



Setting Up the DR Series System on Veeam®

Quest Engineering
June 2017

Revisions

Date	Description
January 2014	Initial release
May 2014	Updated to add note to explain purpose of enabling dedupe on Veeam side.
July 2014	Updated to add workflow specific best practices
April 2015	Updated with Veeam 8.0 screenshots
June 2015	Updated cleaner recommendations
November 2015	Updated with Instant Recovery with DR series
April 2016	Updated with Veeam 9.0 screenshots and its features
November 2016	Updated for the DR Series system release 4.0.
June 2017	Updated for Rapid CIFS support

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2016 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

PRODUCT WARRANTIES APPLICABLE TO THE DELL PRODUCTS DESCRIBED IN THIS DOCUMENT MAY BE FOUND AT: <http://www.dell.com/learn/us/en/19/terms-of-sale-commercial-and-public-sector> Performance of network reference architectures discussed in this document may vary with differing deployment conditions, network loads, and the like. Third party products may be included in reference architectures for the convenience of the reader. Inclusion of such third party products does not necessarily constitute Dell's recommendation of those products. Please consult your Dell representative for additional information.

Trademarks used in this text:

Dell™, the Dell logo, PowerVault™, EqualLogic™ are trademarks of Dell Inc. Other Dell trademarks may be used in this document. Microsoft®, Windows®, Windows Server®, Internet Explorer®, MS-DOS®, Windows Vista® and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Veeam® and Veeam Backup & Replication™ are registered trademarks or trademarks of Veeam Software. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and/or names or their products and are the property of their respective owners. Dell disclaims proprietary interest in the marks and names of others.



Table of contents

Revisions	2
Executive summary	4
1 Installing and configuring the DR Series system	5
2 Creating and configuring target container(s) for Veeam	8
3 Setting up Veeam	11
4 Setting up DR Series native replication and restore from a replication target container	25
4.1 Building a replication relationship between DR Series systems	25
4.2 Restoring data from the target DR Series system	27
5 Configuring Rapid CIFS for Veeam	34
5.1.1 Windows prerequisites	34
5.1.2 Installing Rapid CIFS on a Backup Proxy Server	34
6 Using Instant VM Recovery with the DR Series system	37
6.1 Instant Recovery with ESX	37
6.2 Instant Recovery with Hyper-V Server	43
6.2.1 Enabling Instant Recovery for Hyper-V VM backups	43
6.2.2 Performing Instant Recovery	43
6.3 Finalizing Instant VM recovery	49
6.3.1 Migrating the VM to production	49
6.3.2 Terminating the Instant VM Recovery session	50
7 Backup copy operation	51
8 Setting up the DR Series system cleaner	55
9 Monitoring deduplication, compression, and performance	56



Executive summary

This document provides information about how to set up the DR Series system as a backup target for Veeam[®] Backup & Replication[™] software.

For additional information, see the DR Series system documentation and other data management application best practices whitepapers for your specific DR Series system at:

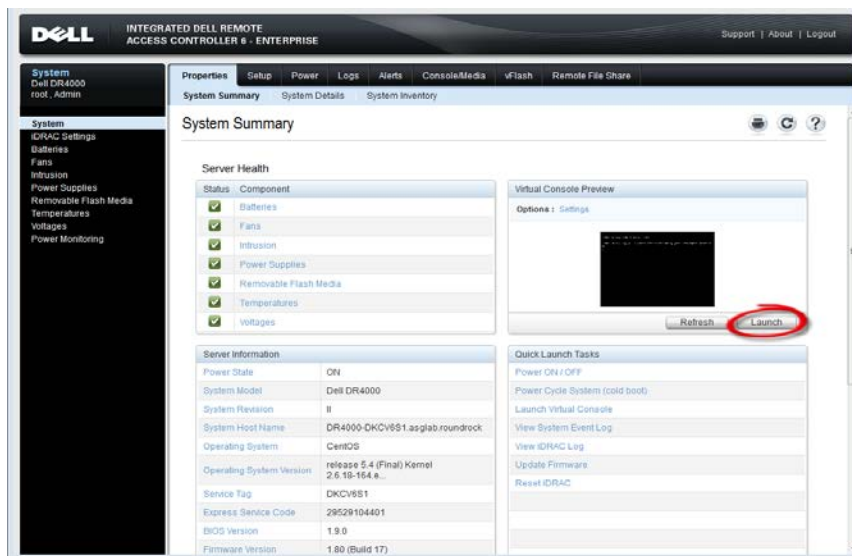
<http://support.quest.com/DR-Series>

Note: The DR Series system and Veeam screenshots used in this document may vary slightly, depending on the DR Series system firmware version and Veeam version you are using.



1 Installing and configuring the DR Series system

1. Rack and cable the DR Series system, and power it on.
2. In the *Dell DR Series System Administrator Guide*, see the following sections for information about using the iDRAC connection and initializing the appliance.
 - “iDRAC Connection”,
 - “Logging in and Initializing the DR Series system”
 - “Accessing iDRAC6/Idrac7 Using RACADM”
3. Log on to iDRAC using the default credentials (username: **root** and password: **calvin**) and either:
 - the default address **192.168.0.120**,
 - or the IP address that is assigned to the iDRAC interface
4. Launch the virtual console.



5. When the virtual console is open, log on to the system as the user **administrator** with the password **St0r@ge!** (The “0” in the password is the numeral zero).



6. Set the user-defined networking preferences as needed.

```
Would you like to use DHCP (yes/no) ?

Please enter an IP address:

Please enter a subnet mask:

Please enter a default gateway address:

Please enter a DNS Suffix (example: abc.com):

Please enter primary DNS server IP address:

Would you like to define a secondary DNS server (yes/no) ?

Please enter secondary DNS server IP address:
```

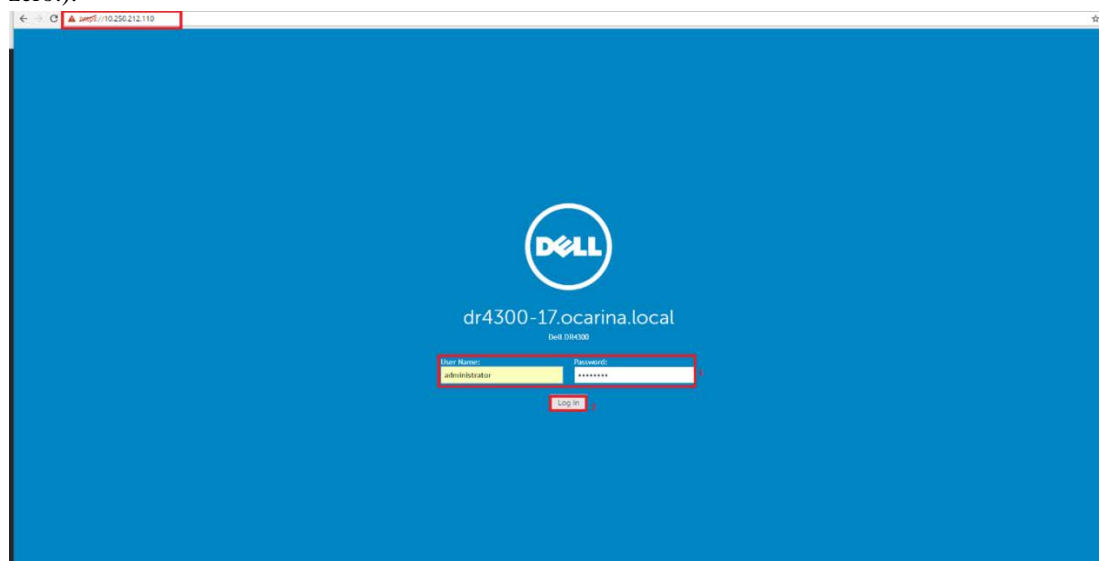
7. View the summary of preferences and confirm that it is correct.

```
=====
                        Set Static IP Address
=====

IP Address       : 10.10.86.108
Network Mask     : 255.255.255.128
Default Gateway  : 10.10.86.126
DNS Suffix       : idmdemo.local
Primary DNS Server : 10.10.86.101
Secondary DNS Server : 143.166.216.237
Host Name        : DR4080-5

Are the above settings correct (yes/no) ? _
```

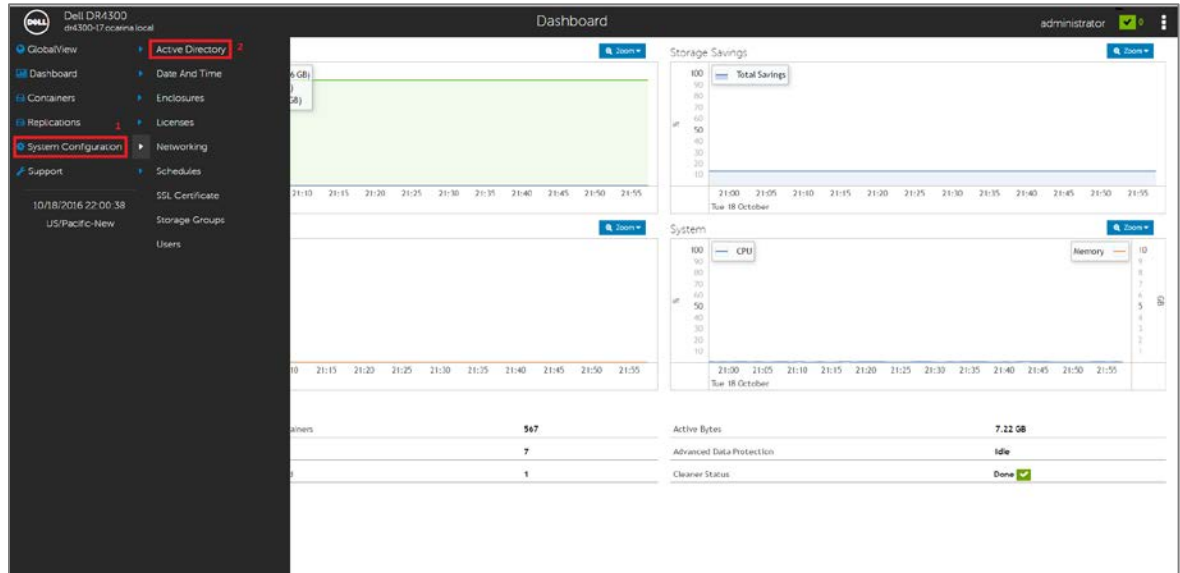
8. Log on to the DR Series System administrator console, using the IP address you just provided for the DR Series System, with the username **administrator** and password **St0r@ge!** (The “0” in the password is the numeral zero.).



9. Join the DR Series system to your Active Directory domain.

Note: If you do not want to add the DR Series system to Active Directory, see the *DR Series System Owner's Manual* for guest logon instructions.

- a. In the left navigation area, click **System Configuration > Active Directory**.
- b. On the Active Directory page, click **Join**.

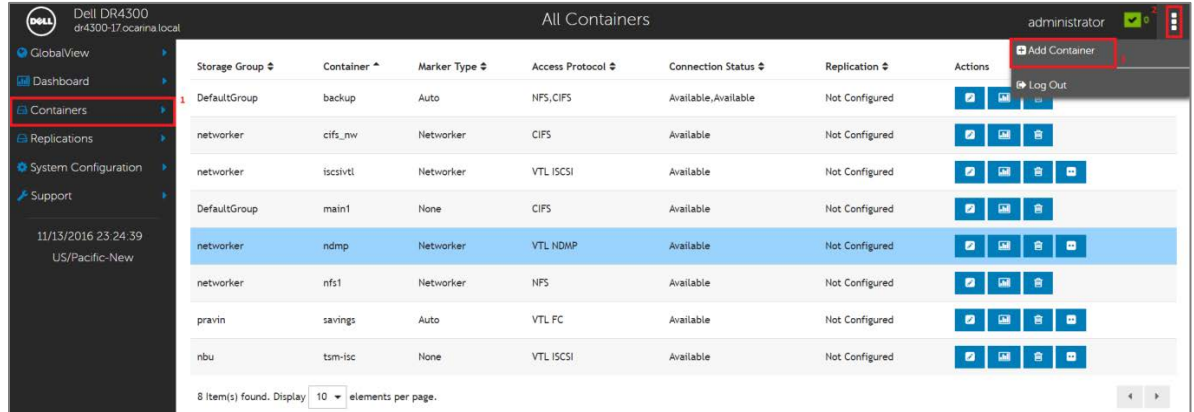


- c. Enter your **Active Directory** credentials and click **Join**.

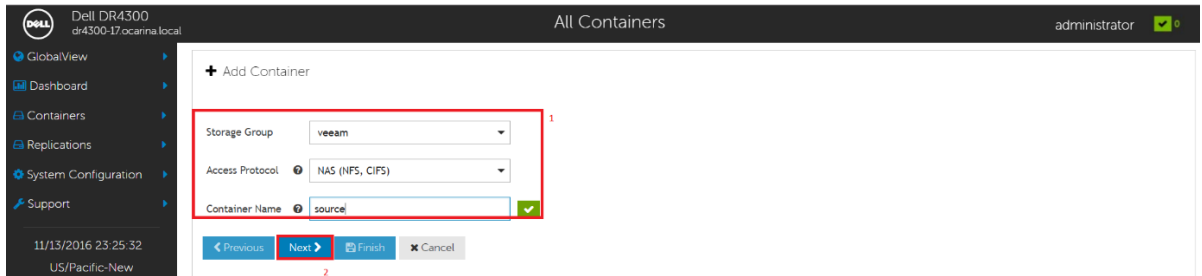


2 Creating and configuring target container(s) for Veeam


1. Select **Containers** in the left navigation area of the DR Series system GUI, and, on the Action Menu in the upper right corner, click **Add Container**.



2. Enter a Container Name, and, for Access Protocol, select **NAS (NFS, CIFS)**. Click **Next**.



3. Select **NFS** or **CIFS** as appropriate, select **Auto** for Marker Type, and click **Next**.



4. Enter the backup container information for NFS options as needed, and then click **Next**.

The screenshot shows the 'Add Container - source' configuration page in the Dell DR4300 interface. The page is titled 'All Containers' and shows the user 'administrator'. The left sidebar contains navigation links: GlobalView, Dashboard, Containers, Replications, System Configuration, and Support. The main content area is titled '+ Add Container - source'. It contains the following fields and options:

- NFS Options:** Radio buttons for 'Read Write Access' (selected) and 'Read Only Access'.
- Map Root To:** A dropdown menu with 'Administrator' selected.
- Client Access:** Radio buttons for 'Open (allow all clients)' (selected) and 'Create Client Access List'.
- Client FQDN or IP Address:** A text input field with a '+' button to add addresses.
- Allow Clients:** A checkbox.

At the bottom, there are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. The 'Next >' button is highlighted with a red box and labeled '2'.

5. Enter backup container information for CIFS options as needed, and then click **Next**.

The screenshot shows the 'Add Container - source' configuration page in the Dell DR4300 interface, specifically for CIFS options. The page is titled 'All Containers' and shows the user 'administrator'. The left sidebar contains navigation links: GlobalView, Dashboard, Containers, Replications, System Configuration, and Support. The main content area is titled '+ Add Container - source'. It contains the following fields and options:

- CIFS Client Access:** Radio buttons for 'Open (allow all clients)' (selected) and 'Create Client Access List'.
- Client FQDN or IP Address:** A text input field with a '+' button to add addresses.
- Allow Clients:** A checkbox.

At the bottom, there are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. The 'Next >' button is highlighted with a red box and labeled '2'.

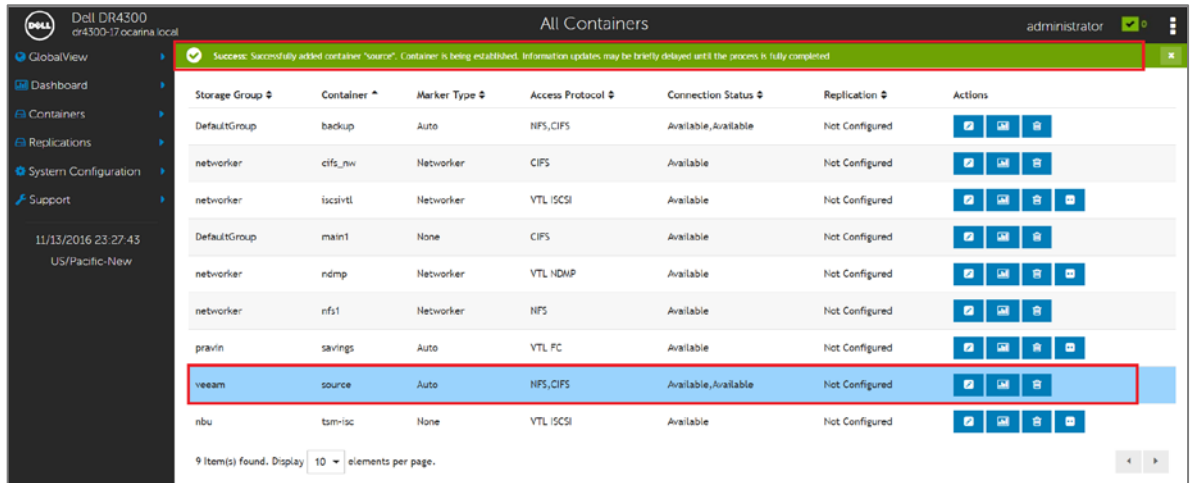
6. Confirm the settings and click **Save**.

The screenshot shows the 'Add Container - source' configuration page in the Dell DR4300 interface, displaying the final settings. The page is titled 'All Containers' and shows the user 'administrator'. The left sidebar contains navigation links: GlobalView, Dashboard, Containers, Replications, System Configuration, and Support. The main content area is titled '+ Add Container - source'. It contains the following fields and options:

- Storage Access Protocol:** A section with the following settings:
 - Storage Group: veeam
 - Access Protocol: NAS (NFS, CIFS)
 - Container Name: source
- Configure NAS Access & Marker:** A section with the following settings:
 - NAS Access Protocol: NFS, CIFS
 - Marker Type: Auto
- Configure NFS Client Access:** A section with the following settings:
 - NFS Options: Read Write Access
 - Map Root To: Administrator
 - Client Access: Open (allow all clients)
- Configure CIFS Client Access:** A section with the following settings:
 - Client Access: Open (allow all clients)

At the bottom, there are navigation buttons: '< Previous', 'Next >', 'Save', and 'Cancel'. The 'Save' button is highlighted with a red box.

7. Confirm that the container is added.



Success: Successfully added container "source". Container is being established. Information updates may be briefly delayed until the process is fully completed.

Storage Group	Container	Marker Type	Access Protocol	Connection Status	Replication	Actions
DefaultGroup	backup	Auto	NFS,CIFS	Available, Available	Not Configured	[Icons]
networker	cifs_nw	Networker	CIFS	Available	Not Configured	[Icons]
networker	iscsi_vtl	Networker	VTI iSCSI	Available	Not Configured	[Icons]
DefaultGroup	main1	None	CIFS	Available	Not Configured	[Icons]
networker	ndmp	Networker	VTI NDMP	Available	Not Configured	[Icons]
networker	nfs1	Networker	NFS	Available	Not Configured	[Icons]
pravin	savings	Auto	VTI FC	Available	Not Configured	[Icons]
veeam	source	Auto	NFS,CIFS	Available, Available	Not Configured	[Icons]
nbu	tsm-isc	None	VTI iSCSI	Available	Not Configured	[Icons]

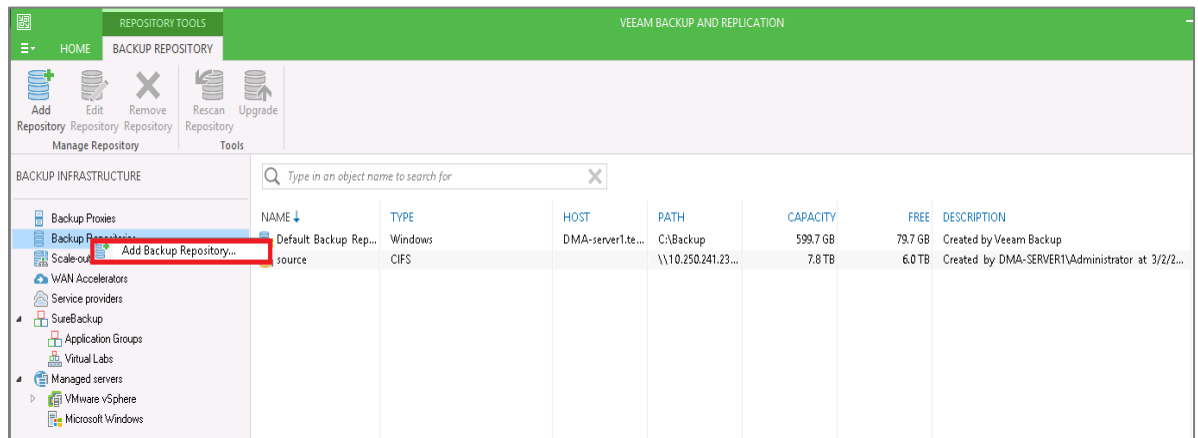
9 item(s) found. Display 10 elements per page.

3 Setting up Veeam

Notes:

To maximize the DR Series system and Veeam deduplication savings, you should use the exact settings in this guide for all the data being backed up. The backup data will change format completely when backup settings are changed. Hence, to get accurate savings numbers, all the data should be backed up with the same settings.

1. Open the **Veeam Backup & Replication** console.
2. In the **Backup Infrastructure** section, right-click **Backup Repositories**, and click **Add Backup Repository**.



3. Enter a name for the DR Series system container repository and click **Next**.

Edit Backup Repository

Type in a name and description for this backup repository.

Name

Name: source

Description: Created by RAMATEJA-W12-V6\Administrator at 9/24/2015 9:53 AM.

< Previous Next > Finish Cancel

4. For a CIFS container, select **Shared folder** as the type of backup repository, and click **Next**.

New Backup Repository

Choose type of backup repository you want to create.

Type

☐ **Microsoft Windows server (recommended)**
Microsoft Windows server with internal or directly attached storage. Data mover process running directly on the server allows for improved backup efficiency, especially over slow links.

☐ **Linux server (recommended)**
Linux server with internal, directly attached, or mounted NFS storage. Data mover process running directly on the server allows for more efficient backups, especially over slow links.

☒ **Shared folder**
CIFS (SMB) share. When backing up over slow links, we recommend that you specify a gateway server located in the same site with the shared folder.

☐ **Deduplicating storage appliance**
Advanced integration with EMC Data Domain, ExaGrid and HP StoreOnce. For basic integration, use the Shared folder option above.

< Previous Next > Finish Cancel

5. (For CIFS) In the **Shared folder** field, enter the DR Series system container share UNC path (or TCP/IP address to replace hostname), select the Gateway Server, and click **Next**.

The screenshot shows the 'Edit Backup Repository' dialog box with the 'Share' tab selected. The 'Name' field is empty. The 'Type' field is set to 'Share'. The 'Shared folder' field contains the UNC path '\\10.250.241.229\source'. The checkbox 'This share requires access credentials' is checked. The 'Credentials' field shows 'Administrator (Administrator, last edited: 9/15/17)'. The 'Gateway server' section has 'Automatic selection' selected. The 'Finish' button is highlighted.

Note: The Veeam Server is supported on the Windows platform only. To configure an NFS container from the DR Series system as a backup repository, you need to add the Linux server where the NFS container would be mounted.

6. (For NFS) Select **Linux Server (recommended)** as the type of Backup Repository, then click **Next**.

The screenshot shows the 'New Backup Repository' dialog box with the 'Type' tab selected. The 'Name' field is empty. The 'Type' field is set to 'Linux server (recommended)'. The 'Server' field is empty. The 'Repository' field is empty. The 'vPower NFS' field is empty. The 'Review' field is empty. The 'Apply' field is empty. The 'Finish' button is highlighted.

7. Add the New Repository server (Linux) or select the server from the list if added already.

The screenshot shows the 'New Backup Repository' wizard at the 'Server' step. The left sidebar has a tree view with 'Server' selected. The main area has a 'Repository server:' dropdown menu showing '10.250.213.24 (Created by DMA-SERVER1\Administrator at 3/22/2016 5:56 AM)' and an 'Add New...' button. Below this is a table with columns 'Path', 'Capacity', and 'Free', and a 'Populate' button. At the bottom are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

8. Mount the DR Series system NFS Container on this Linux server, and enter the container mount path as the Backup Repository.

The screenshot shows the 'New Backup Repository' wizard at the 'Repository' step. The left sidebar has a tree view with 'Repository' selected. The main area has a 'Location' section with a 'Path to folder:' text box containing '/mnt/nfs' and a 'Browse...' button. Below this are 'Capacity' and 'Free space' fields, both showing '...', and a 'Populate' button. The 'Load control' section has a note about performance and two checkboxes: 'Limit maximum concurrent tasks to:' (checked) with a value of '4', and 'Limit combined data rate to:' (unchecked) with a value of '...' and a unit of 'MB/s'. At the bottom right is an 'Advanced...' button. At the bottom are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

9. To customize the repository settings, click **Advanced**.

Edit Backup Repository

Repository
Type in path to the folder where backup files should be stored, and set repository load control options.

Name

Type

Share

Repository

vPower NFS

Review

Apply

Location

Path to folder:
\\10.250.241.229\source

Capacity: ...
Free space: ...

Populate

Load control

Running too many concurrent jobs against the same repository reduces overall performance, and may cause storage I/O operations to timeout. Control repository saturation with the following

☒ Limit maximum concurrent tasks to: 4

☐ Limit combined data rate to: MB/s

Click Advanced to customize repository settings

Advanced...

< Previous Next > Finish Cancel

Note: Refer to the *DR Series System Interoperability Guide* for information on **maximum concurrent jobs** supported for CIFS/NFS for your DR Series model. The maximum concurrent tasks also depend on the number of CPU cores of the Veeam Server. To run more tasks in parallel, you can add more **Backup proxy servers** to the Veeam server.

10. Select the option, **Decompress backup data blocks before storing**.

Note: Clearing the selection for the **Decompress backup data blocks before storing** option can increase your overall deduplication storage capacity usage. It is not recommended to switch these settings after the data has been written to the DR Series system.

Edit Backup Repository

Repository
Type in path

Name

Type

Share

Repository

Mount Server

Review

Apply

Storage Compatibility Settings

☐ **Align backup file data blocks**
Allows to achieve better deduplication ratio on deduplicating storage devices leveraging constant block size deduplication. Increases the backup size when backing up to raw disk storage.

☒ **Decompress backup data blocks before storing**
VM data is compressed by backup proxy according to the backup job compression settings to minimize LAN traffic. Uncompressing the data before storing allows for achieving better deduplication ratio on most deduplicating storage appliances at the cost of backup performance.

☐ **This repository is backed by rotated hard drives**
Backup jobs pointing to this repository will tolerate the disappearance of previous backup files by creating new full backup, clean up backup files no longer under retention on the newly inserted hard drives, and track backup repository location across unintended drive letter changes.

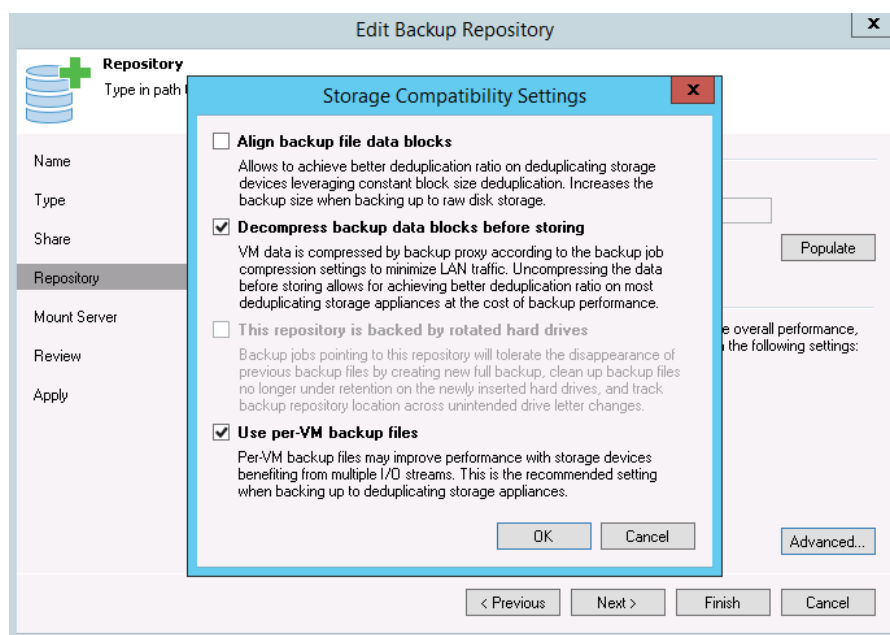
☒ **Use per-VM backup files**
Per-VM backup files may improve performance with storage devices benefiting from multiple I/O streams. This is the recommended setting when backing up to deduplicating storage appliances.

OK **Cancel** **Advanced...**

< Previous Next > Finish Cancel

Warning: Do not change the **Align backup file data blocks** setting after backups have been taken as this can impact deduplication savings for further backups.

11. Select the option, **Use Per-VM Backup Files Chains**, for a new backup repository option, which makes any backup job that is writing to a repository store each VM's restore point in a dedicated backup file. If you decide to create separate backup files for VMs in the job, make sure that you enable parallel data processing.



Note: This enables multiple write streams within a single job with parallel processing enabled. Enabling multiple streams dramatically improves overall job backup performance. So it is recommended to use per-VM backup files options for better backup throughput.

12. Click **OK** and **Next**.

13. Select the option, **Enable vPower NFS Server** for Instant Recovery to work.
For CIFS, enable the option as shown below:

The screenshot shows the 'Edit Backup Repository' dialog box with the 'vPower NFS' tab selected. The 'Name' field is 'vPower NFS'. The 'Type' is 'vPower NFS server (recommended)'. The 'Share' is 'This server'. The 'Repository' section shows the 'Folder' as 'F:\veeam2'. There are 'Manage...' and 'Ports...' buttons at the bottom right. Navigation buttons '< Previous', 'Next >', 'Finish', and 'Cancel' are at the bottom.

For NFS, enable the option as shown below:

The screenshot shows the 'New Backup Repository' dialog box with the 'Mount Server' tab selected. The 'Name' field is 'Mount server:'. The 'Type' is 'DMA-server1.testad.ocarina.local (Backup server)'. The 'Server' section shows the 'Folder' as 'C:\veeamdata'. There are 'Ports...' and 'Finish' buttons at the bottom right. Navigation buttons '< Previous', 'Next >', 'Finish', and 'Cancel' are at the bottom.

14. On the review page, verify the settings, and click **Next** to apply changes.

The screenshot shows the 'Edit Backup Repository' window with the 'Review' tab selected. The window title is 'Edit Backup Repository'. The 'Review' tab is highlighted in the left sidebar. The main content area displays the following information:

- Backup repository properties:**
 - Repository type: **CIFS**
 - Mount host: **This server**
 - Account: **Administrator**
 - Backup folder: **\\10.250.241.229\source**
 - Write throughput: **Not limited**
 - Max parallel tasks: **4**
- The following components will be processed on server This server:**
 - Installer: **already exists**
 - vPower NFS: **already exists**
- ☐ Import existing backups automatically
- ☐ Import guest file system index

At the bottom, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

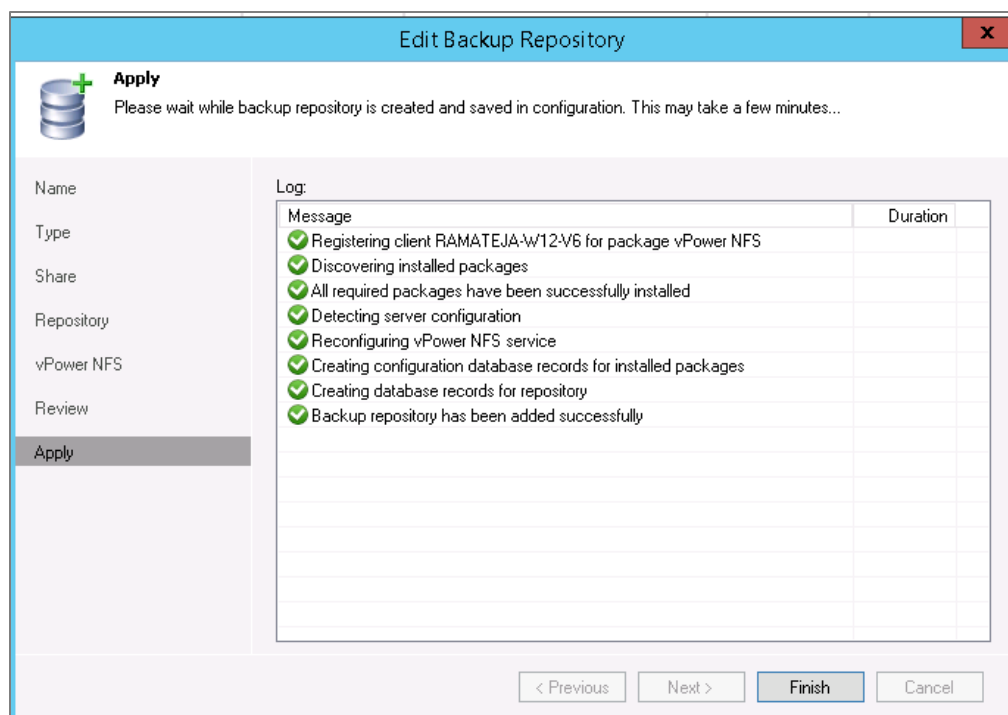
For an NFS Container Repository, the Review page will look like the following:

The screenshot shows the 'New Backup Repository' window with the 'Review' tab selected. The window title is 'New Backup Repository'. The 'Review' tab is highlighted in the left sidebar. The main content area displays the following information:

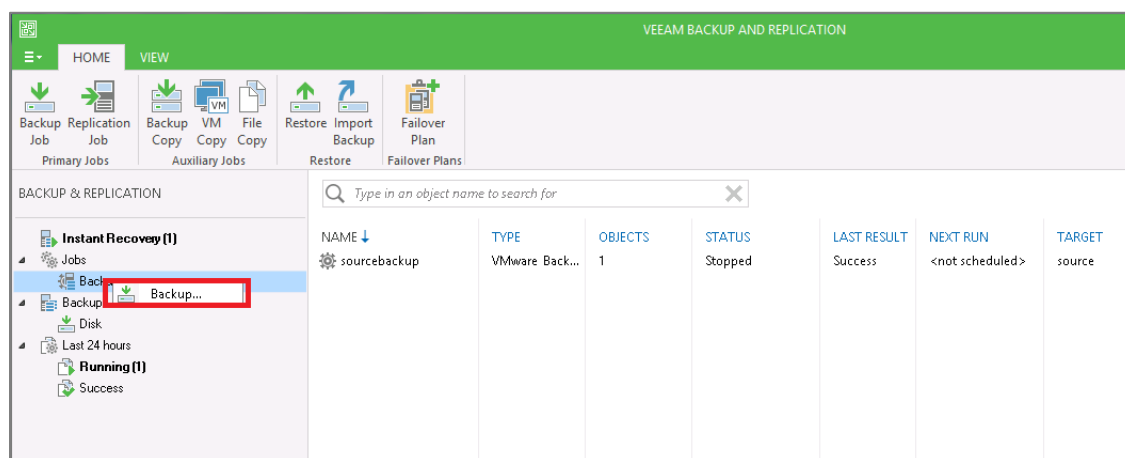
- Backup repository properties:**
 - Repository type: **Linux**
 - Mount host: **DMA-server1.testad.ocarina.local**
 - Account: **root**
 - Backup folder: **/mnt/veeam**
 - Write throughput: **Not limited**
 - Max parallel tasks: **4**
- The following components will be processed on server DMA-server1.testad.ocarina.local:**
 - Transport: **already exists**
 - vPower NFS: **already exists**
 - Mount Server: **already exists**
- ☐ Import existing backups automatically
- ☐ Import guest file system index

At the bottom, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

15. Click **Finish**.



16. On the Backup & Replication menu, go to **Jobs > Backup**, and right-click **Backup** to create a new backup job.



17. Enter the Name for the Backup job and click **Next**.

The screenshot shows the 'New Backup Job' dialog box. On the left is a sidebar with a 'vm' icon and a list of categories: Name, Virtual Machines, Storage, Guest Processing, Schedule, and Summary. The 'Name' category is selected. The main area has a 'Name' field with the text 'sourcebackup' and a 'Description' field with the text 'Created by DMA-SERVER1\Administrator at 3/2/2016 1:17 AM.' At the bottom right, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. The 'Next >' button is highlighted with a red box.

18. Select one or more virtual machines, data stores, resource pools, vApps, SCVMM clusters, etc. for backup and then click **Add**.

The screenshot shows the 'Add Objects' dialog box. On the left is a sidebar with a 'vm' icon and a list of categories: Name, Virtual Machines, Storage, Guest Processing, Schedule, and Summary. The 'Virtual Machines' category is selected. The main area shows a tree view of objects. The tree is expanded to show 'Hosts and Clusters' and 'veeam-test'. Under 'Hosts and Clusters', there is a folder '10.250.240.226' which contains 'Hyd-Datacenter', 'Production VMs', 'SAN', and 'Template'. Under 'veeam-test', there is a folder '10.250.213.25' which contains 'VirtualLab'. Under 'VirtualLab', there are three objects: 'dma-rhel7-v1', 'dma-server1', and 'ramtej-w12-v1'. The 'Add...' button is highlighted with a red box. At the bottom, there are two buttons: 'Add' and 'Cancel'. The 'Add' button is highlighted with a red box.

19. Select the DR Series system container share as the Backup Repository for this job, and click **Advanced**.

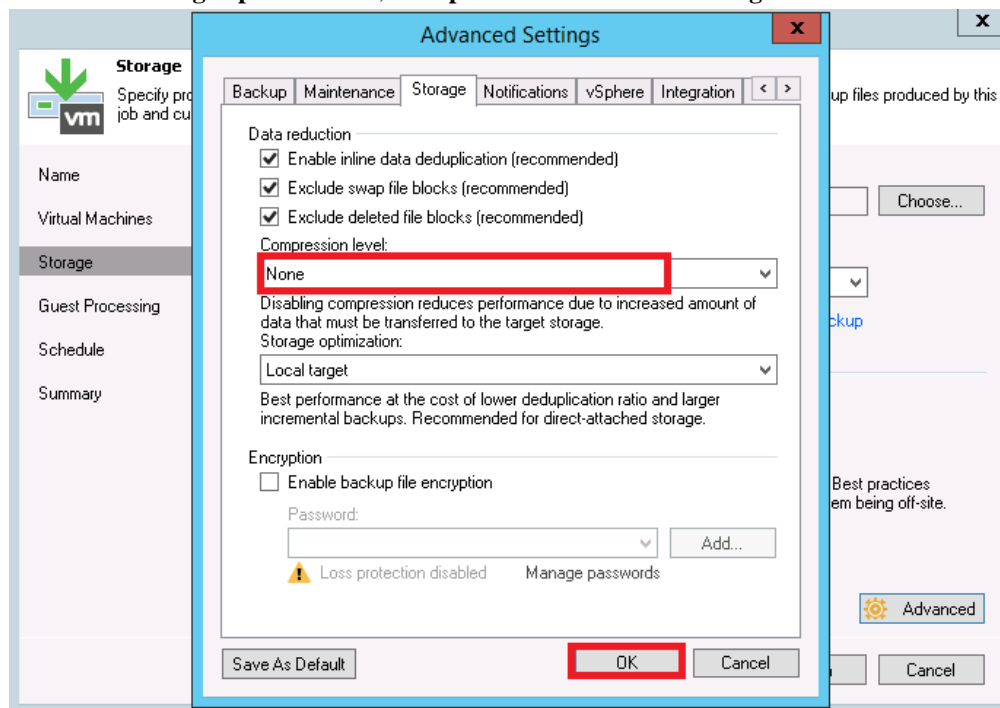
The screenshot shows the 'New Backup Job' dialog box with the 'Storage' tab selected. The 'Backup proxy' is set to 'Automatic selection'. The 'Backup repository' is set to 'source [Created by DMA-SERVER1\Administrator at 3/2/2016 1:09 AM.]', which is highlighted with a red box. Below the repository name, it shows '6.0 TB free of 7.8 TB' and a 'Map backup' link. The 'Retention policy' is set to '14' days. There is a checkbox for 'Configure secondary destinations for this job'. At the bottom, there is a red box around the 'Advanced' button. Navigation buttons at the bottom include '< Previous', 'Next >', 'Finish', and 'Cancel'.

20. On the **Backup** tab, make sure **Incremental** is selected.

Note: It is recommended to enable **Active Full backups** once a week. The active full backup produces a full backup of a VM just as if you ran the backup job for the first time. The active full backup resets the chain of incremental backups. All subsequent incremental backups use the last active full backup as a new starting point. A previously used full backup file remains on disk until it is automatically deleted according to the backup retention policy.

The screenshot shows the 'Advanced Settings' dialog box with the 'Storage' tab selected. The 'Backup mode' section has two options: 'Reverse incremental (slower)' and 'Incremental (recommended)'. The 'Incremental (recommended)' option is selected and highlighted with a red box. Below this, there are checkboxes for 'Create synthetic full backups periodically' and 'Transform previous backup chains into rollbacks'. The 'Active full backup' section has a checkbox for 'Create active full backups periodically' which is also highlighted with a red box. Below this, there are options for 'Monthly on: First' and 'Monday' with a 'Months...' button, and 'Weekly on selected days: Saturday' with a 'Days...' button. At the bottom, there are buttons for 'Save As Default', 'OK', and 'Cancel'.

21. On the **Storage** tab, do the following:
- Under **Deduplication**, select **Enable inline data deduplication**.
 - Under **Compression**, set the **Level** to **None**.
 - Under **Storage optimizations**, set **Optimization** to **LOCAL target**.



Note: For Advanced Settings, between backup performance and deduplication savings, if overall space/storage savings is the focus, it is recommended to choose the options for all of the backup jobs.

For Veeam deduplication: Normally, turning off encryption, compression, and deduplication is recommended for most backup software. However, with Veeam, you should enable deduplication as Veeam runs deduplication for data block sizes 512 KB or above; and, deduplication of these large block sizes does not heavily impact DR Series duplication results. In addition, this reduces network bandwidth utilization when Veeam sends data to the DR Series system, so it benefits the backup practice overall.

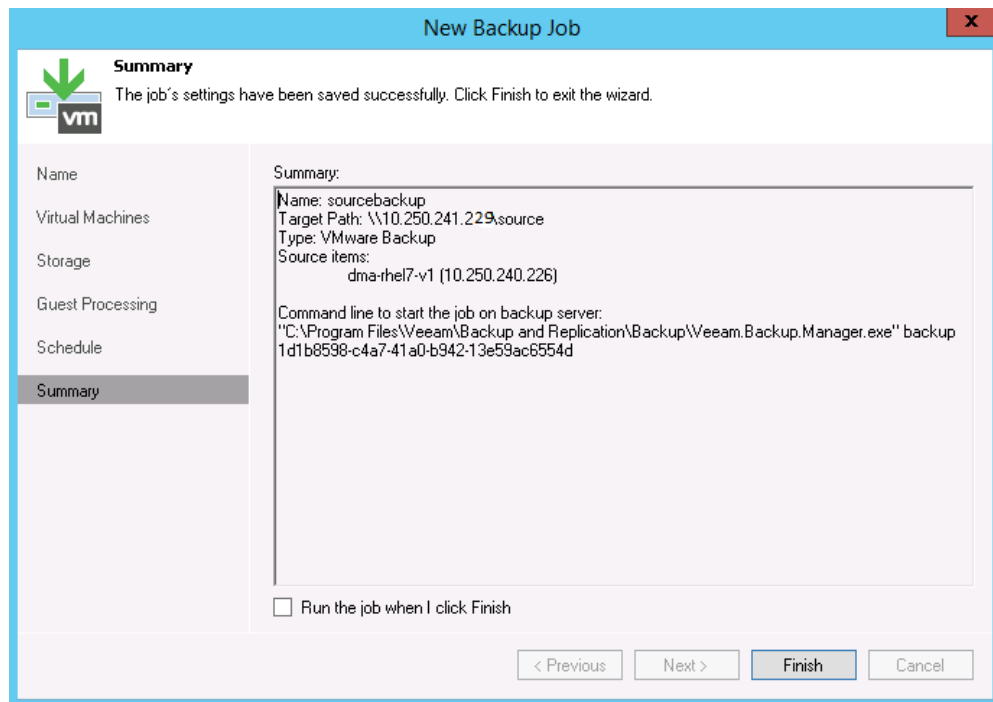
22. Click **Next**.

The screenshot shows the 'New Backup Job' wizard with the 'Guest Processing' step selected. The left sidebar contains links for Name, Virtual Machines, Storage, Guest Processing (highlighted), Schedule, and Summary. The main area is titled 'Guest Processing' and includes a green arrow icon with a 'vm' label. Below the title is the instruction: 'Choose guest OS processing options available for running VMs.' The main content area contains several options: 'Enable application-aware processing' (unchecked), 'Enable guest file system indexing' (unchecked), 'Guest OS credentials' (a dropdown menu with an 'Add...' button), 'Guest interaction proxy' (set to 'Automatic selection' with a 'Choose...' button), and 'Test Now' button. At the bottom are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

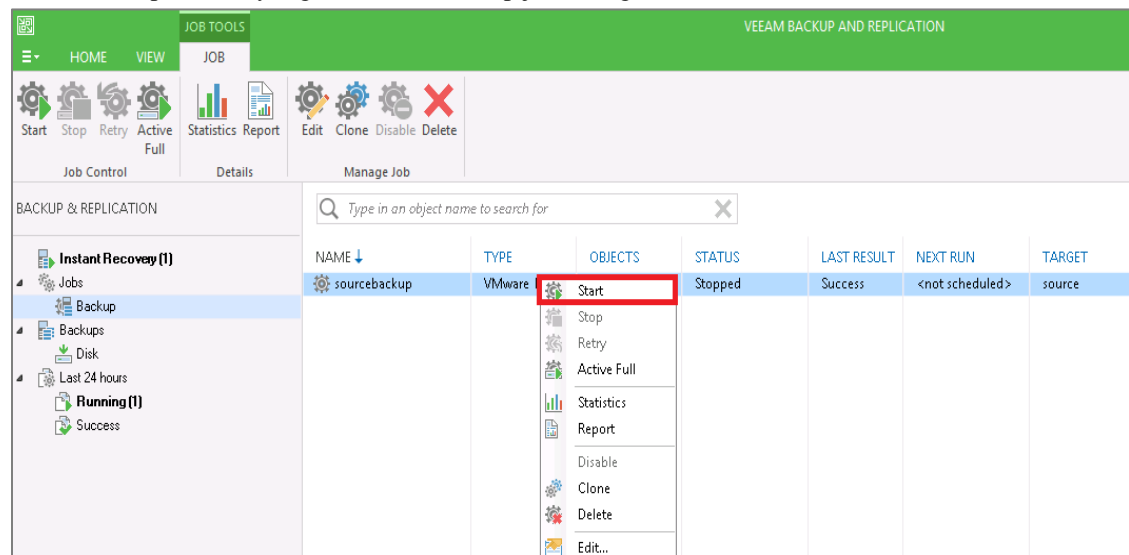
23. Schedule the backup and click **Create**.

The screenshot shows the 'New Backup Job' wizard with the 'Schedule' step selected. The left sidebar contains links for Name, Virtual Machines, Storage, Guest Processing, Schedule (highlighted), and Summary. The main area is titled 'Schedule' and includes a green arrow icon with a 'vm' label. Below the title is the instruction: 'Specify the job scheduling options. If you do not set the schedule, the job will need to be controlled manually.' The main content area contains several options: 'Run the job automatically' (checked), 'Daily at this time' (selected) with a time of 10:00 PM and frequency of Everyday, 'Monthly at this time' (unselected) with a time of 10:00 PM and frequency of Fourth, 'Periodically every' (unselected) with a frequency of 1 Hours, and 'After this job' (unselected) with a selection of 'vcifs1 (Created by TESTAD\administrator at 2/17/2015 4:05 AM.)'. Below these are 'Automatic retry' options: 'Retry failed VMs processing' (checked) with a count of 3 times and 'Wait before each retry attempt for' (10 minutes). At the bottom are navigation buttons: '< Previous', 'Create', 'Finish', and 'Cancel'.

24. Click **Finish**.




25. To run backup manually, right-click the backup job configured, and select **Start**.



4 Setting up DR Series native replication and restore from a replication target container

4.1 Building a replication relationship between DR Series systems

1. Create a target container on the target DR Series system.




2. On the source DR Series system, go to the **Replication** menu, and, in the Action menu in the upper right corner, click **Add Replication**.



3. Choose **Replication** type and click **Next**.



4. Select the Source Container for replication and click **Next**.



5. Select the Encryption Type for the source container and click **Next**.

The screenshot shows the 'Add Replication' wizard in the Dell DR4300 interface. The left sidebar contains navigation links: GlobalView, Dashboard, Containers, Replications, System Configuration, and Support. The main area is titled 'All Replications' and shows the 'Add Replication' step. A red box highlights the 'Source Container => Replica Container' section, with a '1' annotation. Below it, the 'Encryption' section has three radio buttons: 'Not Enabled' (selected), 'AES 128-bit', and 'AES 256-bit'. At the bottom, the 'Next' button is highlighted with a red box and a '2' annotation. Other buttons include 'Previous', 'Finish', and 'Cancel'.

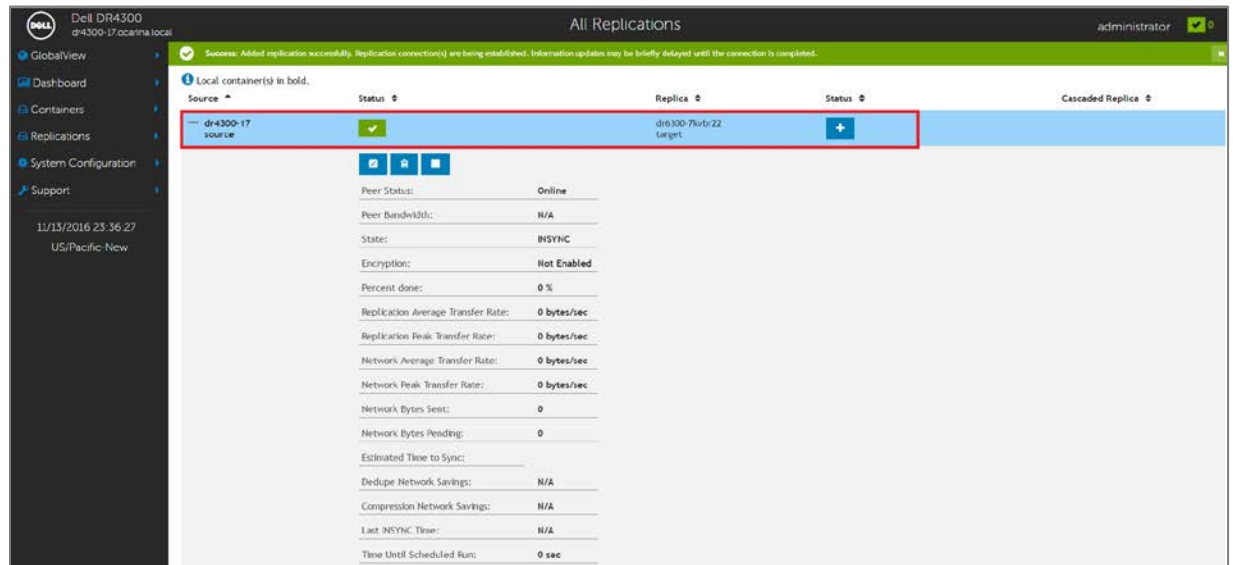
6. Select **Container from remote system**, enter the target DR Series system related information, click **Retrieve Remote Containers**, select a populated target container from the list, and click **Next**.

The screenshot shows the 'Add Replication' wizard in the Dell DR4300 interface, Step 2: Replica Container. The left sidebar is the same as in the previous screenshot. The main area is titled 'All Replications' and shows the 'Add Replication' step. A red box highlights the 'Replica Container' section, with a '1' annotation. Inside this section, 'Select container location' has 'Local' and 'Remote' radio buttons, with 'Remote' selected. Below this, the 'Username' field is 'administrator', the 'Password' field is masked with '*****', and the 'Remote system' field is '10.250.212.190'. A red box highlights the 'Retrieve Remote Container(s)' button, with a '2' annotation. Below this, the 'Select remote container' dropdown shows 'target' selected, with a '3' annotation. At the bottom, the 'Next' button is highlighted with a red box and a '4' annotation. Other buttons include 'Previous', 'Finish', and 'Cancel'.

7. Verify the **Summary** and click **Finish**.

The screenshot shows the 'Add Replication' wizard in the Dell DR4300 interface, Step 3: Summary. The left sidebar is the same as in the previous screenshots. The main area is titled 'All Replications' and shows the 'Add Replication' step. The 'Summary' section displays the following information: 'Source Container' (Location: local, Name: source), 'Source Container -> Replica Container' (Encryption: Not Enabled), 'Replica Container' (Location: remote, Remote System: 10.250.212.190, name: target). At the bottom, the 'Finish' button is highlighted with a red box. Other buttons include 'Previous', 'Next', and 'Cancel'.

8. Verify that the replication relationship was created successfully.



Notes:

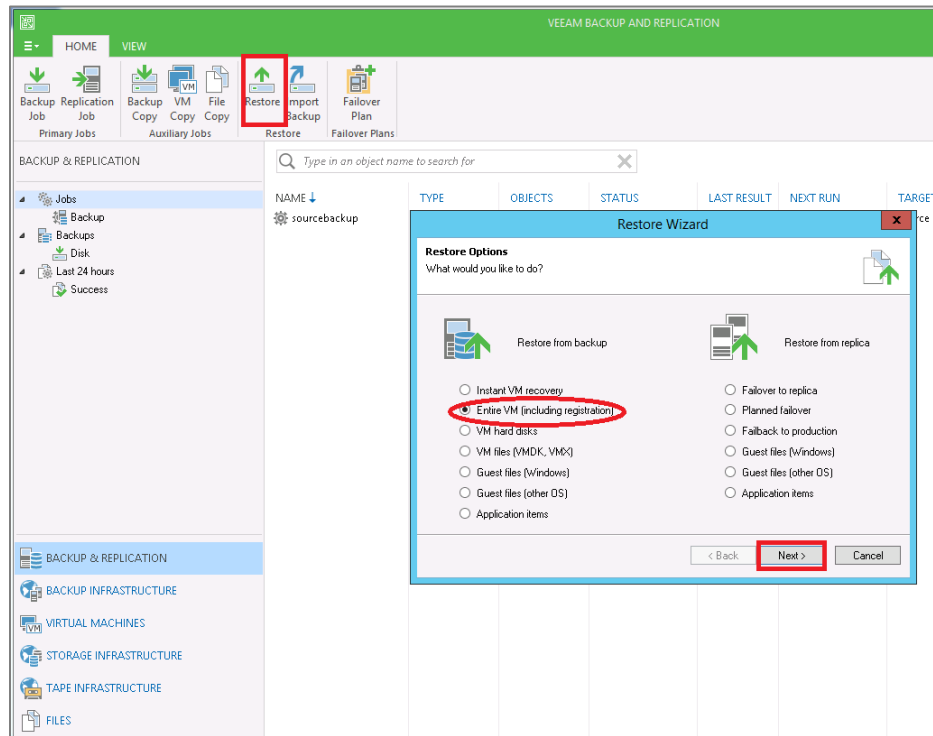
- Make sure the replication session has **Peer Status** as **Online**. If restore from replication target is needed,
- Make sure the replication is in **INSYNC** state from Replication Statistics menu, and Stop or Delete the replication.
- Make sure the replication target has **CIFS/NFS** connection(s) enabled when restoring from it.

4.2 Restoring data from the target DR Series system

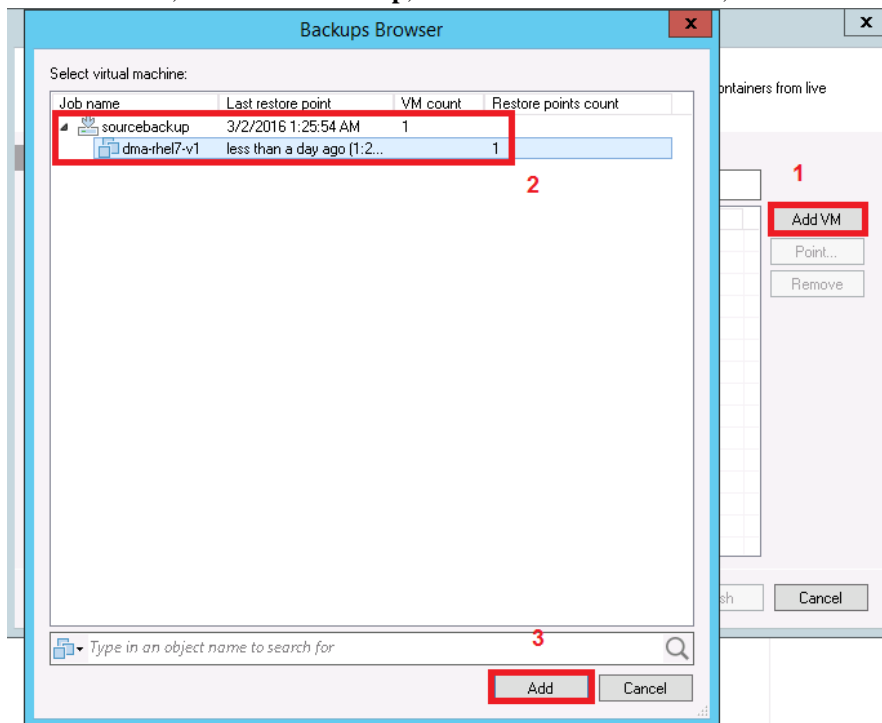
Note: Before restoring from the target DR Series system, make sure that the replication session state is **INSYNC** on the DR Series system GUI **Replication Statistics** menu. **Stop** or **Delete** the replication session, and make sure that the target DR Series system container has the CIFS/NFS connection(s) enabled.

1. Add the target DR Series system container to the Veeam repository. For instructions, see the section, “Setting up Veeam.”
2. Update all backup jobs that use the source DR Series system container as a repository and change them to use the target DR Series system container as the backup repository.


3. For Backup & Replication, click **Restore** to create a restore job. Select the appropriate option under Restore from backup.



4. Click **Add VM**, select **From backup**, select the VM to be restored, and then click **Add**.



5. Select the Restore Mode and click **Next**.



Full VM Restore Wizard

Virtual Machines

Restore Mode

Host

Resource Pool

Datastore

Folder

Network

Reason

Summary

Restore Mode

Specify whether selected VMs should be restored back to the original location, or to a new location or with different settings.

☐

Restore to the original location

Quickly initiate restore of selected VMs to the original location, and with the original name and settings. This option minimizes the chance of user input error.

☒

Restore to a new location, or with different settings

Customize restored VM location, and change its settings. The wizard will automatically populate all controls with the original VM settings as the default settings.

[Pick proxy to use](#)

☐

Restore VM tags

Select this option to restore VM tags that were assigned to the VM when backup was taken.

☐

Quick rollback (restore changed blocks only)

Allows for quick VM recovery in case of guest OS software problem, or user error. Do not use this option when recovering from disaster caused by hardware or storage issue, or power loss.

< Previous

Next >

Finish

Cancel

6. Provide the required host details and click **Next**.

7. Select the resource pool, and click **Next**.

Full VM Restore Wizard

Resource Pool

By default, original resource pool is selected as restore destination for each VM. You can change resource pool by selecting desired VM and clicking Pool. Use multi-select (Ctrl-click and Shift-click) to select multiple VMs at once.

Virtual Machines

Restore Mode

Host

Resource Pool

Datastore

Folder

Network

Reason

Summary

VM resource pool:

Name	Resource Pool
dma-rhel7-v1	Resources

Select multiple VMs and click Pool to apply changes in bulk.

Pool...

< Previous


Next >

Finish

Cancel

8. Select the datastore disk type, and click **Next**.

Full VM Restore Wizard



Datastore

By default, original datastore and disk type are selected for each VM file. You can change them by selecting desired VM file, and clicking Datastore or Disk Type. Use multi-select (Ctrl-click and Shift-click) to select multiple VMs at once.

X

Virtual Machines

Restore Mode

Host

Resource Pool

Datastore

Folder

Network

Reason

Summary

Files location:

File	Size	Datastore	Disk type
<div style="display: flex; align-items: center;"> <div style="margin-right: 5px;">▲</div> <div style="display: flex; align-items: center;"> <div style="margin-right: 5px;">■</div> dma-rhel7-v1 </div> </div>			
<div style="display: flex; align-items: center;"> <div style="margin-right: 5px;">■</div> <div style="display: flex; align-items: center;"> <div style="margin-right: 5px;">■</div> Configuration files </div> </div>		datastore1 [616.9 GB free]	
<div style="display: flex; align-items: center;"> <div style="margin-right: 5px;">■</div> <div style="display: flex; align-items: center;"> <div style="margin-right: 5px;">■</div> Hard disk 1 (dma... </div> </div>	160.0 ...	datastore1 [616.9 GB free]	Same as source

Select multiple VMs to apply settings in bulk.

Datastore...
Disk Type...

< Previous
Next >
Finish
Cancel

9. Provide the new name for the restored VM, click **OK**, and then click **Next**.

Full VM Restore Wizard

Folder

By default, original VM folder is selected as restore destination for each VM. You can change folder by selecting desired VM and clicking Folder. Use multi-select (Ctrl-click and Shift-click) to select multiple VMs at once.

Virtual Machines

Restore Mode

Host

Resource Pool

Datastore

Folder

Network

Reason

Summary

Change Name

Specify how selected VM name should be changed:

Set name to:
dma-rhel7-v1

☒ Add prefix:
new_

☒ Add suffix:
_restored

OK Cancel

Select multiple VMs to apply settings change in bulk.

Name... Folder...

< Previous **Next >** Finish Cancel

10. Select the network location and click **Next**.

Full VM Restore Wizard

Network

By default, restored VM is connected to the same virtual networks as the original VM. If you are restoring to a different location, specify how original location's networks map to new location's networks.

Virtual Machines

Restore Mode

Host

Resource Pool

Datastore

Folder

Network

Reason

Summary

Network connections:

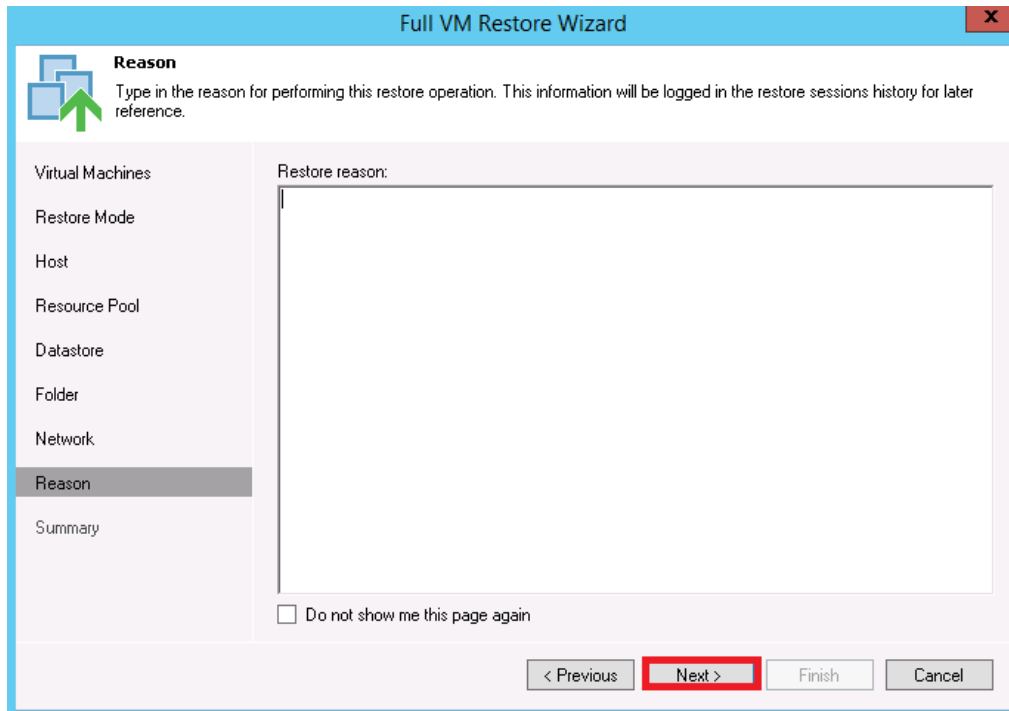
Source	Target
dma-rhel7-v1	VM Network

Select multiple VMs to apply settings change in bulk.

Network... Disconnected

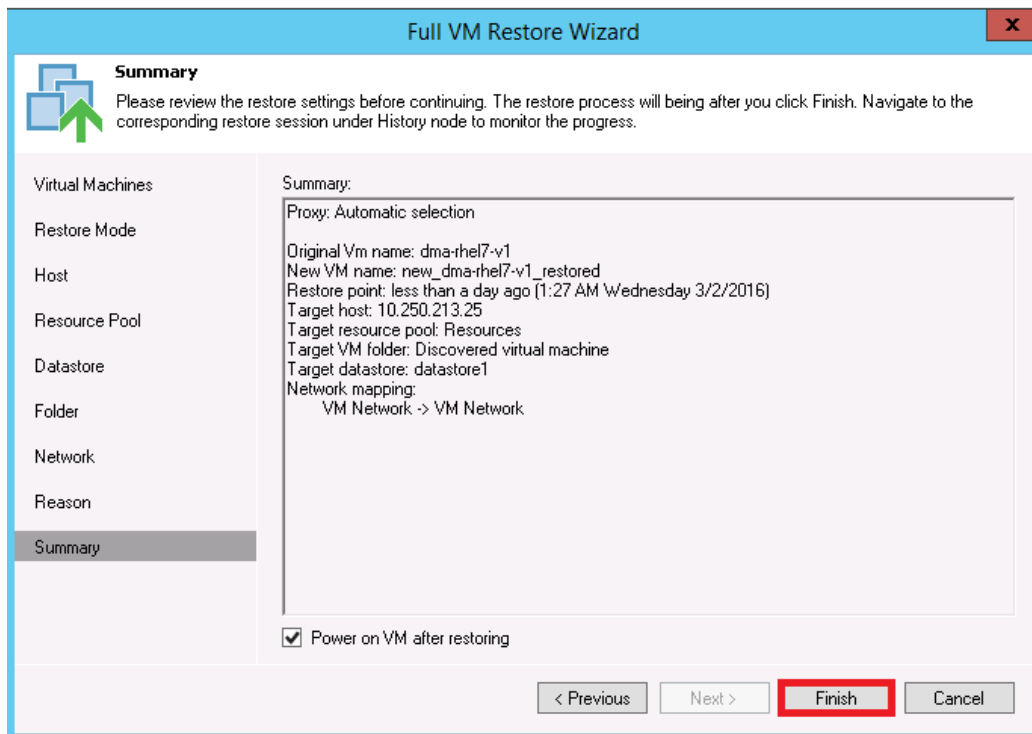
< Previous **Next >** Finish Cancel

11. Provide the reason for the restore and click **Next**.



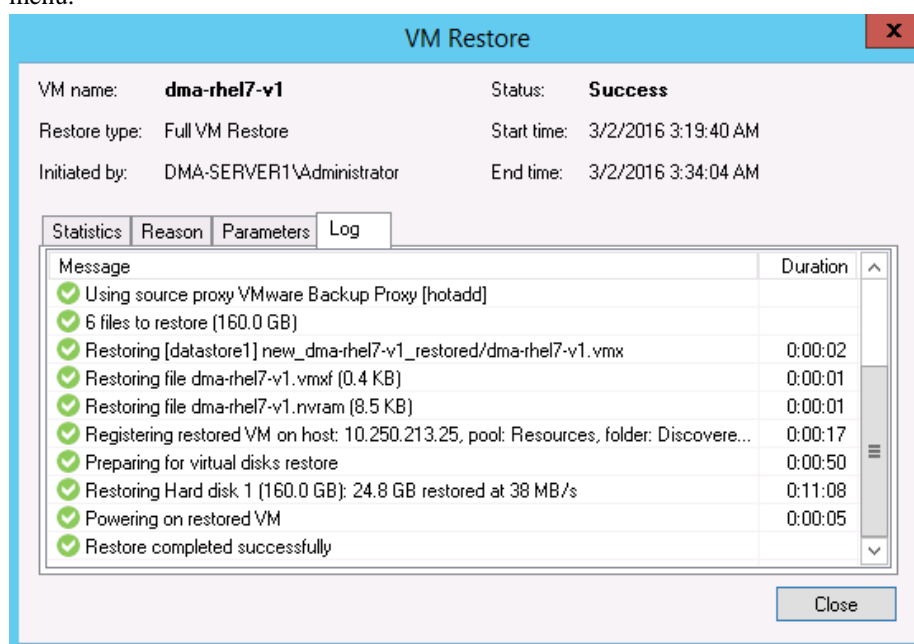
The screenshot shows the 'Full VM Restore Wizard' window, specifically the 'Reason' step. The window has a blue title bar with the text 'Full VM Restore Wizard' and a close button. On the left, there is a sidebar with a list of steps: Virtual Machines, Restore Mode, Host, Resource Pool, Datastore, Folder, Network, Reason (highlighted), and Summary. The main area is titled 'Reason' and contains a text box for 'Restore reason:'. Below the text box is a checkbox labeled 'Do not show me this page again'. At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted with a red box), 'Finish', and 'Cancel'.

12. Review the summary and click **Finish**.



The screenshot shows the 'Full VM Restore Wizard' window, specifically the 'Summary' step. The window has a blue title bar with the text 'Full VM Restore Wizard' and a close button. On the left, there is a sidebar with a list of steps: Virtual Machines, Restore Mode, Host, Resource Pool, Datastore, Folder, Network, Reason, and Summary (highlighted). The main area is titled 'Summary' and contains a text box with the following text: 'Proxy: Automatic selection', 'Original Vm name: dma-rhel7-v1', 'New VM name: new_dma-rhel7-v1_restored', 'Restore point: less than a day ago (1:27 AM Wednesday 3/2/2016)', 'Target host: 10.250.213.25', 'Target resource pool: Resources', 'Target VM folder: Discovered virtual machine', 'Target datastore: datastore1', and 'Network mapping: VM Network -> VM Network'. Below the text box is a checkbox labeled 'Power on VM after restoring' which is checked. At the bottom, there are four buttons: '< Previous', 'Next >', 'Finish' (highlighted with a red box), and 'Cancel'.

13. After the restore job has been created, you can run the job and monitor it from the **Backup & Replication** menu.



5 Configuring Rapid CIFS for Veeam

Rapid CIFS enable write operation acceleration on clients that use CIFS file system protocols. These accelerators allow for better coordination and integration between DR Series systems backup, restore, and optimized duplication operations with Data Management Applications (DMAs).

Rapid CIFS is a Windows-certified filter driver that also ensures that only unique data is written to the DR Series system. All chunking and hash computations are done at the client level where this filter driver is installed.

As the Veeam Backup **Proxy server** does all the operations related to backup and Restore, it is recommended to install Quest Rapid CIFS filter driver software on Backup **Proxy server**.

5.1.1 Windows prerequisites


- The Veeam Backup Proxy Server OS must be the 64-bit version of Windows 2008 R2 or Windows 2012.
- It should be with a minimum of 4 CPU cores running at a minimum of 4 GHz cumulative processing power and 2 GB memory
- The DR container share must be mapped on the Veeam Backup Proxy Server.

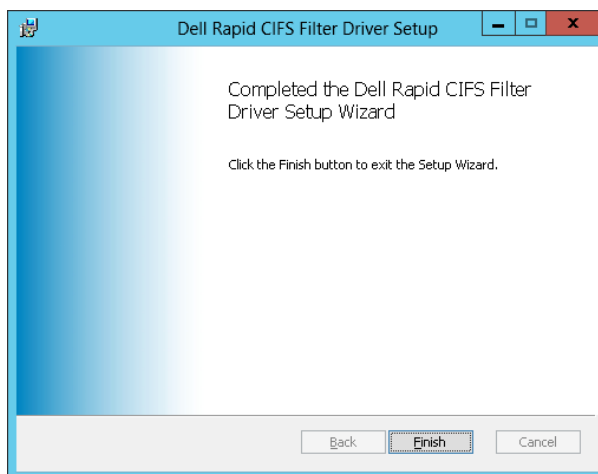

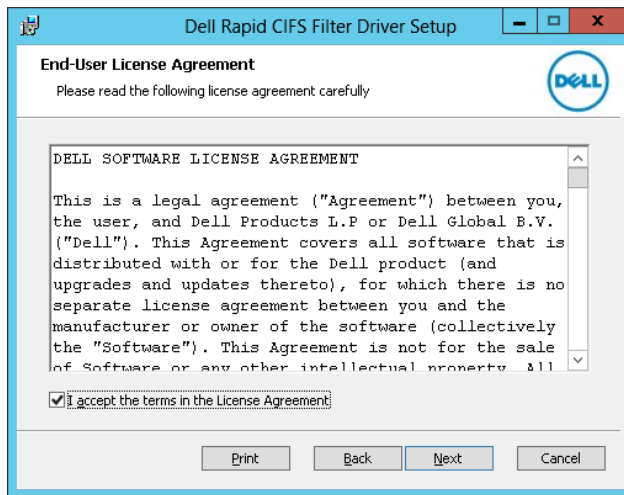
5.1.2 Installing Rapid CIFS on a Backup Proxy Server

Follow these steps to install Rapid CIFS -

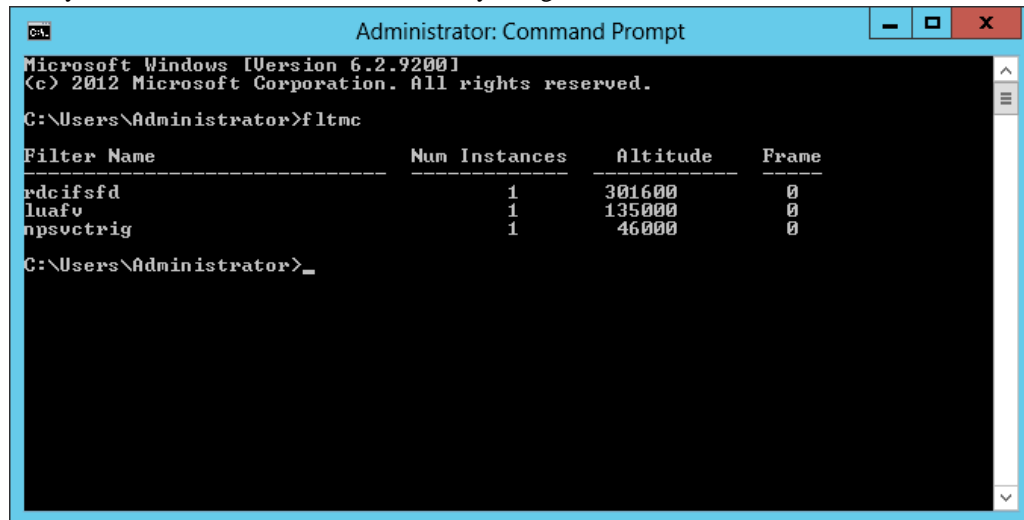
Note: Rapid CIFS should only be installed on a **Backup Proxy Server**.

1. Download the MSI to the Backup Proxy server by doing the following:
 - a. Go to support.quest.com/DR-Series and select your specific product (such as DR4100, DR6000, etc).
 - b. On the support page for your product, click **Software Downloads**.
 - c. For the RDCIFS plugin for your DR Series system OS version, click the Download icon to download the installer package (.msi file).
2. Run the MSI and follow the instructions in the installation wizard as shown in the screenshots below. Click **Finish** when installation is complete.





3. Verify that the “**rdcifsfd**” driver is loaded by using the command **fltmc**.



```
Administrator: Command Prompt
Microsoft Windows [Version 6.2.9200]
(c) 2012 Microsoft Corporation. All rights reserved.

C:\Users\Administrator>fltmc

Filter Name                Num Instances  Address      Frame
-----
rdcifsfd                   1              301600      0
luafv                     1              135000      0
npsvcctrig                 1              46000       0

C:\Users\Administrator>
```

To check for Client Side optimization, enter the following command:

```
C:\Program Files\Dell\Rapid CIFS>rdcifsctl.exe stats -s
```

Aggregate Statistics:

Total Bytes Written: 2,411,298,816

Total Bytes Sent: 1,378,067,343

Total Network Savings: 42.8496

Note: For more information, such as about troubleshooting and logging, refer to the *DR Series System Administrator Guide*.



6 Using Instant VM Recovery with the DR Series system

Veeam's Instant VM Recovery immediately restores a virtual machine (VM) back into your production environment by running it directly from the backup file. Instant VM Recovery uses patented vPower® technology to mount a VM image to a production VMware vSphere or Microsoft Hyper-V host directly from a compressed and deduplicated backup file.

By default, all changes to virtual disks that take place while the VM is running are logged to auxiliary redo logs residing on the NFS server (Veeam backup server or backup repository). These changes are discarded as soon as a restored VM is removed, or merged with the original VM data when VM recovery is finalized, that is, when VM is migrated back to production storage.

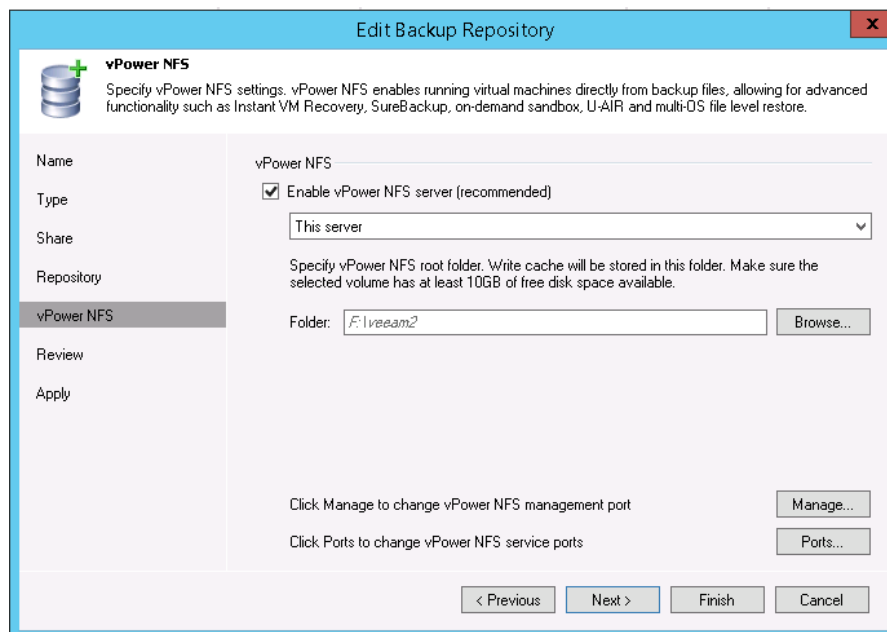
Veeam vPower NFS service is a Windows service that runs on a windows backup repository server and enables it to act as NFS server

6.1 Instant Recovery with ESX

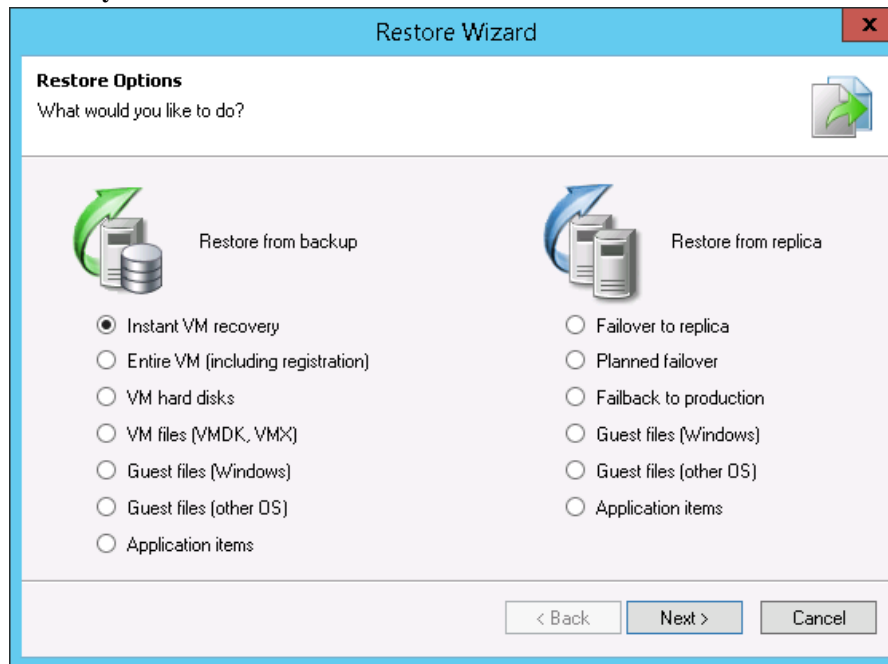
Follow these steps to enable and perform Instant Recovery for ESX VM backups:

1. Create a backup job for the required VM as described in Section 3 of this document with the only difference as follows:
 - a. On the vPower NFS tab, select the checkbox **Enable vPower NFS Server**, and select the appropriate folder as the NFS Datastore.

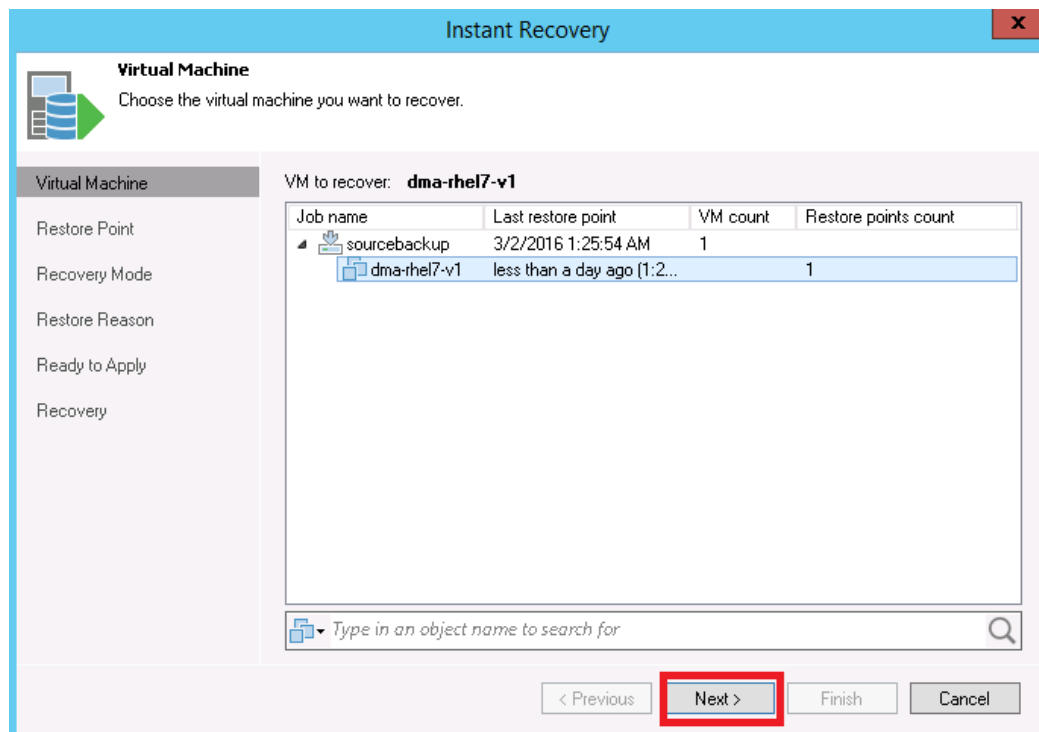
Note: You can also configure the NFS Datastore on a different Windows server if required by selecting the drop-down and adding the host information and credentials.



2. On the Veeam Server console, click the **Restore Wizard** option, select **VMware**, and then select **Instant VM Recovery**.



3. Select the Virtual Machine to be recovered and click **Next**.



4. At the Restore point step, select the point to which you want to restore the VM and click **Next**.

The screenshot shows the 'Instant Recovery' wizard at the 'Restore Point' step. The left sidebar contains a list of steps: Virtual Machine, Restore Point (selected), Recovery Mode, Restore Reason, Ready to Apply, and Recovery. The main area displays VM details: VM name: dma-rhel7-v1, Original host: 10.250.240.226, and VM size: 76.6 GB. Below this, a table titled 'Available restore points:' shows one entry: 'less than a day ago (1:27 AM Wednesday 3/2/2016)' with a 'Full' type. At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted with a red box), 'Finish', and 'Cancel'.

Created	Type
less than a day ago (1:27 AM Wednesday 3/2/2016)	Full

5. At the Restore Mode step, select the option, **Restore to a new location, or with different settings**.

The screenshot shows the 'Instant Recovery' wizard at the 'Restore Mode' step. The left sidebar contains a list of steps: Virtual Machine, Restore Point, Restore Mode (selected), Destination, Datastore, Restore Reason, Ready to Apply, and Recovery. The main area displays two radio button options. The first option is 'Restore to the original location' with a description: 'Quickly initiate restore of selected VMs to the original location, and with the original name and settings. This option minimizes the chance of user input error.' The second option is 'Restore to a new location, or with different settings' (selected) with a description: 'Customize restored VM location, and change its settings. The wizard will automatically populate all controls with the original VM settings as the default settings.' At the bottom, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

6. At the Destination step, do the following:
 - a. Select the ESX host on which the VM should be restored instantly.
 - b. In the Resource pool box, select the resource pool to which the restored VM should belong.
 - c. In the Restored VM name field, add the `_restored` suffix to the VM name.

Instant Recovery

Destination
Choose ESX server to run the recovered virtual machine on. You can choose to power on VM automatically, unless you need to adjust VM settings first (such as change VM network).

Virtual Machine: Host: 10.250.213.25 [Choose...]

Restore Point: VM folder: Discovered virtual machine [Choose...]

Recovery Mode: [Choose...]

Destination

Restored VM name: dmarhel7-v1_IP

Resource pool: Resources

Resources

VirtualLab

< Previous **Next >** Finish Cancel

7. On the Datastore tab, ensure that checkbox, **Redirect virtual disk updates**, is not selected. This ensures that you use Storage vMotion to migrate the VM to production after the VM recovers from the backup. Click **Next**.

Instant Recovery

Datastore
By default, virtual disk changes of recovered VM are stored on vPower NFS server. You can redirect these changes to a different datastore. This improves I/O performance, but prevents Storage vMotion on vSphere versions prior to vSphere 5.0

Virtual Machine: ☐ Redirect virtual disk updates

Restore Point: Datastore: Click Choose to pick the datastore [Choose...]

Restore Mode: Datastore info

Destination: Capacity: <Datastore not set>

Datastore: Free space: <Datastore not set>

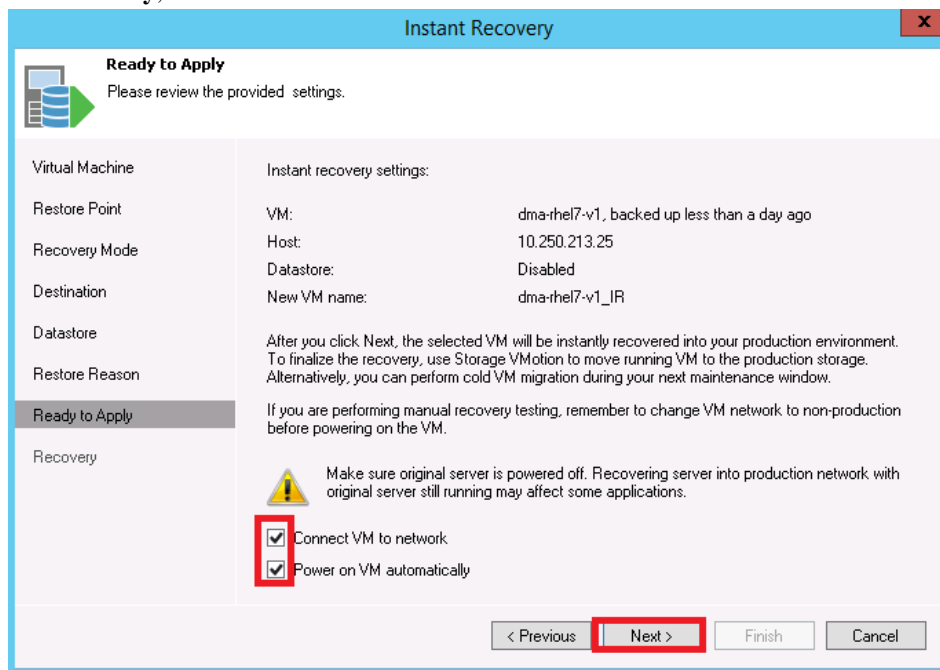
Restore Reason: [Choose...]

Ready to Apply: [Choose...]

Recovery: [Choose...]

< Previous **Next >** Finish Cancel

8. In the Ready to Apply dialog box, select the options, **Connect VM to network** and **Power on VM automatically**, and click **Next**.



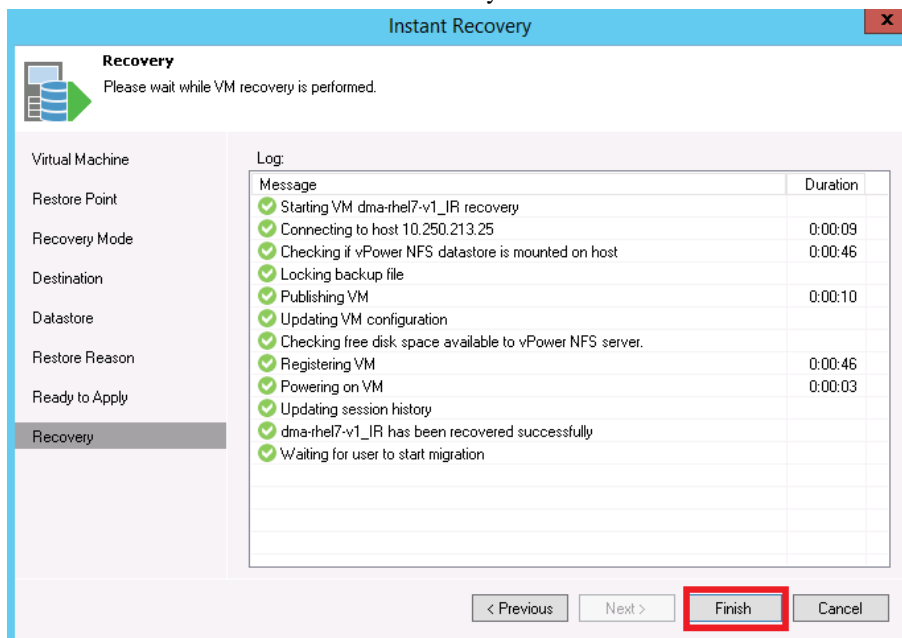
The 'Instant Recovery' dialog box is titled 'Ready to Apply'. It contains a sidebar on the left with options: Virtual Machine, Restore Point, Recovery Mode, Destination, Datastore, Restore Reason, Ready to Apply (selected), and Recovery. The main area displays 'Instant recovery settings:' with the following details: VM: dma-rhel7-v1, backed up less than a day ago; Host: 10.250.213.25; Datastore: Disabled; New VM name: dma-rhel7-v1_IR. Below this, a warning message states: 'After you click Next, the selected VM will be instantly recovered into your production environment. To finalize the recovery, use Storage VMotion to move running VM to the production storage. Alternatively, you can perform cold VM migration during your next maintenance window.' A second warning message says: 'If you are performing manual recovery testing, remember to change VM network to non-production before powering on the VM.' At the bottom, there are two checked checkboxes: 'Connect VM to network' and 'Power on VM automatically'. The 'Next >' button is highlighted with a red box.

Virtual Machine	Instant recovery settings:
Restore Point	VM: dma-rhel7-v1, backed up less than a day ago
Recovery Mode	Host: 10.250.213.25
Destination	Datastore: Disabled
Datastore	New VM name: dma-rhel7-v1_IR
Restore Reason	After you click Next, the selected VM will be instantly recovered into your production environment. To finalize the recovery, use Storage VMotion to move running VM to the production storage. Alternatively, you can perform cold VM migration during your next maintenance window.
Ready to Apply	If you are performing manual recovery testing, remember to change VM network to non-production before powering on the VM.
Recovery	Make sure original server is powered off. Recovering server into production network with original server still running may affect some applications.

☒ Connect VM to network
☒ Power on VM automatically

< Previous **Next >** Finish Cancel

9. Click **Finish** to start the Instant VM Recovery.



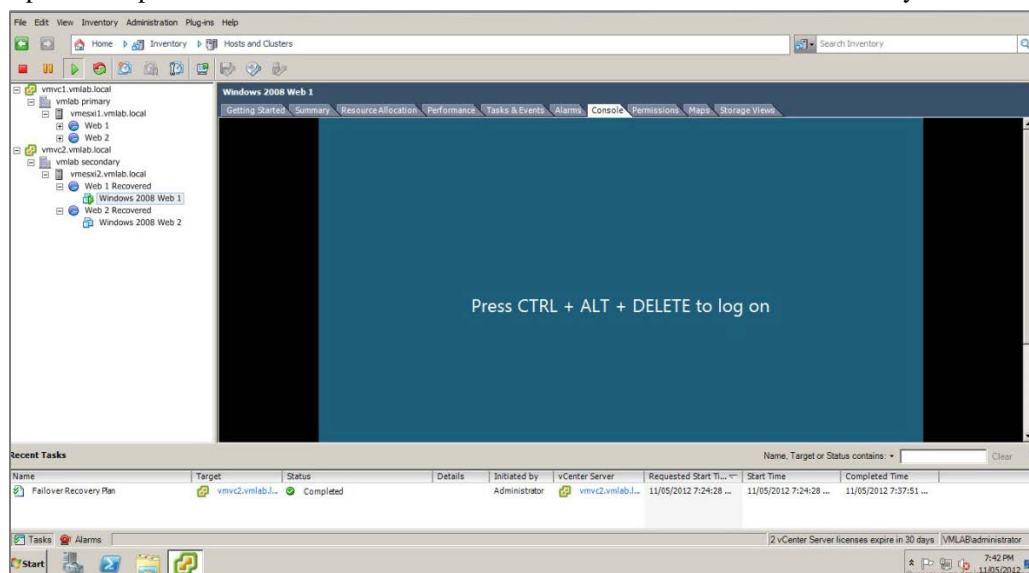
The 'Instant Recovery' dialog box is titled 'Recovery'. It contains a sidebar on the left with options: Virtual Machine, Restore Point, Recovery Mode, Destination, Datastore, Restore Reason, Ready to Apply, and Recovery (selected). The main area displays a 'Log:' table with the following entries:

Message	Duration
Starting VM dma-rhel7-v1_IR recovery	
Connecting to host 10.250.213.25	0:00:09
Checking if vPower NFS datastore is mounted on host	0:00:46
Locking backup file	
Publishing VM	0:00:10
Updating VM configuration	
Checking free disk space available to vPower NFS server.	
Registering VM	0:00:46
Powering on VM	0:00:03
Updating session history	
dma-rhel7-v1_IR has been recovered successfully	
Waiting for user to start migration	

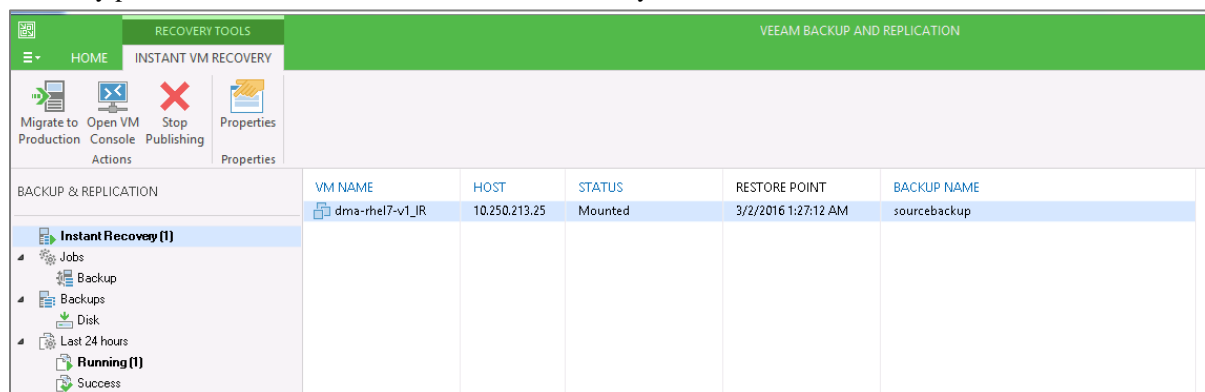
The 'Finish' button is highlighted with a red box.

< Previous Next > **Finish** Cancel

10. Open the vSphere client and ensure that the restored VM is started on the ESX host you selected.



11. In Veeam Backup & Replication, open the Backup & Replication view, select the Instant Recovery node in the inventory pane and make sure that the Instant VM Recovery session is available and mounted.



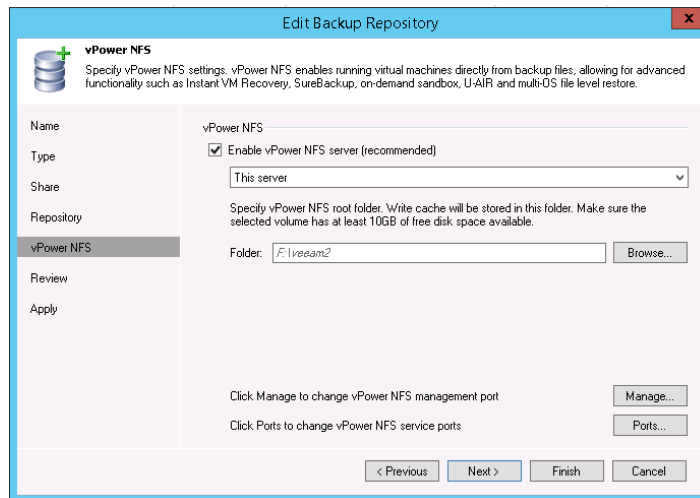
6.2 Instant Recovery with Hyper-V Server

6.2.1 Enabling Instant Recovery for Hyper-V VM backups

Do the following to enable Instant Recovery for HyperV VM backups:

1. Create a backup job for the required VM as described previously in this document (in Section 3) with the only difference as specified in the following step.
2. On the vPower NFS tab, select the option, **Enable vPower NFS Server**.

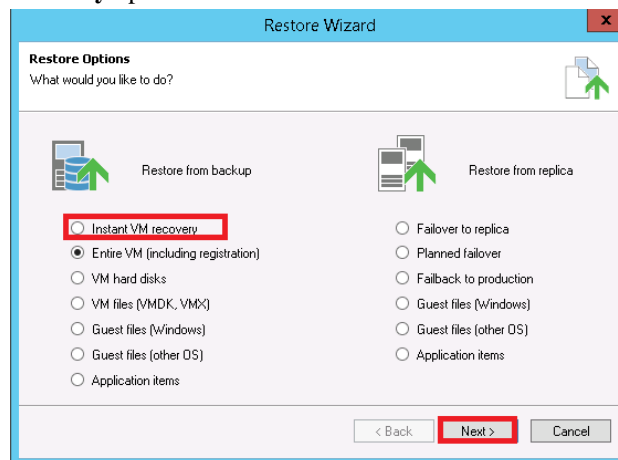
Note: You do not have to provide a folder for the NFS Datastore. In the case of the Hyper-V server, the cache data is directly stored at the Hyper-V server's datastore location and not to the NFS datastore path.




6.2.2 Performing Instant Recovery

Follow these steps to perform instant recovery for Hyper-V.

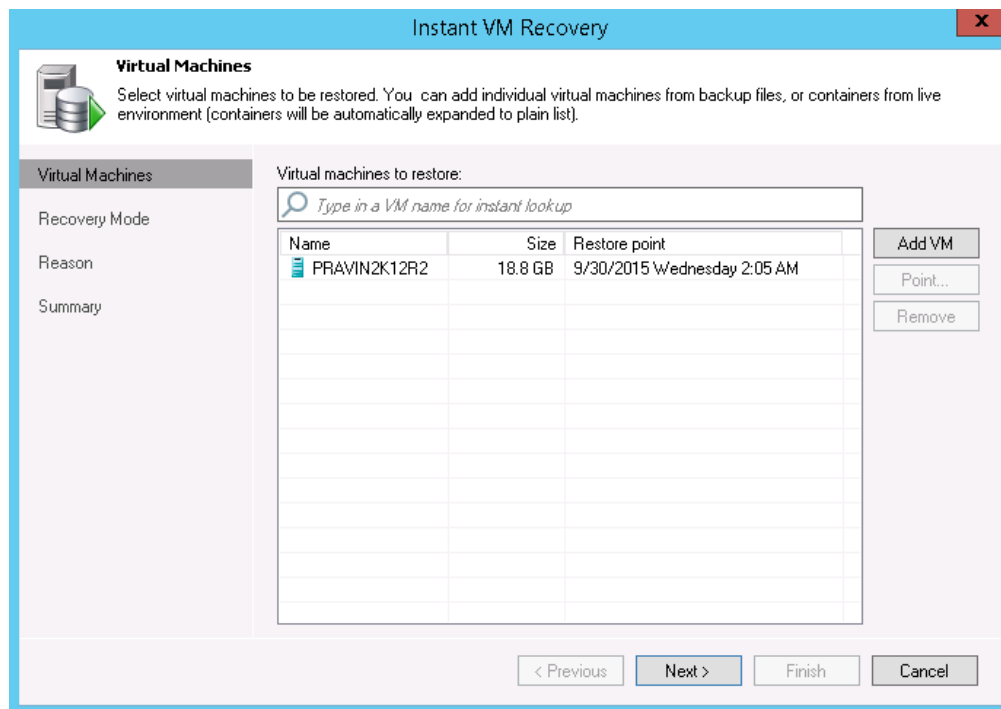
1. In the Veeam console, click the Restore Wizard option, select Hyper-V, and then select the **Instant VM recovery** option. Click **Next**.



2. Select the Virtual Machine to be recovered and then click **Add**.



3. Add the VM that needs to be recovered and click **Next**.



- At the Restore Mode step, select the option, **Restore to a new location or with different settings**, and then click **Next**.

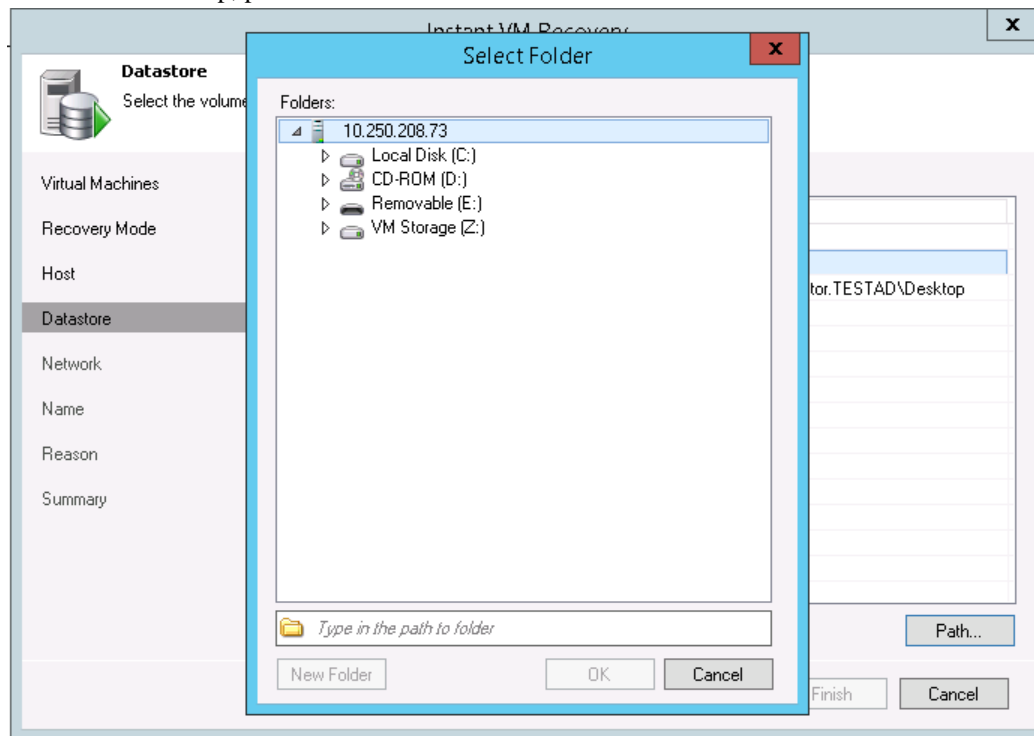
The screenshot shows the 'Instant Recovery' dialog box with the 'Restore Mode' tab selected. The left sidebar lists steps: Virtual Machine, Restore Point, **Restore Mode**, Destination, Datastore, Restore Reason, Ready to Apply, and Recovery. The main area has two radio button options: 'Restore to the original location' (unselected) and 'Restore to a new location, or with different settings' (selected). The selected option has a description: 'Customize restored VM location, and change its settings. The wizard will automatically populate all controls with the original VM settings as the default settings.' At the bottom are buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

- Select the Host to which to recover the VM and click **Next**.

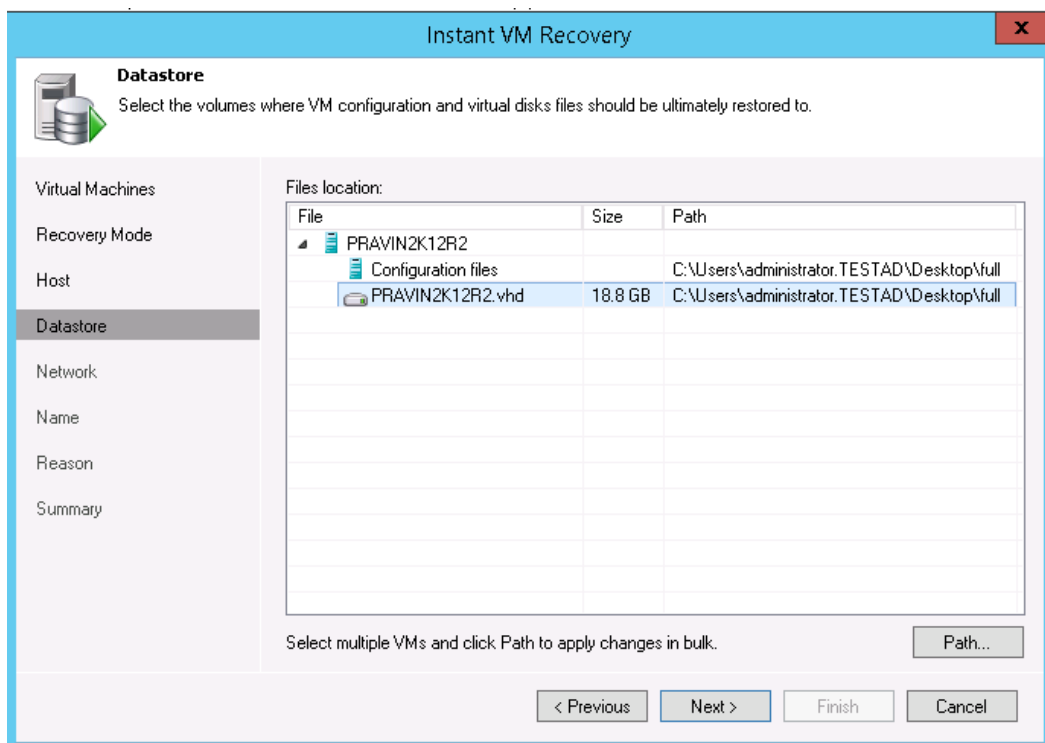
The screenshot shows the 'Instant VM Recovery' dialog box with the 'Host' tab selected. The left sidebar lists steps: Virtual Machines, Recovery Mode, **Host**, Datastore, Network, Name, Reason, and Summary. The main area is titled 'Select the host to recover VM to.' and contains a table labeled 'VM location:'. The table has three columns: Name, Host, and Cluster Resource. One row is populated with 'PRAVIN2K12R2' under Name and '10.250.208.73' under Host. Below the table is the text 'Select multiple VMs and click Host to apply changes in bulk.' and two buttons: 'Host...' and 'Resource...'. At the bottom are buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Name	Host	Cluster Resource
PRAVIN2K12R2	10.250.208.73	

6. At the Datastore step, provide the details of the cache data needed to be stored.




7. Provide the details of the Path where the VM cache data is stored.



8. In Network, select the Virtual Networks map to each other between original and new VM locations.

Instant VM Recovery X



Network

Select how virtual networks map to each other between original and new VM locations.

Virtual Machines

Recovery Mode

Host

Datastore

Network

Name

Reason

Summary

Network connections:

Source	Target
<div style="display: flex; align-items: center;"> <div style="width: 15px; height: 15px; background-color: #0070c0; margin-right: 5px;"></div> <div> <div style="font-size: 0.8em; font-weight: bold;">PRAVIN2K12R2</div> <div style="font-size: 0.7em; font-weight: bold;">Virtual Switch II</div> </div> </div>	Virtual Switch II

Select multiple VMs to apply settings change in bulk.

Network...

Disconnected

< Previous

Next >

Finish

Cancel

9. In the Restored VM name field, add the `_restored` suffix to the VM name and click **Next**.

Instant VM Recovery

Name

Specify the new virtual machine name, and whether you would like unique identifier preserved.

Virtual Machines

Recovery Mode

Host

Datastore

Network

Name

Reason

Summary

Virtual machines:

Name	New Name	VM UUID
PRAVIN2K12R2	PRAVIN2K12R2_restored	Create new

Select multiple VMs to apply settings change in bulk.

Name...

VM UUID...

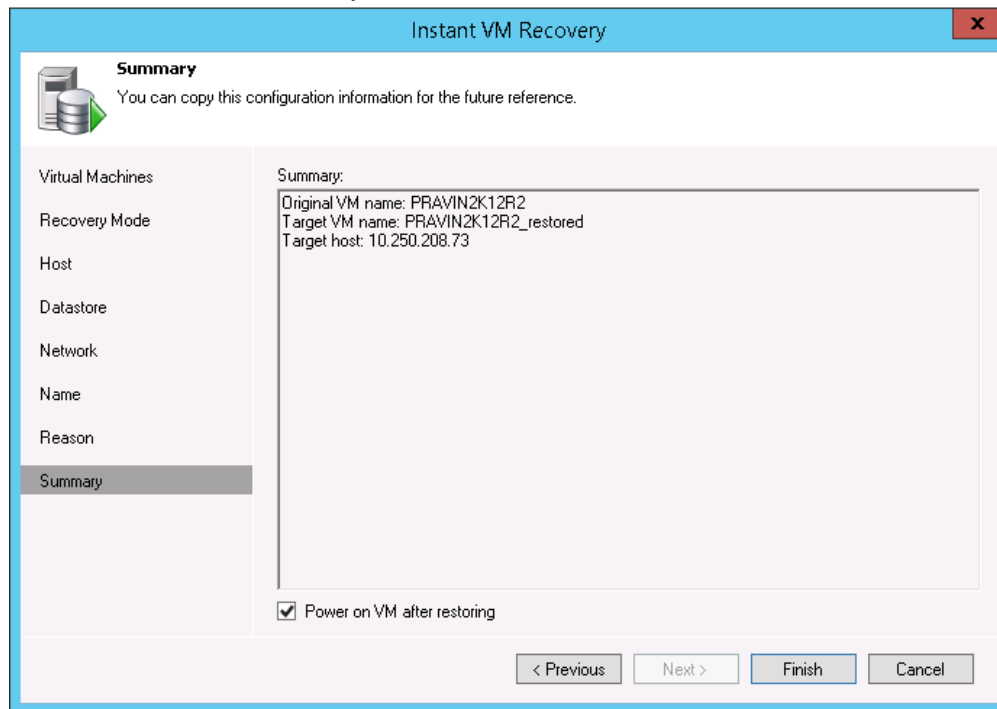
< Previous

Next >

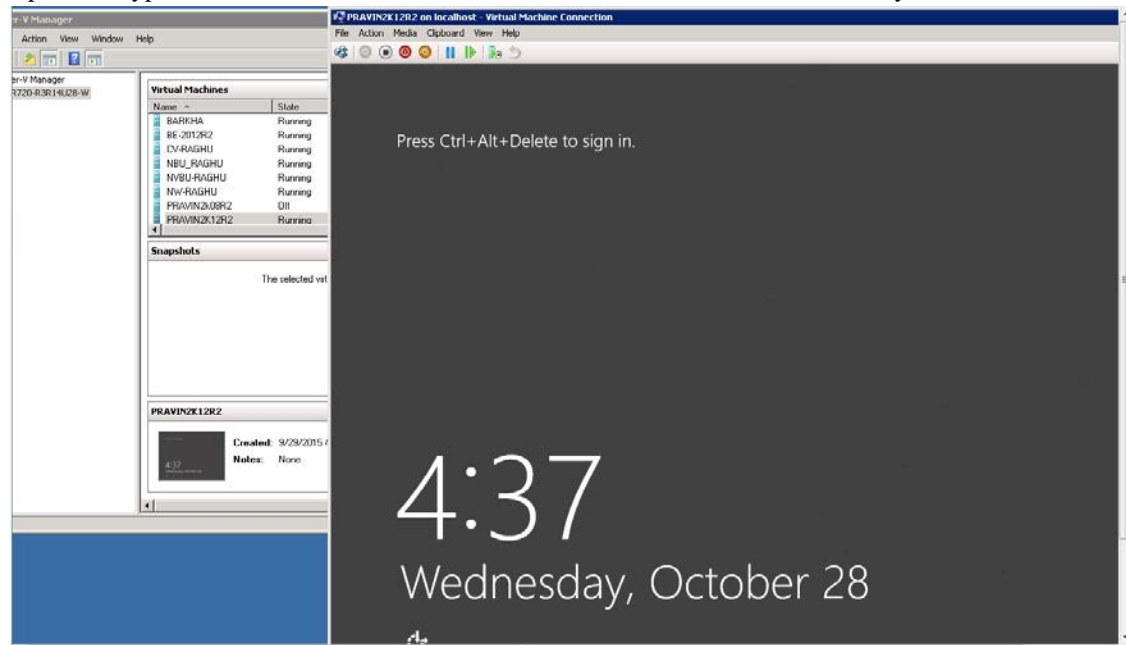
Finish

Cancel

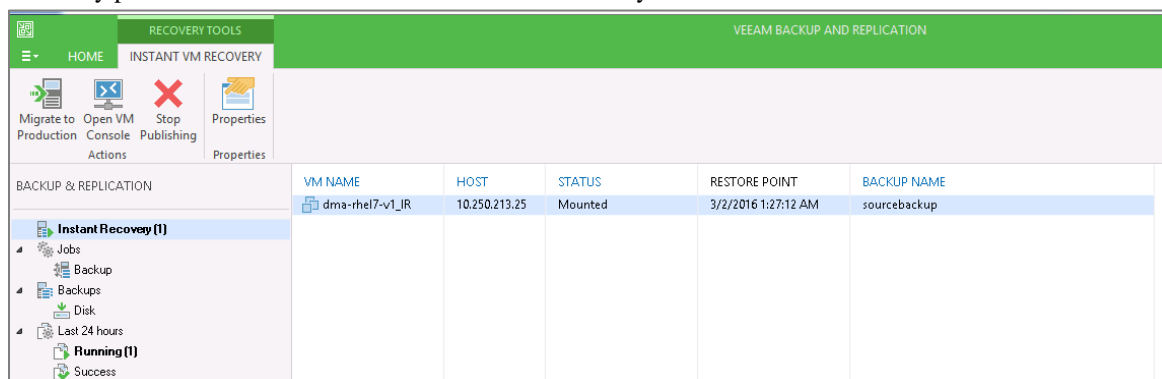
10. Click **Finish** to start the recovery.



11. Open the Hyper-v Client and make sure that the restored VM is started on the host you selected.



12. In Veeam Backup & Replication, open the Backup & Replication view, select the Instant Recovery node in the inventory pane and make sure that the Instant VM Recovery session is available and mounted.



6.3 Finalizing Instant VM recovery

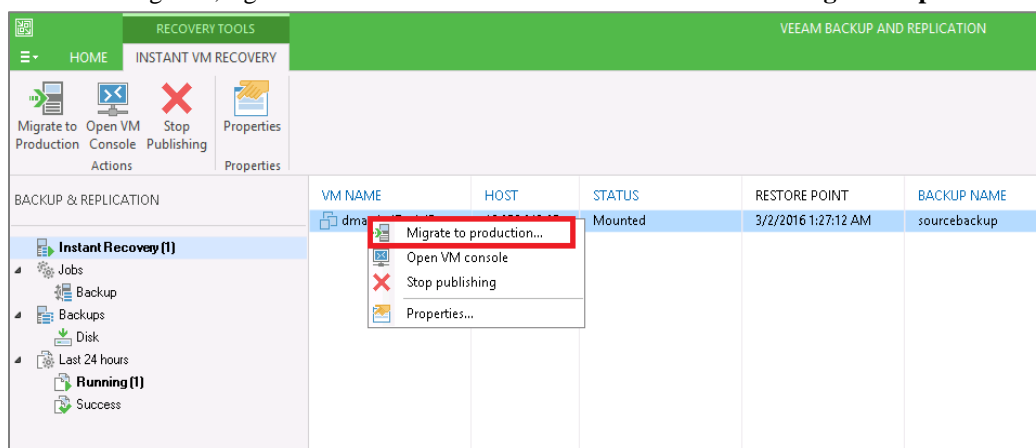
After Instant VM recovery is successfully completed, you can do one of the following:

- **Migrate VM to production** – Use this scenario if you have recovered a failed VM to the production ESX(i) host and want to permanently move the VM files to the production storage.
- **Terminate the Instant VM recovery session** – Use this scenario if you have recovered a VM for testing purpose and want to power it off and remove after testing is completed.

6.3.1 Migrating the VM to production

For VM migration, you can use VMware Storage vMotion, replicate or copy a VM to production with Veeam Backup & Replication, or use Veeam's Quick Migration. When you migrate the VM to production, you move the VM contents from the backup file to the production storage. The VM data is pulled from the backup and consolidated with changes made to the VM (redo logs). As a result, you get the VM in the latest state in your production environment. To migrate the restored VM with Quick Migration, follow these steps:

1. Open the **Backup & Replication** view in Veeam Backup & Replication.
2. In the inventory pane, select **Instant Recovery**.
3. In the working area, right-click the name of the recovered VM and select **Migrate to production**.



6.3.2 Terminating the Instant VM Recovery session

When you terminate the Instant VM Recovery session, the VM is unpublished from the ESX(i) host and redo logs are cleared from the vPower NFS datastore.

To terminate the current Instant VM recovery session, follow these steps:

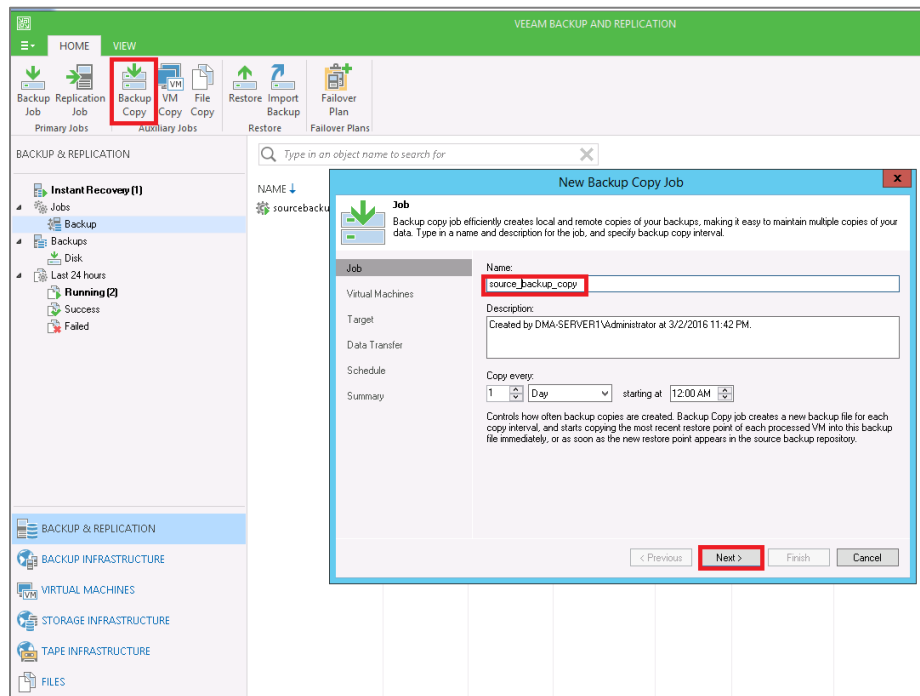
1. Open the **Backup & Replication** view in Veeam Backup & Replication
2. In the Inventory pane, select **Instant Recovery**
3. In the working area, right-click the name of the recovered VM and select **Stop publishing** as seen in the above screen (3rd Option).



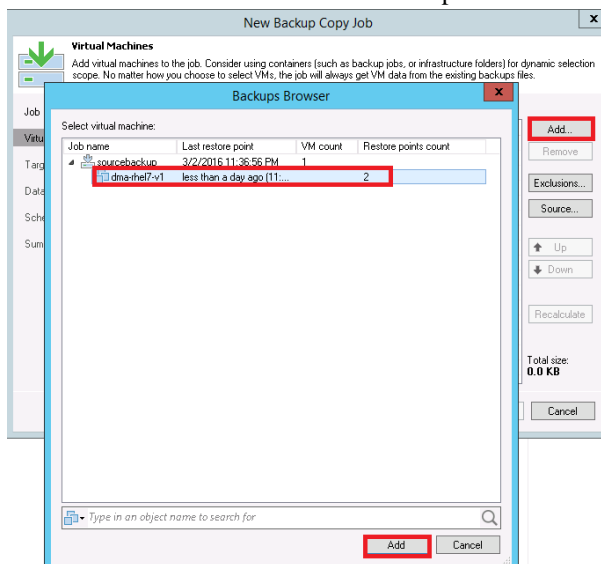
Backup copy operation

The main backup purpose is to protect your data against disasters and VM failures. However, having just one backup does not provide the necessary level of safety. Your primary backup may get destroyed along with your production data, leaving you with no backup to restore from. The backup copy job is a separate task that needs to be set apart from the backup job. Veeam Backup copy allows user to copy backup data to secondary storage with the help of Backup copy Jobs. To create a backup copy job, follow these steps:

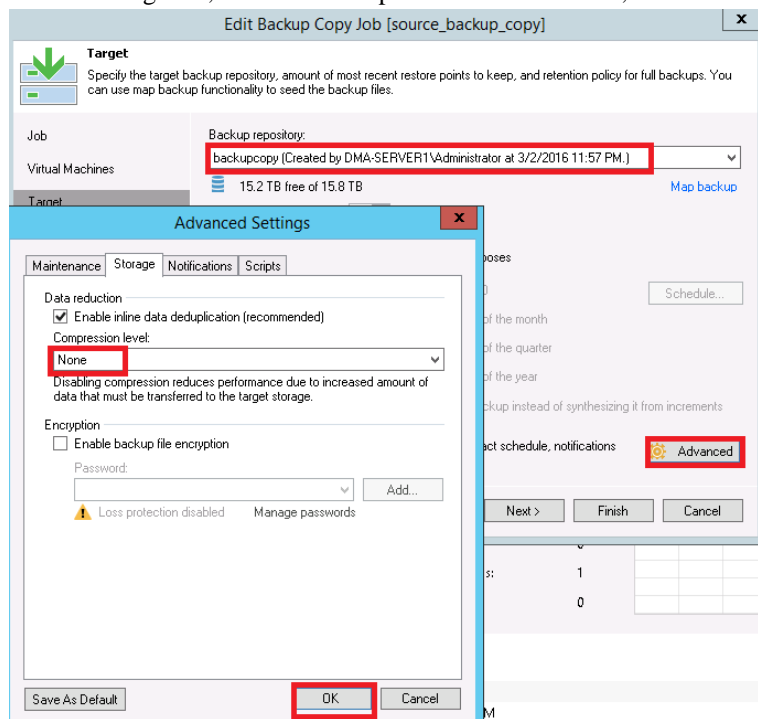
1. Click the backup copy in Auxiliary jobs, provide the name of the job, and click **Next**.



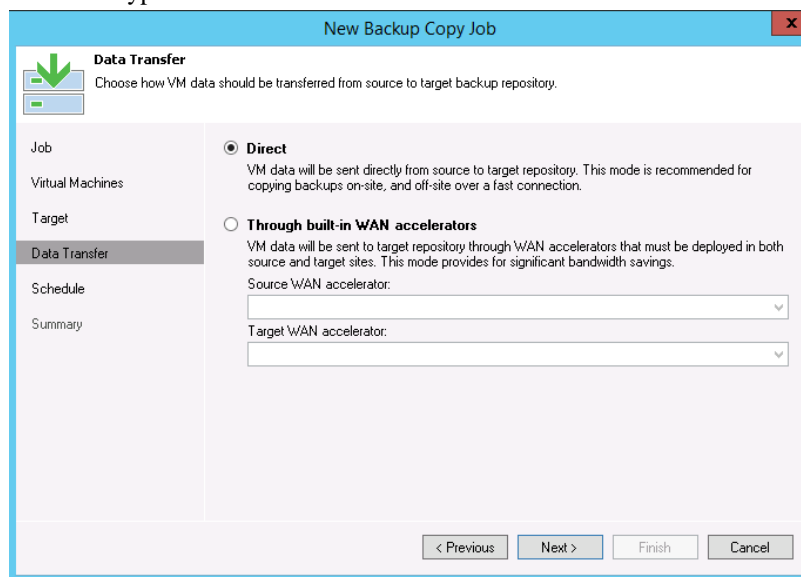
2. Add the Virtual Machine from the backup Jobs.



3. Provide the backup repository and click **Advanced**.
4. On the Storage tab, select the Compression level as **None**, and click **OK**.



5. Select the type of data transfer and click **Next**.



6. Schedule the job as needed.

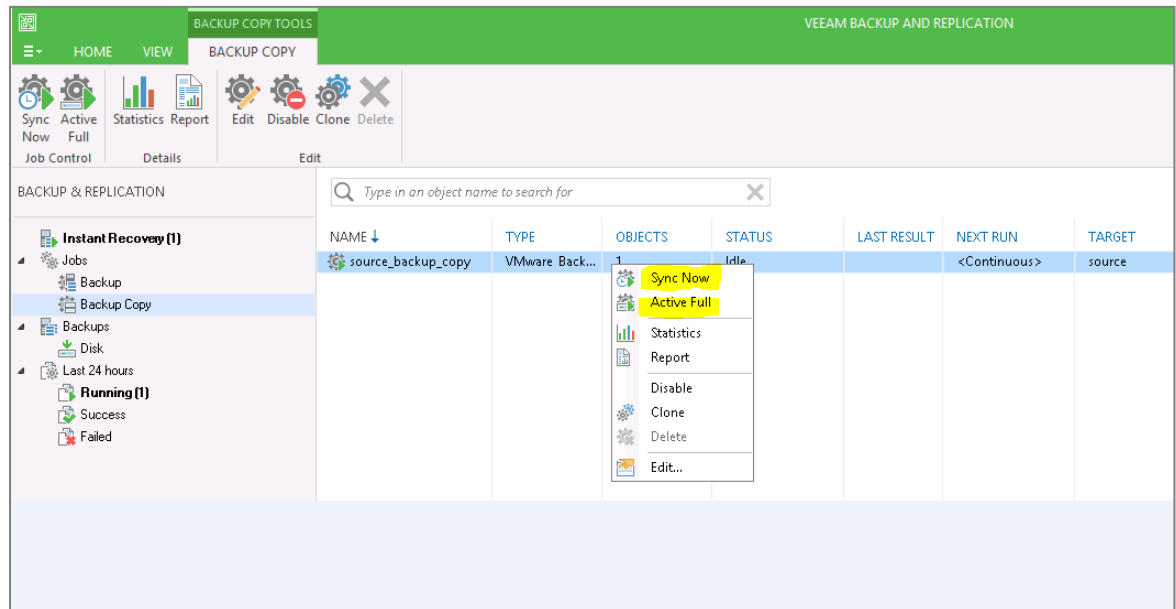
The screenshot shows the 'New Backup Copy Job' wizard with the 'Schedule' tab selected. The left sidebar contains links for Job, Virtual Machines, Target, Data Transfer, Schedule (selected), and Summary. The main area is titled 'Schedule' and includes a description: 'Specify when this job is allowed to transfer data over the network. Backup copy jobs run continuously, starting data transfers according to copy interval and/or as the new VM restore points appear.' Below this, there are two radio buttons: 'Any time (continuously)' (selected) and 'During the following time periods only:'. The latter option is accompanied by a calendar grid showing days of the week and time slots (12, 2, 4, 6, 8, 10, 12). To the right of the calendar is a legend with 'Enable' (represented by a grey square) and 'Disable' (represented by a white square). At the bottom, there are four buttons: '< Previous', 'Create', 'Finish', and 'Cancel'.

7. Click Finish

The screenshot shows the 'New Backup Copy Job' wizard with the 'Summary' tab selected. The left sidebar contains links for Job, Virtual Machines, Target, Data Transfer, Schedule, and Summary (selected). The main area is titled 'Summary' and includes a description: 'Review the settings, and click Finish to save and exit the wizard.' Below this, there is a 'Summary:' section with a text box containing the following information: 'Name: source_backup_copy', 'Target Path: \\10.250.241.231\\source', 'Type: VMware Backup Copy', and 'Source items: dma-hel7-v1'. At the bottom, there is a checkbox labeled 'Enable the job when I click Finish' which is checked. At the bottom right, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

8. Select one of the following operations as needed:

- **Sync Now** – Traditional Veeam backup copy job where it syncs the restore points from source storage to target.
- **Active Full** – This added feature in Veeam 9.0 helps improve local (on-site) backup copy performance and reduces the load on deduplication appliances by eliminating the data rehydration required to process the backup copy job retention policy, or to create a new GFS (Grandfather-Father-Son) restore point. Enabling this option will disable a full backup transformation (oldest incremental backups will no longer be merged into the full backup file for retention processing). Instead, GFS full backup files will be created by copying the most recent VM state data from the primary backup storage in its entirety.



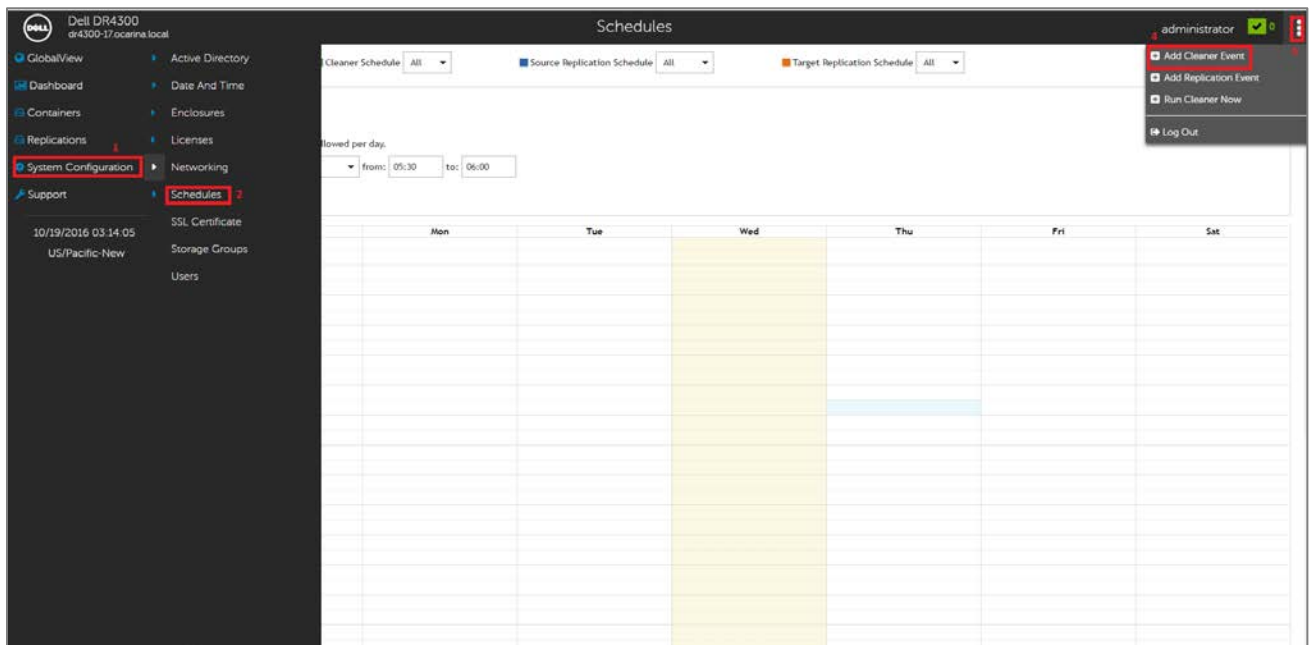
8 Setting up the DR Series system cleaner

Performing scheduled disk space reclamation operations are recommended as a method for recovering disk space from system containers in which files were deleted as a result of deduplication.

The cleaner runs during idle time. If your workflow does not have a sufficient amount of idle time on a daily basis, then you should consider scheduling the cleaner to force it to run during a scheduled time.

If necessary, you can perform the procedure shown in the following screenshot to force the cleaner to run. After all of the backup jobs are set up, the DR Series system cleaner can be scheduled. The DR Series system cleaner should run at least 40 hours per week when backups are not taking place and generally after a backup job has completed.

Dell recommends scheduling the cleaner at a time separate from backup and replication jobs.



9 Monitoring deduplication, compression, and performance

After backup jobs have run, the DR Series system tracks capacity, storage savings, and throughput on the DR Series system dashboard. This information is valuable in understanding the benefits of the DR Series system.

Note: Deduplication ratios increase over time. It is not uncommon to see a 2-4x reduction (25-50% total savings) on the initial backup. As additional full backup jobs are completed, the ratios will increase. Backup jobs with a 12-week retention will average a 15x ratio, in most cases.

