Troubleshooting

Contents

Overview

This chapter 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 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, see the Support and Warranty booklet shipped with the switch.

Troubleshooting Approaches

Use these approaches to diagnose switch problems:

- **Check the ProCurve web site** – the web site may have software updates or other information to help solve your problem:
  http://www.procurve.com
- **Check the switch LEDs** – The LEDs on the switch are a fundamental diagnostic tool. They provide indications of proper switch operation and of any hardware faults that may have occurred:
  - 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.

    See the *Installation Guide* shipped with the switch for a description of the LED behavior and information on using the LEDs for trouble-shooting.
- **Check the network topology/installation** – See the *Installation Guide* shipped with the switch for topology information.

- **Check the network cables** – Cabling problems are a frequent cause of network faults. Check the 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. See the *Installation Guide* shipped with the switch for correct cable types and connector pin-outs.
- **Use the software tools:**
  - **Web Browser Interface** – Use the Port Utilization Graph and Alert Log in the web browser interface included in the switch to help isolate problems. See Chapter 5, "Using the 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)
  - **Switch Console** – For help in isolating problems, use the easy-to-access switch console built into the switch or Telnet to the switch console. See chapter 2, "Using the Menu Interface" and chapter 3, "Using the Command Line Interface (CLI)" for console operation information. 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)
  - **ProCurve Manager / ProCurve Manager +** – Use ProCurve Manager to help isolate problems and recommend solutions.

Chassis Over-Temperature Detection

If a Switch 2800 Series device reaches an over-temperature condition, it generates a chassis-module Warning message in the Event Log and in any optionally configured debug destinations (console session and SyslogD servers). If the switch later returns to its acceptable temperature range, it signals this event with a chassis module Information message to the same destinations. These messages include the number of times the switch has detected the events since the last reboot. For example, suppose that you notice the following three messages at the end of the current Event Log message listing:

```
W 08/17/03 11:28:05 chassis: Over-temperature detected. Failures: 1
I 08/17/03 11:33:23 chassis: Temperature back to normal. Failures: 1
W 08/17/03 12:03:18 chassis: Over-temperature detected. Failures: 2
```

Figure C-1. Chassis Over-Temperature Messaging

The above messages indicate that the switch detected the following chassis conditions since the last reboot:

1. An over-temperature condition occurred on August 17, 2003 at 11:28:05, meaning the switch was operating above its acceptable, internal temperature range. The Failure value of "1" indicates this is the first over-temperature condition to occur since the last reboot.
2. The switch returned to its acceptable temperature range at 11:33:23 on the same day. (To determine this temperature range, refer to the *Installation and Getting Started Guide* shipped with the switch.)
3. Another over-temperature condition occurred on August 17th at 12:03:18 and the switch is currently operating in this condition. The Failure value of "2" indicates this is the second over-temperature condition to occur since the last reboot.

CAUTION

If an over-temperature condition occurs in a Switch Series 2800 device, continued operation can result in damage to the device.

- Check the event log for fan failure warnings. If the switch has experienced a fan failure, remove power from the switch and contact your ProCurve service and support representative.
- If there are no fan failures, ensure that the ambient temperature in the switch's operating area is not causing the over-temperature condition. If the condition persists, remove power from the switch until you can find the cause and apply an effective remedy.

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, see 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. See 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:

- 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, see 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, see 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 the 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.

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 [e] < *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 a static **Trunk** port. To enable LACP on a static-trunked port:
first use the **no trunk [e] < *port-number* >** command to disable the static trunk assignment, and then execute **interface [e] < *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-13.

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. See "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. See "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 the "Note on Supplicant Statistics" 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-2. Example of a Port Remaining Open After Being Configured with "Control Unauthorized"

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.

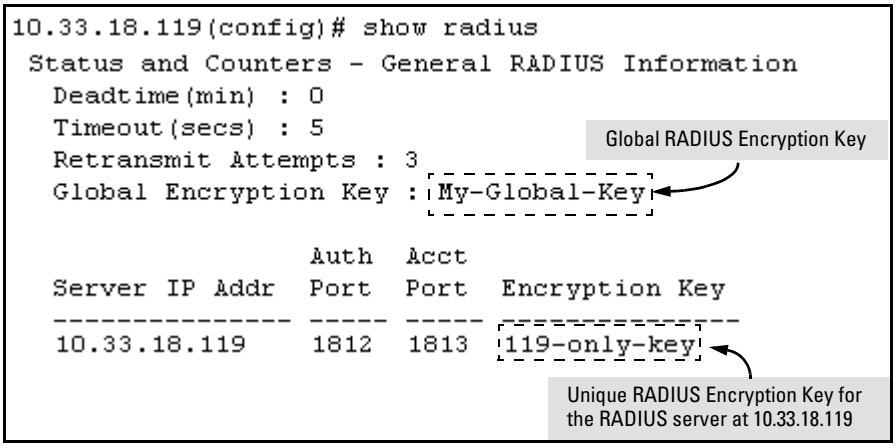

Figure C-3. Example of How To List the Global and Server-Specific Radius Encryption Keys

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.

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.

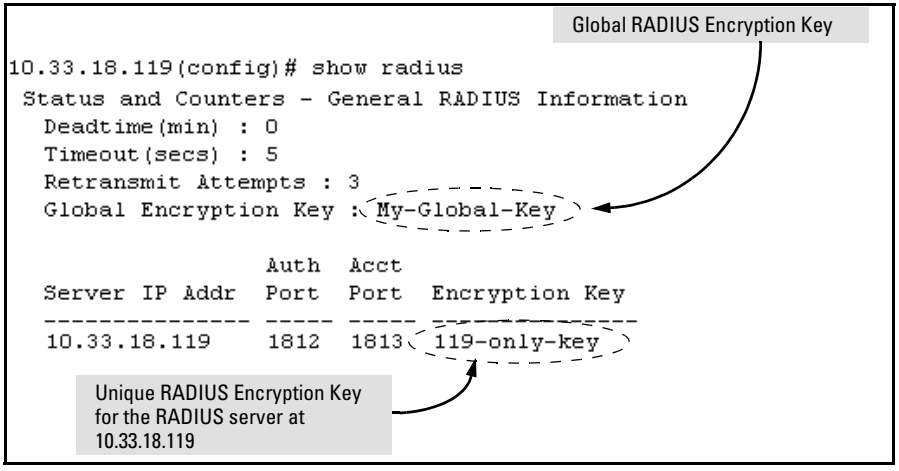

Figure C-4. Examples of Global and Unique Encryption Keys

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

Caution

If you enable STP, it is recommended that you leave the remainder of the STP parameter settings at their default values until you have had an opportunity to evaluate STP performance in your network. Because incorrect STP settings can adversely affect network performance, you should avoid making changes without having a strong understanding of how STP operates. To learn the details of STP operation, refer to the IEEE 802.1D 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 STP 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 devices such as the switches covered by this guide, STP 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. See the chapter on VLANs in the *Advanced Traffic Management Guide*..

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

Problem sources can include:

- Fast-Uplink is configured on a switch that is the STP 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 STP 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 STP root switch has ports configured for fast-uplink STP 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 *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".

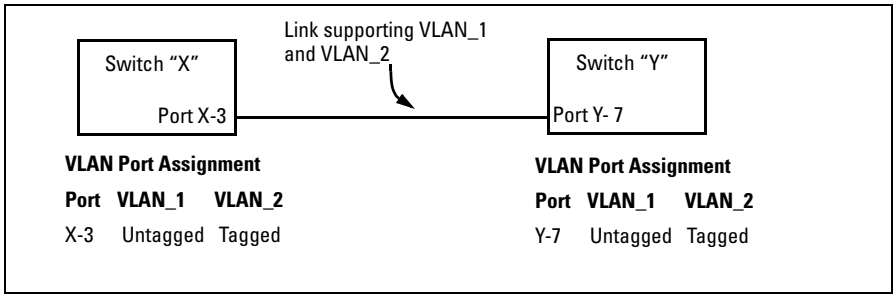

Figure C-5. Example of Correct VLAN Port Assignments on a Link

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 switch operates 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 STP 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 switch has 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.

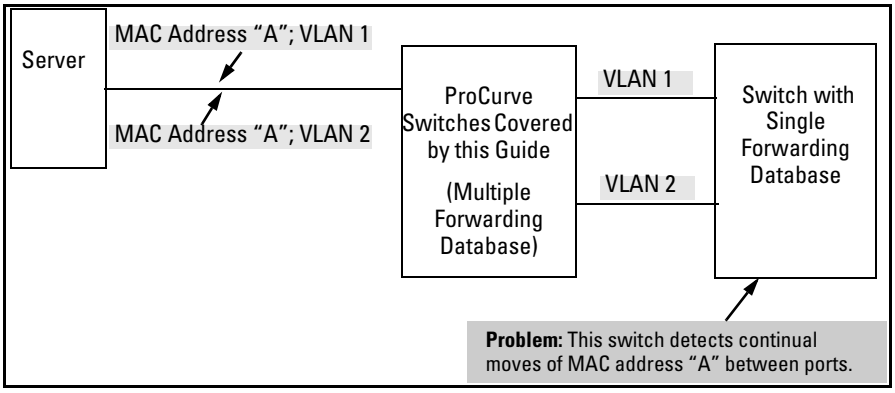

Figure C-6. Example of Duplicate MAC Address

Using Logging To Identify Problem Sources

Event Log Operation

The Event Log records operating events as single-line entries listed in chronological order, and serves as a tool for isolating problems. Each Event Log entry is composed of five fields:

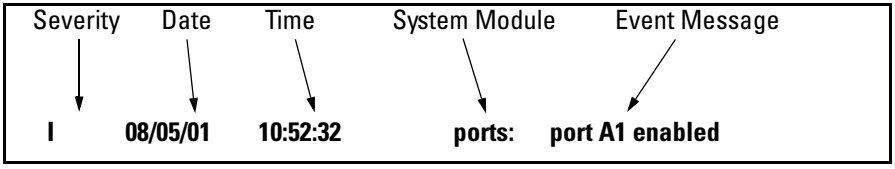

Figure C-7. Anatomy of an Event Log Message

Severity is one of the following codes:

- **I** (information) indicates routine events.
- **W** (warning) indicates that a service has behaved unexpectedly.
- **C** (critical) indicates that a severe switch error has occurred.
- **D** (debug) reserved for internal diagnostic information.

Date is the date in *mm/dd/yy* format that the entry was placed in the log.

Time is the time in *hh:mm:ss* format that the entry was placed in the log.

System Module is the internal module (such as "ports" for port manager) that generated the 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 on page C-24 lists the individual modules.

Event Message is a brief description of the operating event.

The event log holds up to 1000 lines in chronological order, from the oldest to the newest. Each line consists of one complete event message. Once the log has received 1000 entries, it discards the current oldest line each time a new line is received. The event log window contains 14 log entry lines and can be positioned to any location in the log.

The event log will be *erased* if power to the switch is interrupted.

(The event log is *not* erased by using the **Reboot Switch** command in the Main Menu.)

Table C-1.Event Log System Modules

Module	Event Description	Module	Event Description
addrMgr	Address table	mgr	Console management
chassis	switch hardware	ports	Change in port status; static trunks
bootp	bootp addressing	snmp	SNMP communications
console	Console interface	stack	Stacking
dhcp	DHCP addressing	stp	Spanning Tree
download	file transfer	sys, system	Switch management
FFI	Find, Fix, and Inform -- available in the console event log and web browser interface alert log	telnet	Telnet activity
garp	GARP/GVRP	tcp	Transmission control
igmp	IP Multicast	tftp	File transfer for new OS or config.
ip	IP-related	timep	Time protocol
ipx	Novell Netware	vlan	VLAN operations
lacp	Dynamic LACP trunks	Xmodem	Xmodem file transfer

Menu: Entering and Navigating in the Event Log

From the Main Menu, select **Event Log**.

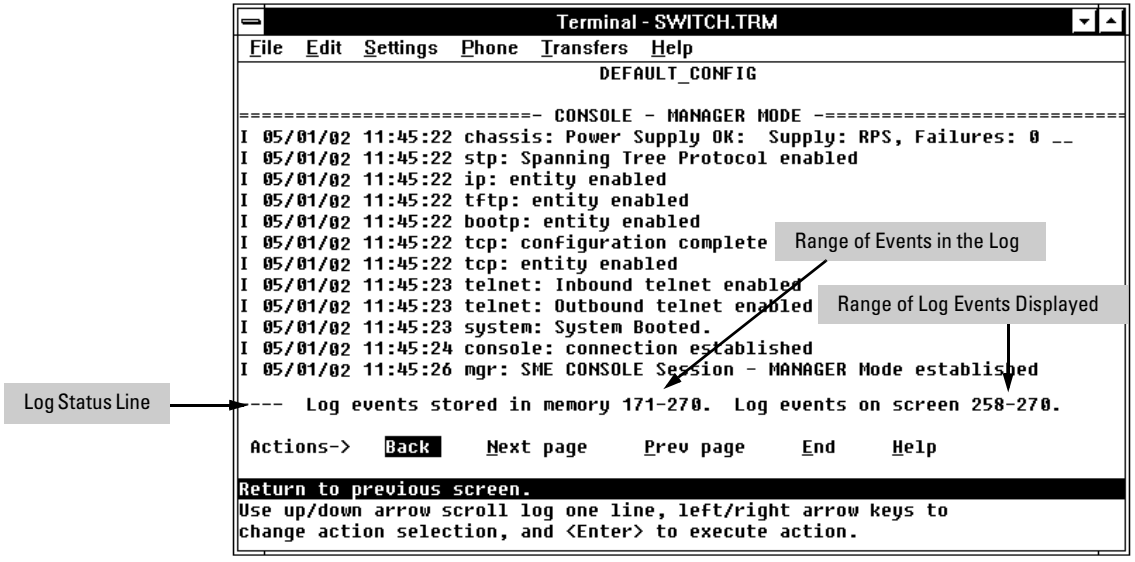

Figure C-8. Example of an Event Log Display

The *log status line* at the bottom of the display identifies where in the sequence of event messages the display is currently positioned.

To display various portions of the Event Log, either preceding or following the currently visible portion, use either the actions listed at the bottom of the display (**Next page**, **Prev page**, or **End**), or the keys described in the following table:

Table C-2. Event Log Control Keys

Key	Action
[N]	Advance the display by one page (next page).
[P]	Roll back the display by one page (previous page).
↓	Advance display by one event (down one line).
↑	Roll back display by one event (up one line).
[E]	Advance to the end of the log.
[H]	Display Help for the event log.

CLI:

Using the CLI, you can list

- Events recorded since the last boot of the switch
- All events recorded
- Event entries containing a specific keyword, either since the last boot or all events recorded

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

```
ProCurve> show logging
```
Lists recorded log messages since last reboot.

```
ProCurve> show logging -a
```
Lists all recorded log messages, including those before the last reboot.

```
ProCurve> show logging -a system
```
Lists log messages with "system" in the text or module name.

```
ProCurve> show logging system
```
Lists all log messages since the last reboot that have "system" in the text or module name.

Debug and Syslog Operation

You can direct switch debug (Event log) messages to these destinations:

- Up to six SyslogD servers
- One management-access session through:
  - A direct-connect RS-232 console CLI session
  - A Telnet session
  - An SSH session

```
ProCurve(config)# debug destination session
ProCurve(config)# EVNT I 01/01/90 05:03:45 ports: port A17 is now off-line
EVNT I 01/01/90 05:03:45 vlan: VLAN_20 virtual LAN disabled
EVNT I 01/01/90 05:03:45 ip: VLAN_20: network disabled on 18.255.120.1
EVNT I 01/01/90 05:03:47 ports: port A18 is Blocked by LACP
EVNT I 01/01/90 05:03:49 ports: port A18 is now on-line
EVNT I 01/01/90 05:03:49 vlan: VLAN_20 virtual LAN enabled
EVNT I 01/01/90 05:03:50 ip: VLAN_20: network enabled on 18.255.120.1
```

Figure C-9. Example of Debug Output to a Console CLI Session

Debug logging requires a logging destination (SyslogD server and/or a session type), and involves the **logging** and **debug destination** commands. Actions you can perform with Debug and Syslog operation include:

- Configure the switch to send Event Log messages to one or more SyslogD servers. Included is the option to send the messages to the **user** log facility (default) on the configured servers, or to another log facility.

Note

As of August, 2003, the **logging facility <** ***facility-name*** **>** option (described on page C-29) is available on these switch models:

- Switch 2600/2600-PWR Series and the Switch 6108 (software release H.07.30 or greater)
- Switch 2800 Series

For the latest feature information on ProCurve switches, visit the ProCurve 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, this includes the Syslog server list.
- Display the current Syslog server list when Syslog logging is disabled.

Debug Types. This section describes the types of debug messages the switch can send to configured debug destinations.

Syntax: [no] debug < debug-type >

all

Configures the switch to send all debug types to the configured debug destination(s). (Default: Disabled)

event

Configures the switch to send Event Log messages to the configured debug destination(s). ***Note:*** *This has no effect on event notification messages the switch routinely sends to the Event Log itself. Also, this debug type is automatically enabled in these cases:*

- *If there is currently no Syslog server address configured and you use* **logging < *ip-addr* >** *to configure an address.*
- *If there is currently at least one Syslog server address configured and the switch is rebooted or reset.*

(Default: Disabled)

port-access-auth

If 802.1x authentication is configured, this option shows the various communication messages sent between the switch, client, and RADIUS server.

(Default: Disabled)

Configuring the Switch To Send Debug Messages to One or More SyslogD Servers. Use the logging command to configure the switch to send Syslog messages to a SyslogD server, or to remove a SyslogD server from the switch configuration.

Syntax: [no] logging < syslog-ip-address | facility < facility-name >>

< ***syslog-ip-address*** >

If there are no SyslogD servers configured, **logging** *enters a SyslogD server IP address* ***and*** *automatically enables Syslog logging to the server. If at least one SyslogD server is already configured and Syslog logging has been disabled, you can still use* **logging** < syslog-ip-addr > *to add another SyslogD server, but Syslog logging remains disabled until you re-enable it with the* **debug destination logging** *command. While Syslog logging is enabled, the switch attempts to send Syslog messages to all configured SyslogD server addresses, and operates regardless of whether session logging is also enabled. To configure multiple SyslogD servers, repeat the command once for each server IP address. (****Default:*** *none;* ***Range:*** *Up to six IP addresses*)

facility < facility-name >

Specifies the destination subsystem the SyslogD server(s) must use. (All SyslogD servers must use the same subsystem.) ProCurve recommends the default (user) subsystem unless your application specifically requires another subsystem. Options include:

user *(the default) - Various 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 *through* **sys14** *- Reserved for system use*
local0 *through* **local7** *- Reserved for system use*

(Some switches covered by this manual do not offer the **facility** *option. Refer to the* **Note** *on page C-27.)*

For example, on a switch where there are no SyslogD servers configured, you would do the following to configure SyslogD servers 18.120.38.155 and 18.120.43.125 and automatically enable Syslog logging (with **user** as the default logging facility):

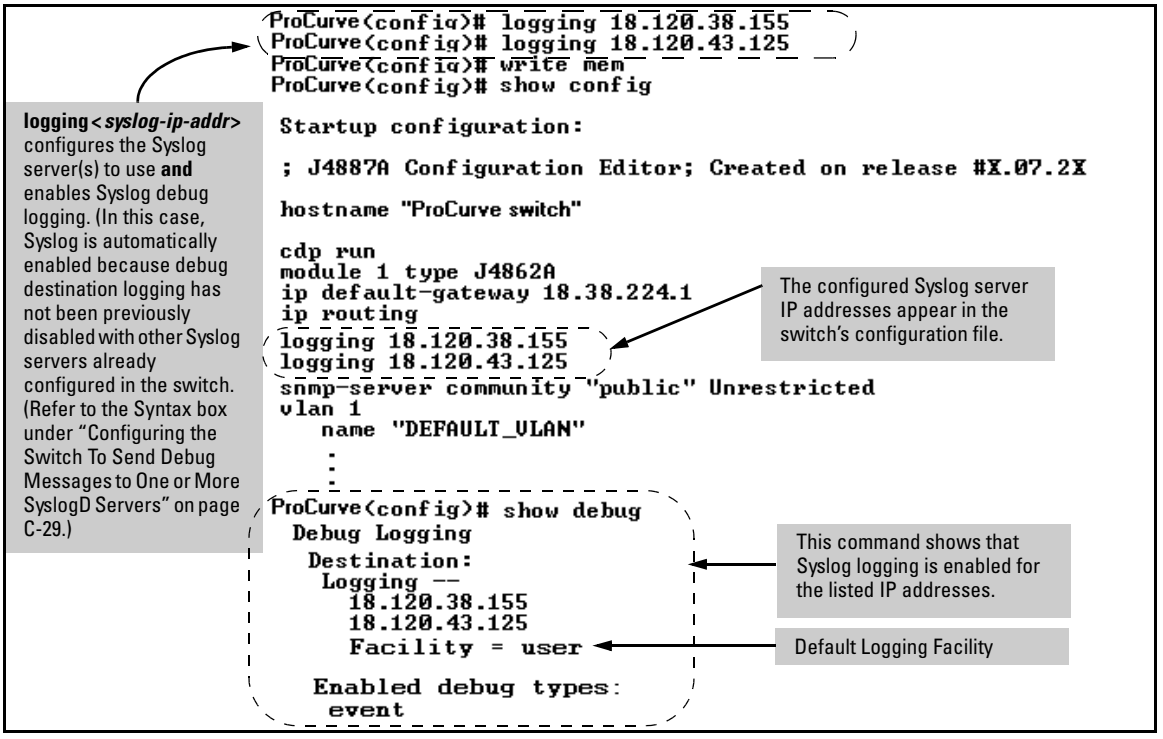

Figure C-10. Example of Configuring and Enabling Syslog Logging

To use a non-default logging facility, such as **lpr**, in the same operation as in figure C-10, you would use this command set:

```
ProCurve(config)# logging 18.120.38.155
ProCurve(config)# logging 18.120.43.125
ProCurve(config)# logging facility lpr
```

Enabling or Disabling Logging to Management Sessions and SyslogD Servers. Use this command when you want to do any of the following:

- Disable Syslog logging on all currently configured SyslogD servers without removing the servers from the switch configuration.
- Re-enable Syslog logging if it is disabled and there is at least one SyslogD server currently configured in the switch.
- Enable or disable logging output to the current management-access session.

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

logging

The **no** *form of the command disables Syslog logging, but retains the currently configured SyslogD server addresses in the switch configuration.When Syslog logging is currently disabled with one or more SyslogD servers configured, this command enables Syslog logging on the switch. The* **show config** *command output includes the SyslogD server IP addresses currently configured in the startup-config file.*

session

Enables and disables debug logging to the current session. The "current session" is the session that most recently executed **debug destination session** *on the switch (since the last reboot). This makes it easy to move session logging from one session to another.*

For example, figure C-11 shows the process for checking the current Syslog status and then disabling Syslog logging.

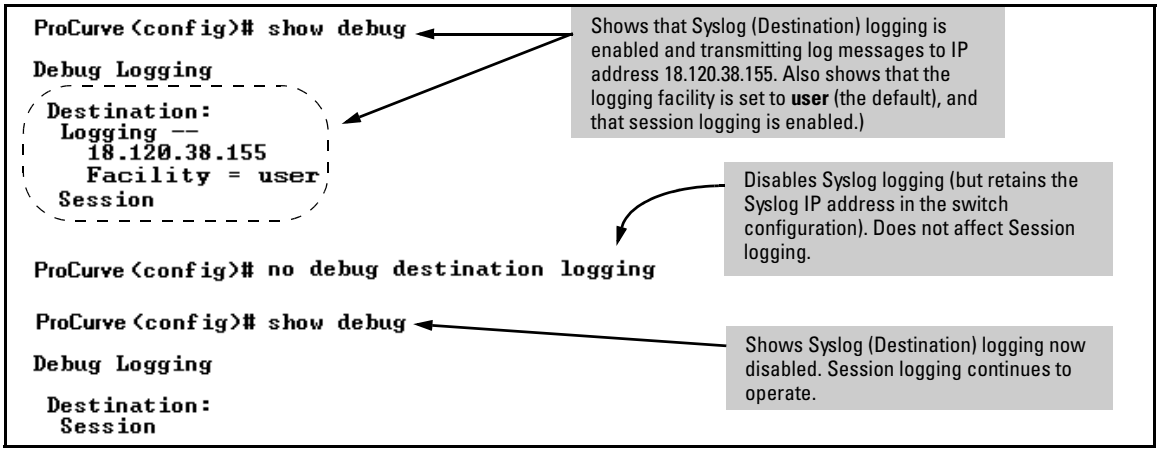

Figure C-11. Example of Disabling Syslog Operation

Viewing Debug (Syslog and Session) Status. Use these commands to determine the current debug configuration and status:

Syntax: show < config | running >

Lists the current startup-config or running-config file, with any currently configured IP addresses for SyslogD servers.

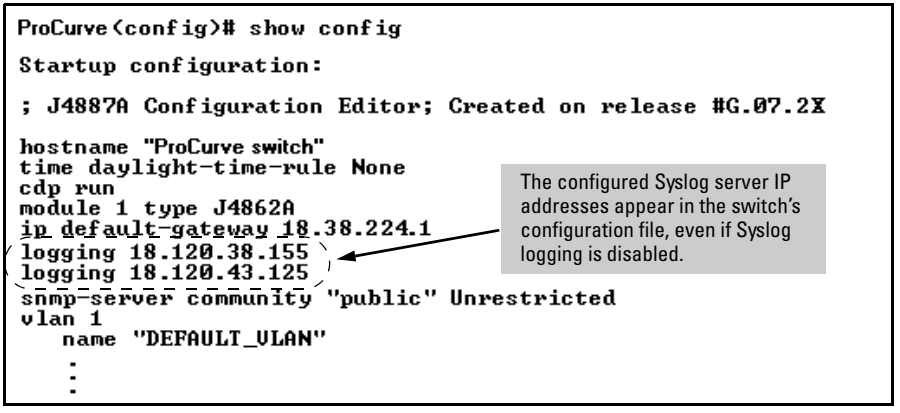

Figure C-12. Example of Show Config Output with SyslogD Servers Configured

Syntax: show debug

List the current debug status for both Syslog logging and Session logging.

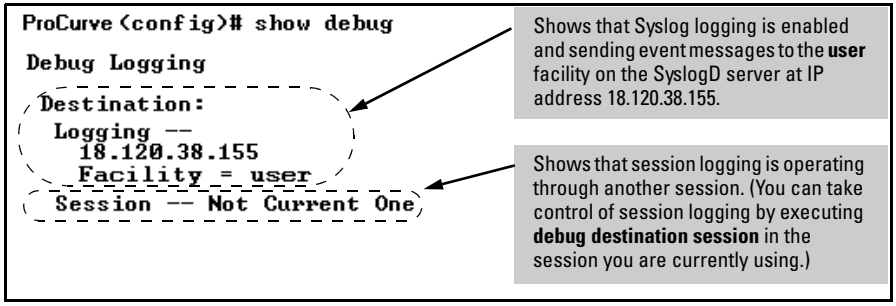

Figure C-13. Example of Show Debug Status

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

Debug Option	Effect of a Reboot or Reset
logging (destination)	If any SyslogD server IP addresses are 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 (destination)	Disabled.
All (event type)	Disabled.
Event (event type)	If a Syslog server is configured in the startup-config file, resets to enabled, regardless of prior setting. Disabled if no Syslog server is configured.
port-access-auth (event type)	Disabled

- **Debug commands do not affect message output to the Event Log.** As a separate option, invoking debug with the **event** option causes the switch to send Event Log messages to whatever debug destination(s) you configure (session and/or logging), as well as to the Event Log.

- **Ensure that your Syslog server(s) will accept Debug messages.** All Syslog messages the switch generates carry the configured facility. All Syslog messages resulting from debug operation carry a "debug" severity. If you configure the switch to transmit debug messages to a SyslogD 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.)
- **A reboot temporarily suspends Syslog logging.** After a reboot, the switch suspends configured Syslog logging for 30 seconds.

Diagnostic Tools

Diagnostic Features

Feature	Default	Menu	CLI	Web
Port Autonegotiation	n/a	n/a	n/a	n/a
Ping Test	n/a	—	page C-37	page C-36
Link Test	n/a	—	page C-37	page C-36
Display Config File	n/a	—	page C-39	page C-39
Admin. and Troubleshooting Commands	n/a	—	page C-42	—
Factory-Default Config	page C-43 (Buttons)	—	page C-43	—
Port Status	n/a	pages B-9 and B-10	pages B-9 and B-10	pages B-9 and B-10

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. See Chapter 10, “Port Status and Basic 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).

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

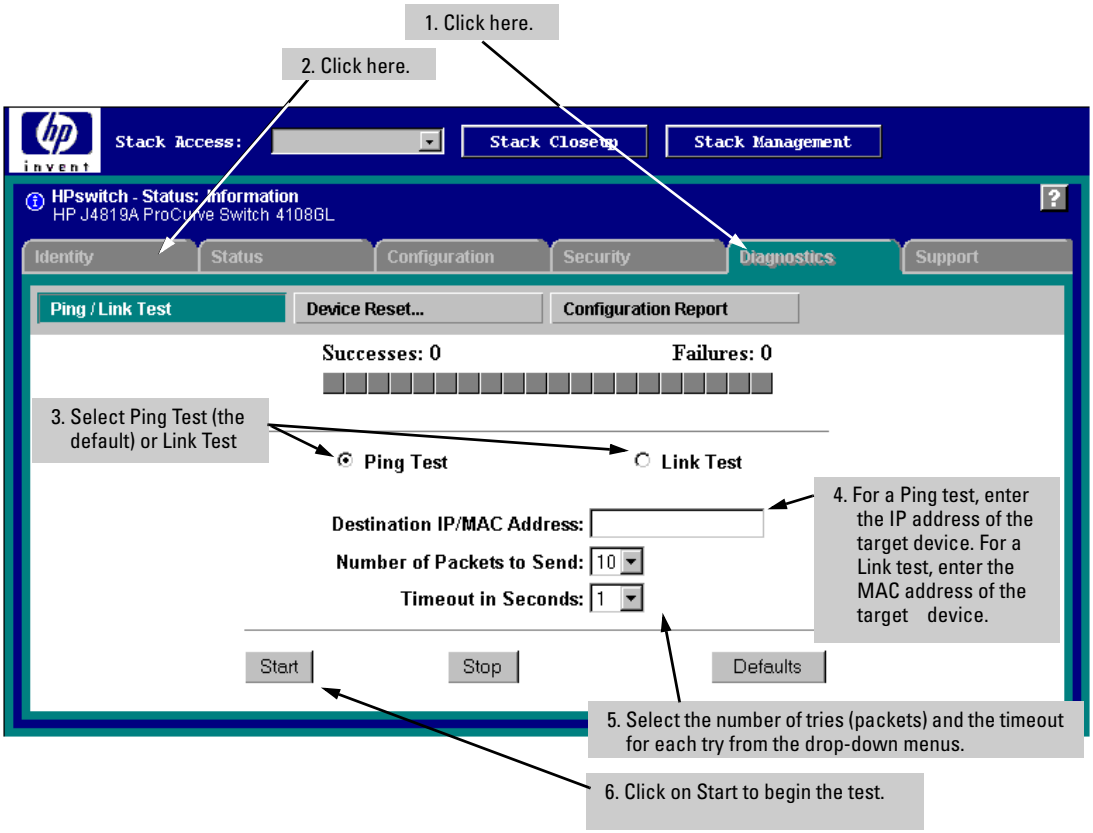

Figure C-14.Link and Ping Test Screen on the Web Browser Interface

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 >]

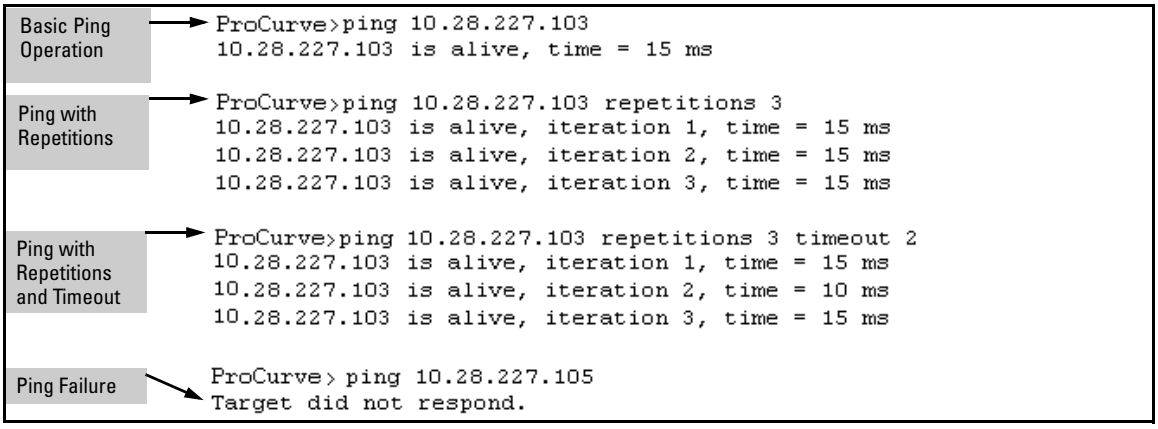

Figure C-15. Examples of Ping Tests

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

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 >]

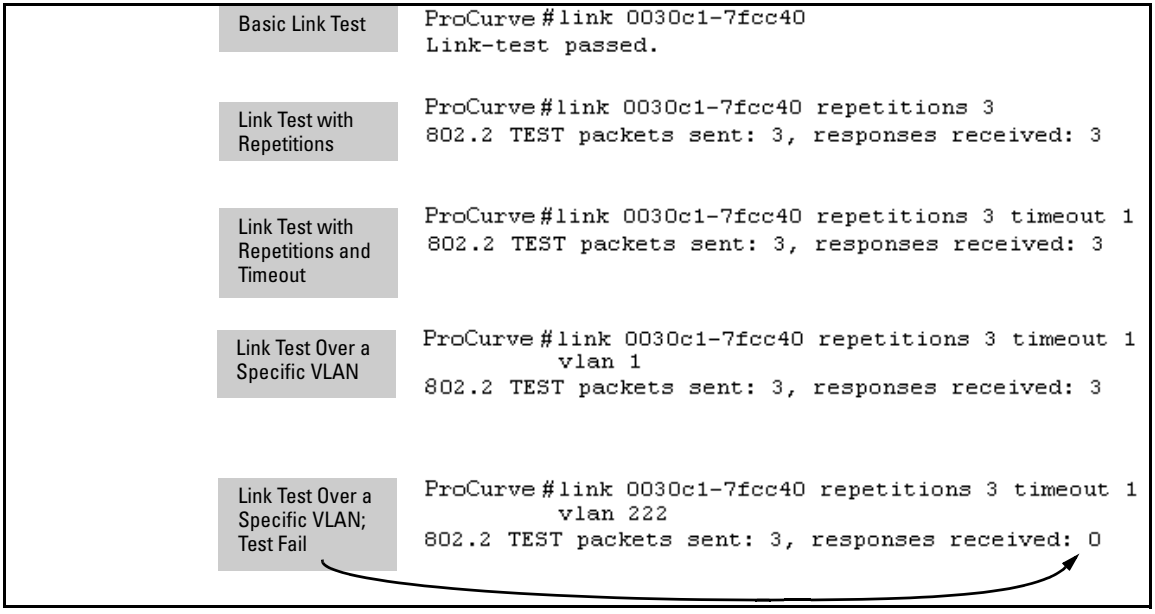

Figure C-16. Example of Link Tests

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, see appendix C, "Switch Memory and Configuration".)

Syntax: write terminal
Displays the running-config file.

show running-config
Displays the running-config file.

show config
Displays the startup-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 for Help in Troubleshooting

Release G.04.05 and greater includes the **show tech** command. This 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)
- Stacking status — this switch
- Stacking status — all

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...**

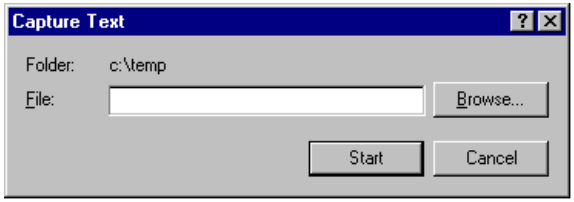

Figure C-17. The Capture Text window of the Hypertext Application Used with Microsoft Windows Software

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

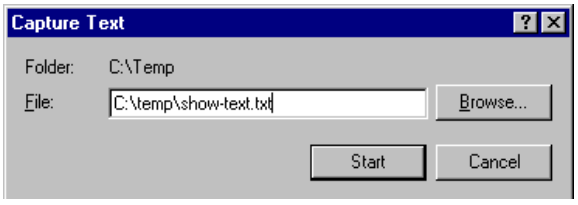

Figure C-18. Example of a Path and Filename for Creating a Text File from show tech Output

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 "Using the Command Line Interface (CLI)" on page 4-1.

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.

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.

Using the CLI

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.

Using the Clear/Reset Buttons

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
   - No parity
   - 8 Bits
   - 1 stop bit
   - No flow control

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:

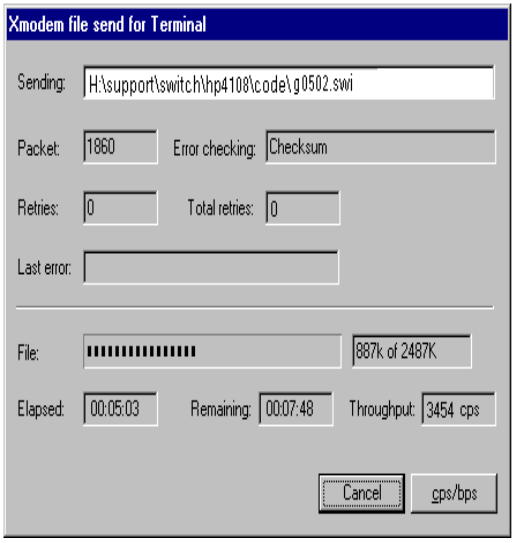

Figure C-19. Example of Xmodem Download in Progress

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.