



Understanding and Configuring 802.1x Port-Based Authentication

This chapter describes how to configure IEEE 802.1x port-based authentication to prevent unauthorized client devices from gaining access to the network.

This chapter contains the major sections:

- [Understanding 802.1x Port-Based Authentication, page 28-1](#)
- [How to Configure 802.1x, page 28-10](#)
- [Displaying 802.1x Statistics and Status, page 28-21](#)



Note

For complete syntax and usage information for the switch commands used in this chapter, refer to the *Catalyst 4500 Series Switch Cisco IOS Command Reference* and related publications at <http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/index.htm>

Understanding 802.1x Port-Based Authentication

To configure 802.1x port-based authentication, you need to understand the concepts in these sections:

- [Device Roles, page 28-2](#)
- [Authentication Initiation and Message Exchange, page 28-3](#)
- [Ports in Authorized and Unauthorized States, page 28-4](#)
- [Using 802.1x with the VLAN Assignment, page 28-5](#)
- [Using 802.1x Authentication for Guest VLANs, page 28-6](#)
- [Using 802.1x with Port Security, page 28-6](#)
- [802.1x RADIUS Accounting, page 28-7](#)
- [Supported Topologies, page 28-9](#)



Note

802.1x support requires an authentication server that is configured for Remote Authentication Dial-In User Service (RADIUS). 802.1x authentication does not work unless the network access switch can route packets to the configured authentication RADIUS server. To verify that the switch can route packets, you must ping the server from the switch.


802.1x defines 802.1x port-based authentication as a client-server based access control and authentication protocol that restricts unauthorized clients from connecting to a LAN through publicly accessible ports. An authentication server validates each supplicant (client) connected to an authenticator (network access switch) port before making available any services offered by the switch or the LAN.

Until a client is authenticated, only Extensible Authentication Protocol over LAN (EAPOL) traffic is allowed through the port to which the client is connected. Once authentication succeeds, normal traffic can pass through the port.

Device Roles

With 802.1x port-based authentication, network devices have specific roles. Figure 28-1 shows the roles of each device.

Figure 28-1 802.1x Device Roles



- **Client**—The workstation that requests access to the LAN, and responds to requests from the switch. The workstation must be running 802.1x-compliant client software.



Note

For more information on 802.1x-compliant client application software such as Microsoft's Windows 2000 Professional or Windows XP, refer to the Microsoft Knowledge Base article at this URL: <http://support.microsoft.com>

- **Authenticator**—Controls physical access to the network based on the authentication status of the client. The switch acts as an intermediary between the client and the authentication server, requesting identity information from the client, verifying that information with the authentication server, and relaying a response to the client. The switch encapsulates and decapsulates the Extensible Authentication Protocol (EAP) frames and interacts with the RADIUS authentication server.

When the switch receives EAPOL frames and relays them to the authentication server, the Ethernet header is stripped and the remaining EAP frame is reencapsulated in the RADIUS format. The EAP frames are not modified or examined during encapsulation, and the authentication server must support EAP within the native frame format. When the switch receives frames from the authentication server, the server's frame header is removed, leaving the EAP frame, which is then encapsulated for Ethernet and sent to the client.

Cisco devices that are capable of functioning as an 802.1x network access point include Catalyst 4500 series switches, the Catalyst 3550 multilayer switch, the Catalyst 2950 switch, and a Cisco Aironet series wireless access point. These devices must be running software that supports the RADIUS client and 802.1x.

- Authentication server—Performs the actual authentication of the client. The authentication server validates the identity of the client and notifies the switch that the client is authorized to access the LAN and switch services. (The only supported authentication server is the RADIUS authentication server with EAP extensions; it is available in Cisco Secure Access Control Server version 3.2 and later.)

Authentication Initiation and Message Exchange

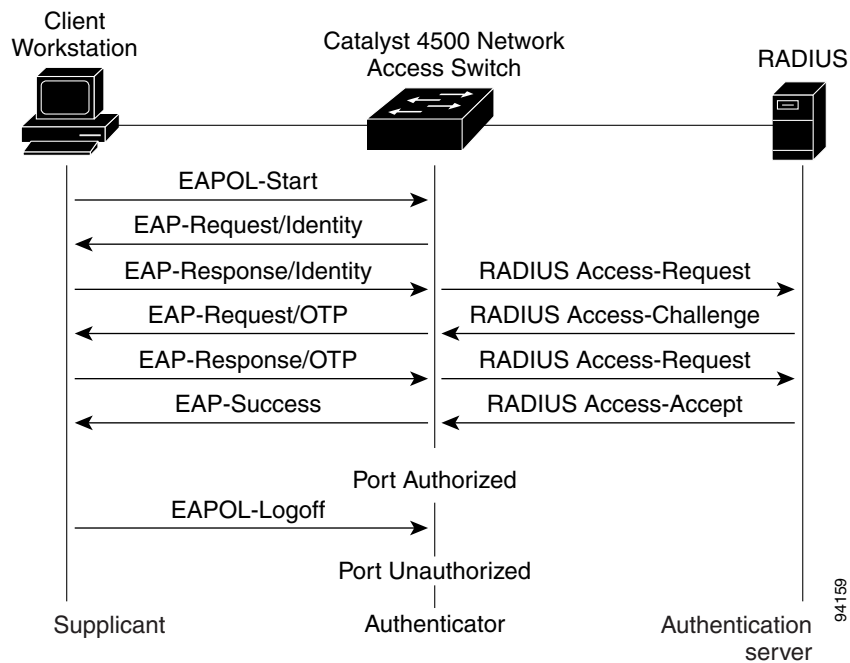
The switch or the client can initiate authentication. If you enable authentication on a port by using the **dot1x port-control auto** interface configuration command, the switch must initiate authentication when it determines that the port link state has changed. It then sends an EAP-request/identity frame to the client to request its identity (typically, the switch sends an initial identity/request frame followed by one or more requests for authentication information). Upon receipt of the frame, the client responds with an EAP-response/identity frame.

However, if during bootup, the client does not receive an EAP-request/identity frame from the switch, the client can initiate authentication by sending an EAPOL-start frame, which prompts the switch to request the client's identity.

If 802.1x is not enabled or supported on the network access device, any EAPOL frames from the client are dropped. If the client does not receive an EAP-request/identity frame after three attempts to start authentication, the client transmits frames as if the port is in the authorized state. A port in the authorized state means that the client has been successfully authenticated. When the client supplies its identity, the switch begins its role as the intermediary, passing EAP frames between the client and the authentication server until authentication succeeds or fails. If the authentication succeeds, the switch port becomes authorized.

The specific exchange of EAP frames depends on the authentication method being used. [Figure 28-2](#) shows a message exchange that is initiated by the client using the One-Time Password (OTP) authentication method with an authentication server.

Figure 28-2 Message Exchange



Ports in Authorized and Unauthorized States

The switch port state determines whether or not the client is granted access to the network. The port starts in the unauthorized state. While in this state, the port disallows all ingress and egress traffic except for 802.1x protocol packets. When a client is successfully authenticated, the port transitions to the authorized state, allowing all traffic for the client to flow normally.

If a client that does not support 802.1x is connected to an unauthorized 802.1x port, the switch requests the client's identity. In this situation, the client does not respond to the request, the port remains in the unauthorized state, and the client is not granted access to the network. If the guest VLAN is configured for a port that connects to a client that does not support 802.1x, the port is placed in the configured guest VLAN and in the authorized state. For more information, see the [“Using 802.1x Authentication for Guest VLANs”](#) section on page 28-6.

In contrast, when an 802.1x-enabled client connects to a port that is not running the 802.1x protocol, the client initiates the authentication process by sending the EAPOL-start frame. When no response is received, the client sends the request a fixed number of times. Because no response is received, the client begins sending frames as if the port is in the authorized state.

You can control the port authorization state with the **dot1x port-control** interface configuration command and these keywords:

- **force-authorized**—Disables 802.1x authentication and causes the port to transition to the authorized state without any authentication exchange required. The port transmits and receives normal traffic without 802.1x-based authentication of the client. This setting is the default.
- **force-unauthorized**—Causes the port to remain in the unauthorized state, ignoring all attempts by the client to authenticate. The switch cannot provide authentication services to the client through the interface.

- **auto**—Enables 802.1x authentication and causes the port to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the port. The authentication process begins when the link state of the port transitions from down to up or when an EAPOL-start frame is received. The switch requests the identity of the client and begins relaying authentication messages between the client and the authentication server. The switch can uniquely identify each client attempting to access the network by the client's MAC address.

If the client is successfully authenticated (receives an Accept frame from the authentication server), the port state changes to authorized, and all frames from the authenticated client are allowed through the port. If authentication fails, the port remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the switch can retransmit the request. If no response is received from the server after the specified number of attempts, authentication fails and network access is not granted.

If the link state of a port transitions from up to down, or if an EAPOL-logoff frame is received by the port, the port returns to the unauthorized state.

Using 802.1x with the VLAN Assignment

You can use the VLAN assignment to limit network access for certain users. With the VLAN assignment, 802.1x-authenticated ports are assigned to a VLAN based on the username of the client connected to that port. The RADIUS server database maintains the username-to-VLAN mappings. After successful 802.1x authentication of the port, the RADIUS server sends the VLAN assignment to the switch.



Note

To enable the guest VLAN feature in Release 12.1(19)EW and later releases, the port must be statically configured as an access port.

When configured on the switch and the RADIUS server, 802.1x with VLAN assignment has these characteristics:

- If no VLAN is supplied by the RADIUS server, the port is configured in its access VLAN when authentication succeeds.
- If the authentication server provides invalid VLAN information, the port remains unauthorized. This situation prevents ports from appearing unexpectedly in an inappropriate VLAN due to a configuration error.

Configuration errors might occur if you specify a VLAN for a routed port, a malformed VLAN ID, or a nonexistent or internal (routed port) VLAN ID. Similarly, an error might occur if you make an assignment to a voice VLAN ID.

- If the authentication server provides valid VLAN information, the port is authorized and placed in the specified VLAN when authentication succeeds.
- If the multiple-hosts mode is enabled, all hosts are in the same VLAN as the first authenticated user.
- If 802.1x is disabled on the port, the port is returned to the configured access VLAN.

To configure VLAN assignment you need to perform these tasks:

- Enable AAA authorization with the **network** keyword to allow interface configuration from the RADIUS server. For an illustration of how to apply the **aaa authorization network group radius** command, refer to the section “Enabling 802.1x Authentication” on page 12.
- Enable 802.1x. (The VLAN assignment feature is automatically enabled when you configure 802.1x on an access port.)

- Assign vendor-specific tunnel attributes in the RADIUS server. To ensure proper VLAN assignment, the RADIUS server must return these attributes to the switch:
 - Tunnel-Type = VLAN
 - Tunnel-Medium-Type = 802
 - Tunnel-Private-Group-ID = VLAN NAME

Using 802.1x Authentication for Guest VLANs

You can use guest VLANs to enable non-802.1x capable hosts to access networks that use 802.1x authentication. For example, you can use guest VLANs while you are upgrading your system to support 802.1x authentication.

Guest VLANs are supported on a per-port basis, and you can use any VLAN (except a private VLAN) as a guest VLAN. If a port is already forwarding on the guest VLAN and you enable 802.1x support on the network interface of the host, the port is immediately moved out of the guest VLAN and the authenticator waits for authentication to occur.

Enabling 802.1x authentication on a port starts the 802.1x protocol. If the host fails to respond to the packets from the authenticator within a certain amount of time, the authenticator puts the port in the guest VLAN.

Usage Guidelines for Using 802.1x Authentication with Guest VLANs on Windows-XP Hosts

The usage guidelines for using 802.1x authentication with guest VLANs on Windows-XP hosts are as follows:

- If the host fails to respond to the authenticator, the port attempts to connect three times (with a 30 second timeout between each attempt). After this time, the login/password window does not appear on the host, so you must unplug and reconnect the network interface cable.
- Hosts responding with an incorrect login/password fail authentication. Hosts failing authentication are not put in the guest VLAN. The first time that a host fails authentication, the quiet-period timer starts, and no activity occurs for the duration of the quiet-period timer. When the quiet-period timer expires, the host is presented with the login/password window. If the host fails authentication for the second time, the quiet-period timer starts again, and no activity will occur for the duration of the quiet-period timer. The host is presented with the login/password window a third time. If the host fails authentication the third time, the port is placed in the unauthorized state, and you must disconnect and reconnect the network interface cable.

Using 802.1x with Port Security

You can enable port security on an 802.1x port in either single- or multiple-host mode. (To do so, you must configure port security with the **switchport port-security** interface configuration command. Refer to the “Configuring Port Security” chapter in this guide.) When you enable port security and 802.1x on a port, 802.1x authenticates the port, and port security manages the number of MAC addresses allowed on that port, including that of the client. Hence an 802.1x port with port security enabled can be used to limit the number or group of clients that can access the network.

These examples describe the interaction between 802.1x and port security on the switch:

- When a client is authenticated, and the port security table is not full, the client’s MAC address is added to the port security list of secure hosts. The port then proceeds to come up normally.

When a client is authenticated and manually configured for port security, it is guaranteed an entry in the secure host table (unless port security static aging has been enabled).

A security violation occurs if an additional host is learned on the port. The action taken depends on which feature (802.1x or port security) detects the security violation:

- If 802.1x detects the violation, the action is to err-disable the port.
- If port security detects the violation, the action is to shutdown or restrict the port (the action is configurable).

The following describes when port security and 802.1x security violations occur:

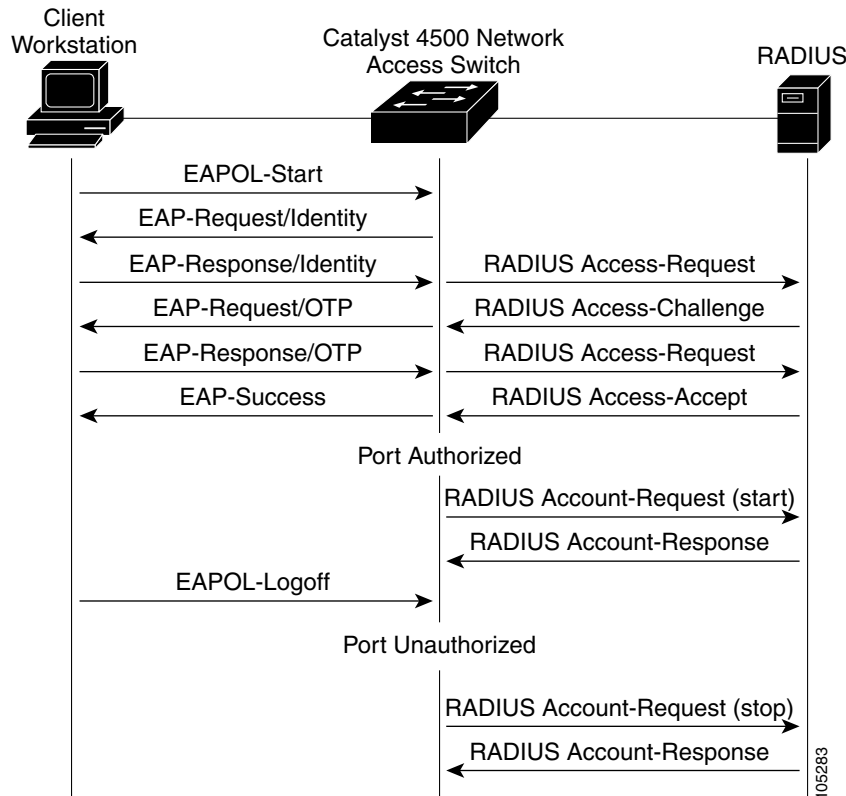
- in single host mode, after the port is authorized, any MAC address received other than the client's will cause a 802.1x security violation.
- in single host mode, if installation of an 802.1x client's MAC address fails because port security has already reached its limit (due to a configured secure MAC addresses), a port security violation is triggered.
- in multi host mode, once the port is authorized, any additional MAC addresses that cannot be installed because the port security has reached its limit will trigger a port security violation.
- When an 802.1x client logs off, the port transitions back to an unauthenticated state, and all dynamic entries in the secure host table are cleared, including the entry for the client. Normal authentication then ensues.
- If you administratively shut down the port, the port becomes unauthenticated, and all dynamic entries are removed from the secure host table.
- Only 802.1x can remove the client's MAC address from the port security table. Note that in multi host mode, with the exception of the client's MAC address, all MAC addresses that are learned by port security can be deleted using port security CLIs.
- Whenever port security ages out a 802.1x client's MAC address, 802.1x attempts to reauthenticate the client. Only if the reauthentication succeeds will the client's MAC address be retained in the port security table.
- All of the 802.1x client's MAC addresses are tagged with (dot1x) when you display the port security table by using CLI.

802.1x RADIUS Accounting

802.1x RADIUS accounting relays important events to the RADIUS server (such as the client's connection session). This session is defined as the difference in time from when client is authorized to use the port and when the client stops using the port.

Figure 28-3 shows the 802.1x device roles.

Figure 28-3 Radius Accounting

**Note**

You must configure the 802.1x client to send an EAP-logoff (Stop) message to the switch when the user logs off. If you do not configure the 802.1x client, an EAP-logoff message is not sent to the switch and the accompanying accounting Stop message will not be sent to the authentication server. Refer to the Microsoft Knowledge Base article at the URL: <http://support.microsoft.com>. Also refer to the Microsoft article at the URL:

<http://www.microsoft.com/technet/treeview/default.asp?url=/technet/columns/cableguy/cg0703.asp>, and set the SupplicantMode registry to 3 and the AuthMode registry to 1.

The client uses EAP to authenticate itself with the RADIUS server. The switch relays EAP packets between the client and the RADIUS server.

After the client is authenticated, the switch sends accounting-request packets to the RADIUS server, which responds with accounting-response packets to acknowledge the receipt of the request.

A RADIUS accounting-request packet contains one or more Attribute-Value pairs to report various events and related information to the RADIUS server. The following events are tracked:

- User successfully authenticates
- User logs-off
- Link-down occurs on a 802.1x port
- Reauthentication succeeds
- Reauthentication fails

When the port's state transitions between authorized and unauthorized, the RADIUS messages are transmitted to the RADIUS server.

The switch does not log any accounting information. Instead, it sends such information to the RADIUS server, which must be configured to log accounting messages.

The 802.1x authentication, authorization and accounting process is as follows:

-
- | | |
|---------------|---|
| Step 1 | A user connects to a port on the switch. |
| Step 2 | Authentication is performed, for example, using the username/password method. |
| Step 3 | VLAN assignment is enabled, as appropriate, per RADIUS server configuration. |
| Step 4 | The switch sends a start message to an accounting server. |
| Step 5 | Reauthentication is performed, as necessary. |
| Step 6 | The switch sends an interim accounting update to the accounting server that is based on the result of reauthentication. |
| Step 7 | The user disconnects from the port. |
| Step 8 | The switch sends a stop message to the accounting server. |
-

To configure 802.1x accounting, you need to do the following tasks:

- Enable logging of “Update/Watchdog packets from this AAA client” in your RADIUS server’s Network Configuration tab.
- Enable “Logging>CVS RADIUS Accounting” in your RADIUS server’s System Configuration tab.
- Enable 802.1x accounting on your switch.
- Enable AAA accounting by using the **aaa system accounting** command. Refer to the [“Enabling 802.1x Accounting” section on page 28-15](#).

Enabling AAA system accounting along with 802.1x accounting allows system reload events to be sent to the accounting RADIUS server for logging. By doing this, the accounting RADIUS server can infer that all active 802.1x sessions are appropriately closed.

Because RADIUS uses the unreliable transport protocol UDP, accounting messages may be lost due to poor network conditions. If the switch does not receive the accounting response message from the RADIUS server after a configurable number of retransmissions of an accounting request, the following system message appears:

```
Accounting message %s for session %s failed to receive Accounting Response.
```

When the stop message is not transmitted successfully, the following message appears:

```
00:09:55: %RADIUS-3-NOACCOUNTINGRESPONSE: Accounting message Start for session  
172.20.50.145 sam 11/06/03 07:01:16 11000002 failed to receive Accounting Response.
```

Use the **show radius statistics** command to display the number of RADIUS messages that do not receive the accounting response message.

Supported Topologies


The 802.1x port-based authentication supports two topologies:

- Point to point
- Wireless LAN

In a point-to-point configuration (see [Figure 28-1 on page 28-2](#)), only one client can be connected to the 802.1x-enabled switch port when the multi-host mode is not enabled (the default). The switch detects the client when the port link state changes to the up state. If a client leaves or is replaced with another client, the switch changes the port link state to down, and the port returns to the unauthorized state.

[Figure 28-4](#) illustrates 802.1x port-based authentication in a wireless LAN. You must configure the 802.1x port as a multiple-host port that is authorized as a wireless access point once the client is authenticated. (See the [“Enabling Multiple Hosts” section on page 28-20](#).) When the port is authorized, all other hosts that are indirectly attached to the port are granted access to the network. If the port becomes unauthorized (reauthentication fails or an EAPOL-logoff message is received), the switch denies access to the network for all wireless access point-attached clients. In this topology, the wireless access point is responsible for authenticating clients attached to it, and the wireless access point acts as a client to the switch.

Figure 28-4 Wireless LAN Example



How to Configure 802.1x

These sections describe how to configure 802.1x:

- [Default 802.1x Configuration, page 28-11](#)
- [802.1x Configuration Guidelines, page 28-12](#)
- [Enabling 802.1x Authentication, page 28-12 \(required\)](#)
- [Configuring Switch-to-RADIUS-Server Communication, page 28-14 \(required\)](#)
- [Enabling 802.1x Accounting, page 28-15](#)
- [Configuring 802.1x with Guest VLANs, page 28-16](#)
- [Enabling Periodic Reauthentication, page 28-16 \(optional\)](#)
- [Manually Reauthenticating a Client Connected to a Port, page 28-17 \(optional\)](#)
- [Changing the Quiet Period, page 28-17 \(optional\)](#)
- [Changing the Switch-to-Client Retransmission Time, page 28-18 \(optional\)](#)
- [Setting the Switch-to-Client Frame-Retransmission Number, page 28-19 \(optional\)](#)
- [Enabling Multiple Hosts, page 28-20 \(optional\)](#)
- [Resetting the 802.1x Configuration to the Default Values, page 28-20 \(optional\)](#)

Default 802.1x Configuration

Table 28-1 shows the default 802.1x configuration.

Table 28-1 Default 802.1x Configuration

Feature	Default Setting
Authentication, authorization, and accounting (AAA)	Disabled
RADIUS server <ul style="list-style-type: none"> • IP address • UDP authentication port • Key 	<ul style="list-style-type: none"> • None specified • 1812 • None specified
Per-interface 802.1x protocol enable state	Disabled (force-authorized) The port transmits and receives normal traffic without 802.1x-based authentication of the client.
Periodic reauthentication	Disabled
Time between reauthentication attempts	3600 sec
Quiet period	60 sec Number of seconds that the switch remains in the quiet state following a failed authentication exchange with the client.
Retransmission time	30 sec Number of seconds that the switch should wait for a response to an EAP request/identity frame from the client before retransmitting the request.
Maximum retransmission number	2 Number of times that the switch will send an EAP-request/identity frame before restarting the authentication process.
Multiple host support	Disabled
Client timeout period	30 sec When relaying a request from the authentication server to the client, the amount of time that the switch waits for a response before retransmitting the request to the client.
Authentication server timeout period	30 sec When relaying a response from the client to the authentication server, the amount of time that the switch waits for a reply before retransmitting the response to the server. This setting is not configurable.

802.1x Configuration Guidelines

This section describes the guidelines for configuring 802.1x authentication:

- The 802.1x protocol is supported on both Layer 2 static-access ports and Layer 3 routed ports, but it is not supported on the following port types:
 - Trunk port—If you try to enable 802.1x on a trunk port, an error message appears, and 802.1x is not enabled. If you try to change the mode of an 802.1x-enabled port to trunk, the port mode is not changed.
 - Default ports—All ports default as dynamic-access ports (auto). Use the **no switchport** command to access a router port.
 - Dynamic ports—A port in dynamic mode can negotiate with its neighbor to become a trunk port. If you try to enable 802.1x on a dynamic port, an error message appears, and 802.1x is not enabled. If you try to change the mode of an 802.1x-enabled port to dynamic, the port mode is not changed.
 - EtherChannel port—Before enabling 802.1x on the port, you must first remove it from the EtherChannel. If you try to enable 802.1x on an EtherChannel or on an active port in an EtherChannel, an error message appears, and 802.1x is not enabled. If you enable 802.1x on a not-yet active port of an EtherChannel, the port does not join the EtherChannel.
 - Switched Port Analyzer (SPAN) destination port—You can enable 802.1x on a port that is a SPAN destination port; however, 802.1x is disabled until the port is removed as a SPAN destination. You can enable 802.1x on a SPAN source port.

If you are planning to use either 802.1x accounting or VLAN assignment, be aware that both features utilize general AAA commands. For information how to configure AAA, refer to “Enabling 802.1x Authentication” on page 12 and “Enabling 802.1x Accounting” on page 15. Alternatively, you can refer to the Cisco IOS security documentation.

Refer to the following Cisco IOS security documentation for information on how to configure AAA system accounting:

- http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_c/index.htm
- http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_r/index.htm

Enabling 802.1x Authentication

To enable 802.1x port-based authentication, you first enable AAA, then specify the authentication method list. A method list describes the sequence and authentication methods that must be queried to authenticate a user.

The software uses the first method listed in the method list to authenticate users; if that method fails to respond, the software selects the next authentication method in the list. This process continues until there is successful communication with a listed authentication method or until all defined methods are exhausted. If authentication fails at any point in this cycle, the authentication process stops, and no other authentication methods are attempted.

To allow VLAN assignment, you must enable AAA authorization to configure the switch for all network-related service requests.

To configure 802.1x port-based authentication, perform this task:

	Command	Purpose
Step 1	Switch# configure terminal	Enters global configuration mode.
Step 2	Switch(config)# aaa new-model	Enables AAA.
Step 3	Switch(config)# aaa authentication dot1x {default} method1 [method2...]	Creates an 802.1x authentication method list. To create a default list that is used when a named list is not specified in the authentication command, use the default keyword followed by the methods that are to be used in default situations. The default method list is automatically applied to all interfaces. Enter at least one of these keywords: <ul style="list-style-type: none"> • group radius—Use the list of all RADIUS servers for authentication. • none—Use no authentication. The client is automatically authenticated by the switch without using the information supplied by the client.
Step 4	Switch(config)# aaa authorization network {default} group radius	(Optional) Configure the switch for user RADIUS authorization for all network-related service requests, such as VLAN assignment.
Step 5	Switch(config)# interface interface-id	Enters interface configuration mode and specifies the interface to be enabled for 802.1x authentication.
Step 6	Switch(config-if)# dot1x port-control auto	Enables 802.1x authentication on the interface. For feature interaction information with trunk, dynamic, dynamic-access, EtherChannel, secure, and SPAN ports, see the “802.1x Configuration Guidelines” section on page 28-12.
Step 7	Switch(config-if)# end	Returns to privileged EXEC mode.
Step 8	Switch # show dot1x all	Verifies your entries. Check the Status column in the 802.1x Port Summary section of the display. An enabled status means that the port-control value is set either to auto or to force-unauthorized .
Step 9	Switch# show running-config	Verifies your entries.
Step 10	Switch# copy running-config startup-config	(Optional) Saves your entries in the configuration file.

To disable AAA, use the **no aaa new-model** global configuration command.

To disable 802.1x AAA authentication, use the **no aaa authentication dot1x {default | list-name} method1 [method2...]** global configuration command.

To disable 802.1x authentication, use the **dot1x port-control force-authorized** or the **no dot1x port-control** interface configuration command.

This example shows how to enable AAA and 802.1x on Fast Ethernet port 2/1:

```
Switch# configure terminal
Switch(config)# aaa new-model
Switch(config)# aaa authentication dot1x default group radius
Switch(config)# interface fastethernet2/1
Switch(config-if)# dot1x port-control auto
Switch(config-if)# end
```

Configuring Switch-to-RADIUS-Server Communication

A RADIUS security server is identified by its host name or IP address, host name and specific UDP port number, or IP address and specific UDP port numbers. The combination of the IP address and UDP port number creates a unique identifier, which enables RADIUS requests to be sent to multiple UDP ports on a server at the same IP address. If two different host entries on the same RADIUS server are configured for the same service—for example, authentication—the second host entry configured acts as the fail-over backup to the first one. The RADIUS host entries are tried in the order they were configured.

To configure the RADIUS server parameters on the switch, perform this task:

	Command	Purpose
Step 1	Switch# configure terminal	Enters global configuration mode.
Step 2	Switch(config)# radius-server host {hostname ip-address} auth-port port-number [acct-port port-number] key string	<p>Configures the RADIUS server parameters on the switch.</p> <p>For <i>hostname ip-address</i>, specify the hostname or IP address of the remote RADIUS server.</p> <p>For auth-port <i>port-number</i>, specify the UDP destination port for authentication requests. The default is 1812.</p> <p>For acct-port <i>port-number</i>, specify the UDP destination port for accounting requests. The default is 1813.</p> <p>For key string, specify the authentication and encryption key used between the switch and the RADIUS daemon running on the RADIUS server. The key is a text string that must match the encryption key used on the RADIUS server.</p> <p>Note Always configure the key as the last item in the radius-server host command syntax because leading spaces are ignored, but spaces within and at the end of the key are used. If you use spaces in the key, do not enclose the key in quotation marks unless the quotation marks are part of the key. This key must match the encryption used on the RADIUS daemon.</p> <p>If you want to use multiple RADIUS servers, reenter this command.</p>
Step 3	Switch(config-if)# ip radius source-interface m/p	Establishes the IP address to be used as the source address for all outgoing RADIUS packets.
Step 4	Switch(config)# end	Returns to privileged EXEC mode.
Step 5	Switch# show running-config	Verifies your entries.
Step 6	Switch# copy running-config startup-config	(Optional) Saves your entries in the configuration file.

To delete the specified RADIUS server, use the **no radius-server host** {hostname | ip-address} global configuration command.

This example shows how to specify the server with IP address 172.20.39.46 as the RADIUS server. The first command specifies port 1612 as the authorization port, sets the encryption key to rad123. The second command dictates that key matches will be performed on the RADIUS server:

```
Switch(config)# radius-server host 172.120.39.46 auth-port 1612 key rad123
Switch(config)# ip radius source-interface m/p
```

You can globally configure the timeout, retransmission, and encryption key values for all RADIUS servers by using the **radius-server host** global configuration command. If you want to configure these options on a per-server basis, use the **radius-server timeout**, **radius-server retransmit**, and the **radius-server key** global configuration commands.

You also need to configure some settings on the RADIUS server. These settings include the IP address of the switch and the key string to be shared by both the server and the switch.

Refer to the following Cisco IOS security documentation for information on how to configure AAA system accounting:

- http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_c/index.htm
- http://www.cisco.com/univercd/cc/td/doc/product/software/ios122/122cgcr/fsecur_r/index.htm

Enabling 802.1x Accounting



Note

If you plan to implement system-wide accounting, you should also configure 802.1x accounting. Moreover, you need to inform the accounting server of the system reload event when the system is reloaded. Doing this, ensures that the accounting server knows that all outstanding 802.1x sessions on this system are closed.

Once you configure 802.1x authentication and switch-to-RADIUS server communication, perform this task to enable 802.1x accounting:

	Command	Purpose
Step 1	Switch # configure terminal	Enters global configuration mode.
Step 2	Switch(config)# aaa accounting dot1x default start-stop group radius	Enables 802.1x accounting, using the list of all RADIUS servers.
Step 3	Switch(config)# clock timezone PST -8	Sets the time zone for the accounting event-time stamp field.
Step 4	Switch(config)# clock calendar-valid	Enables the date for the accounting event-time stamp field.
Step 5	Switch(config-if)# aaa accounting system default start-stop group radius	(Optional) Enables system accounting (using the list of all RADIUS servers) and generates system accounting reload event messages when the switch reloads.
Step 6	Switch(config-if)# end	Returns to privileged EXEC mode.
Step 7	Switch # show running-config	Verifies your entries.
Step 8	Switch # copy running-config startup-config	(Optional) Saves your entries in the configuration file.

This example shows how to configure 802.1x accounting. The first command configures the RADIUS server, specifying 1813 as the UDP port for accounting:

```
Switch(config)# radius-server host 172.120.39.46 auth-port 1812 acct-port 1813 key rad123
Switch(config)# aaa accounting dot1x default start-stop group radius
Switch(config)# aaa accounting system default start-stop group radius
```

**Note**

You must configure the RADIUS server to perform accounting tasks, such as logging start, stop, and interim-update messages and time stamps. To turn on these functions, enable logging of “Update/Watchdog packets from this AAA client” in your RADIUS server Network Configuration tab. Next, enable “CVS RADIUS Accounting” in your RADIUS server System Configuration tab.

Configuring 802.1x with Guest VLANs

**Note**

When a port is put into a guest VLAN, it is automatically placed into multihost mode, and an unlimited number of hosts can connect through the port. Changing the multihost configuration does not effect a port in a guest VLAN.

To configure 802.1x with guest-VLAN, perform this task:

	Command	Purpose
Step 1	Switch # configure terminal	Enters global configuration mode.
Step 2	Switch(config-if)# interface <i>interface-id</i>	Enters interface configuration mode and specifies the interface to be enabled for 802.1x authentication.
Step 3	Switch(config-if)# dot1x port-control auto]	Enables 802.1x authentication on the interface. For feature interaction information with trunk, dynamic, dynamic-access, EtherChannel, secure, and SPAN ports, see the “802.1x Configuration Guidelines” section on page 28-12 .
Step 4	Switch(config-if)# dot1x guest-vlan <i>vlan-id</i>	Enables a guest VLAN on a particular interface.
Step 5	Switch(config-if)# end	Returns to configuration mode.
Step 6	Switch(config)# end	Returns to privileged EXEC mode.

To disable the guest VLAN feature on a particular port, use the **no dot1x guest-vlan** interface configuration command.

This example shows how to enable a guest VLAN on Fast Ethernet interface 4/3:

```
Switch# configure terminal
Switch(config)# interface fastethernet4/3
Switch(config-if)# dot1x port-control auto
Switch(config-if)# dot1x guest-vlan 50
Switch(config-if)# end
Switch(config)# end
Switch#
```

Enabling Periodic Reauthentication

You can enable periodic 802.1x client reauthentication and specify how often it occurs. If you do not specify a time value before enabling reauthentication, the interval between reauthentication attempts is 3600 seconds.

Automatic 802.1x client reauthentication is a per-interface setting and can be set for clients connected to individual ports. To manually reauthenticate the client connected to a specific port, see the “[Manually Reauthenticating a Client Connected to a Port](#)” section on page 28-17.

To enable periodic reauthentication of the client and to configure the number of seconds between reauthentication attempts, perform this task:

	Command	Purpose
Step 1	Switch# configure terminal	Enters global configuration mode.
Step 2	Switch(config)# interface <i>interface-id</i>	Enters interface configuration mode and specifies the interface to be enabled for periodic reauthentication.
Step 3	Switch(config-if)# dot1x re-authentication	Enables periodic reauthentication of the client, which is disabled by default.
Step 4	Switch(config)# dot1x timeout reauth-period <i>seconds</i>	Specifies the number of seconds between reauthentication attempts. The range is 1 to 65,535; the default is 3600 seconds. This command affects the behavior of the switch only if periodic reauthentication is enabled.
Step 5	Switch(config)# end	Returns to privileged EXEC mode.
Step 6	Switch# show dot1x all	Verifies your entries.
Step 7	Switch(config)# copy running-config startup-config	(Optional) Saves your entries in the configuration file.

To disable periodic reauthentication, use the **no dot1x re-authentication** interface configuration command. To return to the default number of seconds between reauthentication attempts, use the **no dot1x timeout reauth-period** global configuration command.

This example shows how to enable periodic reauthentication and set the number of seconds between reauthentication attempts to 4000:

```
Switch(config)# dot1x timeout reauth-period 4000
Switch(config)# dot1x re-authentication
```

Manually Reauthenticating a Client Connected to a Port

You can manually reauthenticate a client connected to a specific port at any time by entering the **dot1x re-authenticate interface** *interface-id* privileged EXEC command. If you want to enable or disable periodic reauthentication, see the “[Enabling Periodic Reauthentication](#)” section on page 28-16.

This example shows how to manually reauthenticate the client connected to Fast Ethernet port 1/1:

```
Switch# dot1x re-authenticate interface fastethernet1/1
Starting reauthentication on FastEthernet1/1
```

Changing the Quiet Period

When the switch cannot authenticate the client, the switch remains idle for a set period of time, and then tries again. The idle time is determined by the **quiet-period** value. A failed authentication of the client might occur because the client provided an invalid password. You can provide a faster response time to the user by entering a number smaller than the default.

To change the quiet period, perform this task:

	Command	Purpose
Step 1	Switch# configure terminal	Enters global configuration mode.
Step 2	Switch(config)# interface <i>interface-id</i>	Enters interface configuration mode and specifies the interface to be enabled for timeout quiet-period .
Step 3	Switch(config)# dot1x timeout quiet-period <i>seconds</i>	Sets the number of seconds that the switch remains in the quiet-period following a failed authentication exchange with the client. The range is 0 to 65,535 seconds; the default is 60.
Step 4	Switch(config)# end	Returns to privileged EXEC mode.
Step 5	Switch# show dot1x all	Verifies your entries.
Step 6	Switch# copy running-config startup-config	(Optional) Saves your entries in the configuration file.

To return to the default quiet-period, use the **no dot1x timeout quiet-period** configuration command. This example shows how to set the **quiet-period** on the switch to 30 seconds:

```
Switch(config)# dot1x timeout quiet-period 30
```

Changing the Switch-to-Client Retransmission Time

The client responds to the EAP-request/identity frame from the switch with an EAP-response/identity frame. If the switch does not receive this response, it waits a set period of time (known as the retransmission time) and then retransmits the frame.



Note

You should change the default value of this command only to adjust for unusual circumstances, such as unreliable links or specific behavioral problems with certain clients and authentication servers.

To change the amount of time that the switch waits for client notification, perform this task:

	Command	Purpose
Step 1	Switch# configure terminal	Enters global configuration mode.
Step 2	Switch(config)# interface <i>interface-id</i>	Enters interface configuration mode and specifies the interface to be enabled for timeout tx-period.
Step 3	Switch(config-if)# dot1x timeout tx-period <i>seconds</i>	Sets the number of seconds that the switch waits for a response to an EAP-request/identity frame from the client before retransmitting the request. The range is 1 to 65,535 seconds; the default is 30.
Step 4	Switch(config)# end	Returns to privileged EXEC mode.
Step 5	Switch# show dot1x all	Verifies your entries.
Step 6	Switch# copy running-config startup-config	(Optional) Saves your entries in the configuration file.

To return to the default retransmission time, use the **no dot1x timeout tx-period** interface configuration command.

This example shows how to set the retransmission time to 60 seconds:

```
Switch(config)# dot1x timeout tx-period 60
```

Setting the Switch-to-Client Frame-Retransmission Number

In addition to changing the switch-to-client retransmission times, you can change the number of times that the switch sends EAP-Request/Identity and other EAP-Request frames to the client before restarting the authentication process. The number of EAP-Request/Identity retransmissions is controlled by the **dot1x max-reauth-req** command; the number of retransmissions for other EAP-Request frames is controlled by the **dot1x max-req** command.



Note

You should change the default values of these commands only to adjust for unusual circumstances such as unreliable links or specific behavioral problems with certain clients and authentication servers.

To set the switch-to-client frame-retransmission numbers, perform this task:

	Command	Purpose
Step 1	Switch# configure terminal	Enters global configuration mode.
Step 2	Switch(config)# interface <i>interface-id</i>	Enters interface configuration mode and specifies the interface to be enabled for max-reauth-req and/or max-req .
Step 3	Switch(config-if)# dot1x max-req <i>count</i>	Specifies the number of times that the switch retransmits an EAP-request frame of a type other than EAP-request/identity to the client before restarting the authentication process. The range for <i>count</i> is 1 to 10; the default is 2.
	or Switch(config-if)# dot1x max-req <i>count</i>	
Step 4	Switch(config)# end	Returns to privileged EXEC mode.
Step 5	Switch# show dot1x all	Verifies your entries.
Step 6	Switch# copy running-config startup-config	(Optional) Saves your entries in the configuration file.

To return to the default retransmission number, use the **no dot1x max-req** and **no dot1x max-reauth-req** global configuration command.

This example shows how to set 5 as the number of times that the switch retransmits an EAP-request/identity request before restarting the authentication process:

```
Switch(config)# dot1x max-reauth-req 5
```

Enabling Multiple Hosts

You can attach multiple hosts to a single 802.1x-enabled port as shown in [Figure 28-4 on page 28-10](#). In this mode, only one of the attached hosts must be successfully authorized for all hosts to be granted network access. If the port becomes unauthorized (reauthentication fails or an EAPOL-logoff message is received), all attached clients are denied access to the network.

To allow multiple hosts (clients) on an 802.1x-authorized port that has the **dot1x port-control** interface configuration command set to **auto**, perform this task:

	Command	Purpose
Step 1	Switch# configure terminal	Enters global configuration mode.
Step 2	Switch(config)# interface interface-id	Enters interface configuration mode and specifies the interface to which multiple hosts are indirectly attached.
Step 3	Switch(config-if)# dot1x multiple-hosts	Allows multiple hosts (clients) on an 802.1x-authorized port. Make sure that the dot1x port-control interface configuration command set is set to auto for the specified interface.
Step 4	Switch(config-if)# end	Returns to privileged EXEC mode.
Step 5	Switch# show dot1x all interface interface-id	Verifies your entries.
Step 6	Switch# copy running-config startup-config	(Optional) Saves your entries in the configuration file.

To disable multiple hosts on the port, use the **no dot1x multiple-hosts** interface configuration command.

This example shows how to enable 802.1x on Fast Ethernet interface 0/1 and to allow multiple hosts:

```
Switch(config)# interface fastethernet0/1
Switch(config-if)# dot1x port-control auto
Switch(config-if)# dot1x multiple-hosts
```

Resetting the 802.1x Configuration to the Default Values

To reset the 802.1x configuration to the default values, perform this task:

	Command	Purpose
Step 1	Switch# configure terminal	Enters global configuration mode.
Step 2	Switch(config)# dot1x default	Resets the configurable 802.1x parameters to the default values.
Step 3	Switch(config)# end	Returns to privileged EXEC mode.
Step 4	Switch# show dot1x all	Verifies your entries.
Step 5	Switch# copy running-config startup-config	(Optional) Saves your entries in the configuration file.

Displaying 802.1x Statistics and Status

To display 802.1x statistics for all interfaces, use the **show dot1x statistics** privileged EXEC command. To display 802.1x statistics for a specific interface, use the **show dot1x statistics interface *interface-id*** privileged EXEC command.

To display the 802.1x administrative and operational status for the switch, use the **show dot1x all** privileged EXEC command. To display the 802.1x administrative and operational status for a specific interface, use the **show dot1x interface *interface-id*** privileged EXEC command.

