CHAPTER 26

Managing System Data

One of the roles of an administrator is to manage Prime Infrastructure's network data collection and retention so that it:

- Scales to fit the real needs of the system's users.
- Minimizes the burden on monitored devices, applications, and network bandwidth.
- Survives hardware failures .

The following sections explain how to achieve these goals, and how to perform other data management tasks.

Scaling the System

Your Prime Infrastructure system implementation should match the recommendations on appropriate OVA sizes given in the System Requirements section of the *Cisco Prime Infrastructure 1.2 Quick Start Guide*.

Note that the device, interface, and flow record recommendations given in the Quick Start Guide are all maximums, so an OVA of a given size has been tuned to handle up to this number of devices, interfaces and flows per second. Also note that the system requirements for RAM, disk space and processors are all minimums, so you can increase any of these resources and either store more data for a longer period, or process incoming flows more quickly.

As your network grows, you will approach the maximum device/interface/flow rating for your OVA. You will want to check on this from time to time (for details, see and Checking on Device and Interface Usage, page 26-2 and Checking on System Disk Usage, page 26-3).

If you find Prime Infrastructure is managing 80 per cent or more of your system resources or of the counts recommended for the size of OVA you have installed, Cisco recommends that you address this using one or more of the following approaches, as appropriate for your needs:

- **Add disk** —VMWare OVA technology enables you to add disk space to an existing server easily. You will need to shut down the Prime Infrastructure server and then follow the instructions VMWare provides on expanding physical disk space. Once you restart the virtual appliance, Prime Infrastructure will make use of the additional disk space automatically.

- **Limit collection**—Not all data Prime Infrastructure is capable of collecting will be of interest to you. For example: If you are not using the product to report on wireless radio performance statistics, you need not collect or retain that data, and can disable the Radio Performance collection task. Alternatively, you may decide that you need only the aggregated Radio Performance data, and can disable retention of raw performance data. For details on how to do this, see Controlling Background Data Collection Tasks, page 26-3.
- **Shorten retention**—Prime Infrastructure defaults set generous retention periods for all of the data it persists and for the reports it generates. You may find that some of these periods exceed your needs, and that you can reduce them without negative effects. For details on this approach, see Controlling Report Storage and Cleanup, page 26-4 and Controlling Data Retention, page 26-4.
- **Off load backups and reports** —You can save space on the Prime Infrastructure server by saving reports and backups to a remote server. For details, see Setting Up Remote Repositories, page 26-10.
- **Migrate to a new server** —Set up a new server that meets at least the minimum RAM, disk space and processor requirements of the next higher level of OVA. Backup your existing system, then restore it to a VM on the higher-rated server. For details, see Restoring From Backups, page 26-11.

Related Topics

- Controlling Report Storage and Cleanup, page 26-4
- Checking on System Disk Usage, page 26-3
- Controlling Background Data Collection Tasks, page 26-3
- Controlling Report Storage and Cleanup, page 26-4
- Controlling Data Retention, page 26-4

Checking on Device and Interface Usage

You can quickly check on the total number of devices and interfaces Prime Infrastructure is managing using **Administration > Licenses**.

Step 1 Choose **Administration > Licenses**.

Step 2 Under **Licenses,**, Prime Infrastructure displays:

- Device Limit: The maximum number of devices you are licensed to monitor using the product.
- Interface Limit: The maximum number of interfaces you are licensed to monitor.
- Device Count: The actual number of devices being monitored.
- Interface Count: The actual number of interfaces being monitored.
- % Used: The percentage of device licenses actually in use.
- % Used (PAM): The percentage of interface licenses actually in use.

Checking on System Disk Usage

You can quickly check on the total system disk space usage using the **Appliance Status** tab under **Administration > Appliance**.

Step 1 Choose **Administration > Appliance > Appliance Status**.

Under **Disk Usage**, Prime Infrastructure displays the current storage allocation and percentage of use for each of the main disk volumes it uses.

Controlling Background Data Collection Tasks

Prime Infrastructure executes scheduled data collection tasks in the background an a regular basis. You can enable or disable these collection tasks, change the interval at which each task is executed, or change the retention period for the data (raw or aggregated) collected during each execution of each task.

Disabling or limiting these background data collection tasks can have a direct impact on how you use Prime Infrastructure, especially for reporting. To help you consider these impacts, take note of the reports this data is used in, which are listed in the Collection Set Details for each task.

Step 1 Choose **Administration > Background Tasks**.

Step 2 Under **Data Collection Tasks,**, under the **Task** heading in the table, click on the name of the task you want to change.

Step 3 Under **Collection Set Details**, edit the following fields as appropriate:

- Collection Status: Indicates whether the collection task is enabled or not. Check the box to enable it, uncheck it to disable it..
- Interval: The number of minutes between executions of this task.

Step 4 If the collected data is aggregated after collection, edit the following fields as appropriate::

- Non-Aggregation Data Retain Period: The number of days for which the aggregated data is persisted in the database.
- Retain Aggregation Raw Data: Indicates whether the raw data collected (before aggregation) is retained or not. Check the box to enable this, uncheck it to disable it. The raw data will be retained for the same period as the aggregated data.

Step 5 Click Save.

To enable or disable background data collection tasks in bulk:

Step 1 Choose **Administration > Background Tasks**.

Step 2 Under **Data Collection Tasks,**, click the check box next to each Task you want to enable or disable.

Step 3 In the **Go** menu, select **Enable Tasks** or **Disable Tasks.**

Controlling Report Storage and Cleanup

All scheduled reports are stored in the Scheduled Reports Repository. You will want to ensure that scheduled reports are retained in the report repository for reasonable lengths of time only, and deleted on a regular basis.

The default retention scheme is to retain generated reports for a maximum of 31 days. You can customize this retention period following the steps below.

Step 1 Select **Administration > System Settings**.

Step 2 Select **Report**

Step 3 In **Repository Path**, specify the report repository path as needed.

Step 4 Under **File Retain Period**, change the scheduled report retention period as needed

Step 5 Click **Save**.

Related Topics

- Scaling the System, page 26-1
- Controlling Alarm, Event, and Syslog Retention, page 26-5

Controlling Data Retention

In addition to the retention period for background data collection (see Controlling Background Data Collection Tasks, page 26-3), you can also control the retention periods for broad categories of data used in most Prime Infrastructure dashlets and reports, as explained in the following sections:

- Controlling Alarm, Event, and Syslog Retention, page 26-5
- Controlling Health Data Retention, page 26-5
- Controlling Trend Data Retention, page 26-6
- Controlling Performance Data Retention, page 26-6
- Controlling Network Audit Data Retention, page 26-7

Controlling Alarm, Event, and Syslog Retention

As part of managing your system data, you will want to ensure that raw alarm, event and syslog data are retained for reasonable lengths of time only, and deleted on a regular basis.

Data cleanup tasks run nightly to delete old alarms. In addition to the data cleanup task, Prime Infrastructure has an hourly task to check alarm table size. When the alarm table size exceeds 300,000 records, the task deletes the oldest cleared alarms until the alarm table size is within that limit.

The default retention scheme is to retain active alarms, cleared security alarms, events, and syslogs for no more than 30 days. Cleared non-security alarms are retained for 7 days only.

You can customize these retention periods following the steps below.

Step 1 Select **Administration > System Settings**.

Step 2 Select **Alarms and Events**.

Step 3 Under **Alarm and Event Cleanup Options**, change the alarm and event retention periods, as needed.

Step 4 Under **Syslog Cleanup Options**, change the syslog retention period as needed.

Step 5 Click Save.

Related Topics

- Scaling the System, page 26-1

Controlling Health Data Retention

Device Health and System Health data are retained on an hourly, daily and weekly basis.

For Device Health data, the default retention scheme is to retain hourly data for 31 days, daily data for 90 days, and weekly data for 54 weeks.

For System Health data, the default retention scheme is to retain hourly data for 1 day, daily data for 7 days, and weekly data for 54 weeks.

You can customize these retention periods following the steps below.

Step 1 Select **Administration > System Settings**.

Step 2 Select **Data Retention**.

Step 3 Under **Device Health Data Retain Periods**, change the hourly, daily and weekly retention period, as needed.

Step 4 Under **System Health Data Retain Periods**, change the hourly, daily and weekly retention period, as needed.

Step 5 Click **Save**.

Related Topics

- Scaling the System, page 26-1

Controlling Trend Data Retention

Trend data is aggregated and retained on an hourly, daily and weekly basis. The default retention scheme is to retain hourly data for 31 days, daily data for 90 days, and weekly data for 54 weeks.

You can customize the retention period following the steps below.

Step 1 Select **Administration > System Settings**.

Step 2 Select **Data Retention**.

Step 3 Under **Trend Data Retain Periods**, change the hourly, daily and weekly aggregated data retention period, as needed.

Step 4 Click **Save**.

Related Topics

- Scaling the System, page 26-1

Controlling Performance Data Retention

Performance data is retained on a short-, medium- and long-term basis . The default retention scheme is to retain short-term data for 7 days, medium-term data for 31 days, and long-term data for 378 days.

You can customize these retention periods following the steps below.

Step 1 Select **Administration > System Settings**.

Step 2 Select **Data Retention**.

Step 3 Under **Performance Data Retain Periods**, change the short, medium and long-term data retention period, as needed.

Step 4 Click Save.

Related Topics

- Scaling the System, page 26-1

Controlling Network Audit Data Retention

Network audit data is normally deleted after 90 days.

You can customize this retention period following the steps below.

Step 1 Select **Administration > System Settings**.

Step 2 Select **Data Retention**.

Step 3 Under **Network Audit Data Retain Period**, change the retention period, as needed.

Step 4 Click **Save**.

Related Topics

- Scaling the System, page 26-1

Handling Backups

As with any other system upon which your organization relies, you will need to ensure that Prime Infrastructure is backed up regularly, so it can be restored in case of hardware failure.

Backups are always stored in a repository. You may specify remote or local backup repositories.

Backups are saved as encrypted .tar.gpg files in the specified backup repository.

Related Topics

- Running Backups On Demand, page 26-8
- Running Backups From the Command Line, page 26-8
- Scheduling Automatic Backups, page 26-9
- Creating Backup Repositories, page 26-9
- Setting Up Remote Repositories, page 26-10
- Restoring From Backups, page 26-11

Running Backups On Demand

If you want to execute an immediate system backup using the Prime Infrastructure interface, follow the steps below.

You can also run an on-demand backup from the command line (see Running Backups From the Command Line, page 26-8).

Step 1 Choose **Administration > Background Tasks**.

Step 2 Under **Other Background Tasks**, find the **NCS Server Backup** task.

Step 3 If you want to change the backup repository and maximum number of backups, click the **NCS Server Backup** link and adjust these values, then click **Save**.

Step 4 Check the **NCS Server Backup** task checkbox.

Step 5 In the Go menu at the top of the page, select the command **Execute Now**.

Step 6 Click **Go**.

Step 7 Click **Refresh** to see the current status of the task.

Running Backups From the Command Line

If you want to execute an immediate system backup using the command line, follow the steps below. Executing a backup from the command line allows you to specify the backup file name.

You can also run an on-demand backup Prime Infrastructureuser interface (see Running Backups On Demand, page 26-8).

Step 1 At the Prime Infrastructure virtual appliance, exit to the command line.

Step 2 At the command line, log in using the administrator ID and password used to install Prime Infrastructure.

Step 3 Enter the following command to display the list of backups:

```
#show repository repositoryName
```

Where *repositoryName* is the repository alias on which you want to create the backup. (for example: `RemoteFTP)`.

Step 4 Enter the following command to back up the application:

```
#backup filename repository repositoryName application NCS
```

Where:

- *filename* is the name of the backup file (for example: `myBackup)`. The date and time of the backup, as well as the .tar.gpg filename extension, will be appended to the filename you specify.
- *repositoryName* is the name of the repository (for example: `RemoteFTP)`.

Scheduling Automatic Backups

You can schedule regular application backups through the Prime Infrastructureuser interface. This method ensures that time- and processor-intensive backup processes occur at relatively low-traffic periods of the day.

If you want to back up to a new local or remote location, you must first create it. You can create a local backup repository from the Prime Infrastructureuser interface. To create a remote repository, you must use both the interface and the command line. For details on all of these tasks, see Creating Backup Repositories, page 26-9 and Setting Up Remote Repositories, page 26-10.

To schedule automatic backups of the Prime Infrastructure application, follow these steps:

Step 1 Choose **Administration > Background Tasks**.

Step 2 Under **Other Background Tasks**, click **NCS Server Backup**.

Step 3 Complete the fields as follows:

a. **Enabled**: Ensure the box is checked. Uncheck it to disable automatic backups.

b. **Max backups to keep**: Enter the maximum number of backups to keep (the default is 2).

c. **Backup Repository**: Select the backup repository.

d. **Interval**: Enter the number of days between each backup. For example, 1 = a daily backup, 2 = a backup every other day, 7 = a weekly backup, and so on. The default is 7, the minimum is 1, the maximum is 360.

e. **Time of Day**: Enter the time of day when you want the backup to start. Use this format: *hh:mm AM/PM* (for example: `03:00 AM`).

Backing up affects the performance of the server. You should schedule backups to run when the server is less active (for example, in the middle of the night).

Step 4 Click **Save**.

Creating Backup Repositories

You can create new backup repositories as needed. You can then specify this new backup repository when scheduling an automatic backup or before performing an on-demand backup.

If you are only creating a new local backup repository, entering a new alias in the Name field and click Submit will be sufficient to create the new repository under the alias you specified. If the repository you are creating is located on a remote FTP server, you should first set up the server so Prime Infrastructure can write to it. For details, see Setting Up Remote Repositories, page 26-10

Step 1 Choose **Administration > Background Tasks**.

Step 2 Under **Other Background Tasks**, click **NCS Server Backup**.

Step 3 Next to **Backup Repository**, click **Create**.

Step 4 In the **Name** field, enter a unique alias for the new backup repository. This alias is shown in the Backup Repository dropdown list, which you pick from when scheduling an automatic backup, or performing an on-demand backup.

Step 5 If you want the new backup repository to be located on a remote FTP server, complete the following additional fields:

- **Type**: Ensure the **FTP Repository** box is checked
- **FTP Location**: Enter the complete URL of the FTP repository location, including the FTP server address or hostname, and the repository path (for example: `ftp://192.198.110.100/RemoteFTP`).
- **Username**: Enter the name of a user with write privileges on the remote FTP server.
- **Password**: Enter the corresponding password.

Step 6 Click **Submit**.

Related Topics

- Running Backups On Demand, page 26-8
- Scheduling Automatic Backups, page 26-9
- Setting Up Remote Repositories, page 26-10

Setting Up Remote Repositories

Follow the steps below to set up a symbolic link to a remote FTP backup repository so that it can be accessed from the Prime Infrastructure server.

You can locate the FTP server anywhere in your network, as long as it is accessible from the Prime Infrastructure server location. You will need to create a user with write access to the server, and a subdirectory on the server that matches the repository alias you create.

Note Although you are not required to perform this task before creating an FTP repository using the procedure in Creating Backup Repositories, page 26-9, you must do so before the first on-demand or scheduled backup to this repository takes place. If you do not, the backup will fail.

Step 1 At the Prime Infrastructure virtual appliance, exit to the command line.

Step 2 At the command line, log in using the administrator ID and password used to install Prime Infrastructure.

Step 3 Enter the following command to enter server configuration mode:

```
#config t
```

Step 4 Enter the following commands to configure a symbolic link to the remote FTP server:

```
#repository repositoryName
#url ftp://serverIPorHost
#user name password plain userPassword
```

Where:

- *repositoryName* is the repository alias as entered in the Name field when creating the new backup repository in the Prime Infrastructure user interface (for example: `RemoteFTP)`.
- *serverIPorHost* is the IP address or hostname of the remote FTP server ((for example: `ftp://192.198.110.100/`).
- *name* is the name of a user with write privileges to the repository on the FTP server.
- *userPassword* is the corresponding password for that user.

Step 5 When you are finished, press Ctrl+Z to exit configuration mode.

Step 6 To verify creation of the symbolic link, enter the following command:

```
#show repository repositoryName
```

Restoring From Backups

Follow the steps below to restore Prime Infrastructure from a backup using the command line. You cannot restore from the Prime Infrastructure user interface.

You can restore to the same host machine you were using, or to a different host. Remember, however, that licenses are node-locked. If you restore to the same host, you must reapply your license files once the restore is complete. If you restore to a different host, you will need to contact Cisco license support team to re-host your licenses, then apply the re-hosted licenses.

Note that you can restore the backup from a Small OVA implementation to a Large or XLarge OVA. You cannot restore from a larger OVA to a smaller OVA.

Backup files created using the Prime Infrastructure user interface (either on demand or as a scheduled background task) are assigned generic filenames of the format backup-*yymmdd-hhmm*.tar.gpg (for example: `backup-120806-1748.tar.gpg`). Backups created via the command line will have the filename the user specified with the timestamp from the generic format appended to the filename.

Prime Infrastructure server backups are complete application backups, containing all of the application code and all of the data it maintains. However, most machine-specific settings are not included in the backup. If you restore the backup to a different device, you will need to manually re-create these settings. Machine-specific settings include: FTP enable and disable, the FTP port, the FTP root directory, TFTP enable and disable, the TFTP port, the TFTP root directory, HTTP forward enable and disable, the HTTP port, the HTTPS port, the report repository directory, and all high availability settings.

Step 1 At the Prime Infrastructure virtual appliance, exit to the command line.

Step 2 At the command line, log in using the administrator ID and password used to install Prime Infrastructure.

Step 3 Enter the following command to display the list of backups:

```
#show repository repositoryName
```

Where *repositoryName* is the repository alias from which you want to pull the backup. (for example: `RemoteFTP)`.

Step 4 Identify the backup file you want to restore and then enter the following command to restore from that file:

```
#restore filename repository repositoryName application NCS
```

Where:

- *filename* is the name of the backup file (for example: `backup-20120801.tar.gpg)`.
- *repositoryName* is the repository alias as entered in the Name field when creating the new backup repository in the Prime Infrastructure user interface (for example: `RemoteFTP)`.

Enabling Data Deduplication

Data Deduplication allows you to identify authoritative sources for each of the following classes of application data:

- Application Response Time (for TCP applications
- Voice/Video (for RTP applications)

Whenever Prime Infrastructure receives duplicate data about the same network elements and protocols from two or more data sources, it will resolve all such conflicts in the authoritative source's favor.

Step 1 Choose **Administration > System Settings > Data Deduplication**.

Step 2 Click **Enable Deduplication**.