



System Configuration: Advanced

This chapter addresses the ACS internal database replication and RDBMS synchronization features in the System Configuration section of Cisco Secure Access Control Server Release 4.0 for Windows, hereafter referred to as ACS.

This chapter contains the following sections:

- [ACS Internal Database Replication, page 9-1](#)
- [RDBMS Synchronization, page 9-17](#)
- [IP Pools Server, page 9-28](#)
- [IP Pools Address Recovery, page 9-33](#)

ACS Internal Database Replication

This section provides information about the ACS internal database replication feature, including procedures for implementing this feature and configuring the ACSs involved.



Note

ACS does not support distributed deployments in a NAT environment. If a Primary or Secondary address is NATed, the database replication file will indicate shared secret mismatch.

This section contains the following topics:

- [About ACS Internal Database Replication, page 9-2](#)
 - [Replication Process, page 9-3](#)
 - [Replication Frequency, page 9-5](#)
- [Important Implementation Considerations, page 9-5](#)
- [Database Replication Versus Database Backup, page 9-6](#)
- [Database Replication Logging, page 9-7](#)
- [Replication Options, page 9-7](#)
 - [Replication Components Options, page 9-7](#)
 - [Outbound Replication Options, page 9-9](#)
 - [Inbound Replication Options, page 9-10](#)
- [Implementing Primary and Secondary Replication Setups on ACSs, page 9-10](#)

- [Configuring a Secondary ACS, page 9-11](#)
- [Replicating Immediately, page 9-12](#)
- [Scheduling Replication, page 9-14](#)
- [Disabling ACS Database Replication, page 9-15](#)
- [Configuring Automatic Change Password Replication, page 9-16](#)
- [Database Replication Event Errors, page 9-16](#)

About ACS Internal Database Replication

Database replication creates mirror systems of ACSs by duplicating parts of the primary ACS setup to one or more secondary ACSs. You can configure your AAA clients to use these secondary ACSs if the primary ACS fails or is unreachable. With a secondary ACS whose ACS internal database is a replica of the ACS internal database on the primary ACS, if the primary ACS goes out of service, incoming requests are authenticated without network downtime, provided that your AAA clients are configured to fail over to the secondary ACS.

You can use database replication to:

- Select the parts of the primary ACS configuration to be replicated.
- Control the timing of the replication process, including creating schedules.
- Export selected configuration items from the primary ACS.
- Securely transport selected configuration data from the primary ACS to one or more secondary ACSs.
- Update the secondary ACSs to create matching configurations.

The following items cannot be replicated:

- IP pool definitions (for more information, see [About IP Pools Server, page 9-28](#)).
- ACS certificate and private key files.
- Unknown user group mapping configuration.
- Dynamically-mapped users.
- Settings on the ACS Service Management page in the System Configuration section.
- RDBMS Synchronization settings.
- Third-party software, such as Novell Requestor or RSA ACE client software.



Tip

For a list of the various components and what database replication encompasses, see [Replication Components Options, page 9-7](#).

With regard to database replication, we make the following distinctions about ACSs:

- **Primary ACS**—An ACS that sends replicated ACS internal database components to other ACSs.
- **Secondary ACS**—An ACS that receives replicated ACS internal database components from a primary ACS. In the web interface, these are identified as replication partners.

An ACS can be a primary ACS and a secondary ACS, provided that it is not configured to be a secondary ACS to an ACS for which it performs as a primary ACS.

**Note**

Bidirectional replication, wherein an ACS sends database components to and receives database components from the same remote ACS, is not supported. Replication fails if an ACS is configured to replicate to and from the same ACS.

**Note**

All ACSs that are involved in replication must run the same release of the ACS software. For example, if the primary ACS is running ACS version 3.2, all secondary ACSs should be running ACS version 3.2. Because patch releases can introduce significant changes to the ACS internal database, we strongly recommend that ACSs involved in replication use the same patch level.

Replication Process

This topic describes the process of database replication, including the interaction between a primary ACS and each of its secondary ACSs. The following steps occur in database replication:

1. The primary ACS determines if its database has changed since the last successful replication. If it has, replication proceeds. If it has not, replication is aborted. No attempt is made to compare the databases of the primary and secondary ACSs.

**Tip**

You can force replication to occur by making one change to a user or group profile, such as changing a password or modifying a RADIUS attribute.

2. The primary ACS contacts the secondary ACS. In this initial connection, the following four events occur:
 - a. The two ACSs perform mutual authentication based on the shared secret of the primary ACS. If authentication fails, replication fails.

**Note**

On the secondary ACS, the AAA Servers table entry for the primary ACS must have the same shared secret that the primary ACS has for itself in its own AAA Servers table entry. The shared secret of the secondary ACS is irrelevant.

- b. The secondary ACS verifies that it is not configured to replicate to the primary ACS. If it is, replication is aborted. ACS does not support bidirectional replication, wherein an ACS can act as a primary and a secondary ACS to the same remote ACS.
 - c. The primary ACS verifies that the version of ACS that the secondary ACS is running is the same as its own version of ACS. If not, replication fails.
 - d. The primary ACS compares the list of database components that it is configured to send with the list of database components that the secondary ACS is configured to receive. If the secondary ACS is not configured to receive any of the components that the primary ACS is configured to send, the database replication fails.
3. After the primary ACS has determined which components to send to the secondary ACS, the replication process continues on the primary ACS:
 - a. The primary ACS stops its authentication and creates a copy of the ACS internal database components that it is configured to replicate. During this step, if AAA clients are configured properly, those that usually use the primary ACS fail over to another ACS.

- b. The primary ACS resumes its authentication service. It also compresses and encrypts the copy of its database components for transmission to the secondary ACS.
 - c. The primary ACS transmits the compressed, encrypted copy of its database components to the secondary ACS. This transmission occurs over a TCP connection by using port 2000. The TCP session uses a 128-bit encrypted, Cisco-proprietary protocol.
4. After the preceding events on the primary ACS, the database replication process continues on the secondary ACS:
 - a. The secondary ACS receives the compressed, encrypted copy of the ACS internal database components from the primary ACS. After transmission of the database components is complete, the secondary ACS decompresses the database components.
 - b. The secondary ACS stops its authentication service and replaces its database components with the database components that it received from the primary ACS. During this step, if AAA clients are configured properly, those that usually use the secondary ACS fail over to another ACS.
 - c. The secondary ACS resumes its authentication service.


ACS can act as a primary ACS and a secondary ACS. [Figure 9-1](#) shows a cascading replication scenario. Server 1 acts only as a primary ACS, replicating to servers 2 and 3, which act as secondary ACSs. After replication from server 1 to server 2 has completed, server 2 acts as a primary ACS while replicating to servers 4 and 5. Similarly, server 3 acts as a primary ACS while replicating to servers 6 and 7.

**Note**

If you intend to use cascading replication to replicate network configuration device tables, you must configure the primary ACS with all ACSs that will receive replicated database components, regardless of whether they receive replication directly or indirectly from the primary ACS. In [Figure 9-1](#), server 1 must have an entry in its AAA Servers table for each of the other six ACSs. If this is not done, after replication, servers 2 and 3 do not have servers 4 through 7 in their AAA Servers tables and replication will fail.

If server 2 were configured to replicate to server 1 in addition to receiving replication from server 1, replication to server 2 would fail. ACS cannot support such a configuration, known as bidirectional replication. To safeguard against this, a secondary ACS aborts replication when its primary ACS appears on its Replication list.

Figure 9-1 Cascading Database Replication



Replication Frequency

The frequency with which your ACSs replicate can have important implications for overall AAA performance. With shorter replication frequencies, a secondary ACS is more up to date with the primary ACS. This frequency produces a more current secondary ACS if the primary ACS fails.

Frequent replications incur a cost. The more frequent the replication, the higher the load on a multi-ACS architecture and your network environment. If you schedule frequent replications, network traffic is much higher. Also, processing load on the replicating systems is increased. Replication consumes system resources and briefly interrupts authentication; thus the more often replication occurs, the greater the impact on the AAA performance of the ACS. And because service is momentarily interrupted on both servers, NAS failovers can occur.

**Note**

Regardless of how frequently replication is scheduled to occur, it only occurs when the database of the primary ACS has changed since the last successful replication.

This issue is more apparent with databases that are large or frequently change. Database replication is a nonincremental, destructive backup. In other words, it completely replaces the database and configuration on the secondary ACS every time it runs. Therefore, a large database results in substantial amounts of data being transferred, and the processing overhead can also be large.

Important Implementation Considerations

Several important points to consider when you implement the ACS Database Replication feature are:

- ACS only supports database replication to other ACSs. All ACSs participating in ACS internal database replication must run the same version of ACS. We strongly recommend that ACSs that are involved in replication use the same patch level, too.
- You must ensure correct configuration of the AAA Servers table in all ACSs that are involved in replication.
 - In its AAA Servers table, a primary ACS must have an accurately configured entry for each secondary ACS.

**Note**

If you intend to use cascading replication to replicate network configuration device tables, you must configure the primary ACS with all ACSs that will receive replicated database components, regardless of whether they receive replication directly or indirectly from that primary ACS. For example, if the primary ACS replicates to two secondary ACSs, which, in turn, each replicate to two more ACSs, the primary ACS must have AAA server configurations for all six ACSs that will receive replicated database components.

- In its AAA Servers table, a secondary ACS must have an accurately configured entry for each of its primary ACSs.
 - On a primary ACS and all its secondary ACSs, the AAA Servers table entries for the primary ACS must have identical shared secrets.
- Only suitably configured, valid ACSs can be secondary ACSs. To configure a secondary ACS for database replication, see [Configuring a Secondary ACS, page 9-11](#).

- Replication only occurs when the database of the primary ACS has changed since the last successful replication, regardless of how frequently replication is scheduled to occur. When a scheduled or manually started replication begins, the primary ACS automatically aborts replication if its database has not changed since the last successful replication.

**Tip**

You can force replication to occur by making one change to a user or group profile, such as changing a password or modifying a RADIUS attribute.

- Replication to secondary ACSs occurs sequentially in the order listed in the Replication list under Replication Partners on the ACS Database Replication page.
- You must configure a secondary ACS that is receiving replicated components to accept database replication from the primary ACS. To configure a secondary ACS for database replication, see [Configuring a Secondary ACS, page 9-11](#).
- ACS does not support bidirectional database replication. The secondary ACS that receives the replicated components verifies that the primary ACS is not on its Replication list. If not, the secondary ACS accepts the replicated components. If so, it rejects the components.
- For all components (except for Network Access Profiles), if you replicate user accounts, ensure that you name external database configurations identically on primary and secondary ACSs. A replicated user account retains its association with the database that is assigned to provide authentication or posture validation service, regardless of whether a database configuration of the same name exists on the secondary ACS. For example, if user account is associated with a database named **WestCoast LDAP**; on the primary ACS, the replicated user account on all secondary ACSs remains associated with an external user database named **WestCoast LDAP** even if you have not configured an LDAP database instance of that name.
- For all components (except for Network Access Profiles), in order to replicate user and group settings that use user-defined RADIUS vendor and VSAs, you must manually add the user-defined RADIUS vendor and VSA definitions on primary and secondary ACSs, ensuring that the RADIUS vendor slots that the user-defined RADIUS vendors occupy are identical on each ACS. After you have done so, replication of settings by using user-defined RADIUS vendors and VSAs is supported. For more information about user-defined RADIUS vendors and VSAs, see [Custom RADIUS Vendors and VSAs, page 9-19](#).

Database Replication Versus Database Backup

Do not confuse database replication with system backup. Database replication does *not* replace System Backup. While both features protect against partial or complete server loss, each feature addresses the issue in a different way:

- System Backup archives data into a format that you can later use to restore the configuration if the system fails or the data becomes corrupted. The backup data is stored on the local hard drive, and can be copied and removed from the system for long-term storage. You can store several generations of database backup files.
- You use ACS Database Replication to copy various components of the ACS internal database to other ACSs. This method can help you plan a failover AAA architecture, and reduce the complexity of your configuration and maintenance tasks. While unlikely, it is possible that ACS Database Replication can propagate a corrupted database to the ACSs that generate your backup files.

**Caution**

Because the possibility of replicating a corrupted database always exists, we strongly recommend that you implement a backup plan, especially in mission-critical environments. For more information about backing up the ACS internal database, see [ACS Backup, page 8-7](#) and [Appendix D, “CSUtil Database Utility”](#).

Database Replication Logging

ACS logs all replication events—regardless of whether they are successful—in two files. The:

- Windows Event Log
- Database Replication report

To view the Windows Event Log, use the Windows administration utilities. You can view recent reports in the Reports and Activity section of ACS.

For more information about ACS reports, see [Chapter 1, “Overview”](#).

Replication Options

The ACS web interface provides three sets of options for configuring ACS Database Replication.

This section contains the following topics:

- [Replication Components Options, page 9-7](#)
- [Outbound Replication Options, page 9-9](#)
- [Inbound Replication Options, page 9-10](#)

Replication Components Options

You can specify the ACS internal database components that an ACS sends as a primary ACS and the components that it receives as a secondary ACS.

For increased security, you might want to have one ACS always be the sender and the other ACSs always be the receivers. You can use this method to ensure that all your ACSs are configured identically.

**Note**

The ACS internal database components that a secondary ACS receives *overwrite* the ACS internal database components on the secondary ACS. Any information that is unique to the overwritten database component is lost. For example, if the Receive checkbox is selected for the User and Group Database, any existing user or group records are lost on replication when the new ACS internal database is received.

Table 9-1 describes the Replication Components table on the ACS Database Replication page and describes the component options that are replicated.

Table 9-1 Replication Component Descriptions

Component	What Gets Replicated?
User and group database	Groups and users. Using this option excludes the use of the Group database only option.
Group database only	Groups, but not for users. Using this option excludes the use of the User and group database option.
Network Configuration Device tables ¹	AAA Servers tables and the AAA Clients tables in the Network Configuration section. This also controls whether Network Device Groups (NDG) are replicated.
Distribution table	Proxy Distribution Table in the Network Configuration section.
Interface configuration	Advanced Options settings, RADIUS settings, and TACACS+ settings from the Interface Configuration section.
Interface security settings	Administrators and security information for the ACS web interface.
Password validation settings	Password validation settings.
EAP-FAST master keys and policies	Active and retired master keys and policies for EAP-FAST.
Network Access Profiles ²	A collaboration of configuration settings. These include: Network Access Profiles, Posture Validation settings, AAA clients and hosts, external user database configuration, global authentication configuration, NDGs, user-defined RADIUS dictionaries, shared profile components and additional logging attributes.

1. If you intend to use cascading replication to replicate network configuration device tables, you must configure the primary ACS with all ACSs that will receive replicated database components, regardless of whether they receive replication directly or indirectly from the primary ACS. For example, if the primary ACS replicates to two secondary ACSs that, in turn, each replicate to two more ACSs, the primary ACS must have AAA server configurations for all six ACSs that will receive replicated database components.
2. Replication of Network Access Profiles contradicts the replication of Network Configuration Device tables, therefore do not check both of these components at the same time. NAP settings will override all other settings. Dynamically-mapped users are not replicated, only statically-added users are replicated.

If mirroring the entire database might send confidential information to the secondary ACS, such as the Proxy Distribution Table, you can configure the primary ACS to send only a specific category of database information.

Outbound Replication Options

In the Outbound Replication table on the ACS Database Replication page, you can schedule outbound replication and specify the secondary ACSs for this primary ACS.

Table 9-2 Outbound Replication Options

Option	Description
Scheduling Options	
Manually	ACS does not perform automatic database replication.
Automatically Triggered Cascade ¹	ACS performs database replication to the configured list of secondary ACSs when database replication from a primary ACS completes. You use this option to build a propagation hierarchy of ACS, relieving a primary ACS from the burden of propagating the replicated components to every other ACS. For an illustration of cascade replication, see Figure 9-1 .
Every X minutes	ACS performs, on a set frequency, database replication to the configured list of secondary ACSs. The unit of measurement is minutes, with a default update frequency of 60 minutes.
At specific times	ACS performs, at the time that is specified in the day and hour graph, database replication to the configured list of secondary ACSs. The minimum interval is one hour, and the replication occurs on the hour selected.
Partner Options	
AAA Server	Represents the secondary ACSs that this primary ACS <i>does not</i> send replicated components to.
Replication	Represents the secondary ACSs that this primary ACS <i>does</i> send replicated components to.
Replication Timeout	Specifies the number of minutes that this primary ACS continues replicating to a secondary ACS. When the timeout value is exceeded, ACS terminates replication to the secondary ACS it was attempting to replicate to and then it restarts the CSAuth service. The replication timeout feature helps prevent loss of AAA services due to stalled replication communication, which can occur when the network connection between the primary and secondary ACS is abnormally slow or when a fault occurs within either ACS. The default value is five minutes.

1. If you intend to use cascading replication to replicate network configuration device tables, you must configure the primary ACS with all ACSs that will receive replicated database components, regardless of whether they receive replication directly or indirectly from the primary ACS. For example, if the primary ACS replicates to two secondary ACSs which, in turn, each replicate to two more ACSs, the primary ACS must have AAA server configurations for all six ACSs that will receive replicated database components.

The items in the AAA Server and Replication lists reflect the AAA servers configured in the AAA Servers table in Network Configuration. To make a particular ACS available as a secondary ACS, you must first add that ACS to the AAA Servers table of the primary ACS.

**Tip**

The size of the components replicated affects the time required for replication. For example, replicating a user database containing 80,000 user profiles takes more time than replicating a user database containing 500 user profiles. You may need to monitor successful replication events to determine a reasonable timeout value for your implementation.

ACS does not support bidirectional database replication. A secondary ACS receiving replicated components verifies that the primary ACS is not on its Replication list. If not, the secondary ACS accepts the replicated components. If so, it rejects the components.

Inbound Replication Options

You can specify the primary ACSs from which a secondary ACS accepts replication. This option appears in the Inbound Replication table on the ACS Database Replication page.

The **Accept replication from** list controls which ACSs the current ACS does accept replicated components from. The list contains:

- **Any Known ACS Server**—If this option is selected, ACS accepts replicated components from any ACS configured in the AAA Servers table in Network Configuration.
- **Other AAA servers**—The list displays all the AAA servers configured in the AAA Servers table in Network Configuration. If a specific AAA server name is selected, ACS accepts replicated components only from the ACS specified.

**Note**

ACS does not support bidirectional database replication. A secondary ACS receiving replicated components verifies that the primary ACS is not on its Replication list. If not, the secondary ACS accepts the replicated components. If so, it rejects the components.

For more information about the AAA Servers table in Network Configuration, see [AAA Server Configuration, page 4-15](#).

Implementing Primary and Secondary Replication Setups on ACSs

If you implement a replication scheme that uses cascading replication, the ACS configured to replicate only when it has received replicated components from another ACS acts as a primary ACS and as a secondary ACS. First, it acts as a secondary ACS while it receives replicated components, and then it acts as a primary ACS while it replicates components to other ACSs. For an illustration of cascade replication, see [Figure 9-1](#).

To implement primary and secondary replication setups on ACSs:

Step 1

On each secondary ACS:

- a. In the Network Configuration section, add the primary ACS to the AAA Servers table.
For more information about adding entries to the AAA Servers table, see [AAA Server Configuration, page 4-15](#).
- b. Configure the secondary ACS to receive replicated components. For instructions, see [Configuring a Secondary ACS, page 9-11](#).

- Step 2** On the primary ACS:
- In the Network Configuration section, add each secondary ACS to the AAA Servers table.



Note If you intend to use cascading replication to replicate network configuration device tables, you must configure the primary ACS with all ACSs that will receive replicated database components, regardless of whether they receive replication directly or indirectly from the primary ACS. For example, if the primary ACS replicates to two secondary ACSs which, in turn, each replicate to two more ACSs, the primary ACS must have AAA server configurations for all six ACSs that will receive replicated database components.

For more information about adding entries to the AAA Servers table, see [AAA Server Configuration, page 4-15](#).

- If you want to replicate according to a schedule, at intervals, or whenever the primary ACS has received replicated components from another ACS, see [Scheduling Replication, page 9-14](#).
- If you want to initiate replication immediately, see [Replicating Immediately, page 9-12](#).

Configuring a Secondary ACS



Note

If this feature does not appear, choose **Interface Configuration > Advanced Options**. Then, check the **ACS Database Replication** check box. Check the **Distributed System Settings** check box if not already selected.

The ACS Database Replication feature requires that you configure specific ACSs to act as secondary ACSs. The components that a secondary ACS is to receive must be explicitly specified, as must be its primary ACS.

Replication is always initiated by the primary ACS. For more information about sending replication components, see [Replicating Immediately, page 9-12](#) or [Scheduling Replication, page 9-14](#).



Caution

The ACS internal database components received by a secondary ACS *overwrite* the ACS internal database components on the secondary ACS. Any information unique to the overwritten database component is lost.

Before You Begin

Ensure correct configuration of the AAA Servers table in the secondary ACS. This secondary ACS must have an entry in its AAA Servers table for each of its primary ACSs. Also, the AAA Servers table entry for each primary ACS must have the same shared secret that the primary ACS has for its own entry in its AAA Servers table. For more information about the AAA Servers table, see [AAA Server Configuration, page 4-15](#).

To configure an ACS to be a secondary ACS:

- Step 1** Log in to the web interface on the secondary ACS.
- Step 2** In the navigation bar, click **System Configuration**.

Step 3 Click **Database Replication**.

The Database Replication Setup page appears.

Step 4 In the Replication Components table, select the **Receive** check box for each database component to be received from a primary ACS.

For more information about replication components, see [Replication Components Options, page 9-7](#).

Step 5 Make sure that no ACS that the secondary ACS is to receive replicated components from is included in the Replication list. If so, select the primary ACS in the Replication list and click the <-- (left arrow) to move it to the AAA Servers list.

Note ACS does not support bidirectional database replication. A secondary ACS receiving replicated components verifies that the primary ACS is not on its Replication list. If not, the secondary ACS accepts the replicated components. If so, it aborts replication.

Step 6 If the secondary ACS is to receive replication components from *only one* primary ACS, from the Accept replication from list, select the name of the primary ACS.

The primary ACSs available in the Accept replication from list are determined by the AAA Servers table in the Network Configuration section. For more information about the AAA Servers table, see [AAA Server Configuration, page 4-15](#).



Note On the primary ACS and all secondary ACSs, the AAA Servers table entries for the primary ACS must have identical shared secrets.

Step 7 If the secondary ACS is to receive replication components from *more than one* primary ACS, from the Accept replication from list, select **Any Known ACS Server**.

The Any Known ACS Server option is limited to the ACSs listed in the AAA Servers table in Network Configuration.



Note For each primary ACS for this secondary ACS, on the primary and secondary ACS, the AAA Servers table entries for the primary ACS must have identical shared secrets.

Step 8 Click **Submit**.

ACS saves the replication configuration, and at the frequency or times that you specified, ACS begins accepting the replicated components from the other ACSs you specified.

Replicating Immediately

You can manually start database replication.



Note Replication cannot occur until you have configured at least one secondary ACS. For more information about configuring a secondary ACS, see [Configuring a Secondary ACS, page 9-11](#).

Before You Begin

Ensure correct configuration of the primary and secondary ACSs. For detailed steps, see [Implementing Primary and Secondary Replication Setups on ACSs, page 9-10](#).

For each secondary ACS that this ACS is to send replicated components to, make sure that you have completed the steps in [Configuring a Secondary ACS, page 9-11](#).

To initiate database replication immediately:

-
- Step 1** Log in to the web interface on the primary ACS.
- Step 2** In the navigation bar, click **System Configuration**.
- Step 3** Click **Database Replication**.



Note If this feature does not appear, choose **Interface Configuration > Advanced Options**. Then, check the **ACS Database Replication** check box. Check the **Distributed System Settings** check box if not already selected.

The Database Replication Setup page appears.

- Step 4** For each ACS internal database component you want to replicate to a secondary ACS, under Replication Components, select the corresponding **Send** check box.
- Step 5** For each secondary ACS that you want the primary ACS to replicate its select components to, select the secondary ACS from the AAA Servers list, and then click --> (right arrow button).



Tip If you want to remove a secondary ACSs from the Replication list, select the secondary ACS in the Replication list, and then click <-- (left arrow button).



Note ACS does not support bidirectional database replication. A secondary ACS receiving replicated components verifies that the primary ACS is not on its Replication list. If not, the secondary ACS accepts the replicated components. If so, it rejects the components.

- Step 6** In the **Replication timeout** text box, specify how long this ACS will perform replication to each of its secondary ACS before terminating the replication attempt and restarting the CSAuth service.

- Step 7** At the bottom of the browser window, click **Replicate Now**.

ACS saves the replication configuration. ACS immediately begins sending replicated database components to the secondary ACSs you specified.



Note Replication only occurs when the database of the primary ACS has changed since the last successful replication. You can force replication to occur by making one change to a user or group profile, such as changing a password or RADIUS attribute.

Scheduling Replication

You can schedule when a primary ACS sends its replicated database components to a secondary ACS. For more information about replication scheduling options, see [Outbound Replication Options, page 9-9](#).


Note

Replication cannot occur until the secondary ACSs are configured properly. For more information, see [Configuring a Secondary ACS, page 9-11](#).

Before You Begin

Ensure correct configuration of the primary and secondary ACSs. For detailed steps, see [Implementing Primary and Secondary Replication Setups on ACSs, page 9-10](#).

For each secondary ACS of this primary ACS, ensure that you have completed the steps in [Configuring a Secondary ACS, page 9-11](#).

To schedule when a primary ACS replicates to its secondary ACSs:

-
- Step 1** Log in to the web interface on the primary ACS.
- Step 2** In the navigation bar, click **System Configuration**.
- Step 3** Click **ACS Database Replication**.


Note

If this feature does not appear, choose **Interface Configuration > click Advanced Options**. Then, check the **ACS Database Replication** check box. Check the **Distributed System Settings** check box if not already selected.

The Database Replication Setup page appears.

- Step 4** To specify which ACS internal database components the primary ACS should send to its secondary ACSs, under Replication Components, select the corresponding **Send** check box for each database component to be sent.
- For more information about replicated database components, see [Replication Components Options, page 9-7](#).
- Step 5** To have the primary ACS send replicated database components to its secondary ACSs at regular intervals, under Replication Scheduling, select the **Every X minutes** option and in the X box type the length of the interval at which ACS should perform replication (up to 7 characters).


Note

Because ACS is momentarily shut down during replication, a short replication interval may cause frequent failover of your AAA clients to other ACSs. If AAA clients are not configured to failover to other ACSs, the brief interruption in authentication service may prevent users from authenticating. For more information, see [Replication Frequency, page 9-5](#).

- Step 6** If you want to schedule times at which the primary ACS sends its replicated database components to its secondary ACSs:
- a. In the Outbound Replication table, select the **At specific times** option.
 - b. In the day and hour graph, click the times at which you want ACS to perform replication.



Tip Clicking times of day on the graph selects those times; clicking again clears them. At any time you can click **Clear All** to clear all hours, or you can click **Set All** to select all hours.

Step 7 If you want to have this ACS send replicated database components immediately on receiving replicated database components from another ACS, select the **Automatically triggered cascade** option.



Note If you specify the Automatically triggered cascade option, you must configure another ACS to act as a primary ACS to this ACS; otherwise, this ACS never replicates to its secondary ACSs.

Step 8 You must specify the secondary ACSs that this ACS should replicate to. To do so:



Note ACS does not support bidirectional database replication. A secondary ACS receiving replicated database components verifies that the primary ACS is not on its Replication list. If not, the secondary ACS accepts the replicated database components. If so, it rejects the components. For more information about replication partners, see [Inbound Replication Options, page 9-10](#).

- a. In the Outbound Replication table, from the AAA Servers list, select the name of a secondary ACS to which you want the primary ACS to send its selected replicated database components.



Note The AAA Servers table in Network Configuration determines which secondary ACSs are available in the AAA Servers list. For more information about the AAA Servers table, see [AAA Server Configuration, page 4-15](#).

- b. Click --> (right arrow button).
The selected secondary ACS moves to the Replication list.
- c. Repeat Step a and Step b for each secondary ACS to which you want the primary ACS to send its selected replicated database components.

Step 9 In the **Replication timeout** text box, specify how long this ACS will perform replication to each of its secondary ACS before terminating the replication attempt and restarting the CSAuth service.

Step 10 Click **Submit**.

ACS saves the replication configuration you created.

Disabling ACS Database Replication

You can disable scheduled ACS database replications without losing the schedule itself. This allows you to cease scheduled replications temporarily and later resume them without having to re-enter the schedule information.

To disable ACS database replication:

Step 1 Log in to the web interface on the primary ACS.

Step 2 In the navigation bar, click **System Configuration**.

- Step 3** Click **Database Replication**.
The Database Replication Setup page appears.
- Step 4** In the Replication Components table, clear all check boxes.
- Step 5** In the Outbound Replication table, select the **Manually** option.
- Step 6** Click **Submit**.
ACS does not permit any replication to or from this ACS server.
-

Configuring Automatic Change Password Replication

To configure automatic change password replication and other local password management options:

- Step 1** Log in to the web interface on the primary ACS.
- Step 2** In the navigation bar, click **System Configuration**.
- Step 3** Click **Local Password Management**.
- Step 4** Select the options to configure:

Field	Description
Password Validation Options	<ul style="list-style-type: none"> • Character length • May not contain username • Different from previous value • Alphanumeric
Remote Change Password	<ul style="list-style-type: none"> • Disable Telnet password on this ACS and return desired message to users Telnet session. • Upon remote password change, immediately propagate the change to selected replication partners.
Password Change Log File Management	<ul style="list-style-type: none"> • Set frequency to generate new password change log file • Set log file deletion based on selected options

- Step 5** Click **Submit**.
-

Database Replication Event Errors

The Database Replication report contains messages indicating errors that occur during replication. For more information about the Database Replication report, see [ACS System Logs, page 11-8](#).



Tip

Brief descriptions of errors are reported to the replication report, however sometimes more detailed errors are written to the CSAuth service log file, auth.log.

RDBMS Synchronization

This section provides information about the RDBMS Synchronization feature, including procedures for implementing this feature, within ACS and the external data source involved.

This section contains the following topics:

- [About RDBMS Synchronization, page 9-17](#)
 - [Users, page 9-18](#)
 - [User Groups, page 9-18](#)
 - [Network Configuration, page 9-19](#)
 - [Custom RADIUS Vendors and VSAs, page 9-19](#)
- [RDBMS Synchronization Components, page 9-19](#)
 - [About CSDBSync, page 9-20](#)
 - [About the accountActions Table, page 9-20](#)
- [ACS Database Recovery Using the accountActions Table, page 9-21](#)
- [Reports and Event \(Error\) Handling, page 9-22](#)
- [Preparing to Use RDBMS Synchronization, page 9-22](#)
- [Configuring a System Data Source Name for RDBMS Synchronization, page 9-23](#)
- [RDBMS Synchronization Options, page 9-24](#)
 - [RDBMS Setup Options, page 9-24](#)
 - [Synchronization Scheduling Options, page 9-25](#)
 - [Synchronization Partners Options, page 9-25](#)
- [Performing RDBMS Synchronization Immediately, page 9-25](#)
- [Scheduling RDBMS Synchronization, page 9-26](#)
- [Disabling Scheduled RDBMS Synchronizations, page 9-28](#)

About RDBMS Synchronization

The RDBMS Synchronization feature enables you to update the ACS internal database with information from an ODBC-compliant data source. The ODBC-compliant data source can be the RDBMS database of a third-party application. It can also be an intermediate file or database that a third-party system updates. Regardless of where the file or database resides, ACS reads the file or database via the ODBC connection. You can also regard RDBMS Synchronization as an API—much of what you can configure for a user, group, or device through the ACS web interface, you can alternatively maintain through this feature. RDBMS Synchronization supports addition, modification, and deletion for all data items it can access.

You can configure synchronization to occur on a regular schedule. You can also perform synchronizations manually, updating the ACS internal database on demand.

Synchronization performed by a single ACS can update the internal databases of other ACSs, so that you only need configure RDBMS Synchronization on one ACS. ACSs listen on TCP port 2000 for synchronization data. RDBMS Synchronization communication between ACSs is encrypted using 128-bit encrypted, proprietary algorithm.

The topics in this section provide an overview of the kinds of configuration that RDBMS Synchronization can automate. You specify the actions in a relational database table or text file named `accountActions`. For more information about `accountActions`, see [About the accountActions Table, page 9-20](#). For specific information about all actions that RDBMS Synchronization can perform, see [Appendix F, “RDBMS Synchronization Import Definitions”](#).

Users

Among the user-related configuration actions that RDBMS Synchronization can perform are:

- Adding users.
- Deleting users.
- Setting passwords.
- Setting user group memberships.
- Setting Max Sessions parameters.
- Setting network usage quota parameters.
- Setting network usage quota parameters.
- Configuring command authorizations.
- Configuring network access restrictions.
- Configuring time-of-day/day-of-week access restrictions.
- Assigning IP addresses.
- Specifying outbound RADIUS attribute values.
- Specifying outbound TACACS+ attribute values.

**Note**

For specific information about all actions that RDBMS Synchronization can perform, see [Appendix F, “RDBMS Synchronization Import Definitions”](#).

User Groups

Among the group-related configuration actions that RDBMS Synchronization can perform are:

- Setting Max Sessions parameters.
- Setting network usage quota parameters.
- Configuring command authorizations.
- Configuring network access restrictions.
- Configuring time-of-day/day-of-week access restrictions.
- Specifying outbound RADIUS attribute values.
- Specifying outbound TACACS+ attribute values.

**Note**

For specific information about all actions that RDBMS Synchronization can perform, see [Appendix F, “RDBMS Synchronization Import Definitions”](#).

Network Configuration

Among the network device-related configuration actions that RDBMS Synchronization can perform are:

- Adding AAA clients.
- Deleting AAA clients.
- Setting AAA client configuration details.
- Adding AAA servers.
- Deleting AAA servers.
- Setting AAA server configuration details.
- Adding and configuring Proxy Distribution Table entries.

**Note**

For specific information about all actions that RDBMS Synchronization can perform, see [Appendix F, “RDBMS Synchronization Import Definitions”](#).

Custom RADIUS Vendors and VSAs

RDBMS Synchronization enables you to configure custom RADIUS vendors and VSAs. In addition to supporting a set of predefined RADIUS vendors and vendor-specific attributes (VSAs), ACS supports RADIUS vendors and VSAs that you define. Vendors you add must be IETF-compliant; therefore, all VSAs that you add must be sub-attributes of IETF RADIUS attribute number 26.

You can define up to ten custom RADIUS vendors. ACS allows only one instance of any given vendor, as defined by the unique vendor IETF ID number and by the vendor name.

**Note**

If you intend to replicate user-defined RADIUS vendor and VSA configurations, user-defined RADIUS vendor and VSA definitions to be replicated must be identical on the primary and secondary ACSs, including the RADIUS vendor slots that the user-defined RADIUS vendors occupy. For more information about database replication, see [ACS Internal Database Replication, page 9-1](#).

For specific information about all actions that RDBMS Synchronization can perform, see [Appendix F, “RDBMS Synchronization Import Definitions”](#).

RDBMS Synchronization Components


The RDBMS Synchronization feature comprises two components:

- **CSDBSync**—A dedicated Windows service that performs automated user and group account management services for ACS.
- **accountActions Table**—The data object that holds information used by CSDBSync to update the ACS internal database.

About CSDBSync

The CSDBSync service uses an ODBC system data source name (DSN) to access the accountActions table. See [Figure 9-2](#). This service looks specifically for a table named **accountActions**. Synchronization events fail if CSDBSync cannot access the accountActions table.

Figure 9-2 RDBMS Synchronization



CSDBSync reads each record from the accountActions table and updates the ACS internal database as specified by the action code in the record. For example, a record could instruct CSDBSync to add a user or change a user password. In a distributed environment, a single ACS, known as the senior synchronization partner, accesses the accountActions table and sends synchronization commands to its synchronization partners. In [Figure 9-2](#), Access Control Server 1 is the senior synchronization partner and the other two ACSs are its synchronization partners.



Note

The senior synchronization partner must have AAA configurations for each ACS that is a synchronization partners. In turn, each of the synchronization partners must have a AAA server configuration for the senior partner. Synchronization commands from the senior partner are ignored if the ACS receiving the synchronization commands does not have a AAA server configuration for the senior partner.

CSDBSync reads and writes (deletes records) in the accountActions table. After CSDBSync processes each record, it deletes the record from the table. This requires that the database user account that you configure the system DSN to use must have read and write privileges.

For more information about CSDBSync or other Windows services used by ACS, see [Chapter 1](#), “Overview”.

About the accountActions Table

The accountActions table contains a set of rows that define actions CSDBSync is to perform in the ACS internal database. Each row in the accountActions table holds user, user group, or AAA client information. Each row also contains an action field and several other fields. These fields provide CSDBSync with the information it needs to update the ACS internal database. For full details of the accountActions table format and available actions, see [Appendix F](#), “RDBMS Synchronization Import Definitions”.

The database containing the `accountActions` table must support a multi-user ODBC driver. This is required to prevent problems if ACS and the third-party system attempt to access the `accountActions` table simultaneously.

ACS includes files to help you create your `accountActions` table for several common formats. You can find these files on the ACS in the following location, assuming a default installation of ACS:

C:\Program Files\CiscoSecure ACS vx.x\CSDBSync\Databases

The Databases directory contains the following subdirectories:

- **Access**—Contains the file `CiscoSecure Transactions.mdb`.

The `CiscoSecure Transactions.mdb` database contains a preconfigured `accountActions` table. When you install ACS, the installation routine creates a system DSN named `CiscoSecure DBSync`. This system DSN is configured to communicate with `CiscoSecure Transactions.mdb`.



Note By default, the username and password for the `CiscoSecure Transactions.mdb` database are set to null. To increase the security of RDBMS synchronizations performed using this database, change the username and password, in the `CiscoSecure Transactions.mdb` database and in ACS. Any other processes that access the `CiscoSecure Transactions.mdb` database should be changed to use the new username and password, too.

- **CSV**—Contains the files `accountactions` and `schema.ini`.

The `accountactions` file is the `accountActions` table in a comma-separated value file. The `schema.ini` file provides the Microsoft ODBC text file driver with the information it needs to access the `accountactions` file.

- **Oracle 7**—Contains the files `accountActions.sql` and `testData.sql`.

The `accountActions.sql` file contains the Oracle 7 SQL procedure needed to generate an `accountActions` table. The `testData.sql` file contains Oracle 7 SQL procedures for updating the `accountActions` table with sample transactions that `CSDBSync` can process.

- **Oracle 8**—Contains the files `accountActions.sql` and `testData.sql`.

The `accountActions.sql` file contains the Oracle 8 SQL procedure needed to generate an `accountActions` table. The `testData.sql` file contains Oracle 8 SQL procedures for updating the `accountActions` table with sample transactions that `CSDBSync` can process.

- **SQL Server 6.5**—Contains the files `accountActions.sql` and `testData.sql`.

The `accountActions.sql` file contains the Microsoft SQL Server 6.5 SQL procedure needed to generate an `accountActions` table. The `testData.sql` file contains Microsoft SQL Server 6.5 SQL procedures for updating the `accountActions` table with sample transactions that `CSDBSync` can process.

ACS Database Recovery Using the `accountActions` Table

Because the RDBMS Synchronization feature deletes each record in the `accountActions` table after processing the record, the `accountActions` table can be considered a transaction queue. The RDBMS Synchronization feature does not maintain a transaction log/audit trail. If a log is required, the external system that adds records to the `accountActions` table must create it. Unless the external system can recreate the entire transaction history in the `accountActions` table, we recommend that you construct a transaction log file for recovery purposes. To do this, create a second table that is stored in a safe location and backed up regularly. In that second table, mirror all the additions and updates to records in the `accountActions` table.

If the database is large, it is not practical to replay all transaction logs to synchronize the ACS internal database with the third-party system. Instead, you should create regular backups of the ACS internal database and replay the transaction logs from the time of most recent backup to resynchronize the ACS internal database with the third-party system. For information on creating backup files, see [ACS Backup, page 8-7](#).

Replaying transaction logs that slightly predate the checkpoint does not damage the ACS internal database, although some transactions might be invalid and reported as errors. As long as the entire transaction log is replayed, the ACS internal database is consistent with the database of the external RDBMS application.

Reports and Event (Error) Handling

The CSDBSync service provides event and error logging. For more information about the RDBMS Synchronization log, see [ACS System Logs, page 11-8](#). For more information about the CSDBSync service log, see [Service Logs, page 11-23](#).

During manual synchronizations, ACS provides visual alerts to notify you of problems that occurred during synchronization.

Preparing to Use RDBMS Synchronization

Synchronizing the ACS internal database by using data from the accountActions table requires that you complete several steps that are external to ACS before you configure the RDBMS Synchronization feature within ACS. If you are planning to use a CSV file as your accountActions table, also see [Configuring a System Data Source Name for RDBMS Synchronization, page 9-23](#).



Note

The `schema.ini` file must be located in the same folder as the `accountactions.csv` file.

To prepare to use RDBMS Synchronization:

-
- Step 1** Determine where you want to create the accountActions table and in what format. For more information about the accountActions table, see [About the accountActions Table, page 9-20](#). For details on the format and content of the accountActions table, see [Appendix F, “RDBMS Synchronization Import Definitions”](#).
 - Step 2** Create your accountActions table.
 - Step 3** Configure your third-party system to generate records and update the accountActions table with them. This effort will most likely involve creating stored procedures that write to the accountActions table at a triggered event; however, the mechanism for maintaining your accountActions table is unique to your implementation. If the third-party system that you are using to update the accountActions table is a commercial product, for assistance, refer to the documentation from your third-party system vendor.
For information about the format and content of the accountActions table, see [Appendix F, “RDBMS Synchronization Import Definitions”](#).
 - Step 4** Validate that your third-party system updates the accountActions table properly. Rows that are generated in the accountActions table must be valid. For details on the format and content of the accountActions table, see [Appendix F, “RDBMS Synchronization Import Definitions”](#).



Note After testing that the third-party system updates the accountActions table properly, discontinue updating the accountActions table until after you have completed [Step 6](#) and [Step 7](#).

- Step 5** If you have a distributed AAA environment and want to synchronize multiple ACSs:
- Determine which ACS you want to use to communicate with the third-party system. This ACS is the senior synchronization partner, which you will later configure to send synchronization data to its synchronization partners, which are the other ACSs needing synchronization.
 - On the senior synchronization partner, verify that there is a AAA server configuration for each synchronization partner. Add a AAA server configuration for each missing synchronization partner. For detailed steps about adding a AAA server, see [Adding AAA Servers, page 4-16](#).
 - On all the other synchronization partners, verify that a AAA server configuration exists for the senior synchronization partner. If no AAA server configuration for the senior synchronization partner exists, create one. For detailed steps about adding a AAA server, see [Adding AAA Servers, page 4-16](#).

Synchronization between the senior synchronization partner and the other synchronization partners is enabled.

- Step 6** Set up a system DSN on the senior synchronization partner (the ACS that will communicate with the third-party system). For steps, see [Configuring a System Data Source Name for RDBMS Synchronization, page 9-23](#).
- Step 7** Schedule RDBMS synchronization on the senior synchronization partner. For steps, see [Scheduling RDBMS Synchronization, page 9-26](#).
- Step 8** Configure your third-party system to begin updating the accountActions table with information that will be imported into the ACS internal database.
- Step 9** Confirm that RDBMS synchronization is operating properly by monitoring the RDBMS Synchronization report in the Reports and Activity section. For more information about the RDBMS Synchronization log, see [ACS System Logs, page 11-8](#).
- Also, monitor the CSDBSync service log. For more information about the CSDBSync service log, see [Service Logs, page 11-23](#).

Configuring a System Data Source Name for RDBMS Synchronization

On the ACS, a system DSN must exist for ACS to access the accountActions table. If you plan to use the `CiscoSecure Transactions.mdb` Microsoft Access database that is provided with ACS, you can use the `CiscoSecure DBSync` system DSN, rather than create one.



Tip

Everything ACS does with ODBC requires System DSNs. User DSNs will not work. Confusing the two DSNs is an easy mistake to make when configuring the datasources in the ODBC control panel applet. Ensure your System DSN is set properly.

For more information about the `CiscoSecure Transactions.mdb` file, see [Preparing to Use RDBMS Synchronization, page 9-22](#).

To create a system DSN for use with RDBMS synchronization:

Step 1 From Windows Control Panel, open the ODBC Data Source Administrator window.



Tip In Windows 2000 and new Microsoft operating systems, the ODBC Data Sources icon is located in the Administrative Tools folder.

Step 2 In the ODBC Data Source Administrator window, click the **System DSN** tab.

Step 3 Click **Add**.

Step 4 Select the driver to use with your new DSN, and then click **Finish**.

A dialog box displays fields requiring information that is specific to the selected ODBC driver.

Step 5 In the **Data Source Name** box, type a descriptive name for the DSN.

Step 6 Complete the other fields that the selected ODBC. These fields may include information such as the IP address of the server on which the ODBC-compliant database runs.

Step 7 Click **OK**.

The name that you assigned to the DSN appears in the System Data Sources list.

Step 8 Close the **ODBC** window and **Windows Control Panel**.

On your ACS, you create the system that ACS uses to access your accountActions table.

RDBMS Synchronization Options

The RDBMS Synchronization Setup page, which is available from System Configuration, provides control of the RDBMS Synchronization feature. It contains three tables whose options are described in this section.

This section contains the following topics:

- [RDBMS Setup Options, page 9-24](#)
- [Synchronization Scheduling Options, page 9-25](#)
- [Synchronization Partners Options, page 9-25](#)

RDBMS Setup Options

The RDBMS Setup table defines how ACS accesses the accountActions table. It contains:

- **Data Source**—Specifies which of all the system DSNs that are available on the ACS is used to access the accountActions table.
- **Username**—Specifies the username that ACS should use to access the database that contains the accountActions table.



Note The database user account that the username specifies must have sufficient privileges to read and write to the accountActions table.

- **Password**—Specifies the password that ACS uses to access the database that contains the accountActions table.

Synchronization Scheduling Options

The Synchronization Scheduling table defines when synchronization occurs. It contains:

- **Manually**—ACS does not perform automatic RDBMS synchronization.
- **Every X minutes**—ACS performs synchronization on a set frequency. The unit of measurement is minutes, with a default update frequency of 60 minutes.
- **At specific times**—ACS performs synchronization at the time that is specified in the day and hour graph. The minimum interval is one hour, and the synchronization occurs on the hour that you selected.

Synchronization Partners Options

The Synchronization Partners table defines which ACSs are synchronized with data from the accountActions table. It provides:

- **AAA Server**—This list represents the AAA servers that are configured in the AAA Servers table in Network Configuration for which the ACS *does not* perform RDBMS synchronization.
- **Synchronize**—This list represents the AAA servers that are configured in the AAA Servers table in Network Configuration for which the ACS *does* perform RDBMS synchronization. The AAA servers on this list are the synchronization partners of this ACS. During synchronization, communication between this ACS and its synchronization partners is 128-bit encrypted with a Cisco-proprietary protocol. The synchronization partners receive synchronization data on TCP port 2000.



Note Each synchronization partner *must* have a AAA server configuration in its Network Configuration section that corresponds to this ACS; otherwise, the synchronization commands this ACS sends to it are ignored.

For more information about the AAA Servers table in Network Configuration, see [AAA Server Configuration, page 4-15](#).

Performing RDBMS Synchronization Immediately

You can manually start an RDBMS synchronization event.

To perform manual RDBMS synchronization:

-
- Step 1** In the navigation bar, click **System Configuration**.
- Step 2** Click **RDBMS Synchronization**.



Note If this feature does not appear, click **Interface Configuration > Advanced Options**, then check **RDBMS Synchronization**.

The RDBMS Synchronization Setup page appears.

Step 3 To specify options in the RDBMS Setup table:



Note For more information about RDBMS setup, see [RDBMS Setup Options, page 9-24](#).

- a. From the Data Source list, select the system DSN that you configured to communicate with the database that contains your accountActions table.
For more information about configuring a system DSN for use with RDBMS Synchronization, see [Configuring a System Data Source Name for RDBMS Synchronization, page 9-23](#).
- b. In the **Username** box, type the username for a database user account that has read-write access to the accountActions table.
- c. In the **Password** box, type the password for the username that was specified in the Step b.

ACS has the information with which to access the accountActions table.



Note You do *not* have to select Manually under Replication Scheduling. For more information, see [Disabling Scheduled RDBMS Synchronizations, page 9-28](#).

Step 4 For each ACS that you want this ACS to update with data from the accountActions table, select the ACS in the AAA Servers list, and then click --> (right arrow button).

The selected ACS appears in the Synchronize list.

Step 5 To remove ACSs from the Synchronize list, select the ACS in the Synchronize list, and then click <-- (left arrow button).

The selected ACS appears in the AAA Servers list.

Step 6 At the bottom of the browser window, click **Synchronize Now**.

ACS immediately begins a synchronization event. To check the status of the synchronization, view the RDBMS Synchronization report in Reports and Activity.

Scheduling RDBMS Synchronization

You can schedule when an ACS performs RDBMS synchronization.

To schedule when an ACS performs RDBMS synchronization:

Step 1 In the navigation bar, click **System Configuration**.

Step 2 Click **RDBMS Synchronization**.



Note If this feature does not appear, click **Interface Configuration > Advanced Options**, then click **RDBMS Synchronization**.

The RDBMS Synchronization Setup page appears.

Step 3 To specify options in the RDBMS Setup table:



Note For more information about RDBMS setup, see [RDBMS Setup Options, page 9-24](#).

- a. From the Data Source list, select the system DSN that you configured to communicate with the database that contains your accountActions table.

For more information about configuring a system DSN for use with RDBMS Synchronization, see [Configuring a System Data Source Name for RDBMS Synchronization, page 9-23](#).

- b. In the **Username** box, type the username for a database user account that has read-write access to the accountActions table.
- c. In the **Password** box, type the password for the username that was specified in the Step b.

Step 4 To have this ACS perform RDBMS synchronization at regular intervals, under Synchronization Scheduling, select the **Every X minutes** option and in the **X** box type the length of the interval in minutes at which ACS should perform synchronization (up to 7 characters).

Step 5 To schedule times at which this ACS performs RDBMS synchronization:

- a. Under Synchronization Scheduling, select the **At specific times** option.
- b. In the day and hour graph, click the times at which you want ACS to perform replication.



Tip Clicking times of day on the graph selects those times; clicking again clears them. At any time you can click **Clear All** to clear all hours, or you can click **Set All** to select all hours.

Step 6 For each ACS that you want to synchronize with data from the accountActions table:



Note For more information about synchronization targets, see [Inbound Replication Options, page 9-10](#).

- a. In the Synchronization Partners table, from the AAA Servers list, select the name of an ACS that you want this ACS to update with data from the accountActions table.



Note The AAA Servers table in Network Configuration determines which ACSs are available in the AAA Servers list, with the addition of the name of the current ACS server. For more information about the AAA Servers table, see [AAA Server Configuration, page 4-15](#).

- b. Click --> (right arrow button).

The selected ACS moves to the Synchronize list.



Note At least one ACS must be in the Synchronize list. This includes the server on which you are configuring RDBMS Synchronization. RDBMS Synchronization does not automatically include the internal database of the current server.

Step 7 Click **Submit**.

ACS saves the RDBMS synchronization schedule that you created.

Disabling Scheduled RDBMS Synchronizations

You can disable scheduled RDBMS synchronization events without losing the schedule itself. You can use this ability to end scheduled synchronizations and resume them later without having to recreate the schedule.

To disable scheduled RDBMS synchronizations:

Step 1 In the navigation bar, click **System Configuration**.

Step 2 Click **RDBMS Synchronization**.

The RDBMS Synchronization Setup page appears.

Step 3 Under Synchronization Scheduling, select the **Manually** option.

Step 4 Click **Submit**.

ACS does not perform scheduled RDBMS synchronizations.

IP Pools Server

This section provides information about the IP Pools feature, including procedures for creating and maintaining IP pools.

This section contains the following topics:

- [About IP Pools Server, page 9-28](#)
- [Allowing Overlapping IP Pools or Forcing Unique Pool Address Ranges, page 9-29](#)
- [Refreshing the AAA Server IP Pools Table, page 9-30](#)
- [Adding a New IP Pool, page 9-30](#)
- [Editing an IP Pool Definition, page 9-31](#)
- [Resetting an IP Pool, page 9-32](#)
- [Deleting an IP Pool, page 9-32](#)

About IP Pools Server

If you are using VPNs you may have to overlap IP address assignments; that is, it may be advantageous for a PPTP tunnel client within a given tunnel to use the same IP address that another PPTP tunnel client in a different tunnel is using. You can use the IP Pools Server feature to assign the same IP address to multiple users, provided that the users are being tunnelled to different home gateways for routing beyond the boundaries of your own network. You can, therefore, conserve your IP address space without having

to resort to using illegal addresses. When you enable this feature, ACS dynamically issues IP addresses from the IP pools that you have defined by number or name. You can configure up to 999 IP pools, for approximately 255,000 users.

If you are using IP pooling and proxy, all accounting packets are proxied so that the ACS that is assigning the IP addresses can confirm whether an IP address is already in use.

**Note**

The CiscoSecure Database Replication feature does not replicate IP pool definitions; however, user and group assignments to IP pools are replicated. By not replicating IP pool definitions, ACS avoids inadvertently assigning an IP address that a replication partner has already assigned to a different workstation. To support IP pools in a AAA environment that uses replication, you must manually configure each secondary ACS to have IP pools with names that are identical to the IP pools that are defined on the primary ACS.

To use IP pools, the AAA client must have network authorization (in IOS, **aaa authorization network**) and accounting (in IOS, **aaa accounting**) enabled.

**Note**

To use the IP Pools feature, you must set up your AAA client to perform authentication and accounting by using the same protocol; TACACS+ or RADIUS.

For information on assigning a group or user to an IP pool, see [Setting IP Address Assignment Method for a User Group, page 6-21](#) or [Assigning a User to a Client IP Address, page 7-7](#).

Allowing Overlapping IP Pools or Forcing Unique Pool Address Ranges

ACS provides automated detection of overlapping pools.

**Note**

To use overlapping pools, you must be using RADIUS with VPN, and you cannot be using the Dynamic Host Configuration Protocol (DHCP).

You can determine whether overlapping IP pools are allowed by checking which button appears below the AAA Server IP Pools table:

- **Allow Overlapping Pool Address Ranges**—Overlapping IP pool address ranges are *not allowed*. Clicking this button allows IP address ranges to overlap between pools.
- **Force Unique Pool Address Range**—Overlapping IP pool address ranges are *allowed*. Clicking this button prevents IP address ranges from overlapping between pools.

To allow overlapping IP pools or to force unique pool address ranges:

Step 1 In the navigation bar, click **System Configuration**.

Step 2 Click **IP Pools Server**.

**Note**

If this feature does not appear, click **Interface Configuration > Advanced Options**, then click **IP Pools**.

The AAA Server IP Pools table lists any IP pools you have configured, their address ranges, and the percentage of pooled addresses in use.

- Step 3** To allow overlapping IP pool address ranges:
- If the **Allow Overlapping Pool Address Ranges** button appears, click it.
ACS allows overlapping IP pool address ranges.
 - If the **Force Unique Pool Address Range** button appears, do nothing.
ACS already allows overlapping IP pool address ranges.
- Step 4** To deny overlapping IP pool address ranges:
- If the **Allow Overlapping Pool Address Ranges** button appears, do nothing.
ACS already does not permit overlapping IP pool address ranges.
 - If the **Force Unique Pool Address Range** button appears, click it.
ACS does not permit overlapping IP pool address ranges.
-

Refreshing the AAA Server IP Pools Table

You can refresh the AAA Server IP Pools table to get the latest usage statistics for your IP pools.

To refresh the AAA Server IP Pools table:

-
- Step 1** In the navigation bar, click **System Configuration**.
- Step 2** Click **IP Pools Server**.
- The AAA Server IP Pools table lists any IP pools that you have configured, their address ranges, and the percentage of pooled addresses in use.
- Step 3** Click **Refresh**.
- ACS updates the percentages of pooled addresses in use.
-

Adding a New IP Pool

You can define up to 999 IP address pools.

To add an IP pool:

-
- Step 1** In the navigation bar, click **System Configuration**.
- Step 2** Click **IP Pools Server**.
- The AAA Server IP Pools table lists any IP pools that you have already configured, their address ranges, and the percentage of pooled addresses in use.
- Step 3** Click **Add Entry**.
- The New Pool table appears.
- Step 4** In the **Name** box, type the name (up to 31 characters) to assign to the new IP pool.
- Step 5** In the **Start Address** box, type the lowest IP address (up to 15 characters) of the range of addresses for the new pool.



Note All addresses in an IP pool must be on the same Class C network; so the first three octets of the start and end addresses must be the same. For example, if the start address is 192.168.1.1, the end address must be between 192.168.1.2 and 192.168.1.254.

Step 6 In the **End Address** box, type the highest IP address (up to 15 characters) of the range of addresses for the new pool.

Step 7 Click **Submit**.

The new IP pool appears in the AAA Server IP Pools table.

Editing an IP Pool Definition

To edit an IP pool definition:

Step 1 In the navigation bar, click **System Configuration**.

Step 2 Click **IP Pools Server**.

The AAA Server IP Pools table lists any IP pools that you have configured, their address ranges, and the percentage of pooled addresses in use.

Step 3 Click the name of the IP pool to edit.

The *name* pool table appears, where *name* is the name of the IP pool that you selected. The In Use field displays how many IP addresses in this pool are allocated to a user. The Available field displays how many IP addresses are unallocated to users.

Step 4 To change the name of the pool, in the **Name** box, type the name (up to 31 characters) to which to change the IP pool.

Step 5 To change the starting address of the pool range of IP addresses, in the **Start Address** box, type the lowest IP address (up to 15 characters) of the new range of addresses for the pool.



Note All addresses in an IP pool must be on the same Class C network, so the first three octets of the start and end addresses must be the same. For example, if the start address is 192.168.1.1, the end address must be between 192.168.1.2 and 192.168.1.254.

Step 6 To change the ending address of the pool range of IP addresses, in the **End Address** box, type the highest IP address (up to 15 characters) of the new range of addresses for the pool.

Step 7 Click **Submit**.

The edited IP pool appears in the AAA Server IP Pools table.

Resetting an IP Pool

The Reset function recovers IP addresses within an IP pool when there are dangling connections. A dangling connection occurs when a user disconnects and ACS does not receive an accounting stop packet from the applicable AAA client. If the Failed Attempts log in Reports and Activity shows a large number of `Failed to Allocate IP Address For User` messages, consider using the Reset function to reclaim all allocated addresses in this IP pool.



Note

Using the Reset function to reclaim all allocated IP addresses in a pool can result in users being assigned addresses that are already in use.

To reset an IP pool and reclaim all its IP addresses:

Step 1 In the navigation bar, click **System Configuration**.

Step 2 Click **IP Pools Server**.

The AAA Server IP Pools table lists any IP pools that you have configured, their address ranges, and the percentage of pooled addresses in use.

Step 3 Click the name of the IP pool to reset.

The *name* pool table appears, where *name* is the name of the IP pool you selected. The In Use field displays how many IP addresses in this pool are assigned to a user. The Available field displays how many IP addresses are not assigned to users.

Step 4 Click **Reset**.

ACS displays a dialog box indicating the possibility of assigning user addresses that are already in use.

Step 5 To continue resetting the IP pool, click **OK**.

The IP pool is reset. All its IP addresses are reclaimed. In the In Use column of the AAA Server IP Pools table, zero percent of the IP pool addresses are assigned to users.

Deleting an IP Pool



Note

If you delete an IP pool that has users assigned to it, those users cannot authenticate until you edit the user profile and change their IP assignment settings. Alternatively, if the users receive their IP assignment based on group membership, you can edit the user group profile and change the IP assignment settings for the group.

To delete an IP pool:

Step 1 In the navigation bar, click **System Configuration**.

Step 2 Click **IP Pools Server**.

The AAA Server IP Pools table lists any IP pools that you have configured, their address ranges, and the percentage of pooled addresses in use.

Step 3 Click the name of the IP pool to delete.

The *name* pool table appears, where *name* is the name of the IP pool you selected. The In Use column displays how many IP addresses in this pool are assigned to a user. The Available column displays how many IP addresses are not assigned to users.

Step 4 Click **Delete**.

ACS displays a dialog box to confirm that you want to delete the IP pool.

Step 5 To delete the IP pool, click **OK**.

The IP pool is deleted. The AAA Server IP Pools table does not list the deleted IP pool.

IP Pools Address Recovery

You use the IP Pools Address Recovery feature to recover assigned IP addresses that have not been used for a specified period of time. You must configure an accounting network on the AAA client for ACS to reclaim the IP addresses correctly.

Enabling IP Pool Address Recovery

To enable IP pool address recovery:

Step 1 In the navigation bar, click **System Configuration**.

Step 2 Click **IP Pools Address Recovery**.



Note If this feature does not appear, click **Interface Configuration > Advanced Options**, then click **IP Pools**.

The IP Address Recovery page appears.

Step 3 Select the **Release address if allocated for longer than X hours** check box and in the *X* box type the number of hours (up to 4 characters) after which ACS should recover assigned, unused IP addresses.

Step 4 Click **Submit**.

ACS implements the IP pools address recovery settings you made.
