next to the trunk group you want to delete.
6. Select the Save link at the bottom of the dialog. Select Yes when prompted to save the configuration change to the startup-config file on the device's flash memory.
7. Click on the plus sign next to Command in the tree view to list the command options.
8. Select the Reload link and select Yes when the Web management interface asks you whether you really want to reload the software.

NOTE: If the other end of the trunk group is a Routing Switch, log in to the other system and follow the applicable steps above.

Server Trunk Group Load Sharing Enhancements and Options (Release 07.7.00 and Higher)

Software release 07.7.00 introduced the following server trunk load balancing enhancements and options on HP's EP devices:

- Load balancing IP traffic based on TCP and UDP application ports
- Ability to configure load balancing per packet
- Ability to configure the maximum number of hash buckets per server trunk

These new features improve the overall performance of server trunk load balancing for Layer 2 and Layer 3 IP traffic.

NOTE: These enhancements apply to server trunks only.

NOTE: These enhancements apply to IP traffic only.

Server Trunk Load Balancing Based on Application Ports

In software release 07.7.00 and later, HP's EP devices load balance IP traffic on server trunks based on source and destination TCP and UDP application ports (Layer 4 information), as well as on source and destination IP addresses (Layer 3 information). Previous releases supported server trunk load balancing based on Layer 3 information only. Adding Layer 4 information to the load balancing scheme enables the HP device to efficiently forward traffic to server trunk group ports.

You do not need to perform any configuration steps to enable this support. If you configure the device for server trunking, it automatically load balances traffic over server trunks based on Layer 3 and Layer 4 information.

For information about configuring server trunks, see the "Configuring Trunk Groups and Dynamic Link Aggregation" chapter in the *Advanced Configuration and Management Guide for ProCurve 9300/9400 Series Routing Switches*.

Note Regarding Fragmented Packets

The descriptions above apply to non-fragmented packets. For fragmented packets, the HP device uses the source and destination IP addresses only, so that all fragments in a session are forwarded over the same port.

Configuring Server Trunk Load Balancing Per Packet

NOTE: The configuration commands in this section apply to incoming (trunk) ports only.

Starting in release 07.7.00, you can configure the ports on an HP device to load balance IP traffic based on individual packets received on the interface. When you enable this feature, the device uses the IP packet headers to load balance the traffic among all the ports in the trunk group.

You configure this feature at the Interface level of the CLI, and not globally (on the entire device). When you configure this feature on the primary port of the trunk group, the software automatically applies it to the other ports in the trunk group.

To enable this feature, enter commands such as the following:

```
ProCurveRS(config)# interface e 1/1
ProCurveRS(config-if-e10000-1/1)# serv-trunk-per-pkt-lb
```

When interface e 1 receives IP packets destined for a trunk port, it uses information in the IP packet header to select the trunk port on which to forward the traffic.

Syntax: [no] serv-trunk-per-pkt-lb

Configuring the Maximum Number of Hash Buckets for Server Trunks

NOTE: This section applies to ProCurve Routing Switches only.

Server trunks use hash buckets to implement packet forwarding and load balancing. The hash buckets enable forwarding of packets in hardware, as opposed to forwarding them in software (sending them to the CPU). Packets forwarded in hardware travel faster in comparison to packets sent to the CPU for processing.

When the HP device learns that a specific packet has to go through an outgoing port, it places an entry in the hash bucket. The entry defines the data path from the incoming port to the outgoing port. When the device receives subsequent packets destined for the same path, it retrieves the entry in the hash bucket and forwards the packets accordingly.

In releases prior to 07.7.00, the HP device allocates a fixed number of hash buckets for each server trunk. This number is not configurable.

In software release 07.7.00, depending on the number of server trunks configured on the HP device, you can specify the maximum number of hash buckets per server trunk, up to a maximum of 256. In addition, the total number of hash buckets for all server trunks combined has increased. The 9315M supports a total of 1024 hash buckets, and the 9304M and 9308M support a total of 8192 hash buckets.

Increasing the number of hash buckets per server trunk enhances the speed and efficiency at which the HP device forwards and load balances IP packets on server trunk ports.

Table 7.4 shows the hash bucket configurations supported on the 9315M.

Table 7.4: Configurable Hash Buckets on the 9315M

Maximum Number of Hash Buckets per Server Trunk	Number of Server Trunks	Total Number of Hash Buckets for all Server Trunks Combined
16 – default value	15	240
32	15	480
64	15	960
128	8	1024
256	4	1024

Table 7.5 shows the hash bucket configurations supported on the 9304M and 9308M.

Table 7.5: Configurable Hash Buckets on the 9304M and 9308M

Maximum Number of Hash Buckets per Server Trunk	Number of Server Trunks	Total Number of Hash Buckets for all Server Trunks Combined
16	32	512
32	32	1024
64	32	2048
128	32	4096
256 – default value	32	8192

To configure the maximum number of hash buckets per server trunk, enter commands such as the following:

```
ProCurveRS(config)# system hash-per-server-trunk 64
ProCurveRS(config)# write mem
ProCurveRS(config)# end
ProCurveRS # reload
```

NOTE: You must reload the software to place this configuration in effect.

Syntax: [no] system hash-per-server-trunk <maximum number of hash buckets>

where **maximum number of hash buckets** can be 32, 64, 128, or 256. On the 9315M, the default is 16. On the 9304M and 9308M, the default is 256.

Enabling Optimized Server Trunk Load Balancing (T-Flow only)

You can optimize individual ports for server trunk load balancing. An optimized port load balances based on source and destination IP address but uses a smaller session table, which enables the port to more quickly forward traffic received on the port to the server trunk group ports.

NOTE: This enhancement applies only to the T-Flow.

NOTE: This enhancement applies to server trunk groups only, not to switch trunk groups.

Without optimization, the device performs the following types of load balancing for IP traffic.

Layer 2

The load balancing occurs at Layer 2 if the traffic is being forwarded in hardware. IP traffic on a server trunk group is load balanced as follows:

- On a Routing Switch:
  - IP traffic received on a 10/100 port or Gigabit port is load balanced based on destination IP address.

Layer 3

If any of the following features are enabled on a port, load balancing occurs in software using the entries in the session table. In this case, the IP traffic is load balanced based on source and destination IP address.

- ACLs
- Rate limiting (Fixed Rate Limiting or Adaptive Rate Limiting)
- NetFlow
- sFlow Export

- Network Address Translation (NAT)
- Policy-Based Routing (PBR)

If you do not have any of these features enabled on the port but you still want to load balance the traffic based on source and destination IP address, you can do so by enabling the server trunk load balancing optimization feature. Even if you do have one of the features above configured on the port, you can enhance load balancing performance by enabling the optimization feature. The optimization feature uses a smaller session table, which allows forwarding to occur more quickly.

NOTE: When you enable the server trunk load balancing optimization feature on a port, the feature listed above are disabled on that port. This occurs because the features use the session table, but the optimization feature uses a smaller session table than the other features. The configuration information for the other features is retained in the device's configuration file, but the features are disabled.

Example of Server Trunk Load Balancing at Layer 3

Figure 7.6 shows an example of how IP traffic is load balanced to server trunk ports when the traffic is forwarded at Layer 3. In this example, server trunk load balancing based on source and destination IP addresses is enabled on a Gigabit Ethernet port connected to a network containing multiple clients. Four other Ethernet ports are configured in a server trunk group that is connected to a multi-homed server. The server can have multiple network adapters or a single adapter with multiple ports that have unique MAC and IP addresses.

Figure 7.6 Server trunk load balancing based on source and destination IP addresses

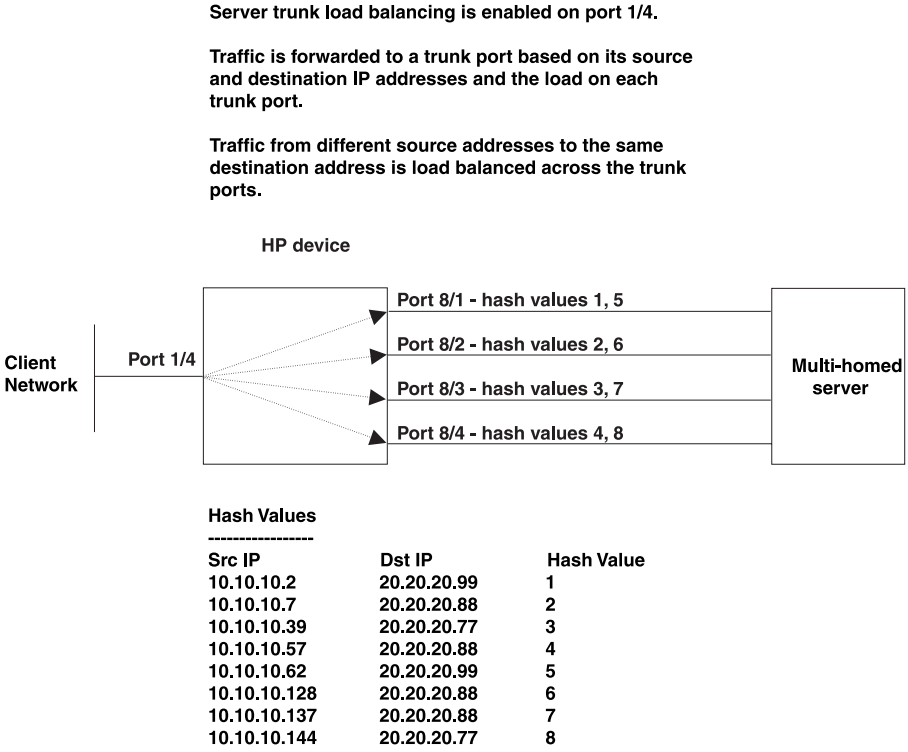

When the port connected to the client network receives traffic that needs to be forwarded to the server, the HP device selects one of the ports in the trunk group, and forwards the traffic on the selected port.

The HP device selects the trunk port based on a hash value, which can be a number from 1 – 256. The HP device calculates a hash value for traffic that enters the device through the server trunk load balancing port and exits the device through a trunk group. The hash value is calculated based on the source and destination IP addresses in the traffic.

After the HP device calculates the hash value for the traffic, the device examines the trunk ports connected to the destination address and selects the port with the fewest hash values already assigned. After calculating a hash value and assigning the value to a port, the device always uses the same port to forward traffic for the same source and destination IP addresses.

For example, the first time the HP device receives traffic from 10.10.10.7 addressed to 20.20.20.88, the device calculates the hash value 2 for the traffic. The device then checks the trunk ports to see whether a port is assigned to hash value 2.

- If a trunk port is assigned to hash value 2, the device uses that port to forward the traffic.
- Otherwise, the device assigns hash value 2 to the trunk port with the fewest hash values already assigned to it. The device continues to use this port for traffic with hash value 2, until a state change occurs on a trunk port or a trunk port is added or removed.

Trunk ports keep the hash values that are assigned to them until a trunk port's state changes or a trunk port is added or removed. When any of these changes occurs, the HP device clears the hash values from all of the trunk ports and begins calculating and assigning hash values again for new traffic.

Configuration Considerations

- You can enable the server trunk load balancing optimization feature on an individual port basis only. You cannot enable the feature on a virtual routing interface basis. This is true even if you have assigned a virtual routing interface to the trunk ports.
- Each TSP CPU has a separate hash bucket for the ports managed by the CPU. The buckets are independent of one another. Thus, if you enable the feature on more than one port and the ports are not managed by the same CPU, it is possible for the same hash values to be assigned to more than one trunk port, because the values are assigned separately by each CPU.
- When you enable the server trunk load balancing optimization feature on a port, the following features are disabled on the port:
  - ACLs
  - Rate limiting (Fixed Rate Limiting or Adaptive Rate Limiting)
  - NetFlow
  - sFlow Export
  - Network Address Translation (NAT)
  - Policy-Based Routing (PBR)

  The features are disabled because the server trunk load balancing optimization feature uses a simpler session table whose forwarding entries are keyed by source and destination IP addresses only. The features listed above require use of the standard session table, which also includes keys for the IP protocol and the source and destination TCP or UDP application ports (when the IP protocol is TCP or UDP).

  The configuration information for these features remains in the device's configuration file but the features are disabled on the port.

Enabling Server Trunk Load Balancing Optimization (T-Flow ModuleOnly)

To enable server trunk load balancing optimization, you enable the feature on the ports that will receive the traffic that needs to be load balanced. To enable the optimization feature on a port, enter the following command at the configuration level for the port:

```
ProCurveRS(config-if-e1000-1/4)# stlb
```

Syntax: [no] stlb

Displaying Server Trunk Load Balancing Information

To display the current hash assignments for server trunk ports, log on to the TSP CPU that is managing the ports, then enter the **show trunk** command. Here is an example.

```
ProCurveRS# rconsole 2 1
ProCurveRS2/1 # show trunk
ProCurveRS2/1 #Number of trunk groups: 1
Note: Value in () is for server trunk hashing.

TRUNK ID: 71    server:1        multi-slot:0
        configured ports: 8/1  8/2  8/3  8/4
        active ports    : 8/1 (2) 8/2 (2) 8/3 (2) 8/4 (1)
ProCurveRS2/1 # rconsole-exit
```

The **rconsole 2 1** command logs on to TSP CPU 1 on the T-Flow module in slot 2.

The **show trunk** command displays the trunk information for the ports managed by the CPU. The server trunk load balancing information is shown in bold type in this example. The number in parentheses indicates how many hash values are assigned to the port. The CPU assigns the hash values evenly to the trunk ports managed by the CPU. In this example, the next time the device needs to assign a hash value, the device will assign the value to port 8/4.

The **rconsole-exit** command logs out of the TSP CPU.

Syntax: show trunk

For information about the T-Flow, including how the module distributes management of the ports in the chassis, see “Using the T-Flow Redundant Management Module” on page 4-1.

Displaying Trunk Group Configuration Information

To display configuration information for the trunk groups configured on the Routing Switch, use one of the following methods. Each method displays information for configured trunk groups and operational trunk groups. A configured trunk group is one that has been configured in the software but has not been placed into operation by a reset or reboot. An operational trunk group is one that has been placed into operation by a reset or reboot.

USING THE CLI

Enter the following command at any CLI level:

```
ProCurveRS(config)# show trunk
Configured trunks:
Trunk Type   Ports
  1   Switch 1/1 1/2 1/3 1/4 2/1 2/2 2/3 2/4
Operational trunks:
Trunk Type   Ports                                    Duplex Speed Tag Priority
  1   Switch 1/1 1/2 1/3 1/4 2/1 2/2 2/3 2/4         None   None  No  level0
```

Syntax: show trunk [ethernet <portnum> to <portnum>]

The following table describes the information displayed by the **show trunk** command.

Table 7.6: CLI Trunk Group Information

This Field...	Displays...
Trunk	The trunk group number. The software numbers the groups in the display to make the display easy to use.
Type	The type of trunk group, which can be one of the following: • Server – The trunk group is connected to a server. • Switch – The trunk group is connected to another Routing Switch.

Table 7.6: CLI Trunk Group Information (Continued)

This Field...	Displays...
Ports	The ports in the trunk group.
Duplex	The mode of the port, which can be one of the following: • None – The link on the primary trunk port is down. • Full – The primary port is running in full-duplex. • Half – The primary port is running in half-duplex. **Note**: This field and the following fields apply only to operational trunk groups.
Speed	The speed set for the port. The value can be one of the following: • None – The link on the primary trunk port is down. • 10 – The port speed is 10 Mbps. • 100 – The port speed is 100 Mbps. • IG – The port speed is 1000 Mbps.
Tag	Indicates whether the ports have 802.1q VLAN tagging. The value can be Yes or No.
Priority	Indicates the Quality of Service (QoS) priority of the ports. The priority can be a value from 0 – 7.

To display trunk group information for specific ports, enter a command such as the following:

```
ProCurveRS(config)# show trunk ethernet 1/1 to 1/8

Configured trunks:

Trunk ID: 1
Type: Switch
Ports_Configured: 8
Primary Port Monitored: Jointly

Ports        1/1      1/2      1/3      1/4      1/5      1/6      1/7      1/8
Port Names   none     none     none     none     none     longna   test     none
Port_Status  enable   enable   enable   enable   disable  disable  enable   enable
Monitor      on       on       off      on       off      off      off      off
Mirror Port  3/3      3/4      N/A      3/5      N/A      N/A      N/A      N/A
Monitor Dir  both     in       N/A      out      N/A      N/A      N/A      N/A

Operational trunks:

Trunk ID: 1
Type: Switch
Duplex: Full
Speed: 1G
Tag: No
Priority: level0
Active Ports: 6

Ports           1/1      1/2      1/3      1/4      1/5    1/6     1/7      1/8
Link_Status     active   active   active   active   down   down    active   active
LACP_Status     ready    ready    ready    expired  down   down    ready    ready
Load Sharing
 Mac Address    3        2        2        2        0      0       6        1
 IP             0        0        0        0        0      0       0        0
 IPX            0        2        1        0        0      0       0        1
 Apple Talk     1        2        0        4        0      0       0        3
```

The display is divided into sections for configured trunks and operational trunks. A configured trunk group is one that has not been activated yet.

Table 7.7 describes the information displayed by the **show trunk** command.

Table 7.7: CLI Trunk Group Information

This Field...	Displays...
Trunk ID	The trunk group number. The software numbers the groups in the display to make the display easy to use.
Type	The type of trunk group, which can be one of the following: • Server – The trunk group is connected to a server. • Switch – The trunk group is connected to another Routing Switch.

Table 7.7: CLI Trunk Group Information (Continued)

This Field...	Displays...
Duplex	The mode of the port, which can be one of the following: • None – The link on the primary trunk port is down. • Full – The primary port is running in full-duplex. • Half – The primary port is running in half-duplex. **Note**: This field and the following fields apply only to operational trunk groups.
Speed	The speed set for the port. The value can be one of the following: • None – The link on the primary trunk port is down. • 10 – The port speed is 10 Mbps. • 100 – The port speed is 100 Mbps. • IG – The port speed is 1000 Mbps.
Tag	Indicates whether the ports have 802.1q VLAN tagging. The value can be Yes or No.
Priority	Indicates the Quality of Service (QoS) priority of the ports. The priority can be a value from 0 – 7.
Active Ports	The number of ports in the trunk group that are currently active.
Ports	The ports in the trunk group.
Link_Status	The link status or each port in the trunk group.
LACP_Status	This field appears in software releases 07.6.04 and later. For more information about this feature, see the section “Displaying and Determining the Status of Aggregate Links” on page 7-37. • Ready - The port is functioning normally in the trunk group and is able to transmit and receive LACP packets. • Expired - The time has expired (as determined by timeout values) and the port has shut down because the port on the other side of the link has stopped transmitting packets. • Down - The port’s physical link is down.
Load Sharing	The number of traffic flows currently being load balanced on the trunk ports. All traffic exchanged within the flow is forwarded on the same trunk port. For information about trunk load sharing, see “Trunk Group Load Sharing” on page 7-8.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write access. The System configuration panel is displayed.
2. Click on the plus sign next to Configure in the tree view to display the configuration options.
3. Select the Trunk link to display a table listing the configured trunk groups.

This display shows the following information.

Table 7.8: Web Management Trunk Group Information

This Field...	Displays...
Connection Type	The type of trunk group, which can be one of the following: • Server – The trunk group is connected to a server. • Switch – The trunk group is connected to another Routing Switch.
Port Members	The ports in the trunk group.

Dynamic Link Aggregation

The software supports the IEEE 802.3ad standard for link aggregation. This standard describes the Link Aggregation Control Protocol (LACP), a mechanism for allowing ports on both sides of a redundant link to configure themselves into a trunk link (aggregate link), without the need for manual configuration of the ports into trunk groups.

When you enable link aggregation on a group of HP ports, the HP ports can negotiate with the ports at the remote ends of the links to establish trunk groups.

Usage Notes

- You cannot use 802.3ad link aggregation on a port configured as a member of a static trunk group.
- This feature is supported only for 10/100 and Gigabit Ethernet ports.
- When the feature dynamically adds or changes a trunk group, the **show trunk** command displays the trunk as both configured and active. However, the **show running-config** or **write terminal** command does not contain a trunk command defining the new or changed trunk group.
- If the feature places a port into a trunk group as a secondary port, all configuration information except information related to link aggregation is removed from the port. For example, if port 1/3 has an IP interface, and the link aggregation feature places port 1/3 into a trunk group consisting of ports 1/1 – 1/4, the IP interface is removed from the port.
- If you use this feature on a Routing Switch that is running OSPF or BGP4, the feature causes these protocols to reset when a dynamic link change occurs. The reset includes ending and restarting neighbor sessions with OSPF and BGP4 peers, and clearing and relearning dynamic route entries and forwarding cache entries. Although the reset causes a brief interruption, the protocols automatically resume normal operation.
- You can enable link aggregation on 802.1q tagged ports (ports that belong to more than one port-based VLAN) in software release 07.7.00 and later.
- Dynamic Operation of Allocation Keys (section 43.6.2 in the 802.3ad specification) is supported in release 07.7.00 and later.

Configuration Rules

HP ports follow the same configuration rules for dynamically created aggregate links as they do for statically configured trunk groups. See “Trunk Group Rules” on page 7-3 and “Trunk Group Load Sharing” on page 7-8.

Figure 7.7 on page 7-28 shows some examples of valid aggregate links.

Figure 7.7 Examples of valid aggregate links

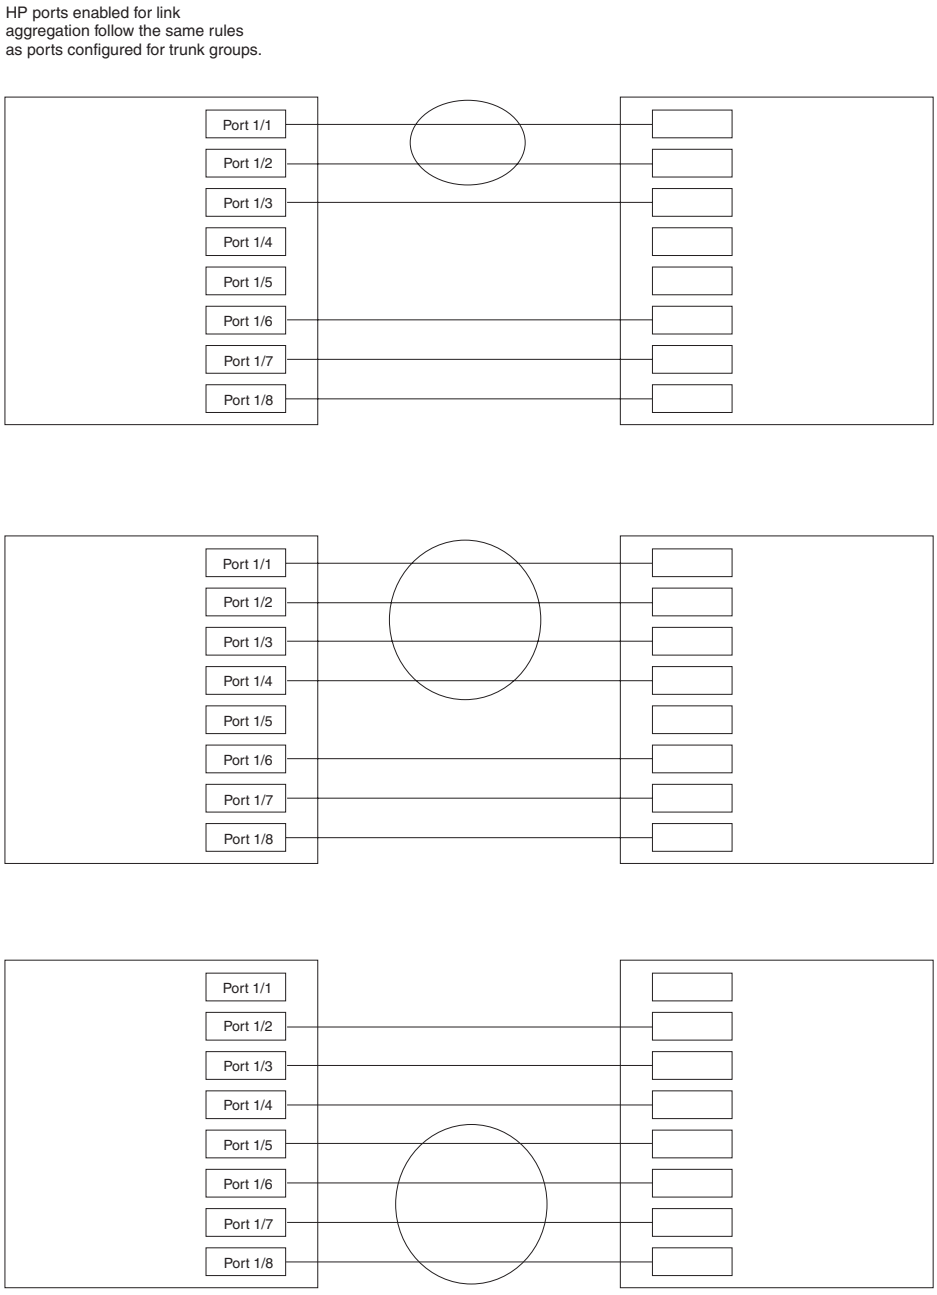

In this example, assume that link aggregation is enabled on all of the links between the HP device on the left and the device on the right (which can be either an HP device or another vendor's device). Notice that some ports are not able to join an aggregate link even though link aggregation is enabled on them. The ports that are not members of aggregate links in this example are not following the configuration rules for trunk links on HP devices.

The HP rules apply to an HP device even if the device at the other end is from another vendor and uses different rules. See "Trunk Group Rules" on page 7-3.

The link aggregation feature automates trunk configuration but can coexist with HP's trunk group feature. Link aggregation parameters do not interfere with trunk group parameters.

NOTE: Use the link aggregation feature only if the device at the other end of the links you want to aggregate also supports IEEE 802.3ad link aggregation. Otherwise, you need to manually configure the trunk links.

Link aggregation support is disabled by default. You can enable the feature on an individual port basis, in active or passive mode.

- Active mode – When you enable a port for active link aggregation, the HP port can exchange standard LACP Protocol Data Unit (LACPDU) messages to negotiate trunk group configuration with the port on the other side of the link. In addition, the HP port actively sends LACPDU messages on the link to search for a link aggregation partner at the other end of the link, and can initiate an LACPDU exchange to negotiate link aggregation parameters with an appropriately configured remote port.
- Passive mode – When you enable a port for passive link aggregation, the HP port can exchange LACPDU messages with the port at the remote end of the link, but the HP port cannot search for a link aggregation port or initiate negotiation of an aggregate link. Thus, the port at the remote end of the link must initiate the LACPDU exchange.

NOTE: HP recommends that you disable or remove the cables from the ports you plan to enable for dynamic link aggregation. Doing so prevents the possibility that LACP will use a partial configuration to talk to the other side of a link. A partial configuration does not cause errors, but does sometimes require LACP to be disabled and re-enabled on both sides of the link to ensure that a full configuration is used. It's easier to disable a port or remove its cable first. This applies both for active link aggregation and passive link aggregation.

802.3ad Enhancements in Release 07.6.04

Software release 07.6.04 contains the following enhancements to 802.3ad support:

- Adaptation to trunk disappearance. The HP device will tear down an aggregate link if the device at the other end of the link reboots or brings all the links down. Tearing the aggregate link down prevents a mismatch if the other device has a different trunk configuration following the reboot or re-establishment of the links.
- The criteria for being eligible to be in an aggregate link are more flexible. A range of ports can contain down ports and still be eligible to become an aggregate link.

Adaptation to Trunk Disappearance

Release 07.6.04 prevents trunk mismatches caused when one device changes the number of ports in group of ports that has become part of an 802.3 aggregate link. In 07.6.04 and later, if a device changes the number of ports in an active aggregate link, the HP device on the other end of the link tears down the link. Once the other device recovers, 802.3 can renegotiate the link without a mismatch.

In previous releases, it is possible for a trunk mismatch to occur between two devices that have established an aggregate link. This can occur if one of the devices reboots or brings the trunk links down, then re-establishes the links but with a different number of trunk ports. Figure 7.8 shows an example.

Figure 7.8 Trunk port mismatch

Four ports on each device are eligible for link aggregation. The device negotiates a four-port trunk using the ports.

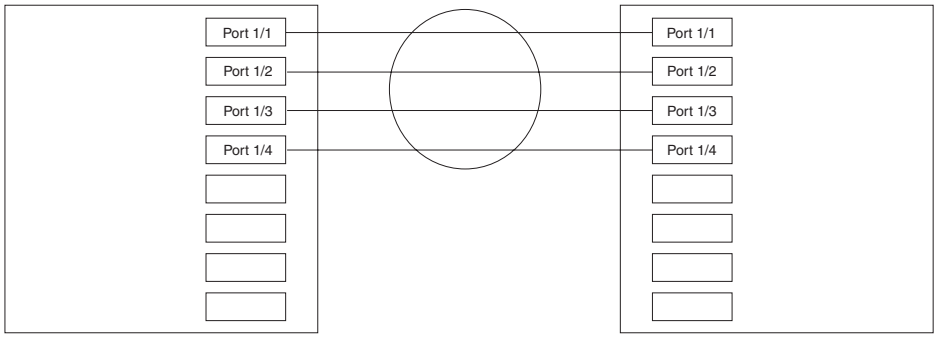

One device reloads, after which only two of its ports are eligible for link aggregation.

However, the first device is still configured with the four-port trunk group. The trunks are mismatched.

This type of mismatch does not occur in release 06.7.01 and later.

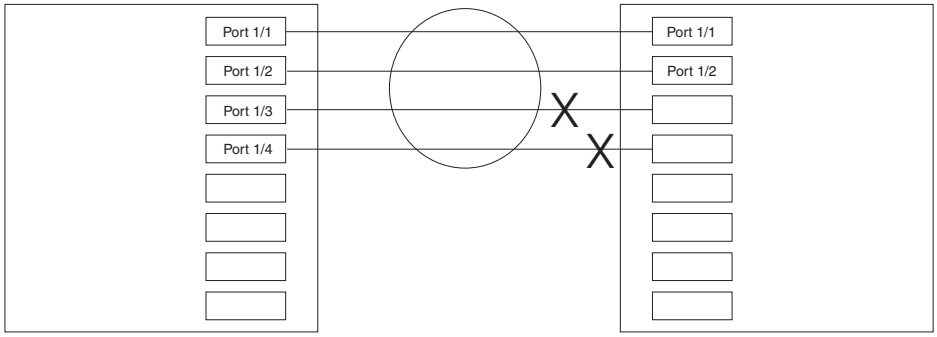

Flexible Trunk Eligibility

Software release 07.6.01b also increases the tolerance for down ports during link negotiation. In previous releases, all the ports in a valid trunk configuration (2-port, 4-port, or 8-port trunk starting on a valid primary port number) need to be up. Thus, in previous releases, if you enable link aggregation on four ports but one of the ports is down, the device will negotiate based only on a valid two-port trunk group consisting of two of the up ports. For example, if you enable link aggregation on ports 1/1 - 1/4 and port 1/3 is down, 802.3ad will negotiate only for a two-port link consisting of ports 1/1 and 1/2.

In release 07.6.01b and later, the device groups the device's ports into 2-port groups consisting of an odd-numbered port and the next even-numbered port. For example, ports 1/1 and 1/2 are a two-port group, as are ports 1/3 and 1/4, 9/1 and 9/10, and do on. If either of the ports in a two-port group is up, the device considers both ports to be eligible to be in an aggregate link.

Figure 7.9 shows an example of 2-port groups in a range of eight ports on which link aggregation is enabled. Based on the states of the ports, some or all of them will be eligible to be used in an aggregate link.

Figure 7.9 Two-port groups used to determine aggregation eligibility

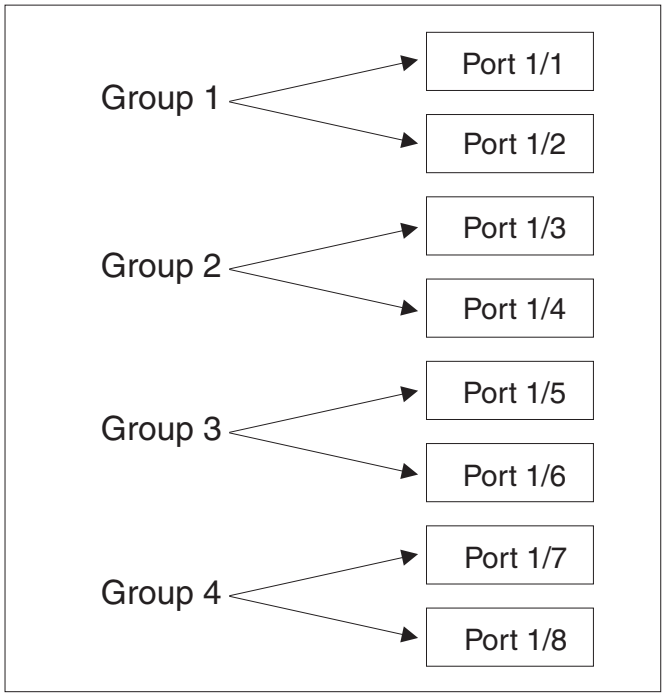

Table 7.9 shows examples of the ports from Figure 7.9 that will be eligible for an aggregate link based on individual port states.

Table 7.9: Port Eligibility for Link Aggregation

	Port Group 1		Port Group 2		Port Group 3		Port Group 4		Trunk Eligibility
	1/1	1/2	1/3	1/4	1/5	1/6	1/7	1/8	
Link State	Up	Up	Up	Up	Up	Up	Up	Up	8-port 1/1 – 1/8
	Up	Up	Up	Up	Up	Down	Up	Up	8-port 1/1 – 1/8
	Up	Up	Up	Up	Up	Down	Up	Down	8-port 1/1 – 1/8
	Up	Up	Up	Up	Down	Down	Down	Up	4-port 1/1 – 1/4
	Down	Down	Down	Up	Up	Up	Up	Up	4-port 1/5 – 1/8
	Up	Down	Down	Down	Up	Down	Down	Down	2-port 1/1 – 1/2

As shown in these examples, all or a subset of the ports within a port range will be eligible for formation into an aggregate link based on port states. Notice that the sets of ports that are eligible for the aggregate link must be valid static trunk configurations. For example, a 4-port link consisting of ports 1/4 – 1/7 is not valid because this port configuration is not valid for static trunk groups on the HP device.

Enabling Link Aggregation

By default, link aggregation is disabled on all ports. To enable the feature, use one of the following CLI methods.

USING THE CLI

To enable link aggregation on a set of ports, enter commands such as the following at the interface configuration level of the CLI.

NOTE: Configuration commands for link aggregation differ depending on whether you are using the default link aggregation key automatically assigned by the software, or if you are assigning a different, unique key. Follow the commands below, according to the type of key you are using. For more information about keys, see “Key” on page 7-33.

Using the Default Key Assigned by the Software

```
ProCurveRS(config)# interface ethernet 1/1
ProCurveRS(config-if-e1000-1/1)# link-aggregate active
ProCurveRS(config)# interface ethernet 1/2
ProCurveRS(config-if-e1000-1/2)# link-aggregate active
```

The commands in this example enable the active mode of link aggregation on ports 1/1 and 1/2. The ports can send and receive LACPDU messages. Note that these ports will use the default key, since one has not been explicitly configured.

Assigning a Unique Key

```
ProCurveRS(config)# interface ethernet 1/1
ProCurveRS(config-if-e1000-1/1)# link-aggregate configure key 10000
ProCurveRS(config-if-e1000-1/1)# link-aggregate active
ProCurveRS(config)# interface ethernet 1/2
ProCurveRS(config-if-e1000-1/2)# link-aggregate configure key 10000
ProCurveRS(config-if-e1000-1/2)# link-aggregate active
```

The commands in this example assign the key 10000 and enable the active mode of link aggregation on ports 1/1 and 1/2. The ports can send and receive LACPDU messages.

NOTE: As shown in this example, when configuring a key, it is pertinent that you assign the key prior to enabling link aggregation.

The following commands enable passive link aggregation on ports 1/5 – 1/8:

```
ProCurveRS(config)# interface ethernet 1/5 to 1/8
ProCurveRS(config-mif-1/5-1/8)# link-aggregate passive
```

The commands in this example enable the passive mode of link aggregation on ports 1/5 – 1/8. These ports wait for the other end of the link to contact them. After this occurs, the ports can send and receive LACPDU messages.

To disable link aggregation on a port, enter a command such as the following:

```
ProCurveRS(config-if-e1000-1/8)# link-aggregate off
```

Syntax: [no] link-aggregate active | passive | off

Syntax: [no] link-aggregate configure [system-priority <num>] | [port-priority <num>] | [key <num>] | [type server | switch]

NOTE: For more information about keys, including details about the syntax shown above, see “Key” on page 7-33.

Link Aggregation Parameters

You can change the settings for the following link aggregation parameters, on an individual port basis:

- System priority
- Port priority
- Link type
- Key

System Priority

The system priority specifies the HP device’s link aggregation priority relative to the devices at the other ends of the links on which link aggregation is enabled. A higher value indicates a lower priority. You can specify a priority from 0 – 65535. The default is 1.

NOTE: If you are connecting the HP device to another vendor’s device and the link aggregation feature is not working, set the system priority on the HP device to a lower priority (a higher priority value). In some cases, this change allows the link aggregation feature to operate successfully between the two devices.

Link Type

The link type specifies whether the trunk is connecting to a server (server link) or to another networking device (switch link). The default link type is switch.

Key

Every port that is 802.3ad-enabled has a key. The key identifies the group of potential trunk ports to which the port belongs. Ports with the same key are called a key group and are eligible to be in the same trunk group.

When you enable link-aggregation on a tagged or untagged port, HP’s software assigns a default key to the port. The default key is based on the position of the port within an eight-port group (the maximum number of ports in a trunk group on a Routing Switch). The software assigns the keys in ascending numerical order, beginning with key 0 for the first group of eight ports. For example, a 24-port module in chassis slot 1 contains keys 0, 1, and 2 by default. Ports 1/1 – 1/8 have key 0, ports 1/9 – 1/16 have key 1, and so on.

All ports within an aggregate link must have the same key. However, if the device has ports that are connected to two different devices, and the port groups allow the ports to form into separate aggregate links with the two devices, then each group of ports can have the same key while belonging to separate aggregate links with different devices. Figure 7.10 on page 7-34 shows an example.

Figure 7.10 Ports with the same key in different aggregate links

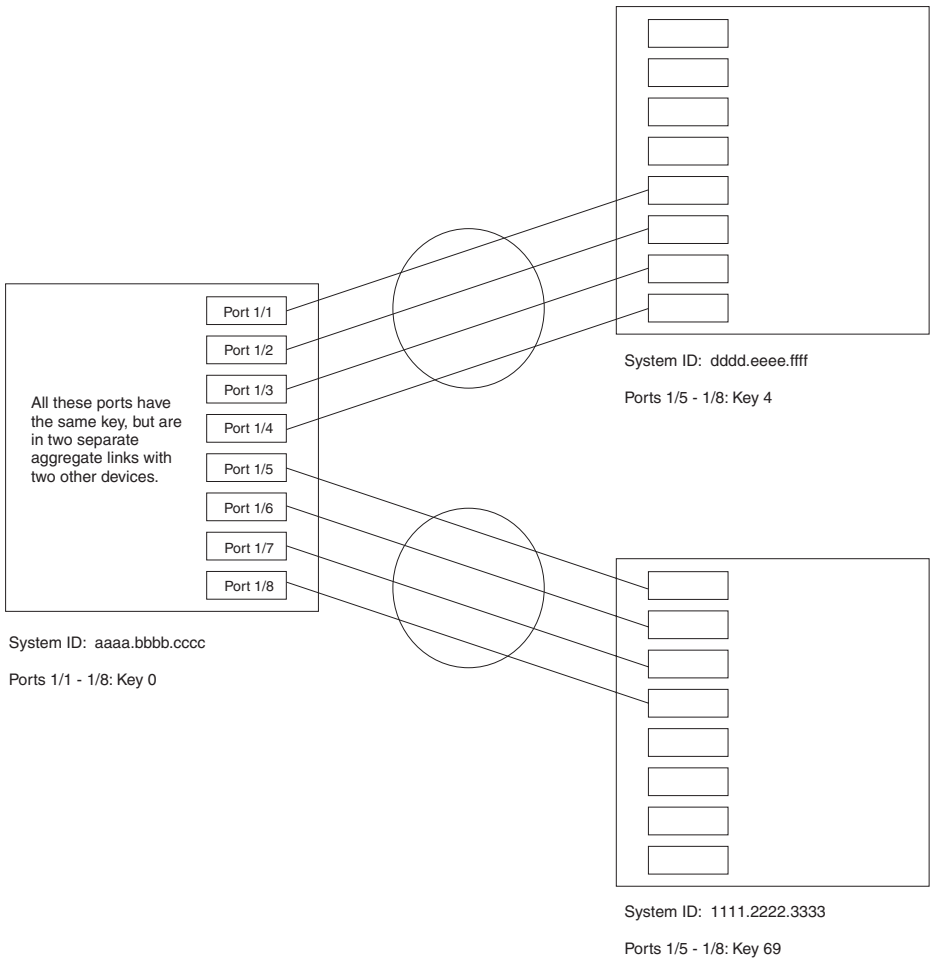

Notice that the keys between one device and another do not need to match. The only requirement for key matching is that all the ports within an aggregate link on a given device must have the same key.

Devices that support multi-slot trunk groups can form multi-slot aggregate links using link aggregation. However, the link aggregation keys for the groups of ports on each module must match. For example, if you want to allow link aggregation to form an aggregate link containing ports 1/1 – 1/4 and 3/5 – 3/8, you must change the link aggregation key on one or both groups of ports so that the key is the same on all eight ports. Figure 7.11 on page 7-35 shows an example.

Figure 7.11 Multi-slot aggregate link

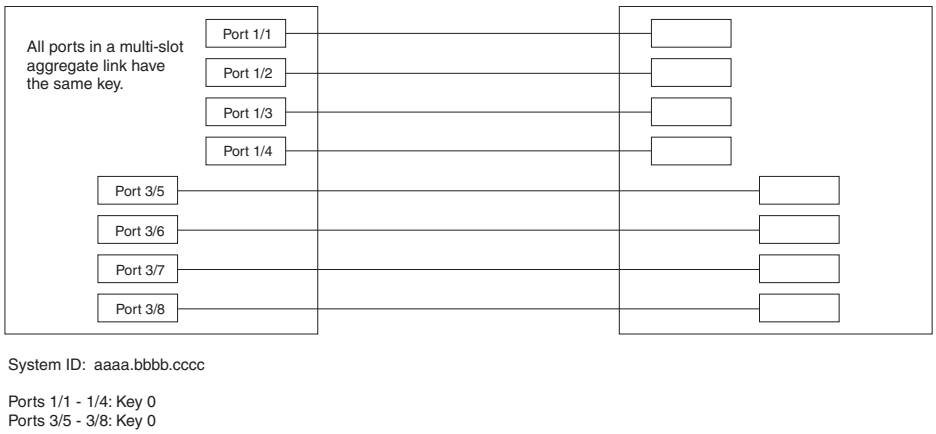

By default, the device's ports are divided into 4-port groups. The software dynamically assigns a unique key to each 4-port group. If you need to divide a 4-port group into two 2-port groups, change the key in one of the groups so that the two 2-port groups have different keys. For example, if you plan to use ports 1/1 and 1/2 in VLAN 1, and ports 1/3 and 1/4 in VLAN 2, change the key for ports 1/3 and 1/4.

NOTE: If you change the key for a port group, HP recommends that you use the value 10000 or higher, to avoid potential conflicts with dynamically created keys.

Dynamic Operation of Allocation Keys

Starting with software release 07.7.00, the HP device dynamically changes a port's key based on changes to the port's VLAN membership.

When you change a port's VLAN membership, the device searches through existing key groups for a port with matching port properties. Specifically, it searches for a match on all three of the following properties:

- VLAN ID
- default key
- port tag type (tagged or untagged)

If it finds a match, the port (whose VLAN membership you are changing) gets the matching port's key. If it does not find a match, the port gets a new key.

NOTE: For multi-slot trunk groups, you must manually configure the keys in the trunk group(s) to match. For instructions on configuring keys manually, see "Configuring Keys For Ports with Link Aggregation Enabled" on page 7-37.

How Changing a Port's VLAN Membership Affects Trunk Groups and Dynamic Keys

When you change a port's VLAN membership, and the port is currently a member of a trunk group, the following changes occur:

- The HP device tears down the existing trunk group.
- All ports in the trunk group get a new key.
- The new key group aggregates into a new trunk group.

When you change a port's VLAN membership, and the port is not a member of a trunk group, the following changes occur:

- The port gets a new key depending on changes to the port's VLAN tag type, as follows:
  - Tagged to Tagged VLAN – The primary port of the trunk group gets a new key.

  - Tagged to Untagged VLAN –The port gets the default key for untagged ports.
  - Untagged to Tagged VLAN – If the HP device finds a port with matching port properties, the port gets that port's key. If it doesn't find one, the port gets a new key.
  - Untagged to Untagged VLAN – The port gets a new key depending on whether it's in the default VLAN or not. If there is a trunk group associated with the key, it is not affected.
- All other ports keep their existing key.
- The new key groups try to aggregate into trunk groups.

Viewing Keys for Tagged Ports

To display link aggregation information, including the key for a specific port, enter a command such as the following at any level of the CLI:

```
ProCurveRS# show link-aggregation ethernet 1/1

System ID: 00e0.52a9.bb00
Port  [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp]
1/1        0        0       0   No    L   No   No   No   No   No   No
```

The command in this example shows the key and other link aggregation information for port 1/1.

To display link aggregation information, including the key for all ports on which link aggregation is enabled, enter the following command at any level of the CLI:

```
ProCurveRS# sh link-agg

System ID: 0004.8055.b200
Long  timeout: 90, default: 90
Short timeout: 3, default: 3

Port  [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp][Ope]
1/1        1        1   10000   Yes   S   Agg  Syn  Col  Dis  Def  No   Dwn
1/2        1        1   10000   Yes   S   Agg  Syn  Col  Dis  Def  No   Dwn
2/1        1        1   10000   Yes   S   Agg  Syn  Col  Dis  Def  No   Dwn
2/2        1        1   10000   Yes   S   Agg  Syn  Col  Dis  Def  No   Dwn
4/1        1        1     480   Yes   S   Agg  Syn  Col  Dis  Def  No   Dwn
4/2        1        1     480   Yes   S   Agg  Syn  Col  Dis  Def  No   Dwn
4/3        1        1     480   Yes   S   Agg  Syn  Col  Dis  Def  No   Dwn
4/4        1        1     480   Yes   S   Agg  Syn  Col  Dis  Def  No   Dwn
4/17       1        1     481   Yes   S   Agg  Syn  Col  Dis  Def  No   Ope
4/18       1        1     481   Yes   S   Agg  Syn  Col  Dis  Def  No   Ope
4/19       1        1     481   Yes   S   Agg  Syn  Col  Dis  Def  No   Ope
4/20       1        1     481   Yes   S   Agg  Syn  Col  Dis  Def  No   Ope
```

For information about the fields in this display, see Table 7.10 on page 7-39.

Syntax: show link-aggregation [ethernet <portnum>]

Possible values: N/A

Default value: N/A

Configuring Link Aggregation Parameters

You can configure one or more parameters on the same command line, and you can enter the parameters in any order.

NOTE: For key configuration only, configuration commands differ depending on whether or not link aggregation is enabled on the port(s). Follow the appropriate set of commands below, according to your system's configuration.

For example, to change a port group's key from the one assigned by the software to another value, enter commands such as the following:

NOTE: Use this command sequence to change the key for ports that do not have link aggregation enabled, and for all other link aggregation parameters (i.e., system priority, port priority, and link type).

```
ProCurveRS(config)# interface ethernet 1/1 to 1/4
ProCurveRS(config-mif-1/1-1/4)# link-aggregate configure key 10000
ProCurveRS(config-mif-1/1-1/4)# interface ethernet 3/5 to 3/8
ProCurveRS(config-mif-3/5-3/8)# link-aggregate configure key 10000
```

Configuring Keys For Ports with Link Aggregation Enabled

NOTE: As shown in this command sequence, to change the key on ports that already have link aggregation enabled, you must first turn OFF link aggregation, configure the new key, then re-enable link aggregation.

```
ProCurveRS(config)# interface ethernet 1/1 to 1/4
ProCurveRS(config-mif-1/1-1/4)# link-aggregate off
ProCurveRS(config-mif-1/1-1/4)# link-aggregate configure key 10000
ProCurveRS(config-mif-1/1-1/4)# link-aggregate active
ProCurveRS(config-mif-1/1-1/4)# interface ethernet 3/5 to 3/8
ProCurveRS(config-mif-3/5-3/8)# link-aggregate off
ProCurveRS(config-mif-3/5-3/8)# link-aggregate configure key 10000
ProCurveRS(config-mif-3/5-3/8)# link-aggregate active
```

These commands change the key for ports 1/1 – 1/4 and 3/5 – 3/8 to 10000. Since all ports in an aggregate link must have the same key, the command in this example enables ports 1/1 – 1/4 and 3/5 – 3/8 to form a multi-slot aggregate link.

Syntax: [no] link-aggregate configure [system-priority <num>] | [port-priority <num>] | [key <num>] | [type server | switch]

The **system-priority** <num> parameter specifies the HP device's link aggregation priority. A higher value indicates a lower priority. You can specify a priority from 0 – 65535. The default is 1.

The **port-priority** <num> parameter specifies an individual port's priority within the port group. A higher value indicates a lower priority. You can specify a priority from 0 – 65535. The default is 1.

The **key** <num> parameter identifies the group of ports that are eligible to be aggregated into a trunk group. The software automatically assigns a key to each group of ports. The software assigns the keys in ascending numerical order, beginning with 0. You can change a port group's key to a value from 0 – 65535.

NOTE: If you change the key for a port group, HP recommends that you use the value 10000 or higher, to avoid potential conflicts with dynamically created keys.

The **type server | switch** parameter specifies whether the port group is connected to a server (**server**) or to another networking device (**switch**). The default is **switch**.

You can enter one or more of the command's parameters on the same command line, in any order.

Displaying and Determining the Status of Aggregate Links

Software release 07.6.04 and later provides the ability to determine the status of ports that are members of an aggregate link, and whether or not LACPDU messages are being transmitted between the ports. In releases prior to 07.6.04, this level of detail was not readily available. With the link aggregation enhancement, the **show link-aggregation** command provides the ability to view the status of dynamic links.

The following section provides details about the events that can affect the status of ports in an aggregate link and the status of LACP messages exchanged between the ports. Later sections provide instructions for viewing these status reports.

About Blocked Ports

HP devices can block traffic on a port or shut down a port that is part of a trunk group or aggregate link for the following reasons:

- For the purpose of link aggregation, the ports on HP devices are grouped into pairs of two; one odd-numbered port, and the next even-numbered port. When you configure link aggregation on a port (for instance, on an odd-numbered port), this port will be blocked and unable to join a trunk group until you configure the adjacent port (the even-numbered port) as part of the aggregate link. When you configure both ports with link aggregation and assign both ports the same key, both ports are able to join a trunk group. Once the ports become part of a trunk group, they can transmit and receive LACP packets.

NOTE: Ports that are configured as part of an aggregate link must also have the same key. For more information about assigning keys, see the section titled “Configuring Link Aggregation Parameters” in the *Installation and Basic Configuration Guide for ProCurve 9300 Series Routing Switches*.

- When a port joins a trunk group and the port on the other end of the link shuts down or stops transmitting LACP packets, the HP device blocks the port. Depending on the timeout value set on the port, the link aggregation information expires.

NOTE: For more information about timeout values, see the section titled “Displaying Link Aggregation Information” in the *Installation and Basic Configuration Guide for ProCurve 9300 Series Routing Switches*.

If either of these events occur, the HP device shuts down the port and notifies all the upper layer protocols that the port is down.

HP devices can also block traffic on a port that is initially configured with link aggregation. The port is blocked until it joins a trunk group. In this case, traffic is blocked, but the port is still operational.

A port remains blocked until one of the following events occur:

- Link aggregation is enabled on the adjacent port (the paired port) and both ports have the same key
- LACP brings the port back up
- The port joins a trunk group

Displaying Link Aggregation and Port Status Information

Use the **show link-aggregation** command to determine the operational status of ports associated with aggregate links.

To display the link aggregation information for a specific port, enter a command such as the following at any level of the CLI:

```
ProCurveRS(config-mif-1/1-1/8)# show link-aggregation ethernet 1/1
System ID: 00e0.52a9.bb00
Port  [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp] [Ope]
1/1        0        0       0   No    L   No   No   No   No   No   No    Ope
```

The command in this example shows the link aggregation information for port 1/1.

NOTE: The **Ope** column displays in software releases 07.6.04 and later.

To display the link aggregation information for all ports on which link aggregation is enabled, enter the following command at any level of the CLI:

```
ProCurveRS(config)# show link-aggregation

System ID: 00e0.52a9.bb00
Port  [Sys P] [Port P] [ Key ] [Act][Tio][Agg][Syn][Col][Dis][Def][Exp][Ope]
1/1        1        1       0   No    L   Agg  Syn  No   No   Def  Exp  Ope
1/2        1        1       0   No    L   Agg  Syn  No   No   Def  Exp  Ina
1/3        1        1       0   No    L   Agg  Syn  No   No   Def  Exp  Ina
1/4        1        1       0   No    L   Agg  Syn  No   No   Def  Exp  Blo
1/5        1        1       1   No    L   Agg  No   No   No   Def  Exp  Ope
1/6        1        1       1   No    L   Agg  No   No   No   Def  Exp  Ope
1/7        1        1       1   No    L   Agg  No   No   No   Def  Exp  Dwn
1/8        1        1       1   No    L   Agg  No   No   No   Def  Exp  Dwn
```

NOTE: The **Ope** column displays in software releases 07.6.04 and later.

Syntax: show link-aggregation [ethernet <portnum>]

Use **ethernet <portnum>** to display link-aggregation information for a specific port.

NOTE: Ports that are configured as part of an aggregate link must also have the same key. For more information about assigning keys, see the section titled “Configuring Link Aggregation Parameters” in the *Installation and Basic Configuration Guide for ProCurve 9300 Series Routing Switches*.

The **show link aggregation** command shows the following information.

Table 7.10: CLI Display of Link Aggregation Information

This Field...	Displays...
System ID	Lists the base MAC address of the device. This is also the MAC address of port 1 (or 1/1).
Port	Lists the port number.
Sys P	Lists the system priority configured for this port.
Port P	Lists the port’s link aggregation priority.
Key	Lists the link aggregation key.
Act	Indicates the link aggregation mode, which can be one of the following: • No – The mode is passive or link aggregation is disabled (off) on the port. If link aggregation is enabled (and the mode is passive), the port can send and receive LACPDU messages to participate in negotiation of an aggregate link initiated by another port, but cannot search for a link aggregation port or initiate negotiation of an aggregate link. • Yes – The mode is active. The port can send and receive LACPDU messages.

Table 7.10: CLI Display of Link Aggregation Information (Continued)

This Field...	Displays...
Tio	Indicates the timeout value of the port. The timeout value can be one of the following: • L – Long. The trunk group has already been formed and the port is therefore using a longer message timeout for the LACPDU messages exchanged with the remote port. Typically, these messages are used as confirmation of the health of the aggregate link. • S – Short. The port has just started the LACPDU message exchange process with the port at the other end of the link. The S timeout value also can mean that the link aggregation information received from the remote port has expired and the ports are starting a new information exchange.
Agg	Indicates the link aggregation state of the port. The state can be one of the following: • Agg – Link aggregation is enabled on the port. • No – Link aggregation is disabled on the port.
Syn	Indicates the synchronization state of the port. The state can be one of the following: • No – The port is out of sync with the remote port. The port does not understand the status of the LACPDU process and is not prepared to enter a trunk link. • Syn – The port is in sync with the remote port. The port understands the status of the LACPDU message exchange process, and therefore knows the trunk group to which it belongs, the link aggregation state of the remote port, and so on.
Col	Indicates the collection state of the port, which determines whether the port is ready to send traffic over the trunk link. • Col – The port is ready to send traffic over the trunk link. • No – The port is not ready to send traffic over the trunk link.
Dis	Indicates the distribution state of the port, which determines whether the port is ready to receive traffic over the trunk link. • Dis – The port is ready to receive traffic over the trunk link. • No – The port is not ready to receive traffic over the trunk link.
Def	Indicates whether the port is using default link aggregation values. The port uses default values if it has not received link aggregation information through LACP from the port at the remote end of the link. This field can have one of the following values: • Def – The port has not received link aggregation values from the port at the other end of the link and is therefore using its default link aggregation LACP settings. • No – The port has received link aggregation information from the port at the other end of the link and is using the settings negotiated with that port.

Table 7.10: CLI Display of Link Aggregation Information (Continued)

This Field...	Displays...
Exp	Indicates whether the negotiated link aggregation settings have expired. The settings expire if the port does not receive an LACPDU message from the port at the other end of the link before the message timer expires. This field can have one of the following values: • Exp – The link aggregation settings this port negotiated with the port at the other end of the link have expired. The port is now using its default link aggregation settings. • No – The link aggregation values that this port negotiated with the port at the other end of the link have not expired, so the port is still using the negotiated settings.
Ope	• Ope (operational) - The port is operating normally. • Ina (inactive) - The port is inactive because the port on the other side of the link is down or has stopped transmitting LACP packets. • Blo (blocked) - The port is blocked because the adjacent port is not configured with link aggregation or because it is not able to join a trunk group. To unblock the port and bring it to an operational state, enable link aggregation on the adjacent port and ensure that the ports have the same key.

Displaying Trunk Group and LACP Status Information

Use the **show trunk** command to determine the status of LACP. See “Displaying Trunk Group Configuration Information” on page 7-22.

Clearing the Negotiated Link Aggregations

When a group of ports negotiates a trunk group configuration, the software stores the negotiated configuration in a table. You can clear the negotiated link aggregation configurations from the software. When you clear the information, the software does not remove link aggregation parameter settings you have configured. Only the configuration information negotiated using LACP is removed.

NOTE: The software automatically updates the link aggregation configuration based on LACPDU messages. However, clearing the link aggregation information can be useful if you are troubleshooting a configuration.

To clear the link aggregation information, use the following CLI method.

USING THE CLI

To clear the link aggregation information, enter the following command at the Privileged EXEC level of the CLI:

```
ProCurveRS# clear link-aggregate
```

Syntax: clear link-aggregate