



RADIUS Server Load Balancing

Last Updated: July 18, 2011

The RADIUS Server Load Balancing feature distributes authentication, authorization, and accounting (AAA) authentication and accounting transactions across servers in a server group. These servers can then share the transaction load, resulting in faster responses to incoming requests by optimally using available servers.

- [Finding Feature Information, page 1](#)
- [Prerequisites for RADIUS Server Load Balancing, page 1](#)
- [Restrictions for RADIUS Server Load Balancing, page 2](#)
- [Information About RADIUS Server Load Balancing, page 2](#)
- [How to Configure RADIUS Server Load Balancing, page 4](#)
- [Configuration Examples for RADIUS Server Load Balancing, page 9](#)
- [Additional References, page 16](#)
- [Feature Information for RADIUS Server Load Balancing, page 18](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the Feature Information Table at the end of this document.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Prerequisites for RADIUS Server Load Balancing

- AAA must be configured on your RADIUS server.

- RADIUS must be configured for functions such as authentication, accounting, or static route download.
- AAA RADIUS server groups must be established.

Restrictions for RADIUS Server Load Balancing

- Load balancing is not supported on proxy RADIUS servers.
- Incoming RADIUS requests, such as Packet of Disconnect (POD) requests, are not supported.
- Load balancing is not supported for private server-groups.

Information About RADIUS Server Load Balancing

- [How RADIUS Server Load Balancing Works, page 2](#)
- [How Transactions Are Load-Balanced Across RADIUS Server Groups, page 3](#)
- [RADIUS Server Status and Automated Testing, page 3](#)

How RADIUS Server Load Balancing Works

Load balancing distributes batches of transactions to servers within a server group. It assigns each batch of transactions to the server with the lowest number of outstanding transactions in its queue. The process of assigning a batch of transactions is as follows:

- The first transaction is received for a new batch.
- All server transaction queues are checked.
- The server with the lowest number of outstanding transactions is identified.
- The identified server is assigned the next batch of transactions.

Batch size is a user configured parameter. Changes in batch size may impact CPU load and network throughput. As batch size increases, CPU load decreases and network throughput increases. However, if a large batch size is used, all available server resources may not be fully utilized. As batch size decreases, CPU load increases, and network throughput decreases. It is recommended that the default batch size, 25, be used because it is optimal for high throughput, without adversely impacting CPU load.



Note

There is no set number for large or small batch sizes. As a frame of reference, a batch size greater than 50 is considered large and a batch size less than 25 is considered small.



Note

If you have ten or more servers in a server group, it is recommended that a high batch size be set in order to reduce CPU load.

How Transactions Are Load-Balanced Across RADIUS Server Groups

You can configure load balancing either per named RADIUS server group or for the global RADIUS server group. This server group must be referred to as “radius” in the AAA method lists. All public servers that are part of this server group will then be load balanced.

Authentication and accounting can be configured to use the same server or different servers. In some cases, the same server is used for preauthentication, authentication, or accounting transactions for a session. The preferred server, which is an internal setting and set as default, tells AAA to use same server for the start and stop record for a session regardless of server cost. When using the preferred server setting, it is expected that the server used for the initial transaction (for example, authentication), the preferred server, should also be part of any other server group that is used for a subsequent transaction (for example, accounting).

The preferred server is used unless one of the following states is true:

- The **ignore-preferred-server** keyword is used.
- The preferred server is dead.
- The preferred server is in quarantine.
- The want server flag has been set, overriding the preferred server setting.

The want server flag, an internal setting, is used when the same server must be used for all stages of a multistage transaction regardless of server cost. If the want server is not available, the transaction fails.

You may want to use the **ignore-preferred-server** keyword if you have either of the following configurations:

- Dedicated authentication server and a separate dedicated accounting server.
- Network where you can track all call record statistics and call record details, including start- and stop-records, and those records are stored on separate servers.

Also, if you have a configuration where your authentication servers are a superset of your accounting servers, then the preferred server will not be used.

RADIUS Server Status and Automated Testing

The RADIUS Server Load Balancing feature takes server status into account when assigning batches. Only servers that are verified alive are sent transaction batches. It is recommended that you test the status all RADIUS load-balanced servers, including low usage servers (for example, backup servers).

Transactions are not sent to a server that is marked dead. A server is marked dead until its timer expires, at which time it is in quarantine. A server is in quarantine until it is verified alive by the RADIUS automated tester functionality.

The RADIUS automated tester uses the following steps to determine if a server is alive and available to process transactions:

- A request is sent periodically to the server for a test user ID.
- If an Access-Reject message is returned from the server, the server is alive.
- If no message is returned from the server, it is not alive; that is, the server is either dead or quarantined.

If transactions have been sent to a server that is not responding, before it is marked dead, that transaction is failed over to the next available server. It is recommended that the retry reorder mode for failed transactions be used.

When using the RADIUS automated tester, verify that the test packets being sent by the network access server (NAS) to the AAA servers are being responded to. If the servers are not configured correctly, the packets may be dropped and the server erroneously marked dead.

**Caution**

It is recommended that a test user, one that is not defined on the RADIUS server, be used for RADIUS server automated testing to protect against security issues that may arise if the test user is not correctly configured.

**Note**

If you want to check load balancing transactions at a specific point in time, you can use the **test aaa group** command.

How to Configure RADIUS Server Load Balancing

- [Enabling Load Balancing for Named RADIUS Server Group, page 4](#)
- [Enabling Load Balancing for Global RADIUS Server Group, page 5](#)
- [Troubleshooting RADIUS Server Load Balancing, page 7](#)

Enabling Load Balancing for Named RADIUS Server Group

Use the following task to enable RADIUS Server Load Balancing for a named server group.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **radius-server host** *{hostname| ip-address}* [**test username** *user-name*] [**auth-port** *port-number*] [**ignore-auth-port**] [**acct-port** *port-number*] [**ignore-acct-port**] [**idle-time** *seconds*]
4. **aaa group server radius** *group-name*
5. **load-balance method least-outstanding** [**batch-size** *number*] [**ignore-preferred-server**]

DETAILED STEPS

Command or Action	Purpose
Step 1 enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.

Command or Action	Purpose
<p>Step 2 <code>configure terminal</code></p> <p>Example:</p> <pre>Router# configure terminal</pre>	<p>Enters global configuration mode.</p>
<p>Step 3 <code>radius-server host {hostname ip-address} [test username user-name] [auth-port port-number] [ignore-auth-port] [acct-port port-number] [ignore-acct-port] [idle-time seconds]</code></p> <p>Example:</p> <pre>Router(config)# radius-server host 192.0.2.1 test username test1 idle-time 1</pre>	<p>Enables RADIUS automated testing.</p> <ul style="list-style-type: none"> • The test username keyword must be used to enable RADIUS automated testing, followed by the value for the <i>user-name</i> argument. • By default, auth-port is tested using port 1645. • Use ignore-auth-port to turn off testing of the authentication port. • By default, acct-port is tested using port 1645. • Use ignore-acct-port to turn off testing of the accounting port. • By default, the idle-time is 3600 seconds. The range is 1 - 35791.
<p>Step 4 <code>aaa group server radius group-name</code></p> <p>Example:</p> <pre>Router(config)# aaa group server radius rad-sg</pre>	<p>Enters server group configuration mode.</p>
<p>Step 5 <code>load-balance method least-outstanding [batch-size number] [ignore-preferred-server]</code></p> <p>Example:</p> <pre>Router(config-sg)# load-balance method least-outstanding batch-size 30</pre>	<p>Enables least-outstanding load balancing for a server group.</p> <ul style="list-style-type: none"> • By default, the batch-size is set to 25. A range of 1 - 2147483647 may be used. • By default, the preferred server is enabled. • If you want to disable the preferred-server setting, use the keyword ignore-preferred-server.

Enabling Load Balancing for Global RADIUS Server Group

Use the following task to enable RADIUS Server Load Balancing for the global RADIUS server group. This is the group referred to as “radius” in the AAA method lists.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **radius-server host** *{hostname| ip-address}* [**test username** *user-name*] [**auth-port** *port-number*] [**ignore-auth-port**] [**acct-port** *port-number*] [**ignore-acct-port**] [**idle-time** *seconds*]
4. **radius-server load-balance method least-outstanding** [**batch-size** *number*] [**ignore-preferred-server**]
5. **load-balance method least-outstanding** [**batch-size** *number*][**ignore-preferred-server**]

DETAILED STEPS

Command or Action	Purpose
Step 1 enable Example: <pre>Router> enable</pre>	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2 configure terminal Example: <pre>Router# configure terminal</pre>	Enters global configuration mode.
Step 3 radius-server host <i>{hostname ip-address}</i> [test username <i>user-name</i>] [auth-port <i>port-number</i>] [ignore-auth-port] [acct-port <i>port-number</i>] [ignore-acct-port] [idle-time <i>seconds</i>] Example: <pre>Router(config)# radius-server host 192.0.2.1 test username test1 idle-time 1</pre>	Enables RADIUS automated testing. <ul style="list-style-type: none"> • The test username keyword must be used to enable RADIUS automated testing, followed by the value for the <i>user-name</i> argument. • By default, auth-port is tested using port 1645. • Use ignore-auth-port to turn off testing of the authentication port. • By default, acct-port is tested using port 1645. • Use ignore-acct-port to turn off testing of the accounting port. • By default, the idle-time is 3600 seconds. The range is 1 - 35791.

Command or Action	Purpose
<p>Step 4 <code>radius-server load-balance method least-outstanding [batch-size number] [ignore-preferred-server]</code></p> <p>Example:</p> <pre>Router(config)# radius-server load-balance method least-outstanding</pre>	<p>Enables least-outstanding load balancing for the global RADIUS server group and enters server group configuration mode.</p> <ul style="list-style-type: none"> By default, the batch-size is set to 25. A range of 1 - 2147483647 may be used. <p>Note Batch size may impact throughput and CPU load. It is recommended that the default batch size, 25, be used because it is optimal for high throughput, without adversely impacting CPU load.</p> <ul style="list-style-type: none"> By default, the preferred server is enabled. If you want to disable the preferred server setting, use the ignore-preferred-server keyword.
<p>Step 5 <code>load-balance method least-outstanding [batch-size number][ignore-preferred-server]</code></p> <p>Example:</p> <pre>load-balance method least-outstanding batch-size 5</pre>	<p>Enables RADIUS server load balancing for a named RADIUS server group.</p> <ul style="list-style-type: none"> By default, the batch-size is set to 25. A range of 1 - 2147483647 may be used. By default, the preferred server is enabled. If you want to disable the preferred server setting, use the ignore-preferred-server keyword.

Troubleshooting RADIUS Server Load Balancing

After configuring the RADIUS Server Load Balancing feature, you may monitor the idle timer, dead timer, load balancing server selection, or issue a manual test command to verify server status.

Use the following commands as appropriate for troubleshooting the RADIUS Server Load Balancing feature:

- The **debug aaa test** command can be used to determine when the idle timer or dead timer has expired, when test packets are sent, the status of the server, or to verify server state.
- The **debug aaa sg-server selection** command can be used to examine which server is being selected for load balancing.
- The **test aaa group** command can be used to manually verify RADIUS load-balanced server status.

SUMMARY STEPS

- The idle timer is used to check the server status and is updated with or without any incoming requests. It is useful to monitor the idle timer to determine if there are nonresponsive servers and to keep your RADIUS server status updated in order to efficiently utilize your available resources. For instance, an updated idle timer would help ensure that incoming requests are being sent to servers that are alive.
- For example, the following debug output shows 5 access requests being sent to a server group with a batch size of 3:
- The following example shows the response from a load-balanced RADIUS server that is alive when the username “test” does not match a user profile. The server is verified alive when it issues an Access-Reject response to a AAA packet generated by the **test aaa group** command.

DETAILED STEPS

Step 1 The idle timer is used to check the server status and is updated with or without any incoming requests. It is useful to monitor the idle timer to determine if there are nonresponsive servers and to keep your RADIUS server status updated in order to efficiently utilize your available resources. For instance, an updated idle timer would help ensure that incoming requests are being sent to servers that are alive.

The dead timer is used either to determine that a server is dead or to update a dead server's status appropriately.

Monitoring server selection can help you determine how often the server selection changes. This is effective in analyzing if there is a bottleneck, a large number of queued up requests, or if only specific servers are processing incoming requests.

For example, the following debug output shows when the idle-timer has expired:

Example:

```
Router# debug aaa test
Jul 16 00:07:01: AAA/SG/TEST: Server (192.0.2.245:1700,1701) quarantined.
Jul 16 00:07:01: AAA/SG/TEST: Sending test request(s) to server (192.0.2.245:1700,1701)
Jul 16 00:07:01: AAA/SG/TEST: Sending 1 Access-Requests, 1 Accounting-Requests in current batch.
Jul 16 00:07:01: AAA/SG/TEST(Req#: 1): Sending test AAA Access-Request.
Jul 16 00:07:01: AAA/SG/TEST(Req#: 1): Sending test AAA Accounting-Request.
Jul 16 00:07:01: AAA/SG/TEST: Obtained Test response from server (192.0.2.245:1700,1701)
Jul 16 00:07:01: AAA/SG/TEST: Obtained Test response from server (192.0.2.245:1700,1701)
Jul 16 00:07:01: AAA/SG/TEST: Necessary responses received from server (192.0.2.245:1700,1701)
Jul 16 00:07:01: AAA/SG/TEST: Server (192.0.2.245:1700,1701) marked ALIVE. Idle timer set for 60
sec(s).
Jul 16 00:07:01: AAA/SG/TEST: Server (192.0.2.245:1700,1701) removed from quarantine.
```

Step 2 For example, the following debug output shows 5 access requests being sent to a server group with a batch size of 3:

Example:

```
Router# debug aaa sg-server selection
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: Obtaining least loaded server.
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: [3] transactions remaining in batch. Reusing server.
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: Obtaining least loaded server.
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: [2] transactions remaining in batch. Reusing server.
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: Obtaining least loaded server.
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: [1] transactions remaining in batch. Reusing server.
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: Obtaining least loaded server.
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: No more transactions in batch. Obtaining a new server.
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: Obtaining a new least loaded server.
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: Server[0] load: 3
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: Server[1] load: 0
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: Server[2] load: 0
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: Selected Server[1] with load 0
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: [3] transactions remaining in batch.
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: Obtaining least loaded server.
Jul 16 03:15:05: AAA/SG/SERVER_SELECT: [2] transactions remaining in batch. Reusing server.
```

Step 3 The following example shows the response from a load-balanced RADIUS server that is alive when the username "test" does not match a user profile. The server is verified alive when it issues an Access-Reject response to a AAA packet generated by the `test aaa group` command.

Example:

```
Router# test aaa group SG1 test lab new-code
```



```

00:06:07: RADIUS/ENCODE(00000000):Orig. component type = INVALID
00:06:07: RADIUS/ENCODE(00000000): dropping service type, "radius-server attribute 6 on-for-login-
auth" is off
00:06:07: RADIUS(00000000): Config NAS IP: 192.0.2.4
00:06:07: RADIUS(00000000): sending
00:06:07: RADIUS/ENCODE: Best Local IP-Address 192.0.2.141 for Radius-Server 192.0.2.176
00:06:07: RADIUS(00000000): Send Access-Request to 192.0.2.176:1645 id 1645/1, len 50
00:06:07: RADIUS:   authenticator CA DB F4 9B 7B 66 C8 A9 - D1 99 4E 8E A4 46 99 B4
00:06:07: RADIUS:   User-Password      [2]  18  *
00:06:07: RADIUS:   User-Name          [1]  6  "test"
00:06:07: RADIUS:   NAS-IP-Address    [4]  6  192.0.2.141
00:06:07: RADIUS: Received from id 1645/1 192.0.2.176:1645, Access-Reject, len 44
00:06:07: RADIUS:   authenticator 2F 69 84 3E F0 4E F1 62 - AB B8 75 5B 38 82 49 C3
00:06:07: RADIUS:   Reply-Message    [18] 24
00:06:07: RADIUS:   41 75 74 68 65 6E 74 69 63 61 74 69 6F 6E 20 66 [Authentication f]
00:06:07: RADIUS:   61 69 6C 75 72 65 [failure]
00:06:07: RADIUS(00000000): Received from id 1645/1
00:06:07: RADIUS/DECODE: Reply-Message fragments, 22, total 22 bytes
Router#

```

Configuration Examples for RADIUS Server Load Balancing

- [Global RADIUS Server Group Examples, page 9](#)
- [Named RADIUS Server Group Examples, page 11](#)
- [Idle Timer Monitoring Examples, page 14](#)
- [Preferred Server with the Same Authentication and Authorization Server Example, page 15](#)
- [Preferred Server with Different Authentication and Authorization Servers Example, page 15](#)
- [Preferred Server with Overlapping Authentication and Authorization Servers Example, page 15](#)
- [Preferred Server with Authentication Servers As a Subset of Authorization Servers Example, page 16](#)
- [Preferred Server with Authentication Servers As a Superset of Authorization Servers Example, page 16](#)

Global RADIUS Server Group Examples

The following example shows how to enable load balancing for global RADIUS server groups. It is shown in three parts: the current configuration of RADIUS command output, debug output, and AAA server status information. You can use the delimiting characters to display only the relevant parts of the configuration.

- [Server Configuration and Enabling Load Balancing for Global RADIUS Server Group Example, page 9](#)
- [Debug Output for Global RADIUS Server Group Example, page 10](#)
- [Server Status Information for Global RADIUS Server Group Example, page 11](#)

Server Configuration and Enabling Load Balancing for Global RADIUS Server Group Example

The following shows the relevant RADIUS configuration.

```
Router# show running-config | include radius
```

```

aaa authentication ppp default group radius
aaa accounting network default start-stop group radius
radius-server host 192.0.2.238 auth-port 2095 acct-port 2096 key cisco
radius-server host 192.0.2.238 auth-port 2015 acct-port 2016 key cisco
radius-server load-balance method least-outstanding batch-size 5

```

The lines in the current configuration of RADIUS command output above are defined as follows:

- The **aaa authentication ppp** command authenticates all PPP users using RADIUS.
- The **aaa accounting** command enables the sending of all accounting requests to the AAA server once the client is authenticated and after the disconnect using the keyword **start-stop**.
- The **radius-server host** command defines the IP address of the RADIUS server host with the authorization and accounting ports specified and the authentication and encryption key identified.
- The **radius-server load-balance** command enables load balancing for the global radius server groups with the batch size specified.

Debug Output for Global RADIUS Server Group Example

The debug output below shows the selection of preferred server and processing of requests for the configuration above.

```

Router# show debug
General OS:
  AAA server group server selection debugging is on
#
<sending 10 pppoe requests>
Router#
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT(00000014):No preferred server available.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:Obtaining least loaded server.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:No more transactions in batch. Obtaining a new
server.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:Obtaining a new least loaded server.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:Server[0] load:0
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:Server[1] load:0
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:Selected Server[0] with load 0
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:[5] transactions remaining in batch.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT(00000014):Server (192.0.2.238:2095,2096) now
being used as preferred server
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT(00000015):No preferred server available.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:Obtaining least loaded server.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:[4] transactions remaining in batch. Reusing
server.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT(00000015):Server (192.0.2.238:2095,2096) now
being used as preferred server
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT(00000016):No preferred server available.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:Obtaining least loaded server.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:[3] transactions remaining in batch. Reusing
server.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT(00000016):Server (192.0.2.238:2095,2096) now
being used as preferred server
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT(00000017):No preferred server available.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:Obtaining least loaded server.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:[2] transactions remaining in batch. Reusing
server.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT(00000017):Server (192.0.2.238:2095,2096) now
being used as preferred server
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT(00000018):No preferred server available.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:Obtaining least loaded server.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:[1] transactions remaining in batch. Reusing
server.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT(00000018):Server (192.0.2.238:2095,2096) now
being used as preferred server
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT(00000019):No preferred server available.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:Obtaining least loaded server.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:No more transactions in batch. Obtaining a new
server.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:Obtaining a new least loaded server.

```

```
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:Server[1] load:0
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:Server[0] load:5
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:Selected Server[1] with load 0
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT:[5] transactions remaining in batch.
*Feb 28 13:40:32.199:AAA/SG/SERVER_SELECT(00000019):Server (192.0.2.238:2015,2016) now
being used as preferred server.
```

Server Status Information for Global RADIUS Server Group Example

The output below shows the AAA server status for the global RADIUS server group configuration example.

```
Router# show aaa server
RADIUS:id 4, priority 1, host 192.0.2.238, auth-port 2095, acct-port 2096
  State:current UP, duration 3175s, previous duration 0s
  Dead:total time 0s, count 0
  Quarantined:No
  Authen:request 6, timeouts 1
    Response:unexpected 1, server error 0, incorrect 0, time 1841ms
    Transaction:success 5, failure 0
  Author:request 0, timeouts 0
    Response:unexpected 0, server error 0, incorrect 0, time 0ms
    Transaction:success 0, failure 0
  Account:request 5, timeouts 0
    Response:unexpected 0, server error 0, incorrect 0, time 3303ms
    Transaction:success 5, failure 0
  Elapsed time since counters last cleared:2m
RADIUS:id 5, priority 2, host 192.0.2.238, auth-port 2015, acct-port 2016
  State:current UP, duration 3175s, previous duration 0s
  Dead:total time 0s, count 0
  Quarantined:No
  Authen:request 6, timeouts 1
    Response:unexpected 1, server error 0, incorrect 0, time 1955ms
    Transaction:success 5, failure 0
  Author:request 0, timeouts 0
    Response:unexpected 0, server error 0, incorrect 0, time 0ms
    Transaction:success 0, failure 0
  Account:request 5, timeouts 0
    Response:unexpected 0, server error 0, incorrect 0, time 3247ms
    Transaction:success 5, failure 0
  Elapsed time since counters last cleared:2m
```

The output shows the status of two RADIUS servers. Both servers are up and, in the last 2 minutes, have successfully processed:

- 5 out of 6 authentication requests
- 5 out of 5 accounting requests

Named RADIUS Server Group Examples

The following example shows load balancing enabled for a named RADIUS server group. It is shown in three parts: the current configuration of RADIUS command output, debug output, and AAA server status information.

- [Server Configuration and Enabling Load Balancing for Named RADIUS Server Group Example, page 12](#)
- [Debug Output for Named RADIUS Server Group Example, page 12](#)
- [Server Status Information for Named RADIUS Server Group Example, page 13](#)

Server Configuration and Enabling Load Balancing for Named RADIUS Server Group Example

The following shows the relevant RADIUS configuration.

```
Router# show running-config
.
.
.
aaa group server radius server-group1
  server 192.0.2.238 auth-port 2095 acct-port 2096
  server 192.0.2.238 auth-port 2015 acct-port 2016
  load-balance method least-outstanding batch-size 5
!
aaa authentication ppp default group server-group1
aaa accounting network default start-stop group server-group1
.
.
.
```

The lines in the current configuration of RADIUS command output above are defined as follows:

- The **aaa group server radius** command shows the configuration of a server group with two member servers.
- The **load-balance** command enables load balancing for the global radius server groups with the batch size specified.
- The **aaa authentication ppp** command authenticates all PPP users using RADIUS.
- The **aaa accounting** command enables the sending of all accounting requests to the AAA server once the client is authenticated and after the disconnect using the **start-stop** keyword.

Debug Output for Named RADIUS Server Group Example

The debug output below shows the selection of preferred server and processing of requests for the configuration above.

```
Router#
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT(0000002C):No preferred server available.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT:Obtaining least loaded server.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT:No more transactions in batch. Obtaining a new server.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT:Obtaining a new least loaded server.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT:Server[0] load:0
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT:Server[1] load:0
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT:Selected Server[0] with load 0
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT:[5] transactions remaining in batch.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT(0000002C):Server (192.0.2.238:2095,2096) now being used as preferred server
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT(0000002D):No preferred server available.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT:Obtaining least loaded server.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT:[4] transactions remaining in batch. Reusing server.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT(0000002D):Server (192.0.2.238:2095,2096) now being used as preferred server
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT(0000002E):No preferred server available.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT:Obtaining least loaded server.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT:[3] transactions remaining in batch. Reusing server.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT(0000002E):Server (192.0.2.238:2095,2096) now being used as preferred server
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT(0000002F):No preferred server available.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT:Obtaining least loaded server.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT:[2] transactions remaining in batch. Reusing
```

```

server.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT(0000002F):Server (192.0.2.238:2095,2096) now
being used as preferred server
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT(00000030):No preferred server available.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT:Obtaining least loaded server.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT:[1] transactions remaining in batch. Reusing
server.
*Feb 28 13:51:16.019:AAA/SG/SERVER_SELECT(00000030):Server (192.0.2.238:2095,2096) now
being used as preferred server
*Feb 28 13:51:16.023:AAA/SG/SERVER_SELECT(00000031):No preferred server available.
*Feb 28 13:51:16.023:AAA/SG/SERVER_SELECT:Obtaining least loaded server.
*Feb 28 13:51:16.023:AAA/SG/SERVER_SELECT:No more transactions in batch. Obtaining a new
server.
*Feb 28 13:51:16.023:AAA/SG/SERVER_SELECT:Obtaining a new least loaded server.
*Feb 28 13:51:16.023:AAA/SG/SERVER_SELECT:Server[1] load:0
*Feb 28 13:51:16.023:AAA/SG/SERVER_SELECT:Server[0] load:5
*Feb 28 13:51:16.023:AAA/SG/SERVER_SELECT:Selected Server[1] with load 0
*Feb 28 13:51:16.023:AAA/SG/SERVER_SELECT:[5] transactions remaining in batch.
*Feb 28 13:51:16.023:AAA/SG/SERVER_SELECT(00000031):Server (192.0.2.238:2015,2016) now
being used as preferred server
*Feb 28 13:51:16.023:AAA/SG/SERVER_SELECT(00000032):No preferred server available.
*Feb 28 13:51:16.023:AAA/SG/SERVER_SELECT:Obtaining least loaded server.
*Feb 28 13:51:16.023:AAA/SG/SERVER_SELECT:[4] transactions remaining in batch. Reusing
server.
.
.
.

```

Server Status Information for Named RADIUS Server Group Example

The output below shows the AAA server status for the named RADIUS server group configuration example.

```

Router# show aaa servers
RADIUS:id 8, priority 1, host 192.0.2.238, auth-port 2095, acct-port 2096
State:current UP, duration 3781s, previous duration 0s
Dead:total time 0s, count 0
Quarantined:No
Authen:request 0, timeouts 0
      Response:unexpected 0, server error 0, incorrect 0, time 0ms
      Transaction:success 0, failure 0
Author:request 0, timeouts 0
      Response:unexpected 0, server error 0, incorrect 0, time 0ms
      Transaction:success 0, failure 0
Account:request 0, timeouts 0
      Response:unexpected 0, server error 0, incorrect 0, time 0ms
      Transaction:success 0, failure 0
Elapsed time since counters last cleared:0m
RADIUS:id 9, priority 2, host 192.0.2.238, auth-port 2015, acct-port 2016
State:current UP, duration 3781s, previous duration 0s
Dead:total time 0s, count 0
Quarantined:No
Authen:request 0, timeouts 0
      Response:unexpected 0, server error 0, incorrect 0, time 0ms
      Transaction:success 0, failure 0
Author:request 0, timeouts 0
      Response:unexpected 0, server error 0, incorrect 0, time 0ms
      Transaction:success 0, failure 0
Account:request 0, timeouts 0
      Response:unexpected 0, server error 0, incorrect 0, time 0ms
      Transaction:success 0, failure 0
Elapsed time since counters last cleared:0m

```

The output shows the status of two RADIUS servers. Both servers are alive, and no requests have been processed since the counters were cleared 0 minutes ago.

Idle Timer Monitoring Examples

The following example shows idle timer and related server state for load balancing enabled for a named RADIUS server group. It is shown in two parts: the current configuration of RADIUS command output and debug output.

- [Server Configuration and Enabling Load Balancing for Idle Timer Monitoring Example, page 14](#)
- [Debug Output for Idle Timer Monitoring Example, page 14](#)

Server Configuration and Enabling Load Balancing for Idle Timer Monitoring Example

The following shows the relevant RADIUS configuration.

```
Router# show running-config | include radius
aaa group server radius server-group1
radius-server host 192.0.2.238 auth-port 2095 acct-port 2096 test username junk1 idle-
time 1 key cisco
radius-server host 192.0.2.238 auth-port 2015 acct-port 2016 test username junk1 idle-
time 1 key cisco
radius-server load-balance method least-outstanding batch-size 5
```

The lines in the current configuration of RADIUS command output above are defined as follows:

- The **aaa group server radius** command shows the configuration of a server group.
- The **radius-server host** command defines the IP address of the RADIUS server host with the authorization and accounting ports specified and the authentication and encryption key identified.
- The **radius-server load-balance** command enables load balancing for the radius server with the batch size specified.

Debug Output for Idle Timer Monitoring Example

The debug output below shows the test requests being sent to servers. The response to the test request sent to the server is received, the server is removed from quarantine as appropriate, marked alive, and then the idle timer is reset.

```
Router#
*Feb 28 13:52:20.835:AAA/SG/TEST:Server (192.0.2.238:2015,2016) quarantined.
*Feb 28 13:52:20.835:AAA/SG/TEST:Sending test request(s) to server (192.0.2.238:2015,2016)
*Feb 28 13:52:20.835:AAA/SG/TEST:Sending 1 Access-Requests, 1 Accounting-Requests in
current batch.
*Feb 28 13:52:20.835:AAA/SG/TEST(Req#:1):Sending test AAA Access-Request.
*Feb 28 13:52:20.835:AAA/SG/TEST(Req#:1):Sending test AAA Accounting-Request.
*Feb 28 13:52:21.087:AAA/SG/TEST:Obtained Test response from server
(192.0.2.238:2015,2016)
*Feb 28 13:52:22.651:AAA/SG/TEST:Obtained Test response from server
(192.0.2.238:2015,2016)
*Feb 28 13:52:22.651:AAA/SG/TEST:Necessary responses received from server
(192.0.2.238:2015,2016)
*Feb 28 13:52:22.651:AAA/SG/TEST:Server (192.0.2.238:2015,2016) marked ALIVE. Idle timer
set for 60 secs(s).
*Feb 28 13:52:22.651:AAA/SG/TEST:Server (192.0.2.238:2015,2016) removed from quarantine.
.
.
.
```

Preferred Server with the Same Authentication and Authorization Server Example

The following example shows an authentication server group and an authorization server group that use the same servers, 209.165.200.225 and 209.165.200.226. Both server groups have the preferred server flag enabled.

```
aaa group server radius authentication-group
server 209.165.200.225 key radkey1
server 209.165.200.226 key radkey2
aaa group server radius accounting-group
server 209.165.200.225 key radkey1
server 209.165.200.226 key radkey2
```

Once a preferred server is selected for a session, all transactions for that session will continue to use the original preferred server. The servers 209.165.200.225 and 209.165.200.226 will be load balanced based on sessions rather than transactions.

Preferred Server with Different Authentication and Authorization Servers Example

The following example shows an authentication server group that uses servers 209.165.200.225 and 209.165.200.226 and an authorization server group that uses servers 209.165.201.1 and 209.165.201.2. Both server groups have the preferred server flag enabled.

```
aaa group server radius authentication-group
server 209.165.200.225 key radkey1
server 209.165.200.226 key radkey2
aaa group server radius accounting-group
server 209.165.201.1 key radkey3
server 209.165.201.2 key radkey4
```

The authentication server group and the accounting server group do not share any common servers. A preferred server will never be found for accounting transactions, therefore, authentication and accounting servers will be load balanced based on transactions. Start and stop records will be sent to the same server for a session.

Preferred Server with Overlapping Authentication and Authorization Servers Example

The following example shows an authentication server group that uses servers 209.165.200.225, 209.165.200.226, and 209.165.201.1 and an accounting server group that uses servers 209.165.201.1 and 209.165.201.2. Both server groups have the preferred server flag enabled.

```
aaa group server radius authentication-group
server 209.165.200.225 key radkey1
server 209.165.200.226 key radkey2
server 209.165.201.1 key radkey3
aaa group server radius accounting-group
server 209.165.201.1 key radkey3
server 209.165.201.2 key radkey4
```

If all servers have equal transaction processing capability, one-third of all authentication transactions will be directed towards server 209.165.201.1. Therefore, one-third of all accounting transactions will also be

directed towards server 209.165.201.1. The remaining two-thirds accounting transactions will be load balanced equally between servers 209.165.201.1 and 209.165.201.2. The server 209.165.201.1 will receive fewer authentication transactions since server 209.165.201.1 will have outstanding accounting transactions.

Preferred Server with Authentication Servers As a Subset of Authorization Servers Example

The following example shows an authentication server group that uses servers 209.165.200.225 and 209.165.200.226 and an authorization server group that uses servers 209.165.200.225, 209.165.200.226, and 209.165.201.1. Both server groups have the preferred server flag enabled.

```
aaa group server radius authentication-group
  server 209.165.200.225 key radkey1
  server 209.165.200.226 key radkey2
aaa group server radius accounting-group
  server 209.165.200.225 key radkey1
  server 209.165.200.226 key radkey2
  server 209.165.201.1 key radkey3
```

One-half of all authentication transactions will be sent to server 209.165.200.225 and the other half to server 209.165.200.226. Servers 209.165.200.225 and 209.165.200.226 will be the preferred servers for authentication and accounting transaction, therefore there will be an equal distribution of authentication and accounting transactions across servers 209.165.200.225 and 209.165.200.226. Server 209.165.201.1 will be relatively unused.

Preferred Server with Authentication Servers As a Superset of Authorization Servers Example

The following example shows an authentication server group that uses servers 209.165.200.225 and 209.165.200.226, and 209.165.201.1 and an authorization server group that uses servers 209.165.200.225 and 209.165.200.226. Both server groups have the preferred server flag enabled.

```
aaa group server radius authentication-group
  server 209.165.200.225 key radkey1
  server 209.165.200.226 key radkey2
  server 209.165.201.1 key radkey3
aaa group server radius accounting-group
  server 209.165.200.225 key radkey1
  server 209.165.200.226 key radkey2
```

Initially, one-third of authentication transactions will be assigned to each server in the authorization server group. As accounting transactions are generated for more sessions, the accounting transactions will only be sent to servers 209.165.200.225 and 209.165.200.226, since the preferred server flag is on. As servers 209.165.200.225 and 209.165.200.226 begin to process more transactions, authentication transactions will start to be sent to server 209.165.201.1. The transaction requests authenticated by server 209.165.201.1, will not have any preferred server setting and will be split between servers 209.165.200.225 and 209.165.200.226, which negates the use of the preferred server flag. This configuration should be used cautiously.

Additional References

The following sections provide references related to the RADIUS Server Load Balancing feature.

Related Documents

Related Topic	Document Title
AAA and RADIUS	<i>Cisco IOS Security Configuration Guide: Securing User Services</i> , Release 12.4T
AAA Server Groups and RADIUS Configuration	“ Configuring RADIUS ” module.
Failover retry reorder mode	“ RADIUS Server Reorder on Failure ” module.

Standards

Standards	Title
None	--

MIBs

MIBs	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

RFCs

RFCs	Title
No new or modified RFCs are supported by this feature.	--

Technical Assistance

Description	Link
<p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p>	http://www.cisco.com/techsupport

Feature Information for RADIUS Server Load Balancing

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1 Feature Information for RADIUS Server Load Balancing

Feature Name	Releases	Feature Information
RADIUS Server Load Balancing	12.2(28)SB 12.4(11)T 12.2(33)SRC	<p>The RADIUS Server Load Balancing feature distributes authentication, authorization, and accounting (AAA) authentication and accounting transactions across servers in a server group. These servers can then share the transaction load, resulting in faster responses to incoming requests by optimally using available servers.</p> <p>This feature was integrated into Cisco IOS Release 12.2(28)SB.</p> <p>This feature was integrated into Cisco IOS Release 12.4(11)T.</p> <p>This feature was integrated into Cisco IOS Release 12.2(33)SRC.</p> <p>The following commands were introduced or modified: debug aaa sg-server selection, debug aaa test, load-balance (server-group), radius-server host, radius-server load-balance, test aaa group.</p>
RADIUS Server Load Balancing porting	Cisco IOS XE Release 2.1	This feature was introduced on Cisco ASR 1000 series routers.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.