



RADIUS Server State Behavior

This appendix provides an explanation of RADIUS server states and the commands that affect them. It also provides a list of triggers that change servers in a "Down" state to "Active".

- [Understanding RADIUS Server States and Commands, on page 1](#)

Understanding RADIUS Server States and Commands

Server States

The system defines three server states for connected RADIUS servers:

- **Active:** The server is believed to be operational.
- **Not Responding:** The server has failed to respond to a message from the system a configured number of times (retries).
- **Down:** The system is no longer sending requests to the server.

RADIUS Server Commands

RADIUS server states are controlled by parameters set in the RADIUS Server Group Configuration Mode. The commands are:

- **detect-dead-server:** Configures how the system determines that a RADIUS server is not functioning. One or both of the following parameters should be set:
 - **consecutive-failures:** Configures the consecutive number of times the RADIUS server is unreachable by any single aaamgr on the system based on the **max-retries** command. If this command is enabled, each time the maximum number of retries is exceeded, this counter increments by one for the particular aaamgr and server. When any aaamgr exceeds this counter for a specific RADIUS server, the server's state is changed to "Down" and the deadtime timer is started. The default is enabled and 4.
 - **response-timeout:** Configures a specific delay, in seconds, in receiving a response from the RADIUS server before the server's state is changed to "Down" and the deadtime timer is started. The default is disabled.



Note If **response-timeout** is configured and **consecutive-failures** is not, the system will only wait for the specified period of time before changing the server's state to "Down", ignoring other settings such as **radius timeout**, and **max-retries**.

If **response-timeout** is configured and **consecutive-failures** is not, **consecutive-failures** is removed entirely from the system, including default configuration. If both parameters are configured, then both conditions must be met to change a RADIUS server's state to "Down".

- **deadtime**: Configure the maximum amount of time, in minutes, that must elapse after a context has exceeded one or both of the **detect-dead-server** parameters, depending on which parameter is configured. Once this timer has elapsed, the system reclassifies the RADIUS server as "Active" and subsequent requests to it can be made. If **radius deadtime** is not explicitly configured, the default value of 10 minutes is used.



Note Configuring deadtime as 0 disables the feature and the server is never marked as DOWN.

- **max-retries**: Configures maximum number of times the system attempts to retry communication with a RADIUS server. Once exceeded, the system changes the state of the server to "Not Responding", increments the **detect-dead-server consecutive-failures** counter (if configured), and attempts to communicate with another RADIUS server. The default value for this parameter is 5.
- **max-transmissions**: Configures the maximum number of times the system transmits authentication requests across all configured/enabled servers before it fails the authentication due to lack of response. The absolute maximum number of transmissions is equal to $NS * (N + 1)$, where NS is the number of configured authentication servers, and N is the setting for **radius max-retries**. The default for this command is disabled.
- **timeout**: Specifies how many seconds the system waits for a response from a RADIUS server before re-transmitting the request.

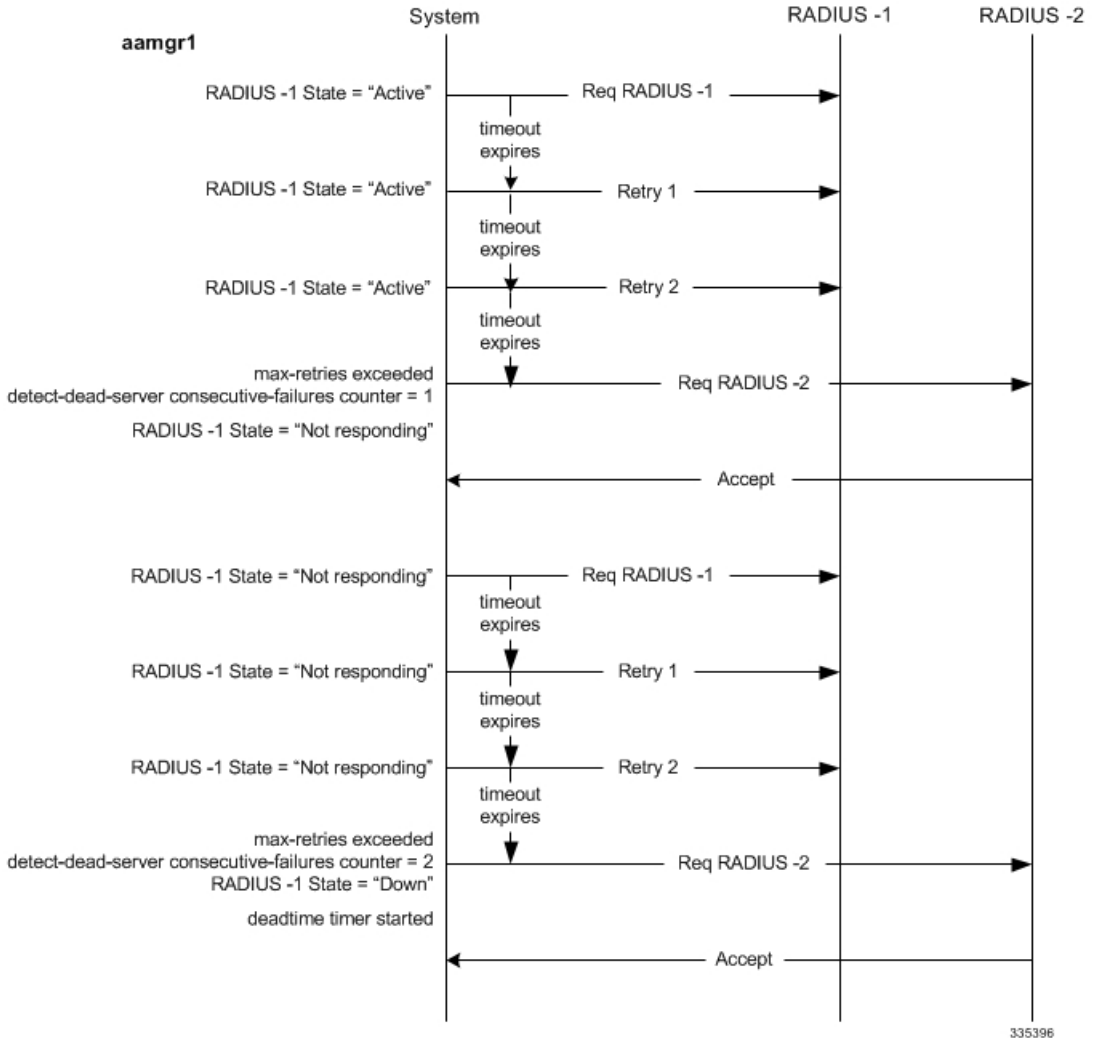
More information regarding each command can be found in the *Command Line Interface Reference*.

The following figure shows a simple flow of events and how the system reacts based on configured parameters.

Figure 1: Sample RADIUS Communication Flow

Configuration

radius timeout = 3 seconds
 radius max-retries = 2
 radius detect-dead-server consecutive failures = 2
 deadtime = 10 minutes



Server State Triggers

A number of triggers, events, and conditions can occur that change the state of a RADIUS server from "Down" to "Active" as defined by the system. They are:

- When the timer, based on the RADIUS Server Group Configuration Mode command: **deadtime** has expired, the server's state on the system is returned to "Active".



Note This parameter should be set to allow enough time to solve the issue that originally caused the server's state to be changed to "Down". After the deadtime timer expires, the system returns the server's state to "Active" regardless of whether or not the issue has been fixed.

- When a RADIUS authentication server is configured, the server state is initialized as "Active".
- When a RADIUS accounting server is configured and after receiving response for Acct-On message, the server state is made "Active".
- When a RADIUS accounting server is configured and after the Acct-On message exceeds the max retries setting and times-out, the server state is made "Active".
- When a RADIUS accounting server is configured with Acct-On disabled, the server state is made "Active".
- When a response from a RADIUS server is received, the server state is made "Active".



Note These triggers, events and conditions are applicable for each individual AAAMgr instance and the state change will be propagated throughout the system. The state of the server could be set to "Down" even if a single AAAMgr instance is affected and satisfies the **detect-dead-server** parameter criteria. However, even if any one of the non-affected AAAMgr instances receives a response from the RADIUS server, the state of the server is changed back to "Active", so that the affected AAAMgr does not impact all the other working ones.

- When a RADIUS server responds to the Exec Mode command **radius test**, the server state is made "Active".
- When a RADIUS probe is enabled and the probe response is received, the server state is made "Active".
- When a RADIUS probe request times-out after max retries, the server state is made "Active".
- If only one RADIUS authentication server is "Active" and goes down, all RADIUS authentication servers are made "Active".
- If only one RADIUS accounting server is "Active" and goes down, all RADIUS accounting servers are made "Active".
- In releases prior to 18.0, whenever a chassis boots up or when a new RADIUS accounting server or RADIUS mediation-device accounting server is configured with Acct-On configuration enabled, the state of the RADIUS server in all the AAA manager instances is initialized to "Waiting-for-response-to-Accounting-On". The Acct-On transmission and retries are processed by the Admin-AAAMgr.

When the Acct-On transaction is complete (i.e., when a response for Acct-On message is received or when Acct-On message is retried and timed-out), Admin-AAAMgr changes the state of the RADIUS accounting server to Active in all the AAA manager instances. During the period when the state of the server is in "Waiting-for-response-to-Accounting-On", any new RADIUS accounting messages which are generated as part of a new call will not be transmitted towards the RADIUS accounting server but it

will be queued. Only when the state changes to Active, these queued up messages will be transmitted to the server.

During ICSR, if the interface of the radius nas-ip address is srp-activated, then in the standby chassis, the sockets for the nas-ip will not be created. The current behavior is that if the interface is srp-activated Accounting-On transaction will not happen at ICSR standby node and the state of the RADIUS server in all the AAAmgr instances will be shown as "Waiting-for-response-to-Accounting-On" till the standby node becomes Active.

In 18.0 and later releases, whenever the chassis boots up or when a new RADIUS accounting server or RADIUS mediation-device accounting server is configured with Acct-On configuration enabled, the state of the RADIUS server will be set to Active for all the non-Admin-AAAmgr instances and will be set to "Waiting-for-response-to-Accounting-On" for only Admin-AAAmgr instance. The Accounting-On transaction logic still holds good from Admin-AAAmgr perspective. However, when any new RADIUS accounting messages are generated even before the state changes to Active in Admin-AAAmgr, these newly generated RADIUS accounting messages will not be queued at the server level and will be transmitted to the RADIUS server immediately.

During ICSR, even if the interface of radius nas-ip address is srp-activated, the state of the RADIUS accounting server will be set to Active in all non-Admin-AAAmgr instances and will be set to "Waiting-for-response-to-Accounting-On" in Admin-AAAmgr instance.



Note The system uses the above triggers to mark RADIUS servers as "Active", however, this does not necessarily mean that the actual server is functional. When the system changes a server state, a trap is automatically sent to the management station. Action should be taken to identify the cause of the failure.
