CHAPTER 15

# Managing System Operations and Configuration in the Monitoring and Report Viewer

This chapter describes the tasks that you must perform to configure and administer the Monitoring and Report Viewer. The Monitoring Configuration drawer allows you to:

- Manage data—The Monitoring and Report Viewer handles large volumes of data from ACS servers. Over a period of time, the performance and efficiency of the Monitoring and Report Viewer depends on how well you manage the data.

  To do so efficiently, you must back up the data and transfer it to a remote repository on a periodic basis. You can automate this task by scheduling jobs to run periodically. See Configuring Data Purging and Incremental Backup, page 15-3 for more information on data backup.

- View log collections—The Monitoring and Report Viewer collects log and configuration data from ACS servers in your deployment, stores the data in the Monitoring and Report Viewer server, and processes it to generate reports and alarms. You can view the details of the logs collected from any of the servers in your deployment. See Viewing Log Collections, page 15-8 for more information.

- Recovering Log Messages—The Monitoring and Report Viewer recovers the logging entries that are missed during the log collection. The log messages are missed when the Monitoring and Report Viewer server is down or the connectivity between the Monitoring and Report Viewer and ACS server is broken.

  When connectivity is regained, the Monitoring and Report Viewer discovers the entries that were missed, and notifies the ACS server. When the ACS server receives this notification, it resends the entries to the Monitoring and Report Viewer. See Recovering Log Messages, page 15-12 for more information.

- View scheduled jobs—The Monitoring and Report Viewer allows you to schedule tasks that you must perform periodically.

  For example, you can schedule an incremental or full backup to be run at regular intervals. You can use the Scheduler to view the details of these tasks. See Viewing Scheduled Jobs, page 15-13 for more information on the Scheduler.

- View process status—You can view the status of the various processes that run in the Monitoring and Report Viewer. See Viewing Process Status, page 15-14 for more information on the various processes that run in the Monitoring and Report Viewer.

- View data upgrade status—After you upgrade from ACS 5.4 or 5.5 to ACS 5.6 through the CLI, you must ensure that the Monitoring and Report Viewer data upgrade is complete. You can view the Monitoring and Report Viewer data upgrade status through the web interface and switch the

Monitoring and Report Viewer database if upgrade is complete. See Viewing Data Upgrade Status, page 15-15 for more information.

- Configure and edit failure reasons—The Monitoring and Report Viewer allows you to configure the description of the failure reason code and provide instructions to resolve the problem. See Viewing Failure Reasons, page 15-15 for more information on how to edit the failure reason description and instructions for resolution.
- Configure e-mail settings—You can configure the e-mail server and administrator e-mail address. See Specifying E Mail Settings, page 15-16 for more information.
- Configure collection filters—The Monitoring and Report Viewer provides you the option to filter data that is not used for monitoring or troubleshooting purposes. The data that is filtered is not stored in the database and hence saves much needed disk space. See Understanding Collection Filters, page 15-18 for more information on how to configure collection filters.
- Configure system alarms—System alarms notify you of critical conditions encountered during the execution of the ACS Monitoring and Reporting viewer. You can configure if and how you would like to receive notification of system alarms. See Configuring System Alarm Settings, page 15-20 for more information.
- Configure Syslog targets—If you have configured the Monitoring and Report Viewer to send system alarm notifications as Syslog messages, then you must configure a Syslog target to receive the notification. See Configuring Alarm Syslog Targets, page 15-20 for more information.
- Export Monitoring and Report Viewer data—You can configure a remote database, which could either be an Oracle SID or Microsoft AD to which you can export the Monitoring and Report Viewer data.

  You can create and run custom reporting applications using the data in your remote database. See Configuring Remote Database Settings, page 15-20 for more information on how to configure a remote database with the Monitoring and Report Viewer.

ACS provides you the option to schedule jobs in the Monitoring and Report Viewer. By scheduling jobs, you can automate the monitoring tasks to be run at specified intervals. You can view the status of the scheduled jobs, control events, and intervene whenever necessary. You can schedule the following jobs:

- Data Purge
- Backup
- Event notification (system and threshold alarms)
- Export of Monitoring and Report Viewer data to a remote database

This chapter contains the following sections:

- Configuring SNMP Preferences, page 15-18
- Understanding Collection Filters, page 15-18
- Configuring System Alarm Settings, page 15-20
- Configuring Alarm Syslog Targets, page 15-20
- Configuring Remote Database Settings, page 15-20

# Configuring Data Purging and Incremental Backup

The Monitoring and Report Viewer database handles large volumes of data. When the database size becomes too large, it slows down all the processes. You do not need all the data all the time. Therefore, to efficiently manage data and to make good use of the disk space, you must back up your data regularly and purge unwanted data that uses up necessary disk space. Purging data deletes it from the database.

Since the Monitoring and Report Viewer database size is large, the backup process takes a long time to complete. The incremental backup option enables you to take a complete backup of your Monitoring and Report Viewer database once and then to back up data incrementally (that is, only the updates are backed up and stored separately) from the next time onwards.

An incremental backup performs a full database backup the first time it is run, and subsequently only backs up the updates that are made to the database. Incremental backups are therefore much faster and make efficient use of disk space. You can also configure the frequency and time of incremental backups.

With incremental backups, multiple backup files are stored in the repository. However, when you restore data from an incremental backup, ACS restores data from all the backup files starting from the full backup and continuing until the latest incremental backup.

**Note** If you disable incremental backup for some reason, ensure that you run a full backup the next time before you can continue with incremental backups again.

You can also configure a full database backup and define its frequency and time.

ACS also allows you to run an immediate backup of the full Monitoring and Report Viewer database. However, you cannot concurrently run an incremental backup, full backup, and data purge. If any of these jobs are running, you must wait for a period of 90 minutes before you can begin the next job.

**Timesaver** We recommend that you take a full backup the first time and then incrementally back up your data instead of running full backups every time.

**Note** It is highly recommended that you schedule a incremental backup daily and a full backup monthly or weekly. Otherwise the database purge process fails to purge data, which in turn leads to disk space issues. The monthly scheduled backups occur on the last day of the month and the weekly scheduled backups occur on the last day of the week.

**Note** To ensure that your data is backed up before the purge, configure a data repository via the CLI or the ACS web interface (**System Administration > Operations > Software Repositories**). Refer to the *CLI Reference Guide for Cisco Secure Access Control System 5.6* for more information on configuring a repository.

If you enable incremental backup, data is purged daily at 4:00 a.m. at the local time zone where the ACS instance that runs the View process is located.

In ACS 5.6, the view database is allocated based on the opt partition size. ACS View database is 42 percent of opt partition size.

The following database limitations apply for purging:

- If the database disk usage is greater than 60 percent of the allocated view database size, an alarm is sent to the dashboard.
- If the database disk usage is greater than 80 percent of the allocated view database size, a backup is run immediately followed by a purge until the database disk usage is below 60 percent of the allocated view database size. If the backup fails, check the database disk usage again. The Monitoring and Report Viewer data is purged from the database. The oldest data is purged first.
  - – If the database disk usage is greater than 60 percent of the allocated view database size, a backup is run immediately followed by a purge until the database disk usage is below 60 percent of the allocated view database size.
  - – If the backup fails and the database disk usage is greater than 60 percent of the allocated view database size, the Monitoring and Report Viewer decides to wait.

  For example:

  - If you specify that you want to preserve one month of data, and the database size is greater than 100 percent of the allocated view database size within a month, the purge deletes the data on a weekly basis until the database size reaches 80 percent of the allocated view database size.
  - If you specify that you want to preserve more than one month (for example, 5 months of data) but the database size is over 80 percent of the allocated view database size, a purge occurs. If the database size remains over 80 percent of the allocated view database size after the purge, an additional month of data is purged, which results in 4 months of data preserved. Before the purge, the database is backed up.
- If the database size is over 100 percent of the allocated view database size, a purge occurs regardless of whether or not a database backup has occurred. If the database size remains over 80 percent of the allocated view database size, additional purges occur until the database is 80 percent of the allocated view database size.

**Note** If the Incremental backup is configured as ON with no repository configured, database backup will fail and Incremental backup mode will be changed to OFF.

**Note** When incremental backup is disabled, data is purged at the end of every month (Local time).

You can use the Data Purging and Incremental Backup page to:

- Configure purge window size
- Purge data from the database

- Assign a data repository backup location to manage backup (of the purge job)
- Configure incremental and full backup schedules
- Configure immediate backup.

The ACS Database needs to be compressed as a part of maintenance operation. You can run the `acsview-db-compress` command from acs-config mode to reduce the physical size of the view database when there is a difference between the physical size and actual size of the view database. ACS 5.6 stops only the log collector services during compress operation and will be up and running after the compress operation is completed. You need to enable the log recovery feature to recover the log messages that are received during the database compress operation.

In ACS 5.6, database compress operation is automated. You can check the **Enable ACS View Database Compress** check box to compress the ACS View database automatically every day at 5 A.M. The database compress operation is run everyday automatically at 5 A.M whenever there is a need.

**Note** You need to enable the log recovery option to recover the log messages that may be received during the database compress operation. If the log recovery feature is not enabled, then ACS sends an alert message to enable the log recovery feature.

The following database limitations apply for ACS database compress:

- An automatic database compress operation is started the forthcoming day at 5 A.M as soon as the database size is greater than 80 percent of allocated view database size.
- ACS displays an alert message when the difference between the physical and actual size of the view database is greater than 7 percent of the allocated view database size and less than 36 percent of the allocated view database size. Also, an automatic database compress operation is triggered when the size of the database exceeds 80 percent of allocated view database size to avoid disk space issues.
- ACS displays an alert message when the difference between the physical and actual size of the view database is greater than 36 percent of the allocated view database size.
  - If the log recovery feature is not enabled and the ACS view database compress option is enabled, an automatic database compress operation is triggered only after enabling the log recovery feature when the size of the database exceeds 80 percent of allocated view database size to avoid disk space issues.
  - If the log recovery feature and the ACS view database compress option are enabled, an automatic database compress operation is started to avoid disk space issues. The log collector services are shut down during this operation and will be up and running after the compress operation is completed. Since you have log recovery feature enabled already, any log messages that are received during the database compress operation are recovered after the log collector services are up and running.
  - If the log recovery feature and the ACS view database compress options are not enabled, ACS does not trigger any database compress operation. But, if the size of the database exceeds 80 percent of the allocated view database, an automatic database compress operation is triggered only after enabling the log recovery feature to avoid disk space issues.
  - If the log recovery feature is enabled, and the ACS view database compress option is not enabled, an automatic database compress operation is started when the size of the database exceeds 80 percent of allocated view database size limit to avoid disk space issues. The log collector services are shut down during this operation and will be up and running after the compress operation is completed. Since you have log recovery feature enabled already, any log messages that are received during the database compress operation are recovered after the log collector services are up and running.

**Note** It is recommended to perform database compress during the maintenance hours. DB compress may take long time depends on the database size. Database compress should be done after the purge operation gets completed.

From the Monitoring and Report Viewer, select **Monitoring Configuration > System Operations > Data Management > Removal and Backup**.

*Table 15-1 Data Purging and Incremental Backup Page*

| Option | Description |
|---|---|
| **Data Purging** | |
| Data Repository | Use the drop-down list box to select the data repository backup location to be used during data purging.<br>See the *CLI Reference for ACS 5.6* to add a data repository. |
| Maximum Stored Data Period *num* months. | Use the drop-down list box to indicate the number of months, where *num* is the number of months of data you want to retain in the Monitoring and Report Viewer database. |
| Enable ACS View Database Compress | Check the **Enable ACS View Database Compress** check box to compress the ACS View database automatically every day at 5 A.M. |
| **On-Demand Data Purge** | |
| Purge Now | Click **Purge Now** to purge the data. This purge overrides the purge limits that are already set.<br>**Note** It is recommended that you make a full backup before doing an on-demand purge. |
| **View Full Database Backup Now** | |
| Data Repository | Use the drop-down list box to select the data repository backup location to store the full database backup. |
| Backup Now | Click **Backup Now** to start a full Monitoring and Report Viewer database backup. |
| **Incremental Backup** | |
| On | Click the **On** radio button to enable incremental backup. If incremental backup is enabled, the delta is backed up. |
| Off | Click the **Off** radio button to disable incremental backup. |
| **Configure Incremental View Database Backup** | |
| Data Repository | Use the drop-down list box to select a data repository for the backup files. |
| Schedule | Use the drop-down list boxes to select the time of the day when you want the incremental backup to run. |
| Frequency | Use the drop-down list box to choose the frequency at which you want the incremental backup to run. Valid options are:<br>• Daily<br>• Weekly—Typically occurs at the end of every week.<br>• Monthly—Typically occurs at the end of every month. |
| **Configure Full View Database Backup** | |
| Data Repository | Use the drop-down list box to select a data repository to store the backup files. |

*Table 15-1 Data Purging and Incremental Backup Page (continued)*

| Option | Description |
|---|---|
| Schedule | Use the drop-down list boxes to select the time of the day when you want the full View database backup to run. |
| Frequency | Use the drop-down list box to choose the frequency at which you want the full View database backup to run. Valid options are:<br>• Daily<br>• Weekly—Typically occurs at the end of every week.<br>• Monthly—Typically occurs at the end of every month. |

## Configuring NFS Staging

If the utilization of **/opt** exceeds 30 percent, then you are required to use NFS staging with a remote repository to take successful view database backups and generate support bundles. NFS staging uses a Network File System (NFS) share as a staging area of additional disk space during a backup or support bundle request, because these operations are disk space intensive. You can enable NFS staging through ACS CLI using the **backup-staging-url** command. You must provide full permission to NFS directory when you configure the NFS location using the **backup-staging-url** command in ACS 5.6 to perform a successful On Demand Backup. For more information on the **backup-staging-url** command, see the *CLI Reference Guide for Cisco Secure Access Control System 5.6.*

**Note** This section is not applicable to ACS backup operation, as it does not suffer from the same disk space limitations as the View backup and support bundle generation.

**Note** You cannot back up any data when the staging server is down. When the staging server is down, you cannot perform backup and restore operations using any of the configured repositories as they use the same staging server to create the backup file. You have to bring the staging server up or delete the backup staging URL so that the repositories work properly. The backup.tar.gpg file is created under /opt during backup operation when the NFS staging URL is not configured. So, before deleting the backup staging URL, you need to make sure that you have enough space in the /opt location. The backup operation will fail if ACS does not have enough space in /opt location.

**Related Topic**

Restoring Data from a Backup, page 15-7

# Restoring Data from a Backup

Use this page to restore data from the View database that was backed up earlier. You can restore data from an incremental or full backup. If you choose to restore incremental backup data, ACS restores the full View data backup and then the rest of the incremental backups one at a time in the correct sequence.

To restore data from a backup:

**Step 1** Choose **Monitoring Configuration > System Operations > Data Management > Restore**.

The Incremental Backup Restore page appears, displaying the Available Backups to Restore table. Table 15-2 describes the columns in the table.

*Table 15-2 Incremental Backup Restore Page*

| Column | Description |
|---|---|
| Skip View Database backup before Restore | Check this check box to skip the Monitoring and Report Viewer database backup before restoring data from a backup. This option, when checked, hastens the restore process.<br>We recommend that you uncheck this check box because your current data might be lost if a failure occurs during the restore process. |
| Name | Name of the backup file. The backup filename includes the time stamp; for example, ACSViewBackup-20090618_003400.<br>For an incremental backup, click the **Expand** icon to view the associated full and incremental backups. |
| Date | Date on which the backup is run. |
| Repository | Name of the repository that contains the backup file. |
| Type | The type of backup, Incremental or Full. |

**Step 2** Choose a backup file that you want to restore.

**Note** If you choose an incremental backup file to restore, ACS restores all previously associated incremental and full backups. This restore process restores only the Monitoring and Report Viewer data.

**Step 3** Click **Restore** to restore the backup file.

**Related Topic**

Configuring Data Purging and Incremental Backup, page 15-3

# Viewing Log Collections

Use this page to view the recently collected logs from ACS servers.

From the Monitoring and Report Viewer, select **Monitoring Configuration > System Operations > Log Collection**.

**Note** You can use the refresh symbol to refresh the contents of the page.

*Table 15-3 Log Collection Page*

| Option | Description |
|---|---|
| ACS Server | Name of the ACS server. Click to open the Log Collection Details page and view recently collected logs. |
| Last Syslog Message | *Display only.* Indicates the arrival time of the most recent syslog message, in the format *Ddd Mmm dd hh:mm:ss timezone yyyy*, where:<br>• Ddd = Sun, Mon, Tue, Wed, Thu, Fri, Sat.<br>• Mmm = Jan, Feb, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec.<br>• dd = A two-digit numeric representation of the day of the month, from 01 to 31.<br>• hh = A two-digit numeric representation of the hour of the day, from 00 to 23.<br>• mm = A two-digit numeric representation of the minute of the hour, from 00 to 59.<br>• ss = A two-digit numeric representation of the second of the minute, from 00 to 59.<br>• timezone = The time zone. In a distributed environment, the time zone displayed for all secondary servers corresponds to the time zone of the server in which the view is active.<br>If your primary instance has a time zone of PDT and the secondary instance is in UTC, the secondary instance displays the time zone and timestamp of syslog messages with PDT, which corresponds to the time zone of the primary instance.<br>• yyyy = A four-digit representation of the year. |
| Last Error | *Display only.* Indicates the name of the most recent error message. |
| Last Error Time | *Display only.* Indicates the arrival time of the most recent error message, in the format *Ddd Mmm dd hh:mm:ss timezone yyyy*, where:<br>• Ddd = Sun, Mon, Tue, Wed, Thu, Fri, Sat.<br>• Mmm = Jan, Feb, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec.<br>• dd = A two-digit numeric representation of the day of the month, from 01 to 31.<br>• hh = A two-digit numeric representation of the hour of the day, from 00 to 23.<br>• mm = A two-digit numeric representation of the minute of the hour, from 00 to 59.<br>• ss = A two-digit numeric representation of the second of the minute, from 00 to 59.<br>• timezone = The time zone. In a distributed environment, the timezone displayed for all secondary servers corresponds to the timezone of the server in which the view is active.<br>If your primary instance has a timezone of PDT and the secondary instance is in UTC, the secondary instance displays the timezone and timestamp of syslog messages with PDT, which corresponds to the timezone of the primary instance.<br>• yyyy = A four-digit representation of the year. |
| Get Details | Click to view recently collected logs for a selected ACS server. |

**Related Topic**

Log Collection Details Page, page 15-10

## Log Collection Details Page

Use this page to view the recently collected log names for an ACS server.

**Step 1** From the Monitoring and Report Viewer, select **Monitoring Configuration > System Operations > Log Collection**.

**Step 2** Do one of the following:

- Click the name of an ACS server.
- Select the radio button of the ACS server name that you want to use to view recently collected logs, and click **Get Details**.

**Note** You can use the refresh symbol to refresh the contents of the page.

*Table 15-4 Log Collection Details Page*

| Option | Description |
|---|---|
| Log Name | Name of the log file. |
| Last Syslog Message | *Display only*. Indicates the arrival time of the most recent syslog message, in the format *Ddd Mmm dd hh:mm:ss timezone yyyy*, where:<br>• Ddd = Sun, Mon, Tue, Wed, Thu, Fri, Sat.<br>• Mmm = Jan, Feb, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec.<br>• dd = A two-digit numeric representation of the day of the month, from 01 to 31.<br>• hh = A two-digit numeric representation of the hour of the day, from 00 to 23.<br>• mm = A two-digit numeric representation of the minute of the hour, from 00 to 59.<br>• ss = A two-digit numeric representation of the second of the minute, from 00 to 59.<br>• timezone = The time zone. In a distributed environment, the timezone displayed for all secondary servers corresponds to the timezone of the server in which the view is active.<br>If your primary instance has a timezone of PDT and the secondary instance is in UTC, the secondary instance displays the timezone and timestamp of syslog messages with PDT, which corresponds to the timezone of the primary instance.<br>• yyyy = A four-digit representation of the year. |
| Last Error | *Display only*. Indicates the name of the most recent error message. |
| Last Error Time | *Display only*. Indicates the arrival time of the most recent error message, in the format *Ddd Mmm dd hh:mm:ss timezone yyyy*, where:<br>• Ddd = Sun, Mon, Tue, Wed, Thu, Fri, Sat.<br>• Mmm = Jan, Feb, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec.<br>• dd = A two-digit numeric representation of the day of the month, from 01 to 31.<br>• hh = A two-digit numeric representation of the hour of the day, from 00 to 23.<br>• mm = A two-digit numeric representation of the minute of the hour, from 00 to 59.<br>• ss = A two-digit numeric representation of the second of the minute, from 00 to 59.<br>• timezone = The time zone. In a distributed environment, the timezone displayed for all secondary servers corresponds to the timezone of the server in which the view is active.<br>If your primary instance has a timezone of PDT and the secondary instance is in UTC, the secondary instance displays the timezone and timestamp of syslog messages with PDT, which corresponds to the timezone of the primary instance.<br>• yyyy = A four-digit representation of the year. |
| Back | Click to return to the Log Collection page. |
| Refresh | Click to refresh the data in this page. |

**Related Topic**

- Viewing Log Collections, page 15-8

# Recovering Log Messages

ACS server sends syslog messages to the Monitoring and Report Viewer for the activities such as passed authentication, failed attempts, authorization, accounting, and so on.

The syslog messages have a sequence number attached. If the Monitoring and Report Viewer goes down or if it is not able to receive messages from ACS, then the Monitoring and Report Viewer retries those missed logs from ACS, using the logging recovery mechanism.

The Monitoring and Report Viewer processes the syslog messages, and identifies any discrepancies in the sequence. In this way, it finds the messages that have been missed.

The Monitoring and Report Viewer then notifies the ACS server to resend the missing log messages. ACS server processes the messages stored in its local store and resends them to the Monitoring and Report Viewer.

**Note** For the Recovering Log Messages feature to work as desired, you must enable the Log to Local Target option for the relevant logging categories in ACS under **System Administration > Configuration > Log Configuration > Logging Categories > Global**.

To enable Recovering Log Messages, from the Monitoring and Report Viewer, select **Monitoring Configuration > System Operations > Log Message Recovery.**

*Table 15-5 Log Message Recovery Page*

| Option | Description |
|---|---|
| **Log Message Recovery Option** | |
| On | Enable the log message recovery feature. |
| Off | Disable the log message recovery feature. |
| **Configure Log Message Recovery Intervals** | |
| Run Every Minute(s) | Set the duration in minutes, at which the recovery should happen. |
| Run Every Hour(s) | Set the duration in hours, at which the recovery should happen. |
| **Configure Missing Entry count to be re-sent by Collector** | |
| No.of Missing Entries to be re-sent by Collector during recovery at a time | Maximum number of missing entries that can be sent by the ACS server at a time.The default limit is 1000 and the maximum limit is 9999. If you set value higher than this, ACS performance might go down. |

**Note** View logging recovery will not retrieve the missed logs when the View Logging Recovery feature is disabled and the view is down.

# Viewing Scheduled Jobs

Use this page to view the scheduled jobs.

From the Monitoring and Report Viewer, select **Monitoring Configuration > System Operations > Scheduler**.

*Table 15-6 Scheduler Status Page*

| Option | Description |
|---|---|
| Name | *Display only*. Name of the job. |
| Type | *Display only*. Type of associated job; for example, Incremental Backup Utility, Session Termination, DB Aggregation Event, Database Purge Utility, and so on. This list includes both system- and user-defined jobs. |
| Owner | *Display only*. Owner of the associated job—System. |
| Last Run Time | *Display only*. Time of the associated job, in the format *Ddd Mmm dd hh:mm:ss timezone yyyy*, where:<br>• Ddd = Sun, Mon, Tue, Wed, Thu, Fri, Sat.<br>• Mmm = Jan, Feb, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec.<br>• dd = A two-digit numeric representation of the day of the month, from 01 to 31.<br>• hh = A two-digit numeric representation of the hour of the day, from 00 to 23.<br>• mm = A two-digit numeric representation of the minute of the hour, from 00 to 59.<br>• ss = A two-digit numeric representation of the second of the minute, from 00 to 59.<br>• *timezone* = The time zone.<br>• yyyy = A four-digit representation of the year. |
| Last Run Result | *Display only*. The result of the last run of the associated job. |
| Status | *Display only*. The status of the associated job. |

**Note** When you change any schedule through the ACS web interface, for the new schedule to take effect, you must manually restart the Job Manager process. For more information on the CLI command to restart processes, see *CLI Reference Guide for Cisco Secure Access Control System 5.6*.

# Viewing Process Status

Use this page to view the status of processes running in your ACS environment.

From the Monitoring and Report Viewer, select **Monitoring Configuration > System Operations > Process Status**.

**Note** You can click the refresh symbol to refresh the contents of the page.

*Table 15-7 Process Status Page*

| **Option** | **Description** |
|---|---|
| Process Name | *Display only*. Name of the process. Options can be:<br>• Database<br>• Management (ACS management subsystem)<br>• Ntpd<br>• Runtime (ACS runtime subsystem)<br>• View-alertmanager<br>• View-collector<br>• View-database<br>• View-jobmanager<br>• View-logprocessor |
| Status | *Display only*. Indicates the status of the associated process. |
| CPU Utilization | *Display only*. Indicates the CPU utilization of the associated process. |
| Memory Utilization | *Display only*. Indicates the memory utilization of the associated process. |
| Uptime | *Display only*. Indicates the time that the process was started successfully, in the format *Ddd Mmm dd hh:mm:ss timezone yyyy*, where:<br>• Ddd = Sun, Mon, Tue, Wed, Thu, Fri, Sat.<br>• Mmm = Jan, Feb, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec.<br>• dd = A two-digit numeric representation of the day of the month, from 01 to 31.<br>• hh = A two-digit numeric representation of the hour of the day, from 00 to 23.<br>• mm = A two-digit numeric representation of the minute of the hour, from 00 to 59.<br>• ss = A two-digit numeric representation of the second of the minute, from 00 to 59.<br>• *timezone* = The time zone.<br>• yyyy = A four-digit representation of the year. |

# Viewing Data Upgrade Status

After you upgrade to ACS 5.6, ensure that the Monitoring and Report Viewer database upgrade is complete.

You can do this through the ACS web interface. Refer to the *Installation Guide for Cisco Secure Access Control System 5.6* for more information on the upgrade process.

To view the status of Monitoring and Report Viewer data upgrade:

**Step 1** From the Monitoring and Report Viewer, select **Monitoring Configuration > System Operations > Data Upgrade Status**.

**Step 2** The Data Upgrade Status page appears with the following information:

Status—Indicates whether or not the Monitoring and Report Viewer data upgrade is complete.

**Note** It is recommended not to upgrade ACS during aggregation time. If you upgrade ACS during the aggregation time, ACS View upgrade will fail.

# Viewing Failure Reasons

Use this page to view failure reasons.

From the Monitoring and Report Viewer, select **Monitoring Configuration > System Configuration > Failure Reasons Editor**.

Table 15-8 lists the field in the Failure Reasons page.

***Table 15-8 Failure Reasons Page***

| Option | Description |
|---|---|
| Failure Reasons | Description of the possible failure reasons. Click a failure reason name to open the Failure Reasons Editor page. |

**Related Topic**

- Editing Failure Reasons, page 15-15

# Editing Failure Reasons

Use this page to edit failure reasons and include possible resolution steps to assist administrators when they encounter failures.

**Step 1** From the Monitoring and Report Viewer, select **Monitoring Configuration > System Configuration > Failure Reasons Editor**.

**Step 2** Click:

- The name of the failure reason you want to edit.
- The radio button associated with the failure reason you want to edit, then click **Edit**.

The Failure Reason Editor Page appears as described in Table 15-9.

*Table 15-9 Failure Reasons Editor Page*

| Option | Description |
|---|---|
| Failure Reason | Display only. The error code and associated failure reason name. |
| Description | Enter a free text description of the failure reason to assist administrators; use the text tools as needed. |
| Resolution Steps | Enter a free text description of possible resolution steps for the failure reason to assist administrators; use the text tools as needed. |

**Related Topic**

Viewing Failure Reasons, page 15-15

# Specifying E Mail Settings

Use this page to specify the email server and administrator email address.

From the Monitoring and Report Viewer, choose **Monitoring Configuration > System Configuration > Email Settings**.

*Table 15-10 Email Settings Page*

| Option | Description |
|---|---|
| Mail Server | Enter a valid IPv4 or IPv6 email host server. |
| Mail From | Enter the email address name that users will see when they receive email from the system. |

# SNMP Traps

SNMP traps helps you to monitor the status of ACS processes. If you do not have access to an ACS server, but want to monitor the ACS processes, then you can request that the ACS administrator to configure a MIB browser as an SNMP host in the ACS server. After the MIB browser is configured as an SNMP server in ACS, you can monitor the ACS process status from the MIB browser.

ACS 5.4 sends the following generic system traps if you configure the SNMP host from the ACS CLI:

- Cold start—if the device is reloaded.
- Linkup—when Ethernet interface is up.
- Linkdown—when Ethernet interface is down.
- Authentication failure—if the community strings do not match.

In ACS 5.6, this feature is enhanced to send traps for ACS process status to the SNMP manager if you configure an SNMP host from the ACS CLI. ACS 5.6 uses cron job to trigger these traps. After you configure the SNMP host in the ACS CLI, a cron job starts running every minute and monitors the ACS processes. The first time after you configure the SNMP host, you can see that separate traps are received in the SNMP server for each process that is running in ACS, irrespective of its status. The administrator can verify that the configured SNMP server is able to receive the traps that are sent from ACS. After that, the traps are sent from ACS only when there is a change in the ACS process status. You can view the SNMP traps using the traps receiver in a MIB browser.

ACS 5.6 sends traps using the OID of hrSWRunName that belongs to the HOST-RESOURCES MIB and sets the OID value as < ACS PROCESS NAME > - < PROCESS STATUS >.

For instance, runtime - running.

The kron job retrieves the ACS process status from the monit binary. ACS 5.6 supports both SNMPv1 and SNMPv2c.

ACS sends traps for the following status to the configured SNMP server :

- Process Start (monitored state)
- Process Stop (not monitored state)
- Execution Failed
- Does not exists

In the SNMP server, for every object, a unique object ID is generated and a value is assigned to the OID. You can find the object with its OID value in the SNMP server. The OID value for a running trap is "running," and the OID value for not monitored, does not exist, and execution failed traps is "stopped."

To stop ACS from sending SNMP traps to the SNMP server, remove the SNMP configuration from the ACS CLI. This operation stops sending SNMP traps and polling from the SNMP manager.

To configure an SNMP server to receive traps from ACS:

---

**Step 1** Log in to the ACS CLI using the CLI username and password.

**Step 2** Enter **su admin** to enter EXEC mode.

**Step 3** Enter **config t** to enter configuration mode.

**Step 4** Enter the command **snmp-server host <*host_ipaddress*> version <snmp*version*> <*communitystring*>**.

For more information on this command, see the *CLI Reference Guide for Cisco Secure Access Control System.*

**Note** You must configure both the host and the community string to send traps from ACS to a configured SNMP host.

The SNMP server is now configured. The configured SNMP host will receive the traps from ACS.

---

# Configuring SNMP Preferences

You can configure SNMP preferences to authenticate access to MIB objects. The text string that you enter for SNMP preference functions as an embedded password.

To configure SNMP preferences:

**Step 1** From the Monitoring and Report Viewer, choose **Monitoring Configuration > System Configuration > SNMP Settings**.

The SNMP Preferences page appears.

**Step 2** Enter a password in the SNMP V2 Read Community String field to authenticate MIB objects.

**Step 3** Click **Submit**.

# Understanding Collection Filters

You can create collection filters that allow you to filter and drop syslog events that are not used for monitoring or troubleshooting purposes. When you configure collection filters, the Monitoring and Report Viewer does not record these events in the database and thus saves much needed disk space.

**Note** ACS 5.6 supports collecting syslog messages from IPv6 sources.

This section contains the following topics:

- Creating and Editing Collection Filters, page 15-18
- Deleting Collection Filters, page 15-19

## Creating and Editing Collection Filters

Use this page to create or edit collection filters. To do this:

**Step 1** From the Monitoring and Report Viewer, choose **Monitoring Configuration > System Configuration > Collection Filters**.

The Collection Filters page appears.

**Step 2** In the Filters area, do one of the following:

- Click **Create** to create a collection filter.
- Check the check box of the syslog attribute that you want to edit, then click **Edit**.
- Check the check box of the syslog attribute that you want to delete, then click **Delete**.

The Add or Edit Collection Filters page described in Table 15-11 appears.

*Table 15-11 Add or Edit Collection Filters Page*

| Option | Description |
|---|---|
| Syslog Attribute | • In the Add Filter page, choose any one of the following syslog attributes:<br>– NAS IP Address—IPv4 and IPv6 addresses are supported.<br>– Access Service<br>– MAC Address<br>– User<br>• In the Edit Filter page, this field is Display only. |
| Value | Enter the value of the syslog attribute:<br>• NAS IP Address—Enter the IP address of the NAS that you want to filter.<br>• Access Service—Enter the name of the access service that you want to filter.<br>• MAC Address—Enter the MAC address of the machine that you want to filter.<br>• User—Enter the username of the user you want to filter. |

**Step 3** Click **Submit**.

**Related Topics**

- Creating and Editing Collection Filters, page 15-18
- Deleting Collection Filters, page 15-19

# Deleting Collection Filters

To delete a collection filter:

**Step 1** Choose **Monitoring Configuration > System Configuration > Collection Filters**.

The Collection Filters page appears.

**Step 2** Check the check box of the collection filter or filters that you want to delete, then click **Delete**.

The following message appears:

```
Are you sure you want to delete the selected item(s)?
```

**Step 3** Click **Yes**.

The Collection Filters page appears without the deleted collection filter.

# Configuring System Alarm Settings

See Configuring System Alarm Settings, page 12-37 for a description of how to configure system alarm settings.

# Configuring Alarm Syslog Targets

See Understanding Alarm Syslog Targets, page 12-38 for a description of how to configure the syslog targets.

# Configuring Remote Database Settings

Use this page to configure a remote database to which you can export the Monitoring and Report Viewer data. ACS exports data to this remote database at specified intervals. You can schedule the export job to be run once every 1, 2, 4, 6, 8, 12, or 24 hours. You can also schedule the export job to run every 20 or 40 minutes. You can create custom reporting applications that interact with this remote database. ACS supports the following databases:

- Oracle SQL Developer 12c
- Microsoft SQL Server 2008 R2

**Note** ACS does not support remote database with cluster setup.

To configure a remote database:

**Step 1** From the Monitoring and Report Viewer, choose **Monitoring Configuration > System Configuration > Remote Database Settings**.

The Remote Database Settings Page appears as described in Table 15-12.

*Table 15-12 Remote Database Settings Page*

| Option | Description |
|---|---|
| Publish to Remote Database | Check the check box for ACS to export data to the remote database periodically. By default, ACS exports data to the remote database every 4 hours. |
| Server | Enter the IP address of the remote database. |
| Port | Enter the port number of the remote database. The default port for Microsoft database is 1433 and the default port for Oracle database is 1521. To change the port number for Oracle database, see Changing the Port Numbers for Oracle Database, page 15-21. |
| Username | Enter the username for remote database access. |
| Password | Enter the password for remote database access. |

*Table 15-12 Remote Database Settings Page*

| Option | Description |
|---|---|
| Export Every Minutes | Choose a time interval from the drop-down list box for ACS to use to export data. Valid options are 20 and 40 minutes. The default interval is 20 minutes.<br>**Note** If you choose the time interval as 40 minutes, ACS starts the remote database export operation immediately for the first time and it continues to do the operation every 40 minutes from then. |
| Export Every Hours | Choose a time interval from the drop-down list box for ACS to use to export data. Valid options are 1, 2, 4, 6, 8, 12, and 24 hours. The default interval is 4 hours. |
| Database Type | The type of remote database that you want to configure:<br>• Click Microsoft Database radio button to configure a Microsoft database, and enter the name of the remote database.<br>• Click Oracle SID radio button to configure an Oracle database, and enter the Oracle service name for the Oracle database. |
| Download Remote Database schema files | Click this link to download the remote database schema files. The following two schema files are downloaded:<br>• acsview_microsoft_schema.sql<br>• acsview_oracle_schema.sql |

**Step 2** Click **Submit** to configure the remote database.

**Note** Special characters are not supported in remote database names.

**Note** You can view the status of your export job in the Scheduler. See Viewing Scheduled Jobs, page 15-13 for more information.

**Note** If there are two log collector servers that have been configured to export data to a remote database, only one log collector server can export data to the remote database at a time. If a second log collector is pointed to the same remote database, it can cause issues such as over-writing of existing entries in the tables.

## Changing the Port Numbers for Oracle Database

To change the port number for Oracle database, complete the following steps:

**Step 1** Log in to Oracle database.

**Step 2** Open the command prompt.

**Step 3** Run the command **cd** *C:\oraclexe\app\oracle\product\10.2.0\server\BIN*.

**Step 4** Run the command **LSNRCTL status** to find the status of the listener service.

**Step 5** Run the command **LSNRCTL Stop** to stop the listerner service

**Step 6** Go to *C:\oraclexe\app\oracle\product\10.2.0\server\NETWORK\ADMIN* folder and edit the oracle database port numbers in listener.ora and tnsnames.ora files. You should update the same port number in ACS web interface.

**Step 7** Run the command **LSNRCTL Start** to start the listerner service.

**Step 8** Log in to ACS web interface.

**Step 9** From the Monitoring and Report Viewer, choose **Monitoring Configuration > System Configuration > Remote Database Settings** to change the oracle database port number**.**

**Step 10** Enter the new oracle database port number.

ACS displays the following message:

```
This will require view database restart. Are you sure you want to do this?
```

**Step 11** Click **OK**.

For more information, see Configuring Remote Database Settings, page 15-20.