CHAPTER 5

Configuring RADIUS

This chapter describes how to configure RADIUS protocol on Cisco NX-OS devices.

This chapter includes 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.

This section includes the following topics:

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 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 to the and authenticate to an NX-OS device using RADIUS, the following happens:

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. See Figure 5-1.

Figure 5-1 RADIUS Server States

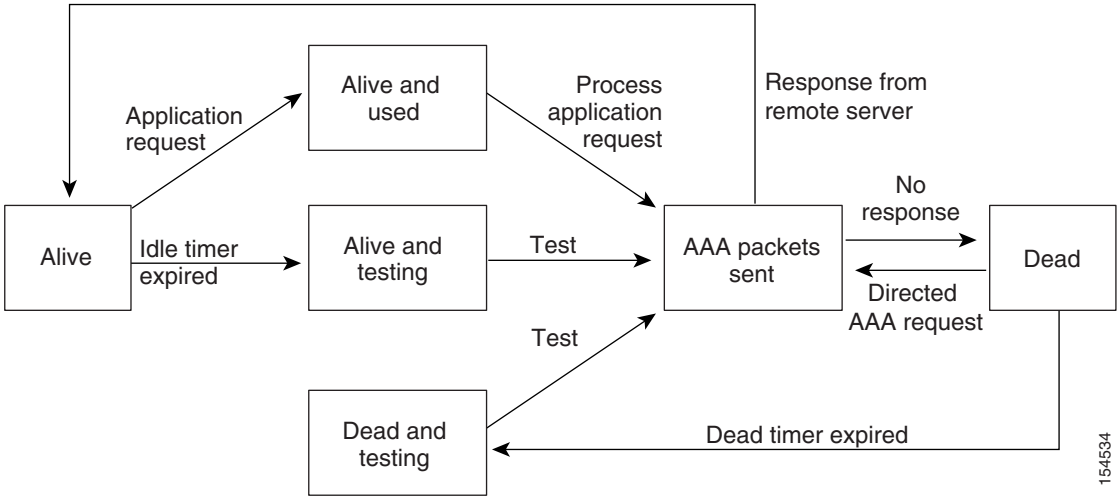

Note The monitoring interval for alive servers and dead servers are different and can be configured by the user. The RADIUS server monitoring is performed by sending a test authentication request to the RADIUS server.

Vendor-Specific Attributes

The Internet Engineering Task Force (IETF) draft standard specifies a method for communicating 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, 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, along with authentication results. This authorization information is specified through VSAs.

The following are supported VSA protocol options:

- 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 are supported attributes:

- 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 only be used with the shell protocol value. The following examples show the roles attribute as supported by Cisco 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\""
  ```

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

RADIUS has the following prerequisites:

- 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

RADIUS has the following guidelines and limitations:

- You can configure a maximum of 64 RADIUS servers.

Configuring RADIUS Servers

This section includes the following topics:

Note Be aware that the Cisco NX-OS commands for this feature may differ from those used in Cisco IOS.

Configuring RADIUS Server Hosts

Use this procedure to configure the IP address or the hostname for each RADIUS server to be used for authentication.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.
- You can configure up to 64 RADIUS servers.
- All RADIUS server hosts are automatically added to the default RADIUS server group.

SUMMARY STEPS

1. **config t**
2. **radius-server host** {*ipv4-address* | *host-name*}
3. **exit**
4. **show radius-server**
5. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	**config t** **Example:** n1000v# config t n1000v(config)#	Places you into the CLI Global Configuration mode.
Step 2	**radius-server host** {*ipv4-address* | *host-name*} **Example:** n1000v(config)# radius-server host 10.10.1.1	Defines the IP address or hostname for the RADIUS server.
Step 3	**exit** **Example:** n1000v(config)# exit n1000v#	Returns you to the CLI EXEC mode.
Step 4	**show radius-server** **Example:** n1000v# show radius-server	(Optional) Displays the RADIUS server configuration.
Step 5	**copy running-config startup-config** **Example:** n1000v# copy running-config startup-config	(Optional) Saves this change in the running configuration to the startup configuration.

Deleting a RADIUS Server Host

You can delete a RADIUS server host from a server group.

Configuring the Global RADIUS Key

Use this procedure to configure the key that is used by all RADIUS servers to authenticate with the Cisco Nexus 1000V.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.
- You know the global key that is used for RADIUS server authentication.

SUMMARY STEPS

1. **config t**
2. **radius-server key** [**0** | **7**] *key-value*
3. **exit**

4. **show radius-server**
5. **copy running-config startup-config**

DETAILED STEPS

To configure a global preshared key, follow these steps:

	Command	Purpose

Step 1	**config t**
	
	Example:
	`n1000v# config t`
	`n1000v(config)#`
		Places you into the CLI Global Configuration mode.

Step 2	**radius-server key** [**0** | **7**] *key-value* **Example:** `n1000v(config)# radius-server key 0 QsEfThUkO`	Specifies a preshared key for all RADIUS servers. You can specify a clear text (**0**) or encrypted (**7**) preshared key. The default format is clear text. The maximum length is 63 characters. By default, no preshared key is configured.

Step 3	**exit**
	
	Example:
	`n1000v(config)# exit`
	`n1000v#`
		Returns you to the CLI EXEC mode.

Step 4	**show radius-server** **Example:** `n1000v# 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	**copy running-config startup-config** **Example:** `n1000v# copy running-config startup-config`	(Optional) Saves this change in the running configuration to the startup configuration.

Configuring a RADIUS Server Key

Use this procedure to configure a key for a single RADIUS server host.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.
- You have the key to be used for the remote RADIUS host.

SUMMARY STEPS

1. **config t**
2. **radius-server host** {*ipv4-address* | *host-name*} **key** *key-value*
3. **exit**

4. **show radius-server**
5. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose

Step 1	**config t**
	
	Example:
	`n1000v# config t`
	`n1000v(config)#`
		Places you into the CLI Global Configuration mode.

Step 2	**radius-server host** {*ipv4-address* | *host-name*} **key** [**0** | **7**] *key-value* **Example:** `n1000v(config)# radius-server host 10.10.1.1 key 0 PlIjUhYg`	Specifies a preshared key for a specific RADIUS server. You can specify a clear text (**0**) or encrypted (**7**) preshared key. The default format is clear text. The maximum length is 63 characters. This preshared key is used instead of the global preshared key.

Step 3	**exit**
	
	Example:
	`n1000v(config)# exit`
	`n1000v#`
		Returns you to the CLI EXEC mode.

Step 4	**show radius-server** **Example:** `n1000v# 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	**copy running-config startup-config** **Example:** `n1000v# copy running-config startup-config`	(Optional) Saves this change in the running configuration to the startup configuration.

Configuring RADIUS Server Groups

Use this procedure to configure a RADIUS server group whose member servers share authentication functions.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.
- All servers in a RADIUS server group must belong to the RADIUS protocol.
- The servers in the group are tried in the same order in which you configure them.

SUMMARY STEPS

1. **config t**
2. **aaa group server radius** *group-name*
3. **server** {*ipv4-address* | *server-name*}
4. **deadtime** *minutes*
5. **use-vrf** *vrf-name*
6. **exit**
7. **show radius-server groups** [*group-name*]
8. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose

Step 1	**config t**
	
	Example:
	`n1000v# config t`
	`n1000v(config)#`
		Places you into the CLI Global Configuration mode.

Step 2	**aaa group server radius** *group-name*
	
	Example:
	`n1000v(config)# aaa group server radius RadServer`
	`n1000v(config-radius)#`
		Creates a RADIUS server group and enters the RADIUS server group configuration submode for that group. The *group-name* argument is a case-sensitive alphanumeric string with a maximum length of 127 characters.

Step 3	**server** {*ipv4-address* | *server-name*} **Example:** `n1000v(config-radius)# server 10.10.1.1`	Configures the RADIUS server as a member of the RADIUS server group. **Tip** If the specified RADIUS server is not found, configure it using the **radius-server host** command and retry this command.
Step 4	**deadtime** *minutes* **Example:** `n1000v(config-radius)# deadtime 30`	(Optional) Configures the monitoring dead time. The default is 0 minutes. The range is from 1 through 1440. **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 (see the "Configuring the Global Dead-Time Interval" section on page 5-19).
Step 5	**use-vrf** *vrf-name* **Example:** `n1000v(config-radius)# use-vrf vrf1`	(Optional) Specifies the VRF to use to contact the servers in the server group.

Step 6	**exit**
	
	Example:
	`n1000v(config-radius)# exit`
	`n1000v(config)#`
		Returns you to the CLI EXEC mode.

	Command	Purpose
Step 7	**show radius-server groups** [*group-name*] **Example:** n1000v(config)# show radius-server group	(Optional) Displays the RADIUS server group configuration.
Step 8	**copy running-config startup-config** **Example:** n1000v(config)# copy running-config startup-config	(Optional) Saves this change in the running configuration to the startup configuration.

Enabling RADIUS Server Directed Requests

Use this procedure to let users 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 VRF to use and *hostname* is the name of a configured RADIUS server.

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

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.
- Directed requests are disabled by default.

SUMMARY STEPS

1. **config t**
2. **radius-server directed-request**
3. **exit**
4. **show radius-server directed-request**
5. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	**config t** **Example:** n1000v# config t n1000v(config)#	Places you into the CLI Global Configuration mode.
Step 2	n1000v(config)# radius-server directed-request **Example:** n1000v(config)# radius-server directed-request	Enables directed requests. The default is disabled.

	Command	Purpose

Step 3	**exit**
	
	Example:
	`n1000v(config)# exit`
	`n1000v#`
		Returns you to the CLI EXEC mode.

Step 4	**show radius-server directed-request** **Example:** `n1000v# show radius-server directed-request`	(Optional) Displays the directed request configuration.
Step 5	**copy running-config startup-config** **Example:** `n1000v# copy running-config startup-config`	(Optional) Copies the running configuration to the startup configuration.

Setting the Global Timeout for All RADIUS Servers

Use this procedure to configure the global timeout interval specifying how long to wait for a response from a RADIUS server before declaring a timeout failure.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.
- The timeout specified in the "Setting the Timeout Interval for a Single RADIUS Server" procedure on page 5-13 overrides the global RADIUS timeout.

SUMMARY STEPS

1. **config t**
2. **radius-server timeout** *seconds*
3. **exit**
4. **show radius-server**
5. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose

Step 1	**config t**
	
	Example:
	`n1000v# config t`
	`n1000v(config)#`
		Places you into the CLI Global Configuration mode.

Step 2	**radius-server timeout** *seconds* **Example:** `n1000v(config)# radius-server timeout 10`	Specifies the transmission timeout interval for RADIUS servers. The default timeout interval is 5 seconds and the allowable range is from 1 to 60 seconds.

	Command	Purpose

Step 3	**exit**
	
	Example:
	`n1000v(config)# exit`
	`n1000v#`
		Returns you to the CLI EXEC mode.

Step 4	**show radius-server** **Example:** `n1000v# show radius-server`	(Optional) Displays the RADIUS server configuration.
Step 5	**copy running-config startup-config** **Example:** `n1000v# copy running-config startup-config`	(Optional) Saves this change in the running configuration to the startup configuration.

Configuring a Global Retry Count for All RADIUS Servers

Use this procedure to 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.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.
- 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 the "Configuring Retries for a Single RADIUS Server" procedure on page 5-14, overrides this global setting.

SUMMARY STEPS

1. **config t**
2. **radius-server retransmission** *count*
3. **radius-server timeout** *seconds*
4. **exit**
5. **show radius-server**
6. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	**config t** **Example:** n1000v# config t n1000v(config)#	Places you into the CLI Global Configuration mode.
Step 2	**radius-server retransmit** *count* **Example:** n1000v(config)# radius-server retransmit 3	Defines the number of retransmits allowed before reverting to local authentication. This is a global setting that applies to all RADIUS servers. The default number of retransmits is 1 and the range is from 0 to 5.
Step 3	**exit** **Example:** n1000v(config)# exit n1000v#	Returns you to the CLI EXEC mode.
Step 4	**show radius-server** **Example:** n1000v# show radius-server	(Optional) Displays the RADIUS server configuration.
Step 5	**copy running-config startup-config** **Example:** n1000v# copy running-config startup-config	(Optional) Saves this change in the running configuration to the startup configuration.

Setting the Timeout Interval for a Single RADIUS Server

Use this procedure to configure how long to wait for a response from a RADIUS server before declaring a timeout failure.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.
- The timeout specified for a single RADIUS server overrides the timeout defined in the "Setting the Global Timeout for All RADIUS Servers" procedure on page 5-11.

SUMMARY STEPS

1. **config t**
2. **radius-server host** {*ipv4-address* | *host-name*} **timeout** *seconds*
3. **exit**
4. **show radius-server**
5. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	**config t** **Example:** n1000v# config t n1000v(config)#	Places you into the CLI Global Configuration mode.
Step 2	**radius-server host** {*ipv4-address* | *host-name*} **timeout** *seconds* **Example:** n1000v(config)# radius-server host server1 timeout 10	Specifies the timeout interval for the specified server. The default timeout interval is 5 seconds and the allowable range is from 1 to 60 seconds. **Note** The timeout specified for a single RADIUS server overrides the global RADIUS timeout.
Step 3	**exit** **Example:** n1000v(config)# exit n1000v#	Returns you to the CLI EXEC mode.
Step 4	**show radius-server** **Example:** n1000v# show radius-server	(Optional) Displays the RADIUS server configuration.
Step 5	**copy running-config startup-config** **Example:** n1000v# copy running-config startup-config	(Optional) Saves this change in the running configuration to the startup configuration.

Configuring Retries for a Single RADIUS Server

Use this procedure to configure the maximum number of times to retry transmitting to a RADIUS server before reverting to local authentication.This setting is applied to a single RADIUS server and takes precedence over the global retry count.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.
- 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. **config t**
2. **radius-server host** {*ipv4-address* | *host-name*} **retransmit** *count*
3. **exit**

4. **show radius-server**
5. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose

Step 1	**config t**
	
	Example:
	`n1000v# config t`
	`n1000v(config)#`
		Places you into the CLI Global Configuration mode.

Step 2	**radius-server host** {*ipv4-address* | *host-name*} **retransmit** *count* **Example:** `n1000v(config)# radius-server host server1 retransmit 3`	Specifies the retransmission count for a specific server. The default is the global value. **Note** This retransmit count for a single RADIUS server overrides the global setting for all RADIUS servers.

Step 3	**exit**
	
	Example:
	`n1000v(config)# exit`
	`n1000v#`
		Returns you to the CLI EXEC mode.

Step 4	**show radius-server** **Example:** `n1000v# show radius-server`	(Optional) Displays the RADIUS server configuration.
Step 5	**copy running-config startup-config** **Example:** `n1000v# copy running-config startup-config`	(Optional) Saves this change in the running configuration to the startup configuration.

Configuring a RADIUS Accounting Server

Use this procedure to configure a server to perform accounting functions.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.
- By default, RADIUS servers are used for both accounting and authentication.
- You know the destination UDP port number for RADIUS accounting messages.

SUMMARY STEPS

1. **config t**
2. **radius-server host** {*ipv4-address* | *host-name*} **acct-port** *udp-port*
3. **radius-server host** {*ipv4-address* | *host-name*} **accounting**
4. **exit**

5. **show radius-server**
6. **copy running-config startup-config**

DETAILED STEPS

To configure the authentication and accounting attributes for RADIUS servers, follow these steps:

	Command	Purpose

Step 1	**config t**
	
	Example:
	`n1000v# config t`
	`n1000v(config)#`
		Places you into the CLI Global Configuration mode.

Step 2	**radius-server host** {*ipv4-address* | *host-name*} **acct-port** *udp-port* **Example:** `n1000v(config)# radius-server host 10.10.1.1 acct-port 2004`	(Optional) Associates a specific host with the UDP port that receives RADIUS accounting messages. The default UDP port is 1812. The range is from 0 to 65535.
Step 3	**radius-server host** {*ipv4-address* | *host-name*} **accounting** **Example:** `n1000v(config)# radius-server host 10.10.1.1 accounting`	(Optional) Designates the specific RADIUS host as an accounting server. The default is both accounting and authentication.

Step 4	**exit**
	
	Example:
	`n1000v(config)# exit`
	`n1000v#`
		Returns you to the CLI EXEC mode.

Step 5	**show radius-server** **Example:** `n1000v(config)# show radius-server`	(Optional) Displays the RADIUS server configuration.
Step 6	**copy running-config startup-config** **Example:** `n1000v# copy running-config startup-config`	(Optional) Saves this change in the running configuration to the startup configuration.

Configuring a RADIUS Authentication Server

Use this procedure to configure a server to perform authentication functions.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.
- By default, RADIUS servers are used for both accounting and authentication.
- You know the destination UDP port number for RADIUS authentication messages.

SUMMARY STEPS

1. **config t**
2. **radius-server host** {*ipv4-address* | *host-name*} **auth-port** *udp-port*
3. **radius-server host** {*ipv4-address* | *host-name*} **authentication**
4. **exit**
5. **show radius-server**
6. **copy running-config startup-config**

DETAILED STEPS

To configure the authentication and accounting attributes for RADIUS servers, follow these steps:

	Command	Purpose

Step 1	**config t**
	
	Example:
	`n1000v# config t`
	`n1000v(config)#`
		Places you into the CLI Global Configuration mode.

Step 2	**radius-server host** {*ipv4-address* | *host-name*} **auth-port** *udp-port* **Example:** `n1000v(config)# radius-server host 10.10.2.2 auth-port 2005`	(Optional) Associates a specific host with the UDP port that receives RADIUS authentication messages. The default UDP port is 1812. The range is from 0 to 65535.
Step 3	**radius-server host** {*ipv4-address* | *host-name*} **authentication** **Example:** `n1000v(config)# radius-server host 10.10.2.2 authentication`	(Optional) Designates the specific RADIUS host as an authentication server. The default is both accounting and authentication.

Step 4	**exit**
	
	Example:
	`n1000v(config)# exit`
	`n1000v#`
		Returns you to the CLI EXEC mode.

Step 5	**show radius-server** **Example:** `n1000v(config)# show radius-server`	(Optional) Displays the RADIUS server configuration.
Step 6	**copy running-config startup-config** **Example:** `n1000v# copy running-config startup-config`	(Optional) Saves this change in the running configuration to the startup configuration.

Configuring Periodic RADIUS Server Monitoring

Use this procedure to configure the monitoring of RADIUS servers.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.
- The test idle timer specifies the interval of time that elapses before a test packet is sent to a nonresponsive RADIUS server.

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

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

SUMMARY STEPS

1. **config t**
2. **radius-server host** {*ipv4-address* | *host-name*} **test** {**idle-time** *minutes* | **password** *password* [**idle-time** *minutes*] | **username** *name* [**password** *password* [**idle-time** *minutes*]]}
3. **radius-server dead-time** *minutes*
4. **exit**
5. **show radius-server**
6. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	**config t** **Example:** n1000v# config t n1000v(config)#	Places you into the CLI Global Configuration mode.
Step 2	**radius-server host** {*ipv4-address* | *host-name*} **test** {**idle-time** *minutes* | **password** *password* [**idle-time** *minutes*] | **username** *name* [**password** *password* [**idle-time** *minutes*]]} **Example:** n1000v(config)# radius-server host 10.10.1.1 test username user1 password Ur2Gd2BH idle-time 3	Specifies parameters for server monitoring. The default username is test and the default password is test. The default value for the idle timer is 0 minutes. The valid range is 0 to 1440 minutes. **Note** For periodic RADIUS server monitoring, you must set the idle timer to a value greater than 0.

	Command	Purpose
Step 3	**radius-server dead-time** *minutes* **Example:** `n1000v(config)# radius-server dead-time 5`	Specifies the number of minutes to wait before sending a test packet to a RADIUS server that was declared dead. The default value is 0 minutes. The valid range is 1 to 1440 minutes.

Step 4	**exit**
	
	Example:
	`n1000v(config)# exit`
	`n1000v#`
		Returns you to the CLI EXEC mode.

Step 5	**show radius-server** **Example:** `n1000v# show radius-server`	(Optional) Displays the RADIUS server configuration.
Step 6	**copy running-config startup-config** **Example:** `n1000v# copy running-config startup-config`	(Optional) Saves this change in the running configuration to the startup configuration.

Configuring the Global Dead-Time Interval

Use this procedure to 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. The default value is 0 minutes.

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 (see the "Configuring RADIUS Server Groups" section on page 5-8).

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.

SUMMARY STEPS

1. **config t**
2. **radius-server deadtime** *minutes*
3. **exit**
4. **show radius-server**
5. **copy running-config startup-config**

DETAILED STEPS

To configure the RADIUS dead-time interval, follow these steps:

	Command	Purpose

Step 1	**config t**
	
	Example:
	`n1000v# config t`
	`n1000v(config)#`
		Places you into the CLI Global Configuration mode.

Step 2	`n1000v(config)#` **radius-server deadtime** *minutes* **Example:** `n1000v(config)# radius-server deadtime 5`	Configures the dead-time interval. The default value is 0 minutes. The range is from 1 to 1440 minutes.

Step 3	**exit**
	
	Example:
	`n1000v(config)# exit`
	`n1000v#`
		Returns you to the CLI EXEC mode.

Step 4	**show radius-server** **Example:** `n1000v# show radius-server`	(Optional) Displays the RADIUS server configuration.
Step 5	**copy running-config startup-config** **Example:** `n1000v# copy running-config startup-config`	(Optional) Saves this change in the running configuration to the startup configuration.

Manually Monitoring RADIUS Servers or Groups

Use this procedure to manually send a test message to a RADIUS server or to a server group.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.

SUMMARY STEPS

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

DETAILED STEPS

	Command	Purpose
Step 1	**test aaa server radius** {*ipv4-address* | *server-name*} [**vrf** *vrf-name*] *username password* **Example:** `n1000v# test aaa server radius 10.10.1.1 user1 Ur2Gd2BH`	Sends a test message to a RADIUS server to confirm availability.
Step 1	**test aaa group** *group-name username password* **Example:** `n1000v# test aaa group RadGroup user2 As3He3CI`	Sends a test message to a RADIUS server group to confirm availability.

Deleting a RADIUS Server Host

Use this procedure to delete a RADIUS server host from a RADIUS server group.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.
- You have added one or more RADIUS server hosts using the "Configuring RADIUS Server Hosts" section on page 5-5.

Verifying RADIUS Configuration

Use the following commands to display RADIUS configuration information:

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 [*server-name* | *ipv4-address*] [**directed-request** | **groups** | **sorted** | **statistics**]	Displays all configured RADIUS server parameters.

For detailed information about the fields in the output from this command, see the *Cisco Nexus 1000V Command Reference, Release 4.0(4)SV1(2)*.

Displaying RADIUS Server Statistics

Use this procedure to display the statistics that the NX-OS device maintains for RADIUS server activity.

BEFORE YOU BEGIN

Before beginning this procedure, you must know or do the following:

- You are logged in to the CLI in EXEC mode.

SUMMARY STEPS

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

DETAILED STEPS

	Command	Purpose

Step 1	`n1000v# show radius-server statistics {hostname | ipv4-address}`
	
	Example:
	`n1000v# show radius-server statistics 10.10.1.1`
		Displays the RADIUS statistics.

For detailed information about the fields in the output from this command, see the *Cisco Nexus 1000V Command Reference, Release 4.0(4)SV1(2).*

Example RADIUS Configuration

The following example shows how to configure RADIUS:

```
radius-server key 7 "ToIkLhPpG"
radius-server host 10.10.1.1 key 7 "ShMoMhTl" authentication accounting
aaa group server radius RadServer
    server 10.10.1.1
```

Default Settings

Table 5-1 lists the RADIUS default settings.

Table 5-1 Default RADIUS Parameters

Parameters	Default
Server roles	Authentication and accounting
Dead timer interval	0 minutes
Retransmission count	1
Retransmission timer interval	5 seconds

Table 5-1 Default RADIUS Parameters (continued)

Parameters	Default
Idle timer interval	0 minutes
Periodic server monitoring username	test
Periodic server monitoring password	test

Additional References

For additional information related to implementing RADIUS, see the following sections:

- Related Documents, page 5-23
- Standards, page 5-23

Related Documents

Related Topic	Document Title
Command reference	*Cisco Nexus 1000V Command Reference, Release 4.0(4)SV1(2)*

Standards

Standards	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

Feature History for RADIUS

This section provides the RADIUS release history.

Feature Name	Releases	Feature Information
RADIUS	4.0	This feature was introduced.