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 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, see 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.

    See 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. See 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. See 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. 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)
- For help in isolating problems, use the easy-to-access switch console built into the switch or Telnet to the switch console. See chapter 4, "Using the Switch Console Interface" for operating 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)

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. 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:

- 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. Refer to chapter 16, "IP Routing Features", for more information.
- 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, 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.

ACL Problems

Series 5300xl Switches Only: ACLs are properly configured and assigned to VLANs, but the switch is not using the ACLs to filter IP layer 3 packets.

1. The 5300xl switch may be running with IP routing disabled. To ensure that IP routing is enabled, execute **show running** and look for the IP routing statement in the resulting listing. For example:

```
ProCurve(config)# show running

Running configuration:

; J4850A Configuration Editor; Created on release #E.08.01

hostname " ProCurve"
cdp run
module 1 type J4820A
ip default-gateway 10.30.248.1
ip routing
logging 10.28.227.2
snmp-server community "public" Unrestricted
ip access-list extended "Controls for VLAN 20"
   permit tcp 0.0.0.0 255.255.255.255 10.10.20.98 0.0.0.0 eq 80
   permit tcp 0.0.0.0 255.255.255.255 10.10.20.21 0.0.0.0 eq 80
   deny tcp 0.0.0.0 255.255.255.255 10.10.20.1 0.0.0.255 eq 80
   deny tcp 10.10.20.17 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
   deny tcp 10.10.20.23 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
   deny tcp 10.10.20.40 0.0.0.0 10.10.10.100 0.0.0.0 eq 23 log
   permit ip 10.10.20.1 0.0.0.255 10.10.10.100 0.0.0.0
   deny ip 10.10.30.1 0.0.0.255 10.10.10.100 0.0.0.0
   permit ip 10.10.30.1 0.0.0.255 10.10.10.1 0.0.0.255
   exit
-- MORE --, next page: Space, next line: Enter, quit: Control-C
```

Indicates that routing is enabled; a requirement for ACL operation. (There is an exception. See the **Note**, below.)

Figure C-1. Indication that Routing Is Enabled

Note

If an ACL assigned to a VLAN includes an ACE referencing an IP address on the switch itself as a packet source or destination, the ACE screens traffic to or from this switch address regardless of whether IP routing is enabled. This is a security measure designed to help protect the switch from unauthorized management access.

If you need to configure IP routing, execute the **ip routing** command.

2. ACL filtering on the 5300xl switches applies only to routed packets and packets having a destination IP address (DA) on the switch itself. Also, the switch applies assigned ACLs only at the point where traffic enters or leaves the switch on a VLAN. Ensure that you have correctly applied your ACLs ("in" and/or "out") to the appropriate VLAN(s).

The switch does not allow management access from a device on the same VLAN.

The implicit **deny any** function that the switch automatically applies as the last entry in any ACL always blocks packets having the same DA as the switch's IP address on the same VLAN. That is, bridged packets with the switch itself as the destination are blocked as a security measure. To preempt this action, edit the ACL to include an ACE that permits access to the switch's DA on that VLAN from the management device.

Error (Invalid input) when entering an IP address.

When using the "host" option in the command syntax, ensure that you are not including a mask in either dotted decimal or CIDR format. Using the "host" option implies a specific host device and therefore does not permit any mask entry.

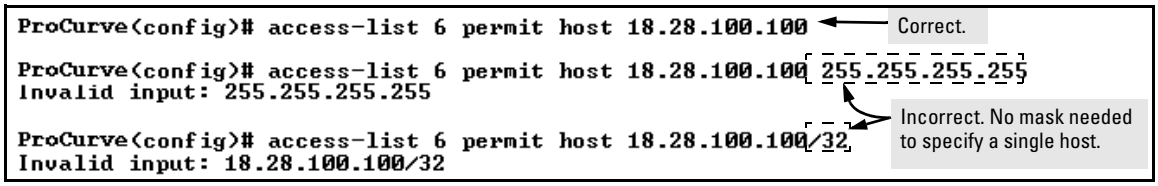

Figure C-2. Examples of Correctly and Incorrectly Specifying a Single Host

Apparent failure to log all "Deny" Matches.

Where the **log** statement is included in multiple ACEs configured with a "deny" option, a large volume of "deny" matches generating logging messages in a short period of time can impact switch performance. If it appears that the switch is not consistently logging all "deny" matches, try reducing the number of logging actions by removing the **log** statement from some ACEs configured with the "deny" action.

The switch does not allow any routed access from a specific host, group of hosts, or subnet.

The implicit **deny any** function that the switch automatically applies as the last entry in any ACL may be blocking all access by devices not specifically permitted by an entry in an ACL affecting those sources. If you are using the ACL to block specific hosts, a group of hosts, or a subnet, but want to allow any access not specifically permitted, insert **permit any** as the last explicit entry in the ACL.

The switch is not performing routing functions on a VLAN

Two possible causes of this problem are:

- Routing is not enabled. If **show running** indicates that routing is not enabled, use the **ip routing** command to enable routing.
- *On a Series 5300xl switch*, an ACL may be blocking access to the VLAN. Ensure that the switch's IP address on the VLAN is not blocked by one of the ACE entries in an ACL applied to that VLAN. A common mistake is to either not explicitly permit the switch's IP address as a

DA or to use a wildcard ACL mask in a deny statement that happens to include the switch's IP address. For an example of this problem, refer to the section titled "General ACL Operating Notes" in the "Access Control Lists (ACLs)" chapter of the Advanced Traffic Management Guide for your switch.

Routing Through a Gateway on the Switch Fails

Configuring a "deny" ACE that includes a gateway address can block traffic attempting to use the gateway as a next-hop.

Remote Gateway Case on a Series 5300xl Switch. For example, configuring ACL "101" (below) and applying it outbound on VLAN 1 in figure C-4 includes the router gateway (10.0.8.1) needed by devices on other networks. This can prevent the switch from sending ARP and other routing messages to the gateway router to support traffic from authorized remote networks.

In figure C-4, this ACE denies access to the 10 Net's 10.0.8.1 router gateway needed by the 20 Net.

(*Subnet* mask is 255.255.255.0.)

```
ProCurve(config)# show access-list config

ip access-list extended "101"
   deny ip 0.0.0.0 255.255.255.255 10.0.8.30 0.0.0.255
   permit ip 0.0.0.0 255.255.255.255 0.0.0.0 255.255.255.255
   exit
```

Figure C-3. Example of ACE Blocking an Entire Subnet

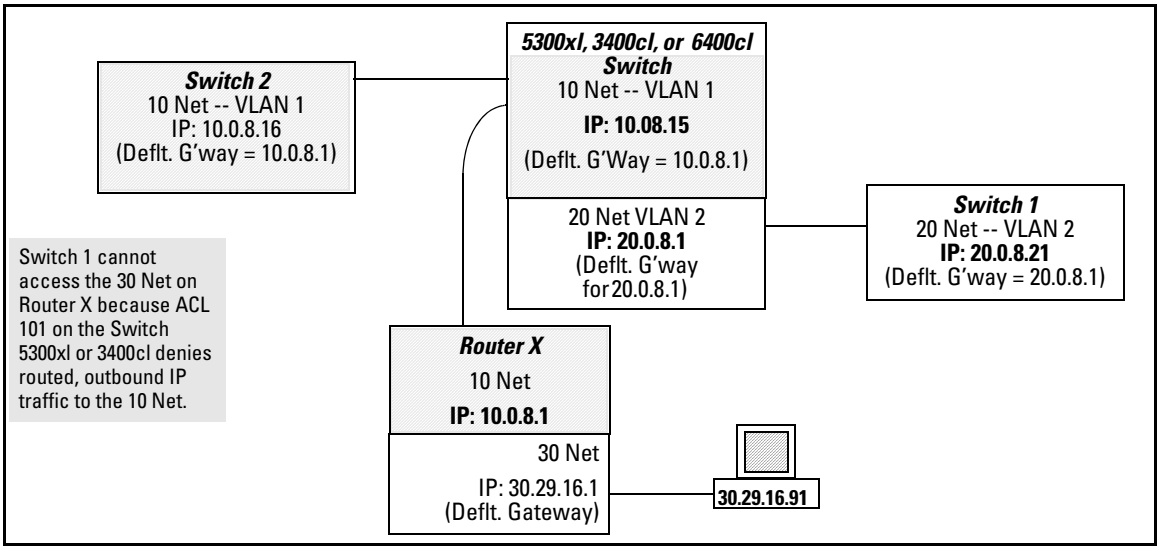

Figure C-4. Example of Inadvertently Blocking a Gateway

To avoid inadvertently blocking the remote gateway for authorized traffic from another network (such as the 20 Net in this example):

1. Configure an ACE that specifically permits authorized traffic from the remote network.
2. Configure narrowly defined ACEs to block unwanted IP traffic that would otherwise use the gateway. Such ACEs might deny traffic for a particular application, particular hosts, or an entire subnet.
3. Configure a "permit any" ACE to specifically allow any IP traffic to move through the gateway.

Local Gateway Case. If you use the switch as a gateway for traffic you want routed between subnets, use these general steps to avoid blocking the gateway for authorized applications:

1. Configure gateway security first for routing with specific permit and deny statements.
2. Permit authorized traffic.
3. Deny any unauthorized traffic that you have not already denied in step 1.

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.

Mesh-Related Problems

Traffic on a dynamic VLAN does not get through the switch mesh .

GVRP enables dynamic VLANs. Ensure that all switches in the mesh have GVRP enabled. (Note that ProCurve 1600M/2400M/2424M/4000M/8000M switches do not offer GVRP. Thus, if there are any of these switches in the mesh, GVRP must be disabled for any switch in the mesh.)

The Switch Mesh Does Not Allow A ProCurve Switch 1600M/2400M/2424M/4000M/8000M Port To Join the Mesh . One of the switches in the mesh domain has detected a duplicate MAC address on multiple switches. For example:

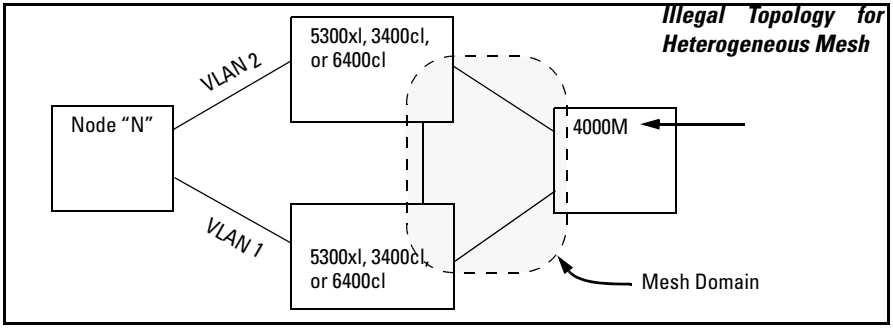

Figure C-5. Example of Illegal Topology for Heterogeneous Mesh

Changing the topology can solve this problem. Also, the duplicate MAC address must age out before the Switch 1600M/2400M/2424M/4000M/8000M port can join the mesh. Refer to the following two topics in the "Switch Meshing" chapter of the *Advanced Traffic Management Guide* for your switch:

- The section titled "Using a Heterogeneous Switch Mesh"
- The bulleted item titled "Compatibility with Older Switches" in the section titled "Requirements and Restrictions".

Duplicate MAC Addresses on Different Switches. In a switch mesh that includes any 1600M/2400M/2424M/4000M/8000M switches, duplicate MAC addresses on different switches are not allowed. (The 1600M/2400M/2424M/4000M/8000M switches do not recognize multiple instances of a particular MAC address on different VLANs.) Refer to "The Switch Mesh Does Not Allow A ProCurve Switch 1600M/2400M/2424M/4000M/8000M Port To Join the Mesh" on page C-14.

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-18.

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 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-6. 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

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

Global RADIUS Encryption Key

Unique RADIUS Encryption Key for the RADIUS server at 10.33.18.119

Figure C-7. 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.

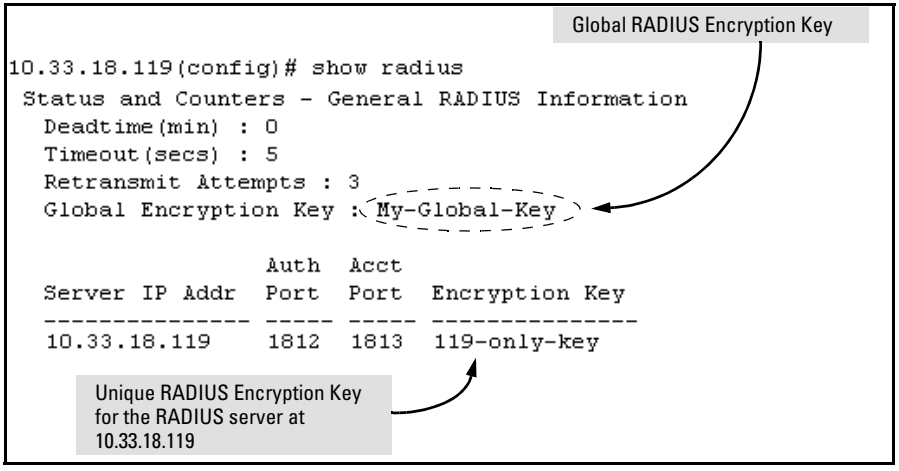

Figure C-8. Examples of Global and Unique Encryption Keys

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 switches 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. 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 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 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.

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".

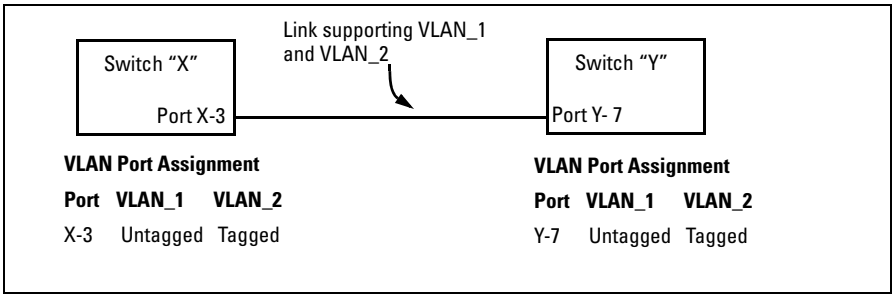

Figure C-9. Example of Correct VLAN Port Assignments on a Link

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 by 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 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 switches covered by this guide 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.

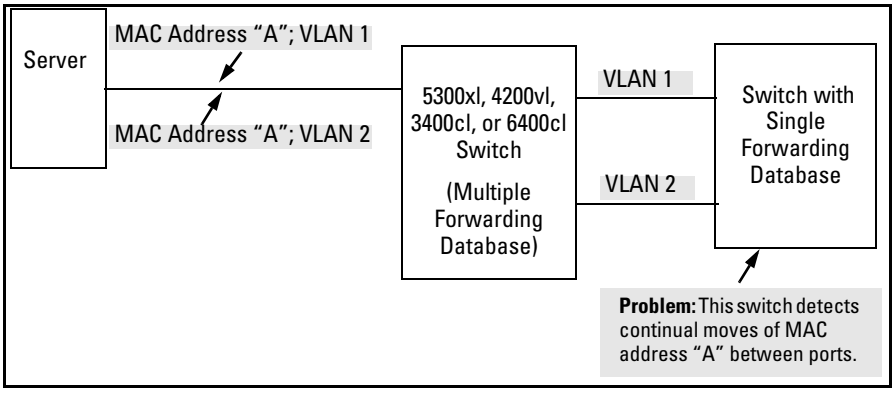

Figure C-10. Example of Duplicate MAC Address

Using the Event Log To Identify Problem Sources

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:

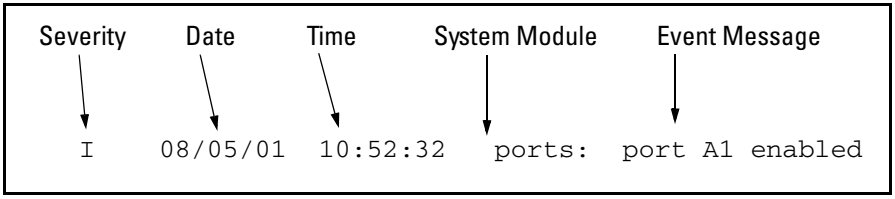

Severity is one of the following codes:

- **I** (information) indicates routine events.
- **W** (warning) indicates that a service has behaved unexpectedly.
- **M** (major) indicates that a severe switch error has occurred.
- **D** (debug) reserved for ProCurve 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-28 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	timep	Time protocol
chassis	switch hardware	udpf	UDP broadcast forwarder
bootp	bootp addressing	vlan	VLAN operations
connfilt	Connection-Rate filtering	RateLim	Rate-limiting
console	Console interface		
dhcp	DHCP addressing		
download	file transfer		
FFI	Find, Fix, and Inform -- available in the console Event Log and web browser interface alert log		
garp	GARP/GVRP		
igmp	IP Multicast		
ip	IP-related		
ipx	Novell Netware		
lacp	Dynamic LACP trunks		
ldbal	Load-Balance Protocol (meshing)		
lldp	Link-Layer Discovery Protocol		
maclock	MAC lockdown and MAC lockout		
mgr	Console management		
PIM	Protocol-Independent multicast		
ports	Change in port status; static trunks		
radius	RADIUS authentication		
snmp	SNMP communications		
ssh	Secure-Shell status		
ssl	Secure sockets layer status		
stp	Spanning Tree		
sys, system	Switch management		
telnet	Telnet activity		
tcp	Transmission control		
tftp	File transfer for new OS or config.		

Menu: Entering and Navigating in the Event Log

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

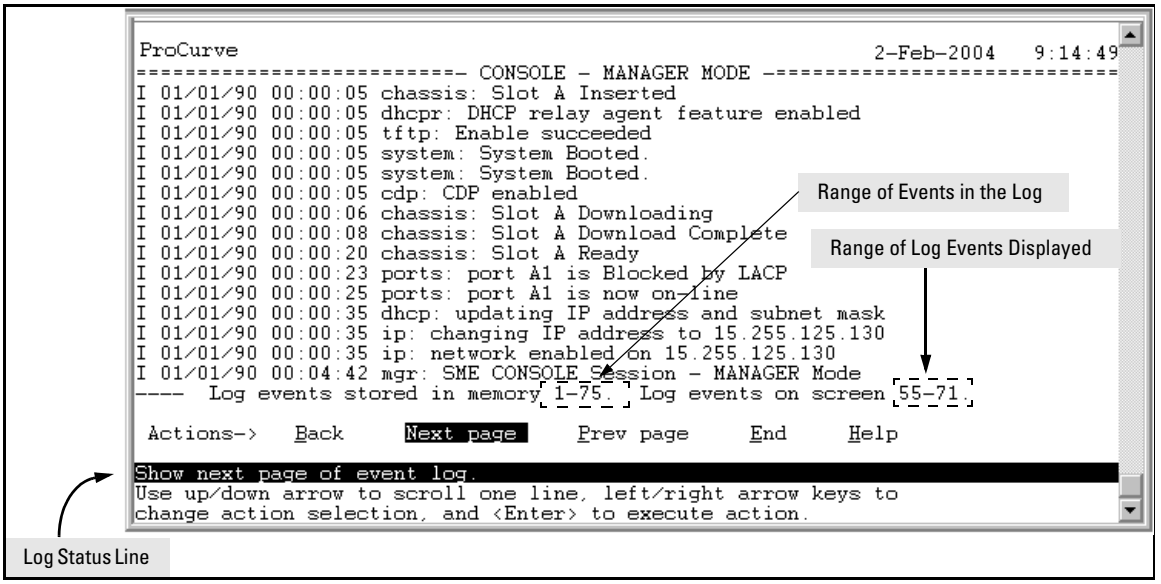

Figure C-11. Example of an Event Log Display

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: Listing Events

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

Uses the CLI to 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*

show logging

Lists recorded log messages since last reboot.

show logging -a

Lists all recorded log messages, including those before the last reboot.

show logging -a system

Lists log messages with "system" in the text or module name.

show logging system

Lists all log messages since the last reboot that have "system" in the text or module name.

Reducing Duplicate Event Log and SNMP Trap Messages

Note

This feature is available with all software releases for the Series 3400/6400cl switches, Series 4200vl switches and with software release E.08.xx and greater on the Series 5300xl switches. Initially it applies only to Event Log messages and SNMP traps generated by the PIM software module.

A recurring event can generate a series of duplicate Event Log messages and SNMP traps in a relatively short time. This can flood the Event Log and any configured SNMP trap receivers with excessive, exactly identical messages. To help reduce this problem, the switch uses *log throttle periods* to regulate (throttle) duplicate messages for a given recurring event, and maintains a counter to record how many times it detects duplicates of a particular event since the last system reboot. That is, 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, then 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 one instance of the corresponding message in the Event Log for each successive log throttle period applied to recurrences of that event. 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.

The log throttle period for an event depends on the event's severity level:

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

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

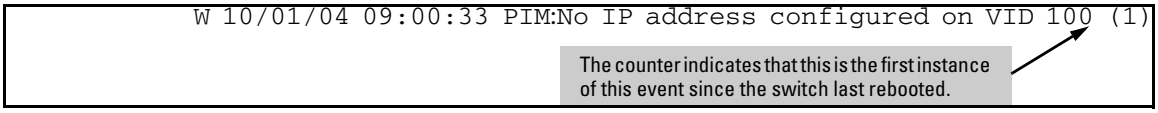

Figure C-12. Example of the First Instance of an Event Message and Counter

If PIM 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.

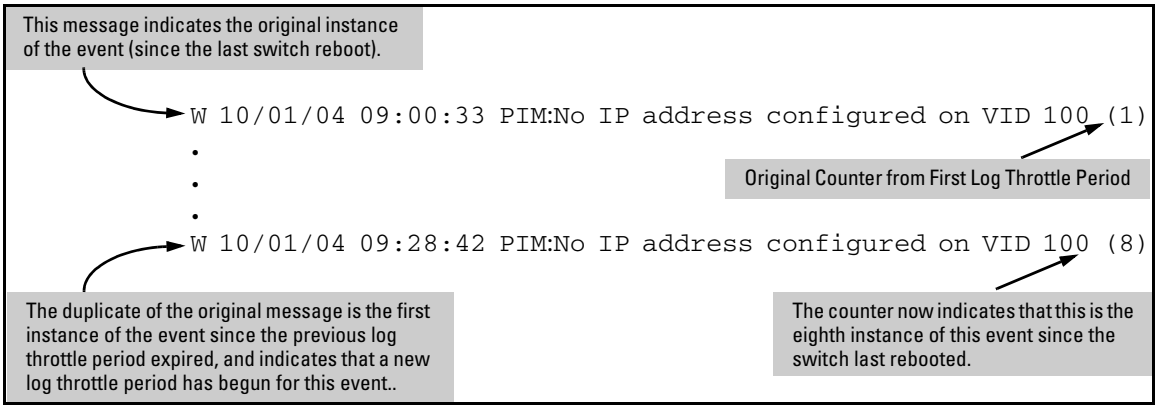

Figure C-13. Example of Duplicate Messages Over Multiple Log Throttling Periods

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 PIM operation simultaneously detected that VLANs 100 and 205 were configured without IP addresses, you would see log messages similar to the following:

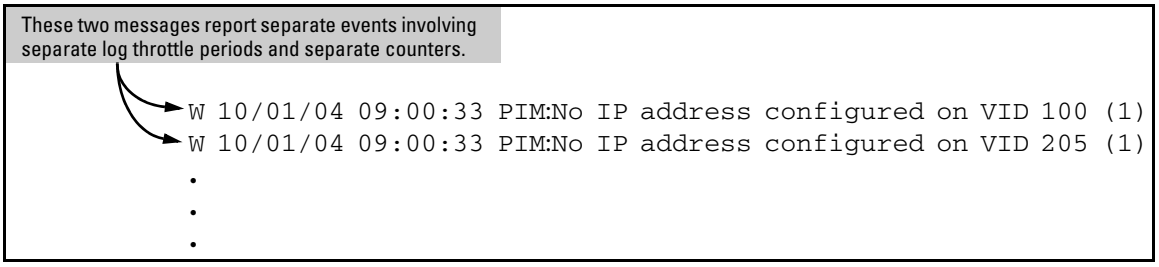

Figure C-14. Example of Log Messages Generated by Unrelated Events of the Same Type

Example of Event Counter Operation. Suppose the switch detects the following after a reboot:

- Three duplicate instances of the PIM "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 and Syslog Messaging Operation

The switch's Event Log records switch-level progress, status, and warning messages. The Debug/System-Logging (*Syslog*) feature provides a method for recording messages you can use to help in debugging network-level problems, such as routing misconfigurations and other network protocol details.

Debug enables you to specify the types of event notification messages to send to external devices. Debug messaging reports on these event types:

- ACL "deny" matches
- Selected IP routing events
- Events that generate messages for the switch's Event Log

You can configure the switch to send debug messages to these destinations:

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

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 file.
Optional **Debug** Commands	**all**		Assigns debug logging to the configured debug destination(s) for all ACL, Event Log, IP-OSPF, and IP-RIP options.
	acl		Assigns ACL Syslog logging to the debug destination(s). When there is a match with a "deny" ACE, directs the resulting message to the configured debug destination(s).
	event		Assigns standard Event Log messages to the debug destination(s). (The same messages are also sent to the switch's Event Log, regardless of whether you enable this option.)
	IP		
		ospf	Assigns OSPF event logging to the debug destination(s).
		rip	Assigns RIP event logging to the debug destination(s).
	lldp		Assigns LLDP debug logging to the debug destination(s).
Debug Destinations	**Destination**	**logging**	Used to disable or re-enable Syslog logging if one or more Syslog servers are already configured by the separate **logging < *ip-addr*>** command. Optionally, also specifies the destination (facility) subsystem the Syslog servers must use.
		session	Assigns or re-assigns destination status to the terminal device most recently using this command to request debug output.

Figure C-15. Event Messaging Structure

Debug logging requires a logging destination (Syslog 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 Syslog servers. Included is the option to send the messages to the User log facility (default) on the configured server(s) or to another log facility.

Note

As of September 2004, the **logging facility** <*facility-name*> option (described on page C-40) is available on these switch models:

- Series 5300xl switches (software release E.08.xx or greater)
- Series 4200vl switches
- Series 4100gl switches (software release G.07.50 or greater)
- Series 3400cl switches
- Series 6400cl 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, this includes the Syslog server list.
- Display the current Syslog server list when Syslog logging is disabled.

Debug Command Operation

As shown in figure C-15, the **debug** command performs two main functions:

- Specifies the type(s) of event messaging to send to a destination.
- Specifies the destination(s) of the selected message types.

Except as noted below, rebooting the switch returns the debug destination and debug message types to their default settings (disabled).

Note Using the **logging < *dest-ip-addr* >** command to configure a Syslog server address creates an exception to the above general operation. Refer to "Syslog Operation" on page C-39.

Debug Types

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

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

acl

When a match occurs on an ACL "deny" Access Control Entry (with **log** *configured), the switch sends an ACL message to the configured debug destination(s). For more on ACLs, refer to the chapter titled "Access Control Lists" in the Advanced Traffic Management Guide for your switch. (Default: Disabled)*

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.*

ip

Enables all IP-OSPF message types for the configured destinations.

lldp

Enables all LLDP message types for the configured destinations.

Syntax: [no] debug < *debug-type* > *(Continued)*

ip [ospf < adj | event | flood | lsa-generation | packet | retransmission | spf >]

For the configured debug destination(s):

ospf < adj | event | flood | lsa-generation | packet | retransmission | spf > — *Enables the specified IP-OSPF message type.*

adj — *Adjacency changes.*
event — *OSPF events.*
flood — *Information on flood messages.*
lsa-generation — *New LSAs added to database.*
packet — *Packets sent/received.*
retransmission — *Retransmission timer messages.*
spf — *Path recalculation messages.*

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.*

(Default: Event (log) message type.)

Debug Destinations

Debug enables you to disable and re-enable Syslog messaging to configured servers, and to designate a CLI session to receive messaging of any debug type.

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

logging

This command enables Syslog logging to the configured Syslog server(s). That is, the switch sends the debug message types (specified by the **debug < *debug-type* >** *command in the preceding section) to the configured Syslog server(s). (Default: Logging disabled)*

(To configure a Syslog server IP address, refer to "Syslog Operation" on page C-39.)

Note: *Debug messages from Series 5300xl switches running software release E.07.21 or greater, Series 4200vl switches, and any 3400cl/6400cl switches carry a "debug" severity level. Because some Syslog servers, in their default configuration, ignore Syslog messages with this severity level, you should ensure that the Syslog servers you intend to receive debug messages are configured to accept the "debug" severity level. For more information, refer to* "Operating Notes for Debug and Syslog" on page C-44.

session

*Enables or disables 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 those specified by the **debug < *debug-type* >** *command. (Refer to* "Debug Types" on page C-36.*)*

Syslog Operation

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. Use of Syslog requires that you set up a Syslog server application on a networked host accessible to the switch. (Refer to the documentation for the Syslog server application you select.) Except as described below, you must use the **debug** command to specify the message types the switch sends to the configured Syslog server(s).

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. Thus, at a minimum, the switch begins sending Event Log messages to the configured Syslog server(s). The ACL, IP-OSPF, and/or IP-RIP message types will also be sent to the Syslog server(s) if they are currently enabled debug types. (Refer to "Debug Types" on page C-36*.)*

no logging *removes all currently configured Syslog logging destinations from the switch.*

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

*If you use the "***no***" form of the command to delete the only remaining logging address, debug destination logging is disabled on the switch, but the Event debug type is not changed from its current setting.*

To block messages to the configured Syslog server(s) from any currently enabled debug type, use **no debug < *debug-type* >**. *(Refer to* "Debug Types" on page C-36*.)*

To disable Syslog logging on the switch without deleting the server addresses, use **no debug destination logging**. *Note that, unlike the case where there are no Syslog servers configured, if one or more Syslog servers are already configured, but Syslog messaging is disabled, adding a new server address to those already configured does not re-enable Syslog messaging. In this case, you must use* **debug destination logging** *to re-enable Syslog messaging.*

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

The logging facility specifies the destination subsystem the Syslog server(s) must use. (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 *(the 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*

For a listing of applicable ProCurve switches, refer to the Note on page C-35.

Viewing the Debug Configuration

Syntax: show debug

This command displays the currently configured debug logging destination(s) and type(s). For examples of show debug output, refer to figure C-16 on page C-42.

Steps for Configuring Debug and Syslog Messaging

1. Skip this step if you do not want to use a Syslog server.

   If you want a Syslog server as a destination for debug messaging:

   a. Use this command to configure the Syslog server IP address and enable Syslog logging:

      ```
      ProCurve(config)# logging < ip-addr >
      ```

Using this command when there are no Syslog server IP addresses already configured enables both debug messaging to a Syslog server and the Event debug-type, which means that the switch begins sending Event Log messages to the server, regardless of other debug types that may be configured.

b. Use the command in step "a" to configure any additional Syslog servers you want to use, up to a total of six. (When multiple server IP addresses are configured, the switch sends the selected debug message types to all such addresses.)

c. If you want Event Log messages sent to the Syslog server, skip this step. Otherwise, use this command to block Event Log messages to the server:

```
ProCurve# no debug event
```

2. If you do not want a CLI session for a destination, skip this step.

   Otherwise, from the device to which you want the switch to send debug messages:

   a. Use a serial, Telnet, or SSH connection to access the switch's CLI.

   b. Execute this command:

```
ProCurve# debug destination session
```

3. Enable the debug types for which you want messages sent to the Syslog server(s) and/or the current session device:

   `ProCurve# debug` < acl | all | event | ip [ospf-opt]>

   Repeat this step if necessary to enable multiple debug types.

Example: Suppose that there are no Syslog servers configured on the switch (the default). Configuring one Syslog server enables debug logging to that server and also enables Event Log messages to be sent to the server.

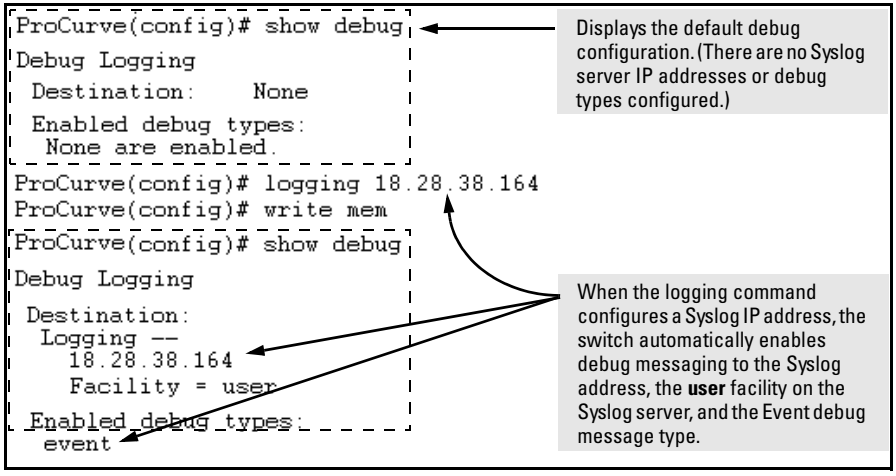

Figure C-16. Example of Configuring Basic Syslog Operation

Note that after you enable Syslog logging, if you do not want Event Log messages sent to the Syslog server(s), you can block such messages by executing **no debug event**. (This has no effect on standard logging of messages in the switch's Event Log.)

Example. Suppose that you want to:

- Configure Syslog logging of ACL and IP-OSPF packet messages on a Syslog server at 18.38.64.164 (with **user** as the default logging facility).
- Also display these messages in the CLI session of your terminal device's management access to the switch.
- Prevent the Switch's standard Event Log messages from going to the Syslog server and the CLI.

Assuming the debug/Syslog feature is disabled on the switch, you would use the commands shown in figure C-17 to configure the above operation.

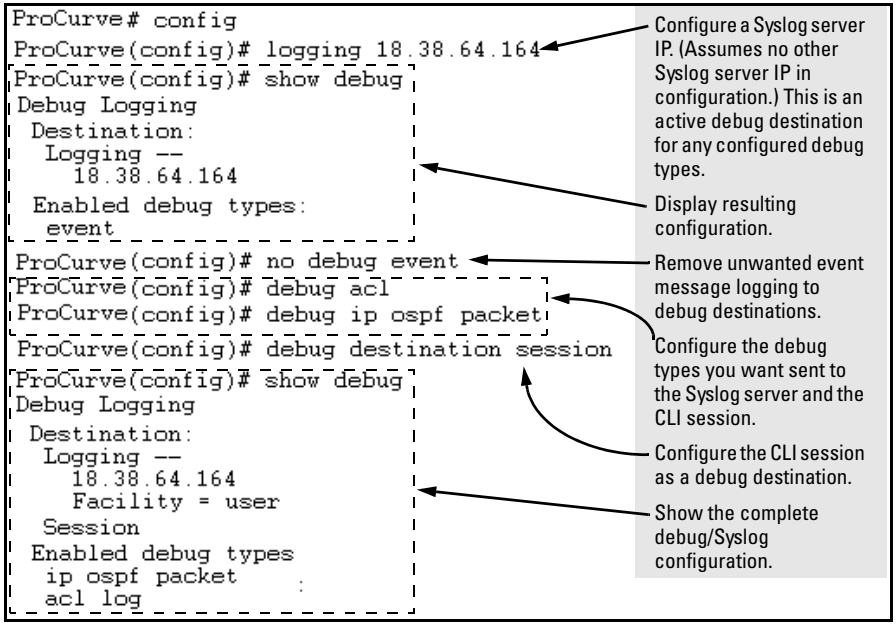

Figure C-17. Example Debug/Syslog Configuration for Multiple Types and Destinations

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 (destination)	If any Syslog 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.
ACL (event type)	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.
IP (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 resulting from debug operation carry a “debug” severity. If you configure the switch to transmit 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-48	page C-47
Link Test	n/a	—	page C-48	page C-47
Display Config File	n/a	—	page C-50	page C-50
Admin. and Troubleshooting Commands	n/a	—	page C-52	—
Factory-Default Config	page C-57 (Buttons)	—	page C-57	—
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. Refer to 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

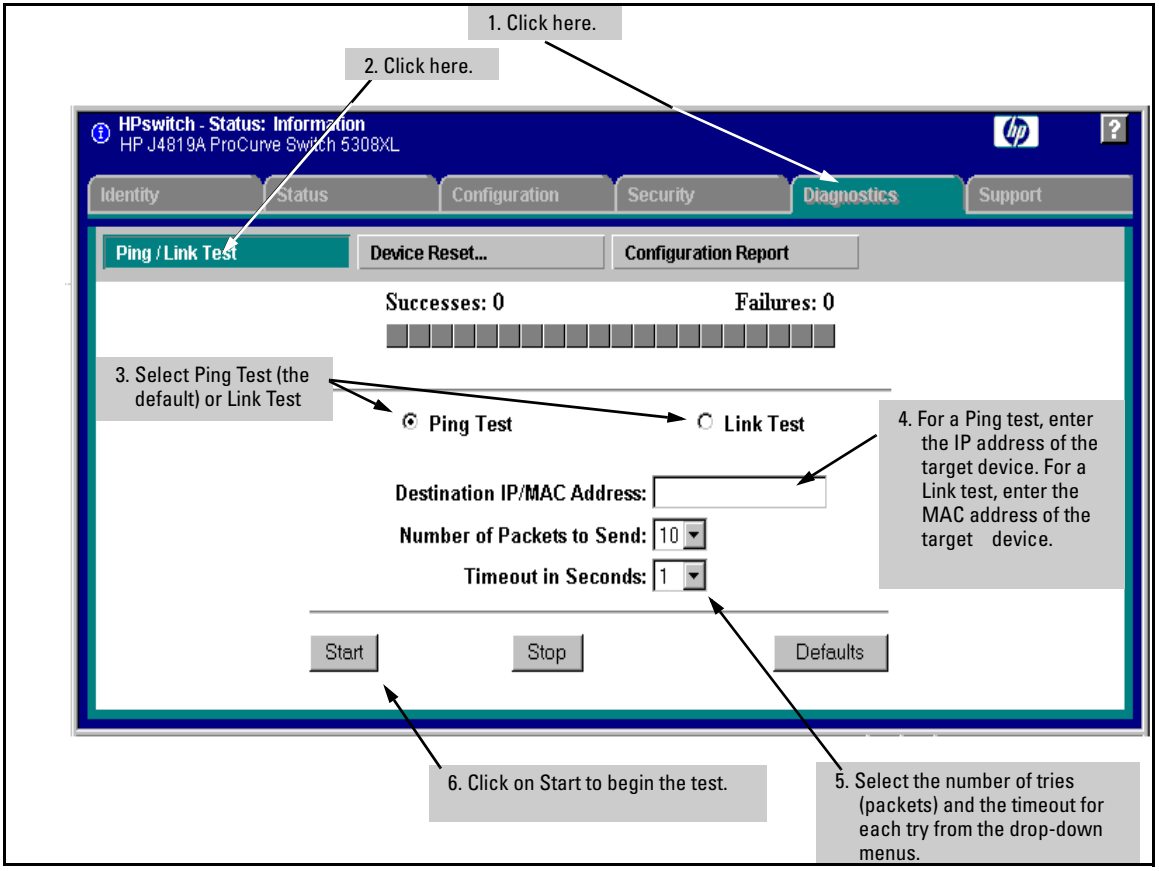

Figure C-18. Link and Ping Test Screen on the Web Browser Interface

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 >]

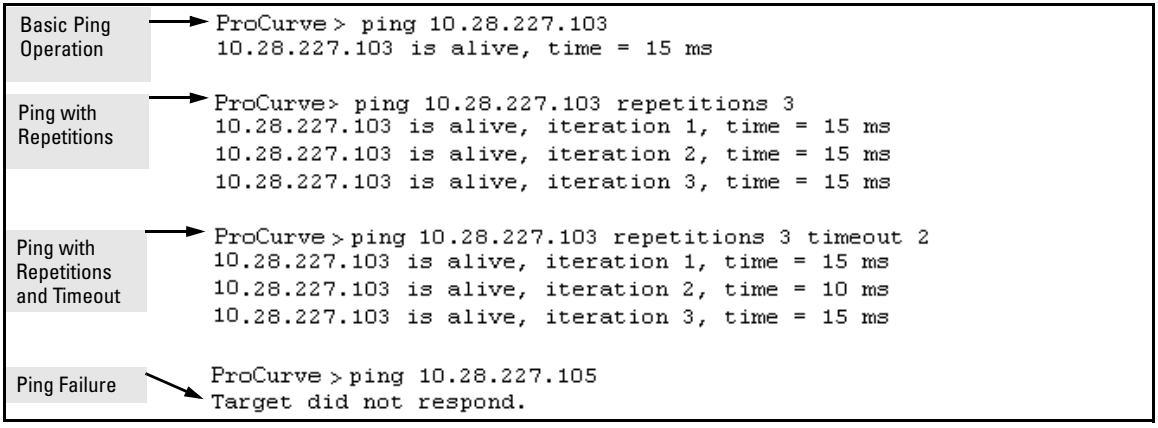

Figure C-19. Examples of Ping Tests

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* >]

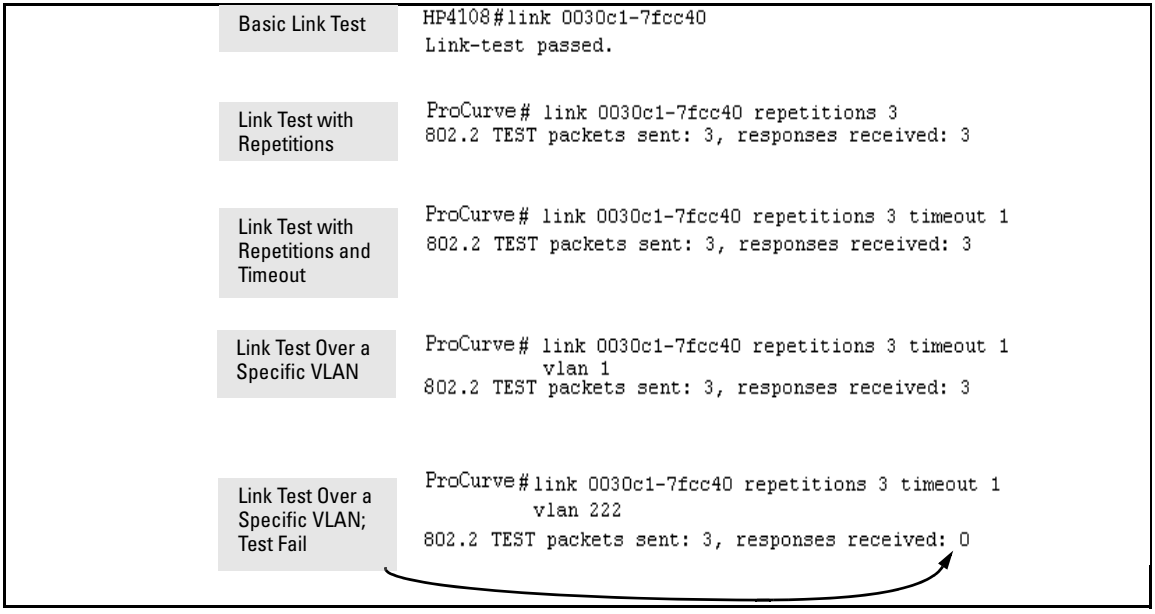

Figure C-20. Example of Link Tests

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 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...**

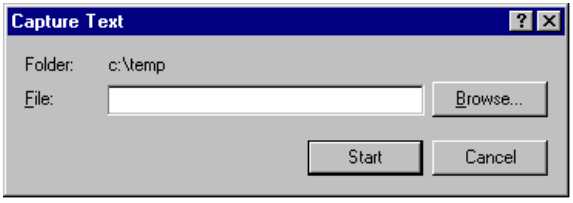

Figure C-21. The Capture Text window of the Hyperterminal Application

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

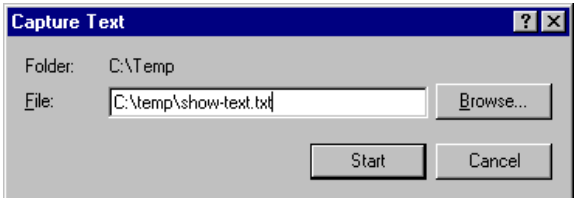

Figure C-22. Example of a Path and Filename for Creating a Text File from show tech Output

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 3, "Using the Command Line Reference (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.

Syntax: show version

[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:

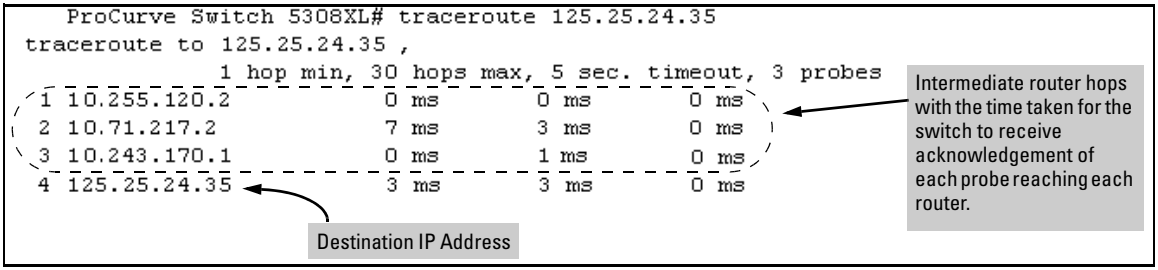

Figure C-23. Example of a Completed Traceroute Enquiry

Continuing from the previous example (figure C-23, above), executing **traceroute** with an insufficient **maxttl** for the actual hop count produces an output similar to this:

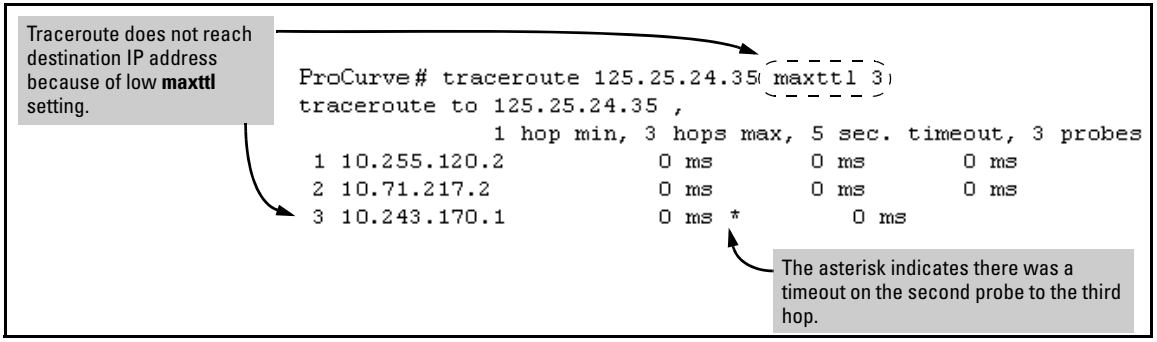

Figure C-24. Example of Incomplete Traceroute Due to Low Maxttl Setting

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; see figure C-24, 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:

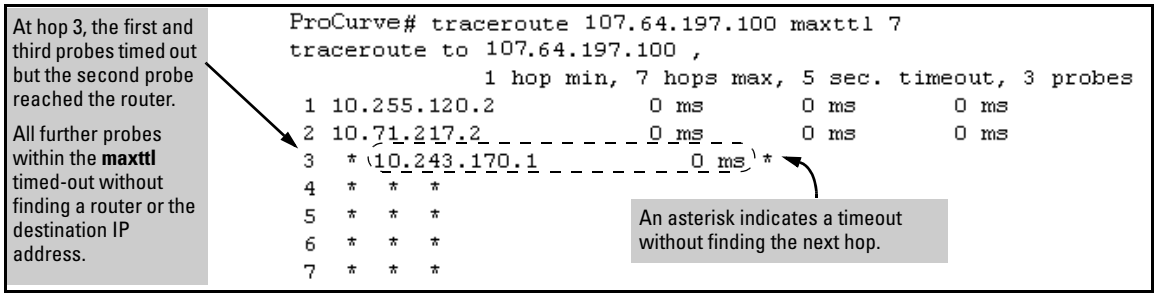

Figure C-25. Example of Traceroute Failing to Reach the Destination Address

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
   - 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:

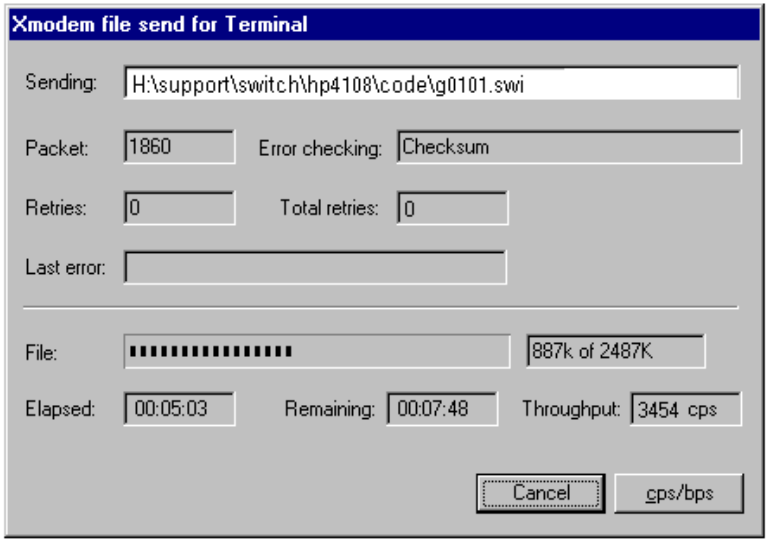

Figure C-26. Example of Xmodem Download in Progress

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.