

Arcserve® Unified Data Protection Agent for Linux

User Guide

Version 6.5

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.

© 2017 Arcserve, including its affiliates and subsidiaries. All rights reserved. Any third party trademarks or copyrights are the property of their respective owners.

Contents

Chapter 1: Understanding Arcserve UDP Agent (Linux)	11
Introduction	12
Chapter 2: Installing/Uninstalling Arcserve UDP Agent (Linux)	14
How to Install Arcserve UDP Agent (Linux)	15
Installation Considerations	16
Install Arcserve UDP Agent (Linux)	17
Install Arcserve UDP Agent (Linux) in AWS Cloud	20
Verify the Installation	23
How to Uninstall Arcserve UDP Agent (Linux)	24
Review the Uninstallation Considerations	25
Uninstall Arcserve UDP Agent (Linux)	26
Verify the Uninstallation	27
How to Upgrade Arcserve UDP Agent (Linux)	28
Upgrade Considerations	29
Upgrade Arcserve UDP Agent (Linux)	30
Verify the Upgrade	32
How to Migrate 32-bit Linux Backup Server to a 64-bit Server	33
Chapter 3: User Interface	35
How to Navigate the Arcserve UDP Agent (Linux) User Interface	36
Access the Backup Server	38
Understanding the Menu Bar	39
Understanding the Status Pane	43
Understanding the Backup Servers Pane	47
Understanding the Help	48
Register Arcserve UDP	50
Chapter 4: Using Arcserve UDP Agent (Linux)	52
How to Manage the Licenses	53
Access the License Manager	54
Understanding the License Management Dialog	55
Manage the Licenses	57
How to Manage Jobs	58
Review the Prerequisites to Manage Jobs	59
Modify Jobs	60
Cancel Jobs	61

Delete Jobs	62
How to Back up Linux Nodes	63
Review the Backup Prerequisites and Considerations	65
Do You Want To Back Up More Than 200 Nodes	70
Add Linux Nodes for Backup	75
(Optional) Enroll Arcserve UDP Public Key for Secure Boot	77
(Optional) Prepare the iSCSI Volume as the Backup Storage	79
Configure the Backup Settings and Run Backup Job	81
Verify the Backup was Successful	104
How to Modify and Rerun a Backup Job	105
Review the Prerequisites for Modifying a Backup Job	106
Do You Want to Add Nodes to an Existing Job	107
Add Nodes to an Existing Job	108
Rerun an Existing Backup Job	109
Verify the Backup was Successful	111
How to Perform a File-Level Recovery for Linux Nodes	112
Review the Prerequisites	113
Specify the Recovery Point for Host-Based Agentless Backup	114
Specify the Recovery Point for Agent-Based Backup	118
Specify the Target Machine Details	122
Specify the Advanced Settings	125
Create and Run the Restore Job	129
Verify that Files are Restored	130
How to Create a Bootable Live CD	131
Review the Live CD Prerequisites	133
Install the Restore-Utility Package	134
Create and Verify the Bootable Live CD	135
How to Use Live CD as a Linux Backup Server	136
How to Create a CentOS-Based Live CD	137
Review the Live CD Prerequisites and Considerations	139
Install the Restore-Utility Package	141
Create and Verify the CentOS-based Live CD	142
How to Perform a Bare Metal Recovery (BMR) for Linux Machines	144
Create a Configuration template Using Command Line	146
Review the BMR Prerequisites	150
Get the IP Address of the Target Machine Using the Live CD	151

(Optional) Recover Data to the iSCSI Volume of the Target Machine	153
(Optional) Recover Data from the iSCSI Volume to the Target Machine	155
Review the Backup Server	157
Specify the Recovery Points	158
Specify the Target Machine Details	160
Specify the Advanced Settings	162
Create and Run the Restore Job	167
Verify that the Target Node is Restored	175
How to Perform a Bare Metal Recovery (BMR) for Linux Machines in AWS Cloud ...	176
Review the BMR Prerequisites	177
Launch an Instance Using the Arcserve UDP Agent Live CD	178
Review the Backup Server Instance	180
Specify the Recovery Points	181
Specify the Target Instance Details	183
Specify the Advanced Settings	185
Create and Run the Restore Job	189
Verify that the Target Instance is Restored	197
How to Perform a Bare Metal Recovery (BMR) for Linux Machines in Azure Cloud ...	198
Review the BMR Prerequisites	199
Create a new Machine on Microsoft Azure as BMR target	200
Review the Backup Server Virtual Machine	201
Specify the Recovery Points	202
Specify the Target Virtual Machine Details	203
Specify the Advanced Settings	205
Create and Run the Restore Job	206
Verify that the Target Virtual Machine is Restored	207
How to Perform a Migration BMR for Linux Machines	208
Review the Prerequisites for Migration BMR	209
Perform a BMR to the Temporary Machine	210
Perform a Migration BMR	212
Verify that the Target Node is Restored	213
How to Perform a Migration BMR for Linux Machines from Amazon EC2 to local ...	214
Review the Prerequisites for Migration BMR	215
Perform a BMR Migration from Amazon EC2 to the local Machine	216
Verify that the Target Node is Restored	218
How to Automatically Recover a Virtual Machine	219

Review the Prerequisites and Considerations	220
Create a Configuration Template	222
(Optional) Create a Global Configuration File	227
Modify the Configuration Template and File	229
Submit a Job Using the d2drestorevm Utility	230
Verify the VM is Recovered	231
How to Integrate and Automate Arcserve D2D for Linux with the Existing IT Environment	232
Review the Automation Prerequisites	234
Understanding the Scripting Utilities	235
Manage Pre/Post Scripts for Automation	245
Create the Backup Storage Alert Script	252
Discover Nodes Using a Script	253
Create the Scripts to Back Up Oracle Database	254
Create the Scripts to Back Up MySQL Database	256
Customize the Job Schedule	258
Run a BMR Batch Job	260
Replicate and Manage Backup Sessions	262
Verify the Recovery Points are Usable	265
How to Manage the Backup Server Settings	271
Review the Prerequisites to Manage the Backup Server	272
Configure the Job History and Activity Log Retention Settings	273
Configure the Debug Log Retention Settings	274
Configure the UI Timeout Duration	275
Change the SSH Port Number of the Backup Server	276
Manage the Recovery Sets	277
Disable the BOOTPD and TFTP Services	278
Improve the Query Performance for Job History and Activity Log	279
Skip CIFS and NFS Module Verification	280
Skip CIFS and NFS Validation on Linux Backup Server	281
Configure the Default Temporary Folder	282
Configure the Snapshot Path for Backup Node	283
Configure the Hyper-V Server Connection Information for Instant VM	284
How to Manage the Linux Backup Server from the Command Line	286
Review the Backup Server Prerequisites	288
Start, Stop, or Release the Backup Server	289
Change the Web Service Port Number of the Backup Server	291

Configure the Private Key and Public Key Authentication	292
Change the Backup Server Protocol	294
Avoid the SSL Certificate Error While Opening Arcserve UDP Agent (Linux)	295
Configure the System Settings When the Host Name or IP Address is Changed	297
How to Add a User to Linux Backup Server Console Using Command Line	302
Review the Prerequisites	303
Add a User to Linux Backup Server Console Using Command Line	304
How to Manage the Non-Root Users	306
Review the Prerequisites	307
Grant Login Permissions to the Non-Root Users	308
Display the Default User in the Login Dialog	309
Enable the Non-Root Users to Add Nodes	310
How to Configure Sudo User Account for Linux Nodes	312
Review the Prerequisites	313
Modify the Default Sudo Settings in SUSE	314
Configure sudo in Debian	315
Configure Sudo for Authorization Without Password When Using SSH Public Key Authentication	316
Configure Sudo to Allow Only Backup Agent Process	317
How to Restore Volumes on a Target Node	318
Review the Prerequisites and Considerations	319
Verify the d2drestorevol Utility is Installed	320
Verify the Volume Details in the Session	321
Submit the Volume Restore Job	324
Cancel the Volume Restore Job	328
Verify the Restored Volume	329
How to Restore an Oracle Database Using Arcserve UDP Agent (Linux)	330
Perform a Bare Metal Recovery (BMR) of an Oracle Server	331
Perform an Instant Recovery of an Oracle Database	335
Perform Granular Recovery of an Oracle Database	339
How to Run Assured Recovery Test from the Command Line	345
Review the Prerequisites and Considerations	346
Create a Configuration Template	347
Modify the Configuration Template and File	352
Submit a Job Using the d2dar Utility	353
How to Mount Recovery Point	354
Review the Prerequisites	355

Specify the Recovery Point for Mount Recovery Point	356
Specify the Settings for Mount Recovery Point	359
Create and Run the Mount Recovery Point Job	361
Mount NFS or WebDAV Share on Linux Server	362
How to Enable Support for the Latest Ubuntu Kernel	365
Review the Prerequisites	366
Deploy the Updated Ubuntu Driver Package Manually	367
(Optional) Using Staging Server for Updating Drivers	368
(Optional) Using Staging Server for Updating Drivers	369
Chapter 5: Troubleshooting	370
Arcserve UDP Agent (Linux) Fails to Install on Supported Servers	371
Arcserve UDP Agent (Linux) Displays an Operation Timeout Error	373
All Scheduled Jobs Fail When the System Time Is Changed to an Already Passed Value	374
Arcserve UDP Agent (Linux) Fails to Mount Linux Software RAID Devices	375
Arcserve UDP Agent (Linux) Fails to Download and Deploy Updated Ubuntu Drivers on SLES 11 and RHEL 6	376
A Paravirtual Machine (PVM) Displays a Black Screen on the Virtual Network Computing (VNC) Client Window When Booted Using a Live CD	377
The Backup Job Fails to Collect the BMR-related Information or the BMR Job Fails to Create a Disk Layout	378
How to Adjust the Disk Boot Sequence After a BMR Job on an Oracle VM Server ...	379
How to Restore the Previous Version of Backup Server	381
How to Start the Linux Backup Server Automatically on SLES15	382
How to Backup Debian 9.X EC2 Instances in AWS Cloud	383
SLES 10.X Does Not Start Successfully After a BMR	384
The d2drestorevm and d2dverify Jobs Fail on Oracle VM Server	385
ESXi Virtual Machine Fails to Start After BMR From a Physical Machine	386
Failed to Mount CIFS on the Server or Target Node	387
File-level restore in a host-based Linux VM fail due to an unsupported file system ..	388

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.

[Contact 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.

Chapter 1: Understanding Arcserve UDP Agent (Linux)

This section contains the following topics:

Introduction	12
------------------------------------	----

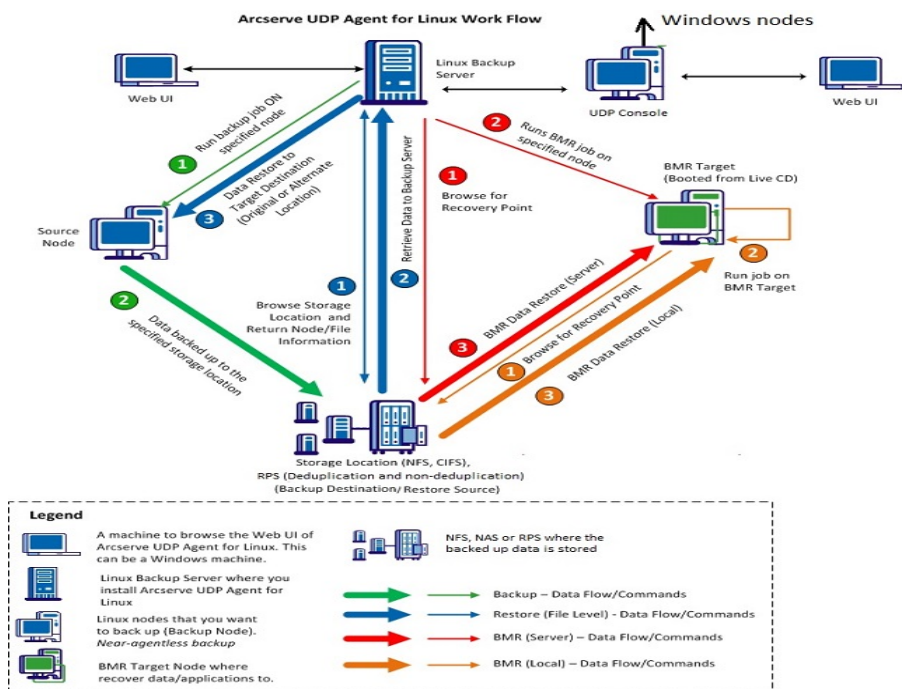
Introduction

Arcserve D2D 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 while performing a backup of data, Arcserve UDP Agent (Linux) 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):



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

This section contains the following topics:

How to Install Arcserve UDP Agent (Linux)	15
How to Uninstall Arcserve UDP Agent (Linux)	24
How to Upgrade Arcserve UDP Agent (Linux)	28
How to Migrate 32-bit Linux Backup Server to a 64-bit Server	33

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.

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

- [Installation Considerations](#)
- [Install Arcserve UDP Agent \(Linux\)](#)
- [Install Arcserve UDP Agent \(Linux\) In AWS Cloud](#)
- [Verify the Installation](#)

Installation Considerations

Consider the following points before you begin the installation:

- When you perform a Preboot Execution Environment (PXE)-based BMR, the Arcserve D2D 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.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.
- The unattended or silent installation is not supported.

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.

Notes:

- Ensure that the Linux server has minimum 2 GB memory. For more information about the system requirements for a Linux server, see the *Arcserve UDP Release Notes v6.5*.
- Use sudo to install Linux Server on Microsoft Azure.
- For Debian/Ubuntu system, the root is not permitted to log into ssh by default. To grant non-root user the permission to log into Linux Backup server UI, see [Grant Login Permissions to non-root user](#).

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) to the root folder.

Important! When you download the installation package file 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. Run the following command to start installation:

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

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.

Notes:

- ◆ If a non-English operating system is detected, you are prompted to select the applicable language before continuing with the installation process.
 - ◆ To support Korean language when you upgrade a build, perform the following steps:
 - a. Modify the following configuration file in the Arcserve UDP Agent (Linux) server: `/opt/Arcserve/d2dserver/nls/nls.cfg`
 - b. set `D2D_LANG= ko_KR`.
 - c. Restart the d2dserver using the following command: `#!/opt/Arcserve/d2dserver/bin/d2dserver restart`.
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/Arcserve/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)
- 8016 (Instant BMR service)
- 8021 (Backup service)

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 BMR target nodes.

Note: For details about ports, view [Communication Ports Used by Arcserve UDP](#).

7. (Optional) To install Linux backup server to a VM on Amazon EC2 or Azure, perform the following steps to create a D2D user:

Note: When the server starts, a message prompt asks you to create a D2D user that is used to log in at Arcserve UDP Agent (Linux) web UI.

- a. Enter a user name to create.
- b. Set the password and confirm by entering again.
- c. Select if you want the user account as default login user for Arcserve UDP Agent (Linux) web UI
Default: Y (yes)
- d. Decide how many consecutive login failures can get the user account locked.
Default: 3

Arcserve UDP Agent (Linux) is successfully installed.

Install Arcserve UDP Agent (Linux) in AWS Cloud

Compared to traditional installation in a Linux machine, you can launch an Arcserve UDP Agent (Linux) instance directly using Amazon Machine Image (AMI) in AWS Cloud. After the Arcserve UDP Agent (Linux) instance is launched, you can open the user interface from any computer using a web browser and the server is referred as Backup Server.

Follow these steps:

1. Log into the EC2 management console with your account and select Launch Instance.

The Launch Instance Wizard appears displaying 7 tabs.

2. From the first tab **Choose AMI**, select the Arcserve UDP Agent (Linux) AMI in **Community AMIs** for **Step 1: Choose an Amazon Machine Image (AMI)** and click **Next: Choose an Instance Type**.

You can search the Arcserve UDP Agent (Linux) AMI using *Arcserve_Unified_Data_Protection_Agent_Linux* in Community AMIs.

Note: Select an Arcserve UDP Agent (Linux) AMI with the latest version to launch the instance.

The second tab **Choose Instance Type** appears.

3. Select an Instance Type based upon your requirement to complete **Step 2: Choose an Instance Type** and click **Next: Configure Instance Details**.

Note: Verify that the Instance type is at least t2.medium and has a minimum 4 GB memory. For more information about the system requirements for a Linux server, view *Arcserve UDP v6.5 Release Notes*.

The third tab **Configure Instance** appears.

4. Select details for fields such as Network, Subnet, Auto-assign Public IP or not, and others to complete **Step 3: Configure Instance details**, and then click **Next: Add Storage**.

The fourth tab **Add Storage** appears.

5. Allocate storage for the Instance to complete **Step 4: Add Storage** and click **Next: Add Tags**.

Note: You can adjust the disk size based on your business requirement. Verify that the Linux Instance disk has a minimum 40 GB size.

The fifth tab **Add tags** appears.

5. Enter tags for the AMI target Instance to complete **Step 5: Add tags** and click **Next: Configure Security Group**.

The sixth tab **Configure Security Groups** appears.

6. Perform the following steps to assign security groups for the AMI target instance to complete **Step 6: Configure the security group**, and then click **Review and Launch**:

Follow these steps:

- a. Create a new security group for SSH and Arcserve UDP Agent (Linux).
- b. Verify that Port 22 is enabled for **Type SSH** and configure the **Source** to *Anywhere*.
- c. Verify that Port 8014 used by tomcat is enabled for the *Custom TCP Rule Type* and configure the **Source** to *Anywhere*.
- d. Verify that Port 8016 used by d2dds and 8021 used by cresvc are enabled for the Custom TCP Rule **Type** and configure the **Source** of the rule to Custom.

Note: You can specify the custom source with CIDR format to let d2dds and cresvc serve the Linux instances that are in the same subnet with Arcserve UDP Agent (Linux) but not accessible by other internet machines. For example, if the subnet CIDR is 102.31.16.0/20, you can also specify the source to 102.31.16.0/20.

The seventh tab **Review** appears.

7. Verify details by selecting or creating a key pair to connect your instance to complete **Step 7: Review Instance Launch**, and then click **Launch Instance**.
8. From the launched Arcserve UDP Agent (Linux) instance, set a new password for `udpuser` as below:

```
#sudo /opt/Arcserve/d2dserver/bin/d2duser --action=passwd --username=udpuser
```

Note: The default username of Arcserve UDP Agent (Linux) management UI is `udpuser`.

9. (optional) If you want to switch to other language, you can modify the configuration file in the Arcserve UDP Agent (Linux) server:

```
/opt/Arcserve/d2dserver/nls/nls.cfg
```

And then set `D2D_LANG=$OTHER_LANGUAGE` and restart the `d2dserver` with the below command:

```
#!/opt/Arcserve/d2dserver/bin/d2dserver restart
```

Note: English is the default language of Arcserve UDP Agent (Linux).

Now, Arcserve UDP Agent (Linux) is ready to use in AWS Cloud and the URL to browse the Linux Backup Server is [https://\\$INSTANCE_IP:8014](https://$INSTANCE_IP:8014).

Arcserve UDP Agent (Linux) is successfully installed in AWS Cloud.

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.

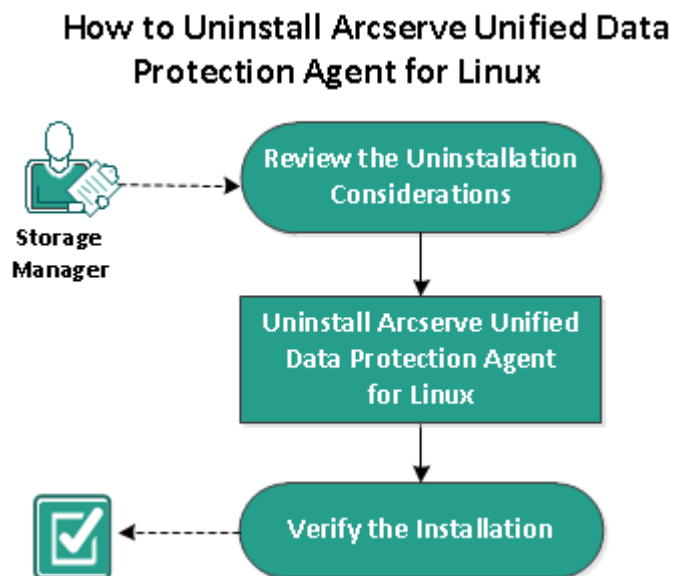
The Arcserve UDP Agent (Linux) user interface opens.

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](#)
- [Uninstall Arcserve UDP Agent \(Linux\)](#)
- [Verify the Uninstallation](#)

Review the Uninstallation Considerations

Consider the following points before you begin the uninstallation:

- You have the root login credentials to the Backup Server.
- Review the [Compatibility Matrix](#) that 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 D2D for Linux is installed with the following command:

```
# cd /opt/Arcserve/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:

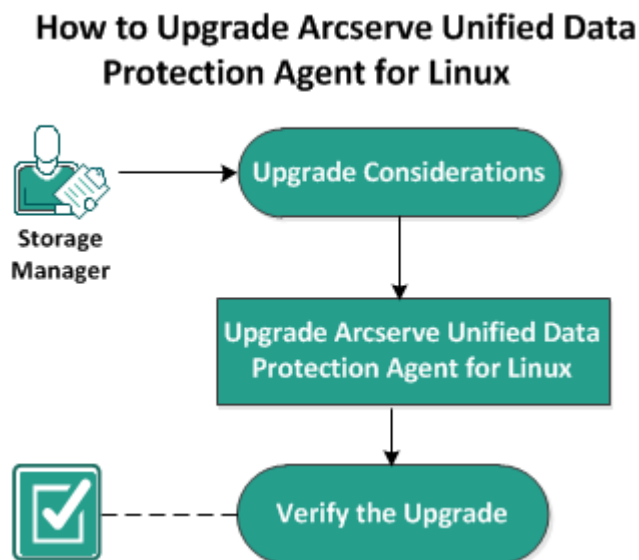
```
/opt/Arcserve/d2dserver
```

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](#)
- [Upgrade Arcserve UDP Agent \(Linux\)](#)
- [Verify the Upgrade](#)

Upgrade Considerations

Consider the following points before you begin the upgrade:

- Ensure that you schedule your upgrade when the backup jobs are not running.
- Review the [Compatibility Matrix](#) that 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 (for example, 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) to the root folder.

Important! When you download the installation package file 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. Run the following command to start installation:

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

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.
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/Arcserve/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/Arcserve/d2dserver.bak`

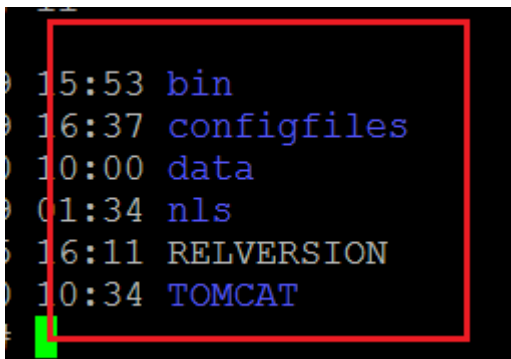
Arcserve UDP Agent (Linux) is successfully upgraded and verified.

How to Migrate 32-bit Linux Backup Server to a 64-bit Server

From Version 6 onwards, Arcserve UDP Agent (Linux) will not support a 32-bit server for Linux backup Server. To use Arcserve UDP Agent (Linux) Version 6, migrate the 32-bit Linux server to a 64-bit Linux server.

Follow these steps:

1. Reserve the following files and folders in the Arcserve UDP Agent (Linux) installation folder:



A typical installation folder for Arcserve UDP Agent (Linux) Version 5 was '/opt/CA/d2dserver/'

Note: If the TOMCAT folder is a large folder, reserve only TOMCAT/conf folder.

2. Copy the reserved files and folders to some other location, such as '/opt/d2dserver_32bit/'.
3. Package the reserved files and folders at the following location:

```
tar -czf UDP_LINUX_AGENT.tar.gz /ultraconservative
```
4. Copy the packaged file from the 32-bit Linux OS to the 64-bit Linux OS using scp or ftp.
5. Create a folder on the 64-bit OS server using the following command:

```
mkdir -p /opt/CA/d2dserver
```
6. Extract the packaged file on the 64-bit Linux OS using the following command:

```
tar -xzf UDP_LINUX_AGENT.tar.gz
```
7. Copy the reserved files and folders to the following location:

`/opt/CA/d2dserver`

For example: `cp -Rp /opt/d2dserver_32bit/* /opt/CA/d2dserver`

8. Run the Arcserve UDP Agent (Linux) Version 6.0 installation package on the 64-bit Linux server.
9. The Linux Backup Server upgrades automatically.

Note: If the host name or IP address is changed, see [Configure the System Settings When the Host Name or IP Address is Changed](#).

Chapter 3: User Interface

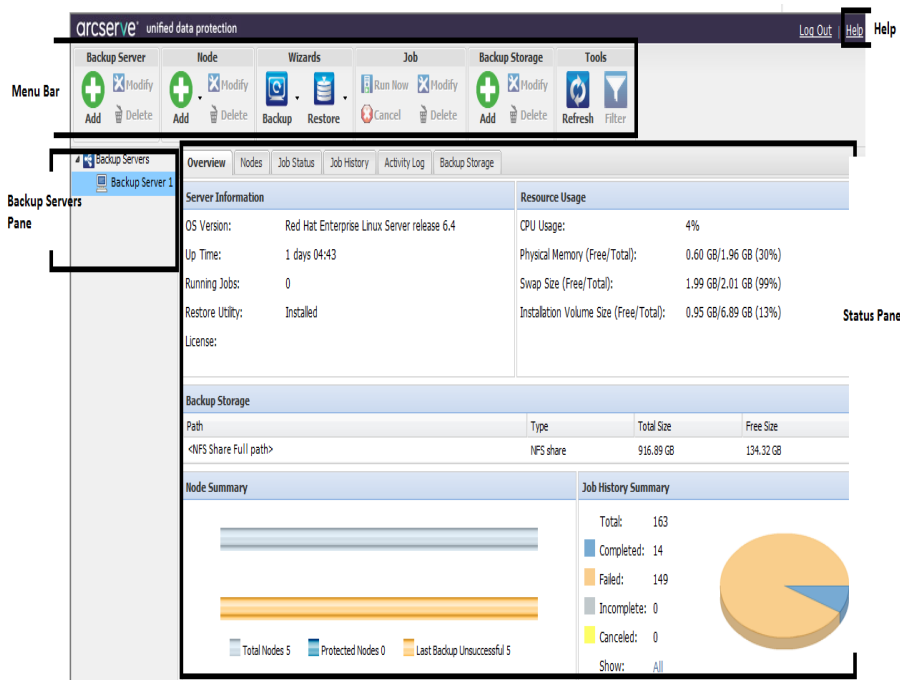
This section contains the following topics:

How to Navigate the Arcserve UDP Agent (Linux) User Interface	36
Register Arcserve UDP	50

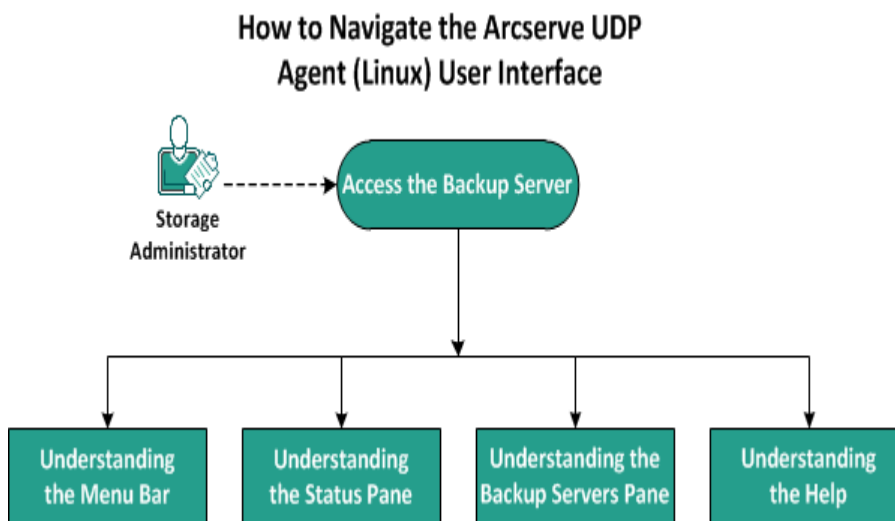
How to Navigate the Arcserve UDP Agent (Linux) User Interface

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 Home page 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](#)
- [Understanding the Menu Bar](#)
- [Understanding the Status Pane](#)
- [Understanding the Backup Servers Pane](#)
- [Understanding the Help](#)

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.

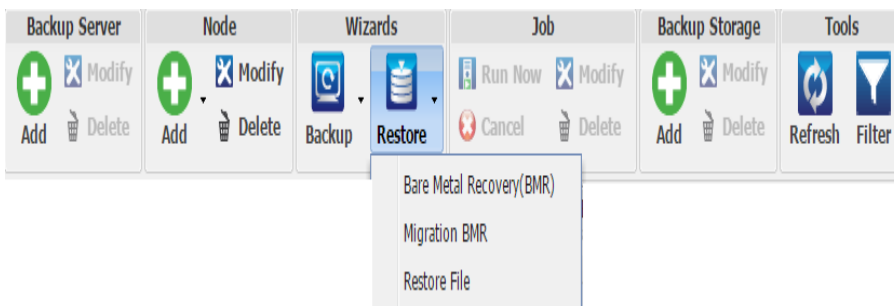
Back Up Selected Nodes

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 three 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.

Migration BMR

Use this option to perform a migration BMR.

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, Local, or RPS server. Local is a local path in the Backup Server. RPS server is Recovery Point Server. RPS is installed when you install Arcserve UDP. In RPS, you create data stores where you store the recovery points. When you add an RPS server, you also have to specify the data store.

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.

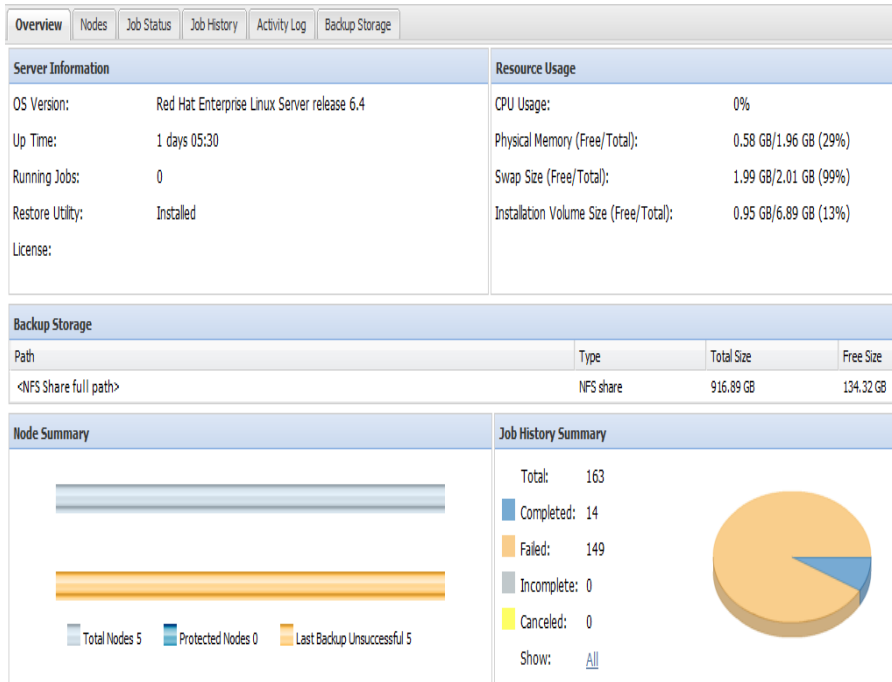
The following screen displays the filters applied to the Activity Log:



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 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 drop-down list:



The following options are available in the Help drop-down list:

Knowledge Center

Lets you access the bookshelf.

Online Support

Lets you access the Arcserve Support website.

Solutions Guide

Lets you access the HTML version of the Arcserve UDP Agent Solutions Guide.

Agent for Linux User Guide

Lets you access the HTML version of the user guide.

Ask Support: Live Chat

Lets you open a chat window and contact an Arcserve Support executive for a live chat.

Provide Feedback

Lets you access the Arcserve Support website and provide feedback to the development team.

Videos

Lets you access online tutorials and videos related to Arcserve UDP Agent (Linux).

Manage Licenses

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

Product Improvement Program

Lets you provide suggestions to improve Arcserve product.

About

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

Register Arcserve UDP

After you install Arcserve UDP, you must register the product from the Console. This registration allows Arcserve to automatically collect the usage details and statistics of the Console.

Important! Arcserve does not collect any personal or business critical information such as node name, IP address, login credentials, domain name, and network names.

If you have not registered the Console, you will receive the following notification in the **Messages** tab of the Console.

Your copy of Arcserve Unified Data Protection has not been registered in the Arcserve Product Improvement Program.
Register.

Follow these steps:

1. From the Console, click **Help, Product Improvement Program**.

The **Arcserve Product Improvement Program** dialog opens.

2. Select the **Participate in Arcserve's Product Improvement Program** check box.
3. Specify the following details:

Name

Specify your name.

Company

Specify the name of your company.

Phone Number

Specify your phone number in the following format:

Country code - Phone number. For example: 000-1122334455

Email Address

Specify your email address. This is a mandatory field. The verification email will be sent to this email address.

Fulfillment Number

Specify the fulfillment number. You must have received this number in an email when you downloaded Arcserve UDP.

4. Click **Send Verification Email**.

The verification email is sent to the email address that you have mentioned on the **Arcserve Product Improvement Program** dialog.

5. Log into the email account and open the received email.

6. Click the verification link provided in the email.

You have successfully registered Arcserve UDP.

After you have registered, the **Cancel Participation** button is activated.

To cancel your registration, click **Cancel Participation**.

If you want to update the email address, you have to register again. To register again, follow the same process as described in this topic.

Chapter 4: Using Arcserve UDP Agent (Linux)

This section contains the following topics:

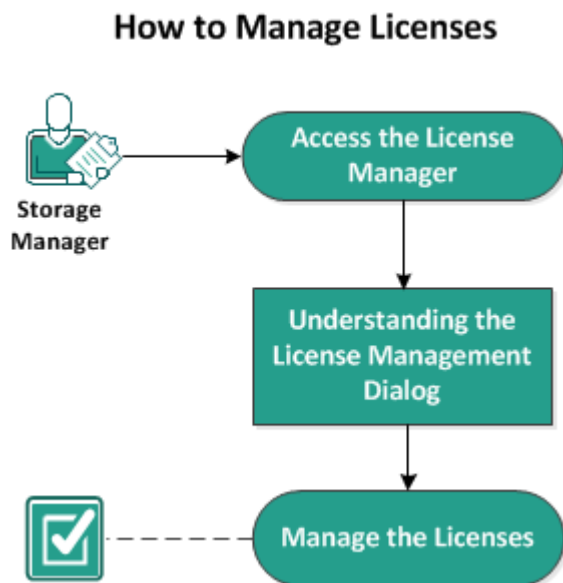
How to Manage the Licenses	53
How to Manage Jobs	58
How to Back up Linux Nodes	63
How to Modify and Rerun a Backup Job	105
How to Perform a File-Level Recovery for Linux Nodes	112
How to Create a Bootable Live CD	131
How to Create a CentOS-Based Live CD	137
How to Perform a Bare Metal Recovery (BMR) for Linux Machines	144
How to Perform a Bare Metal Recovery (BMR) for Linux Machines in AWS Cloud	176
How to Perform a Bare Metal Recovery (BMR) for Linux Machines in Azure Cloud	198
How to Perform a Migration BMR for Linux Machines	208
How to Perform a Migration BMR for Linux Machines from Amazon EC2 to local	214
How to Automatically Recover a Virtual Machine	219
How to Integrate and Automate Arcserve D2D for Linux with the Existing IT Environment	232
How to Manage the Backup Server Settings	271
How to Manage the Linux Backup Server from the Command Line	286
How to Add a User to Linux Backup Server Console Using Command Line	302
How to Manage the Non-Root Users	306
How to Configure Sudo User Account for Linux Nodes	312
How to Restore Volumes on a Target Node	318
How to Restore an Oracle Database Using Arcserve UDP Agent (Linux)	330
How to Run Assured Recovery Test from the Command Line	345
How to Mount Recovery Point	354
How to Enable Support for the Latest Ubuntu Kernel	365

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 D2D 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](#)
- [Understanding the License Management Dialog](#)
- [Manage the Licenses](#)

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.

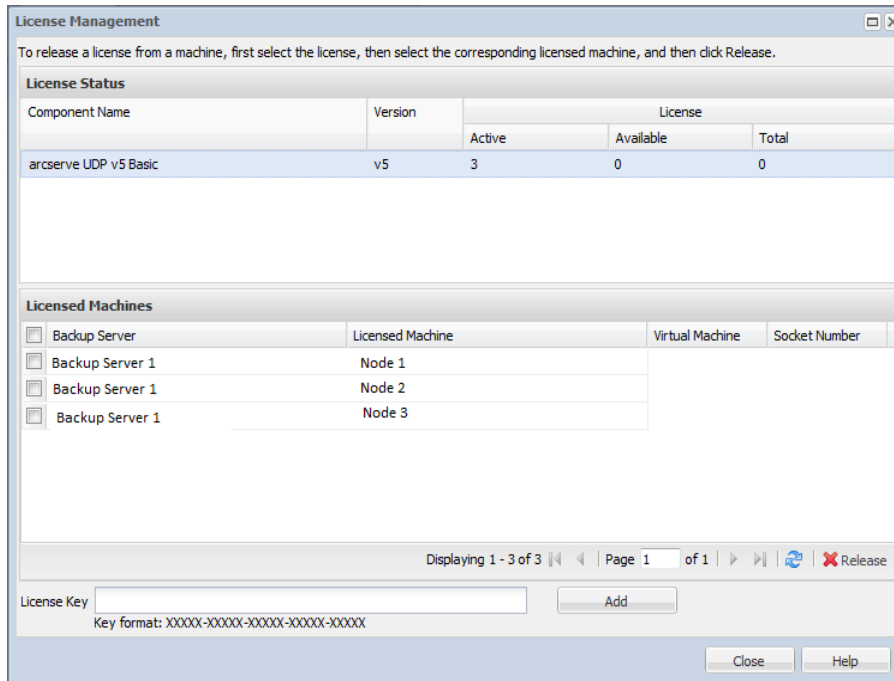
The License Management dialog opens.

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 back up a machine, you can release the license from that machine.

To add a license, follow these steps:

- a. Using Arcserve License Portal, generate the license key. For details, see [How to Generate Arcserve License Keys for Standalone Agents](#).
- b. Enter the license key in the License Key field of the License Management dialog and click Add.
- c. Close and open the License Management dialog.

The license is added and listed in the License Status area.

To release a license, follow these steps:

- a. Select the license from the License Status area of the License Management dialog.
- 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](#)
- [Modify Jobs](#)
- [Cancel Jobs](#)
- [Delete Jobs](#)

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](#) that 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.

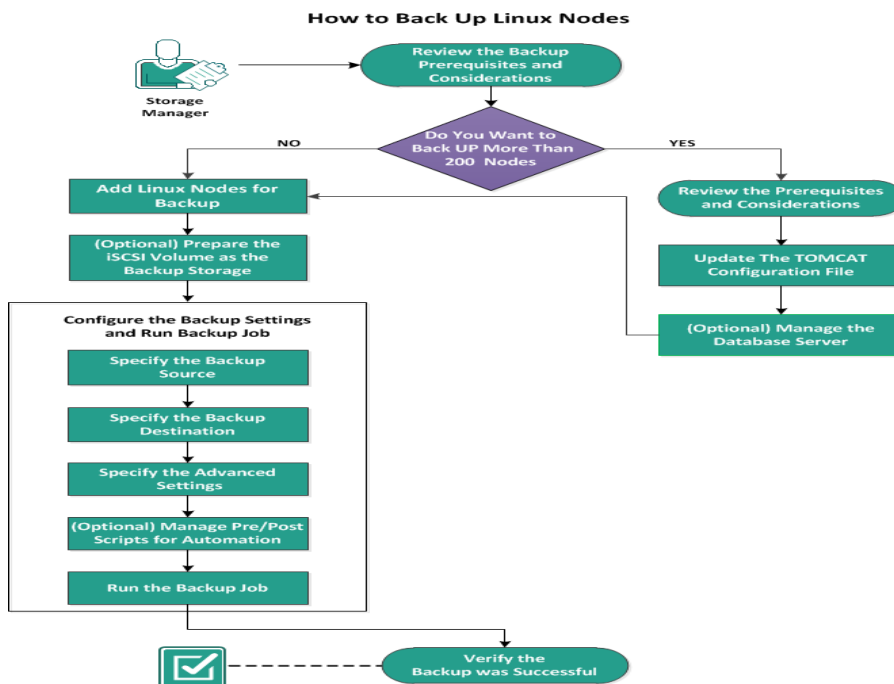
How to Back up Linux Nodes

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.

Note: If you restart the backup source node, the next backup is converted to a Verify Backup (for non-deduplication backup) or Full Backup (for deduplication backup).

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](#)
- [Do You Want to Back Up More Than 200 Nodes](#)
 - ◆ [Review the Prerequisites and Considerations](#)
 - ◆ [Update the TOMCAT Configuration File](#)
 - ◆ [Manage the Database Server](#)

- [Add Linux Nodes for Backup](#)
- [\(Optional\) Enroll the Arcserve Public Key for Secure Boot](#)
- [\(Optional\) Prepare the iSCSI Volume as the Backup Storage](#)
- [Configure the Backup Settings and Run Backup Job](#)
 - ◆ [Specify the Backup Source](#)
 - ◆ [Specify the Backup Destination](#)
 - ◆ [Specify the Advanced Settings](#)
 - ◆ [\(Optional\) Manage Pre/Posts Scripts for Automation](#)
 - ◆ [Run the Backup Job](#)
- [Verify the Backup was Successful](#)

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 *Arcserve UDP 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.
- 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.

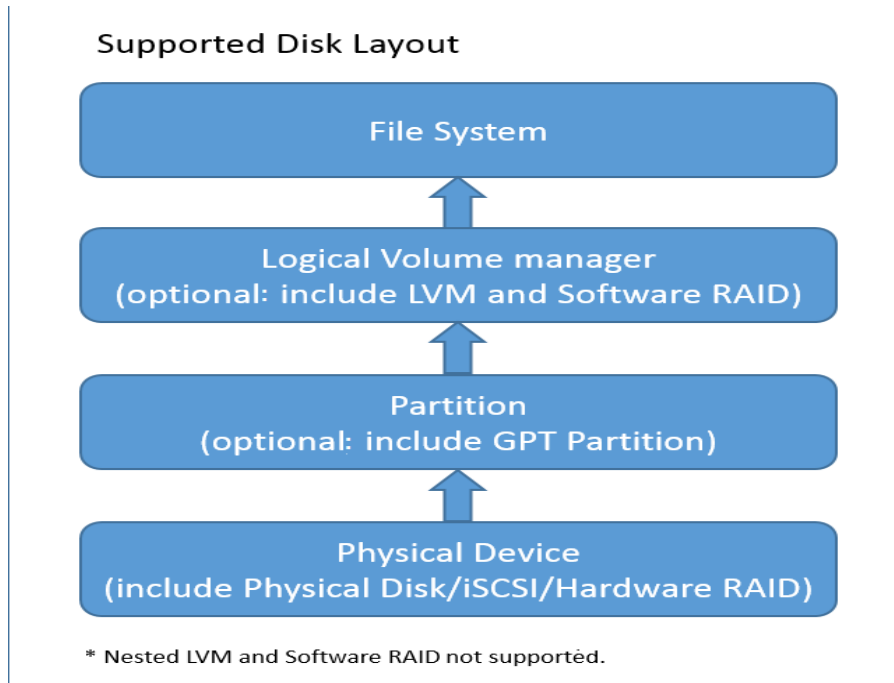
Review the following backup considerations:

- To optimize the management of your Recovery Points, you should consider the following recommendation 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).

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 Layout Supported by Arcserve UDP Agent (Linux)

The following illustration shows the disk layout that is supported by Arcserve UDP Agent (Linux) backup source:



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:		
■ GPT (GUID Partition Table) Data Disk	Yes	Yes
■ GPT (GUI Partition Table) Boot Disk	Yes	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
■ XFS *3	Yes	Yes
■ Btrfs *4	Yes	Yes
Shared Volume:		
■ Windows Shared Volume (CIFS share)	Not Applicable	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).
*2	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.
*3	File level restore for higher version of XFS is not supported on a Linux backup server that has a lower version of XFS. Such as, performing file-level restore for XFS on RHEL7.X is not supported on RHEL6.x as the backup server. But, you can instead use Live CD as a temporary Backup server to perform file-level restore.
*4	File level restore on source machine is not supported (For example, install Linux Backup Server on machine A, back up machine A, and then run restore from A's recovery point on machine A). File/folder filter is not supported. File system balance/scrub process will

	<p>be canceled at the beginning of backup.</p> <p>BTRFS RAID support: RAID-0 and RAID-1.</p> <p>Volume filter UI: Only primary volume is displayed. This is not a limitation, but expected behavior.</p>
--	--

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 d2d-jobhistory 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/Arcserve/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.

```
/opt/Arcserve/d2dserver/TOMCAT/conf/
```

4. 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.

5. Stop the Backup Server.

```
./d2dserver stop
```

6. 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.

```
<Machine Name> :/opt/Arcserve/d2dserver/bin # ./d2dserver stop
The Arcserve UDP Agent(Linux) is stopped.
<Machine Name> :/opt/Arcserve/d2dserver/bin # ./pgmgr init
The installation process has started for Postgresql database. The debug log is placed at the following locati
e/d2dserver/logs/pginit.log.
The Postgresql database has been successfully installed.
Data has been successfully migrated to the new database.
The 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 image shows a dialog box titled "Add Node" with a close button in the top right corner. It contains the following fields and buttons:

- Hostname/IP Address: xxx.xxx.xxx.xxx:123
- User Name: root
- Password: *****
- Description: (empty text area)
- Buttons: Add and More, Add and Close, 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) Enroll Arcserve UDP Public Key for Secure Boot

When running under Secure Boot, the backup source node needs manual installation of Arcserve public key for the backup driver to be trusted. Only when the key is enrolled, the node management and backup feature function properly. This topic describes how to enroll the public key of Arcserve for Secure Boot enabled node.

Prerequisites:

- Verify that you have access to Arcserve public key.
- Verify if your system has the related package of MokManager.efi file that is located at the below folder:

RedHat: /boot/efi/EFI/redhat folde

CentOS: /boot/efi/EFI/centos

Ubuntu: /boot/efi/EFI/ubuntu

SLES: /boot/efi/EFI/SLES12

Follow these steps:

1. Log into the shell environment of the backup source node.
2. Locate Arcserve public key at the following location:
`/tmp/arcserve_public_key_for_secureboot.der`
3. From the document of the running Linux distribution to add the public key to the UEFI MOK list perform the following steps as explained in the below example:

- a. Import the certification to MOK:

```
mokutil [--root-pw] --import
```

```
/tmp/arcserve_public_key_for_secureboot.der
```

The `--root-pw` option enables usage of the root user directly. The root password is required to enroll the key after restarting the system.

- b. Specify a password for the certification when the `--root-pw` option is not available.

This password is required to enroll the key after restarting the system.

- c. Verify the list of certificates that are prepared to be enrolled from mokutil:

```
mokutil --list-new>
```

The list must have Arcserve public key.

- d. Restart the system.

The system launches shim UEFI key management tool.

Note: If the shim UEFI key management tool is not launched, the system may not have the MokManager.efi file.

- e. Enter the password that you specified while importing Arcserve public key to enroll the certification to the MOK list.
- f. Verify if the newly imported key appears enrolled after the system starts up:

```
mokutil --list-enrolled
```

The list must have Arcserve public key.

4. Add or back up the node again to verify the Arcserve public key is successfully enrolled.

The Secure Boot enabled node is ready to be protected by the Arcserve UDP Agent (Linux).

(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 into 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 into 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 below 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 created 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.

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	User name	Status	Volume Filter	Priority
---------------------	-----------	--------	---------------	----------

Volumes to be filtered 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) Enter the volumes in **Volumes to be filtered for all listed nodes**.

Select either Include or Exclude from the drop-down list. Include specifies that only the specified volumes will be included for backup. Any volume that is not specified will not be backed up. Excluded specifies that the volumes will be excluded from the backup.

6. (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.

7. 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.
- ◆ If you have selected Amazon S3, then enter the Backup destination detail in the following format:

```
//S3_Region_ID/S3_bucket_name
```

Notes:

- *///.* could be used as shortcut for Amazon global account. For example, *///./Global_bucket_name*
- *///China/* could be used as shortcut for Amazon cloud China account. For example, *///China/China_bucket_name*

- If you want to export the Amazon S3 bucket as CIFS share, you can click the Enable CIFS client access check box. The default port is 8017.

This feature has the following config file:

/opt/Arcserve/d2dserver/configfiles/ofs.cfg

Do not modify its original content. You can add below content:

- ♦ PROXY_HOST= (If you want to use proxy, enter the proxy name here.)
- ♦ PROXY_USERNAME= (Proxy user name)
- ♦ PROXY_PASSWORD_ENC= (Proxy password, which needs to be encrypted)
- ♦ PROXY_PORT= (Proxy port)
- ♦ WRITE_THROUGHPUT= (If you want to limit the write throughput, Unit: KB/s)
- ♦ HTTPS = yes/no (default is yes)
- ♦ S3_STORAGE_CLASS = STANDARD/STANDARD_IA/REDUCED_REDUNDANCY (default is STANDARD)
- ♦ DEBUG_LEVEL= (Debug log level: 0,1,2,3, 3 will print most log)

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.

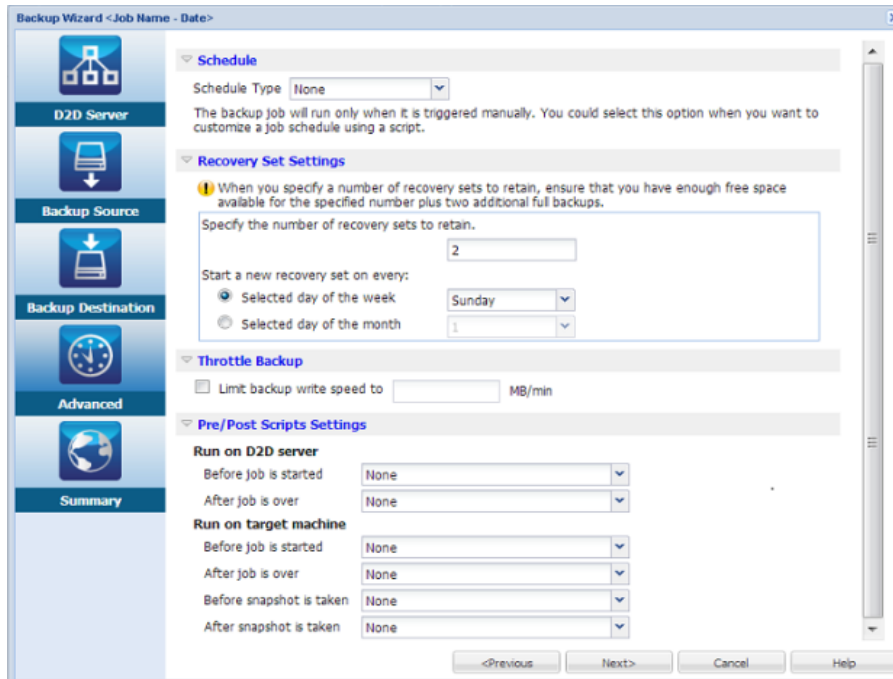
The Advanced page opens.

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.

Important! Set the same timezone between UDP Server and Linux Backup Server. After you change the timezone in both the servers, you must restart the UDP management service or Linux Backup Server for the changes to take effect.
- 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 drop-down list:

Simple

The **Simple** schedule type is not available when you create a new schedule. However, if you are modifying an old backup job that had Simple schedule, you can configure the Simple schedule.

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 Simple

Set start date and time
Specify the scheduled starting date and time for the full, incremental, and verify backups.
Start Date 6/30/13 Start Time 3 : 07 PM

Incremental Backup
Incrementally backs up only the data that has changed since the last successful backup.
 Repeat Every 1 Hours

Full Backup
Backs up all selected data from the machine.
 Repeat Every 1 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 1 Days
 Never

Custom

Select the Custom 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 Custom schedule. When you click Clear, all the Custom backup schedules are deleted from the Custom Schedule Tray.

Schedule

Schedule Type Custom

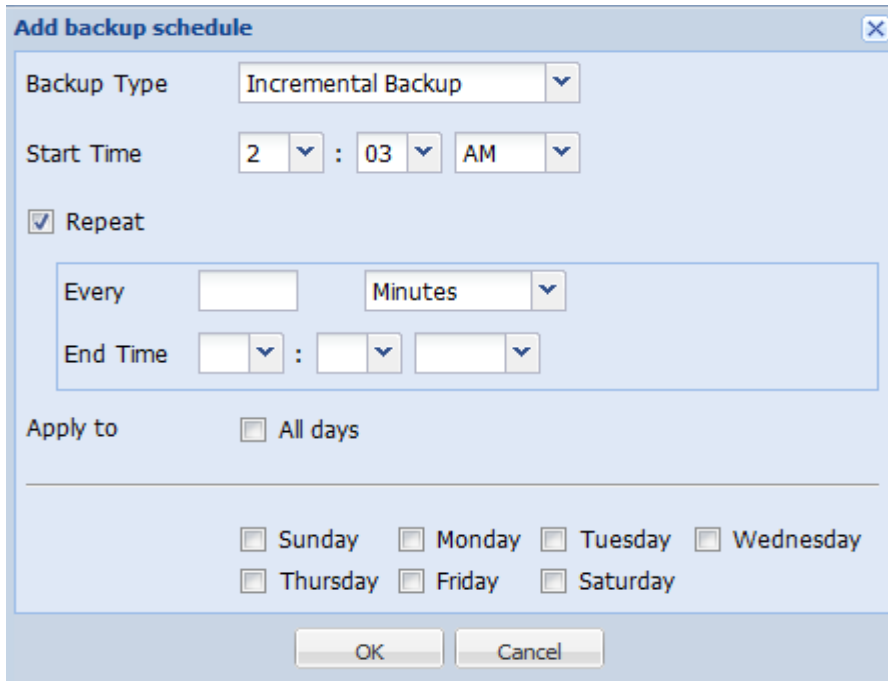
Start Date 1/26/16 + Add Modify Delete Clear

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		

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 Custom 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 to 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/Arcserve/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 custom 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 D2D 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 D2D 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 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-target

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

pre-job-target

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

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.

D2D_RECOVERYPOINT

Identifies the recovery point created by the backup job. This value is applicable in the post backup script only.

D2D_RPSSCHEDULETYPE

Identifies the schedule type when backing up to a Data Store on RPS. The following values identify the D2D_RPSSCHEDULETYPE variable:

daily

Identifies the schedule as a daily backup.

weekly

Identifies the schedule as a weekly backup.

monthly

Identifies the schedule as a monthly backup.

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/Arcserve/d2dserver/usr/prepost
```

Follow these steps:

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

```
/opt/Arcserve/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.

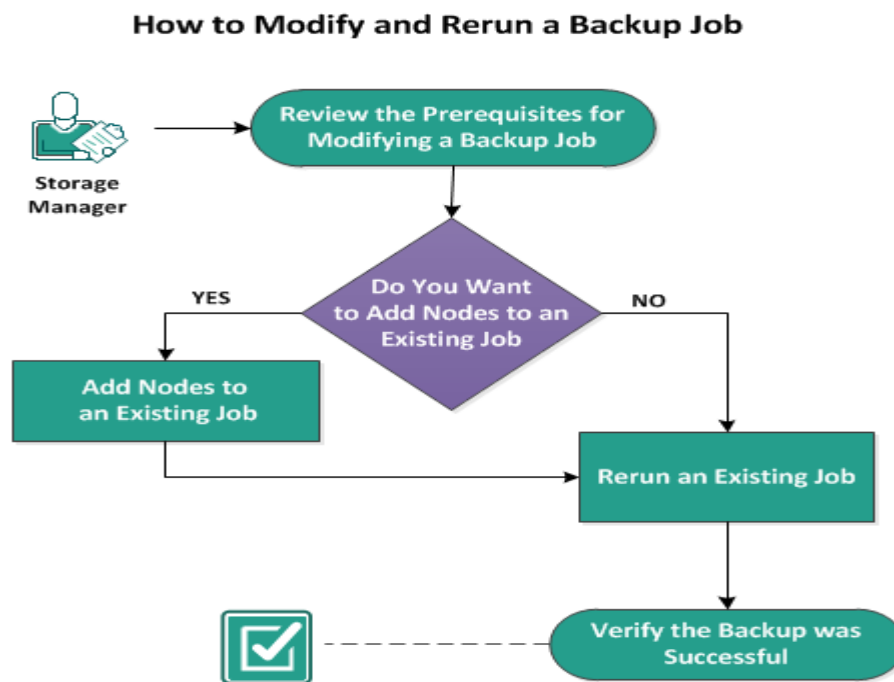
The backup data is successfully verified.

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:



Perform these tasks to modify and rerun a backup job:

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

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](#) that 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 Run backup job now dialog 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 for 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.

Perform these tasks for a file-level recovery:

- [Review the Restore Prerequisites](#)
- [Specify the Recovery Point for Host-Based Agentless Backup](#)
- [Specify the Recovery Point for Agent-Based Backup](#)
- [Specify the Target Machine Details](#)
- [Specify the Advanced Settings](#)
 - ◆ [\(Optional\) Manage Pre/Post Scripts for Automation](#)
- [Create and Run the Restore Job](#)
- [Verify that Files are Restored](#)

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.
- When the backup destination of a backup job is source local, then to perform a file-level restore job from the destination you need to export the source local destination through NFS or CIFS and specify the recovery point as available at NFS share or CIFS share.
- You have verified that the Linux Backup Server supports the file system that you want to restore.

For example, RedHat 7.x does not support the *reiserfs* file system. If the operating system of the Backup Server is RedHat 7.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.

- You have installed the following packages on the Linux Backup Server:
 - ◆ mdadm
 - ◆ kpartx
 - ◆ lvm2
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Specify the Recovery Point for Host-Based Agentless Backup

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.

Follow these steps:

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

- ◆ From Arcserve UDP:

- a. Click the **resources** tab.
- b. Select **All Nodes** in the left pane.

All the added nodes are displayed in the center pane.

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

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

- e. 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.

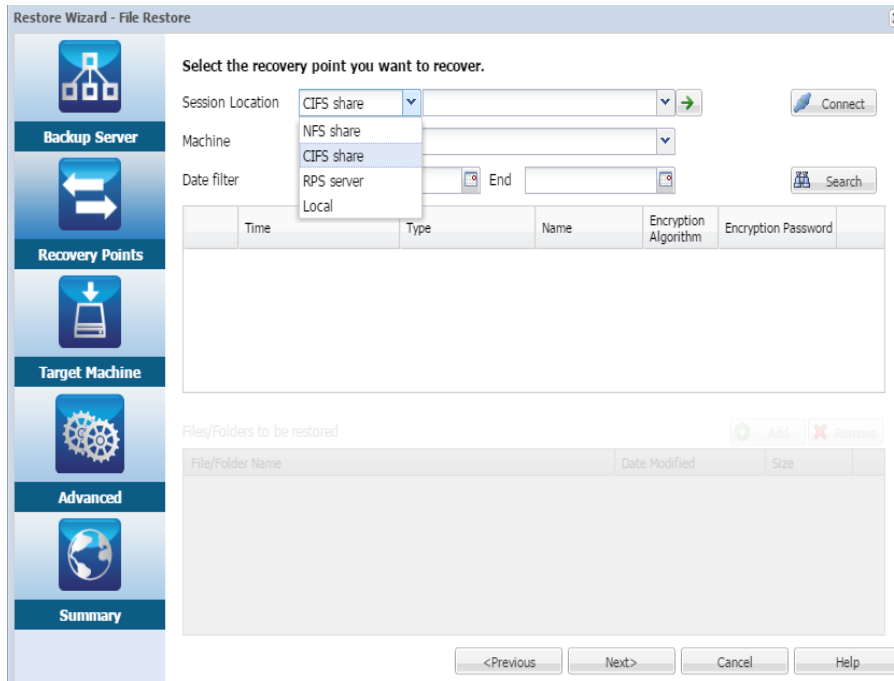
You can see the Backup Server in the **Backup Server** page of the **Restore Wizard**.

You cannot select any option from the **Backup Server** drop-down list.

2. Click **Next**.

The **Recovery Points** page of the **Restore Wizard** opens.

Important! If you have opened the Wizard from the Console, the session location and machine details are automatically displayed. You can skip to Step 5.



3. Select either a **CIFS share** or **RPS server** from the **Session Location** drop-down list.

Note: You cannot select NFS share or Local for restoring host-based agentless backup sessions.

4. Follow one of the following steps depending on your session location:

For CIFS share

- a. Specify the full path of the CIFS share and click **Connect**.
- b. Specify the username and password to connect to the CIFS share and click **OK**.

For RPS server

- a. Select RPS server and click Add.
The **Recovery Point Server** Information dialog opens.
- b. Provide the RPS details and click Load

- c. Select the data store from the drop-down list and click **Yes**.

The **Recovery Point Server Information** dialog closes and you see the wizard.

- d. Click **Connect**.

All the machines are listed in the Machine drop-down list.

- e. Select the machine from the drop-down list.

All the recovery points from the selected machine are displayed below the **Date Filter** option.

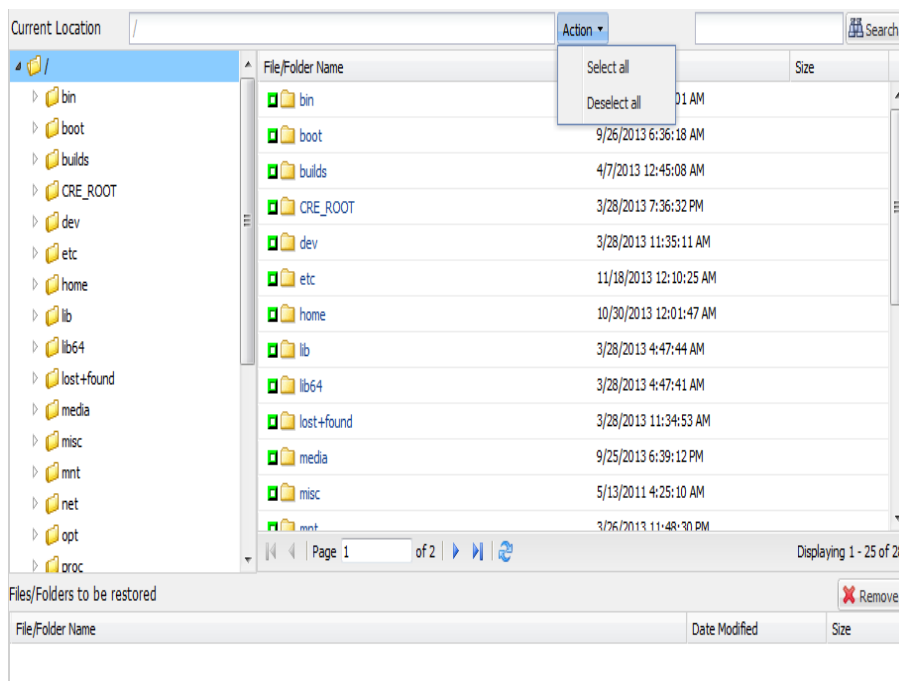
- 5. Apply the date filter to display the recovery points that are generated between the specified date and click **Search**.

Default: Recent two weeks.

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

- 6. 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.



Important! If you see the warning message, "The files/folders are displayed under device file. Click for more information." on the Console, refer the following Note for resolution.

Note: For some complex disk layout, the file system is shown by the device file. The change in the file system display behavior does not affect the function of host-based Linux VM file-level restore. You can browse the file system under the device file. Also, you can use the search function to search specific file or directory.

7. 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**.

8. Click **Next**.

The **Target Machine** page opens.

The recovery point is specified.

Specify the Recovery Point for Agent-Based Backup

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.

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** drop-down 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.

Important! If you have opened the Wizard from the Console, the session location and machine details are automatically displayed. You can skip to Step 5.



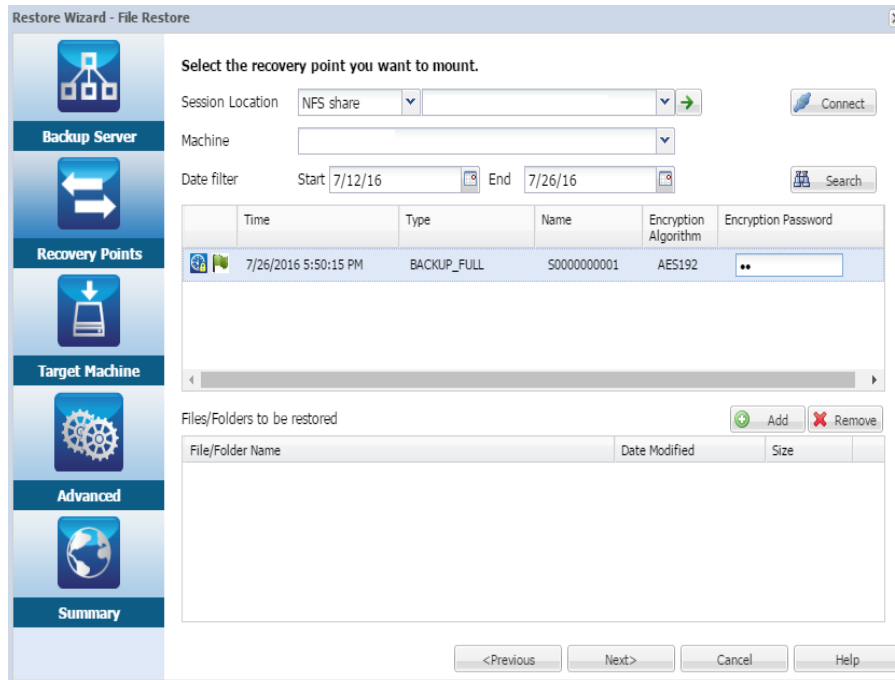
3. Select either a **CIFS share/NFS share/RPS server/Local** from the **Session Location** drop-down list.
4. Follow one of the following steps depending on your session location:

For CIFS share/NFS share/Local

- a. Specify the full path of the CIFS share/NFS share/Local and click **Connect**.

All the machines are listed in the **Machine** drop-down list.

Note: If you select the **CIFS share** option, specify the user name and password.



For RPS server

- a. Select RPS server and click **Add**.

The **Recovery Point Server Information** dialog opens.

- b. Provide the RPS details and click **Yes**.
- c. Select the data store from the drop-down list and click **Yes**.

The **Recovery Point Server Information** dialog closes and you see the wizard.

- d. Click **Connect**.

All the machines are listed in the Machine drop-down list.

- e. Select the machine from the drop-down list.

All the recovery points from the selected machine are appear below the **Date Filter** option.


5. 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.

6. 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.



Important! If you see the warning message, "The files/folders are displayed under device file. Click for more information." on the Console, refer the following Note for resolution.

Note: For some complex disk layout, the file system is shown by the device file. The change in the file system display behavior does not affect the function of host-based Linux VM file-level restore. You can browse the file system under the device file. Also, you can use the search function to search specific file or directory.

7. 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**.

8. 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

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 *.d2dduplicate<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.

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: _____

Destination: _____

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 *.d2d-duplicate<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/Arcserve/d2dserver/usr/prepost
```

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

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/Arcserve/d2dserver/usr/prepost`

Follow these steps:

1. Place the file in the following location of the Backup Server:
`/opt/Arcserve/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

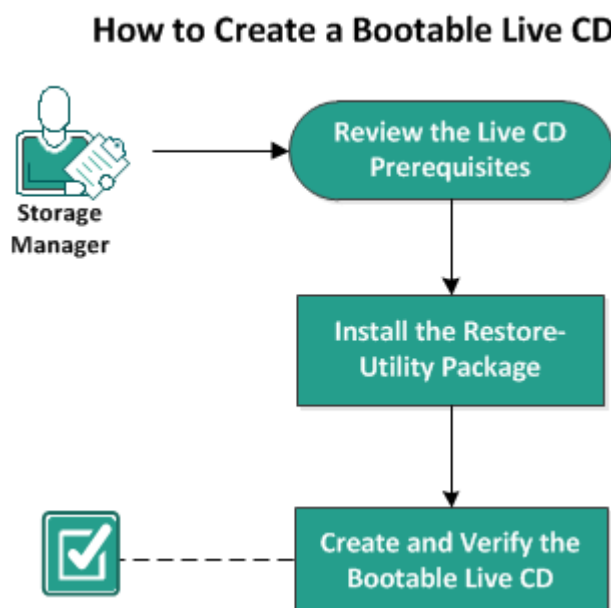
As a storage manager, you can create a bootable Live CD. When created, this bootable Live CD contains a complete read-only image of the computer operating system, and can be used to provide a temporary operating system functionality. This Live CD includes all your system settings and operating system 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](#)
- [Install the Restore-Utility Package](#)
- [Create and Verify the Bootable Live CD](#)

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](#) that 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 Use Live CD as a Linux Backup Server

You can use a live CD as a linux backup server.

Follow these steps:

1. Create a live CD from your Linux backup server.

To create the live CD, from home page

- ♦ Click Restore, Bare Metal Recovery (BMR).
- ♦ From the Restore Wizard - BMR, click the link **Click here to download Live CD**, and save as your live CD.

2. Start a Virtual machine or Physical machine with the live CD.

Note: We recommend 4GB memory for this machine.

When machine is started with live CD, you can view the following message:

*Use the following URL to access and manage this Arcserve UDP Agent (Linux):
<https://xxx.xxx.xxx.xxx:8014>.*

xxx.xxx.xxx.xxx refers to the current URL that the machine is using.

3. Enter the URL <https://xxx.xxx.xxx.xxx:8014> into your browser.

The Linux backup server home page appears.

4. Use the Linux backup server functions to perform a job.

For example: click Restore, Restore File, then find the backup session location, and perform the file level restore job.

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.

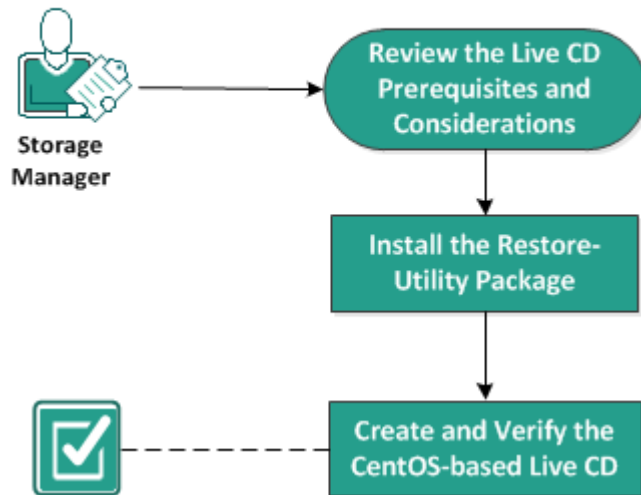
Note: The recovery points that you want to restore does not include the device drivers for the storage system of the target BMR machine. As a result, Arcserve UDP Agent (Linux) will block any attempt to perform a BMR job at an early stage.

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](#)
- [Install the Restore-Utility Package](#)
- [Create and Verify the CentOS-based Live CD](#)

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 1.1 GB.	Approximately 1.9 GB.
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, 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](#) that 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:
3. Run the following command to install the restore-utility package:

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

```
# ./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-7-x86_64-LiveGNOME.ISO

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

Important! When creating a bootable Live CD for CentOS 7, download and use the *CentOS-7-x86_64-LiveGNOME.ISO* image instead of *CentOS-7-x86_64-LiveCD.ISO* from the Cent OS 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 64-bit CentOS 7.0 or later Live CD to the /tmp folder on the Backup Server.

The *CentOS-7-x86_64-LiveGNOME.ISO* file is downloaded.

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 D2D for Linux Live CD based on CentOS and stores the ISO image file at the following location:

```
/opt/Arserve/d2dserver/packages/  
CentOS-LiveCD-for-UDP_  
Agent_Linux.iso
```

4. Navigate to the packages folder and verify that the CentOS-LiveCD-for-UDP_Agent_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. test

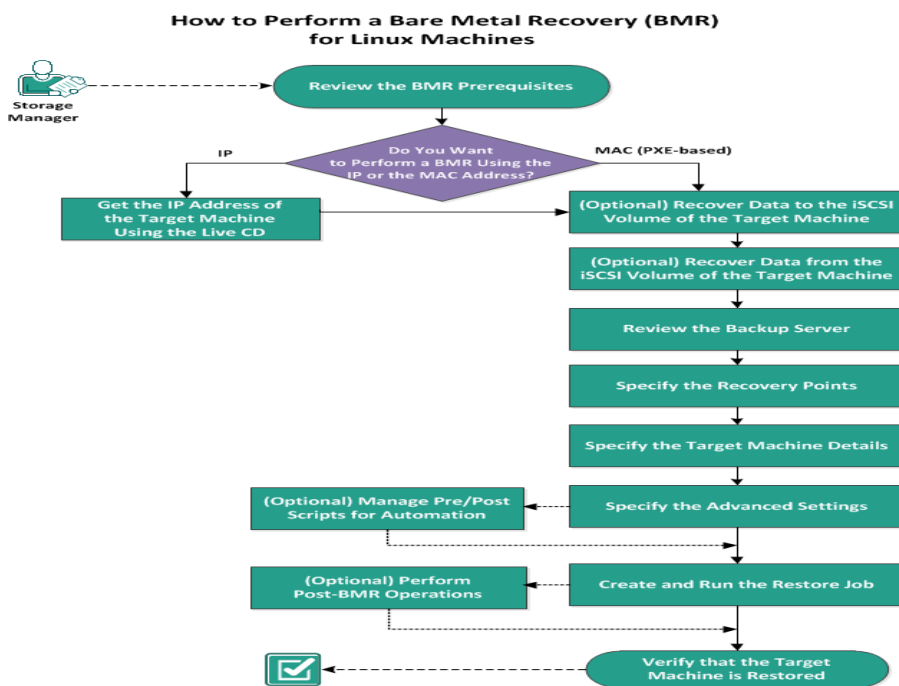
A complete BMR is possible because when you back up data, the backup also captures information related to the operating system, installed applications, drivers, and so on.

You can perform a BMR using either of the following options:

- Using the Command line option. For details, view [Create a Configuration template Using Command Line](#).
- 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.

Note: Machine can boot up. Only one NIC is configured.

The following diagram displays the process to perform a BMR using the IP or MAC address:



Complete the following tasks to perform a BMR:

- [Review the BMR Prerequisites](#)
- [Get the IP Address of the Target Machine Using the Live CD](#)
- [\(Optional\) Recover Data to the iSCSI Volume of the Target Machine](#)
- [\(Optional\) Recover Data from the iSCSI Volume to the Target Machine](#)
- [Review the Backup Server](#)
- [Specify the Recovery Points](#)
- [Specify the Target Machine Details](#)
- [Specify the Advanced Settings](#)
- [\(Optional\) Manage Pre/Posts Scripts for Automation](#)
- [Create and Run the Restore Job](#)
- [\(Optional\) Perform Post-BMR Operations](#)
- [Verify that the Target Machine is Restored](#)

Create a Configuration template Using Command Line

Create a configuration file so that the `d2dbmr` command can restore VMs based on the parameters specified in the file. The `d2dbmr` file gathers all the specifications from the file and performs the restore based on the specifications. The `d2dbmr` command is used to run BMR from the command line.

Syntax

```
d2dbmr --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 `--pwd-file=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. Create the configuration template using the following command:

```
d2dbmr --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.

storage_location_type

Specifies the type of the storage location of the session. The storage location can be CIFS, NFS, or RPS.

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 CIFS as the storage location.

storage_password

Specifies the password when you use CIFS as the storage location. The password is encrypted using the d2dutil encryption utility.

rps_server

Specifies the Recovery Point Server name when **storage_location_type** is RPS.

rps_server_username

Specifies the username of Recovery Point Server when **storage_location_type** is RPS.

rps_server_password

Specifies the password of Recovery Point Server when **storage_location_type** is RPS. The password is encrypted using the d2dutil encryption utility.

rps_server_protocol

Specifies the protocol of Recovery Point Server when **storage_location_type** is RPS.

rps_server_port

Specifies the port of Recovery Point Server when **storage_location_type** is RPS.

rps_server_datastore

Specifies the data store name of Recovery Point Server when **storage_location_type** is RPS.

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'.

exclude_volumes

Specifies the volume(s) to exclude for the target VM.

Do not exclude volume '/'. Use ':' to separate multiple volumes.

include_volumes

Specifies the volume(s) to include for the target VM.

Must include following volumes: / , /boot , /boot/efi , /home , /usr , /usr/local. Use ':' to separate multiple volumes.

restore_target

Specifies the IP/MAC address of the restore target.

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 be 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

Specifies the gateway address when you specify the static IP.

guest_dns

Specifies the DNS 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.

enable_instant_restore

(Optional) Specifies to enable instant restore. The values are yes and no.

auto_restore_data

(Optional) Specifies to restore data automatically. The values are yes and no.

script_pre_job_server

(Optional) Specifies the script to run before the job is executed on the server.

script_post_job_server

(Optional) Specifies the script to run after the job is executed on the server.

script_pre_job_client

(Optional) Specifies the script to run before the job is executed on the client.

script_post_job_client

(Optional) Specifies the script to run after the job is executed on the client.

script_ready_to_use

(Optional) Specifies the script to run when the target machine is ready to use and when the parameter **enable_instant_restore** value is Yes.

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.

5. Submit a job using d2dbmr template using the following command:

```
./d2dbmr -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 available, return to the shell environment immediately after submitting the job.

The restore job is submitted.

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.
- When the backup destination of the backup job is source local, then to perform a BMR job from the destination, you need to export the source local destination through NFS or CIFS and specify the recovery point as available at NFS share or CIFS share.
- The recovery point must be from the Linux agent-based backup.
- Review the [Compatibility Matrix](#) that 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>
```

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

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.

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

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 node.

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

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
```

```
# mkdir /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. Click the **resources** tab.
- b. Select **All Nodes** in the left pane.

All the added nodes are displayed in the center pane.

- c. In the center pane, select the node and click **Actions**.
- d. 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.

- e. 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.

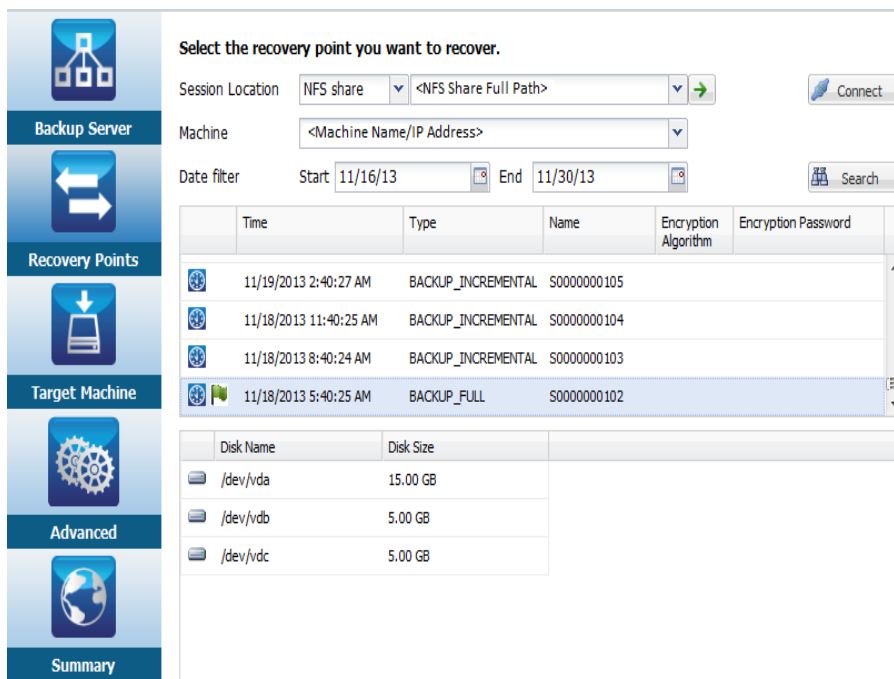
Important! To perform a BMR from a recovery point, the root volume and the boot volume must be present in the recovery point.

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.



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.

6. Apply the Volume filter settings for the selected recovery point and click **OK**.

All the available volumes present on that node are displayed. You can either include / exclude volumes based on the requirement.

Note: Do not exclude following volumes: / , /boot, /boot/efi, /home, /usr, /usr/local.

7. 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 **Enable instant BMR** option so that you can use the target machine instantly.

When you enable this option, Arcserve UDP Agent (Linux) first recovers all the necessary data that is required to start the machine. The remaining data are recovered after the target machine is started. The network connection must be constantly available during instant BMR.

Example: If you have 100-GB data and you want to perform a BMR and you *do not* select this option, first all 100-GB data will be recovered and then you can use the target machine. However, only around 1-GB data is required to start the machine. When you enable this option, first the required 1-GB data is recovered so that you

can start and use the machine. After the machine is started, the remaining 99-GB data is automatically recovered.

Note: The necessary data that is required to start the machine depends on the operating system configuration. You can also pause or resume the auto recovery of data if the **Do not recover data automatically after machine is started** option is not selected.

5. (Optional) Select the **Do not recover data automatically when machine is started** option to stop the automatic recovery of data when the target machine is started.

When you select the **Enable instant BMR** option, the default behavior is to recover the necessary data first and start the machine. After the machine starts, the remaining data gets recovered automatically. If you update any source data during the recovery, then by selecting this option, the data will be recovered until the point before they are updated.

6. 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/Arcserve/d2dserver/usr/prepost
```

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

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. (Optional) Select the **Reboot** option to automatically restart the target node after the BMR is complete.
9. 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 the job as a full backup.

backup.incremental

Identifies the job as the job as an incremental backup.

backup.verify

Identifies the job as 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 after 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/Arcserve/d2dserver/usr/prepost
```

Follow these steps:

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

```
/opt/Arcserve/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.

```
# lvs /dev/mapper/VGTest-LVTest
# mount /dev/sdb1 /data
--- Logical volume ---
LV Name                /dev/VGTest/LVTest
VG Name                VGTest
LV UUID                udoBIx-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): pDisk /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
/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
/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.

```
# vgchange -a y VGTest# 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 /lvm
# lvdisplay -m /dev/mapper/VGTest-LVTest
---Logical volume---
LV Name                /dev/VGTest/LVTest
VG Name                VGTest
LV UUID                GTP0a1-kUL7-WUL8-bpbM-9eTR-
SVz1-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 Perform a Bare Metal Recovery (BMR) for Linux Machines in AWS Cloud

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 of the target Linux Instance in Amazon EC2. If you boot the target Linux Instance using the Arcserve UDP Agent (Linux) AMI, then you can get the private IP address of the Instance.

The process to perform a BMR for Linux Instances in Amazon EC2 is almost the same with Linux machines in the local.

Complete the following tasks to perform a BMR:

- [Review the BMR Prerequisites](#)
- [Launch an Instance Using the Arcserve UDP Agent Live CD](#)
- [Review the Backup Server Instance](#)
- [Specify the Recovery Points](#)
- [Specify the Target Instance Details](#)
- [Specify the Advanced Settings](#)
- [Create and Run the Restore Job](#)
- [Verify that the Target Instance is Restored](#)

Review the BMR Prerequisites

Consider the following options before performing a BMR for Linux Instances in Amazon EC2:

- You have a valid recovery point and the encryption password, if any, for restore.
- When the backup destination of the backup job is source local, then to perform a BMR job from the destination, you need to export the source local destination through NFS or CIFS and specify the recovery point as available at NFS share or CIFS share.
- The recovery point must be from the Linux agent-based backup.
- You have an Arcserve UDP Agent for Linux Instance in Amazon EC2.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Launch an Instance Using the Arcserve UDP Agent Live CD

Before performing a BMR for Linux Instances in Amazon EC2, you need to launch a BMR target Instance using Arcserve UDP Agent Live CD. When the target BMR instance is ready, then you can get the IP address of the Instance and perform a BMR job with IP address.

Follow these steps:

1. Log into the EC2 management console with your account and select **Launch Instance**.
2. Select an Amazon Machine Image (AMI) in Community AMIs.

You can search the Live CD AMI with *Arcserve_UDP_Agent_Linux-LiveCD* in Community AMIs.

Notes:

- If PVM is the backup source node that you want to restore, select *Arcserve_UDP_Agent_Linux-LiveCD-PVM-UDP\$version* AMI to launch the instance.
 - If HVM or other target machine is the backup source node that you want to restore, select *Arcserve_UDP_Agent_Linux-LiveCD-HVM-UDP\$version* AMI to launch the instance.
3. From Launch Instance Wizard, select the required Instance Type.
 4. Configure the Instance details as you launch other Instances. For example: including the Network, Subnet, Auto-assign Public IP or not, and so on.
 5. Add storage for the Instance, using the following steps:
 - a. Get the disk information including the disk number and the disk size of the backup source node that you want to restore. You can get the disk information when you select a recovery point in the Restore wizard to perform a BMR job.
 - b. Extend the root volume size to match the root disk size of the backup source node. You can add other disks, if the backup source node has more disks.
 6. Add tags for the BMR target Instance.
 7. Configure the security group for the BMR target instance, using the following steps:
 - a. Create a new security group for SSH Type.
 - b. To make the BMR target instance safer, select the **Custom** mode for the source that will determine the traffic to reach the BMR target instance in the

newly created rule. Specify the custom source with the CIDR format so that the BMR target instance is accessible to the Arcserve UDP Agent for Linux Server but not accessible to other internet machines.

For example, if the IP of the Arcserve UDP Agent for Linux Server is 172.31.X.X, specify the source as 172.31.0.0/16 or 172.0.0.0/8.

8. Review instance details and then click **Launch**.

The **Select an existing key pair or create a new pair** dialog is displayed.

9. From the dialog, select the **Proceed without a key pair** option, and click **Launch Instances**.

10. Acquire the private IP in the instance description, when the BMR target instance is ready to use.

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 BMR target instance details.

Review the Backup Server Instance

When you open the **Restore Wizard**, review the Backup Server Instance 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. Click the **resources** tab.
- b. Select **All Nodes** in the left pane.

All the added nodes are displayed in the center pane.

- c. In the center pane, select the node and click **Actions**.
- d. Click **Restore** from the **Actions** drop-down menu.

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

- e. 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. 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..



Select the recovery point you want to recover.

Session Location: NFS share <NFS Share Full Path>

Machine: <Machine Name/IP Address>

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

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 **BMR Target Instance** page opens.

The recovery point is specified.

Specify the Target Instance Details

Specify the BMR target instance details to restore data to that machine. A target instance is a bare metal machine where you perform a BMR. You need the IP address of the BMR target instance that you previously recorded at the start of this process.

Follow these steps:

1. Enter the IP address of the BMR target instance in the **MAC/IP Address** field.
2. Enter a name in the **Host Name** field.

The BMR target instance uses 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 **Enable instant BMR** option so that you can use the target machine instantly.

When you enable this option, Arcserve UDP Agent (Linux) first recovers all the necessary data that is required to start the machine. The remaining data are recovered after the target machine is started. The network connection must be constantly available during instant BMR.

Example: If you have 100-GB data and you want to perform a BMR and you *do not* select this option, first all 100-GB data will be recovered and then you can use the target machine. However, only around 1-GB data is required to start the machine. When you enable this option, first the required 1-GB data is recovered so that you can start and use the machine. After the machine is started, the remaining 99-GB data is automatically recovered.

Note: The necessary data that is required to start the machine depends on the operating system configuration. You can also pause or resume the auto recovery of data if the **Do not recover data automatically after machine is started** option is not selected.

5. (Optional) Select the **Do not recover data automatically when machine is started** option to stop the automatic recovery of data when the target machine is started.

When you select the **Enable instant BMR** option, the default behavior is to recover the necessary data first and start the machine. After the machine starts, the remaining data gets recovered automatically. If you update any source data during the recovery, then by selecting this option, the data will be recovered until the point before they are updated.

6. Click **Next**.

The **Advanced** page opens.

The BMR target instance 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 BMR target instance.

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/Arcserve/d2dserver/usr/prepost
```

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

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 BMR target instance from participating in the recovery process.
7. (Optional) Select the **Reboot** option to automatically restart the target node after the BMR is complete.
8. Click **Next**.

The **Summary** page opens.

The advanced settings are specified.

(Optional) Manage Pre/Post Scripts for Automation in AWS Cloud

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 the job as a full backup.

backup.incremental

Identifies the job as the job as an incremental backup.

backup.verify

Identifies the job as 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 BMR target instance after the job starts.

post-job-target

Identifies the script that runs on the BMR target instance after the job completes.

pre-snapshot

Identifies the script that runs on the BMR target instance before capturing the snapshot.

post-snapshot

Identifies the script that runs on the BMR target instance 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/Arcserve/d2dserver/usr/prepost
```

Follow these steps:

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

```
/opt/Arcserve/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.

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:

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 BMR target instance 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.

```
# lvs /dev/mapper/VGTest-LVTest
# mount /dev/sdb1 /data

--- Logical volume ---
LV Name                /dev/VGTest/LVTest
VG Name                VGTest
LV UUID                udoBIx-XKBS-1Wky-3FVQ-mxMf-
FayO-tpfP18
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): pDisk /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.

```
# vgchange -a y VGTest# 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 /lvm  
# lvs -m /dev/mapper/VGTest-LVTest  
---Logical volume---  
  
LV Name                /dev/VGTest/LVTest  
VG Name                VGTest  
LV UUID                GTP0a1-kUL7-WUL8-bpbM-9eTR-  
SVz1-WgA1lh  
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 Instance is Restored

After the completion of restore job, verify that the target instance is restored with relevant data.

Follow these steps:

1. Navigate to the BMR target instance that you restored.
2. Verify that the BMR target instance has all the information that you backed up.

The target instance is successfully verified.

Note: When the BMR target instance is ready to use, you can modify the newly created security group according to your business requirement.

The BMR is successfully performed for Linux Machines.

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

A BMR restores the operating system and software applications, and recovers all the backed-up data. 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 of the target Linux virtual machine in Microsoft Azure. The process to perform a BMR for Linux Instances in Azure Cloud is slightly different than the process with Linux machines in the local.

Complete the following tasks to perform a BMR:

- [Review the BMR Prerequisites](#)
- [Create a New Machine on Microsoft Azure as BMR target](#)
- [Review the Backup Server Virtual Machine](#)
- [Specify the Recovery Points](#)
- [Specify the Target Virtual Machine Details](#)
- [Specify the Advanced Settings](#)
- [Create and Run the Restore Job](#)
- [Verify that the Target Instance is Restored](#)

Review the BMR Prerequisites

Consider the following options before performing a BMR for Linux Instances in Microsoft Azure:

- You have a valid recovery point and the encryption password, if any, for restore.
- When the backup destination of the backup job is source local, then to perform a BMR job from the destination, you need to export the source local destination through NFS or CIFS and specify the recovery point as available at NFS share or CIFS share.
- The recovery point must be from the Linux agent-based backup.
- You have an Arcserve UDP Agent for Linux Instance in Microsoft Azure.
- BMR to the target Linux virtual machine should have the same operation system as the source Linux node.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Create a new Machine on Microsoft Azure as BMR target

For BMR on Azure, user could perform a BMR for virtual machine to a Linux virtual machine with same Linux System on Azure directly instead of launch the target node with Arcserve UDP Agent Live CD.

Firstly, create a new virtual machine on Azure as BMR target node. Please verify below prerequisites.

- Prepare a new virtual machine with the same operation system as the virtual machine that wants to perform a BMR.
- Configure the Authentication Type as Password for the virtual machine. Remember the User name and Password of the virtual machine.
- Configure the Resource Group like the group at the Linux Backup Server that performs BMR. Otherwise, BMR fails to create the SSH connection between Linux Backup Server and the target virtual machine.

Review the Backup Server Virtual Machine

For more information, see [Review the Backup Server](#).

Specify the Recovery Points

For more information, see [Specify the Recovery Points](#).

Specify the Target Virtual Machine Details

Specify the BMR target virtual machine details to restore data to that machine. A target virtual machine is a bare metal machine where you perform a BMR. You need the IP address, User name, and Password of the BMR target virtual machine that you previously recorded at the start of this process.

Follow these steps:

1. From the Restore BMR - Wizard screen, enter the following details:
 - Enter the IP address of the BMR target virtual machine in the IP Address field.
 - Enter User name and Password of the target virtual machine that you have created on Azure.

2. For Machine details:

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

The BMR target virtual machine uses this name as the host name after the restore process is complete.

- Verify if DHCP is selected by default as Network Settings.

Note: Only DHCP is available on Azure. IP address is automatically configured.

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.

3. (Optional) Select the **Enable instant BMR** option so that you can use the target machine instantly.

When you enable this option, Arcserve UDP Agent (Linux) first recovers all the necessary data that is required to start the machine. The remaining data are recovered after the target machine is started. The network connection must be constantly available during instant BMR.

Example: If you have 100-GB data and you want to perform a BMR and you *do not* select this option, first all 100-GB data will be recovered and then you can use the target machine. However, only around 1-GB data is required to start the machine. When you enable this option, first the required 1-GB data is recovered so that you

can start and use the machine. After the machine is started, the remaining 99-GB data is automatically recovered.

Note: The necessary data that is required to start the machine depends on the operating system configuration. You can also pause or resume the auto recovery of data if the **Do not recover data automatically after machine is started** option is not selected.

4. (Optional) Select the **Do not recover data automatically when machine is started** option to stop the automatic recovery of data when the target machine is started.

When you select the **Enable instant BMR** option, the default behavior is to recover the necessary data first and start the machine. After the machine starts, the remaining data gets recovered automatically. If you update any source data during the recovery, then by selecting this option, the data will be recovered until the point before they are updated.

5. Click **Next**.

The **Advanced** page opens.

The BMR target instance details are specified.

Specify the Advanced Settings

For details, view [Specify the advanced settings](#).

Create and Run the Restore Job

For details, view [Create and Run the Restore Job](#).

Verify that the Target Virtual Machine is Restored

For details, view [Verify that the Target Node is Restored](#).

How to Perform a Migration BMR for Linux Machines

A migration BMR is a two part process where the data is first restored to a temporary machine and then to the actual machine. A BMR with instant BMR option enabled lets you recover data to a temporary machine. You can use the temporary machine until the actual machine is ready. When you have the actual machine, a migration BMR lets you migrate data from the temporary machine to the actual machine. When performing a migration BMR, any data that you create on the temporary machine gets migrated to the actual machine.

Note: You can perform Migration BMR with an agent-based backup only. An agent-less backup does not support Migration BMR.

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.

Note: Machine can boot up. Only one NIC is configured.

Complete the following tasks to perform a Migration BMR:

- [Review the Prerequisites for Migration BMR](#)
- [Perform a BMR to the Temporary Machine](#)
- [Perform a Migration BMR](#)
- [Verify that the Target Machine is Restored](#)

Review the Prerequisites for Migration BMR

Consider the following options before performing a migration 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.
- The recovery point must be from the Linux agent-based backup.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Perform a BMR to the Temporary Machine

Before you perform a migration BMR, you have to restore data from the source to a temporary machine. To restore the data temporarily, you can perform a BMR to the temporary machine. After the temporary machine is ready to use, you can continue working on the temporary machine.

When the actual machine is ready, you can perform a migration BMR from the temporary machine to the actual machine.

Note: For more information on performing a BMR, see How to Perform a Bare Metal Recovery (BMR) for Linux Machines.

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.

- b. Log in to Arcserve UDP Agent (Linux).

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

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

3. Provide all the details in the **Restore Wizard - BMR** and save the wizard.
4. Ensure that you select the **Enable instant BMR** check box on the **Target Machine** page of the wizard.
5. Ensure that you select the **Do not recover data automatically after machine is started** check box on the **Target Machine** page of the wizard.
6. Run the BMR job.

The temporary machine is recovered using the BMR, with the instant BMR option enabled. You can use the temporary machine until the actual machine is ready.

Perform a Migration BMR

When the actual machine is ready, perform a migration BMR. Migration BMR restores the original data from the backup session and the new data from the temporary machine to the actual machine.

Follow these steps:

1. Click **Restore** from the **Wizard** menu and select **Migration BMR**.

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

2. Provide all the details in the **Restore Wizard - Migration BMR**.

Note: For more information on performing a BMR, see How to Perform a Bare Metal Recovery (BMR) for Linux Machines.

3. Ensure that the following information is provided on the **Backup Server** page of the wizard.
 - a. Select the instant VM recovery job or the Instant BMR job.

Local Server

Specifies that the Backup Server is locally managed. The BMR job for the temporary machine is run on the local server.

Remote Server

Specifies that the Backup Server is remotely managed. The BMR job for the temporary machine is run on the remote server. You have to provide the remote server details to connect to the remote server.

- b. Select the restore job from the Job Name drop-down list.

The list displays the Instant VM recovery job or Instant BMR job, which is in the Ready to use job phase or Power off job phase, once it is ready to use.

4. Save the BMR job.

In the home page, the **Job Phase** on the **Job Status** tab changes to **Click here to migrate data**.

5. (Optional) Boot the temporary machine using a Live CD when the selected job type is Instant BMR.
6. From the **Job Status** tab, click **Click here to migrate data**.

The data migration begins.

You have successfully performed a migration BMR.

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 from the temporary machine, including any new data that you created on the temporary machine.

The target machine is successfully verified.

The migration BMR is successfully performed for agent-based Linux machines.

How to Perform a Migration BMR for Linux Machines from Amazon EC2 to local

A migration BMR is a two part process where the data is first restored to a temporary machine and then to the actual machine. A BMR with instant BMR option enabled lets you recover data to a temporary machine. You can use the temporary machine until the actual machine is ready. When you have the actual machine, a migration BMR lets you migrate data from the temporary machine to the actual machine. When performing a migration BMR, any data that you create on the temporary machine gets migrated to the actual machine.

You may encounter an issue at Linux server locally that needs some downtime. Then, you can use backup session to create Instant VM on Amazon EC2, and use that server to provide continues services. When the local issue is fixed, migration BMR helps you to migrate all the data from Amazon EC2 to local, and the local server is restored to provide required service again.

Note: You can perform Migration BMR with an agent-based backup only. An agent-less backup does not support Migration BMR.

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.

Note: Machine can boot up. Only one NIC is configured.

Complete the following tasks to perform a Migration BMR:

- [Review the Prerequisites for Migration BMR](#)
- [Perform a Migration BMR](#)
- [Verify that the Target Machine is Restored](#)

Review the Prerequisites for Migration BMR

Consider the following options before performing a migration 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.
- The recovery point must be from the Linux agent-based backup.
- Review the [Compatibility Matrix](#) that provides the supported operating systems, databases, and browsers.

Perform a BMR Migration from Amazon EC2 to the local Machine

Before you perform a migration BMR from Amazon EC2, you have to restore data from the source to an EC2 instance. To restore the data temporarily, you can perform an Instant VM to the EC2 instance. After the EC2 instance is ready to use, you can continue working on the instance.

When the actual local machine is ready, you can perform a migration BMR from the Amazon EC2 instance to the actual local machine.

Note: For more information about performing a BMR, see [How to Perform a Bare Metal Recovery \(BMR\) for Linux Machines](#).

Follow these steps:

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

- ◆ From Arcserve UDP:

- a. Log on 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.

- b. Log in to Arcserve UDP Agent (Linux).

2. Click **Restore** from the **Wizard** menu and select **Migration BMR**.

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

3. Perform the following steps, and then click **Next**:

- a. Select **Remote server** as Server Location.
- b. Specify the Linux Backup Server on Amazon EC2 to connect to the server.
- c. Enter Host Name, User Name, Password, Protocol, and Port for the Linux Backup Server.
- d. Click **Refresh**, select the restore job from the **Job Name** drop-down list.

The list displays the Instant VM recovery job that is in the **Ready to use** job phase or **Power off** job phase, when ready to use.

The Recovery Points section appears.

4. From the **Recovery Points** section, perform the following steps, and click **Connect**.

- Select the **RPS server** that was created on Amazon EC2 (or NFS CIFS share on Amazon EC2).
- Select the corresponding data store.

The machine is automatically loaded according to the Instant VM job.

- Select the session and click **Next**.

You are automatically taken to the **Target Machine** tab.

5. On the Target Machine section, enter the MAC/IP address, and click **Next**.

Note: You can boot a local machine with LiveCD to get the MAC/IP address.

You are led to the Advanced section.

6. In the **Advanced** section, configure the Pre/Post Scripts, and then click Next.

The **Summary** section appears.

7. Specify the job name and click **Submit**.

A BMR job is performed on the machine booted with LiveCD.

8. From the Linux agent home page, navigate to the **Job Status** tab and click **Click here to migrate data**.

The data on Amazon EC2 VM is migrated to your local machine.

You have successfully performed a migration BMR.

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 from the temporary machine, including any new data that you created on the temporary machine.

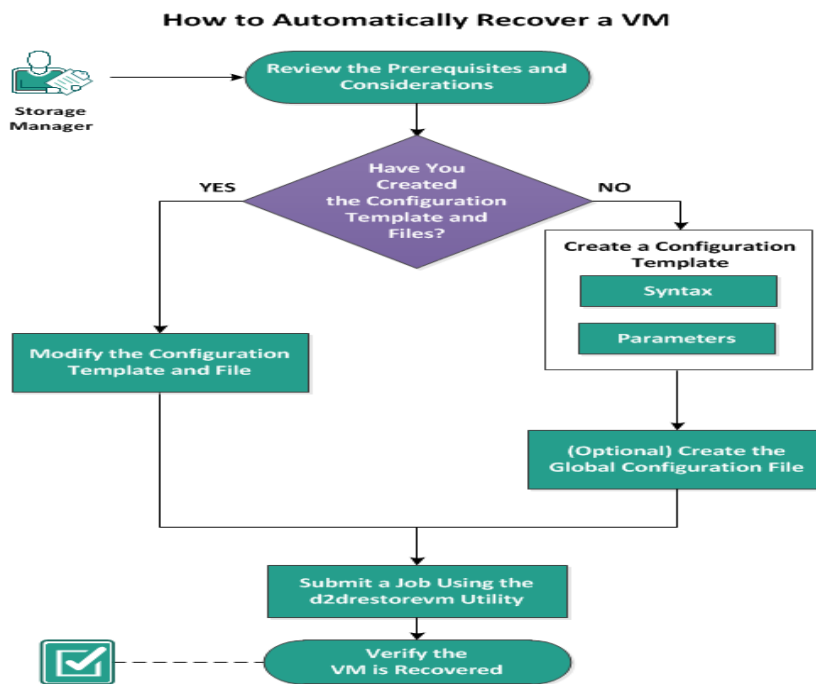
The target machine is successfully verified.

The migration BMR is successfully performed for agent-based 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 or Instant 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](#)
- [Create a Configuration Template](#)
- [\(Optional\) Create the Global Configuration File](#)
- [Modify the Configuration Template and File](#)
- [Submit a Job Using the `d2drestorevm` Utility](#)
- [Verify the VM is Recovered](#)

Review the Prerequisites and Considerations

Review the following prerequisites before you restore the VM:

- The following versions of hypervisors are supported for BMR or Instant VM using the d2drestorevm utility:
 - ♦ XenServer 6.0 and later (restore VM using regular BMR method)
 - ♦ OVM 3.2 (restore VM using regular BMR method)
 - ♦ VMware vCenter/ESX(i) 5.0 or later (submit Instant VM job)
 - ♦ Windows Hyper-V server 2012 or later (submit Instant VM job)
- 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).
- 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.
- 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:



- 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 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 when restoring to Xen or OVM:
`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](#) that 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 D2D 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.
- 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 `--pwd-file=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 and OVM.

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.

vm_svr_password

Specifies the password of the hypervisor. The password is encrypted using the d2dutil encryption utility.

vm_sub_server

Specifies the ESX server name when restoring to vCenter.

vm_svr_protocol

Specifies the protocol of the hypervisor when restoring to vCenter/ESX(i).

vm_svr_port

Specifies the port of the hypervisor when restoring to vCenter/ESX(i).

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.

vm_memory

Specifies the memory of the virtual machine in MB when restoring to vCenter/ESX(i) or Hyper-V.

vm_cpu_count

Specifies the CPU count of the virtual machine when restoring to vCenter/ESX(i) or Hyper-V.

vm_resource_pool

Specifies the resource pool of the hypervisor when restoring to vCenter/ESX(i).

vm_datastore

Specifies the datastore of the hypervisor when restoring to vCenter/ESX(i).

storage_location_type

Specifies the type of the storage location of the session. The storage location can be CIFS, NFS, or RPS.

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.

rps_server

Specifies the Recovery Point Server name when **storage_location_type** is RPS.

rps_server_username

Specifies the username of Recovery Point Server when **storage_location_type** is RPS.

rps_server_password

Specifies the password of Recovery Point Server when **storage_location_type** is RPS. The password is encrypted using the d2dutil encryption utility.

rps_server_protocol

Specifies the protocol of Recovery Point Server when **storage_location_type** is RPS.

rps_server_port

Specifies the port of Recovery Point Server when **storage_location_type** is RPS.

rps_server_datastore

Specifies the data store name of Recovery Point Server when **storage_location_type** is RPS.

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

Specifies the gateway address when you specify the static IP.

guest_dns

Specifies the DNS 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.

enable_instant_restore

(Optional) Specifies to enable instant restore. The values are yes and no.

auto_restore_data

(Optional) Specifies to restore data automatically. The values are yes and no.

script_pre_job_server

(Optional) Specifies the script to run before the job is executed on the server.

script_post_job_server

(Optional) Specifies the script to run after the job is executed on the server.

script_pre_job_client

(Optional) Specifies the script to run before the job is executed on the client.

script_post_job_client

(Optional) Specifies the script to run after the job is executed on the client.

script_ready_to_use

(Optional) Specifies the script to run when the target machine is ready to use and when the parameter **enable_instant_restore** value is Yes.

force

Specifies whether to force restore the VM. The values are yes and no.

Default: no

exclude_volumes

Specifies the volume(s) to exclude for the target VM.

Do not exclude volume '/'. Use ':' to separate multiple volumes.

include_volumes

Specifies the volume(s) to include for the target VM.

Must include following volumes: / , /boot , /boot/efi , /home , /usr , /usr/local. Use ':' to separate multiple volumes.

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/Arcserve/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.

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.

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/Arcserve/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/Arcserve/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:

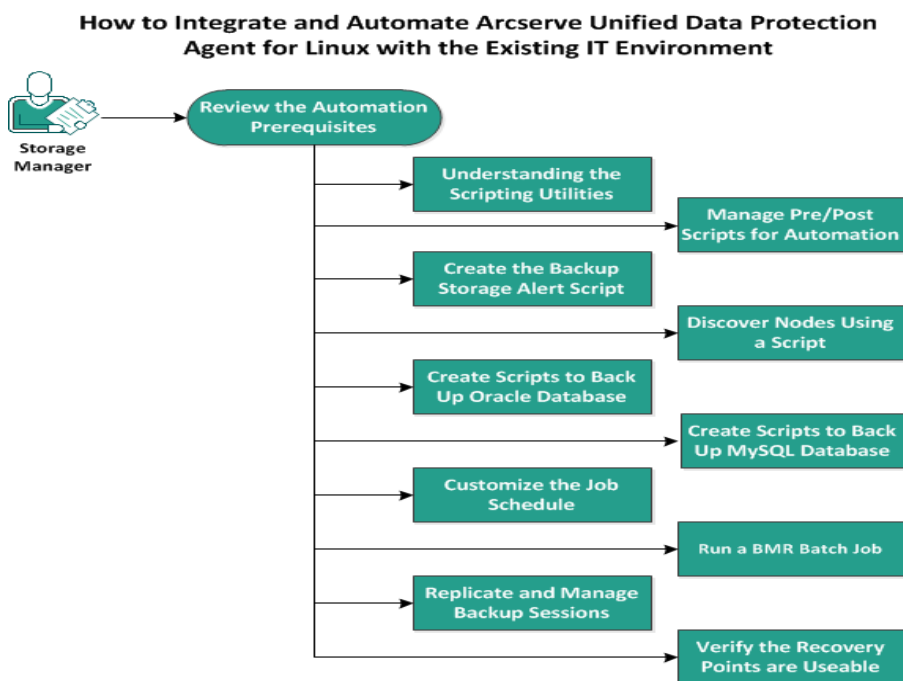
1. Navigate to the VM that you restored.
2. Verify that the VM has all the information that you backed up.

The VM is successfully verified.

How to Integrate and Automate Arcserve D2D 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:



Perform the following tasks to automate and manage Arcserve UDP Agent (Linux):

- [Review the Automation Prerequisites](#)
- [Understanding the Scripting Utilities](#)
- [Manage Pre/Post Scripts for Automation](#)
- [Create the Backup Storage Alert Script](#)
- [Discover Nodes Using a Script](#)
- [Create the Scripts to Back Up Oracle Database](#)
- [Create the Scripts to Back Up MySQL Database](#)

- [Customize the Job Schedule](#)
- [Run a BMR Batch Job](#)
- [Replicate and Manage Backup Sessions](#)
- [Verify the Recovery Points are Usable](#)

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](#) that 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/Arcserve/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 `--pwd-file=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/Arcserve/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 --listLists all the nodes that are managed by
the Backup Server.

# ./d2dnode --add=nodename/ip --user=username --pass-
word=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 --pass-
word=password --rootuser=rootaccount --rootp-
wd=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-rot 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 192.168.100.150

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 submit a file restore job:

```
d2drestorefile --createtemplate=file
```

Specifies to create a template. After the template is created, you can modify the template. This template is used by the `d2drestorefile` command. You can set values in this template. The `d2drestorefile` reads from the template and provides the result as specified in the template.

```
d2drestorefile --template=restore_template [--wait]
```

Specifies to submit the file restore job. If you include the `[--wait]` parameter in the command, the status message is displayed only after the completion of the restore job.

5. 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 --recoverysetstart --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.

The `--recoverysetstart` parameter is optional. If this option is specified, the current backup is converted to Full backup and marked as the first recovery point of the Recovery Set if Recovery Set is not available.

```
# ./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-e=filepath>
```

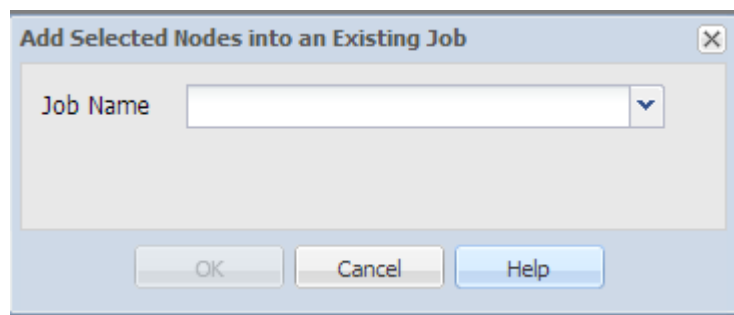
Exports multiple jobs from the Backup Server to a file. If you want similar backup configurations in multiple Backup Servers, you can export the backup jobs to a file, and then import the file to other Backup Servers.

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
```

6. 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=  
e=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=  
e=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=r-  
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 storage path.

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 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.

```
# ./d2djobhistory --starttime=yyyymmdd --endtime=yyyymmdd --headers=column_name1, column_name2,...column_name_n --width=width_value --format=column/csv/html
```

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.

D2D_RECOVERYPOINT

Identifies the recovery point created by the backup job. This value is applicable in the post backup script only.

D2D_RPSSCHEDULETYPE

Identifies the schedule type when backing up to a Data Store on RPS. The following values identify the D2D_RPSSCHEDULETYPE variable:

daily

Identifies the schedule as a daily backup.

weekly

Identifies the schedule as a weekly backup.

monthly

Identifies the schedule as a monthly backup.

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/Arcserve/d2dserver/usr/prepost
```

Follow these steps:

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

```
/opt/Arcserve/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.

Example of Creating User-defined Scripts

The environment variable `D2D_JOBSTAGE`, which has four different stages, is very important to write script. At the stage `pre_share`, you could do some preparations or implement the access method. At the stage `post_share`, you can also implement the access method here and do some other things. The difference between the two stages is that `D2D_SHARE_PATH` indicated path is available at `post_share` stage. The stages `pre_cleanup` and `post_cleanup` give you the chance to clean up resources that you allocate or break connections to your shared path. The difference of the two stages is that `D2D_SHARE_PATH` indicated path is available at stage `pre_cleanup` and not available at stage `post_cleanup`.

Notes:

- You can read pass word, that you set for the user in web UI, from standard input.
- Your codes are executed by different process at different stage. So if you want to share data at different stage, you have to use global resource such as temporary file or database.

Example: Create User-defined Scripts

Note: SFTP script is used as example in the `examples/sharerp` directory.

```
#!/bin/bash

function pre_sftp_share()
{
    local share_path=${D2D_SHARE_PATH}
    local user_name=${D2D_SHARE_USER}
    local pass_word=""

    # Read pass word from standard input.
    read -s pass_word

    # Check user whether exist.
    if grep $user_name /etc/passwd >/dev/null 2>&1; then
        return 1
    fi

    # Add new user.
    useradd $user_name -d $share_path >/dev/null 2>&1
```

```
[ $? -ne 0 ] && return 2

# Set pass word for the user.
echo -e "$pass_word\n$pass_word"|passwd "$user_name" >/dev/null 2>&1
[ $? -ne 0 ] && return 3

return 0
}

function post_sftp_share()
{
return 0
}

function pre_sftp_cleanup()
{
return 0
}

function post_sftp_cleanup()
{
local user_name=${D2D_SHARE_USER}

# Delete the user.
userdel $user_name >/dev/null 2>&1

return 0
}

# Main
#####
ret=0
stage=${D2D_JOBSTAGE}
case $stage in
```

```
pre_share)
pre_sftp_share
ret=$?
;;

post_share)
post_sftp_share
ret=$?
;;

pre_cleanup)
pre_sftp_cleanup
ret=$?
;;

post_cleanup)
post_sftp_cleanup
ret=$?
;;
esac

exit $ret
```

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/Arcserve/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/Arcserve/d2dserver/examples/discovery
```

You can copy the sample script to the following location and modify that script per your requirement:

```
/opt/Arcserve/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/Arcserve/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/Arcserve/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 \(Linux\)](#).

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_mysqlldb() {
(
echo "flush tables with read lock;"
sleep 5
) | mysql -u$dbuser -p$dbpwd ${ARGUMENTS}
}
lock_mysqlldb &
PID="/tmp/mysql-plock.$!"
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/Arcserve/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/Arcserve/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/Arcserve/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 drop-down 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/Arcserve/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.

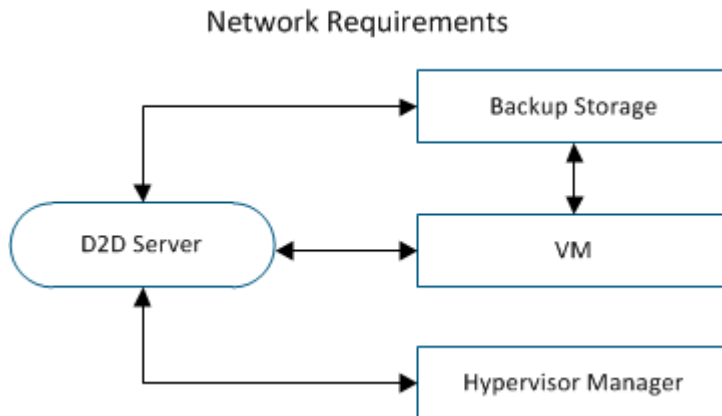
- Review the network requirements
- Identify a network in which the VMs will be connected.

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
- OVM 3.2

Arguments:

--template

Identifies the template that includes the parameters to run the d2dverify utility.

--createtemplate

Creates an empty template 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, or ovm.

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

Perform the following tasks to manage the Backup Server settings:

- [Review the Prerequisites to Manage the Backup Server](#)
- [Configure the Job History and Activity Log Retention Settings](#)
- [Configure the Debug Log Retention Settings](#)
- [Configure the UI Timeout Duration](#)
- [Change the SSH Port Number of the Backup Server](#)
- [Manage the Recovery Sets](#)
- [Disable the BOOTPD and TFTP Services](#)
- [Improve the Query Performance for Job History and Activity Log](#)
- [Skip CIFS and NFS Client Verification](#)
- [Skip CIFS and NFS Validation on Linux Backup Server](#)
- [Configure the Default Temporary Folder](#)
- [Configure the Snapshot Path for Backup Node](#)
- [Configure the Hyper-V Server Connection Information for Instant VM](#)

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](#) that 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/Arcserve/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/Arcserve/d2dserver/configfiles/server.cfg
```

Note: By default, the Job History and Activity Logs are retained for 90 days.

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/Arcserve/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>
```

Example:

The value must be in minutes. The maximum limit for the UI timeout value is 60.

```
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/Arcserve/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/Arcserve/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_recoverysset_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/Arcserve/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/Arcserve/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/Arcserve/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.



The backup job may fail because it depends on the type of session at the destination (NFS or CIFS). The following modules are needed, but are not installed: NFS client, CIFS client

Follow these steps:

1. Log in to the Backup Server.
2. Open the `server.cfg` file:

```
/opt/Arcserve/d2dserver/configfiles/server.cfg
```

3. Add the following parameter:

```
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.

Skip CIFS and NFS Validation on Linux Backup Server

When you add or modify backup storage, the Backup Server validates if the CIFS or NFS is accessible on the Linux Backup server. If you want to skip this validation on the Linux Backup server, you can configure the `server.env` file.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Open the `server.env` file:

```
/opt/Arcserve/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:

```
export skip_validate_backup_storage_on_server=true
```

4. Restart the Backup Server.

Configure the Default Temporary Folder

When you back up Linux Nodes, the default **/tmp** folder is used to store the required binary, temporary snapshot data, and debug logs. The **/tmp** folder must have enough empty space and the necessary permissions to execute the binaries. To change the default path on Linux Nodes, you can configure the `server.env` file and specify the new paths.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Open the `server.env` file:

```
/opt/Arcserve/d2dserver/configfiles/server.env
```

Note: If the file is not present, create the `server.env` file.

3. To configure the Linux Node agent execution path, add the following line in the `server.env` file:

```
export TARGET_BOOTSTRAP_DIR=<path>
```

Example: To deploy the Linux agent under the **/d2dagent** path, enter the following line:

```
export TARGET_BOOTSTRAP_DIR=/d2dagent
```

Note: By default, the agent is deployed and executed under the **/tmp** folder.

4. To configure the Linux Node debug log and temporary snapshot data store path, add the following line in the `server.env` file:

```
export TARGET_WORK_DIR=<path>
```

Example: To configure the debug logs and temporary snapshot data under **/d2d-agentlogs** path, enter the following line:

```
export TARGET_WORK_DIR=/d2dagentlogs
```

Note: By default, the agent is deployed and executed under the **/tmp** folder.

5. Restart the Backup Server.

```
/opt/Arcserve/d2dserver/bin/d2dserver restart
```

The default temporary folder is configured.

Configure the Snapshot Path for Backup Node

When you back up Linux Nodes, the default **/tmp** folder is used to store the disk snapshot file. The **/tmp** folder must have sufficient empty space. To change the snapshot path on Linux Nodes, you can configure a node-specific file and specify the new path.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Navigate to the **node** folder:

```
/opt/Arcserve/d2dserver/configfiles/node
```

Note: If the folder is not present, create the folder.

The **node** folder contains the `<node_name>.cfg` file. Each node has its own `cfg` file.

3. To configure the Linux Node snapshot path, add the following line in the specific `<node_name>.cfg` file:

```
target_snapshot_dir=<path>
```

Note: If the `<node_name>.cfg` file is not present, create the file.

Example: If the node name is **d2dbackupnode** and you want to store the snapshot under the **/d2dsnapshot** path, then open the following `cfg` file:

```
/opt/Arcserve/d2dserver/configfiles/node/d2dbackupnode.cfg
```

Add the following line:

```
target_snapshot_dir=/d2dsnapshot
```

The snapshot folder on target node is configured.

Configure the Hyper-V Server Connection Information for Instant VM

When you submit Instant VM jobs for Linux Nodes, the Backup Server tries to detect the Hyper-V server automatically. But, if the process fails, you can verify to ensure that the correct Hyper-V server connection information is used.

Linux IVM supports Hyper-V with SMB 2.0 or above to avoid the vulnerabilities of SMB 1.0.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Navigate to to the following Hyper-V folder:

```
/opt/Arcserve/d2dserver/configfiles/hyperv
```

Note: If the folder is not present, create the folder. The Hyper-V folder contains the <upper_case_hyperv_server_name>.cfg file. Each Hyper-V server has its own cfg file.

3. To configure the Hyper-V connection information, add the following lines in the specific <upper_case_hyperv_server_name>.cfg file:

```
protocol=<HTTP/HTTPS>
```

```
port=<number>
```

Note: If the <upper_case_hyperv_server_name>.cfg file is not present, create the file.

For protocol and port number, access the target Hyper-V Server using the following command line:

```
winrm enumerate winrm/Config/Listener
```

For example: The target Hyper-V server name is ivm-hyperv and the WinRM on Hyper-V server is configured as HTTPS listening on port 5986, then open the following cfg file:

```
/opt/Arcserve/d2dserver/configfiles/hyperv/IVM-HYPERV.cfg
```

Add the following lines:

```
protocol=HTTPS
```

```
port=5986
```

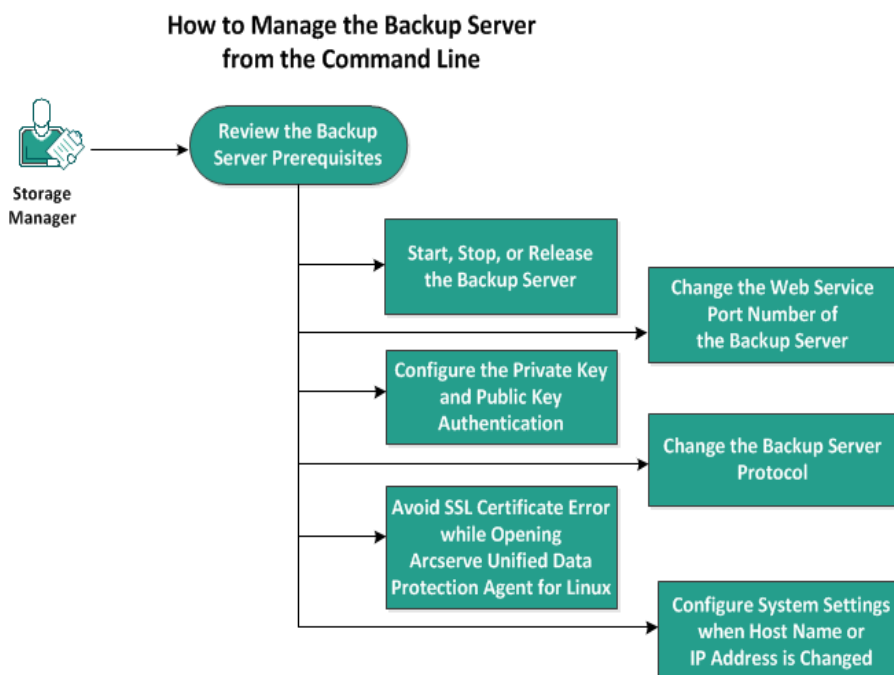
The connection information for Hyper-V server is configured.

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](#)
- [Start, Stop, or Release the Backup Server](#)
- [Change the Web Service Port Number of the Backup Server](#)
- [Configure the Private Key and Public Key Authentication](#)
- [Change the Backup Server Protocol](#)

- [Avoid the SSL Certificate Error While Opening Arcserve UDP Agent \(Linux\)](#)
- [Configure the System Settings When the Host Name or IP Address is Changed](#)

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](#) that 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/Arcserve/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/Arcserve/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/Arcserve/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" cli-
entAuth="false" sslProtocol="TLS" key-
storeFile="{catalina.home}/conf/server.keystore"
keystorePass="LinuxD2D"/>
```

3. Run the following command to restart the Backup Server:

```
/opt/Arcserve/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.
- When using sudo authentication, see [How to Configure Sudo User Account for Linux Nodes](#) for specific configuration.

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/Arcserve/d2dserver/configfiles/server_pub.key
```

4. Copy the private key file server to the following location:

```
/opt/Arcserve/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/Arcserve/d2dserver/configfiles/key.pass
```

6. Change the permission for the key.pass file using the following command:

```
chmod 600 /opt/Arcserve/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:

```
/<user_home>/.ssh/authorized_keys
```

Example: For a backup_admin, user_home is */home/backup_admin*

Example: */home/backup_admin/.ssh/authorized_keys*

9. (Optional) Run the following command in the node if SELinux blocks the authentication:

```
restorecon /<user_home>/.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/Arcserve/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/Arcserve/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.

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/Arcserve/d2dserver/bin/setenv

/opt/Arcserve/d2dserver/sbin/sqlite3 /opt/Arcserve/d2dserver/data/ARCserveLinuxD2D.db "update D2DServer set Name=' New_Hostname' where IsLocal=1"

/opt/Arcserve/d2dserver/sbin/sqlite3 /opt/Arcserve/d2dserver/data/License.db "update LicensedMachine set ServerName = 'New_Hostname' where ServerName = 'Old_Hostname' "
```

3. Rename the keystore file:

```
mv /opt/Arcserve/d2dserver/TOMCAT/conf/server.keystore /opt/Arcserve/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 /op-
```

```
t/Arcserve/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 /op-
t/Arcserve/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" cli-
entAuth="false" sslProtocol="TLS" key-
storeFile="${catalina.home}/conf/server.keystore"
keystorePass="YOUR_VALUE"/>
```

Example:

```
<Connector port="8014" protocol="HTTP/1.1" SSLEnabled="true"
maxThreads="150" scheme="https" secure="true" cli-
entAuth="false" sslProtocol="TLS" key-
storeFile="${catalina.home}/conf/server.keystore"
keystorePass="LinuxD2D"/>
```

6. Restart the central Backup Server.

```
/opt/Arcserve/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/Arcserve/d2dserver/bin/setenv
```

```
/opt/Arcserve/d2dserver/sbin/sqlite3 /opt/Arcserve/d2dserver/data/ARCserveLinuxD2D.db "update D2DServer set Name=' New_Hostname' where IsLocal=1"
```

3. Rename the keystore file:

```
mv /opt/Arcserve/d2dserver/TOMCAT/conf/server.keystore /opt/Arcserve/d2dserver/TOMCAT/conf/server.keystore.old
```

4. Create a keystore file using the following keytool Java command.

```
keytool -genkey -alias tomcat -keyalg DSA -keypass LinuxD2D -storepass LinuxD2D -keystore /opt/Arcserve/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/Arcserve/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/Arcserve/d2dserver/bin/d2dserver restart
```

7. Log into the central Arcserve D2D 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/Arcserve/d2dserver/bin/setenv  
  
/opt/Arcserve/d2dserver/sbin/sqlite3 /opt-  
t/Arcserve/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.

4. To update the host name, enter the following commands:

```
source /opt/Arcserve/d2dserver/bin/setenv

/opt/Arcserve/d2dserver/bin/d2drp --storagepath=Backup
Destination --node=New_Hostname

/opt/Arcserve/d2dserver/sbin/sqlite3 /opt/
t/Arcserve/d2dserver/data/ARCserveLinuxD2D.db "update Tar-
getMachine set Name=' New_Hostname' where Name=' Old_Hostname' "

/opt/Arcserve/d2dserver/sbin/sqlite3 /opt/
t/Arcserve/d2dserver/data/ARCserveLinuxD2D.db "update
JobQueue set TargetName=' New_Hostname' where JobType in
(1,3,4,5) and TargetName=' 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/Arcserve/d2dserver/bin/d2drp --storagepath=<mount
point> --node=New_Hostname
```

Note: If you use Local share, then the command is:

```
/opt/Arcserve/d2dserver/bin/d2drp --storagepath=<local path>
--node=New_Hostname
```

5. Log into the central Backup Server as a root user.
6. To update the license information, enter the following command:

```
/opt/Arcserve/d2dserver/sbin/sqlite3 /opt/
t/Arcserve/d2dserver/data/License.db "update LicensedMachine
set MachineName =' New_Hostname' where MachineName =' Old_Host-
name' "
```

The host name is configured to perform a backup without any error.

How to Add a User to Linux Backup Server Console Using Command Line

Using Arcserve UDP Agent for Linux, using the command line you can create a user who can act as the substitute of the root user on the Linux server. You can use the command line: `d2duser` to add such a user who can act when the root user is disabled.

The root user is disabled due to multiple reasons. For example, when you create the virtual machine on AWS EC2, by default the root is disabled.

- [Review the Prerequisites](#)
- [Add a User to Linux Backup Server Console Using Command Line](#)

Review the Prerequisites

Consider the following prerequisite or consideration before you add the user:

- You have the root login credentials to the Backup Server.
- Only the root user can execute the command line: `d2duser`.

Add a User to Linux Backup Server Console Using Command Line

You can use the command line `d2duser` to add a user who can act as replacement of the root user, when required.

Follow these steps:

1. Log on to the Backup Server as a root user.
2. Navigate to `/opt/Arcserve/d2dserver/configfiles`, and open the file: `server.cfg`.

Note: If a file with that name does not exist, create a new file by that name and add the following content to the file:

ui_login_use_udp_user= true|false

Lets you create the user that acts as the default user in the absence of root user when you login to the server. You can select **true** for this option.

ui_login_user_password_min_length = 6

Lets you decide the minimum password length. You can modify the default value 6, if required.

login_failure_time_to_lock_user = 9

Lets you decide after how many consecutive login failures, the user account is locked. You can modify the default value 9, if required.

3. Navigate to `/opt/Arcserve/d2dserver/bin`, and locate the `d2duser` command line.
4. Enter `./d2duser` to view the usage for this command line:

```
d2duser --action=<add|delete|lock|unlock|passwd> --user-name=<username>
```

5. Enter the following details on the `d2duser` command line:

d2duser --action=add --username=arcserve

Lets you add a user with the name `arcserve`. When you press enter, you are asked to enter a password, and then enter again to confirm.

d2duser --action=delete --username=arcserve

Lets you delete the user `arcserve`.

d2duser --action=lock --username=arcserve

Lets you lock the user `arcserve`.

d2duser --action=unlock --username=arcserve

Lets you unlock the user `arcserve`.

d2duser --action=passwd --username=arcserve

Lets you change the password for the user arcserve.

d2duser --action=list

Lets you view the list all the users.

6. From the browser, open the Linux Backup Server console page.
7. Verify if the displayed default user is the one that you just added.
8. Login using that username and password.

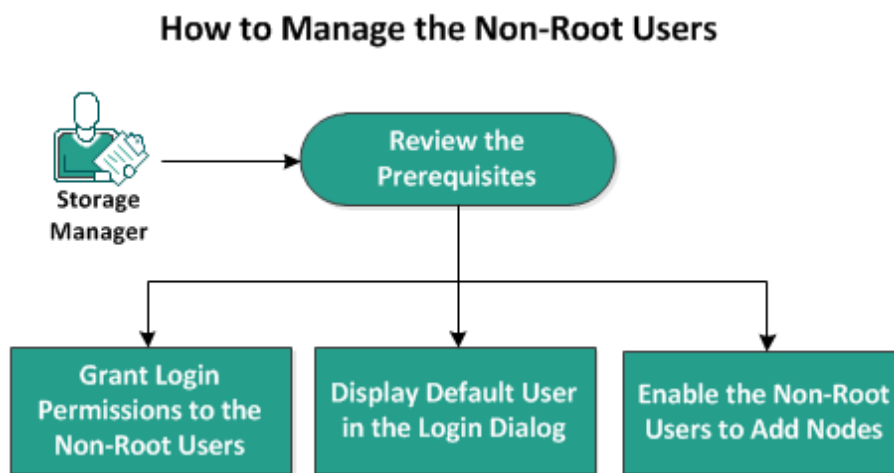
Successful login confirms that the user is created.

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](#)
- [Grant Login Permissions to the Non-Root Users](#)
- [Display the Default User in the Login Dialog](#)
- [Enable the Non-Root Users to Add Nodes](#)

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](#) that 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/Arcserve/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/Arcserve/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.



Follow these steps:

1. Log into the Backup Server as a root user.
2. Open the *server.cfg* file from the following location:

```
/opt/Arcserve/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 Configure Sudo User Account for Linux Nodes

You can use sudo to configure regular user accounts for performing backup and restore tasks. For sudo accounts, all configurations are related to Linux nodes. When the sudo account is properly configured, you can use the sudo account similar to a normal root account in all user interfaces. Using the sudo account, you can perform tasks such as add nodes, backup nodes, restore files. Configure the sudo according to the specific Linux Distribution document.

Perform these tasks to manage the sudo users:

- [Review the Prerequisites](#)
- [Modify the Default Sudo Settings in SUSE](#)
- [Configure sudo in Debian](#)
- [Configure Sudo for Authorization Without Password When Using SSH Public Key Authentication](#)
- [Configure Sudo to Allow Only Backup Agent Process](#)

Review the Prerequisites

Consider the following prerequisites before you manage the non-root users:

- You have the root login credentials of the Linux Node.
- You have properly configured the sudo permission for the desired user.
 - ♦ Verify that the sudo user is allowed to run at least these programs: d2d_ea and ln. For example, if the user name is backupadmin, the sudo configuration example is: *backupadmin ALL=(ALL) /usr/bin/d2d_ea,/usr/bin/ln.*
 - ♦ Verify if the sudo user is allowed to preserve at least following environment variables:

HOSTNAME	USERNAME	LANG	LC_ADDRESS
LC_CTYPE	LC_COLLATE	LC_IDENTIFICATION	LC_MEASUREMENT
LC_MESSAGES	LC_MONETARY	LC_NAME	LC_NUMERIC
LC_TIME	LC_ALL LANGUAGE	SSH_CONNECTION	CRE_ROOT_PATH
CRE_LOG_BASE_DIR	TARGET_BOOTSTRAP_DIR	TARGET_WORK_DIR	jobID

For example, if the user name is backupadmin, the sudo configuration examples are:

Default: backupadmin env_keep += "HOSTNAME USERNAME LANG LC_ADDRESS LC_CTYPE"

Default: backupadmin env_keep += "LC_COLLATE LC_IDENTIFICATION LC_MEASUREMENT"

Default: backupadmin env_keep += "LC_MESSAGES LC_MONETARY LC_NAME LC_NUMERIC LC_TIME LC_ALL LANGUAGE"

Default: backupadmin env_keep += "SSH_CONNECTION CRE_LOG_BASE_DIR jobID TARGET_BOOTSTRAP_DIR CRE_ROOT_PATH TARGET_WORK_DIR"

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

Modify the Default Sudo Settings in SUSE

By default, SUSE requires root password instead of the user password for authorization. A Sudo authentication does not work in Linux Backup Server because the Backup Server uses the user credentials for authorization. You can modify the default sudo settings to authorize the use of user credentials.

Follow these steps:

1. Log in to the Linux node as root user.
2. Open the `/etc/sudoer` file or run the `visudo` command.
3. Type a comment on the settings as shown in the following example:

Example:

```
#Defaults targetpw # ask for the password of the target user
i.e. root

#ALL ALL=(ALL) ALL # WARNING! Only use this together with
'Defaults targetpw'!
```

4. Verify the sudo command line now requires a user password instead of the root password for authorization.

You have successfully modified the default sudo settings.

Configure sudo in Debian

By default, root account is not enabled to login to Debian. As a result, a sudo authentication is needed when you add Debian Linux as a Linux Node.

Follow these steps:

1. Log in to the Linux node and switch to root using `su` command.
2. If sudo is not installed, install sudo package using the following command:

```
apt-get install sudo
```

3. Add an existing user with `id=user` to `group=sudo`:

Example:

```
adduser user sudo
```

or create a new user with sudo

```
adduser user
```

```
adduser user sudo
```

4. Log into the user shell, and type the following command to verify that the user is authorized:

```
sudo -v
```

You have successfully configured sudo in Debian.

Configure Sudo for Authorization Without Password When Using SSH Public Key Authentication

When using SSH public key authentication, the Linux Backup Server does not store the user credentials. You can configure Sudo to allow authorization without any password.

Follow these steps:

1. Log in to the Linux node as root user.
2. Open the `/etc/sudoer` file or run `visudo` to edit the configuration file.
3. Navigate to the configuration line for the specified user and add the 'NOPASSWD' option.

For example, if the user name is backupadmin, add the 'NOPASSWD' option as shown in the following example:

Example: backupadmin ALL=(ALL) NOPASSWD: /usr/bin/d2d_ea,/user/bin/ln

4. Log in to the user shell and type the following command to verify that the authorization does not require any password:

```
sudo -v
```

You have successfully configured Sudo for authorization without password when you are using the SSH public key configuration.

Configure Sudo to Allow Only Backup Agent Process

When the user is only permitted to use limited commands under sudo, manual installation of backup agent program is required. For backup jobs to run, sudo permission is required for the *d2d_ea* process.

Follow these steps:

1. Log in to the Linux node as root user.
2. Open the **/etc/sudoer** file or run *visudo* to edit the configuration file.
3. Navigate to the configuration line for the specified user and add `'/usr/bin/d2d_ea'` to the allowed commands configuration item.

For example, if the user name is backupadmin, add `'/usr/bin/d2d_ea'` as shown in the following example:

Example: backupadmin ALL=(ALL) /usr/bin/d2d_ea

4. Determine whether the backup source node is 32 or 64 bit and locate the correct binary on the backup agent server:
5. Copy the determined binary from step 4 to the backup source node as *d2d_ea* and then place at `'/usr/bin/d2d_ea'`.

For 32 bit: `/opt/Arcserve/d2dserver/sbin/ea.32`

For 64 bit: `/opt/Arcserve/d2dserver/sbin/ea.64`

6. Run the following command to verify execution permission:

```
chmod +x /usr/bin/d2d_ea
```

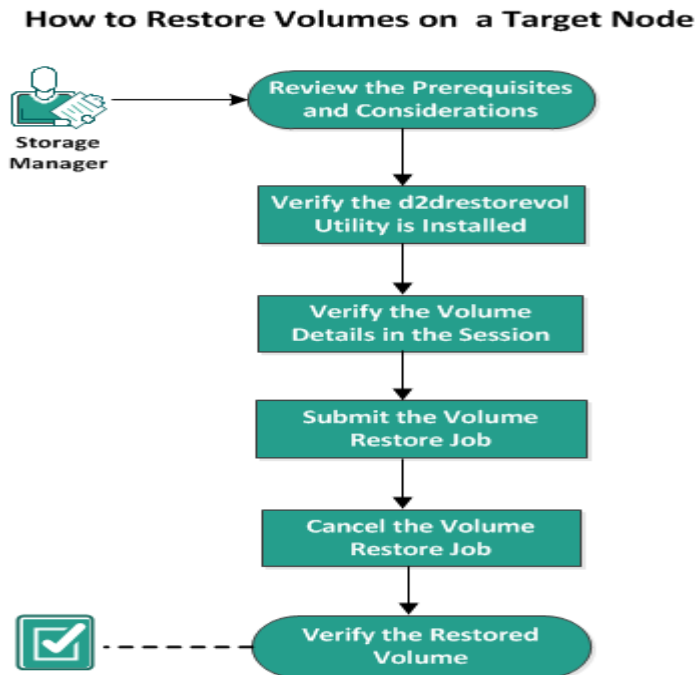
You have successfully configured Sudo for allowing only the backup agent process.

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](#)
- [Verify the d2drestorevol Utility is Installed](#)
- [Verify the Volume Details in the Session](#)
- [Submit the Volume Restore Job](#)
- [Cancel the Volume Restore Job](#)
- [Verify the Restored Volume](#)

Review the Prerequisites and Considerations

Review the following prerequisites before restoring volumes:

- You have a valid backup session to perform a restore.
- Volume restore supports session generated by Linux agent-based plans or jobs.
- 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. If the session is backed up to a RPS data store, first find the shared path in data store details. Then, mount the shared path to a local mount point and 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](#) that 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/Arcserve/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.cfg file is not present, create the file.

```
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. If the recovery points are in a local or shared folder, 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.

3. If the recovery points are in a RPS data store, use the following command to verify the volume information:

```
d2drestorevol --command=info --storage-path=<rps_path> --
node="<node_name>[UUID_number]" --rp=<recovery_point> --rps-
host=<host_name> --rps-user=<user_name> --rps-pw=<rps_pass-
word> --rps-protocal=<internet_secutity_protocol> --rps-port-
t=<port_number> --rps-dedup
```

The following command is an example for a deduplication enabled data store:

```
d2drestorevol --command=info --storage-path=/root/rpsshare -
-node="xx.xx.xx.xx[11111aa-22bb-33cc-yyyy-4c4c4c4c]" --
rp=VStore/S0000000001 --rps-host=machine_name --rps-user-
```

```
=administrator --rps-pw=***** --rps-protocol=https --rps-  
port=8014 --rps-dedup
```

--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 in the following format:

<node name>[<uuid>]

--rp

Specifies the recovery point or recovery session that you want to restore from an RPS data store. Typically, a recovery point session from a RPS data store must be specified in the following format:

VStore/S00000000X, where X is a numeric value

-- rps-host

Specifies the host name of the RPS where the recovery session are stored.

-- rps-user

Specifies the user name to access the RPS host.

-- rps-pw

Specifies the password to access the RPS host.

-- rps-protocol

Specifies the protocol for the RPS host. The protocol is ether http or https.

-- rps-port

Specifies the port number of the RPS host.

-- rps-dedup

Specifies the data store has deduplication enabled. This parameter is required only when the data store has enabled deduplication.

-- ds-share-folder

Specifies the shared path of the data store. This parameter is required only when the data store has disabled deduplication.

-- ds-user

Specifies the user name for accessing the shared path of the data store.

-- ds-user-pw

Specifies the user name for accessing the shared path of the data store.

-- ds-pw

Specifies the data encryption password if the data store also has enabled encryption.

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. If the recovery points are in a local folder or shared network, 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. If the recovery points are in a RPS data store, submit the restore job using the following command:

```
d2drestorevol --command=restore --storage-path=<local_path>  
--node=<node_name> --rp=<recovery_point> --source-volume=  
e=<source_volume> --target-volume=<target_volume> [--encrypt-  
ion-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 in the following format:

<node name>[<uuid>]

--rp

Specifies the recovery point or recovery session that you want to restore from a data store in RPS. Typically, a recovery point session from a RPS data store must be specified in the following format:

VStore/S00000000X, where X is a numeric value

--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

-- rps-host

Specifies the host name of the RPS where the recovery sessions are stored.

-- rps-user

Specifies the user name to access the RPS host.

-- rps-pw

Specifies the password to access the RPS host.

-- rps-protocol

Specifies the protocol for the RPS host. The protocol is ether http or https.

-- rps-port

Specifies the port number of the RPS host.

-- rps-dedup

Specifies the data store has enabled deduplication. This parameter is required only when the data store has enabled deduplication.

-- ds-share-folder

Specifies the shared path of the data store. This parameter is required only when the data store has disabled deduplication.

-- ds-user

Specifies the user name for accessing the shared path of the data store.

-- ds-user-pw

Specifies the password for accessing the shared path of the data store.

-- ds-pw

Specifies the data encryption password if the data store also has enabled encryption.

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.

4. (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.

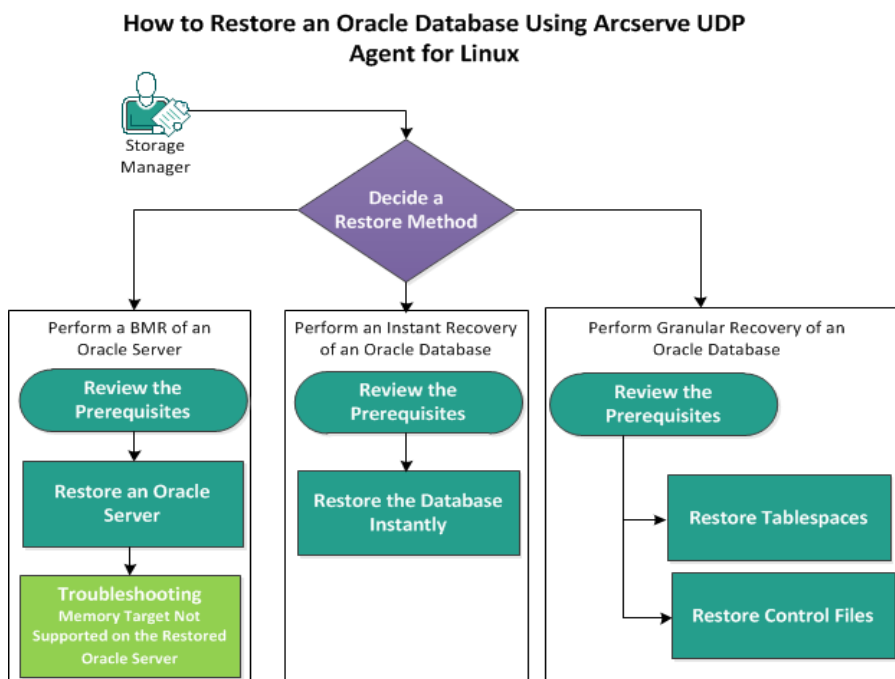
The volume restore job is verified.

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](#)
- [Perform an Instant Recovery of an Oracle Database](#)
- [Perform Granular Recovery of an Oracle Database](#)

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](#) that 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](#) that 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=e-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;
```

Example: #. /d2drestorevol --command=restore --storage-path=/CRE_ROOT --node-e=rh63-v2 --rp=last --source-volume=/opt/oracle --target-volume=/dev/sdb1 --mount-target=/opt/oracle --quick-recovery

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](#) that 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:
 - o Select Restore to original location.
 - o Enter the hostname or IP address of the target Oracle Server.
 - o Enter the root user name and the password of the target Oracle Server.
 - o 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.
```

10. Media recovery complete.
11. Open the database with the RESETLOGS clause after completing the recovery process.

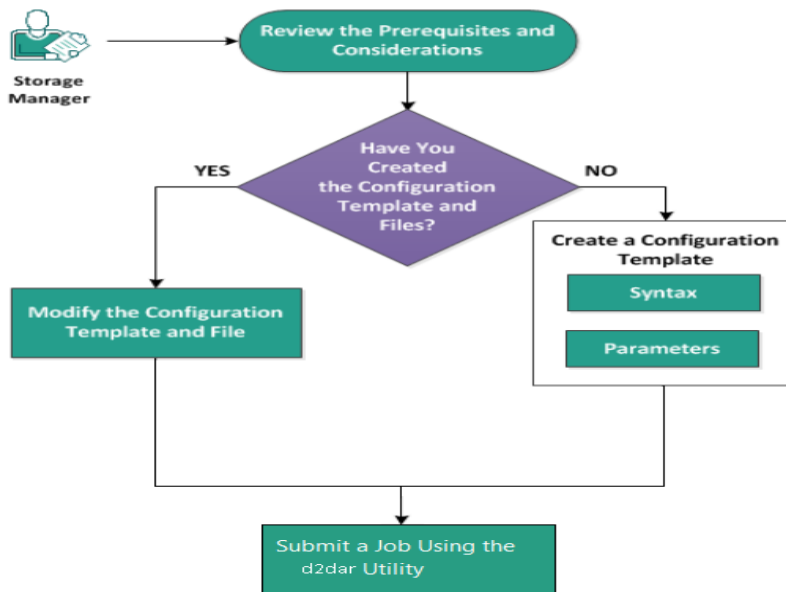
```
SQL>alter database open resetlogs;
```

The control files are successfully recovered.

How to Run Assured Recovery Test from the Command Line

You can run the Assured Recovery Test from the command line of the Backup Server using d2dar utility. The d2dar utility automates the process of performing an Assure Recovery Testing for specified backed up sessions.

The following diagram displays the process to run the Assured Recovery test from the command line using the d2dar utility:



Perform these tasks to run Assured Recovery test:

- [Review the Prerequisites and Considerations](#)
- [Create a Configuration Template](#)
- [Modify the Configuration Template and File](#)
- [Submit a Job Using the d2dar Utility](#)

Review the Prerequisites and Considerations

Review the following consideration before you perform the Assured Recovery test:

- The following versions of hypervisors are supported for the Assured Recovery test using the d2dar utility:
 - ◆ VMware vCenter/ESX(i) 5.0 or later
 - ◆ Windows Hyper-v Server 2012 or later

Note: To know more about supported Linux virtual machines on Hyper-v, click the [link](#).

- The Assured Recovery test is performed only from the command line. The option is not available on the user interface.

Create a Configuration Template

You can create a configuration file to allow the `d2dar` command run the Assured Recovery test according to the parameters specified in the file.

Syntax

```
d2dar --createtemplate=<cfg_file_path>
```

The `d2dutil --encrypt` utility encrypts the password and provides an encrypted password. You must use this utility to encrypt all your passwords.

Method 1

```
echo 'string' | ./d2dutil --encrypt
```

string is the password that you specify.

Method 2

Type the `d2dutil --encrypt` command and specify your password. Press **Enter** and you can view the result on your screen. In this method, the password that you enter is not echoed on the screen.

Follow these steps:

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

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

3. Create the configuration template using the following command:

```
#!/d2dar --createtemplate=<cfg_file_path>
```

<cfg_file_path> indicates the location where the configuration template is created.

4. Open the configuration template and update the following parameters in the configuration template:

job_name

Specifies the name of the Assured Recovery job.

vm_name_prefix

Specifies the prefix for the VM that is created for the Assured Recovery job. The name of Assured Recovery VM is `vm_name_prefix + node name + timestamp`.

vm_type

Specifies the type of the hypervisor where you perform the Assured Recovery test. The valid types of hypervisors are hyper-v and esx.

vm_server

Specifies the address of the hypervisor server. The address is either the hostname or IP address.

vm_svr_username

Specifies the username of the hypervisor.

vm_svr_password

Specifies the password of the hypervisor. The password is encrypted using the d2dutil encryption utility.

vm_svr_protocol

Specifies the protocol of the hypervisor when you perform Assured Recovery on vCenter/ESX(i).

vm_svr_port

Specifies the port of the hypervisor when perform Assured Recovery on vCenter/ESX(i).

vm_sub_server

Specifies the ESX server name when you perform Assured Recovery on vCenter.

vm_datastore

Specifies the storage location for VM used by Assured Recovery test. The location is the datastore on the ESX(i) server when you perform Assured Recovery test on vCenter/ESX(i). The location should be local path on the Hyper-V server when you perform Assured Recovery on Hyper-V.

vm_resource_pool

Specifies the Resource Pool name when you perform Assured Recovery on vCenter/ESX(i)

timeout

Specifies the time for the Assured Recovery job during rebooting till the VM is ready to use. The unit time is in second.

vm_memory

Specifies the VM memory size. The unit size is in MB, and in multiples of 4.

vm_cpu_count

Specifies the VM CPU number.

run_after_backup

Specifies that the Assured Recovery job runs once or every time for the backup job that is defined by the parameter `backup_job_name`. The Assured Recovery job runs immediately for the specified backup job, when set as **no**, and runs every time after the specified backup job finishes, when set as **yes**.

Default: no

backup_job_name

Specifies the backup job name of nodes to perform the Assured Recovery job.

storage_type

Specifies the storage type for the backed up session. The valid types of the storage are `cifs`, `nfs`, and `rps`.

storage_location

Specifies the NFS or CIFS location.

storage_username

Specifies the user name for the CIFS location.

storage_password

Specifies the password for the CIFS location. The password is encrypted using the `d2dutil` encryption utility.

rps_protocol

Specifies the protocol of the recovery point server when you perform the Assured Recovery job for sessions in the recovery point server.

rps_hostname

Specifies the hostname of the recovery point server. The address is either the hostname or the IP address.

rps_username

Specifies the user name of the recovery point server.

rps_password

Specifies the password of the recovery point server. The password is encrypted using the `d2dutil` encryption utility.

rps_port

Specifies the port of the recovery point server.

Default value: 8014.

rps_datastore

Specifies the datastore name on the recovery point server.

encryption_password

Specifies the encrypted session password. The password is encrypted using the d2dutil encryption utility.

node_name_list

Specifies the name(s) for the node(s) where the Assured Recovery test runs. The names are separated by using ';'. If a name is not specified or left empty, all nodes with the same backup job name or in the same location run the Assured Recovery test.

recovery_point_date_filter

Specifies the date of the recovery point. The Assured Recovery test runs for the last recovery point before the specified date. If the date is not specified or left empty, the latest backed up session runs the Assured Recovery test.

gateway_vm_network

Specifies the VM network for the gateway server. The VM and the Backup server are in the same network.

gateway_guest_network

Specifies the network IP address type for the gateway server. The network is either dhcp or static.

gateway_guest_ip

Specifies the IP address for the gateway server, if you provide the static IP.

gateway_guest_netmask

Specifies the netmask for the gateway server, if you provide the static IP.

gateway_guest_gateway

Specifies the gateway for the gateway server, if you specify the static IP.

script_post_job_server

(Optional) Specifies the script to run after the job finishes on the Backup server.

script_ready_to_use

(Optional) Specifies the script to run when the target machine is ready to use on the Assured Recovery VM.

run_script_ready_to_use_timeout

Specifies the time for running the ready-to-use script specified by the `script_ready_to_use`. The unit time is in second.

Note: Parameters for session related information including *storage_type*, *storage_location*, *storage_username*, *storage_password*, *rps_protocol*, *rps_hostname*, *rps_username*, *rps_password*, *rps_port*, and *rps_datastore* are only required when *backup_job_name* is not specified.

5. Click **Save** and close the configuration template.

The configuration template is successfully created.

Modify the Configuration Template and File

If you already have the configuration template file, you can modify the file and run the Assured Recovery test with a different configuration. You do not need to create another configuration template. When you submit the job, a new job is added on the web interface. You can view the activity logs on the web interface.

Follow these steps:

1. Log on 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 according to your requirement.
3. Click **Save** and close the configuration template.
4. Click **Save** and close the global configuration file.

The configuration template is successfully modified.

Submit a Job Using the d2dar Utility

You can use the `d2dar` command to run the Assured Recovery test for the backed up session(s). After submitting, you can view the job from the web interface. During the Assured Recovery process if any requirement is not met, the command line displays an error. You can also view the activity log on the web interface.

Follow these steps:

1. Log on to the Backup Server as a root user.
2. Submit the Assured Recovery job using the following command:

```
#!/d2dar --template=cfg_file_path
```

How to Mount Recovery Point

Mount Recovery Point can share files in a Recovery Point through NFS or WebDAV and you can access these files by mounting the location in the Linux Server.

Perform these tasks for Mount Recovery Point:

- [Review the Prerequisites](#)
- [Specify the Recovery Point for Mount Recovery Point](#)
- [Specify Settings for the Mount Recovery Point](#)
- [Create and Run the Mount Recovery Point Job](#)
- [Mount NFS share or WebDAV share on Linux Server](#)

Review the Prerequisites

Consider the following prerequisites before you Mount Recovery Point:

- You have a valid recovery point and the encryption password, if any.
- If you want to Mount Recovery Point by WebDAV, make sure package davfs2 has been installed in the Linux Server.
- Review the [Compatibility Matrix](#) which provides the supported operating systems, databases, and browsers.

Specify the Recovery Point for Mount Recovery Point

Each time that you perform a backup, a recovery point is created. When you want to access files in the recovery point.

Follow these steps:

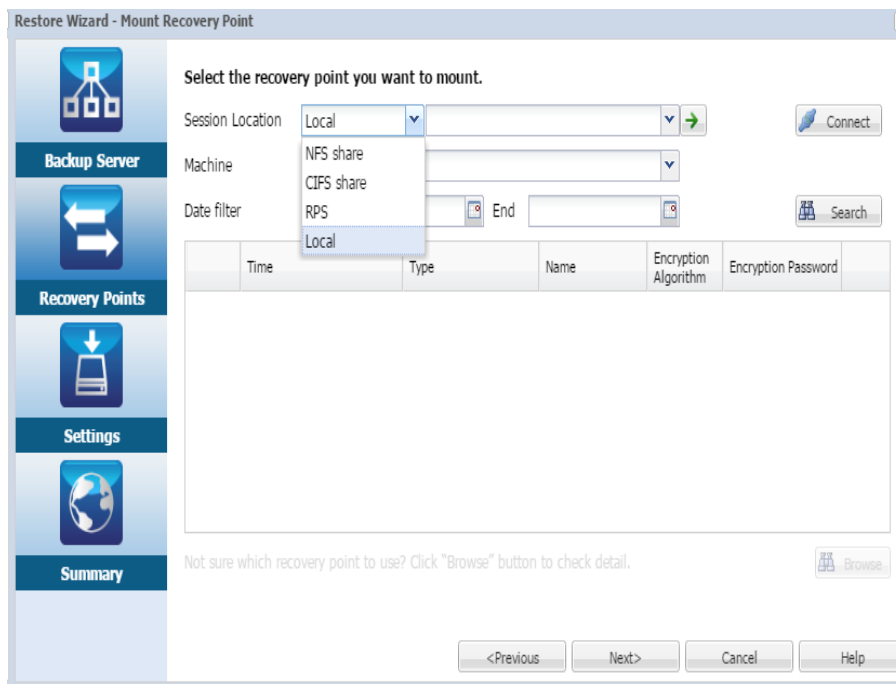
1. Open the Arcserve UDP Agent (Linux) web interface.
2. Click **Restore** from the **Wizard** menu and select **Mount Recovery Point**.

Restore Wizard – Mount Recovery Point 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.

3. Click **Next**.

The **Recovery Points** page of the **Restore Wizard** opens.



4. Select either a **CIFS share/NFS share/RPS server/Local** from the Session Location drop-down list.
5. Follow one of the following steps depending on your session location:

For CIFS share/NFS share/Local

Specify the full path of the CIFS share/NFS share/Local and click **Connect**.

All the machines are listed in the **Machine** drop-down list.

Note: If you select the CIFS share option, specify the user name and password.

For RPS server

- a. Select RPS server and click **Add**.

The **Recovery Point Server Information** dialog opens.

- b. Provide the RPS details and click **Load** button.
- c. Select the data store from the drop-down list and click **Yes**.

The Recovery Point Server Information dialog closes and you see the wizard.

- d. Click **Connect**.

All the machines are listed in the Machine drop-down list.

- e. Select the machine from the drop-down list.

All the recovery points from the selected machine appear below the **Date Filter** option.

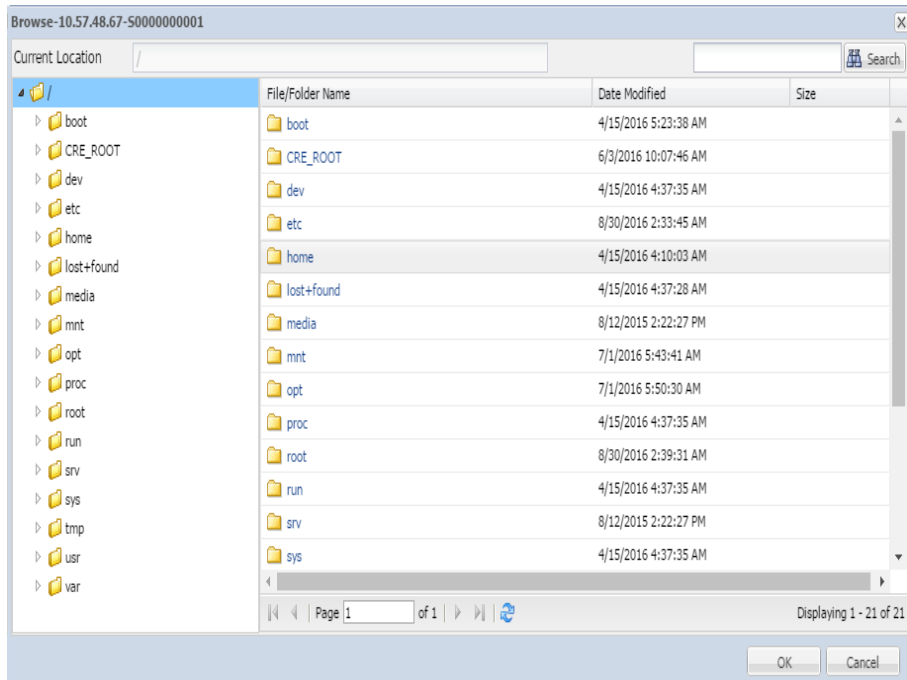
6. Apply the date filter to display the recovery points that are generated between the specified date and click **Search**.

Default: Recent two weeks.

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

7. Click Browse to view the recovery point.

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



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.

8. Click **OK**.

The **Browse-<node name>-<session number>** dialog closes and you return to the Recovery Points page.

9. Click **Next**.

The **Settings** for Mount Recovery Point pages opens.

Specify the Settings for Mount Recovery Point

Specify the settings for Mount Recovery Point to choose the proper share method.

Follow these steps:

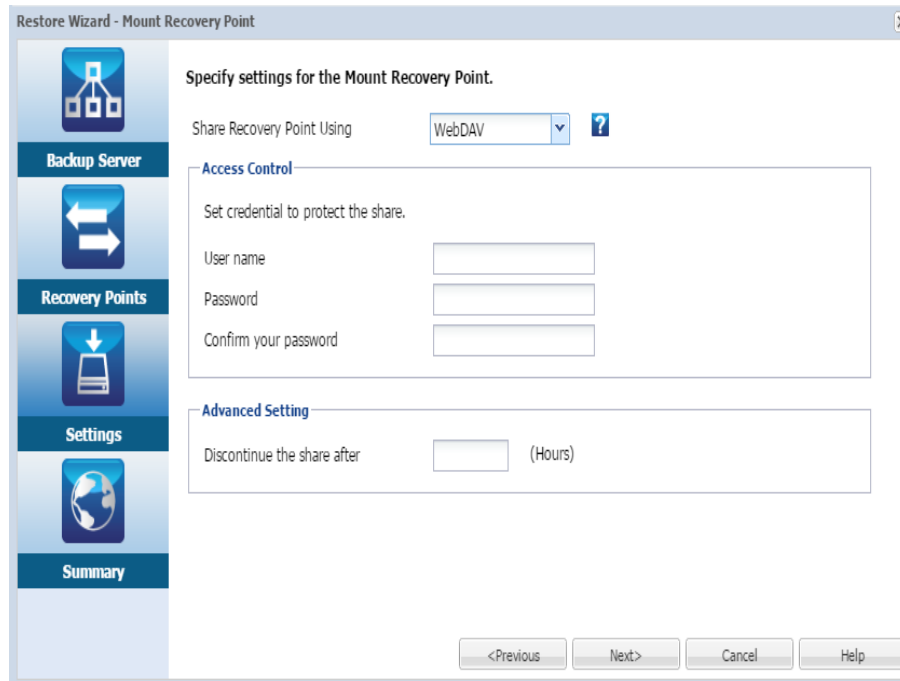
1. To Mount the Recovery Point through NFS, follow these steps:
 - a. Select **NFS** from share method drop-down list.

The files in the recovery point will be shared via NFS. And you can mount the NFS share on any machine which can access the Linux Backup Server.

- b. (Optional) Enter **NFS share option** according to your requirement. Consult man page for exports, candidate options, and valid format. Leave it empty if you don't need access control.
 - c. Enter **Hour** for time that share will be discontinued after the specified hour(s).
If you enter 0 for this field, the share could be access forever.
 - d. Click **Next**.
Summary page for Mount Recovery Point Job opens.
2. To Mount Recovery Point through WebDAV, follow these steps:

- a. Select **WebDAV** from share method drop-down list.

The files in the recovery point will be shared via WebDAV. And you can mount the WebDAV share using mount.davfs. This is the recommended method when you need access the share across internet.



- b. Enter **Username**, **Password**, and re-type your password to **Confirm your password** for Access Control.

Please remember the User name and Password, as they will be required for you access the mounted recovery point.

- c. Enter **Hour** for time that share will be discontinued after the specified hour(s).

If you enter 0 for this field, the share could be access forever.

If the time reaches the specified hour, the mounted recovery point could not be accessed.

- d. Click **Next**.

The Summary page for Mount Recovery Point Job opens.

Create and Run the Mount Recovery Point Job

You can create and run the Mount Recovery Point Job to access the files in the specified Recovery Point. Verify the configuration information before you submit the job. If needed, you can go back and change the settings on the wizard.

Follow these steps:

1. Verify the Mount Recovery Point details on the **Summary** page.
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 Mount Recovery Point job is successfully created and run.

Mount NFS or WebDAV Share on Linux Server

You can access the mounted recovery point after the **Job Phase** in the **Job Status** tab is **Sharing recovery Point**.

Follow these steps:

1. Get the **Job ID / Job Name** of the Mount Recovery Point Job in the **Job Status** tab.
2. Filter the activity logs for the Mount Recovery Point by **Job ID / Job Name** in the **Activity Log** page by using **Filter** tools.

Type	Job Id	Job Name	Time	Node Name	Message
Info	30	NFS	8/30/2016 3:12:40 AM	10.57.61.80	The recovery point is shared successfully.
Info	30	NFS	8/30/2016 3:12:40 AM	10.57.61.80	The recovery point share to continue for 1 hour(s).
Info	30	NFS	8/30/2016 3:12:40 AM	10.57.61.80	The script reports: Please access the share directory using NFS share: 10.57.48.45:/opt/Arcserve/d2dserver/tmp/d2d_share_path30
Info	30	NFS	8/30/2016 3:12:40 AM	10.57.61.80	Run mount recovery point job script NFS completed successfully at stage post_share.
Info	30	NFS	8/30/2016 3:10:08 AM	10.57.61.80	Run mount recovery point job script NFS completed successfully at stage pre_share.
Info	30	NFS	8/30/2016 3:10:07 AM	10.57.61.80	Recovery point is: 10.57.61.80[7d5af3ec-4865-49c2-2a67-75c1d3e1fff7]5000000005.
Info	30	NFS	8/30/2016 3:10:07 AM	10.57.61.80	Backup session location is Arcserve UDP Recovery Point Server (xiaoy02-v12r2-1), data store [1207B6s].
Info	30	NFS	8/30/2016 3:10:07 AM	10.57.61.80	Mount recovery point job name is: NFS.
Info	30	NFS	8/30/2016 3:10:07 AM	10.57.61.80	Mount recovery point job started successfully.

3. Get the shared directory for mounted recovery point displayed in the activity log.

Directory format when mounting via NFS :

`< d2dserver >:/opt/Arcserve/d2dserver/tmp/d2d_share_path<jobid>`

You can access the files in the recovery point by mounting the directory.

Example:

```
mount < d2dserver >:/opt/Arcserve/d2dserver/tmp/d2d_share_path<-
jobid> /mnt
```

Directory format when mounting via WebDAV:

<https://<d2dserver>:8014/share/<User Name>/>

You can access the files in the recovery point with web browser or by mounting the directory.

Example:

mount.dafs <https://<d2dserver>:8014/share/<User Name>/> /mnt

4. Enter username and password that you provided while submitting the Mount Recovery Point job.

Install the davfs package on the Linux Server

You can install the davfs package on the linux server.

- For Red Hat Linux, CentOS Linux or Oracle Linux

Follow these steps:

1. Get Extra Packages for Enterprise Linux (EPEL) for your Linux Server with the matched version from http://fedoraproject.org/wiki/EPEL#How_can_I_use_these_extra_packages.3F
2. Copy the downloaded EPEL package to the target Linux Server.
3. Install EPEL package by command as below

```
# yum install <package_path>/epel-release-<version_information>.rpm
```
4. Install davfs2 package by command as below.

```
# yum install davfs2
```

- For SuSE Linux 12 SP1

Follow these steps:

1. Log into the Linux Server.
2. Install davfs2 package by command as below.

```
# zypper addrepo  
# zypper refresh  
# zypper install davfs2
```

For more information, click the [link](#).

How to Enable Support for the Latest Ubuntu Kernel

Ubuntu updates its kernel on a regular basis resulting into the drivers sent with release being outdated. While turning off the automatic kernel updating process of the Ubuntu system helps, Arcserve also offers support for the updated kernels when required.

Important! Despite the best effort to support the latest Ubuntu kernel, major kernel change may still delay or cancel the corresponding drivers.

As a Storage manager, you can review below scenarios to enable using Arcserve UDP Agent (Linux) with the latest Ubuntu kernel:

- If your Arcserve UDP Agent (Linux) server has active internet connection, the updated drivers is downloaded and deployed unattended. You can use the software without additional efforts.
- If your Arcserve UDP Agent (Linux) server does not have internet access, you can download and deploy the updated driver package manually.
- If you have multiple Arcserve UDP Agent (Linux) servers, you can deploy the updated driver package on one server, and then configure the other server to use as a staging server.

Perform the following steps to deploy the updated driver package:

- [Review the Prerequisites](#)
- [Deploy the updated Ubuntu driver package manually](#)
- [\(Optional\) Using staging server for updating drivers](#)
- [\(Optional\) Configuring HTTP proxy](#)

Review the Prerequisites

Consider the following prerequisites:

- You have the root login credentials is available to log into the Backup Server.
- You have curl or wget is installed on the Backup Server.
- You have gpg is installed on the Backup Server.

Deploy the Updated Ubuntu Driver Package Manually

When your Arcserve UDP Agent (Linux) server does have internet access, you can still update the drivers by downloading and deploying manually.

Follow these steps:

1. Download the driver package and signature, please consult Arcserve support to get the url.
2. Log into the Backup Server as a root user.
3. Navigate to the location containing the downloaded package and initiate deploy using the following command:

```
# /opt/Arcserve/d2dserver/bin/d2dupgradetool deploy <folder containing the downloaded package>
```

The updated driver package is successfully deployed.

(Optional) Using Staging Server for Updating Drivers

When you have multiple Arcserve UDP Agent (Linux) servers that need to support the latest Ubuntu kernel, you can configure them to use one as the staging server. Ensure the staging server has the updated driver already deployed by using active internet connection or follows the instructions in task [Deploy the updated Ubuntu driver package manually](#). You can configure each Backup Server that needs the updated Ubuntu driver package.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Open and edit the configuration file:

```
# /opt/Arcserve/d2dserver/configfiles/auto_upgrade.cfg
```

3. Edit the following configuration items:

```
scheme=<http or https>
```

```
host=<the staging server address>
```

```
port=<agent server port, usually 8014>
```

The automated driver package update is successfully configured.

(Optional) Using Staging Server for Updating Drivers

You can configure proxy for Arcserve UDP Agent (Linux) to access Internet connection.

Follow these steps:

1. Log into the Backup Server as a root user.
2. Open and edit the configuration file:

```
# /opt/Arcserve/d2dserver/configfiles/auto_upgrade.cfg
```

3. Edit the following configuration items:

```
# /opt/Arcserve/d2dserver/configfiles/auto_upgrade.cfg
```

```
http_proxy=<proxy address>
```

```
proxy_user=<username>
```

```
proxy_password=<password>
```

The proxy is successfully configured.

Chapter 5: Troubleshooting

This section contains the following topics:

Arcserve UDP Agent (Linux) Fails to Install on Supported Servers	371
Arcserve UDP Agent (Linux) Displays an Operation Timeout Error	373
All Scheduled Jobs Fail When the System Time Is Changed to an Already Passed Value	374
Arcserve UDP Agent (Linux) Fails to Mount Linux Software RAID Devices	375
Arcserve UDP Agent (Linux) Fails to Download and Deploy Updated Ubuntu Drivers on SLES 11 and RHEL 6	376
A Paravirtual Machine (PVM) Displays a Black Screen on the Virtual Network Computing (VNC) Client Window When Booted Using a Live CD	377
The Backup Job Fails to Collect the BMR-related Information or the BMR Job Fails to Create a Disk Layout	378
How to Adjust the Disk Boot Sequence After a BMR Job on an Oracle VM Server	379
How to Restore the Previous Version of Backup Server	381
How to Start the Linux Backup Server Automatically on SLES15	382
How to Backup Debian 9.X EC2 Instances in AWS Cloud	383
SLES 10.X Does Not Start Successfully After a BMR	384
The d2drestorevm and d2dverify Jobs Fail on Oracle VM Server	385
ESXi Virtual Machine Fails to Start After BMR From a Physical Machine	386
Failed to Mount CIFS on the Server or Target Node	387
File-level restore in a host-based Linux VM fail due to an unsupported file system	388

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                Treate 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/Arcserve/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.

Arcserve UDP Agent (Linux) Fails to Download and Deploy Updated Ubuntu Drivers on SLES 11 and RHEL 6

Valid on some outdated versions of SUSE Linux Enterprise Server (SLES) 11 and Red Hat Enterprise Linux (RHEL) 6

Symptom

When I want to back up the Ubuntu node that has the updated kernel version, the backup job fails and the message in the activity log refers to failure in download and deployment of Ubuntu drivers.

Solution

Update the system packages and verify if curl or wget has the latest version.

Follow these steps:

1. Restart the target machine.
2. Run the following command:
On SUSE: zypper update wget curl
On RHEL: yum update wget curl
3. Run the failed backup job again on the Ubuntu node.

The Ubuntu driver is successfully updated.

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.



Configuration		Networks		Disks	
Name:	oel5.8_pvm_from_iso	Memory (MB):	1024	Processor Cap:	100
Status:	Running	Priority:	50	Mouse Type:	Default
Operating System:	Oracle Linux 5	Domain Type:	Xen PVM	Start Policy:	Start on best server
Keymap:	en-us	High Availability:	No		
Max. Processors:	1				
Processors:	1				
Max. Memory (MB):	1024				
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.

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.

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.

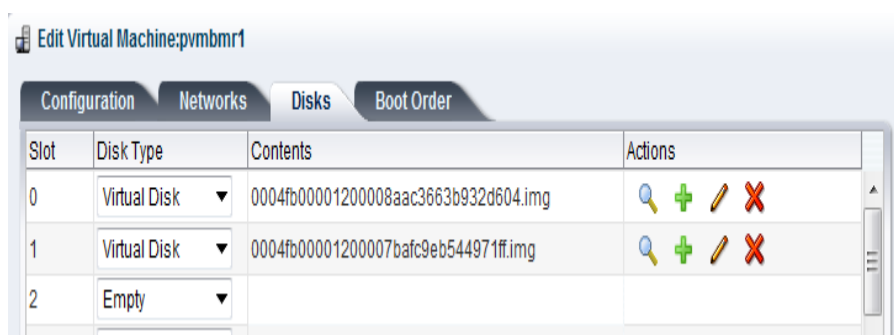
You will use the disk name and the slot number in the later steps.

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.

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/Arcserve/d2dserver.bak
```

Follow these steps:

1. Uninstall the existing Backup Server using the following command:

```
/opt/Arcserve/d2dserver/bin/d2duninstall
```

2. Install the previously installed version of the Backup Server.
3. Stop the Backup Server using the following command:

```
/opt/Arcserve/d2dserver/bin/d2dserver stop
```

4. Copy the old configuration files and database files to the d2dserver folder using the following command:

```
cp -Rpf /opt/Arcserve/d2dserver.bak/* /opt/Arcserve/d2dserver/
```

5. Start the Backup Server using the following command:

```
/opt/Arcserve/d2dserver/bin/d2dserver start
```

The previously installed version of the Backup Server is successfully restored.

How to Start the Linux Backup Server Automatically on SLES15

Symptom

When I install Arcserve UDP Agent (Linux) on SLES15, the Linux Backup Server does not start automatically after rebooting the SLES15 system.

Solution

After completion of installation, the Arcserve UDP Agent (Linux) service must be running on the Linux Backup Server to protect your machines. On SLES15, the service cannot start automatically after rebooting the system.

Use the following command to verify the status of service:

```
/opt/Arcserve/d2dserver/bin/d2dserver status
```

To troubleshoot, perform the following steps:

1. Install the following packages using the *zypper* command:

```
insserv
```

2. Run the following command:

```
systemctl enable start-d2d
```

```
systemctl start start-d2d
```

Note: If Linux Backup Server is still not installed, then first install the insserv package to avoid this issue.

How to Backup Debian 9.X EC2 Instances in AWS Cloud

Symptom

When backup is run for Debian 9.X EC2 instances in AWS cloud, the backup job fails without showing any specific errors.

Solution

When the Debian 9.X instances are created in AWS cloud and added to protect, absence of Perl modules may lead to the error. To resolve, install the packages using the following commands:

```
sudo apt update
```

```
sudo apt install apt-file
```

```
sudo apt-file update
```

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.

Failed to Mount CIFS on the Server or Target Node

Symptom

When I try to backup or restore using CIFS, CIFS fail to mount on the server or target node.

Solution

You must fulfill some requirements when mounting CIFS on a Linux machine.

Follow these steps:

1. Use the mount command on the server or the target node to verify the error.
2. Verify that when using a shared path exported from non-Windows systems, the letter case of the shared path characters matches the original path.
3. If the mount command returns an error, verify whether the time on the server or the target node is synchronized with the CIFS server.
4. If you do not find the error, add some options to the mount command to try again.

For example, add "sec=ntlm" when you receive the Permission denied error.

5. When you diagnose the error, follow these steps:

For failing to mount CIFS on the server

- a. Open the server.env file from the following location:
`/opt/Arcserve/d2dserver/configfiles/server.env`
- b. Add all the options to the file using the following command:
`export D2D_MOUNTOPTION=<options>`
- c. Save the file and restart the service.

For failing to mount CIFS on target node

- a. Open the .bashrc file from the user's home path.
Example: The location for a user is `/home/user/` and for root is `/root/`.
- b. Add all the options to the file using the following command:
`export D2D_MOUNTOPTION=<options>`
- c. Save the file.

Note: The .bashrc file is the recommended file here, but you could also modify other files such as `/ect/profile`, `/etc/bashrc`, and so on.

6. Verify that when using a shared path exported from non-Windows systems, the letter case of the shared path characters matches the original path.

File-level restore in a host-based Linux VM fail due to an unsupported file system

Symptom

When I perform a file-level restore for a host-based Linux VM, the restore wizard displays the following error message:

Not supported: reiserfs file system

The error occurs because you are trying to restore an unsupported file system.

Solution

You can restore the host-based Linux VM using one of the following ways:

- Use Arcserve UDP Agent (Linux) Live CD to perform the file-level restore because Live CD supports all types of file system. This is a convenient, but temporary solution. You can restore using a Live CD if you do not restore this node frequently.
- Another permanent method is that you must install the correct file system driver to support reiserfs or enable the corresponding driver which is already installed in your Backup Server.