CHAPTER 15

Managing ISE Backup and Restore Operations

This chapter describes the Cisco Identity Services Engine (Cisco ISE) database backup and restore operations, which include Cisco ISE application configuration and Cisco Application Deployment Engine operating system (ADE operating system) configuration. This chapter does not cover the Monitoring and Troubleshooting database backup and restore procedures. For information on the Monitoring and Troubleshooting database backup and restore, see Chapter 24, "Monitoring and Troubleshooting."

Note Backup and restore is not available for Inline Posture nodes in Cisco ISE Release 1.1. For more information on this and other known issues, refer to the *Release Notes for the Cisco Identity Services Engine, Release 1.1.x*.

This chapter contains the following sections:

Overview of ISE Backup and Restore

Cisco ISE allows you to back up data only from the primary or standalone Administration ISE node. Backup can be done either from the Cisco ISE command-line interface (CLI) or Cisco ISE user interface. The restore operation can only be done through the CLI.

Cisco ISE allows you to back up the following data:

- Application-specific configuration data—Contains only Cisco ISE configuration data from the Cisco ISE database

- Application and ADE operating system data—Contains both application-specific and Cisco ADE operating system configuration data

Backup and restore operation can be performed with the backup files of the previous versions of the Cisco ISE and restored on a later version. For example, if you have a backup that is taken from an ISE node (Cisco ISE, Release 1.0) before an upgrade, you can restore it on Cisco ISE, Release 1.1.

Cisco ISE allows you to restore Cisco ISE application and ADE operating system data on a primary or standalone administration node. After you restore data on the primary administration node, the changes are replicated to the secondary nodes in your deployment.

If you obtain the backup from your primary Administration ISE node in one timezone and try to restore it on another ISE node in another timezone, the restore process might fail. This failure happens if the timestamp in the backup file is later than the system time on the ISE node on which the backup is restored. If you restore the same backup a day after it was obtained, then the timestamp in the backup file is in the past and the restore process succeeds.

Note We recommend that you do not change the system timezone after the initial Cisco ISE installation and setup.

Note After you obtain the backup from your standalone ISE node or primary Administration ISE node, if you change the certificate configuration on one or more nodes in your deployment, you must obtain another backup to restore the data. Otherwise, if you try to restore data using the older backup, the communication between the nodes might fail.

Typically, you would need the application-specific backup to be scheduled more frequently, and the whole system backup infrequently. The whole system backup is required in case of a hardware failure that requires you to reimage your hardware.

You need a data repository, which is the location where Cisco ISE saves your backup file. You must create a repository before you can run an on-demand or scheduled backup.

Note If you have a standalone administration node that fails, then you must run the full system backup to restore it. If your primary Administration ISE node fails, you can use the distributed setup to promote your secondary Administration ISE node to become the primary, and restore data on your primary Administration ISE node after it comes up.

You can perform a backup either through the CLI or through the Cisco ISE user interface.

Refer to the *Cisco Identity Services Engine CLI Reference Guide, Release 1.1.x* for more information on the CLI backup commands.

Note Cisco ISE also provides another CLI command, **backup-logs**, that you can use to collect log and configuration files for troubleshooting purposes. For more information, refer to the *Cisco Identity Services Engine CLI Reference Guide, Release 1.1.x*.

Supported Scenarios for Backup, Restore, and Upgrade

For details on supported approaches to a previous backup on a newer build and upgrade scenarios, refer to the “Upgrading Cisco ISE” chapter of the *Cisco Identity Services Engine Upgrade Guide, Release1.1.x*.

Configuring Repositories

Cisco ISE allows you to create and delete repositories through the Cisco ISE user interface. You can use these repositories for various operations such as backup, restore, and so on. You can create the following types of repositories:

- DISK
- FTP
- SFTP
- TFTP
- NFS
- CDROM
- HTTP
- HTTPS

The Repositories page allows you to manage repositories from the Cisco ISE administrative user interface. You can create, and delete repositories through the administrative user interface.

Note We recommend that you have a repository size of 10 GB for small deployments (100 endpoints or less), 100 GB for medium deployments, and 200 GB for large deployments.

This section contains the following topics:

- Creating Repositories
- Deleting Repositories

Creating Repositories

Prerequisite:

Every Cisco ISE administrator account is assigned one or more administrative roles. To perform the operations described in the following procedure, you must have one of the following roles assigned: Super Admin or System Admin. See Cisco ISE Admin Group Roles and Responsibilities for more information on the various administrative roles and the privileges associated with each of them.

To create a repository, complete the following steps:

Step 1 Choose **Administration > System > Maintenance**.

Step 2 From the Operations navigation pane on the left, click **Repository**.

The Repository List page appears with a list of configured repositories. This page will be blank when you create repositories for the first time.

Step 3 Click **Add** to add a new repository.

The Repository Configuration page appears.

Step 4 Enter the values as described:

- Repository—(Required) Name of the repository. Alphanumeric characters are allowed and the maximum length is 80 characters.

Note You cannot edit the name of a repository.

- Protocol—(Required) From the drop-down list, choose one of the protocols.
- Path—(Required) Enter the path to your repository in this field. This value must start with a forward slash (/).

  The path must be valid and must exist at the time you create the repository. The following three fields are required depending on the protocol that you have chosen.

  - ServerName—(Required for TFTP, HTTP, HTTPS, FTP, SFTP, and NFS) Enter the hostname or IPv4 address of the server where you want to create the repository.
  - Username—(Required for FTP, SFTP, and NFS) Enter the username that has write permission to the specified server. Only alphanumeric characters are allowed.
  - Password—(Required for FTP, SFTP, and NFS) Enter the password that will be used to access the specified server. Passwords can consist of the following characters: 0 through 9, a through z, A through Z, -, ., |, @, #,$, %, ^, &, *, (,), +, and =.

Step 5 Click **Submit** to create the repository.

A message similar to the following one appears:

Repository is created successfully.

Step 6 Click **Repository** in the Operations navigation pane on the left or click the **Repository List** link at the top of this page to go to the repository listing page.

Next Steps:

1. Ensure that the repository that you created is working by executing the following command from the Cisco ISE command-line interface:

   show repository *repository_name*

   where *repository_name* is the name of the repository that you have created. For more information, see the *Cisco Identity Services Engine CLI Reference Guide, Release 1.1.x*.

Note If the path that you provided while creating the repository does not exist, then you will get the following error: %Invalid Directory.

2. Run an on-demand backup or schedule a backup. See Running On-Demand Backup and Scheduling a Backup for more information.

Deleting Repositories

Prerequisite:

Every Cisco ISE administrator account is assigned one or more administrative roles. To perform the operations described in the following procedure, you must have one of the following roles assigned: Super Admin or System Admin. See Cisco ISE Admin Group Roles and Responsibilities for more information on the various administrative roles and the privileges associated with each of them.

To delete a repository, complete the following steps:

Step 1 Choose **Administration > System > Maintenance**.

Step 2 From the Operations navigation pane on the left, click **Repository**.

The repositories listing page appears.

Step 3 Click the radio button next to the repository that you want to delete, then click **Delete**.

Cisco ISE prompts you with the following message:

Are you sure you want to delete this repository?

Step 4 Click **OK** to delete the repository.

The following message appears:

Repository was deleted successfully.

The Repository List page appears and the repository that you deleted will no longer be listed in this page.

On-Demand Backup

Cisco ISE provides an option to obtain an on-demand backup of the primary administration node. You can obtain a backup of the Cisco ISE application-specific configuration data, or application and Cisco ADE operating system data.

Running On-Demand Backup

Prerequisites:

1. Before you perform this task, you should have a basic understanding of the Backup and Restore operations in Cisco ISE.
2. Ensure that you have configured repositories. See the "Configuring Repositories" section on page 15 -3 for more information.
3. Every Cisco ISE administrator account is assigned one or more administrative roles. To perform the operations described in the following procedure, you must have one of the following roles assigned: Super Admin or System Admin. See Cisco ISE Admin Group Roles and Responsibilities for more information on the various administrative roles and the privileges associated with each of them.

Note For backup and restore operations, you cannot choose the CDROM, HTTP, or HTTPS options because these are read-only repositories.

To perform an on-demand backup, complete the following steps:

Step 1 Choose **Administration > System > Maintenance**.

Step 2 From the Operations navigation pane on the left, choose **Data Management > Administration Node > Full Backup On Demand**.

The Backup On Demand page appears.

Step 3 Enter the name of your backup file.

Step 4 Select the repository where your backup file should be saved.

You cannot enter a repository name here. You can only choose an available repository from the drop-down list. Ensure that you create the repository before you run a backup.

Step 5 Check the **Application-Only Backup, Excludes OS System Data** check box to obtain a Cisco ISE application data backup. Uncheck this check box if you want the Cisco ADE operating system data as well.

Step 6 Enter the **Encryption Key**. This key is used to encrypt and decrypt the backup file.

Step 7 Click **Backup Now** to run your backup.

Note In a distributed deployment, do not change the role of a node or promote a node when the backup is running. Changing node roles will shut down all the processes and might cause some inconsistency in data if backup is running concurrently. Wait for the backup to complete before you make any node role changes.

Step 8 Your page is refreshed and the following message appears in the lower right corner of the page, if you are viewing the Backup On Demand page:

Backup is done successfully.

If you have moved to other pages in the Cisco ISE user interface, to check the status of your backup, you must go to the Backup History page. See the "Viewing Backup History" section on page 15 -10 for more information.

Cisco ISE appends the backup filename with the timestamp and stores this file in the specified repository. Check if your backup file exists in the repository that you have specified.

For more information:

This procedure backs up your Cisco ISE application and Cisco ADE operating system data. To back up Monitoring and Troubleshooting database data, see the "Backing Up and Restoring the Monitoring Database" section on page 24 -49. You can also schedule backup jobs that runs periodically. See the "Scheduled Backups" section on page 15 -6 for more information.

Scheduled Backups

Cisco ISE allows you to schedule your system-level backup operations. You can schedule a backup to be run periodically (daily, weekly, monthly), and specify the time of the day when the backup should be run. Backup operations usually take some amount of time and the scheduling option allows you to configure backups at a convenient time. The Scheduled Backup page lists the backups that you have scheduled.

You can schedule a backup from the Cisco ISE CLI or through the Cisco ISE user interface. To schedule a job from the CLI, you must use the **kron** CLI command.

Refer to the *Cisco Identity Services Engine CLI Reference Guide, Release 1.1.x* for more information on the **kron** command.

The following is an example of the **kron policy-list** *policylistname* command:

```
ise/admin(config)# kron policy-list policylistname
ise/admin(config-Policy List)# cli backup backupfilename repository repositoryname
application ise
ise/admin(config-Policy List)# kron occurrence backup_occur_backupfilename
ise/admin(config-Occurrence)# at 10:00 Sunday
ise/admin(config-Occurrence)# recurring
ise/admin(config-Occurrence)# policy-list policylistname
ise/admin(config-Occurrence)# exit
ise/admin(config)# exit
ise/admin#
```

To create a kron job, you must define a policy list. This policy list will also be created when you schedule a backup through the Cisco ISE user interface.

Note If you promote your secondary Administration ISE node to become the primary Administration ISE node, you must reconfigure your scheduled backups on the new primary Administration ISE node because scheduled backup configurations are not replicated from the primary to secondary Administration ISE nodes.

Note After you upgrade from Cisco ISE Release 1.0.3.377 or Cisco ISE Maintenance Release 1.0.4.573 to Cisco ISE, Release 1.1, the scheduled backup jobs need to be recreated, as the older jobs will not work properly.

Scheduling a Backup

Prerequisites:

1. Before you perform this task, you should have a basic understanding of the Backup and Restore, On-Demand Backup, and Scheduled Backups operations in Cisco ISE.
2. Ensure that you have configured repositories. See the "Configuring Repositories" section on page 15 -3 for more information.
3. Every Cisco ISE administrator account is assigned one or more administrative roles. To perform the operations described in the following procedure, you must have one of the following roles assigned: Super Admin or System Admin. See Cisco ISE Admin Group Roles and Responsibilities for more information on the various administrative roles and the privileges associated with each of them.

Note For backup and restore operations, you cannot choose the CDROM, HTTP, or HTTPS options because these are read-only repositories.

To schedule a backup from the Cisco ISE user interface, complete the following steps:

Step 1 Choose **Administration > System > Maintenance**.

Step 2 From the Operations navigation pane on the left, choose **Data Management > Administration Node > Scheduled Backup**.

The Scheduled Backup List page appears. This page provides the following information:

- Name—Name of the scheduled backup job.
- Type—The frequency of recurrence, whether it is daily, weekly, or monthly.
- Time:Date—The time at which the backup will be run, the day of the week if the schedule is weekly, and the date if the schedule is monthly.
- Policy—Name of the policy list.
- Recurring—Indicates whether the backup should be repeated at the specified date and time or just performed once.

Step 3 Click **Add** to add a scheduled backup.

The Scheduled Backup Configuration page appears as shown in Figure 15-1.

Figure 15-1 Scheduled Backup: Create Page

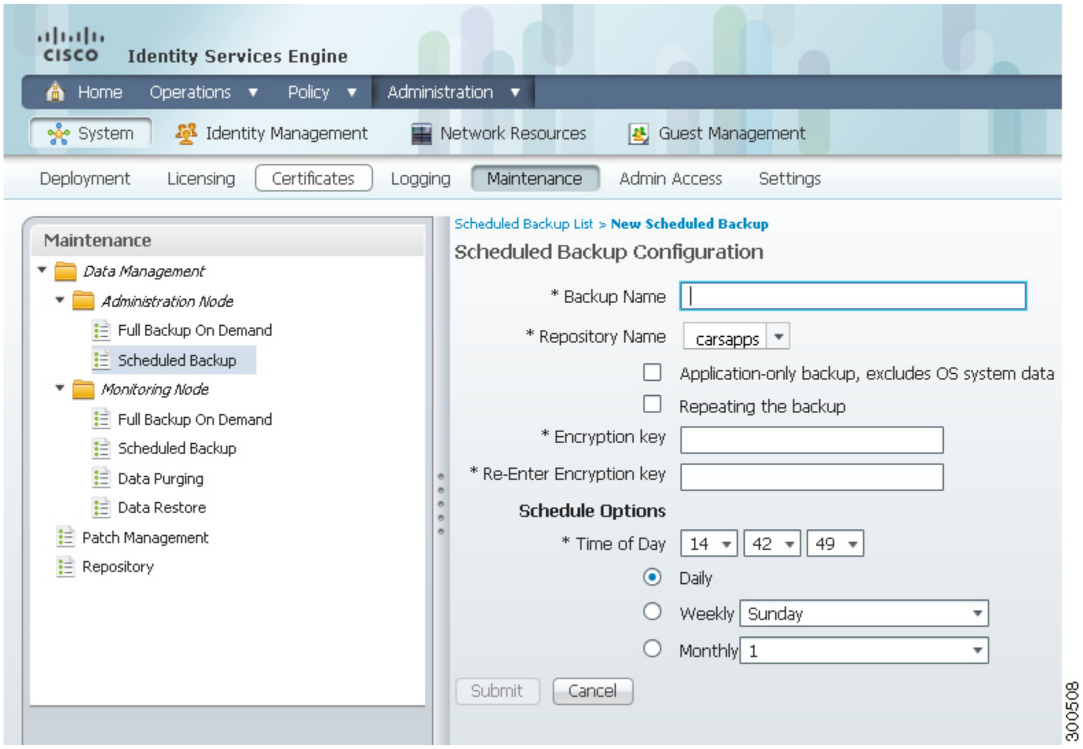

Scheduled Backup Configuration

* Backup Name

* Repository Name carsapps

□ Application-only backup, excludes OS system data

□ Repeating the backup

* Encryption key

* Re-Enter Encryption key

Schedule Options

* Time of Day 14 42 49

● Daily

○ Weekly Sunday

○ Monthly 1

Step 4 Enter a name for your backup file.

You can enter a descriptive name of your choice. Cisco ISE appends the timestamp to the backup filename and stores it in the repository. You will have unique backup filenames even if you configure a series of backups.

Note On the Scheduled Backup list page, the backup filename will be prepended with "backup_occur" to indicate that the file is a **kron** occurrence job.

Step 5 Choose a repository from the Repository Name drop-down list.

You cannot enter a repository name. You have to create a repository from the Cisco ISE user interface or through the Cisco ISE CLI. See the "Configuring Repositories" section on page 15 -3 for information on how to create repositories. Ensure that you create a repository before you schedule a backup job.

Step 6 Check the **Application-Only Backup, Excludes OS System Data** check box to back up only the Cisco ISE application data. Uncheck this check box if you want to include the Cisco ADE operating system data in the backup as well.

Step 7 Check the **Repeating the Backup** check box if you want the scheduled backup to recur at the specified date and time. Uncheck this check box if you are scheduling the backup to be run only once.

Step 8 Enter the **Encryption Key**. This key is used to encrypt and decrypt the backup file.

Step 9 In the Schedule Options group box:

- Choose the time of the day when you want the backup to run.
- Choose any one of the following:
  - Daily—If you want the backup to be run at a specified time every day.
  - Weekly—Choose the day of the week from the drop-down list for the backup to be run on the specified day and time every week.
  - Monthly—Choose any date of the month (from 1 to 28) on which the backup will be run at the specified time.

Step 10 Click **Submit** to schedule the backup.

Click the **Scheduled Backup List** link at the top of this page to return to the Scheduled Backup Listing page.

For more information:

The scheduled backup will be listed in the Scheduled Backup page. To see the status of your previously scheduled jobs, see the "Viewing Backup History" section on page 15 -10. This procedure schedules a backup job that backs up the Cisco ISE application and the Cisco ADE operating system data. To schedule a Monitoring and Troubleshooting database backup job, see the "Backing Up and Restoring the Monitoring Database" section on page 24 -49.

Deleting a Scheduled Backup

Cisco ISE allows you to delete an existing backup schedule and create a new schedule. There is no option to edit a scheduled backup job in Cisco ISE.

Prerequisite:

Every Cisco ISE administrator account is assigned one or more administrative roles. To perform the operations described in the following procedure, you must have one of the following roles assigned: Super Admin or System Admin. See Cisco ISE Admin Group Roles and Responsibilities for more information on the various administrative roles and the privileges associated with each of them.

To delete a scheduled backup job, complete the following steps:

Step 1 Choose **Administration > System > Maintenance**.

Step 2 From the Operations navigation pane on the left, choose **Data Management > Administration Node > Scheduled Backup**.

The Scheduled Backup List page appears with a list of scheduled jobs.

Step 3 Click the radio button next to the scheduled backup job that you want to delete, and click **Delete**.

Step 4 The following message appears:

Are you sure you want to delete this scheduled backup?

Step 5 Click **OK** to delete the scheduled backup.

Viewing Backup History

For scheduled backups, you can obtain information about the backup, backup events, and status (when the backup was performed, whether it was successful or not, and so on) from the Backup History page.

Prerequisite:

Every Cisco ISE administrator account is assigned one or more administrative roles. To perform the operations described in the following procedure, you must have one of the following roles assigned: Super Admin or Monitoring Admin or Helpdesk Admin. See Cisco ISE Admin Group Roles and Responsibilities for more information on the various administrative roles and the privileges associated with each of them.

To view the backup history, complete the following steps:

Step 1 Choose **Operations > Reports > System**.

Step 2 From the System navigation pane on the left, choose **Data Management > Administration Node > Backup History**.

The Backup History page appears with information about all backups that were run on the Cisco ISE node as shown in Figure 15-2.

Figure 15-2 Backup History Page

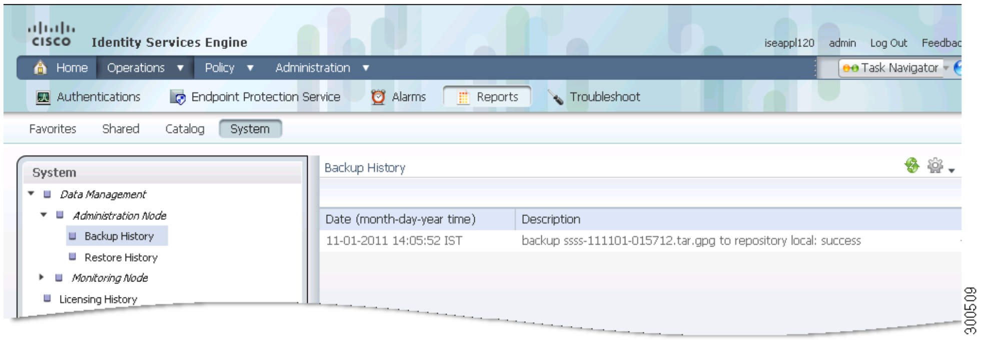

Date (month-day-year time)	Description
11-01-2011 14:05:52 IST	backup ssss-111101-015712.tar.gpg to repository local: success

The Backup History page provides basic information about the scheduled backups that were run. For failed backups, you must run the **backup-logs** command from the Cisco ISE CLI and look at the ADE.log for more information.

Note The backup history is stored along with the Cisco ADE operating system configuration data. After an application upgrade, backup history is not lost and the Backup History page lists all the backups that were run. The backup history will be removed only when you reimage the primary administration node.

Restoring Data from a Backup

You can restore data only through the Cisco ISE CLI.

- To restore the application data, from the Cisco ISE CLI, enter the following command:

  restore *backupfilename.tar.gpg* **repository** *repositoryname* **application** *application name* **encryption-key hash | plain** *encryption-key name*

- To restore the application and Cisco ADE operating system data, from the Cisco ISE CLI, enter the following command:

  restore *backupfilename.tar.gpg* **repository** *repositoryname* **encryption-key hash | plain** *encryption-key name*

where

- *backupfilename.tar.gpg* is the name of the backup file that you want to restore
- *repositoryname* is the repository that contains your backup file
- *encryption-key name* is the key that was used while creating the backup file. Encryption-key is optional while restoring data. To support restoring earlier backups where you have not provided encryption-keys, you can use the restore command without the encryption-key.

After you restore data, you must wait until all the application server processes are up and running. To verify if the Cisco ISE application server processes are running, enter the following command from the Cisco ISE CLI:

show application status ise

For more information, refer to the *Cisco Identity Services Engine CLI Reference Guide, Release 1.1.x*.

Note You can restore data only on the same version of Cisco ISE. If your Cisco ISE database backup was obtained from Cisco ISE Release 1.0 with patches 1, 2, and 3 installed, then you can only restore it on a Cisco ISE node that has Release 1.0 and patch 3 (highest of the patches) installed.

To check for the status of your restore job, see the "Viewing Restore History" section on page 15 -12.

Note If the sync status and replication status after application restore for any secondary node is *Out of Sync*, you have to reimport the certificate of that secondary node to the primary administration node and perform a manual synchronization. See Synchronizing Primary and Secondary Nodes in a Distributed Environment, page 15-12 for the procedure to perform manual synchronization.

Viewing Restore History

You can obtain information about all restore operations, restore log events, and statuses (when the restore was done, whether it was successful or not, and so on) from the Restore History page.

Prerequisite:

Every Cisco ISE administrator account is assigned one or more administrative roles. To perform the operations described in the following procedure, you must have one of the following roles assigned: Super Admin or Monitoring Admin or Helpdesk Admin. See Cisco ISE Admin Group Roles and Responsibilities for more information on the various administrative roles and the privileges associated with each of them.

To view the restore history, complete the following steps:

Step 1 Choose **Operations > Reports > System**.

Step 2 From the System navigation pane on the left, choose **Data Management > Administration Node > Restore History**.

The Restore History page appears with information about all the restore operations that were performed on the Cisco ISE node.

Note Similar to the Backup History page, the Restore History page provides basic information on the restore job. For troubleshooting information, you have to run the **backup-logs** command from the Cisco ISE CLI and look at the ADE.log file.

Synchronizing Primary and Secondary Nodes in a Distributed Environment

In a distributed environment, after restoring a backup file on your primary administration node, sometimes the Cisco ISE database in the primary and secondary nodes are not synchronized automatically. At such times, you can manually force a full replication from the primary administration node to your secondary ISE nodes. You can force a synchronization only from a primary to secondary nodes. During the sync-up operation, you cannot make any configuration changes. Once a sync-up operation starts, a progress bar appears displaying the progress of the forced replication. Cisco ISE allows you to navigate to other Cisco ISE user interface pages and make any configuration changes only after the synchronization is complete.

Prerequisite:

Every Cisco ISE administrator account is assigned one or more administrative roles. To perform the operations described in the following procedure, you must have one of the following roles assigned: Super Admin or System Admin. See Cisco ISE Admin Group Roles and Responsibilities for more information on the various administrative roles and the privileges associated with each of them.

To synchronize your secondary Cisco ISE nodes with your primary Cisco ISE node, complete the following steps:

Step 1 Choose **Administration > System > Deployment**.

Step 2 From the Deployment navigation pane on the left, click **Deployment**.

The Deployment Nodes page appears.

Step 3 Check the check boxes next to the secondary ISE nodes whose Replication Status is *Out of Sync*.

Step 4 Click **Syncup**.

The nodes are synchronized with the primary administration node. You will have to wait until this process is complete before you can access the Cisco ISE user interface again.

Result

When all the nodes are synchronized, the following message appears:

Sync up is done for all the nodes.

An error message appears if Cisco ISE cannot force a full replication.

Recovering Lost Nodes in Standalone and Distributed Deployments

This section provides troubleshooting information that you can use to recover lost nodes in standalone and multinode deployments. Some of the following use cases use the backup and restore functionality and others use the replication feature to recover lost data:

- Loss of All Nodes in a Distributed Setup, Recovery Using Existing IP Addresses and Hostnames, page 15-13
- Loss of All Nodes in a Distributed Deployment, Recovery Using New IP Addresses and Hostnames, page 15-14
- Standalone Deployment, Recovery Using Existing IP Address and Hostname, page 15-15
- Standalone Deployment, Recovery Using New IP Address and Hostname, page 15-15
- Configuration Rollback, page 15-16
- Primary Node Failure in a Distributed Deployment, page 15-16
- Secondary Node Failure in a Distributed Deployment, page 15-16

Loss of All Nodes in a Distributed Setup, Recovery Using Existing IP Addresses and Hostnames

In a distributed deployment setup, there is a natural disaster leading to the loss of all the nodes. After recovery, you want to use the existing addresses and hostnames.

Scenario

You have two nodes: N1 (primary Administration node) and N2 (secondary Administration node) and a backup of the N1 node is available that was taken at time t1. Later, both N1 and N2 nodes fail because of a natural disaster.

Assumption

All Cisco ISE nodes in the deployment were destroyed. The new hardware was imaged using the same hostnames and IP addresses.

Resolution Steps

1. You have to replace both N1 and N2 nodes. See "Replacing the ISE Appliance Hardware" section on page 9 -28 for more information. N1 and N2 nodes will now have a standalone configuration.
2. You must then restore the backup on the replaced N1 node. See "Restoring Data from a Backup" section on page 15 -11 for more information. The restore script will try to sync the data on N2, but N2 is now a standalone node and the sync will fail. Data on N1 will be reset to time t1.
3. You must log in to the N1 user interface to delete and reregister the N2 node. See the following for more information:
   - "Removing a Node from Deployment" section on page 9 -26
   - "Registering and Configuring a Secondary Node" section on page 9 -13

   Both the N1 and N2 nodes will now have data reset to time t1.

Loss of All Nodes in a Distributed Deployment, Recovery Using New IP Addresses and Hostnames

In a distributed setup, all the nodes in the deployment are destroyed because of a natural disaster. The new hardware is reimaged at a new location and requires new IP addresses and hostnames.

Scenario

You have two ISE nodes: N1 (primary Administration node) and N2 (secondary Policy Service node) and a backup of N1 node is available that was taken at time t1. Later, both N1 and N2 nodes fail because of a natural disaster. ISE nodes are replaced at a new location and the new hostnames are N1A (primary Administration node) and N2A (secondary Policy Service node). N1A and N2A are standalone nodes at this point in time.

Assumptions

All Cisco ISE nodes in the deployment were destroyed. The new hardware was imaged at a different location using different hostnames and IP addresses.

Resolution Steps

1. Obtain the N1 backup and restore it on N1A. See "Restoring Data from a Backup" section on page 15 -11 for more information. The restore script will identify the hostname change and domain name change, and will update the hostname and domain name in the deployment configuration based on the current hostname.
2. You must generate a new self-signed certificate. See "Generating a Self-Signed Certificate" section on page 13 -7 for more information.
3. You must log in to the Cisco ISE user interface on N1A, choose **Administration > System > Deployment**, and do the following:

a. Delete the old N2 node. See "Removing a Node from Deployment" section on page 9 -26 for more information.

b. Register the new N2A node as a secondary node. See "Registering and Configuring a Secondary Node" section on page 9 -13 for more information. Data from the N1A node will be replicated to the N2A node.

Standalone Deployment, Recovery Using Existing IP Address and Hostname

There is a standalone Administration node that goes down.

Scenario

You have a standalone Administration node, N1, and a backup of the N1 database that was taken at time t1 is available. The N1 node goes down because of a physical failure and must be reimaged or a new hardware is required. The N1 node must be brought back up with the same IP address and hostname.

Assumptions

This deployment is a standalone deployment and the new or reimaged hardware has the same IP address and hostname.

Resolution Steps

Once the N1 node is back up after a reimage or you have introduced a new ISE node with the same IP address and hostname, you must restore the backup taken from the old N1 node. You do not have to make any role changes. See "Restoring Data from a Backup" section on page 15 -11 for more information.

Standalone Deployment, Recovery Using New IP Address and Hostname

There is a standalone Administration node that goes down.

Scenario

You have a standalone administration node, N1, and a backup of the N1 database that was taken at time t1 is available. The N1 node goes down because of a physical failure and will be replaced by a new hardware at a different location with a different IP address and hostname.

Assumptions

This deployment is a standalone deployment and the replaced hardware has a different IP address and hostname.

Resolution Steps

1. Replace the N1 node with a new hardware. See "Replacing the ISE Appliance Hardware" section on page 9 -28 for more information. This node will be in a standalone state and the hostname is N1B.
2. You can restore the backup on the N1B node. See "Restoring Data from a Backup" section on page 15 -11 for more information. No role changes are required.

Configuration Rollback

There may be instances where you inadvertently make configuration changes that you later determine were incorrect. For example, you may delete several NADs or modify some RADIUS attributes incorrectly and realize this issue several hours later. In this case, you can revert back to the original configuration by restoring a backup that was taken before you made the changes.

Scenario

There are two nodes: N1 (primary Administration node) and N2 (secondary Administration node) and a backup of the N1 node is available. You made some incorrect configuration changes on N1 and want to remove the changes.

Resolution Steps

Obtain a backup of the N1 node that was taken before the incorrect configuration changes were made. Restore this backup on the N1 node. See "Restoring Data from a Backup" section on page 15 -11 for more information. Restore script will sync the data from N1 to N2.

Primary Node Failure in a Distributed Deployment

In a multinode deployment, the primary Administration node fails.

Scenario

You have two ISE nodes, N1 (primary Administration node) and N2 (secondary Administration node). N1 fails because of hardware issues.

Assumptions

Only the primary node in a distributed deployment has failed.

Resolution Steps

1. Log in to the N2 user interface. Choose **Administration > System > Deployment** and configure N2 as your primary node. See "Configuring Administration Cisco ISE Nodes for High Availability" section on page 9 -15 for more information.

   The N1 node is replaced with a new hardware, reimaged, and is in the standalone state.

2. From the N2 user interface, register the new N1 node as a secondary node. See "Registering and Configuring a Secondary Node" section on page 9 -13 for more information.

   Now, the N2 node becomes your primary node and the N1 node becomes your secondary node.

If you wish to make the N1 node the primary node again, log in to the N1 user interface and make it the primary node. N2 automatically becomes a secondary server. There is no data loss.

Secondary Node Failure in a Distributed Deployment

In a multinode deployment, a single secondary node has failed. No restore is required.

Scenario

You have multiple nodes: N1 (primary Administration node), N2 (secondary Administration node), N3 (secondary Policy Service node), N4 (secondary Policy Service node). One of the secondary nodes, N3, fails.

Resolution Steps

1. Reimage the new N3A node to the default standalone state.
2. Log in to the N1 user interface and delete the N3 node. See "Removing a Node from Deployment" section on page 9 -26 for more information.
3. Reregister the N3A node. See "Registering and Configuring a Secondary Node" section on page 9 -13 for more information.

   Data is replicated from N1 to N3A. No restore is required.