

Arcserve® Unified Data Protection Agent for Linux

User Guide

Version 5.0

arcserve®

This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the "Documentation") is for your informational purposes only and is subject to change or withdrawal by Arcserve at any time. This Documentation is proprietary information of Arcserve and may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of Arcserve.

If you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all Arcserve copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to Arcserve that all copies and partial copies of the Documentation have been returned to Arcserve or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, ARCSERVE PROVIDES THIS DOCUMENTATION "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL ARCSERVE BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF ARCSERVE IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice.

The manufacturer of this Documentation is Arcserve.

Provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

© 2015 Arcserve, including its affiliates and subsidiaries. All rights reserved. Any third party trademarks or copyrights are the property of their respective owners.

Contact Arcserve Support

The Arcserve Support team offers a rich set of resources for resolving your technical issues and provides easy access to important product information.

www.arcserve.com/support

With Arcserve Support:

- You can get in direct touch with the same library of information that is shared internally by our Arcserve Support experts. This site provides you with access to our knowledge-base (KB) documents. From here you easily search for and find the product-related KB articles which contain field-tested solutions for many top issues and common problems.
- You can use our Live Chat link to instantly launch a real-time conversation between you and the Arcserve Support team. With Live Chat, you can get immediate answers to your concerns and questions, while still maintaining access to the product.
- You can participate in the Arcserve Global User Community to ask and answer questions, share tips and tricks, discuss best practices and participate in conversations with your peers.
- You can open a support ticket. By opening a support ticket online, you can expect a callback from one of our experts in the product area you are inquiring about.
- You can access other helpful resources appropriate for your Arcserve product.

Contents

Chapter 1: Understanding Arcserve UDP Agent (Linux)	9
Introduction	9
Chapter 2: Installing/Uninstalling Arcserve UDP Agent (Linux)	13
How to Install Arcserve UDP Agent (Linux)	13
Installation Considerations	14
Install Arcserve UDP Agent (Linux).....	14
Verify the Installation.....	17
How to Uninstall Arcserve UDP Agent (Linux).....	17
Review the Uninstallation Considerations	18
Uninstall Arcserve UDP Agent (Linux)	18
Verify the Uninstallation	19
How to Upgrade Arcserve UDP Agent (Linux)	20
Upgrade Considerations.....	20
Upgrade Arcserve UDP Agent (Linux)	21
Verify the Upgrade	23
Chapter 3: User Interface	25
How to Navigate the Arcserve UDP Agent (Linux) User Interface.....	25
Access the Backup Server.....	26
Understanding the Menu Bar	27
Understanding the Status Pane	29
Understanding the Backup Servers Pane	33
Understanding the Help	34
Chapter 4: Using Arcserve UDP Agent (Linux)	37
How to Manage the Licenses	37
Access the License Manager	38
Understanding the License Management Dialog	39
Manage the Licenses.....	40
How to Manage Jobs	41
Review the Prerequisites to Manage Jobs	41
Modify Jobs	42
Cancel Jobs.....	42
Delete Jobs.....	43

How to Back up Linux Nodes	43
Review the Backup Prerequisites and Considerations	45
Do You Want To Back Up More Than 200 Nodes.....	47
Add Linux Nodes for Backup	50
(Optional) Prepare the iSCSI Volume as the Backup Storage	52
Configure the Backup Settings and Run Backup Job	54
Verify the Backup was Successful	73
How to Modify and Rerun a Backup Job	73
Review the Prerequisites for Modifying a Backup Job	75
Do You Want to Add Nodes to an Existing Job.....	75
Add Nodes to an Existing Job	75
Rerun an Existing Backup Job.....	76
Verify the Backup was Successful	77
How to Perform a File-Level Recovery on Linux Nodes	77
Review the Prerequisites	79
(Optional) Recover Data from the iSCSI Volume to the Target Machine	80
Specify the Recovery Point.....	81
Specify the Target Machine Details.....	85
Specify the Advanced Settings	89
Create and Run the Restore Job.....	93
Verify that Files are Restored.....	94
How to Create a Bootable Live CD	94
Review the Live CD Prerequisites.....	95
Install the Restore-Utility Package	96
Create and Verify the Bootable Live CD	96
How to Create a CentOS-Based Live CD	97
Review the Live CD Prerequisites and Considerations.....	98
Install the Restore-Utility Package	100
Create and Verify the CentOS-based Live CD.....	100
How to Perform a Bare Metal Recovery (BMR) for Linux Machines	101
Review the BMR Prerequisites.....	103
Get the IP Address of the Target Machine Using the Live CD	104
(Optional) Recover Data to the iSCSI Volume of the Target Machine	105
(Optional) Recover Data from the iSCSI Volume to the Target Machine	106
Review the Backup Server.....	107
Specify the Recovery Points	108
Specify the Target Machine Details.....	110
Specify the Advanced Settings	111
Create and Run the Restore Job.....	115
Verify that the Target Node is Restored	122
How to Automatically Recover a Virtual Machine	122
Review the Prerequisites and Considerations.....	124

Create a Configuration Template	127
(Optional) Create a Global Configuration File	130
Modify the Configuration Template and File	132
Submit a Job Using the d2drestorevm Utility	132
Verify the VM is Recovered	133
How to Integrate and Automate Arcserve Unified Data Protection Agent for Linux with the Existing IT Environment	133
Review the Automation Prerequisites	135
Understanding the Scripting Utilities	136
Manage Pre/Post Scripts for Automation	145
Create the Backup Storage Alert Script	149
Discover Nodes Using a Script	149
Create the Scripts to Back Up Oracle Database	150
Create the Scripts to Back Up MySQL Database	152
Customize the Job Schedule	154
Run a BMR Batch Job	155
Replicate and Manage Backup Sessions	157
Verify the Recovery Points are Usable	159
How to Manage the Backup Server Settings	164
Review the Prerequisites to Manage the Backup Server	166
Configure the Job History and Activity Log Retention Settings	166
Configure the Debug Log Retention Settings	167
Configure the UI Timeout Duration	168
Change the SSH Port Number of the Backup Server	168
Manage the Recovery Sets	169
Disable the BOOTPD and TFTP Services	170
Improve the Query Performance for Job History and Activity Log	170
Skip CIFS and NFS Module Verification	171
How to Manage the Linux Backup Server from the Command Line	172
Review the Backup Server Prerequisites	173
Start, Stop, or Release the Backup Server	173
Change the Web Service Port Number of the Backup Server	174
Configure the Private Key and Public Key Authentication	175
Change the Backup Server Protocol	176
Avoid the SSL Certificate Error While Opening Arcserve UDP Agent (Linux)	177
Configure the System Settings When the Host Name or IP Address is Changed	179
How to Manage the Non-Root Users	183
Review the Prerequisites	184
Grant Login Permissions to the Non-Root Users	184
Display the Default User in the Login Dialog	185
Enable the Non-Root Users to Add Nodes	186
How to Restore Volumes on a Target Node	187

Review the Prerequisites and Considerations.....	189
Verify the d2drestorevol Utility is Installed	189
Verify the Volume Details in the Session	190
Submit the Volume Restore Job.....	191
Cancel the Volume Restore Job	193
Verify the Restored Volume	193
How to Restore an Oracle Database Using Arcserve UDP Agent (Linux)	193
Perform a Bare Metal Recovery (BMR) of an Oracle Server	195
Perform an Instant Recovery of an Oracle Database	196
Perform Granular Recovery of an Oracle Database	198

Chapter 5: Troubleshooting

203

Chapter 1: Understanding Arcserve UDP Agent (Linux)

This section contains the following topics:

[Introduction](#) (see page 9)

Introduction

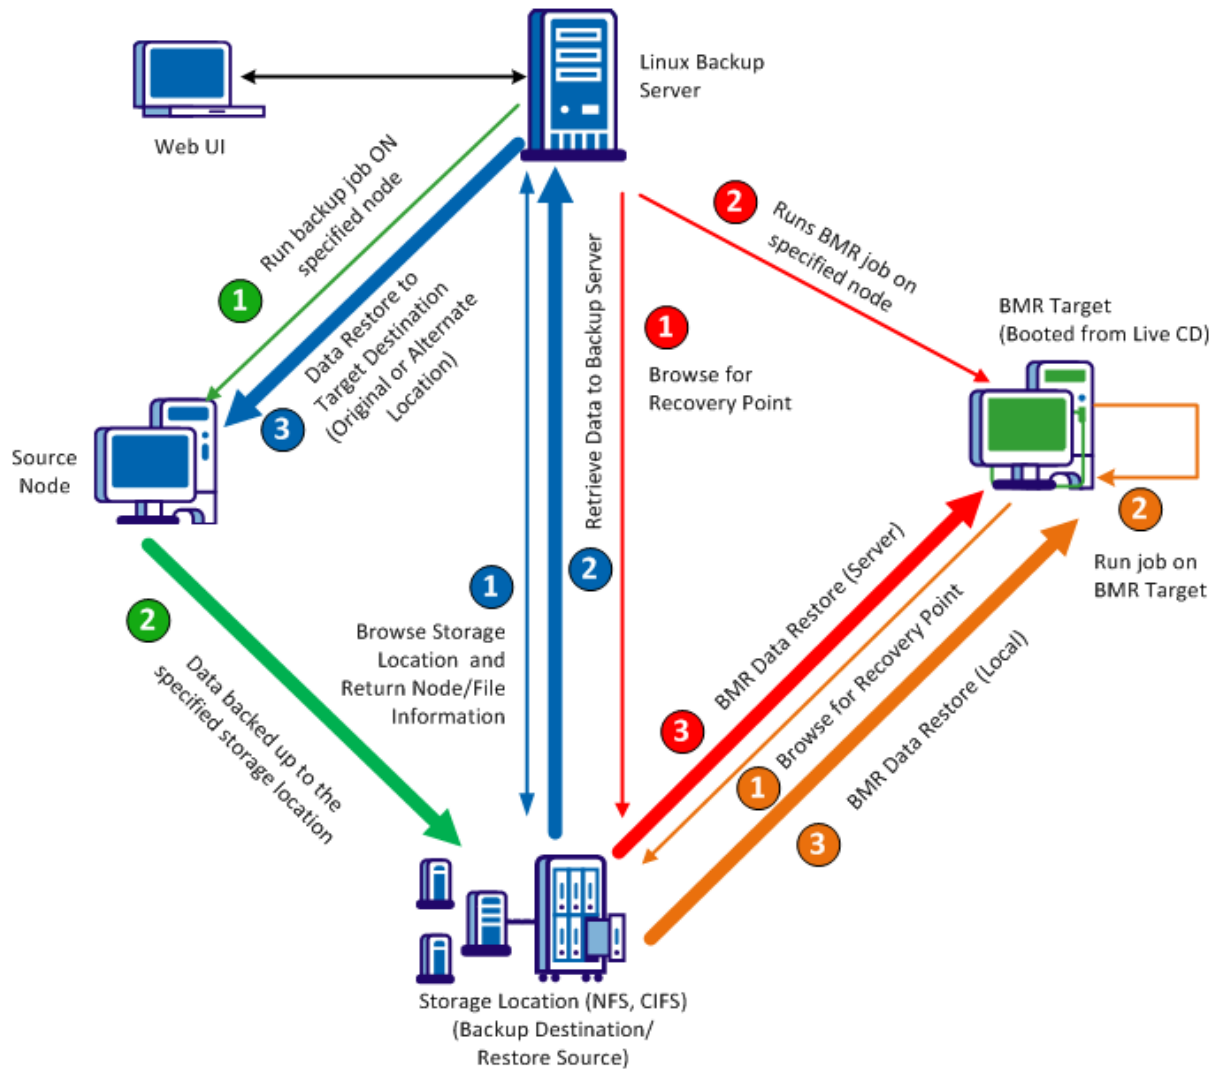
Arcserve Unified Data Protection Agent for Linux (Arcserve UDP Agent (Linux)) is a disk-based backup product that is designed for Linux operating systems. It provides a fast, simple, and reliable way to protect and recover critical business information. Arcserve UDP Agent (Linux) tracks changes on a node at the block level and then backs up only those changed blocks in an incremental process. As a result, Arcserve UDP Agent (Linux) lets you perform frequent backups, reducing the size of each incremental backup (and the backup window) and providing a more up-to-date backup. Arcserve UDP Agent (Linux) also provides the capability to restore files or folders and perform a bare metal recovery (BMR) from a single backup. You can store the backup information either on a Network File System (NFS) share, Common Internet File System (CIFS) share, or in the backup source node.

BMR is the process of restoring a computer system from *bare metal*. Bare metal is a computer without any operating system, drivers, and software applications. The restoration includes installing the operating system, software applications, drivers, and then restoring the data and settings. BMR is possible because when Arcserve UDP Agent (Linux) performs a backup of data, it also captures information that is related to the operating system, installed applications, drivers, and so on. After a BMR is complete, the target node has the same operating system and data as the production node.










Arcserve UDP Agent (Linux) uses a near-agentless approach to enable fast and flexible protection of all your Linux clients. The feature totally eliminates the need to manually install agents on each client node, thereby fully automating the detection, configuration, and protection of all your Linux clients. You can install Arcserve UDP Agent (Linux) to help protect your entire Linux production environment. The server where you install Arcserve UDP Agent (Linux) is known as Backup Server. After you install Arcserve UDP Agent (Linux), you can then connect to the Backup Server over a network and can open the user interface using a web browser.

The following diagram shows the overall work flow of Arcserve UDP Agent (Linux):

Arcserve UDP Agent for Linux Work Flow



Legend

	A machine to browse the Web UI of Arcserve UDP Agent for Linux. This can be a Windows machine.		NFS or NAS where the backed up data is stored
	Linux Backup Server where you install Arcserve UDP Agent for Linux		Backup – Data Flow/Commands
	Linux nodes that you want to back up (Backup Node). <i>Near-agentless backup</i>		Restore (File Level) - Data Flow/Commands
	BMR Target Node where recover data/applications to.		BMR (Server) – Data Flow/Commands
			BMR (Local) – Data Flow/Commands

Chapter 2: Installing/Uninstalling Arcserve UDP Agent (Linux)

This section contains the following topics:

[How to Install Arcserve UDP Agent \(Linux\)](#) (see page 13)

[How to Uninstall Arcserve UDP Agent \(Linux\)](#) (see page 17)

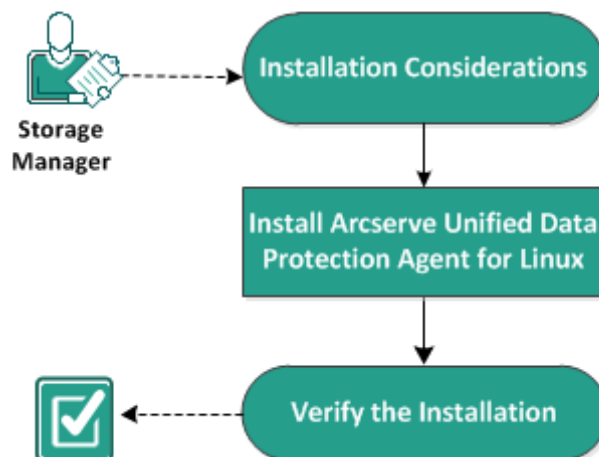
[How to Upgrade Arcserve UDP Agent \(Linux\)](#) (see page 20)

How to Install Arcserve UDP Agent (Linux)

Install Arcserve UDP Agent (Linux) on a Linux server to protect and manage all your backup source nodes from one UI. It is not necessary to install this software on the backup source nodes.

The following diagram displays the process to install Arcserve UDP Agent (Linux):

How to Install Arcserve Unified Data Protection Agent for Linux



Perform these tasks to install Arcserve UDP Agent (Linux):

- [Install Arcserve UDP Agent \(Linux\)](#) (see page 14)
- [Verify the Installation](#) (see page 17)
- [Installation Considerations](#) (see page 14)

Installation Considerations

Consider the following points before you begin the installation:

- When you perform a Preboot Execution Environment (PXE)-based BMR, the Arcserve Unified Data Protection Agent for Linux server and the production source nodes have to be in the same subnet. If they are not in the same subnet, ensure that there is a gateway to forward the PXE broadcast packets across subnets.
- If the backup destination is an NFS server, verify that the NFS server supports *lock*. Also, verify that the root user has write access on the Linux nodes.
- To use an NFS server as the backup destination, install the NFS client package on the Linux nodes.
- Perl and sshd (SSH Daemon) are installed on the Linux server and the Linux nodes that you want to back up.
- The unattended or silent installation is not supported.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Install Arcserve UDP Agent (Linux)

Install Arcserve UDP Agent (Linux) on a Linux server to manage backup and restore operations. After you install Arcserve UDP Agent (Linux), you can open the user interface from any computer using a web browser and the server is referred as Backup Server.

At the beginning of installation, the installation script verifies if some of the mandatory applications are installed on the Linux server and the applications are running.

The following mandatory applications are required for the installation file to work:

- sshd (SSH Daemon)
- Perl

The installation file also verifies the following optional applications at the beginning of the installation:

- rpc.statd - This application is used by the NFS server to implement the file lock.
- mkisofs - Arcserve UDP Agent (Linux) uses this application to create a Live CD.
- mount.nfs - Arcserve UDP Agent (Linux) uses this application to mount the NFS server.

- mount.cifs - Arcserve UDP Agent (Linux) uses this application to mount the CIFS server.
- ether-wake - Arcserve UDP Agent (Linux) uses this application to send the wake-on-lan request.

Note: Ensure that the Linux server has a minimum 1-GB memory. For more information about the system requirements for a Linux server, see the *Release Notes*.

Follow these steps:

1. Log in to the Linux server as a root user.
2. Download the Arcserve UDP Agent (Linux) installation package (*.bin file) and the restore-utility package file to the root folder.

Important! When you download the two installation package files to a local folder, the full path of this local folder must not contain any special characters except blank spaces and the path should only include the following characters: a-z, A-Z, 0-9, -, and _.

3. Provide the execution permission to the installation package.

4. Perform one of the following steps depending on the location of the installation package and the restore-utility package:

- If the installation package and the restore-utility package are in the same folder, then run the following command to begin the installation:

```
./<linux_installation_file_name>.bin
```

Note: If you rename the restore-utility package, the package name must include the restore utility characters for the installation command to automatically find the restore-utility package and install it. If the package name does not have the restore utility characters, then you must provide the full path of the restore-utility package.

The installation package verifies the supported platform and displays a confirmation message.

If a non-supported platform is detected, type Y and press Enter to confirm the non-supported platform installation.

- If the installation package and the restore-utility package are in different folders, then provide the path of the restore-utility package in the first parameter:

```
./<linux_installation_file_name>.bin --path=<path_of_the_restore-utility_package>
```

The installation package verifies the supported platform and displays a confirmation message.

If a non-supported platform is detected, type Y and press Enter to confirm the non-supported platform installation.

Note: If a non-English operating system is detected, you are prompted to select the applicable language before continuing with the installation process.

5. Type Y and press Enter to confirm the installation.

The installation package displays the licensing agreement information.

6. Type Y and press Enter to accept the licensing agreement.

The Arcserve UDP Agent (Linux) installation process begins.

When the restore-utility package installation is complete, the Live CD build information is displayed.

The Live CD is built at the following location:

```
/opt/CA/d2dserver/packages
```

Note: Live CD is required to get the IP address of the target node when you perform a Bare Metal Recovery (BMR).

Arcserve UDP Agent (Linux) is installed and the URL to browse the Linux Backup Server is displayed.

Note: Ensure that the following incoming ports are enabled on your firewall for the Backup Server:

- TCP port 22 (SSH Server)
- Broadcast port 67 (Boot Server)
- 8014 (Agent Web Service)
- User Datagram Protocol (UDP) port 69 (TFTP Server)

Ensure that the following incoming port is enabled on your firewall for the client nodes that you want to back up:

- TCP port 22 (SSH Server)

Ensure that the required outgoing port for NFS, CIFS, or both backup destinations are enabled on your firewall for the Linux Backup Server and client nodes.

Arcserve UDP Agent (Linux) is successfully installed.

Verify the Installation

Verify that the installation is complete after you have installed Arcserve UDP Agent (Linux).

Follow these steps:

1. Open a web browser from any Windows computer.
2. Enter the URL of the Linux Backup Server that is displayed on the install screen.

Example: `https://hostname:8014`

The Arcserve UDP Agent (Linux) login page opens.

3. Enter your root login credentials and click Login.

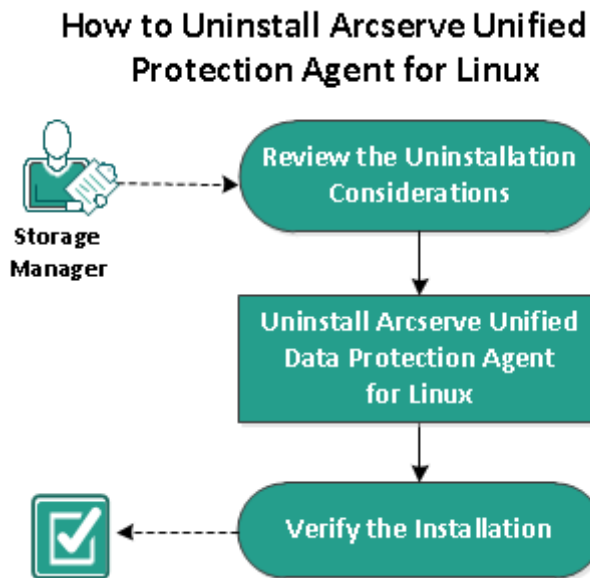
The Arcserve UDP Agent (Linux) user interface opens.

Arcserve UDP Agent (Linux) is successfully installed and verified.

How to Uninstall Arcserve UDP Agent (Linux)

Uninstall Arcserve UDP Agent (Linux) from the Linux Backup Server to stop protecting all your nodes.

The following flowchart shows the uninstallation process of Arcserve UDP Agent (Linux):



Perform these tasks to uninstall Arcserve UDP Agent (Linux):

- [Review the Uninstallation Considerations](#) (see page 18)
- [Uninstall Arcserve UDP Agent \(Linux\)](#) (see page 18)
- [Verify the Uninstallation](#) (see page 19)

Review the Uninstallation Considerations

Consider the following points before you begin the uninstallation:

- You have the root login credentials to the Backup Server.
- You do not have any running jobs. If a job is running, you cannot uninstall Arcserve UDP Agent (Linux).
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Uninstall Arcserve UDP Agent (Linux)

You can uninstall Arcserve UDP Agent (Linux) from the command line of the Backup Server. The uninstall process removes all the files and directories that are created during the installation of the software.

Follow these steps:

1. Log in to the Backup Server as a root user.
2. Navigate to the *bin* folder where Arcserve Unified Data Protection Agent for Linux is installed with the following command:

```
# cd /opt/CA/d2dserver/bin/
```

3. Run the following command to uninstall Arcserve UDP Agent (Linux):

```
# ./d2duninstall
```

A message displays after the uninstallation is complete.

Arcserve UDP Agent (Linux) is uninstalled from the server.

Verify the Uninstallation

Verify that the Arcserve UDP Agent (Linux) is removed from the server after the uninstallation process is complete.

Navigate to the following folder and verify that Arcserve UDP Agent (Linux) is removed:

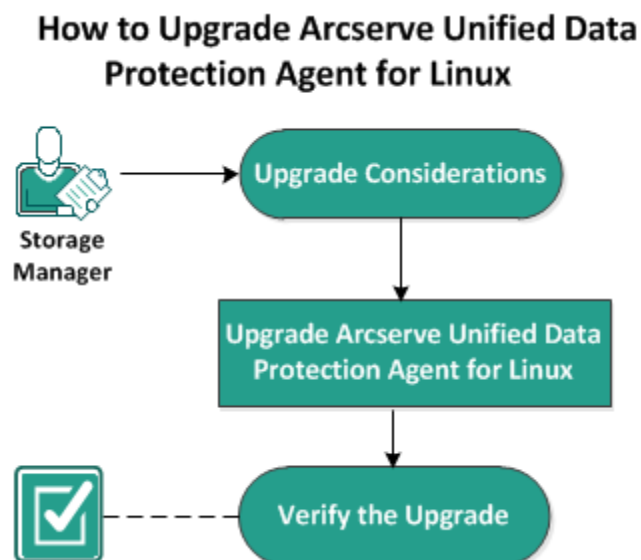
```
/opt/CA/d2dserver
```

You have verified the uninstallation of Arcserve UDP Agent (Linux). Arcserve UDP Agent (Linux) is removed from the Linux server.

How to Upgrade Arcserve UDP Agent (Linux)

Upgrade Arcserve UDP Agent (Linux) to the next release to avail several modifications and enhancements on the features and performance of Arcserve UDP Agent (Linux).

The following diagram displays the process to upgrade Arcserve UDP Agent (Linux):



Perform these tasks to upgrade Arcserve UDP Agent (Linux):

- [Upgrade Considerations](#) (see page 20)
- [Upgrade Arcserve UDP Agent \(Linux\)](#) (see page 21)
- [Verify the Upgrade](#) (see page 23)

Upgrade Considerations

Consider the following points before you begin the upgrade:

- Ensure that you schedule your upgrade when there are no backup jobs running.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Upgrade Arcserve UDP Agent (Linux)

Upgrade Arcserve UDP Agent (Linux) to the next release to avail several modifications and enhancements on the features and performance of Arcserve UDP Agent (Linux).

When you install the upgrade, Arcserve UDP Agent (Linux) tries to detect an existing installation.

- If Arcserve UDP Agent (Linux) detects an existing installation, it automatically performs the upgrade process. All existing configurations (e.g. configuration files, database) are saved and upgraded.
- If Arcserve UDP Agent (Linux) does not detect any existing installation, it automatically performs a new installation.

Follow these steps:

1. Log in to the Backup Server as a root user.
2. Download the Arcserve UDP Agent (Linux) installation package (*.bin file) and the restore-utility package file to the root folder.
Important! When you download the two installation package files to a local folder, the full path of this local folder must not contain any special characters except blank spaces and the path should only include the following characters: a-z, A-Z, 0-9, - and _.
3. Provide the execution permission to the installation package.

4. Perform one of the following steps depending on the location of the installation package and the restore-utility package:

- If the installation package and the restore-utility package are in the same folder, then run the following command to begin the installation:

```
./<linux_installation_file_name>.bin
```

Note: If you rename the restore-utility package, the package name must include the restore utility characters for the installation command to automatically find the restore-utility package and install it. If the package name does not have the restore utility characters, then you must provide the full path of the restore-utility package.

The installation package verifies the supported platform and displays a confirmation message.

If a non-supported platform is detected, type Y and press Enter to confirm the non-supported platform installation.

- If the installation package and the restore-utility package are in different folders, then provide the path of the restore-utility package in the first parameter:

```
./<linux_installation_file_name>.bin --path=/<path_of_the_restore-utility_package>
```

The installation package verifies the supported platform and displays a confirmation message.

If a non-supported platform is detected, type Y and press Enter to confirm the non-supported platform installation.

The installation package detects an existing installation and displays a confirmation message for upgrade.

5. (Optional) Type Y and press Enter to confirm application dependencies.

The installation package verifies application dependencies.

6. Type Y and press Enter to confirm the installation.

The installation package displays the licensing agreement information.

7. Type Y and press Enter to accept the licensing agreement.

The Arcserve UDP Agent (Linux) installation process begins.

When the restore-utility package installation is complete, the Live CD build information is displayed.

The Live CD is built at the following location:

```
/opt/CA/d2dserver/packages
```

Note: Live CD is required to get the IP address of the target node when you perform a Bare Metal Recovery (BMR).

Arcserve UDP Agent (Linux) is successfully upgraded.

Verify the Upgrade

Verify that the upgrade is complete after you have upgraded Arcserve UDP Agent (Linux) to the next release. Backup Server stores a backup of the existing configurations files. After the verification is complete, delete the backup of the existing configurations files.

Follow these steps:

1. Open a web browser from any Windows computer.
2. Enter the URL of the Backup Server.
Example: `https://hostname:8014`
The Arcserve UDP Agent (Linux) login page opens.
3. Enter your root login credentials and click Login.
The Arcserve UDP Agent (Linux) user interface opens.
4. Verify that Backup Server is working properly.
5. Log into the Backup Server as a root user.
6. Navigate to the `d2dserver.bak` folder and delete the folder.

`/opt/CA/d2dserver.bak`

Arcserve UDP Agent (Linux) is successfully upgraded and verified.

Chapter 3: User Interface

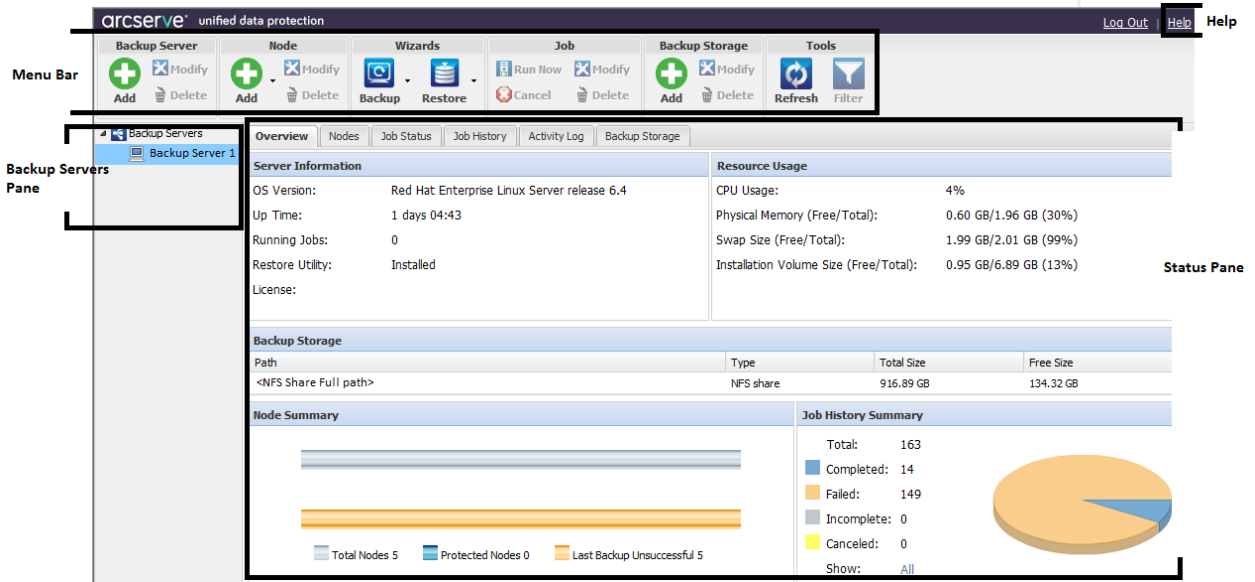
This section contains the following topics:

[How to Navigate the Arcserve UDP Agent \(Linux\) User Interface](#) (see page 25)

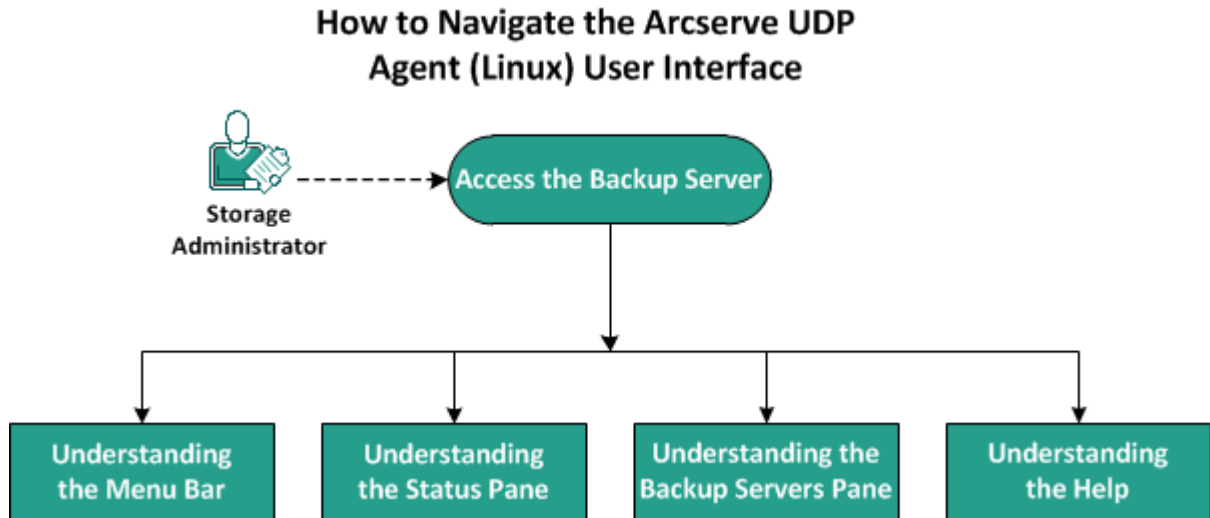
How to Navigate the Arcserve UDP Agent (Linux) User Interface

Before you start using Arcserve UDP Agent (Linux), you must be familiar with the user interface (UI). From the interface, you can manage nodes, manage backup storage locations, manage backup and restore jobs, and access the help topics.

The homepage interface includes four main areas: Menu bar, Status pane, Backup Servers pane, and Help.



The following diagram displays the process to navigate the Arcserve UDP Agent (Linux) interface:



Perform these tasks to get started with the Backup Server interface:

- [Access the Backup Server](#) (see page 26)
- [Understanding the Menu Bar](#) (see page 27)
- [Understanding the Status Pane](#) (see page 29)
- [Understanding the Backup Servers Pane](#) (see page 33)
- [Understanding the Help](#) (see page 34)

Access the Backup Server

As a storage manager, you can access the Backup Server using the web interface. Log in with root or non-root credentials to access the Backup Server. Use the IP address that was received during the installation of Arcserve UDP Agent (Linux) to log in to the server. If you have recorded the host name of the server, you can log in to the server using that host name.

Note: For more information about providing the login permission to the non-root users, see Grant Login Permissions to the Non-Root Users.

Follow these steps:

1. Open a web browser and type the IP address of the Backup Server.

Note: By default, the Backup Server follows https protocol and uses the 8014 port.

2. Enter the login credentials and click Login.

The Backup Server interface opens.


The Backup Server is successfully accessed.

Understanding the Menu Bar

The menu bar lets you perform the following tasks:

- Manage Backup Servers
- Manage nodes
- Manage backup jobs
- Manage restore jobs
- Manage backup storage locations
- Filter searches
- Refresh pages

The following screen displays the menu bar:



The menu bar includes the following options:

Backup Server

Lets you add, modify, and delete servers that have Arcserve UDP Agent (Linux) installed. You can install Arcserve UDP Agent (Linux) to multiple servers and can manage all installed servers from a central UI. The nodes that are managed by the selected server are displayed in the Status pane. All the added servers are displayed in the Backup Servers pane. You cannot modify or delete the central server. A central server is the first server that is displayed in the Backup Servers pane. You can modify and delete other servers from the Backup Servers pane. The Modify button lets you update only the Port Number of the servers.

Node

Lets you add, modify, and delete nodes that you want to back up. Nodes are the machines that you want to back up. You can add multiple nodes to be backed up. You can also discover nodes that are present in your network using a script. You can add the maximum of 200 nodes for each server.

If you delete a node, the Backup Server clears all information about the node from the database, including the backup job information. The Backup Server also deletes the drivers from the node. It may take some time to delete the drivers completely.

Wizards

Lets you launch the Backup Wizard and the Restore Wizard to help guide you through the backup and restore process.

- The Backup Wizard contains a drop-down list with three available options:

Back Up

Use this option if you have not previously added any nodes to be backed up. Selecting this option launches the Backup Wizard and lets you add your nodes during the process.

Back Up Selected Nodes

Use this option if you have previously added your nodes before launching the Backup Wizard. If you click Back Up Selected Nodes without adding any nodes or selecting the existing nodes, you get an error message. To avoid this error, select the node from the Nodes tab and then select Back Up Selected Nodes.

Add Selected Nodes to an Existing Job

Use this option if you have an existing backup job and you want to apply the same backup settings to new nodes. You do not have to configure the Backup Wizard.

- The Restore Wizard contains a drop-down list with two available options:

Bare Metal Recovery (BMR)

Use this option to perform a BMR. You can perform a BMR using the IP address or MAC address of the bare-metal computer to be recovered.

Restore File

Use this option to perform a file-level restore. You can select specific files from a recovery point and restore those files.

Job

Lets you manage jobs that you create. A job is an instance of a backup or a restore operation. After you create a backup job for a node, you do not have to create another job to run a backup for the same node next time. However, you have to create a restore job each time you want to perform a BMR.

Backup Storage

Lets you add and manage the backup storage locations. The backup storage location could be Network File System (NFS) share, Common Internet File System (CIFS) share, or Local. Local is a local path in the Backup Server.

When you add a backup storage location, you have to provide your credentials for the selected backup storage location. You can only modify the user name and password of the CIFS share. You cannot modify any details of the NFS share. Select the Run script when free space is less than checkbox to run the *backup_storage_alert.sh* script when the free space is less than the specified value. This value can be a percentage of the total space at the backup destination or a minimum amount of space (in MB) at the backup destination. The *backup_storage_alert.sh* script can be configured to send an alert when the available free space becomes less than the specified value.

Note: For more information about configuring the *backup_storage_alert.sh* script, see *How to Integrate and Automate Arcserve UDP Agent (Linux) with the Existing IT Environment*.

After you add a backup storage location, you can view the corresponding total file size and the empty space in the Status pane. Select a backup storage location to see the recovery sets and recovery points, and the used space for each node that are backed up in that backup storage location. The added storage destinations are also displayed in the Backup Destination page of the Backup Wizard and in the Recovery Points page of the Restore Wizard.

Tools

The tools menu includes the Refresh button and the Filter button.


Refresh

Lets you refresh the selected display area in the Status pane, including the Activity Log to view the latest backup or restore status messages.

Filter

Lets you filter information displayed in the Status pane based on your input. The Filter button acts like a switch so that you can show and hide filters using the same button. When you show filters, the search fields are displayed in the Status pane. When you hide filters, the search fields are removed from the Status pane.

The following screen displays the filters applied to the Activity Log:



Understanding the Status Pane

The Status pane is the area that displays all the information in the UI. The Status pane includes six tabs that let you view information based on the selected tab.

The following screen displays the Status pane:

The screenshot shows the Status pane with the following sections:

- Overview** (selected tab): Overview, Nodes, Job Status, Job History, Activity Log, Backup Storage
- Server Information**: OS Version: Red Hat Enterprise Linux Server release 6.4; Up Time: 1 days 05:30; Running Jobs: 0; Restore Utility: Installed; License: (empty)
- Resource Usage**: CPU Usage: 0%; Physical Memory (Free/Total): 0.58 GB/1.96 GB (29%); Swap Size (Free/Total): 1.99 GB/2.01 GB (99%); Installation Volume Size (Free/Total): 0.95 GB/6.89 GB (13%)
- Backup Storage**: Table with columns Path, Type, Total Size, Free Size. Row: <NFS Share full path>, NFS share, 916.89 GB, 134.32 GB
- Node Summary**: Horizontal bar chart showing Total Nodes 5 (blue), Protected Nodes 0 (green), Last Backup Unsuccessful 5 (red)
- Job History Summary**: Total: 163; Completed: 14; Failed: 149; Incomplete: 0; Canceled: 0. Includes a pie chart and a 'Show: All' link.

The Status pane includes the following tabs:

Overview

Provides a summary of the following items:

Server Information

Displays the operating system version, time elapsed since the server started, and the licensing information for Arcserve UDP Agent (Linux). It also displays whether the restore utility is installed on this server.

Resource Usage

Displays the usage of CPU, total and available physical memory, and swap size. It also displays the installation volume size.

Backup Storage

Displays all backup session locations that you have added and the available space in each location. This information helps you plan your next backup location depending on the available storage space.

Node Summary

Displays a graphical representation of the nodes that are protected and nodes with last unsuccessful backups. Node Summary includes the following categories:

Total Nodes display the number of nodes that are included in Arcserve UDP Agent (Linux), regardless of the backup status.

Protected Nodes displays the number of nodes that the most recent backup was successful and are considered protected in case a recovery is necessary.

Last Backup Unsuccessful displays the number of nodes that the most recent backup was not successful (failed, canceled, incomplete). Depending on the cause of the unsuccessful backup, some of these nodes are unprotected in case a recovery is necessary.

Job History Summary

Displays a pie chart that summarizes the history of all jobs. The summary does not include the running jobs.

The following fields are not self-explanatory:

Incomplete displays the number of jobs that ran successfully with minor changes. For example, when you restore files from Red Hat 6 to Red Hat 5, the files are restored successfully but some attributes are missing in the restored files.

Other displays the number of jobs that you canceled.

Nodes

Displays all nodes that you have added to the Backup Server. You can apply filters to the Nodes tab to search for the required nodes. The Nodes tab also includes a context menu. The context menu lets you search the job status or the job history for the selected node. The context menu also lets you restore data. You can filter the job history or the job status using either the job name or the node name. If you search the job history for the selected node, then the Job History tab opens with the search filter applied to the tab. Similarly, if you search the job status, then the Job Status tab opens with the search filter applied to the tab. The Restore option lets you perform BMR or File-level restore. It opens the Restore Wizard and displays all the recovery points of the selected node.

Overview	Nodes	Job Status	Job History	Activity Log	Backup Storage	
Node Name	Username	Backup Job	Recovery Point Count	Last Result	Operating System	Description
Node 1	root			✖	Oracle Linux Server release 6.1	
Node 2	root			N/A	Oracle Linux Server release 6.1	
Node 3	root			N/A	SUSE Linux Enterprise Server 11 SP2	
Node 4	root			✖	Oracle Linux Server release 6.4	

Job Status

Displays the list of backup and restore jobs that are created, including the status of each job. Use this tab to run a backup or restore job and rerun a backup job. You can see the progress of backup or restore jobs that you run. You can apply filters to the Job Status tab to search for the required jobs. The Job Status tab also includes a context menu. The context menu lets you search the job history for the selected job. You can filter the job history using either the job name or the node name. If you search the job history for the selected job, then the Job History tab opens with the search filter applied to the tab.

The following screen displays the context menu in the Job Status tab:

Overview	Nodes	Job Status	Job History	Activity Log	Backup Storage
Job Name	Job Id	Job Type	Node Name	Job Phase	Status
Backup - 7/1/2013 8:33:00 PM		Backup			Ready

Job History

Displays the list of backup and restore jobs that were previously run. You can apply filters to the Job History tab to search for the required job history. When you select a job, the status of that job is displayed at the bottom of the page.

Activity Log

Displays a list of processing messages and status messages for backup and restore jobs. Refresh the Activity Log to get the latest messages for recent backup and restore jobs. You can apply filters to the Activity Log tab to search for required activity logs.

Backup Storage

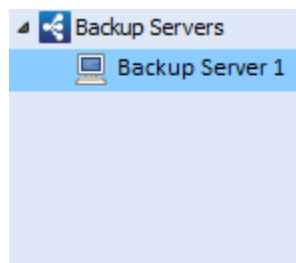
Displays the backup destination that you have added from the menu bar. You can view the free storage space and manage your backup destination. This option is useful if you want to know the available free space at any particular backup destination to plan your backup. When you add a storage destination, this destination appears in the Backup Wizard.

Understanding the Backup Servers Pane

The Backup Servers pane displays the list of Backup Servers that are managed by the current server. You can add servers from the menu bar and can manage all the servers from one interface. If you have added multiple servers, then the Status pane displays the status of the selected server. Each server can manage at least 200 client nodes.

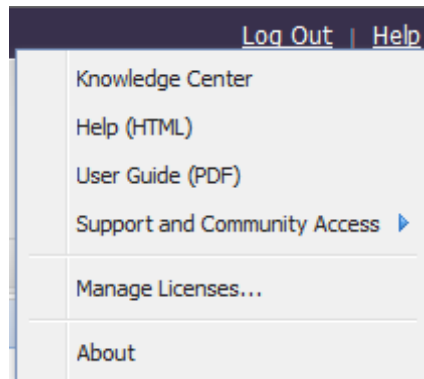
Typically the first server displayed in the Backup Servers pane is the central Backup Server and other servers are member servers. If you are managing multiple servers from a central server then verify that the version of the central server and member servers are same.

The following screen displays the Backup Servers pane:



Understanding the Help

The Help dialog lets you access the Help topics of Arcserve UDP Agent (Linux). You can perform the following tasks from the Help dropdown list:



The following options are available in the Help dropdown list:

Knowledge Center

Lets you access the bookshelf.

Help (HTML)

Lets you access the HTML version of the Arcserve UDP Agent (Linux) User Guide.

User Guide

Lets you access the PDF version of the Arcserve UDP Agent (Linux) User Guide.

Support and Community Access

Lets you access the Arcserve UDP Agent (Linux) support site and ARCserve community sites. You can perform the following tasks from Support and Community Access:

- View product specific
- Access the ARCserve official website
- Provide feedback to the development team
- Read tips and comments from the ARCserve experts

- Chat with a support executive
- Subscribe to RSS feeds

Manage License

Lets you access the License Management dialog and manage all your licenses from a central interface.

About

Lets you view the product information (version number and build number) and access the Release Notes of Arcserve UDP.

Chapter 4: Using Arcserve UDP Agent (Linux)

This section contains the following topics:

[How to Manage the Licenses](#) (see page 37)

[How to Manage Jobs](#) (see page 41)

[How to Back up Linux Nodes](#) (see page 43)

[How to Modify and Rerun a Backup Job](#) (see page 73)

[How to Perform a File-Level Recovery on Linux Nodes](#) (see page 77)

[How to Create a Bootable Live CD](#) (see page 94)

[How to Create a CentOS-Based Live CD](#) (see page 97)

[How to Perform a Bare Metal Recovery \(BMR\) for Linux Machines](#) (see page 101)

[How to Automatically Recover a Virtual Machine](#) (see page 122)

[How to Integrate and Automate Arcserve Unified Data Protection Agent for Linux with the Existing IT Environment](#) (see page 133)

[How to Manage the Backup Server Settings](#) (see page 164)

[How to Manage the Linux Backup Server from the Command Line](#) (see page 172)

[How to Manage the Non-Root Users](#) (see page 183)

[How to Restore Volumes on a Target Node](#) (see page 187)

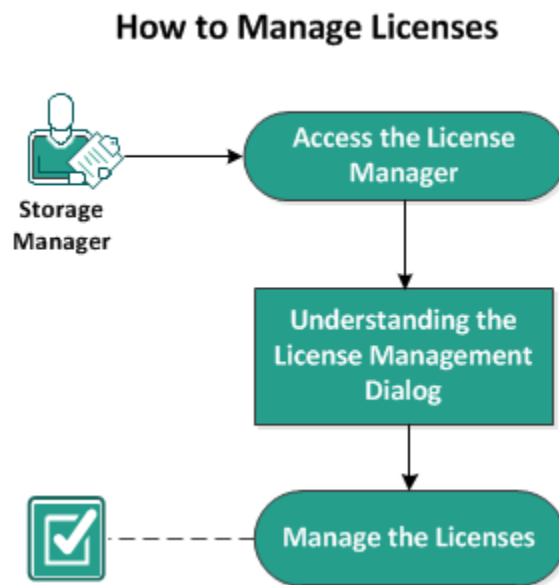
[How to Restore an Oracle Database Using Arcserve UDP Agent \(Linux\)](#) (see page 193)

How to Manage the Licenses

Arcserve UDP Agent (Linux) requires you to license your product to receive authorized and uninterrupted access to the related components. In addition, if you want to deploy Arcserve Unified Data Protection Agent for Linux to remote locations, you must license these remote sites to take advantage of the benefits Arcserve UDP Agent (Linux) provides.

Arcserve UDP Agent (Linux) will function for a trial period of 30 days after you begin using it. Then, apply an appropriate license key to continue using it. Arcserve UDP Agent (Linux) lets you manage the licenses for all of your Linux Backup Servers from a central interface.

The following diagram displays the process to manage licenses:



Complete the following tasks to manage the licenses:

- [Access the License Manager](#) (see page 38)
- [Understanding the License Management Dialog](#) (see page 39)
- [Manage the Licenses](#) (see page 40)

Access the License Manager

You must access the License Management dialog from the Arcserve UDP Agent (Linux) web interface to manage all your licenses.

Follow these steps:

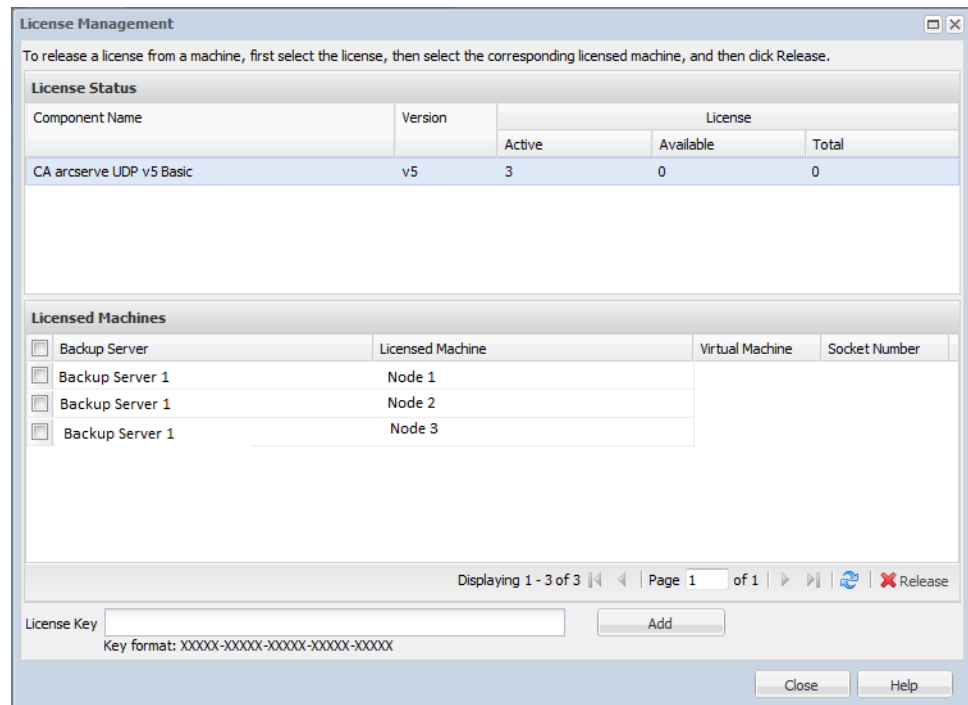
1. Log in to the Arcserve UDP Agent (Linux) web interface.
2. From the home page, click Help, Manage License.
The License Management dialog opens.

The license manager is accessed.

Understanding the License Management Dialog

The License Management dialog lets you manage all your licenses for Arcserve UDP Agent (Linux). You can manage the licenses for multiple Linux Backup Servers from a single interface.

The following screen displays the License Management dialog:



The License Management dialog is divided into two sections: License Status and Licensed Machines.

License Status

Component Name

Identifies the name of the license.

Version

Identifies the release number of the license.

Active

Identifies the number of licenses that are currently being used to back up the nodes.

Available

Identifies the number of licenses that are still available in the license pool and that can be used to back up Linux machines.

Total

Identifies the total number of licenses that have been obtained to back up the machine. Total is the sum of Active and Available licenses.

Licensed Machines

Backup Server

Identifies the Linux server where you have installed Arcserve UDP Agent (Linux).

Licensed Machines

Identifies the Linux machines for which you have applied a license to protect those machines.

Manage the Licenses

You can add and release licenses from the License Management dialog. The added license is displayed in the License Management dialog. If you do not want to backup any machine anymore, you can release the license from that machine.

Follow these steps:

- To add a license, follow these steps:
 - a. Check the license key on your media case or on your License certificate.
 - b. Enter the license key in the License Key field and click Add.
 - c. Close and open the License Management dialog.

The license is added and is listed in the License Status area.
- To release a license, follow these steps:
 - a. Select the license from the License Status area.
 - b. Select the Backup Server from Licensed Machines and click Release.
 - c. Close and open the License Management dialog.

The license is released from the machine.

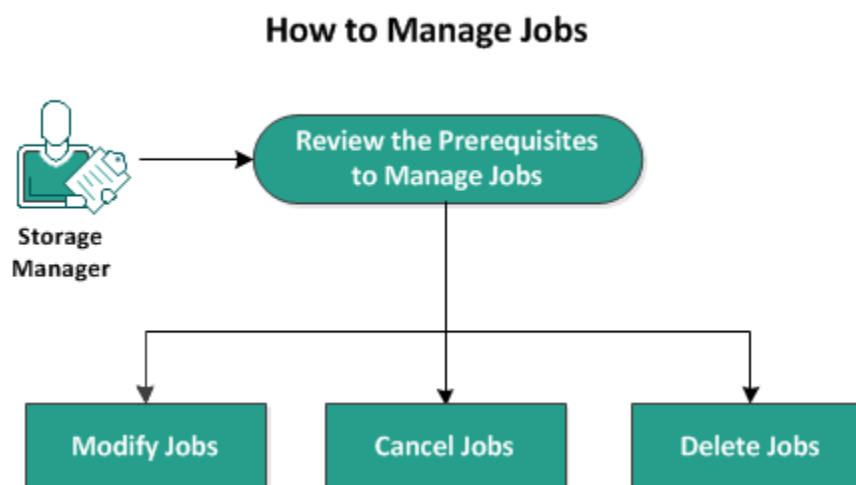
The licenses are successfully managed.

How to Manage Jobs

After you create a backup or a restore job, you can manage all your jobs from the Job menu. Managing a job includes the following tasks:

- Modifying a job
- Canceling a job
- Deleting a job

The following diagram displays the process to manage jobs:



Perform these tasks to manage your jobs:

- [Review the Prerequisites](#) (see page 41)
- [Modify Jobs](#) (see page 42)
- [Cancel Jobs](#) (see page 42)
- [Delete Jobs](#) (see page 43)

Review the Prerequisites to Manage Jobs

Consider the following prerequisites before you manage your jobs:

- You have a valid existing job to manage
- You have the appropriate permission to manage jobs.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Modify Jobs

You can open any existing job and modify the settings for the job from the web interface. For example, if you want to change the backup destination for an already protected machine, you do not have to create a new job. You can open the existing job that protects the machine and modify only the backup destination section. Your all other settings remain unchanged except the backup destination settings.

Follow these steps:

1. Select a job from the Job Status tab.
2. Click Modify from the Job menu.

The wizard for the selected job opens.

3. Modify your settings in the wizard.
4. Click Submit on the Summary page of the wizard.

The job is submitted and the job runs depending on your settings.

The job is successfully modified.

Cancel Jobs

You can cancel a running job from the web interface of Arcserve UDP Agent (Linux).

Follow these steps:

1. Select a job from the Job Status tab.
2. Click Cancel from the Job menu.

The Cancel job dialog opens.

3. Select one of the following options from the Cancel job for dropdown list:

Selected node

Specifies that the job is canceled only for the selected node.

All nodes protected by the selected job

Specifies that the job is canceled for all the nodes protected by the selected job.

4. Click OK.

The job is canceled.

Delete Jobs

You can delete a job when you do not want to protect or restore a machine anymore. You can also delete a job that protects a group of nodes. When you delete a job, the previously backed up recovery points still remain available in the specified backup destination. You can use those recovery points to restore your data.

For a running job, the Delete option is inactive. You have to cancel the running job and then delete the job.

Follow these steps:

1. Select a job from the Job Status tab.
2. Click Delete from the Job menu.

The Delete job dialog opens.

3. Select one of the following options from the Delete job for dropdown list:

Selected node

Specifies that the job is deleted only for the selected node.

All nodes protected by the selected job

Specifies that the job is deleted for all the nodes protected by the selected job.

4. Click OK.

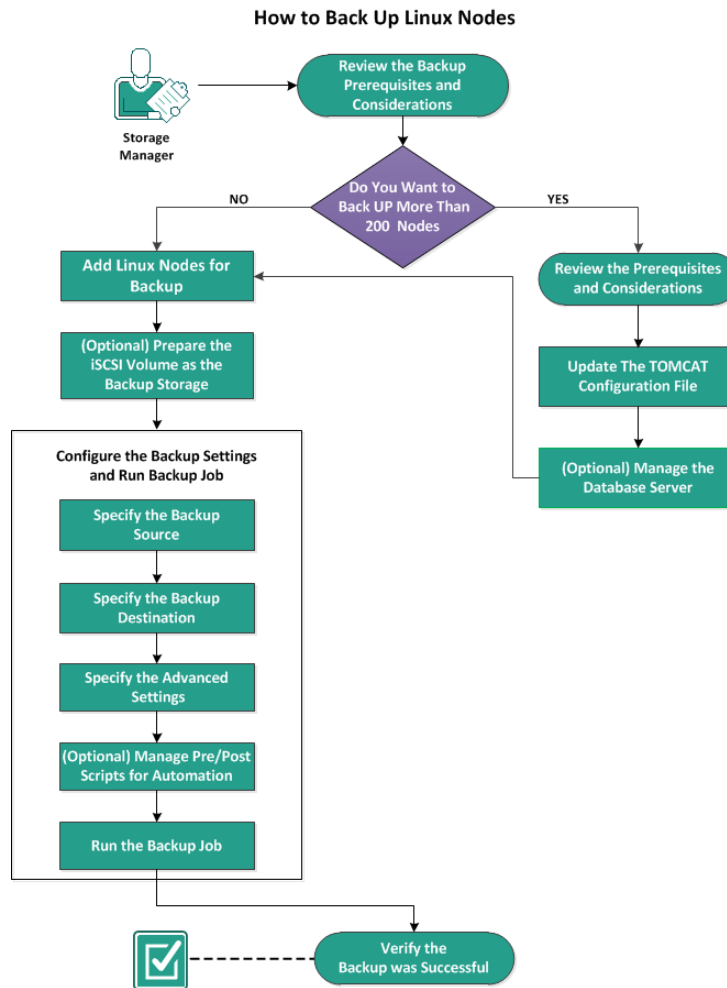
The job is deleted.

How to Back up Linux Nodes

Arcserve UDP Agent (Linux) lets you back up Linux nodes and data that are stored in it. You can also back up the Backup Server itself like any other Linux node. The Backup Server can back up a maximum of 200 nodes.

When Arcserve UDP Agent (Linux) performs a backup of data, it also captures information that is related to the operating system, installed applications, drivers, and so on, from the production node. As a result, when you restore the backed up data, you can perform a BMR or you can restore files specific to your need.

The following diagram displays the process to back up Linux nodes:



Perform these tasks to back up a Linux node:

- [Review the Backup Prerequisites and Considerations](#) (see page 45)
- [Do You Want to Back Up More Than 200 Nodes](#) (see page 47)
 - [Review the Prerequisites and Considerations](#) (see page 48)
 - Update the TOMCAT Configuration File
 - Manage the Database Server
- Add [Linux Nodes for Backup](#) (see page 50)
- [\(Optional\) Prepare the iSCSI Volume as the Backup Storage](#) (see page 52)
- [Configure the Backup Settings and Run Backup Job](#) (see page 54)
 - [Specify the Backup Source](#) (see page 54)
 - [Specify the Backup Destination](#) (see page 57)
 - [Specify the Advanced Settings](#) (see page 60)
 - [\(Optional\) Manage Pre/Posts Scripts for Automation](#) (see page 69)
 - [Run the Backup Job](#) (see page 72)
- [Verify the Backup was Successful](#) (see page 73)

Review the Backup Prerequisites and Considerations

Verify the following requirements before performing a backup:

- You have the supported hardware and software requirements for the backup node.
Note: For more information about the supported hardware and software requirements, see the *Release Notes*.
- You have a valid destination to store your backed up data.
- You have the user names and passwords of nodes that you want to back up.
- The */tmp* folder in the backup node has a minimum of 300 MB space. The */tmp* folder is used to process the accumulation of incremental blocks.
- Perl and sshd (SSH Daemon) are installed on the nodes that you want to back up.
- mount.nfs is installed on the nodes that you want to back up.

- mount.cifs is installed on the nodes that you want to back up.
- The backup node can access your backup destination and you have the write permission.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

To rerun a backup job, verify that you have backed up the node before and you have a valid backup job.

Review the following backup considerations:

- To optimize the management of your Recovery Points, you should consider the following recommendations when scheduling the frequency of your backups:
 - For systems that are protected with Incremental Backups performed every 15 minutes, you should schedule a Full Backup every week (to refresh your base image).
 - For systems that are protected with Incremental Backups performed every hour, you should schedule a Full Backup every month (to refresh your base image).

Note: If the amount of space used to store your backup images is a concern, you should consider scheduling your Full Backups less frequently to consume less storage space.

Disk Supported by Arcserve UDP Agent (Linux)

Different types of disks are supported for Arcserve UDP Agent (Linux) backup source and backup disks. The following matrix lists the types of disks that are supported for each function.

Backup and BMR Support		
Disk (Volume) Type	As Backup Source	As Backup Destination
Mounted Volume (Traditional Disk Partition and LVM *2)	Yes	Yes
RAW Volume (Not formatted)	No	No
Swap	No	Not applicable
GPT Disk:		

Backup and BMR Support		
■ GPT (GUID Partition Table) Data Disk	Yes	Yes
■ GPT (GUI Partition Table) Boot Disk	No	Not applicable
RAID Disk *1:		
■ Software RAID (RAID-0 (Stripe))	Yes	Yes
■ Software RAID (RAID-1 (Mirrored))	Yes	Yes
■ Software RAID-5	Yes	Yes
■ Hardware RAID (include Embedded RAID)	Yes	Yes
File System:		
■ EXT2	Yes	Yes
■ EXT3	Yes	Yes
■ EXT4	Yes	Yes
■ Reiserfs Version 3	Yes	Yes
Shared Volume:		
■ Windows Shared Volume	No	Yes
■ Windows NFS 3.0 Share Volume	No	Yes
■ Linux Shared Volume (samba shared)	No	Yes
■ Linux NFS Share	No	Yes
Device Type:		
■ Removable Disk (Ex. Memory Stick, RDX)	Yes	Yes
*1	Fake RAID, also called Embedded RAID, provided by the BIOS on the motherboard is not supported by Arcserve UDP Agent (Linux).	
	LVM (Logical Volume Manager) is not supported on SUSE Linux Enterprise Server (SLES) 10, but is supported on SLES 10 SP1 to SP4. Embedded LVM is not supported.	

Do You Want To Back Up More Than 200 Nodes

A Backup Server can manage a maximum of 200 nodes by default. If you have more than 200 nodes to back up, you can set up member Backup Servers. Then use a central Backup Server to manage all your member servers.

If you have one dedicated Backup Server and you have more than 200 nodes to manage, you can enable specific settings and manage more than 200 nodes.

Review the Prerequisites and Considerations

Verify the following prerequisites before you back up more than 200 Linux nodes:

- Only 64-bit Linux is supported for the Backup Server
- The Backup Server must be a dedicated server. Arcserve UDP Agent (Linux) modifies the system settings to meet the high scalability requirement of the server.
- The server must meet the following minimum hardware requirements. If you have larger number of nodes, then the hardware specifications must be greater than the minimum requirements.
 - 8-GB memory
 - 10-GB free disk space for the /opt folder

Review the following considerations:

- When you enable Arcserve UDP Agent (Linux) to back up more than 200 nodes, a new database (postgresql) is used by the server to meet the high scalability requirement. All existing node and job information in the old database (sqlite) are migrated to the new database, except the job history and the activity log. You cannot revert to the old database (sqlite) after the migration.
- After the migration, the output is displayed in a different format for the d2djobhistory command.
- As a best practice, one backup job should back up less than 1000 nodes.

Update the TOMCAT Configuration File

When you upgrade to Arcserve UDP Agent (Linux) from the previous version, such as r16.5 SP1, update the TOMCAT configuration file to support the high scalability requirement of the Backup Server. This update lets you back up more than 200 nodes using one Backup Server.

Follow these steps:

1. Log in to the Backup Server as a root user.
2. Navigate to the bin folder:

```
/opt/CA/d2dserver/bin
```

3. Verify that there are no running jobs and then stop the Backup Server using the following command:

```
d2dserver stop
```

If there are jobs running, then wait for the completion of the jobs before you stop the Backup Server.

4. Open the server.xml file from the following location:

```
/opt/CA/d2dserver/TOMCAT/conf/
```

5. Update the following parameters.

If https is used, then update the following parameters:

```
<Connector port="8014" connectionTimeout="180000" protocol="HTTP/1.1" SSLEnabled="true"
maxThreads="300" acceptCount="200" scheme="https" secure="true" clientAuth="false"
sslProtocol="TLSv1, TLSv1.1, TLSv1.2" keystoreFile="{catalina.home}/conf/server.keystore"
keystorePass="LinuxD2D"/>
```

If http is used, then update the following parameters:

```
<Connector connectionTimeout="180000" port="8014" maxThreads="300" acceptCount="200"
protocol="HTTP/1.1"/>
```

The TOMCAT configuration file is successfully updated.

6. Stop the Backup Server.

```
./d2dserver stop
```

7. Run the following command to start the Backup Server:

```
./pgmgr init
```

The command verifies that all the necessary changes are completed and starts the Backup Server.

```
[root@<Machine Name> bin]# ./d2dserver stop
arcserve UDP Agent(Linux) is stopped.
[root@<Machine Name> bin]# ./pgmgr init
The installation process has started for Postgresql database. The debug log
is placed at the following location: /opt/CA/d2dserver/logs/pginit.log.
The Postgresql database has been successfully installed.
Data has been successfully migrated to the new database.
arcserve UDP Agent(Linux) is started.
```

The Backup Server and the database server are successfully started.

Manage the Database Server

The *d2dserver start* command typically starts the database server along with the Backup Server. If there are no jobs in progress, then the *d2dserver stop* command typically stops both the servers.

If you want to start and stop the database server manually, you can run the following commands:

pgmgr start

Starts the database server.

pgmgr stop

Stops the database server.

pgmgr status

Displays the status of the database server. It displays whether the database server is running or is it stopped.

Note: If the database is loaded with excessive data, the Arcserve UDP Agent (Linux) Console takes longer time to load data for job history and activity log. To improve the data querying, see [Improve the Query Performance for Job History and Activity Log](#).

Add Linux Nodes for Backup

Add Linux nodes so that you can back up those nodes to a backup storage location. Linux nodes are your machines that you want to back up. You can either add nodes manually or you can run a script to discover and add nodes.

Follow these steps:

1. Enter the URL of the Backup Server in a web browser to open the user interface.

Note: During the installation of Arcserve UDP Agent (Linux), you received the URL to access and manage the server.

2. Perform the following tasks if you want to discover nodes using a script:

- a. Click Add from the Node menu and select Discovery.

The Node Discovery dialog opens.

- b. Select a script from the Script drop-down list.

Note: For more information about creating the node discovery script, see Discover Nodes Using a Script in How to Integrate and Automate Arcserve UDP Agent (Linux) with the Existing IT Environment.

- c. Specify the Schedule and click OK.

The Node Discovery dialog closes and the node discovery process starts. The Activity Log tab is updated with a new message.

3. Perform the following tasks if you want to add each node manually:

- a. Click Add from the Node menu and select Hostname/IP Address.

The Add Node dialog opens.

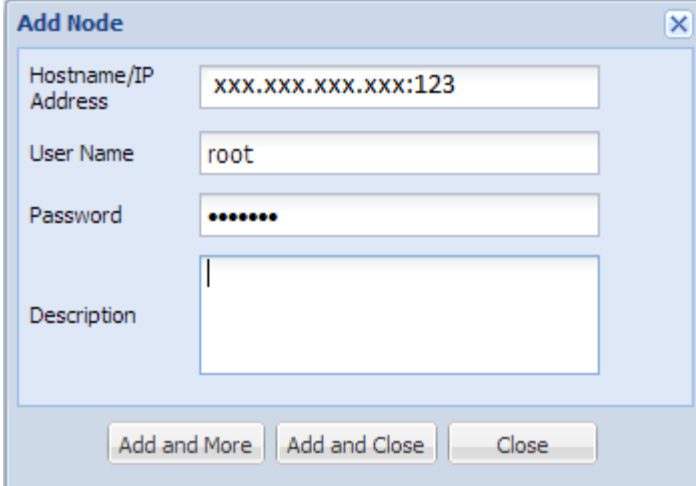
- b. Enter the host name or IP address of the Linux node, the user name that has the root permission, and the password.

Note: If the default ssh port of the node is changed, you can add the node as following:

<IP Name>:Port Number

Example: xxx.xxx.xxx.xxx:123

Where, xxx.xxx.xxx.xxx is the IP address and 123 is the port number.



The screenshot shows a dialog box titled "Add Node" with a close button in the top right corner. The dialog contains the following fields and values:

- Hostname/IP Address: xxx.xxx.xxx.xxx:123
- User Name: root
- Password:
- Description: (empty text area)

At the bottom of the dialog, there are three buttons: "Add and More", "Add and Close", and "Close".

- c. (Optional) Enter a description for the node to assist you in locating the node.
- d. Select one of the following options.

Add and More

Lets you add multiple nodes one at a time. After you finish adding your nodes, click Add and Close or Close to close the Add Node dialog.

Add and Close

Lets you add one node and then the Add Node dialog closes.

Close

Closes the dialog without adding any nodes.

4. Click the Nodes tab and verify that new nodes are listed in it.

Linux nodes are added for backup.

(Optional) Prepare the iSCSI Volume as the Backup Storage

You can store your recovery points to an Internet Small Computer System Interface (iSCSI) volume. iSCSI is used to manage data transfer and storage over a network using the IP standard.

Verify that you have the latest release of the iSCSI-initiator software installed on your Backup Server. The initiator software on RHEL systems is packaged as `iscsi-initiator-utils`. The initiator software on SLES systems is packaged as `open-iscsi`.

Follow these steps:

1. Log in to the shell environment of the backup source node.
2. Run one of the following commands to start the iSCSI initiator daemon.

- For RHEL systems:

```
/etc/init.d/iscsid start
```

The service on RHEL systems is named `iscsid`.

- For SLES systems:

```
/etc/init.d/open-iscsi start
```

The service on SLES systems is named `open-iscsi`.

3. Run a discovery script to discover the iSCSI target host.

```
iscsiadm -m discovery -t sendtargets -p <ISCSI-SERVER-IP-ADDRESS>:<Port_Number>
```

The default port value of iSCSI target host is 3260.

4. Make a note of the iSCSI Qualified Name (IQN) of the iSCSI target host found by the discovery script before you manually log into the discovered target.

5. List the available block device of the backup source node.

```
#fdisk -l
```

6. Log in to the discovered target.

```
iscsiadm -m node -T <iSCSI Target IQN name> -p <iSCSI-SERVER-IP-ADDRESS>:<Port_Number> -l
```

You can see a block device in the /dev directory of the backup source node.

7. Run the following command to obtain the new device name:

```
#fdisk -l
```

You can see an additional device named /dev/sd<x> on the backup source node.

For example, consider the name of the device is /dev/sdc. This device name is used to create a partition and a file system in the following steps.

8. Format and mount the iSCSI Volume.

9. Create a partition and a file system on the backup source node using the following commands.

```
# fdisk /dev/sdc
```

If you have create only one partition, then use the following command to create a file system for the single partition:

```
# mkfs.ext3 /dev/sdc1
```

10. Mount the new partition using the following commands:

```
# mkdir /iscsi
```

```
# mount /dev/sdc1 /iscsi
```

The new partition is mounted and the iSCSI volume is ready to be used as a backup storage in a backup job.

11. (Optional) Add the following record to the /etc/fstab folder so that the iSCSI volume automatically connects with the Backup Server after you restart the server.

```
/dev/sdc1 /iscsi ext3 _netdev 0 0
```

The iSCSI volume is ready to be used as the backup storage.

Configure the Backup Settings and Run Backup Job

Configure the backup settings using the Backup Wizard. You can back up your data to a Network File System (NFS) location, Network-attached storage (NAS), Common Internet File System (CIFS), or to a source local location. A source local location is a location in the backup source node where the backed up data is stored. The backup process is initiated by a backup job. The Backup Wizard creates the backup job and runs the job. Each time that you perform a successful backup, a recovery point is created. A recovery point is a point-in-time copy of the backup node.

Specify the Backup Source

Specify the backup source nodes in the Backup Wizard so that you can back up those nodes to a desired location. The Backup Source page of the Backup Wizard displays the node that you want to backup. Use the Add button on this page to add more nodes for backup.

Note: If you open the Backup Wizard using the Back Up Selected Nodes button, then all the selected nodes are listed in the wizard page. If you open the Backup Wizard using the Back Up button, then nodes are not listed in the wizard page. You have to add nodes using the Add button in the wizard page.


Follow these steps:

1. Select the nodes that you want to back up from the Nodes tab.
2. Click Backup, and select the Back Up Selected Nodes option from the Wizard menu.


The Backup Server page of the Backup Wizard opens. The Backup Server page displays the server name.

3. Click Next.


The Backup Source page opens. The previously selected nodes are displayed on this page.




Backup Server




Backup Source



Backup Destination



Advanced



Summary

Set information for the target nodes you want to back up.

You can enter information for multiple nodes. All those nodes will share one backup job.

You can select the backup sources from the "Nodes" page or add them manually by clicking the "Add" button.

Hostname/IP Address	Username	Status	Exclude Volumes
Node 1	root	Click here to get more node information.	
Node 2	root	Click here to get more node information.	
Node 3	root	Click here to get more node information.	

Volumes to be excluded for all listed nodes:

Files/Folders to be excluded for all listed nodes:

4. (Optional) Click Add in the Backup Source page to add more nodes and provide the details in the Add Node dialog.

5. (Optional) Click the Exclude Volumes icon.

The Exclude Volume Settings dialog opens which includes all the volumes of that node.

6. (Optional) Select the check box for each volume that you do not want to back up and click OK.

The Exclude Volume Settings dialog closes.

Note: To exclude a specific volume from all the backup nodes, enter the volume mount point names in **Volumes to be excluded for all listed nodes**. If you exclude the `/` volume or the `/boot` volume of a node, then you cannot perform a BMR of that node.

7. (Optional) Enter the files/folders in **Files/Folders to be excluded for all listed nodes**.

The files/folders should be specified with an absolute path name and separated with a colon (:). Wildcard characters, such as `*` and `?` are supported and should be used after the last slash of the absolute path name. If the files/folders name after the last slash is enclosed in parentheses, these files/folders will be excluded recursively, otherwise the files/folders will be excluded directly.

For example:

```
/home/user/a/foo*:/home/user/b/(foo*)
```

The first part (`home/user/a/foo*`) will exclude only files/folders that match `foo*` under `"/home/user/a"`, but it will back up sub-directories within. The second part (`/home/user/b/(foo*)`) will exclude all the files/folders that match `foo*` under `"/home/user/b"`, including all sub-folders.

Notes:

- If many files/folders are excluded from a volume, it is recommended to exclude the relevant volume.
- If many files/folders are excluded, the job phase and status may stay "Backing up volume" and "Active" for a long time, when the backup job is launched.
- If the value of **Files/Folders to be excluded for all listed nodes** is changed, the backup job will be converted to a full backup.

If certain system files are excluded from the backup, then Linux OS may not boot, and the BMR function does not work properly. Such system files include, but not limited to:

- Files and folders under `/bin`, `/sbin`, `/usr`, `/etc`, `/lib`, `/lib64`, `/boot`, `/var`
- Folder `/proc`, `/sys`, `/dev`, `/tmp`

If you exclude the system files, then it is recommended to verify the BMR function and confirm whether the Linux OS boots properly.

- Click Next.

The Backup Destination page opens.

The backup source is specified.

Specify the Backup Destination

Specify a location to store the backed up data (recovery points) in the Backup Destination page of the Backup Wizard. The backup destination could be an NFS share, a CIFS share, or Source local. Source local is the backup source node. If your backup destination is Source local, then the backup data is written to its own local disk directly.

Specify the storage location for your backup data.

Backup Destination

NFS share NFS Share Full path

Specify the storage options for your backup data.

Compression

Using compression will reduce the amount of space required on your destination.

Standard Compression

Encryption Algorithm

Encryption Algorithm No Encryption

Encryption Password

Re-type Password

If a physical disk includes two logical volumes, then you can specify one volume as the backup source and the other volume as the backup destination.

Note: If you select Source local as your backup destination, the Backup Server cannot manage the recovery points. To manage the recovery sets, see Manage the Recovery Sets in How to Manage the Backup Server Settings.

Follow these steps:

1. Select a destination from the Backup Destination drop-down list and enter the complete path of the storage location.
 - If you have selected NFS share, then type the Backup Destination detail in the following format:
IP address of the NFS Share:full path of the storage location

Note: Some versions of Data Domain NAS do not support the file locking mechanism of NFS. As a result, such NFS share cannot be used as a backup destination. For more information about this issue, see Compatibility Issues with Arcserve UDP Agent (Linux) in the [Release Notes](#).
 - If you have selected CIFS share, then type the Backup Destination detail in the following format:
//hostname/share_folder

Note: The shared folder name cannot contain any spaces.
 - If you have selected Source local, then you have to modify some settings so that the Backup Server can manage the recovery points. For example, consider server-A as the host name of the Backup Server and node-B as the host name of the source node. Now, follow these steps to modify the settings of node-B:
 - Ensure that the NFS server is running. You can run the following command to verify the NFS server status:

`service nfs status`
 - If the NFS server is not running, run the following command to start the NFS server:

`service nfs start`
 - If your backup destination folder on node-B is `/backup/test`, then add the following line to `/etc/exports`:

`/backup/test server-A(rw,no_root_squash)`

Now, run the following command:

`exportfs -a`
 - On the Backup Server UI, add `node-B:/backup/test` as a backup storage location. The Source local storage location is displayed in the Backup Destination drop-down list.

2. Click the arrow button to validate the Backup Destination information.

If the backup destination is invalid, an error message is displayed.

3. Select a compression level from the Compression drop-down list to specify a type of compression that is used for backup.

The available options for Compression are:

Standard Compression

Specifies that this option provides a good balance between the CPU usage and the disk space usage. This compression is the default setting.

Maximum Compression

Specifies that this option provides the highest CPU usage (lowest speed), but also has the lowest disk space usage for your backup image.

4. Select an algorithm from the Encryption Algorithm drop-down list and type the encryption password, if necessary.

- a. Select the type of encryption algorithm that you want to use for backups.

Data encryption is the translation of data into a form that is unintelligible without a deciphering mechanism. Arcserve UDP Agent (Linux) data protection uses secure, AES (Advanced Encryption Standard) encryption algorithms to achieve the maximum security and privacy of your specified data.

The available format options are No Encryption, AES-128, AES-192, and AES-256. (To disable encryption, select No Encryption).

- A full backup and all its related incremental backups must use the same encryption algorithm.
- If the encryption algorithm for an incremental backup has changed, you must perform a full backup.

For example, if you change the algorithm format and then you run an incremental backup, then the backup type automatically converts to a full backup.

- b. When an encryption algorithm is selected, you must provide (and confirm) an encryption password.

- The encryption password is limited to a maximum of 23 characters.
- A full backup and all its related incremental backups use the same password to encrypt data.

5. Click Next.

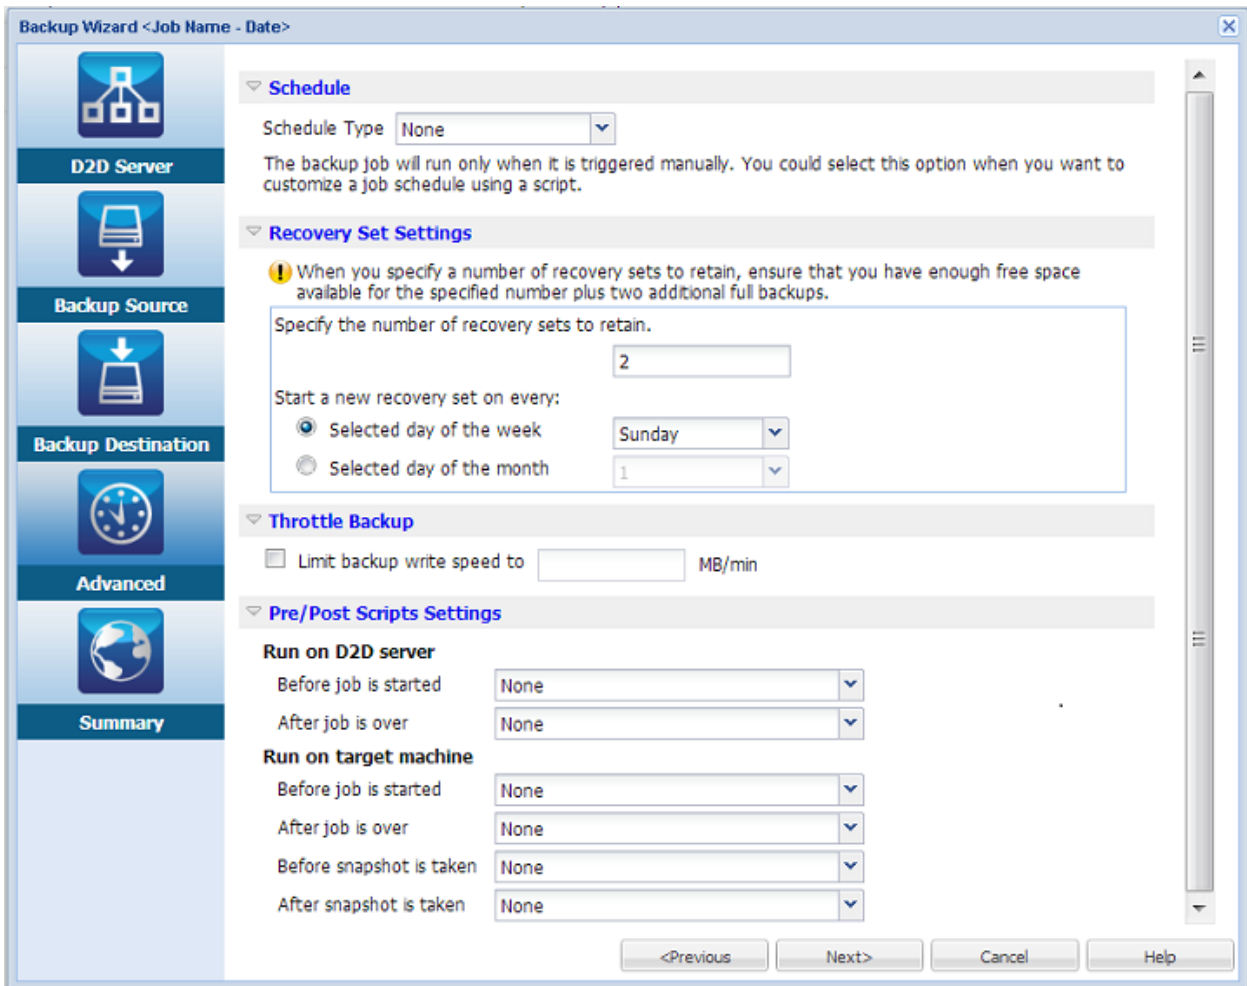
The Advanced page opens.

The backup destination is specified.

Specify the Advanced Settings

Specify the backup schedule, recovery set settings, and pre-backup and post-backup settings on the Advanced page.

The following diagram displays the Advanced page of the Backup Wizard. In this diagram, None option is selected for the Schedule Type.



The following settings are available on the Advanced page:

- The Schedule settings ensure that the backup job runs periodically at a specified time.
- The Recovery Set Settings ensures periodic maintenance of the recovery sets. If the number of recovery sets exceeds the specified number, then the oldest recovery set is deleted to maintain the specified number all the time.
- The Throttle Backup setting lets you enable and specify the maximum speed (MB/min) at which the backups are written.
- The Pre/Post Scripts Settings defines the scripts that can be run on the Backup Server and the target node. You can configure the scripts to take specific actions before the start of a job, during the job is running, or after the completion of the job.

To optimize the management of your Recovery Points, you should consider the following recommendations when scheduling the frequency of your backups:

- For systems that are protected with Incremental Backups performed every 15 minutes, you should schedule a Full Backup every week (to refresh your base image).
- For systems that are protected with Incremental Backups performed every hour, you should schedule a Full Backup every month (to refresh your base image).

Note: If the amount of space used to store your backup images is a concern, you should consider scheduling your Full Backups less frequently to consume less storage space.

Follow these steps:

1. Set the start date and time by selecting one of the following options from the Schedule Type dropdown list:

Simple

Select the Simple option to schedule the Incremental Backup, Full Backup, and Verify Backup per the specified Start Date and Start Time. For each type of backup, you can also specify the repeat duration for a backup or never repeat a backup. The start date and time is fixed for all types of backup. So, you cannot specify a different start date and time for different types of backup.

Note: For more information about the backup types, see *Understanding the Types of Backup*.

Schedule Type

Set start date and time
Specify the scheduled starting date and time for the full, incremental, and verify backups.
Start Date Start Time :

Incremental Backup
Incrementally backs up only the data that has changed since the last successful backup.
 Repeat Every

Full Backup
Backs up all selected data from the machine.
 Repeat Every Days
 Never

Verify Backup
Performs a confidence check to compare data from the last successful backup with data from the source, and then incrementally backs up (resynchronizes) only the differences.
 Repeat Every Days
 Never

Advanced

Select the Advanced option to specify multiple backup schedules each day of the week. You can specify different start date and time for different types of backup. You can add, modify, delete, and clear the Advanced schedule. When you click Clear, all the Advanced backup schedules are deleted from the Advanced Schedule Tray.

▼ Schedule

Schedule Type

Start Date


Time	Backup Type	Repeat
☰ Sunday		
☰ Monday		
📅 10:00 PM	Incremental Backup	Never
☰ Tuesday		
📅 10:00 PM	Incremental Backup	Never
☰ Wednesday		
📅 10:00 PM	Incremental Backup	Never
☰ Thursday		
📅 10:00 PM	Incremental Backup	Never
☰ Friday		

Advanced Schedule Tray

To add a backup schedule, follow these steps:

- a. Click Add.

The Add backup schedule dialog opens.



- b. Specify your backup schedule options and click OK.

The specified backup schedule displays on the Advanced Schedule Tray.

None

Select the None option to create the backup job and stores the job in the Job Status tab. This option will not run the job because there is no specified schedule. When you submit the job, the status of the job changes to Ready. When you want to run the job, you have to select the job and click Run Now from the Job menu. Each time you want to run the job, you must run the job manually. You can also write a script to run this job at your own customized schedule.

- 2. Specify your recovery set settings.

Note: For more information about the recovery sets, see *Understanding the Recovery Sets*.

Specify the number of recovery sets to retain

Specifies the number of recovery sets retained.

Start a new recovery set on every:**Selected day of the week**

Specifies the day of the week selected to start a new recovery set.

Selected day of the month

Specifies the day of the month selected to start a new recovery set. Specify 1 through 30, or the last day of the month.

Note: The Backup Server checks for the number of recovery sets in the configured backup storage every 15 minutes and deletes any extra recovery set from the backup storage location.

3. Specify the throttle backup value.

You can specify the maximum speed (MB/min) at which backups are written. You can throttle the backup speed to reduce CPU or network use. However, limiting the backup speed has an adverse effect on the backup window. As you lower the maximum backup speed, it increases the amount of time of perform the backup. For a backup job, the Job Status tab displays the average Read and Write speed of the job in progress and the configured throttle speed limit.

Note: By default, the Throttle Backup option is not enabled and backup speed is not being controlled.

4. Specify your pre-backup settings and post-backup settings in Pre/Post Scripts Settings.

These scripts run script commands for actions to take before the start of the job and/or upon the completion of the job.

Note: The Pre/Post Script Settings fields are populated only if you already created a script file and placed it at the following location:

```
/opt/CA/d2dserver/usr/prepost
```

Note: For more information about creating the pre/post scripts, see *Manage Pre/Post Scripts for Automation*.

5. Click Next.

The Summary page opens.

The advanced schedule is specified.

Note: If at a given time there is more than one type of backup scheduled to be performed simultaneously, the type of backup that will be performed is based on the following priorities:

- Priority 1 - Full backup
- Priority 2 - Verify backup
- Priority 3 - Incremental backup

For example, if you schedule all three types of backups to be performed simultaneously, Arcserve UDP Agent (Linux) will perform the Full Backup. If there is no Full Backup scheduled, but you scheduled a Verify Backup and Incremental Backup to be performed simultaneously, Arcserve UDP Agent (Linux) will perform the Verify Backup. A scheduled Incremental Backup is performed only if there is no conflict with any other type of backup.

Understanding the Types of Backup

You can specify the following types of backup in the Advanced page of the Backup Wizard:

Incremental Backup

Backs up only those blocks that have changed since the last successful backup. The advantages of Incremental Backup are that it is a fast backup and it produces a small backup image. Arcserve Unified Data Protection Agent for Linux uses a driver to monitor the changed blocks in the source node since the last successful backup.

The available options are Repeat and Never. If you select the Repeat option, you must also specify the elapsed time period (in minutes, hours, or days) between backup attempts.

Minimum: 15 minutes

Default: 1 day

Full Backup

Backs up the entire source node. Depending on the volume size of the backup node, Full Backup produces a large backup image and usually takes a longer time to complete. The available options are Repeat and Never.

If you select the Repeat option, you must also specify the elapsed time period (in minutes, hours, or days) between backup attempts.

Minimum: 1 day

Default: Never (no scheduled repeat)

Verify Backup

Verifies that the protected data is valid and complete by performing a confidence check of the stored backup image to the original backup source. If necessary, the image is resynchronized. A Verify Backup looks at the most recent backup of each individual block and compares the content and information to the source. This comparison verifies that the latest backed up blocks represent the corresponding information at the source. If the backup image for any block does not match the source (possibly because of changes in the system since the last backup), Arcserve Unified Data Protection Agent for Linux refreshes (resynchronizes) the backup of the block that does not match. You can also use a Verify Backup (very infrequently) to get the guarantee of full backup without using the space required for a full backup.

Advantages: Produces a small backup image when compared to full backup because only the changed blocks (blocks that do not match the last backup) are backed up.

Disadvantages: Backup time is long because all source blocks are compared with the blocks of the last backup.

The available options are Repeat and Never. If you select the Repeat option, you must also specify the elapsed time period (in minutes, hours, or days) between backup attempts.

Minimum: 1 day

Default: Never (no scheduled repeat)

The type of backup that runs depends on the following situations:

- If you run the backup job for the first time for the selected nodes, the first backup is always a Full Backup.
- If you run the backup job again for the same set of nodes and the backup destination is also same, the backup type is Incremental Backup.
- If you run the backup job for the same set of nodes but the backup destination is different, the backup type is Full Backup. This is because you have changed the backup destination and for the new destination this is the first backup. So, the first backup is always a Full Backup.
- If you delete your node and then add the same node again but you do not change the backup destination, the backup will be a Verify Backup. This is because you have already backed up that node in your previous backup jobs. When you delete the node and then you add the node again, the backup job verifies all the blocks of that node with the last backup image. When the backup job decides that it is the same node then it backs up only the changed blocks. If the backup job does not find any backup image of that node in the backup destination, the backup type is a Full Backup.

Understanding the Recovery Sets

A recovery set is a storage setting where a group of recovery points backed-up over a specified period is stored as one set. A recovery set includes a series of backups, starting with a full backup, and then followed by a number of incremental, verify, or full backups. You can specify the number of recovery sets to retain.

The Recovery Set Settings ensures periodic maintenance of recovery sets. When the specified limit is exceeded, the oldest recovery set is deleted. The following values define the default, minimum, and maximum recovery sets in Arcserve UDP Agent (Linux):

Default: 2

Minimum: 1

Maximum number of recovery sets: 100

Maximum number of recovery points (Including one Full backup): 1344

Note: If you want to delete a recovery set to save backup storage space, reduce the number of retained sets and Backup Server automatically deletes the oldest recovery set. Do not attempt to delete the recovery set manually.

Example Set 1:

- Full
- Incremental
- Incremental
- Verify
- Incremental

Example Set 2:

- Full
- Incremental
- Full
- Incremental

A full backup is required to start a new recovery set. The backup that starts the set will be automatically converted to a full backup, even if there is no full backup configured or scheduled to be performed at that time. After the recovery set setting is changed (for example, changing the recovery set starting point from the first backup of Monday to the first backup of Thursday), the starting point of existing recovery sets will not be changed.

Note: An incomplete recovery set is not counted when calculating an existing recovery set. A recovery set is considered complete only when the starting backup of the next recovery set is created.

Example 1 - Retain 1 Recovery Set:

- Specify the number of recovery sets to retain as 1.

Backup Server always keeps two sets to keep one complete set before starting the next recovery set.

Example 2 - Retain 2 Recovery Sets:

- Specify the number of recovery sets to retain as 2.

Backup Server deletes the first recovery set when the fourth recovery set is about to start. This ensures that when the first backup is deleted and the fourth is starting, you still have two recovery sets (recovery set 2 and recovery set 3) available on disk.

Note: Even if you choose to retain only one recovery set, you will need space for at least two full backups.

Example 3 - Retain 3 Recovery Sets:

- The backup start time is 6:00 AM, August 20, 2012.
- An incremental backup runs every 12 hours.
- A new recovery set starts at the last backup on Friday.
- You want to retain 3 recovery sets.

With the above configuration, an incremental backup will run at 6:00 AM and 6:00 PM every day. The first recovery set is created when the first backup (must be a full backup) is taken. Then the first full backup is marked as the starting backup of the recovery set. When the backup scheduled at 6:00 PM on Friday is run, it will be converted to a full backup and marked as the starting backup of the recovery set.

(Optional) Manage Pre/Post Scripts for Automation

Pre/Post scripts let you run your own business logic at specific stages of a running job. You can specify when to run your scripts in **Pre/Post Script Settings** of the **Backup Wizard** and the **Restore Wizard** in the UI. The scripts can be run on the Backup Server depending on your setting.

Managing the pre/post script is a two part process, consisting of creating the pre/post script and placing the script in the prepost folder.

Create Pre/Post Scripts

Follow these steps:

1. Log into the Backup Server as a root user.
2. Create a script file using the environment variables in your preferred scripting language.

Pre/Post Script Environment Variables

To create your script, use the following environment variables:

D2D_JOBNAME

Identifies the name of the job.

D2D_JOBID

Identifies the job ID. Job ID is a number provided to the job when you run the job. If you run the same job again, you get a new job number.

D2D_TARGETNODE

Identifies the node that is being backed up or restored.

D2D_JOBTYPE

Identifies the type of the running job. The following values identify the D2D_JOBTYPE variable:

backup.full

Identifies the job as a full backup.

backup.incremental

Identifies the job as an incremental backup.

backup.verify

Identifies the job as a verify backup.

restore.bmr

Identifies the job as a bare-metal recovery (BMR). This is a restore job.

restore.file

Identifies the job as a file-level restore. This is a restore job.

D2D_SESSIONLOCATION

Identifies the location where the recovery points are stored.

D2D_PREPOST_OUTPUT

Identifies a temp file. The content of the first line of the temp file is displayed in the activity log.

D2D_JOBSTAGE

Identifies the stage of the job. The following values identify the D2D_JOBSTAGE variable:

pre-job-server

Identifies the script that runs on the Backup Server before the job starts.

post-job-server

Identifies the script that runs on the Backup Server after the job completes.

pre-job-target

Identifies the script that runs on the target machine before the job starts.

post-job-target

Identifies the script that runs on the target machine after the job completes.

pre-snapshot

Identifies the script that runs on the target machine before capturing the snapshot.

post-snapshot

Identifies the script that runs on the target machine after capturing the snapshot.

D2D_TARGETVOLUME

Identifies the volume that is backed up during a backup job. This variable is applicable for pre/post snapshot scripts for a backup job.

D2D_JOBRESULT

Identifies the result for a post job script. The following values identify the D2D_JOBRESULT variable:

success

Identifies the result as successful.

fail

Identifies the result as unsuccessful.

D2DSVR_HOME

Identifies the folder where Backup Server is installed. This variable is applicable for the scripts that run on the Backup Server.

The script is created.

Note: For all scripts, a return value of zero indicates success and a nonzero return value indicates failure.

Place the Script in the Prepost Folder and Verify

All the pre/post scripts for a Backup Server are centrally managed from the prepost folder at the following location:

```
/opt/CA/d2dserver/usr/prepost
```

Follow these steps:

1. Place the file in the following location of the Backup Server:

```
/opt/CA/d2dserver/usr/prepost
```

2. Provide the execution permission to the script file.
3. Log into the Arcserve UDP Agent (Linux) web interface.
4. Open the **Backup Wizard** or the **Restore Wizard** and navigate to the **Advanced** tab.
5. Select the script file in the **Pre/Post Script Settings** drop-down list and then submit the job.
6. Click **Activity Log** and verify that the script is executed to the specified backup job.
The script is executed.

The pre/post scripts are successfully created and placed in the prepost folder.

Run the Backup Job

Run the backup job so that a recovery point is created. You can use this recovery point to restore data.

On the Summary page, review the summary of the backup details and provide a job name to distinguish it from other jobs.

Follow these steps:

1. Review the summary and enter a job name.
The Job Name field has a default name initially. You can enter a new job name of your choice but you cannot leave this field empty.
2. (Optional) Click Previous to modify any settings on any wizard pages.
3. Click Submit.

The backup process starts. In the Job Status tab, the job is added and the backup status is displayed.

The backup job is created and run.

Verify the Backup was Successful

After the backup job is complete, verify that the recovery point is created at the specified destination.

Follow these steps:

1. Navigate to the specified destination where you have stored your backup data.
2. Verify that the backup data is present in that destination.

For example, if the backup job name is *Demo* and the backup destination is `xxx.xxx.xxx.xxx:/Data`, then navigate to the backup destination and verify that a new recovery point is generated.

The backup data is successfully verified.

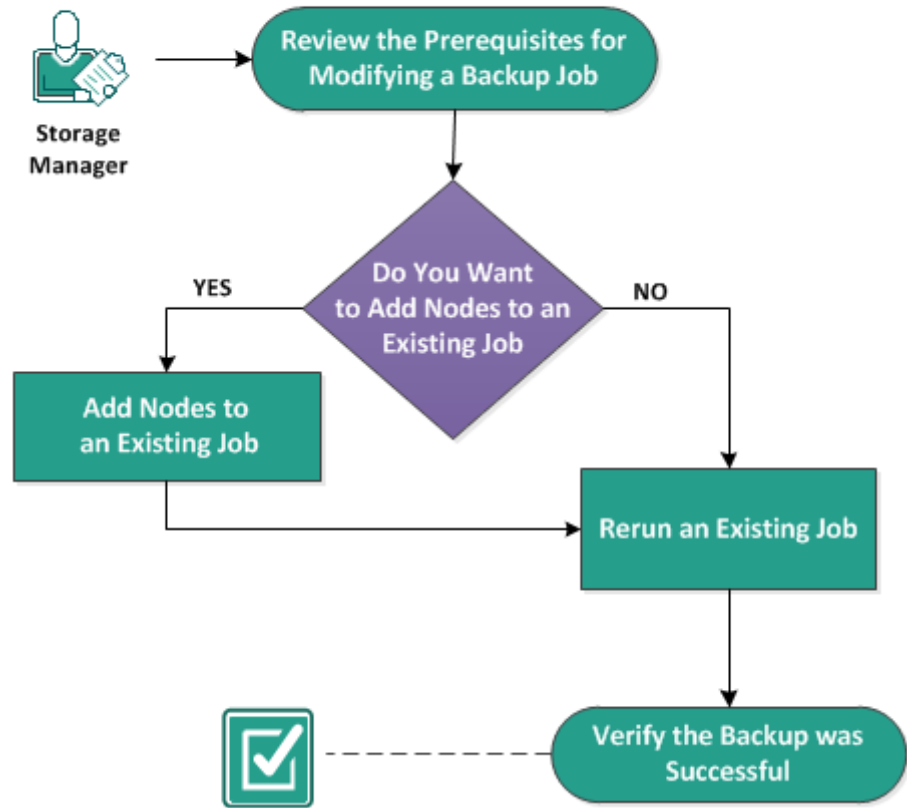
The Linux nodes are successfully backed up.

How to Modify and Rerun a Backup Job

If you have already created a job for a node, you can modify it and rerun the job multiple times. You do not have to create another job to protect the same node. If you do not want to make any changes to the job, you can also run the job without modifying it. Modifying a job includes adding a node to an existing job, configuring the job settings, or both.

The following diagram displays the process to modify and rerun a backup job:

How to Modify and Rerun a Backup Job



Perform these tasks to modify and rerun a backup job:

- [Review the Prerequisites for Modifying a Backup Job](#) (see page 75)
- [Do You Want to Add Nodes to an Existing Job](#) (see page 75)
- [Add Nodes to an Existing Job](#) (see page 75)
- [Rerun an Existing Job](#) (see page 76)
- [Verify the Backup was Successful](#) (see page 77)

Review the Prerequisites for Modifying a Backup Job

Verify the following requirements before you modify and rerun a backup job:

- You have a valid backup job.
- You have added the nodes to Arcserve UDP.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Do You Want to Add Nodes to an Existing Job

If you already have a backup job and you want to protect new nodes with the same backup settings, you can add nodes to an existing job. After you add the nodes, you can also modify the backup settings and run the job.

Add Nodes to an Existing Job

You can add new nodes to an existing backup job and can run the job. All the settings of the selected job are applied to the new node and you do not have to configure any new backup settings. Use this option if you want to keep the same backup settings for all the nodes.

Follow these steps:

1. Select all the new nodes from the Nodes tab in the Status pane.
2. From the Wizard menu, click Backup and select Add Selected Nodes to an Existing Job.

The Add Selected Nodes into an Existing Job dialog opens.

3. Select a job from the Job Name drop-down list and click OK.

The node is added to the selected backup job and the Protected column in the Nodes tab changes to Yes.

Nodes are added to an existing job.

Rerun an Existing Backup Job

Rerun the backup job to take another backup of the specified nodes. A recovery point is created after each successful backup. If you have already backed up a node, you do not have to create another backup job to back up that node again. All of the previous jobs are listed in the Job Status tab in the Status pane.

When you rerun a backup job, specify the type of job that you want to rerun.

Note: If you update any information in the Backup Destination page of the Backup Wizard before rerunning a job, the job type automatically changes to *Full Backup*.

Follow these steps:

1. Enter the URL of Arcserve UDP Agent (Linux) in a web browser to open the user interface.

Note: During the installation of Arcserve UDP Agent (Linux), you received the URL to access and manage the server.

2. Click the Job Status tab and select the job that you want to run.
3. Verify that the status of the selected job is Done or Ready.

Done implies that the job is not scheduled and Ready implies that the job is scheduled.

4. Perform one of the following steps:

- To run the job without any changes,

- a. Click Run Now from the Job menu.

The Run backup job now dialog opens.

- b. Select the Backup Type.

- c. Select an option from the Run job for dropdown list:

Selected Node

Specifies that the backup job runs for only the selected node.

All nodes protected by the selected job

Specifies that the backup job runs for all the nodes protected by the selected job.

- d. Click OK.

The Run backup job now dialog closes. The status of the job changes to Active in the Job Status tab and the same job runs again.

- To modify the job before you run the job, follow these steps:

- a. Select a job and click Modify.

The Backup Wizard opens.

- b. Update the required field in the Backup Wizard.
- c. Click Submit.

The job runs again depending on the job schedule.

The backup job successfully reruns.

Verify the Backup was Successful

After the backup job is complete, verify that the recovery point is created at the specified destination.

Follow these steps:

1. Navigate to the specified destination where you have stored your backup data.
2. Verify that the backup data is present in that destination.

For example, if the backup job name is *Demo* and the backup destination is `xxx.xxx.xxx.xxx:/Data`, then navigate to the backup destination and verify that a new recovery point is generated.

The backup data is successfully verified.

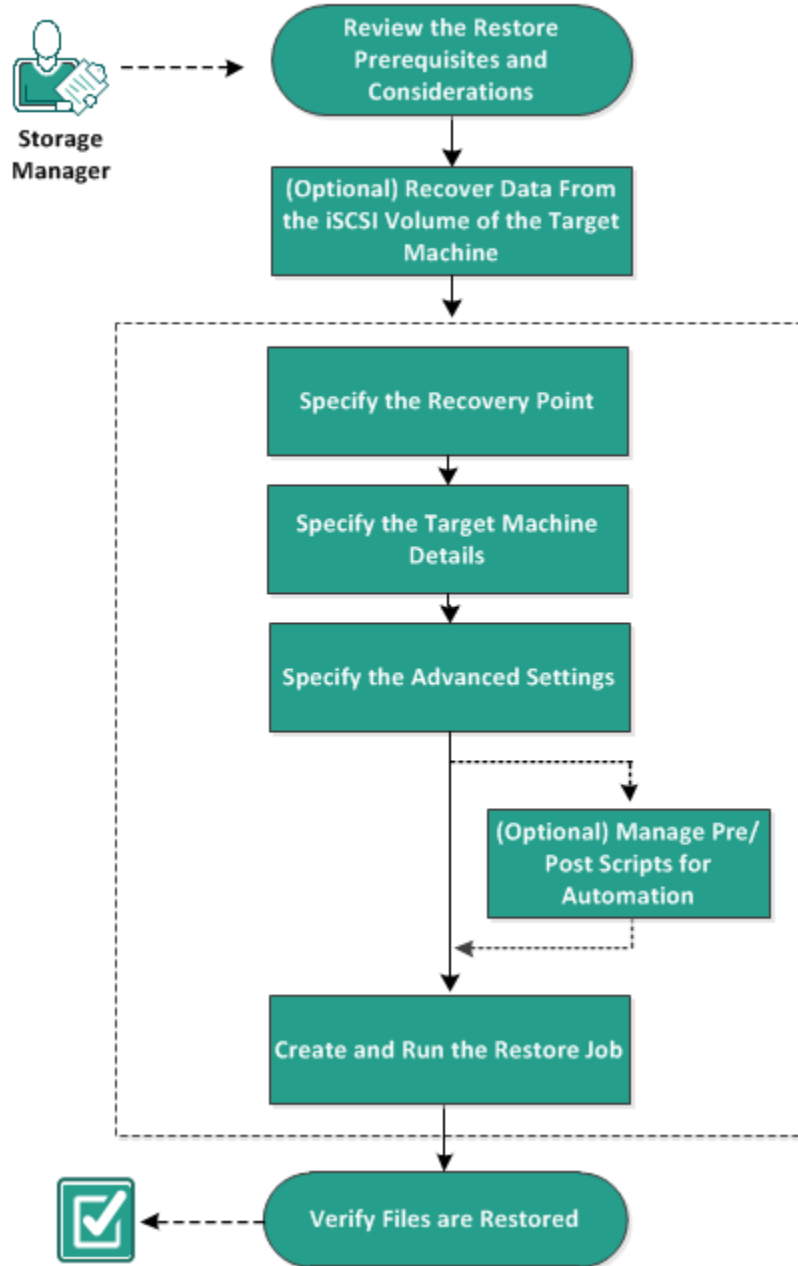
The backup job is successfully modified and rerun.

How to Perform a File-Level Recovery on Linux Nodes

A file-level recovery restores individual files and folders from a recovery point. You can restore as minimum as one file from the recovery point. This option is useful if you want to restore selected files and not the entire recovery point.

The following diagram displays the process to perform a file-level recovery:

How to Perform a File-Level Recovery



Perform these tasks for a file-level recovery:

- [Review the Restore Prerequisites](#) (see page 79)
- [\(Optional\) Recover Data from the iSCSI Volume to the Target Machine](#) (see page 80)
- [Specify the Recovery Point](#) (see page 81)
- [Specify the Target Machine Details](#) (see page 85)
- [Specify the Advanced Settings](#) (see page 89)
- [\(Optional\) Manage Pre/Post Scripts for Automation](#) (see page 90)
- [Create and Run the Restore Job](#) (see page 93)
- [Verify that Files are Restored](#) (see page 94)

Review the Prerequisites

Consider the following options before you perform a file-level recovery:

- You have a valid recovery point and the encryption password, if any.
- You have a valid target node to recover data.
- You have verified that the Linux Backup Server supports the file system that you want to restore.

For example, RedHat 5.x does not support the *reiserfs* file system. If the operating system of the Backup Server is RedHat 5.x and you want to restore the reiserfs file system, you must install the file system driver to support reiserfs. You can also use Arcserve UDP Agent (Linux) Live CD to perform the file-level restore because Live CD supports all types of file system.

- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

(Optional) Recover Data from the iSCSI Volume to the Target Machine

If you have stored your data in an iSCSI target volume, you can connect to the iSCSI volume and recover data. The iSCSI volume lets you manage data and transfer data over a network.

Verify that you have the latest release of the iSCSI-initiator software installed on your Backup Server. The initiator software on RHEL systems is packaged as `iscsi-initiator-utils`. The initiator software on SLES systems is packaged as `open-iscsi`.

Follow these steps:

1. Log into the shell environment of the Backup Server.
2. Run one of the following commands to start the iSCSI initiator daemon.
 - For RHEL systems:

```
/etc/init.d/iscsid start
```

The service on RHEL systems is named `iscsid`.
 - For SLES systems:

```
/etc/init.d/open-iscsi start
```

The service on SLES systems is named `open-iscsi`.
3. Run a discovery script to discover the iSCSI target host.

```
iscsiadm -m discovery -t sendtargets -p <ISCSI-SERVER-IP-ADDRESS>:<Port_Number>
```

The default port value of iSCSI target host is 3260.
4. Make a note of the iSCSI Qualified Name (IQN) of the iSCSI target host found by the discovery script before you manually log into the discovered target.
5. List the available block device of the Backup Server.

```
#fdisk -l
```
6. Log in to the discovered target.

```
iscsiadm -m node -T <iSCSI Target IQN name> -p <ISCSI-SERVER-IP-ADDRESS>:<Port_Number> -l
```

You can see a block device in the `/dev` directory of the Backup Server.

7. Run the following command to obtain the new device name:

```
#fdisk -l
```

You can see an additional device named `/dev/sd<x>` on the Backup Server.

For example, consider the name of the device is `/dev/sdc`. This device name is used to create a partition and a file system in the following steps.

8. Mount the iSCSI volume using the following commands:

```
# mkdir /iscsi
```

```
# mount /dev/sdc1 /iscsi
```

Note: When you specify the session location in the Restore Wizard, you need to select Local and enter the path `/iscsi`.

Example: `<path>/iscsi`

9. (Optional) Add the following record to the `/etc/fstab` file so that the iSCSI volume automatically connects with the Backup Server after you restart the server.

```
/dev/sdc1 /iscsi ext3 _netdev 0 0
```

The Backup Server can now connect to the iSCSI volume and can recover data from the iSCSI volume.

Specify the Recovery Point

Each time that you perform a backup, a recovery point is created. Specify the recovery point information in the **Restore Wizard** so that you can recover the exact data that you want. You can restore specific files or all files depending on your requirement.

Note: If you have selected **Source local** as your backup destination, the Backup Server cannot connect to the Source local directly. To access the Source local, you have to perform additional configurations.

To restore files from Source local, follow these steps:

- a. Share the backup destination (Source local) and ensure that the Backup server can connect to the backup destination.
- b. Add the shared destination as the backup storage location to the Backup server.

Now, Source local behaves as an NFS backup storage location and you can restore files from the share.

Follow these steps:

1. Access the Restore Wizard in one of the following ways:

- From Arcserve UDP:

- a. Log in to Arcserve UDP.
- b. Click the **resources** tab.
- c. Select **All Nodes** in the left pane.

All the added nodes are displayed in the center pane.

- d. In the center pane, select the node and click **Actions**.
- e. Click **Restore** from the **Actions** dropdown menu.

The Arcserve UDP Agent (Linux) web interface opens. The restore type selection dialog is displayed in the Agent UI.

- f. Select the restore type and click **OK**.

Note: You are automatically logged in to the agent node and the **Restore Wizard** opens from the agent node.

- From Arcserve UDP Agent (Linux):
 - a. Open the Arcserve UDP Agent (Linux) web interface.

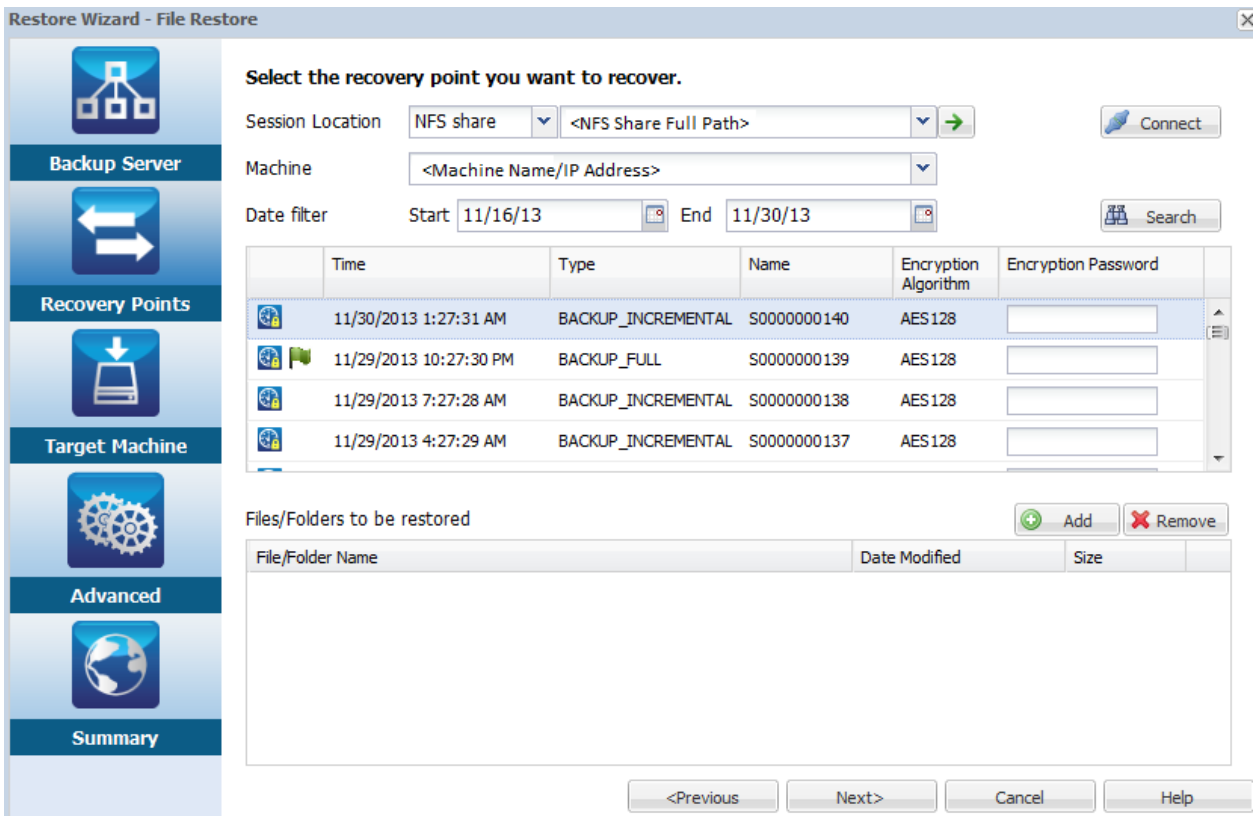
Note: During the installation of Arcserve UDP Agent (Linux), you received the URL to access and manage the server. Log in to Arcserve UDP Agent (Linux).
 - b. Click **Restore** from the **Wizard** menu and select **Restore File**.

Restore Wizard - File Restore opens.

You can see the Backup Server in the **Backup Server** page of the **Restore Wizard**. You cannot select any option from the **Backup Server** drop-down list.

2. Click **Next**.

The **Recovery Points** page of the **Restore Wizard** opens. The recent recovery point is selected.



3. Select a session from the **Session Location** drop-down list, if you want to restore another session, and enter the full path of the share.

For example, consider the Session Location as NFS share, xxx.xxx.xxx.xxx as the IP address of the NFS share, and the folder name is *Data*. You would enter xxx.xxx.xxx.xxx:/Data as the NFS share location.

Note: If the backend up data is stored in Source local, then you must first convert the source node to an NFS server, and then share the session location.

4. Click **Connect**.

All the nodes that have been backed up to this location get listed in the **Machine** drop-down list.

5. Select the node that you want to restore from the **Machine** drop-down list.

All the recovery points of the selected node get listed.

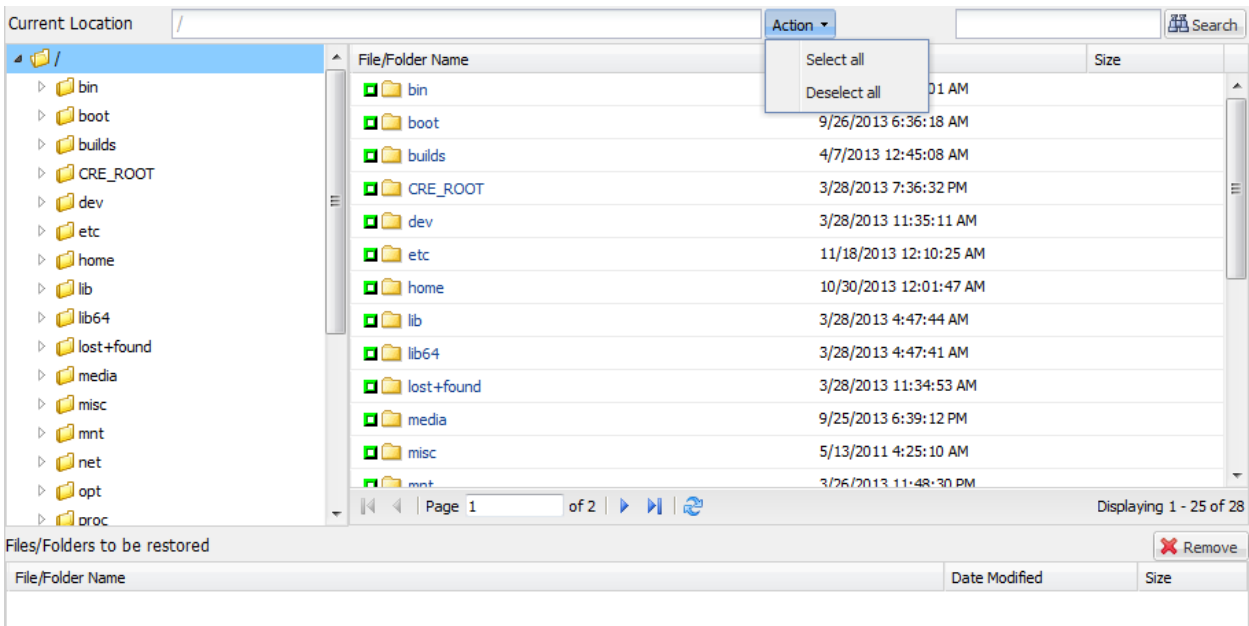
6. Apply the date filter to display the recovery points that are generated between the specified date and click **Search**.

Default: Recent two weeks.

All the recovery points available between the specified dates are displayed.

7. Select the recovery point that you want to restore and click **Add**. If the recovery point is encrypted, enter the encryption password to restore data.

The **Browse-<node name>** dialog opens.



8. Select the files and folders that you want to restore and click **OK**.

Note: If you try to locate a file or folder using the **Search** field, ensure that you select the highest folder in the hierarchy. The search is conducted on all the child folders of the selected folder.

The **Browse-<node name>** dialog closes and you return to the **Recovery Points** page. The selected files and folders are listed under **Files/Folders to be restored**.

9. Click **Next**.

The **Target Machine** page opens.

The recovery point is specified.

Specify the Target Machine Details

Specify the target node details so that data is restored to that node. You can restore the selected files or folders to the source node or to a new node.

Follow these steps:

- To restore to the node from where the data was backed up, follow these steps:
 1. Select **Restore to original location** on the **Target Machine** page.

The **Host Name** field in **Target Machine Settings** gets populated with the name of the source node.

Specify the target machine information for the File Restore.

Restore to original location Restore to alternative location

Target Machine Settings

Host Name/IP	Machine Name/IP Address
Username	
Password	

Resolving Conflicts
How should arcserve UDP Agent(Linux) resolve conflicting files

Overwrite existing files
 Rename files
 Skip existing files

Directory Structure
Whether to create root directory during restore

Create root directory

2. Enter the user name and the password of the node.
3. Select one of the following options to resolve conflicting files:

Overwrite existing files

Specifies that if the file exists in the target machine then the backup file from the recovery point replaces the existing file.

Rename files

Specifies that if the file exists in the target machine, then a new file is created with the same file name and *.d2duplicate<x>* file extension. *<x>* specifies the number of times the file is restored. All the data is restored to the new file.

Skip existing files

Specifies that if the same file exists in the target machine, then those files are not restored from the recovery point.

4. Click **Next**.

The **Advanced** page opens.

- To restore to a new node, follow these steps:
 1. Select **Restore to alternative location** on the **Target Machine** page.

Specify the target machine information for the File Restore.

Restore to original location Restore to alternative location

Target Machine Settings

Host Name/IP	<input type="text" value="<Machine Name/IP Address>"/>
Username	<input type="text"/>
Password	<input type="password"/>
Destination	<input type="text"/> <input type="button" value="Browse"/>

Resolving Conflicts

How should arcserve UDP Agent(Linux) resolve conflicting files

Overwrite existing files

Rename files

Skip existing files

Directory Structure

Whether to create root directory during restore

Create root directory

2. Enter the host name or the IP address of the target node.
3. Enter the user name and the password of the node.
4. Enter the path where the data is restored, or click **Browse** to select the folder where the data is restored and click **OK**.

5. Select one of the following options to resolve conflicting files:

Overwrite existing files

Specifies that if the file exists in the target machine then the backup file from the recovery point replaces the existing file.

Rename files

Specifies that if the file exists in the target machine, then a new file is created with the same file name and `.d2duplicate<x>` file extension. `<x>` specifies the number of times the file is restored. All the data is restored to the new file.

Skip existing files

Specifies that if the same file exists in the target machine then those files are not restored from the recovery point.

6. (Optional) Select **Create root directory**.

7. Click **Next**.

The **Advanced** page opens.

The target machine details are specified.

Specify the Advanced Settings

Specify the advanced settings to perform a scheduled recovery of your data. Scheduled recovery ensures that your data is recovered at the specified time even in your absence.

Follow these steps:

1. Set the start date and time by selecting one of the following options:

Run Now

Starts the file-level restore job as soon as you submit the job.

Set Starting Date and Time

Starts the file-level restore job at the specified date and time after submitting the job.

2. (Optional) Select **Estimate file size**.
3. (Optional) Select a script from the **Pre/Post Scripts Settings** option.

These scripts run script commands for actions to take before the start of the job and/or upon the completion of the job.

Note: The **Pre/Post Script Settings** fields are populated only if you already created a script file and placed it at the following location:

```
/opt/CA/d2dserver/usr/prepost
```

Note: For more information about creating the pre/post scripts, see *Manage Pre/Post Scripts for Automation* (see page 145).

4. Click **Next**.

The **Summary** page opens.

The advanced settings are specified.

(Optional) Manage Pre/Post Scripts for Automation

Pre/Post scripts let you run your own business logic at specific stages of a running job. You can specify when to run your scripts in **Pre/Post Script Settings** of the **Backup Wizard** and the **Restore Wizard** in the UI. The scripts can be run on the Backup Server depending on your setting.

Managing the pre/post script is a two part process, consisting of creating the pre/post script and placing the script in the prepost folder.

Create Pre/Post Scripts

Follow these steps:

1. Log into the Backup Server as a root user.
2. Create a script file using the environment variables in your preferred scripting language.

Pre/Post Script Environment Variables

To create your script, use the following environment variables:

D2D_JOBNAME

Identifies the name of the job.

D2D_JOBID

Identifies the job ID. Job ID is a number provided to the job when you run the job. If you run the same job again, you get a new job number.

D2D_TARGETNODE

Identifies the node that is being backed up or restored.

D2D_JOBTYPE

Identifies the type of the running job. The following values identify the D2D_JOBTYPE variable:

backup.full

Identifies the job as a full backup.

backup.incremental

Identifies the job as an incremental backup.

backup.verify

Identifies the job as a verify backup.

restore.bmr

Identifies the job as a bare-metal recovery (BMR). This is a restore job.

restore.file

Identifies the job as a file-level restore. This is a restore job.

D2D_SESSIONLOCATION

Identifies the location where the recovery points are stored.

D2D_PREPOST_OUTPUT

Identifies a temp file. The content of the first line of the temp file is displayed in the activity log.

D2D_JOBSTAGE

Identifies the stage of the job. The following values identify the D2D_JOBSTAGE variable:

pre-job-server

Identifies the script that runs on the Backup Server before the job starts.

post-job-server

Identifies the script that runs on the Backup Server after the job completes.

pre-job-target

Identifies the script that runs on the target machine before the job starts.

post-job-target

Identifies the script that runs on the target machine after the job completes.

pre-snapshot

Identifies the script that runs on the target machine before capturing the snapshot.

post-snapshot

Identifies the script that runs on the target machine after capturing the snapshot.

D2D_TARGETVOLUME

Identifies the volume that is backed up during a backup job. This variable is applicable for pre/post snapshot scripts for a backup job.

D2D_JOBRESULT

Identifies the result for a post job script. The following values identify the D2D_JOBRESULT variable:

success

Identifies the result as successful.

fail

Identifies the result as unsuccessful.

D2DSVR_HOME

Identifies the folder where Backup Server is installed. This variable is applicable for the scripts that run on the Backup Server.

The script is created.

Note: For all scripts, a return value of zero indicates success and a nonzero return value indicates failure.

Place the Script in the Prepost Folder and Verify

All the pre/post scripts for a Backup Server are centrally managed from the prepost folder at the following location:

```
/opt/CA/d2dserver/usr/prepost
```

Follow these steps:

1. Place the file in the following location of the Backup Server:

```
/opt/CA/d2dserver/usr/prepost
```

2. Provide the execution permission to the script file.
3. Log into the Arcserve UDP Agent (Linux) web interface.
4. Open the **Backup Wizard** or the **Restore Wizard** and navigate to the **Advanced** tab.
5. Select the script file in the **Pre/Post Script Settings** drop-down list and then submit the job.
6. Click **Activity Log** and verify that the script is executed to the specified backup job.

The script is executed.

The pre/post scripts are successfully created and placed in the prepost folder.

Create and Run the Restore Job

Create and run the restore job so that you can initiate the file-level recovery. Verify the recovery point information before you restore the files. If needed, you can go back and can change the restore settings on the wizard.

Follow these steps:

1. Verify the restore details on the **Summary** page of the **Restore Wizard**.
2. (Optional) Click **Previous** to modify the information that you have entered on any page of the **Restore Wizard**.
3. Enter a job name and click **Submit**.

The **Job Name** field has a default name initially. You can enter a new job name of your choice but you cannot leave this field empty.

The **Restore Wizard** closes. You can see the status of the job in the **Job Status** tab.

The restore job is successfully created and run.

Verify that Files are Restored

After the completion of restore job, verify that all the files are restored in the target node. Check the **Job History** and **Activity Log** tabs in the **Status** pane to monitor the progress of the restore process.

Follow these steps:

1. Navigate to the target machine where you restored data.
2. Verify that the required data from the recovery point is restored.

The files are successfully verified.

The file-level recovery is successfully performed.

How to Create a Bootable Live CD

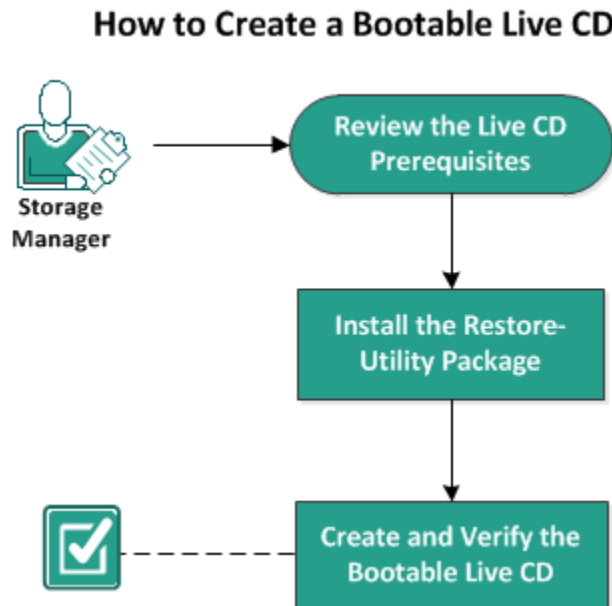
As a storage manager, you can create a bootable Live CD. When created, this bootable Live CD contains a complete read-only image of the computer operating system, and can be used to provide a temporary operating system functionality. This Live CD includes all your system settings and operating system files and can be used to perform the following functions:

- You can use Arcserve UDP Agent (Linux) without actually installing the product. This allows you to experience and evaluate the product without installing it or making any changes to the existing hard drive of your computer.
- You can install Arcserve UDP Agent (Linux) (to multiple servers) using only one setup package. Without a Live CD, you must install two separate files (.bin file and restore-utility package) to install Arcserve UDP Agent (Linux). The restore-utility package is included in the same Live CD setup package.
- You can perform a Bare Metal Recovery (BMR). You can use this Live CD to get the IP address of the target machine (which is required during the BMR).

The bin folder contains the scripts that you can run from the command line to create a bootable Live CD. The bin folder is located in the following path:

```
# /opt/Arcserve/d2dserver/bin
```

The following diagram displays the process to create a bootable Live CD:



The following list describes each task to create a bootable Live CD:

- [Review the Live CD Prerequisites](#) (see page 95)
- [Install the Restore-Utility Package](#) (see page 96)
- [Create and Verify the Bootable Live CD](#) (see page 96)

Review the Live CD Prerequisites

Consider the following prerequisites before you create a Live CD:

- You have the root login credentials to log into the Backup Server.
- You have read the Release Notes to understand the functions of a Live CD.
- You have knowledge of Linux scripting.
- You have installed the *mkisofs* tool in the Backup Server. The Backup Server uses the *mkisofs* tool to create the Live CD.iso file.
- You have at least 1024 MB free memory on your machine to boot and run the Live CD.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Install the Restore-Utility Package

You must install the restore-utility package to perform any restore operations. If you do not install the restore-utility package, you cannot perform file-level restore or BMR. You can install the restore-utility package during the installation of Arcserve UDP Agent (Linux). You can also download and install the restore-utility package any time after Arcserve UDP Agent (Linux) is installed.

After you install the restore-utility package, you can create a Live CD.

Follow these steps:

1. Log in to the Backup Server as a root user.
2. Navigate to the bin folder using the following command:

```
# cd /opt/Arcserve/d2dserver/bin
```
3. Run the following command to install the restore-utility package:

```
# ./configutility
```

A message is displayed prompting you to provide the path of the restore-utility package.

4. Provide the full path where you have downloaded the restore-utility package.
The installation begins.

The restore-utility package is installed.

Create and Verify the Bootable Live CD

Live CD creates the Linux Backup Server's environment without installing the software. Live CD facilitates BMR using IP in a private network.

Live CD is a complete bootable computer operating system which runs in the computer's memory, rather than loading from the hard disk. Live CD allows you to experience and evaluate an operating system without installing it or changing the existing operating system on the computer.

Follow these steps:

1. Navigate to the bin folder using the following command:

```
# cd /opt/Arcserve/d2dserver/bin
```

2. Run the following command to create a Live CD:

```
# ./makelivecd
```

3. Navigate to the following location and verify that the LiveCD.iso file is created:

```
/opt/Arcserve/d2dserver/packages
```


You have successfully created and verified the bootable Live CD. If you want to use the Live CD on a virtual network, you can directly mount the LiveCD.iso file to the virtual machine. If you want to use you the Live CD on a physical machine, then you must burn the LiveCD.iso image on a media file (CD or DVD) and then use the media file to boot your machine.

How to Create a CentOS-Based Live CD

As a storage manager, you can create a bootable CentOS-based Live CD. CentOS Live CD is an in-memory computing environment based on CentOS. The purpose of this Live CD is to provide users the capability to experience the CentOS functionality without installing CentOS. The Live CD runs in memory without impacting the hard disk. The changes that you make in the Live CD runtime environment are lost after you restart the machine.

This Live CD includes all your system settings and operating system files and can be used to perform the following functions:

- You can use Arcserve UDP Agent (Linux) without actually installing the product. This allows you to experience and evaluate the product without installing it or making any changes to the existing hard drive of your computer.
- You can perform a Bare Metal Recovery (BMR). You can use this Live CD to get the IP address of the target machine (which is required during the BMR).

When to use the CentOS-based Live CD:

When the default Live CD cannot identify storage device and network device because of the lack of the device driver.

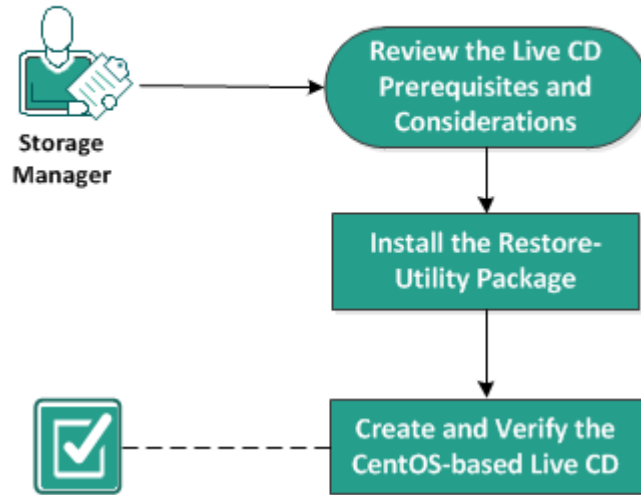
Note: The recovery points that you want to restore does not include the device drivers for the storage system of the target BMR machine. As a result, Arcserve UDP Agent (Linux) will block any attempt to perform a BMR job at an early stage.

The bin folder contains the scripts that you can run from the command line to create a bootable Live CD. The bin folder is located in the following path:

```
# /opt/Arcserve/d2dserver/bin
```

The following diagram displays the process to create a CentOS-based Live CD:

How to Create a CentOS-Based Live CD



Perform the following tasks to create a CentOS-based Live CD:

- [Review the Live CD Prerequisites and Considerations](#) (see page 98)
- [Install the Restore-Utility Package](#) (see page 100)
- [Create and Verify the CentOS-based Live CD](#) (see page 100)

Review the Live CD Prerequisites and Considerations

Before you create a CentOS-based Live CD, review the following table which compares the default Live CD to the CentOS-based Live CD:

Parameters	Default Live CD	CentOS-based Live CD
Backup Server Installation Media	Supported	Not Supported
Desktop UI	Not supported. Users must use a browser on a Windows machine to browse the Backup Server web UI.	Supported. CentOS-based Live CD includes a browser. Users do not need any additional browser to browse the Backup Server web UI..
Image size	Approximately 400 MB. The image can be burned on a CD.	More than 800 MB. The image must be burned on a DVD.

Parameters	Default Live CD	CentOS-based Live CD
Additional device driver for the Live CD	Not supported	Supported
Local BMR (Recover machine without installing another Backup Server)	Supported	Supported
PXE Boot Image	Supported	Not supported
Remove CD/ISO from the BMR target machine after the machine is booted	Supported	Not supported. DVD/ISO must be mounted on the BMR target machine all the time during the recovery process, until the BMR job is finished and the machine is rebooted.
Live CD operating system environment in English	Yes	Yes. Desktop UI is also in English
Localized language for the Backup Server Web UI	Yes	Yes
Node type support	Support physical machine, VMWare ESX server, RHEV, OVM, Citrix Xen VM	Only support physical machine and VMware ESX server VM

Consider the following prerequisites before you create a CentOS-based Live CD:

- Verify that you have installed the following software packages on the Backup Server:
 - genisoimage
 - squashfs-tools
- The CentOS-based Live CD can boot from a physical machine and ESX server VM only. It does not support other virtualization solutions.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Install the Restore-Utility Package

You must install the restore-utility package to perform any restore operations. If you do not install the restore-utility package, you cannot perform file-level restore or BMR. You can install the restore-utility package during the installation of Arcserve UDP Agent (Linux). You can also download and install the restore-utility package any time after Arcserve UDP Agent (Linux) is installed.

After you install the restore-utility package, you can create a Live CD.

Follow these steps:

1. Log in to the Backup Server as a root user.
2. Navigate to the bin folder using the following command:

```
# cd /opt/Arcserve/d2dserver/bin
```
3. Run the following command to install the restore-utility package:

```
# ./configutility
```

A message is displayed prompting you to provide the path of the restore-utility package.

4. Provide the full path where you have downloaded the restore-utility package.
The installation begins.

The restore-utility package is installed.

Create and Verify the CentOS-based Live CD

You can use this Live CD to boot a BMR target machine and then run the BMR job. The following files are used to create the CentOS-based Live CD:

makelivecd.centos

A script used to remaster the CentOS Live CD.

CentOS-6.X-i386-LiveCD.iso

A CentOS Live CD ISO image. The image can be downloaded from the CentOS website.

The recovery point that is restored does not contain device driver for target BMR machine's storage system. Arcserve UDP Agent (Linux) blocks such BMR job at an early stage.

Follow these steps:

1. Prepare the device drivers (*.ko and *.rpm files) for CentOS and store them in a folder.

Example: Store the device drivers in the /tmp/drivers folder.

Note: You must provide the device driver that matches the kernel version of the CentOS Live CD.

2. Access the CentOS website and download the 32-bit CentOS 6.0 or later Live CD to the /tmp folder on the Backup Server.

The CentOS-6.X-i386-LiveCD.iso file is downloaded.

Important: If you use CentOS 6.5 Live CD to restore from network locations, manually install mount.nfs or mount.cifs on the Live CD server.

3. Navigate to the bin folder (/opt/Arcserve/d2dserver/bin) and run the following command:

```
makelivecd.centos <full_path_to_CentOS_live_cd> <path_where_device_drivers_are_stored>
```

Example: makelivecd.centos <full_path_to_CentOS_live_cd> /tmp/drivers

The script creates the Arcserve Unified Data Protection Agent for Linux Live CD based on CentOS and stores the ISO image file at the following location:

```
/opt/Arcserve/d2dserver/packages/CentOS-LiveCD-for-UDP_Agent_Linux.iso
```

4. Navigate to the packages folder and verify that the CentOS-LiveCD-for-UDP_Aget_Linux.iso file is included in the folder.

The CentOS-based Live CD is created and verified.

You have successfully created a CentOS-based Live CD.

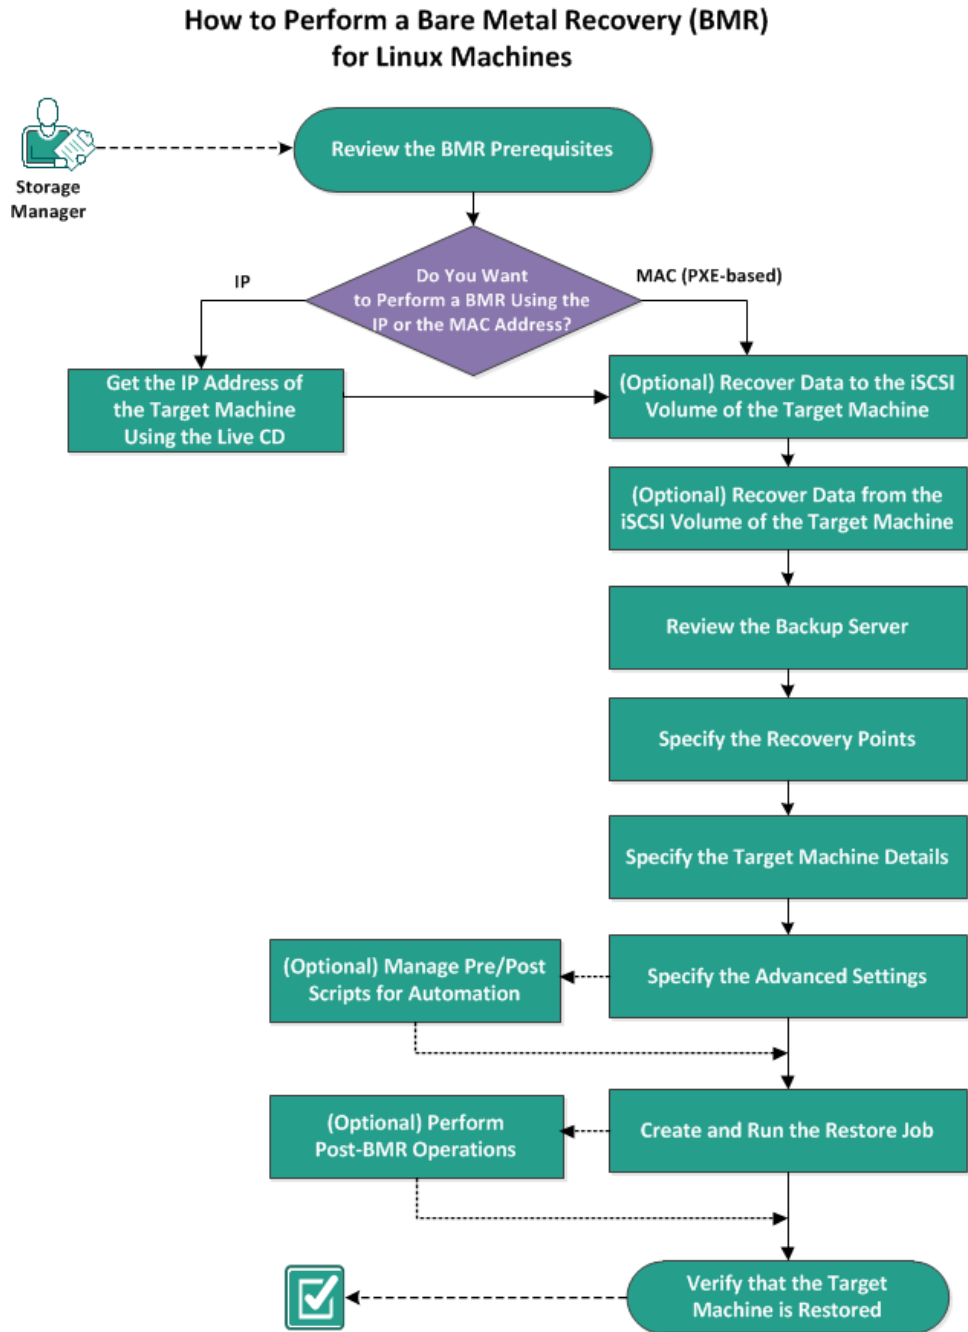
How to Perform a Bare Metal Recovery (BMR) for Linux Machines

A BMR restores the operating system and software applications, and recovers all the backed-up data. BMR is the process of restoring a computer system from *bare metal*. Bare metal is a computer without any operating system, drivers, and software applications. After the restore is complete, the target machine automatically reboots in the same operating environment as the backup source node and all data is restored.

A complete BMR is possible because when you back up data, the backup also captures information related to the operating system, installed applications, drivers, and so on.

You can perform a BMR using the IP address or the Media Access Control (MAC) address of the target machine. If you boot the target machine using the Arcserve UDP Agent (Linux) Live CD, you can get the IP address of the target machine.

The following diagram displays the process to perform a BMR:



Complete the following tasks to perform a BMR:

- [Review the BMR Prerequisites](#) (see page 103)
- [Get the IP Address of the Target Machine Using the Live CD](#) (see page 104)
- [\(Optional\) Recover Data to the iSCSI Volume of the Target Machine](#) (see page 105)
- [\(Optional\) Recover Data from the iSCSI Volume to the Target Machine](#) (see page 106)
- [Review the Backup Server](#) (see page 107)
- [Specify the Recovery Points](#) (see page 108)
- [Specify the Target Machine Details](#) (see page 110)
- [Specify the Advanced Settings](#) (see page 111)
- [\(Optional\) Manage Pre/Posts Scripts for Automation](#) (see page 112)
- [Create and Run the Restore Job](#) (see page 115)
- [\(Optional\) Perform Post-BMR Operations](#) (see page 116)
- [Verify that the Target Machine is Restored](#) (see page 122)

Review the BMR Prerequisites

Consider the following options before performing a BMR:

- You have a valid recovery point and the encryption password, if any, for restore.
- You have a valid target machine for BMR.
- You have created the Arcserve UDP Agent (Linux) Live CD.
- If you want to perform a BMR using the IP address, you must get the IP address of the target machine using the Live CD.
- If you want to perform a PXE-based BMR using the MAC address, you must have the MAC address of the target machine.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Get the IP Address of the Target Machine Using the Live CD

Before performing a BMR using the IP address, you need to get the IP address of the target machine. A bare-metal machine does not have any IP address initially. So, you have to boot the bare-metal machine using the default Live CD, which is Arcserve UDP Agent (Linux) Live CD, or the CentOS-based Live CD to get the IP address. After you get the IP address of the target machine, you can configure the static IP of the target machine.

Follow these steps:

1. Insert the Live CD or mount the .iso file of the Live CD into the CD-ROM drive of the target node.
2. Boot the target machine from CD-ROM.

The target machine boots into the Arcserve UDP Agent (Linux) Live CD environment. On the screen, the IP address of the target machine is displayed.

3. To configure the static IP of the target machine using the default Live CD, follow these steps:

- a. On the target machine's screen, press Enter to enter the shell environment.
- b. Run the following command to configure the static IP:

```
ifconfig <NIC name> <static IP address> netmask <netmask>  
route add default gw <gateway IP address> <NIC name>
```

Note: The Network Interface Card (NIC) name depends on your hardware. For example, the typical NIC names are eth0 or em0.

4. To configure the static IP of the target machine using the CentOS-based Live CD, follow these steps:
 - a. Open a terminal window on the target machine by clicking Applications, System Tools, Terminal.
 - b. Run the following commands:

```
sudo ifconfig <NIC name> <static IP address> netmask <netmask>  
sudo route add default gw <gateway IP address> <NIC name>
```

The static IP is configured.

The IP address of the target machine is acquired.

Important! Maintain a record of this IP address as it is used in the **Restore Wizard** when you have to specify the target machine details.

(Optional) Recover Data to the iSCSI Volume of the Target Machine

You can integrate the iSCSI volume to the target machine and make that volume a part of the target machine. Then you can restore data to the iSCSI volume of the target machine. By doing so, you can manage data and transfer data over a network.

Important! When you integrate the iSCSI volume with the target machine, you will lose all the existing data from the iSCSI volume.

Follow these steps:

1. Insert the Arcserve UDP Agent (Linux) Live CD or mount the iso file of the Arcserve UDP Agent (Linux) Live CD into the CD-ROM drive of the target machine.
2. Boot the target machine from the CD-ROM.

The target machine boots into the Arcserve UDP Agent (Linux) Live CD environment. On the screen, the IP address of the target machine is displayed.

3. Enter the shell environment of the target machine.
4. Run the following command to start the iSCSI initiator daemon:

```
/etc/init.d/iscsid start
```

5. Run a discovery script to discover the iSCSI target host.

```
iscsiadm -m discovery -t sendtargets -p <ISCSI-SERVER-IP-ADDRESS>:<Port_Number>
```

The default port value of iSCSI target host is 3260.

6. Make a note of the iSCSI Qualified Name (IQN) of the iSCSI target host found by the discovery script before you manually log into the discovered target.
7. List the available block device of the target node.

```
#fdisk -l
```

8. Log in to the discovered target.

```
iscsiadm -m node -T <iSCSI Target IQN name> -p <ISCSI-SERVER-IP-ADDRESS>:<Port_Number> -l
```

You can see a block device in the /dev directory of the target node.

9. Run the following command to obtain the new device name:

```
#fdisk -l
```

You can see an additional device named /dev/sd<x> on the target node.

The iSCSI volume is integrated with the target volume.

(Optional) Recover Data from the iSCSI Volume to the Target Machine

If you have stored your data in an iSCSI target volume, you can connect to the iSCSI volume and recover data. The iSCSI volume lets you manage data and transfer data over a network.

Follow these steps:

1. Insert the Arcserve UDP Agent (Linux) Live CD or mount the iso file of the Arcserve UDP Agent (Linux) Live CD into the CD-ROM drive of the target machine.
2. Boot the target machine from the CD-ROM.

The target machine boots into the Arcserve UDP Agent (Linux) Live CD environment. On the screen, the IP address of the target machine is displayed.

3. Enter the shell environment of the target machine.
4. Run the following command to start the iSCSI initiator daemon:

```
/etc/init.d/iscsid start
```

5. Run a discovery script to discover the iSCSI target host.

```
iscsiadm -m discovery -t sendtargets -p <ISCSI-SERVER-IP-ADDRESS>:<Port_Number>
```

The default port value of iSCSI target host is 3260.

6. Make a note of the iSCSI Qualified Name (IQN) of the iSCSI target host found by the discovery script before you manually log into the discovered target.
7. List the available block device of the target node.

```
#fdisk -l
```

8. Log in to the discovered target.

```
iscsiadm -m node -T <iSCSI Target IQN name> -p <ISCSI-SERVER-IP-ADDRESS>:<Port_Number> -l
```

You can see a block device in the /dev directory of the target node.

9. Run the following command to obtain the new device name:

```
#fdisk -l
```

You can see an additional device named /dev/sd<x> on the target node.

For example, consider the name of the device is /dev/sdc. This device name is used to create a partition and a file system in the following steps.

10. Mount the iSCSI volume using the following commands:

```
# mkdir /iscsi
```

```
# mount /dev/sdc1 /iscsi
```

Note: When you specify the session location in the Restore Wizard, you need to select Local and enter the path /iscsi.

Example: <path>/iscsi

The target machine can now connect to the iSCSI volume and can recover data from the iSCSI volume.

Review the Backup Server

When you open the **Restore Wizard**, review the Backup Server where you want to perform the restore operation.

Follow these steps:

1. Access the Restore Wizard in one of the following ways:

- From Arcserve UDP:

- a. Log in to Arcserve UDP.
- b. Click the **resources** tab.
- c. Select **All Nodes** in the left pane.
All the added nodes are displayed in the center pane.
- d. In the center pane, select the node and click **Actions**.
- e. Click **Restore** from the **Actions** dropdown menu.

The Arcserve UDP Agent (Linux) web interface opens. The restore type selection dialog is displayed in the Agent UI.

- f. Select the restore type and click **OK**.

Note: You are automatically logged in to the agent node and the **Restore Wizard** opens from the agent node.

- From Arcserve UDP Agent (Linux):

- a. Open the Arcserve UDP Agent (Linux) web interface.

Note: During the installation of Arcserve UDP Agent (Linux), you received the URL to access and manage the server. Log in to Arcserve UDP Agent (Linux).

- b. Click **Restore** from the **Wizard** menu and select **Bare Metal Recovery (BMR)**.

The **Backup Server** page of the **Restore Wizard - BMR** opens.

2. Verify the server from the **Backup Server** drop-down list in the **Backup Server** page.
You cannot select any option from the **Backup Server** drop-down list.
3. Click **Next**.
The **Recovery Points** page of the **Restore Wizard - BMR** opens.

The Backup Server is specified.

Specify the Recovery Points

Each time that you perform a backup, a recovery point is created. Specify the recovery point information in the **Restore Wizard** so that you can recover the exact data that you want. You can restore specific files or all files depending on your requirement.

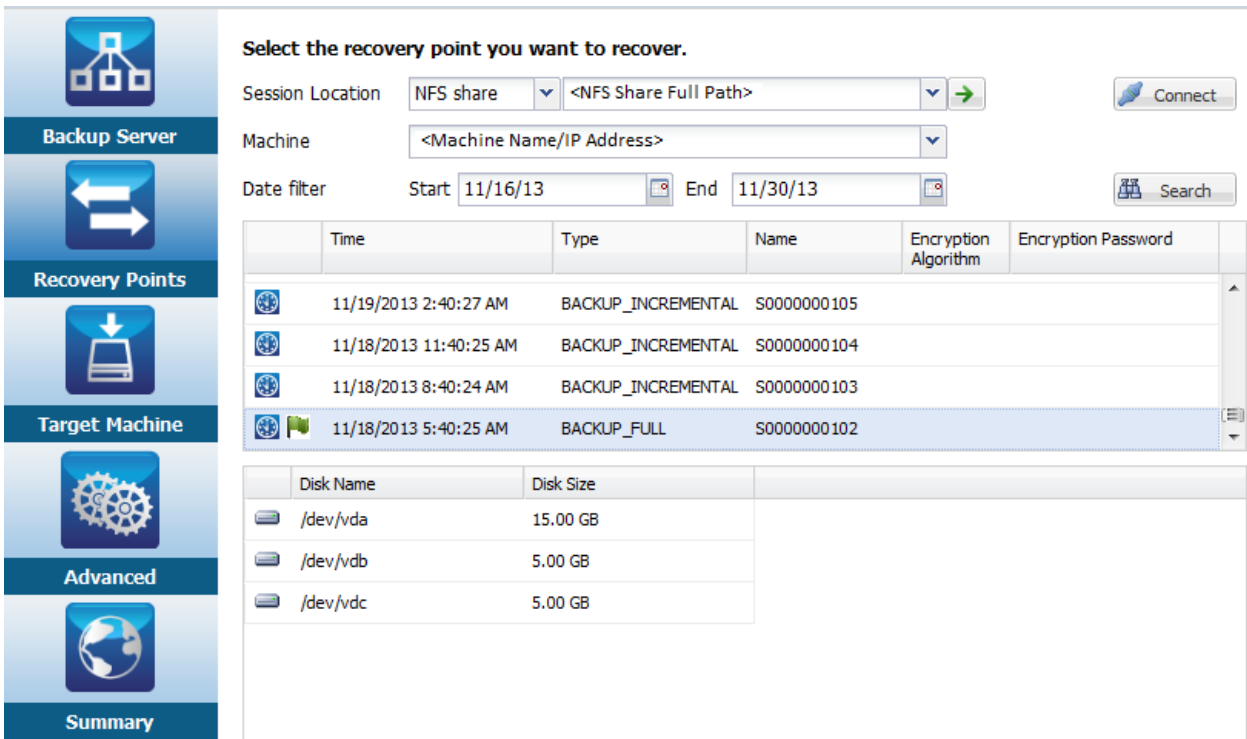
Important! To perform a BMR from a recovery point, the root volume and the boot volume must be present in the recovery point.

Follow these steps:

1. Perform one of the following steps depending on your backup storage.
 - Perform the following steps to access the recovery points if the recovery points are stored on a mobile device:
 - a. Start the target machine using the Live CD.
 - b. Log into the Arcserve UDP Agent (Linux) web interface from the Live CD.
 - c. Open the **BMR Wizard**.
 - d. Navigate to the **Recovery Points** page.
 - e. Select **Local** as the **Session Location** on the **Recovery Points** page of the **BMR Wizard**.
 - Perform the following steps if the session location is NFS share or CIFS share:
 - a. Select a session from the **Session Location** drop-down list and enter the full path of the share.

 For example, consider the Session Location as NFS share, xxx.xxx.xxx.xxx as the IP address of the NFS share, and the folder name is *Data*. You would enter xxx.xxx.xxx.xxx:/Data as the NFS share location.

Note: If the backed up data is stored in Source local, then you must first convert the source node to an NFS server, and then share the session location.



Select the recovery point you want to recover.

Session Location: NFS share <NFS Share Full Path> Connect

Machine: <Machine Name/IP Address>

Date filter: Start 11/16/13 End 11/30/13 Search

	Time	Type	Name	Encryption Algorithm	Encryption Password
	11/19/2013 2:40:27 AM	BACKUP_INCREMENTAL	S0000000105		
	11/18/2013 11:40:25 AM	BACKUP_INCREMENTAL	S0000000104		
	11/18/2013 8:40:24 AM	BACKUP_INCREMENTAL	S0000000103		
	11/18/2013 5:40:25 AM	BACKUP_FULL	S0000000102		

Disk Name	Disk Size
/dev/vda	15.00 GB
/dev/vdb	5.00 GB
/dev/vdc	5.00 GB

2. Click **Connect**.

All the nodes that have been backed up to this location get listed in the **Machine** drop-down list.

3. Select the node that you want to restore from the **Machine** drop-down list.

All the recovery points of the selected node get listed.

4. Apply the date filter to display the recovery points that are generated between the specified date and click **Search**.

Default: Recent two weeks.

All the recovery points available between the specified dates are displayed.

5. Select the recovery point that you want to restore and click **Next**.

The **Target Machine** page opens.

The recovery point is specified.

Specify the Target Machine Details

Specify the target machine details so that data is restored to that machine. A target machine is a bare metal machine where you will perform a BMR. If you restore using the IP address, you need the IP address of the target machine that you previously recorded at the beginning of this process. If you restore using the Media Access Control (MAC) address, you need the MAC address of the target machine.

Follow these steps:

1. Enter the MAC address or the IP address of the target machine in the **MAC/IP Address** field.

2. Enter a name in the **Host Name** field.

The target machine will use this name as the host name after the restore process is complete.

3. Select one of the following options as the network:

DHCP

Automatically configures the IP address. This is the default option. Use this option if you have a Dynamic Host Configuration Protocol (DHCP) server to restore with the DHCP network.

Static IP

Manually configures the IP address. If you select this option then, enter the **IP Address**, **Subnet Mask**, and **Default Gateway** of the target machine.

Important! Ensure that the Static IP is not used by any other machines on the network during the restore process.

4. (Optional) Select the **Reboot** option to automatically restart the target node after the BMR is complete.

5. Click **Next**.

The **Advanced** page opens.

The target machine details are specified.

Specify the Advanced Settings

Specify the advanced settings to perform a scheduled BMR of your data. Scheduled BMR ensures that your data is recovered at the specified time even in your absence.

Follow these steps:

1. Set the start date and time by selecting one of the following options:

Run Now

Starts the restore job as soon as you submit the job.

Set Special Time

Starts the restore job at the specified time after submitting the job.

2. (Optional) Select a script from the **Pre/Post Scripts Settings** option for the Backup Server and the target machine.

These scripts run script commands for actions to take before the start of the job and/or upon the completion of the job.

Note: The **Pre/Post Script Settings** fields are populated only if you already created a script file and placed it at the following location:

```
/opt/CA/d2dserver/usr/prepost
```

Note: For more information about creating the pre/post scripts, see *Manage Pre/Post Scripts for Automation* (see page 145).

3. (Optional) Click **Show More Settings** to display more settings for BMR.
4. (Optional) Reset the password for the specified user name for the recovered target machine.
5. (Optional) Enter the full path of the backup storage location of the recovery points in **Recover Point Local Access**.

6. (Optional) Enter the full name of the disk in the **Disks** field to exclude those disks on the target machine from participating in the recovery process.
7. (Optional) Select **Enable Wake-on-LAN** if you are performing Preboot Execution Environment (PXE) BMR.

Note: The **Enable Wake-on-LAN** option is applicable only for physical machines. Ensure whether you have enabled the Wake-on-LAN settings in the BIOS settings of your physical machine.

8. Click **Next**.

The **Summary** page opens.

The advanced settings are specified.

(Optional) Manage Pre/Post Scripts for Automation

Pre/Post scripts let you run your own business logic at specific stages of a running job. You can specify when to run your scripts in **Pre/Post Script Settings** of the **Backup Wizard** and the **Restore Wizard** in the UI. The scripts can be run on the Backup Server depending on your setting.

Managing the pre/post script is a two part process, consisting of creating the pre/post script and placing the script in the prepost folder.

Create Pre/Post Scripts

Follow these steps:

1. Log into the Backup Server as a root user.
2. Create a script file using the environment variables in your preferred scripting language.

Pre/Post Script Environment Variables

To create your script, use the following environment variables:

D2D_JOBNAME

Identifies the name of the job.

D2D_JOBID

Identifies the job ID. Job ID is a number provided to the job when you run the job. If you run the same job again, you get a new job number.

D2D_TARGETNODE

Identifies the node that is being backed up or restored.

D2D_JOBTYPE

Identifies the type of the running job. The following values identify the D2D_JOBTYPE variable:

backup.full

Identifies the job as a full backup.

backup.incremental

Identifies the job as an incremental backup.

backup.verify

Identifies the job as a verify backup.

restore.bmr

Identifies the job as a bare-metal recovery (BMR). This is a restore job.

restore.file

Identifies the job as a file-level restore. This is a restore job.

D2D_SESSIONLOCATION

Identifies the location where the recovery points are stored.

D2D_PREPOST_OUTPUT

Identifies a temp file. The content of the first line of the temp file is displayed in the activity log.

D2D_JOBSTAGE

Identifies the stage of the job. The following values identify the D2D_JOBSTAGE variable:

pre-job-server

Identifies the script that runs on the Backup Server before the job starts.

post-job-server

Identifies the script that runs on the Backup Server after the job completes.

pre-job-target

Identifies the script that runs on the target machine before the job starts.

post-job-target

Identifies the script that runs on the target machine after the job completes.

pre-snapshot

Identifies the script that runs on the target machine before capturing the snapshot.

post-snapshot

Identifies the script that runs on the target machine after capturing the snapshot.

D2D_TARGETVOLUME

Identifies the volume that is backed up during a backup job. This variable is applicable for pre/post snapshot scripts for a backup job.

D2D_JOBRESULT

Identifies the result for a post job script. The following values identify the D2D_JOBRESULT variable:

success

Identifies the result as successful.

fail

Identifies the result as unsuccessful.

D2DSVR_HOME

Identifies the folder where Backup Server is installed. This variable is applicable for the scripts that run on the Backup Server.

The script is created.

Note: For all scripts, a return value of zero indicates success and a nonzero return value indicates failure.

Place the Script in the Prepost Folder and Verify

All the pre/post scripts for a Backup Server are centrally managed from the prepost folder at the following location:

```
/opt/CA/d2dserver/usr/prepost
```

Follow these steps:

1. Place the file in the following location of the Backup Server:

```
/opt/CA/d2dserver/usr/prepost
```
2. Provide the execution permission to the script file.
3. Log into the Arcserve UDP Agent (Linux) web interface.
4. Open the **Backup Wizard** or the **Restore Wizard** and navigate to the **Advanced** tab.
5. Select the script file in the **Pre/Post Script Settings** drop-down list and then submit the job.
6. Click **Activity Log** and verify that the script is executed to the specified backup job.
The script is executed.

The pre/post scripts are successfully created and placed in the prepost folder.

Create and Run the Restore Job

Create and run the restore job so that you can initiate the process of BMR. Verify the recovery point information before you perform a BMR. If needed, you can go back and can change the restore settings.

Follow these steps:

1. Verify the restore details on the **Summary** page of the **Restore Wizard**.
2. (Optional) Click **Previous** to modify the restore settings on any of the **Restore Wizard** pages.
3. Enter a job name and click **Submit**.

The **Job Name** field has a default name initially. You can enter a new job name of your choice but you cannot leave this field empty.

The **Restore Wizard** closes. You can see the job in the **Job Status** tab. If you use the IP address for the BMR, the target machine automatically reboots to the same operating system as the backup source after the BMR process.

If you use the MAC address for BMR, the status in the **Job Status** tab changes to *Waiting for target node startup*.

4. (Optional) For BMR using the MAC address, start the target machine when you see the *Waiting for target node startup* message in the **Job Status** tab.

Note: If the target machine is already started before you submit the restore job, you must restart the target machine. Ensure that BIOS is configured to boot from the network.

The status in the **Job Status** column changes to **Restoring volume**. This indicates the restore is in progress. After the restore job is complete, the target machine automatically reboots with the same operating system as the backup source.

The restore job was successfully created and run.

(Optional) Perform Post-BMR Operations

The following topics are optional configuration settings that you may have to perform after a BMR:

Configure X Windows

When you perform a BMR across a dissimilar hardware, X Windows of the restored OS does not function properly and the target node displays an error dialog. The error dialog appears because the display configuration has changed. To resolve this error, follow the instructions in the error dialog to configure the graphic card. After that, you can see the X Windows and the desktop UI.

Configure the System Fully Qualified Domain Name (FQDN)

When you need an FQDN, then you must configure the FQDN. The BMR process does not automatically configure the FQDN.

Maximum character count for FQDN: 63

Follow these steps to configure the FQDN:

1. Edit the */etc/hosts* file and provide the IP Address, the FQDN name, and the server name.

```
#vi /etc/hosts
```

```
ip_of_system servername.domainname.com servername
```

2. Restart the network service.

```
#/etc/init.d/network restart
```

3. Verify the host name and the FQDN name.

```
#hostname
```

```
servername
```

```
#hostname -f
```

```
servername.domainname.com
```

The FQDN is configured.

Extend the Data Volume after a BMR on Dissimilar Disks

When you perform a BMR to a larger disk than the disk on the original node, some disk space is left unused. The BMR operation does not automatically process the unused disk space. You can format the disk space to a separate partition or resize the existed partition with the unused disk space. The volume that you want to resize must be unused, so you must avoid resizing a system volume. In this section, we will focus on how to extend a data volume with the unused disk space.

Note: To avoid data loss, resize your volumes immediately after the BMR process. You can also back up the node before starting the volume resizing task.

When the target machine successfully restarts after the BMR, you can extend the data volume.

Raw partition volume

For example, a 2-GB disk in the session is restored to a 16-GB disk named `/dev/sdb` with only one partition. The `/dev/sdb1` raw partition is directly mounted on the `/data` directory.

This example is used to explain the procedure of extending Raw partition volume.

Follow these steps:

1. Check the status of the `/dev/sdb1` volume.

```
# df -h /dev/sdb1
/dev/sdb1      2.0G 40M 1.9G 3% /data
```

2. Umount the `/dev/sdb1` volume.

```
# umount /data
```

3. Resize `/dev/sdb1` to occupy the entire disk space using the `fdisk` command.

To perform this operation, first delete your existing partition and then recreate it with the same start sector number. The same start sector number is responsible for avoiding the data loss.

```
# fdisk -u /dev/sdb
Command (m for help): p
Disk /dev/sdb: 17.1 GB, 17179869184 bytes
255 heads, 63 sectors/track, 2088 cylinders, total 33554432 sectors
Units = sectors of 1 * 512 = 512 bytes
Device Boot  Start    End  Blocks  Id System
/dev/sdb1    63  4192964  2096451  83  Linux
Command (m for help): d
```

Selected partition 1

Command (m for help): n

Command action

e extended

p primary partition (1-4)

p

Partition number (1-4): 1

First sector (63-33554431, default 63):

Using default value 63

Last sector or +size or +sizeM or +sizeK (63-33554431, default 33554431):

Using default value 33554431

Command (m for help): p

Disk /dev/sdb: 17.1 GB, 17179869184 bytes

255 heads, 63 sectors/track, 2088 cylinders, total 33554432 sectors

Units = sectors of 1 * 512 = 512 bytes

Device	Boot	Start	End	Blocks	Id	System
--------	------	-------	-----	--------	----	--------

/dev/sdb1		63	33554431	16777184+	83	Linux
-----------	--	----	----------	-----------	----	-------

Command (m for help): w

The partition changes to the same start sector number as the original partition and the end sector number is 33554431.

4. Resize the volume using `resize2fs` command. If necessary, first run the `e2fsck` command.

```
# e2fsck -f /dev/sdb1
```

```
# resize2fs /dev/sdb1
```

5. Mount the volume to the mount point and check the volume status again.

```
# mount /dev/sdb1 /data
```

```
# df -h /dev/sdb1
```

```
/dev/sdb1      16G  43M 16G  1% /data
```

The volume is extended to 16 GB and is ready for use.

LVM volume:

For example, an 8-GB disk in the session is restored to a 16-GB disk named */dev/sdc* with only one partition. The */dev/sdc1* raw partition is used as the only physical volume of the */dev/mapper/VGTest-LVTest* LVM logical volume whose mount point is */lvm*.

This example is used to explain the procedure of extending LVM volume.

Follow these steps:

1. Check the status of the */dev/mapper/VGTest-LVTest* volume.

```
# lvsdisplay -m /dev/mapper/VGTest-LVTest
--- Logical volume ---
LV Name      /dev/VGTest/LVTest
VG Name      VGTest
LV UUID      udoBlx-XKBS-1Wky-3FVQ-mxMf-FayO-tpfPl8
LV Write Access  read/write
LV Status     available
# open        1
LV Size       7.88 GB
Current LE    2018
Segments      1
Allocation    inherit
Read ahead sectors  0
Block device  253:2
---Segments---
Logical extent 0 to 2017:
Type          linear
Physical volume /dev/sdc1
Physical extents 0 to 2017
```

The physical volume is */dev/sdc1*, the volume group is *VGTest*, and the logical volume is */dev/VGTest/LVTest* or */dev/mapper/VGTest-LVTest*.

2. Umount the */dev/mapper/VGTest-LVTest* volume.

```
# umount /lvm
```

3. Disable the volume group in which the `/dev/sdc1` physical volume is located.

```
# vgchange -a n VGTest
```

4. Create a partition to occupy the unused disk space using the `fdisk` command.

```
# fdisk -u /dev/sdc
```

```
Command (m for help): p
```

```
Disk /dev/sdc: 17.1 GB, 17179869184 bytes
```

```
255 heads, 63 sectors/track, 2088 cylinders, total 33554432 sectors
```

```
Units = sectors of 1 * 512 = 512 bytes
```

```
Device Boot Start End Blocks Id System
```

```
/dev/sdc1 63 16777215 8388576+ 83 Linux
```

```
Command (m for help): n
```

```
Command action
```

```
e extended
```

```
p primary partition (1-4)
```

```
p
```

```
Partition number (1-4): 2
```

```
First sector (16777216-33554431, default 16777216):
```

```
Using default value 16777216
```

```
Last sector or +size or +sizeM or +sizeK (16777216-33554431, default 33554431):
```

```
Using default value 33554431
```

```
Command (m for help): p
```

```
Disk /dev/sdc: 17.1 GB, 17179869184 bytes
```

```
255 heads, 63 sectors/track, 2088 cylinders, total 33554432 sectors
```

```
Units = sectors of 1 * 512 = 512 bytes
```

```
Device Boot Start End Blocks Id System
```

```
/dev/sdc1 63 16777215 8388576+ 83 Linux
```

```
/dev/sdc2 16777216 33554431 8388608 83 Linux
```

```
Command (m for help): w
```

```
The /dev/sdc2 partition is created.
```


5. Create a new physical volume.

```
# pvcreate /dev/sdc2
```

6. Extend the volume group size.

```
# vgextend VGTest /dev/sdc2
```

7. Enable the volume group that you have already disabled.

```
# vgchange -a y VGTest
```

8. Extend the logical volume size using the lvextend command.

```
# lvextend -L +8G /dev/VGTest/LVTest
```

9. Resize the volume using the resize2fs command. If necessary, first run the e2fsck command.

```
# e2fsck -f /dev/mapper/VGTest-LVTest
```

```
# resize2fs /dev/mapper/VGTest-LVTest
```

10. Mount the volume to the mount point and check the volume status again.

```
# mount /dev/mapper/VGTest-LVTest /vm
```

```
# lvsdisplay -m /dev/mapper/VGTest-LVTest
```

```
---Logical volume---
```

```
LV Name      /dev/VGTest/LVTest
```

```
VG Name      VGTest
```

```
LV UUID      GTP0a1-kUL7-WUL8-bpbM-9eTR-SVzl-WgA11h
```

```
LV Write Access  read/write
```

```
LV Status      available
```

```
# open        0
```

```
LV Size        15.88 GB
```

```
Current LE     4066
```

```
Segments       2
```

```
Allocation     inherit
```

```
Read ahead sectors  0
```

```
Block device    253:2
```

```
--- Segments ---
```

```
Logical extent 0 to 2046:
```

```
Type          linear
Physical volume /dev/sdc1
Physical extents 0 to 2046
Logical extent 2047 to 4065:
```

```
Type          linear
Physical volume /dev/sdc2
Physical extents 0 to 2018
```

The LVM volume extends to 16 GB and is ready for use.

Verify that the Target Node is Restored

After the completion of restore job, verify that the target node is restored with relevant data.

Follow these steps:

1. Navigate to the target machine that you restored.
2. Verify that the target machine has all the information that you backed up.

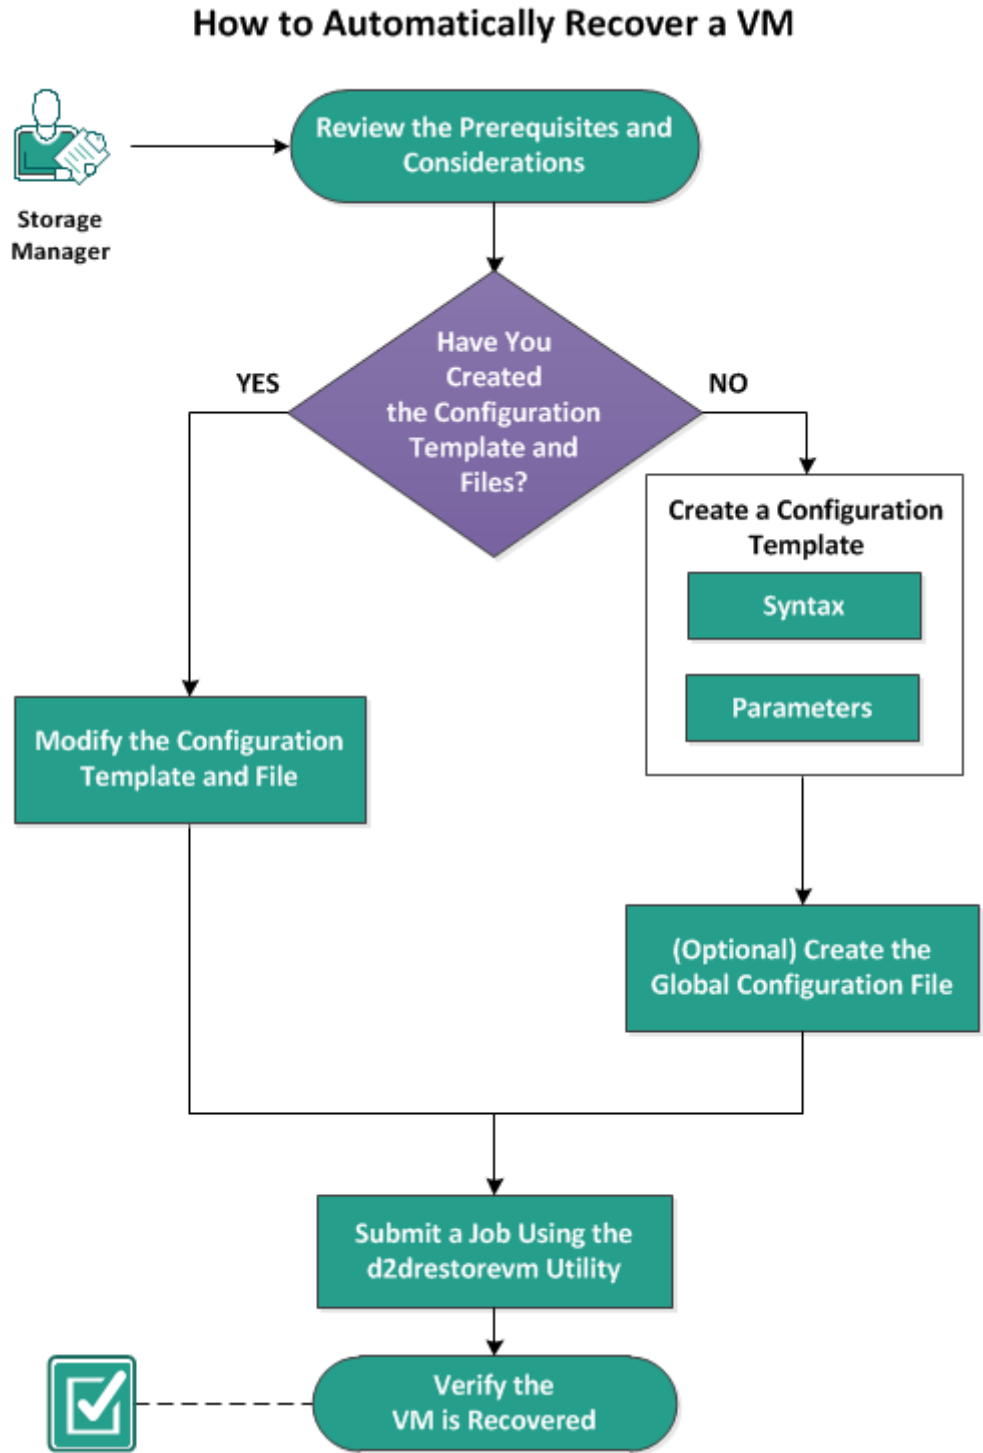
The target machine is successfully verified.

The BMR is successfully performed for Linux Machines.

How to Automatically Recover a Virtual Machine

You can recover a virtual machine (VM) from the command line of the Backup Server using the `d2drestorevm` utility. The `d2drestorevm` utility automates the process of performing a BMR without the need to manually boot the VM using a Live CD.

The following diagram displays the process to recover a virtual machine from the command line using the d2drestorevm utility:



Perform these tasks to automatically recover a VM:

- [Review the Prerequisites and Considerations](#) (see page 124)
- [Create a Configuration Template](#) (see page 127)
- [\(Optional\) Create the Global Configuration File](#) (see page 130)
- [Modify the Configuration Template and File](#) (see page 132)
- [Submit a Job Using the d2drestorevm Utility](#) (see page 132)
- [Verify the VM is Recovered](#) (see page 133)

Review the Prerequisites and Considerations

Review the following prerequisites before you restore the VM:

- The following versions of hypervisors are supported for BMR using the d2drestorevm utility:
 - XenServer 6.0 and later
 - RHEV 3.0 and later
 - OVM 3.2
- The VM restore option can be performed from the command line only. This option is not available on the user interface.
- You can use the user interface to monitor the job status and activity logs. You can use the user interface to pause, delete, and rerun the restore VM job. However, you cannot modify the restore VM job.
- Before you restore a VM, you have to manually set up the VM on Xen, Oracle Virtual Machine (OVM), or Red hat Enterprise Virtualization (RHEV).
- When restoring to Xen and OVM virtual machines, the NFS server is required to be installed and running on the Backup Server. Verify that the firewall is not blocking the NFS service and the hypervisor has proper access and permission to use the NFS service on the Backup Server.

- When you restore to RHEV, the Perl interpreter version 5.10.0 and later, and the following modules must be installed on the Backup Server:

XML::Simple

URI::Escape

WWW::Curl

File::Copy

File::Temp

The modules can be installed with the system package manager. You can also use the Perl package manager, CPAN, to install the modules.

Example: Install "XML::Simple" using CPAN

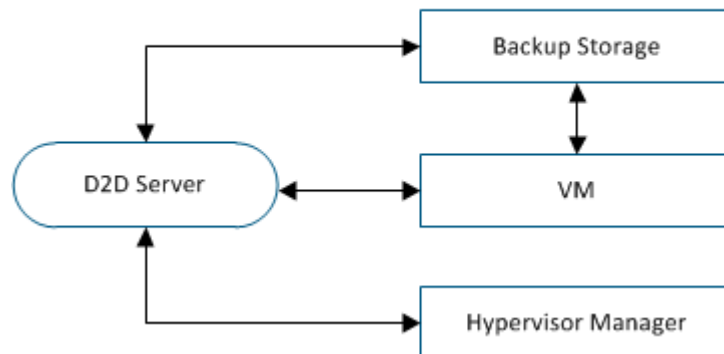
```
# perl -MCPAN -e "install XML::Simple"
```

Run the following command to verify the installation:

```
# perldoc -l "XML::Simple"
```

- To perform a successful VM restore, both the hypervisor and the target VM must have a valid network connection with the Backup Server. The following diagram displays the network requirement:

Network Requirements



The Backup Server will attempt to automatically discover and set up a virtual NIC for the VM. However, sometimes a valid network may not be selected for the NIC. The `vm_network` parameter allows you to specify certain network that the NIC should be connected with. The following considerations are for different virtual platforms:

- On XenServer, after an installation the default network is displayed as Network 0 in XenCenter, which is not the actual network. Any network with name "Pool-wide network associated with xxx" is displayed as "Network 0" on XenCenter. In such cases, rename the default network and use the new value for the `vm_network` parameter.
- On RHEV, when the `vm_network` parameter is not specified, the default rhevm network will have a higher priority.
- On OVM, it is recommended to manually set the `vm_network` parameter when there is more than one network available.
- When using the CIFS share as a backup (session) location, consider the following points:
 - Use the character / instead of \.
 - The `storage_username` and `storage_password` parameters are required to verify the credentials for CIFS shares.
- At least one of the following parameters must be specified for the `d2drestorevm` to work:
`vm_name`
`vm_uuid`
If both parameters are provided, then these parameters must belong to the same virtual machine. If the parameters belong to different virtual machines, you will get an error.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Review the following considerations before you restore the VM:

- It is recommended that you restore the sessions from the previous release of Arcserve UDP Agent (Linux) or Arcserve Unified Data Protection Agent for Linux to the original VMs.
- When you restore a VM in a XenServer PV and the restored VM displays a blank screen but the SSH and other services are active, then verify the `'console=kernel` parameter is set correctly in the boot arguments.
- PV sessions can only be restored to a PV target VM on XenServer and OVM.
- Non-PV sessions can be restored to Non-PV target VMs, such as XenServer HVM, OVM HVM, and RHEV.
- HVM of RHEL 6 series and its derivatives (RHEL 6, CentOS 6, and Oracle Linux6) can be restored to PV VM.

Create a Configuration Template

Create a configuration file so that the `d2drestorevm` command can restore VMs based on the parameters specified in the file. The `d2drestorevm` file gathers all the specifications from the file and performs the restore based on the specifications.

Syntax

```
d2drestorevm --createtemplate=[save path]
```

The `d2dutil --encrypt` utility encrypts the password and provides an encrypted password. You must use this utility to encrypt all your passwords. If you use the `--pwdfile=pwdfilepath` parameter, then you must encrypt the password. You can use the utility one of the following methods:

Method 1

```
echo "string" | d2dutil --encrypt
```

string is the password that you specify.

Method 2

Type the "`d2dutil --encrypt`" command and then specify your password. Press Enter and you will see the result on your screen. In this method, the password that you enter is not echoed on the screen.

Follow these steps:

1. Log in to the Backup Server as a root user.
2. Create the configuration template using the following command:

```
d2drestorevm --createtemplate=[save path]
```

[save path] indicates the location where the configuration template is created.

3. Open the configuration template and update the following parameters in the configuration template:

job_name

Specifies the name of the restore job.

vm_type

Specifies the type of the hypervisor where you restore the VM. The valid types of hypervisors are xen, ovm, and rhev.

vm_server

Specifies the address of the hypervisor server. The address could be either the hostname or IP address.

vm_svr_username

Specifies the username of the hypervisor. The username must be in the following format for RHEV:

[username]@[domain]

The username must be in the following format for OVM and Xen:

[username]

vm_svr_password

Specifies the password of the hypervisor. The password is encrypted using the d2dutil encryption utility.

vm_name

Specifies the name of the target VM that is displayed in the hypervisor.

Important! The `vm_name` parameter must not contain any special characters except blank spaces and should only include the following characters: a-z, A-Z, 0-9, - and _.

vm_uuid

Specifies the uuid of the target VM.

vm_network

(Optional) Specifies the network name that you want to use. If you do not provide the network name, then the default network is auto-selected.

storage_location

Specifies the storage server location of the session. The storage location can be CIFS or NFS.

storage_username

Specifies the username when you use the CIFS as the storage location.

storage_password

Specifies the password when you use the CIFS as the storage location. The password is encrypted using the d2dutil encryption utility.

encryption_password

Specifies the session encryption password. The password is encrypted using the d2dutil encryption utility.

source_node

Specifies the node name of the source whose recovery point is used to restore.

recovery_point

Specifies the session that you want to restore. Typically, a recovery session is in the following format: S00000000X, where X is a numeric value. If you want to restore the most recent session, specify the keyword 'last'.

guest_hostname

Specifies the host name that you want to provide after you restore the VM.

guest_network

Specifies the network type that you want to configure. The network could either dhcp or static.

guest_ip

Specifies the IP address when you specify the Static IP.

guest_netmask

Specifies the network mask when you specify the static IP.

guest_gateway

Specify the gateway address when you specify the static IP.

guest_reboot

(Optional) Specifies whether the target VM should be restarted after the VM is restored. The values are yes and no.

Default: no

guest_reset_username

(Optional) Specifies to reset the password to the value you provide in the `guest_reset_password` parameter.

guest_reset_password

(Optional) Specifies to reset the password to the specified value. The password is encrypted using the `d2dutil` encryption utility.

force

Specifies whether to force restore the VM. The values are `yes` and `no`.

Default: `no`

4. Save and close the configuration template.

The configuration template is successfully created.

(Optional) Create a Global Configuration File

The global configuration file (`vm.cfg`) has parameters and values related to storage locations where the VM virtual disks are created. The values for storage locations are auto-detected during the restore process. The `vm.cfg` file overrides the values related to storage locations and other parameters. If you want to specify your own storage location instead of the auto-detected value, you can use the `vm.cfg` file.

The global configuration file is at the following location:

```
/opt/CA/d2dserver/configfiles/vm.cfg
```

The following parameters can be configured in the `vm.cfg` file:

General Parameters

D2D_VM_PORT

Allows you to specify a custom port to communicate with the hypervisor server.

- For OVM, the `d2drestorevm` command requires the OVM CLI interface and the default port is 10000.
- For XenServer, the `d2drestorevm` command communicates with the server using the SSH and the default port is 22.

- For RHEV, the `d2drestorevm` command utilizes the Representational State Transfer (REST) API to communicate with the server using HTTPS.

OVM Specific Parameters

OVM_ISO_REPOSITORY

Lets you manually set the repository to upload the Arcserve UDP Agent (Linux) Live CD.

OVM_ISO_UPLOAD_SERVER

Lets you manually specify the repository server to upload the Arcserve UDP Agent (Linux) Live CD.

OVM_DISK_REPOSITORY

Lets you use specific OVM repository to create virtual disks.

Note: The `d2drestorevm` utility uses the ID for the OVM specific parameters.

RHEV Specific Parameters

RHEV_DISK_STORAGE_DOMAIN

Lets you use specific RHEV storage domain to create virtual disks. The `d2drestorevm` utility uses the lexical file name for the RHEV specific parameters.

Xen Specific Parameters

XEN_DISK_SR

Lets you use specific Xen storage repository to create virtual disks. The `d2drestorevm` utility uses the lexical file name for the Xen specific parameters.

Follow these steps:

1. Log in to the Backup Server.
2. Create the global configuration file and name it as `vm.cfg`.
3. Open the global configuration file and update the parameters in the file.
4. Save and close the file.
5. Place the file at the `configfiles` folder:

```
/opt/CA/d2dserver/configfiles/vm.cfg
```

The global configuration file is successfully created.

Modify the Configuration Template and File

If you already have the configuration template and the global configuration file, you can modify the files and restore another VM. You do not have to create another configuration templates and files each time you restore a VM. When you submit the job, a new job is added on the web UI. You can see the activity logs on the web UI.

Follow these steps:

1. Log in to the Backup Server as a root user.
2. Open the configuration template from the location where you have saved the file and modify the parameters per your requirement.
3. Save and close the configuration template.
4. (Optional) Open the global configuration file from the following location and modify the parameters per your requirement:

```
/opt/CA/d2dserver/configfiles/vm.cfg
```

5. Save and close the global configuration file.

The configuration template and file are successfully modified.

Submit a Job Using the d2drestorevm Utility

Run the d2drestorevm command to restore the VM. The command verifies the target VM and submits a restore job. The restore job can be seen from the web UI. During the restore process if any requirement is not met, you will get an error. You can view the activity log on the web UI.

Follow these steps:

1. Log in to the Backup Server as a root user.
2. Submit the restore job for the VM using the following command:

```
d2drestorevm --template=cfg_file_path [--wait]
```

Note: The --wait switch lets you return to the shell environment after the restore job is complete. If the --wait switch is not present, you return to the shell environment immediately after submitting the job.

The restore job is submitted.

Verify the VM is Recovered

After the completion of restore job, verify that the target node is restored with relevant data.

Follow these steps:

1. Navigate to the VM that you restored.
2. Verify that the VM has all the information that you backed up.

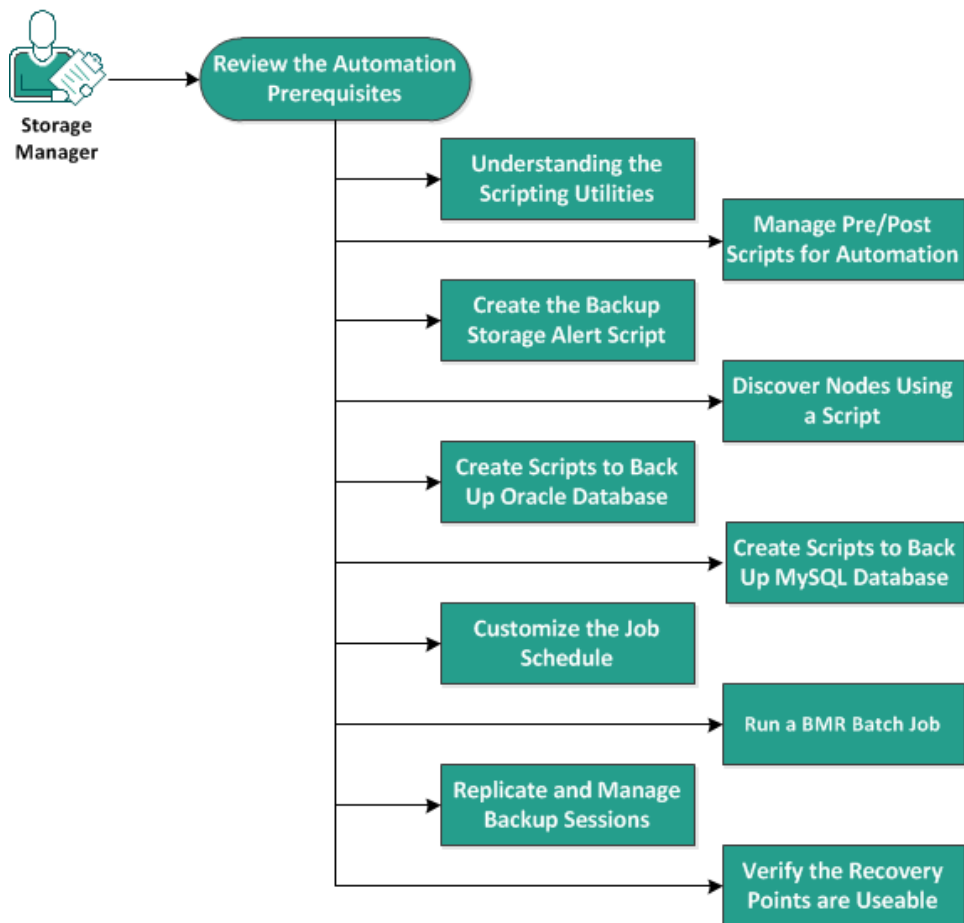
The VM is successfully verified.

How to Integrate and Automate Arcserve Unified Data Protection Agent for Linux with the Existing IT Environment

As a Storage Manager, you can create scripts and can automate tasks to integrate Arcserve UDP Agent (Linux) with your existing IT environment. Scripts reduce the manual intervention and decrease the dependency on the web interface of the Backup Server to perform any task. Arcserve UDP Agent (Linux) also provides the interface and utilities to perform the job management, node management, and activity log management tasks.

The following diagram displays the process to integrate and automate Arcserve UDP Agent (Linux) with the existing IT environment:

How to Integrate and Automate Arcserve Unified Data Protection Agent for Linux with the Existing IT Environment



Perform the following tasks to automate and manage Arcserve UDP Agent (Linux):

- [Review the Automation Prerequisites](#) (see page 135)
- [Understanding the Scripting Utilities](#) (see page 136)
- [Manage Pre/Post Scripts for Automation](#) (see page 145)
- [Create the Backup Storage Alert Script](#) (see page 149)
- [Discover Nodes Using a Script](#) (see page 149)
- [Create the Scripts to Back Up Oracle Database](#) (see page 150)
- [Create the Scripts to Back Up MySQL Database](#) (see page 152)
- [Customize the Job Schedule](#) (see page 154)
- [Run a BMR Batch Job](#) (see page 155)
- [Replicate and Manage Backup Sessions](#) (see page 157)
- [Verify the Recovery Points are Usable](#) (see page 159)

Review the Automation Prerequisites

Consider the following prerequisites before you automate and manage Arcserve UDP Agent (Linux):

- You have the root login credentials to the Backup Server.
- You have knowledge of Linux scripting.
- You have a better understanding of the Arcserve UDP Agent (Linux) web interface.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Understanding the Scripting Utilities

Arcserve UDP Agent (Linux) provides scripting utilities to help you create your automation script. These utilities are merely for scripting so their output is scripting friendly. The utilities are used to manage nodes, jobs, replicate backup destinations, and manage activity logs.

All the utilities are contained in the *bin* folder at the following location:

```
/opt/CA/d2dserver/bin
```

The `d2dutil --encrypt` utility encrypts the password and provides an encrypted password. You must use this utility to encrypt all your passwords. If you use the `--pwdfile=pwdfilepath` parameter, then you must encrypt the password. You can use the utility one of the following methods:

Method 1

```
echo "string" | d2dutil --encrypt
```

string is the password that you specify.

Method 2

Type the "`d2dutil --encrypt`" command and then specify your password. Press Enter and you will see the result on your screen. In this method, the password that you enter is not echoed on the screen.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Navigate to the *bin* folder using the following command:

```
# cd/opt/CA/d2dserver/bin
```


3. Run the following commands to manage nodes:

```
# ./d2dnode
```

Displays a list of available commands to help you manage all related Linux nodes. Using this command, you can add, delete, modify and import nodes. You can also add nodes using the non-root credentials.

Note: You can use all the parameters of the `d2dnode` command, when the Backup Server is a standalone Linux agent. When the Backup Server is managed by UDP Console, the `d2dnode` command lets you perform only the list, add, modify, and import parameters. The list, add, modify, or import parameters will update the node on the UDP Console. For example, the `./d2dnode --list` command will list all the Linux nodes that are added to the UDP Console.

```
# ./d2dnode --list
```

Lists all the nodes that are managed by the Backup Server.

```
# ./d2dnode --add=nodename/ip --user=username --password=password --description="the description of that node" --attach=jobname --force
```

Adds the specific node to the Backup Server. If you are a root user, use this command to add nodes.

Note: If you change the port number of the node, then you must specify the new port number in the `--add` parameter as shown in the following example.

Example: `./d2dnode --add=nodename/ip:new_port --user=username --password=password --description="the description of that node" --attach=jobname --force`

--attach=jobname

Adds a new node to an existing backup job.

--force

Adds the node forcefully even if the node is managed by another Backup Server. If you remove the *force* parameter, then the node is not added to this server if it is managed by another Backup Server.

```
# ./d2dnode --add=nodename --user=username --password=password --rootuser=rootaccount  
--rootpwd=rootpassword --pwdfile=pwdfilepath --description=description --attach=jobname --force
```

Adds the specific node to the Backup Server. If you are a non-root user, use this command to add nodes.

Note: If you change the port number of the node, then you must specify the new port number in the --add parameter as shown in the following example.

Example: # ./d2dnode --add=nodename/ip:new_port --user=username
--password=password --rootuser=rootaccount --rootpwd=rootpassword
--pwdfile=pwdfilepath --description=description --attach=jobname --force

--user=username

Specifies the username of the non-root user.

--password=password

Specifies the password of the non-root user. If the --pwdfile=pwdfilepath parameter is provided, then you do not have to specify this parameter.

--rootuser=rootaccount

Specifies the username of the root user.

--rootpwd=rootpassword

Specifies the password of the root user. If the --pwdfile=pwdfilepath parameter is provided, then you do not have to specify this parameter.

--pwdfile=pwdfilepath

(Optional) Specifies the password of the root user and non-root user. This is an optional parameter that you use if you have stored the passwords of the root user and non-root users in a separate file. The password file includes the following parameters: --password=password and --rootpwd=rootpassword. For added security, the password must be encrypted using the d2dutil --encrypt utility. After you encrypt the password, replace the old password with the encrypted password in the --pwdfile parameter.

```
# ./d2dnode --node=nodename --attach=jobname
```

Adds the specified node to an existing backup job.

```
# ./d2dnode --modify=nodename/ip --user=username --password=newpassword  
--description=newdescription
```

Modifies the username, password, or the description of the added node. If you are a root user, use this command to modify nodes.

```
# ./d2dnode --modify=nodename --user=username --password=newpassword --rootuser=rootaccount  
--rootpwd=newrootpassword --pwdfile=pwdfilepath --description=newdescription
```

Modifies the username, password, or the description of the added node. If you are a non-root user, use this command to modify nodes.

--user=username

Specifies the username of the non-root user.

--password=newpassword

Specifies the new password of the non-root user.

--rootuser=rootaccount

Specifies the username of the root user.

--rootpwd=newrootpassword

Specifies the new password of the root user.

--pwdfile=pwdfilepath

(Optional) Specifies the password of the root user and non-root user. This is an optional parameter that you use if you have stored the passwords of the root user and non-root users in a separate file. The password file includes the following parameters: --password=newpassword and --rootpwd=newrootpassword.

```
# ./d2dnode --delete=nodename1,nodename2,nodename3
```

Deletes the specified nodes from the Backup Server. To delete multiple nodes, use a comma (,) as a delimiter.

```
# ./d2dnode --import=network --help
```

Imports nodes from the network. When you import nodes, you get the following options to configure:

--netlist

Specifies the IP v4 IP address list. For multiple entries, the list should be a comma separated entries.

Example

192.168.1.100 : Imports the node that has the IP address 192.168.1.100

192.168.1.100-150 : Import all the nodes that belong to the scope (range) between 192.168.1.100 and 192.168.1.150

192.168.1.100- : Imports all the nodes that belong to the scope (range) between 192.168.1.100 and 192.168.1.254. Here you not have to mention the end range.

192.168.1.100-150,192.168.100.200-250 : Imports multiple nodes that belong to two different scopes. The first scope (range) between 192.168.1.100 and 192.168.1.150, and the second scope between 192.168.100.200 and 192.168.100.250. Each entry is separated by a comma.

--joblist

Specifies the job name list. A job name must not include a comma. After a node is successfully imported, the node is added to the job. For multiple jobs, the list should be a comma separated entries.

Example: --joblist=jobA,jobB,jobC

In this example, each job entry is separated by a comma.

Note: This option is only supported by the Arcserve UDP Agent (Linux) standalone version.

--user

Specifies the user name to import and add the nodes.

--password

Specifies the password to import and add nodes.

--rootuser

Specifies the user name of the root user. If a non-root user is added, then use this parameter to specify the root user credential.

--rootpwd

Specifies the password of the root user. If a non-root user is added, then use this parameter to specify the root user credential.

--pwdfile

(Optional) Specifies the password of the root user and non-root user. This is an optional parameter that you use if you have stored the passwords of the root user and non-root users in a separate file. The password file includes the following parameters: --password=newpassword and --rootpwd=newrootpassword.

--prefix

Specifies the prefix given to a host name. Use this parameter to filter nodes that includes the prefix in the host name.

--blacklistfile

Specifies a file that includes a list of node hostname that you do not want to add to the Backup Server. You must provide one node per line in the file.

--force

Adds the node forcefully even if the node is managed by another Backup Server. If you remove the *force* parameter, then the node is not added to this server if it is managed by another Backup Server.

--verbose

Displays more information about the node import process. Use this parameter for debugging or automation scripting purpose.

--help

Displays the help screen.

Notes:

- The import function uses the SSH server to detect whether a node is a Linux node. If your SSH server uses non-default port, then configure the server to use the non-default port. For more information on configuring the SSH port number, see [Change the SSH Port Number of the Backup Server](#).
- When the password is not provided, SSH key authentication method is used.

4. Run the following commands to manage jobs:

```
# ./d2djob
```

Displays a list of commands to help you manage jobs. Using this command, you can run, cancel, and delete jobs.

```
# ./d2djob --delete=jobname
```

Deletes the specified job from the Job Status tab.

```
# ./d2djob --run=jobname --jobtype=1 --wait
```

Runs the specified job. The `--jobtype` parameter is optional. The `d2djob` command automatically identifies the job type from the job name that you specify. If the command identifies a restore job, the restore job starts. If the command identifies a backup job and you do not provide any value for the `--jobtype` parameter, then an incremental backup job starts. The Incremental backup is the default job type.

If you want to specify the job type for a backup job, then the values are 0, 1, and 2, where 0 indicates a Full backup job, 1 indicates an Incremental backup job, and 2 indicates a Verify backup job.

```
# ./d2djob --cancel=jobname --wait
```

Cancels a job that is in progress.

If you include `--wait` in the command, the job status is displayed after the job is canceled. If you do not include `--wait` in the command, the job status is displayed immediately after submitting the cancellation request.

```
# ./d2djob --newrestore=restoreJobName --target=macaddress/ipaddress --hostname=hostname  
--network=dhcp/staticip --staticip=ipaddress --subnet=subnetMask --gateway=gateway --runnow --wait
```

Runs a restore job for a new target machine based on an existing restore job. This command lets you use the same restore settings as the existing restore job and only the target machine details are different. If you use this command, you do not have to create multiple restore jobs for different target machines.

You must provide a value for `--newrestore`, `--target`, `--hostname`, and `--network`.

If the value for `--network` is `staticip`, then you must provide a value for `--staticip`, `--subnet`, and `--gateway`. If the value for `--network` is `dhcp`, then you do not have to provide any value for `--staticip`, `--subnet`, and `--gateway`.

If you include `--runnow` in the command, the job runs immediately after you submit the job, irrespective of the job schedule.

If you include the `--wait` parameter in the command, the status message is displayed after the completion of the job. If you do not include `--wait` in the command, the status message is displayed immediately after submitting the job.

```
# ./d2djob <--export=jobname1,jobname2,jobname3> <--file=filepath>
```

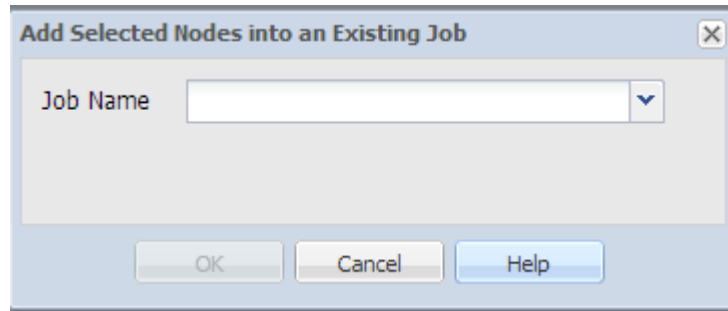
Exports multiple jobs from the Backup Server to a file. If you want similar backup configurations in multiple Backup Servers, you can export the backup jobs to a file, and then import the file to other Backup Servers.

Note: If the Linux Backup Server is managed by Arcserve UDP console, the export function is not supported.

```
# ./d2djob <--import=filepath>
```

Imports the file containing the backup job information to a Backup Server. You can also import the file to Arcserve UDP, if the Backup Server is managed by Arcserve UDP.

If the backup job is imported to a Backup Server, then you can select the job from the following dialog:



You can also use the following command line utility to add nodes to this job:

```
./d2dnode --attach=jobname
```

5. Run the following commands to create or update the recovery points configuration file. Arcserve UDP Agent (Linux) uses the configuration file to manage and display the recovery points in the UI.

```
# ./d2drp
```

Creates or updates the recovery points configuration files based on the recovery points detail. Using this command, you can create or update the configuration files.

```
# ./d2drp --build --storagepath=/backupdestination --node=node_name
```

Verifies all recovery points that belong to *node_name* and update all the recovery points configuration files. If the recovery point configuration files are not present, this command create the files automatically. The `--build` parameter creates the configuration files of recovery points.

```
# ./d2drp --build --storagepath=/backupdestination --node=node_name --rp=recovery_point
```

Verifies the specified session name and update all the recovery points configuration files. If the recovery point configuration files are not present, this command create the files automatically. Specify the keyword 'last' for the `--rp` parameter to get the most recent recovery point.

```
# ./d2drp --show --storagepath=path --node=nodeName --rp=recovery_point --user=username  
--password=password
```

Displays system information for the specified recovery point.

--rp=recovery_point

Specifies the recovery point that you want to access. Specify the keyword 'last' to get the most recent recovery point.

--user=username

Specifies the username to access the storage location or backup destination.

--password=password

Specifies the password to access the storage location or backup destination.

Note: For the `--build` parameter, d2drp does not support the NFS share or the CIFS share. If you want to use the NFS share or the CIFS share, you must first mount the share to the local host and then use the mount point as the storagepath.

6. Run the following command to register Backup Server to Arcserve UDP. When you register the Backup Server with Arcserve UDP, you can manage the Backup Server from Arcserve UDP. You can also import nodes and jobs that were previously managed by Backup Server to Arcserve UDP.

```
# ./d2dreg <--reg=servername> <--user=username> <--port=port> <--protocol=http/https>  
[--password=password]
```

Registers the Backup Server to Arcserve UDP so that the Backup Server can be managed from Arcserve UDP Console.

Note: The d2dreg command uses the host name of the Backup Server to identify the server. If Arcserve UDP Console cannot connect to the Backup Server using the host name, then change the host name to IP Address in the **Update Node** dialog.

7. Run the following commands to manage activity logs:

```
# ./d2dlog
```

Displays the format that helps you get the activity logs for the specified job id in the specified format.

```
# ./d2dlog --show=jobid --format=text/html
```

Displays the activity log of the specified job. The format value is optional because the default value is text.

8. Run the following commands to manage the job history:

```
# ./d2djobhistory
```

Displays the job history based on the filters you specify. You can filter the job history by days, weeks, months, and start and end date.


```
# /d2djobhistory --day=n --headers=column_name1,column_name2,...column_name_n  
--width=width_value --format=column/csv/html
```

Displays the recent job history based on the specified days.

--headers=column_name1,column_name2,...column_name_n

(Optional) Specifies the columns that you want to view in the job history. This is an optional parameter. The predefined columns are ServerName, TargetName, JobName, JobID, JobType, DestinationLocation, EncryptionAlgoName, CompressLevel, ExecuteTime, FinishTime, Throughput, WriteThroughput, WriteData, ProcessedData, and Status.

--width=width_value

(Optional) Specifies the number of characters that you want to display for each column. This is an optional parameter. Each column has its own default width. You can update the width value for each column, where each width value is separated by a comma (,).

--format=column/csv/html

Specifies the display format of the job history. The available formats are column, csv, and html. You can specify only one format at a time.

```
# /d2djobhistory --week=n --headers=column_name1,column_name2,...column_name_n  
--width=width_value --format=column/csv/html
```

Displays the recent job history based on the specified weeks.

```
# /d2djobhistory --month=n --headers=column_name1,column_name2,...column_name_n  
--width=width_value --format=column/csv/html
```

Displays the recent job history based on the specified months.

```
# /d2djobhistory --starttime=yyyymmdd --endtime=yyyymmdd  
--headers=column_name1,column_name2,...column_name_n --width=width_value  
--format=column/csv/html
```

Displays the recent job history based on the specified start and end date.

The scripting utilities have been used to successfully manage nodes, jobs, and activity logs.

Manage Pre/Post Scripts for Automation

Pre/Post scripts let you run your own business logic at specific stages of a running job. You can specify when to run your scripts in **Pre/Post Script Settings** of the **Backup Wizard** and the **Restore Wizard** in the Console. The scripts can be run on the Backup Server depending on your setting.

Managing the pre/post script is a two part process, consisting of creating the pre/post script and placing the script in the prepost folder.

Create Pre/Post Scripts

Follow these steps:

1. Log into the Backup Server as a root user.
2. Create a script file using the environment variables in your preferred scripting language.

Pre/Post Script Environment Variables

To create your script, use the following environment variables:

D2D_JOBNAME

Identifies the name of the job.

D2D_JOBID

Identifies the job ID. Job ID is a number provided to the job when you run the job. If you run the same job again, you get a new job number.

D2D_TARGETNODE

Identifies the node that is being backed up or restored.

D2D_JOBTYPE

Identifies the type of the running job. The following values identify the D2D_JOBTYPE variable:

backup.full

Identifies the job as a full backup.

backup.incremental

Identifies the job as an incremental backup.

backup.verify

Identifies the job as a verify backup.

restore.bmr

Identifies the job as a bare-metal recovery (bmr). This is a restore job.

restore.file

Identifies the job as a file-level restore. This is a restore job.

D2D_SESSIONLOCATION

Identifies the location where the recovery points are stored.

D2D_PREPOST_OUTPUT

Identifies a temp file. The content of the first line of the temp file is displayed in the activity log.

D2D_JOBSTAGE

Identifies the stage of the job. The following values identify the D2D_JOBSTAGE variable:

pre-job-server

Identifies the script that runs on the Backup Server before the job starts.

post-job-server

Identifies the script that runs on the Backup Server after the job completes.

pre-job-target

Identifies the script that runs on the target machine before the job starts.

post-job-target

Identifies the script that runs on the target machine after the job completes.

pre-snapshot

Identifies the script that runs on the target machine before capturing the snapshot.

post-snapshot

Identifies the script that runs on the target machine after capturing the snapshot.

D2D_TARGETVOLUME

Identifies the volume that is backed up during a backup job. This variable is applicable for pre/post snapshot scripts for a backup job.

D2D_JOBRESULT

Identifies the result for a post job script. The following values identify the D2D_JOBRESULT variable:

success

Identifies the result as successful.

fail

Identifies the result as unsuccessful.

D2DSVR_HOME

Identifies the folder where Backup Server is installed. This variable is applicable for the scripts that run on the Backup Server.

The script is created.

Note: For all scripts, a return value of zero indicates success and a nonzero return value indicates failure.

Place the Script in the Prepost Folder and Verify

All the pre/post scripts for a Backup Server are centrally managed from the prepost folder at the following location:

`/opt/CA/d2dserver/usr/prepost`

Follow these steps:

1. Place the file in the following location of the Backup Server:

`/opt/CA/d2dserver/usr/prepost`

2. Provide the execution permission to the script file.
3. Log into the Arcserve UDP Agent (Linux) web interface.
4. Open the **Backup Wizard** or the **Restore Wizard** and navigate to the **Advanced** tab.
5. Select the script file in the **Pre/Post Script Settings** drop-down list and then submit the job.
6. Click **Activity Log** and verify that the script is executed to the specified backup job.

The script is executed.

The pre/post scripts are successfully created and placed in the prepost folder.

Create the Backup Storage Alert Script

Create the backup storage alert script so that you can run the script when your backup storage space is less than the specified value. When you add a backup storage location in the UI, you have the option to select the Send alert checkbox. When you select the checkbox, Arcserve UDP Agent (Linux) monitors the available storage space every 15 minutes. Every time the storage space is less than the specified value, Arcserve UDP Agent (Linux) runs the *backup_storage_alert.sh* script. You can configure the *backup_storage_alert.sh* script to perform any task for you when the backup storage space is less.

Example 1: You can configure the script to automatically send you an email alert to remind you of the decreasing storage space.

Example 2: You can configure the script to automatically delete some data from the backup storage space when the storage space is less than the specified value.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Create the *backup_storage_alert.sh* script using the following variables:

backupstoragename

Defines the name of backup storage location. For example, NFS or CIFS.

freesize

Defines the available free space in the backup storage location.

3. Place the script at the following location:

`/opt/CA/d2dserver/usr/alert/backup_storage_alert.sh`

The *backup_storage_alert.sh* script is created.

Discover Nodes Using a Script

Arcserve UDP Agent (Linux) provides the capability to run a script that discovers nodes in your network. You can write a script to discover nodes in your network and then place the script in the *discovery* folder.

You can configure the node discovery setting on the web interface and set the frequency of running the script. In the script, you can specify the utilities to discover nodes in your network. After the script discovers a node, use the *d2dnode* command to add that node to Arcserve UDP Agent (Linux). There is an activity log for every time the script is run.

Note: For all scripts, a return value of zero indicates success and a nonzero return value indicates failure.

If you want to print something into the Activity Log regarding your node discovery script, you can use the following special environment variable:

```
echo "print something into activity log" > "$D2D_DISCOVER_OUTPUT"
```

A sample script is placed in the *discovery* folder at the following location that can discover the Linux nodes in a sub network.

```
/opt/CA/d2dserver/examples/discovery
```

You can copy the sample script to the following location and modify that script per your requirement:

```
/opt/CA/d2dserver/usr/discovery
```

Follow these steps:

1. Log into the Backup Server as a root user.
2. Create a node discovery script and place the script in the *discovery* folder at the following location:

```
/opt/CA/d2dserver/usr/discovery
```

3. Provide the necessary execution permission to the script file.
4. Log into the web interface.
5. Configure the node discovery settings in the Node menu to run your script.
6. Click Activity Log and verify that the script is executed.

The Activity Log displays a list of all discovered nodes.

Nodes are successfully discovered using the script.

Create the Scripts to Back Up Oracle Database

You can create scripts that you use to back up your Oracle Database. You do not have to stop your database to perform a backup. Verify that the database is in the archive log mode. If it is not in the archive log mode, then change the database to the archive log mode before you back up the database. You create the following two scripts to back up Oracle Database:

- **pre-db-backup-mode.sh** - This script prepares and keeps the entire database in the backup mode.
- **post-db-backup-mode.sh** - This script removes the database from the backup mode.

You can specify the scripts to run on the Oracle Database nodes in Pre/Post Scripts Settings of the Backup Wizard.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Create the *pre-db-backup-mode.sh* script using to the following code:

```
#!/bin/bash
orauser="oracle"
orasid="orcl"
su - ${orauser} << BOF 2>&1
export ORACLE_SID=${orasid}
sqlplus /nolog << EOF 2>&1
connect / as sysdba
alter database begin backup;
exit;
EOF
BOF
```

Note: Specify the value for *orauser* and *orasid* variables as defined in your Oracle Database.

3. Create the *post-db-backup-mode.sh* script using the following code:

```
#!/bin/bash
orauser="oracle"
orasid="orcl"
su - ${orauser} << BOF 2>&1
export ORACLE_SID=${orasid}
sqlplus /nolog << EOF 2>&1
connect / as sysdba
alter database end backup;
exit;
EOF
BOF
```

Note: Specify the value for *orauser* and *orasid* variables as defined in your Oracle Database.

4. Provide the execution permission to both the scripts.
5. Place both scripts at the following location:

```
/opt/CA/d2dserver/usr/prepost/
```

6. Log into the Arcserve UDP Agent (Linux) web interface.
7. Open the Backup Wizard and navigate to the Advanced tab.
8. In the Pre/Post Scripts Settings option, select the *pre-db-backup-mode.sh* script file from the "Before snapshot is taken" dropdown list.
9. In the Pre/Post Scripts Settings option, select the *post-db-backup-mode.sh* script file from the "After snapshot is taken" dropdown list.
10. Submit the backup job.

The backup job is submitted.

The scripts are created to back up Oracle Database.

Note: Arcserve UDP Agent (Linux) supports the volume level snapshot. To ensure the data consistency, all data files of the database must be on one volume.

To restore the Oracle database, see [How to Restore an Oracle Database Using Arcserve UDP Agent \(Linux\)](#) (see page 193).

Create the Scripts to Back Up MySQL Database

You can create scripts that you use to back up MySQL Database. You do not have to stop your database to perform a backup. You create the following two scripts to back up MySQL Database:

- **pre-db-backup-mode.sh** - This script closes all open tables, and it locks all the tables for all the databases with a global read lock.
- **post-db-backup-mode.sh** - This script releases all the locks.

You can specify the scripts to run on the MySQL Database nodes in Pre/Post Scripts Settings of the Backup Wizard.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Create the *pre-db-backup-mode.sh* script using to the following code:

```
#!/bin/bash
dbuser=root
dbpwd=rootpwd
lock_mysqldb(){
(
echo "flush tables with read lock;"
sleep 5
) | mysql -u$dbuser -p$dbpwd ${ARGUMENTS}
}
lock_mysqldb &
PID="/tmp/mysql-plock.$!"
touch ${PID}
```

Note: Specify the value for *dbuser* and *dbpwd* variables as defined in your MySQL Database.

3. Create the *post-db-backup-mode.sh* script using the following code:

```
#!/bin/bash
killcids(){
pid="$1"
cids=`ps -ef|grep ${pid}|awk '{if("$pid"==$3){print $2}}'`
for cid in ${cids}
do
echo ${cid}
kill -TERM ${cid}
done
echo -e "\n"
}
mysql_lock_pid=`ls /tmp/mysql-plock.* | awk -F . '{print $2}'`
[ "$mysql_lock_pid" != "" ] && killcids ${mysql_lock_pid}
rm -fr /tmp/mysql-plock.*
```

4. Provide the execution permission to both the scripts.
5. Place both the scripts in the following location:

```
/opt/CA/d2dserver/usr/prepost/
```
6. Log into the Arcserve UDP Agent (Linux) web interface.
7. Open the Backup Wizard and navigate to the Advanced tab.
8. In the Pre/Post Scripts Settings option, select the *pre-db-backup-mode.sh* script file from the "Before snapshot is taken" dropdown list.
9. In the Pre/Post Scripts Settings option, select the *post-db-backup-mode.sh* script file from the "After snapshot is taken" dropdown list.
10. Submit the backup job.
The backup job is submitted.

The scripts are created to back up MySQL Database.

Note: Arcserve UDP Agent (Linux) supports the volume level snapshot. To ensure the data consistency, all data files of the database must be on one volume.

Customize the Job Schedule

Arcserve UDP Agent (Linux) provides the capability to define your own schedule using a script to run a job. If you require to run a job periodically and you cannot schedule using the web UI, you can create a script to define such schedule. For example, you want to run a backup at 10:00 PM on the last Saturday of every month. You cannot define such schedule using the web interface, but you can create a script to define such schedule.

You can submit a backup job without specifying any schedule (using the None option on the Advanced page). Use the Linux Cron scheduler to define your customized schedule and run the *d2djob* command to run the job.

Note: The following procedure assumes that you have submitted a backup job without specifying any schedule and you want to run a backup at 10:00 PM on the last Saturday of every month.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Create a script file and enter the following command to run a backup at 10:00 PM on the last Saturday of every month:

```
#!/bin/bash
LAST_SAT=$(cal | awk '$7!=""{t=$7} END {print t}')
TODAY=$(date +%d)
if [ "$LAST_SAT" = "$TODAY" ]; then
    source /opt/CA/d2dserver/bin/setenv
    d2djob --run=your_job_name --jobtype=your_job_type #run your backup job here
fi
```

Note: You must provide the necessary execution permission to the file.

3. Navigate to the crontab folder and add the following command to your system crontab (/etc/crontab):

```
00 22 * * Saturday root runjob.sh
```

Cron runs the runjob.sh script at 10:00 PM every Saturday. In runjob.sh, it first checks if today is last Saturday of the month. If yes, it uses d2djob to run the backup job.

The job schedule is customized to run a backup at 10:00 PM on the last Saturday of every month.

Run a BMR Batch Job

If you want to perform a BMR on multiple machines and you want to install the same operating environment on all the machines, you can perform a batch BMR. You do not have to create a job for each BMR job. You can save time and effort, and you can reduce the risk of any error while configuring the BMR machines.

Note: You must have a valid recovery point of the source machine that you want to restore. If you do not have a valid recovery point, you must first back up the source machine and then submit a restore job.

You first define all your BMR settings in a template BMR job and then change the target machine's address (IP or MAC), hostname, and network configuration using the following command:

```
d2djob
```

Follow these steps:

1. Create a BMR job named BMR-TEMPLATE and run the job for one machine of your multiple machines.

Note: You can provide any name for the BMR job. You must provide the same job name in the batch BMR script.

2. Log in to the Backup Server as a root user.

3. Create a batch BMR script based on the BMR-TEMPLATE job to automatically submit multiple BMR jobs. Use the following script to create a batch BMR script:

```
#!/bin/sh
prename=lab-server
serverList[0]="<MAC_Address>"
serverList[1]=" <MAC_Address>"
serverList[2]=" <MAC_Address>"
.
.
.
serverList[300]=" <MAC_Address>"
for((i=0;i<${#serverList[@]};i=i+1))
do
./d2djob --newrestore="BMR-TEMPLATE" --target=${serverList[i]} --hostname=$prename$i --network=dhcp
done
```

4. Run the batch BMR script.

The script runs. Multiple BMR jobs are created in the UI.

A batch of BMR job is run.

Replicate and Manage Backup Sessions

You can create a script to replicate your backup sessions so that you can recover your data when your original backup data is corrupted. The backup sessions include all the recovery points that were backed up. You can protect your backup sessions by replicating your backup sessions to a replication destination.

After you have replicated your backup sessions, you can then manage your replication destination by adding the destination to the Arcserve UDP Agent (Linux) interface.

Replicating and managing backup sessions is a three part process. It includes the following three parts:

- Replicating the backup sessions to the replication destination
- Creating or updating the recovery points configuration files so that recovery points can be managed and displayed on the Arcserve UDP Agent (Linux) web interface
- Adding the replication destination to the Arcserve UDP Agent (Linux) web interface

Replicating the Backup Sessions

You can leverage the Pre/Post Scripts Settings feature in the Backup Wizard to replicate the backup sessions to the replication destination. You can choose any option, such as File Transfer Protocol (FTP), Secure Copy (SCP), or the cp command, to replicate the backup session.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Create a pre/post script to replicate the backup sessions.
3. Place the script at the following location:
`/opt/CA/d2dserver/usr/prepost`
4. Log into the Arcserve UDP Agent (Linux) web interface.
5. Open the Backup Wizard and navigate to the Advanced page.
6. In the Pre/Post Scripts Settings option for Run on backup server, select the replication script from the After job is over dropdown list.
7. Submit the backup job.

The backup session is replicated to the backup destination.

Create or Update the Recovery Point Configuration Files

After you replicate the backup sessions, you create and configure the recovery points configuration file. This file is used to identify the recovery points when you perform the restore operation from the Arcserve UDP Agent (Linux) interface.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Navigate to the following location:

```
/opt/CA/d2dserver/bin
```

3. Enter the following command to create or update the recovery points configuration file:

```
./d2drp --storagepath=/backupdestination --node=node_name --session=session_name
```

If you provide only the `--storagepath` and `--node` information, then the command updates all the backup sessions for the selected node. If you provide the `--session` information, then the command updates the specific session information.

Note: For more information about the `d2drp` command, see *Understanding the Scripting Utilities*.

The recovery points configuration file is created or updated depending on the status of the file.

Add the Replication Destination

Add the replication destination to the Arcserve UDP Agent (Linux) interface to manage the destination. After you add the replication destination, you can see the available free space in that destination and manage your data accordingly.

Follow these steps:

1. Log into the replication destination.
2. Create a file named Settings and enter the following code in the Settings file:

```
RecoverySetLimit=n
```

n indicates the number of recovery sets that you want to retain in the replication destination.

3. Place the file in the node folder of the replication destination.

For example, `/backup_destination/node_name/Settings`

4. Log into the Arcserve UDP Agent (Linux) web interface.
5. Add the replication destination from the Backup Storage menu.

The replication destination is added to the Arcserve UDP Agent (Linux) web interface.

The backup sessions are successfully replicated and managed.

Verify the Recovery Points are Usable

The d2dverify utility helps to verify the recovery points from various backup sessions are usable. Typically, backup jobs run every day and when you have multiple recovery points you may not be sure if the recovery points are usable for data recovery during a system failure. To avoid such situations, you can perform BMR jobs periodically to verify if the backups are useable. The d2dverify utility helps you automate the task of verifying the usability of the recovery points.

After you set up the required parameters, the d2dverify utility submits the BMR job and recovers data to the specified VM. Then d2dverify starts the VM and runs a script to verify if the applications in the VM function properly. You can also create a schedule to run the d2dverify utility periodically using system utilities such as Linux Cron. For example, you can run the d2dverify utility after the last backup of a recovery set. In such case, d2dverify verifies all recovery points in that recovery set.

Note: To know more about scheduling a job using the Linux Cron scheduler, see [Customize the Job Schedule](#).

The d2dverify utility can also be used in the following scenarios:

- You can use the d2dverify utility to migrate the backups of several physical machines to virtual machines.
- After a hypervisor is recovered, you can use the d2dverify utility to restore all the VMs to the new hypervisor.

Consider the following prerequisites before you use the d2dverify utility:

- Identify the source nodes whose backup you want to verify.
- Identify a hypervisor on which VMs will be created.
- Create VMs for each node that you want to verify. Assign the VM name in the following format:

`verify_<node name>`

Note: You do not require to attach virtual hard disks for these VMs. And you may not attach virtual network to these VMs if you specify "vm_network" parameters.

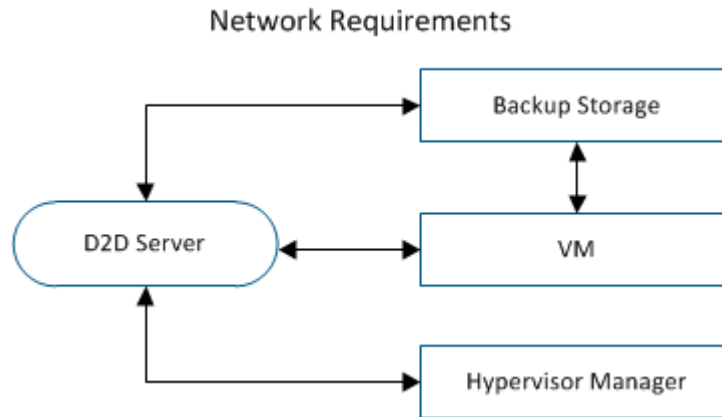
- Review the network requirements
- Identify a network in which the VMs will be connected.

Note: The d2dverify utility supports the static IP network only.

Important! If the database has the node account information related to a non-root user, then d2dverify will reset the password of the non-root user to 'CAAd2d@2013 for the target VM.

Network Requirements:

When you use d2dverify, it is recommended to keep the target VMs in an isolated virtual network to avoid any conflict with the production environment. In such cases, the target VMs must be connected to the Backup Server and the backup storage both.



Hypervisor Support:

d2dverify depends on the d2drestorevm utility to perform the restore. d2dverify supports the following versions of hypervisors:

- XenServer 6.0 and above
- RHEV 3.0 and above
- OVM 3.2

Arguments:

--template

Identifies the template that includes the parameters to run the d2dverify utility.

--createtemplate

Creates an empty template that that includes the parameters to run the d2dverify utility.

Follow these steps:

1. Log in to the Backup Server as a root user.
2. Create the template that is used by the d2dverify utility using the following command:

```
d2dverify --createtemplate=file_path
```

3. Open the template and update the following parameters:

node_list

Specifies a list of nodes or a query criteria that queries information from the database of the Backup Server. Each node is separated by a comma, such as Node1,Node2,Node3.

Notes: If the ssh port number is not the default port 22, then the format to specify each node is: Node1:new_port,Node2:new_port,Node3:new_port. The VM name is assigned as verify_<node name>, where node name does not include the port number.

Example: Node1:222,Node2:333,Node4:333

The following list is an example of query criteria:

[node=prefix]

Finds the node name that contains the defined prefix.

[desc=prefix]

Finds the node description that contains the defined prefix.

guest_ip_list =

Specifies the list of IP address that is applied to each target node respectively. Each IP address is separated with a comma, such as IP1,IP2,IP3. If there is only one IP address available but there are multiple nodes in the node_list parameter, then the fourth segment of the IP address is increased by one for each node. The d2dverify utility verifies if an IP address has been used. If yes, that IP address is skipped.

For example, if you have three nodes, Node 1, Node 2, and Node 3, and one IP address, xxx.xxx.xxx.xx6, then the IP address is applied as shown in the following list:

Node 1: xxx.xxx.xxx.xx6

Node 2: xxx.xxx.xxx.xx7

Node 3: xxx.xxx.xxx.xx8

vm_type

Specifies the type of the hypervisor. The following three types of hypervisors are valid: xen, ovm, or rhev.

vm_server

Specifies the host name or IP address of the hypervisor manager.

vm_svr_username

Specifies the user name of the hypervisor manager.

vm_svr_password

Specifies the password of the hypervisor manager. The password must be encrypted using the d2dutil --encrypt utility.

The following command is used to encrypt the password:

```
echo "password" | d2dutil --encrypt
```

vm_network

Specifies the virtual network that is used by the target VM. It is recommended to specify this parameter when your target VM is connected to multiple virtual networks.

guest_gateway

Specifies the network gateway that is used by the guest operating system (OS) of the target VM.

guest_netmask

Specifies the net mask that is used by the guest OS of the target VM.

guest_username

Specifies the username that is used to connect to the recovered VM. The password is reset to the password specified in the guest_password parameter. The guest_username parameter is ignored when you use the d2dverify utility to query information from the Backup Server database. In such cases, the VM guest password is reset to the node's password stored in database.

guest_password

Specifies the password for the guest_username parameter. The password must be encrypted using the d2dutil --encrypt utility. The guest_password parameter is ignored when you use the d2dverify utility to query information from the Backup Server database.

storage_location

Specifies the network path of the backup storage location. You do not have to specify the storage location if the nodes in the node_list parameter are in the Backup Server database. If the storage location is a CIFS share, use the following format to specify the location:

//hostname/path

storage_username

Specifies the user name to access the backup storage location. This parameter is not required for an NFS share.

For a Windows domain user, use the following format to specify the location:

domain_name/username

storage_password

Specifies the password to access the backup storage location. The password must be encrypted using the d2dutil --encrypt utility. This parameter is not required for an NFS share.

recovery_point = last

Specifies the session that you want to restore. Typically, a recovery session is in the following format: S00000000X, where X is a numeric value. S00000000X is the folder name of the recovery points. If you want to restore the most recent session, specify the keyword 'last'.

encryption_password

Specifies the encryption password for the recovery point. The password must be encrypted using the d2dutil --encrypt utility.

script

Specifies the script that you want to run. The script runs on the target machine after a successful recovery. If this parameter is not provided, the d2dverify utility runs the 'ls /proc' command on the target machine.

email_to_address

Specifies the email address of the recipients who will receive reports in an email. You can specify more than one email address, separated by a comma.

email_subject

Specifies the subject line of the email.

report_format

Specifies the format of the report that you will receive in an email. The format could be either text (.txt) or html.

Default: html

node_not_in_db

Specifies the nodes from the node_list parameters that are not in the Backup Server database. You must specify the storage_* related parameters.

Value: yes

stop_vm_after_recovery

Specifies that the target VM stops after a successful recovery and verification. The values for this parameter are yes and no.

Default: yes

4. Save and close the template.
5. Run the d2dverify utility using the following command:

```
d2dverify --template=file_path
```

Note: The d2dverify utility fails if the nodes in the node_list parameter are added using the public/private key. To resolve this issue, configure the environment variable 'export D2D_SSH_IGNORE_PWD=yes' in the shell environment where you run the d2dverify utility.

The usability of recovery points has been successfully verified.

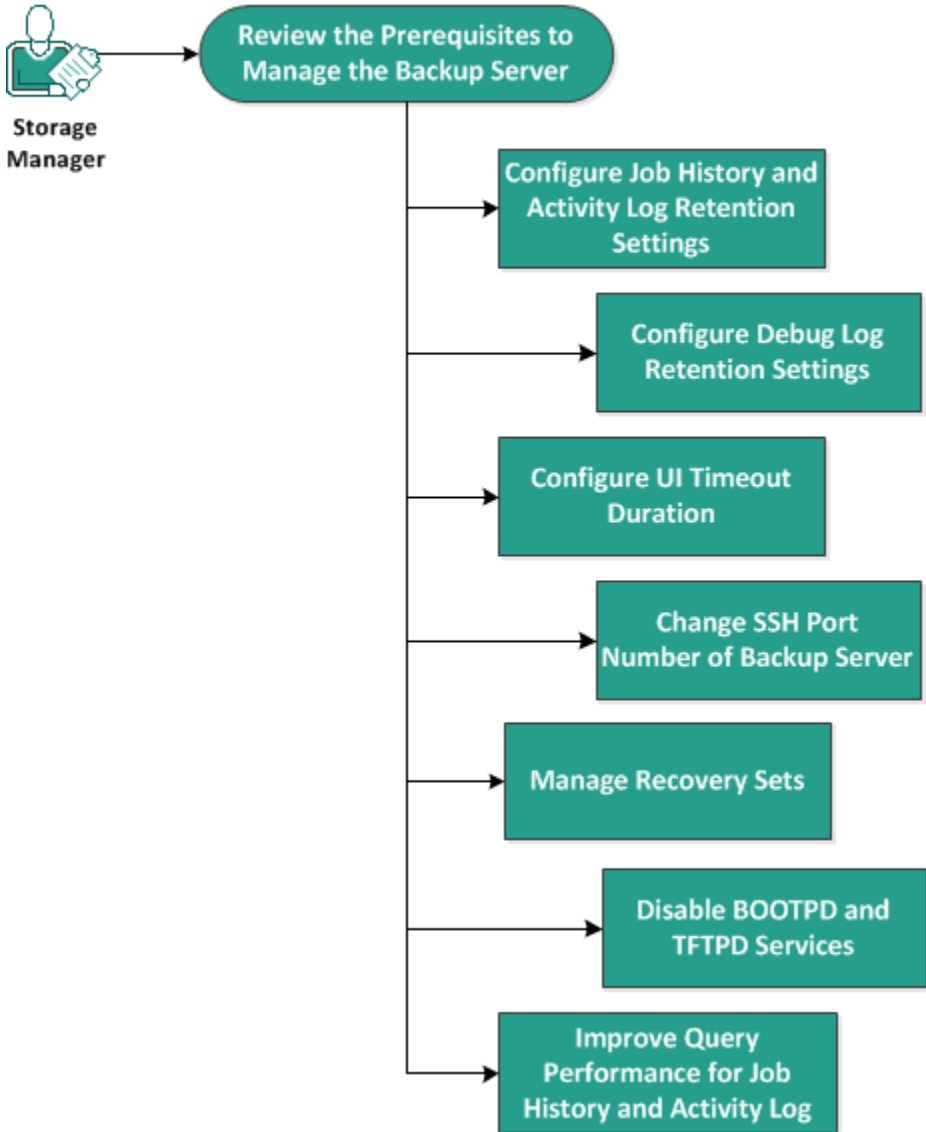
How to Manage the Backup Server Settings

You can perform the following tasks to manage the Backup Server:

- Configure the duration to retain the Job History and Activity Logs
- Configure the duration to retain the debug logs
- Change the secure shell (SSH) port number of the Backup Server

The following diagram displays the process to manage the Backup Server:

How to Manage the Backup Server Settings



Perform the following tasks to manage the Backup Server settings:

- [Review the Prerequisites to Manage the Backup Server](#) (see page 166)
- [Configure the Job History and Activity Log Retention Settings](#) (see page 166)
- [Configure the Debug Log Retention Settings](#) (see page 167)
- [Configure the UI Timeout Duration](#) (see page 168)
- [Change the SSH Port Number of the Backup Server](#) (see page 168)
- [Manage the Recovery Sets](#) (see page 169)
- [Disable the BOOTPD and TFTP Services](#) (see page 170)
- [Improve the Query Performance for Job History and Activity Log](#) (see page 170)
- [Skip CIFS and NFS Client Verification](#) (see page 171)

Review the Prerequisites to Manage the Backup Server

Consider the following prerequisites before you manage the Backup Server:

- You have the root login credentials to the Backup Server.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Configure the Job History and Activity Log Retention Settings

You can configure the duration to retain the Job History and Activity Logs. If you want to retain the Activity Logs and Job History for a longer duration, you have to configure the server file.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Open the server.cfg file:

```
/opt/CA/d2dserver/configfiles/server.cfg
```

Note: If the file is not present, create the server.cfg file.

3. Add the following line in the server.cfg file:

```
job_history_activity_log_keep_day=<number of days>
```

Example: To retain the Job History and Activity Log for 30 days, enter the following line:

```
job_history_activity_log_keep_day=30
```

Note: By default, the Job History and Activity Logs are retained for 90 days.

The Job History and Activity Log is retained for the specified time.

Configure the Debug Log Retention Settings

You can configure the duration to retain the debug logs. If you want to retain the debug logs for a longer duration, you have to configure the server file.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Open the server.cfg file:

```
/opt/CA/d2dserver/configfiles/server.cfg
```

Note: If the file is not present, create the server.cfg file.

3. Add the following line in the server.cfg file:

```
d2d_log_keep_day =<number of days>
```

Example: To retain the debug log for 30 days, enter the following line:

```
d2d_log_keep_day =30
```

Note: By default, the Debug Logs are retained for 90 days.

The Arcserve UDP Agent (Linux) debug log is retained for the specified period.

Configure the UI Timeout Duration

You can configure the webserver configuration file so that you are logged out of the UI when the UI is inactive. After you configure the file, if you do not perform any activity on the UI for the specified duration, you are logged out automatically. You can log in again and resume your activity.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Open the server.cfg file from the following location:

```
/opt/CA/d2dserver/configfiles/server.cfg
```

Note: If the server.cfg file is not present, create the file.

3. Add the following line in the server.cfg file:

```
ui_timeout=<value>
```

The value must be in minutes. The maximum limit for the UI timeout value is 60.

Example:

```
ui_timeout=40
```

The example indicates that if the Backup Server does not detect any activity on the UI for 40 minutes, it logs out the user.

4. Refresh the web browser to implement the changes.

The duration for the UI timeout is configured.

Change the SSH Port Number of the Backup Server

Backup Server uses the default secure shell (SSH) port 22 to connect to the nodes. If you want to change the default port to a different port, you can configure the server.env file to specify the new port.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Open the server.env file.

```
/opt/CA/d2dserver/configfiles/server.env
```

Note: If the file is not present, create the server.env file.

3. Add the following line in the server.env file and save the file:

```
export D2D_SSH_PORT=new_port_number
```

The new_port_number must be a numeric value.

4. Restart the Backup Server.

After you configure the `server.env` file, all jobs, except the BMR job, use the new port number to connect to the target node. The BMR job uses the default port.

The SSH port number of the Backup Server is successfully changed.

Manage the Recovery Sets

Managing the recovery sets include deleting the recovery sets. You should manage your recovery sets regularly so that you are aware of the available free space. You can plan the storage of the recovery sets accordingly. There are two ways to manage the recovery sets:

- **Method 1:** Manage using a dedicated Backup storage. In this method the backup storage manages the recovery sets every 15 minutes. You can manage only those backup storages that the Backup Server can access. If you choose source local as the backup destination, you have to share the local folder.
- **Method 2:** Manage using a Backup Job. In this method the backup job manages the recovery sets. The recovery sets are managed after the backup job is over. You can manage the recovery sets that are stored in source local.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Open the `server.cfg` file.

```
/opt/CA/d2dserver/configfiles/server.cfg
```

Note: If the file is not present, create the `server.cfg` file.

3. Add the following line in the `server.cfg` file and save the file:

```
manage_recoveryset_local=0 or 1
```

The value 0 indicates that the file uses Method 1.

The value 1 indicates that the file uses Method 2.

4. Restart the Backup Server.

The recovery sets are managed from the command line of the Backup Server.

Disable the BOOTPD and TFTP Services

You can disable the BOOTPD and TFTP services if you do not require the PXE BMR function.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Open the server.env file.

```
/opt/CA/d2dserver/configfiles/server.env
```

3. Update the following parameter in the server.env file and save the file:

```
export D2D_DISABLE_PXE_SERVICE=yes
```

4. Restart the Backup Server

```
/opt/CA/d2dserver/bin/d2dserver restart
```

The BOOTPD and TFTP services are successfully disabled.

Improve the Query Performance for Job History and Activity Log

If you have a larger database file, then querying Job History and Activity Log takes a lot of time. You can improve the query time for Job History and Activity Log using specific switches and get your output in a short time.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Open the server.cfg file:

```
/opt/CA/d2dserver/configfiles/server.cfg
```

Note: If the file is not present, create the server.cfg file.

3. Add the following lines in the server.cfg file:

- To improve Job History query performance, add the following line:

```
skip_getting_job_history_count=true
```

- To improve Activity Log query performance, add the following line:

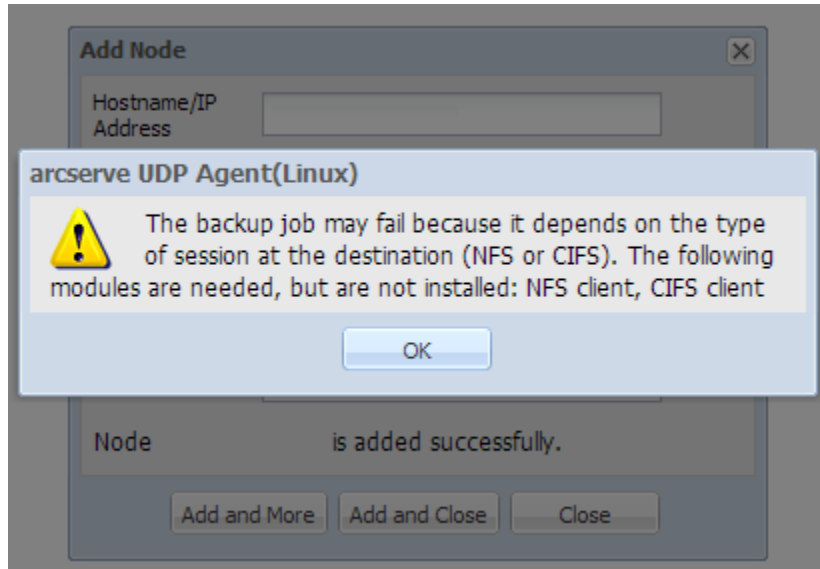
```
skip_getting_activity_log_count=true
```

4. Save the server.cfg file.

The query time for Job History and Activity Log has been successfully improved.

Skip CIFS and NFS Module Verification

When you add or modify a node, the Backup Server verifies the CIFS and NFS modules on the target node. If any of the module is not installed, a warning dialog opens. You can hide this dialog by configuring the server.cfg file.



Follow these steps:

1. Log in to the Backup Server.
2. Open the server.cfg file:
3. Add the following parameter:

```
/opt/CA/d2dserver/configfiles/server.cfg
```

```
skip_client_check=nfs,cifs
```

The given example skips the verification of both NFS and CIFS modules on the target node. When you provide both the modules, then the verification is skipped for both the modules. When you provide only one module, then the verification is skipped for only that module.

4. Save the server.cfg file.


The verification is skipped for CIFS and NFS modules.

How to Manage the Linux Backup Server from the Command Line

The Linux Backup Server performs all the processing tasks of Arcserve UDP Agent (Linux). For the smooth functioning of Arcserve UDP Agent (Linux), you must ensure that the Backup Server is running all the time. You can log into the Backup Server and manage the server using some commands.

For example, if you want to access the web interface of Arcserve UDP Agent (Linux), you must ensure that the web server is running. You can verify the running status of the web server from the Backup Server and ensure proper functioning of Arcserve UDP Agent (Linux).

The following diagram displays the process to manage the Backup Server from the command line:



Perform the following tasks to manage the Backup Server:

- [Review the Backup Server Prerequisites](#) (see page 173)
- [Start, Stop, or Release the Backup Server](#) (see page 173)
- [Change the Web Service Port Number of the Backup Server](#) (see page 174)
- [Configure the Private Key and Public Key Authentication](#) (see page 175)
- [Change the Backup Server Protocol](#) (see page 176)
- [Avoid the SSL Certificate Error While Opening Arcserve UDP Agent \(Linux\)](#) (see page 177)
- [Configure the System Settings When the Host Name or IP Address is Changed](#) (see page 179)

Review the Backup Server Prerequisites

Consider the following prerequisites before you manage the Backup Server:

- You have the root login credentials to the Backup Server.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Start, Stop, or Release the Backup Server

Manage your Backup Server to know the running status of the Backup Server. You can verify whether your Backup Server has stopped or is still running and then manage the server accordingly. Arcserve UDP Agent (Linux) supports the following command-line functions:

- Start the Backup Server
- Stop the Backup Server
- Release the Backup Server

Follow these steps:

1. Navigate to the bin folder using the following command:

```
# cd/opt/CA/d2dserver/bin
```

You gain access to the bin folder.

2. From the bin folder, run the following commands depending on the task that you want to perform on the server:

Note: If any command is not successful, an error message is displayed explaining the reason.

```
# ./d2dserver start
```

Starts the Backup Server.

If you are successful, a message is displayed informing you that the server has started.

```
# ./d2dserver stop
```

Stops the Backup server.

If you are successful, a message is displayed informing you that the server has stopped.

```
# ./d2dserver restart
```

Restarts the Backup server.

If you are successful, a message is displayed informing you that the server has restarted.

```
# ./d2dserver status
```

Displays the status of the Backup server.

```
# /opt/CA/d2dserver/bin/d2dreg --release
```

Releases the remaining Backup Servers that are managed by the main server.

For example, if Backup Server A manages two other servers, Backup Server B and Backup Server C, then when you uninstall Backup Server A you cannot access Backup Server B and Backup Server C. You can release Backup Server B and Backup Server C using this script and can access those servers.

The Backup Server is successfully managed from the command line.

Change the Web Service Port Number of the Backup Server

Arcserve UDP Agent (Linux) uses port 8014 by default. If the 8014 port number is used by other application, Arcserve UDP Agent (Linux) will not function properly. In such situations, you must change the Arcserve UDP Agent (Linux) default port number to a different port number.

Follow these steps:

1. Open the server.xml file from the following location:

```
/opt/CA/d2dserver/TOMCAT/conf/server.xml
```

2. Search the following string in the file and change the port number 8014 to your desired port number:

```
<Connector port="8014" protocol="HTTP/1.1" SSLEnabled="true" maxThreads="150" scheme="https"
secure="true" clientAuth="false" sslProtocol="TLS" keystoreFile="{catalina.home}/conf/server.keystore"
keystorePass="LinuxD2D"/>
```

3. Run the following command to restart the Backup Server:

```
/opt/CA/d2dserver/bin/d2dserver restart
```

The default port number is changed to your desired port number.

Configure the Private Key and Public Key Authentication

The public key and the private key allow you to securely connect to the nodes when you do not provide the password. Each time the Backup Server creates an SSH connection with the nodes, the Backup Server verifies the public key and private key for the respective nodes. If the keys do not match, you get an error message.

Note:

- Only the users having the root permission are supported to use the public key and private key authentication. It is not necessary to have the user name as root. The non-root users are not supported to use the public key and private key authentication. The non-root users must provide the user name and password authentication.
- Public key and private key authentication takes effect when the password is not provided. The username is still required and it must match the owner of the key.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Generate a public/private key using the following ssh-keygen command:

```
ssh-keygen -t rsa -f server
```

Two files are generated, namely server.pub and server.

3. Copy the public key file server.pub to the following location:

```
/opt/CA/d2dserver/configfiles/server_pub.key
```

4. Copy the private key file server to the following location:

```
/opt/CA/d2dserver/configfiles/server_pri.key
```

5. (Optional) Run the following command if you have provided the passphrase while generating the private and public keys:

```
echo "passphrase" | d2dutil encrypt > /opt/CA/d2dserver/configfiles/key.pass
```

6. Change the permission for the key.pass file using the following command:

```
chmod 600 /opt/CA/d2dserver/configfiles/key.pass
```

7. Log into the source node.

8. Copy the content from the server_pub.key file in the Backup Server to the following location in the node:

```
/root/.ssh/authorized_keys
```

The private key and the public key are successfully configured. You can connect to the source nodes using the public key and private key.

Change the Backup Server Protocol

Arcserve UDP Agent (Linux) is installed with the https protocol. You can change your protocol if you do not want to transfer data with encryption. We recommend you to use https because all the data transferred with https are encrypted. The data transferred with http are plaintext.

Follow these steps:

1. Open the server.xml file from the following location:

```
/opt/CA/d2dserver/TOMCAT/conf/server.xml
```

2. Search the following string in the server.xml file:

```
<!--<Connector connectionTimeout="180000" port="8014" protocol="HTTP/1.1"/>-->
```

3. Remove the <!-- and --> string characters as shown in the following example:

Example: The following string is the desired output after removing the <!-- and --> string characters:

```
<Connector connectionTimeout="180000" port="8014" protocol="HTTP/1.1"/>
```

4. Search the following string in the server.xml file:

```
<Connector port="8014" protocol="HTTP/1.1" SSLEnabled="true" maxThreads="150" scheme="https"
secure="true" clientAuth="false" sslProtocol="TLS" keystoreFile="${catalina.home}/conf/server.keystore"
keystorePass="LinuxD2D"/>
```


5. Add the <!-- and --> string characters as shown in the following example:

Example: The following string is the desired output after adding the <!-- and --> string characters:

```
<!--<Connector port="8014" protocol="HTTP/1.1" SSLEnabled="true" maxThreads="150" scheme="https"
secure="true" clientAuth="false" sslProtocol="TLS" keystoreFile="{catalina.home}/conf/server.keystore"
keystorePass="LinuxD2D"/>-->
```

6. Run the following command to restart the Backup Server:

```
/opt/CA/d2dserver/bin/d2dserver restart
```

The Backup Server protocol is changed from https to http.

Avoid the SSL Certificate Error While Opening Arcserve UDP Agent (Linux)

Remove the custom SSL certificate so that you do not get the certificate error when you open the Arcserve UDP Agent (Linux) web interface. Once you configure the SSL certificate, you do not get the certificate error again.

Follow these steps:

- Use the certificate generated by Arcserve UDP Agent (Linux) for the Firefox browser.
 1. Open Arcserve UDP Agent (Linux) in Firefox.
 2. Click I Understand the Risks and then click Add Exception.
The Add Security Exception dialog opens.
 3. Click View to review the certificate.
The Certificate Viewer dialog opens.
 4. Review the certificate details and click Close.
You do not have to perform any action on the Certificate Viewer dialog.
 5. On the Add Security Exception dialog, select the Permanently store this exception check box.
 6. Click Confirm Security Exception.
The certificate is added.

- Use the certificate generated by Arcserve UDP Agent (Linux) for the Internet Explorer (IE) or the Chrome browser.
 1. Open Arcserve UDP Agent (Linux) in IE or Chrome.
 2. Click Continue to this website (not recommended).

The address bar is displayed in red and a Certificate Error message is displayed in the security status bar.
 3. Click Certificate Error.

The Untrusted Certificate dialog appears.
 4. Click View certificates.

The Certificate dialog opens.
 5. On the General tab, click Install Certificate.

The Certificate Import Wizard opens.
 6. Click Next.
 7. On the Certificate Store page, select Place all certificates in the following store and then click Browse.

The Select Certificate Store dialog opens.
 8. Select Trusted Root Certification Authorities and click OK.

The Certificate Store page of the Certificate Import Wizard opens.
 9. Click Next and then click Finish.

The Security Warning dialog opens.
 10. Click Yes.
 11. Restart IE or Chrome.

The certificate is added.

Note: After you add the certificate, the Chrome browser still shows the error icon for the SSL certificate in the address bar. This is a reminder that the certificate is not identified by the certificate authorities but the certificate is trusted by Chrome and all the data transferred in the network is encrypted.
- Perform the following steps to use a signed certificate:
 1. Use the certificate signed by a certificate authority.
 2. Import the signed certificate using the keytool command.

The certificate is added.

The ssl certificate error is resolved.

Configure the System Settings When the Host Name or IP Address is Changed

If you change the host name or the IP address of the Backup Server or the client node (backup node), you have to configure the system settings. You configure the system settings to help ensure the following items:

- To ensure that the communication between the central server and the member server is good. A member server is a Backup Server that you manage from the central Backup Server. To manage the member server from the central server UI, you must add the member server in the central server UI.
- To ensure that after you change the host name or IP address of the client node you can back up the client node without any error

When the Host Name of the Central Backup Server is Changed

When you change the host name of the central Backup Server, you must configure the server so that you can use Arcserve UDP Agent (Linux) without any problem.

Follow these steps:

1. Log into the central Backup Server as a root user.
2. To update the host name and the license information, enter the following commands:

```
source /opt/CA/d2dserver/bin/setenv
```

```
/opt/CA/d2dserver/sbin/sqlite3 /opt/CA/d2dserver/data/ARCserveLinuxD2D.db "update D2DServer set Name='New_Hostname' where IsLocal=1"
```

```
/opt/CA/d2dserver/sbin/sqlite3 /opt/CA/d2dserver/data/License.db "update LicensedMachine set ServerName ='New_Hostname' where ServerName ='Old_Hostname'"
```

3. Rename the keystore file:

```
mv /opt/CA/d2dserver/TOMCAT/conf/server.keystore /opt/CA/d2dserver/TOMCAT/conf/server.keystore.old
```

4. Create a keystore file using the following keytool Java command.

```
keytool -genkey -alias tomcat -keyalg DSA -keypass <YOUR_VALUE> -storepass <YOUR_VALUE> -keystore /opt/CA/d2dserver/TOMCAT/conf/server.keystore -validity 3600 -dname "CN=<New Hostname>"
```

Note: Update the YOUR_VALUE field according to your requirement. Typically, the value is your password.

Example:

```
keytool -genkey -alias tomcat -keyalg DSA -keypass LinuxD2D -storepass LinuxD2D -keystore /opt/CA/d2dserver/TOMCAT/conf/server.keystore -validity 3600 -dname "CN=New Hostname"
```

5. Open the server.xml TOMCAT configuration file and change the keystoreFile value and the keystorePass value according to the keystore file that you just created:

```
<Connector port="8014" protocol="HTTP/1.1" SSLEnabled="true" maxThreads="150" scheme="https" secure="true" clientAuth="false" sslProtocol="TLS" keystoreFile="{catalina.home}/conf/server.keystore" keystorePass="YOUR_VALUE"/>
```

Example:

```
<Connector port="8014" protocol="HTTP/1.1" SSLEnabled="true" maxThreads="150" scheme="https" secure="true" clientAuth="false" sslProtocol="TLS" keystoreFile="{catalina.home}/conf/server.keystore" keystorePass="LinuxD2D"/>
```

6. Restart the central Backup Server.

```
/opt/CA/d2dserver/bin/d2dserver restart
```

The central Backup Server is configured.

When the Host Name or IP Address of the Member Server is Changed

When you change the host name or the IP address of the member Backup Server, configure the member server to manage it from the central server. If you do not configure the member server, then you will have an error when you try to manage it from the central server. A member server is a server that you have added to the central Backup Server web interface.

Follow these steps:

1. Log into the member Backup Server as a root user:

2. To change the host name, enter the following commands:

```
source /opt/CA/d2dserver/bin/setenv
```

```
/opt/CA/d2dserver/sbin/sqlite3 /opt/CA/d2dserver/data/ARCServeLinuxD2D.db "update D2DServer set Name='New_Hostname' where IsLocal=1"
```

3. Rename the keystore file:

```
mv /opt/CA/d2dserver/TOMCAT/conf/server.keystore /opt/CA/d2dserver/TOMCAT/conf/server.keystore.old
```

4. Create a keystore file using the following keytool Java command.

```
keytool -genkey -alias tomcat -keyalg DSA -keypass <YOUR_VALUE> -storepass <YOUR_VALUE> -keystore /opt/CA/d2dserver/TOMCAT/conf/server.keystore -validity 3600 -dname "CN=<New Hostname>"
```

Note: Update the YOUR_VALUE field according to your requirement. Typically, the value is your password.

Example:

```
keytool -genkey -alias tomcat -keyalg DSA -keypass LinuxD2D -storepass LinuxD2D -keystore /opt/CA/d2dserver/TOMCAT/conf/server.keystore -validity 3600 -dname "CN=New Hostname"
```

5. Open the server.xml TOMCAT configuration file and change the keystoreFile value and the keystorePass value according to the keystore file.

```
<Connector port="8014" protocol="HTTP/1.1" SSLEnabled="true" maxThreads="150" scheme="https" secure="true" clientAuth="false" sslProtocol="TLS" keystoreFile="{catalina.home}/conf/server.keystore" keystorePass="YOUR_VALUE"/>
```

Example:

```
<Connector port="8014" protocol="HTTP/1.1" SSLEnabled="true" maxThreads="150" scheme="https" secure="true" clientAuth="false" sslProtocol="TLS" keystoreFile="{catalina.home}/conf/server.keystore" keystorePass="LinuxD2D"/>
```

6. Restart the member Backup Server.

```
/opt/CA/d2dserver/bin/d2dserver restart
```

7. Log into the central Arcserve Unified Data Protection Agent for Linux web interface.
8. From the Backup Servers pane, select the old host name server.
9. From the Backup Server menu, click Delete.
10. In the Delete dialog, click OK.

The old host name server is deleted.

11. From the Backup Server menu, click Add.

The Add Server dialog opens.

12. Enter the new host name details in the dialog and click OK.

The Add Server dialog closes and the member server with the new host name is added to the UI.

13. Log into the central Backup Server that manages the member Backup Server.

14. To update the license information, enter the following commands:

```
source /opt/CA/d2dserver/bin/setenv

/opt/CA/d2dserver/sbin/sqlite3 /opt/CA/d2dserver/data/License.db "update LicensedMachine set
ServerName ='New_Hostname' where ServerName ='Old_Hostname'"
```

The member Backup Server is configured.

When the Host Name or the IP Address of the Client Node is Changed

If you change the host name or the IP address of a node, you can configure the host name or the IP address in the system settings so that you can back up that node without any error.

Follow these steps:

1. Log into the backup destination.
2. Locate the folder named "**Old_Hostname**" in the backup destination of this node and rename it to "**New_Hostname**".

For example, consider the old host name for node1 is First_Node. The backup destination for node1 is //Backup_Destination/LinuxBackup. After the first successful backup, a folder named First_Node is created in //Backup_Destination/LinuxBackup. Now, you have modified the old host name to Second_Node. Locate the First_Node folder in //Backup_Destination/LinuxBackup and rename the folder to Second_Node.

3. Log into the Backup server as a root user.

- To update the host name, enter the following commands:

```
source /opt/CA/d2dserver/bin/setenv
```

```
/opt/CA/d2dserver/bin/d2drp --storagepath=Backup Destination --node=New_Hostname
```

```
/opt/CA/d2dserver/sbin/sqlite3 /opt/CA/d2dserver/data/ARCserveLinuxD2D.db "update JobQueue set TargetName='New_Hostname' where JobType in (1,3,4,5) and TargetName='Old_Hostname'"
```

```
/opt/CA/d2dserver/sbin/sqlite3 /opt/CA/d2dserver/data/ARCserveLinuxD2D.db "update TargetMachine set Name='New_Hostname' where Name='Old_Hostname'"
```

Note: If you use NFS share or CIFS share as the backup destination, you should mount it to Local share.

Example: If your mount point is /mnt/backup_destination.

```
/opt/CA/d2dserver/bin/d2drp --storagepath=<mount point> --node=New_Hostname
```

Note: If you use Local share, then the command is:

```
/opt/CA/d2dserver/bin/d2drp --storagepath=<local path> --node=New_Hostname
```

- Log into the central Backup Server as a root user.
- To update the license information, enter the following command:

```
/opt/CA/d2dserver/sbin/sqlite3 /opt/CA/d2dserver/data/License.db "update LicensedMachine set MachineName = 'New_Hostname' where MachineName = 'Old_Hostname'"
```

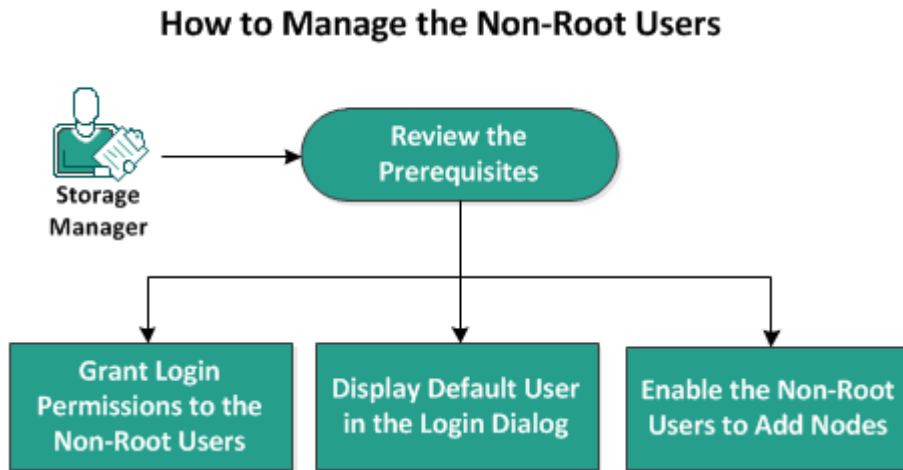
The host name is configured to perform a backup without any error.

How to Manage the Non-Root Users

You can manage all your non-root users that access Arcserve UDP Agent (Linux) and can define the permissions for the non-root users to limit the access level for Arcserve UDP Agent (Linux). You can manage the non-root users by modifying the webserver configuration file (server.cfg file).

Note: If your backup source node is configured with pam_wheel, then use the 'use_uid' option to configure pam_wheel. For more information about pam_wheel, see pam_wheel man page.

The following diagram displays the process to manage the non-root users:



Perform these tasks to manage the non-root users:

- [Review the Prerequisites](#) (see page 184)
- [Grant Login Permissions to the Non-Root Users](#) (see page 184)
- [Display the Default User in the Login Dialog](#) (see page 185)
- [Enable the Non-Root Users to Add Nodes](#) (see page 186)

Review the Prerequisites

Consider the following prerequisites before you manage the non-root users:

- You have the root login credentials to the Backup Server.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Grant Login Permissions to the Non-Root Users

A root user can grant permissions to non-root users to log into the Backup Server. If non-root users get the permission to log into the Backup Server, they can use Arcserve UDP Agent (Linux) to perform all the data protection and recovery tasks.

Note: To grant login permissions to the non-root users, connect to the Backup Server as a root user using the SSH connection.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Open the server.cfg file from the following location:

```
/opt/CA/d2dserver/configfiles/server.cfg
```

Note: If the server.cfg file is not present, create the file.

3. Add the following code to the server.cfg file:

```
allow_login_users=user1 user2
```

Note: Use blank spaces to distinguish multiple users.

The code is added.

4. Verify the non-root user can connect to the Backup Server using the SSH connection.

The login permission is granted to the non-users to access the Backup Server.

Display the Default User in the Login Dialog

You can manage your users and change the name that is displayed in the login dialog of Arcserve UDP Agent (Linux). The default user that is displayed in the login dialog is root. If you do not have root users accessing the product, you can change the default name to any non-root user name. You achieve this by modifying the server.cfg that is located in the Backup Server.

Note: To modify the server.cfg file, connect to the Backup Server as a root user using the SSH connection.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Open the server.cfg file from the following location:

```
/opt/CA/d2dserver/configfiles/server.cfg
```

Note: If the server.cfg file is not present, create the file.

3. Add the following code to the server.cfg file:
show_default_user_when_login=false|true

4. Log into the Arcserve UDP Agent (Linux) web interface.
 - If you have added the `allow_login_users` command, the login dialog displays the first user added in the `allow_login_users` command.
 - If you have not added the `allow_login_users` command, the login dialog displays the root user.

The default user is displayed in the login dialog of Arcserve UDP Agent (Linux)

Enable the Non-Root Users to Add Nodes

If the SSH Server disables the root user login, you can enable the non-root user login to add any nodes. When you enable the non-root user login credentials, the Add Node dialog changes and displays the Root Credential option.

Note: If you change the client node credential from a root user to a non-root user, it is recommended that you clear the */tmp* folder on the client node before you run the backup job.

The screenshot shows a dialog box titled "Add Node" with a close button in the top right corner. The dialog contains the following fields and controls:

- Hostname/IP Address: [Text Input Field]
- User Name: [Text Input Field]
- Password: [Text Input Field]
- Specify root credential if the user above is not root user.
- Root Credential**
- User Name: [Text Input Field]
- Password: [Text Input Field]
- Description: [Text Area]
- Buttons: Add and More, Add and Close, Close

Follow these steps:

1. Log into the Backup Server as a root user.
2. Open the server.cfg file from the following location:

`/opt/CA/d2dserver/configfiles/server.cfg`

Note: If the server.cfg file is not present, create the file.

3. Add the following line in the server.cfg file to enable the non-root user function:

`enable_non_root_user=true`

The non-root user function is enabled.

4. (Optional) Add the following line in the `server.cfg` file to disable the non-root user function:

```
enable_non_root_user=false
```

The non-root user function is disabled.

The non-root users are enabled to add nodes.

Note: If you change the password for the root user or the non-root user and then you modify the node, you must reenter both the root password and the non-root password in their respective fields in the Modify Node dialog.

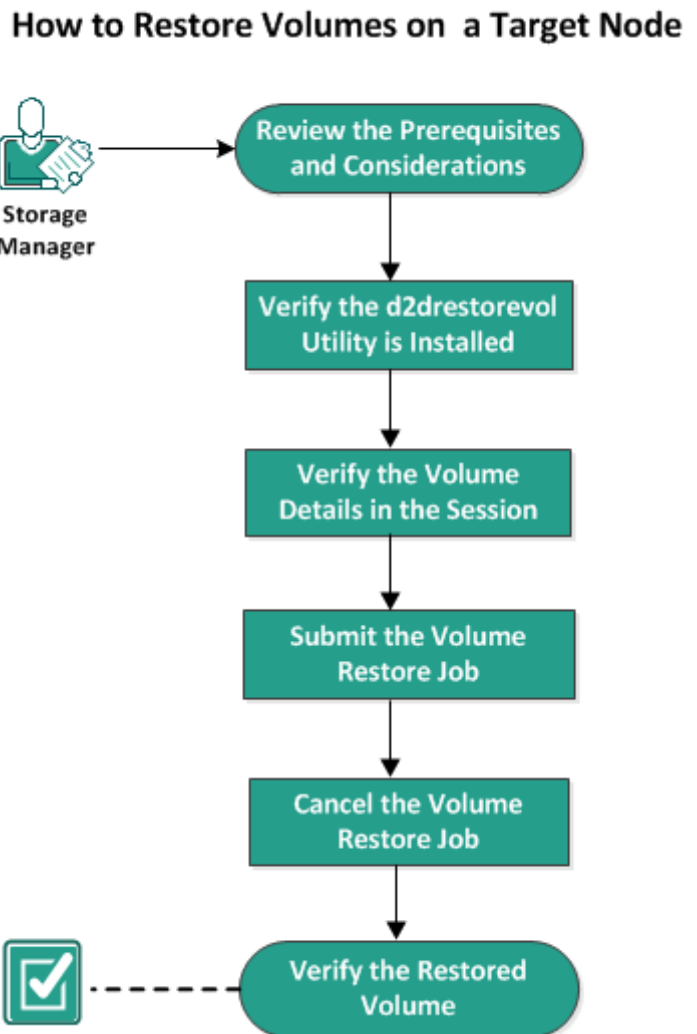
Note: The non-root users cannot manage nodes using the `d2dnode` command from the command line.

How to Restore Volumes on a Target Node

You can restore individual volumes on the target node without performing a full BMR. The target node can be a Backup Server or a protected node.

Restoring individual volumes utilizes less resources and provides a better performance.

The following diagram displays the process to restore volumes:



Perform the following steps to restore volumes:

- [Review the Prerequisites and Considerations](#) (see page 189)
- [Verify the d2drestorevol Utility is Installed](#) (see page 189)
- [Verify the Volume Details in the Session](#) (see page 190)
- [Submit the Volume Restore Job](#) (see page 191)
- [Cancel the Volume Restore Job](#) (see page 193)
- [Verify the Restored Volume](#) (see page 193)

Review the Prerequisites and Considerations

Review the following prerequisites before restoring volumes:

- You have a valid backup session to perform a restore.
- Backup sessions must be accessed locally on the target node. If the session location is on the local volume of the target node, use the exact directory path as the session location. If the session location is on a network share, first mount the network share to a local mount point and then use the mount point path as the session location.
- Target volumes that you want to restore must be un-mounted, using the `umount` command:

Example: `umount /dev/sda2`
- The target volume must be equal or larger than the source volume.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Review the following considerations before you restore volumes:

- When you restore, any existing data on the target volume is erased. Perform a backup of your existing data from the target volume before you restore.

Verify the `d2drestorevol` Utility is Installed

The `d2drestorevol` utility restores volume to the target node. The target node can be a Backup Server or any other Linux node (client). If the `restorevol` utility is not installed on the target node, you must manually install the utility.

Restore on a Backup Server

If the target node is a Backup Server, the utility is already installed with the installation package. Verify if the utility is present at the `bin` folder.

Follow these steps:

1. Log into the Backup Server.
2. Verify that the utility is located at the following location:

```
/opt/CA/d2dserver/bin/d2drestorevol
```

The utility is installed and verified.

Restore on a client

A client node will not have the utility installed in it. You have to manually install the utility on the client.

Important! The utility must be downloaded from the Backup Server as described in the following steps. If you manually copy the utility from a Backup Server to a client, the utility may not work properly.

Follow these steps:

1. Log into the client.
2. Locate the d2drestorevol utility download path from the command line.

```
http[s]://[Backup-Server-address]:[port]/d2drestorevol
```

3. Download the script using a command line tool, such as wget.

```
wget http://192.168.1.1:8014/d2drestorevol -O d2drestorevol
```

Note: If the server uses the https protocol, you may have to include the '--no-check-certificate' parameter.

```
wget https://192.168.1.1:8014/d2drestorevol -O d2drestorevol --no-check-certificate
```

4. Provide the execution permission to the utility using the following command:

```
chmod +x d2drestorevol
```

The permission is provided.

The d2drestorevol is installed and verified.

Verify the Volume Details in the Session

Verify the volume details in the session that you want to restore. You can see the source volume, file system, file size, and mount information in the output.

Follow these steps:

1. Log into the target node.
2. Use the following command to verify the volume information:

```
d2drestorevol --command=info --storage-path=<local_path> --node=<node_name> --rp=<recovery_point>
```

--command=info

Specifies that the volume details of the session will be displayed.

--storage-path

Specifies the path we determined in the Prerequisites topic. For more information, see Review the Prerequisites and Considerations.

--node

Specifies the source node that was backed up.

--rp

Specifies the recovery point or recovery session that you want to restore.

Typically, a recovery point is in the following format: S00000000X, where X is a numeric value.

The output is displayed.

The volume details are verified.

Submit the Volume Restore Job

Submit the volume restore job to start restoring your volume on the target node.

Follow these steps:

1. Log into the target node.
2. Submit the restore job using the following command:

```
d2drestorevol --command=restore --storage-path=<local_path> --node=<node_name>
--rp=<recovery_point> --source-volume=<source_volume> --target-volume=<target_volume>
[--encryption-password=<encryption_password>] [--mount-target=<mount_point> [--quick-recovery]]
```

--command=restore

Specifies that the volume restore job is submitted.

--storage-path

Specifies the path we determined in the Prerequisites topic. For more information, see Review the Prerequisites and Considerations.

--node

Specifies the source node that was backed up.

--rp

Specifies the recovery point or recovery session that you want to restore.

Typically, a recovery point is in the following format: S00000000X, where X is a numeric value.

--encryption-password

Specifies the session password. This option is required if the session is encrypted. If the session is encrypted but this option is not present, you will be prompted to enter the password from the terminal.

--source-volume

Specifies the source volume. You can get the source volume using the *command=info* parameter as described in the Verify the Volume Details in the Session topic, or the source volume can be the it can be the mount point from the source system.

--target-volume

Specifies the device file path of the target node.

Example: /dev/sda2

--mount-target

Specifies the mount point where the restored volume should be mounted.

Example: /mnt/volrestore

--quick-recovery

When used along with ‘--mount-target’, the target volume will be mounted as soon as possible. You can use the data on the target volume while the data is getting restored.

After the restore job is over, the restore process quits automatically and you can continue using the data without any interruption.

Note: When a volume restore job, and a backup job run at the same time, then:

- If --quick-recovery is used, then the job (volume restore or backup) that starts later does not run.
- If --quick-recovery is not used, then the backup job will back up only those volumes that are not being restored.

The restore job is submitted and a screen is opened that displays the progress. If you want to submit other jobs, you can either wait for the current job to complete or press Q to exit the screen and then submit a new job.

3. (Optional) Use the following command to review the progress of the volume restore job:

```
d2drestorevol --command=monitor
```

The progress details, such as volume name, elapsed time, progress, speed, status, and time remaining, are displayed on a screen.

The screen exits when the job completes. You can also press Q to manually exit the screen. Manually exiting the screen does not interrupt the running restore job.

The volume restore job is submitted.

Cancel the Volume Restore Job

You can cancel the volume restore job from the command line of the target node. Use the following command to cancel the volume restore job.

```
d2drestorevol --command=cancel --target-volume=<target_volume>
```

--command=cancel

Specifies that the volume restore job is cancelled.

--target-volume

Specifies the device file path of the target node. The value must be identical with the value used to submit the restore job.

Important: Canceling a volume restore job will make the target volume unusable. In such cases you can retry to perform the volume restore job or you can restore the lost data, if you have a backup.

Verify the Restored Volume

Verify the data when the volume is restored.

Follow these steps:

1. Log into the target node.
2. Review the progress screen to verify the completion status.
3. (Optional) Review the *d2drestvol_activity_[target volume].log* file to see all the logs of the restore job.
4. Mount the restored volume and verify the data is restored.

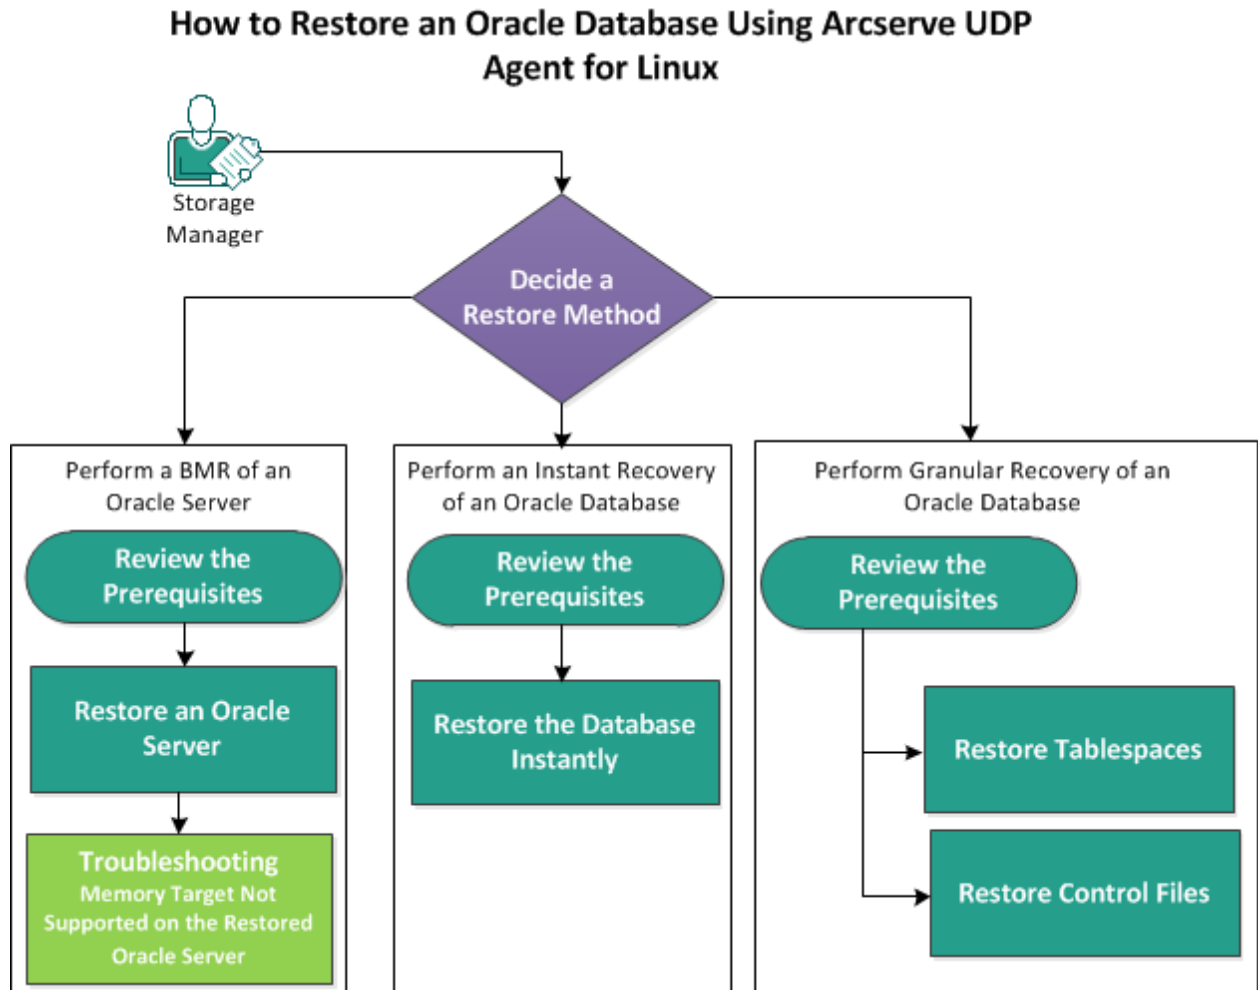
The volume restore job is verified.

The volume is successfully restored.

How to Restore an Oracle Database Using Arcserve UDP Agent (Linux)

You can restore the entire Oracle database, or restore specific files from the database. You can also perform a Bare Metal Recovery (BMR) of an Oracle server when the source server is not functioning properly. If you have lost the database and you want it available immediately, you can perform an instant recovery. Read the prerequisites for each type of restore, before you begin the restore process.

The following diagram illustrates the process to restore an Oracle database using Arcserve UDP Agent (Linux):



Perform the following steps to restore an Oracle database using Arcserve UDP Agent (Linux):

- [Perform a BMR of an Oracle Server](#) (see page 195)
- [Perform an Instant Recovery of an Oracle Database](#) (see page 196)
- [Perform Granular Recovery of an Oracle Database](#) (see page 198)

Perform a Bare Metal Recovery (BMR) of an Oracle Server

A BMR restores the operating system and software applications, and recovers all the backed-up data. BMR is the process of restoring a computer system from bare metal. Bare metal is a computer without any operating system, drivers, and software applications. After the restore is complete, the target machine automatically reboots in the same operating environment as the backup source node and all data is restored.

You can perform a BMR using the IP address or the Media Access Control (MAC) address of the target machine. If you boot the target machine using the Arcserve UDP Agent (Linux) Live CD, you can get the IP address of the target machine.

Review the Prerequisites

Review the following prerequisites before you restore the Oracle database:

- You have a valid recovery point and the encryption password, if any, for restore.
- You have a valid target machine for BMR.
- You have created the Arcserve UDP Agent (Linux) (Linux) Live CD.
- If you want to perform a BMR using the IP address, you must get the IP address of the target machine using the Live CD.
- If you want to perform a PXE-based BMR using the MAC address, you must have the MAC address of the target machine.
- Oracle database stores all the database-related files (data files, redo logs, archive logs, pfile, spfile, backups) on ext2, ext3, ext4, and ReiserFS. The database does not recognize Oracle Cluster File System (OCFS/OCFS2), Raw disks, or Automatic Storage Management (ASM) file systems.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Restore an Oracle Server

If the Oracle server is corrupted, you can restore the entire server by performing a BMR.

Follow these steps:

1. Log in to the Linux Backup Server Console as a root user.
2. Perform a BMR using the Restore Wizard. For more information on the restore process, see [How to Perform a Bare Metal Recovery \(BMR\) for Linux Machines](#).
3. After the BMR job is complete, log in to the target machine and verify that the database is restored.

The Oracle server is successfully recovered.

Memory Target Not Supported on the Restored Oracle Server

Symptom

I have performed a bare metal recovery of an Oracle server. The memory size of the target machine is less than the source Oracle server and the Oracle database uses AMM (Automatic Memory Management). After BMR, when I start the Oracle database instance, I get the following error:

SQL> startup

ORA-00845: MEMORY_TARGET not supported on this system

Solution

To resolve this error, increase the size of the shared memory virtual file system.

Follow these steps:

1. Log in to the target machine as a root user.
2. Open the command prompt and verify the size of the shared memory virtual file system.

```
# df -k /dev/shm
Filesystem      1K-blocks    Used Available Use% Mounted on
tmpfs           510324      88 510236   1% /dev/shm
```

3. Enter the following command and specify the required size of the shared memory:
mount -o remount,size=1200m /dev/shm
4. Navigate to the "/etc/fstab" folder and update the tmpfs setting:
tmpfs /dev/shm tmpfs size=1200m 0 0

Note: The shared memory virtual file system size should be big enough to accommodate the MEMORY_TARGET and MEMORY_MAX_TARGET values. For more information on the variables, refer to the Oracle documentation.

Perform an Instant Recovery of an Oracle Database

You can instantly recover an Oracle database without performing a full BMR. You can recover the database by using specific commands from the command line.

Review the Prerequisites

Review the following prerequisites before you restore the Oracle database:

- You have a valid recovery point and the encryption password, if any, for restore.
- Backup sessions must be accessed locally on the target node. If the session location is on the local volume of the target node, use the exact directory path as the session location. If the session location is on a network share, first mount the network share to a local mount point and then use the mount point path as the session location.
- Target volumes that you want to restore cannot be a root volume and must be un-mounted, using the `umount` command.

Example: `umount /dev/sda1`

- The target volume must be equal or larger than the source volume.
- Oracle database stores all the database-related files (data files, redo logs, archive logs, pfile, spfile, backups) on ext2, ext3, ext4, and ReiserFS. The database does not recognize Oracle Cluster File System (OCFS/OCFS2), Raw disks, or Automatic Storage Management (ASM) file systems.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Restore the Database Instantly

When you recover the database instantly, the database is available for immediate use. However, the recovery process runs in the backend and all the files are available only after the database is recovered completely.

Note: For more information on volume restore, see How to restore Volumes on a Target Node.

Follow these steps:

1. Log in to the target machine as a root user.
2. Open a command prompt as a root user.
3. Verify the target volume /dev/sdb1 is not mounted.

```
# df | grep 'target_volume'
```

Example: # df | grep '/dev/sdb1'

4. Mount the remote NFS share to the local path.

```
#mount <nfs_session_path>:/nfs <session_location_on_local>
```

Example: #mount xxx.xxx.xxx.xxx:/nfs /CRE_ROOT

5. Enter the following command to start the restore job:

```
#. /d2drestorevol --command=restore --storage-path=<session_location_on_local> --node=<oracle_server>
--rp=last --source-volume=<mount_point_for_oracle_data_volume>
--target-volume=<restore_target_volume_name> --mount-target=<mount_point_for_oracle_data_volume>
--quick-recovery
```

Example: #. /d2drestorevol --command=restore --storage-path=/CRE_ROOT
--node=rh63-v2 --rp=last --source-volume=/opt/oracle --target-volume=/dev/sdb1
--mount-target=/opt/oracle --quick-recovery

You can start the Oracle database immediately after the restore job starts. You do not have to wait for the completion of the database recovery.

6. Open another command prompt and log in with the Oracle user name and password.

```
$sqlplus / as sysdba
SQL>startup;
```

The Oracle database opens and you can perform the regular database operations such as query, insert, delete, update data, and so on.

The Oracle database is instantly recovered.

Perform Granular Recovery of an Oracle Database

You can restore specific files related to the Oracle database. These files may be control files, or data files of tablespaces.

Review the Prerequisites

Review the following prerequisites before you restore the Oracle database:

- You have a valid recovery point and the encryption password, if any.
- You have a valid target node to recover data.
- You have verified that the Linux Backup Server supports the file system that you want to restore.
- Oracle database stores all the database-related files (data files, redo logs, archive logs, pfile, spfile, backups) on ext2, ext3, ext4, and ReiserFS. The database does not recognize Oracle Cluster File System (OCFS/OCFS2), Raw disks, or Automatic Storage Management (ASM) file systems.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Restore Tablespaces

If a database tablespace is lost or corrupted, you can restore it by performing a File-Level Recovery. After the file-level recovery is successful, you have to manually recover the tablespace.

Follow these steps:

1. Log in to the target machine as a root user.
2. Make sure that the database is available.
3. Take the required tablespace offline.

Example: Consider that the name of the tablespace is MYTEST_DB. Enter the following command to take the tablespace offline:

```
$ sqlplus "/ as sysdba"
SQL> alter tablespace MYTEST_DB offline;
```

4. List all data files for the specified tablespace MYTEST_DB.

```
SQL> select file_name, tablespace_name from dba_data_files where tablespace_name='MYTEST_DB';
```

```
FILE_NAME
```

```
-----
```

```
TABLESPACE_NAME
```

```
-----
```

```
/opt/oracle/oradata/lynx/MYTEST_DATA01.dbf
```

```
MYTEST_DB
```

5. Restore the data files of tablespaces using the Restore Wizard. For more information on the restore process, see [How to Perform a File-Level Recovery on Linux Nodes](#).
6. Specify the following information on the Restore Wizard and submit the job:
 - a. When you select the files and folders, enter the required data file name of the tablespace and search.
Example: Enter "MYTEST_DATA01.dbf" of the tablespace "MYTEST_DB" and search.
 - b. On the Target Machine page, enter the following information:
 - Select Restore to original location.
 - Enter the hostname or IP address of the target Oracle Server.
 - Enter the root user name and the password of the target Oracle Server.
 - Select Overwrite existing files for the Resolving Conflicts option.
7. After the data file is restored, recover the tablespace of the Oracle database.
SQL>recover tablespace MYTEST_DB;
Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
Auto
8. Make the specified table space online.
SQL>alter tablespace MYTEST_DB online;

The tablespace is successfully recovered.

Restore Control Files

If database control files are lost or corrupted, you can restore it by performing a File-Level Recovery. After the file-level recovery is successful, you have to manually recover the control files.

Follow these steps:

1. Log in to the target machine as a root user.
2. Shut down the Oracle instance.

```
SQL>shutdown abort
```


3. Start the database in the nomount state.

```
SQL>startup nomount
```

4. List the path for all control files.

```
SQL> show parameter control_files;
```

NAME	TYPE	VALUE
------	------	-------

control_files	string	/opt/oracle/oradata/lynx/control01.ctl, /opt/oracle/flash_recovery_area/lynx/control02.ctl
---------------	--------	---

5. Restore the control files using the Restore Wizard. For more information on the restore process, see How to Perform a File-Level Recovery on Linux Nodes.
6. Specify the following information on the Restore Wizard and submit the job:

- a. When you select the files and folders, enter the required name of the control file and search. Repeat this step until all the control files are selected.

Example: Enter "control01.ctl " and search.

- b. On the Target Machine page, provide the following information:

- Select Restore to original location.
- Enter the hostname or IP address of the target Oracle Server.
- Enter the root user name and the password of the target Oracle Server.
- Select Overwrite existing files for the Resolving Conflicts option.

7. After all the control files are restored, mount the database and open it.

```
$sqlplus / as sysdba
```

```
SQL>alter database mount;
```

8. Recover the database with the RECOVER command and add the USING BACKUP CONTROLFILE clause.

```
SQL> RECOVER DATABASE USING BACKUP CONTROLFILE
```

9. Apply the prompted archived logs.

Note: If the required archived log is missing, then it implies that a necessary redo record is located in the online redo logs. It occurs because unarchived changes are located in the online logs when the instance failed. You can specify the full path of an online redo log file and press Enter (you may have to try this a few times until you find the correct log).

Example:

```
SQL> RECOVER DATABASE USING BACKUP CONTROLFILE
ORA-00279: change 1035184 generated at 05/27/2014 18:12:49 needed for thread 1
ORA-00289: suggestion :
/opt/oracle/flash_recovery_area/LYNX/archivelog/2014_05_27/o1_mf_1_6_%u_.arc
ORA-00280: change 1035184 for thread 1 is in sequence #6
Specify log: {<RET>=suggested | filename | AUTO | CANCEL}
/opt/oracle/oradata/lynx/redo03.log
Log applied.
Media recovery complete.
```

10. Open the database with the RESETLOGS clause after completing the recovery process.

```
SQL>alter database open resetlogs;
```

The control files are successfully recovered.

Chapter 5: Troubleshooting

This section contains the following topics:

[Arcserve UDP Agent \(Linux\) Fails to Install on Supported Servers](#) (see page 203)

[Arcserve UDP Agent \(Linux\) Displays an Operation Timeout Error](#) (see page 204)

[All Scheduled Jobs Fail When the System Time Is Changed to an Already Passed Value](#) (see page 205)

[Arcserve UDP Agent \(Linux\) Fails to Mount Linux Software RAID Devices](#) (see page 206)

[A Paravirtual Machine \(PVM\) Displays a Black Screen on the Virtual Network Computing \(VNC\) Client Window When Booted Using a Live CD](#) (see page 206)

[How to Adjust the Disk Boot Sequence After a BMR Job on an Oracle VM Server](#) (see page 207)

[The Backup Job Fails to Collect the BMR-related Information or the BMR Job Fails to Create a Disk Layout](#) (see page 209)

[Backup Read Throughput is Less When the Backup Node is RHEL, CentOS, or Oracle Linux 5.x On a Xen- based PVM](#) (see page 209)

[How to Restore the Previous Version of Backup Server](#) (see page 210)

[SLES 10.X Does Not Start Successfully After a BMR](#) (see page 211)

[The d2drestorevm and d2dverify Jobs Fail on Oracle VM Server](#) (see page 212)

[ESXi Virtual Machine Fails to Start After BMR From a Physical Machine](#) (see page 212)

Arcserve UDP Agent (Linux) Fails to Install on Supported Servers

Valid on CentOS 6.x, Red Hat Enterprise Linux (RHEL) 6.x, SUSE Linux Enterprise Server (SLES) 11.x, and Oracle Linux Server 6.x

Symptom

When I install Arcserve UDP Agent (Linux), the installation fails with the following Linux warning messages:

mkisofs Create Live CD image

mount.nfs Mount NFS share file system as backup destination and restore source

mount.cifs Mount CIFS share file system as backup destination and restore source

The following processes must be running

Inactive Processes	Affected Function
--------------------	-------------------

rpc.statd	The NFS file locking function does not work
-----------	---

Solution

At the beginning of the installation, Arcserve UDP Agent (Linux) verifies if the Linux OS meets the requirement of the Backup Server. If the Linux OS does not meet the minimum requirements, Arcserve UDP Agent (Linux) displays a warning message to inform you of this problem. The message includes the list of all the packages that are required for Backup Server.

To troubleshoot this Arcserve UDP Agent (Linux) installation problem, perform the following steps:

1. Install the following packages using the *yum* command:

- genisoimage
- nfs-utils
- cifs-utils

2. Run the following two commands:

```
service rpcbind start
```

```
service nfs start
```

3. Run the following command to verify if *rpc.statd* is running:

```
ps -ef|grep rpc.statd
```

4. Reinstall Arcserve UDP Agent (Linux).

Arcserve UDP Agent (Linux) is successfully installed.

Arcserve UDP Agent (Linux) Displays an Operation Timeout Error

Valid on CentOS 6.x, Red Hat Enterprise Linux (RHEL) 6.x, SUSE Linux Enterprise Server (SLES) 11.x, and Oracle Linux Server 6.x

Symptom

I get the following error message:

The operation has timed out. The maximum amount of time to complete the operation has been exceeded. Please try again later.

I get this message frequently when I perform a file-level restore and browse recovery points that have more than 1000 incremental recovery points.

Solution

The default timeout value is 3 minutes. You can troubleshoot the problem by increasing the timeout value.

Perform the following steps to increase the timeout value:

1. Log into the Backup Server as a root-user.
2. Add the following system environment variable:

```
D2D_WEBSVR_TIMEOUT
```

The value for the environment variable is a number. The number must be greater than 3. The unit for the value is minute.

3. Restart the Backup Server.

The timeout value is successfully increased.

All Scheduled Jobs Fail When the System Time Is Changed to an Already Passed Value

Valid on CentOS 6.x, Red Hat Enterprise Linux (RHEL) 6.x, SUSE Linux Enterprise Server (SLES) 11.x, and Oracle Linux Server 6.x

Symptom

When I change the system time to an already passed value, all my scheduled jobs are affected. The scheduled jobs fail to run after I change the system time to a previous time.

Solution

After you change the system time, restart the BACKUP service.

Follow these steps to restart the BACKUP service:

1. Log into the Backup Server as a root user.
 2. Navigate to the bin folder
- ```
/opt/CA/d2dserver/bin/
```
3. Restart the Backup Server using the following command:

```
d2dserver restart
```

The Backup Server restarts.

All the scheduled jobs run per schedule.

## Arcserve UDP Agent (Linux) Fails to Mount Linux Software RAID Devices

Valid on CentOS 6.x, Red Hat Enterprise Linux (RHEL) 6.x, SUSE Linux Enterprise Server (SLES) 11.x, and Oracle Linux Server 6.x

### Symptom

Sometimes the BMR process fails to mount Linux Software RAID devices after the target machine restarts.

### Solution

To solve this problem, restart your target machine.

## A Paravirtual Machine (PVM) Displays a Black Screen on the Virtual Network Computing (VNC) Client Window When Booted Using a Live CD

Valid on PVM on Oracle VM Server

### Symptom

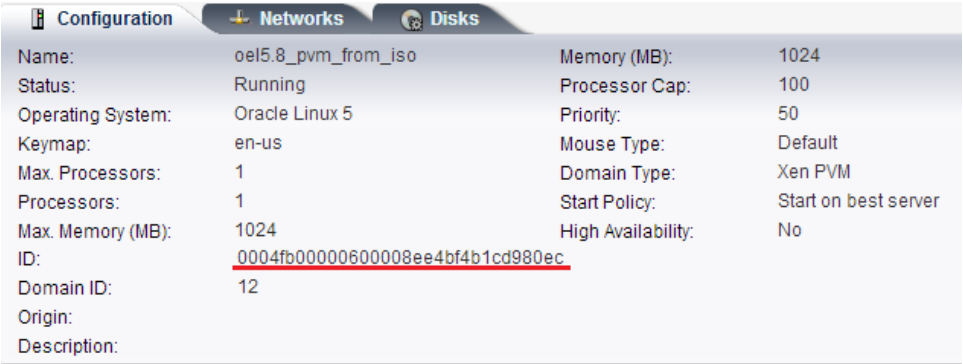
On an Oracle VM Server, when I boot the Paravirtual Machine (PVM) using a Live CD, I see a black screen on the VNC client window.

### Solution

To resolve this issue, log into the Live CD console from the backend.

### Follow these steps:

1. Start the VM using a Live CD.
2. Make a note of the VM's ID that you can access from the Oracle VM Manager.



The screenshot shows the Oracle VM Manager configuration page for a Paravirtual Machine (PVM). The page has three tabs: Configuration, Networks, and Disks. The Configuration tab is active, displaying the following details:

|                   |                                         |                    |                      |
|-------------------|-----------------------------------------|--------------------|----------------------|
| Name:             | oel5.8_pvm_from_iso                     | Memory (MB):       | 1024                 |
| Status:           | Running                                 | Processor Cap:     | 100                  |
| Operating System: | Oracle Linux 5                          | Priority:          | 50                   |
| Keymap:           | en-us                                   | Mouse Type:        | Default              |
| Max. Processors:  | 1                                       | Domain Type:       | Xen PVM              |
| Processors:       | 1                                       | Start Policy:      | Start on best server |
| Max. Memory (MB): | 1024                                    | High Availability: | No                   |
| ID:               | <u>0004fb00000600008ee4bf4b1cd980ec</u> |                    |                      |
| Domain ID:        | 12                                      |                    |                      |
| Origin:           |                                         |                    |                      |
| Description:      |                                         |                    |                      |

3. Log into the Oracle VM Server on which the VM is running using the Secure Shell (ssh).
4. Run the `xm console $ID` command as shown in the following diagram:

```
[root@ ~]# xm console 0004fb00000600008ee4bf4b1cd980ec
```

5. (Optional) Press Enter when prompted to confirm the operation.
6. The console of the Xen PVM booted with the Live CD opens.
7. Configure the network.
8. Exit the console by pressing `ctrl+]` or `ctrl+5`.

The issue is resolved.

## How to Adjust the Disk Boot Sequence After a BMR Job on an Oracle VM Server

### Valid on Oracle VM Server

#### Symptom

When I perform a BMR job to a target node on an Oracle VM Server, I get the following warning message in the Activity log:


The boot volume is restored to disk `/dev/xxx`. Please adjust the disk boot sequence in the BIOS to boot from `/dev/xxx`.

#### Solution


To avoid this problem, swap the disk boot sequence of the BMR target node.

#### Follow these steps:


1. Edit the BMR target node from the Oracle VM Manager and click the Disks tab.



2. Select Slot N disk as the Boot Disk.
3. Make a note of the disk name and the slot number N.  
You will use the disk name and the slot number in the later steps.
4. From the Actions column, select the Virtual Machine Disk button.
5. Select the Leave Slot Empty option and click Save.



6. Select Slot 0 Disk and make a note of the disk name.
7. From the Actions column, select the Virtual Machine Disk button.
8. Select the Leave Slot Empty option and click Save.
9. Attach the selected boot disk image to Slot 0 and the original Slot 0 Disk image to Slot N.



10. Boot the BMR target node.

The disk boot sequence is successfully adjusted.



## The Backup Job Fails to Collect the BMR-related Information or the BMR Job Fails to Create a Disk Layout

### Valid on Oracle VM Server for HVM with LVM volume

#### Symptom

When I perform a backup job for an HVM with LVM volumes on an Oracle VM Server, the backup job fails to collect the BMR-related information. Also, when I perform a BMR job for an HVM with LVM volumes on an Oracle VM Server, the BMR job fails to create the disk layout.

#### Solution

To resolve this issue, disable the PV drivers for the backup source node.

#### Follow these steps:

1. Open the Command Prompt window on the backup source node and enter the following command:

```
sfdisk -s
```

2. Verify if the same disk is displayed twice in the result.

For example, xvdX and hdX are the same disk. Verify if both these disks are shown in the result.

3. If yes, then perform the following steps:
  - a. Add the following line to the `/etc/modprobe.d/blacklist` file on the backup source node:

```
blacklist xen_vbd
```

- b. Restart the backup source node and rerun the backup job.

The backup job runs.

4. If no, then contact the Arcserve Support team.

The issue is resolved.

## Backup Read Throughput is Less When the Backup Node is RHEL, CentOS, or Oracle Linux 5.x On a Xen-based PVM

### Valid on Red Hat Enterprise Linux (RHEL), CentOS, SUSE Linux Enterprise Server (SLES) 11.x, and Oracle Linux Server 6.x on a Xen-based PVM

#### Symptom

When I back up an RHEL, CentOS, SLES, Oracle Linux 6.x server on a Xen-based PVM, the backup read throughput value is less.

### Solution

To resolve this issue, configure the disk input/output (I/O) scheduler of the VM.

#### Follow these steps:

1. Open the Command Prompt window on the backup source node.
2. Run the following command to identify the number of disks the VM has:  

```
ls /dev/xvd*
```
3. Run the following command to identify the I/O scheduler algorithm of the disks:  

```
cat /sys/block/xvda/queue/scheduler
```

The following result is displayed on the VM:

```
[noop] anticipatory deadline cfq
```
4. Run the following command to change the I/O scheduler to cfq:  

```
echo "cfq" > /sys/block/xvda/queue/scheduler
```
5. Verify the I/O scheduler again.  

The following result is displayed on the VM:

```
noop anticipatory deadline [cfq]
```
6. Change the I/O scheduler for every disk.
7. Rerun the backup job.

The backup job runs with an optimum backup read throughput.

## How to Restore the Previous Version of Backup Server

### Valid on Red Hat Enterprise Linux (RHEL) 6.x and CentOS 6.x for Backup Server

#### Symptom

I tried to upgrade the Backup Server but there was an error during the upgrade. The Backup Server is not working as expected. Now I want to restore the previous version of the Backup Server.

#### Solution

When you upgrade to a new release, the Backup Server creates a backup folder that contains all the old configuration files and database files from the previously installed version. The folder is located at the following location:

```
/opt/CA/d2dserver.bak
```

**Follow these steps:**

1. Uninstall the existing Backup Server using the following command:

```
/opt/CA/d2dserver/bin/d2duninstall
```

2. Install the previously installed version of the Backup Server.

3. Stop the Backup Server using the following command:

```
/opt/CA/d2dserver/bin/d2dserver stop
```

4. Copy the old configuration files and database files to the d2dserver folder using the following command:

```
cp -Rpf /opt/CA/d2dserver.bak/* /opt/CA/d2dserver/
```

5. Start the Backup Server using the following command:

```
/opt/CA/d2dserver/bin/d2dserver start
```

The previously installed version of the Backup Server is successfully restored.

## SLES 10.X Does Not Start Successfully After a BMR

### Valid on SUSE Linux Enterprise Server (SLES) 10.X for BMR on old target machines

**Symptom:**

When I perform a BMR using SLES 10.x recovery points to an old target machine, the BMR is successful but the target machine does not start successfully. Also, when I have recovery points from an SLES 10.x old source machine and I try to perform a BMR, the BMR is successful but the target machine does not start successfully.

In both the cases, I get the following error message:

```
No operating system
```

**Solution:**

Modify the boot disk MBR in a Live CD environment and restart the target machine.

**Follow these steps:**

1. Log into the target machine using a Live CD and locate the boot disk

**Example:** /dev/sda

2. Run the following command:

```
echo -en "\x90\x90"|dd of=/dev/sda seek=156 bs=1
```

3. Restart the target machine and verify if the target machine restarts successfully.

The SLES 10.x target machine successfully starts after a BMR.

## The d2drestorevm and d2dverify Jobs Fail on Oracle VM Server

### Valid on Oracle VM Server

#### Symptom

When I start d2drestorevm and d2dverify jobs on an Oracle VM Server, all jobs fail. I get the following error message in the Activity Log:

Failed to import the ISO image to the hypervisor. Check the hypervisor management console or the debug log for more detailed information.

#### Solution

Verify if the Oracle VM Server has hung.

#### Follow these steps:

1. Log into the Oracle VM Server console and navigate to the Jobs tab.
2. Find all jobs which are in progress status, then abort these jobs.
3. Start the d2drestorevm or d2dverify job again.

If the d2drestorevm or d2dverify job fails again and displays the same error message, log into the Oracle VM Server console and verify if there are any jobs that display the status as In Progress. If there are jobs that display the In Progress status, restart that Oracle VM Server.

The d2drestorevm and d2dverify jobs run successfully.

## ESXi Virtual Machine Fails to Start After BMR From a Physical Machine

#### Symptom

I perform a BMR using the recovery points of a physical machine to an ESXi virtual machine. The physical machine uses an older BIOS. The BMR is successful but the ESXi VM does not start successfully.

#### Solution

Modify the SCSI Controller Type of the target ESXi VM, and submit the BMR job again.

#### Follow these steps:

1. Log in to the ESX Server.
2. Right-click the target ESXi VM, and select Edit Settings.
3. From the Hardware tab, select SCSI controller 0, and click the Change Type button.

The Change SCSI Controller Type dialog opens.

4. Select LSI Logic SAS and save the settings.
5. Submit a BMR job to this VM.

The virtual machine starts successfully after the BMR job.