9

Configuring and Monitoring Port Security

Contents

Contents 9-1

Overview 9-2

- Basic Operation 9-2
- Blocking Unauthorized Traffic 9-3
- Trunk Group Exclusion 9-4

Planning Port Security 9-5

Port Security Command Options and Operation 9-6

- Retention of Static MAC Addresses 9-10
- Displaying Current Port Security Settings 9-10
- Configuring Port Security 9-12

MAC Lockdown 9-17

- Differences Between MAC Lockdown and Port Security 9-19
- Deploying MAC Lockdown 9-21

MAC Lockout 9-25

- Port Security and MAC Lockout 9-27

IP Lockdown 9-28

Web: Displaying and Configuring Port Security Features 9-29

Reading Intrusion Alerts and Resetting Alert Flags 9-29

- Notice of Security Violations 9-29
- How the Intrusion Log Operates 9-30
- Keeping the Intrusion Log Current by Resetting Alert Flags 9-31
- Using the Event Log To Find Intrusion Alerts 9-36
- Web: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags 9-36

Operating Notes for Port Security 9-37

Overview

Feature	Default	Menu	CLI	Web
Displaying Current Port Security	n/a	—	page 9-10	page 9-29
Configuring Port Security	disabled	—	page 9-12	page 9-29
Intrusion Alerts and Alert Flags	n/a	page 9-36	page 9-34	page 9-36

Using Port Security, you can configure each switch port with a unique list of the MAC addresses of devices that are authorized to access the network through that port. This enables individual ports to detect, prevent, and log attempts by unauthorized devices to communicate through the switch.

Note

This feature does not prevent intruders from receiving broadcast and multicast traffic.

Basic Operation

Default Port Security Operation. The default port security setting for each port is off, or **continuous**. That is, any device can access a port without causing a security reaction.

Intruder Protection. A port that detects an "intruder" blocks the intruding device from transmitting to the network through that port.

General Operation for Port Security. On a per-port basis, you can configure security measures to block unauthorized devices, and to send notice of security violations. Once you have configured port security, you can then monitor the network for security violations through one or more of the following:

- Alert flags that are captured by network management tools
- Alert Log entries in the switch's web browser interface
- Event Log entries in the console interface
- Intrusion Log entries in either the menu interface, CLI, or web browser interface

For any port, you can configure the following:

- **Authorized (MAC) Addresses:** Specify up to eight devices (MAC addresses) that are allowed to send inbound traffic through the port. This feature:
  - Closes the port to inbound traffic from any unauthorized devices that are connected to the port.
  - Provides the option for sending an SNMP trap notifying of an attempted security violation to a network management station and, optionally, disables the port. (For more on configuring the switch for SNMP management, refer to "Trap Receivers and Authentication Traps" in the *Management and Configuration Guide* for your switch.)

Blocking Unauthorized Traffic

Unless you configure the switch to disable a port on which a security violation is detected, the switch security measures block unauthorized traffic without disabling the port. This implementation enables you to apply the security configuration to ports on which hubs, switches, or other devices are connected, and to maintain security while also maintaining network access to authorized users. For example:

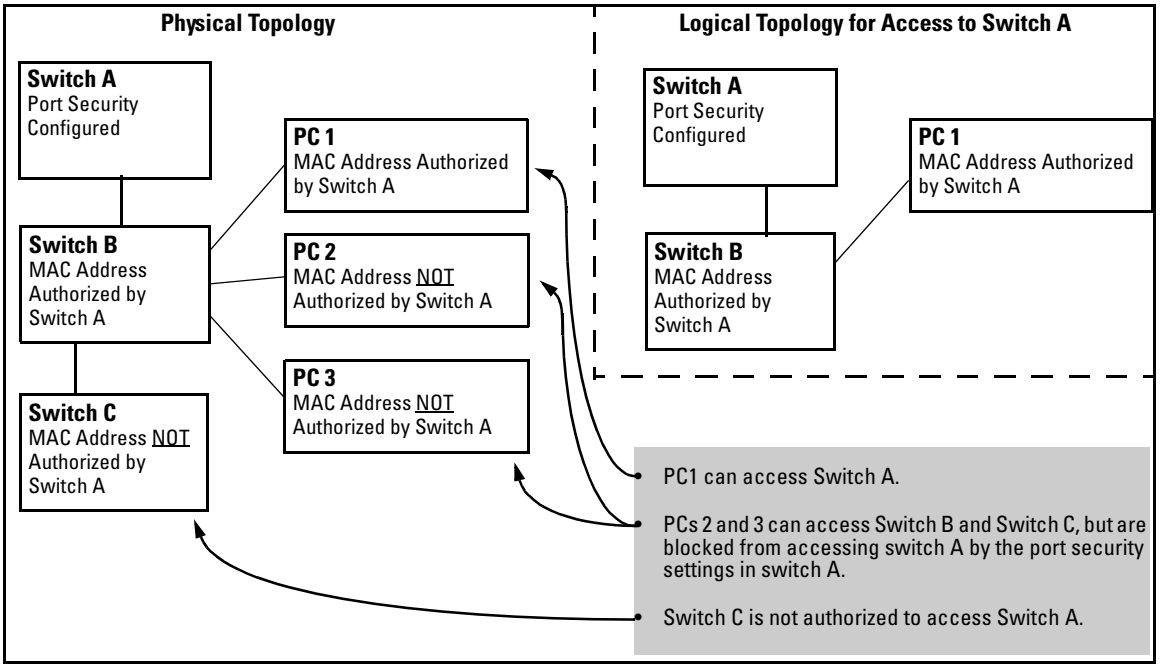

Figure 9-1. Example of How Port Security Controls Access

Note

Broadcast and Multicast traffic is not "unauthorized" traffic, and can be read by intruders connected to a port on which you have configured port security.

Trunk Group Exclusion

Port security does not operate on either a static or dynamic trunk group. If you configure port security on one or more ports that are later added to a trunk group, the switch will reset the port security parameters for those ports to the factory-default configuration. (Ports configured for either Active or Passive LACP, and which are not members of a trunk, can be configured for port security.)

Planning Port Security

1. Plan your port security configuration and monitoring according to the following:
   a. On which ports do you want port security?
   b. Which devices (MAC addresses) are authorized on each port and how many devices do you want to allow per port (up to 8)?
   c. Within the devices-per-port limit, do you want to let the switch automatically accept devices it detects on a port, or do you want it to accept only the devices you explicitly specify? (For example, if you allow three devices on a given port, but specify only one MAC address for that port, do you want the switch to automatically accept the first two additional devices it detects, or not?)
   d. For each port, what security actions do you want? (The switch automatically blocks intruders detected on that port from transmitting to the network.) You can configure the switch to (1) send intrusion alarms to an SNMP management station and to (2) optionally disable the port on which the intrusion was detected.
   e. How do you want to learn of the security violation attempts the switch detects? You can use one or more of these methods:
      – Through network management (That is, do you want an SNMP trap sent to a net management station when a port detects a security violation attempt?)
      – Through the switch's Intrusion Log, available through the CLI, menu, and web browser interface
      – Through the Event Log (in the menu interface or through the CLI **show log** command)
2. Use the CLI or web browser interface to configure port security operating and address controls. The following table describes the parameters.

Port Security Command Options and Operation

Port Security Commands Used in This Section

show port-security	9-11
port-security	9-12
< [ethernet] *port-list* >	9-12
[learn-mode]	9-12
[address-limit]	9-12
[mac-address]	9-12
[action]	9-12
[clear-intrusion-flag]	9-12
no port-security	9-12

This section describes the CLI port security command and how the switch acquires and maintains authorized addresses.

Note

Use the global configuration level to execute port-security configuration commands.

Syntax: port-security [e] < *port-list* >

learn-mode < continuous | static | configured | port-access >

Continuous *(Default)*: *Appears in the factory-default setting or when you execute* **no port-security.** *Allows the port to learn addresses from inbound traffic from any device(s) to which it is connected. In this state, the port accepts traffic from any device(s) to which it is connected. Addresses learned this way appear in the switch and port address tables and age out according to the* **MAC Age Interval** *in the System Information configuration screen of the Menu interface or the* **show system-information** *listing.*

Static: *The static-learn option enables you to use the* **mac-address** *parameter to specify the MAC addresses of the devices authorized for a port, and the* **address-limit** *parameter to specify the number of MAC addresses authorized for the port. You can authorize specific devices for the port, while still allowing the port to accept other, non-specified devices until the port reaches the configured address limit. That is, if you enter fewer MAC addresses than you authorized, the port fills the remainder of the address allowance with MAC addresses it automatically learns. For example, if you specify three authorized devices, but enter only one authorized MAC address, the port adds the one specifically authorized MAC address to its authorized-devices list and the first two additional MAC addresses it detects. If, for example:*

- *You authorize MAC address* **0060b0-880a80** *on port A4.*
- *You allow three devices on port A4, but the port detects these MAC addresses:*

  1. **080090-1362f2** *3.* **080071-0c45a1**

  2. **00f031-423fc1** *4.* **0060b0-880a80** *(the authorized address.)*

Port A4 then has the following list of authorized addresses:

080090-1362f2 *(The first address detected.)*
00f031-423fc1 *(The second address detected.)*
0060b0-880a80 *(The authorized address.)*

The remaining MAC address, **080071-0c45a1**, *is an intruder.* *See also "Retention of Static Addresses" on page 9-10.*

Caution: When you use **learn-mode static** with a device limit greater than the number of MAC addresses you specify with mac-address, an unwanted device can become "authorized". This can occur because the port, in order to fulfill the number of devices allowed by **address-limit**, automatically adds devices it detects until it reaches the specified limit.

Syntax: port-security [e] < *port-list* > *(- **Continued** -)*

learn-mode < continuous | static | configured | port-access > *(- **Continued** -)*

Configured: *The static-configured option operates the same as the static-learn option on the preceding page, except that it does not allow the switch to accept non-specified addresses to reach the address limit. Thus, if you configure an address limit of 3, but only configure two MAC addresses, the switch will handle as intruders all non-specified MAC addresses it detects.*

Note: As of September, 2003, this option is available in the ProCurve Switch 2600 Series and the Switch 6108 running software release H.07.30 (or greater), and the ProCurve Switch 2800 Series. For availability in other switch products, refer to the latest release notes for such products on the ProCurve Networking website. (Refer to "Getting Documentation From the Web" on page 1-9.)

Port-Access: *Enables you to use Port Security with (802.1X) Port-Based Access Control. Refer to "Configuring Port-Based Access Control (802.1X)" on page 8-1.*

address-limit < *integer* >

When Learn Mode is set to **static** (static-learn) *or* **configured** (static-configured), *this parameter specifies the number of authorized devices (MAC addresses) to allow. Default: 1; Range: 1 to 8.*

mac-address < *mac-addr* >

Available for **static** *(static-learn and configured-learn) modes. Allows up to eight authorized devices (MAC addresses) per port, depending on the value specified in the* **address-limit** *parameter.*

- *If you use* **mac-address** *with* **learn-mode configured**, *but enter fewer devices than you specified in the address-limit field, the port accepts only the devices you specified with* **mac-address**. *(See the* **Note**, *above.)*
- *If you use* **mac-address** *with* **learn-mode static**, *but enter fewer devices than you specified in the address-limit field, the port accepts the specified devices AND as many other devices as it takes to reach the device limit.*

Syntax: port-security [e] < *port-list* > *(- **Continued** -)*

action < none | send-alarm | send-disable >

Specifies whether an SNMP trap is sent to a network management station. Operates when:

- *Learn mode is set to* **learn-mode static** *(static-learn) or* **learn-mode configured** *(static-configured) and the port detects an unauthorized device.*
- *Learn mode is set to* **learn-mode continuous** *and there is a MAC address change on a port.*

none *(the default): Prevents an SNMP trap from being sent.*

send alarm: *Causes the switch to send an SNMP trap to a network management station.*

send-disable: *Available only with* **learn-mode configured** and **learn-mode static**. *Causes the switch to send an SNMP trap to a network management station and disable the port. If you subsequently re-enable the port without clearing the port's intrusion flag, the port will block further intruders, but the switch will not disable the port again until you reset the intrusion flag. See the* **Note** *on page 9-31.*

For information on configuring the switch for SNMP management, refer to the Management and Configuration Guide for your switch.

clear-intrusion-flag

Clears the intrusion flag for a specific port. (Refer to "Reading Intrusion Alerts and Resetting Alert Flags" on page 9-29.)

Retention of Static MAC Addresses

Learned MAC Addresses

In the following two cases, a port in Static learn mode (**learn-mode static**) retains a learned MAC address even if you later reboot the switch or disable port security for that port:

- The port learns a MAC address after you configure the port with l**earn-mode static** in both the startup-config file and the running-config file (by executing **write memory**).
- The port learns a MAC address after you configure the port with **learn-mode static** in only the running-config file and, after the address is learned, you execute **write memory** to configure the startup-config file to match the running-config file.

Assigned/Authorized MAC Addresses

If you manually assign a MAC address (using **mac-address < *mac-addr* >**) and then execute **write memory**, the assigned MAC address remains in memory unless removed by one of the methods described below.

Removing Learned and Assigned Static MAC Addresses

To remove a static MAC address, do one of the following:

- Delete the address by using **no port-security < *port-number* > mac-address < *mac-addr* >**.
- Download a configuration file that does not include the unwanted MAC address assignment.
- Reset the switch to its factory-default configuration.

Displaying Current Port Security Settings

The CLI uses the same command to provide two types of port security listings:

- All ports on the switch with their Learn Mode and (alarm) Action
- Only the specified ports with their Learn Mode, Address Limit, (alarm) Action, and Authorized Addresses

Using the CLI To Display Port Security Settings.

Syntax: show port-security
show port-security [e] <*port number*>
show port-security [e] [<*port number*>-<*port number*]. . .[,<*port number*>]

Without port parameters, **show port-security** displays operating control settings for all ports on a switch. For example:

```
ProCurve(config)# show port-security
 Port Security
  Port Learn Mode |          Action
  ---- ---------- + ------------------------
  A1 1  Static     | Send Alarm, Disable Port
  A2 2  Static     | Send Alarm, Disable Port
  A3 3  Static     | Send Alarm
  A4 4  Static     | Send Alarm
  A5 5  Static     | Send Alarm
  A6 6  Static     | Send Alarm
  A7 7  Continuous | None
  A8 8  Continuous | None
```

Figure 9-2. Example Port Security Listing (Ports A7 and A8 Show the Default Setting)

With port numbers included in the command, **show port-security** displays Learn Mode, Address Limit, (alarm) Action, and Authorized Addresses for the specified ports on a switch. The following example lists the full port security configuration for a single port:

```
ProCurve(config)# show port-security A3
 Port Security
  Port : A3

  Learn Mode [Continuous]: Static    Address Limit[1]:

  Action  [None]: Send Alarm

  Authorized Addresses
  --------------------
  00906d-fdcc00
```

Figure 9-3. Example of the Port Security Configuration Display for a Single Port

The following command example shows the option for entering a range of ports, including a series of non-contiguous ports. Note that no spaces are allowed in the port number portion of the command string:

```
ProCurve(config)# show port-security A1-A3,A6,A8
```

Configuring Port Security

Using the CLI, you can:

- Configure port security and edit security settings.
- Add or delete devices from the list of authorized addresses for one or more ports.
- Clear the Intrusion flag on specific ports

Syntax: port-security [e] < *port-list* >
[learn-mode < continuous | static | configured | port-access >]
[address-limit < *integer* >]
[mac-address < *mac-addr* >] [< *mac-addr* > . . . < *mac-addr* >]
[action < none | send-alarm | send-disable >]
[clear-intrusion-flag]

(For the **configured** option, above, refer to the **Note** on page 9-8.

no port-security < *port-list* > mac-address < *mac-addr* > [< *mac-addr* > . . .
< *mac-addr* >]

Specifying Authorized Devices and Intrusion Responses

Learn-Mode Static. This example configures port A1 to automatically accept the first device (MAC address) it detects as the only authorized device for that port. (The default device limit is 1.) It also configures the port to send an alarm to a network management station and disable itself if an intruder is detected on the port.

```
ProCurve(config)# port-security a1 learn-mode static
action send-disable
```

The next example does the same as the preceding example, except that it specifies a MAC address of 0c0090-123456 as the authorized device instead of allowing the port to automatically assign the first device it detects as an authorized device.

```
ProCurve(config)# port-security a1 learn-mode static
mac-address 0c0090-123456 action send-disable
```

This example configures port A5 to:

- Allow two MAC addresses, 00c100-7fec00 and 0060b0-889e00, as the authorized devices.
- Send an alarm to a management station if an intruder is detected on the port.

```
ProCurve(config)# port-security a5 learn-mode static
address-limit 2 mac-address 00c100-7fec00 0060b0-889e00
action send-alarm
```

If you manually configure authorized devices (MAC addresses) and/or an alarm action on a port, those settings remain unless you either manually change them or reset the switch to its factory-default configuration. You can "turn off" device authorization on a port by configuring the port to **continuous** Learn Mode, but subsequently reconfiguring the port to **static** Learn Mode restores the configured device authorization.

Learn-Mode Configured. This option allows only MAC addresses specifically configured with **learn-mode configured mac-address <** *mac-address* **>**, and does not automatically learn non-specified MAC addresses learned from the network. This example configures port A1 to:

- Allow only a MAC address of 0c0090-123456 as the authorized device
- Reserve the option for adding two more specified MAC addresses at a later time without having to change the address-limit setting.
- Send an alarm to a management station if an intruder is detected on the port.

```
ProCurve(config)# port-security A1 learn-mode configured
mac-address 0c0090-123456 address-limit 3 action send-
disable
```

Adding a MAC Address to an Existing Port List

To simply add a device (MAC address) to a port's existing Authorized Addresses list, enter the port number with the **mac-address** parameter and the device's MAC address. *This assumes that Learn Mode is either* **static** *or* **configured** *and the Authorized Addresses list is not already full* (as determined by the current **address-limit** value). For example, suppose port A1 allows two authorized devices, but has only one device in its Authorized Address list:

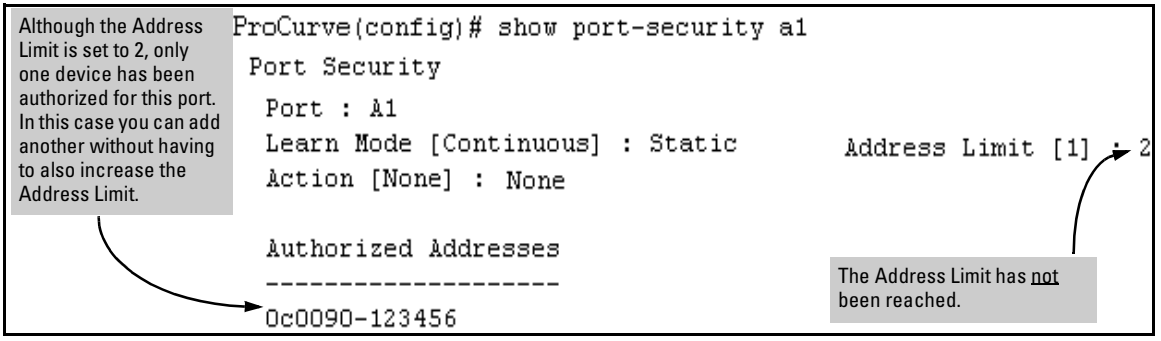

Figure 9-4. Example of Adding an Authorized Device to a Port

With the above configuration for port A1, the following command adds the 0c0090-456456 MAC address as the second authorized address.

```
ProCurve(config)# port-security a1 mac-address 0c0090-
456456
```

After executing the above command, the security configuration for port A1 appears as:

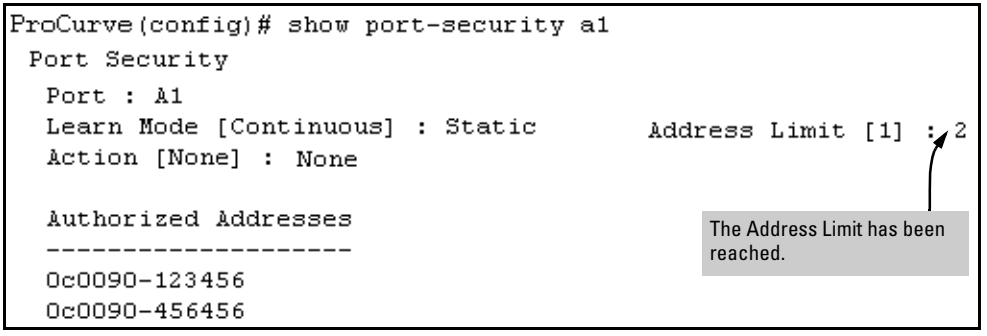

Figure 9-5. Example of Adding a Second Authorized Device to a Port

Note

The message **Inconsistent value** appears if the new MAC address exceeds the current Address Limit or specifies a device that is already on the list. If you change a port from static to continuous learn mode, the port retains in memory any authorized addresses it had while in static mode. If you subsequently attempt to convert the port back to static mode with the same authorized address(es), the `Inconsistent value` message appears because the port already has the address(es) in its "Authorized" list.

If you are adding a device (MAC address) to a port on which the Authorized Addresses list is already full (as controlled by the port's current Address Limit setting), then you must increase the Address Limit in order to add the device, even if you want to replace one device with another. Using the CLI, you can simultaneously increase the limit and add the MAC address with a single command. For example, suppose port A1 allows one authorized device and already has a device listed:

```
ProCurve(config)# show port-security a1
 Port Security
  Port : A1
  Learn Mode [Continuous] : Static   Address Limit [1]:1
  Action [None] : None

  Authorized Addresses
  --------------------
  0c0090-123456
```

Figure 9-6. Example of Port Security on Port A1 with an Address Limit of "1"

To add a second authorized device to port A1, execute a **port-security** command for port A1 that raises the address limit to 2 and specifies the additional device's MAC address. For example:

```
ProCurve(config)# port-security a1 mac-address 0c0090-
456456 address-limit 2
```

Removing a Device From the "Authorized" List for a Port Configured for Learn-Mode Static. This command option removes unwanted devices (MAC addresses) from the Authorized Addresses list. (An Authorized Address list is available for each port for which Learn Mode is currently set to "Static". See the "MAC Address" entry in the table on 9-8.)

Caution

The **address-limit** setting controls how many MAC addresses are allowed in the Authorized Addresses list for a given port. If you remove a MAC address without also reducing the address limit by 1, the port may later detect and accept the same or another MAC address that you do not want in the Authorized Address list. Thus, if you use the CLI to remove a MAC address that is no longer authorized, you should first reduce the Address Limit (**address-limit**) integer by 1, as shown in the next example. This prevents the possibility of the same device or another device on the network from automatically being accepted as "authorized" for that port. (You can prevent the port from learning unauthorized MAC addresses by using the **learn-mode configured** option instead of the **learn-mode static** option. Refer to the **Note** on page 9-8.)

To remove a device (MAC address) from the "Authorized" list and when the current number of devices equals the Address Limit value, you should first reduce the Address Limit value by 1, then remove the unwanted device.

Note

When you have configured the switch for **learn-mode static** operation, you can reduce the address limit below the number of currently authorized addresses on a port. This enables you to subsequently remove a device from the "Authorized" list without opening the possibility for an unwanted device to automatically become authorized. (If you use learn-mode configured instead, the switch cannot automatically add detected devices not included in the **mac-address** configuration. Refer to the **Note** on page 9-8.)

For example, suppose port A1 is configured as shown below and you want to remove 0c0090-123456 from the Authorized Address list:

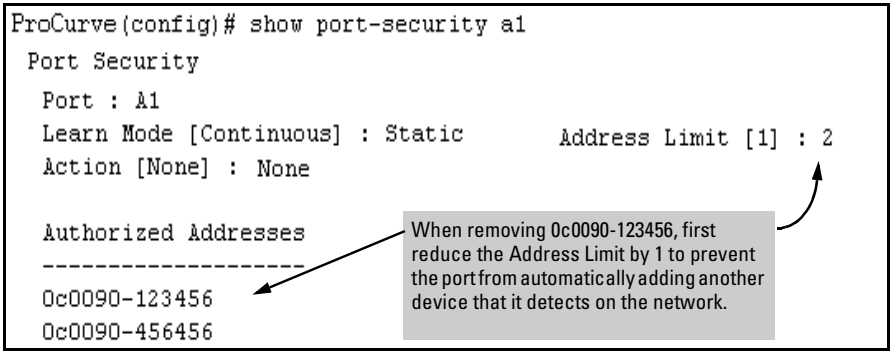

Figure 9-7. Example of Two Authorized Addresses on Port A1

The following command serves this purpose by removing 0c0090-123456 and reducing the Address Limit to 1:

```
ProCurve(config)# port-security a1 address-limit 1
ProCurve(config)# no port-security a1 mac-address 0c0090-
123456
```

The above command sequence results in the following configuration for port A1:

```
ProCurve(config)# show port-sec a1
 Port Security
  Port : A1
  Learn Mode : Static              Address Limit : 1
  Action : None
  Authorized Addresses
  -------------------
  0c0090-456456
```

Figure 9-8. Example of Port A1 After Removing One MAC Address

MAC Lockdown

MAC Lockdown is available on the Series 2600, 2600-PWR, and 2800 switches only.

MAC Lockdown, also known as "static addressing," is the permanent assignment of a given MAC address (and VLAN, or Virtual Local Area Network) to a specific port on the switch. MAC Lockdown is used to prevent station movement and MAC address hijacking. It also controls address learning on the switch. When configured, the MAC Address can only be used on the assigned port and the client device will only be allowed on the assigned VLAN.

Note

Port security and MAC Lockdown are mutually exclusive on a given port. You can either use port security *or* MAC Lockdown, but never both at the same time on the same port.

Syntax: [no] static-mac < *mac-addr* > vlan < *vid* > interface < *port-number* >

You will need to enter a separate command for each MAC/VLAN pair you wish to lock down. If you do not specify a VLAN ID (VID) the switch inserts a VID of "1".

How It Works. When a device's MAC address is locked down to a port (typically in a pair with a VLAN) all information sent to that MAC address must go through the locked-down port. If the device is moved to another port it cannot receive data. Traffic to the designated MAC address goes only to the allowed port, whether the device is connected to it or not.

MAC Lockdown is useful for preventing an intruder from "hijacking" a MAC address from a known user in order to steal data. Without MAC Lockdown, this will cause the switch to learn the address on the malicious user's port, allowing the intruder to steal the traffic meant for the legitimate user.

MAC Lockdown ensures that traffic intended for a specific MAC address can only go through the one port which is supposed to be connected to that MAC address. It does not prevent intruders from transmitting packets with the locked MAC address, but it does prevent responses to those packets from going anywhere other than the locked-down port. Thus TCP connections cannot be established. Traffic sent to the locked address cannot be hijacked and directed out the port of the intruder.

If the device (computer, PDA, wireless device) is moved to a different port on the switch (by reconnecting the Ethernet cable or by moving the device to an area using a wireless access point connected to a different port on that same switch), the port will detect that the MAC Address is not on the appropriate port and will continue to send traffic out the port to which the address was locked.

Once a MAC address is configured for one port, you cannot perform port security using the same MAC address on any other port on that same switch.

You cannot lock down a single MAC Address/VLAN pair to more than one port; however you can lock down multiple different MAC Addresses to a single port on the same switch.

Stations can move from the port to which their MAC address is locked to other parts of the network. They can send, but will not receive data if that data must go through the locked down switch. Please note that if the device moves to a distant part of the network where data sent to its MAC address never goes through the locked down switch, it may be possible for the device to have full two-way communication. For full and complete lockdown network-wide all switches must be configured appropriately.

Other Useful Information. Once you lock down a MAC address/VLAN pair on one port that pair cannot be locked down on a different port.

You cannot perform MAC Lockdown and 802.1x authentication on the same port or on the same MAC address. MAC Lockdown and 802.1x authentication are mutually exclusive.

Lockdown is permitted on static trunks (manually configured link aggregations).

Differences Between MAC Lockdown and Port Security

Because port-security relies upon MAC addresses, it is often confused with the MAC Lockdown feature. However, MAC Lockdown is a completely different feature and is implemented on a different architecture level.

Port security maintains a list of allowed MAC addresses on a per-port basis. An address can exist on multiple ports of a switch. Port security deals with MAC addresses only while MAC Lockdown specifies both a MAC address and a VLAN for lockdown.

MAC Lockdown, on the other hand, is not a “list.” It is a global parameter on the switch that takes precedence over any other security mechanism. The MAC Address will only be allowed to communicate using one specific port on the switch.

MAC Lockdown is a good replacement for port security to create tighter control over MAC addresses and which ports they are allowed to use (only one port per MAC Address on the same switch in the case of MAC Lockdown). (You can still use the port for other MAC addresses, but you cannot use the locked down MAC address on other ports.)

Using only port security the MAC Address could still be used on another port on the same switch. MAC Lockdown, on the other hand, is a clear one-to-one relationship between the MAC Address and the port. Once a MAC address has been locked down to a port it cannot be used on another port on the same switch.

The switch does not allow MAC Lockdown and port security on the same port.

MAC Lockdown Operating Notes

Limits. There is a limit of 500 MAC Lockdowns that you can safely code per switch. To truly lock down a MAC address it would be necessary to use the MAC Lockdown command for every MAC Address and VLAN ID on every switch. In reality few network administrators will go to this length, but it is important to note that just because you have locked down the MAC address and VID for a single switch, the device (or a hacker "spoofing" the MAC address for the device) may still be able to use another switch which hasn't been locked down.

Event Log Messages. If someone using a locked down MAC address is attempting to communicate using the wrong port the "move attempt" generates messages in the log file like this:

```
Move attempt (lockdown) logging:

W 10/30/03 21:33:43 maclock: module A: Move 0001e6-1f96c0
to A15 denied

W 10/30/03 21:33:48 maclock: module A: Move 0001e6-1f96c0
to A15 denied

W 10/30/03 21:33:48 maclock: module A: Ceasing move-denied
logs for 5m
```

These messages in the log file can be useful for troubleshooting problems. If you are trying to connect a device which has been locked down to the wrong port, it will not work but it will generate error messages like this to help you determine the problem.

Limiting the Frequency of Log Messages. The first move attempt (or intrusion) is logged as you see in the example above. Subsequent move attempts send a message to the log file also, but message throttling is imposed on the logging on a per-module basis. What this means is that the logging system checks again after the first 5 minutes to see if another attempt has been made to move to the wrong port. If this is the case the log file registers the most recent attempt and then checks again after one hour. If there are no further attempts in that period then it will continue to check every 5 minutes. If another attempt was made during the one hour period then the log resets itself to check once a day. The purpose of rate-limiting the log messaging is to prevent the log file from becoming too full. You can also configure the switch to send the same messages to a Syslog server. Refer to "Debug and Syslog Messaging Operation" in appendix C of the *Management and Configuration Guide* for your switch.

Deploying MAC Lockdown

When you deploy MAC Lockdown you need to consider how you use it within your network topology to ensure security. In some cases where you are using techniques such as Spanning Tree Protocol (STP) to speed up network performance by providing multiple paths for devices, using MAC Lockdown either will not work or else it defeats the purpose of having multiple data paths.

The purpose of using MAC Lockdown is to prevent a malicious user from "hijacking" an approved MAC address so they can steal data traffic being sent to that address.

As we have seen, MAC Lockdown can help prevent this type of hijacking by making sure that all traffic to a specific MAC address goes only to the proper port on a switch which is supposed to be connected to the real device bearing that MAC address.

However, you can run into trouble if you incorrectly try to deploy MAC Lockdown in a network that uses multiple path technology, like Spanning Tree.

Let's examine a good use of MAC Lockdown within a network to ensure security first.

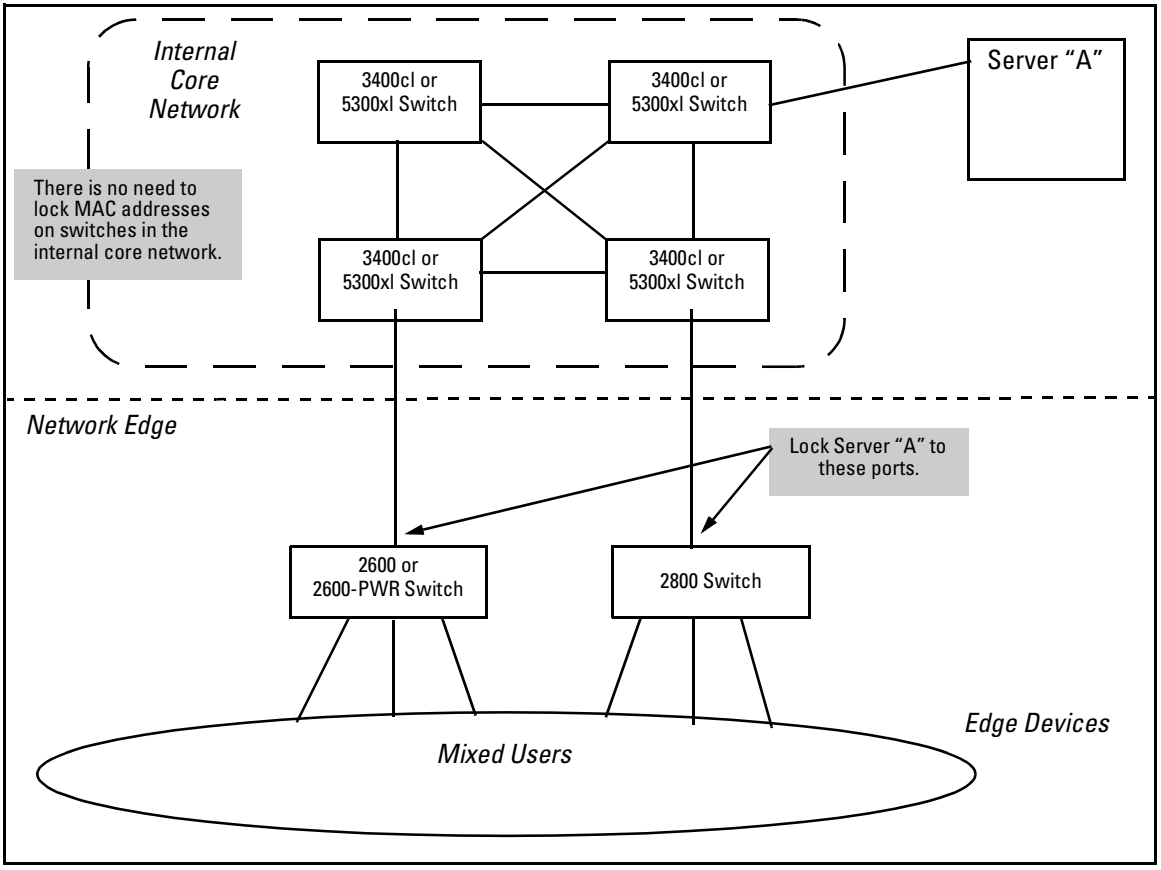

Figure 9-9. MAC Lockdown Deployed At the Network Edge Provides Security

Basic MAC Lockdown Deployment. In the Model Network Topology shown above, the switches that are connected to the edge of the network each have one and only one connection to the core network. This means each switch has only one path by which data can travel to Server A. You can use MAC Lockdown to specify that all traffic intended for Server A's MAC Address must go through the one port on the edge switches. That way, users on the edge can still use other network resources, but they cannot "spoof" Server A and hijack data traffic which is intended for that server alone.

The key points for this Model Topology are:

- The Core Network is separated from the edge by the use of switches which have been "locked down" for security.
- All switches connected to the edge (outside users) each have only one port they can use to connect to the Core Network and then to Server A.
- Each switch has been configured with MAC Lockdown so that the MAC Address for Server A has been locked down to one port per switch that can connect to the Core and Server A.

Using this setup Server A can be moved around within the core network, and yet MAC Lockdown will still prevent a user at the edge from hijacking its address and stealing data.

Please note that in this scenario a user with bad intentions at the edge can still "spoof" the address for Server A and send out data packets that look as though they came from Server A. The good news is that because MAC Lockdown has been used on the switches on the edge, any traffic that is sent *back* to Server A will be sent to the proper MAC Address because MAC Lockdown has been used. The switches at the edge will not send Server A's data packets anywhere but the port connected to Server A. (Data would not be allowed to go beyond the edge switches.)

Caution

Using MAC Lockdown still does not protect against a hijacker *within the core*! In order to protect against someone spoofing the MAC Address for Server A inside the Core Network, you would have to lock down each and every switch inside the Core Network as well, not just on the edge.

Problems Using MAC Lockdown in Networks With Multiple Paths. Now let's take a look at a network topology in which the use of MAC Lockdown presents a problem. In the next figure, Switch 1 (on the bottom-left) is located at the edge of the network where there is a mixed audience that might contain hackers or other malicious users. Switch 1 has two paths it could use to connect to Server A. If you try to use MAC Lockdown here to make sure that all data to Server A is "locked down" to one path, connectivity problems would be the result since both paths need to be usable in case one of them fails.

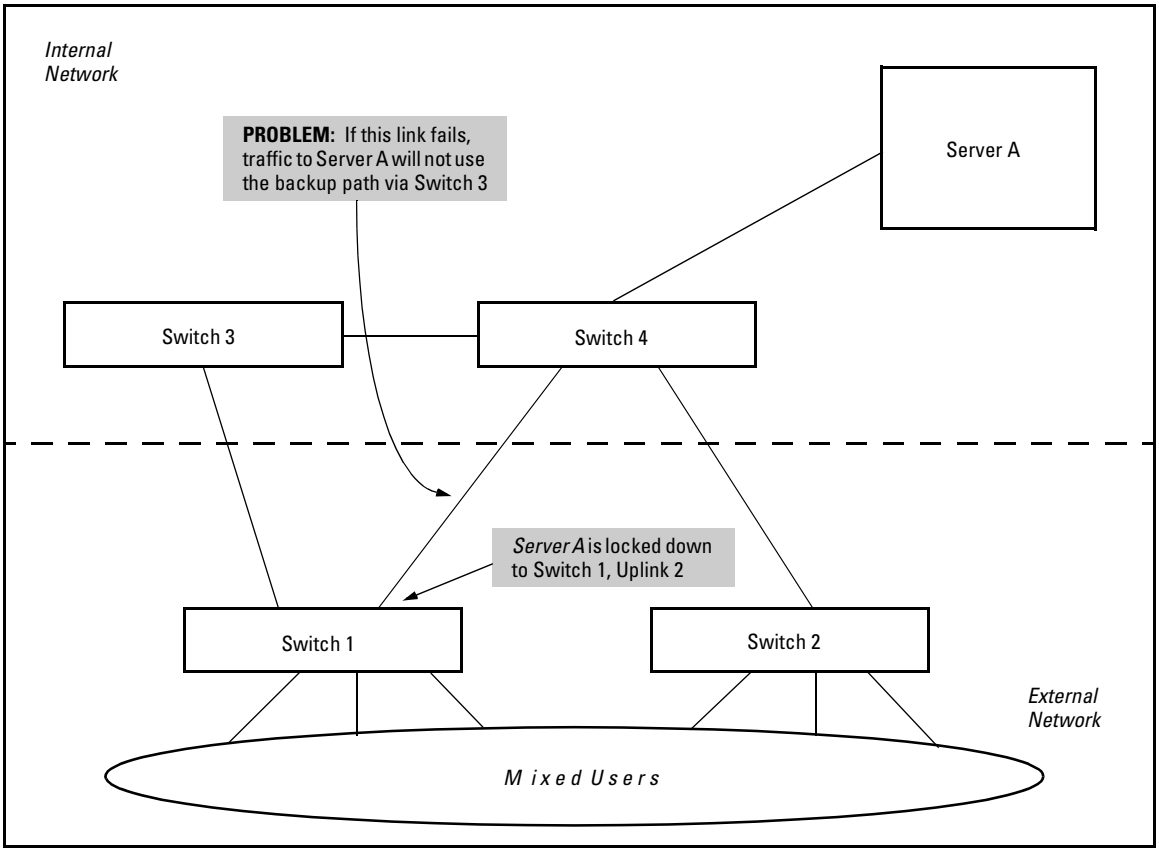

Figure 9-10. Connectivity Problems Using MAC Lockdown with Multiple Paths

The resultant connectivity issues would prevent you from locking down Server A to Switch 1. And when you remove the MAC Lockdown from Switch 1 (to prevent broadcast storms or other connectivity issues), you then open the network to security problems. The use of MAC Lockdown as shown in the above figure would defeat the purpose of using STP or having an alternate path.

Technologies such as STP are primarily intended for an internal campus network environment in which all users are trusted. STP does not work well with MAC Lockdown.

If you deploy MAC Lockdown as shown in the Model Topology in figure 9-9 (page 9-22), you should have no problems with either security or connectivity.

Displaying status. Locked down ports are listed in the output of the **show running-config** command in the CLI. The **show static-mac** command also lists the locked down MAC addresses, as shown below.

```
ProCurve# show static-mac
VLAN  MAC Address Port
   1 001083-34f8fa 9
Number of locked down MAC addresses = 1
ProCurve#
```

Figure 9-11. Listing Locked Down Ports

MAC Lockout

MAC Lockout is available on the Series 2600, 2600-PWR, and 2800 switches only.

MAC Lockout involves configuring a MAC address on all ports and VLANs for a switch so that any traffic to or from the "locked-out" MAC address will be dropped. This means that all data packets addressed to or from the given address are stopped by the switch. MAC Lockout is implemented on a per switch assignment.

You can think of MAC Lockout as a simple blacklist. The MAC address is locked out on the switch and on all VLANs. No data goes out or in from the blacklisted MAC address to a switch using MAC Lockout.

To fully lock out a MAC address from the network it would be necessary to use the MAC Lockout command on all switches.

To use MAC Lockout you must first know the MAC Address you wish to block.

Syntax: [no] lockout-mac < *mac-address* >

How It Works. Let's say a customer knows there are unauthorized wireless clients who should not have access to the network. The network administrator "locks out" the MAC addresses for the wireless clients by using the MAC

Lockout command (**lockout-mac <*mac-address*>**). When the wireless clients then attempt to use the network, the switch recognizes the intruding MAC addresses and prevents them from sending or receiving data on that network.

If a particular MAC address can be identified as unwanted on the switch then that MAC Address can be disallowed on all ports on that switch with a single command. You don't have to configure every single port—just perform the command on the switch and it is effective for all ports.

MAC Lockout overrides MAC Lockdown, port security, and 802.1x authentication.

You cannot use MAC Lockout to lock:

- Broadcast or Multicast Addresses (Switches do not learn these)
- Switch Agents (The switch's own MAC Address)

If someone using a locked out MAC address tries to send data through the switch a message is generated in the log file:

```
Lockout logging format:

W 10/30/03 21:35:15 maclock: module A: 0001e6-1f96c0
detected on port A15
W 10/30/03 21:35:18 maclock: module A: 0001e6-1f96c0
detected on port A15
W 10/30/03 21:35:18 maclock: module A: Ceasing lock-out
logs for 5m
```

As with MAC Lockdown a rate limiting algorithm is used on the log file so that it does not become overclogged with error messages. (Refer to "Limiting the Frequency of Log Messages" on page 9-20.)

Displaying status. Locked out ports are listed in the output of the **show running-config** command in the CLI. The **show lockout-mac** command also lists the locked out MAC addresses, as shown below.

```
ProCurve# show lockout-mac
Locked Out Addresses
 007347-a8fd30
Number of locked out MAC addresses = 1
ProCurve#
```

Figure 9-12. Listing Locked Out Ports

Port Security and MAC Lockout

MAC Lockout is independent of port-security and in fact will override it. MAC Lockout is preferable to port-security to stop access from known devices because it can be configured for all ports on the switch with one command.

It is possible to use MAC Lockout in conjunction with port-security. You can use MAC Lockout to lock out a single address—deny access to a specific device—but still allow the switch some flexibility in learning other MAC Addresses. Be careful if you use both together, however:

- If a MAC Address is locked out and appears in a static learn table in port-security, the apparently "authorized" address will still be locked out anyway.
- MAC entry configurations set by port security will be kept even if MAC Lockout is configured and the original port security settings will be honored once the Lockout is removed.
- A port security static address is permitted to be a lockout address. In that case (MAC Lockout), the address will be locked out (SA/DA drop) even though it's an "authorized" address from the perspective of port security.
- When MAC Lockout entries are deleted, port security will then re-learn the address as needed later on.

IP Lockdown

IP lockdown is available on the Series 2600 and 2800 switches only.

The "IP lockdown" utility enables you to restrict incoming traffic on a port to a specific IP address/subnet, and deny all other traffic on that port.

Operating Rules for IP Lockdown

- Users cannot specify that certain subnets be denied while others are permitted.
- Users cannot filter on protocol or destination IP address.
- The lockdown feature applies to inbound traffic on a port only.
- There is no logging functionality for this feature, i.e. no way to determine if IP address violations occur.
- The same subnet mask must be used for all ports within an 8 port block (1-8, 7-16, etc), for example:
  - If you configure Port 1 with: `ip-lockdown 192.168.0.1/24`
  - Then configure Port 2 with: `ip-lockdown 50.0.0.0/24` This is an acceptable subnet for port 2
  - Then configure Port 3 with: `ip-lockdown 120.15.32.7/32` This command would return an error and not be configured due to the differing subnet mask.

Using the IP Lockdown Command

The IP lockdown command operates as follows:

Syntax: ip-lockdown <*subnet mask/ips* >

Defines the subnet and related IP addresses allowed for incoming traffic on the port.

The following example prevents traffic from all IP addresses other than those specified in subnet 192.168.0.1/24 from entering the switch on interface 1.

```
ProCurve Switch 2626 (config) # interface 1

ProCurve Switch 2626 (eth-1) # ip-lockdown 192.168.0.1/24

ProCurve Switch 2626 (eth-1) # exit
```

Web: Displaying and Configuring Port Security Features

1. Click on the **Security** tab.
2. Click on **[Port Security]**.
3. Select the settings you want and, if you are using the Static Learn Mode, add or edit the Authorized Addresses field.
4. Implement your new data by clicking on **[Apply Changes]**.

To access the web-based Help provided for the switch, click on **[?]** in the web browser screen.

Reading Intrusion Alerts and Resetting Alert Flags

Notice of Security Violations

When the switch detects an intrusion on a port, it sets an "alert flag" for that port and makes the intrusion information available as described below. *While the switch can detect additional intrusions for the same port, it does not list the next chronological intrusion for that port in the Intrusion Log until the alert flag for that port has been reset.*

When a security violation occurs on a port configured for Port Security, the switch responds in the following ways to notify you:

- The switch sets an alert flag for that port. This flag remains set until:
  - You use either the CLI, menu interface, or web browser interface to reset the flag.
  - The switch is reset to its factory default configuration.
- The switch enables notification of the intrusion through the following means:
  - In the CLI:
    - The **show port-security intrusion-log** command displays the Intrusion Log
    - The **log** command displays the Event Log

- In the menu interface:
  - The Port Status screen includes a per-port intrusion alert
  - The Event Log includes per-port entries for security violations
- In the web browser interface:
  - The Alert Log's Status | Overview window includes entries for per-port security violations
  - The Intrusion Log in the Security | Intrusion Log window lists per-port security violation entries
- In an active network management environment via an SNMP trap sent to a network management station

How the Intrusion Log Operates

When the switch detects an intrusion attempt on a port, it enters a record of this event in the Intrusion Log. No further intrusion attempts on that port will appear in the Log until you acknowledge the earlier intrusion event by resetting the alert flag.

The Intrusion Log lists the 20 most recently detected security violation attempts, regardless of whether the alert flags for these attempts have been reset. This gives you a history of past intrusion attempts. Thus, for example, if there is an intrusion alert for port A1 and the Intrusion Log shows two or more entries for port 1, only the most recent entry has not been acknowledged (by resetting the alert flag). The other entries give you a history of past intrusions detected on port A1.

```
ProCurve# show port-security intrusion-log
 Status and Counters - Intrusion Log
  Port  MAC Address          Date / Time
  ---- ------------- -------------------------
  A1   080009-e93d4f           07/03/02 21:09:34
  A1   080009-21ae84           07/03/02 17:26:27
  A1   080009-e93d4f prior to 07/03/02 17:18:43
```

Figure 9-13. Example of Multiple Intrusion Log Entries for the Same Port

The log shows the most recent intrusion at the top of the listing. You cannot delete Intrusion Log entries (unless you reset the switch to its factory-default configuration). Instead, if the log is filled when the switch detects a new intrusion, the oldest entry is dropped off the listing and the newest entry appears at the top of the listing.

Keeping the Intrusion Log Current by Resetting Alert Flags

When a violation occurs on a port, an alert flag is set for that port and the violation is entered in the Intrusion Log. The switch can detect and handle subsequent intrusions on that port, but will not log another intrusion on the port until you reset the alert flag for either all ports or for the individual port.

Note on Send-Disable Operation

On a given port, if the intrusion action is to send an SNMP trap and then disable the port (**send-disable**), and then an intruder is detected on the port, the switch sends an SNMP trap, sets the port's alert flag, and disables the port. If you re-enable the port without resetting the port's alert flag, then the port operates as follows:

- The port comes up and will block traffic from unauthorized devices it detects.
- If the port detects another intruder, it will send another SNMP trap, but will not become disabled again unless you first reset the port's intrusion flag.

This operation enables the port to continue passing traffic for authorized devices while you locate and eliminate the intruder. Otherwise, the presence of an intruder could cause the switch to repeatedly disable the port.

Menu: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags

The menu interface indicates per-port intrusions in the Port Status screen, and provides details and the reset function in the Intrusion Log screen.

1. From the Main Menu select:

   1. Status and Counters
   4. Port Status

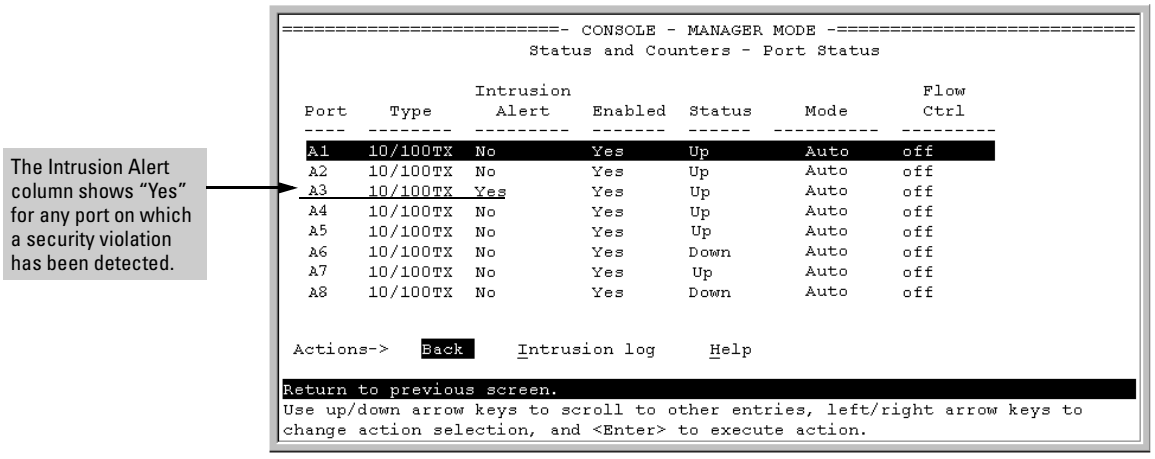

Figure 9-14. Example of Port Status Screen with Intrusion Alert on Port A3

2. Type **[I]** (**Intrusion log**) to display the Intrusion Log.

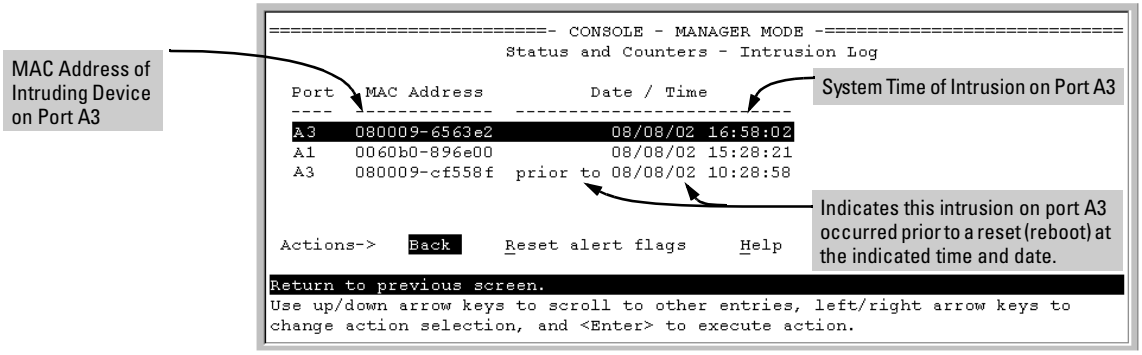

Figure 9-15. Example of the Intrusion Log Display

The above example shows two intrusions for port A3 and one intrusion for port A1. In this case, only the most recent intrusion at port A3 has not been acknowledged (reset). This is indicated by the following:

- Because the Port Status screen (figure 9-14 on page 9-32) does not indicate an intrusion for port A1, the alert flag for the intrusion on port A1 has already been reset.
- Since the switch can show only one uncleared intrusion per port, the older intrusion for port A3 in this example has also been previously reset.

(The intrusion log holds up to 20 intrusion records and deletes an intrusion record only when the log becomes full and a new intrusion is subsequently detected.)

Note also that the "**prior to**" text in the record for the earliest intrusion means that a switch reset occurred at the indicated time and that the intrusion occurred prior to the reset.

3. To acknowledge the most recent intrusion entry on port A3 and enable the switch to enter a subsequently detected intrusion on this port, type **[R]** (for **Reset alert flags**). (Note that if there are unacknowledged intrusions on two or more ports, this step resets the alert flags for all such ports.)

If you then re-display the port status screen, you will see that the Intrusion Alert entry for port A3 has changed to "**No**". That is, your evidence that the Intrusion Alert flag has been acknowledged (reset) is that the Intrusion Alert column in the port status display no longer shows "**Yes**" for the port on which the intrusion occurred (port A3 in this example). (Because the Intrusion Log provides a history of the last 20 intrusions detected by the switch, resetting the alert flags does not change its content. Thus, displaying the Intrusion Log again will result in the same display as in figure 9-15, above.)

CLI: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags

The following commands display port status, including whether there are intrusion alerts for any port(s), list the last 20 intrusions, and either reset the alert flag on all ports or for a specific port for which an intrusion was detected. (The record of the intrusion remains in the log. For more information, refer to "Operating Notes for Port Security" on page 9-37.)

Syntax: show interfaces brief

List intrusion alert status (and other port status information)'.

show port-security intrusion-log

List intrusion log content.

clear intrusion-flags

Clear intrusion flags on all ports.

port-security [e] < port-number > clear-intrusion-flag

Clear the intrusion flag on one or more specific ports.

In the following example, executing **show interfaces brief** lists the switch's port status, which indicates an intrusion alert on port A1.

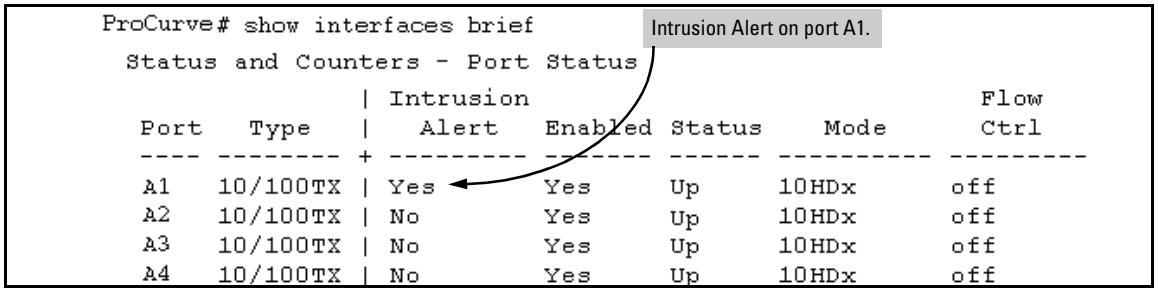

Figure 9-16. Example of an Unacknowledged Intrusion Alert in a Port Status Display

If you wanted to see the details of the intrusion, you would then enter the **show port-security intrusion-log** command. For example:

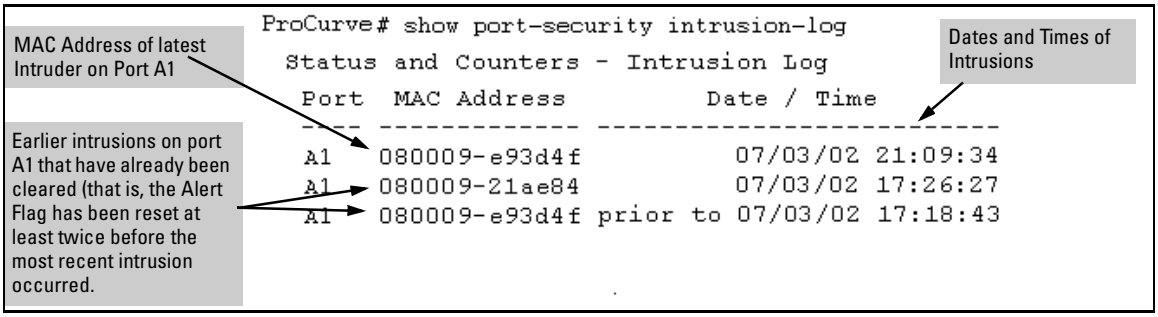

Figure 9-17. Example of the Intrusion Log with Multiple Entries for the Same Port

The above example shows three intrusions for port A1. Since the switch can show only one uncleared intrusion per port, the older two intrusions in this example have already been cleared by earlier use of the **clear intrusion-log** or the **port-security < *port-list* > clear-intrusion-flag** command. (The intrusion log holds up to 20 intrusion records, and deletes intrusion records only when the log becomes full and new intrusions are subsequently added.) The "**prior to**" text in the record for the third intrusion means that a switch reset occurred at the indicated time and that the intrusion occurred prior to the reset.

To clear the intrusion from port A1 and enable the switch to enter any subsequent intrusion for port A1 in the Intrusion Log, execute the port-security **clear-intrusion-flag** command. If you then re-display the port status screen, you will see that the Intrusion Alert entry for port A1 has changed to "**No**". (Executing **show port-security intrusion-log** again will result in the same display as above, and does not include the Intrusion Alert status.)

```
ProCurve(config)# port-security a1 clear-intrusion-flag
ProCurve(config)# show interfaces brief
```

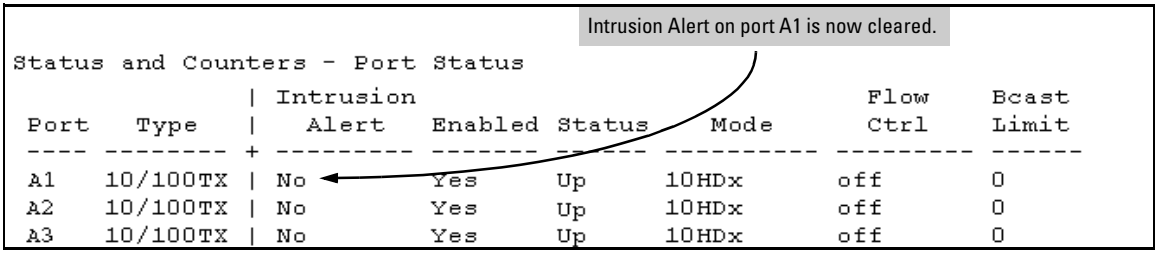

Figure 9-18. Example of Port Status Screen After Alert Flags Reset

For more on clearing intrusions, see "Note on Send-Disable Operation" on page 9-31

Using the Event Log To Find Intrusion Alerts

The Event Log lists port security intrusions as:

```
W MM/DD/YY HH:MM:SS FFI: port A3 - Security Violation
```

where "**W**" is the severity level of the log entry and **FFI** is the system module that generated the entry. For further information, display the Intrusion Log, as shown below.

From the CLI. Type the **log** command from the Manager or Configuration level.

Syntax: log [*search-text*]

For *search-text* , you can use **ffi**, **security**, or **violation**. For example:

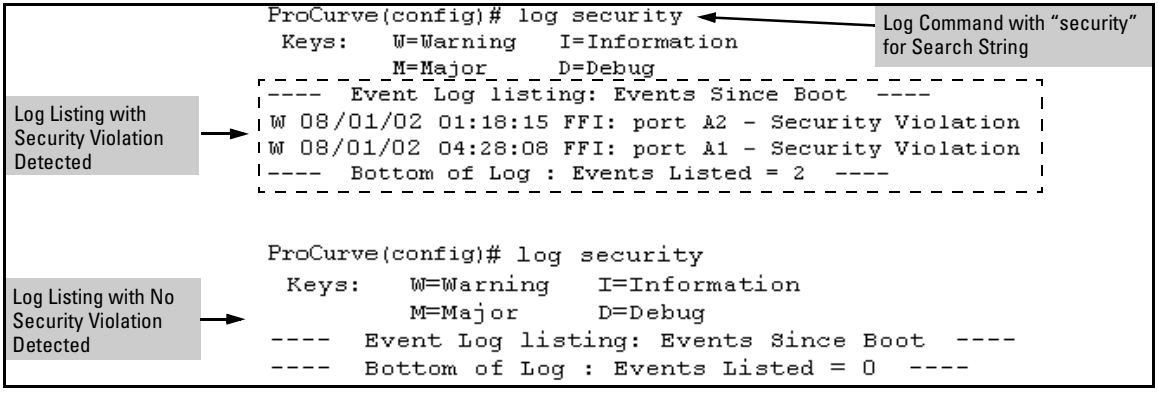

Figure 9-19. Example of Log Listing With and Without Detected Security Violations

From the Menu Interface: In the Main Menu, click on **4. Event Log** and use **Next page** and **Prev page** to review the Event Log contents.

For More Event Log Information. See "Using the Event Log To Identify Problem Sources" in the "Troubleshooting" chapter of the *Management and Configuration Guide* for your switch.

Web: Checking for Intrusions, Listing Intrusion Alerts, and Resetting Alert Flags

1. Check the Alert Log by clicking on the **Status** tab and the **[Overview]** button. If there is a "Security Violation" entry, do the following:

a. Click on the **Security** tab.

b. Click on **[Intrusion Log]**. "Ports with Intrusion Flag" indicates any ports for which the alert flag has not been cleared.

c. To clear the current alert flags, click on **[Reset Alert Flags]**.

To access the web-based Help provided for the switch, click on **[?]** in the web browser screen.

Operating Notes for Port Security

Identifying the IP Address of an Intruder. The Intrusion Log lists detected intruders by MAC address. Proxy Web Servers

If you are using the switch's web browser interface through a switch port configured for Static port security, and your browser access is through a proxy web server, then it is necessary to do the following:

- Enter your PC or workstation MAC address in the port's Authorized Addresses list.
- Enter your PC or workstation's IP address in the switch's IP Authorized Managers list. See chapter 11, "Using Authorized IP Managers".)

Without both of the above configured, the switch detects only the proxy server's MAC address, and not your PC or workstation MAC address, and interprets your connection as unauthorized.

"Prior To" Entries in the Intrusion Log. If you reset the switch (using the Reset button, Device Reset, or Reboot Switch), the Intrusion Log will list the time of all currently logged intrusions as "prior to" the time of the reset.

Alert Flag Status for Entries Forced Off of the Intrusion Log. If the Intrusion Log is full of entries for which the alert flags have not been reset, a new intrusion will cause the oldest entry to drop off the list, but will not change the alert flag status for the port referenced in the dropped entry. This means that, even if an entry is forced off of the Intrusion Log, no new intrusions can be logged on the port referenced in that entry until you reset the alert flags.

LACP Not Available on Ports Configured for Port Security. To maintain security, LACP is not allowed on ports configured for port security. If you configure port security on a port on which LACP (active or passive) is configured, the switch removes the LACP configuration, displays a notice that LACP is disabled on the port(s), and enables port security on that port. For example:

```
ProCurve(config)# port-security e a17 learn-mode static
address-limit 2
LACP has been disabled on secured port(s).
ProCurve(config)#
```

The switch will not allow you to configure LACP on a port on which port security is enabled. For example:

```
ProCurve(config)# int e a17 lacp passive
Error configuring port A17: LACP and port security cannot
be run together.
ProCurve(config)#
```

To restore LACP to the port, you must remove port security and re-enable LACP active or passive.