

Troubleshooting

Contents

Troubleshooting Approaches	C-3
Browser or Telnet Access Problems	C-4
Unusual Network Activity	C-6
General Problems	C-6
802.1Q Prioritization Problems	C-7
IGMP-Related Problems	C-7
LACP-Related Problems	C-8
Port-Based Access Control (802.1X)-Related Problems	C-8
QoS-Related Problems	C-11
Radius-Related Problems	C-12
Spanning-Tree Protocol (MSTP) and Fast-Uplink Problems	C-13
SSH-Related Problems	C-14
Stacking-Related Problems	C-15
TACACS-Related Problems	C-16
TimeP, SNTP, or Gateway Problems	C-18
VLAN-Related Problems	C-18
Fan Failure	C-20
Using the Event Log for Troubleshooting Switch Problems	C-21
Event Log Entries	C-21
Menu: Displaying and Navigating in the Event Log	C-28
CLI: Displaying the Event Log	C-29
CLI: Clearing Event Log Entries	C-30
CLI: Turning Event Numbering On	C-30

Using Log Throttling to Reduce Duplicate Event Log and SNMP Messages	C-30
Log Throttle Periods	C-31
Example of Log Throttling	C-31
Example of Event Counter Operation	C-33
Debug/Syslog Operation	C-34
Debug/Syslog Messaging	C-34
Debug/Syslog Destination Devices	C-34
Debug/Syslog Configuration Commands	C-35
Configuring Debug/Syslog Operation	C-37
Displaying a Debug/Syslog Configuration	C-38
Debug Command	C-40
Debug Messages	C-40
Debug Destinations	C-41
Logging Command	C-42
Configuring a Syslog Server	C-43
Configuring the Severity Level for Event Log Messages Sent to a Syslog Server	C-46
Configuring the System Module Used to Select the Event Log Messages Sent to a Syslog Server	C-47
Operating Notes for Debug and Syslog	C-48
Diagnostic Tools	C-49
Port Auto-Negotiation	C-49
Ping and Link Tests	C-49
Web: Executing Ping or Link Tests	C-51
CLI: Ping or Link Tests	C-52
DNS Resolver	C-54
Terminology	C-54
Basic Operation	C-55
Configuring and Using DNS Resolution with DNS-Compatible Commands	C-56
Configuring a DNS Entry	C-57
Example Using DNS Names with Ping and Traceroute	C-58
Viewing the Current DNS Configuration	C-60
Operating Notes	C-61
Event Log Messages	C-62

Displaying the Configuration File	C-62
CLI: Viewing the Configuration File	C-62
Web: Viewing the Configuration File	C-63
Listing Switch Configuration and Operation Details	C-63
CLI Administrative and Troubleshooting Commands	C-65
Traceroute Command	C-66
Restoring the Factory-Default Configuration	C-69
CLI: Resetting to the Factory-Default Configuration	C-69
Clear/Reset: Resetting to the Factory-Default Configuration .	C-69
Restoring a Flash Image	C-70

Overview

This appendix addresses performance-related network problems that can be caused by topology, switch configuration, and the effects of other devices or their configurations on switch operation. (For switch-specific information on hardware problems indicated by LED behavior, cabling requirements, and other potential hardware-related problems, refer to the *Installation Guide* you received with the switch.)

Note

ProCurve periodically places switch software updates on the ProCurve Networking web site. ProCurve recommends that you check this web site for software updates that may have fixed a problem you are experiencing.

For information on support and warranty provisions, refer to the Support and Warranty booklet shipped with the switch.

Troubleshooting Approaches

Use these approaches to diagnose switch problems:

- Check the ProCurve Networking web site for software updates that may have solved your problem: **www.procurve.com**
- Check the switch LEDs for indications of proper switch operation:
 - Each switch port has a Link LED that should light whenever an active network device is connected to the port.
 - Problems with the switch hardware and software are indicated by flashing the Fault and other switch LEDs.

Refer to the *Installation Guide* shipped with the switch for a description of the LED behavior and information on using the LEDs for troubleshooting.

- Check the network topology/installation. Refer to the *Installation Guide* shipped with the switch for topology information.
- Check cables for damage, correct type, and proper connections. You should also use a cable tester to check your cables for compliance to the relevant IEEE 802.3 specification. Refer to the *Installation Guide* shipped with the switch for correct cable types and connector pin-outs.
- Use ProCurve Manager to help isolate problems and recommend solutions.
- Use the Port Utilization Graph and Alert Log in the web browser interface included in the switch to help isolate problems. Refer to Chapter 5, “Using the ProCurve Web Browser Interface” for operating information. These tools are available through the web browser interface:
 - Port Utilization Graph
 - Alert Log
 - Port Status and Port Counters screens
 - Diagnostic tools (Link test, Ping test, configuration file browser)
- For help in isolating problems, use the easy-to-access switch console built into the switch or Telnet to the switch console. Refer to chapters 3 and 4 for operating information on the Menu and CLI interfaces included in the console. These tools are available through the switch console
 - Status and Counters screens
 - Event Log
 - Diagnostics tools (Link test, Ping test, configuration file browser, and advanced user commands)

Browser or Telnet Access Problems

Cannot access the web browser interface:

- Access may be disabled by the **Web Agent Enabled** parameter in the switch console. Check the setting on this parameter by selecting:

2. Switch Configuration ...

1. System Information

- The switch may not have the correct IP address, subnet mask or gateway. Verify by connecting a console to the switch's Console port and selecting:

2. Switch Configuration ...

5. IP Configuration

Note: If DHCP/Bootp is used to configure the switch, the IP addressing can be verified by selecting:

1. Status and Counters ...

2. Switch Management Address Information

also check the DHCP/Bootp server configuration to verify correct IP addressing.

- If you are using DHCP to acquire the IP address for the switch, the IP address "lease time" may have expired so that the IP address has changed. For more information on how to "reserve" an IP address, refer to the documentation for the DHCP application that you are using.
- If one or more IP-Authorized managers are configured, the switch allows web browser access only to a device having an authorized IP address. For more information on IP Authorized managers, refer to the *Access Security Guide* for your switch.
- Java™ applets may not be running on the web browser. They are required for the switch web browser interface to operate correctly. Refer to the online Help on your web browser for how to run the Java applets.

Cannot Telnet into the switch console from a station on the network:

- Off subnet management stations can lose Telnet access if you enable routing without first configuring a static (default) route. That is, the switch uses the IP default gateway only while operating as a Layer 2 device. While routing is enabled on the switch, the IP default gateway is not used. You can avoid this problem by using the `ip route` command to configure a static (default) route before enabling routing. For more information, refer to the chapter titled “IP Routing Features” in the *Multicast and Routing Guide* for your switch.
- Telnet access may be disabled by the **Inbound Telnet Enabled** parameter in the System Information screen of the menu interface:

2. Switch Configuration

1. System Information

- The switch may not have the correct IP address, subnet mask, or gateway. Verify by connecting a console to the switch’s Console port and selecting:

2. Switch Configuration

5. IP Configuration

Note: If DHCP/Bootp is used to configure the switch, refer to the **Note**, above.

- If you are using DHCP to acquire the IP address for the switch, the IP address “lease time” may have expired so that the IP address has changed. For more information on how to “reserve” an IP address, refer to the documentation for the DHCP application that you are using.
- If one or more IP-Authorized managers are configured, the switch allows inbound telnet access only to a device having an authorized IP address. For more information on IP Authorized managers, refer to the *Access Security Guide* for your switch.

Unusual Network Activity

Network activity that fails to meet accepted norms may indicate a hardware problem with one or more of the network components, possibly including the switch. Such problems can also be caused by a network loop or simply too much traffic for the network as it is currently designed and implemented. Unusual network activity is usually indicated by the LEDs on the front of the switch or measured with the switch console interface or with a network management tool such as ProCurve Manager. Refer to the *Installation Guide* you received with the switch for information on using LEDs to identify unusual network activity.

A topology loop can also cause excessive network activity. The Event Log “FFI” messages can be indicative of this type of problem.

General Problems

The network runs slow; processes fail; users cannot access servers or other devices. Broadcast storms may be occurring in the network. These may be due to redundant links between nodes.

- If you are configuring a port trunk, finish configuring the ports in the trunk before connecting the related cables. Otherwise you may inadvertently create a number of redundant links (i.e. topology loops) that will cause broadcast storms.
- Turn on Spanning Tree Protocol to block redundant links (i.e. topology loops)
- Check for FFI messages in the Event Log.

Duplicate IP Addresses. This is indicated by this Event Log message:

ip: Invalid ARP source: IP address on IP address

where: both instances of *IP address* are the same address, indicating the switch’s IP address has been duplicated somewhere on the network.

Duplicate IP Addresses in a DHCP Network. If you use a DHCP server to assign IP addresses in your network and you find a device with a valid IP address that does not appear to communicate properly with the server or other devices, a duplicate IP address may have been issued by the server. This can occur if a client has not released a DHCP-assigned IP address after the intended expiration time and the server “leases” the address to another device.

This can also happen, for example, if the server is first configured to issue IP addresses with an unlimited duration, then is subsequently configured to issue IP addresses that will expire after a limited duration. One solution is to configure “reservations” in the DHCP server for specific IP addresses to be assigned to devices having specific MAC addresses. For more information, refer to the documentation for the DHCP server.

One indication of a duplicate IP address in a DHCP network is this Event Log message:

```
ip: Invalid ARP source: < IP-address > on <IP-address>
```

where: both instances of *IP-address* are the same address, indicating the IP address that has been duplicated somewhere on the network.

The Switch Has Been Configured for DHCP/Bootp Operation, But Has Not Received a DHCP or Bootp Reply. When the switch is first configured for DHCP/Bootp operation, or if it is rebooted with this configuration, it immediately begins sending request packets on the network. If the switch does not receive a reply to its DHCP/Bootp requests, it continues to periodically send request packets, but with decreasing frequency. Thus, if a DHCP or Bootp server is not available or accessible to the switch when DHCP/Bootp is first configured, the switch may not immediately receive the desired configuration. After verifying that the server has become accessible to the switch, reboot the switch to re-start the process.

802.1Q Prioritization Problems

Ports configured for non-default prioritization (level 1 - 7) are not performing the specified action. If the ports were placed in a trunk group after being configured for non-default prioritization, the priority setting was automatically reset to zero (the default). Ports in a trunk group operate only at the default priority setting.

IGMP-Related Problems

IP Multicast (IGMP) Traffic That Is Directed By IGMP Does Not Reach IGMP Hosts or a Multicast Router Connected to a Port. IGMP must be enabled on the switch and the affected port must be configured for “Auto” or “Forward” operation.

IP Multicast Traffic Floods Out All Ports; IGMP Does Not Appear To Filter Traffic. The IGMP feature does not operate if the switch or VLAN does not have an IP address configured manually or obtained through DHCP/Bootp. To verify whether an IP address is configured for the switch or VLAN, do either of the following:

- **Try Using the Web Browser Interface:** If you can access the web browser interface, then an IP address is configured.
- **Try To Telnet to the Switch Console:** If you can Telnet to the switch, then an IP address is configured.
- **Using the Switch Console Interface:** From the Main Menu, check the Management Address Information screen by clicking on

1. Status and Counters

2. Switch Management Address Information

LACP-Related Problems

Unable to enable LACP on a port with the **interface < port-number > lacp** command. In this case, the switch displays the following message:

Operation is not allowed for a trunked port.

You cannot enable LACP on a port while it is configured as static **Trunk** port. To enable LACP on static-trunked port, first use the **no trunk < port-number >** command to disable the static trunk assignment, then execute **interface < port-number > lacp**.

Caution

Removing a port from a trunk without first disabling the port can create a traffic loop that can slow down or halt your network. Before removing a port from a trunk, ProCurve recommends that you either disable the port or disconnect it from the LAN.

Port-Based Access Control (802.1X)-Related Problems

Note

To list the 802.1X port-access Event Log messages stored on the switch, use **show log 802**.

See also “Radius-Related Problems” on page C-14.

The switch does not receive a response to RADIUS authentication requests. In this case, the switch will attempt authentication using the secondary method configured for the type of access you are using (console, Telnet, or SSH).

There can be several reasons for not receiving a response to an authentication request. Do the following:

- Use **ping** to ensure that the switch has access to the configured RADIUS servers.
- Verify that the switch is using the correct encryption key (RADIUS secret key) for each server.
- Verify that the switch has the correct IP address for each RADIUS server.
- Ensure that the **radius-server timeout** period is long enough for network conditions.

The switch does not authenticate a client even though the RADIUS server is properly configured and providing a response to the authentication request. If the RADIUS server configuration for authenticating the client includes a VLAN assignment, ensure that the VLAN exists as a static VLAN on the switch. Refer to “How 802.1X Authentication Affects VLAN Operation” in the *Access Security Guide* for your switch.

During RADIUS-authenticated client sessions, access to a VLAN on the port used for the client sessions is lost. If the affected VLAN is configured as untagged on the port, it may be temporarily blocked on that port during an 802.1X session. This is because the switch has temporarily assigned another VLAN as untagged on the port to support the client access, as specified in the response from the RADIUS server. Refer to “How 802.1X Authentication Affects VLAN Operation” in the *Access Security Guide* for your switch.

The switch appears to be properly configured as a supplicant, but cannot gain access to the intended authenticator port on the switch to which it is connected. If **aaa authentication port-access** is configured for Local, ensure that you have entered the local *login* (operator-level) username and password of the authenticator switch into the **identity** and **secret** parameters of the supplicant configuration. If instead, you enter the enable (manager-level) username and password, access will be denied.

The supplicant statistics listing shows multiple ports with the same authenticator MAC address. The link to the authenticator may have been moved from one port to another without the supplicant statistics having been cleared from the first port. Refer to “Note on Supplicant Statistics” in the chapter on Port-Based and User-Based Access Control in the *Access Security Guide* for your switch.

The `show port-access authenticator <port-list>` command shows one or more ports remain open after they have been configured with **control unauthorized**. 802.1X is not active on the switch. After you execute `aaa port-access authenticator active`, all ports configured with **control unauthorized** should be listed as **Closed**.

```
ProCurve(config)# show port-access authenticator e A9
Port Access Authenticator Status
Port-access authenticator activated [No] : No
      Access  Authenticator  Authenticator
Port Status Control  State          Backend State
-----
A9  Open  FU          Force Auth  Idle

ProCurve(config)# aaa port-access authenticator active

ProCurve(config)# show port-access authenticator e A9
Port Access Authenticator Status
Port-access authenticator activated [No] : Yes
      Access  Authenticator  Authenticator
Port Status Control  State          Backend State
-----
A9  Closed FU          Force Unauth  Idle
```

Port A9 shows an "Open" status even though Access Control is set to **Unauthorized** (Force Auth). This is because the port-access authenticator has not yet been activated.

Figure C-1. Authenticator Ports Remain "Open" Until Activated

RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch. Use `show radius` to verify that the encryption key (RADIUS secret key) the switch is using is correct for the server being contacted. If the switch has only a global key configured, then it either must match the server key or you must configure a server-specific key. If the switch already has a server-specific key assigned to the server's IP address, then it overrides the global key and must match the server key.

```

10.33.18.119(config)# show radius
Status and Counters - General RADIUS Information
  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key : My-Global-Key

Server IP Addr      Auth  Acct
-----
10.33.18.119      1812 1813 119-only-key
  
```

Figure C-2. Displaying Encryption Keys

Also, ensure that the switch port used to access the RADIUS server is not blocked by an 802.1X configuration on that port. For example, **show port-access authenticator <port-list>** gives you the status for the specified ports. Also, ensure that other factors, such as port security or any 802.1X configuration on the RADIUS server are not blocking the link.

The authorized MAC address on a port that is configured for both 802.1X and port security either changes or is re-acquired after execution of `aaa port-access authenticator <port-list> initialize`. If the port is force-authorized with `aaa port-access authenticator <port-list> control authorized` command and port security is enabled on the port, then executing **initialize** causes the port to clear the learned address and learn a new address from the first packet it receives after you execute **initialize**.

A trunked port configured for 802.1X is blocked. If you are using RADIUS authentication and the RADIUS server specifies a VLAN for the port, the switch allows authentication, but blocks the port. To eliminate this problem, either remove the port from the trunk or reconfigure the RADIUS server to avoid specifying a VLAN.

QoS-Related Problems

Loss of communication when using VLAN-tagged traffic. If you cannot communicate with a device in a tagged VLAN environment, ensure that the device either supports VLAN tagged traffic or is connected to a VLAN port that is configured as **Untagged**.

Radius-Related Problems

The switch does not receive a response to RADIUS authentication requests. In this case, the switch will attempt authentication using the secondary method configured for the type of access you are using (console, Telnet, or SSH).

There can be several reasons for not receiving a response to an authentication request. Do the following:

- Use **ping** to ensure that the switch has access to the configured RADIUS server.
- Verify that the switch is using the correct encryption key for the designated server.
- Verify that the switch has the correct IP address for the RADIUS server.
- Ensure that the **radius-server timeout** period is long enough for network conditions.
- Verify that the switch is using the same UDP port number as the server.

RADIUS server fails to respond to a request for service, even though the server's IP address is correctly configured in the switch. Use **show radius** to verify that the encryption key the switch is using is correct for the server being contacted. If the switch has only a global key configured, then it either must match the server key or you must configure a server-specific key. If the switch already has a server-specific key assigned to the server's IP address, then it overrides the global key and must match the server key.

```
10.33.18.119(config)# show radius
Status and Counters - General RADIUS Information
  Deadtime(min) : 0
  Timeout(secs) : 5
  Retransmit Attempts : 3
  Global Encryption Key : My-Global-Key

  Server IP Addr  Auth  Acct
                Port  Port  Encryption Key
  -----
  10.33.18.119   1812  1813  119-only-key
```

Server IP Addr	Auth Port	Acct Port	Encryption Key
10.33.18.119	1812	1813	119-only-key

Figure C-3. Examples of Global and Unique Encryption Keys

Spanning-Tree Protocol (MSTP) and Fast-Uplink Problems

Caution

If you enable MSTP, it is recommended that you leave the remainder of the MSTP parameter settings at their default values until you have had an opportunity to evaluate MSTP performance in your network. Because incorrect MSTP settings can adversely affect network performance, you should avoid making changes without having a strong understanding of how MSTP operates. To learn the details of MSTP operation, refer to the IEEE 802.1s standard.

Broadcast Storms Appearing in the Network. This can occur when there are physical loops (redundant links) in the topology. Where this exists, you should enable MSTP on all bridging devices in the topology in order for the loop to be detected.

STP Blocks a Link in a VLAN Even Though There Are No Redundant Links in that VLAN. In 802.1Q-compliant switches MSTP blocks redundant physical links even if they are in separate VLANs. A solution is to use only one, multiple-VLAN (tagged) link between the devices. Also, if ports are available, you can improve the bandwidth in this situation by using a port trunk. Refer to “Spanning Tree Operation with VLANs” in the chapter titled “Static Virtual LANs (VLANs)” in the *Advanced Traffic Management Guide* for your switch.

Fast-Uplink Troubleshooting. Some of the problems that can result from incorrect usage of Fast-Uplink MSTP include temporary loops and generation of duplicate packets.

Problem sources can include:

- Fast-Uplink is configured on a switch that is the MSTP root device.
- Either the **Hello Time** or the **Max Age** setting (or both) is too long on one or more switches. Return the **Hello Time** and **Max Age** settings to their default values (2 seconds and 20 seconds, respectively, on a switch).
- A “downlink” port is connected to a switch that is further away (in hop count) from the root device than the switch port on which fast-uplink MSTP is configured.
- Two edge switches are directly linked to each other with a fast-uplink (Mode = **Uplink**) connection.
- Fast uplink is configured on both ends of a link.
- A switch serving as a backup MSTP root switch has ports configured for fast-uplink MSTP and has become the root device due to a failure in the original root device.

SSH-Related Problems

Switch access refused to a client. Even though you have placed the client's public key in a text file and copied the file (using the **copy tftp pub-key-file** command) into the switch, the switch refuses to allow the client to have access. If the source SSH client is an SSHv2 application, the public key may be in the PEM format, which the switch (SSHv1) does not interpret. Check the SSH client application for a utility that can convert the PEM-formatted key into an ASCII-formatted key.

Executing IP SSH does not enable SSH on the switch. The switch does not have a host key. Verify by executing `show ip host-public-key`. If you see the message

```
ssh cannot be enabled until a host key is configured
(use 'crypto' command).
```

then you need to generate an SSH key pair for the switch. To do so, execute **crypto key generate**. (Refer to “2. Generating the Switch's Public and Private Key Pair” in the SSH chapter of the *Access Security Guide* for your switch.)

Switch does not detect a client's public key that does appear in the switch's public key file (show ip client-public-key). The client's public key entry in the public key file may be preceded by another entry that does not terminate with a new line (CR). In this case, the switch interprets the next sequential key entry as simply a comment attached to the preceding key entry. Where a public key file has more than one entry, ensure that all entries terminate with a new line (CR). While this is optional for the last entry in the file, not adding a new line to the last entry creates an error potential if you either add another key to the file at a later time or change the order of the keys in the file.

An attempt to copy a client public-key file into the switch has failed and the switch lists one of the following messages.

```
Download failed: overlength key in key file.
```

```
Download failed: too many keys in key file.
```

```
Download failed: one or more keys is not a valid RSA
public key.
```

The public key file you are trying to download has one of the following problems:

- A key in the file is too long. The maximum key length is 1024 characters, including spaces. This could also mean that two or more keys are merged together instead of being separated by a <CR><LF>.
- There are more than ten public keys in the key file.
- One or more keys in the file is corrupted or is not a valid rsa public key.

Client ceases to respond (“hangs”) during connection phase. The switch does not support data compression in an SSH session. Clients will often have compression turned on by default, but will disable it during the negotiation phase. A client which does not recognize the compression-request FAILURE response may fail when attempting to connect. Ensure that compression is turned off before attempting a connection to prevent this problem.

Stacking-Related Problems

The Stack Commander Cannot Locate any Candidates. Stacking operates on the primary VLAN, which in the default configuration is the DEFAULT_VLAN. However, if another VLAN has been configured as the primary VLAN, and the Commander is not on the primary VLAN, then the Commander will not detect Candidates on the primary VLAN.

TACACS-Related Problems

Event Log. When troubleshooting TACACS+ operation, check the switch's Event Log for indications of problem areas.

All Users Are Locked Out of Access to the Switch. If the switch is functioning properly, but no username/password pairs result in console or Telnet access to the switch, the problem may be due to how the TACACS+ server and/or the switch are configured. Use one of the following methods to recover:

- Access the TACACS+ server application and adjust or remove the configuration parameters controlling access to the switch.
- If the above method does not work, try eliminating configuration changes in the switch that have not been saved to flash (boot-up configuration) by causing the switch to reboot from the boot-up configuration (which includes only the configuration changes made prior to the last **write memory** command.) If you did not use **write memory** to save the authentication configuration to flash, then pressing the Reset button or cycling the power reboots the switch with the boot-up configuration.
- Disconnect the switch from network access to any TACACS+ servers and then log in to the switch using either Telnet or direct console port access. Because the switch cannot access a TACACS+ server, it will default to local authentication. You can then use the switch's local Operator or Manager username/password pair to log on.
- As a last resort, use the Clear/Reset button combination to reset the switch to its factory default boot-up configuration. Taking this step means you will have to reconfigure the switch to return it to operation in your network.

No Communication Between the Switch and the TACACS+ Server Application. If the switch can access the server device (that is, it can **ping** the server), then a configuration error may be the problem. Some possibilities include:

- The server IP address configured with the switch's **tacacs-server host** command may not be correct. (Use the switch's **show tacacs-server** command to list the TACACS+ server IP address.)

- The encryption key configured in the server does not match the encryption key configured in the switch (by using the **tacacs-server key** command). Verify the key in the server and compare it to the key configured in the switch. (Use **show tacacs-server** to list the global key. Use **show config** or **show config running** to list any server-specific keys.)
- The accessible TACACS+ servers are not configured to provide service to the switch.

Access Is Denied Even Though the Username/Password Pair Is Correct. Some reasons for denial include the following parameters controlled by your TACACS+ server application:

- The account has expired.
- The access attempt is through a port that is not allowed for the account.
- The time quota for the account has been exhausted.
- The time credit for the account has expired.
- The access attempt is outside of the time frame allowed for the account.
- The allowed number of concurrent logins for the account has been exceeded

For more help, refer to the documentation provided with your TACACS+ server application.

Unknown Users Allowed to Login to the Switch. Your TACACS+ application may be configured to allow access to unknown users by assigning them the privileges included in a *default user* profile. Refer to the documentation provided with your TACACS+ server application.

System Allows Fewer Login Attempts than Specified in the Switch Configuration. Your TACACS+ server application may be configured to allow fewer login attempts than you have configured in the switch with the **aaa authentication num-attempts** command.

TimeP, SNTP, or Gateway Problems

The Switch Cannot Find the Time Server or the Configured Gateway .

TimeP, SNTP, and Gateway access are through the primary VLAN, which in the default configuration is the DEFAULT_VLAN. If the primary VLAN has been moved to another VLAN, it may be disabled or does not have ports assigned to it.

VLAN-Related Problems

Monitor Port. When using the monitor port in a multiple VLAN environment, the switch handles broadcast, multicast, and unicast traffic output from the monitor port as follows:

- If the monitor port is configured for tagged VLAN operation on the same VLAN as the traffic from monitored ports, the traffic output from the monitor port carries the same VLAN tag.
- If the monitor port is configured for untagged VLAN operation on the same VLAN as the traffic from the monitored ports, the traffic output from the monitor port is untagged.
- If the monitor port is not a member of the same VLAN as the traffic from the monitored ports, traffic from the monitored ports does not go out the monitor port.

None of the devices assigned to one or more VLANs on an 802.1Q-compliant switch are being recognized. If multiple VLANs are being used on ports connecting 802.1Q-compliant devices, inconsistent VLAN IDs may have been assigned to one or more VLANs. For a given VLAN, the same VLAN ID must be used on all connected 802.1Q-compliant devices.

Link Configured for Multiple VLANs Does Not Support Traffic for One or More VLANs. One or more VLANs may not be properly configured as “Tagged” or “Untagged”. A VLAN assigned to a port connecting two 802.1Q-compliant devices must be configured the same on both ports. For example, VLAN_1 and VLAN_2 use the same link between switch “X” and switch “Y”.

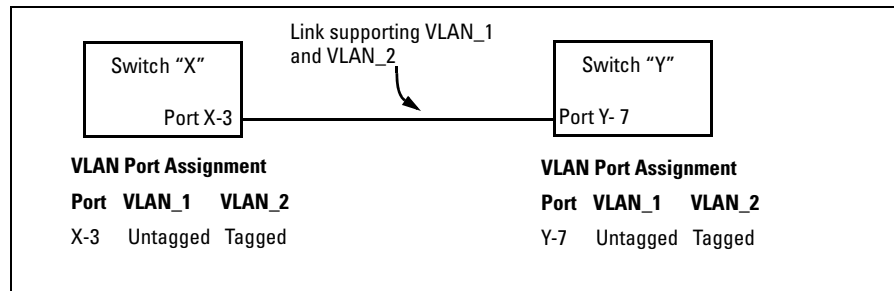


Figure C-4. Example of Correct VLAN Port Assignments on a Link

1. If VLAN_1 (VID=1) is configured as "Untagged" on port 3 on switch "X", then it must also be configured as "Untagged" on port 7 on switch "Y". Make sure that the VLAN ID (VID) is the same on both switches.
2. Similarly, if VLAN_2 (VID=2) is configured as "Tagged on the link port on switch "A", then it must also be configured as "Tagged" on the link port on switch "B". Make sure that the VLAN ID (VID) is the same on both switches.

Duplicate MAC Addresses Across VLANs. The switches covered in this guide operate with multiple forwarding databases. Thus, duplicate MAC addresses occurring on different VLANs can appear where a device having one MAC address is a member of more than one 802.1Q VLAN, and the switch port to which the device is linked is using VLANs (instead of MSTP or trunking) to establish redundant links to another switch. If the other device sends traffic over multiple VLANs, its MAC address will consistently appear in multiple VLANs on the switch port to which it is linked.

Note that attempting to create redundant paths through the use of VLANs will cause problems with some switches. One symptom is that a duplicate MAC address appears in the Port Address Table of one port, and then later appears on another port. While the switches have multiple forwarding databases, and thus does not have this problem, some switches with a single forwarding database for all VLANs may produce the impression that a connected device is moving among ports because packets with the same MAC address but different VLANs are received on different ports. You can avoid this problem by creating redundant paths using port trunks or spanning tree.

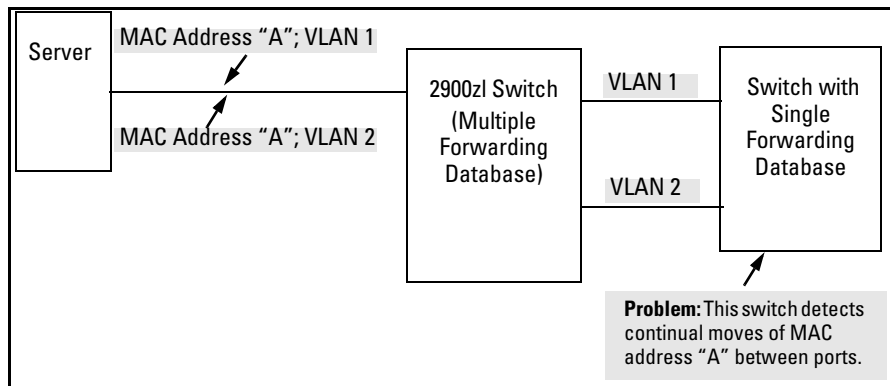


Figure C-5. Example of Duplicate MAC Address

Fan Failure

When two or more fans fail, a two-minute timer starts. After two minutes, the switch is powered down and must be rebooted to restart it. This protects the switch from possible overheating.

ProCurve recommends that you replace a failed fan tray assembly within one minute of removing it.

Using the Event Log for Troubleshooting Switch Problems

The Event Log records operating events in single- or double-line entries and serves as a tool to isolate and troubleshoot problems.

Starting in software release T.13.xxx, the maximum number of entries supported in the Event Log is increased from 1000 to 2000 entries. Entries are listed in chronological order, from the oldest to the most recent.

Once the log has received 2000 entries, it discards the oldest message each time a new message is received. The Event Log window contains 14 log entry lines. You can scroll through it to view any part of the log.

Note

The Event Log is *erased* if power to the switch is interrupted or if you enter the **boot system** command. The contents of the Event Log are *not* erased if you:

- Reboot the switch by choosing the **Reboot Switch** option from the menu interface.
- Enter the **reload** command from the CLI.

Event Log Entries

As shown in Figure C-6, each Event Log entry is composed of five or six fields, depending on whether numbering is turned on or not:

Severity	Date	Time	Event number	System Module	Event Message
I	08/05/06	10:52:32	00063	ports:	port A1 enabled

Figure C-6. Format of an Event Log Entry

Severity is one of the following codes (from highest to lowest severity):

- M** (major) indicates that a fatal switch error has occurred.
- E** (error) indicates that an error condition occurred on the switch.
- W** (warning) indicates that a switch service has behaved unexpectedly.
- I** (information) provides information on normal switch operation.

Troubleshooting

Using the Event Log for Troubleshooting Switch Problems

D (debug) is reserved for ProCurve internal diagnostic information.

Date is the date in the format *mm/dd/yy* when an entry is recorded in the log.

Time is the time in the format *hh:mm:ss* when an entry is recorded in the log.

Event Number is the number assigned to an event. You can turn event numbering on and off with the **[no] log-number** command.

System Module is the internal module (such as “ports:” for port manager) that generated a log entry. If VLANs are configured, then a VLAN name also appears for an event that is specific to an individual VLAN. Table C-1 lists the different system modules with a description of each one.

Event Message is a brief description of the operating event.

Table C-1. Event Log System Modules

System Module	Description	Documented in ProCurve Hardware/Software guide
802.1x	802.1X authentication: Provides access control on a per-client or per-port basis: <ul style="list-style-type: none">• Client-level security that allows LAN access to 802.1X clients (up to 32 per port) with valid user credentials• Port-level security that allows LAN access only on ports on which a single 802.1X-capable client (supplicant) has entered valid RADIUS user credentials	<i>Access Security Guide</i>
addrmgr	Address Table Manager: Manages MAC addresses that the switch has learned and are stored in the switch's address table.	<i>Management and Configuration Guide</i>
auth	Authorization: A connected client must receive authorization through web, AMC, RADIUS-based, TACACS+-based, or 802.1X authentication before it can send traffic to the switch.	<i>Access Security Guide</i>
cdp	Cisco Discovery Protocol: Supports reading CDP packets received from neighbor devices, enabling a switch to learn about adjacent CDP devices. ProCurve switches do not support the transmission of CDP packets to neighbor devices.	<i>Management and Configuration Guide</i>
chassis	Hardware operation, including modules and ports, power supply, fans, transceivers, CPU interrupt errors, switch temperature, and so on. Chassis messages include events on Power Over Ethernet (POE) operation.	<i>Installation Guides</i> <i>Management and Configuration Guide</i>

System Module	Description	Documented in ProCurve Hardware/ Software guide
connfilt	<p>Connection-Rate filtering: Used on the network edge to protect the network from attack by worm-like malicious code by detecting hosts that are generating IP traffic that exhibits this behavior and (optionally) either throttling or dropping all IP traffic from the offending hosts.</p> <p>Connection-Rate filtering messages include events on virus throttling. Virus throttling uses connection-rate filtering to stop the propagation of malicious agents.</p>	<i>Access Security Guide</i>
console	Console interface used to monitor switch and port status, reconfigure the switch, read the event log through an in-band Telnet or out-of-band connection.	<i>Installation and Getting Started Guide</i>
cos	<p>Class of Service (CoS): Provides priority handling of packets traversing the switch, based on the IEEE 802.1p priority carried by each packet.</p> <p>CoS messages also include Quality of Service (QoS) events. The QoS feature classifies and prioritizes traffic throughout a network, establishing an end-to-end traffic priority policy to manage available bandwidth and improve throughput of important data.</p>	<i>Advanced Traffic Management Guide</i>
dca	Dynamic Configuration Arbiter (DCA) determines the client-specific parameters that are assigned in an authentication session.	<i>Access Security Guide</i>
dhcp	Dynamic Host Configuration Protocol (DHCP) server configuration: Switch is automatically configured from a DHCP (Bootp) server, including IP address, subnet mask, default gateway, Timep Server address, and TFTP server address.	<i>Management and Configuration Guide</i>
dhcp v6c	DHCP for IPv6 prefix assignment	<i>IPv6 Management Guide</i>
dhcpr	DHCP relay: Forwards client-originated DHCP packets to a DHCP network server.	<i>Advanced Traffic Management Guide</i>
download	Download operation for copying a software version or files to the switch.	<i>Management and Configuration Guide</i>
dma	Direct Access Memory (DMA): Transmits and receives packets between the CPU and the switch. Not used for logging messages in software release T.13.xx.	—
fault	Fault Detection facility, including response policy and the sensitivity level at which a network problem should generate an alert.	<i>Management and Configuration Guide</i>
ffi	Find, Fix, and Inform: Event or alert log messages indicating a possible topology loop that cause excessive network activity and results in the network running slow. FFI messages include events on transceiver connections with other network devices.	<i>Installation and Getting Started Guide</i> <i>Management and Configuration Guide</i>

Troubleshooting

Using the Event Log for Troubleshooting Switch Problems

System Module	Description	Documented in ProCurve Hardware/ Software guide
garp	Generic Attribute Registration Protocol (GARP), defined in the IEEE 802.1D-1998 standard.	<i>Advanced Traffic Management Guide</i>
gvrp	GARP VLAN Registration Protocol (GVRP): Manages dynamic 802.1Q VLAN operations, in which the switch creates temporary VLAN membership on a port to provide a link to another port in the same VLAN on another device.	<i>Advanced Traffic Management Guide</i>
hpesp	Management module that maintains communication between switch ports.	<i>Installation and Getting Started Guide</i>
idm	Identity-driven Management: Optional management application used to monitor and control access to switch.	<i>Advanced Traffic Management Guide</i>
igmp	Internet Group Management Protocol: Reduces unnecessary bandwidth usage for multicast traffic transmitted from multimedia applications on a per-port basis.	<i>Multicast and Routing Guide</i>
ip	IP addressing: Configures the switch with an IP address and subnet mask to communicate on the network and support remote management access; configures multiple IP addresses on a VLAN; enables IP routing on the switch.	<i>Management and Configuration Guide</i> <i>Multicast and Routing Guide</i>
ipaddrmgr	IP Address Manager: Programs IP routing information in switch hardware.	<i>Multicast and Routing Guide</i>
iplock	IP Lockdown: Prevents IP source address spoofing on a per-port and per-VLAN basis by forwarding only the IP packets in VLAN traffic that contain a known source IP address and MAC address binding for the port.	<i>Access Security Guide</i>
ipx	Novell Netware protocol filtering: On the basis of protocol type, the switch can forward or drop traffic to a specific set of destination ports on the switch.	<i>Access Security Guide</i>
kms	Key Management System: Configures and maintains security information (keys) for all routing protocols, including a timing mechanism for activating and deactivating an individual protocol.	<i>Access Security Guide</i>
lACP	LACP trunks: The switch can either automatically establish an 802.3ad-compliant trunk group or provide a manually configured, static LACP trunk.	<i>Management and Configuration Guide</i>
ldbal	Load balancing in LACP port trunks or 802.1s Multiple Spanning Tree protocol (MSTP) that uses VLANs in a network to improve network resource utilization and maintain a loop-free environment.	<i>Management and Configuration Guide</i> <i>Advanced Traffic Management Guide</i>
lldp	Link-Layer Discovery Protocol: Supports transmitting LLDP packets to neighbor devices and reading LLDP packets received from neighbor devices, enabling a switch to advertise itself to adjacent devices and to learn about adjacent LLDP devices.	<i>Management and Configuration Guide</i>

System Module	Description	Documented in ProCurve Hardware/ Software guide
loop_protect	Loop protection: Detects the formation of loops when an unmanaged device on the network drops spanning tree packets, and provides protection by transmitting loop protocol packets out ports on which loop protection has been enabled.	<i>Advanced Traffic Management Guide</i>
macauth	Web and MAC authentication: Port-based security employed on the network edge to protect private networks and the switch itself from unauthorized access using one of the following interfaces: <ul style="list-style-type: none"> • Web page login to authenticate users for access to the network • RADIUS server that uses a device’s MAC address for authentication 	<i>Access Security Guide</i>
maclock	MAC lockdown and MAC lockout <ul style="list-style-type: none"> • MAC lockdown prevents station movement and MAC address “hijacking” by requiring a MAC address to be used only an assigned port on the switch. MAC Lockdown also restricts the client device to a specific VLAN. • MAC lockout blocks a specific MAC address so that the switch drops all traffic to or from the specified address. 	<i>Access Security Guide</i>
mgr	ProCurve Manager (PCM) and ProCurve Manager Plus (PCM+): Windows-based network management solutions for managing and monitoring performance of ProCurve devices. PCM messages also include events for configuration operations.	<i>Management and Configuration Guide</i>
mld	Multicast Listener Discovery (MLD): IPv6 protocol used by a router to discover the presence of multicast listeners. MLD can also optimize IPv6 multicast traffic flow with the snooping feature.	<i>Multicast and Routing Guide</i>
mtm	Multicast Traffic Manager (MTM): Controls and coordinates L3 multicast traffic for upper layer protocols.	<i>Multicast and Routing Guide</i>
netinet	Network Internet: Monitors the creation of a route or an Address Resolution Protocol (ARP) entry and sends a log message in case of failure.	<i>Advanced Traffic Management Guide</i>
pagp	Ports Aggregation Protocol (PAgP): Obsolete. Replaced by LACP (802.3ad). Not used for logging messages in software release T.13.xx.	—
ports	Port status and port configuration features, including mode (speed and duplex), flow control, broadcast limit, jumbo packets, and security settings. Port messages include events on Power Over Ethernet (POE) operation and transceiver connections with other network devices.	<i>Installation and Getting Started Guide</i> <i>Management and Configuration Guide</i> <i>Access Security Guide</i>

Troubleshooting

Using the Event Log for Troubleshooting Switch Problems

System Module	Description	Documented in ProCurve Hardware/ Software guide
radius	RADIUS (Remote Authentication Dial-In User Service) authentication and accounting: A network server is used to authenticate user-connection requests on the switch and collect accounting information to track network resource usage.	<i>Access Security Guide</i>
sflow	Flow sampling: sFlow is an industry standard sampling technology, defined by RFC 3176, used to continuously monitor traffic flows on all ports providing network-wide visibility into the use of the network.	<i>Management and Configuration Guide</i>
snmp	Simple Network Management Protocol: Allows you to manage the switch from a network management station, including support for security features, event reporting, flow sampling, and standard MIBs.	<i>Management and Configuration Guide</i>
sntp	Simple Network Time Protocol: Synchronizes and ensures a uniform time among interoperating devices.	<i>Management and Configuration Guide</i>
ssh	Secure Shell version 2 (SSHv2): Provides remote access to management functions on a switch via encrypted paths between the switch and management station clients capable of SSH operation. SSH messages also include events from the Secure File Transfer Protocol (SFTP) feature. SFTP provides a secure alternative to TFTP for transferring sensitive information, such as switch configuration files, to and from the switch in an SSH session.	<i>Access Security Guide</i>
ssl	Secure Socket Layer Version 3 (SSLv3), including Transport Layer Security (TLSv1) support: Provides remote web access to a switch via encrypted paths between the switch and management station clients capable of SSL/TLS operation.	<i>Access Security Guide</i>
stack	Stack management: Uses a single IP address and standard network cabling to manage a group (up to 16) of switches in the same IP subnet (broadcast domain), resulting in a reduced number of IP addresses and simplified management of small workgroups for scaling your network to handle increased bandwidth demand.	<i>Advanced Traffic Management Guide</i>
stp	Multiple-instance spanning tree protocol/MSTP (802.1s): Ensures that only one active path exists between any two nodes in a group of VLANs in the network. MSTP operation is designed to avoid loops and broadcast storms of duplicate messages that can bring down the network.	<i>Advanced Traffic Management Guide</i>

System Module	Description	Documented in ProCurve Hardware/ Software guide
system	<p>Switch management, including system configuration, switch bootup, activation of boot ROM image, memory buffers, traffic and security filters.</p> <p>System messages also include events from Management interfaces (menu, CLI, web browser, ProCurve Manager) used to reconfigure the switch and monitor switch status and performance.</p>	<i>Management and Configuration Guide Access Security Guide</i>
tacacs	TACACS+ authentication: A central server is used to control access to the switches (and other TACACS-aware devices) in the network through a switch's console port (local access) or Telnet (remote access).	<i>Access Security Guide</i>
tcp	Transmission Control Protocol: A transport protocol that runs on IP and is used to set up connections.	<i>Advanced Traffic Management Guide</i>
telnet	Session established on the switch from a remote device through the Telnet virtual terminal protocol.	<i>Management and Configuration Guide</i>
tftp	Trivial File Transfer Protocol: Supports the download of files to the switch from a TFTP network server.	<i>Management and Configuration Guide</i>
timep	Time Protocol: Synchronizes and ensures a uniform time among interoperating devices.	<i>Management and Configuration Guide</i>
udld	Uni-directional Link Detection: Monitors a link between two switches and blocks the ports on both ends of the link if the link fails at any point between the two devices.	<i>Access Security Guide</i>
udpf	UDP broadcast forwarding: Supports the forwarding of client requests sent as limited IP broadcasts addressed to a UDP application port on a network server.	<i>Multicast and Routing Guide</i>
update	Updates (TFTP or serial) to ProCurve software and updates to running-config and start-up config files	<i>Management and Configuration Guide</i>
usb	Auxiliary port that allows you to connect external devices to the switch.	<i>Installation and Getting Started Guide</i>
vlan	<p>Static 802.1Q VLAN operations, including port-and protocol-based configurations that group users by logical function instead of physical location</p> <ul style="list-style-type: none"> • A port -based VLAN creates a layer-2 broadcast domain comprised of member ports that bridge IPv4 traffic among themselves. • A protocol-based VLAN creates a layer-3 broadcast domain for traffic of a particular routing protocol, and is comprised of member ports that bridge traffic of the specified protocol type among themselves. <p>VLAN messages include events from Management interfaces (menu, CLI, web browser, ProCurve Manager) used to reconfigure the switch and monitor switch status and performance.</p>	<i>Advanced Traffic Management Guide</i>

Troubleshooting

Using the Event Log for Troubleshooting Switch Problems

System Module	Description	Documented in ProCurve Hardware/ Software guide
wsm	Wireless Edge Services Module: Operation of the Wireless Services application on an installed Wireless Edge Services Module. Messages contain the slot ID in the format: "wsm <slot-letter>"; for example, "wsm A:" for slot A.	<i>Wireless Edge Module Installation and Configuration Guide</i>
xmodem	Xmodem: Binary transfer feature that supports the download of software files from a PC or Unix workstation.	<i>Management and Configuration Guide</i>

Menu: Displaying and Navigating in the Event Log

To display the Event Log from the Main Menu, select **Event Log**. Figure C-7 shows a sample event log display.

```
ProCurve Switch 5406z1                               25-Oct-2007  18:02:52
=====--CONSOLE - MANAGER MODE -----
M 10/25/07 16:30:02 sys: 'Operator cold reboot from CONSOLE session.'
I 10/25/07 17:42:51 00061 system: -----
I 10/25/07 17:42:51 00063 system: System went down: 10/25/07 16:30:02
I 10/25/07 17:42:51 00064 system: Operator cold reboot from CONSOLE session.
W 10/25/07 17:42:51 00374 chassis: WARNING: SSC is out of Date: Load 8.2 or newer
I 10/25/07 17:42:51 00068 chassis: Slot D Inserted
I 10/25/07 17:42:51 00068 chassis: Slot E Inserted
I 10/25/07 17:42:51 00068 chassis: Slot F Inserted
I 10/25/07 17:42:51 00690 udpf: DHCP relay agent feature enabled
I 10/25/07 17:42:51 00433 ssh: Ssh server enabled
I 10/25/07 17:42:52 00400 stack: Stack Protocol disabled
I 10/25/07 17:42:52 00128 tftp: Enable succeeded
I 10/25/07 17:42:52 00417 cdp: CDP enabled

----  Log events stored in memory 1-751.  Log events on screen 690-704.

Actions->   Back      Next page      Prev page      End      Help

Return to previous screen.
Use up/down arrow to scroll one line, left/right arrow keys to
change action selection, and <Enter> to execute action.
```

Figure C-7. Example of an Event Log Display

The *log status line* below the recorded entries states the total number of events stored in the event log and which logged events are currently displayed.

To scroll to other entries in the Event Log, either preceding or following the currently visible portion, press the keys indicated at the bottom of the display (**Back**, **Next page**, **Prev page**, or **End**) or the keys described in Table C-2.

Table C-2. Event Log Control Keys

Key	Action
[N]	Advances the display by one page (next page).
[P]	Rolls back the display by one page (previous page).
[v]	Advances display by one event (down one line).
[^]	Rolls back display by one event (up one line).
[E]	Advances to the end of the log.
[H]	Displays Help for the Event Log.

CLI: Displaying the Event Log

To display messages recorded in the event log from the CLI, enter the **show logging** command. Keyword searches are supported.

Syntax: show logging [-a, -r] [<search-text>]

*By default, the **show logging** command displays the log messages recorded since the last reboot in chronological order.*

***-a** displays all recorded log messages, including those before the last reboot.*

***-r** displays all recorded log messages, with the most recent entries listed first.*

*<search-text> displays all Event Log entries that contain the specified text. Use a <search-text> value with **-a** or **-r** to further filter **show logging** command output.*

Examples. To display all Event Log messages that have “system” in the message text or module name, enter the following command:

```
ProCurve# show logging -a system
```

To display all Event Log messages recorded since the last reboot that have the word, “system”, in the message text or module name, enter:

```
ProCurve# show logging system
```

CLI: Clearing Event Log Entries

Use the **clear logging** command to hide, but not erase, Event Log entries displayed in **show logging** command output. Only new entries generated after you enter the command will be displayed.

To redisplay all hidden entries, including Event Log entries recorded prior to the last reboot, enter the **show logging -a** command.

Syntax: clear logging

Removes all entries from the event log display output.

CLI: Turning Event Numbering On

Syntax: [no] log-number

Turns event numbering on and off

Using Log Throttling to Reduce Duplicate Event Log and SNMP Messages

A recurring event can generate a series of duplicate Event Log messages and SNMP traps in a relatively short time. As a result, the Event Log and any configured SNMP trap receivers may be flooded with excessive, exactly identical messages. To help reduce this problem, the switch uses *log throttle periods* to regulate (throttle) duplicate messages for recurring events, and maintains a counter to record how many times it detects duplicates of a particular event since the last system reboot.

When the first instance of a particular event or condition generates a message, the switch initiates a log throttle period that applies to all recurrences of that event. If the logged event recurs during the log throttle period, the switch increments the counter initiated by the first instance of the event, but does not generate a new message.

If the logged event repeats again after the log throttle period expires, the switch generates a duplicate of the first message, increments the counter, and starts a new log throttle period during which any additional instances of the event are counted, but not logged. Thus, for a particular recurring event, the switch displays only one message in the Event Log for each log throttle period in which the event reoccurs. Also, each logged instance of the event message

includes counter data showing how many times the event has occurred since the last reboot. The switch manages messages to SNMP trap receivers in the same way.

Log Throttle Periods

The length of the log throttle period differs according to an event's severity level:

Severity Level	Log Throttle Period
I (Information)	6000 Seconds
W (Warning)	600 Seconds
D (Debug)	60 Seconds
M (Major)	6 Seconds

Example of Log Throttling

For example, suppose that you configure VLAN 100 on the switch to support an operation, but do not configure an IP address. If the operation attempted to use VLAN 100, the switch would generate the first instance of the following Event Log message and counter.

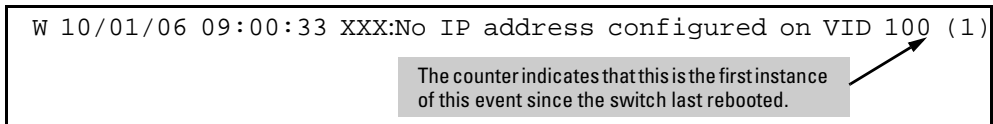


Figure C-8. Example of the First Instance of an Event Message and Counter

Troubleshooting

Using the Event Log for Troubleshooting Switch Problems

If the operation caused the same event to occur six more times during the initial log throttle period, there would be no further entries in the Event Log. However, if the event occurred again after the log throttle period expired, the switch would repeat the message (with an updated counter) and start a new log throttle period.

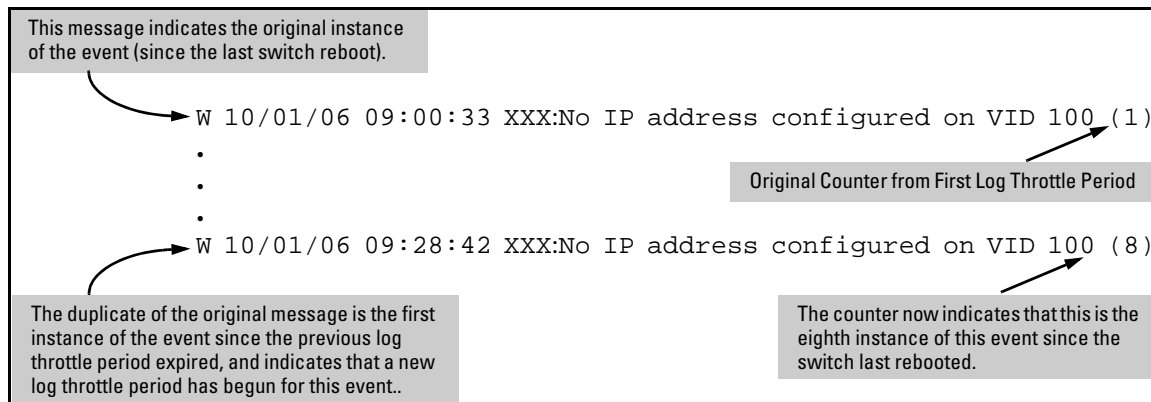


Figure C-9. Example of Duplicate Messages Over Multiple Log Throttling Periods

Note that if the same type of event occurs under different circumstances, the switch handles these as unrelated events for the purpose of Event Log messages. For example, if the operation simultaneously detected that VLANs 100 and 205 were configured without IP addresses, you would see log messages similar to the following:

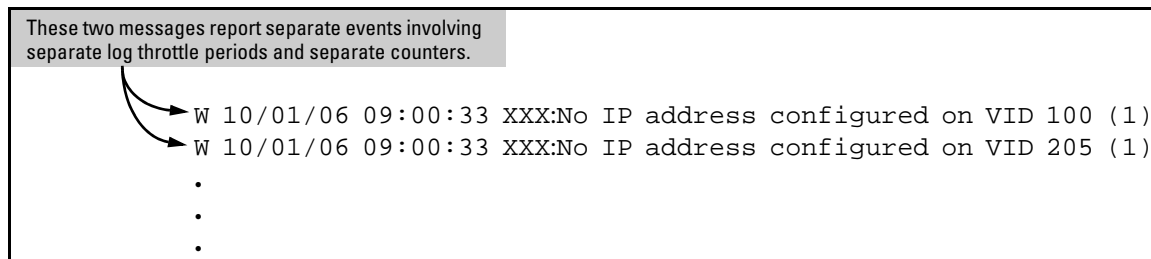


Figure C-10. Example of Log Messages Generated by Unrelated Events of the Same Type

Example of Event Counter Operation

Suppose the switch detects the following after a reboot:

- Three duplicate instances of the “Send error” during the first log throttle period for this event
- Five more instances of the same Send error during the second log throttle period for this event
- Four instances of the same Send error during the third log throttle period for this event

In this case, the duplicate message would appear three times in the Event Log (once for each log throttle period for the event being described), and the Duplicate Message Counter would increment as shown in table C-3. (The same operation would apply for messages sent to any configured SNMP trap receivers.)

Table C-3. How the Duplicate Message Counter Increments

Instances During 1st Log Throttle Period	Instances During 2nd Log Throttle Period	Instances During 3rd Log Throttle Period	Duplicate Message Counter*
3			1
	5		4
		4	9

*This value always comprises the first instance of the duplicate message in the current log throttle period plus all previous occurrences of the duplicate message occurring since the switch last rebooted.

Debug/Syslog Operation

While the Event Log records switch-level progress, status, and warning messages on the switch, the Debug/System Logging (*Syslog*) feature provides a way to record Event Log and debug messages on a remote device. For example, you can send messages about routing misconfigurations and other network protocol details to an external device, and later use them to debug network-level problems.

Debug/Syslog Messaging

The Debug/Syslog feature allows you to specify the types of Event Log and debug messages that you want to send to an external device. As shown in Figure C-11, you can perform the following operations:

- Use the **debug** command to configure messaging reports for the following event types:
 - Events recorded in the switch's Event Log
 - IP routing events
 - LLDP events
- Use the **logging** command to select a subset of Event Log messages to send to an external device for debugging purposes according to:
 - Severity level
 - System module

Debug/Syslog Destination Devices

To use Debug/Syslog messaging, you must configure an external device as the logging destination by using the **logging** and **debug destination** commands. For more information, see “Debug Destinations” on page C-43 and “Configuring a Syslog Server” on page C-45.

A Debug/Syslog destination device can be a Syslog server and/or a console session. You can configure debug and logging messages to be sent to:

- Up to six Syslog servers
- A CLI session through a direct RS-232 console connection, or a Telnet or SSH session

Debug/Syslog Configuration Commands

Event Notification Logging	—	Automatically sends switch-level event messages to the switch's Event Log. Debug and Syslog do not affect this operation, but add the capability of directing Event Log messaging to an external device.
logging Command	<i><syslog-ip-addr></i>	Enables Syslog messaging to be sent to the specified IP address.
	facility	(Optional) The logging facility command specifies the destination (facility) subsystem used on a Syslog server for debug reports.
debug Command	destination	logging Disables or re-enables Syslog logging on one or more Syslog servers configured with the logging < syslog-ip-addr > command.
	destination	session Assigns or re-assigns destination status to the terminal device that was most recently used to request debug output.
debug Command	all	Sends debug logging to configured debug destinations for all Event Log, and IP-RIP options.
	event	Sends standard Event Log messages to configured debug destinations. (The same messages are also sent to the switch's Event Log, regardless of whether you enable this option.)
	ip	rip Sends RIP event logging to the debug destination(s).
	lldp	Sends LLDP debug logging to the debug destination(s).
logging Command	severity	Sends Event Log messages of equal or greater severity than the specified value to configured debug destinations. (The default setting is to send Event Log messages from all severity levels.)
	system-module	Sends Event Log messages from the specified system module to configured debug destinations. The severity filter is also applied to the system-module messages you select. The default setting is to send Event Log messages from all system modules. To restore the default setting, enter the no logging system-module <system-module> or logging system-module all-pass commands.

Figure C-11. Summary of Debug/Syslog Configuration Commands

Using the Debug/Syslog feature, you can perform the following operations:

- Configure the switch to send Event Log messages to one or more Syslog servers. In addition, you can configure the messages to be sent to the User log facility (default) or to another log facility on configured Syslog servers.

Note

As of January 2008, the **logging facility** < *facility-name* > option (described on page C-47) is supported on the following switch models:

- 8212zl Switch
- Series 6400cl switches
- 6200yl Switch
- Series 5400zl switches
- Series 5300xl switches
- Series 4200vl switches
- Series 4100gl switches (software release G.07.50 or greater)
- Series 3500yl switches
- Series 3400cl switches
- Series 2900 switches
- Series 2800 switches
- Series 2600 switches and the Switch 6108 (software release H.07.30 or greater)

For the latest feature information on ProCurve switches, visit the ProCurve Networking web site and check the latest release notes for the switch products you use.

-
- Configure the switch to send Event Log messages to the current management-access session (serial-connect CLI, Telnet CLI, or SSH).
 - Disable all Syslog debug logging while retaining the Syslog addresses from the switch configuration. This allows you to configure Syslog messaging and then disable and re-enable it as needed.
 - Display the current debug configuration. If Syslog logging is currently active, the list of configured Syslog servers is displayed.
 - Display the current Syslog server list when Syslog logging is disabled.

Configuring Debug/Syslog Operation

1. To use a Syslog server as the destination device for debug messaging, follow these steps:
 - a. Enter the **logging** *< syslog-ip-addr >* command at the global configuration level to configure the Syslog server IP address and enable Syslog logging. Optionally, you may also specify the destination subsystem to be used on the Syslog server by entering the **logging facility** command.

If no other Syslog server IP addresses are configured, entering the **logging** command enables both debug messaging to a Syslog server and the Event debug message type. As a result, the switch automatically sends Event Log messages to the Syslog server, regardless of other debug types that may be configured.
 - b. Re-enter the **logging** command in Step “a” to configure additional Syslog servers. You can configure up to a total of six servers. (When multiple server IP addresses are configured, the switch sends the debug message types that you configure in Step 3 to all IP addresses.)

2. To use a CLI session on a destination device for debug messaging:
 - a. Set up a serial, Telnet, or SSH connection to access the switch’s CLI.
 - b. Enter the **debug destination session** command at the manager level.
3. Enable the types of debug messages to be sent to configured Syslog servers and/or the current session device by entering the **debug** *< debug-type >* command:

```
ProCurve# debug < all | event | lldp >
```

Repeat this step if necessary to enable multiple debug message types.

By default, Event Log messages are sent to configured debug destination devices. To block Event Log messages from being sent, enter the **no debug event** command.

4. If necessary, enable a subset of Event Log messages to be sent to configured Syslog servers by specifying a severity level and/or system module using the following commands

```
ProCurve(config)# logging severity < debug | major | error | warning | info >  
ProCurve(config)# logging system-module < system-module >
```

To display a list of valid values for each command, enter **logging severity** or **logging system-module** followed by **?** or pressing the **Tab** key.

The severity levels in order from the highest to lowest severity are: major, error, warning, info, debug. For a list of valid values for the **logging system-module** `<system-module>` command, refer to Table C-1 on page C-24.

5. If you configure system-module and/or severity-level values to filter Event Log messages, when you finish troubleshooting, you may want to reset these values to their default settings so that the switch sends all Event Log messages to configured debug destinations (Syslog servers and/or CLI session).

To remove a configured setting and restore the default values that send all Event Log messages, enter one or both of the following commands:

```
ProCurve(config)# no logging severity <debug|major|error|warning|info>  
ProCurve(config)# no logging system-module <system-module >
```

Caution

If you configure a severity-level, system-module, logging destination, or logging facility value and save the settings to the startup configuration (for example, by entering the **write memory** command), the debug settings are saved after a system reboot (power cycle or reboot) and re-activated on the switch. As a result, after switch startup, one of the following situations may occur:

- Only a partial set of Event Log messages may be sent to configured debug destinations.
- Messages may be sent to a previously configured Syslog server used in an earlier debugging session.

Displaying a Debug/Syslog Configuration

Use the **show debug** command to display the currently configured settings for:

- Debug message types and Event Log message filters (severity level and system module) sent to debug destinations
- Debug destinations (Syslog servers or CLI session) and Syslog server facility to be used

Syntax: show debug

*Displays the currently configured debug logging destinations and message types selected for debugging purposes. (If no Syslog server address is configured with the **logging <syslog-ip-addr>** command, no **show debug** command output is displayed.)*

```
ProCurve(config)# show debug

Debug Logging
Destination:
Logging --
 10.28.38.164
Facility=kern
Severity=warning
System module=all-pass
Enabled debug types:
event
```

Figure C-12. Sample Output of show debug Command

Example: In the following example, no Syslog servers are configured on the switch (default setting). When you configure a Syslog server, debug logging is enabled to send Event Log messages to the server. To limit the Event Log messages sent to the Syslog server, specify a set of messages by entering the **logging severity** and **logging system-module** commands.

```
ProCurve(config)# show debug
Debug Logging
Destination: None
Enabled debug types:
None are enabled

ProCurve(config)# logging 10.28.38.164
ProCurve(config)# write memory
ProCurve(config)# show debug
Debug Logging
Destination:
 10.28.38.164
Facility=user
Severity=debug
System module=all-pass
Enabled debug types:
event

ProCurve(config)# logging severity error
ProCurve(config)# logging system-module iplock
```

Displays the default debug configuration. (No Syslog server IP addresses or debug types are configured.)

When you configure a Syslog IP address with the **logging** command, by default, the switch enables debug messaging to the Syslog address and the **user** facility on the Syslog server, and sends Event Log messages of all severity levels from all system modules.

You can enter the **logging severity** and **logging system-module** commands to specify a subset of Event Log messages to send to the Syslog server.

Figure C-13. Syslog Configuration to Receive Event Log Messages From Specified System Module and Severity Levels

As shown at the top of Figure C-13, if you enter the **show debug** command when no Syslog server IP address is configured, the configuration settings for Syslog server facility, Event Log severity level and system module are not displayed.

However, after you configure a Syslog server address and enable Syslog logging, all debug and logging settings are displayed with the **show debug** command. If you do not want Event Log messages sent to Syslog servers, you can block the messages from being sent by entering the **no debug event** command. (There is no effect on the normal logging of messages in the switch's Event Log.)

Debug Command

At the manager level, use the **debug** command to perform two main functions:

- Specifies the types of event messages to be sent to an external destination.
- Specifies the destinations to which selected message types are sent.

By default, no debug destination is enabled and only Event Log messages are enabled to be sent.

Note

To configure a Syslog server, use the **logging** *<syslog-ip-addr>* command. For more information, see “Configuring a Syslog Server” on page C-45.

Debug Messages

Use the **debug** command to configure the types of debug messages that the switch can send to configured debug destinations.

Syntax: [no] debug *< debug-type >*

all

Configures the switch to send all debug message types (Event Log, IP RIP, and LLDP) to configured debug destination(s). (Default: Disabled - No debug messages are sent.)

event

Configures the switch to send Event Log messages to configured debug destinations.

Note: *This value does not affect the reception of event notification messages in the Event Log on the switch.*

Event Log messages are automatically enabled to be sent to debug destinations in these conditions:

- *If no Syslog server address is configured and you enter the **logging** <syslog-ip-addr> command to configure a destination address.*
- *If at least one Syslog server address is configured in the startup configuration and the switch is rebooted or reset.*

Event log messages are the default type of debug message sent to configured debug destinations.

ip [rip < database | event | trigger >]

rip < database | event | trigger > > — *Enables the specified RIP message type for the configured destination(s).*

database— *Display database changes.*

event— *Display RIP events.*

trigger— *Display trigger messages.*

lldp

Enables all LLDP message types for the configured destinations.

Debug Destinations

Use the **debug destination** command to enable (and disable) Syslog messaging on a Syslog server or to a CLI session for specified types of debug and Event Log messages.

Syntax: [no] debug destination < logging | session | buffer >

logging

*Enables Syslog logging to configured Syslog servers so that the debug message types specified by the **debug** <debug-type> command (see “Debug Messages” on page C-42) are sent. (Default: Logging disabled)*

To configure a Syslog server IP address, refer to “Configuring a Syslog Server” on page C-45.

Note: *Debug messages from the switches covered in this guide have a debug severity level. Because the default configuration of some Syslog servers ignore Syslog messages with the debug severity level, ensure that the Syslog servers you want to use to receive debug messages are configured to accept the debug level. For more information, refer to “Operating Notes for Debug and Syslog” on page C-50.*

session

*Enables transmission of event notification messages to the CLI session that most recently executed this command. The session can be on any one terminal emulation device with serial, Telnet, or SSH access to the CLI at the Manager level prompt (**ProCurve#_**). If more than one terminal device has a console session with the CLI, you can redirect the destination from the current device to another device. Do so by executing **debug destination session** in the CLI on the terminal device on which you now want to display event messages.*

*Event message types received on the selected CLI session are configured with the **debug < debug-type >** command. (Refer to “Debug Messages” on page C-42.)*

buffer

*Enables Syslog logging to send the debug message types specified by the **debug < debug-type >** command to a buffer in switch memory. To view the debug messages stored in the switch buffer, enter the **show debug buffer** command.*

Logging Command

At the global configuration level, the **logging** command allows you to enable debug logging on specified Syslog servers and select a subset of Event Log messages to send for debugging purposes according to:

- Severity level
- System module

By specifying both a severity level and system module, you can use both configured settings to filter the Event Log messages you want to use to troubleshoot switch or network error conditions.

Caution

After you configure a Syslog server and a severity level and/or system module to filter the Event Log messages that are sent, if you save these settings to the startup configuration file by entering the **write memory** command, these debug and logging settings are automatically re-activated after a switch reboot or power recycle. The debug settings and destinations configured in your previous troubleshooting session will then be applied to the current session, which may not be desirable.

After a reboot, messages remain in the Event Log and are not deleted. However, after a power recycle, all Event Log messages are deleted.

If you configure a severity level and/or system module to temporarily filter Event Log messages, be sure to reset the values to their default settings by entering the **no** form of the following commands to ensure that Event Log messages of all severity levels and from all system modules are sent to configured Syslog servers:

```
ProCurve(config)# no logging severity <debug|major|error|warning|info>  
ProCurve(config)# no logging system-module <system-module >
```

Configuring a Syslog Server

Syslog is a client-server logging tool that allows a client switch to send event notification messages to a networked device operating with Syslog server software. Messages sent to a Syslog server can be stored to a file for later debugging analysis.

To use the Syslog feature, you must install and configure a Syslog server application on a networked host accessible to the switch. Refer to the documentation for the Syslog server application for instructions.

To configure a Syslog server, use the **logging** < *syslog-ip-addr* > command as described below.

When you configure a Syslog server, Event Log messages are automatically enabled to be sent to the server. To reconfigure this setting, use the following commands:

- Use **debug** command to specify additional debug message types (see “Debug Messages” on page C-42).
- Use the **logging** command to configure the system module or severity level used to filter the Event Log messages sent to configured Syslog servers (see “Configuring the Severity Level for Event Log Messages Sent to a Syslog Server” on page C-48 and “Configuring the System Module Used to Select the Event Log Messages Sent to a Syslog Server” on page C-49).

To display the currently configured Syslog servers as well as the types of debug messages and the severity-level and system-module filters used to specify the Event Log messages that are sent, enter the **show debug** command (see “Displaying a Debug/Syslog Configuration” on page C-40).

Syntax: [no] logging < syslog-ip-addr >

Enables or disables Syslog messaging to the specified IP address. You can configure up to six addresses. If you configure an address when none are already configured, this command enables destination logging (Syslog) and the Event debug type. Therefore, at a minimum, the switch begins sending Event Log messages to configured Syslog servers. The IP-RIP message types will also be sent to the Syslog server(s) if they are currently enabled as debug types. (Refer to “Debug Messages” on page C-42.)

no logging removes all currently configured Syslog logging destinations from the running configuration.

no logging < syslog-ip-address > removes only the specified Syslog logging destination from the running configuration.

If you use the “no” form of the command to delete the only remaining Syslog server address, debug destination logging is disabled on the switch, but the default Event debug type is not changed.

*Also, removing all configured Syslog destinations with the **no logging** command (or a specified Syslog server destination with the **no logging** < syslog-ip-address > command) does not delete the Syslog server IP addresses stored in the startup configuration. To delete Syslog addresses in the startup configuration, you must enter a **no logging** command followed by the **write memory** command. To verify the deletion of a Syslog server address, display the startup configuration by entering the **show config** command.*

*To block the messages sent to configured Syslog servers from the currently configured debug message type, enter the **no debug** < debug-type > command. (See “Debug Messages” on page C-42.)*

To disable Syslog logging on the switch without deleting configured server addresses, enter the **no debug destination logging** command. Note that, unlike the case in which no Syslog servers are configured, if one or more Syslog servers are already configured and Syslog messaging is disabled, configuring a new server address does not re-enable Syslog messaging. To re-enable Syslog messaging, you must enter the **debug destination logging** command.

Syntax: [no] logging facility < facility-name >

The logging facility specifies the destination subsystem used in a configured Syslog server. (All configured Syslog servers must use the same subsystem.) ProCurve recommends the default (*user*) subsystem unless your application specifically requires another subsystem. Options include:

- user** (default) — Random user-level messages
- kern** — Kernel messages
- mail** — Mail system
- daemon** — System daemons
- auth** — Security/Authorization messages
- syslog** — Messages generated internally by Syslog
- lpr** — Line-Printer subsystem
- news** — Netnews subsystem
- uucp** — uucp subsystem
- cron** — cron/at subsystem
- sys9** — cron/at subsystem
- sys10 - sys14** — Reserved for system use
- local10 - local17** — Reserved for system use

Use the **no** form of the command to remove the configured facility and reconfigure the default (**user**) value. For a list of supported ProCurve switches, refer to the Note on page C-38.

Configuring the Severity Level for Event Log Messages Sent to a Syslog Server

Event Log messages are entered with one of the following severity levels (from highest to lowest):

Major: A fatal error condition has occurred on the switch.

Error: An error condition has occurred on the switch.

Warning: A switch service has behaved unexpectedly.

Information: Information on a normal switch event.

Debug: Reserved for ProCurve internal diagnostic information.

Using the **logging severity** command, you can select a set of Event Log messages according to their severity level and send them to a Syslog server. Messages of the selected and higher severity will be sent. To configure a Syslog server, see “Configuring a Syslog Server” on page C-45.

Syntax: [no] logging severity < major | error | warning | info | debug >

Configures the switch to send all Event Log messages with a severity level equal to or higher than the specified value to all configured Syslog servers.

*Default: **debug** (Reports messages of all severity levels.)*

*Use the **no** form of the command to remove the configured severity level and reconfigure the default value, which sends Event Log messages of all severity levels to Syslog servers.*

Note: *The severity setting does not affect event notification messages that the switch normally sends to the Event Log. All messages remain recorded in the Event Log.*

Configuring the System Module Used to Select the Event Log Messages Sent to a Syslog Server

Event Log messages contain the name of the system module that reported the event. Using the **logging system-module** command, you can select a set of Event Log messages according to the originating system module and send them to a Syslog server. To configure a Syslog server, see “Configuring a Syslog Server” on page C-45.

Using the **logging system-module** command, you can select messages from only one system module to be sent to a Syslog server. You cannot configure messages from multiple system modules to be sent. If you re-enter the command with a different system module name, the currently configured value is replaced with the new one.

Syntax: [no] logging system-module < system-module >

Configures the switch to send all Event Log messages being logged from the specified system module to configured Syslog servers.

Refer to Table C-1 on page C-27 for the correct value to enter for each system module.

*Default: **all-pass** (Reports all Event Log messages.)*

*Use the **no** form of the command to remove the configured system module value and reconfigure the default value, which sends Event Log messages from all system modules to Syslog servers.*

Note: *This setting has no effect on event notification messages that the switch normally sends to the Event Log.*

Operating Notes for Debug and Syslog

- **Rebooting the Switch or pressing the Reset button resets the Debug Configuration.**

Debug Option	Effect of a Reboot or Reset
logging (debug destination)	If Syslog server IP addresses are stored in the startup-config file, they are saved across a reboot and the logging destination option remains enabled. Otherwise, the logging destination is disabled.
session (debug destination)	Disabled.
All (debug type)	Disabled.
event (debug type)	If a Syslog server IP address is configured in the startup-config file, the sending of Event Log messages is reset to enabled , regardless of the last active setting. If no Syslog server is configured, the sending of Event Log messages is disabled .
IP (debug type)	Disabled.

- **Debug commands do not affect normal message output to the Event Log.**

Using the **debug event** command, you can specify that Event Log messages are sent to the debug destinations you configure (CLI session and/or Syslog servers) in addition to the Event Log.

- **Ensure that your Syslog servers accept Debug messages.**

All Syslog messages resulting from a debug operation have a “debug” severity level. If you configure the switch to send debug messages to a Syslog server, ensure that the server’s Syslog application is configured to accept the “debug” severity level. (The default configuration for some Syslog applications ignores the “debug” severity level.)

Diagnostic Tools

Diagnostic Features

Feature	Default	Menu	CLI	Web
Port Auto negotiation	n/a	n/a	n/a	n/a
Ping Test	n/a	—	page C-54	page C-53
Link Test	n/a	—	page C-54	page C-53
Display Config File	n/a	—	page C-64	page C-65
Admin. and Troubleshooting Commands	n/a	—	page C-67	—
Factory-Default Config	page C-71 (Buttons)	—	page C-71	—
Port Status	n/a	pages page B-11 and page B-12	pages page B-11 and page B-12	pages page B-11 and page B-12

Port Auto-Negotiation

When a link LED does not light (indicating loss of link between two devices), the most common reason is a failure of port auto-negotiation between the connecting ports. If a link LED fails to light when you connect the switch to a port on another device, do the following:

1. Ensure that the switch port and the port on the attached end-node are both set to **Auto** mode.
2. If the attached end-node does not have an **Auto** mode setting, then you must manually configure the switch port to the same setting as the end-node port. Refer to Chapter 10, “Port Status and Configuration”.

Ping and Link Tests

The Ping test and the Link test are point-to-point tests between your switch and another IEEE 802.3-compliant device on your network. These tests can tell you whether the switch is communicating properly with another device.

Note

To respond to a Ping test or a Link test, the device you are trying to reach must be IEEE 802.3-compliant.

Ping Test. This is a test of the path between the switch and another device on the same or another IP network that can respond to IP packets (ICMP Echo Requests). To use the **ping** (or **tracert**) command with host names or fully qualified domain names, refer to “DNS Resolver” on page C-56.

Link Test. This is a test of the connection between the switch and a designated network device on the same LAN (or VLAN, if configured). During the link test, IEEE 802.2 test packets are sent to the designated network device in the same VLAN or broadcast domain. The remote device must be able to respond with an 802.2 Test Response Packet.

Web: Executing Ping or Link Tests

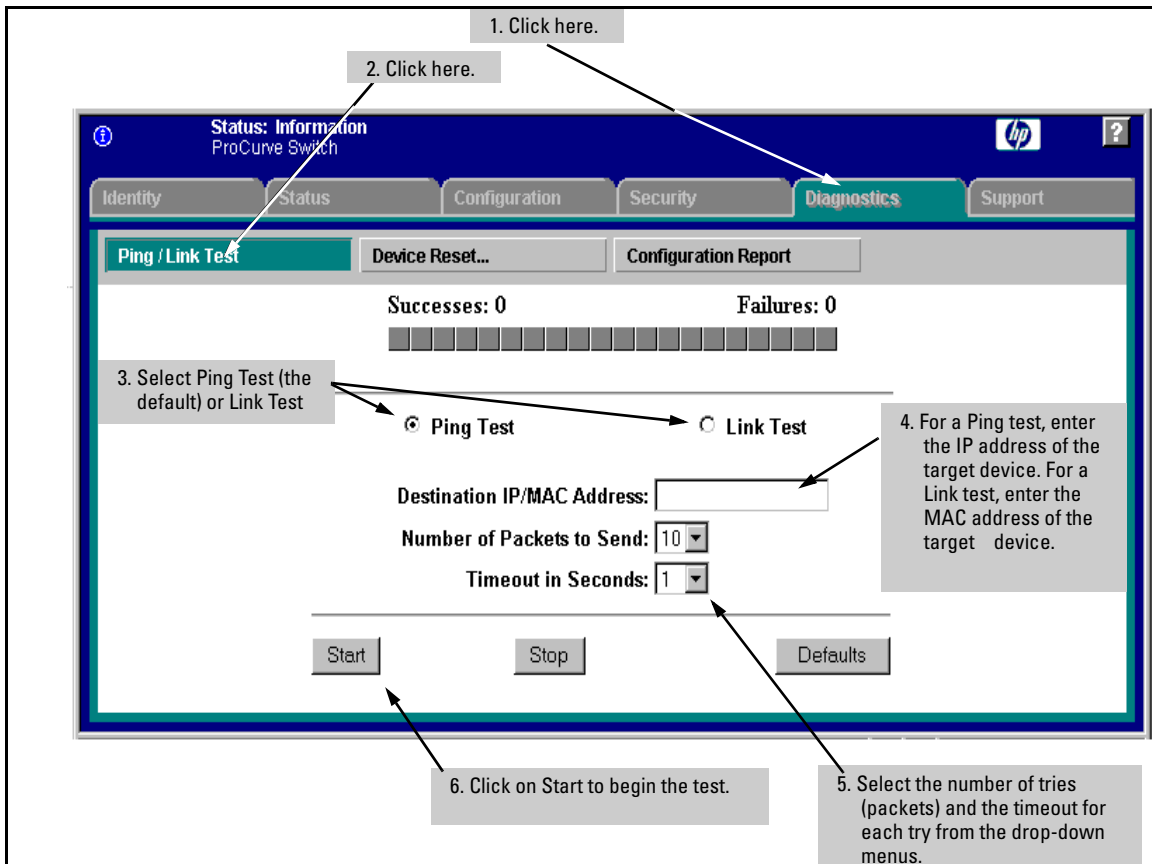


Figure C-14. Link and Ping Test Screen on the Web Browser Interface

Successes indicates the number of Ping or Link packets that successfully completed the most recent test.

Failures indicates the number of Ping or Link packets that were unsuccessful in the last test. Failures indicate connectivity or network performance problems (such as overloaded links or devices).

Destination IP/MAC Address is the network address of the target, or destination, device to which you want to test a connection with the switch. An IP address is in the *X.X.X.X* format where *X* is a decimal number between 0 and 255. A MAC address is made up of 12 hexadecimal digits, for example, 0060b0-080400.

Number of Packets to Send is the number of times you want the switch to attempt to test a connection.

Timeout in Seconds is the number of seconds to allow per attempt to test a connection before determining that the current attempt has failed.

To halt a Link or Ping test before it concludes, click on the Stop button.
To reset the screen to its default settings, click on the Defaults button.

CLI: Ping or Link Tests

Ping Tests. You can issue single or multiple ping tests with varying repetitions and timeout periods. The defaults and ranges are:

- Repetitions: 1 (1 - 999)
- Timeout: 5 seconds (1 - 256 seconds)

Syntax: ping < ip-address > [repetitions < 1 - 999 >] [timeout < 1 - 256 >]

Basic Ping Operation	→	ProCurve > ping 10.28.227.103 10.28.227.103 is alive, time = 15 ms
Ping with Repetitions	→	ProCurve> ping 10.28.227.103 repetitions 3 10.28.227.103 is alive, iteration 1, time = 15 ms 10.28.227.103 is alive, iteration 2, time = 15 ms 10.28.227.103 is alive, iteration 3, time = 15 ms
Ping with Repetitions and Timeout	→	ProCurve > ping 10.28.227.103 repetitions 3 timeout 2 10.28.227.103 is alive, iteration 1, time = 15 ms 10.28.227.103 is alive, iteration 2, time = 10 ms 10.28.227.103 is alive, iteration 3, time = 15 ms
Ping Failure	↘	ProCurve > ping 10.28.227.105 Target did not respond.

Figure C-15. Examples of Ping Tests

To halt a ping test before it concludes, press **[Ctrl] [C]**.

Note

To use the **ping** (or **tracert**) command with host names or fully qualified domain names, refer to “DNS Resolver” on page C-56.

Link Tests. You can issue single or multiple link tests with varying repetitions and timeout periods. The defaults are:

- Repetitions: 1 (1 - 999)
- Timeout: 5 seconds (1 - 256 seconds)

Syntax: link < mac-address > [repetitions < 1 - 999 >] [timeout < 1 - 256 >]
[vlan < vlan-id >]

Basic Link Test	ProCurve# link 0030c1-7fcc40 Link-test passed.
Link Test with Repetitions	ProCurve# link 0030c1-7fcc40 repetitions 3 802.2 TEST packets sent: 3, responses received: 3
Link Test with Repetitions and Timeout	ProCurve# link 0030c1-7fcc40 repetitions 3 timeout 1 802.2 TEST packets sent: 3, responses received: 3
Link Test Over a Specific VLAN	ProCurve# link 0030c1-7fcc40 repetitions 3 timeout 1 vlan 1 802.2 TEST packets sent: 3, responses received: 3
Link Test Over a Specific VLAN; Test Fail	ProCurve# link 0030c1-7fcc40 repetitions 3 timeout 1 vlan 222 802.2 TEST packets sent: 3, responses received: 0

Figure C-16. Example of Link Tests

DNS Resolver

The Domain Name System (DNS) resolver is designed for use in local network domains where it enables use of a host name or fully qualified domain name with DNS-compatible switch CLI commands. (At software release T.13.01, the DNS-compatible commands include **ping** and **traceroute**.)

Beginning with software release T.13.01, DNS operation supports both IPv4 and IPv6 DNS resolution and multiple, prioritized DNS servers. (For information on IPv6 DNS resolution, refer to the latest *IPv6 Configuration Guide* for your switch.)

Terminology

Domain Suffix — Includes all labels to the right of the unique host name in a fully qualified domain name assigned to an IP address. For example, in the fully qualified domain name “device53.evergreen.trees.org”, the domain suffix is “evergreen.trees.org”, while “device53” is the unique (host) name assigned to a specific IP address.

Fully Qualified Domain Name — The sequence of labels in a domain name identifying a specific host (host name) and the domain in which it exists. For example, if a device with an IP address of 10.10.10.101 has a host name of *device53* and resides in the *evergreen.trees.org* domain, then the device’s fully qualified domain name is *device53.evergreen.trees.org* and the DNS resolution of this name is 10.10.10.101.

Host Name — The unique, leftmost label in a domain name assigned to a specific IP address in a DNS server configuration. This enables the server to distinguish a device using that IP address from other devices in the same domain. For example, in the *evergreen.trees.org* domain, if an IPv4 address of 10.10.100.27 is assigned a host name of *accounts015* and another IP address of 10.10.100.33 is assigned a host name of *sales021*, then the switch configured with the domain suffix *evergreen.trees.org* and a DNS server that resolves addresses in that domain can use the host names to reach the devices with DNS-compatible commands. For example:

```
ping accounts015
traceroute accounts015
```

Basic Operation

- When the switch is configured with only the IP address of a DNS server available to the switch, then a DNS-compatible command, executed with a fully qualified domain name, can reach a device found in any domain accessible through the configured DNS server.
- When the switch is configured with both of the following:
 - the IP address of a DNS server available to the switch
 - the domain suffix of a domain available to the configured DNS server

then:

- A DNS-compatible command that includes the host name of a device in the same domain as the configured domain suffix can reach that device.
- A DNS-compatible command that includes a fully qualified domain name can reach a device in any domain that is available to the configured DNS server.

Example. Suppose the switch is configured with the domain suffix **mygroup.procurve.net** and the IP address for an accessible DNS server. If an operator wants to use the switch to ping a target host in this domain by using the DNS name “leader” (assigned by a DNS server to an IP address used in that domain), then the operator can use either of the following commands:

```

ProCurve# ping leader
10.28.229.220 is alive, time = 1 ms

ProCurve# ping leader.mygroup.procurve.net
10.28.229.220 is alive, time = 1 ms
  
```

Figure C-17. Example of Using Either a Host Name or a Fully Qualified Domain Name

In the preceding example, if the DNS server’s IP address is configured on the switch, but a domain suffix is either not configured or is configured for a different domain than the target host, then the fully qualified domain name *must* be used.

Note that if the target host is in a domain *other than* the domain configured on the switch, then:

- The host's domain must be reachable from the switch. This requires that the DNS server for the switch must be able to communicate with the DNS server(s) in the path to the domain in which the target host operates.
- The fully qualified domain name must be used, and the domain suffix must correspond to the domain in which the target host operates, regardless of the domain suffix configured in the switch.

Example. Suppose the switch is configured with the domain suffix **mygroup.procurve.net** and the IP address for an accessible DNS server in this same domain. This time, the operator wants to use the switch to trace the route to a host named "remote-01" in a different domain named **common.group.net**. Assuming this second domain is accessible to the DNS server already configured on the switch, a **traceroute** command using the target's fully qualified DNS name should succeed.

```
ProCurve# traceroute [remote-01.common.group.net] ← Fully Qualified Host Name for the Target Host
[traceroute to 10.22.240.73]
      1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.28.229.3          0 ms      0 ms      0 ms
 2 10.71.217.1         0 ms      0 ms      0 ms
 3 10.0.198.2          1 ms      0 ms      0 ms
 4 10.22.240.73        0 ms      0 ms      0 ms] ← IP Address for Target Host "remote-01"
```

Figure C-18. Example Using the Fully Qualified Domain Name for an Accessible Target in Another Domain

Configuring and Using DNS Resolution with DNS-Compatible Commands

(At software release T.13.01, the DNS-compatible commands include **ping** and **traceroute**.)

1. Determine the following:
 - a. The IP address for a DNS server operating in a domain in your network
 - b. The priority (1 -3) of the selected server, relative to other DNS servers in the domain

- c. The domain name for an accessible domain in which there are hosts you want to reach with a DNS-compatible command. (This is the domain suffix in the fully qualified domain name for a given host operating in the selected domain. Refer to “Terminology” on page C-56.) Note that if a domain suffix is not configured, fully qualified domain names can be used to resolve DNS-compatible commands.
 - d. the host names assigned to target IP addresses in the DNS server for the specified domain
2. Use the data from steps 1a through 1c to configure the DNS entry on the switch.
 3. Use a DNS-compatible command with the host name to reach the target devices.

Configuring a DNS Entry

The switch allows up to three DNS server entries (IP addresses for DNS servers). One domain suffix can also be configured to support resolution of DNS names in that domain by using a host name only. Including the domain suffix enables the use of DNS-compatible commands with a target’s host name instead of the target’s fully qualified domain name.

Syntax: [no] ip dns server-address priority < 1 - 3 > < ip-addr >

Configures the access priority and IP address of a DNS server accessible to the switch. These settings specify:

- *the relative priority of the DNS server when multiple servers are configured*
- *the IP address of the DNS server*

These settings must be configured before a DNS-compatible command can be executed with host name criteria.

*The switch supports three prioritized DNS server entries. Configuring another IP address for a priority that has already been assigned to an IP address is not allowed. To replace one IP address at a given priority level with another address having the same priority, you must first use the **no** form of the command to remove the unwanted address. Also, only one instance of a given server address is allowed in the server list. Attempting to enter a duplicate of an existing entry at a different priority level is not allowed. To change the priority of an existing server address, use the **no** form of the command to remove the entry, then re-enter the address with the new priority.*

*The **no** form of the command replaces the configured IP address with the null setting. (Default: null)*

Syntax: [no] ip dns domain-name < domain-name-suffix >

This optional DNS command configures the domain suffix that is automatically appended to the host name entered with a DNS-compatible command. When the domain suffix and the IP address for a DNS server that can access that domain are both configured on the switch, you can execute a DNS-compatible command using only the host name of the desired target. (For an example, refer to figure C-17 on page C-57.) In either of the following two instances, you must manually provide the domain identification by using a fully qualified DNS name with a DNS-compatible command:

- *If the DNS server IP address is configured on the switch, but the domain suffix is not configured (null)*
- *The domain suffix configured on the switch is not the domain in which the target host exists*

The switch supports one domain suffix entry and three DNS server IP address entries. (Refer to the preceding command description.)

*The **no** form of the command replaces the configured domain suffix with the null setting. (Default: null)*

Example Using DNS Names with Ping and Traceroute

In the network illustrated in figure C-19, the switch at 10.28.192.1 is configured to use DNS names for DNS-compatible commands in the *pubs.outdoors.com* domain. The DNS server has been configured to assign the host name *docservr* to the IP address used by the document server (10.28.229.219).

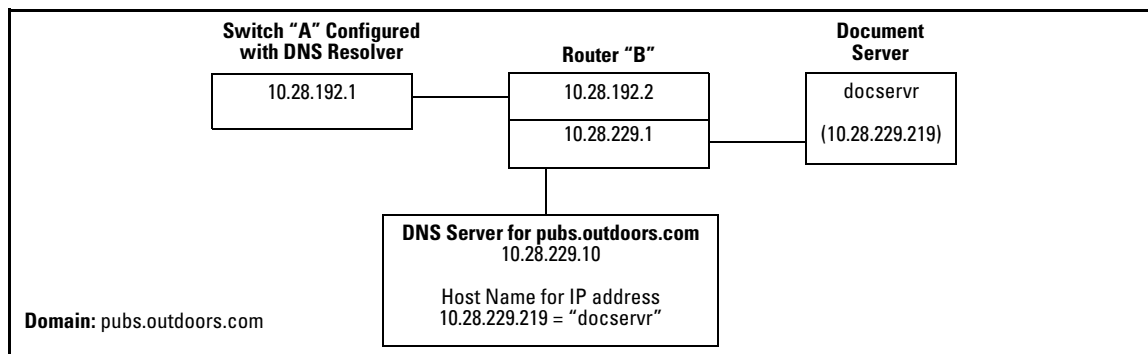


Figure C-19. Example Network Domain

Configuring switch “A” with the domain name and the IP address of a DNS server for the domain enables the switch to use host names assigned to IP addresses in the domain to perform **ping** and **tracert** actions on the devices in the domain. To summarize:

Entity:	Identity:
DNS Server IP Address	10.28.229.10
Domain Name (and Domain Suffix for Hosts in the Domain)	pubs.outdoors.com
Host Name Assigned to 10.28.229.219 by the DNS Server	docservr
Fully Qualified Domain Name for the IP address Used By the Document Server (10.28.229.219)	docservr.pubs.outdoors.com
Switch IP Address	10.28.192.1
Document Server IP Address	10.28.229.219

With the above already configured, the following commands enable a DNS-compatible command with the host name **docservr** to reach the document server at 10.28.229.219.

```
ProCurve(config)# ip dns server-address 10.28.229.10
ProCurve(config)# ip dns domain-name pubs.outdoors.com
```

Figure C-20. Configuring Switch “A” in FigureC-19 To Support DNS Resolution

```
ProCurve# ping docservr
10.28.229.219 is alive, time = 1 ms

ProCurve# traceroute docservr
traceroute to 10.28.229.219
1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.28.192.2 ← 1 ms 0 ms 0 ms
 2 10.28.229.219 ← 0 ms 0 ms 0 ms
```


First-Hop Router (“B”)

Traceroute Target

Figure C-21. Example of Ping and Traceroute Execution for the Network in Figure C-19 on Page C-60

As mentioned under “Basic Operation” on page C-57, if the DNS entry configured in the switch does not include the domain suffix for the desired target, then you must use the target host’s fully qualified domain name with DNS-compatible commands. For example, using the document server in figure C-19 as a target:

```
ProCurve# ping docsrvr.pubs.outdoors.com  
10.28.229.219 is alive, time = 1 ms  
  
ProCurve# traceroute docsrvr.pubs.outdoors.com  
traceroute to 10.28.229.219  
1 hop min, 30 hops max, 5 sec. timeout, 3 probes  
1 10.28.192.2          1 ms          0 ms          0 ms  
2 10.28.229.219       0 ms          0 ms          0 ms
```



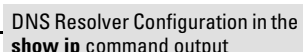
Target's Fully Qualified Domain Name

Figure C-22. Example of Ping and Traceroute Execution When Only the DNS Server IP Address Is Configured

Viewing the Current DNS Configuration

The **show ip** command displays the current domain suffix and the IP address of the highest priority DNS server configured on the switch, along with other IP configuration information. If the switch configuration currently includes a non-default (non-null) DNS entry, it will also appear in the **show run** command output.

```
ProCurve# show ip  
  
Internet (IP) Service  
  
IP Routing : Disabled  
  
Default Gateway : 10.28.192.2  
Default TTL     : 64  
Arp Age        : 20  
Domain Suffix   : pubs.outdoors.com  
DNS server      : 10.28.229.10  
  
VLAN           | IP Config  IP Address   Subnet Mask  
-----+-----  
DEFAULT_VLAN | Manual    10.28.192.1  255.255.255.0
```



DNS Resolver Configuration in the show ip command output

Figure C-23. Example of Viewing the Current DNS Configuration

Operating Notes

- Configuring another IP address for a priority that has already been assigned to an IP address is not allowed. To replace one IP address at a given priority level with another address having the same priority, you must first use the **no** form of the command to remove the unwanted address. Also, only one instance of a given server address is allowed in the server list. Attempting to enter a duplicate of an existing entry at a different priority level is not allowed. To change the priority of an existing server address, use the **no** form of the command to remove the entry, then re-enter the address with the new priority.
- To change the position of an address already configured with priority *x*, you must first use **no ip dns server-address priority *x* < ip-addr >** to remove the address from the configuration, then use **ip dns server-address priority < ip-addr >** to reconfigure the address with the new priority. Also, if the priority to which you want to move an address is already used in the configuration for another address, you must first use the **no** form of the command to remove the current address from the target priority.
- The DNS server(s) and domain configured on the switch must be accessible to the switch, but it is not necessary for any intermediate devices between the switch and the DNS server to be configured to support DNS operation.
- When multiple DNS servers are configured on the switch, they can reside in the same domain or different domains.
- A DNS configuration must include the IP address for a DNS server that is able to resolve host names for the desired domain. If a DNS server has limited knowledge of other domains, then its ability to resolve DNS-compatible command requests is also limited.
- If the DNS configuration includes a DNS server IP address but does not also include a domain suffix, then any DNS-compatible commands should include the target host's fully qualified domain name. Refer to figure C-17 on page C-57.
- Switch-Initiated DNS packets go out through the VLAN having the best route to the DNS server, even if a Management VLAN has been configured.
- The DNS server address must be manually input. It is not automatically determined via DHCP.

Event Log Messages

Message	Meaning
DNS server address not configured	The switch does not have an IP address configured for the DNS server.
DNS server not responding	The DNS server failed to respond or is unreachable. An incorrect server IP address can produce this result.
Unknown host < <i>host-name</i> >	The host name did not resolve to an IP address. Some reasons for this occurring include: <ul style="list-style-type: none">• The host name was not found.• The named domain was not found.• The domain suffix was expected, but has not been configured. (If the server's IP address has been configured in the switch but the domain name has not been configured, then the host's fully qualified domain name must be used.)

Displaying the Configuration File

The complete switch configuration is contained in a file that you can browse from either the web browser interface or the CLI. It may be useful in some troubleshooting scenarios to view the switch configuration.

CLI: Viewing the Configuration File

Using the CLI, you can display either the running configuration or the startup configuration. (For more on these topics, refer to Chapter 6, "Switch Memory and Configuration".)

Syntax: write terminal

Displays the running configuration.

show config

Displays the startup configuration.

show running-config

Displays the running-config file.

Web: Viewing the Configuration File

To display the running configuration, through the web browser interface:

1. Click on the **Diagnostics** tab.
2. Click on **[Configuration Report]**
3. Use the right-side scroll bar to scroll through the configuration listing.

Listing Switch Configuration and Operation Details

The **show tech** command outputs, in a single listing, switch operating and running configuration details from several internal switch sources, including:

- Image stamp (software version data)
- Running configuration
- Event Log listing
- Boot History
- Port settings
- Status and counters — port status
- IP routes
- Status and counters — VLAN information
- GVRP support
- Load balancing (trunk and LACP)

Syntax: show tech

Executing **show tech** outputs a data listing to your terminal emulator. However, using your terminal emulator's text capture features, you can also save **show tech** data to a text file for viewing, printing, or sending to an associate. For example, if your terminal emulator is the Hyperterminal application available with Microsoft® Windows® software, you can copy the show tech output to a file and then use either Microsoft Word or Notepad to display the data. (In this case, Microsoft Word provides the data in an easier-to-read format.)

To Copy show tech output to a Text File. This example uses the Microsoft Windows terminal emulator. To use another terminal emulator application, refer to the documentation provided with that application.

1. In Hyperterminal, click on **Transfer | Capture Text...**




Figure C-24. The Capture Text window of the Hyperterminal Application

2. In the **File** field, enter the path and file name under which you want to store the **show tech** output.




Figure C-25. Example of a Path and Filename for Creating a Text File from show tech Output

3. Click **[Start]** to create and open the text file.
4. Execute **show tech**:

```
ProCurve# show tech
```

 - a. Each time the resulting listing halts and displays -- MORE --, press the Space bar to resume the listing.
 - b. When the CLI prompt appears, the show tech listing is complete. At this point, click on **Transfer | Capture Text | Stop** in HyperTerminal to stop copying data into the text file created in the preceding steps.

Note

Remember to do the above step to stop HyperTerminal from copying into the text file. Otherwise, the text file remains open to receiving additional data from the HyperTerminal screen.

5. To access the file, open it in Microsoft Word, Notepad, or a similar text editor.

CLI Administrative and Troubleshooting Commands

These commands provide information or perform actions that you may find helpful in troubleshooting operating problems with the switch.

Note

For more on the CLI, refer to chapter 4, “Using the Command Line Interface (CLI)”.

Syntax: show version

Shows the software version currently running on the switch, and the flash image from which the switch booted (primary or secondary).

show boot-history

Displays the switch shutdown history.

show history

Displays the current command history.

[no] page

Toggles the paging mode for display commands between continuous listing and per-page listing.

setup

Displays the Switch Setup screen from the menu interface.

repeat

Repeatedly executes the previous command until a key is pressed.

kill

Terminates all other active sessions.

Traceroute Command

The **traceroute** command enables you to trace the route from the switch to a host address.

This command outputs information for each (router) hop between the switch and the destination address. Note that every time you execute **traceroute**, it uses the same default settings unless you specify otherwise for that instance of the command.

Syntax: `traceroute < ip-address >`

Lists the IP address of each hop in the route, plus the time in microseconds for the **traceroute** packet reply to the switch for each hop.

To halt an ongoing traceroute search, press the **[Ctrl] [C]** keys.

`[minttl < 1-255 >]`

*For the current instance of **traceroute**, changes the minimum number of hops allowed for each probe packet sent along the route. If **minttl** is greater than the actual number of hops, then the output includes only the hops at and above the **minttl** threshold. (The hops below the threshold are not listed.) If **minttl** matches the actual number of hops, only that hop is shown in the output. If **minttl** is less than the actual number of hops, then all hops are listed. For any instance of **traceroute**, if you want a **minttl** value other than the default, you must specify that value. (Default: 1)*

`[maxttl < 1-255 >]`

*For the current instance of **traceroute**, changes the maximum number of hops allowed for each probe packet sent along the route. If the destination address is further from the switch than **maxttl** allows, then **traceroute** lists the IP addresses for all hops it detects up to the **maxttl** limit. For any instance of **traceroute**, if you want a **maxttl** value other than the default, you must specify that value. (Default: 30)*

`[timeout < 1-120 >]`

*For the current instance of **traceroute**, changes the timeout period the switch waits for each probe of a hop in the route. For any instance of **traceroute**, if you want a **timeout** value other than the default, you must specify that value. (Default: 5 seconds)*

`[probes < 1-5 >`

*For the current instance of **traceroute**, changes the number of queries the switch sends for each hop in the route. For any instance of **traceroute**, if you want a **probes** value other than the default, you must specify that value. (Default: 3)*

A Low Maxttl Causes Traceroute To Halt Before Reaching the Destination Address. For example, executing **traceroute** with its default values for a destination IP address that is four hops away produces a result similar to this:

```

ProCurve# traceroute 125.25.24.35
traceroute to 125.25.24.35 ,
          1 hop min, 30 hops max, 5 sec. timeout, 3 probes
 1 10.255.120.2      0 ms      0 ms      0 ms
 2 10.71.217.2       7 ms      3 ms      0 ms
 3 10.243.170.1     0 ms      1 ms      0 ms
 4 125.25.24.35    3 ms      3 ms      0 ms
  
```

Intermediate router hops with the time taken for the switch to receive acknowledgement of each probe reaching each router.

Destination IP Address

Figure C-26. Example of a Completed Traceroute Enquiry

Continuing from the previous example (figure C-26, above), executing **traceroute** with an insufficient **maxttl** for the actual hop count produces an output similar to this:

```

ProCurve# traceroute 125.25.24.35 (maxttl 3)
traceroute to 125.25.24.35 ,
          1 hop min, 3 hops max, 5 sec. timeout, 3 probes
 1 10.255.120.2      0 ms      0 ms      0 ms
 2 10.71.217.2       0 ms      0 ms      0 ms
 3 10.243.170.1     0 ms *      0 ms
  
```

Traceroute does not reach destination IP address because of low maxttl setting.

The asterisk indicates there was a timeout on the second probe to the third hop.

Figure C-27. Example of Incomplete Traceroute Due to Low Maxttl Setting

If A Network Condition Prevents Traceroute from Reaching the Destination. Common reasons for Traceroute failing to reach a destination include:

- Timeouts (indicated by one asterisk per probe, per hop; refer to figure C-27, above.)
- Unreachable hosts
- Unreachable networks

- Interference from firewalls
- Hosts configured to avoid responding

Executing traceroute where the route becomes blocked or otherwise fails results in an output marked by timeouts for all probes beyond the last detected hop. For example with a maximum hop count of 7 (**maxttl** = 7), where the route becomes blocked or otherwise fails, the output appears similar to this:

```
ProCurve# traceroute 107.64.197.100 maxttl 7
traceroute to 107.64.197.100 ,
          1 hop min, 7 hops max, 5 sec. timeout, 3 probes
 1 10.255.120.2          0 ms      0 ms      0 ms
 2 10.71.217.2          0 ms      0 ms      0 ms
 3 * (10.243.170.1) *
 4 * * *
 5 * * *
 6 * * *
 7 * * *
```

At hop 3, the first and third probes timed out but the second probe reached the router. All further probes within the **maxttl** timed-out without finding a router or the destination IP address.

An asterisk indicates a timeout without finding the next hop.

Figure C-28. Example of Traceroute Failing to Reach the Destination Address

Restoring the Factory-Default Configuration

As part of your troubleshooting process, it may become necessary to return the switch configuration to the factory default settings. This process momentarily interrupts the switch operation, clears any passwords, clears the console Event Log, resets the network counters to zero, performs a complete self test, and reboots the switch into its factory default configuration including deleting an IP address. There are two methods for resetting to the factory-default configuration:

- CLI
- Clear/Reset button combination

Note

ProCurve recommends that you save your configuration to a TFTP server before resetting the switch to its factory-default configuration. You can also save your configuration via Xmodem, to a directly connected PC.

CLI: Resetting to the Factory-Default Configuration

This command operates at any level *except* the Operator level.

Syntax: erase startup-configuration

Deletes the startup-config file in flash so that the switch will reboot with its factory-default configuration.

Note

The **erase startup-config** command does not clear passwords.

Clear/Reset: Resetting to the Factory-Default Configuration

To execute the factory default reset, perform these steps:

1. Using pointed objects, simultaneously press both the Reset and Clear buttons on the front of the switch.
2. Continue to press the Clear button while releasing the Reset button.

3. When the Self Test LED begins to flash, release the Clear button.

The switch will then complete its self test and begin operating with the configuration restored to the factory default settings.

Restoring a Flash Image

The switch can lose its operating system if either the primary or secondary flash image location is empty or contains a corrupted OS file and an operator uses the **erase flash** command to erase a good OS image file from the opposite flash location.

To Recover from an Empty or Corrupted Flash State. Use the switch's console serial port to connect to a workstation or laptop computer that has the following:

- A terminal emulator program with Xmodem capability, such as the HyperTerminal program included in Windows PC software.
- A copy of a good OS image file for the switch.

Note

The following procedure requires the use of Xmodem, and copies an OS image into primary flash only.

This procedure assumes you are using HyperTerminal as your terminal emulator. If you use a different terminal emulator, you may need to adapt this procedure to the operation of your particular emulator.

1. Start the terminal emulator program.
2. Ensure that the terminal program is configured as follows:
 - Baud rate: 9600
 - 1 stop bit
 - No parity
 - No flow control
 - 8 Bits
3. Use the Reset button to reset the switch. The following prompt should then appear in the terminal emulator:

```
Enter h or ? for help.
```

```
=>
```

4. Since the OS file is large, you can increase the speed of the download by changing the switch console and terminal emulator baud rates to a high speed. For example:

- a. Change the switch baud rate to 115,200 Bps.

=> sp 115200

- b. Change the terminal emulator baud rate to match the switch speed:
 - i. In HyperTerminal, select **Call | Disconnect**.
 - ii. Select **File | Properties**.
 - iii. Click on **Configure**.
 - iv. Change the baud rate to **115200**.
 - v. Click on **[OK]**. In the next window, click on **[OK]** again.
 - vi. Select **Call | Connect**
 - vii. Press **[Enter]** one or more times to display the => prompt.

5. Start the Console Download utility by typing **do** at the => prompt and pressing **[Enter]**:

=> do

6. You will then see this prompt:

```
You have invoked the console download utility.  
Do you wish to continue? (Y/N)>_
```

7. At the above prompt:
 - a. Type **y** (for Yes)
 - b. Select **Transfer | File** in HyperTerminal.
 - c. Enter the appropriate filename and path for the OS image.
 - d. Select the **Xmodem** protocol (and not the 1k Xmodem protocol).
 - e. Click on **[Send]**.

If you are using HyperTerminal, you will see a screen similar to the following to indicate that the download is in progress:

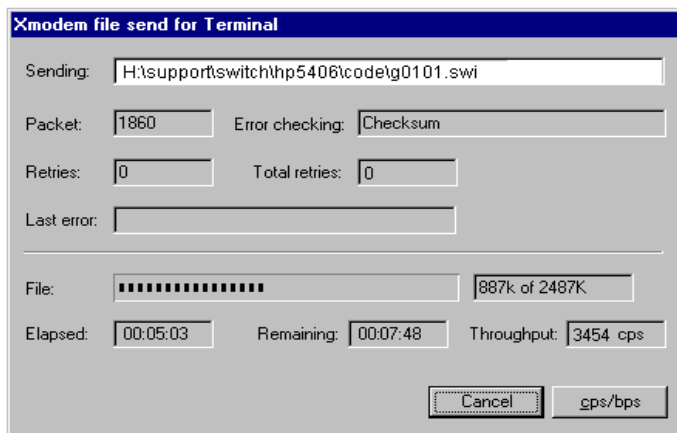


Figure C-29. Example of Xmodem Download in Progress

8. When the download completes, the switch reboots from primary flash using the OS image you downloaded in the preceding steps, plus the most recent startup-config file.