Backup Data Type

Cisco ISE allows you to back up data from the primary PAN and from the Monitoring node. Backup can be done from the CLI or user interface.

Cisco ISE allows you to back up the following type of data:

- Configuration data—Contains both application-specific and Cisco ADE operating system configuration data. Backup can be done via the primary PAN using the GUI or CLI.
- Operational Data—Contains monitoring and troubleshooting data. Backup can be done via the primary PAN GUI or using the CLI for the Monitoring node.

When Cisco ISE is run on VMware, VMware snapshots are not supported for backing up ISE data.

Note Cisco ISE does not support VMware snapshots for backing up ISE data because a VMware snapshot saves the status of a VM at a given point in time. In a multinode Cisco ISE deployment, data in all the nodes are continuously synchronized with current database information. Restoring a snapshot might cause database replication and synchronization issues. Cisco recommends that you use the backup functionality included in Cisco ISE for archival and restoration of data.

Using VMware snapshots or any third-party backup service to back up Cisco ISE data might result in interrupting Cisco ISE services. When a backup is initiated by VMware or any other third-party backup service like CommVault SAN level backup, it quiesces the file system to maintain crash consistency, which can cause your Cisco ISE functionalities to freeze. A reboot is required to resume the services on your Cisco ISE deployment.

Restore operation, can be performed with the backup files of previous versions of Cisco ISE and restored on a later version, as long as the previous versions are in the supported direct upgrade path for the later version.
.

Cisco ISE, Release 2.4 supports restore from backups obtained from Release 2.0 and later.

Note While recreating a deployment after backing up and restoring data, a **Context Visibility Reset** of both Primary PAN and Secondary PAN are required to ensure that data on both the nodes are synced.

- Backup and Restore Repositories, on page 2
- On-Demand and Scheduled Backups, on page 5

Backup and Restore Repositories

Cisco ISE allows you to create and delete repositories through the administrator portal. You can create the following types of repositories:

- DISK
- FTP
- SFTP
- NFS
- CD-ROM
- HTTP
- HTTPS

Note Repositories are local to each device.

We recommend that you have a repository size of minimum 100 GB for all types of deployment (small, medium, and large).

Create Repositories

You can use the CLI and GUI to create repositories. We recommend that you use the GUI due to the following reasons:

- Repositories that are created through the CLI are saved locally and do not get replicated to the other deployment nodes. These repositories do not get listed in the GUI's repository page.
- Repositories that are created on the primary PAN get replicated to the other deployment nodes.

The keys are generated only at the primary PAN on GUI, and so during upgrade you need to generate the keys again at GUI of new primary admin and export it to the SFTP server. If you remove the nodes from your deployment, you need to generate the keys on GUI of non-admin nodes and export it to the SFTP server.

You can configure an SFTP repository in Cisco ISE with RSA public key authentication. Instead of using an administrator-created password to encrypt the database and logs, you can choose the RSA public key authentication that uses secure keys. In case of SFTP repository created with RSA public key, the repositories created through the GUI do not get replicated in the CLI and the repositories created through the CLI do not get replicated in the GUI. To configure same repository on the CLI and GUI, generate RSA public keys on both CLI and GUI and export both the keys to the SFTP server.

Before you begin

- To perform the following task, you must have the privileges of either a Super Admin or System Admin.
- If you want to create an SFTP repository with RSA public key authentication, perform the following steps:
  - Enable RSA public key authentication in the SFTP repository.
  - You must log in as the Admin CLI user. Enter the host key of the SFTP server from the Cisco ISE CLI using the **crypto host_key add** command. The host key string should match the hostname that you enter in the **Path** field of the repository configuration page.
  - Generate the key pairs and export the public key to your local system from the GUI. From the Cisco ISE CLI, generate the key pairs using the **crypto key generate rsa passphrase** *test123* command, where, passphrase must be greater than four letters, and export the keys to any repository (local disk or any other configured repository).
  - Copy the exported RSA public key to the PKI-enabled SFTP server and add it to the "authorized_keys" file.

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

Step 2 Click **Add** to add a new repository.

Step 3 Enter the values as required to set up new repository. See Repository Settings, on page 4 for a description of the fields.

Step 4 Click **Submit** to create the repository.

Step 5 Verify that the repository is created successfully by clicking **Repository** from the **Operations** navigation pane on the left or click the **Repository List** link at the top of **Repository** window to go to the repository listing page.

What to do next

- Ensure that the repository that you have created is valid. You can do so from the **Repository Listing** window. Select the corresponding repository and click **Validate**. Alternatively, you can execute 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.

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

  ```
  %Invalid Directory
  ```

- Run an on-demand backup or schedule a backup.

Repository Settings

Table 1: Repository Settings

Fields	Usage Guidelines
Repository	Enter the name of the repository. Alphanumeric characters are allowed and the maximum length is 80 characters.
Protocol	Choose one of the available protocols that you want to use.
Server Name	(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. **Note** Ensure that the ISE eth0 interface is configured with an IPv6 address if you are adding a repository with an IPv6 address.
Path	Enter the path to your repository. The path must be valid and must exist at the time you create the repository. This value can start with two forward slashes (//) or a single forward slash (/) denoting the root directory of the server. However, for the FTP protocol, a single forward slash (/) denotes the FTP of the local device home directory and not the root directory.
Enable PKI authentication	(Optional; applicable only for SFTP repository) Check this check box if you want to enable RSA Public Key Authentication in SFTP repository.
User Name	(Required for FTP, SFTP, and NFS) Enter the username that has write permission to the specified server. A username can contain alphanumeric and _-. /@\$ characters.
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 to 9, a to z, A to Z, -, ., |, @, #,$, ^, &, *, (,), +, and =.

Related Topics

Backup and Restore Repositories, on page 2

Create Repositories, on page 2

Enable RSA Public Key Authentication in SFTP Repository

In the SFTP server, each node must have two RSA public keys, one each for CLI and for GUI. To enable RSA public key authentication in SFTP repository, perform the following steps:

Note After you enable RSA public key authentication in SFTP repository, you will not be able to log in using SFTP credentials. You can either use PKI-based authentication or credential-based authentication. If you want to use credential-based authentication again, you must remove the public key pair from the SFTP server.

Step 1 Log in to SFTP server with an account that has permission to edit the */etc/ssh/sshd_config.file*.

Note The location of the *sshd_config* file might vary based on the operating system installation.

Step 2 Enter the **vi /etc/ssh/sshd_config** command.

The contents of the sshd_config file is listed.

Step 3 Remove the "#" symbol from the following lines to enable RSA public key authentication:

- RSAAuthentication yes
- PubkeyAuthentication yes

  Note If Public Auth Key is no, change it to yes.

- AuthorizedKeysFile ~/.ssh/authorized_keys

On-Demand and Scheduled Backups

You can configure on-demand backups of the primary PAN and the primary monitoring node. Perform an on-demand backup when you want to back up data immediately.

You can schedule system-level backups to run once, daily, weekly, or monthly. Because backup operations can be lengthy, you can schedule them so they are not a disruption. You can schedule a backup from the Admin portal.

Note If you are using the internal CA, you should use the CLI to export certificates and keys. Backup using in the administration portal does not back up the CA chain.

For more information, see the "Export Cisco ISE CA Certificates and Keys" section in the "Basic Setup" chapter *Cisco Identity Services Engine Administrator Guide* .

Configurational and operational backups on Cisco ISE can overload your system for a short time. This expected behaviour of temporary system overload will depend on the configuration and monitoring database size of your system.

Related Topics

Maintenance Settings

Perform an On-Demand Backup

You can perform an On-demand backup to instantly back up the configuration or monitoring (operational) data. The restore operation restores Cisco ISE to the configuration state that existed at the time of obtaining the backup.

Important When performing a back up and restore, the restore overwrites the list of trusted certificates on the target system with the list of certificates from the source system. It is critically important to note that backup and restore functions do not include private keys associated with the Internal Certificate Authority (CA) certificates.

If you are performing a back up and restore from one system to another, you have to choose from one of these options to avoid errors:

- **Option 1:**

  Export the CA certificates from the source ISE node through the CLI and import them in to the target system through the CLI.

  Pros: Any certificates issued to endpoints from the source system will continue to be trusted. Any new certificates issued by the target system will be signed by the same keys.

  Cons: Any certificates that have been issued by the target system prior to the restore function will not be trusted and will need to be re-issued.

- **Option 2:**

  After the restore process, generate all new certificates for the internal CA.

  Pros: This option is the recommended and clean method, where neither the original source certificates or the original target certificates will be used. Certificates issued by the original source system continues to be trusted.

  Cons: Any certificates that have been issued by the target system prior to the restore function will not be trusted and will need to be re-issued.

Before you begin

- Before you perform an on-demand backup, you should have a basic understanding of the backup data types in Cisco ISE.
- Ensure that you have created repositories for storing the backup files.
- Do not back up using a local repository. You cannot back up the monitoring data in the local repository of a remote Monitoring node.
- Ensure that you perform all certificate-related changes before you obtain the backup.
- To perform the following task, you must be a Super Admin or System Admin.

Note For backup and restore operations, the following repository types are not supported: CD-ROM, HTTP, HTTPS, or TFTP. This is because, either these repository types are read-only or the protocol does not support file listing. To restore a backup, choose the repository and click **Restore**.

Related Topics

Cisco ISE Restore Operation, on page 10

Export Authentication and Authorization Policy Configuration, on page 16

On-Demand Backup Settings

The following table describes the fields on the **On-Demand Backup** window, which you can use to obtain a backup at any point of time. The navigation path for this window is **Administration** > **System** > **Backup & Restore**.

Table 2: On-Demand Backup Settings

Field Name	Usage Guidelines
Type	Choose one of the following: • **Configuration Data Backup**: Includes both application-specific and Cisco ADE operating system configuration data • **Operational Data Backup**: Includes monitoring and troubleshooting data
Backup Name	Enter the name of your backup file.
Repository Name	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.
Encryption Key	This key is used to encrypt and decrypt the backup file.

Related Topics

Backup Data Type, on page 1
On-Demand and Scheduled Backups, on page 5
Backup History, on page 10
Backup Failures, on page 10
Cisco ISE Restore Operation, on page 10
Export Authentication and Authorization Policy Configuration, on page 16
Synchronize Primary and Secondary Nodes in a Distributed Environment, on page 17
Perform an On-Demand Backup, on page 5

Schedule a Backup

You can perform an On-demand backup to instantly back up the configuration or monitoring (operational) data. The restore operation restores Cisco ISE to the configuration state that existed at the time of obtaining the backup.

Important When performing a back up and restore, the restore overwrites the list of trusted certificates on the target system with the list of certificates from the source system. It is critically important to note that backup and restore functions do not include private keys associated with the Internal Certificate Authority (CA) certificates.

If you are performing a back up and restore from one system to another, you will have to choose from one of these options to avoid errors:

- **Option 1:**

  Export the CA certificates from the source ISE node through the CLI and import them in to the target system through the CLI.

  Pros: Any certificates issued to endpoints from the source system will continue to be trusted. Any new certificates issued by the target system will be signed by the same keys.

  Cons: Any certificates that have been issued by the target system prior to the restore function will not be trusted and will need to be re-issued.

- **Option 2:**

  After the restore process, generate all new certificates for the internal CA.

  Pros: This option is the recommended and clean method, where the original source certificates or the original target certificates will be used. Certificates issued by the original source system will continue to be trusted.

  Cons: Any certificates that have been issued by the target system prior to the restore function will not be trusted and will need to be re-issued.

Before you begin

- Before you schedule a backup, you should have a basic understanding of the backup data types in Cisco ISE.
- Ensure that you have configured repositories.
- Do not back up using a local repository. You cannot back up the monitoring data in the local repository of a remote Monitoring node.
- To perform the following task, you must be a Super Admin or System Admin.

Note For backup and restore operations, the following repository types are not supported: CD-ROM, HTTP, HTTPS, or TFTP. This is because, either these repository types are read-only or the protocol does not support file listing.

Step 1 Choose **Administration** > **System** > **Backup and Restore**.

Step 2 Click **Schedule** to schedule a Configuration or an Operational backup.

Step 3 Enter the values as required to schedule a backup.

Step 4 Click **Save** to schedule the backup.

Step 5 Perform one of the following actions:

- From the **Select Repository** drop-down list, choose the required repository.
- Click the **Add Repository** link to add a new repository.

Step 6 Click the **Refresh** link to see the scheduled backup list.

You can create only one schedule at a time for a Configuration or Operational backup. You can enable or disable a scheduled backup, but you cannot delete it.

Scheduled Backup Settings

The following table describes the fields on the Scheduled Backup window, which you can use to restore a full or incremental backup. The navigation path for this window is **Administration** > **System** > **Backup and Restore**.

Table 3: Scheduled Backup Settings

Field Name	Usage Guidelines
Type	Choose one of the following: • **Configuration Data Backup**: Includes both application-specific and Cisco ADE operating system configuration data • **Operational Data Backup**: Includes monitoring and troubleshooting data
Name	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. On the Scheduled Backup list window, the backup filename will be prepended with “backup_occur” to indicate that the file is an occurrence **kron** job.
Description	Enter a description for the backup.
Repository Name	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.
Encryption Key	Enter a key to encrypt and decrypt the backup file.
Schedule Options	Choose the frequency of your scheduled backup and fill in the other options accordingly.

Related Topics

Backup Data Type, on page 1
On-Demand and Scheduled Backups, on page 5
Backup History, on page 10
Backup Failures, on page 10
Cisco ISE Restore Operation, on page 10
Export Authentication and Authorization Policy Configuration, on page 16

Synchronize Primary and Secondary Nodes in a Distributed Environment, on page 17
Backup Using the CLI, on page 10
Schedule a Backup, on page 7

Backup Using the CLI

Although you can schedule backups both from the CLI as well as the GUI, it is recommended to use GUI. However, you can perform operational backup on the secondary monitoring node only from the CLI.

Backup History

Backup history provides basic information about scheduled and on-demand backups. It lists the name of the backup, backup file size, repository where the backup is stored, and time stamp that indicates when the backup was obtained. This information is available in the Operations Audit report and on the Backup and Restore page in the History table.

For failed backups, Cisco ISE triggers an alarm. The backup history page provides the failure reason. The failure reason is also cited in the Operations Audit report. If the failure reason is missing or is not clear, you can run the **backup-logs** command from the Cisco ISE CLI and look at the ADE.log for more information.

While the backup operation is in progress, you can use the **show backup status** CLI command to check the progress of the backup operation.

Backup history is stored along with the Cisco ADE operating system configuration data. It remains there even after an application upgrade and are only removed when you reimage the PAN.

Backup Failures

If backup fails, check the following:

-
- Make sure that no other backup is running at the same time.
- Check the available disk space for the configured repository.
  - Monitoring (operational) backup fails if the monitoring data takes up more than 75% of the allocated monitoring database size. For example, if your Monitoring node is allocated 600 GB, and the monitoring data takes up more than 450 GB of storage, then monitoring backup fails.
  - If the database disk usage is greater than 90%, a purge occurs to bring the database size to less than or equal to 75% of its allocated size.
- Verify if a purge is in progress. Backup and restore operations will not work while a purge is in progress.
- Verify if the repository is configured correctly.

Cisco ISE Restore Operation

You can restore configuration data on a primary or standalone administration node. After you restore data on the Primary PAN, you must manually synchronize the secondary nodes with the Primary PAN.

The process for restoring the operational data is different depending on the type of deployment.

Note The new backup/restore user interface in Cisco ISE makes use of meta-data in the backup filename. Therefore, after a backup completes, you should not modify the backup filename manually. If you manually modify the backup filename, the Cisco ISE backup/restore user interface will not be able to recognize the backup file. If you have to modify the backup filename, you should use the Cisco ISE CLI to restore the backup.

Guidelines for Data Restoration

Following are guidelines to follow when you restore Cisco ISE backup data.

- Cisco ISE allows you to obtain a backup from an ISE node (A) and restore it on another ISE node (B), both having the same host names (but different IP addresses). However, after you restore the backup on node B, do not change the hostname of node B because it might cause issues with certificates and portal group tags.
- If you obtain a backup from the Primary PAN in one timezone and try to restore it on another Cisco 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 Cisco 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.
- When you restore a backup on the Primary PAN with a different hostname than the one from which the backup was obtained, the Primary PAN becomes a standalone node. The deployment is broken and the secondary nodes become nonfunctional. You must make the standalone node the primary node, reset the configuration on the secondary nodes, and reregister them with the primary node. To reset the configuration on Cisco ISE nodes, enter the following command from the Cisco ISE CLI:
    - **application reset-config ise**
- We recommend that you do not change the system timezone after the initial Cisco ISE installation and setup.
- If you changed the certificate configuration on one or more nodes in your deployment, you must obtain another backup to restore the data from the standalone Cisco ISE node or Primary PAN. Otherwise, if you try to restore data using an older backup, the communication between the nodes might fail.
- After you restore the configuration backup on the Primary PAN, you can import the Cisco ISE CA certificates and keys that you exported earlier.

  Note If you did not export the Cisco ISE CA certificates and keys, then after you restore the configuration backup on the Primary PAN, generate the root CA and subordinate CAs on the Primary PAN and Policy Service Nodes (PSNs).

- If you are trying to restore a platinum database without using the correct FQDN (FQDN of a platinum database), you need to regenerate the CA certificates. (choose **Administration > Certificates > Certificate Signing Requests > Replace ISE Root CA certificate chain**).However, If you restore the platinum database with the correct FQDN, note that the CA certificates regenerated automatically.

- 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.
- If you have a standalone administration node that fails, you must run the configuration backup to restore it. If the Primary PAN fails, you can use the distributed setup to promote your Secondary Administration Node to become the primary. You can then restore data on the Primary PAN after it comes up.

Note Cisco ISE also provides the **backup-logs** CLI command that you can use to collect log and configuration files for troubleshooting purposes.

Restoration of Configuration or Monitoring (Operational) Backup from the CLI

To restore configuration data through the Cisco ISE CLI, use the **restore** command in the EXEC mode. Use the following command to restore data from a configuration or operational backup:

restore *filename* **repository** *repository-name* **encryption-key hash|plain** *encryption-key name* **include-adeos**

Syntax Description

restore	Type this command to restore data from a configuration or operational backup.
filename	Name of the backed-up file that resides in the repository. Supports up to 120 alphanumeric characters. **Note** You must add the .tar.gpg extension after the filename (for example, myfile.tar.gpg).
repository	Specifies the repository that contains the backup.
repository-name	Name of the repository you want to restore the backup from.
encryption-key	(Optional) Specifies user-defined encryption key to restore backup.
hash	Hashed encryption key for restoring backup. Specifies an encrypted (hashed) encryption key that follows. Supports up to 40 characters.
plain	Plaintext encryption key for restoring backup. Specifies an unencrypted plaintext encryption key that follows. Supports up to 15 characters.
encryption-key name	Enter the encryption key.
include-adeos	(Optional, applicable only for configuration backup) Enter this command operator parameter if you want to restore ADE-OS configuration from a configuration backup. When you restore a configuration backup, if you do not include this parameter, Cisco ISE restores only the Cisco ISE application configuration data.

Defaults

No default behavior or values.

Command Modes

EXEC

Usage Guidelines

When you use restore commands in Cisco ISE, the Cisco ISE server restarts automatically.

The 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.

Examples

```
ise/admin# restore mybackup-100818-1502.tar.gpg repository myrepository encryption-key plain
 Lab12345
Restore may require a restart of application services. Continue? (yes/no) [yes] ? yes
Initiating restore.  Please wait...
ISE application restore is in progress.
This process could take several minutes. Please wait...
Stopping ISE Application Server...
Stopping ISE Monitoring & Troubleshooting Log Processor...
Stopping ISE Monitoring & Troubleshooting Log Collector...
Stopping ISE Monitoring & Troubleshooting Alert Process...
Stopping ISE Monitoring & Troubleshooting Session Database...
Stopping ISE Database processes...
Starting ISE Database processes...
Starting ISE Monitoring & Troubleshooting Session Database...
Starting ISE Application Server...
Starting ISE Monitoring & Troubleshooting Alert Process...
Starting ISE Monitoring & Troubleshooting Log Collector...
Starting ISE Monitoring & Troubleshooting Log Processor...
Note: ISE Processes are initializing. Use 'show application status ise'
      CLI to verify all processes are in running state.
ise/admin#
```

Related Commands

	Description
backup	Performs a backup (Cisco ISE and Cisco ADE OS) and places the backup in a repository.
backup-logs	Backs up system logs.
repository	Enters the repository submode for configuration of backups.
show repository	Displays the available backup files located on a specific repository.
show backup history	Displays the backup history of the system.
show backup status	Displays the status of the backup operation.
show restore status	Displays the status of the restore operation.

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 PAN and perform a manual synchronization.

Restore Configuration Backups from the GUI

You can restore a configuration backup from the Admin portal.

Before you begin

Ensure that the primary PAN Auto Failover configuration, if enabled in your deployment, is turned off. When you restore a configuration backup, the application server processes are restarted. There might be a delay while these services restart. Due to this delay in restart of services, auto failover of secondary PAN might get initiated.

When your deployment is a dual node deployment at the time configuration backup, ensure the following:

- Source and target nodes for the restore are same as the ones used for the configuration backup, the target node can be either stand-alone or primary.
- Source and target nodes for the restore are different from the ones used in the configuration backup, the target node must be stand-alone.

Note You can restore configuration database backup and regenerate the Root CA on a primary PAN only. However, you cannot restore the configuration database backup on a registered PAN.

Step 1 Choose **Administration** > **System** > **Backup and Restore**.

Step 2 Select the name of the backup from the list of Configurational backup and click **Restore**.

Step 3 Enter the Encryption Key used during the backup.

Step 4 Click **Restore**.

What to do next

If you are using the Cisco ISE CA service, you must:

1. Regenerate the entire Cisco ISE CA root chain.
2. Obtain a backup of the Cisco ISE CA certificates and keys from the primary PAN and restore it on the secondary PAN. This ensures that the secondary PAN can function as the root CA or subordinate CA of an external PKI in case of a Primary PAN failure and you promote the secondary PAN to be the primary PAN.

Restoration of Monitoring Database

The process for restoring the Monitoring database is different depending on the type of deployment. The following sections explain how to restore the Monitoring database in standalone and distributed deployments.

You must use the CLI to restore an on-demand Monitoring database backup from previous releases of Cisco ISE. Restoring a scheduled backup across Cisco ISE releases is not supported.

Note If you attempt to restore data to a node other than the one from which the data was taken, you must configure the logging target settings to point to the new node. This ensures that the monitoring syslogs are sent to the correct node.

Restore a Monitoring (Operational) Backup in a Standalone Environment

The GUI lists only the backups that are taken from the current release. To restore backups that obtained from earlier releases, use the restore command from the CLI.

Before you begin

- Purge the old monitoring data.
- Schedule a backup or perform an on-demand backup.

Step 1 Choose **Administration** > **System** > **Backup and Restore**.

Step 2 Select the name of the backup from the list of Operational backup and click **Restore**.

Step 3 Enter the Encryption Key used during the backup.

Step 4 Click **Restore**.

Restore a Monitoring Backup with Administration and Monitor Personas

You can restore a Monitoring backup in a distributed environment with administration and monitor personas.

Before you begin

- Purge the old monitoring data.
- Schedule a backup or perform an on-demand backup.

Step 1 If you are using a primary and secondary PAN, synchronize the PANs.

When you synchronize the PANs, you must chose a PAN and promote that to be the active primary.

Step 2 Before you deregister the Monitoring node, assign the Monitoring persona to another node in the deployment.

Every deployment must have at least one functioning Monitoring node.

Step 3 Deregister the Monitoring node for backup.

Step 4 Restore the Monitoring backup to the newly deregistered node.

Step 5 Register the newly restored node with the current Administration node.

Step 6 Promote the newly restored and registered node as the active Monitoring node.

Restore a Monitoring Backup with a Monitoring Persona

You can restore a Monitoring backup in a distributed environment with only Monitoring persona.

Before you begin

- Purge the old monitoring data.
- Schedule a backup or perform an on-demand backup.

Step 1 Prepare to deregister the node to be restored. This is done by assigning the monitoring persona to another node in the deployment.

A deployment must have at least one functioning Monitoring node.

Step 2 Deregister the node to be restored.

Note Wait until the deregistration is complete before proceeding with the restore. The node must be in a standalone state before you can continue with the restore.

Step 3 Restore the Monitoring backup to the newly deregistered node.

Step 4 Register the newly restored node with the current Administration node.

Step 5 Promote the newly restored and registered node as the active Monitoring node.

Restore History

You can obtain information about all restore operations, log events, and statuses from the **Operations Audit Report** window.

Note However, the **Operations Audit Report** window does not provide information about the start times corresponding to the previous restore operations.

For troubleshooting information, you have to run the **backup-logs** command from the Cisco ISE CLI and look at the ADE.log file.

While the restore operation is in progress, all Cisco ISE services are stopped. You can use the **show restore status** CLI command to check the progress of the restore operation.

Export Authentication and Authorization Policy Configuration

You can export authentication and authorization policy configuration in the form of an XML file that you can read offline to identify any configuration errors and use for troubleshooting purposes. This XML file includes authentication and authorization policy rules, simple and compound policy conditions, Discretionary Access control Lists (DACLs), and authorization profiles. You can choose to email the XML file or save it to your local system.

Step 1 Choose **Administration** > **System** > **Backup & Restore**.

Step 2 Click **Policy Export**.

Step 3 Enter the values as needed.

Step 4 Click **Export**.

Use a text editor such as WordPad to view the contents of the XML file.

Schedule Policy Export Settings

The following table describes the fields on the **Schedule Policy Export** window. The navigation path for this window is **Administration** > **System** > **Backup and Restore** > **Policy Export**.

Table 4: Schedule Policy Export Settings

Synchronize Primary and Secondary Nodes in a Distributed Environment

In a distributed environment, sometimes the Cisco ISE database in the primary and secondary nodes are not synchronized automatically after restoring a backup file on the PAN. If this happens, you can manually force a full replication from the PAN to the secondary ISE nodes. You can force a synchronization only from the PAN to the secondary nodes. During the sync-up operation, you cannot make any configuration changes. Cisco ISE allows you to navigate to other Cisco ISE Admin portal pages and make any configuration changes only after the synchronization is complete.

Before you begin

To perform the following task, you must be a Super Admin or System Admin.

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

Step 2 Check the check boxes next to the secondary ISE nodes with an Out of Sync replication status.

Step 3 Click **Syncup** and wait until the nodes are synchronized with the PAN. You will have to wait until this process is complete before you can access the Cisco ISE Admin portal again.

Recovery of Lost Nodes in Standalone and Distributed Deployments

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

Recovery of Lost Nodes Using Existing IP Addresses and Hostnames in a Distributed Deployment

Scenario

In a distributed deployment, a natural disaster leads to a loss of all the nodes. After recovery, you want to use the existing IP addresses and hostnames.

For example, you have two nodes: N1 (Primary Policy Administration Node or Primary PAN) and N2 (Secondary Policy Administration Node or Secondary PAN.) A backup of the N1 node, which was taken at time T1, is available. 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 the N1 and N2 nodes. N1 and N2 nodes will now have a standalone configuration.
2. Obtain a license with the UDI of the N1 and N2 nodes and install it on the N1 node.
3. You must then restore the backup on the replaced N1 node. The restore script will try to sync the data on N2, but N2 is now a standalone node and the synchronization fails. Data on N1 will be reset to time T1.
4. You must log in to the N1 Admin portal to delete and reregister the N2 node. Both the N1 and N2 nodes will have data reset to time T1.

Recovery of Lost Nodes Using New IP Addresses and Hostnames in a Distributed Deployment

Scenario

In a distributed deployment, a natural disaster leads to loss of all the nodes. The new hardware is reimaged at a new location and requires new IP addresses and hostnames.

For example, you have two ISE nodes: N1 (primary Policy Administration Node or primary PAN) and N2 (secondary Policy Service Node.) A backup of the N1 node which was taken at time T1, is available. Later, both N1 and N2 nodes fail because of a natural disaster. The Cisco ISE nodes are replaced at a new location and the new hostnames are N1A (primary PAN) 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. 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.
3. You must log in to the Cisco ISE administrator portal on N1A, choose **Administration** > **System** > **Deployment**, and do the following:

   Delete the old N2 node.

   Register the new N2A node as a secondary node. Data from the N1A node will be replicated to the N2A node.

Recovery of a Node Using Existing IP Address and Hostname in a Standalone Deployment

Scenario

A standalone administration node is down.

For example, you have a standalone administration node, N1. A backup of the N1 database was taken at time T1. 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 up after a reimage or you have introduced a new Cisco 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.

Recovery of a Node Using New IP Address and Hostname in a Standalone Deployment

Scenario

A standalone administration node is down.

For example, you have a standalone administration node, N1. A backup of the N1 database taken at time T1 is available. The N1 node is 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 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. This node will be in a standalone state and the hostname is N1B.
2. You can restore the backup on the N1B node. No role changes are required.

Configuration Rollback

Problem

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 to the original configuration by restoring a backup that was taken before you made the changes.

Possible Causes

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

Solution

Obtain a backup of the N1 node that was taken before the incorrect configuration changes were made. Restore this backup on the N1 node. The restore script will synchronize the data from N1 to N2.

Recovery of Primary Node in Case of Failure in a Distributed Deployment

Scenario

In a multinode deployment, the PAN fails.

For example, you have two Cisco ISE nodes, N1 (PAN) 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 administrator portal. Choose **Administration** > **System** > **Deployment** and configure N2 as your primary node.

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

2. From the N2 administrator portal, register the new N1 node as a secondary node.

   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 administrator portal and make it the primary node. N2 automatically becomes a secondary server. There is no data loss.

Recovery of Secondary Node in Case of Failure in a Distributed Deployment

Scenario

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

For example, you have multiple nodes: N1 (primary PAN), N2 (secondary PAN), 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 Admin portal and delete the N3 node.

3. Reregister the N3A node.

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