



Setting Up the Dell™ DR Series System on Veeam

Dell Engineering
April 2016

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 information about Instant Recovery with DR Series.
April 2016	Updated with Veeam 9.0 screenshots and features.

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

Executive summary.....	4
1 Installing and configuring the DR Series system	5
2 Setting up Veeam.....	13
3 Setting up DR Series native replication and restore from a replication target container	29
3.1 Building a replication relationship between DR Series systems.....	29
3.2 Restoring data from a target DR Series system	31
4 Using Veeam Instant VM Recovery with the DR Series system.....	38
4.1 Instant Recovery for ESX VM backups.....	38
4.2 Instant Recovery with Hyper-V Server.....	44
4.3 Finalizing Instant VM Recovery	51
4.3.1 Migrating a VM to production	51
4.3.2 Terminating an Instant VM Recovery session	52
5 Creating a backup copy	53
6 Setting up the DR Series system cleaner	59
7 Monitoring deduplication, compression, and performance	60



Executive summary

This paper provides information about how to set up the Dell 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://www.dell.com/support/home>

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.

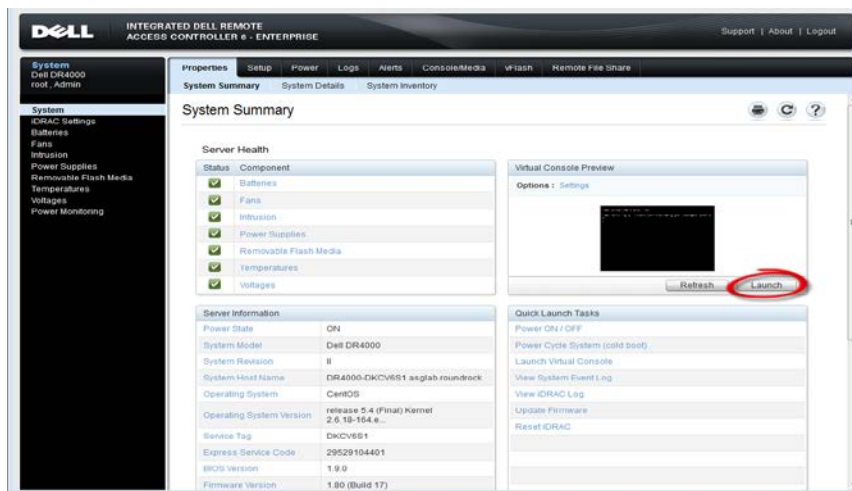


1 Installing and configuring the DR Series system

1. Rack and cable the DR Series system, and power it on.

In the *Dell DR Series System Administrator Guide*, refer to the sections, “iDRAC Connection”, “Logging in and Initializing the DR Series System”, and “Accessing iDRAC6/Idrac7 Using RACADM” for information about using iDRAC connection and initializing the appliance.

2. Log on to iDRAC by using the default address **192.168.0.120**, or the IP address that is assigned to the iDRAC interface. Use the user name and password of “**root/calvin**” and then launch the virtual console.



3. 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).



4. 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:
```

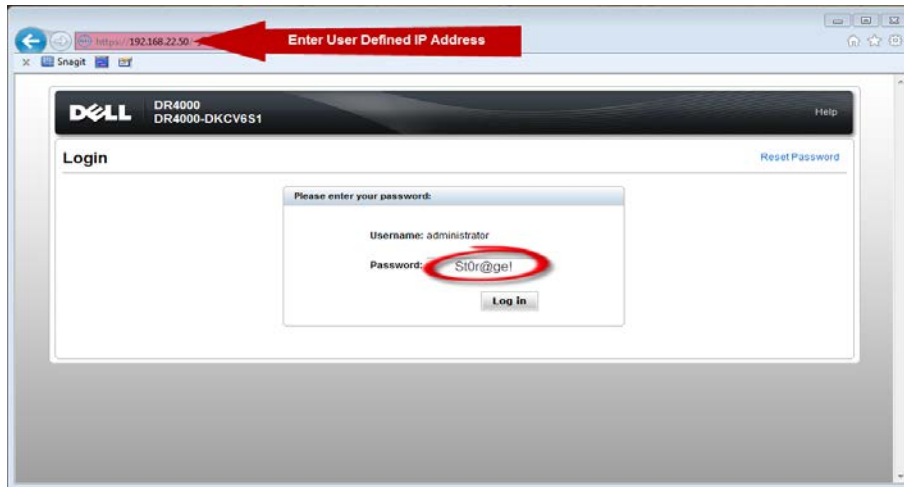
5. View the summary of preferences and confirm that it is correct.

```
=====
                          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       : DR4000-5

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



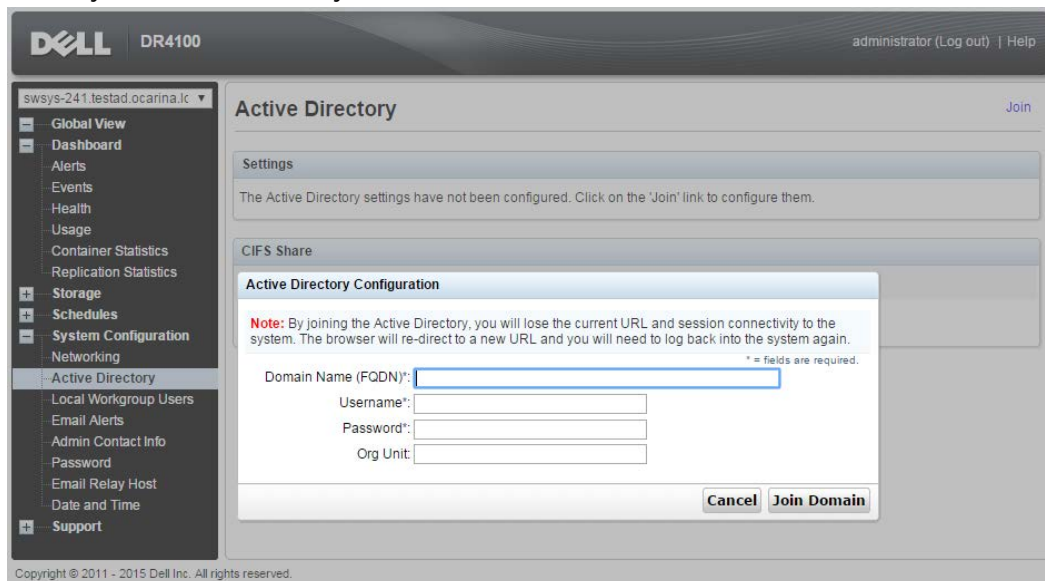
6. Log on to the DR Series System administrator console with the IP address you just provided for the DR Series system, with the username **administrator** and password **St0r@ge!** (The "0" in the password is the numeral zero).



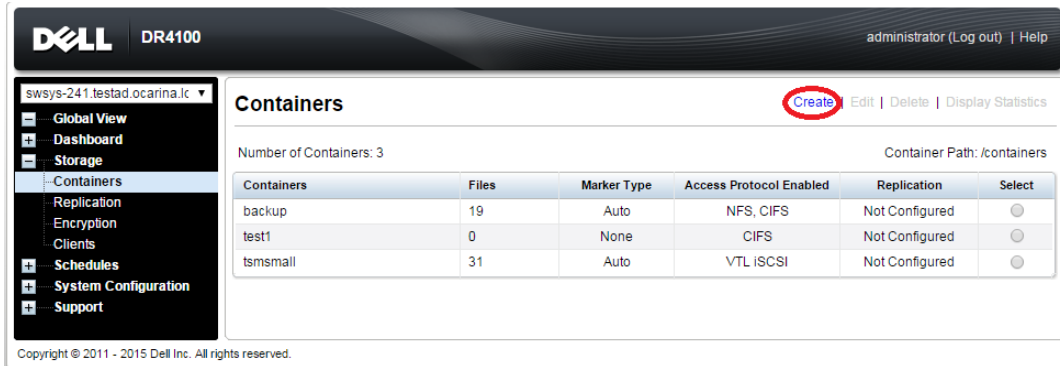
7. Join the DR Series system into the Active Directory domain.

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. Under System Configuration in the left navigation area, click **Active Directory**.
- b. Enter your Active Directory credentials.



- To create the container, in the left navigation area, click **Containers** and then click the **Create** link at the top of the page.

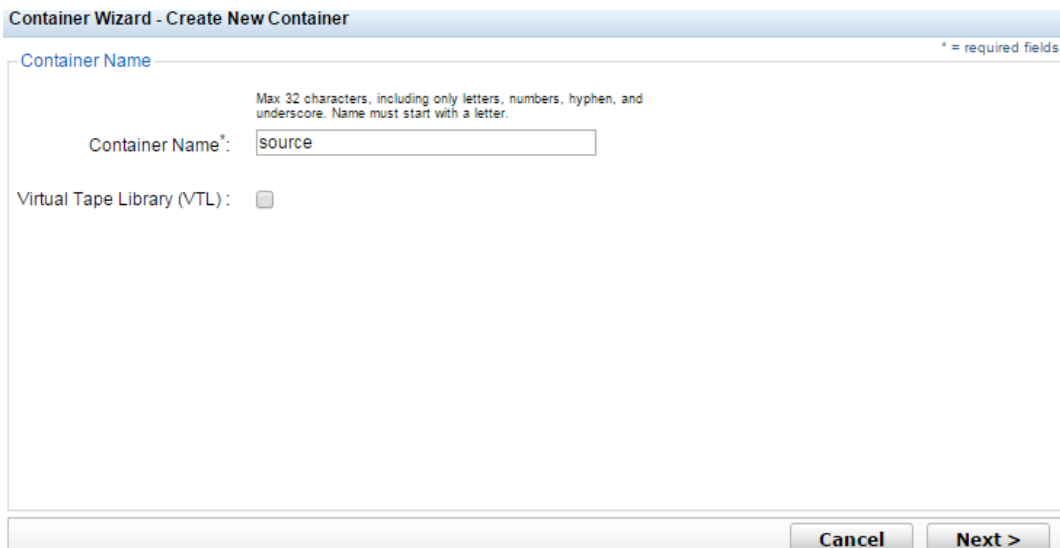


The screenshot shows the Dell DR4100 web interface. The top navigation bar includes the Dell logo, 'DR4100', and user information 'administrator (Log out) | Help'. The left sidebar contains a navigation menu with 'Containers' selected. The main content area is titled 'Containers' and features a 'Create' link circled in red, along with 'Edit', 'Delete', and 'Display Statistics' links. Below this, it states 'Number of Containers: 3' and 'Container Path: /containers'. A table lists the containers:

Containers	Files	Marker Type	Access Protocol Enabled	Replication	Select
backup	19	Auto	NFS, CIFS	Not Configured	<input type="radio"/>
test1	0	None	CIFS	Not Configured	<input type="radio"/>
tsmsmall	31	Auto	VTL iSCSI	Not Configured	<input type="radio"/>

Copyright © 2011 - 2015 Dell Inc. All rights reserved.

- Enter a Container Name, and click **Next**.



The screenshot shows the 'Container Wizard - Create New Container' form. The title bar includes the text '* = required fields'. The form has the following fields:

- Container Name:** A text input field containing 'source'. Above the field is a note: 'Max 32 characters, including only letters, numbers, hyphen, and underscore. Name must start with a letter.'
- Virtual Tape Library (VTL):** A checkbox that is currently unchecked.

At the bottom of the form are two buttons: 'Cancel' and 'Next >'.



10. Select the access protocol (**NAS (NFS, CIFS)**).

Container Wizard - Create New Container

Select Access Protocols

* = required fields

Storage Access Protocol*: Dell Rapid Data Storage (RDS) Symantec OpenStorage (OST) NAS (NFS, CIFS)

Container Name and Type
source

< Back Cancel Next >

11. For CIFS, select the **CIFS** check box, set the marker type as **None**, and then click **Next**.

Edit Container: source

Configure NAS Access

* = required fields

Enable Access Protocols: NFS (Use NFS to backup UNIX or LINUX clients) CIFS (Use CIFS to backup MS Windows clients)

Marker Type*: None Auto Networker Unix Dump BridgeHead Time Navigator

Container Name and Type
source

Access Protocols
NAS (NFS, CIFS)

Cancel Next >

For NFS, select the **NFS** check box and then click **Next**.

The screenshot shows the 'Container Wizard - Create New Container' dialog box, specifically the 'Configure NAS Access' step. The title bar includes the text '* = required fields'. The main area is divided into two sections. On the left, under 'Enable Access Protocols', the 'NFS' checkbox is checked and highlighted with a red box, with the text '(Use NFS to backup UNIX or LINUX clients)' next to it. The 'CIFS' checkbox is unchecked, with the text '(Use CIFS to backup MS Windows clients)' next to it. Below this, the 'Marker Type*' section has a radio button selected for 'None', with a help icon (?) to its right. Other options include 'Auto', 'Networker', 'Unix Dump', 'BridgeHead', and 'Time Navigator'. On the right side, the 'Container Name and Type' section shows 'target_nfs'. Below that, the 'Access Protocols' section lists 'NAS (NFS, CIFS)'. At the bottom of the dialog, there are three buttons: '< Back', 'Cancel', and 'Next >', with the 'Next >' button highlighted by a red box.

12. For CIFS, set the preferred client access credentials, and then click **Next**

The screenshot shows the 'Container Wizard - Create New Container' dialog box, specifically the 'Configure CIFS Client Access' step. The title bar includes the text '* = required fields'. The main area is divided into two sections. On the left, under 'Client Access', the 'Open (allow all clients)' radio button is selected and highlighted with a red box. The 'Create Client Access List' radio button is unselected. Below this, the 'Client FQDN or IP:' field is empty, with an 'Add' button to its right. The 'allow access client(s)' section contains an empty list box with an 'Add' button above it and a 'Remove' button to its right. On the right side, the 'Container Name and Type' section shows 'source'. Below that, the 'Access Protocols' section lists 'NAS (NFS, CIFS)' and 'None'. At the bottom of the dialog, there are three buttons: '< Back', 'Cancel', and 'Next >', with the 'Next >' button highlighted by a red box.



For NFS, set the following preferred client access credentials and then click **Next**

Container Wizard - Create New Container * = required fields

Configure NFS Access

NFS Options *: Read Write Access Read Only Access Insecure

Map root to :

Client Access : Open (allow all clients) Create Client Access List

Client FQDN or IP :

allow access client(s)

Container Name and Type
target_nfs

Access Protocols
NAS (NFS, CIFS)
None

13. Check the configuration summary, and then click **Create a New Container**.

Container Wizard - Create New Container * = required fields

Configuration Summary

Container Name and Type
Container Name: source

Access Protocols
Access Protocol: NAS (NFS, CIFS)
Marker Type: None

CIFS Access
Open (allow all clients):



14. Confirm that the container is successfully added.

The screenshot shows the Dell DR4000 interface. The top header includes the Dell logo, 'DR4000', and 'root (Log out) | Help'. The left sidebar shows a navigation menu with 'Containers' selected. The main content area is titled 'Containers' and includes a 'Message' box with a green checkmark icon and the following text:

- Successfully added container "source".
- Successfully added CIFS connection for container "source".
- Container 'source' has the following marker(s) None.

Below the message, it states 'Number of Containers: 8' and 'Container Path: /containers'. A table lists the containers:

Containers	Files	Marker Type	Access Protocol Enabled	Replication	Select
backup	3	Auto	NFS, CIFS	Not Configured	<input type="radio"/>
NBU	28	None	CIFS	Not Configured	<input type="radio"/>
nbu77-nfs1	36	Auto	NFS	Not Configured	<input type="radio"/>
prvn-canfs	2	None	NFS	Not Configured	<input type="radio"/>
rdcifs	8	None	CIFS	Not Configured	<input type="radio"/>
rdnfs	8	None	NFS	Not Configured	<input type="radio"/>
repNBU	0	None	CIFS	Not Configured	<input type="radio"/>
source	0	None	CIFS	Not Configured	<input type="radio"/>

Copyright © 2011 - 2015 Dell Inc. All rights reserved.

The screenshot shows the Dell DR4000 interface. The top header includes the Dell logo, 'DR4000', and 'root (Log out) | Help'. The left sidebar shows a navigation menu with 'Containers' selected. The main content area is titled 'Containers' and includes a 'Message' box with a green checkmark icon and the following text:

- Successfully added container "target_nfs".
- Successfully added NFS connection for container "target_nfs".
- Container 'target_nfs' has the following marker(s) None.

Below the message, it states 'Number of Containers: 2' and 'Container Path: /containers'. A table lists the containers:

Containers	Files	Marker Type	Access Protocol Enabled	Replication	Select
backup	0	Auto	NFS, CIFS	Not Configured	<input type="radio"/>
target_nfs	33743	None	NFS	Not Configured	<input type="radio"/>




2 Setting up Veeam

Notes:

To maximize the DR Series system and Veeam deduplication savings, Dell recommends to 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 same settings.

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



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

Edit Backup Repository

Name
Type in a name and description for this backup repository.

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, do the following:
- Select **Shared folder** as the type of backup repository, and click **Next**.

New Backup Repository

Type
Choose type of backup repository you want to create.

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

- b. 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 dialog has a blue title bar and a sidebar on the left with options: Name, Type, Share (selected), Repository, vPower NFS, Review, and Apply. The main area contains the following fields and controls:

- Shared folder:** A text box containing the UNC path '\\10.250.241.229\source' and a 'Browse...' button.
- Access Credentials:** A checked checkbox labeled 'This share requires access credentials:'. Below it is a 'Credentials:' field showing 'Administrator (Administrator, last edited: 9/15/16)' and an 'Add...' button. A 'Manage accounts' link is also present.
- Gateway server:** Two radio button options: 'Automatic selection' (selected) and 'The following server:'. Below the second option is a dropdown menu currently showing 'This server'.
- Help text:** Below the gateway server options, it says: 'Use this option to improve performance and reliability of backup to a NAS located in a remote site.'

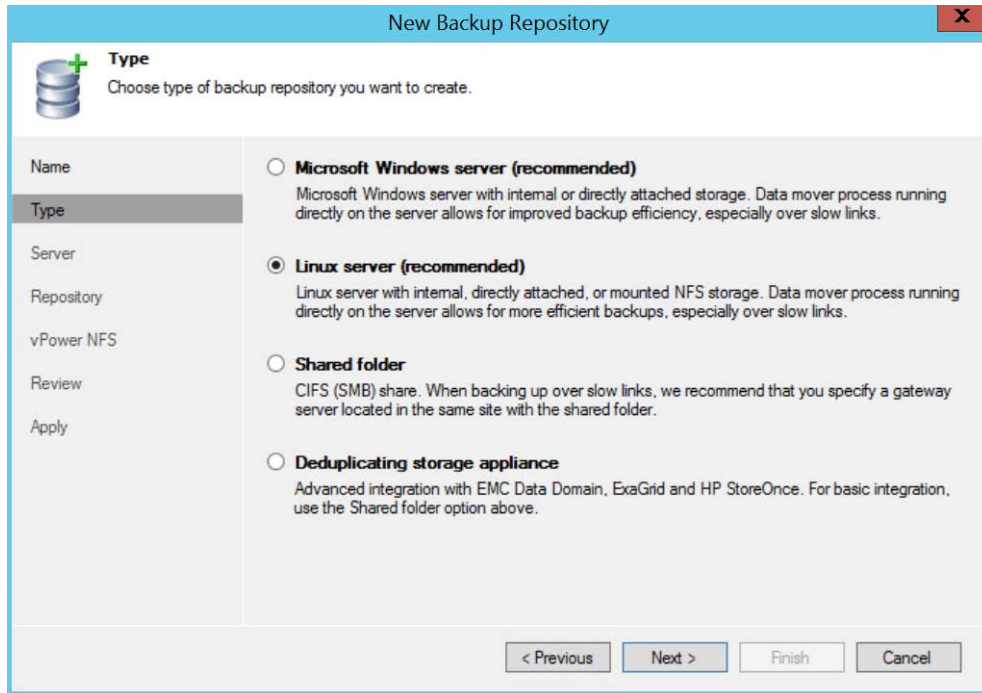
At the bottom of the dialog are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.



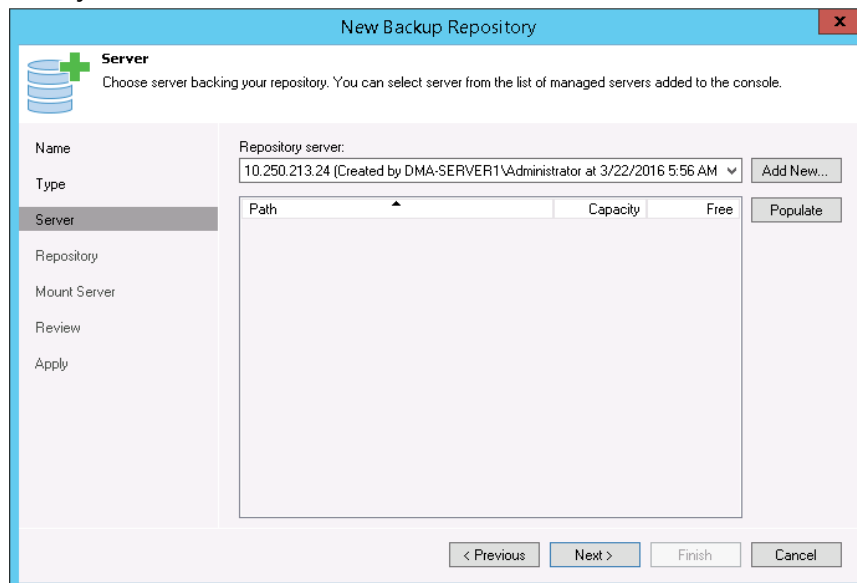
5. For an NFS repository (for a Linux server), do the following:

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

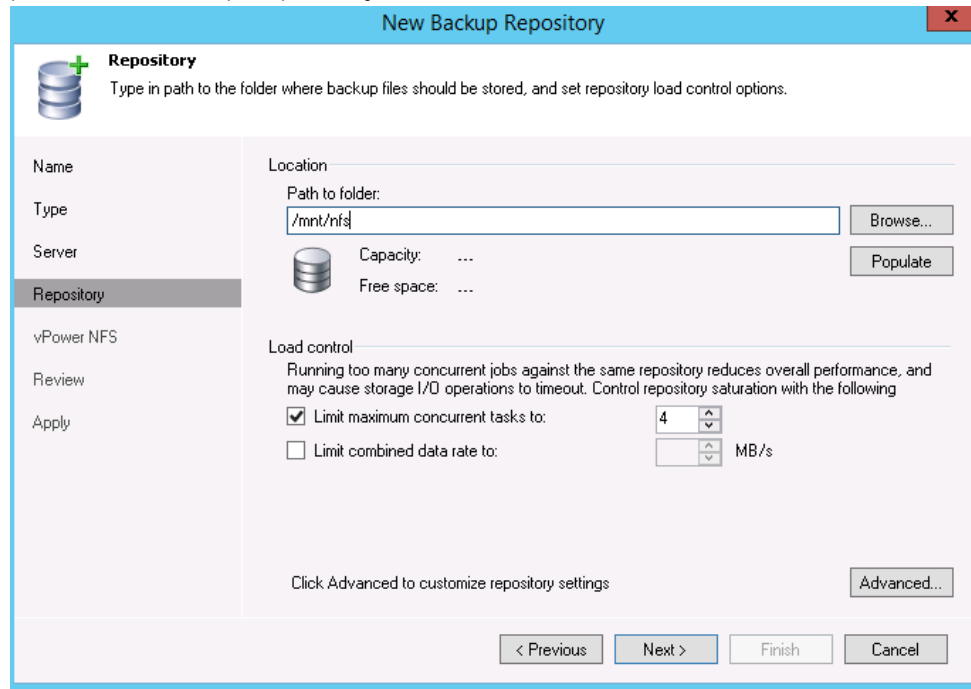
- a. Select **Linux Server (recommended)** as the type of Backup Repository, and then click **Next**.




- b. Add the New Repository server (Linux), or select a server from the list if one has been added already.



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




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



Note: Refer to the latest *Dell DR Series System Interoperability Guide* for the **maximum concurrent jobs** supported for CIFS/NFS based on the DR Series system model. The maximum concurrent jobs setting also depends on the number of CPU cores in the Veeam server. To run more tasks in parallel, you can add more **Backup proxy servers** to the Veeam server as needed.

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

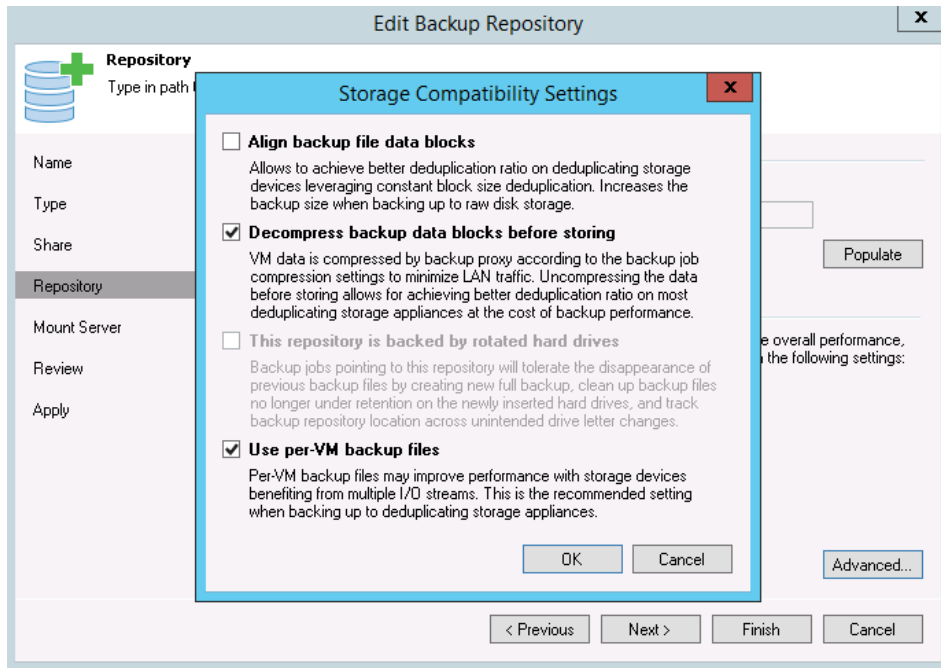
Note: Clearing the selection, **Decompress backup data blocks before storing**, might increase your overall deduplication storage capacity usage. It is not recommended to switch these settings after the data has been written to the DR Series system.



WARNING: Dell recommends that you do **not** change the setting for the option, **Align backup file data blocks**, after backups are taken as this could impact the deduplication savings for further backups.

- To create separate backup files for VMs in the job, select the option, **Use Per-VM Backup Files Chains**. This setting makes any backup job that is writing to a repository store each VM's restore point in a dedicated backup file.

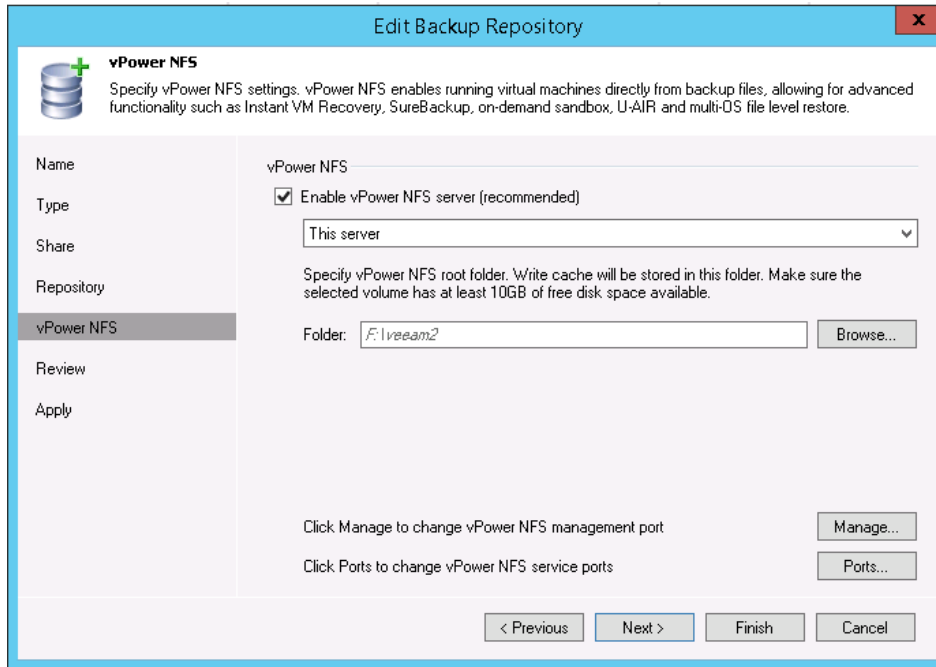
If you decide to create separate backup files for VMs in the job, make sure that you also enable parallel data processing.



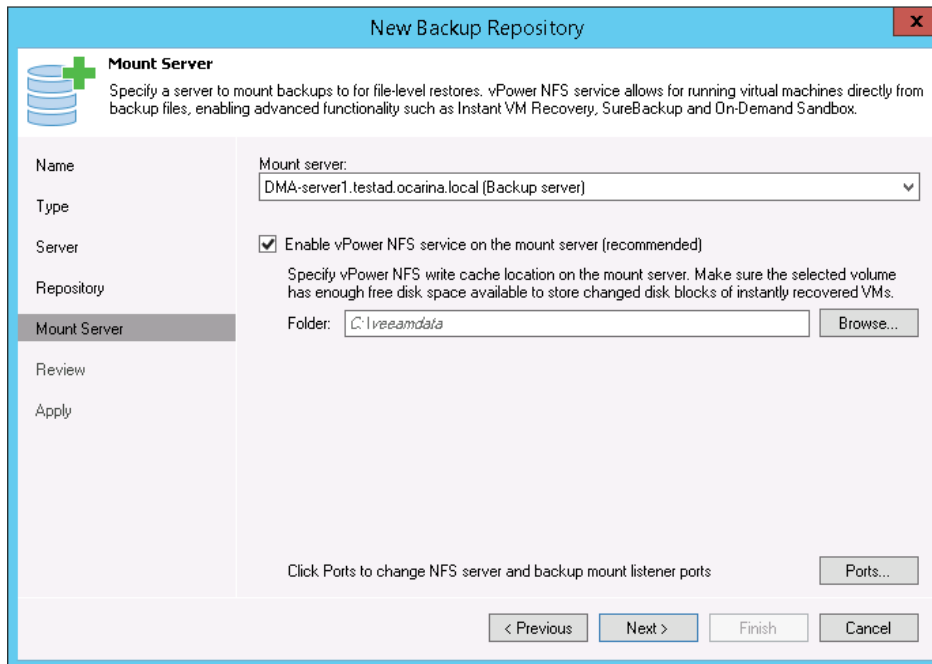
Note: This setting enables multiple write streams within a single job with parallel processing enabled. Enabling multiple streams dramatically improves overall job backup performance. Dell recommends that you use the per-VM backup files option for better backup throughput.

- Click **Next**.

10. For Instant Recovery to work, select the option, **Enable vPower NFS Server**.
For CIFS, enable the option as below -

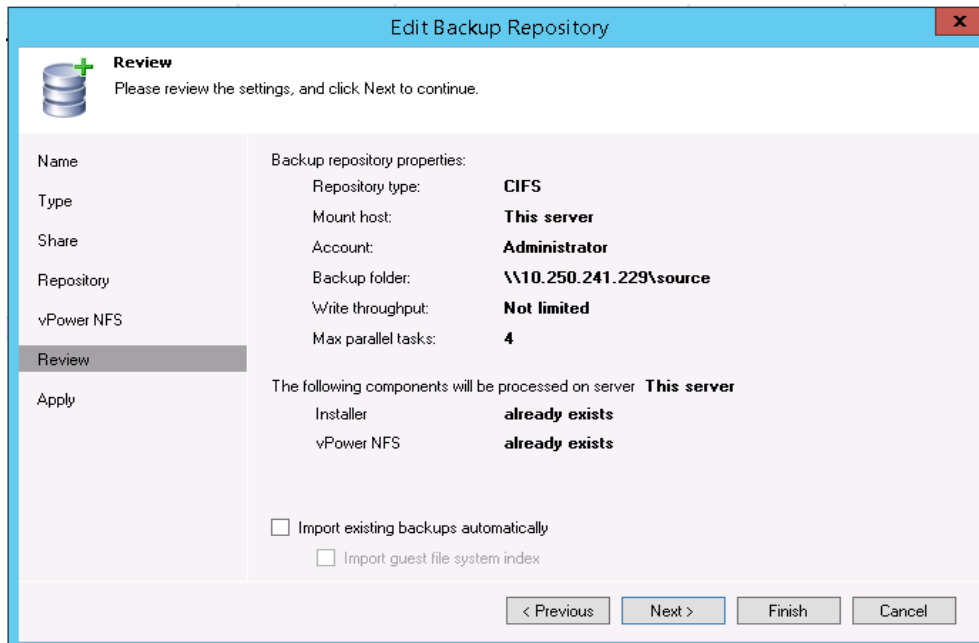


For NFS, enable the option as below -

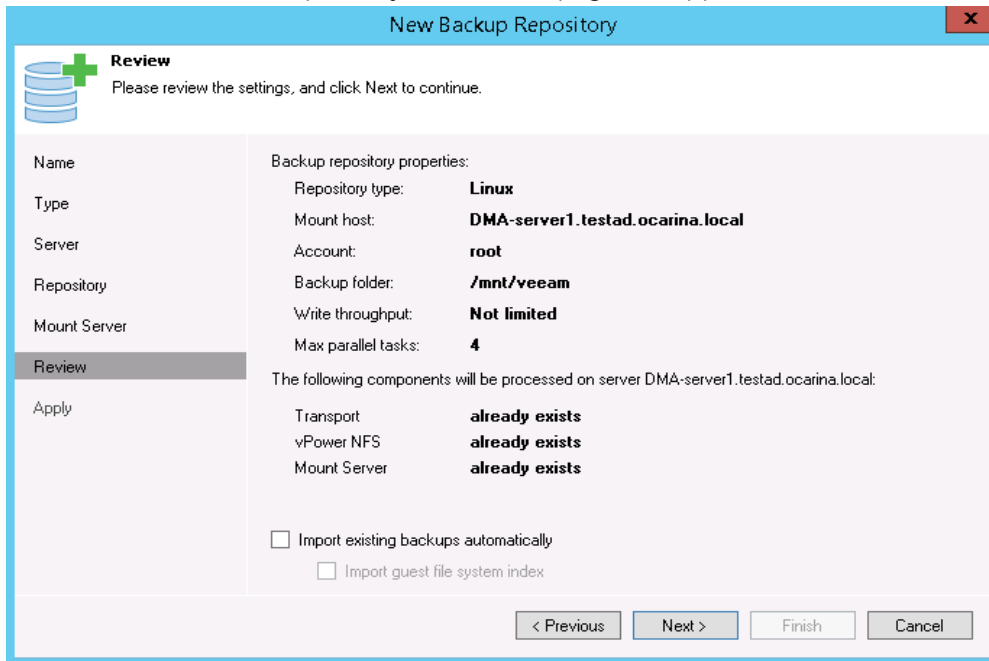


11. On the review page, verify the settings, and click **Next** to apply changes.
For a CIFS Container Repository, the Review page will appear similar to the following example.

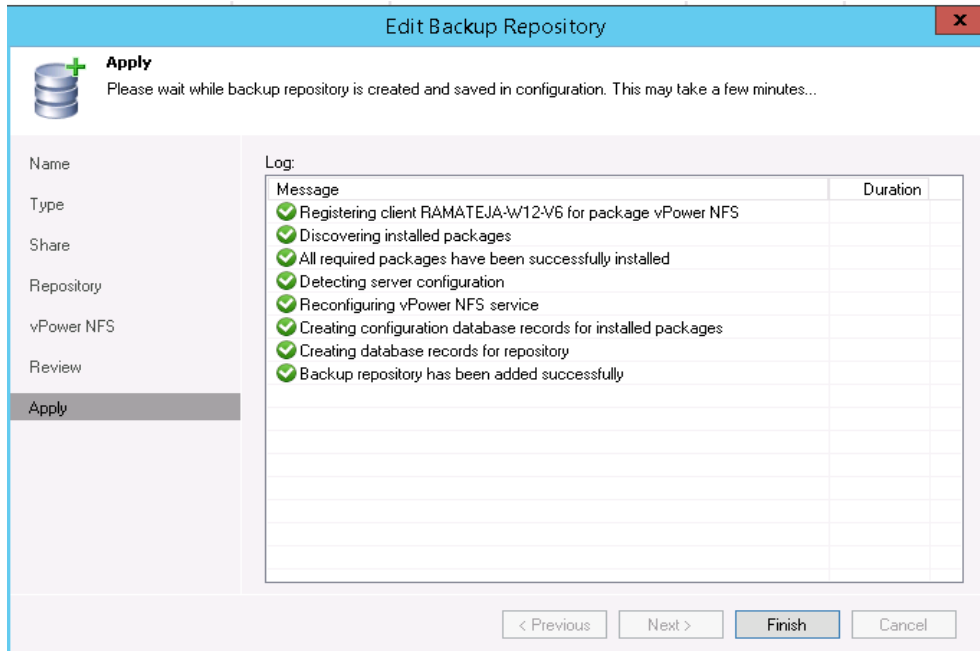




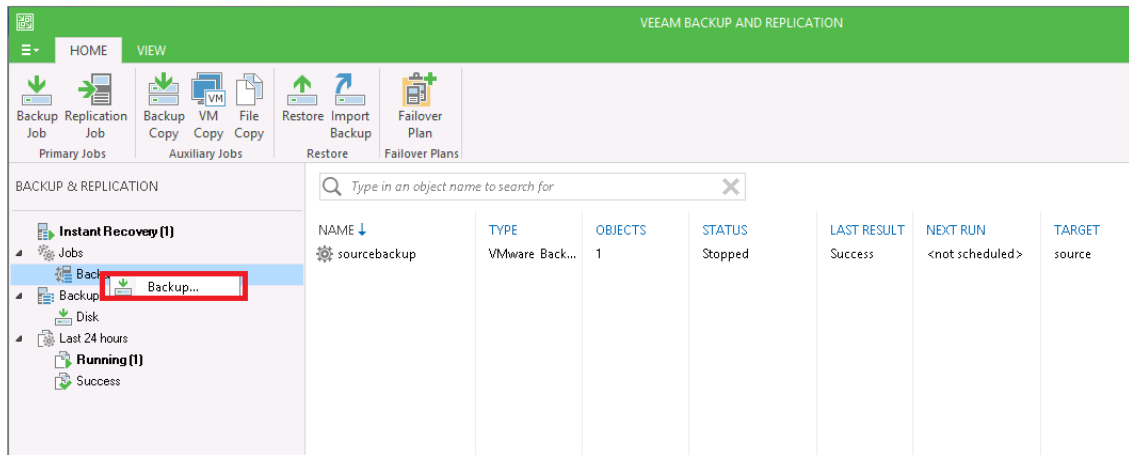
For an NFS Container Repository, the Review page will appear similar to the following example.




12. Click **Finish**.



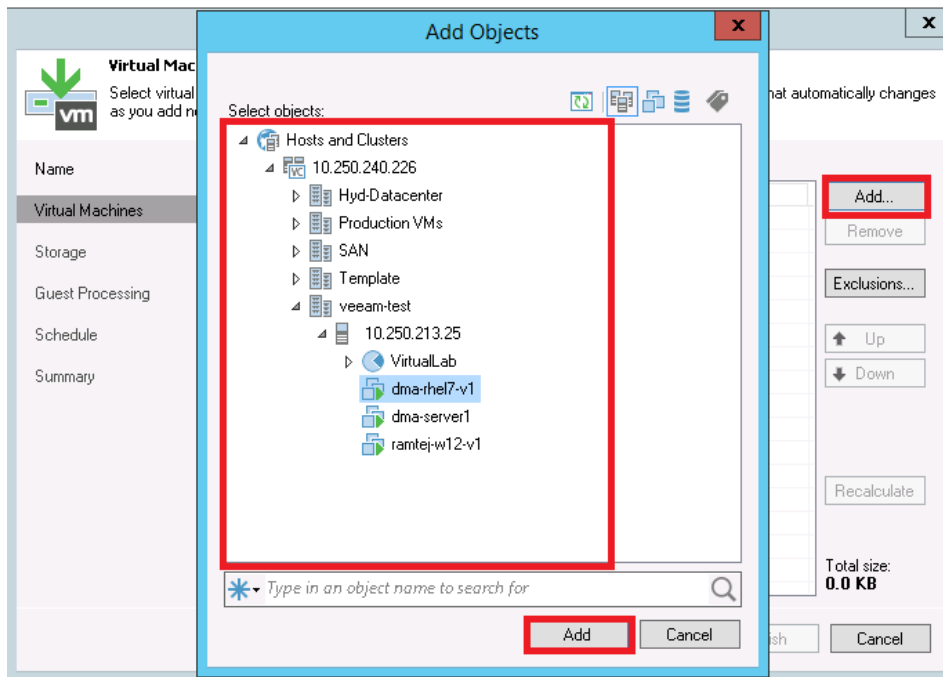
13. On the Backup & Replication menu, go to **Jobs > Backup**, and right-click **Backup** to create a new backup job.



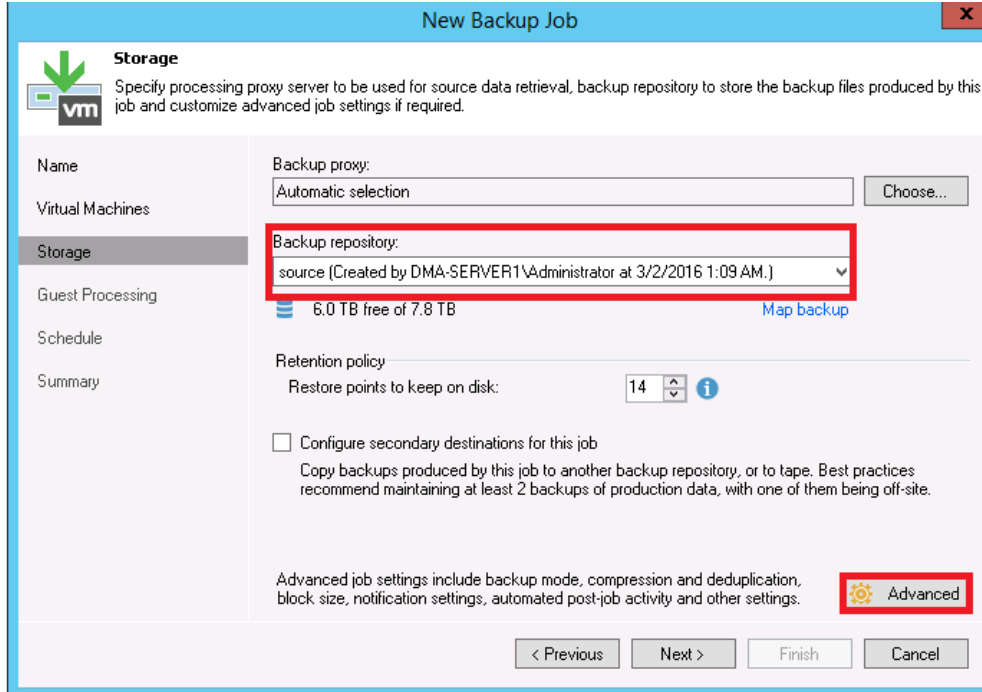
14. Enter the name for the backup job, and click **Next**.



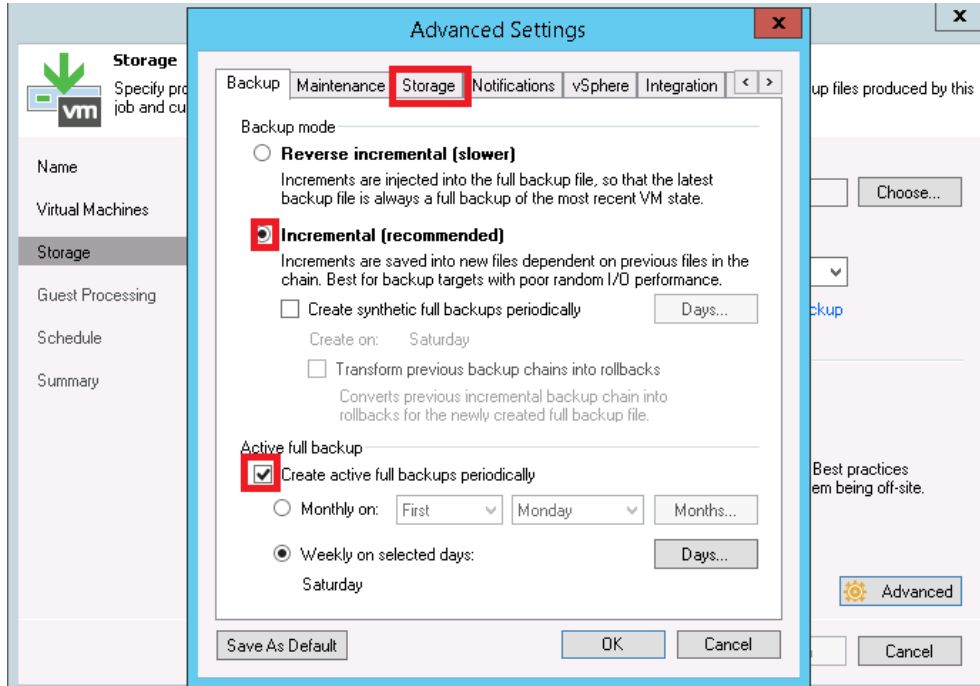
15. For the backup, select one or more virtual machines, data stores, resource pools, vApps, SCVMM clusters, and so on, as needed, and then click **Add**.



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

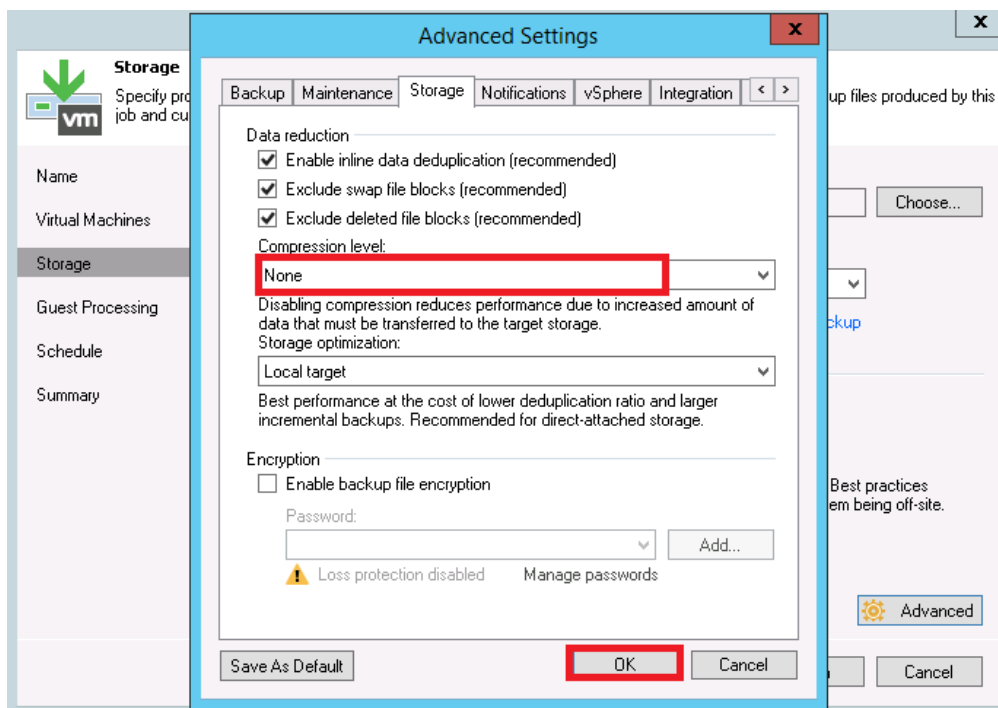


17. On the **Backup** tab, ensure **Incremental** is selected.



Note: Dell recommends that you enable **Active Full backups** once a week. The active full backup produces a full backup of a VM just as if you are running the backup job for the first time. The active full backup resets the chain of increments: all subsequent increments use the latest 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.

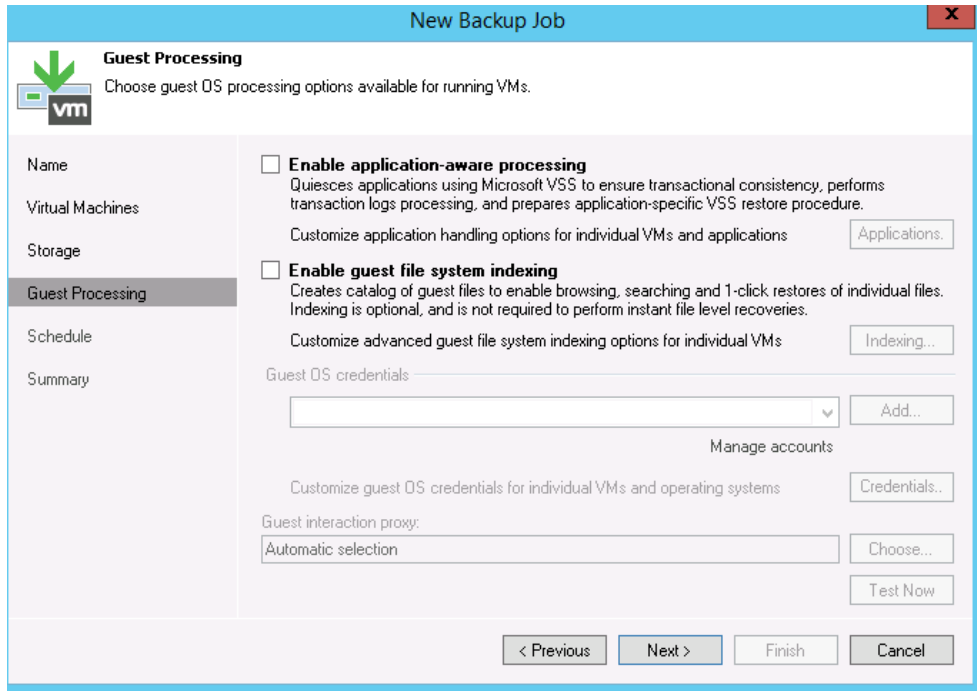
18. On the **Storage** tab, do the following:
 - a. Under **Deduplication**, select **Enable inline data deduplication**.
 - b. Under **Compression**, set the **Level** to **None**.
 - c. Under **Storage optimizations**, set **Optimization** to **LOCAL target**.



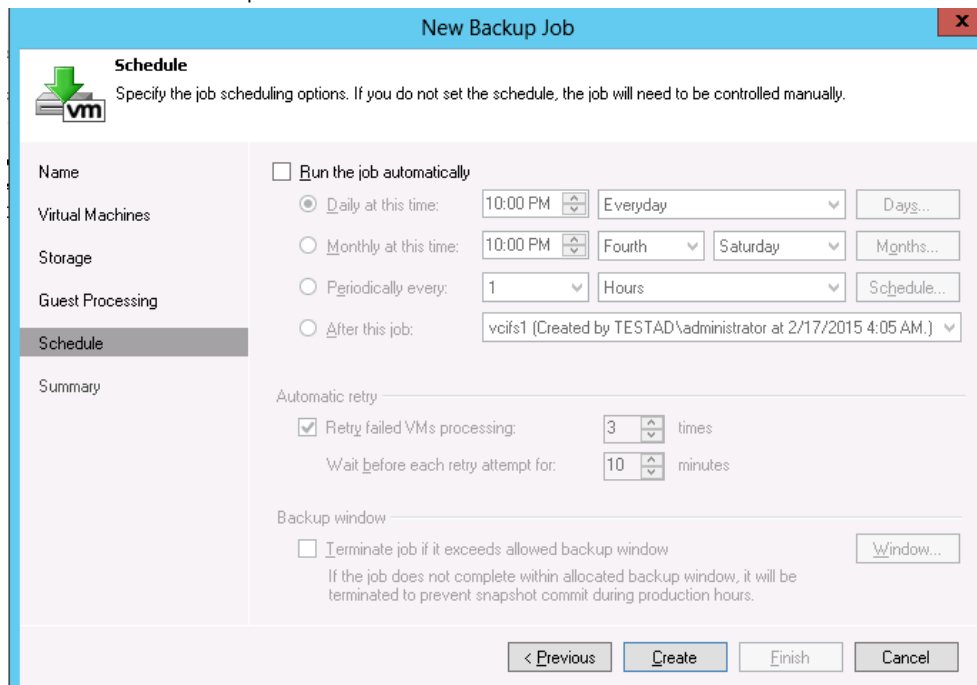
Note: For Advanced Settings, between backup performance and deduplication savings, if overall space/storage savings is the focus, Dell recommends that you select the options for all of the backup jobs.

Generally, Dell recommends turning off encryption, compression, and deduplication settings. However, with Veeam, Dell recommends that you *enable* deduplication. 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.

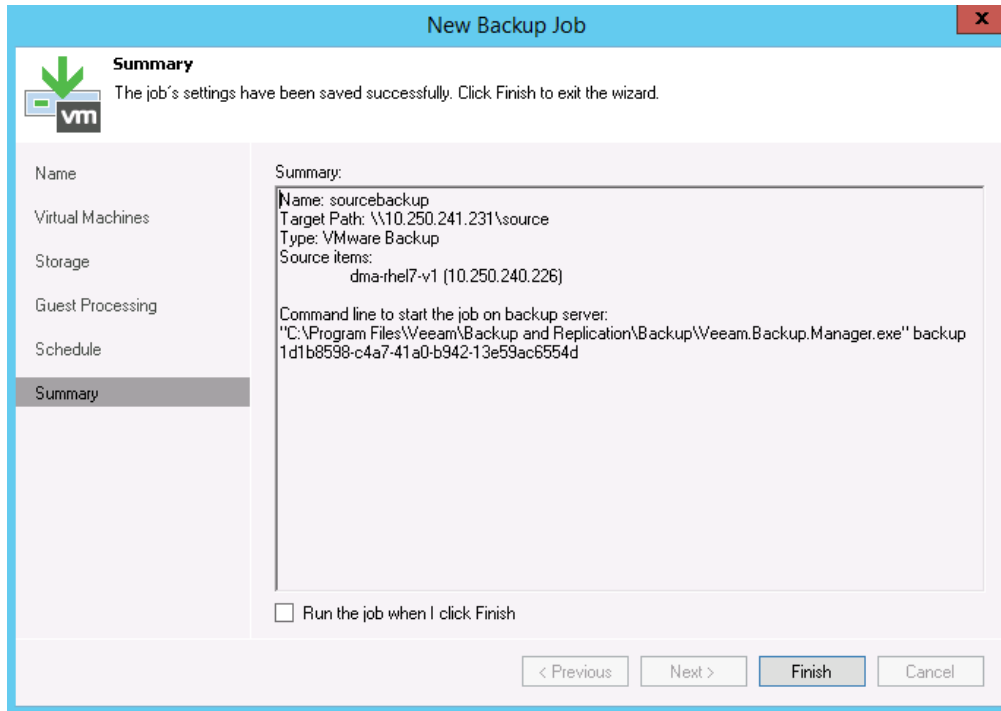
19. Click **Next**.



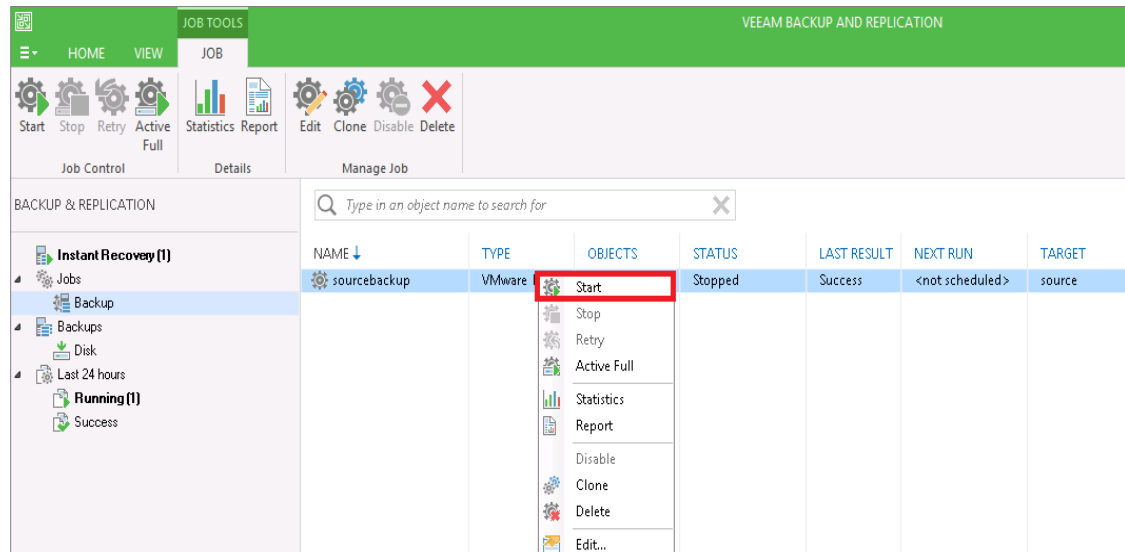
20. Schedule the backup and click **Create**.



21. Click **Finish**.



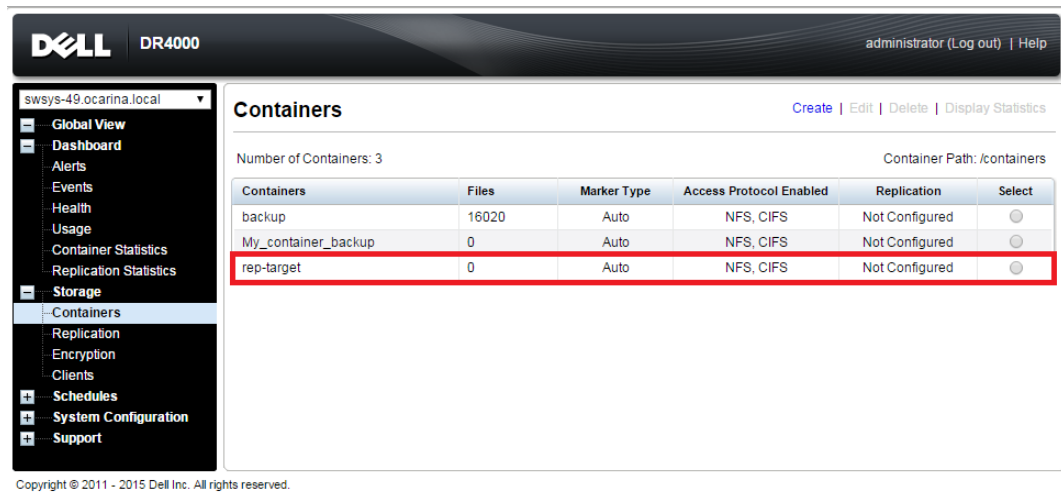
22. To run the backup manually, right-click the configured backup job, and select **Start**.



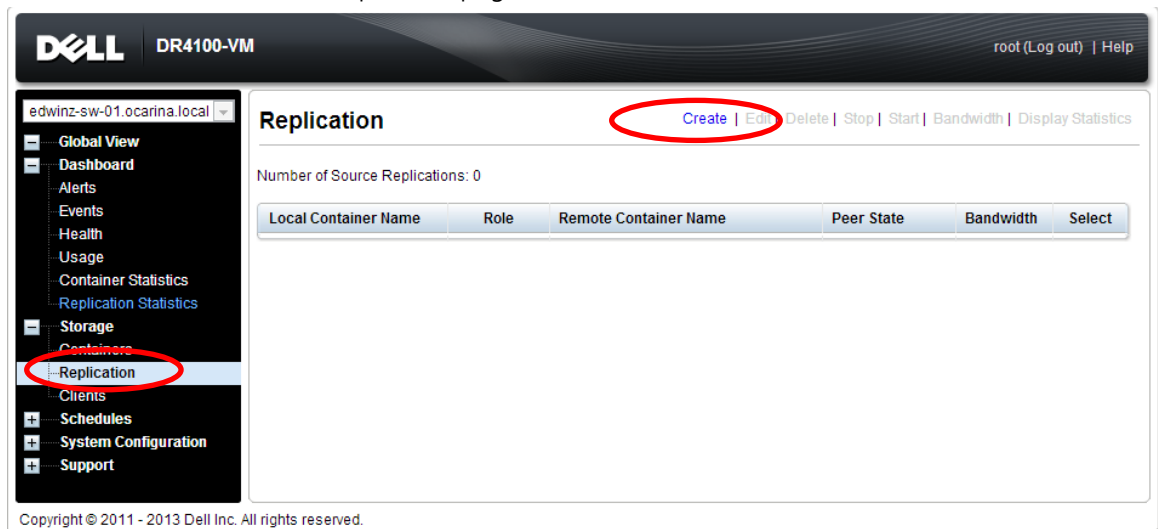
3 Setting up DR Series native replication and restore from a replication target container

3.1 Building a replication relationship between DR Series systems

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



24. On the source DR Series system, in the left navigation area, go to **Storage > Replication**, and then click the **Create** link at the top of the page.



25. Select a local container as the source container. Then, select **Container from remote system**, enter the target DR Series system related information, click **Retrieve Containers**, select a populated target container from the list, and click **Create Replication**.

Edit Replication - (swsys-06 : source ⇒ swsys-54 : target ⇒ Not Configured)

NOTE: Only the bandwidth, encryption settings and remote container's IP address/host name ("Peer System") can be changed, or adding a Cascaded Replica container.

* = required fields

Source container

Select container from local system
source

Select container from remote system

Username*: ?

Password*:

Remote System*: ?

Retrieve Remote Container(s)

- N/A -

Replica Container

Select container from local system
- Select a Container -

Select container from remote system

Username*: ?

Password*:

Remote System*: swsys-54.testad.ocarina.local ?

Retrieve Remote Container(s)

target

Cascaded Replica Container (Optional)

Select container from local system
- Select a Container -

Select container from remote system

Username*: ?

Password*:

Remote System*: ?

Retrieve Remote Container(s)

- N/A -

Source Container ⇒ Replica Container

Encryption: None 128 bit 256 bit

Bandwidth Speed Rate:

Default (not limited)

Kbps Mbps Gbps

Replica ⇒ Cascaded Replica Container

Encryption: None 128 bit 256 bit

Bandwidth Speed Rate:

Default (not limited)

Kbps Mbps Gbps

Cancel Save Replication

26. Verify that the replication is created successfully, and make sure the **Status checkbox** is marked for the replication session.

DELL DR4000 root (Log out) | Help

- Global View
- Dashboard
- Alerts
- Events
- Health
- Usage
- Container Statistics
- Replication Statistics
- Storage
- Containers
- Replication
- Encryption
- Clients
- Schedules
- System Configuration
- Support

Replication

Create | Edit | Delete | Stop | Start | Display Statistics

Message

- Successfully added replication for container 'source' - 'target'.
- NOTE: Replication connection(s) are being established. Information updates may be briefly delayed until the connection is completed.

Number of Replications: 1

Source Container	Status	Replica Container	Status	Cascaded Replica Container	Select
e-shelf1 source	<input checked="" type="checkbox"/>	dr4300-jd6ff42 target	<input checked="" type="checkbox"/>	Not Configured	<input type="checkbox"/>

* Local container(s) in bold.

Copyright © 2011 - 2015 Dell Inc. All rights reserved.

NOTE:

- Make sure the replication session has a **Peer Status** as **Online** if restore from replication target is needed,

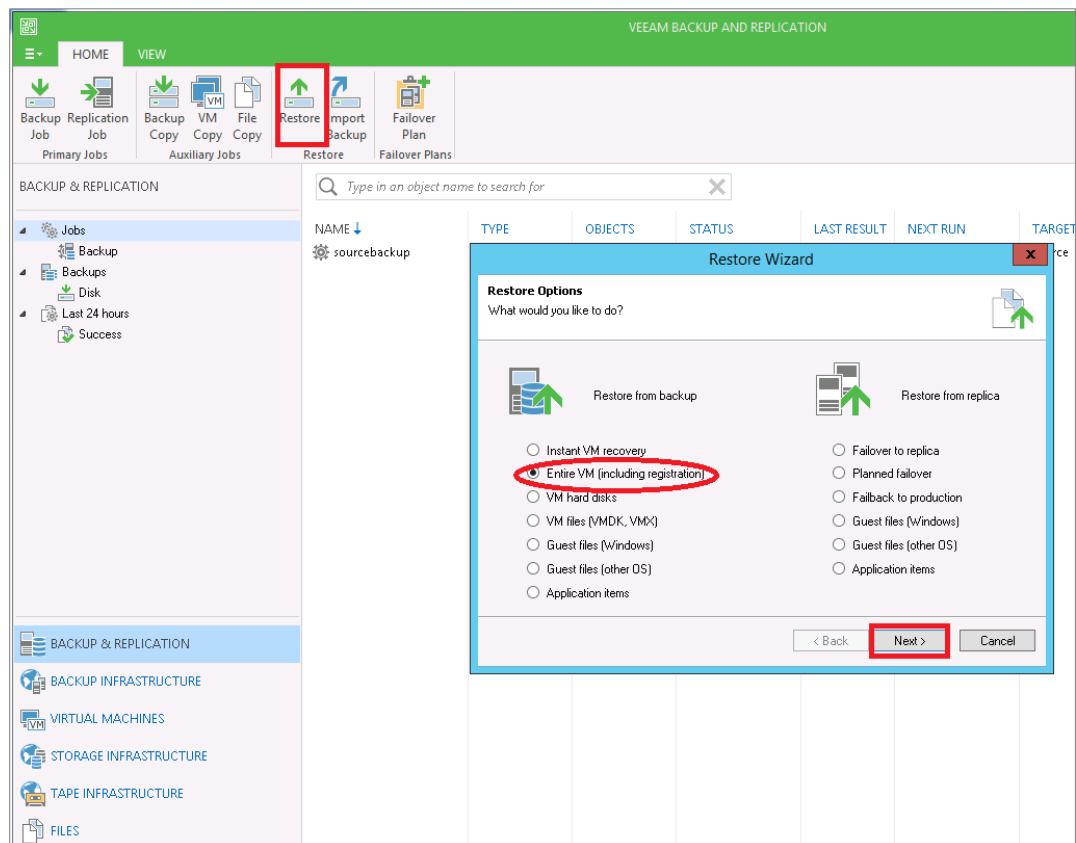


- Make sure the replication is in an **INSYNC** state from Replication Statistics menu, and Stop or Delete the replication.
- Make sure the replication target has **CIFS/NFS** connection(s) enabled when restoring from it.

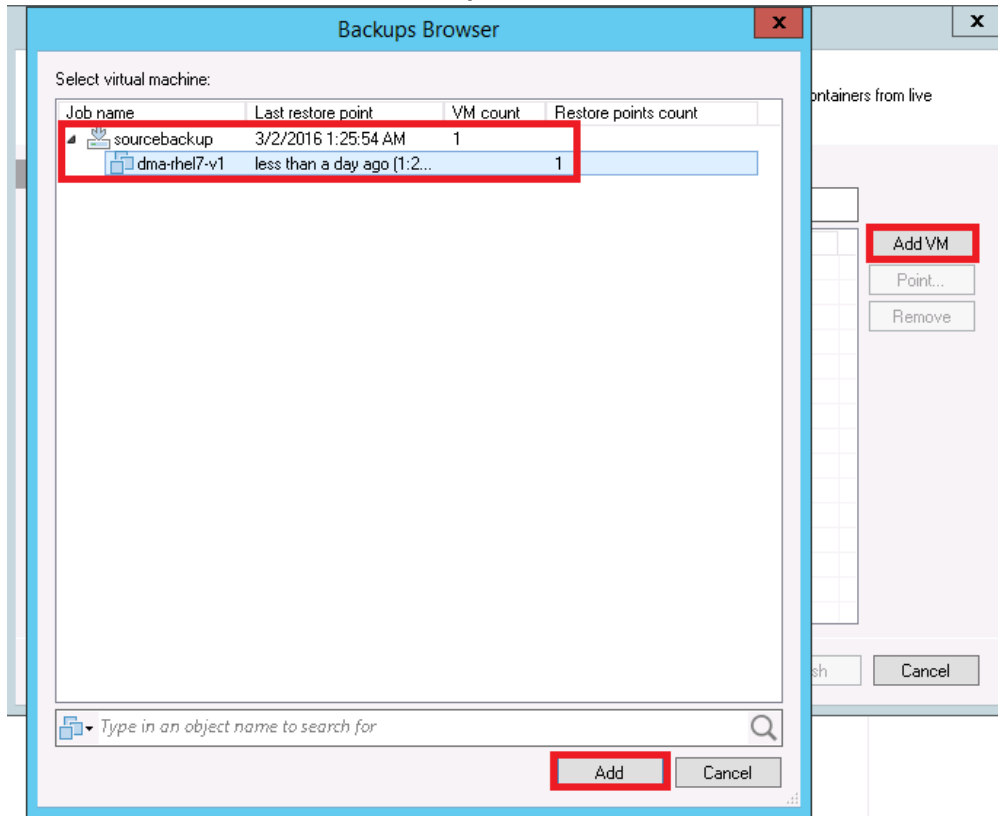
3.2 Restoring data from a target DR Series system

Note: Before restoring from a 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.

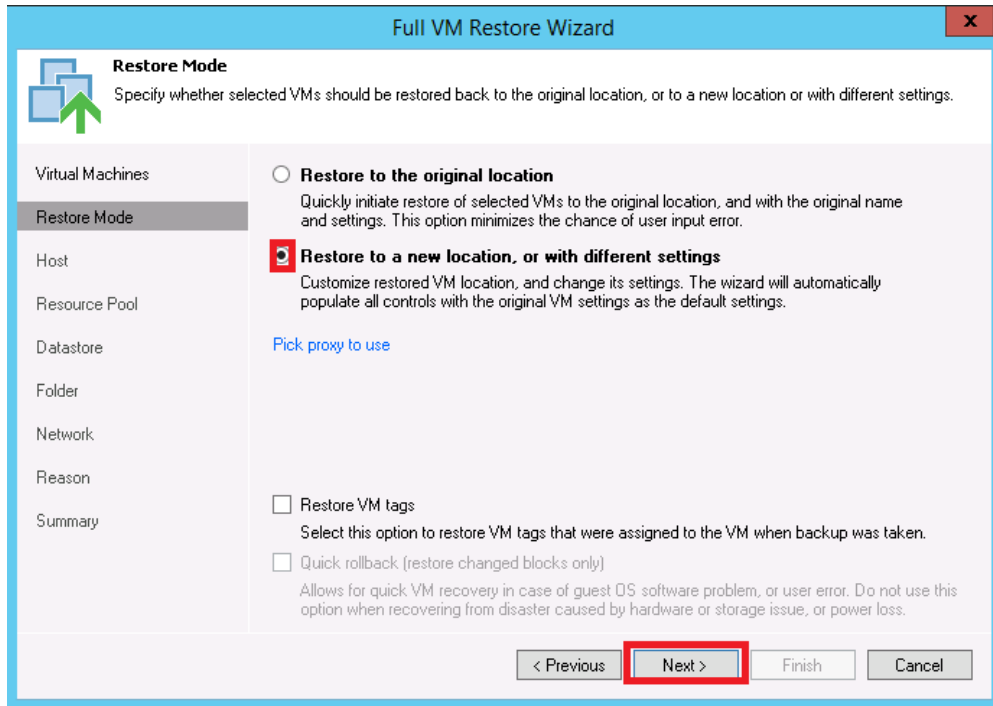
1. Add the target DR Series system container to the Veeam repository. For instructions, see the section, "Setting up Veeam."
2. Update all backup jobs that use the source DR Series system container as a repository, and change them to use the target DR Series system container as the backup repository.
3. Under **Backup & Replication**, click **Restore** to create a restore job. Select the appropriate **Restore from backup** option.



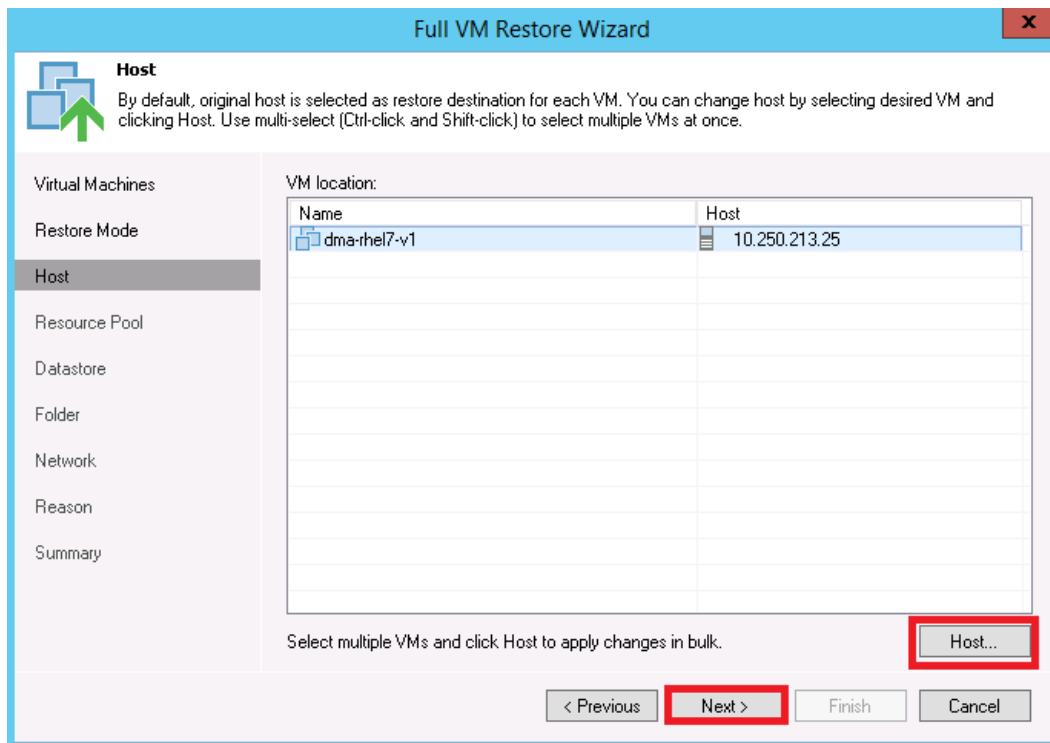
4. Click **Add VM** and select **"From backup"**. Select the VM to be restored and click **Add**.



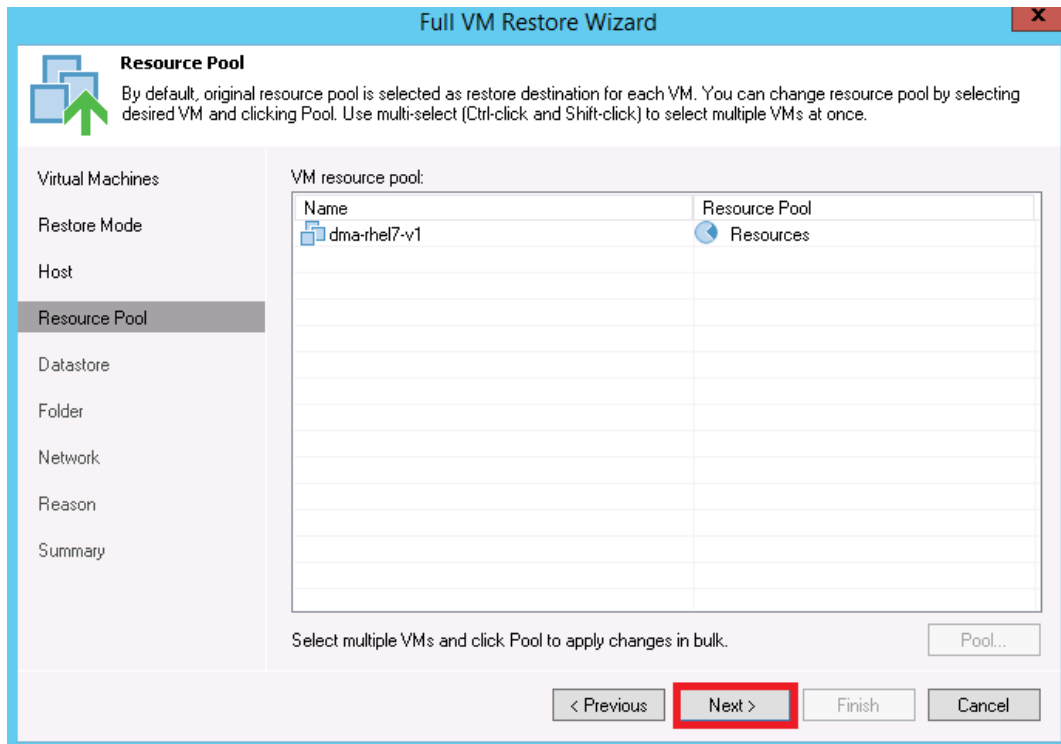
5. Select the Restore Mode and click **Next**.



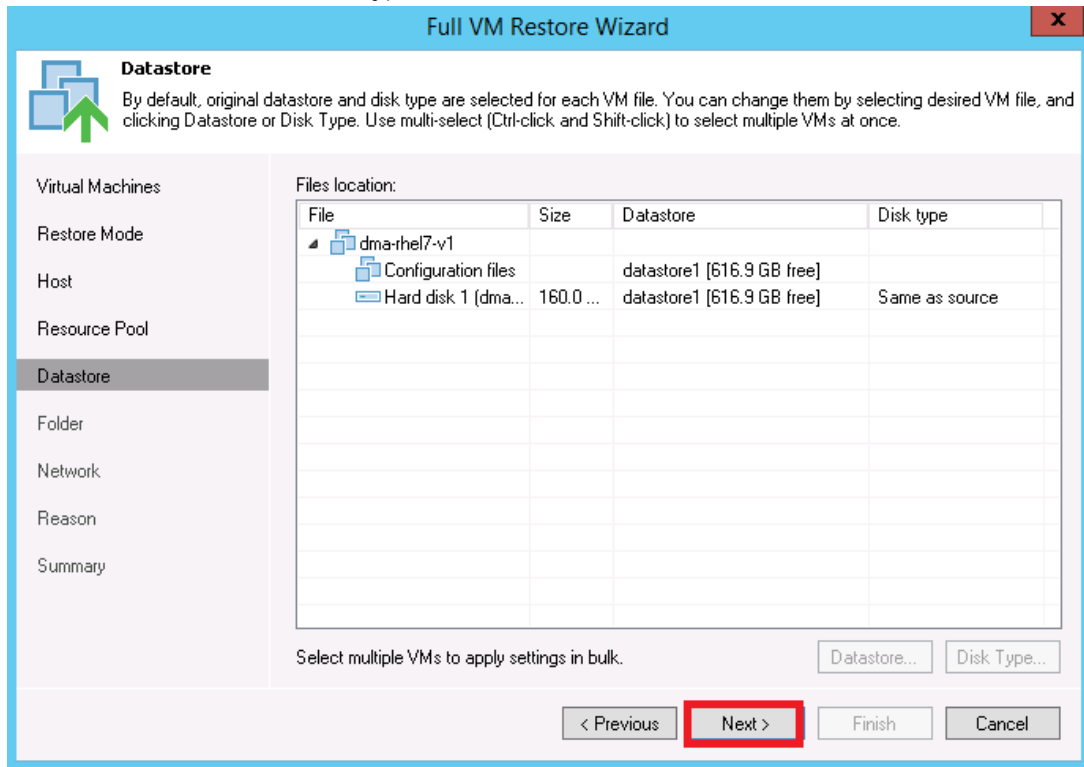
6. Provide the Host details as needed, and click **Next**.



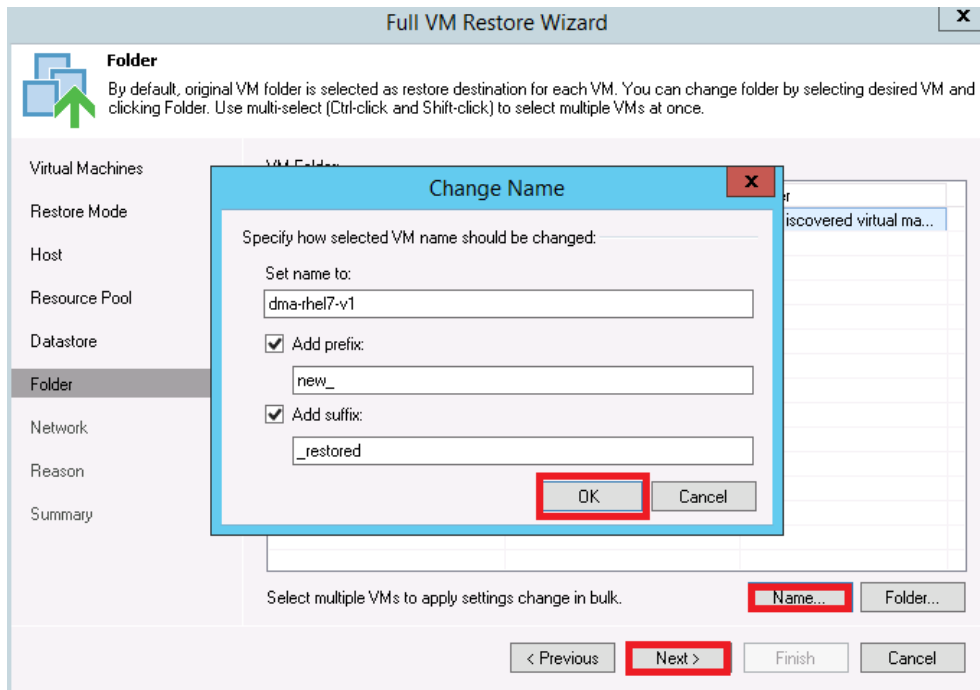
7. Select the resource Pool and click **Next**.



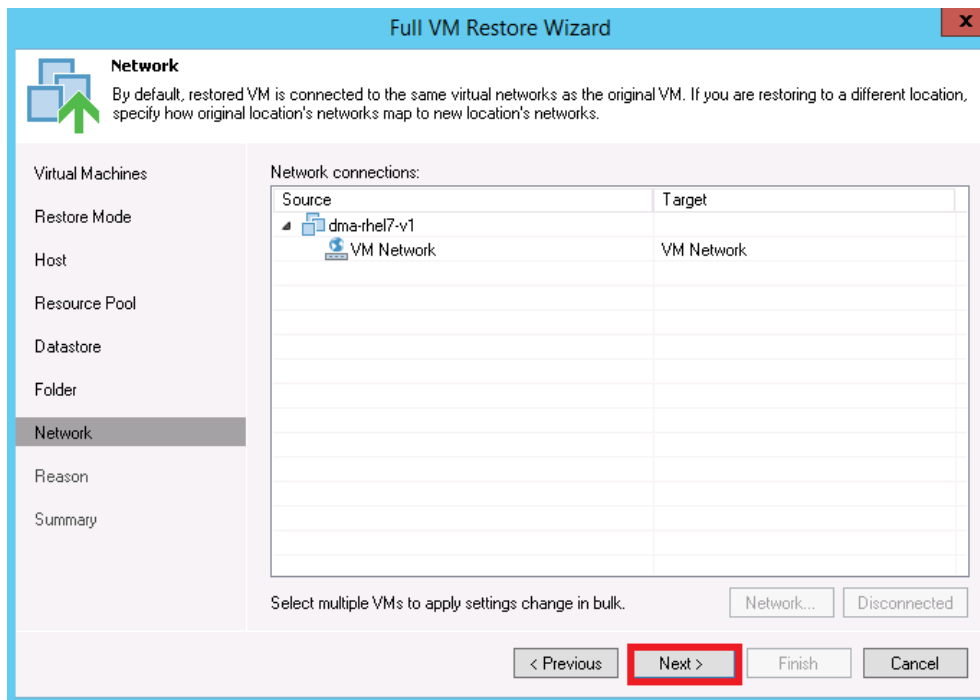
8. Select the datastore and disk type, and click **Next**.




9. Enter the new name for the restored VM and click **Next**.



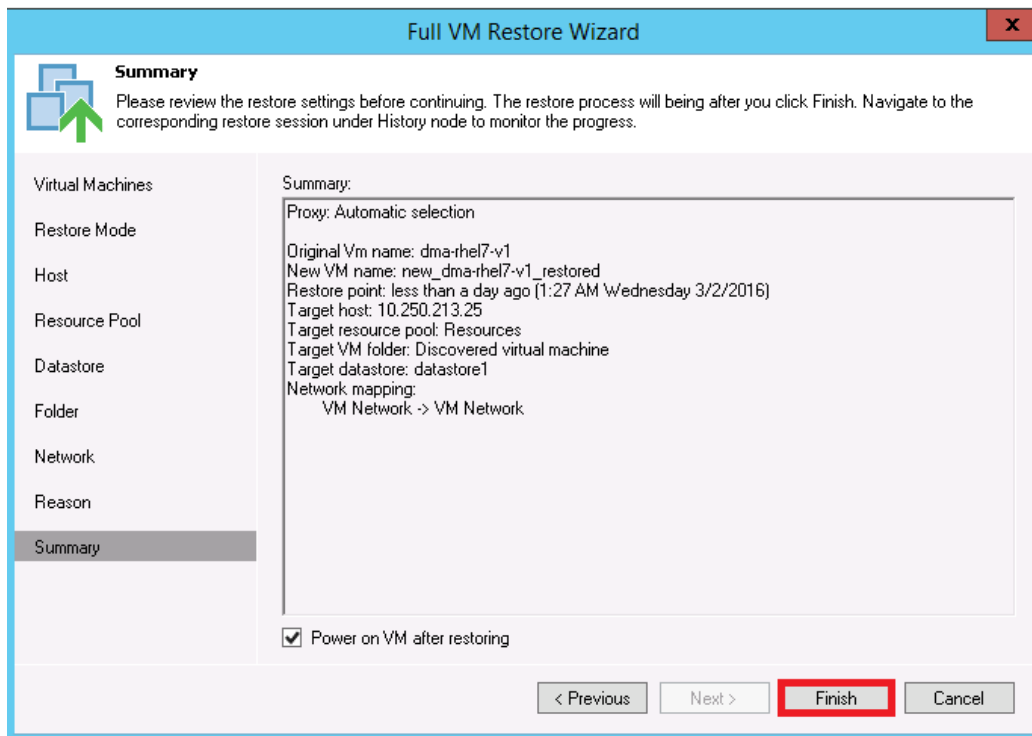
10. Select the network location and click **Next**.



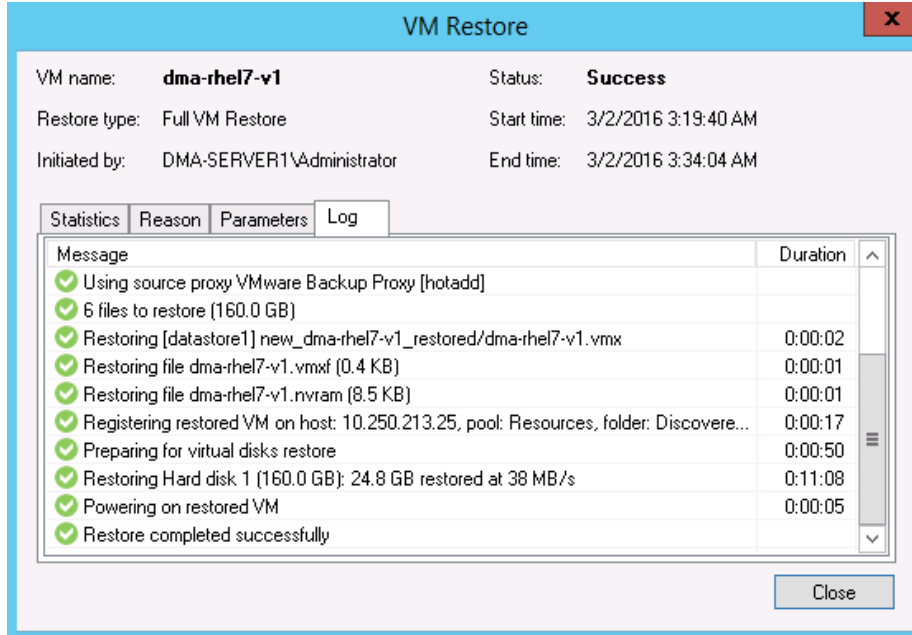
11. Enter text that describes the reason for the restore and click **Next**.



12. Click **Finish**.



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



4 Using Veeam 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.

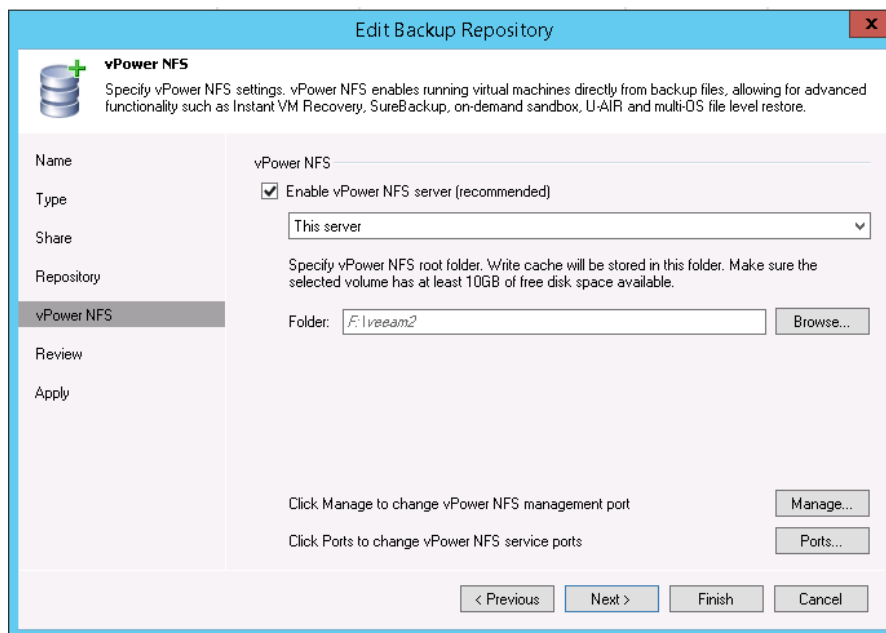
Instant VM Recovery uses patented vPower® technology to mount a VM image to a production VMware vSphere or Microsoft Hyper-V host directly from a compressed and deduplicated backup file.

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 they are merged with the original VM data when VM recovery is finalized, that is, when VM is migrated back to production storage.

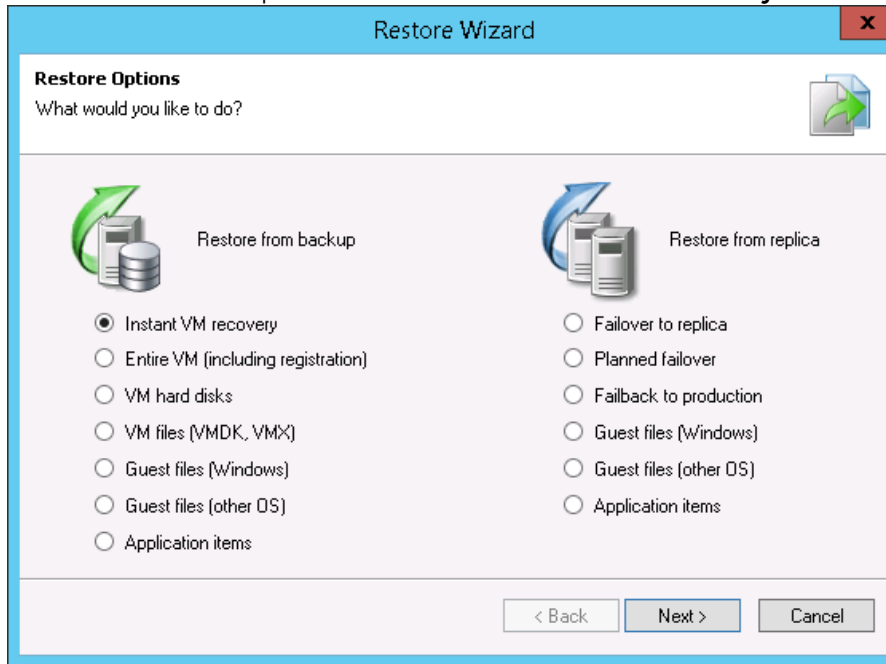
The Veeam vPower NFS service is a Windows service that runs on a windows backup repository server and enables it to act as an NFS server.

4.1 Instant Recovery for ESX VM backups

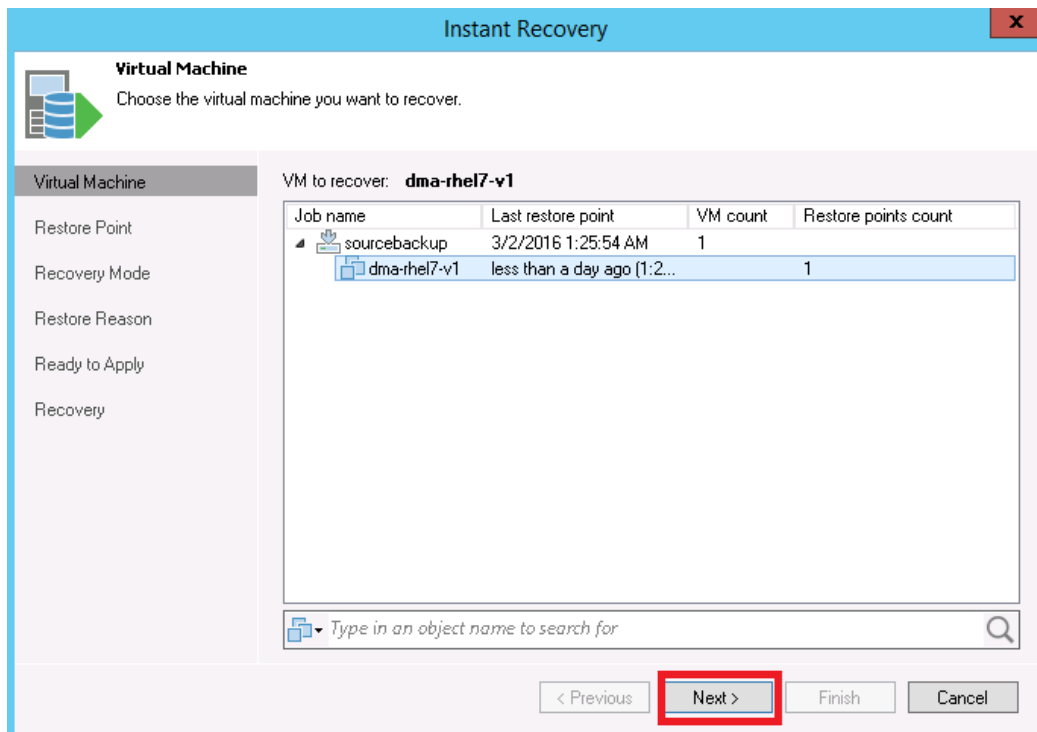
1. To enable instant recovery for ESX VM backups, do the following:
 - a. Create a backup job for the required VM as described previously in Section 3 with the only difference being to set the **vPower NFS Datastore** in the vPower NFS tab.
 - b. On the vPower NFS tab, select the checkbox, **Enable vPower NFS Server**, and select the appropriate folder as the NFS Datastore. You can configure the NFS Datastore on a different Windows server if required. This is done by selecting the drop down and adding the host and associated credentials. Click **Next**.



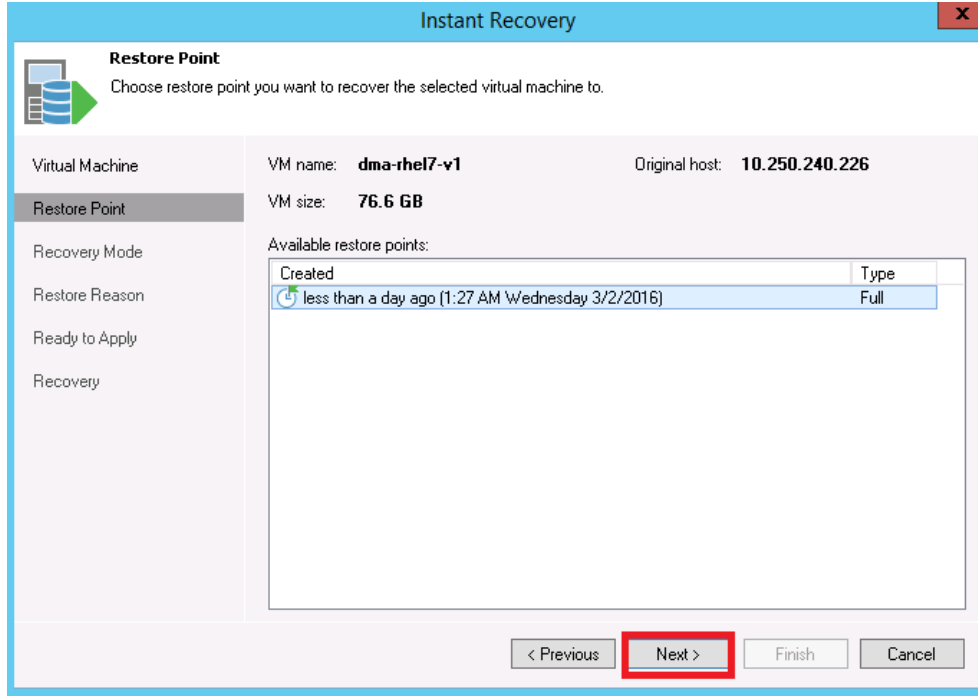
- To perform the instant recovery, click the **Restore Wizard** option. Select the **VMware** option and then select **Instant VM recovery**.



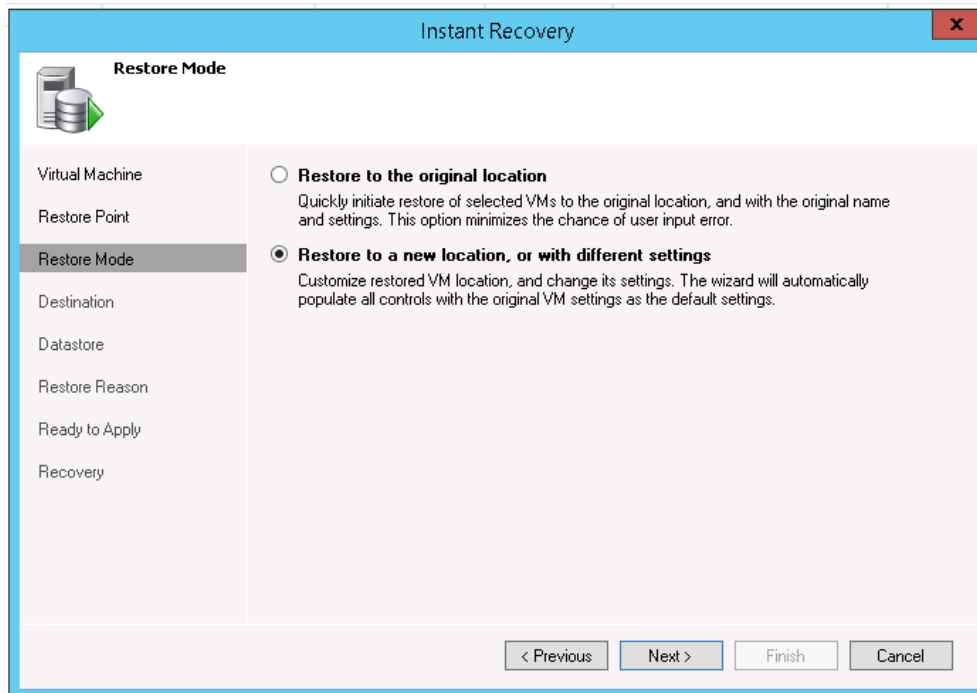
- Select the Virtual Machine to be recovered and click **Next**.



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



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



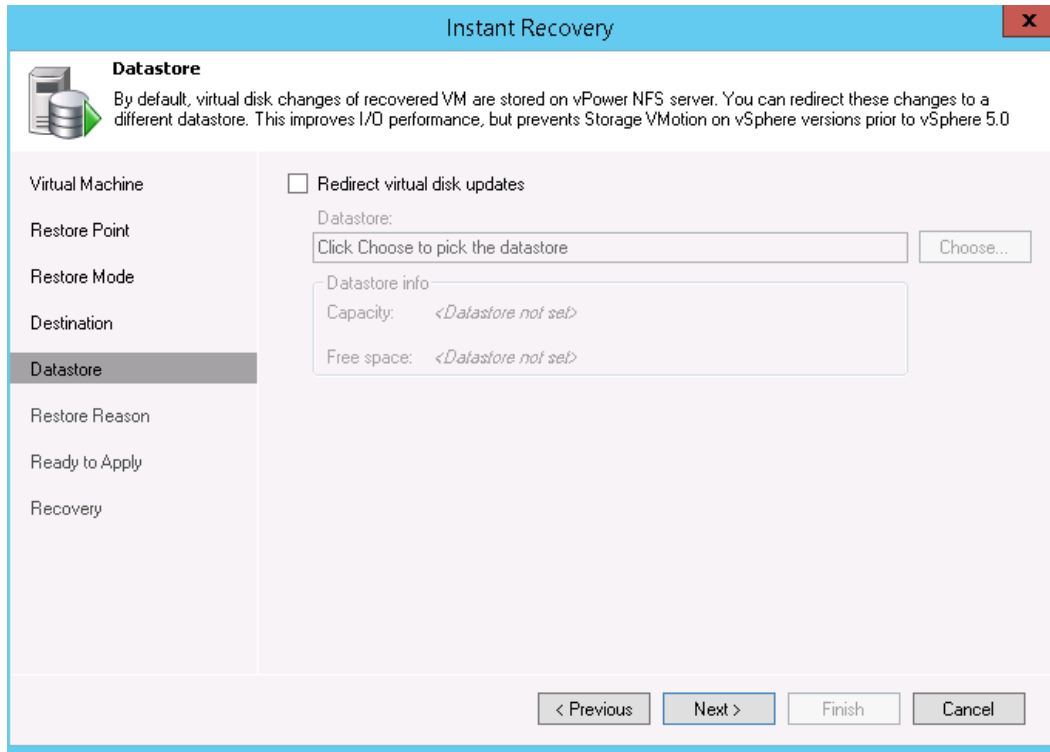
6. At the Destination step, do the following:
 - a. Select the ESX(i) host on which the VM should be restored instantly.
 - b. In the Resource pool box, select the resource pool to which the restored VM should belong.
 - c. In the Restored VM name field, add the _restored suffix to the VM name.
 - d.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

The screenshot shows the 'Instant Recovery' wizard in the 'Destination' step. The window title is 'Instant Recovery'. Below the title bar, there is a 'Destination' section with a sub-header and a description: '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)'. On the left, there is a navigation pane with options: Virtual Machine, Restore Point, Recovery Mode, Destination (selected), Datastore, Restore Reason, Ready to Apply, and Recovery. The main area contains the following fields:

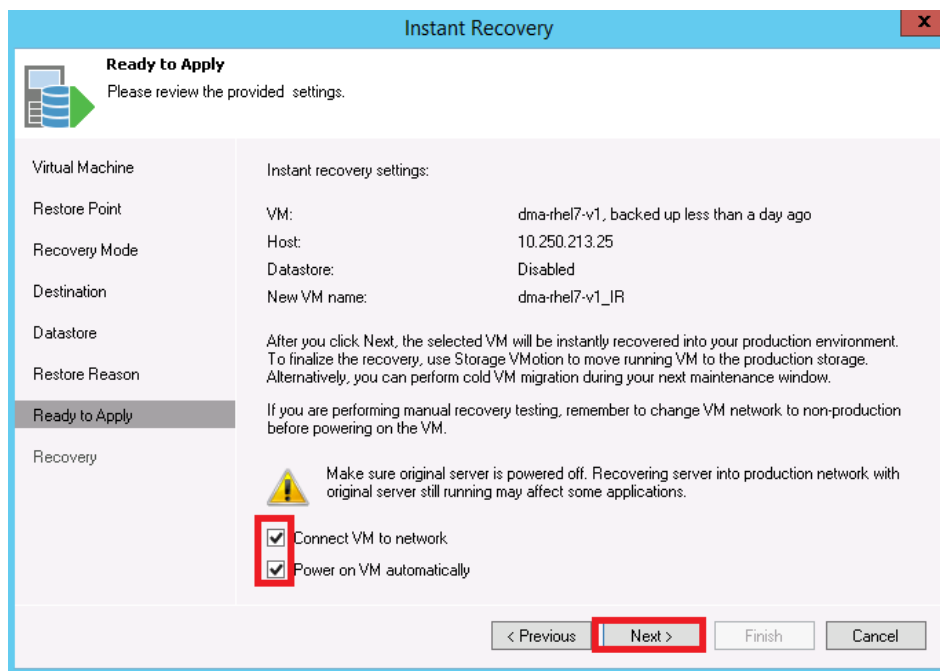
- Host: 10.250.213.25 (with a 'Choose...' button)
- VM folder: Discovered virtual machine (with a 'Choose...' button)
- Restored VM name: dma-rhel7-v1_IR (highlighted with a red box)
- Resource pool: Resources (dropdown menu is open, showing 'Resources' and 'VirtualLab', both highlighted with a red box)

At the bottom, there are four buttons: '< Previous', 'Next >' (highlighted with a red box), 'Finish', and 'Cancel'.

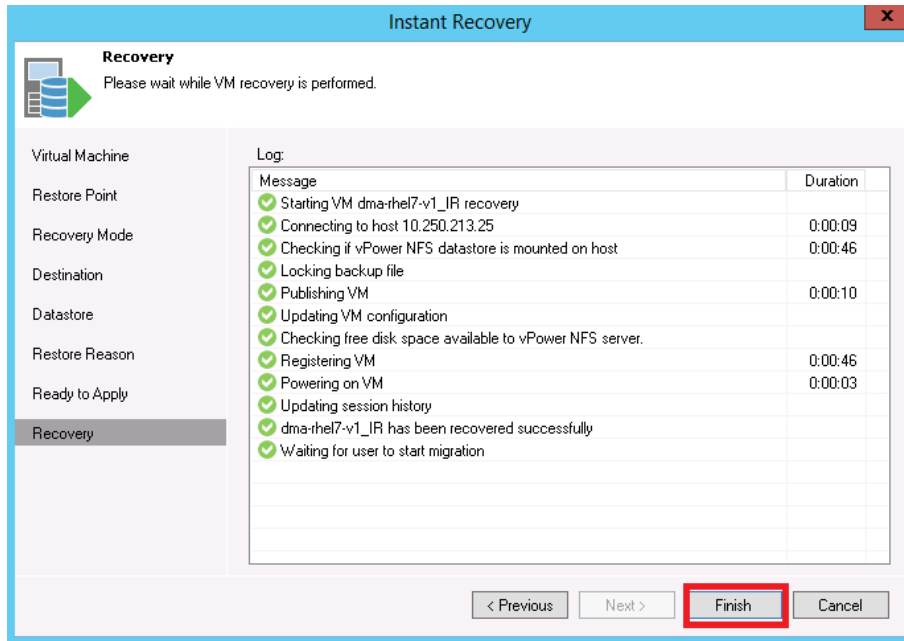
7. On the Datastore tab, ensure the **Redirect virtual disk updates** check box is not selected, and click **Next**. This will let you use Storage vMotion to migrate the VM to production after the VM is recovered from the backup.



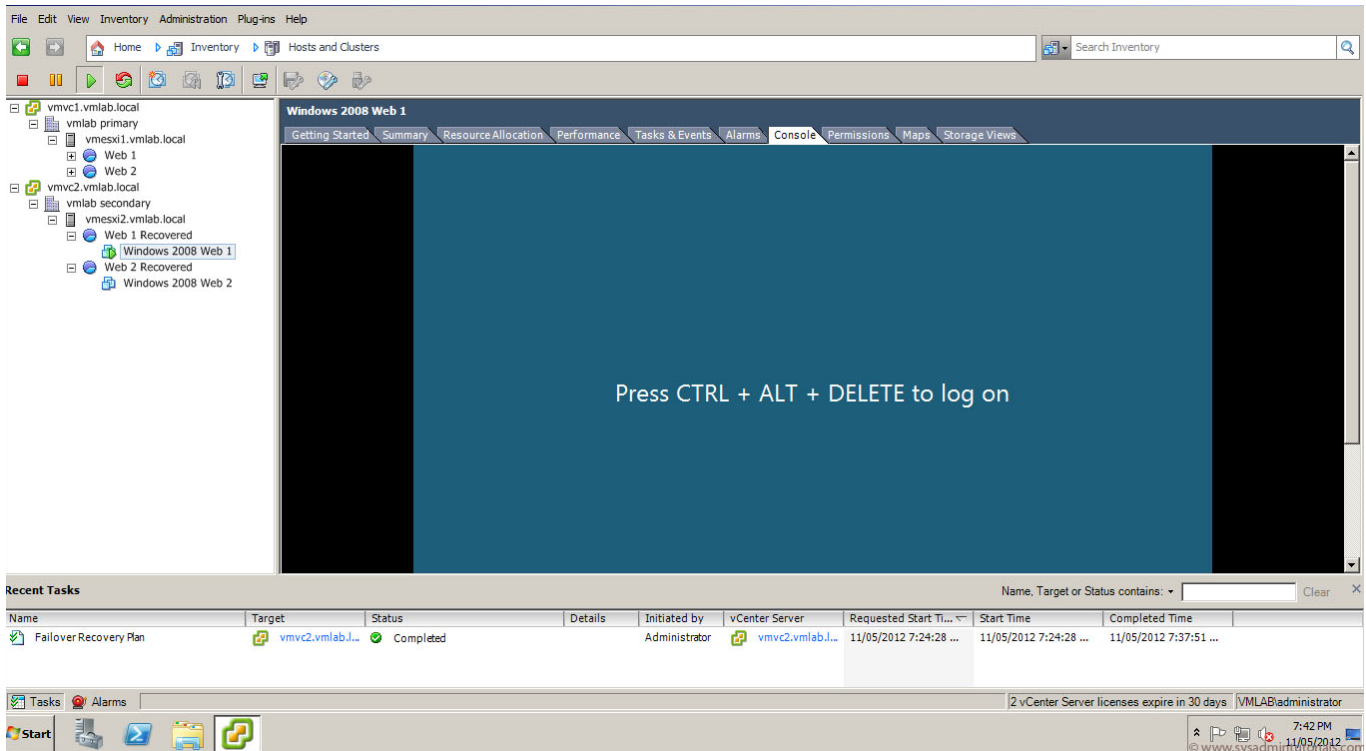
8. Select the checkboxes **Connect VM to network** and **Power on VM automatically**, and then click **Next**.



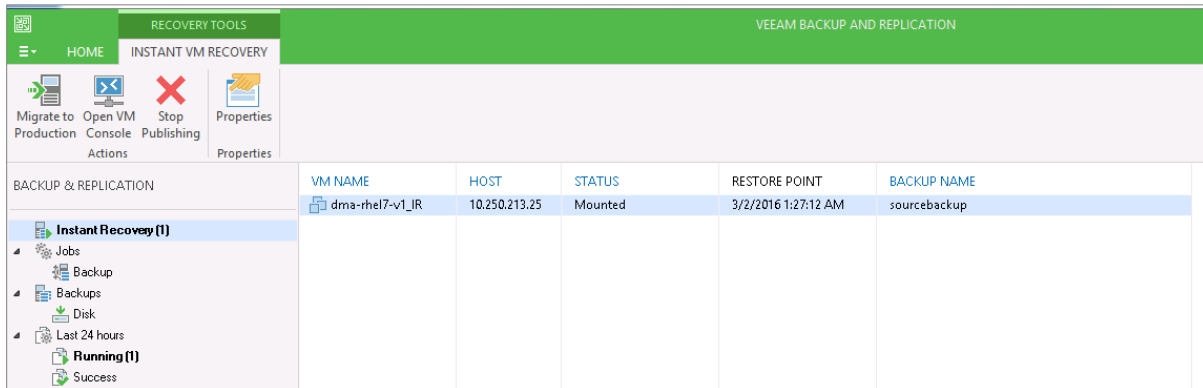
9. Click **Finish** to start the Instant VM Recovery



10. Open the vSphere client and make sure 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.

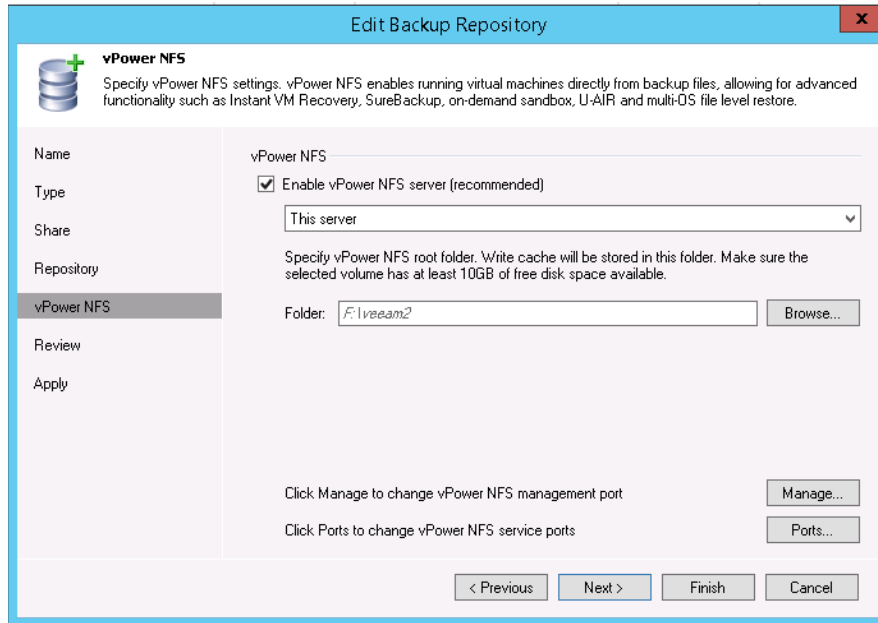


4.2 Instant Recovery with Hyper-V Server

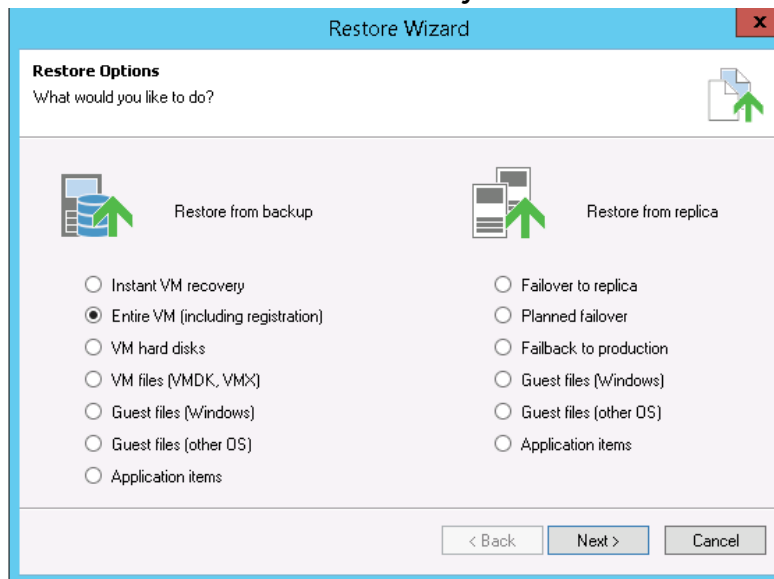
1. To enable Instant Recovery for HyperV VM backups, do the following:
 - a. Create a backup job for the required VM as described previously in Section 3 with the only difference being to set the vPower NFS Datastore option in the “vPower NFS” tab.
 - b. On the vPower NFS tab, select the checkbox, **Enable vPower NFS Server**.

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

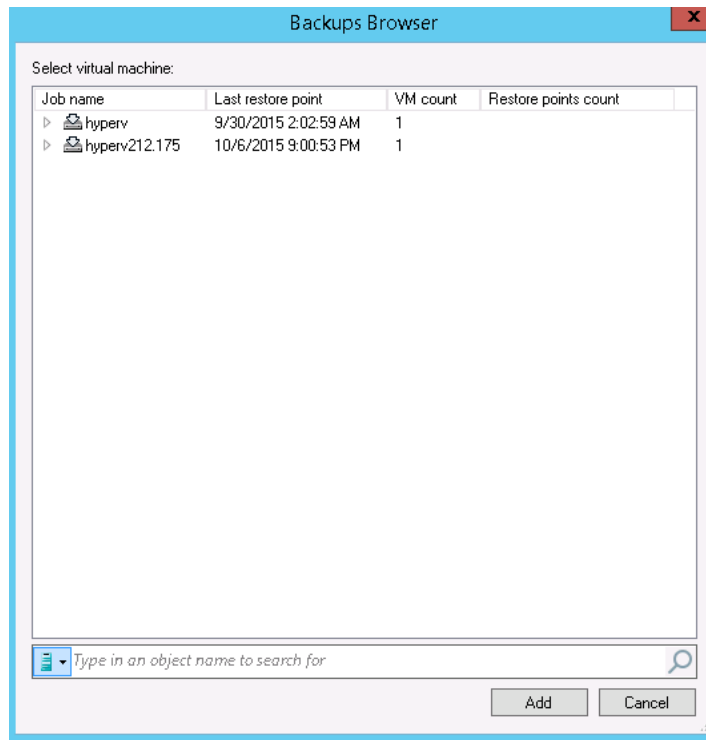




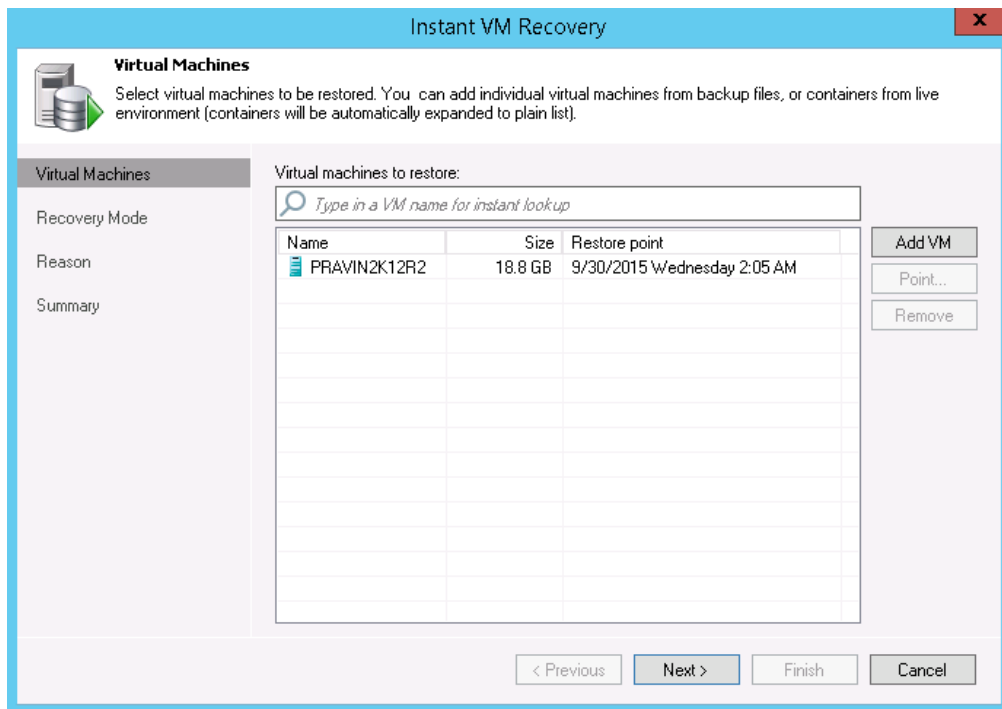
2. To perform Instant Recovery, on the Veeam console, click the **Restore Wizard** option, select **Hyper-V** and then select **Instant VM recovery**.



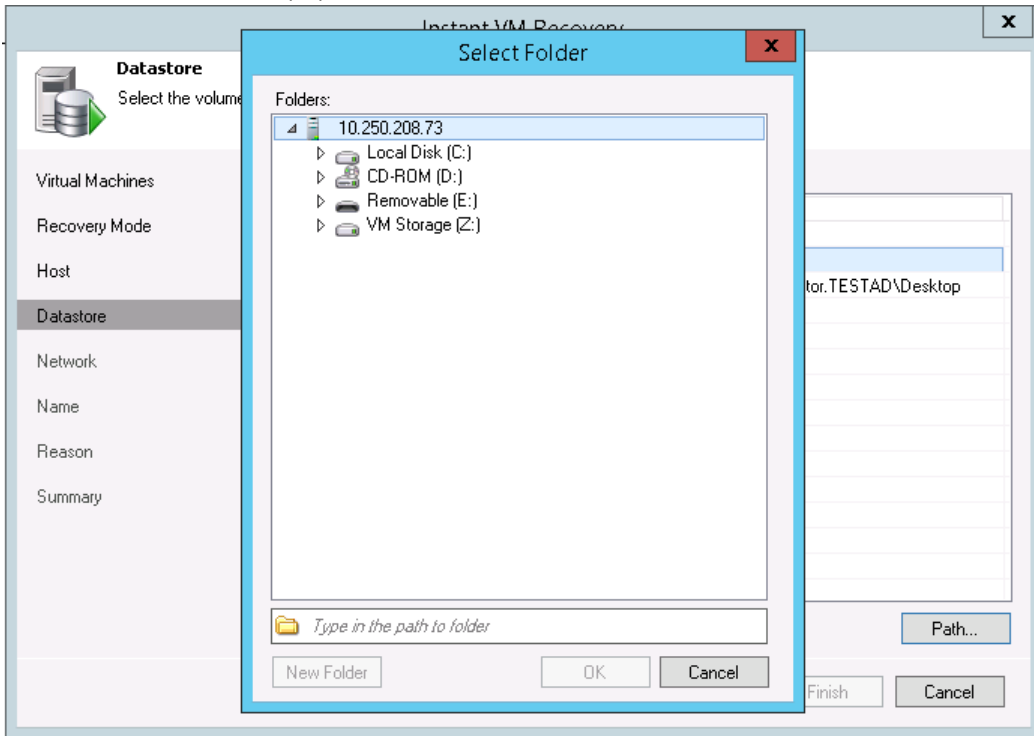
3. Select the Virtual Machine to be recovered.



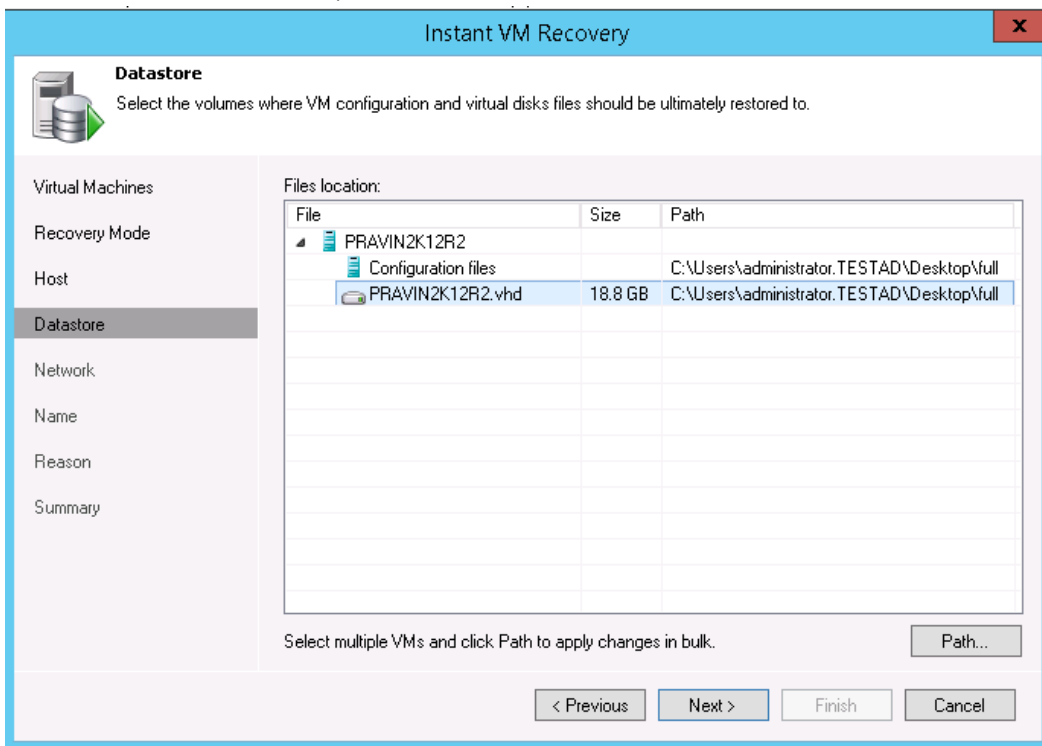
4. Add the VM which need to be recovered.



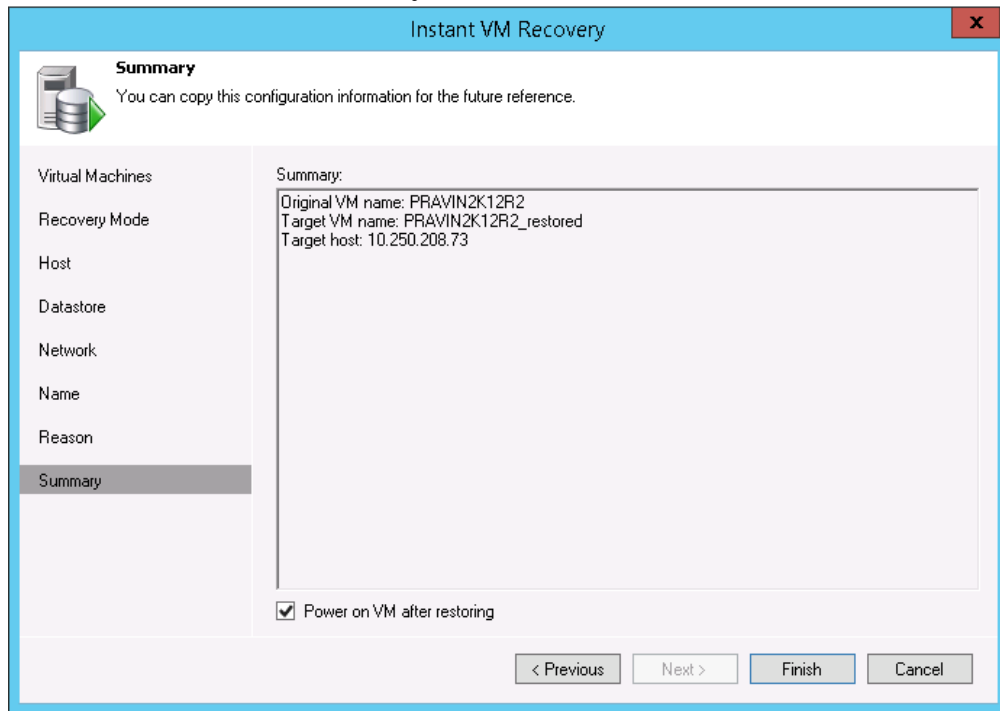
7. At the Datastore step, provide the details of cache data that need to be stored.



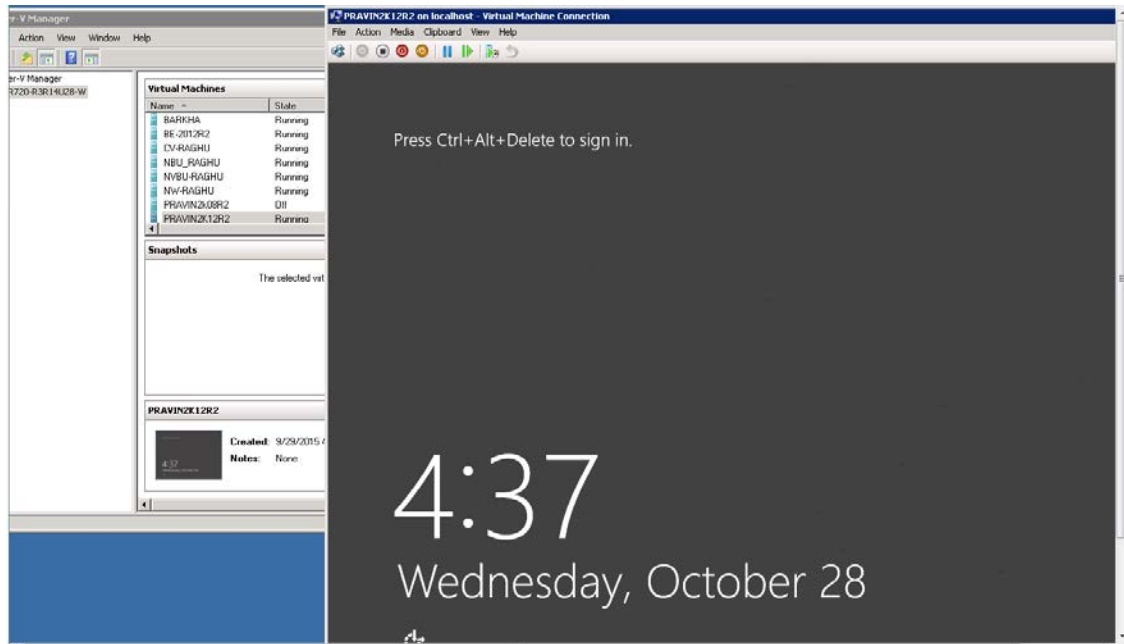
8. Enter the details of the path where VM cache data is stored.



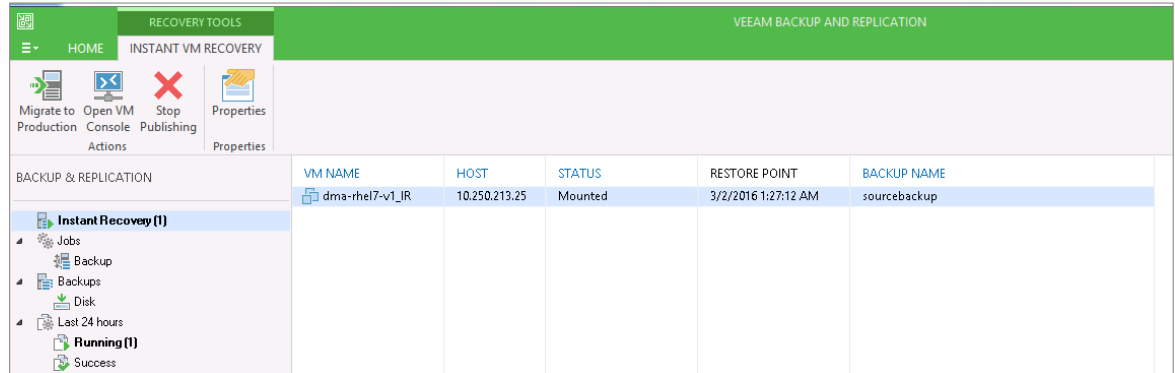
11. Click **Finish** to start the recovery.



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



13. 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.



4.3 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 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.

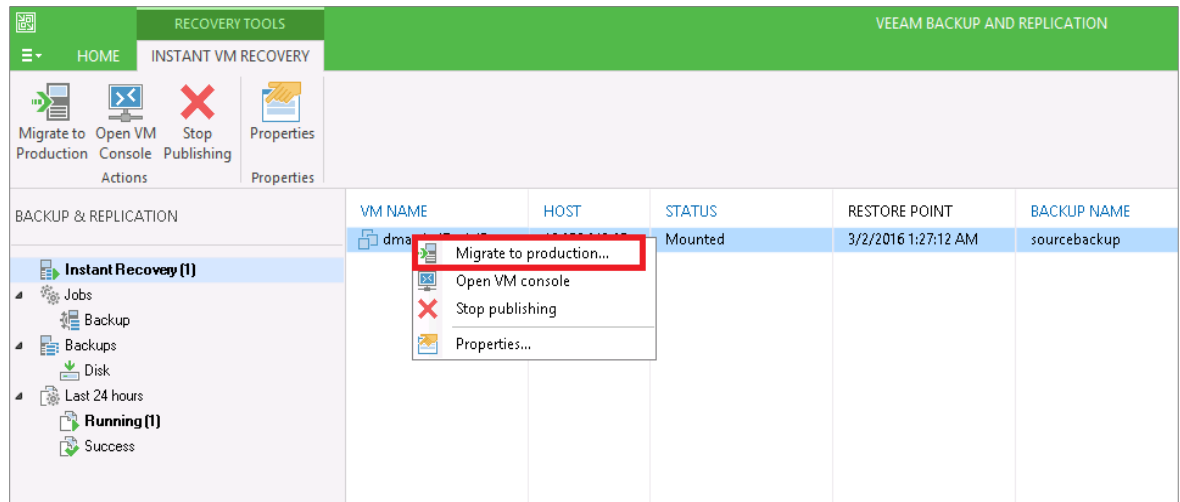
4.3.1 Migrating a VM to production

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 a 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 a 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**.





4.3.2 Terminating an Instant VM Recovery session

When you terminate an 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**.



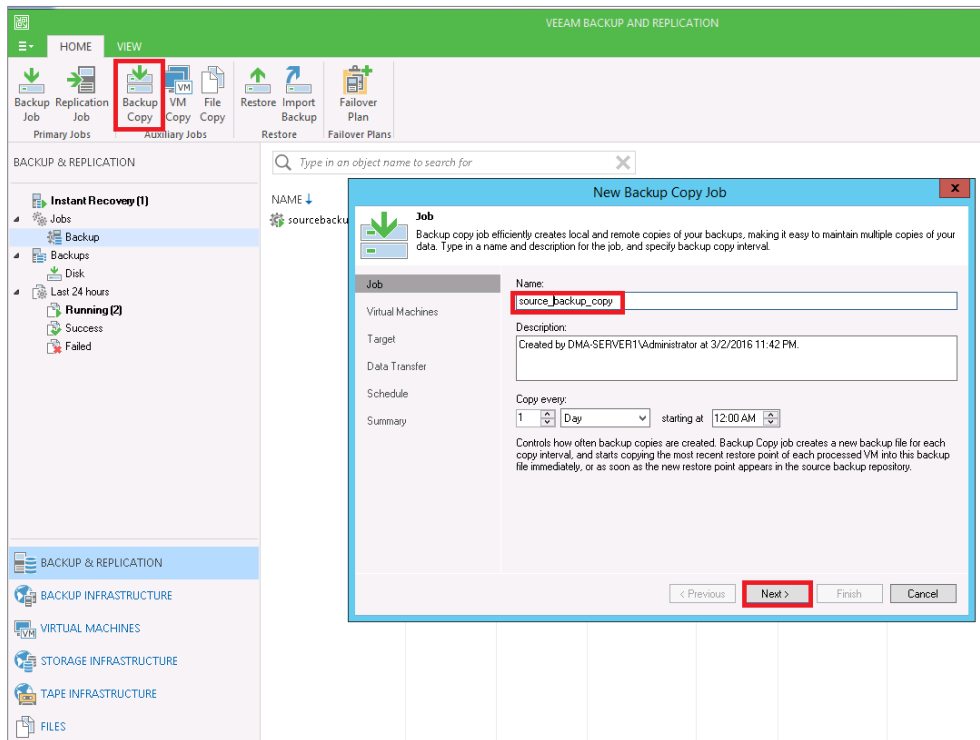
5 Creating a backup copy

The main purpose of a backup 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.

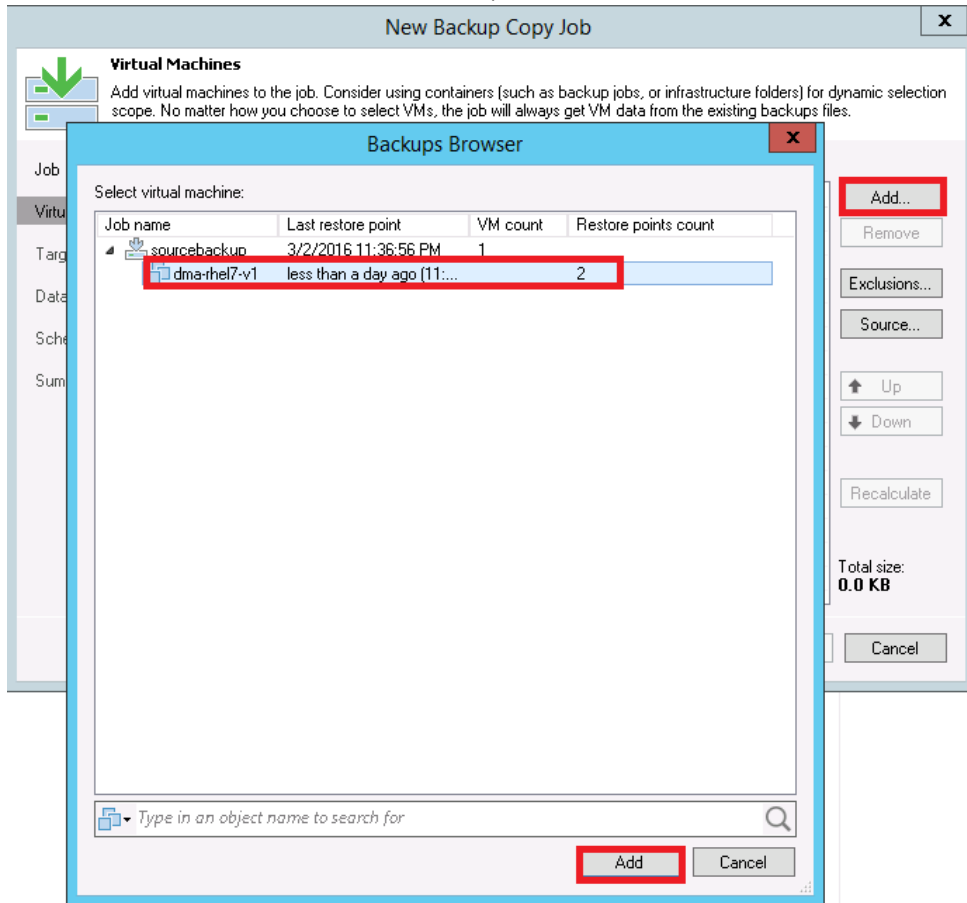
The backup copy job is a separate task that needs to be set apart from the backup job. Veeam Backup Copy allows users to copy backup data to secondary storage.

Follow these steps to create a backup copy job.

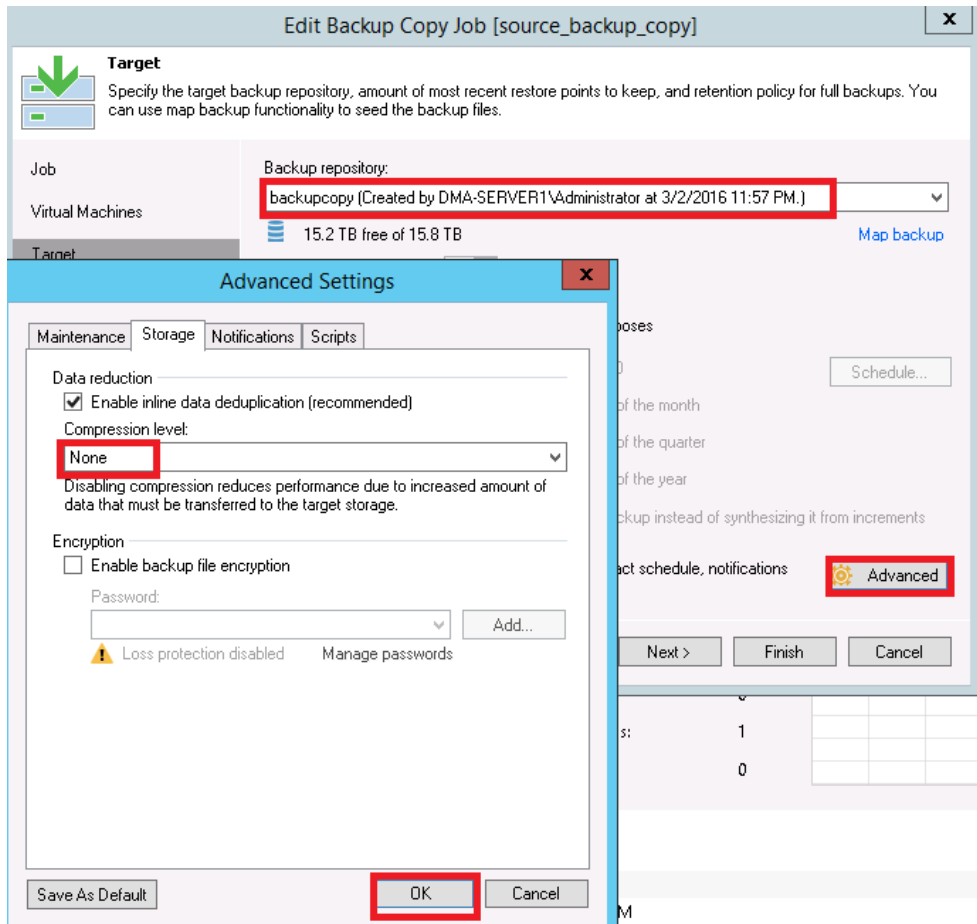
1. Click the backup copy in Auxiliary jobs, enter a name of the job, and click **Next**.



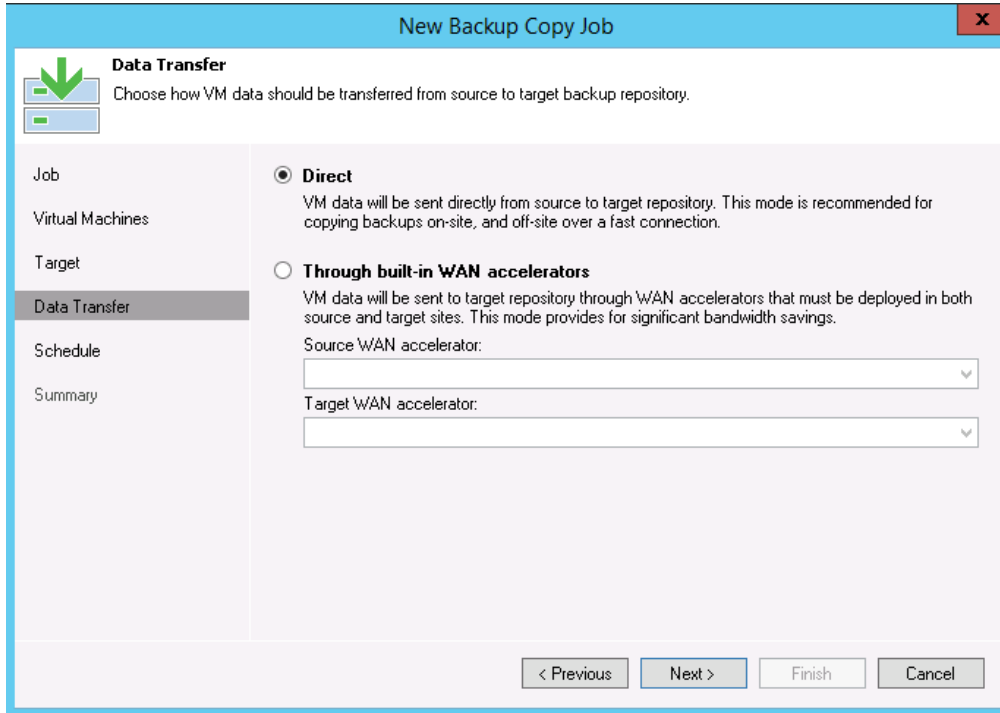
2. Add the Virtual Machine from the backup Jobs.



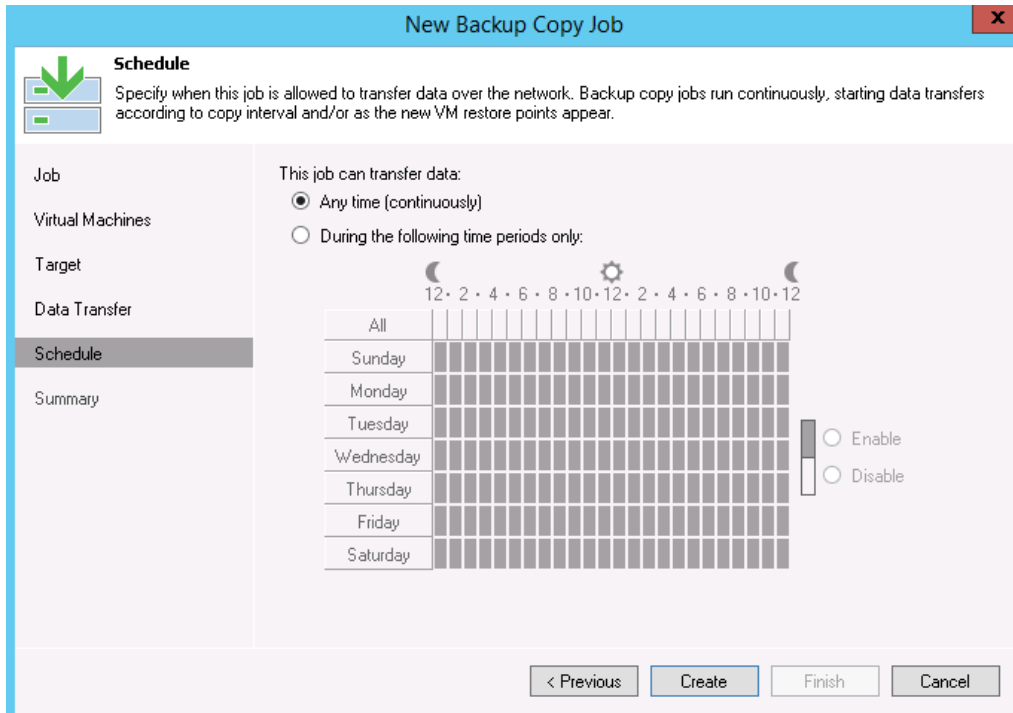
3. Select the backup repository, click **Advanced**, and then do the following:
 - a. Select the **Storage** tab.
 - b. Select the Compression level as **None**.
 - c. Click **OK**.



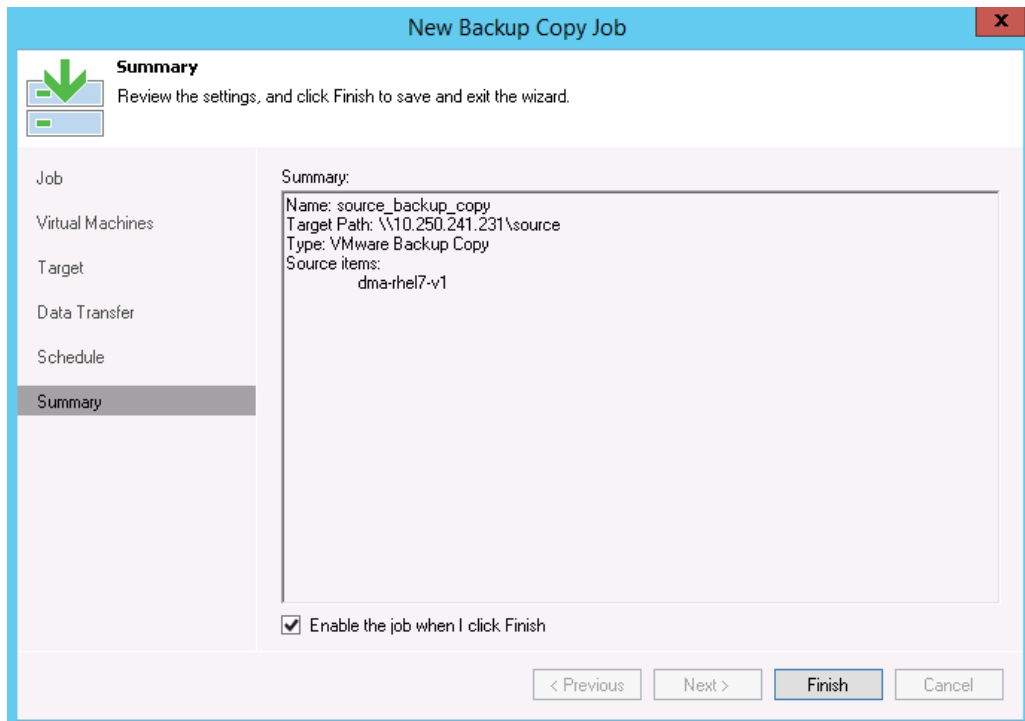
- Select the type of data transfer and click **Next**.



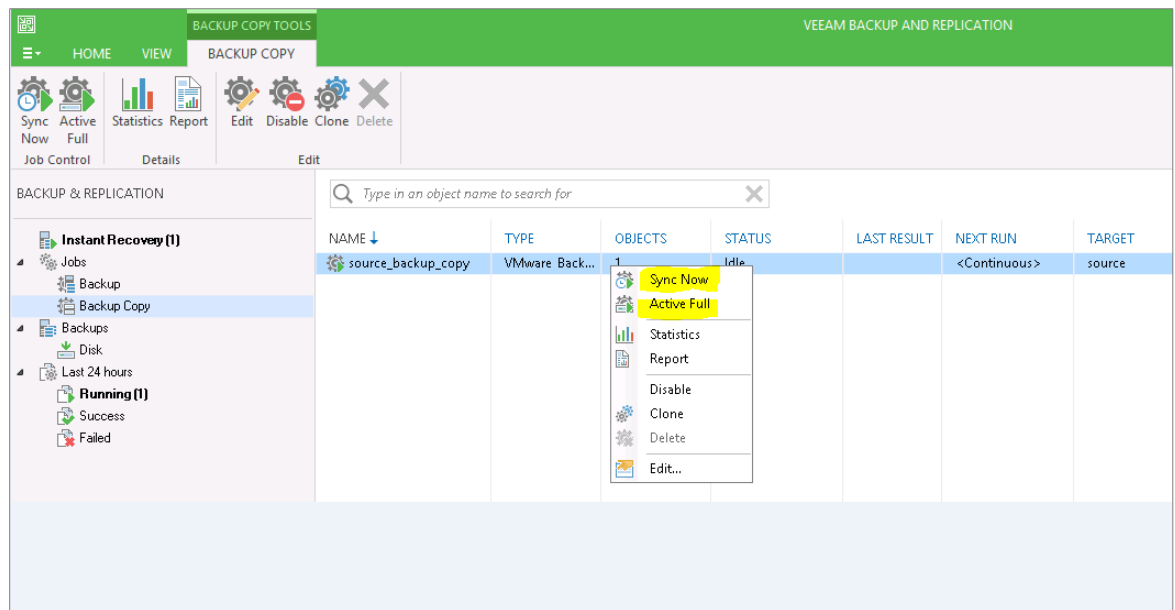
- Schedule the job as needed.



6. Click **Finish**.



7. Select the backup, right-click and select one the following operations as needed:
 - **Sync Now** - Traditional Veeam backup copy job in which the restore points are sync'ed 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.



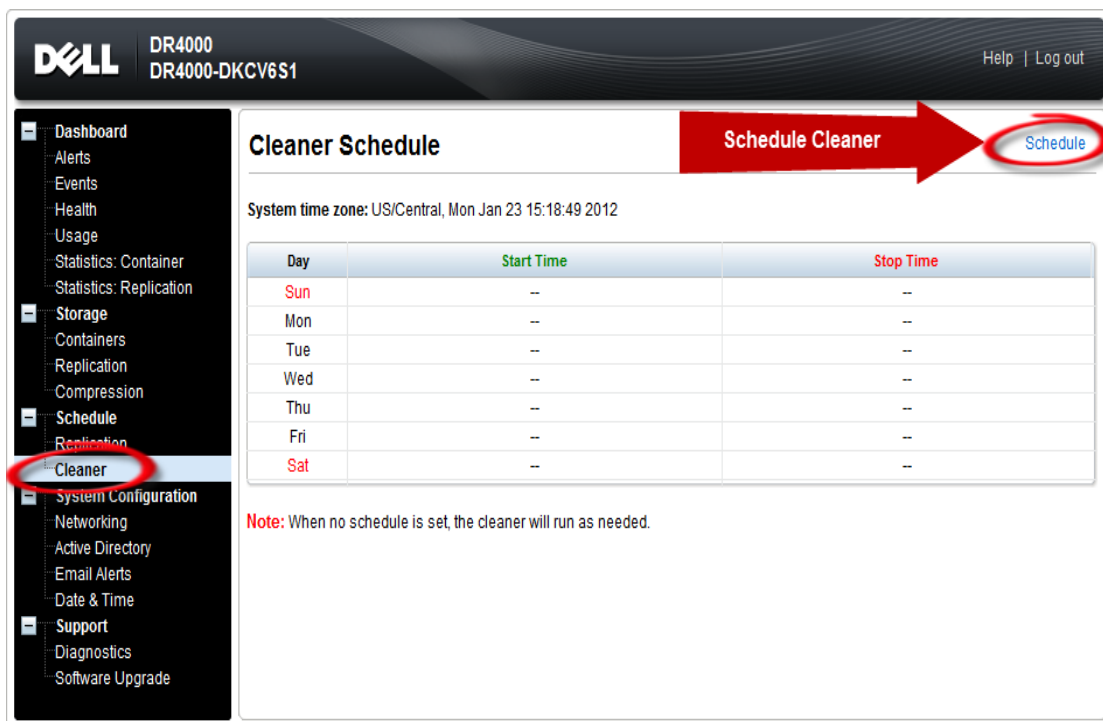
6 Setting up the DR Series system cleaner

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 separate time from backup and replication jobs.



Cleaner Schedule [Schedule Cleaner](#) [Schedule](#)

System time zone: US/Central, Mon Jan 23 15:18:49 2012

Day	Start Time	Stop Time
Sun	--	--
Mon	--	--
Tue	--	--
Wed	--	--
Thu	--	--
Fri	--	--
Sat	--	--

Note: When no schedule is set, the cleaner will run as needed.

7 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.

