



Backup and Restore

- [Backup and Restore Concepts, on page 1](#)
- [Set Up and Manage Repositories, on page 6](#)
- [Set Up Automatic Application Backups, on page 11](#)
- [Perform a Manual Backup, on page 13](#)
- [Restore Prime Infrastructure Data, on page 15](#)
- [How to Manage Disk Space Issues During Backup and Restore, on page 17](#)
- [Backup and Restore with Operations Center, on page 19](#)

Backup and Restore Concepts

- [Backup Types: Application and Appliance, on page 1](#)
- [Backup Scheduling, on page 2](#)
- [Backup Repositories, on page 3](#)
- [Backup Filenames, on page 3](#)
- [Backup Validation Process, on page 4](#)
- [Information That Is Backed Up, on page 4](#)
- [Information That Is Not Backed Up, on page 6](#)

Backup Types: Application and Appliance

Prime Infrastructure supports two types of backups:

- **Application backups**—Contain Prime Infrastructure application data but do not include platform data (host-specific settings, such as the server hostname and IP address). Application backup should be used during Prime Infrastructure upgrade, when you want to move only application data and not the platform/host specific configurations.
- **Appliance backups**—Contain all application data and platform data (host-specific settings, including the hostname, IP address, subnet mask, default gateway, and so on). Appliance backup should be used for disaster recovery (or to recover from platform hardware or software failures). For example, to recover from any disk or filesystem failure, the standard recovery process would be to reinstall Prime Infrastructure and then restore from the appliance backup in order to restore all data as well as platform-specific configurations. You would then need to manually reconstruct the HA configurations as they are not included in the appliance backup.



Note For details on what is considered application data and what is considered platform data, see [Information That Is Backed Up, on page 4](#).

Note the following about application and appliance backups.

- Application and appliance backups can be restored to the same or a new host, as long as the new host has the same hardware and software configuration as the host from which the backup was taken.
- You can only restore an appliance backup to a host running the same version of the Prime Infrastructure server software as the server from which the backup was taken.
- When upgrading to a later version of Prime Infrastructure, application backup and restore can run across different releases, as long as the upgrade path is supported.
- You cannot restore an application backup using the appliance restore command, nor can you restore an appliance backup using the application restore command.

We recommend the following best practices:

- If you are *evaluating* Prime Infrastructure, use the default automatic application backup to the local repository.
- If you are running Prime Infrastructure *in a production environment* as a virtual appliance, take regular application backups to a remote backup server. You can use the application backups to restore your server for all failures except complete failure of the server hardware.

Backup Scheduling

Prime Infrastructure performs automatic scheduled application backups. This feature is enabled by default and creates one application backup file every day in the default local backup repository.

You can change this schedule as needed. You can also take an automatic application backup at any time from the web GUI. Appliance backups can only be taken from the command line.

Automatic application backups can create storage space problems if the backup repository is local to the Prime Infrastructure server. While this is usually acceptable in test implementations, it is not intended to substitute for routine scheduled backups to remote servers in a production environment.

We recommend the following for production environments:

- Set up remote repositories to store the backup files.
- Use the automatic schedule application backup to create backups on the remote repositories on a regular schedule.

Even if you are using scheduled backups, you can still use the command line to create application or appliance backups at any time.



Note By default, two minutes are added to the job execution time for job creation.

Backup Repositories

By default, automatic application backup feature stores backup files in the local backup repository **/localdisk/defaultRepo**. You can use the web GUI to create a new local backup repository and then choose it when you set up automatic application backups. You can also specify a remote repository but you must create the repository first as described in [Set Up and Manage Repositories, on page 6](#).

When taking application or appliance backups using the command line, you must specify the local or remote repository you want the backup to be stored in. In a production environment, this is normally a remote repository that is accessed via NFS, SFTP, or FTP. We recommend you use NFS because it is typically much faster and more reliable than other protocols.

There is no difference between performing an application backup from the command line or performing it from the web GUI. Both actions create the same backup file.

Whenever you use NFS to take backups or restore data from a remote backup, make sure the mounted NFS server remains active throughout the backup or restore operation. If the NFS server shuts down at any point in the process, the backup or restore operation will hang without warning or an error message.

Backup Filenames

Application backups launched from the web GUI—either automatically or manually—are assigned a filename with the following format:

host-yymmdd-hhmm_VERver_BKSZsize_CPUcpus_MEMtarget_RAMram_SWAPswap_APP_CKchecksum.tar.gpg

Application backups launched from the CLI use the same format, except that the file starts with the user-specified filename rather than the server name.

filename-yymmdd-hhmm_VERver_BKSZsize_CPUcpus_MEMtarget_RAMram_SWAPswap_APP_CKchecksum.tar.gpg

Appliance backups launched from the CLI have files that also start with the user-specified filename, but the type is indicated as SYS, not APP.

filename-yymmdd-hhmm_VERver_BKSZsize_CPUcpus_MEMtarget_RAMram_SWAPswap_SYS_CKchecksum.tar.gpg

The following table describes the variables used by the backup files.

Variable	Description
<i>host</i>	Host name of the server from which the backup was taken (for application backups launched from web GUI).
<i>filename</i>	Filename specified by user in command line (for application backups launched from CLI, and for appliance backups)
<i>yymmdd-hhmm</i>	Date and time the backup was taken
<i>ver</i>	Internal version.
<i>size</i>	Total size of the backup
<i>cpus</i>	Total number of CPUs in the server from which the backup was taken
<i>target</i>	Total amount of system memory in the server from which the backup was taken
<i>ram</i>	Total amount of RAM in the server from which the backup was taken

<i>swap</i>	Total size of the swap disk on the server from which the backup was taken
<i>checksum</i>	Backup file checksum

Backup Validation Process

Prime Infrastructure performs the following steps to validate the backup files:

1. Before starting the backup process, validates disk size, fast-recovery area, and control files.
2. Validates the created backup database to ensure that it can be restored.
3. Validates the zipped application data against the files that were backed up.
4. Validates the TAR file to make sure it is correct and complete.
5. Validates the GPG file to ensure that it is correct.

If you manually transfer the backup file, or if you want to verify that the backup file transfer is completed, view the file's md5Checksum and file size.

Another best practice for validating a backup is to restore it to a standalone "test" installation of Prime Infrastructure.

Information That Is Backed Up

The following table describes the information that is contained in backup files. This information is restored to the server from backups.

See [Information That Is Not Backed Up, on page 6](#) for details about data that is not saved by the backup mechanism.



Note

The `/opt/CSCOlumos/conf/Migration.xml` file contains all configuration files and reports that are backed up. This file is included in the backup and is restored.

Data Type	Feature	Information Saved and Restored
-----------	---------	--------------------------------

Application Data	Background job settings	Data in the database
	Configuration archive (device configuration files)	Data in the database
	Configuration templates	<ul style="list-style-type: none"> • Files in /opt/CSCOLumos: <ul style="list-style-type: none"> • /conf/ootb • /xmp_inventory/dar/customized-feature-parts/CONFIGURATION • Data in the database
	Credentials	Data in the database
	Device inventory data	Data in the database
	Licenses	Files in /opt/CSCOLumos/licenses
	Maps	<ul style="list-style-type: none"> • Files in /opt/CSCOLumos/domainmaps • Data in the database
	Reports	<ul style="list-style-type: none"> • Files in /localdisk/ftp: <ul style="list-style-type: none"> • /reports • /reportsOnDemand • Data in the database
	Managed device software image files	Data in the database
	System settings	Data in the database
	User preferences	<ul style="list-style-type: none"> • Files in /opt/CSCOLumos/conf/wap/datastore/webacs/xml/prefs • Data in the database
	Prime Infrastructure users, groups, and roles	Data in the database
	Virtual domains	Data in the database

Platform Data	CLI settings	All CLI information and settings are preserved. This includes the list of backup repositories, the FTP user name, users created using the CLI, AAA information specified via the CLI, and other CLI settings (such as the terminal timeout).
	Credentials	Linux OS credentials file
	Network settings	Files in /opt/CSCOlumos/conf/rfm/classes/com/cisco/packaging/PortResources.xml
	Linux user preferences	Linux data structure
	Linux users, groups, and roles	Linux data structure

Information That Is Not Backed Up

Before performing a backup, make sure that you manually note the following information because it is not saved as part of the backup process. You will need to reconfigure these settings after the data has been restored.

- High availability configurations
- Local customization (for example, report heap size)

Patch history information is also not saved.

For a list of information that is backed up, see [Information That Is Backed Up, on page 4](#).

Set Up and Manage Repositories

Prime Infrastructure supports the following repository types:

- Local repositories
- Remote repositories—NFS, FTP, SFTP

See the following topics for information on how to set up and manage these different types of repositories.

Create a Local Backup Repository

Prime Infrastructure stores automatic backup files in the default local backup repository `/localdisk/defaultRepo`. You can create a different local backup repository and use it if you prefer.

-
- Step 1** Choose **Administration > Dashboards > Job Dashboard**.
 - Step 2** Choose **System Jobs > Infrastructure**.
 - Step 3** In the Jobs list, check the **Server Backup** check box.
 - Step 4** Click **Edit** (the pencil icon) to open the Edit Job Properties dialog box.
 - Step 5** Create the new local repository using the Edit Job properties dialog box.

- a. Click **Create**. The Create Backup Repository dialog box opens.
- b. Enter the name of the local repository you want to create.
- c. Enter the password if you want to make the backup password secured.
Note Make sure you remember the password to restore the backup.
- d. If it is an FTP repository, check the **FTP** check box and enter the location and credentials.
- e. Click **Submit**. The new repository is added to the Backup Repository drop-down list in the Edit Job Properties dialog box.

Step 6 Click **Save**.

Step 7 If you want to use the repository for future automatic application backups, specify it as described in [Specify the Backup Repository for Automatic Backups, on page 12](#).

Use a Remote Backup Repository

In production environments, we recommend that you use remote repositories for backups so that your network management data is protected from hardware and site failures. In most cases, this means you will need to:

1. Create one or more remote repositories to hold Prime Infrastructure backup files. You will need to set these up yourself if your organization does not already have remote backup servers.
2. Specify the remote repository as the destination for automatic application backups.
3. If needed, specify the interval between automatic application backups and time of day to take them. You will need to monitor and manually archive automatic application backups stored on remote repositories (because the **Max backups to keep** setting does not apply to remote repositories).
4. Specify the remote repository as the backup destination when taking an application or appliance backup using the CLI backup commands.

As with any resource that you plan to access remotely, specifying the correct server IP address and login credentials during setup are a requirement for successful use of remote backup repositories with Prime Infrastructure.

Use Remote NFS Backup Repositories

These topics explain how to use remote NFS backup repositories.

Before You Set Up the NFS Backup Configuration

- You know the IP address of the NFS server on which you want to stage and store backups. The staging and storage folders can be on the same NFS server, or on separate NFS servers. If you plan to stage and store on separate NFS servers, you will need IP addresses for both servers.
- You know the path names of the staging and storage folders on the NFS server. If you choose to stage and store on the same NFS server, the staging and storage folders *must* have different names.

You can create backup repositories on a remote NFS server and configure the Prime Infrastructure server to use them. The NFS server hosting your backups can be set up anywhere in your network, as long as the server:

For the NFS server details to appear in the Backup Repository drop down list in UI, you should configure the NFS server using CLI. You can configure the NFS server only using CLI.

SUMMARY STEPS

1. Open a CLI session with the Prime Infrastructure server. (see [How to Connect Via CLI](#)).
2. Enter configuration mode:
3. Configure a symbolic link to the remote NFS server:
4. Verify creation of the symbolic link:
5. When taking backups at the command line, specify the new repository as the repository name in the backup command. For example:

DETAILED STEPS

Step 1 Open a CLI session with the Prime Infrastructure server. (see [How to Connect Via CLI](#)).

Step 2 Enter configuration mode:

```
PIServer/admin# configure terminal
```

Step 3 Configure a symbolic link to the remote NFS server:

```
pi-system-116/admin# conf t
```

```
pi-system-116/admin(config)# backup-staging-url nfs:// RemoteServerIP:/mnt/stagingfolder
```

```
pi-system-116/admin(config)# repository repositoryName
```

```
pi-system-116/admin(config-Repository)# url nfs:// RemoteServerIP:/mnt/sharefolder
```

```
pi-system-116/admin(config-Repository)# user userName password plain userPassword
```

```
pi-system-116/admin(config-Repository)# end
```

- RepositoryName is the name of the repository (for example: MyRepo or PrimeInfrastructure).
- RemoteServerIP is the IP address of the NFS server hosting the staging backup and shared backup folder. Note that the example above specifies an absolute path to the shared folder.

To specify a relative path to the shared folder, use only one slash in the URL. For example:

```
nfs://RemoteServerIP/sharedfolder
```

- Stagingfolder is the name of the staging backup folder on the NFS server, where the initial data will be transferred temporarily to tar the file later.
- Sharedfolder is the name of the shared backup folder on the NFS server, where the backups will be stored
- UserName is the name of a user with write privileges to the repository on the NFS server.
- UserPassword is the corresponding password for that user.

Step 4 Verify creation of the symbolic link:

```
PIServer/admin# show repository repositoryName
```

Step 5 When taking backups at the command line, specify the new repository as the repository name in the backup command. For example:


```
PIServer/admin# backup MyBackupFileName repository MyRepo application NCS
```

If you want to perform backups automatically, select the repository name you created as the repository name in the Prime Infrastructure web interface.

How to Use Remote SFTP Backup Repositories

You can create backup repositories on a remote SFTP server and configure the Prime Infrastructure server to use them.

The SFTP server hosting your backups can be set up anywhere in your network, as long as the server:

- Has an IP address accessible from the Prime Infrastructure server.
- Has a user with write access to the SFTP server disk.
- Has a local shared folder where the backups will be stored.

Other than these requirements, no other configuration is needed on the SFTP backup server.

We recommend using remote NFS repositories.

For the SFTP server details to appear in the Backup Repository drop down list in UI, you should configure the SFTP server using CLI. You can configure the SFTP server only using CLI.

Step 1 Open a CLI session with the Prime Infrastructure server (see [How to Connect Via CLI](#)).

Step 2 Enter configuration mode:

```
PIServer/admin# configure terminal
```

Step 3 Configure a symbolic link to the remote SFTP server:

```
PIServer/admin(config)# repository repositoryName
```

```
PIServer/admin(config-Repository)# url sftp://RemoteServerIP//sharedfolder
```

```
PIServer/admin(config-Repository)# user userName password plain userPassword
```

```
PIServer/admin(config-Repository)# exit
```

```
PIServer/admin(config)# exit
```

Where:

- repositoryName is the name of the repository (for example: MyRepo or PrimeInfrastructure).
- RemoteServerIP is the IP address of the SFTP server hosting the shared backup folder. Note that the example above specifies an absolute path to the shared folder. To specify a relative path to the shared folder, use only one slash in the URL. For example: **url sftp://RemoteServerIP//sharedfolder**
- sharedfolder is the name of the shared backup folder on the SFTP server.
- userName is the name of a user with write privileges to the repository on the SFTP server.
- userPassword is the corresponding password for that user.

Step 4 Verify creation of the symbolic link:
 PIServer/admin# **s how repository repositoryName**

Step 5 When taking backups at the command line, specify the new repository as the repository name in the backup command. For example:

```
PIServer/admin# backup MyBackupFileName repository MyRepo application NCS
```

If you want to perform backups automatically, select the repository name you created as the repository name in the Prime Infrastructure web interface.

Related Topics

- [Use Remote NFS Backup Repositories](#), on page 7
- [Perform an Immediate Application Backup Using the CLI](#), on page 14
- [Perform an Immediate Appliance Backup Using the CLI](#), on page 13
- [Specify the Backup Repository for Automatic Backups](#), on page 12

How to Use Remote FTP Backup Repositories

You can create backup repositories on a remote FTP server and configure the Prime Infrastructure server to use them.

The SFTP server hosting your backups can be set up anywhere in your network, as long as the FTP server:

- Has an IP address accessible from the Prime Infrastructure server.
- Has a user (FTP user) with write access to the FTP server disk.
- Has a local subdirectory that matches the repository name you specify on the Prime Infrastructure server.
- Has a password of 15 characters or less.

Other than these requirements, no other configuration is needed on the FTP backup server.

We recommend using remote NFS repositories.

Step 1 Open a CLI session with the Prime Infrastructure server (see [How to Connect Via CLI](#)).

Step 2 Enter configuration mode:

```
PIServer/admin# configure terminal
```

Step 3 Configure a symbolic link to the remote FTP server:

```
PIServer/admin(config)# repository repositoryName
```

```
PIServer/admin(config-Repository)# url ftp://RemoteServerIP/sharedfolder
```

```
PIServer/admin(config-Repository)# user userName password plain userPassword
```

```
PIServer/admin(config-Repository)# exit
```

```
PIServer/admin(config)# exit
```

Where:

- repositoryName is the name of the repository (for example: MyRepo or PrimeInfrastructure).
- RemoteServerIP is the IP address of the FTP server hosting the shared backup folder.
- sharedfolder is the name of the shared backup folder on the FTP server.
- userName is the name of a user with write privileges to the repository on the FTP server.
- userPassword is the corresponding password for that user. This password must be 15 characters or less.

Step 4 Verify creation of the symbolic link:

```
PIServer/admin# s how repository repositoryName
```

Step 5 When taking backups at the command line, specify the new FTP repository as the repository name in the backup command. For example:

```
PIServer/admin# backup MyBackupFileName repository MyRepo application NCS
```

If you want to perform backups automatically, select the repository name you created as the repository name in the Prime Infrastructure web interface.

Related Topics

[Use Remote NFS Backup Repositories](#), on page 7

[Perform an Immediate Application Backup Using the CLI](#), on page 14

[Perform an Immediate Appliance Backup Using the CLI](#), on page 13

[Specify the Backup Repository for Automatic Backups](#), on page 12

Delete a Local Backup Repository

Use the following procedure to delete a local backup repository. This procedure ensures that the admin interface has the updated information.

Step 1 Log into the server as a Prime Infrastructure CLI admin user (see [Establish an SSH Session With the Prime Infrastructure Server](#)).

Step 2 List the local application backup repositories and identify the one that you want to delete:

```
show running-config | begin repository
```

Step 3 Enter configuration mode and delete the repository:

```
configure terminal  
(config)# no repository repositoryName
```

Step 4 Repeat step 2 to verify that the repository was deleted.

Set Up Automatic Application Backups

Automatic application backups are enabled by default after installation. You can customize the schedule, specify a different backup repository, or adjust the number of backups that are saved.

To check what data is saved by the backup mechanism (and verify whether you need to manually save any data that is not backed up), see these topics:

- [Information That Is Backed Up, on page 4](#)
- [Information That Is Not Backed Up, on page 6](#)

Schedule Automatic Application Backups

Automatic application backups are enabled by default but you can adjust the day and interval at which these backups are performed. Performing a backup is resource-intensive and affects Prime Infrastructure server performance. Avoid scheduling automatic backups to occur at peak traffic times.

If an automatic application backup fails, Prime Infrastructure generates a Backup Failure alarm (with major severity). You can view these alarms just as you do other alarms .



Note After an automatic application backup fails, a pop-up message is displayed before every subsequent login attempt. This message will continue to appear until you acknowledge the corresponding alarm.

-
- Step 1** Choose **Administration > Dashboards > Job Dashboard**.
 - Step 2** Choose **System Jobs > Infrastructure**.
 - Step 3** In the Jobs list, check the **Server Backup** check box, then click **Edit Schedule**. The Schedule dialog box opens.
 - Step 4** In the Schedule dialog box, select a start date, recurrence interval, and optional end time.
 - Step 5** Click **Submit**. These settings will now be used for future automatic application backups.
-

Specify the Backup Repository for Automatic Backups

You can use the Prime Infrastructure interface to specify a different backup repository for automatic application backups. The backup repository can be local or remote. You can also use the interface to create a new local backup repository if it does not already exist.

Before you begin

If you want to use a remote repository for automatic backups, you must create the repository first. Only local repositories can be created using this procedure. See [Set Up and Manage Repositories, on page 6](#).

-
- Step 1** Choose **Administration > Dashboards > Job Dashboard**.
 - Step 2** Choose **System Jobs > Infrastructure**.
 - Step 3** In the list of jobs, check the **Server Backup** check box.
 - Step 4** Click **Edit** (the pencil icon). The Edit Job Properties dialog box opens.
 - Step 5** Select a repository from the Backup Repository drop-down list, then click **Save**. Prime Infrastructure will use the new repository when it performs the next automatic application backup.
-

Change the Number of Automatic Application Backups That Are Saved

Follow this procedure to adjust the number of automatic application backups that are saved on a local repository. When a backup exceeds the number you specify here, Prime Infrastructure deletes the oldest backup from the repository.

The **Max UI backups to keep** setting does not apply if you are using remote repositories for automatic application backups. You must monitor and archive or delete old backups on remote repositories using your own methods.

-
- Step 1** Choose **Administration > Dashboards > Job Dashboard**.
- Step 2** Choose **System Jobs > Infrastructure**.
- Step 3** In the Jobs list, check the **Server Backup** check box.
- Step 4** Click **Edit** (the pencil icon) to open the Edit Job Properties dialog box.
- Step 5** Enter a value in the **Max UI backups to keep** field, then click **Save**. Prime Infrastructure will enforce this setting at the next backup.
-

Perform a Manual Backup

The topics in this section explain how to perform manual application or appliance backups.

To check what data is saved by the backup mechanism (and verify whether you need to manually save any data that is not backed up), see these topics:

- [Information That Is Backed Up, on page 4](#)
- [Information That Is Not Backed Up, on page 6](#)

Perform an Immediate Appliance Backup Using the CLI

-
- Step 1** Open a CLI session with the Prime Infrastructure server (see [How to Connect Via CLI](#)).
- Step 2** Display the list of appliance backups:

```
PIServer/ (admin) #show repository repositoryName
```

where *repositoryName* is the repository on which you want to store the appliance backup.

- Step 3** Back up the appliance:

```
PIServer/ (admin) #backup filename repository repositoryName
```

where *filename* is the name that you want to give the appliance backup file (for example, myBackup). The character length of the file name is 26. Other information is appended to the filename automatically, as explained in [Backup Filenames, on page 3](#)

Perform an Immediate Application Backup Using the Web GUI

Use this procedure to trigger an immediate application backup using the web GUI.

-
- Step 1** Choose **Administration > Dashboards > Job Dashboard**.
 - Step 2** Choose **System Jobs > Infrastructure**.
 - Step 3** In the Jobs list, check the **Server Backup** check box, then click **Run**.
 - Step 4** To view the backup status, scroll to the top of the table to locate the new job, then check its status and results.
-

Perform an Immediate Application Backup Using the CLI

Use this procedure to trigger an immediate application backup using the CLI.

-
- Step 1** Log into the server as a Prime Infrastructure CLI admin user (see [Establish an SSH Session With the Prime Infrastructure Server](#)).
 - Step 2** Display the list of backups, where *repositoryName* is the backup repository:

```
show repository repositoryName
```

- Step 3** Start the remote backup.

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

You will be prompted to enter the password. Enter the password if you want to secure the backup with password, else press **Enter**. You have to remember the password to restore the backup.

where, *filename* is the name that you want to give the application backup file (for example, myBackup). The character length of the file name is 26. Other information is appended to the filename automatically, as explained in [Backup Filenames, on page 3](#).

Perform a Manual Appliance Backup

Use this procedure to perform an appliance backup to a remote repository.

-
- Step 1** Make sure the remote host is available.
 - Step 2** Log into the Prime Infrastructure server as admin (see [Establish an SSH Session With the Prime Infrastructure Server](#)).
 - Step 3** Start the remote backup:

```
(admin)# backup filename repository repositoryName
```

- Step 4** To verify that the backup transfer is complete, view the md5Checksum and file size.
-

Restore Prime Infrastructure Data

All restore operations are performed using the CLI. Data can be restored to the host where the backup is executed (local host), or to a remote host. Backups can only be restored in their entirety; you cannot restore only parts of a backup.

For more information, see the following topics.

- [Restore an Application Backup, on page 15](#)
- [Restore an Appliance Backup, on page 16](#)



Note If you trigger the restore process from clients such as putty or SSH, the server may not be up even after the restoration is complete—if there was any network issue/failure during the restoration process. To avoid the network issue/failure, it is recommended to perform restore for:

- Gen-2 Appliance—from KVM Console
 - ESXi virtual machine—from VM Console
 - Hyper-V Virtual Machine—from Hyper-V Virtual Machine Connection Console
-

Restore an Application Backup



Note To restore an *appliance* backup, use the procedure in [Restore an Appliance Backup, on page 16](#).

When you restore an Application backup, make sure it is being restored to a machine having equal or higher hardware configuration as that of backup machine, else the restore will fail.

Before you begin

If you are using high availability, read the guidelines in [Remove HA During Restore](#) before restoring your data.

- Step 1** Log into the server as a Prime Infrastructure CLI admin user (see [Establish an SSH Session With the Prime Infrastructure Server](#)).
- Step 2** If a previous restoration attempt failed, the database may have been corrupted. Run this command to recreate the database:
- ```
ncs run reset db
```
- Step 3** List the saved application backups and identify the one that you want to restore. *repositoryName* is the repository that contains the backup files.
- ```
show repository repositoryName
```
- Step 4** From the vmWare vSphere client (OVA) or the Cisco IMC server (Bare Metal), restore the data:

```
restore backupFileName repository repositoryName application NCS
```

Note You will be prompted to enter the password. Enter the password if backup is password protected, else press enter.

Step 5 If you are using Cisco Smart Licensing, re-register Prime Infrastructure with the Cisco Smart Software Manager (CSSM) on Cisco.com. See [Register Prime Infrastructure with the Cisco Smart Software Manager](#).

Restore an Appliance Backup



Note To restore an *application* backup, use the procedure in [Restore an Application Backup, on page 15](#).

When you restore an appliance backup, we recommend that you change:

- The restored server's IP address, if the restored host is on the same subnet as the old host, and the old host is still active.
- The restored server's IP address, subnet mask, and default gateway if the restored host is on a different subnet from the old host.

Before you begin

If you are using high availability, read the information in [Remove HA During Restore](#) before restoring your data.

Step 1 Log into the server as a Prime Infrastructure CLI admin user (see [Establish an SSH Session With the Prime Infrastructure Server](#)).

Step 2 If a previous restoration attempt failed, the database may have been corrupted. With the backup stored in an external repository, reinstall the setup using the same release and then retry the restore.

Step 3 List the saved appliance backups and identify the one that you want to restore. *repositoryName* is the repository that contains the backup files.

```
show repository repositoryName
```

Step 4 From the vmWare vSphere client (OVA) or the Cisco IMC server (Bare Metal), restore the data:

```
restore backupFileName repository repositoryName
```

Step 5 Determine whether you should change the IP address, subnet mask, and default gateway.

a) Check if your installation meets the following criteria:

- The restored host is on the same subnet as the old host, and the old host is still active.
- The restored host is on a different subnet from the old host.

If it does, perform the next step.

b) Change the IP address, subnet mask, default gateway and (optionally) the host name on the restored server.

c) Write the changes to the server's running configuration and restart Prime Infrastructure services. For example:


```
configure terminal
(config)# int GigabitEthernet 0
(config-GigabitEthernet)# ip address IPAddress subnetMask
(config-GigabitEthernet)# exit
(config)# ip default-gateway gatewayIP
(config)# hostname hostname
(config)# exit
(admin)# write mem
(admin)# ncs stop
(admin)# ncs start
(admin)# exit
```

- Step 6** If you are using Cisco Smart Licensing, re-register Prime Infrastructure with the Cisco Smart Software Manager (CSSM) on Cisco.com. See [Register Prime Infrastructure with the Cisco Smart Software Manager](#).
-

Recover from Failed Restores

You may sometimes find that a restore does not complete, or reports a failure. Whenever a restore fails, you run the risk of database corruption, which can prevent the further restoration or re-installation. Perform the following steps to restore a corrupted database before attempting another restore or re-installation.

- Step 1** Open a CLI session with the Prime Infrastructure server (see [Establish an SSH Session With the Prime Infrastructure Server](#)).

- Step 2** Enter the following command to reset the corrupted database:

```
ncs run reset db
```

How to Manage Disk Space Issues During Backup and Restore

If you are experiencing issues with disk space *during* a backup or restore, we suggest that you either:

- Use the VMware **Edit Settings** feature to increase the amount of disk space allocated to the virtual machine (see [Modify VM Resource Allocation Using VMware vSphere Client](#)).

If you are using VMware ESXi 5.5 or later, use the vSphereWeb Client to adjust this setting (see [Configuring Virtual Machine Hardware in the vSphere Web Client](#)).

- Use the method explained in [Migrate to Another Virtual Appliance Using Backup and Restore, on page 18](#) (or [Migrate to Another Physical Appliance Using Backup and Restore, on page 18](#)) to move your installation to a server with adequate disk space.

If you are unable to create a backup *after* a restore of your existing system, follow the steps explained in [Compact the Prime Infrastructure Database](#) to free disk space and create a successful backup.

If you are still unable to create a backup after using the **ncs cleanup** command, set up and use a remote repository (using FTP, SFTP, or NFS) for your backups, as explained in [Use a Remote Backup Repository](#).

Related Topics

[Modify VM Resource Allocation Using VMware vSphere Client](#)

[Migrate to Another Physical Appliance Using Backup and Restore, on page 18](#)

[Migrate to Another Virtual Appliance Using Backup and Restore](#), on page 18
[Compact the Prime Infrastructure Database](#)
[Use a Remote Backup Repository](#), on page 7
[How to Manage Disk Space Issues](#)

Migrate to Another Virtual Appliance Using Backup and Restore

You will need to migrate your Prime Infrastructure data from an existing virtual appliance (OVA server installation) to a new one whenever you want to:

- Replace the old server entirely, such as after a catastrophic hardware failure. In this case, you can use your old installation media to re-create the new host on a replacement server, then migrate your application data from the old host to the new host.
- Migrate to a larger or more powerful server, so you can use Prime Infrastructure to manage more of your network. In this case, you will want to ensure that you have the OVA installation file and install it on the new server using the larger installation option before retiring the older, smaller one. You can then migrate your application data from the old host.

In both cases, it is relatively easy to migrate your old data to the new virtual appliance by restoring to the new host an appliance or application backup taken from the old host.

-
- Step 1** If you have not already done so, set up a remote backup repository for the old host, as explained in [Use a Remote Backup Repository, on page 7](#).
- Step 2** Perform an application backup of the old host and save it to the remote repository (see [Perform an Immediate Application Backup Using the CLI, on page 14](#)).
- Step 3** Install the new host (See [Cisco Prime Infrastructure Quick Start Guide](#)).
- Step 4** Configure the new host to use the same remote backup repository as the old host (see [Use a Remote Backup Repository, on page 7](#)).
- Step 5** Restore the application backup on the remote repository to the new host (see [Restore an Application Backup, on page 15](#)).
-

Migrate to Another Physical Appliance Using Backup and Restore

You will need to migrate your Prime Infrastructure data from an existing physical appliance to a new one whenever you want to:

- Replace the old appliance entirely, such as after a catastrophic hardware failure. In this case, you can order a replacement appliance, then migrate your data from the old appliance to the new appliance.
- Migrate to a newly installed appliance.

In both cases, it is relatively easy to migrate your old data to the new appliance by restoring to the new appliance an appliance or application backup from the old host.

-
- Step 1** If the old appliance is still functional:
- a) If you have not already done so, set up a remote backup repository for the old appliance (see “Use a Remote Backup Repositories” in Related Topics).

- b) Take an appliance or application backup of the old appliance on the remote repository (see “Take Appliance Backups” or “Take Application Backups”, as appropriate).

Step 2 Configure the new appliance to use the same remote backup repository as the old appliance (see “Use a Remote Backup Repositories”).

Step 3 Restore the appliance or application backup on the remote repository to the new appliance (see “Restore From Appliance Backups” or “Restore From Application Backups”, as appropriate). Be sure to follow the procedure appropriate for the type of backup you are restoring. For example: If you took an application backup from the old appliance, you must restore it using the procedure for restoring application backups, not appliance backups.

Related Topics

[Use a Remote Backup Repository](#), on page 7

[Perform an Immediate Application Backup Using the CLI](#), on page 14

[Perform an Immediate Appliance Backup Using the CLI](#), on page 13

[Restore an Appliance Backup](#), on page 16

[Restore an Application Backup](#), on page 15

Backup and Restore with Operations Center

Prime Infrastructure instances running Operations Center and Operations center Server can support restores of application backups taken using the CLI from versions 3.7.X, 3.8.x.

You cannot schedule automatic application backups from the Prime Infrastructure instance running Operations Center.

For more details, see [Use a Remote Backup Repository](#) and [Restore an Application Backup](#).

