Configuring RADIUS

This chapter contains the following sections:

Information About RADIUS

The RADIUS distributed client/server system allows you to secure networks against unauthorized access. In the Cisco implementation, RADIUS clients run on Cisco NX-OS devices and send authentication and accounting requests to a central RADIUS server that contains all user authentication and network service access information.

RADIUS Network Environments

RADIUS can be implemented in a variety of network environments that require high levels of security while maintaining network access for remote users.

You can use RADIUS in the following network environments that require access security:

- Networks with multiple-vendor network devices, each supporting RADIUS. For example, network devices from several vendors can use a single RADIUS server-based security database.
- Networks already using RADIUS. You can add a Cisco NX-OS device with RADIUS to the network. This action might be the first step when you make a transition to a AAA server.

- Networks that require resource accounting. You can use RADIUS accounting independent of RADIUS authentication or authorization. The RADIUS accounting functions allow data to be sent at the start and end of services, indicating the amount of resources (such as time, packets, bytes, and so on) used during the session. An Internet service provider (ISP) might use a freeware-based version of the RADIUS access control and accounting software to meet special security and billing needs.
- Networks that support authentication profiles. Using the RADIUS server in your network, you can configure AAA authentication and set up per-user profiles. Per-user profiles enable the Cisco NX-OS device to better manage ports using their existing RADIUS solutions and to efficiently manage shared resources to offer different service-level agreements.

RADIUS Operation

When a user attempts to log in and authenticate to a Cisco NX-OS device using RADIUS, the following occurs:

1. The user is prompted for and enters a username and password.
2. The username and encrypted password are sent over the network to the RADIUS server.
3. The user receives one of the following responses from the RADIUS server:
    - ACCEPT**—**The user is authenticated.
    - REJECT**—**The user is not authenticated and is prompted to reenter the username and password, or access is denied.
    - CHALLENGE**—**A challenge is issued by the RADIUS server. The challenge collects additional data from the user.
    - CHANGE PASSWORD**—**A request is issued by the RADIUS server, asking the user to select a new password.

The ACCEPT or REJECT response is bundled with additional data that is used for EXEC or network authorization. You must first complete RADIUS authentication before using RADIUS authorization. The additional data included with the ACCEPT or REJECT packets consists of the following:

- Services that the user can access, including Telnet, rlogin, or local-area transport (LAT) connections, and Point-to-Point Protocol (PPP), Serial Line Internet Protocol (SLIP), or EXEC services.
- Connection parameters, including the host or client IPv4 address, access list, and user timeouts.

RADIUS Server Monitoring

An unresponsive RADIUS server can cause a delay in processing AAA requests. You can periodically monitor a RADIUS server to check whether it is responding (or alive) to save time in processing AAA requests. Unresponsive RADIUS servers are marked as dead and are not sent AAA requests. Dead RADIUS servers are periodically monitored and returned to the alive state once they respond. This monitoring process verifies that a RADIUS server is in a working state before real AAA requests are sent its way. Whenever a RADIUS server changes to the dead or alive state, a Simple Network Management Protocol (SNMP) trap is generated and an error message is displayed indicating that a failure is taking place.

Note The monitoring interval for alive servers and dead servers are different and can be configured by the user. The RADIUS server monitoring is performed by sending a test authentication request to the RADIUS server.

Figure 1: Radius Server States

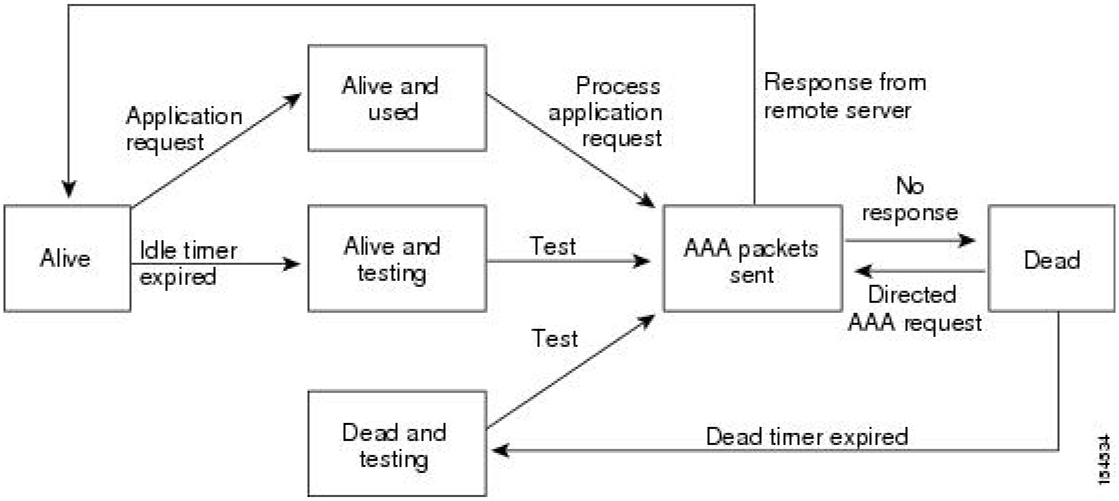

Vendor-Specific Attributes

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating vendor-specific attributes (VSAs) between the network access server and the RADIUS server. The IETF uses attribute 26. VSAs allow vendors to support their own extended attributes that are not suitable for general use. The Cisco RADIUS implementation supports one vendor-specific option using the format recommended in the specification. The Cisco vendor ID is 9, and the supported option is vendor type 1, which is named cisco-av-pair. The value is a string with the following format:

```
protocol : attribute separator value *
```

The protocol is a Cisco attribute for a particular type of authorization. The separator is = (equal sign) for mandatory attributes and * (asterisk) indicates optional attributes.

When you use RADIUS servers for authentication, the RADIUS protocol directs the RADIUS server to return user attributes, such as authorization information, with authentication results. This authorization information is specified through VSAs.

The following VSA protocol options are supported:

- Shell—Protocol used in access-accept packets to provide user profile information.
- Accounting—Protocol used in accounting-request packets. If a value contains any white spaces, you should enclose the value within double quotation marks.

The following attributes are supported:

- roles—Lists all the roles to which the user belongs. The value field is a string that lists the role names delimited by white space. For example, if the user belongs to roles network-operator and vdc-admin, the value field would be “network-operator vdc-admin.” This attribute, which the RADIUS server sends in the VSA portion of the Access-Accept frames, can be used only with the shell protocol value. The following examples show the roles attribute as supported by Cisco Access Control System (ACS):

  ```
  shell:roles=“network-operator vdc-admin”
  ```

  ```
  shell:roles*“network-operator vdc-admin”
  ```

  The following examples show the roles attribute as supported by FreeRADIUS:

  ```
  Cisco-AVPair = ”shell:roles=\“network-operator vdc-admin\””
  ```

  ```
  Cisco-AVPair = “shell:roles*\“network-operator vdc-admin\””
  ```

  If you are using Cisco ACS and intend to use the same ACS group for both Cisco Nexus 1000V and Cisco UCS authentication, use the following roles attribute:

  ```
  cisco-av-pair*shell:roles="network-admin admin"
  ```

  Note When you specify a VSA as shell:roles*"network-operator vdc-admin" or "shell:roles*\"network-operator vdc-admin\"", this VSA is flagged as an optional attribute and other Cisco devices ignore this attribute.

- accountinginfo—Stores accounting information in addition to the attributes covered by a standard RADIUS accounting protocol. This attribute is sent only in the VSA portion of the Account-Request frames from the RADIUS client on the switch. It can be used only with the accounting protocol data units (PDUs).

Prerequisites for RADIUS

- You already know the RADIUS server IP addresses or hostnames.
- You already know the key(s) used to secure RADIUS communication in your network.
- The device is already configured as a RADIUS client of the AAA servers.

Guidelines and Limitations

You can configure a maximum of 64 RADIUS servers.

Default Settings

Table 1: Default RADIUS Parameters

Parameters	Default
Server roles	Authentication and accounting

Parameters	Default
Dead timer interval	0 minutes
Retransmission count	1
Retransmission timer interval	5 seconds
Idle timer interval	0 minutes
Periodic server monitoring username	test
Periodic server monitoring password	test

Configuring RADIUS Servers

Configuring RADIUS Server Hosts

You can configure the IP address or the hostname for each RADIUS server to be used for authentication. You should know the following information:

- You can configure up to 64 RADIUS servers.
- All RADIUS server hosts are automatically added to the default RADIUS server group.

Before You Begin

Log in to the CLI in EXEC mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **radius-server host** {*ipv4-address* | *hostname*}
3. switch(config)# **exit**
4. (Optional) switch# **show radius-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **radius-server host** {*ipv4-address* | *hostname*}	Defines the IP address or hostname for the RADIUS server or the RADIUS server Domain Name Server (DNS) name. *ipv4-address*—The IP address for the RADIUS server.

	Command or Action	Purpose
		hostname—The hostname for the RADIUS server. The hostname is alphanumeric, case sensitive, and has a maximum of 256 characters.
Step 3	switch(config)# **exit**	Returns you to the EXEC mode.
Step 4	switch# **show radius-server**	(Optional) Displays the RADIUS server configuration
Step 5	switch# **copy running-config startup-config**	(Optional) Copies the running configuration to the startup configuration.

The following example configures a RADIUS server host:

```
switch# configure terminal
switch(config)# radius-server host 10.10.1.1
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Configuring the Global RADIUS Key

You can configure the key that is used by all RADIUS servers to authenticate with the Cisco Nexus 1000V.

You must know the global key that is used for RADIUS server authentication.

Before You Begin

Log in to the CLI in EXEC mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **radius-server key** [**0** | **7**] *key-value*
3. switch(config)# **exit**
4. (Optional) switch# **show radius-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **radius-server key** [**0** | **7**] *key-value*	Specifies a preshared key for all RADIUS servers. You can specify a cleartext (0) or encrypted (7) preshared key. The default format is cleartext. *key-value*—The preshared key value. The maximum length is 63 characters. By default, no preshared key is configured.

	Command or Action	Purpose
Step 3	switch(config)# **exit**	Returns you to the EXEC mode.
Step 4	switch# **show radius-server**	(Optional) Displays the RADIUS server configuration. **Note** The preshared keys are saved in encrypted form in the running configuration. Use the show running-config command to display the encrypted preshared keys.
Step 5	switch# **copy running-config startup-config**	(Optional) Copies the running configuration to the startup configuration.

The follow example configures the global RADIUS key:

```
switch# configure terminal
switch(config)# radius-server key 0 QsEfThUkO
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Configuring a RADIUS Server Key

You can configure a key for a single RADIUS server host.

You must have the key for the remote RADIUS host.

Before You Begin

Log in to the CLI in EXEC mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **radius-server host** {*ipv4-address* | *hostname*} **key** [**0** | **7**] *key-value*
3. switch(config)# **exit**
4. (Optional) switch# **show radius-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **radius-server host** {*ipv4-address* | *hostname*} **key** [**0** | **7**] *key-value*	Specifies a preshared key for a specific RADIUS server. You can specify a cleartext (**0**) or encrypted (**7**) preshared key. The default format is cleartext. *ipv4-address*—The IP address for the RADIUS server.

	Command or Action	Purpose
		hostname—The hostname for the RADIUS server. The hostname is alphanumeric, case sensitive, and has a maximum of 256 characters. *key-value*—The preshared key value. The maximum length is 63 characters.
Step 3	switch(config)# **exit**	Returns you to the EXEC mode.
Step 4	switch# **show radius-server**	(Optional) Displays the RADIUS server configuration. **Note** The preshared keys are saved in encrypted form in the running configuration. Use the show running-config command to display the encrypted preshared keys.
Step 5	switch# **copy running-config startup-config**	(Optional) Copies the running configuration to the startup configuration.

The following example configures a RADIUS server key:

```
switch# configure terminal
switch(config)# radius-server host 10.10.1.1 key 0 PlIjUhYg
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Configuring RADIUS Server Groups

You can configure a RADIUS server group whose member servers share authentication functions.

The servers in the group are tried in the same order in which you configure them

Before You Begin

- Log in to the CLI in EXEC mode.
- Know that all servers in a RADIUS server group must belong to the RADIUS protocol.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **aaa group server radius** *group-name*
3. switch(config-radius)# **server** {*ipv4-address* | *server-name*}
4. (Optional) switch(config-radius)# **deadtime** *minutes*
5. (Optional) switch(config-radius)# **use-vrf** *vrf-name*
6. (Optional) switch(config-radius)# **source-interface** {*interface-type*} {*interface-number*}
7. (Optional) switch(config-radius)# **show radius-server groups** [*group-name*]
8. (Optional) switch(config-radius)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **aaa group server radius** *group-name*	Creates a RADIUS server group and enters RADIUS server group configuration mode for that group. *group-name*—The name of the server group. The name is a case-sensitive, alphanumeric string with a maximum length of 127 characters.
Step 3	switch(config-radius)# **server** {*ipv4-address* | *server-name*}	Configures the RADIUS server as a member of the RADIUS server group. *ipv4-address*—The IP address for the RADIUS server. *server-name*—The name of the RADIUS server. The name is alphanumeric, case sensitive, and has a maximum of 256 characters. **Tip** If the specified RADIUS server is not found, configure it using the radius-server host command and retry this command.
Step 4	switch(config-radius)# **deadtime** *minutes*	(Optional) Configures the monitoring dead time. *minutes*—The dead time, in minutes. The range is from 1 to 1440. The default value is 0 minutes. **Note** If the dead-time interval for a RADIUS server group is greater than zero (0), that value takes precedence over the global dead-time value.
Step 5	switch(config-radius)# **use-vrf** *vrf-name*	(Optional) Specifies the virtual routing and forwarding (VFR) to use to contact the servers in the server group.
Step 6	switch(config-radius)# **source-interface** {*interface-type*} {*interface-number*}	(Optional) Specifies a source interface to be used to reach the RADIUS server. *interface-type*—The interface type. *interface-number*—The interface number. The interface types and interface numbers are defines as follows: • loopback—Virtual interface number from 0 to 1023 • mgmt—Management interface 0 • null—Null interface 0 • port-channel—Port channel number from 1 to 4096
Step 7	switch(config-radius)# **show radius-server groups** [*group-name*]	(Optional) Displays the RADIUS server group configuration. *group-name*—The name of the server group. The name is a case-sensitive, alphanumeric string with a maximum length of 127 characters.

	Command or Action	Purpose
Step 8	switch(config-radius)# **copy running-config startup-config**	(Optional) Copies the running configuration to the startup configuration

The following example configures a RADIUS server group:

```
switch# configure terminal
switch(config)# aaa group server radius RadServer
switch(config-radius)# server 10.10.1.1
switch(config-radius)# deadtime 30
switch(config-radius)# use-vrf vrf1
switch(config-radius)# source-interface mgmt0
switch(config-radius)# show radius-server group
total number of groups:2

following RADIUS server groups are configured:
        group Radserver:
                server: 10.10.1.1
                deadtime is 30
        group test:
                deadtime is 30
switch(config-radius)# copy running-config startup-config
```

Enabling RADIUS Server Directed Requests

You can allow users to designate the RADIUS server to send their authentication request to. This is called a directed request.

If you enable this option, a user can log in as username@vrfname:hostname, where *vrfname* is the virtual routing and forwarding (VRF) instance to use and *hostname* is the name of a configured RADIUS server.

Directed requests are disabled by default.

Note User-specified logins are supported only for Telnet sessions.

Before You Begin

Log in to the CLI in EXEC mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **radius-server directed-request**
3. switch(config)# **exit**
4. (Optional) switch(config)# **show radius-server directed-request**
5. (Optional) switch(config)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **radius-server directed-request**	Enables directed requests. The default is disabled.
Step 3	switch(config)# **exit**	Returns you to the EXEC mode.
Step 4	switch(config)# **show radius-server directed-request**	(Optional) Displays the directed request configuration.
Step 5	switch(config)# **copy running-config startup-config**	(Optional) Copies the running configuration to the startup configuration.

The following example enables RADIUS server directed requests:

```
switch# configure terminal
switch(config)# radius-server directed-request
switch(config)# exit
switch# show radius-server directed-request
switch# copy running-config startup-config
```

Setting the Global Timeout for All RADIUS Servers

You can configure the global timeout interval that specifies how long to wait for a response from a RADIUS server before declaring a timeout failure.

The timeout specified in Setting the Timeout Interval for a Single RADIUS Server , on page 13 overrides the global RADIUS timeout.

Before You Begin

Log in to the CLI in EXEC mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **radius-server timeout** *seconds*
3. switch(config-radius)# **exit**
4. (Optional) switch(config-radius)# **show radius-server**
5. (Optional) switch(config-radius)# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **radius-server timeout** *seconds*	Specifies the transmission timeout interval for RADIUS servers. *seconds*—The transmission timeout interval, in seconds. The range is from 1 to 60 seconds. The default value is 5 seconds.
Step 3	switch(config-radius)# **exit**	Returns you to the EXEC mode.
Step 4	switch(config-radius)# **show radius-server**	(Optional) Displays the RADIUS server configuration
Step 5	switch(config-radius)# **copy running-config startup-config**	(Optional) Copies the running configuration to the startup configuration

The following example sets the global timeout for all RADIUS servers:

```
switch# configure terminal
switch(config)# radius-server timeout 101
switch(config-radius)# exit
switch(config-radius)# show radius-server
switch(config-radius)# copy running-config startup-config
```

Configuring a Global Retry Count for All RADIUS Servers

You can configure the maximum number of times to retry transmitting to a RADIUS server before reverting to local authentication. This setting is applied to all RADIUS servers.

By default, retransmission to a RADIUS server is only tried once before reverting to local authentication. You can increase the number of retries up to a maximum of five. The retry count specified for a single RADIUS server in Configuring Retries for a Single RADIUS Server overrides this global setting.

Before You Begin

Log in to the CLI in EXEC mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **radius-server retransmit***count*
3. switch(config)# **exit**
4. (Optional) switch# **show radius-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **radius-server retransmit***count*	Defines the number of retransmits allowed before reverting to local authentication. This global setting applies to all RADIUS servers. *count*—The number of allowed retransmits. The range is from 0 to 5. The default value is 1.
Step 3	switch(config)# **exit**	Returns you to the EXEC mode.
Step 4	switch# **show radius-server**	(Optional) Displays the RADIUS server configuration.
Step 5	switch# **copy running-config startup-config**	(Optional) Copies the running configuration to the startup configuration.

The following example configures the global retry count for all RADIUS servers:

```
switch# configure terminal
switch(config)# radius-server retransmit 31
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Setting the Timeout Interval for a Single RADIUS Server

You can configure how long to wait for a response from a RADIUS server before declaring a timeout failure. The timeout specified for a single RADIUS server overrides the timeout defined in Setting the Global Timeout for All RADIUS Servers, on page 11.

Before You Begin

Log in to the CLI in EXEC mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **radius-server host** {*ipv4-address* | *host-name* } **timeout** *seconds*
3. switch(config)# **exit**
4. (Optional) switch# **show radius-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **radius-server host** {*ipv4-address* | *host-name* } **timeout** *seconds*	Specifies the timeout interval for the specified server. *ipv4-address*—The IP address for the RADIUS server. *hostname*—The hostname for the RADIUS server. The hostname is alphanumeric, case sensitive, and has a maximum of 256 characters. *seconds*—The timeout interval. The range is from 1 to 60 seconds. The default is 5 seconds. **Note** The timeout specified for a single RADIUS server overrides the global RADIUS timeout.
Step 3	switch(config)# **exit**	Returns you to the EXEC mode.
Step 4	switch# **show radius-server**	(Optional) Displays the RADIUS server configuration.
Step 5	switch# **copy running-config startup-config**	(Optional) Copies the running configuration to the startup configuration.

The following example sets the timeout interval for a single RADIUS server:

```
switch# configure terminal
switch(config)# radius-server host server1 timeout 10
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Configuring Retries for a Single RADIUS Server

You can configure the maximum number of times to retry transmitting to a RADIUS server before reverting to local authentication. This setting applies to a single RADIUS server and takes precedence over the global retry count.

Before You Begin

Log in to the CLI in EXEC mode.

Know the following:

- By default, retransmission to a RADIUS server is only tried once before reverting to local authentication.
- You can increase the number of retries up to a maximum of five.
- The retry count specified for a single RADIUS server overrides the global setting made for all RADIUS servers.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **radius-server host** {*ipv4-address* | *host-name*} **retransmit** *count*
3. switch(config)# **exit**
4. (Optional) switch# **show radius-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **radius-server host** {*ipv4-address* | *host-name*} **retransmit** *count*	Specifies the retransmission count for a specific server. *ipv4-address*—The IP address for the RADIUS server. *hostname*—The hostname for the RADIUS server. The hostname is alphanumeric, case sensitive, and has a maximum of 256 characters. *count*—The retransmission count. The default value is the global value. **Note** This retransmit count for a single RADIUS server overrides the global setting for all RADIUS servers.
Step 3	switch(config)# **exit**	Returns you to EXEC mode.
Step 4	switch# **show radius-server**	(Optional) Displays the RADIUS server configuration
Step 5	switch# **copy running-config startup-config**	(Optional) Copies the running configuration to the startup configuration.

The following example configures retries for a single RADIUS server:

```
switch# configure terminal
switch(config)# radius-server host server1 retransmit 3
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Configuring a RADIUS Accounting Server

You can configure a server to perform accounting functions.

By default, RADIUS servers are used for both accounting and authentication.

Before You Begin

- Log in to the CLI in EXEC mode.
- Know the destination UDP port number for RADIUS accounting messages.

SUMMARY STEPS

1. switch# **configure terminal**
2. (Optional) switch(config)# **radius-server host** {*ipv4-address* | *host-name*} **acct-port** *udp-port*
3. (Optional) switch(config)# **radius-server host** {*ipv4-address* | *host-name*} **accounting**
4. switch(config)# **exit**
5. (Optional) switch# **show radius-server**
6. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **radius-server host** {*ipv4-address* | *host-name*} **acct-port** *udp-port*	(Optional) Associates a specific host with the UDP port that receives RADIUS accounting messages. *ipv4-address*—The IP address for the RADIUS server. *hostname*—The hostname for the RADIUS server. The hostname is alphanumeric, case sensitive, and has a maximum of 256 characters. *upd-port*—The UPD port number. The range is from 0 to 65535. The default value is 1812.
Step 3	switch(config)# **radius-server host** {*ipv4-address* | *host-name*} **accounting**	(Optional) Designates the specific RADIUS host as an accounting server. The default is both accounting and authentication. *ipv4-address*—The IP address for the RADIUS server. *hostname*—The hostname for the RADIUS server. The hostname is alphanumeric, case sensitive, and has a maximum of 256 characters.
Step 4	switch(config)# **exit**	Returns you to the EXEC mode.
Step 5	switch# **show radius-server**	(Optional) Displays the RADIUS server configuration
Step 6	switch# **copy running-config startup-config**	(Optional) Copies the running configuration to the startup configuration.

The following example configures a RADIUS accounting server:

```
switch# configure terminal
switch(config)# radius-server host 10.10.1.1 acct-port 2004
switch(config)# radius-server host 10.10.1.1 accounting
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Configuring a RADIUS Authentication Server

You can configure a server to perform authentication functions.

By default, RADIUS servers are used for both accounting and authentication.

Before You Begin

- Log in to the CLI in EXEC mode.
- Know the destination UDP port number for RADIUS authentication messages.

SUMMARY STEPS

1. switch# **configure terminal**
2. (Optional) switch(config)# **radius-server host** {*ipv4-address* | *hostname*} **auth-port** *udp-port*
3. (Optional) switch(config)# **radius-server host** {*ipv4-address* | *host-name*} **authentication**
4. switch(config)# **exit**
5. (Optional) switch# **show radius-server**
6. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **radius-server host** {*ipv4-address* | *hostname*} **auth-port** *udp-port*	(Optional) Associates a specific host with the UDP port that receives RADIUS authentication messages. *ipv4-address*—The IP address for the RADIUS server. *hostname*—The hostname for the RADIUS server. The hostname is alphanumeric, case sensitive, and has a maximum of 256 characters. *upd-port*—The UPD port number. The range is from 0 to 65535. The default value is 1812.
Step 3	switch(config)# **radius-server host** {*ipv4-address* | *host-name*} **authentication**	(Optional) Designates the specific RADIUS host as an authentication server. The default is both accounting and authentication. *ipv4-address*—The IP address for the RADIUS server. *hostname*—The hostname for the RADIUS server. The hostname is alphanumeric, case sensitive, and has a maximum of 256 characters.
Step 4	switch(config)# **exit**	Returns you to the EXEC mode.
Step 5	switch# **show radius-server**	(Optional) Displays the RADIUS server configuration

	Command or Action	Purpose
Step 6	switch# **copy running-config startup-config**	(Optional) Copies the running configuration to the startup configuration.

The following example configures a RADIUS authentication server:

```
switch# configure terminal
switch(config)# radius-server host 10.10.2.2 auth-port 2005
switch(config)# radius-server host 10.10.2.2 authentication
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Configuring Periodic RADIUS Server Monitoring

You can configure the monitoring of RADIUS servers.

The test idle timer specifies the interval of time that elapses before a test packet is sent to a unresponsive RADIUS server

The default idle timer value is 0 minutes. When the idle time interval is 0 minutes, the Cisco NX-OS device does not perform periodic RADIUS server monitoring.

Note For security reasons, do not configure a username that is in the RADIUS database as a test username.

Before You Begin

Log in to the CLI in EXEC mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **radius-server host** {*ipv4-address* | *hostname*} **test** {**idle-time** *minutes* | **password** *password* [**idle-time** *minutes*] | **username** *name* [**password** *password* [**idle-time** *minutes*]]}
3. switch(config)# **radius-server dead-time** *minutes*
4. switch(config)# **exit**
5. (Optional) switch# **show radius-server**
6. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.

	Command or Action	Purpose
Step 2	switch(config)# **radius-server host** {*ipv4-address* | *hostname*} **test** {**idle-time** *minutes* | **password** *password* [**idle-time** *minutes*] | **username** *name* [**password** *password* [**idle-time** *minutes*]]}	Specifies parameters for server monitoring. *ipv4-address*—The IP address for the RADIUS server. *hostname*—The hostname for the RADIUS server. The hostname is alphanumeric, case sensitive, and has a maximum of 256 characters. *minutes*—The idle time, in minutes. The range is from 0 to 1440 minutes. The default value is 0 minutes. *name*—The username to use when connecting to the RADIUS server. The default value is test. *password*—The user's password. The default value is test. **Note** For periodic RADIUS server monitoring, you must set the idle timer to a value greater than 0.
Step 3	switch(config)# **radius-server dead-time** *minutes*	Specifies the number of minutes to wait before sending a test packet to a RADIUS server that was declared dead. *minutes*—The amount of time to wait, in minutes. The range is from 0 to 1440 minutes. The default value is 0 minutes.
Step 4	switch(config)# **exit**	Returns you to the EXEC mode.
Step 5	switch# **show radius-server**	(Optional) Displays the RADIUS server configuration
Step 6	switch# **copy running-config startup-config**	(Optional) Copies the running configuration to the startup configuration.

The following example configures periodic RADIUS server monitoring:

```
switch# configure terminal
switch(config)# radius-server host 10.10.1.1 test username user1 password Ur2Gd2BH idle-time
 3
switch(config)# radius-server dead-time 5
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Configuring the Global Dead-Time Interval

You can configure the dead-time interval for all RADIUS servers. The dead-time interval specifies the time to wait after declaring a RADIUS server dead, before sending out a test packet to determine if the server is now alive.

Note When the dead-time interval is 0 minutes, RADIUS servers are not marked as dead even if they are not responding. You can configure the dead-time interval for a RADIUS server group.

Before You Begin

Log in to the CLI in EXEC mode.

SUMMARY STEPS

1. switch# **configure terminal**
2. switch(config)# **radius-server deadtime** *minutes*
3. switch(config)# **exit**
4. (Optional) switch# **show radius-server**
5. (Optional) switch# **copy running-config startup-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **configure terminal**	Enters global configuration mode.
Step 2	switch(config)# **radius-server deadtime** *minutes*	Configures the dead-time interval. *minutes*—The dead-time interval, in minutes. The range is from 0 to 1440 minutes. The default value is 0 minutes.
Step 3	switch(config)# **exit**	Returns you to the EXEC mode.
Step 4	switch# **show radius-server**	(Optional) Displays the RADIUS server configuration.
Step 5	switch# **copy running-config startup-config**	(Optional) Copies the running configuration to the startup configuration.

The following example configures the global dead-time interval:

```
switch# configure terminal
switch(config)# radius-server deadtime 5
switch(config)# exit
switch# show radius-server
switch# copy running-config startup-config
```

Manually Monitoring RADIUS Servers or Groups

You can manually send a test message to a RADIUS server or to a server group.

Before You Begin

Log in to the CLI in EXEC mode.

SUMMARY STEPS

1. switch# **test aaa server radius** {*ipv4-address* | *hostname*} [**vrf** *vrf-name*] *username password*
2. switch# **test aaa group** *group-name username password*

DETAILED STEPS

	Command or Action	Purpose
Step 1	switch# **test aaa server radius** {*ipv4-address* | *hostname*} [**vrf** *vrf-name*] *username password*	Sends a test message to a RADIUS server to confirm availability. *ipv4-address*—The IP address of the RADIUS server. *hostname*—The hostname of the RADIUS server. The hostname is alphanumeric, case sensitive, and has a maximum of 256 characters. *vrf-name*—The Virtual Routing and Forwarding (VRF) name.
Step 2	switch# **test aaa group** *group-name username password*	Sends a test message to a RADIUS server group to confirm availability. *group-name*—The name of the RADIUS server group. *username*—The username to use when connecting to the RADIUS server group. *password*—The user's password.

The following example manually monitors a RADIUS server and a RADIUS server group:

```
switch# test aaa server radius 10.10.1.1 user1 Ur2Gd2BH
switch# test aaa group RadGroup user2 As3He3CI
```

Verifying the RADIUS Configuration

Use one of the following commands to verify the configuration:

Command	Purpose
show running-config radius [**all**]	Displays the RADIUS configuration in the running configuration.
show startup-config radius	Displays the RADIUS configuration in the startup configuration.
show radius-server [*hostname* | *ipv4-address*] [**directed-request** | **groups** | **sorted** | **statistics**]	Displays all configured RADIUS server parameters. *hostname*—The hostname for the RADIUS server. The hostname is alphanumeric, case sensitive, and has a maximum of 256 characters. *ipv4-address*—The IP address for the RADIUS server.

Displaying RADIUS Server Statistics

Use the following command to display statistics for RADIUS server activity:

show radius-server statistics { *hostname* | *ipv4-address* }

hostname—The hostname for the RADIUS server. The hostname is alphanumeric, case sensitive, and has a maximum of 256 characters.

ipv4-address—The IP address for the RADIUS server.

Configuration Example for RADIUS

The following example shows how to configure a global RADIUS key and a RADIUS server host key:

```
switch# configure terminal
switch(config)# radius-server key 7 "ToIkLhPpG"
switch(config)# radius-server host 10.10.1.1 key 7 "ShMoMhTl" authentication accounting
switch(config)# aaa group server radius RadServer
    server 10.10.1.1
```

Feature History for RADIUS

Feature Name	Releases	Feature Information
RADIUS	Release 5.2(1)SK1(2.1)	This feature was introduced.