Technical Report

OnCommand System Manager 3.0 and 3.1 Workflow Guide

New and Enhanced Workflows in System Manager 3.0 and 3.1

Yuvaraju Balaraman, NetApp

December 2013 | TR-4214

New and Enhanced Workflows in System Manager 3.0 and 3.1

System Manager 3.0 and 3.1 provides enhanced user interfaces and supports new features of the clustered NetApp® Data ONTAP® 8.2.x operating system. This document illustrates and describes new and enhanced System Manager 3.0 and 3.1 workflows for clustered Data ONTAP 8.2.x.

TABLE OF CONTENTS

LIST OF FIGURES

1 Purpose

This document provides an overview of new and enhanced workflows introduced with OnCommand® System Manager 3.0 and 3.1. With System Manager 3.0 and 3.1, new and enhanced workflows are available for clustered Data ONTAP 8.2.x. This release continues to support existing workflows supported in the System Manager 2.2 release.

The guide also details the steps involved in each workflow with graphical screenshots of the application. This guide can be used as a reference document in addition to the user guide to perform effective operations using OnCommand System Manager 3.0 and 3.1.

For information about System Manager 2.2 workflows, refer to TR-4031: “System Manager 2.2 Workflow Guide.” This guide lists the most common workflows used by storage administrators for configuration and ongoing management of storage controllers for both 7-Mode and clustered Data ONTAP operations.

2 Workflows Available in System Manager 3.0 and 3.1 Versions

The following section contains use cases and workflows for managing NetApp storage systems running clustered Data ONTAP 8.2.

2.1 Storage Virtual Machine (SVM) Simplified Setup

With System Manager 3.0 and 3.1, setting up a storage virtual machine (SVM) and configuring data protocols on the SVM have been simplified for clustered Data ONTAP 8.2.x.

Set Up SVM and Configure CIFS, NFS Data Protocols

Use case: A user wants to set up an SVM with FlexVol® volumes and allow CIFS and NFS as the data protocols on the volumes in the SVM.

Note: CIFS and NFS protocols have to be licensed on the cluster.

As part of this workflow, we do the following:

- Set up an SVM with FlexVol volumes.
- Configure CIFS and NFS data protocols on the SVM.
- Set up an SVM administrator account.

Workflow Steps

SVM Setup

1. From the Home tab, double-click the appropriate storage system.
2. Expand the SVM’s hierarchy in the left navigation pane.
3. In the navigation pane, select the cluster.
4. Click Create.
5. In the SVM Setup window, specify the SVM details such as the SVM name, select FlexVol volumes as the volume type, select CIFS and NFS as the protocols allowed, and specify the SVM language, root volume security style, and an aggregate for the root volume.

   The default language setting for an SVM is C.UTF-8.

   By default, the aggregate with the maximum amount of free space is selected as the container for the root volume of the SVM. Based on the protocols selected, the default security style and the root aggregate are selected.

The security style is set to NTFS if you select CIFS protocol or a combination of CIFS protocol and the other protocols. The security style is set to UNIX® if you select NFS, iSCSI, or FC/FCoE or a combination of these protocols.

6. Specify the DNS domain names and the name server IP addresses to configure the DNS services. The default values are selected from the existing SVM configurations.
7. Click Submit & Continue.

   The SVM is created with the specified configuration.

The SVM that you created starts automatically. The root volume name is automatically generated as *SVM name*_root. By default, the *vsadmin* user account is created and is in the locked state.

Figure 1) Storage virtual machine setup.

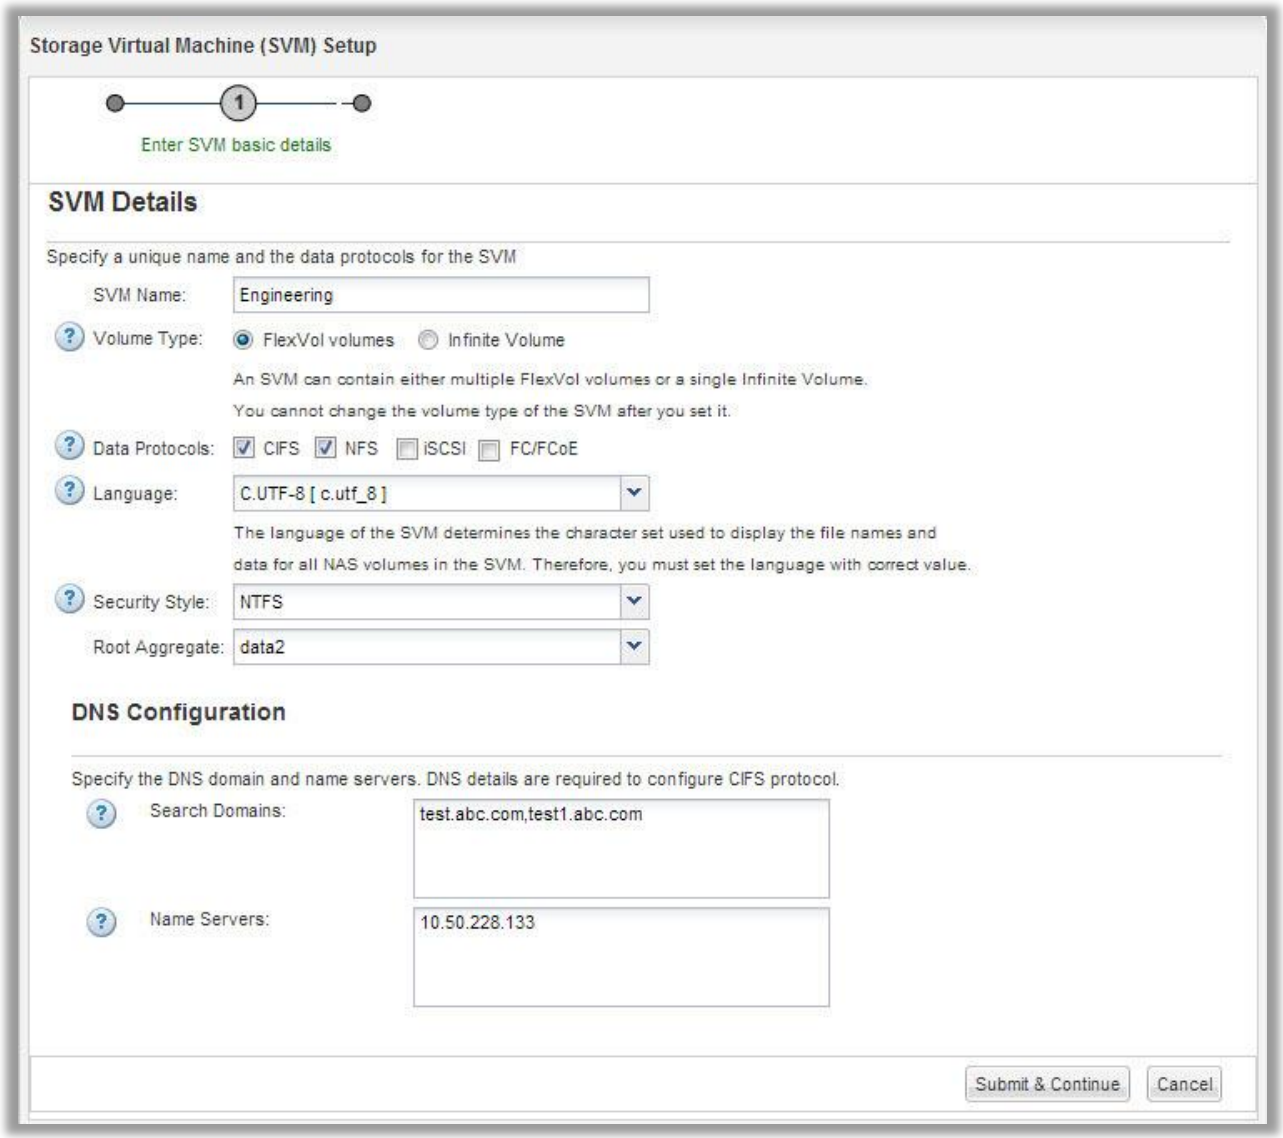

Configuring CIFS and NFS Protocols on an SVM

You can configure CIFS and NFS protocols on the SVM to provide file-level data access for NAS clients. To enable the CIFS protocol, you must create the data LIFs and the CIFS server. To enable the NFS protocol, you can specify the NIS details in addition to the data LIFs.

1. In the Data LIF Configuration section, specify the network details to create data LIFs. You can either retain the same data LIF configuration for both CIFS and NFS or configure a new LIF for each protocol.
2. Specify the following information to create a CIFS server:
   a. CIFS server name
   b. Active Directory® to associate with the CIFS server
   c. Organizational unit (OU) within the Active Directory domain to associate with the CIFS server; by default, this parameter is set to CN=Computers
   d. Credentials of an administrative account that has sufficient privileges to add the CIFS server to the OU
   e. Optional: You can also specify the IP addresses of the NIS servers and NIS domain name to configure NIS services on the SVM
   f. Click Submit & Close

Figure 2) Configure CIFS/NFS protocol.

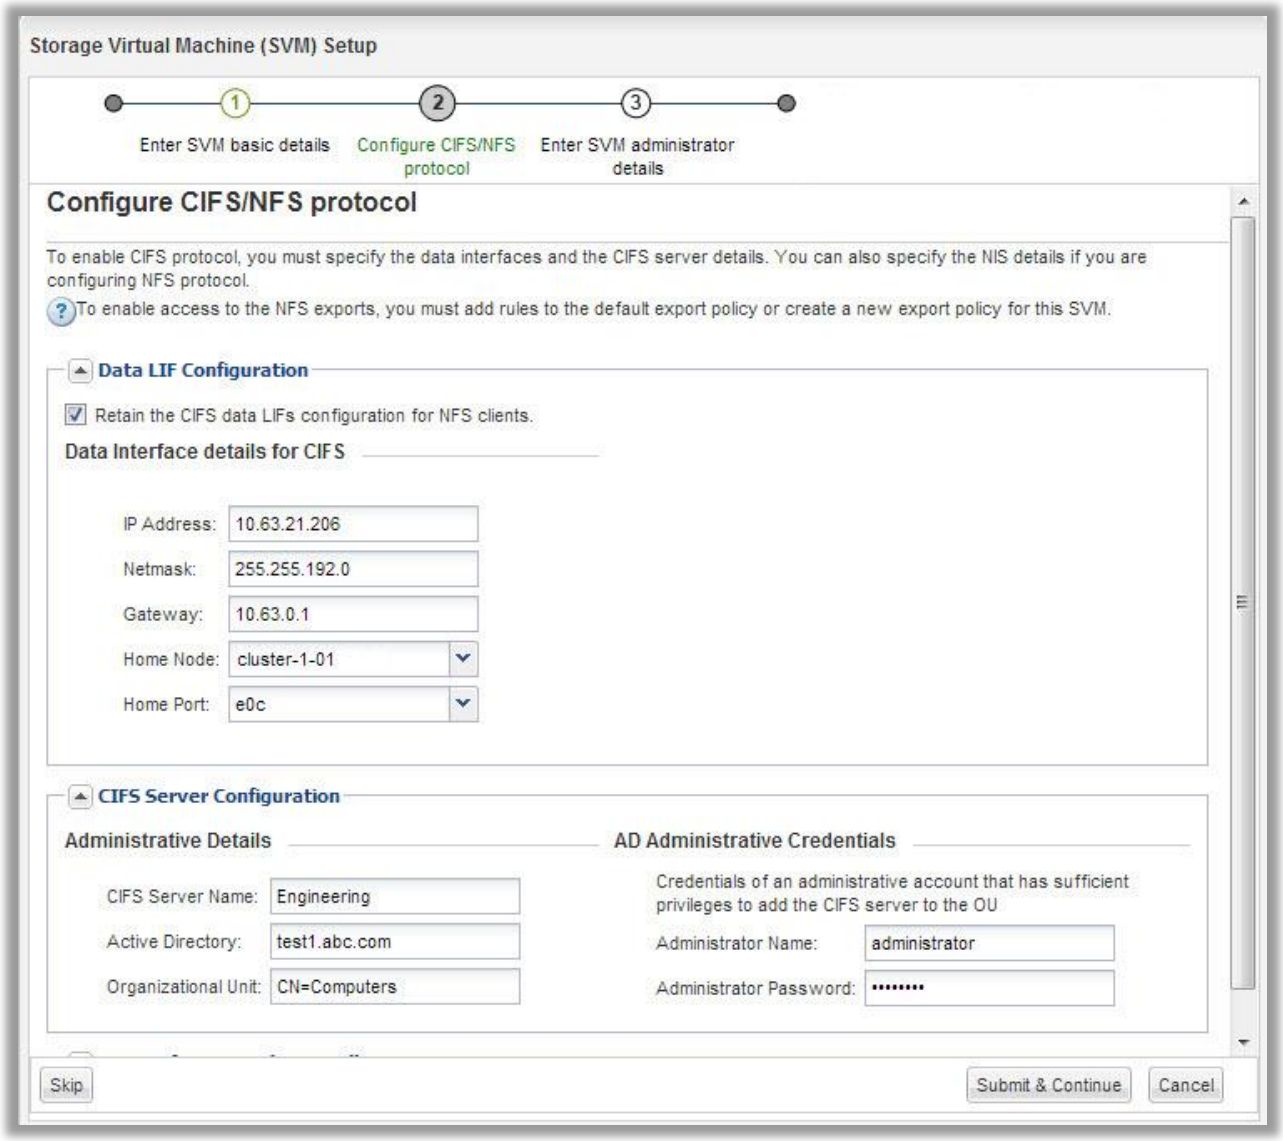

Figure 3) NIS configuration for NFS protocol.

In case you do not have the required information to configure the CIFS or NFS protocol, you can skip configuring the protocol and configure it later.

To configure the protocols, double-click the appropriate storage system in the homepage, expand the SVM hierarchy, select the SVM in the navigation pane, and click the protocol you want to configure.

The CIFS server and NIS domain are configured with the specified configuration. Data LIFs are created. By default, the data LIFs have management access. You can view the configuration details in the Summary page.

Delegating Administration to an SVM Administrator

After setting up a functional SVM or an SVM with basic network configuration, you can optionally delegate the administration of the SVM to an SVM administrator.

1. In the Administrator Details section, set up a password for the vsadmin user account.
2. By default, the data LIFs have management access.

Figure 4) SVM administration.

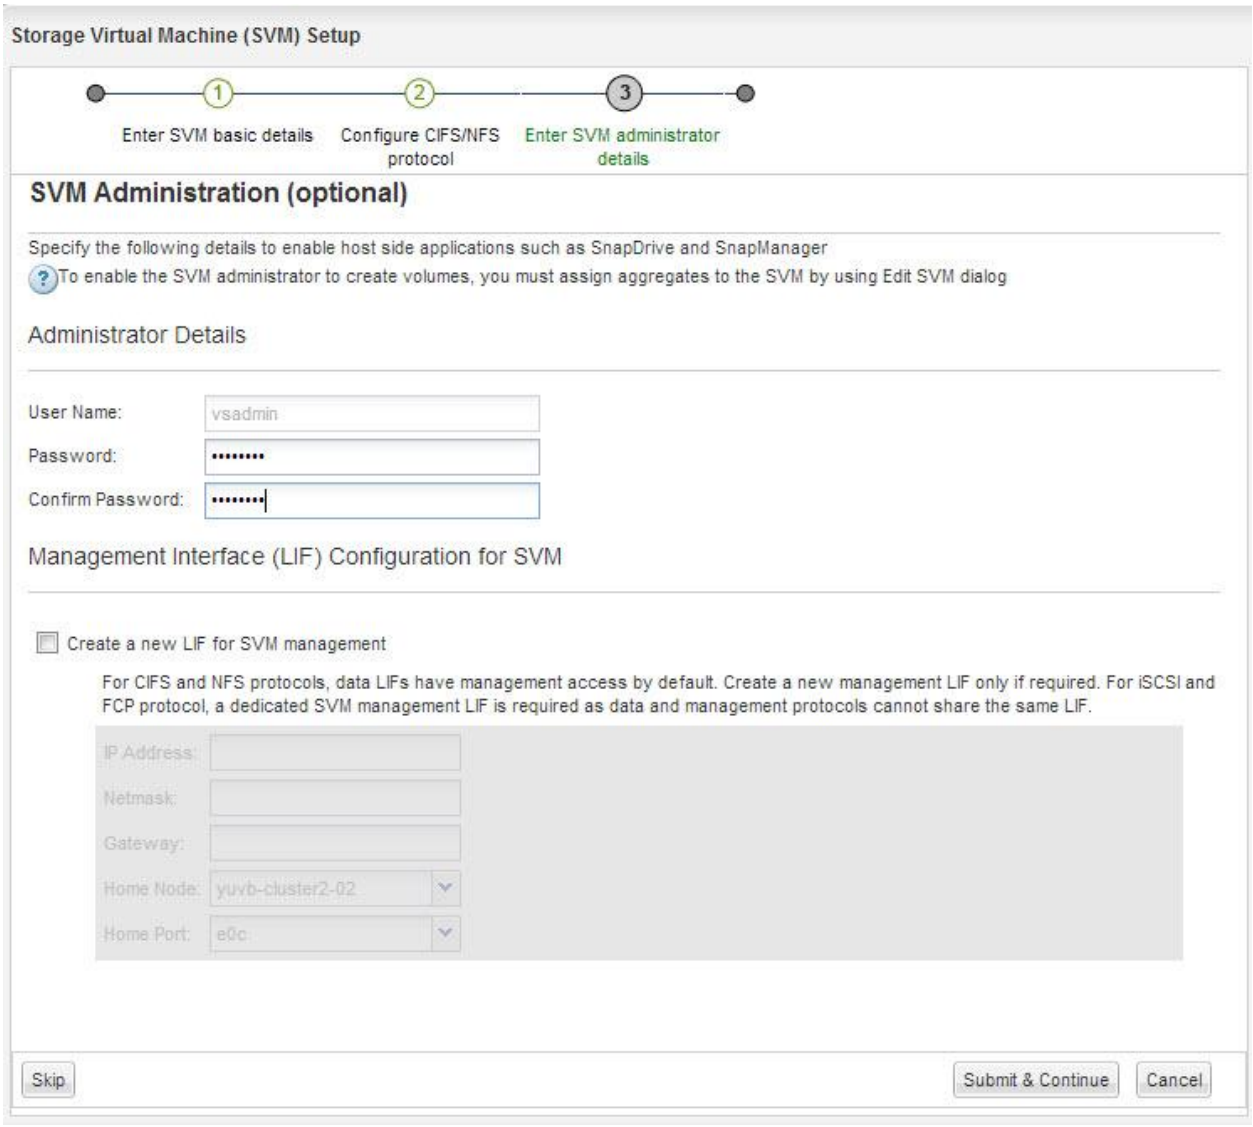

Set Up SVM and Configure iSCSI Data Protocol

Use case: A user wants to set up a storage virtual machine and allow iSCSI as the data protocol on the volumes in the SVM.

Note:

1. iSCSI protocol must be licensed on the cluster.
2. All the nodes in the cluster must be healthy.
3. Each node must have at least two data ports, and the port state must be up.

As part of this workflow, we do the following:

1. Set up an SVM.
2. Configure the iSCSI data protocol on the volumes in the SVM.
3. Delegate administration to an SVM administrator.

Workflow Steps

SVM Setup

1. From the Home tab, double-click the appropriate storage system.
2. Expand the SVM's hierarchy in the left navigation pane.
3. In the navigation pane, select the cluster.
4. Click Create.
5. In the SVM Setup window, specify the SVM details such as the SVM name, select FlexVol volumes as the volume type, and select iSCSI as the protocol allowed, the SVM language, the root volume security style, and its root aggregate.
6. The default language setting for an SVM is C.UTF-8.
7. By default, the aggregate with the maximum amount of free space is selected as the container for the root volume of the SVM. Based on the protocols selected, the default security style and the root aggregate are selected. The security style is set to UNIX if you select iSCSI protocol.
8. Specify the DNS domain names and the name server IP addresses to configure the DNS services. The default values are selected from the existing SVM configurations.
9. Click Submit & Continue.

   The SVM is created with the specified configuration.

Figure 5) SVM setup.

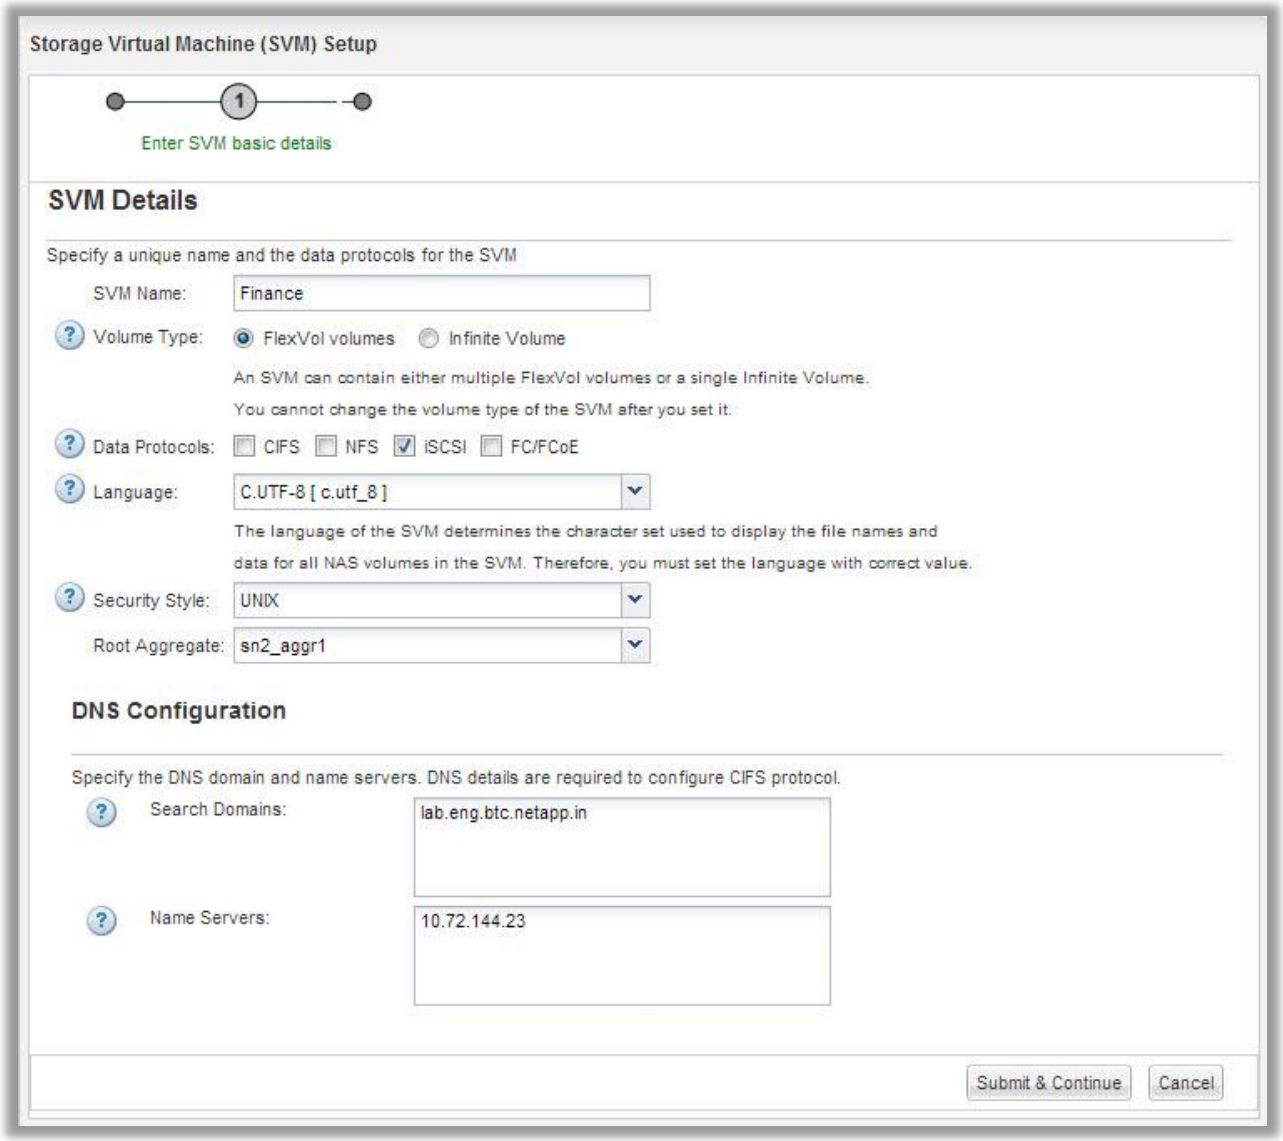

Configuring iSCSI Protocol on an SVM

You can configure the iSCSI protocol on the SVM to provide block-level data access. You can create iSCSI LIFs and portsets and add the LIFs to the portsets. LIFs are created on the most suitable adapters and assigned to portsets for data path redundancy.

1. Specify an alias for the iSCSI target. If you do not specify a target alias, the SVM name is used as an alias.
2. Select the number of iSCSI LIFs that can be assigned to a single node. The minimum number for LIFs per node is one.
3. Specify the network details to create iSCSI LIFs.

   The starting IP address specifies the first of the contiguous addresses in the LIF IP address pool.

Figure 6) Configure iSCSI protocol.

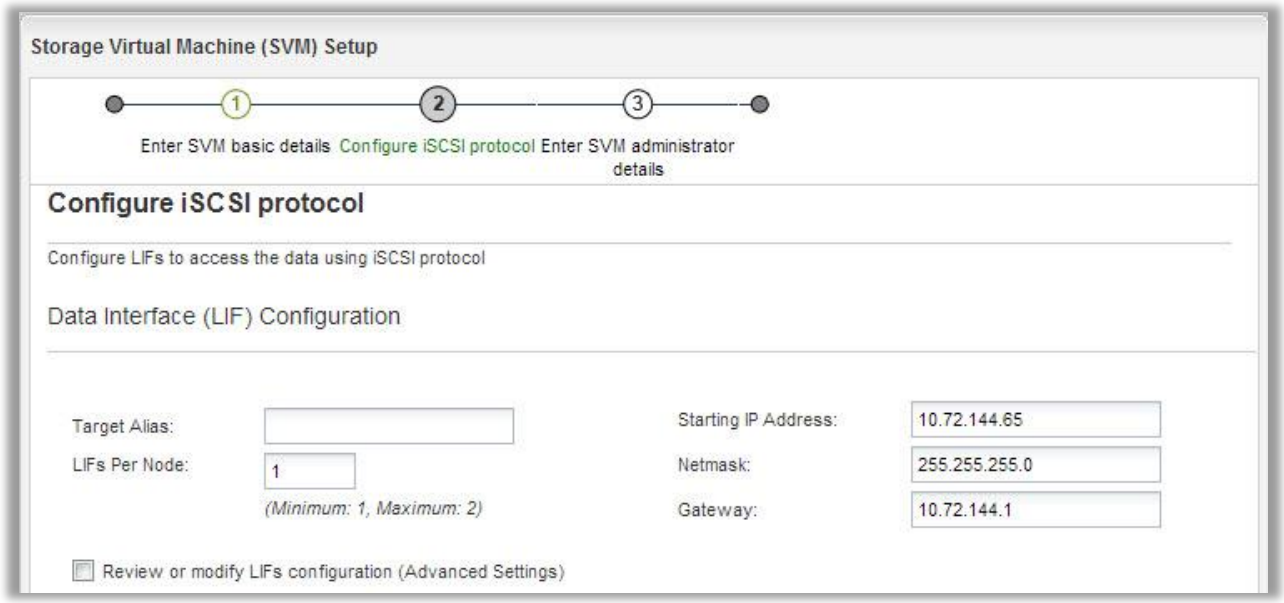

4. If you want to verify or modify the automatically generated iSCSI LIFs configuration, select Review or Modify LIFs configuration (Advanced Settings).

Figure 7) Review LIFs configuration.

`

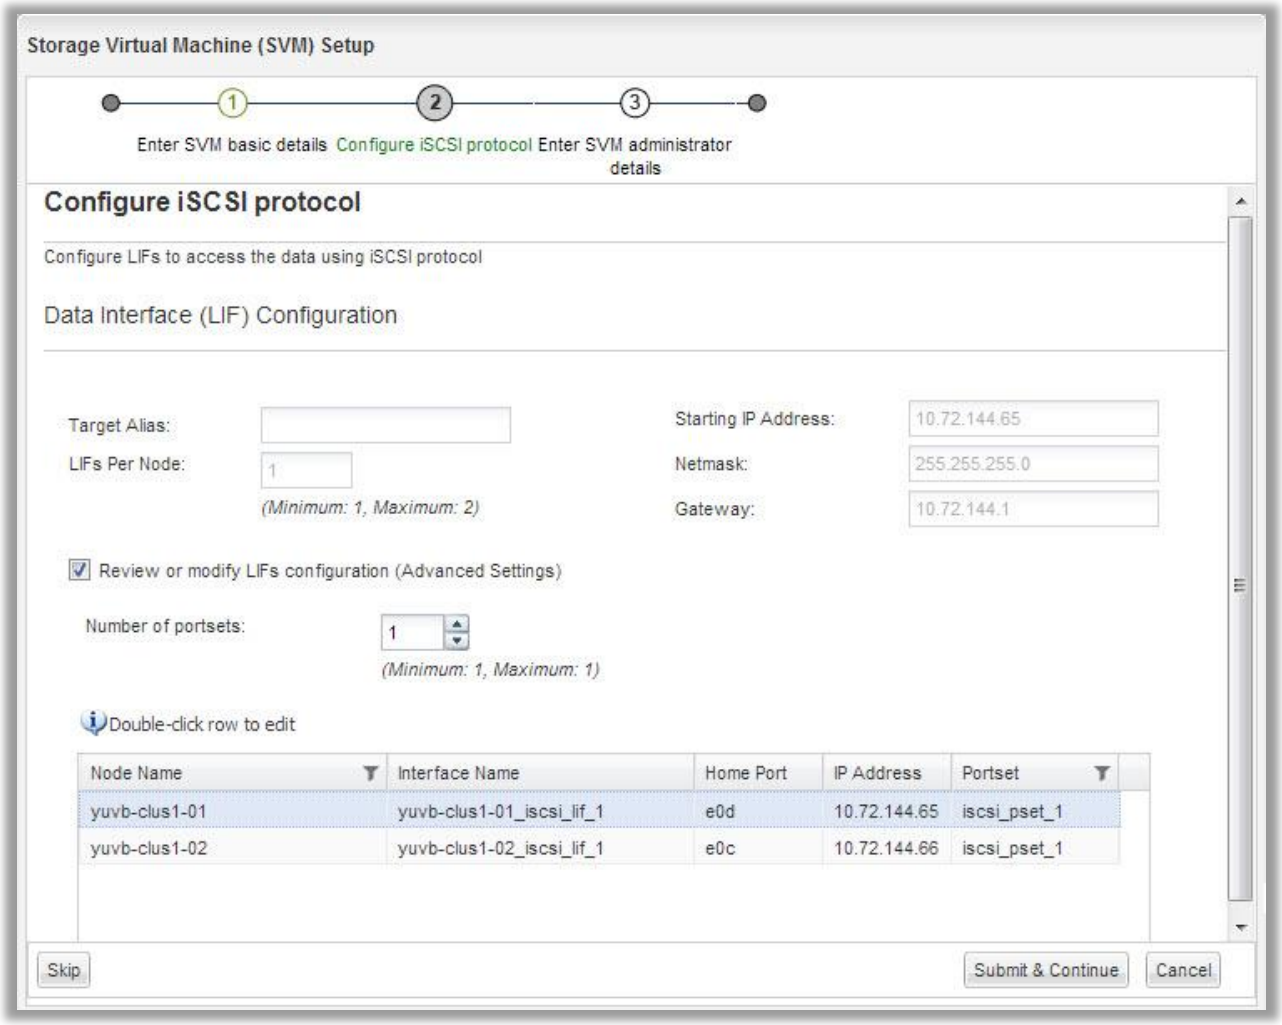

5. You can modify only the LIF name, home port, and LIF IP address. By default, the portsets are set to the minimum value. Do not specify duplicate entries. If you specify duplicate LIF names, System Manager appends numeric values to the duplicate LIF names.

Figure 8) Modify LIFs configuration.

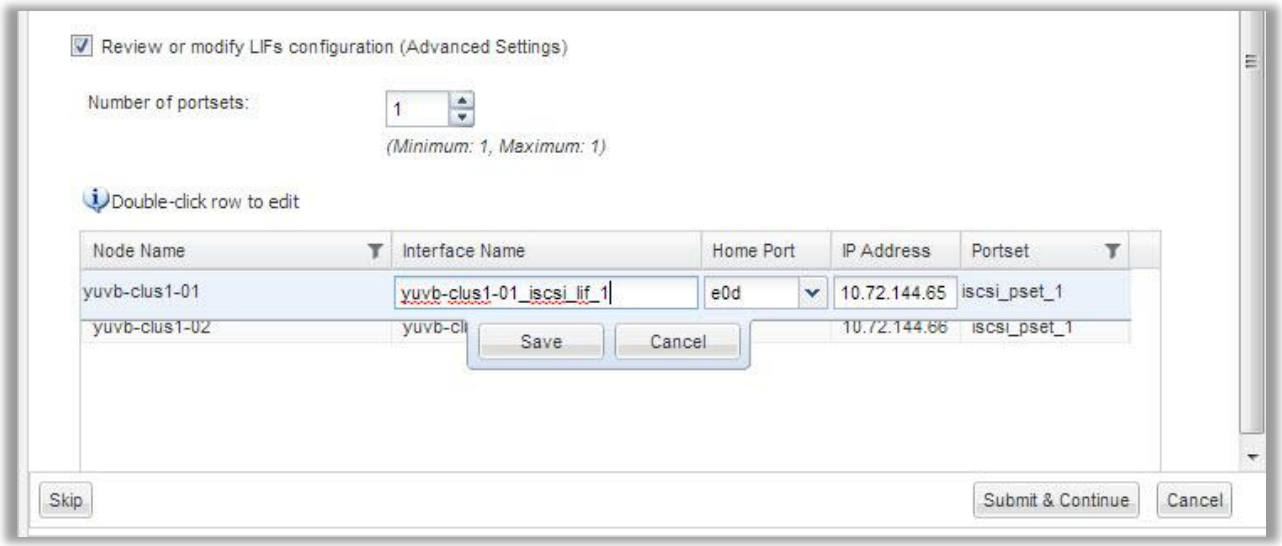

6. Based on the selected portset, the LIFs are distributed across the portsets using a round-robin method for redundancy in case of node or port failure.
7. Click Submit & Continue.

The iSCSI LIFs and portsets are created with the specified configuration. The LIFs are distributed accordingly among the portsets. The iSCSI service is started if all the LIFs are created successfully.

If the LIF creation fails, you can use the Network Interfaces window to create the LIFs, attach the LIFs to the portsets by using the LUNs window, and start the iSCSI service by using the iSCSI window.

Delegating Administration to an SVM Administrator

After setting up a functional SVM or an SVM with basic network configuration, you can optionally delegate the administration of the SVM to an SVM administrator.

1. In the Administrator Details section, set up a password for the vsadmin user account.
2. Select Create a LIF for SVM management and specify the network details. A dedicated SVM management LIF is required for SAN protocols in which data and management protocols cannot share the same LIF. An SVM management LIF can be created only on data ports.

Figure 9) SVM administration.

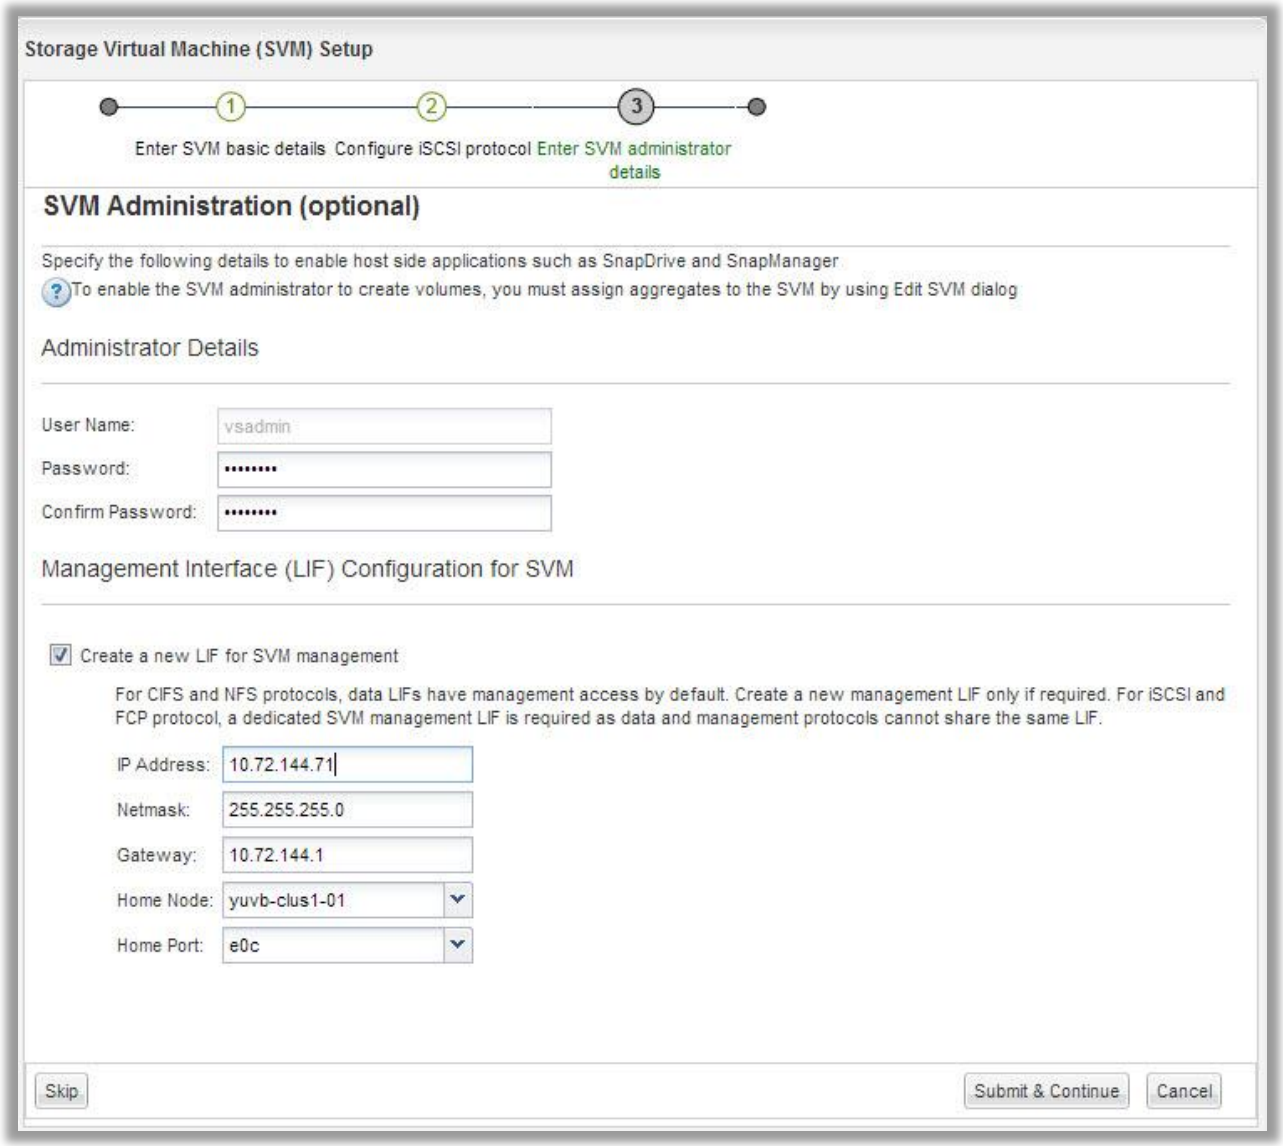

The summary page provides configuration details of the new SVM and the protocols.

Figure 10) SVM summary.

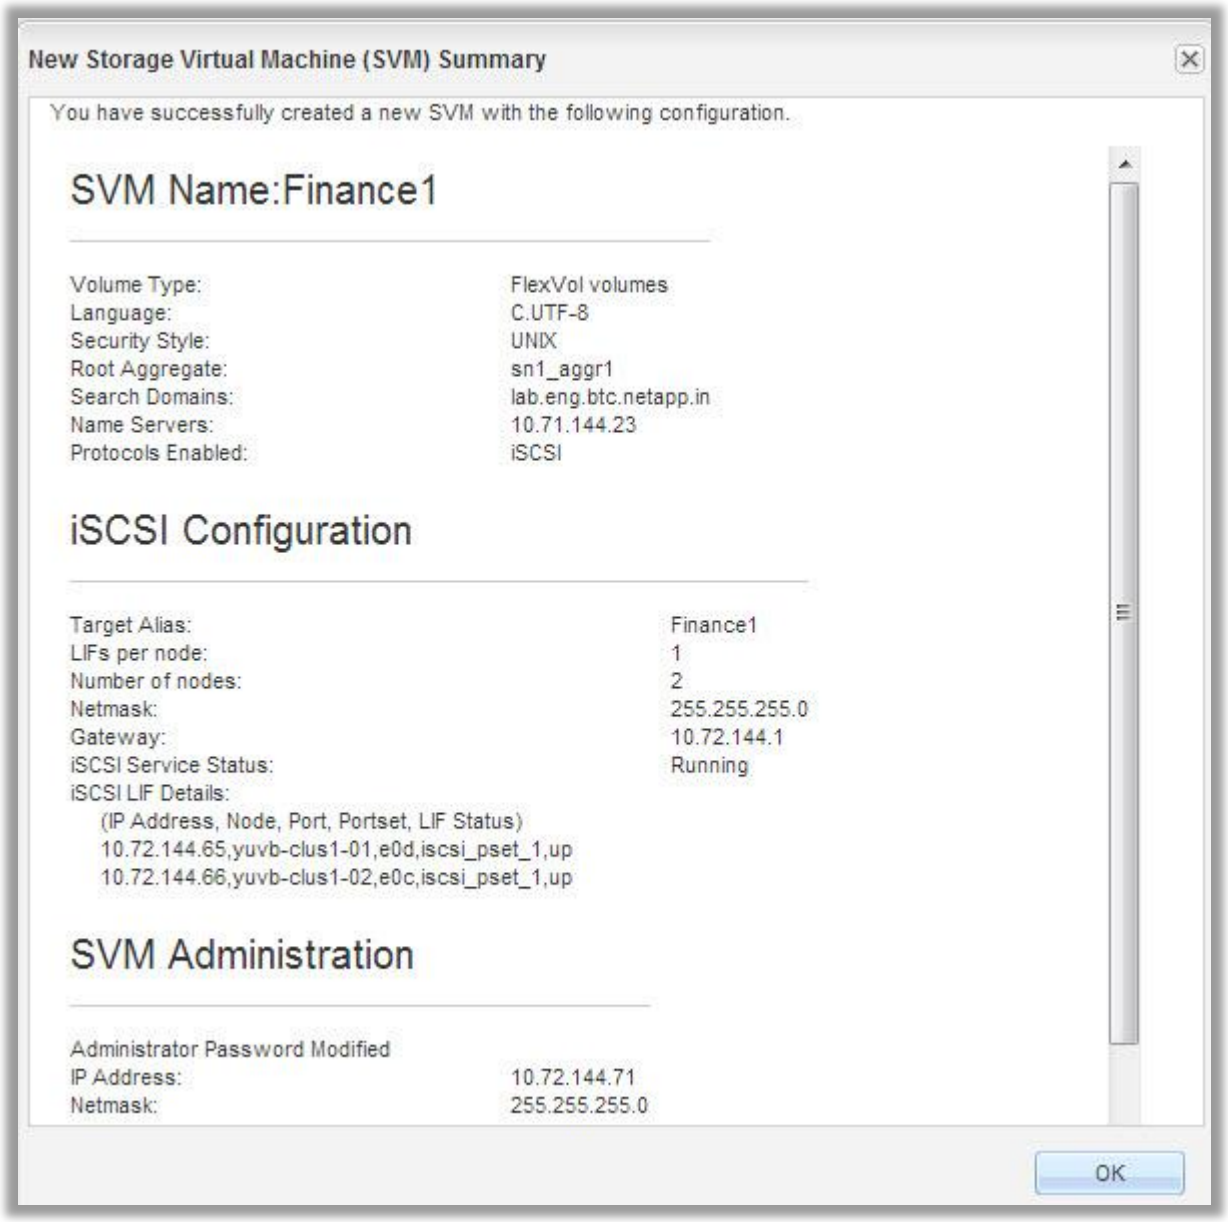

2.2 Creating CIFS Local Users and Groups Workflow

Use case: Allow a business unit's users, for example, HR users, to have access to HR department data using the CIFS data protocol on the volumes in the SVM, using a local user account if the domain controllers are unavailable.

With System Manager 3.0 and 3.1, you can create local users and groups on the SVM. The CIFS server can use local users for CIFS authentication and can use both local users and groups for authorization when determining share, file, and directory access rights.

As part of this workflow, we do the following:

- Create a local group.

- Create local Windows users.
- Add local users to the group.

Creating a Local Windows Group

With System Manager 3.0 and 3.1, you can create local Windows® groups that can be used for authorizing access to data contained in an SVM over an SMB connection. You can also assign privileges that define the user rights or capabilities that a member of the group has while performing administrative activities.

Steps

1. From the homepage, double-click the appropriate storage system.
2. Expand the SVM's hierarchy in the left navigation pane.
3. In the navigation pane, select the SVM and click Configuration > Local Users and Groups > Windows.
4. In the Groups tab, click Create.

Figure 11) Create group window.

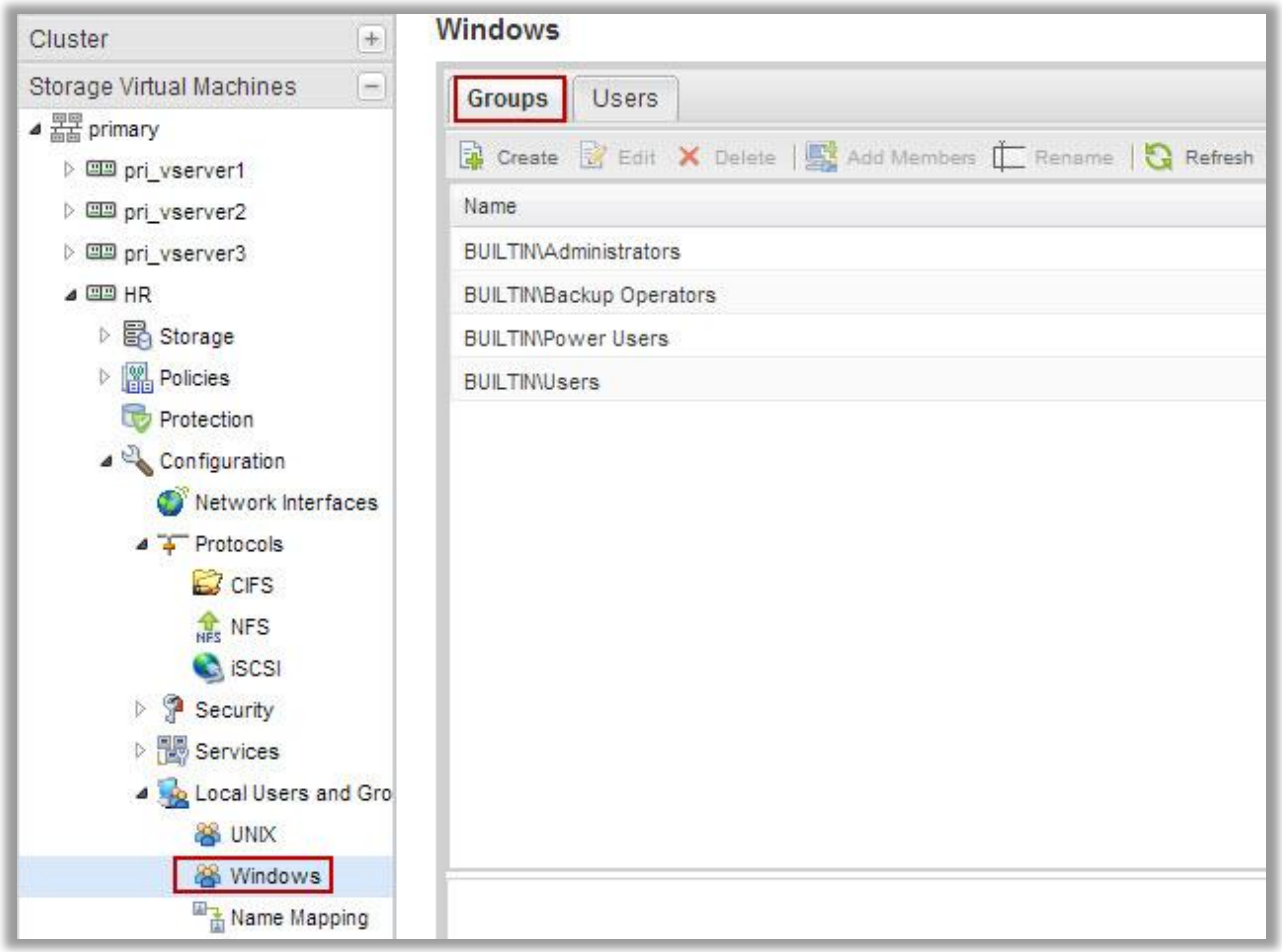

5. Specify a name for the group and a description that helps you identify this new group.

6. Assign a set of privileges to the group. You can select the privileges from the predefined set of supported privileges.

Figure 12) Create group.

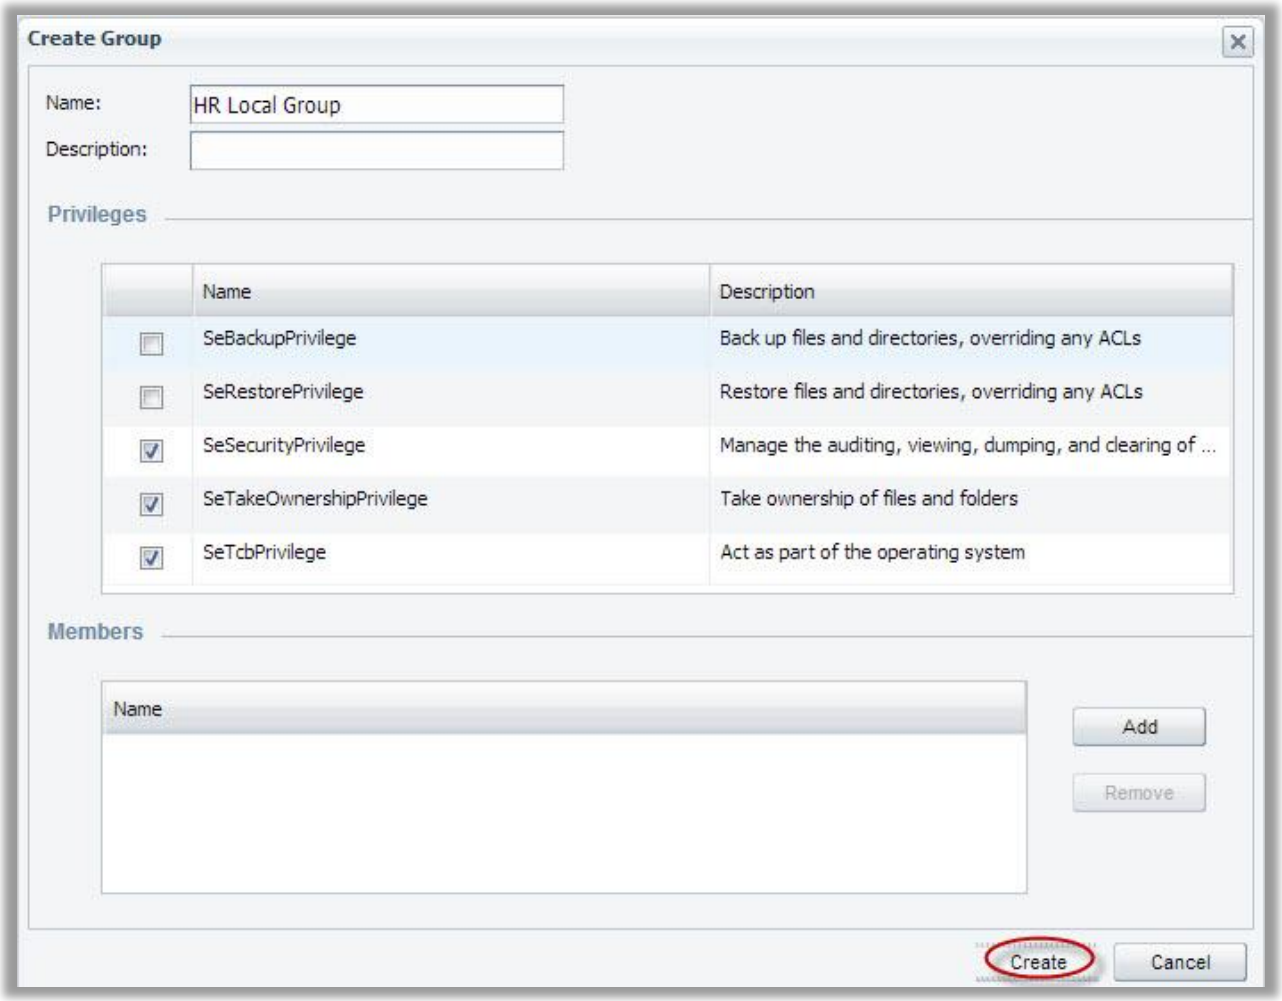

7. Click Create. The local Windows group is created and is listed in the Groups window.

Figure 13) Group window.

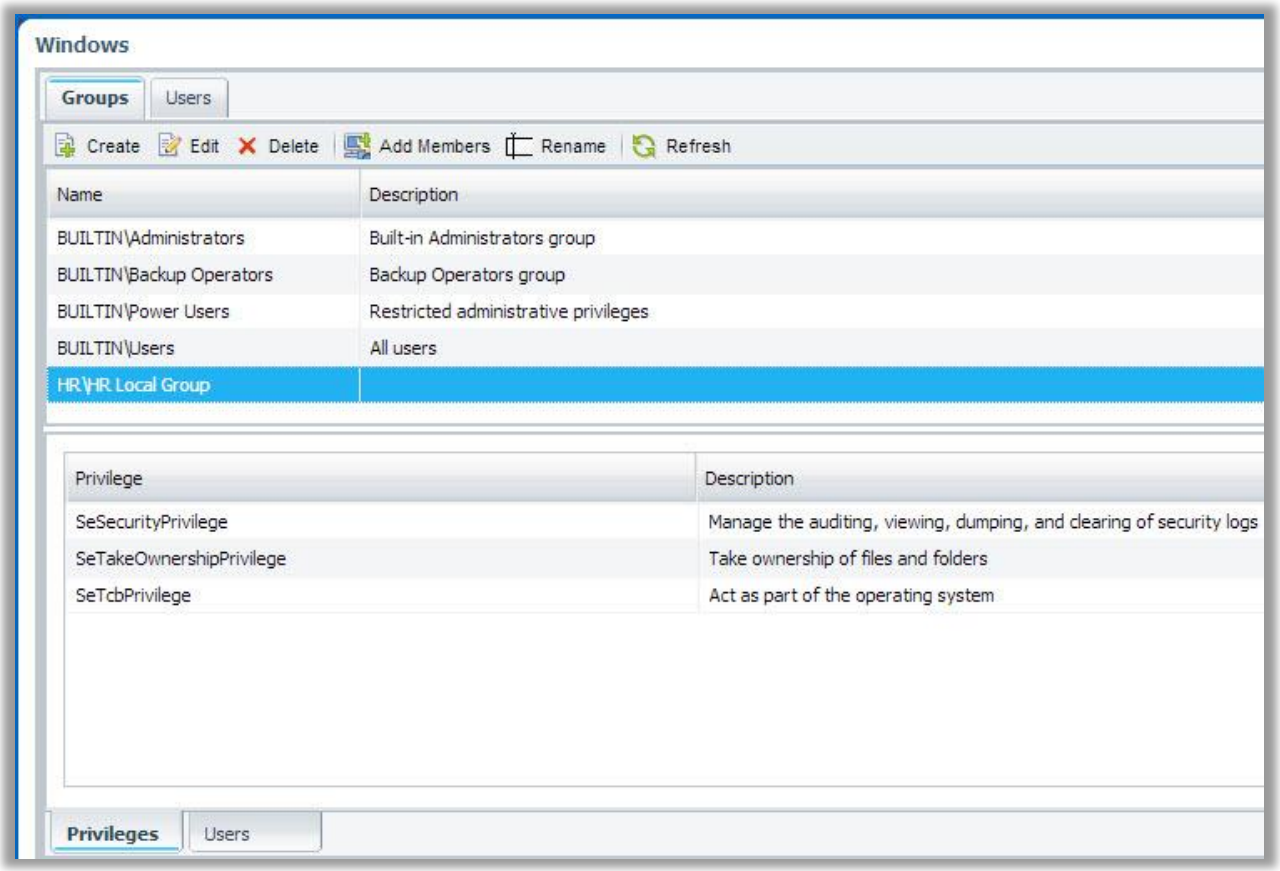

Creating a Local Windows User Account

With System Manager 3.0 and 3.1, you can create a local Windows user account that can be used to authorize access to data contained in the SVM over an SMB connection. You can also use local Windows user accounts for authentication when creating a CIFS session.

Steps

1. From the homepage, double-click the appropriate storage system.
2. Expand the SVM's hierarchy in the left navigation pane.
3. In the navigation pane, select the SVM and click Configuration > Local Users and Groups > Windows.
4. In the Users tab, click Create.

Figure 14) Create user's window.

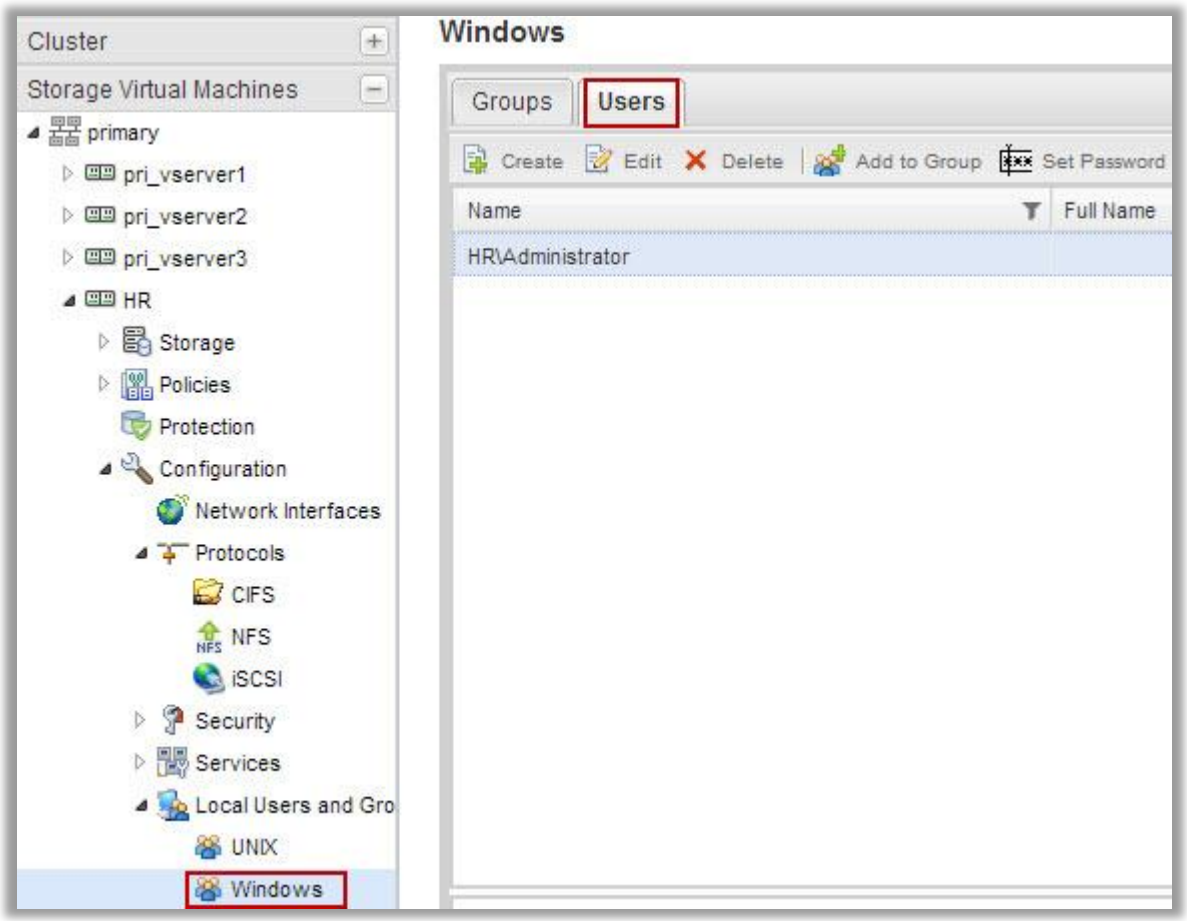

5. Specify a name for the local user.
6. Specify the full name of the local user and a description that helps you identify this new user.
7. Enter a password for the local user and confirm the password.
8. Click Add to assign group memberships to this user.
9. In the Add Groups window, select the groups from the list of available groups in the SVM.
10. Click Create.

Figure 15) Create user.

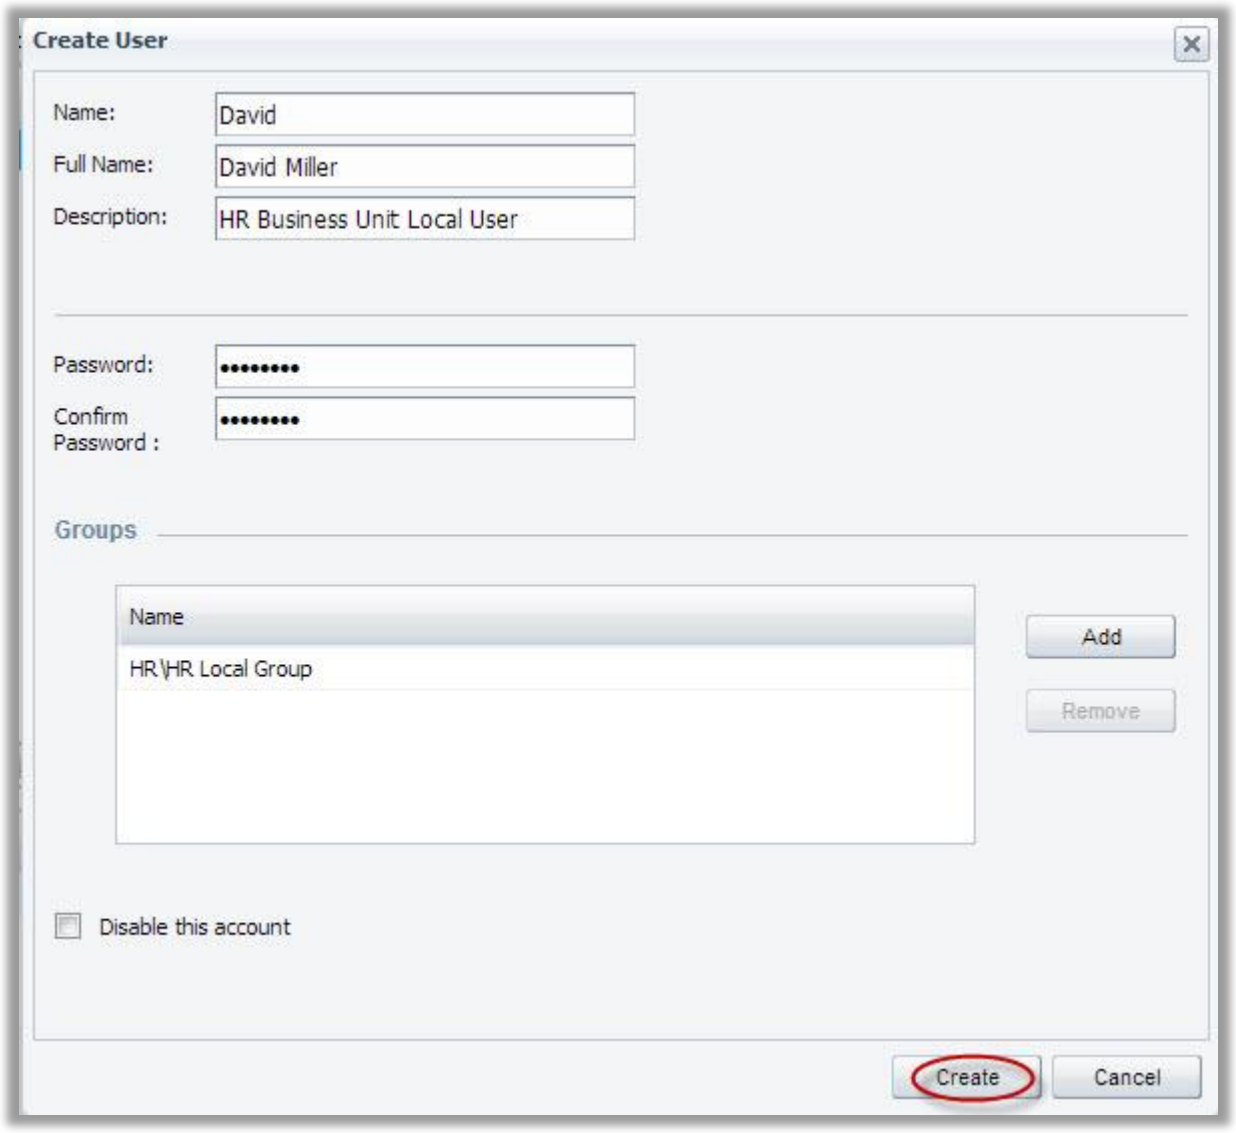

The local Windows user account is created and is assigned membership to the selected groups. The user account is listed in the Users tab.

Figure 16) User's window.

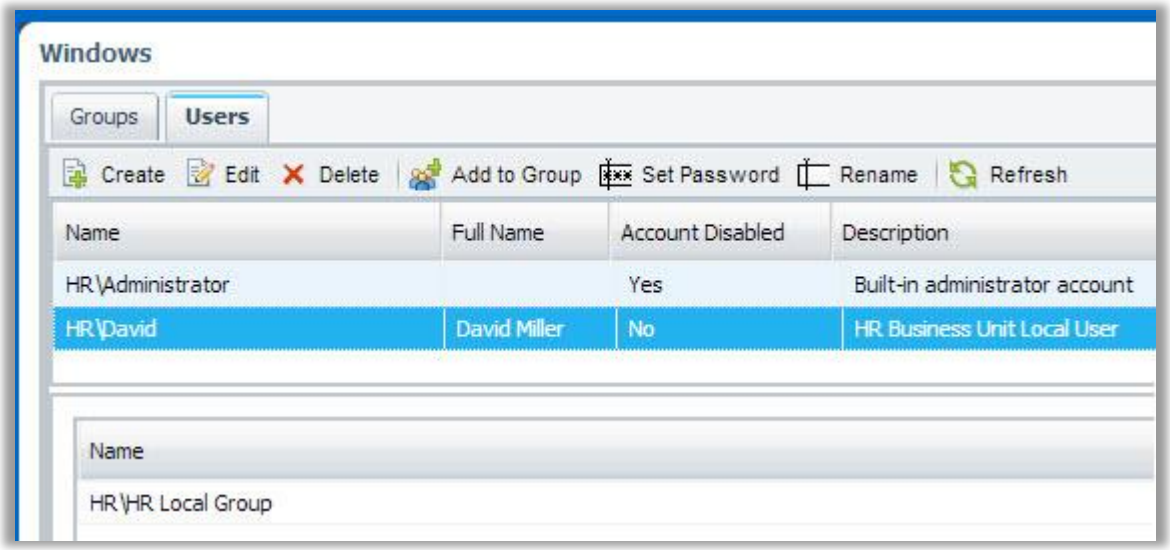

2.3 Continuously Available CIFS Share Workflow

Use case: A storage administrator has received a request to create a continuously available CIFS share.

System Manager 3.0 and 3.1 supports continuously available shares, which enables clients connecting through continuously available SMB 3.0 to continue read/write operations during disruptive events such as takeover and giveback.

Steps

1. From the homepage, double-click the appropriate storage system.
2. Expand the SVM's hierarchy in the left navigation pane.
3. In the navigation pane, select the SVM and click Storage > Shares.
4. Click Create Share.

Figure 17) Create share.

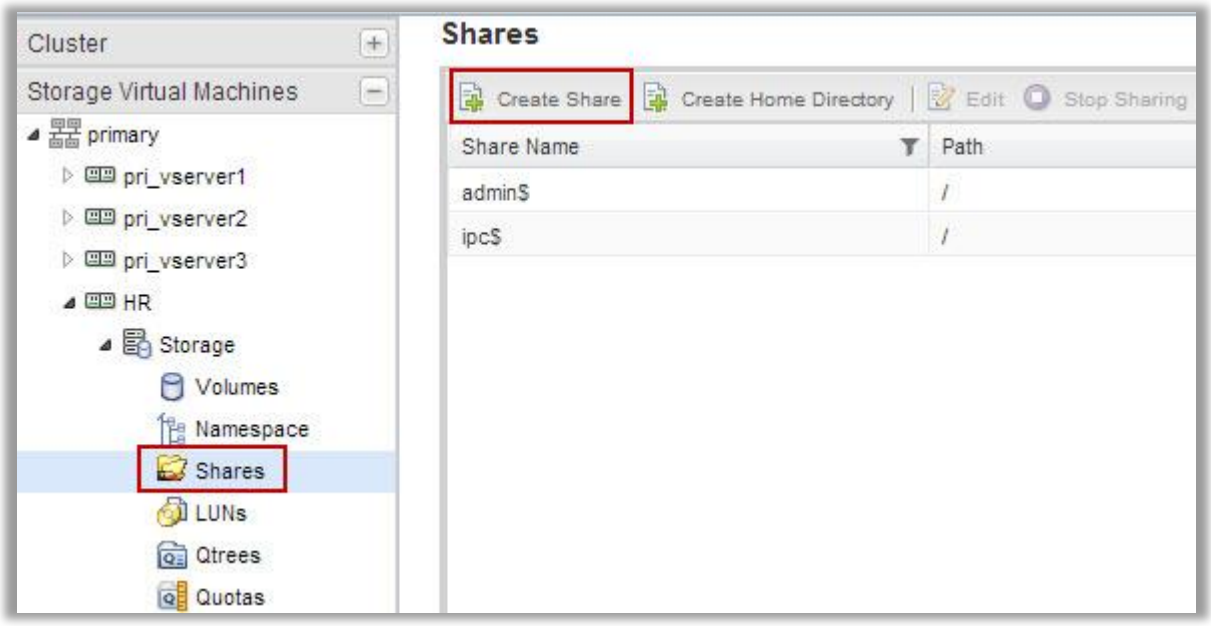

5. Click Browse and select the folder, qtree, or volume that should be shared.
6. Specify a name for the new CIFS share. Select the Enable continuous availability for Hyper-V option to permit SMB 3.0 and later clients that support it to open files persistently during nondisruptive operations.

Figure 18) New share.

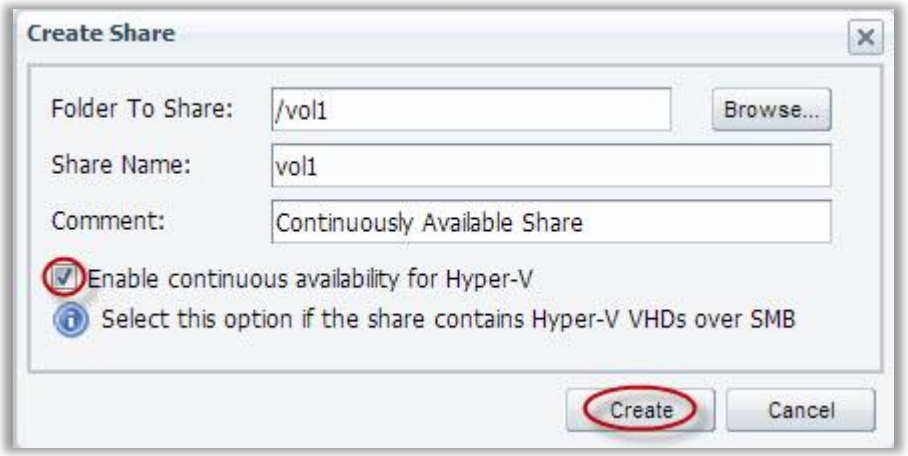

The share is created with the access permissions set to Full Control for Everyone in the group.

Figure 19) Share created.

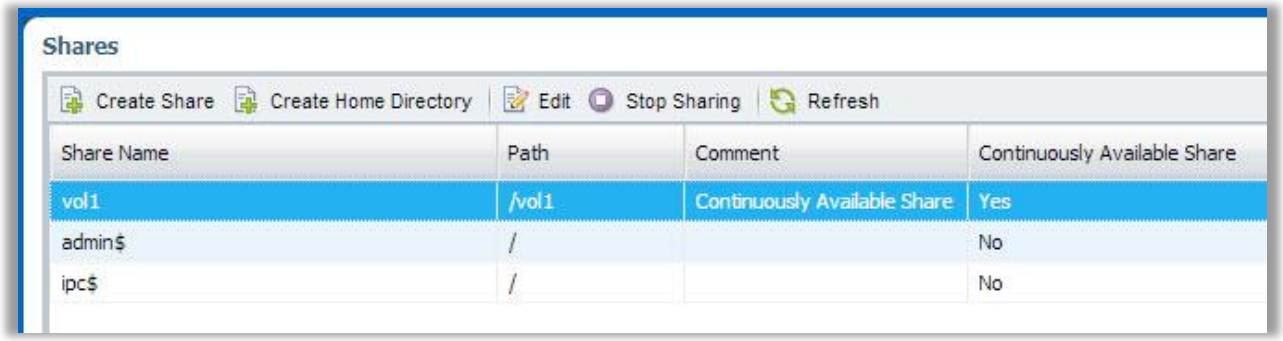

2.4 Simplified Cluster Peer Setup Workflow for Clustered Data ONTAP 8.2.x

Use case: An organization wants to back up data on its primary cluster to a secondary cluster. Primary and secondary clusters have to be peered before a SnapMirror® or SnapVault® relationship can be created.

With System Manager 3.0 and 3.1, the user interface to create cluster peer relationships is simplified for clustered Data ONTAP 8.2.x. Additionally, you can configure intercluster LIFs for one or both of the clusters if the LIFs are not already configured.

Note: A cluster peer relationship can also be set up as part of creating a SnapMirror or SnapVault relationship using System Manager 3.0 and 3.1.

Steps

1. From the homepage, double-click the appropriate storage system.
2. Expand the cluster hierarchy in the left navigation pane.
3. Click Configuration > Peers.
4. Click Create.

Figure 20) Create cluster peer.

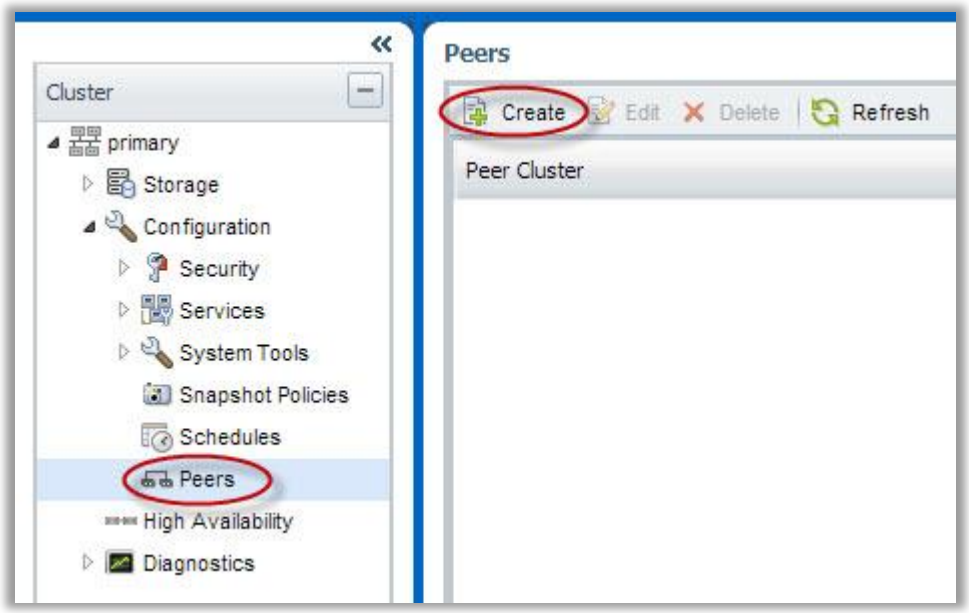

5. In the Create Cluster Peering window, specify the details for the local cluster.

Figure 21) Local cluster.

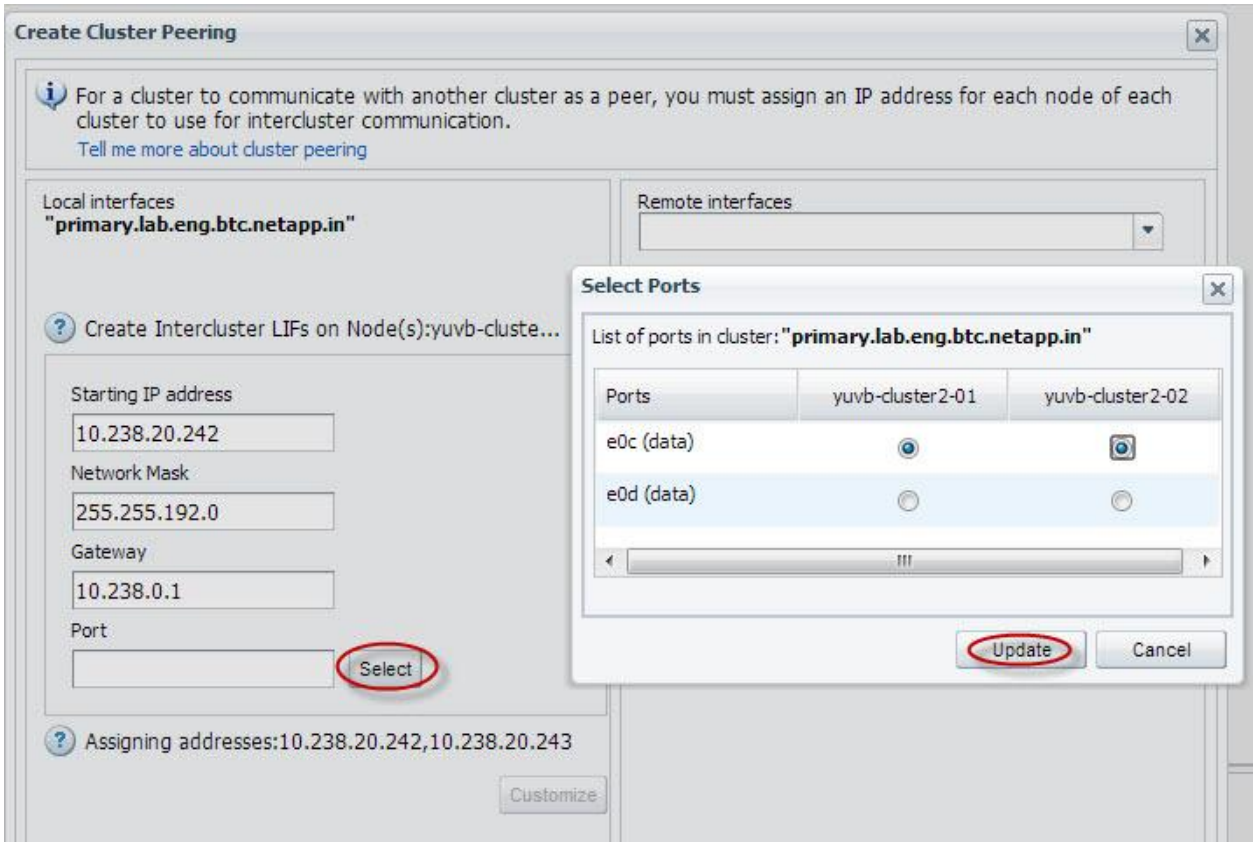

6. Click Customize to modify the interface name, IP address, network mask, gateway, and port details for the other nodes, and click Update.

Figure 22) Customize intercluster LIF.

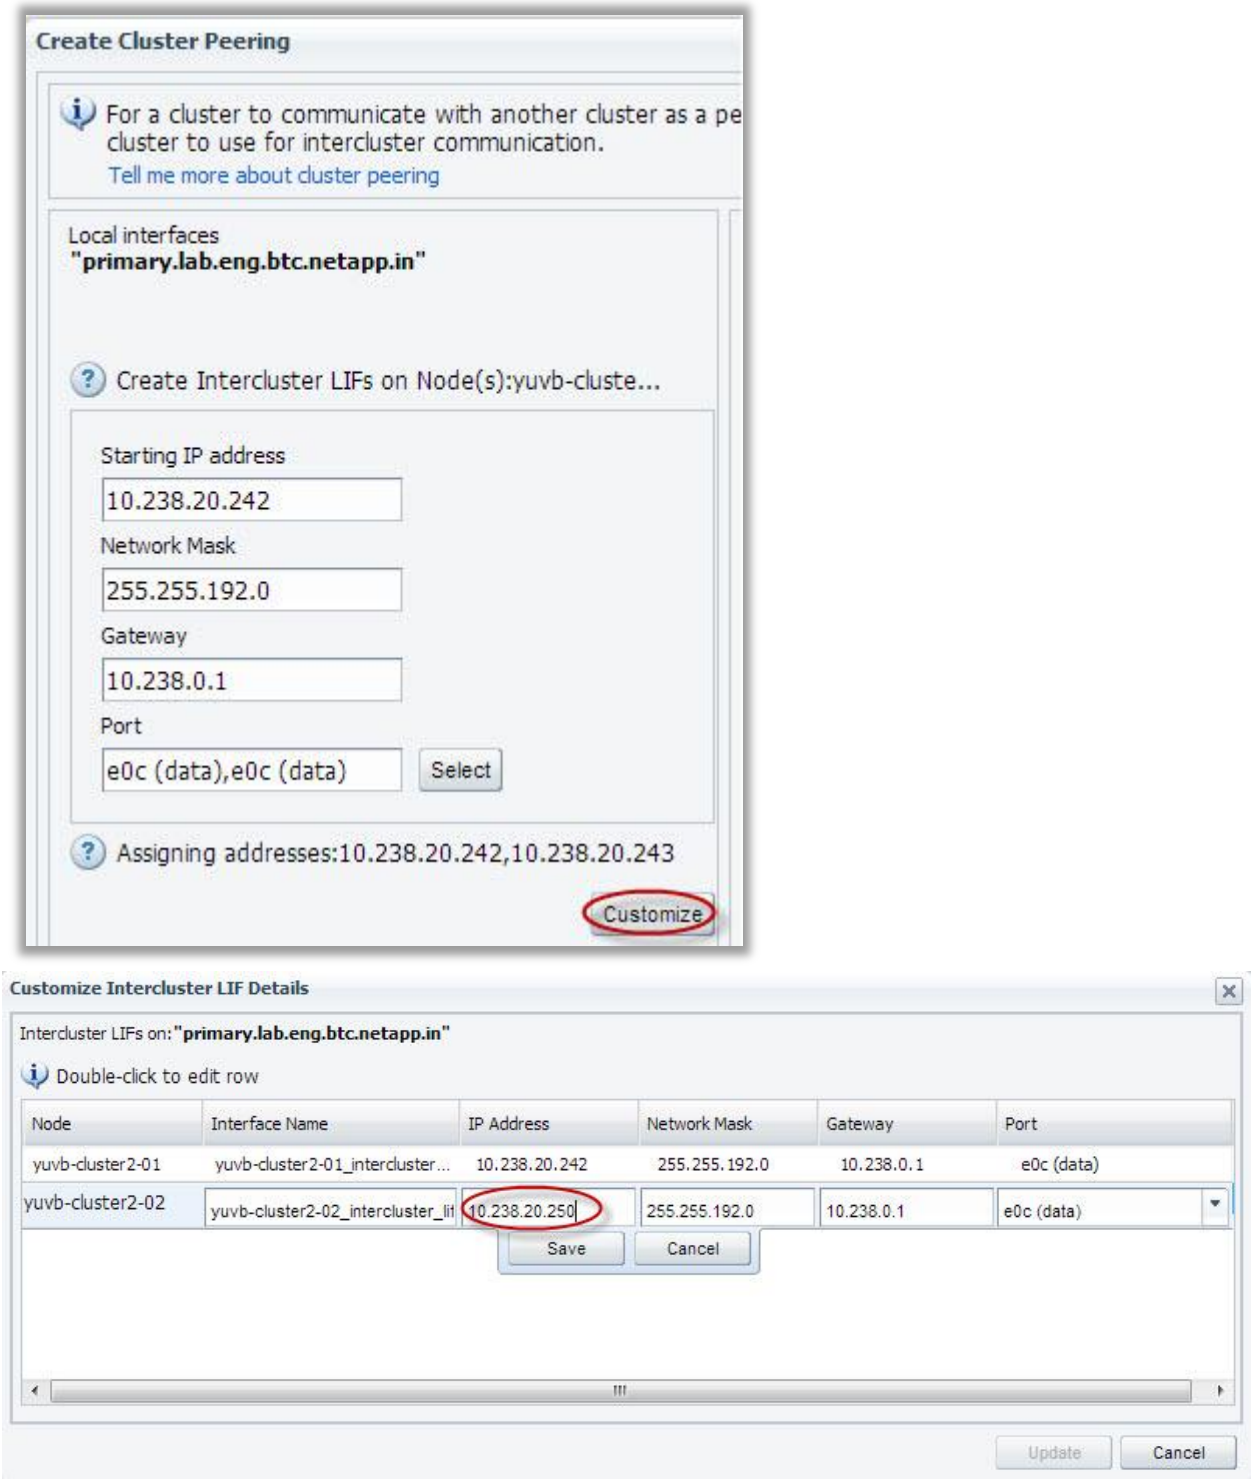

7. In the right pane, select a remote cluster with which you want to create a peer relationship and specify the details for the remote cluster. The list displays remote clusters that are not already peered with the local cluster.

Note: If System Manager is unable to retrieve the credentials of the selected cluster, use the Authenticate link to enter and save the credentials.

8. Click Customize to modify the interface name, IP address, network mask, gateway, and port details for the other nodes, and click Update.
9. Click Create.

Figure 23) Remote cluster.

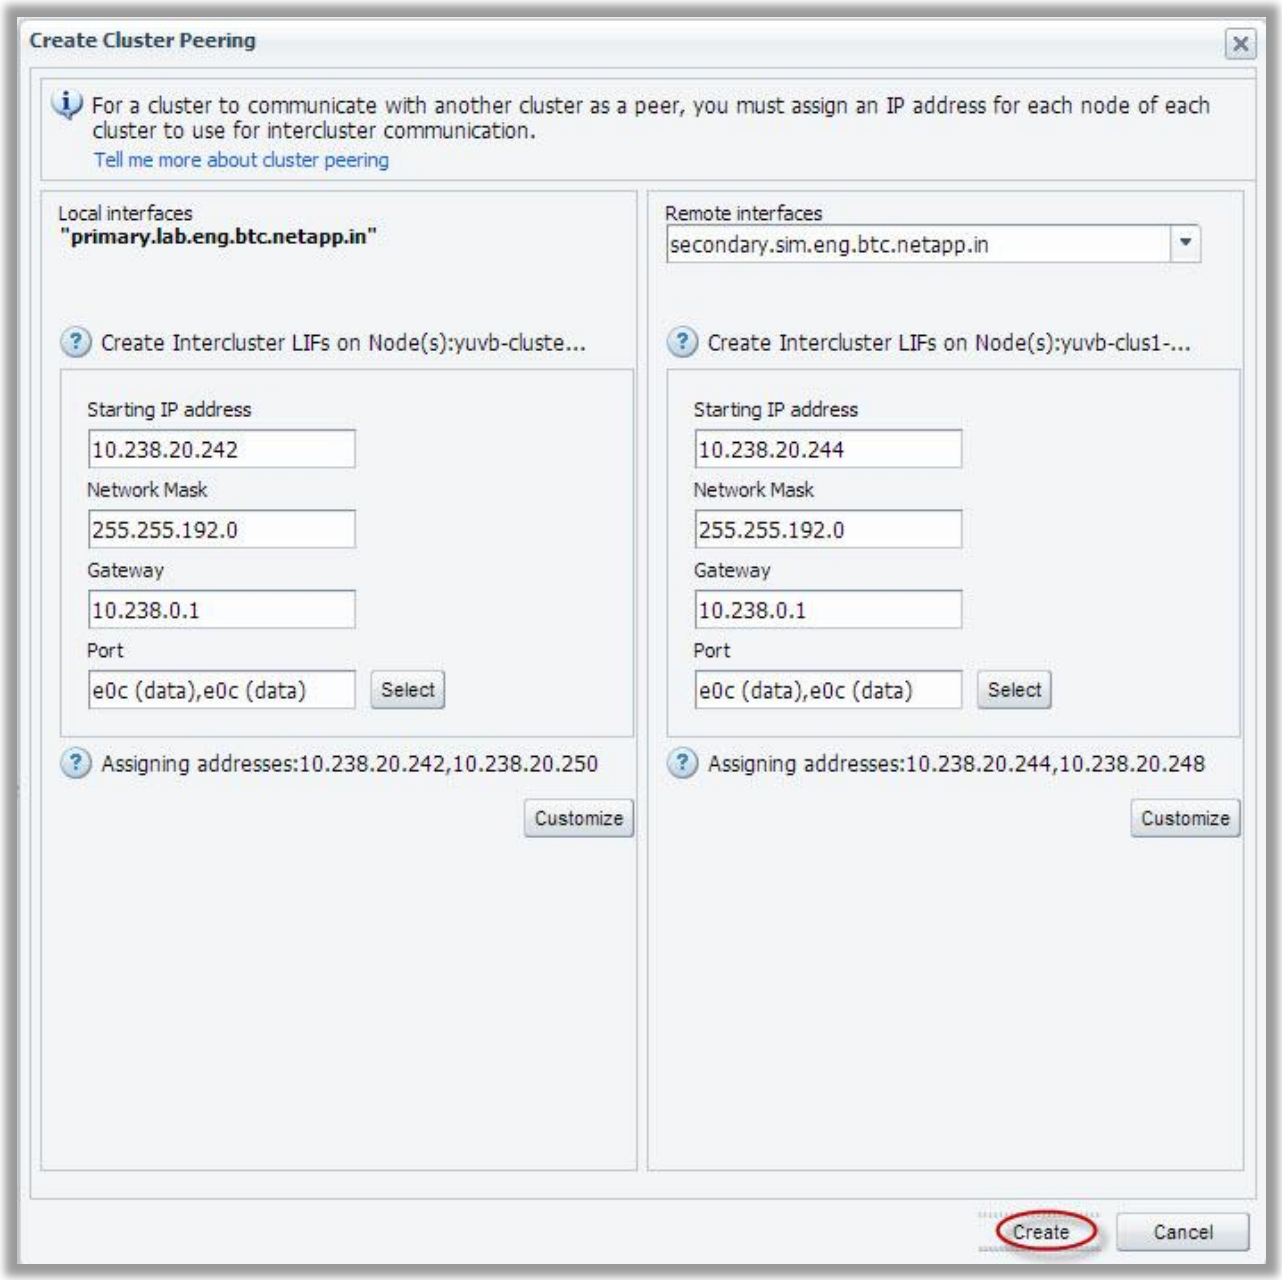

The Details area displays detailed information about the selected cluster peer relationship, including the active IP addresses discovered by the system to set up the intercluster network.

Figure 24) Cluster peers windows.

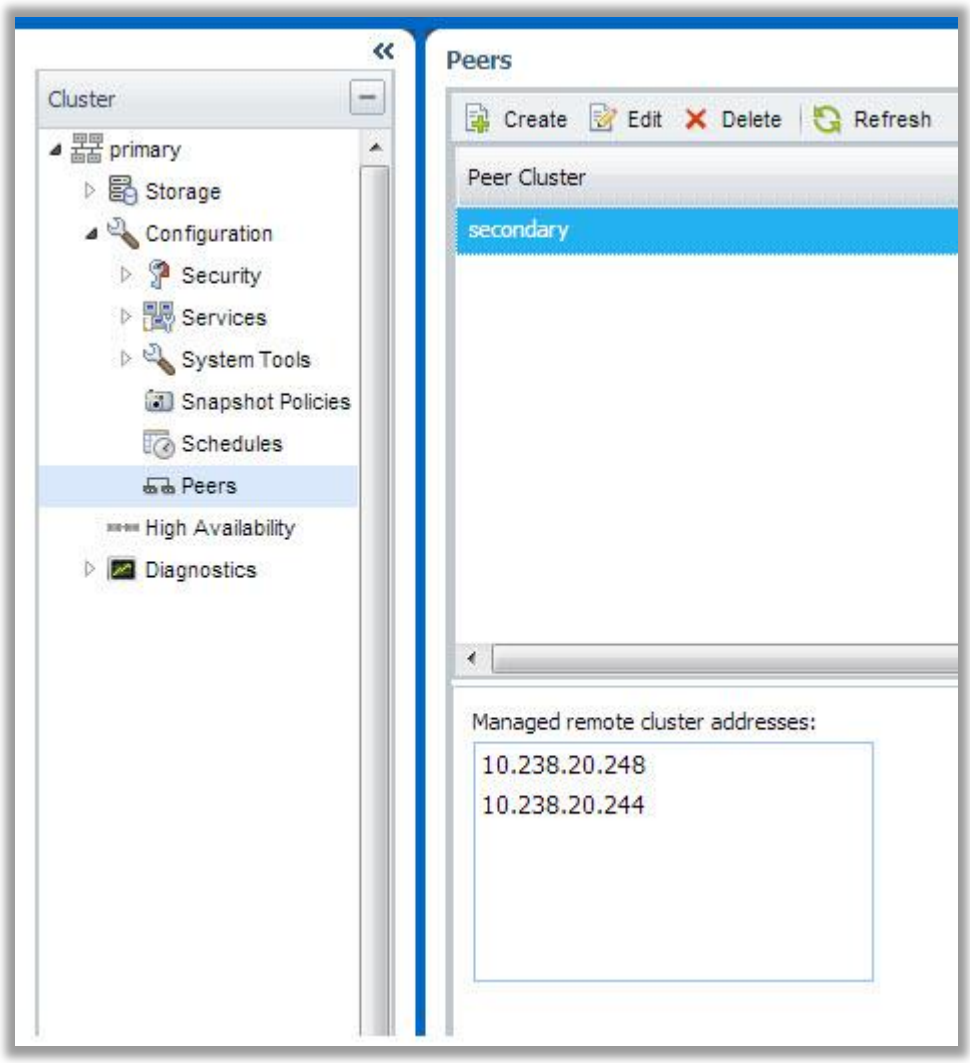

2.5 Enhanced SnapMirror Workflow for Clustered Data ONTAP 8.2.x Clusters

Use case: A storage administrator wants to configure data protection for finance department data. The administrator wants to use mirroring technology to replicate data that is hosted on a clustered Data ONTAP storage system at regular intervals and on demand. The protection requirements indicate that the data must be mirrored every 20 minutes.

With System Manager 3.0 and 3.1, the user interface for SnapMirror relationships and operations is simplified for clustered Data ONTAP 8.2.x. The Protection window provides a simplified user interface that enables you to manage both SnapVault and SnapMirror relationships.

Creating a SnapMirror Relationship

Steps

1. From the homepage, double-click the appropriate storage system.
2. Expand the SVM's hierarchy in the left navigation pane.
3. In the navigation pane, select the SVM and click Protection.
4. In the Protection window, click Create > Mirror.

Figure 25) Protection window.

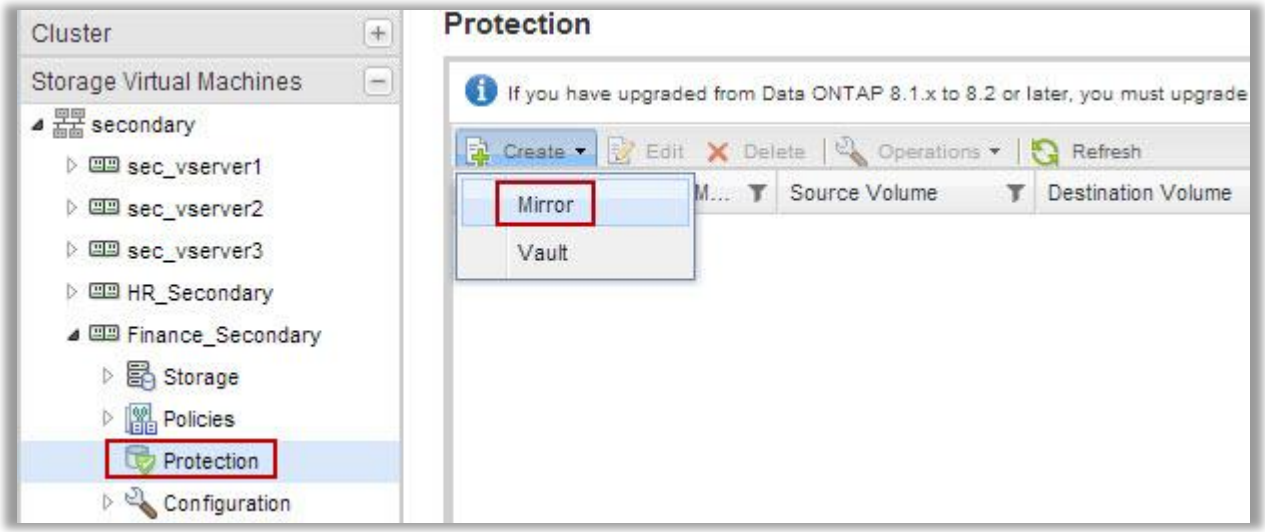

5. In the Create Mirror Relationship from Destination window, specify the source cluster, the SVM, and the source volume.

Note: If System Manager is unable to retrieve the credentials of the selected cluster, use the Authenticate link to enter and save the credentials.

If the source and destination clusters are not in a peer relationship, use the Create Peer link to create a peer relationship between the source and the destination cluster. The Create Peer link is disabled if a cluster peer relationship exists between the source and destination clusters.

Figure 26) Create SnapMirror relationship.

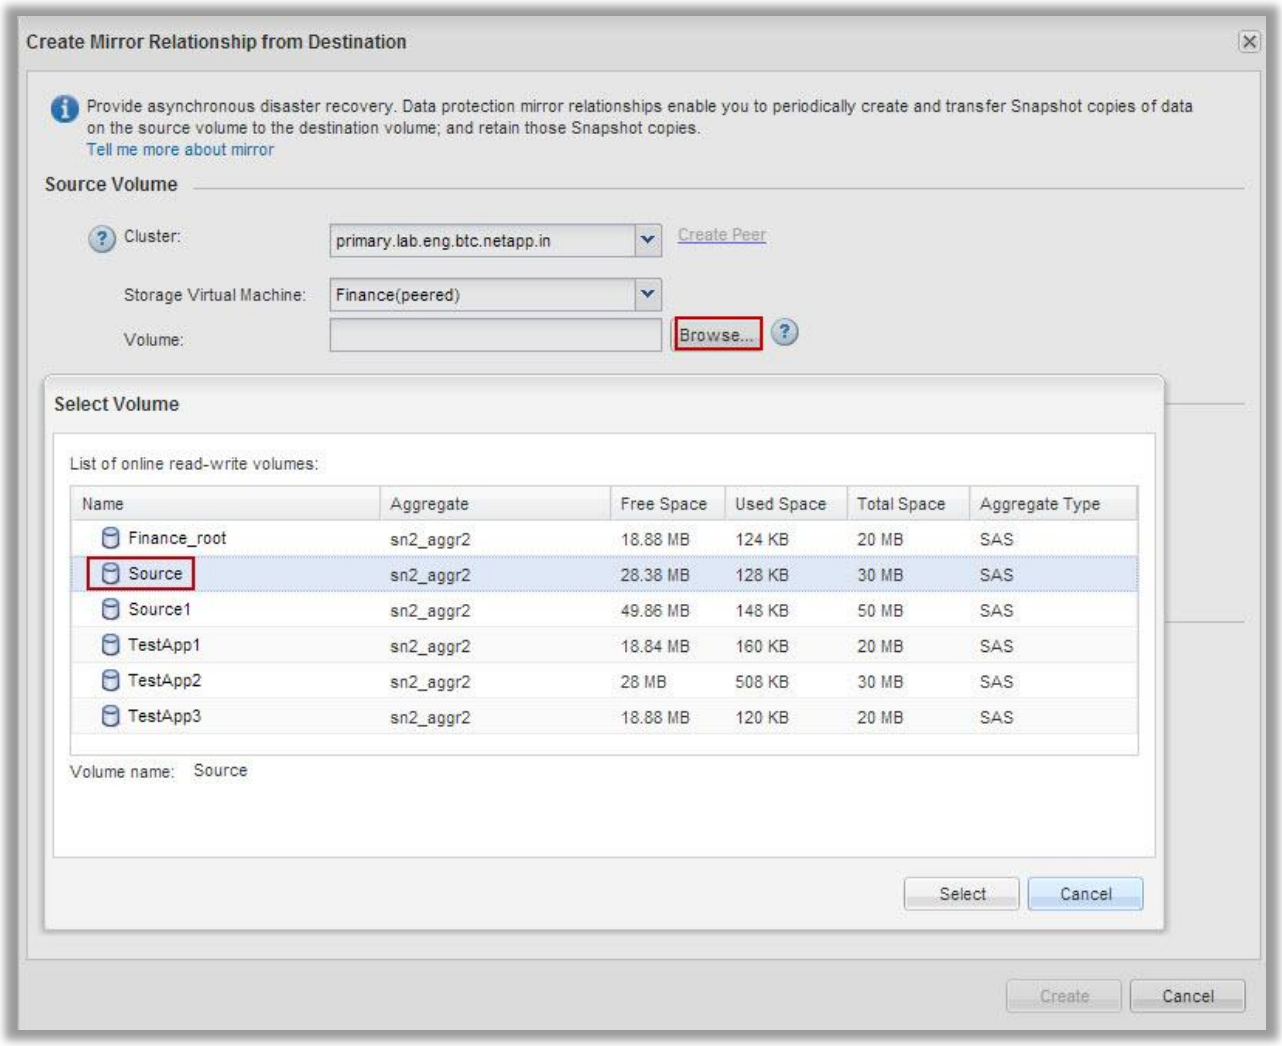

6. Create a new destination volume or select an existing volume of type data protection. The default name is displayed in the format source_SVM_name_source_volume_name_mirror. Specify a new name and select the containing aggregate for the destination volume.

Figure 27) New destination volume.

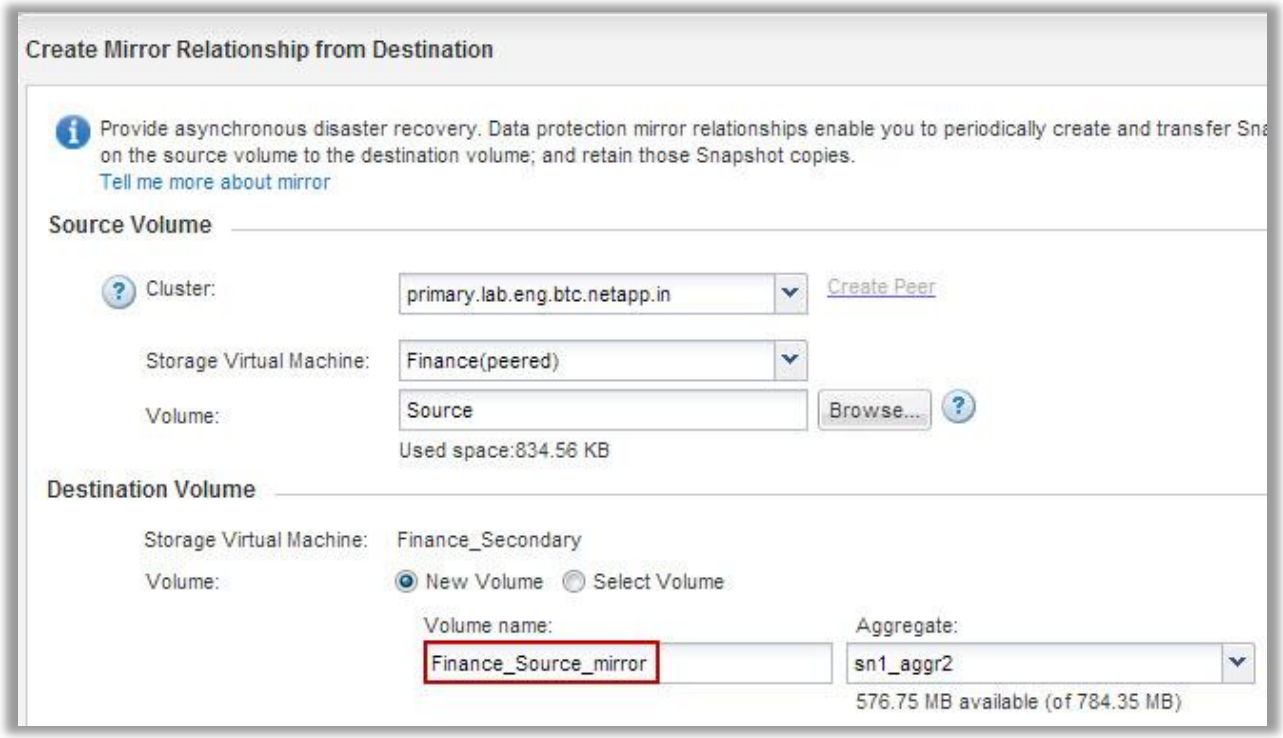

7. Select an existing mirror policy or create a new policy.

Figure 28) SnapMirror policy.

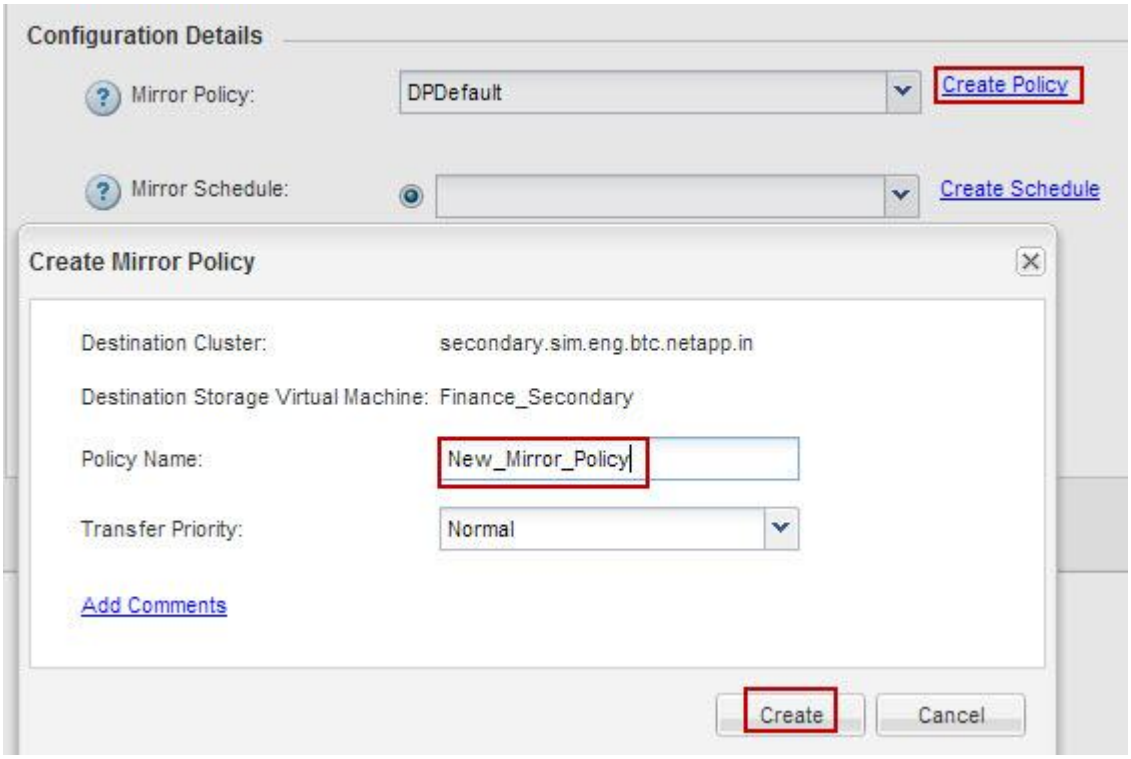

8. Specify a schedule for the relationship. You can either create a new schedule or select an existing schedule from the drop-down list.
9. Select the Initialize Relationship check box to initialize the mirror relationship.
10. Click Create.

Figure 29) Create SnapMirror relationship.

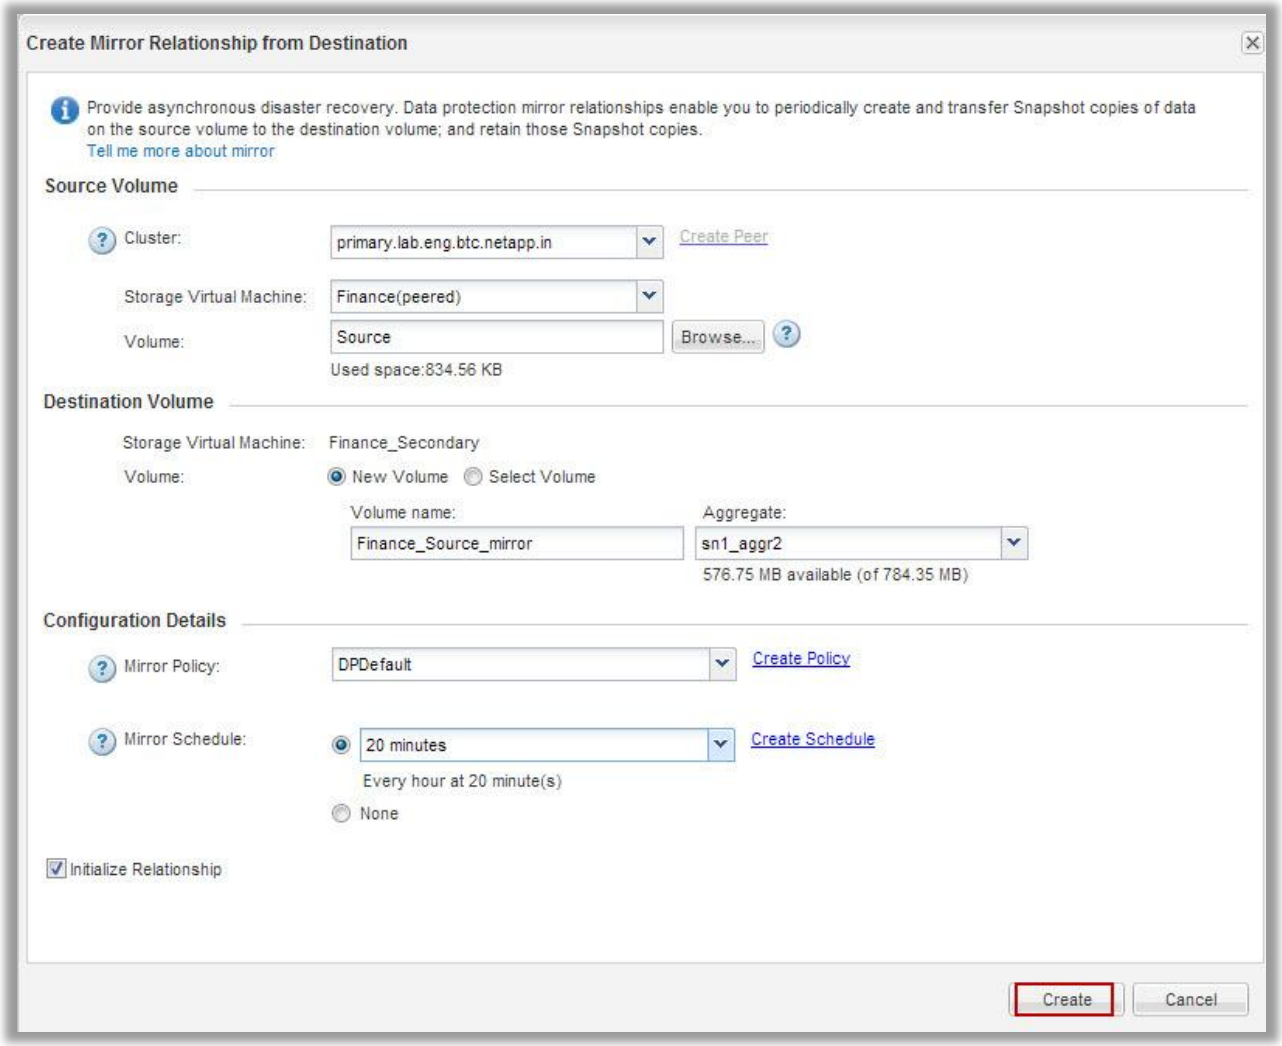

If you choose to create a new destination volume, then a new destination volume of type dp is created with the following default settings: Autogrow is enabled, Compression is disabled, and Language attribute is set to match the language attribute of the source volume.

If the destination volume is on a different SVM compared to the source and if a relationship does not exist, then a peer relationship is created between the two SVMs.

A mirror relationship is created between the source volume and the destination volume. The base Snapshot™ copy is transferred to the destination volume if you have opted to initialize the relationship.

Figure 30) SnapMirror relationship status and summary.

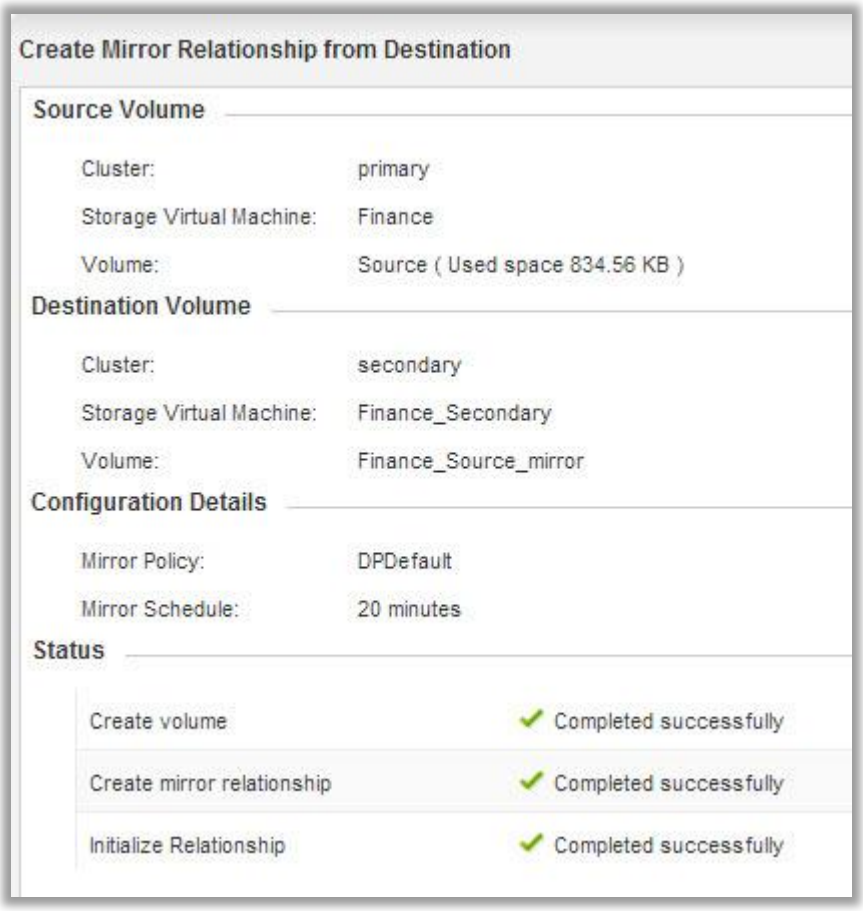

Managing a SnapMirror Relationship

You can use the Protection window to manage SnapMirror relationships.

Figure 31) Protection window.

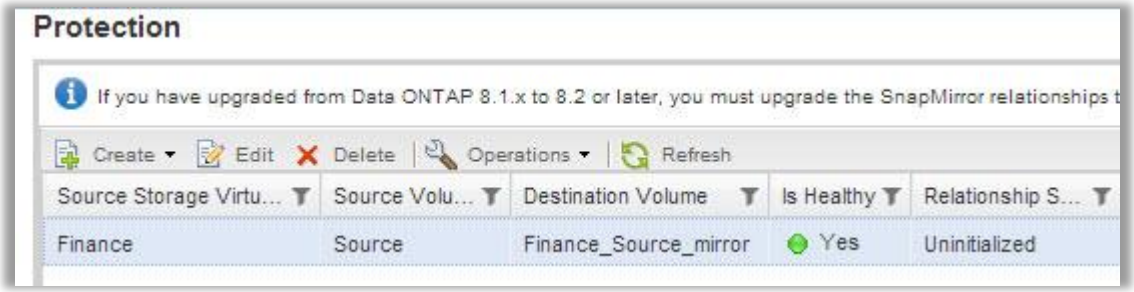

While the transfer is in progress, a user can choose to abort an active transfer by clicking Operations and selecting Abort from the drop-down list. This will abort the current transfer. Note that abort is the only operation a user can perform on an uninitialized mirror relationship.

Figure 32) Abort SnapMirror transfer.

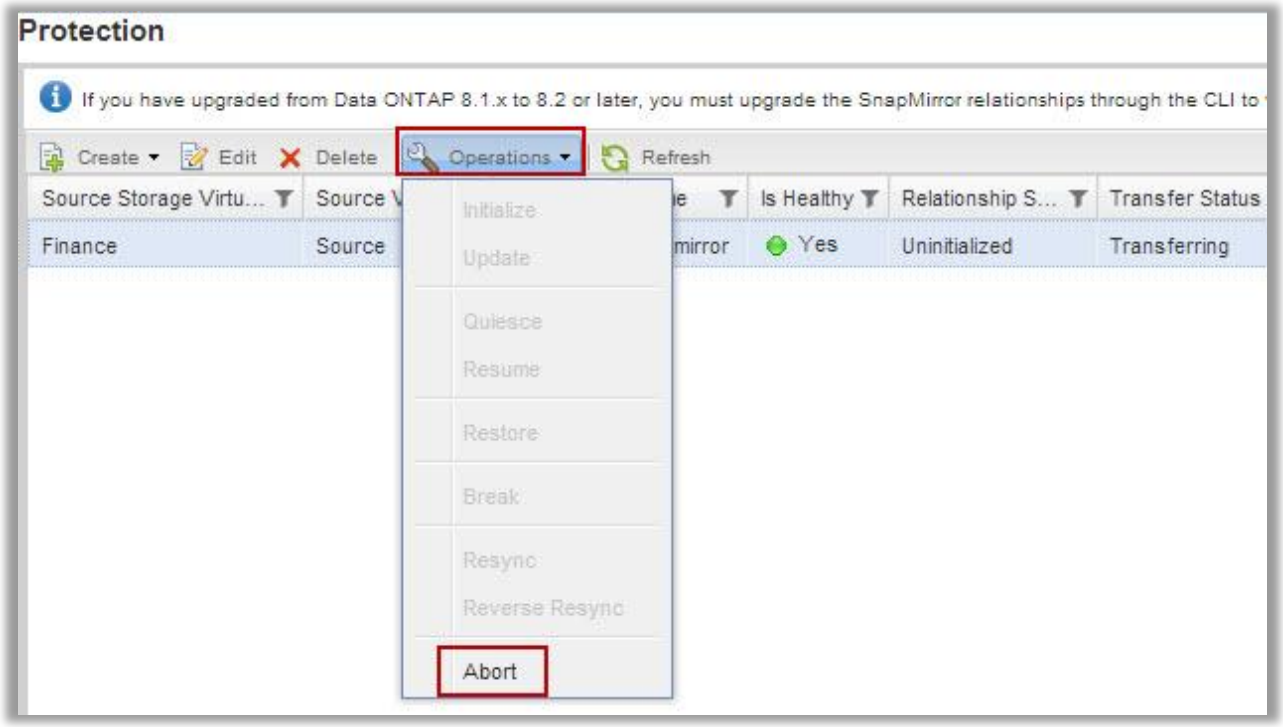

The Details tab in the Protection window provides information about the selected relationship. Information about source and destination SVM, volume, relationship health, relationship state, and transfer status is displayed.

Figure 33) SnapMirror relationship details.

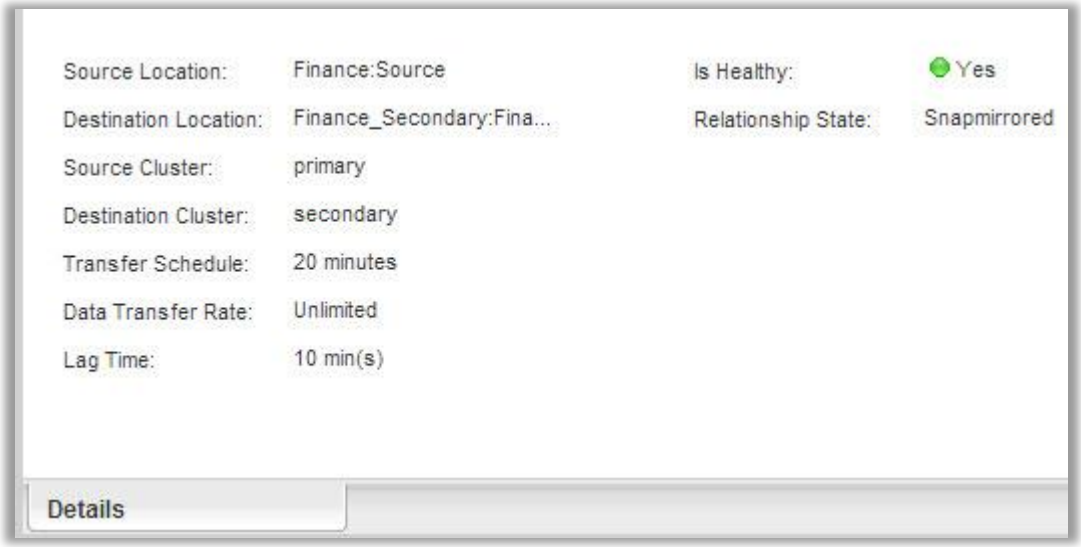

When a relationship is in a SnapMirror state, the operations that can be performed on this relationship are update, quiesce, and break.

Figure 34) SnapMirror state.

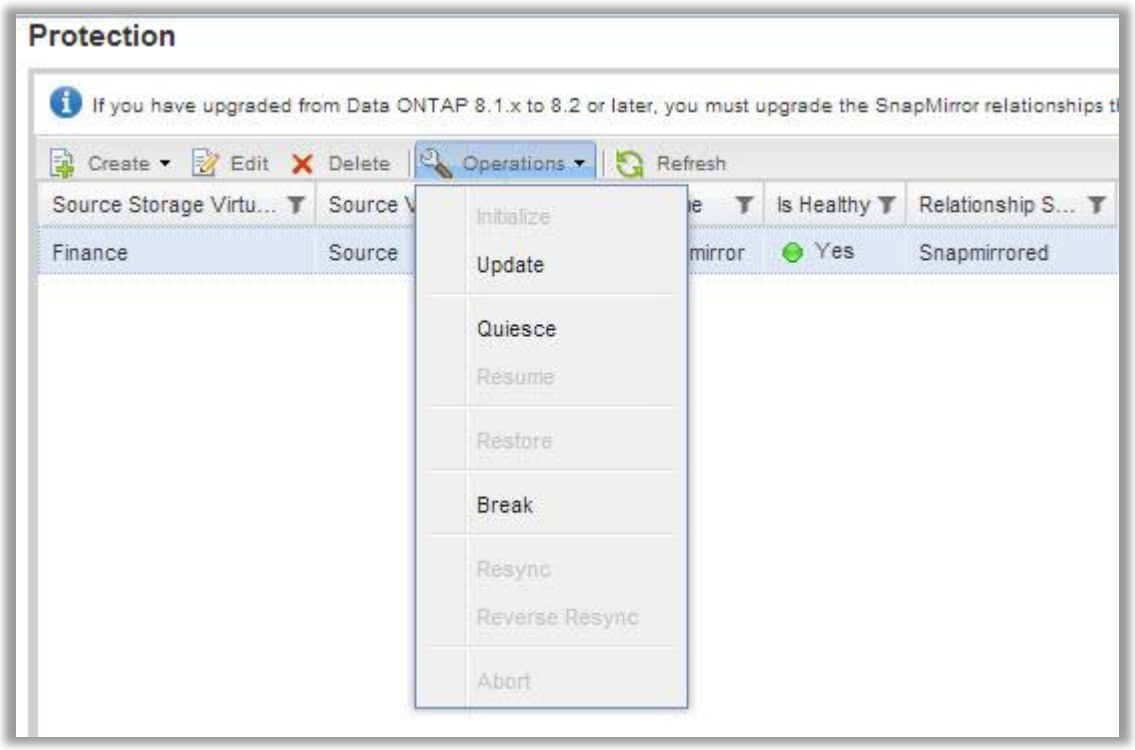

A break operation on a mirror relationship makes a data protection mirror copy destination writable. A user can break a mirror relationship by clicking Operations and selecting Break from the drop-down list. Select the OK to break the selected relationship check box to enable the break option. The relationship status is now displayed as Broken Off.

Figure 35) Break mirror.

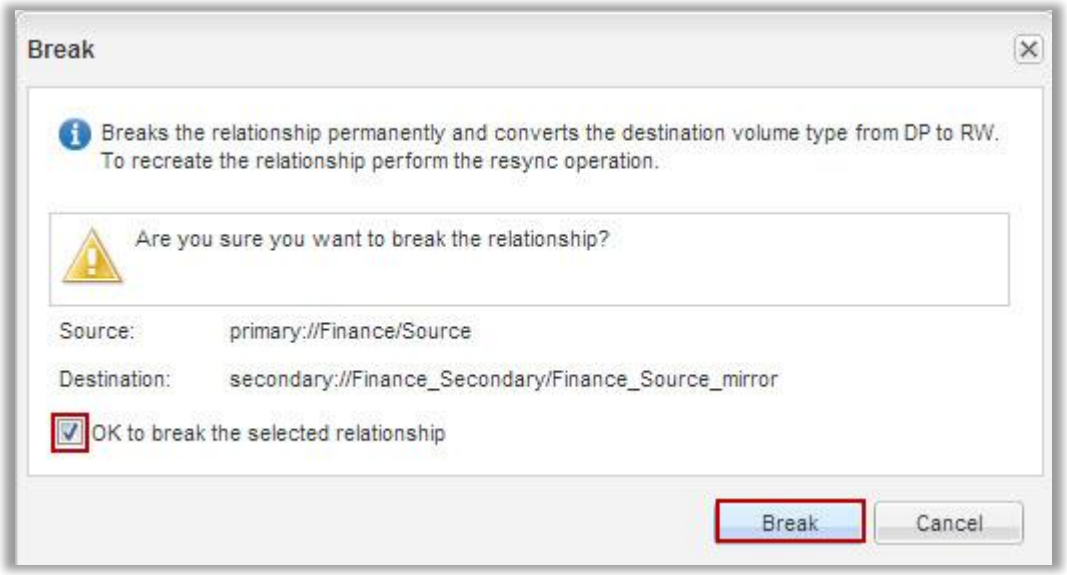

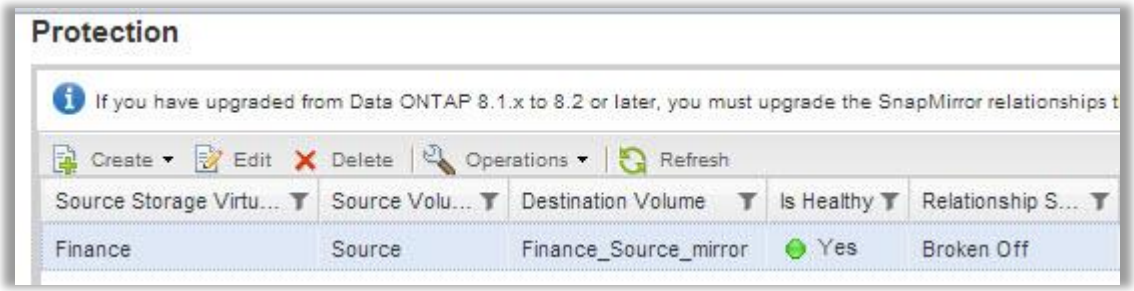

A resynchronize or reverse resynchronize operation can be started on a broken mirror relationship.

Figure 36) Resynchronize SnapMirror relationship.

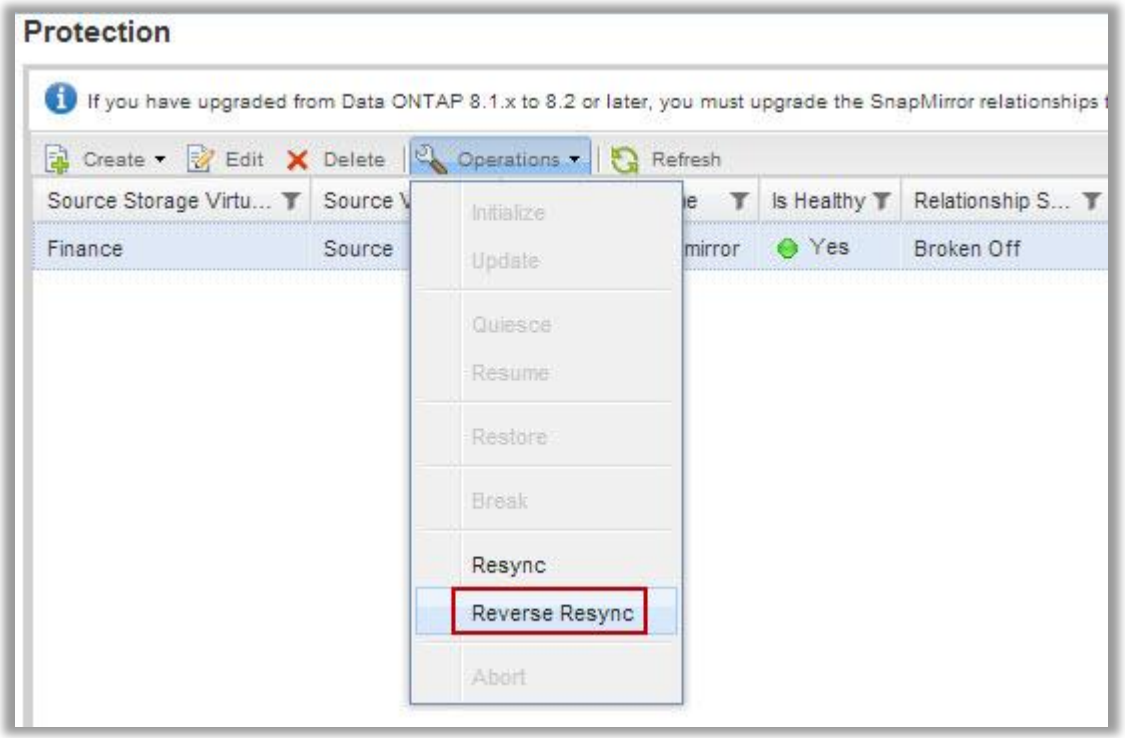

Reverse resynchronizing the relationship deletes the current relationship and converts the source into a destination volume. It also deletes the data in the source volume that is newer than the data in the common Snapshot copy. Selecting the Confirmation check box initiates the reverse resynchronization operation.

Figure 37) Reverse resynchronizing.

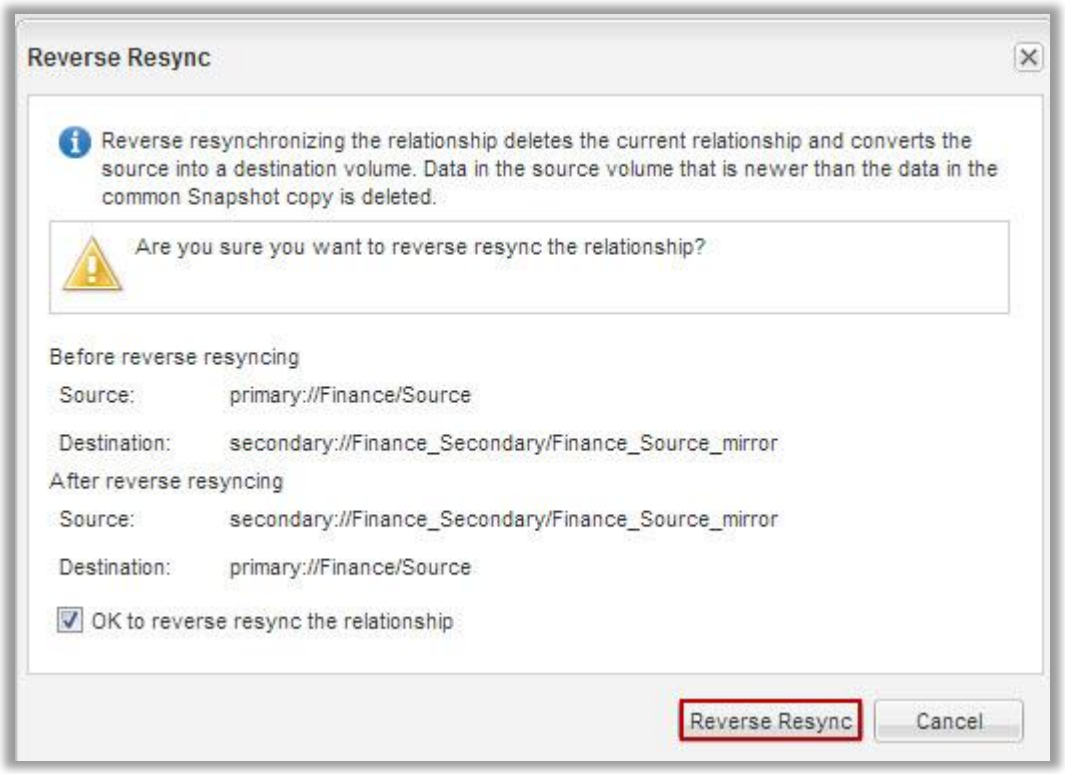

The relationship is no longer displayed under the Protection window in the secondary cluster because this is our new source SVM.

The mirror relationship is listed in the Protection window of the Finance SVM in the primary cluster. The Finance SVM is the new destination SVM, and Finance_Source_data_Protection is our new source SVM as a result of the reverse resynchronization mirror operation.

Figure 38) Protection window on source SVM.

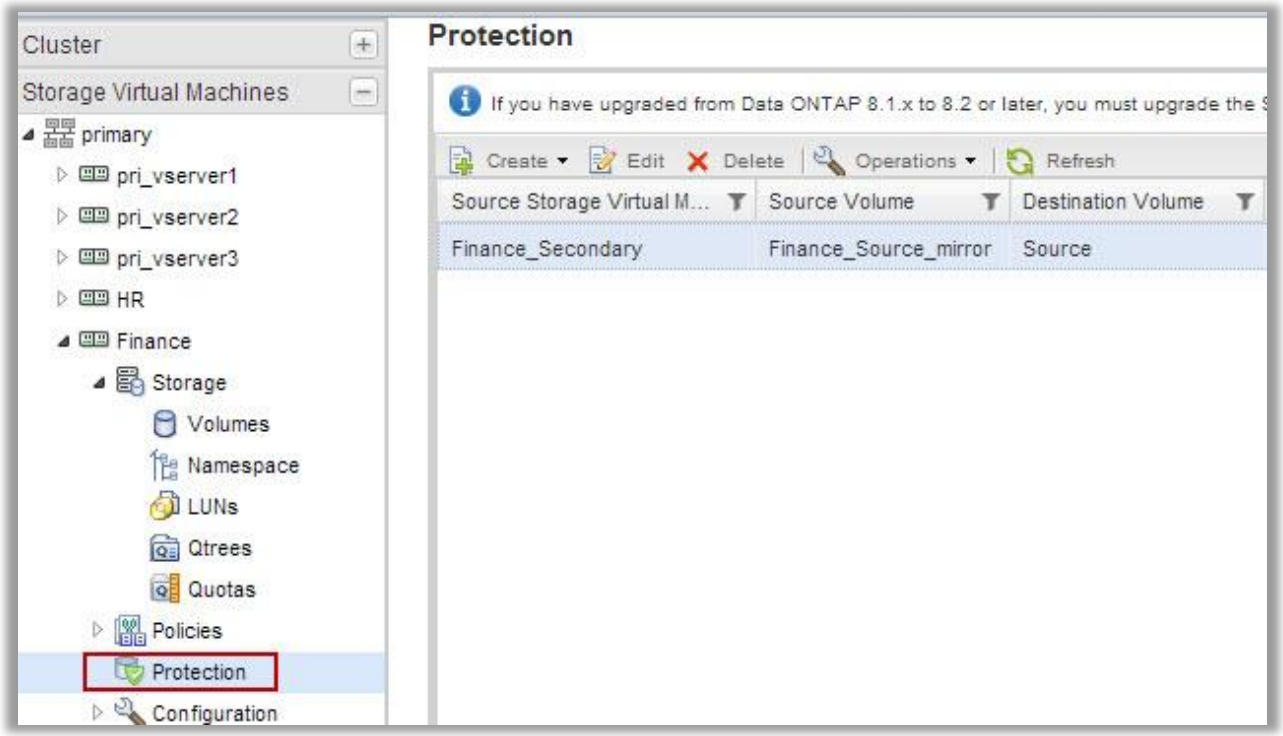

The Delete operation deletes the relationship permanently and releases the base Snapshot copy related to the relationship on the source volume.

Figure 39) Delete SnapMirror relationship.

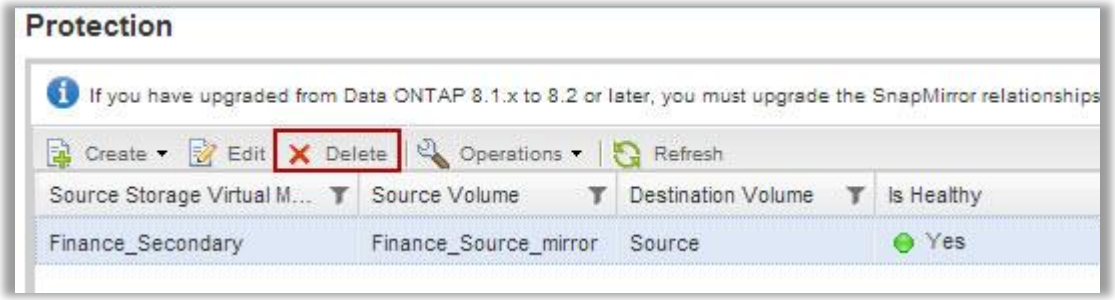

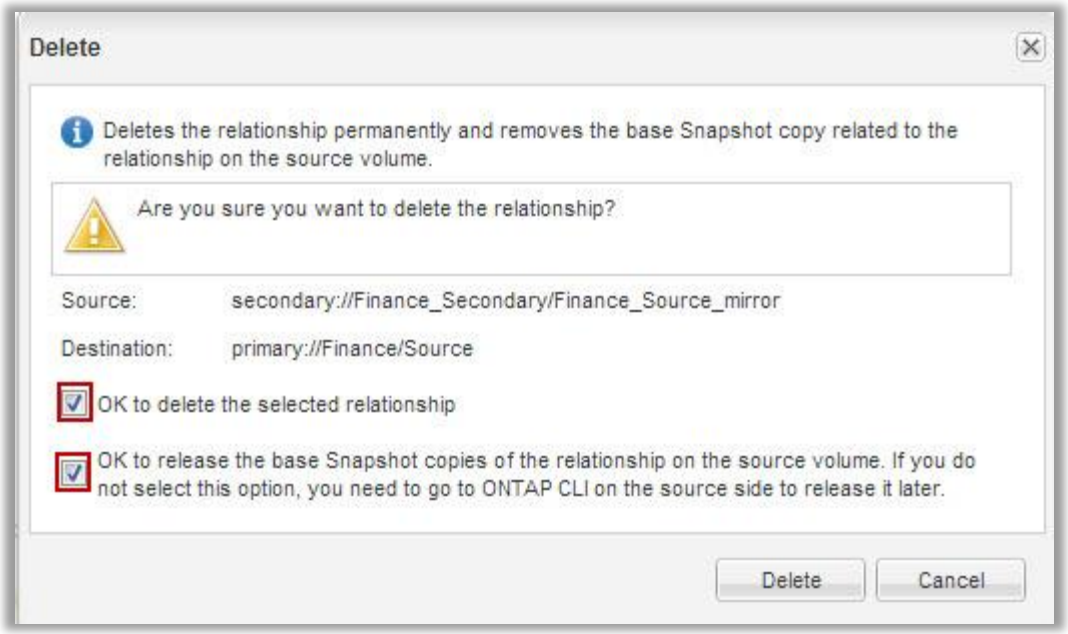

2.6 Simplified SnapVault Workflow for Clustered Data ONTAP 8.2.x Clusters

Use case: A storage administrator wants to configure daily backups of HR business unit data that is hosted in clustered Data ONTAP 8.2.x. The data protection requirements state that the data must be backed up daily at midnight.

With System Manager 3.0 and 3.1, the user interface for creating and managing SnapVault relationships for clustered Data ONTAP 8.2.x is simplified. The Protection window provides a simplified user interface that enables you to manage both SnapVault and SnapMirror relationships.

Creating a SnapVault Relationship

Steps

1. From the homepage, double-click the appropriate storage system.
2. Expand the SVM's hierarchy in the left navigation pane.
3. In the navigation pane, select the SVM and click Protection.
4. In the Protection window, click Create > Vault.

Figure 40) Protection window.

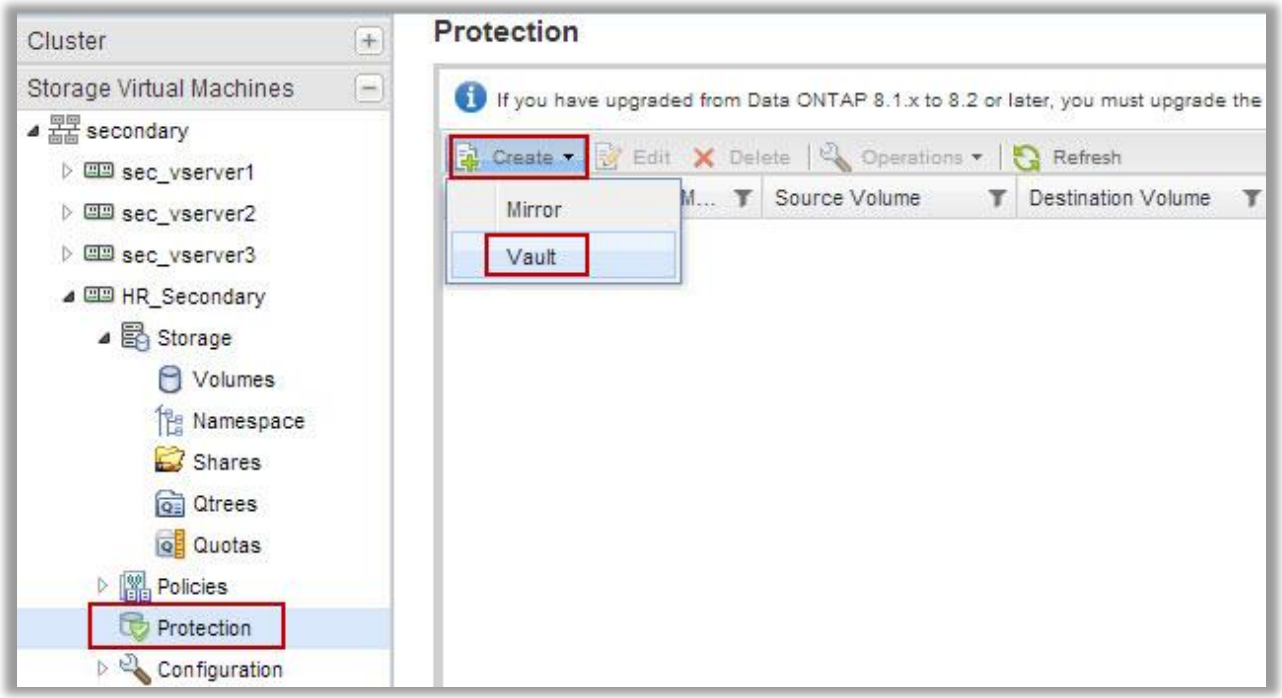

5. In the Create Vault Relationship from Destination window, specify the cluster, the SVM, and the source volume.

Note: If System Manager is unable to retrieve the credentials of the selected cluster, use the Authenticate link to enter and save the credentials.

If the source and destination clusters are not in a peer relationship, use the Create Peer link to create a peer relationship between the source and the destination cluster. Create Peer link is disabled if a cluster peer relationship exists between the source and the destination clusters.

Figure 41) Create SnapVault relationship.

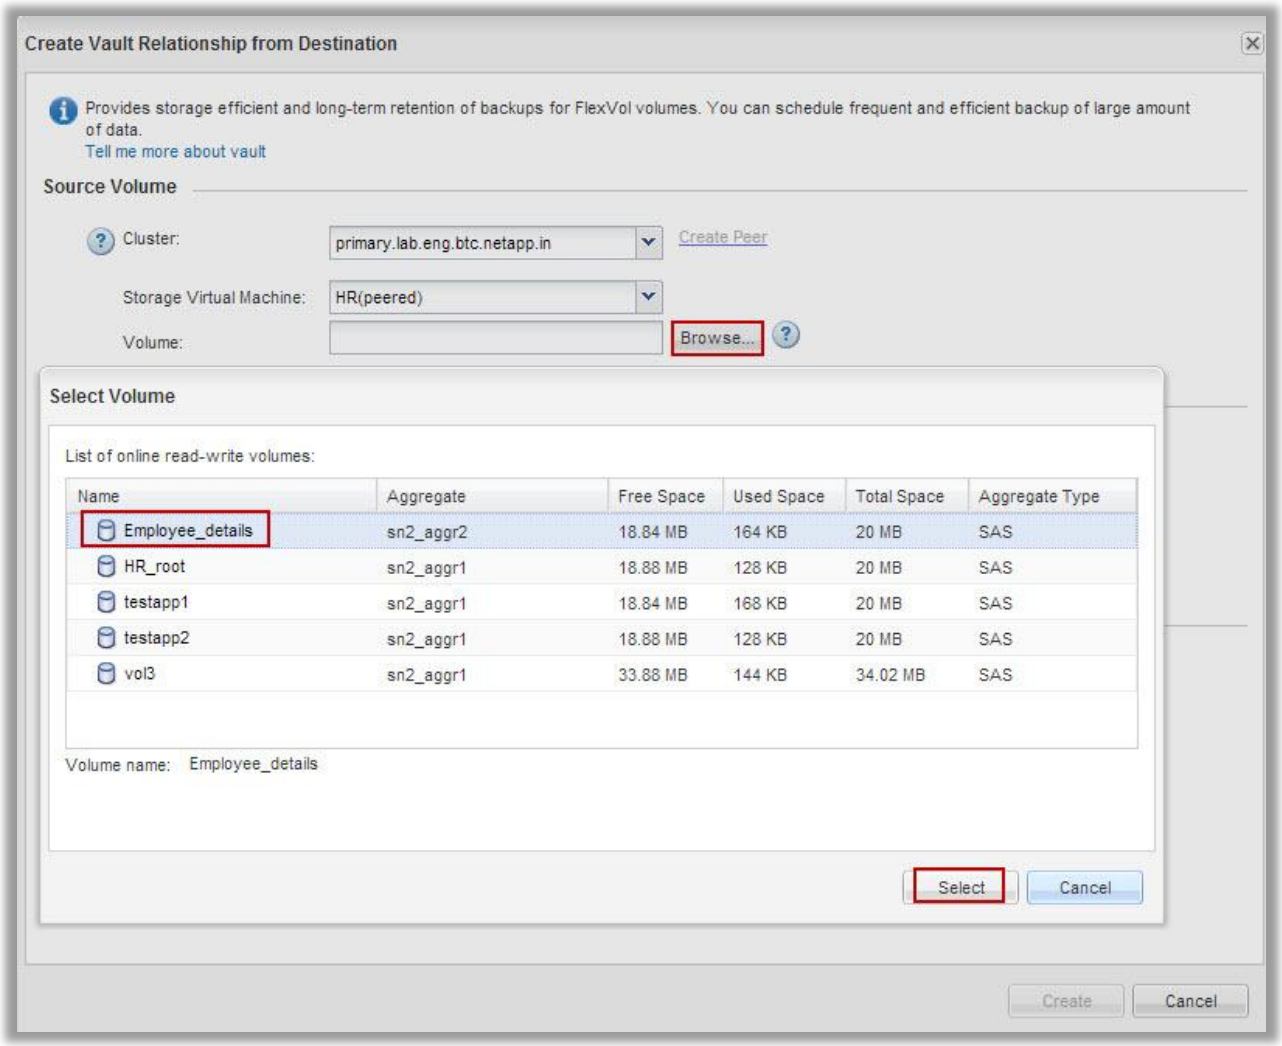

6. Create a new destination volume or select an existing volume.

Note: The default name is displayed in the format primary_SVM_name_source_volume_name_vault. Specify a new name and select the containing aggregate for the destination volume.

Figure 42) Create new secondary volume.

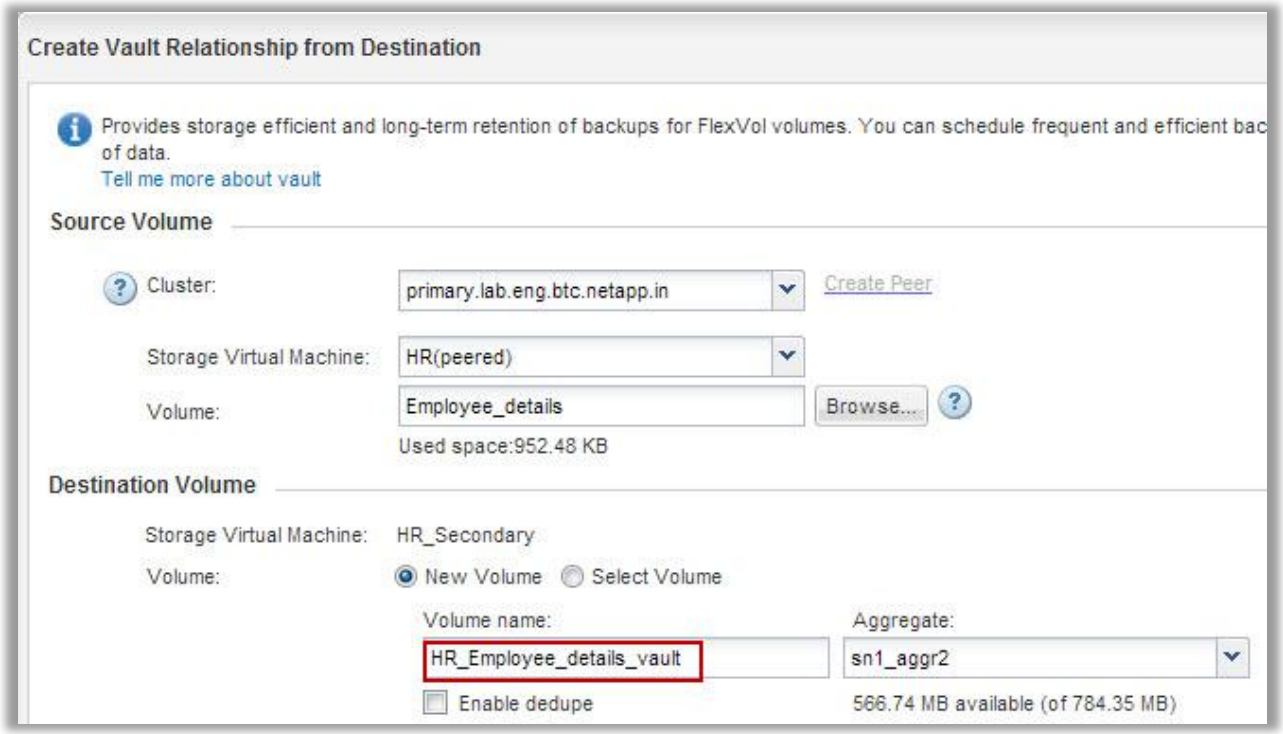

7. Select an existing policy or create a new policy.

Note: Clustered Data ONTAP 8.2 has a predefined XDPDefault SnapVault policy. Assigning this policy to a SnapVault relationship enables a user to back up the primary volume on a daily or weekly basis. In clustered Data ONTAP 8.2, labels are the basis on which a Snapshot copy is picked up for SnapVault transfers. XDPDefault SnapVault policy has daily and weekly Snapshot copy labels in place with a retention count of 52 daily and 7 weekly.

Figure 43) SnapVault policy.

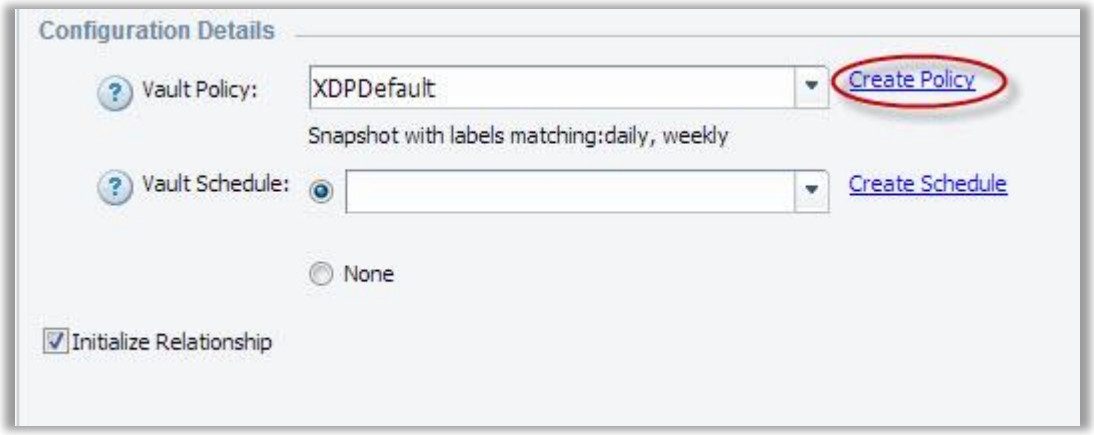

8. To create a new SnapVault policy, click Create Policy and specify a name for the policy.
9. Set the priority for scheduled transfers. Low indicates that the transfer has the least priority and is usually scheduled after normal priority transfers. By default, the priority is set to Normal.

10. System Manager populates the source Snapshot policy details when creating a SnapVault policy. Information about the source Snapshot copy label, the schedule, and the retention count is provided; it can be used as a reference when creating the SnapVault policy.

Figure 44) New SnapVault policy.

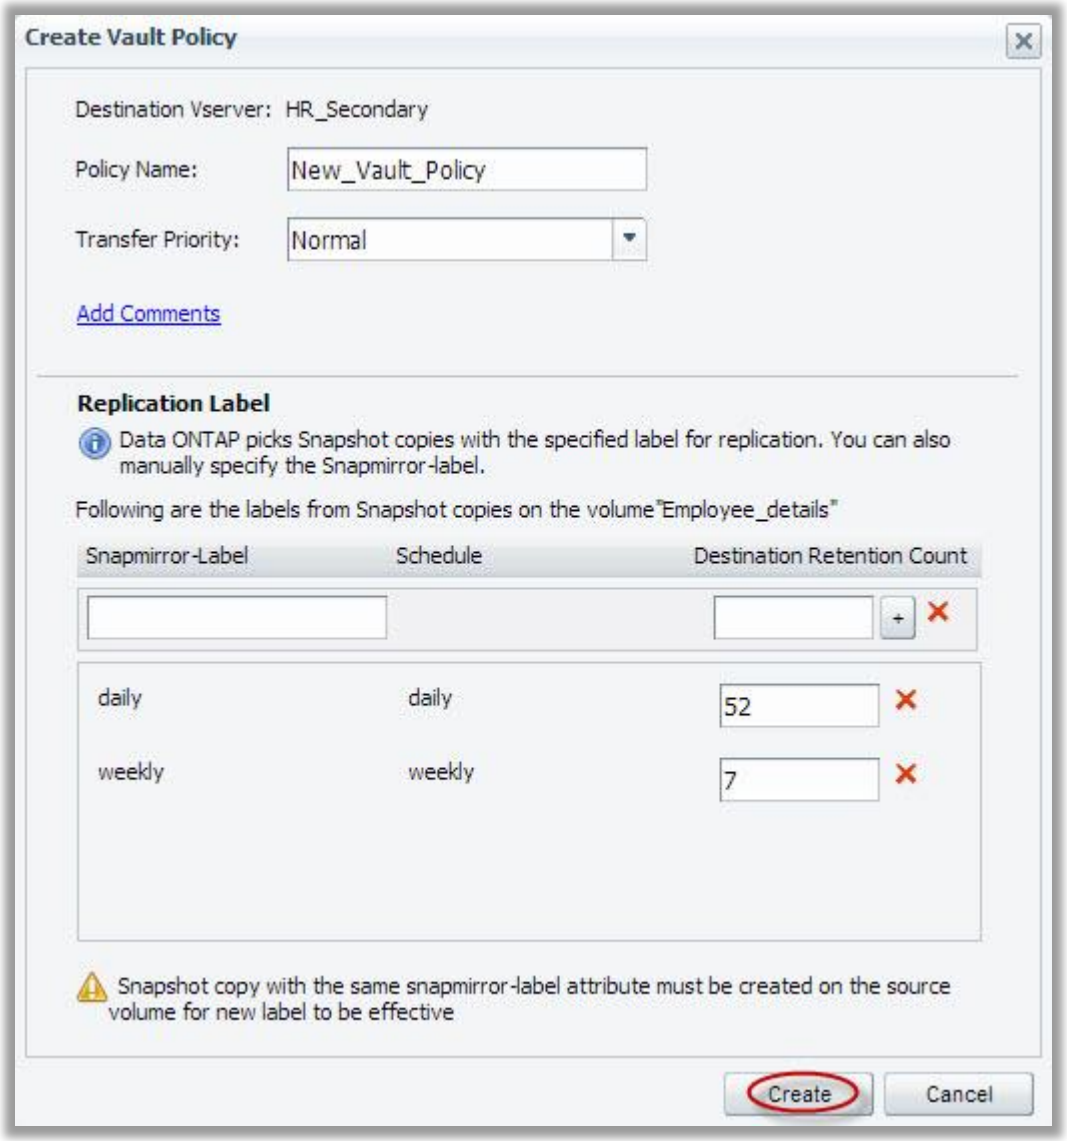

11. If the user's backup requirements require retention of 52 daily and 7 weekly backups, the destination retention count can be modified accordingly.
12. Select the Initialize Relationship check box to initialize the SnapVault relationship.
13. Click Create.

Figure 45) Create SnapVault relationship.

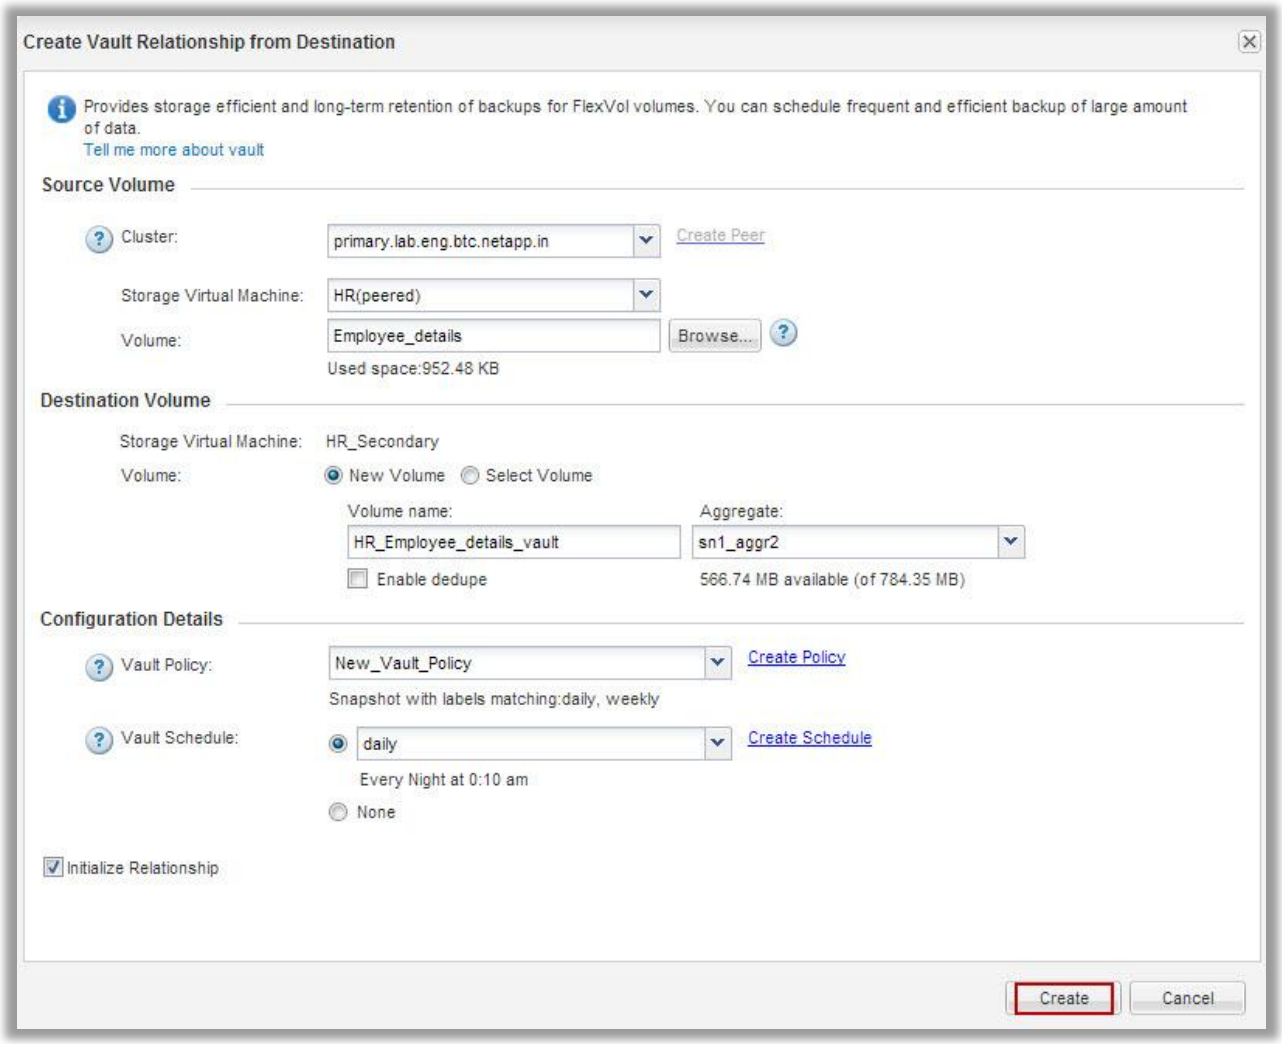

If the destination volume is on a different SVM compared to the source and if a relationship does not exist, then a peer relationship is created between the two SVMs.

A SnapVault relationship is created between the destination volume and the source volume. The base Snapshot copy is transferred to the destination volume if you have opted to initialize the relationship.

Figure 46) SnapVault relationship status and summary.

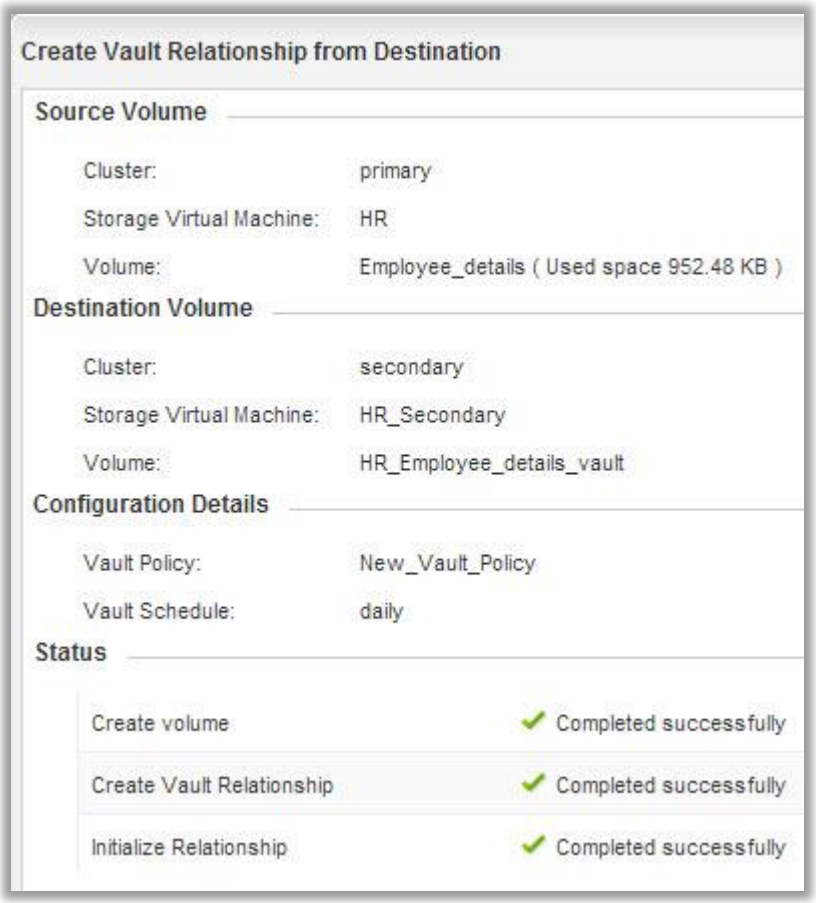

Managing a SnapVault Relationship

The Protection window on the destination SVM can be used to create and manage mirror and vault relationships and to display details about these relationships.

Figure 47) Protection window.

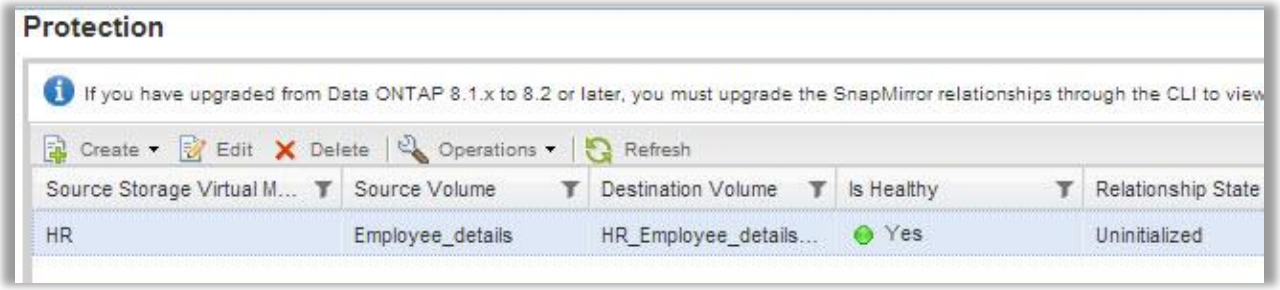

The Details tab provides information about source SVM, source volume, destination volume, health, relationship state, and transfer status.

Figure 48) SnapVault transfer details.

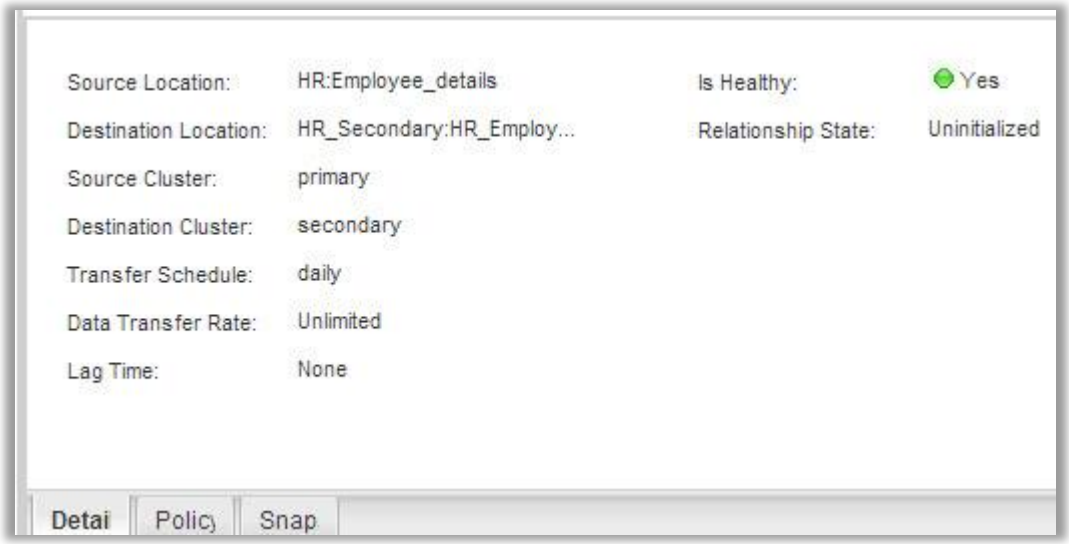

Policy details can be viewed from the Policy Details tab. Policy details also provide information on whether the source has matching Snapshot copy labels.

Figure 49) Policy details.

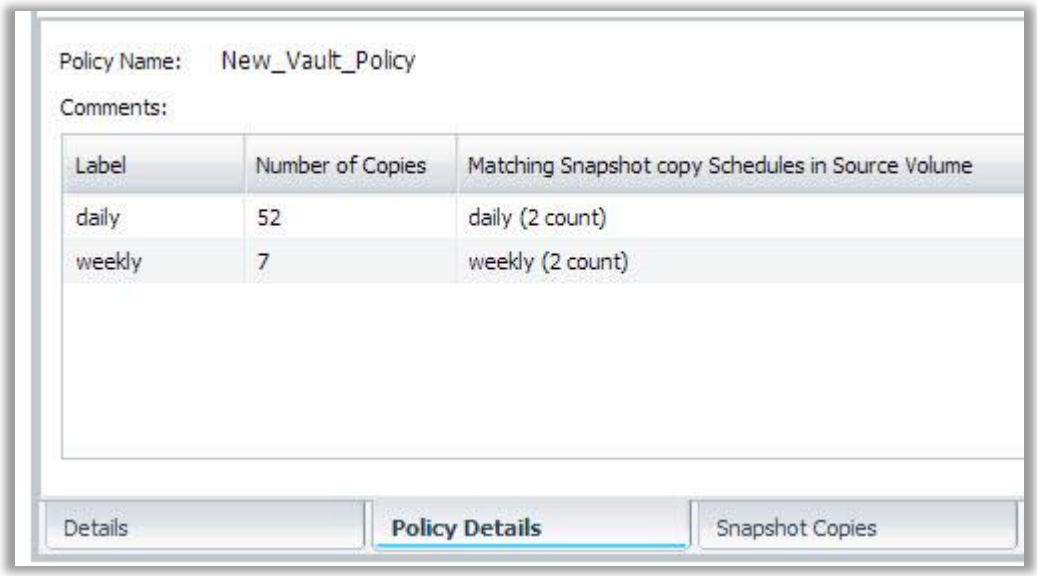

The operations that can be performed on a SnapVault relationship in a SnapMirror state are update, quiesce, and restore.

Figure 50) SnapVault operations.

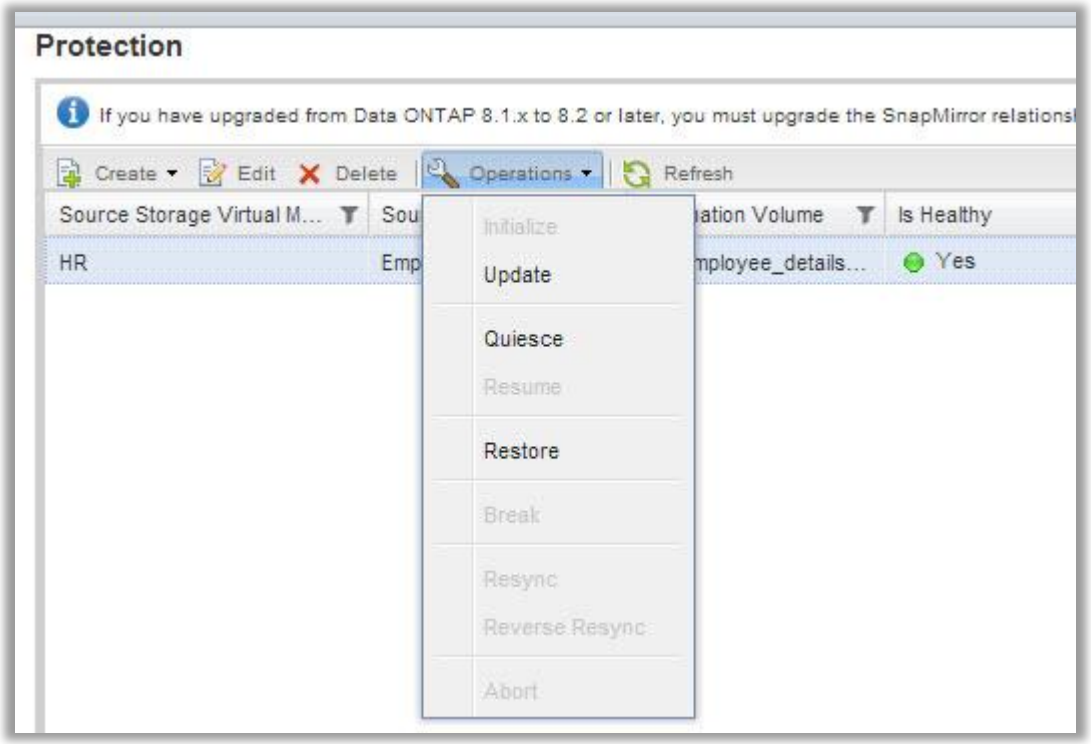

In a SnapVault restore operation, data can be restored to the source or other volumes in the SnapVault relationship if the source data is corrupted and is no longer usable.

Figure 51) Restore to primary volume.

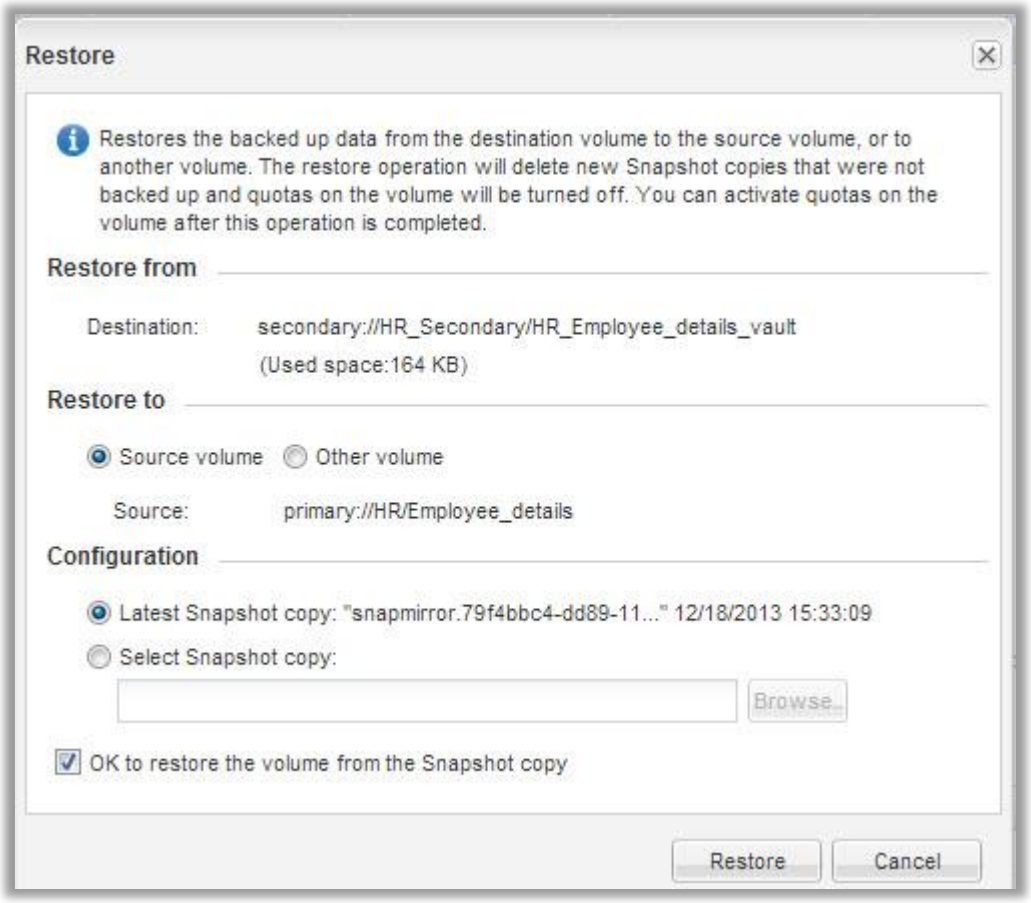

If you choose to restore the data to a volume other than the source volume, create a new volume. Specify the cluster, SVM, and volume name and select the containing aggregate for the volume. In the Configuration section, select the latest Snapshot copy or select a specific Snapshot copy that you want to restore. Click Restore.

Figure 52) Restore to new volume.

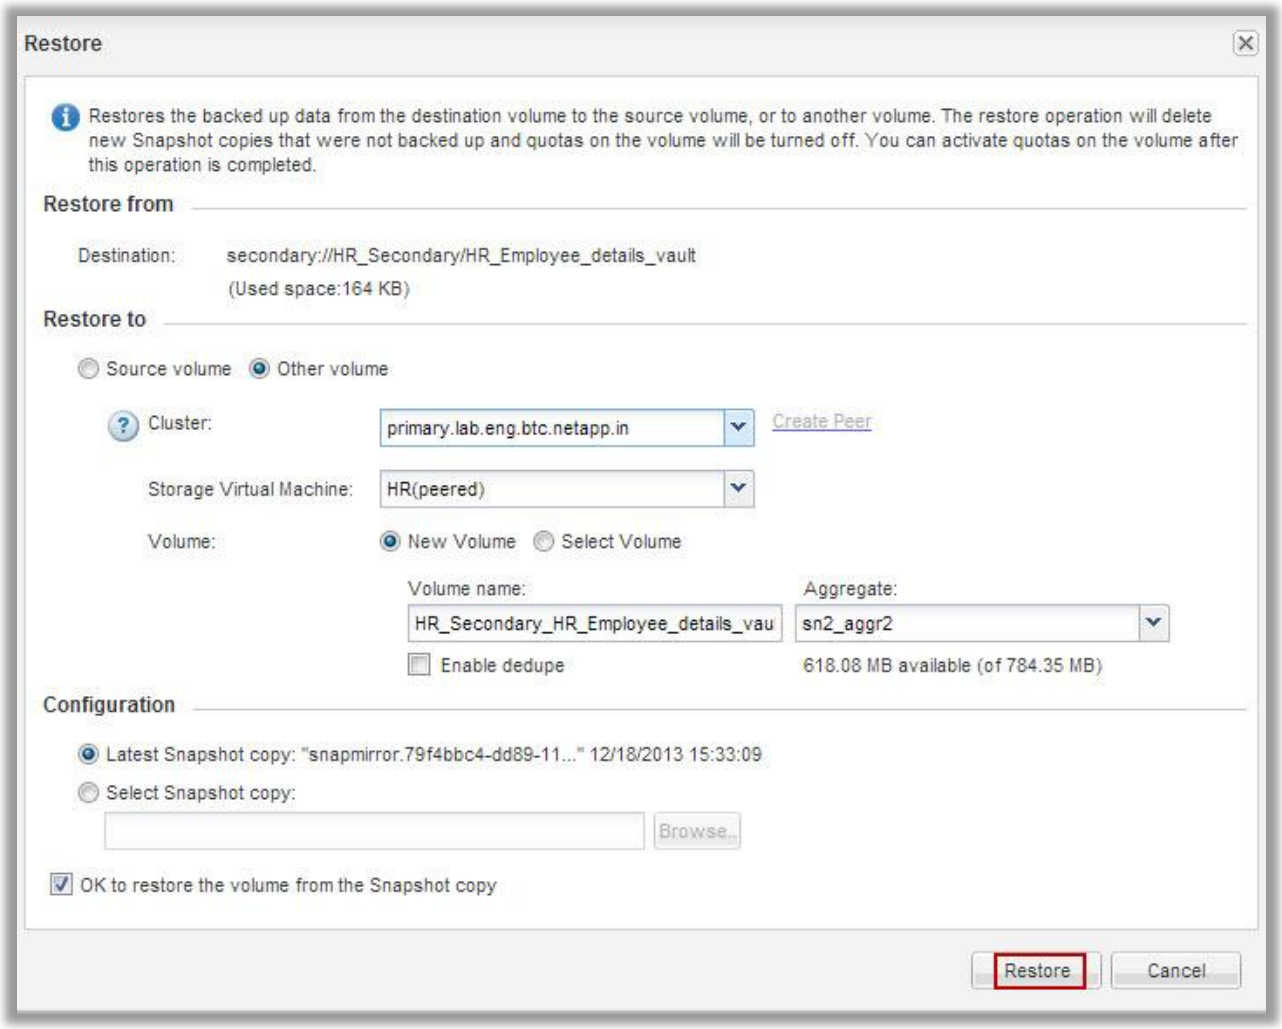

2.7 Applying Clustered Data ONTAP 8.2.x Licensing

Data ONTAP feature licenses are issued as packages, each of which contains multiple features or a single feature. A package requires a license key; installing the key enables you to access all features in the package.

Starting with Data ONTAP 8.2.x, all license keys are 28 characters in length.

If your storage system software was installed at the factory, System Manager automatically adds the software to its list of licenses. If the software was not installed at the factory or if you want to add additional software licenses, you can add the software license through the Add Licenses dialog box.

Steps

1. From the homepage, double-click the appropriate storage system.
2. Expand the cluster hierarchy in the left navigation pane.

3. In the navigation pane, click Configuration > System Tools > Licenses.
4. In the Licenses window, click Add.

Figure 53) License window.

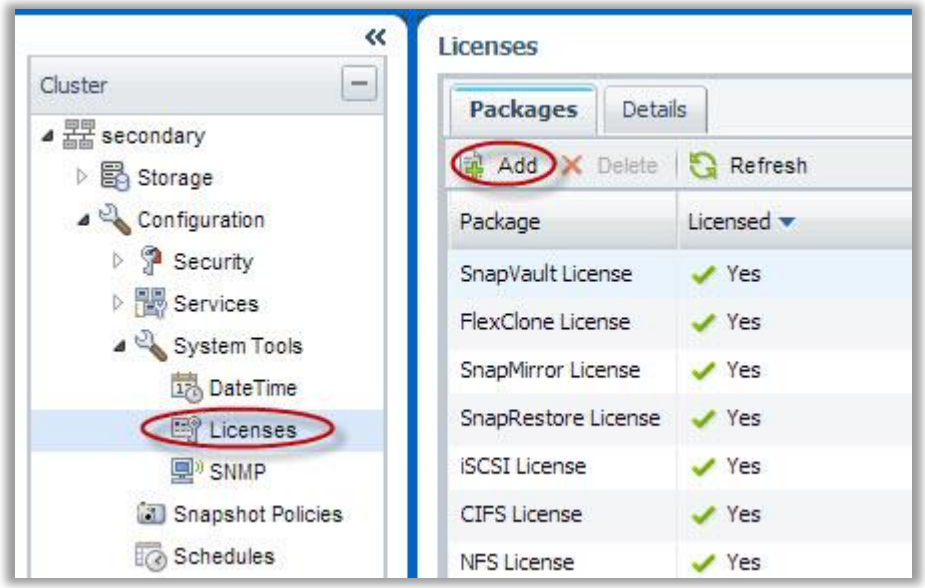

5. In the Add Licenses dialog box, enter the software license key and click Add.
6. For storage systems running Data ONTAP 8.2, you can add multiple licenses by entering the software license keys, separated by commas.

Figure 54) Add licenses.

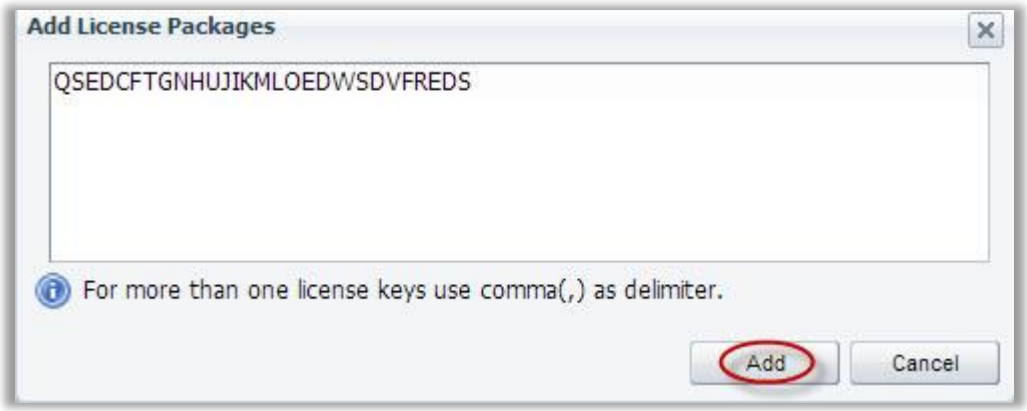

7. The new license is added.
8. For storage systems running Data ONTAP 8.2, the Add licenses status window displays the list of licenses that were added successfully. The window also displays the license keys of the licenses that were not added and the reason why that happened.

Figure 55) Add licenses status.

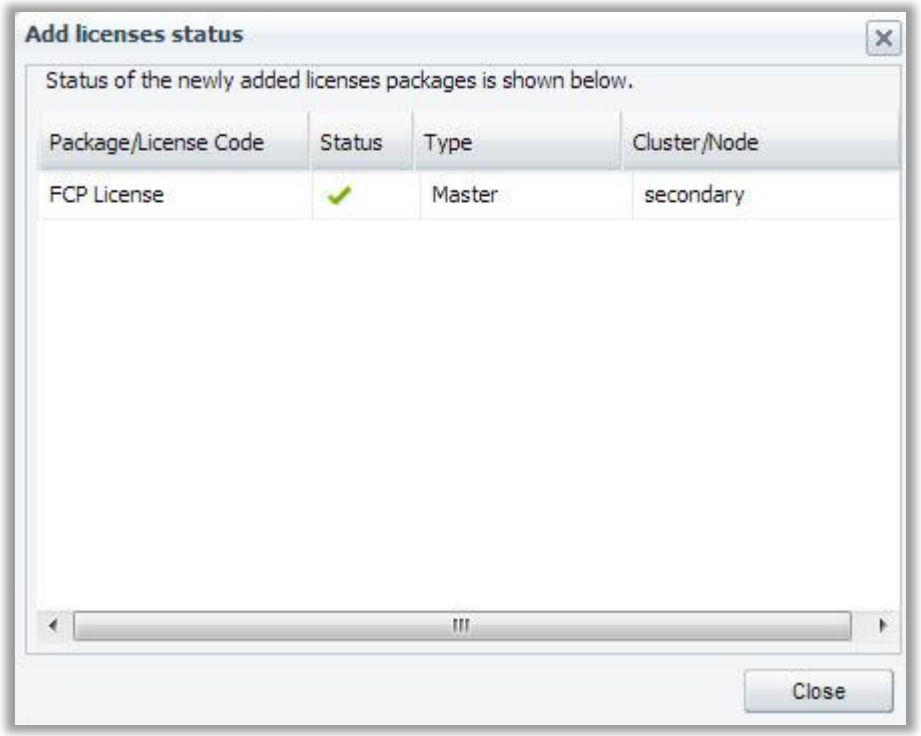

9. Click Close.

The software license is added to your storage system. The newly added license is also included in the list of licenses in the Licenses window.

Figure 56) License details.

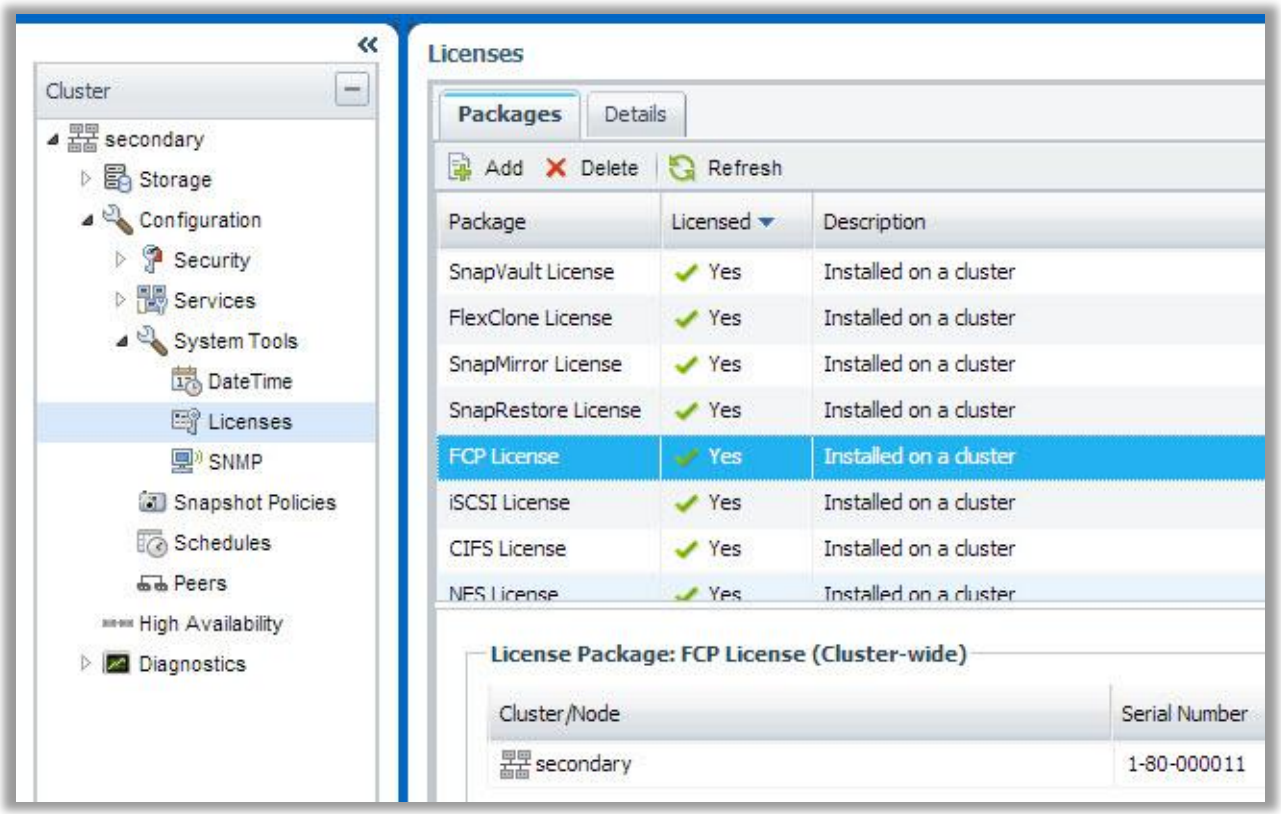

3 New Workflows in System Manager Version 3.1

3.1 Applying Storage Quality of Service Policy Group to Workloads

With System Manager 3.1, a user can manage storage quality of service (QoS) for volumes and LUNs. A user can create QoS policy groups and assign FlexVol volumes or LUNs to new or existing policy groups. The maximum throughput specified for the policy group enables you to manage the workload of storage objects.

For demonstrations about assigning storage quality of service policy group using System Manager 3.1, see the demonstration video on NetApp Community.

Use case: A storage administrator wants to deploy test applications in a production NetApp cluster and wants to make sure that the production workload is not affected.

As part of this workflow, we perform the following tasks:

- Assign new test application workloads to “untested_apps” QoS policy group
- Limit maximum throughput on the “untested_apps” QoS policy group
- When the test application workloads are ready to be moved to production, unassign the test application workloads from the “untested_apps” QoS policy group

Note: This workflow assumes that the user has a few volumes dedicated for test purposes. If not, create them before using this workflow.

Workflow Steps

Assign Test Application Workloads to QoS Policy Group and Limit Maximum Throughput

1. From the homepage, double-click the appropriate storage system.
2. Expand the Storage Virtual Machines hierarchy in the left navigation pane.
3. In the navigation pane, select the SVM and click Storage > Volumes.
4. Select the test application volumes for which you want to assign storage QoS. Click Storage QoS.

Figure 57) Assign workloads to QoS policy group.

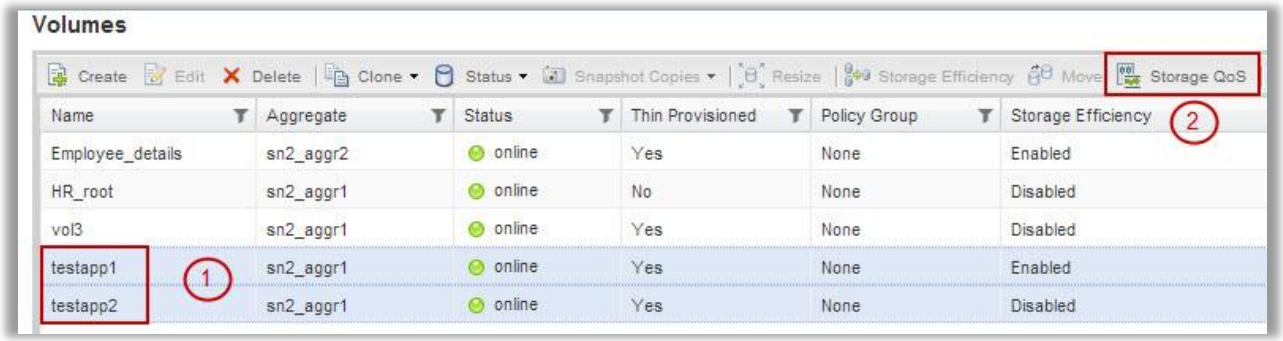

5. In the Quality of Service Details dialog box, the Manage Storage Quality of Service check box is selected to manage the workload performance of the FlexVol volumes.

   Note: If some of the volumes you selected are already assigned to a policy group, the changes that you make will overwrite the existing configuration for these volumes. Click the volumes hyperlink to review the list of selected volumes.

Figure 58) Quality of service details window.

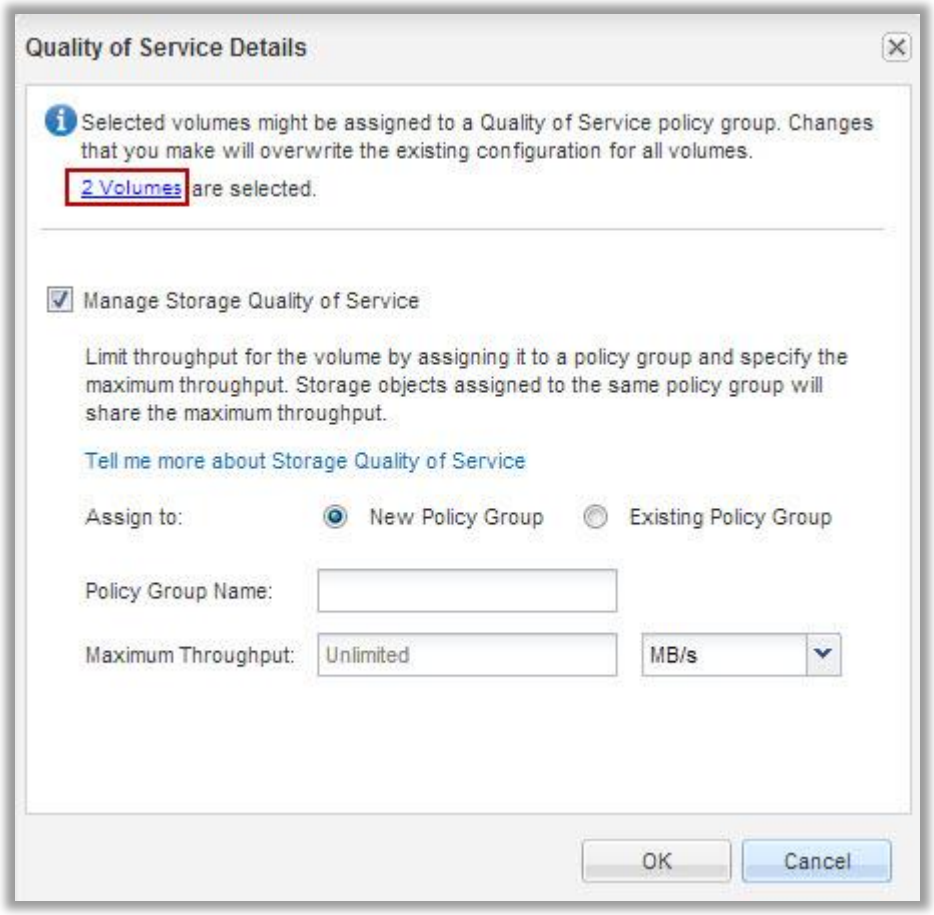

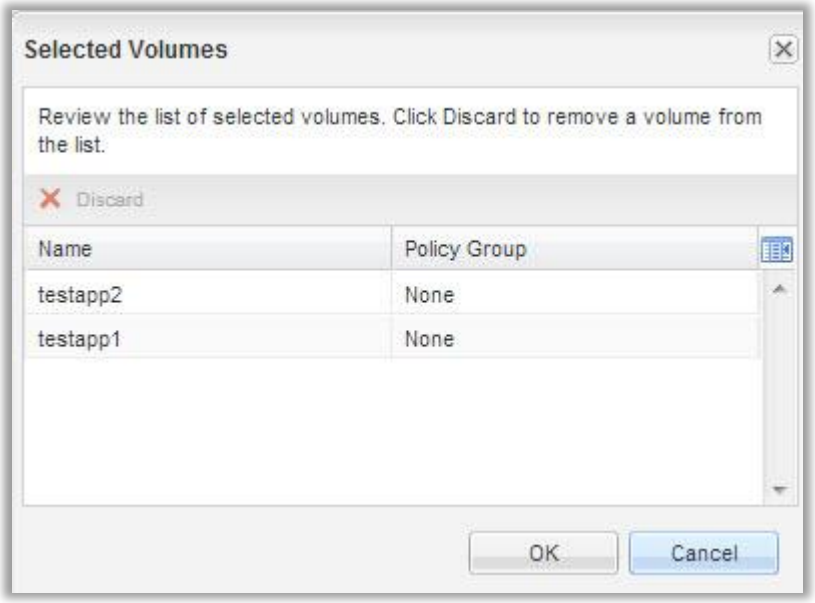

6. Create a new storage QoS policy group to control the input/output (I/O) performance of the FlexVol volumes.

   Specify the policy group name and the maximum throughput limit to make sure that the workload of the objects in the policy group does not exceed the specified throughput limit.

   Note: If you do not specify the maximum throughput limit, the value is set to Unlimited, and the unit that you specify does not affect the maximum throughput.

Figure 59) Assign maximum throughput.

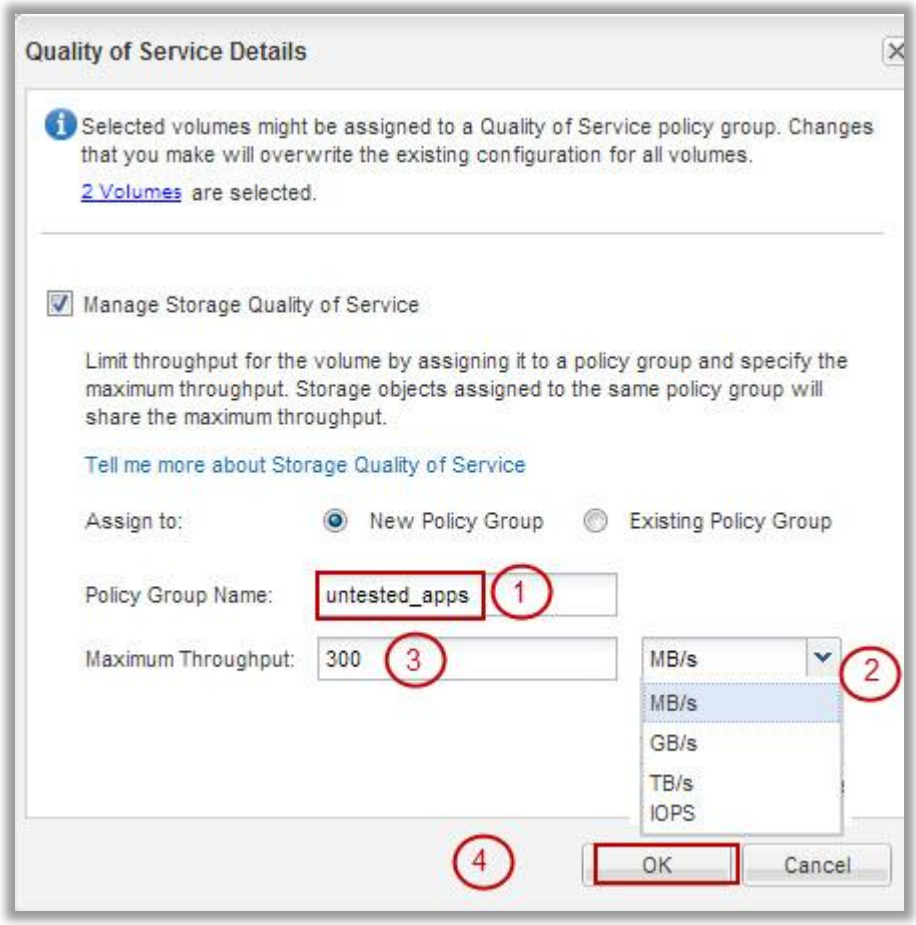

Details about the QoS policy groups can be viewed from the QoS Policy Groups window.

Figure 60) QoS Policy Groups window.

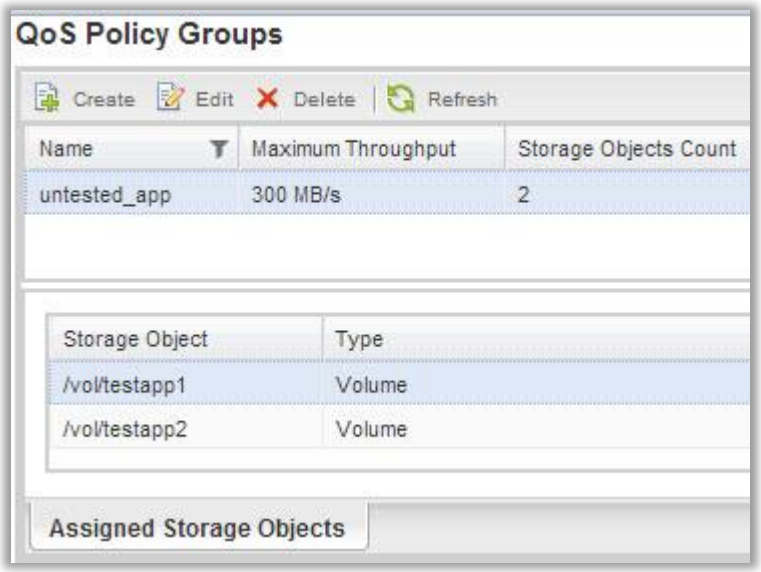

Unassign the Test Application Workload from a QoS Policy Group

When the test application workloads are ready to be moved to production, they can be unassigned from the QoS policy group.

1. Select the test application volumes for which you want to unassign storage QoS. Click Storage QoS.
2. Clear the Manage Storage Quality of Service check box to stop managing the workload performance of the test application volumes.

Figure 55) Unassign QoS.

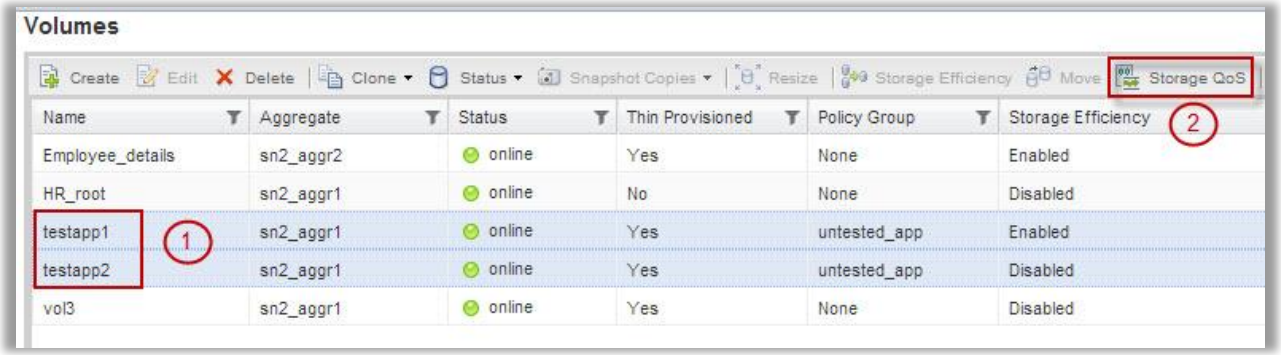

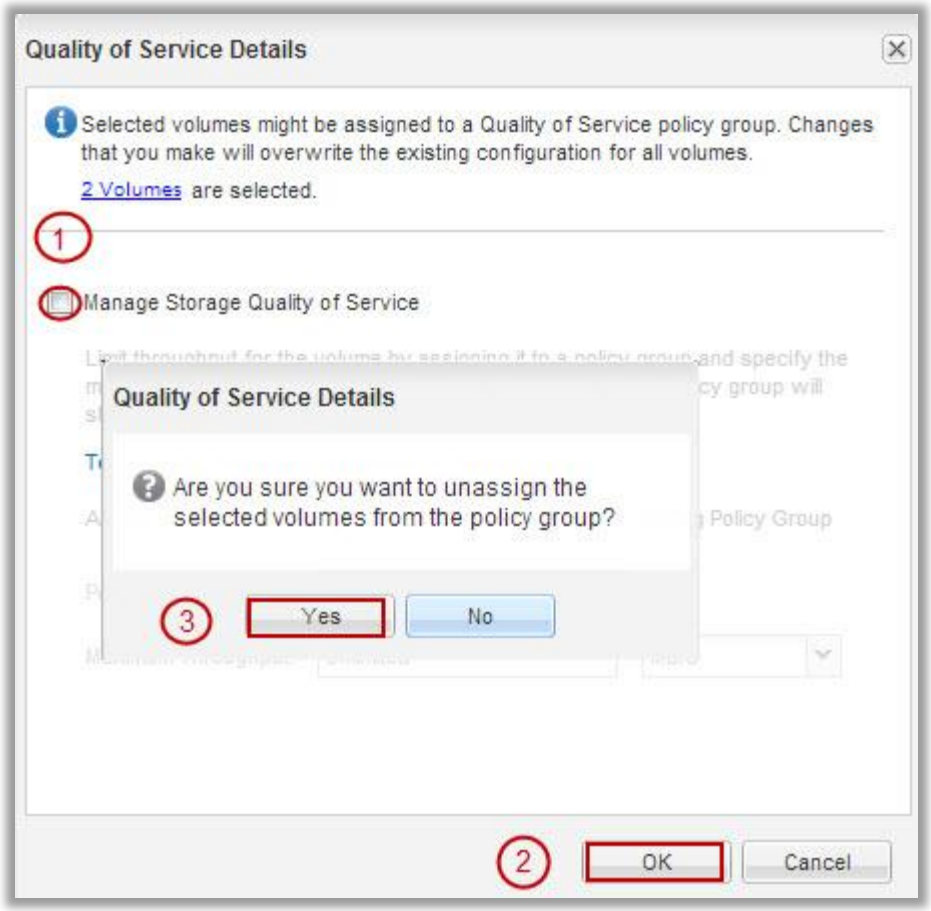

Note:

1. You can assign storage QoS only to read/write (rw) volumes that are online.
2. You cannot assign storage QoS to a volume if the following storage objects are assigned to a policy group:
   a. Parent storage virtual machine (SVM) of the volume
   b. Child LUNs of the volume
   c. Child files of the volume
3. You can assign storage QoS or modify QoS details for a maximum of 10 volumes at the same time.

As part of the assign QoS policy group using System Manager 3.1 workflow, we performed the following tasks:

- We assigned new test application workloads to “untested_apps” QoS policy group.
- We limited maximum throughput on the “untested_apps” QoS policy group.
- When the test application workloads were ready to be moved to production, we unassigned the test application workloads from the “untested_apps” QoS policy group.

3.2 Creating Infinite Volumes for Clustered Data ONTAP 8.2.x

With System Manager 3.1, you can create Infinite Volumes for clustered Data ONTAP 8.2.x to provide a large, scalable data container with a single namespace and a single mount point.

For demonstrations about Infinite Volume support in System Manager 3.1, refer to the demonstration video on NetApp Community.

Use case: A global company wants to create a repository for video assets in the form of raw media files ranging from 100MB to 5GB. These files are written once and read infrequently.

As part of this workflow, we do the following:

- Create SVM with Infinite Volume
- Configure CIFS and NFS protocols on an SVM
- Delegate administration to SVM administrator
- Assign aggregates to SVM with Infinite Volume
- Create and mount an Infinite Volume and then create a share

Workflow Steps

Create a Storage Virtual Machine with Infinite Volume

1. From the homepage, double-click the appropriate storage system.
2. Expand the Storage Virtual Machines hierarchy in the left navigation pane.
3. In the navigation pane, select the cluster name. Click Create.

Figure 56) SVM with Infinite Volume.

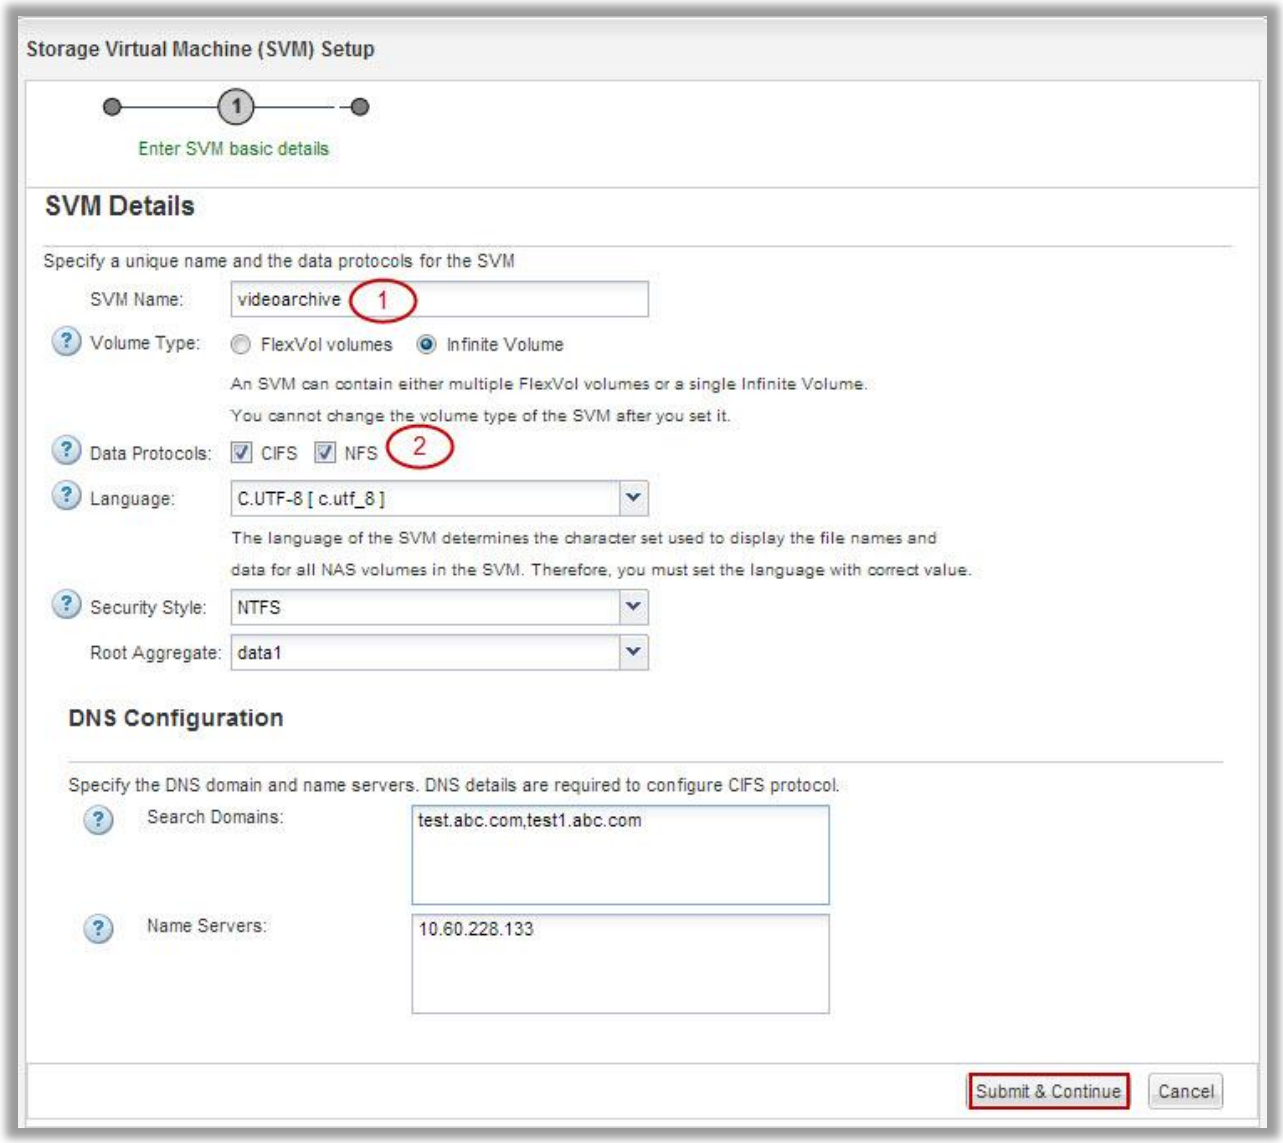

4. In the Storage Virtual Machine (SVM) Setup window, select the SVM volume type as Infinite Volume and specify the SVM details such as the SVM name, volume type, protocols, SVM language, root volume security style, and root aggregate.

   Note: CIFS and NFS are the only data protocols allowed on an SVM with Infinite Volume.

   The default language setting for an SVM is C.UTF-8.

   By default, the aggregate with the maximum free space is selected as the container for the root volume of the SVM. Based on the protocols selected, the default security style and the root aggregate are selected. The security style is set to NTFS if you select CIFS protocol or a combination of CIFS protocol with the other protocols. The security style is set to UNIX only if you select NFS protocol.

5. Specify the DNS domain names and the name server IP addresses to configure the DNS services.

   Note: The default values are selected from the existing SVM configuration.

6. Click Submit & Continue. The SVM is created with the specified configuration.

Configure CIFS and NFS Protocols on an SVM

Figure 57) Configure CIFS and NFS on SVM.

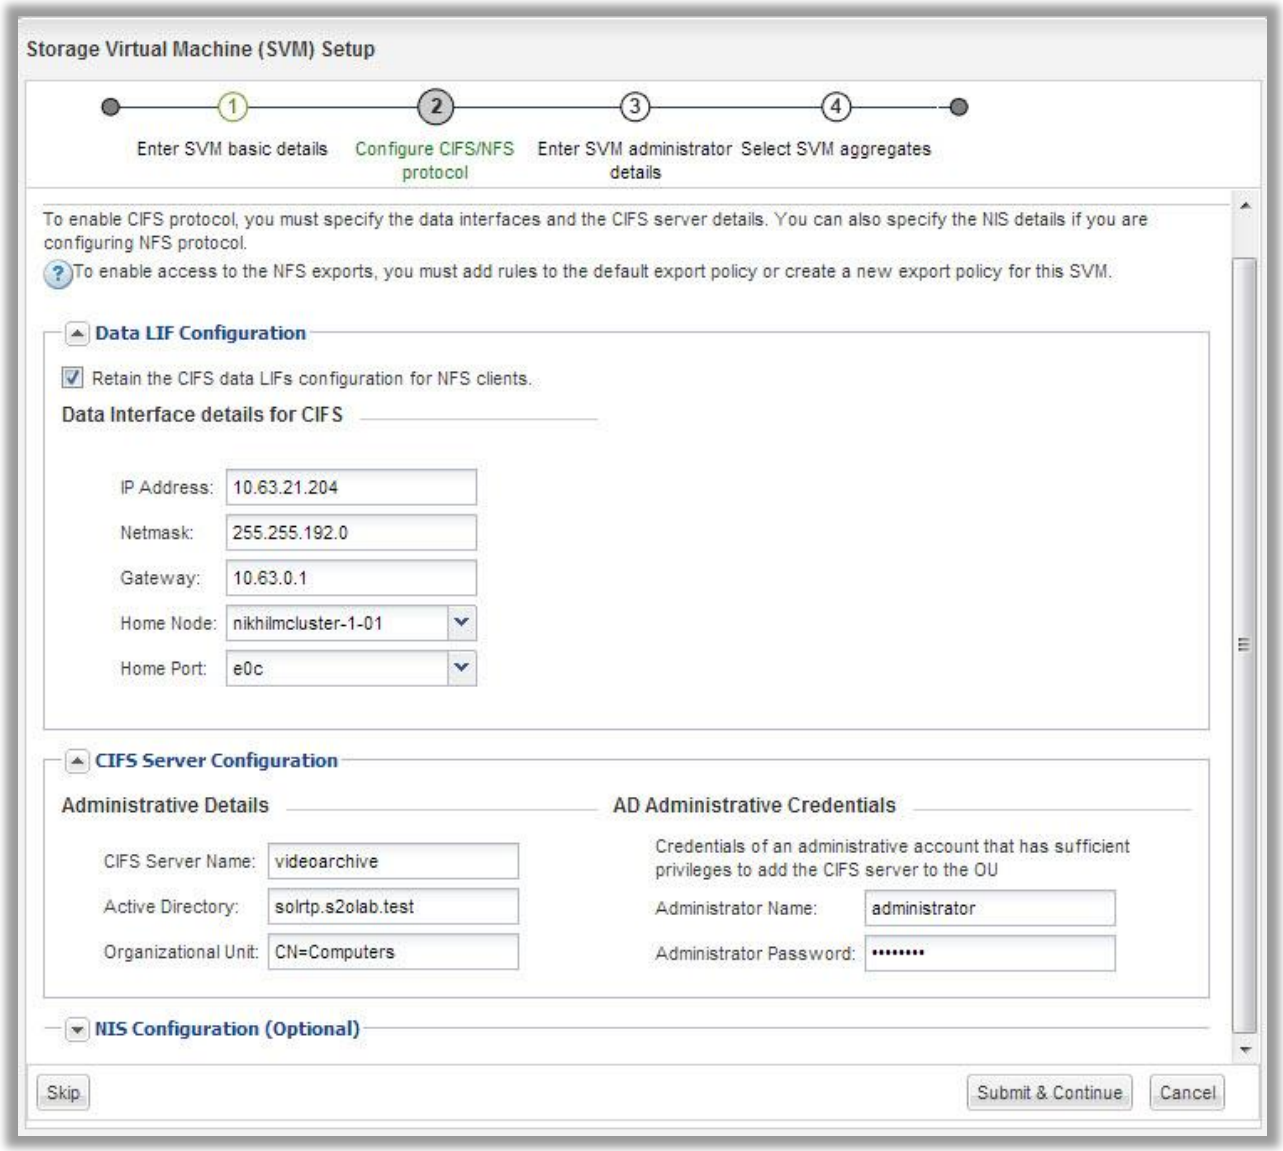

1. In the Data LIF Configuration section, specify the network details to create data LIFs. You can either retain the same data LIF configuration for both CIFS and NFS or configure a new LIF for each protocol.
2. Specify the following information to create a CIFS server:
   a. CIFS server name
   b. Active Directory to associate with the CIFS server
   c. Organizational unit (OU) within the Active Directory domain to associate with the CIFS server
   d. Credentials of an administrative account that has sufficient privileges to add the CIFS server to the OU

   Note: By default, this parameter is set to CN=Computers.

3. **Optional:** You can also specify the IP addresses of the NIS servers and NIS domain name to configure NIS services on the SVM.
4. Click Submit & Close.

Delegating Administration to SVM Administrators

Figure 58) SVM administrator account.

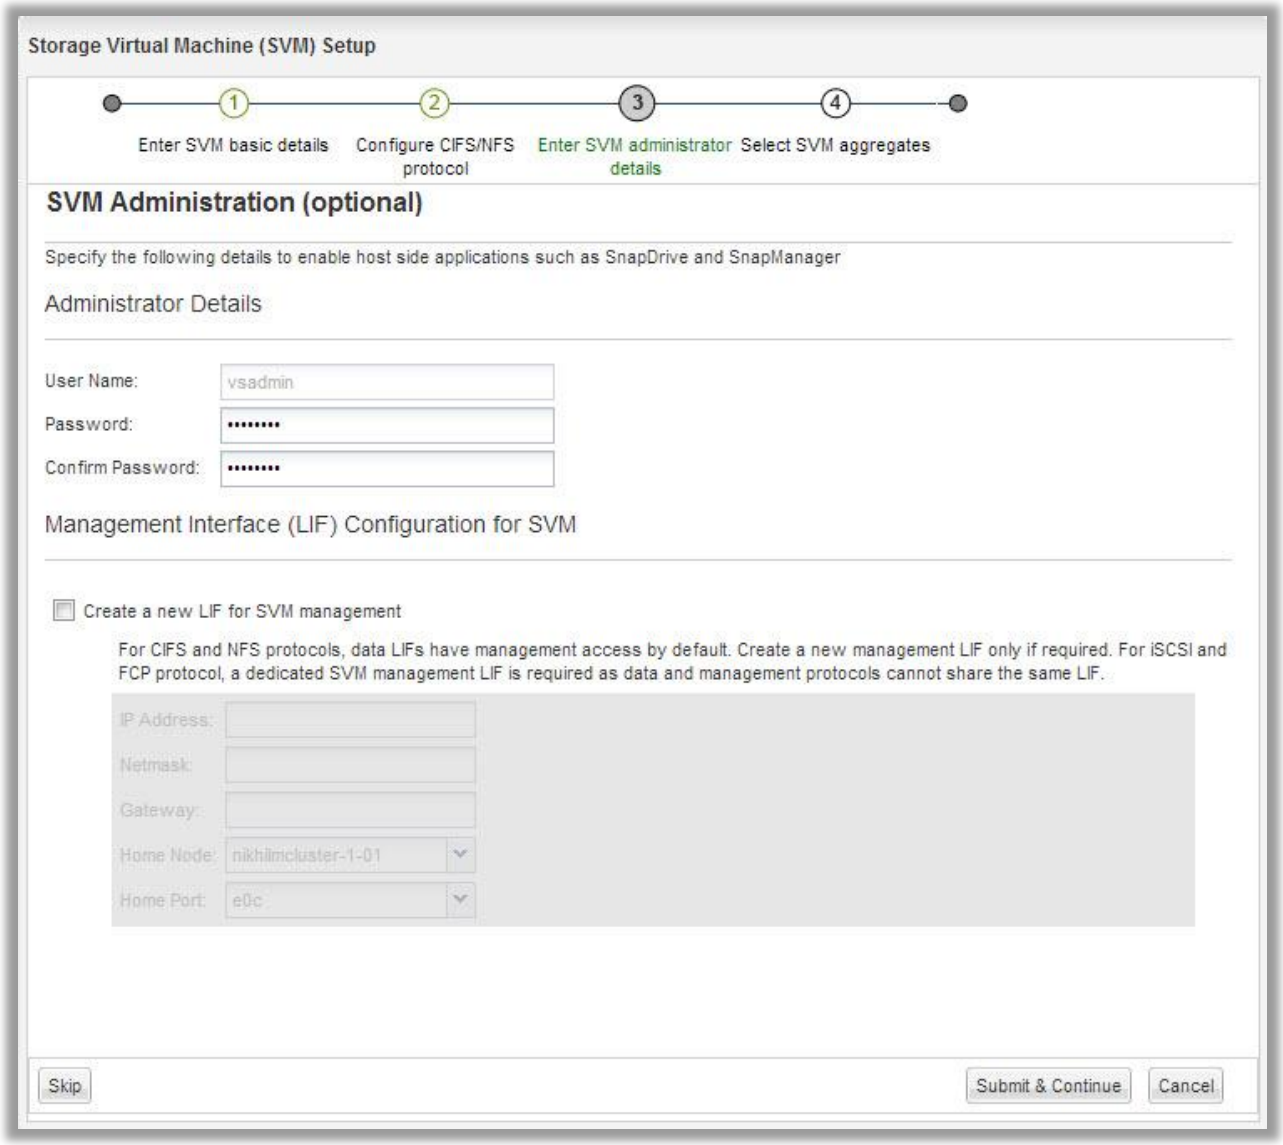

After setting up a functional SVM with Infinite Volume, you can optionally delegate the administration of the SVM to SVM administrators.

1. In the Administrator Details section, set up a password for the vsadmin user account.
2. If you want a dedicated LIF for SVM management, select Create a LIF for SVM management and specify the network details.

Assigning Aggregates to SVMs

After creating an SVM with Infinite Volume, you should assign specific aggregates to it so that the Infinite Volume that you create will use those specific aggregates and not use all the aggregates in the cluster.

Figure 59) Assigning aggregates to SVMs.

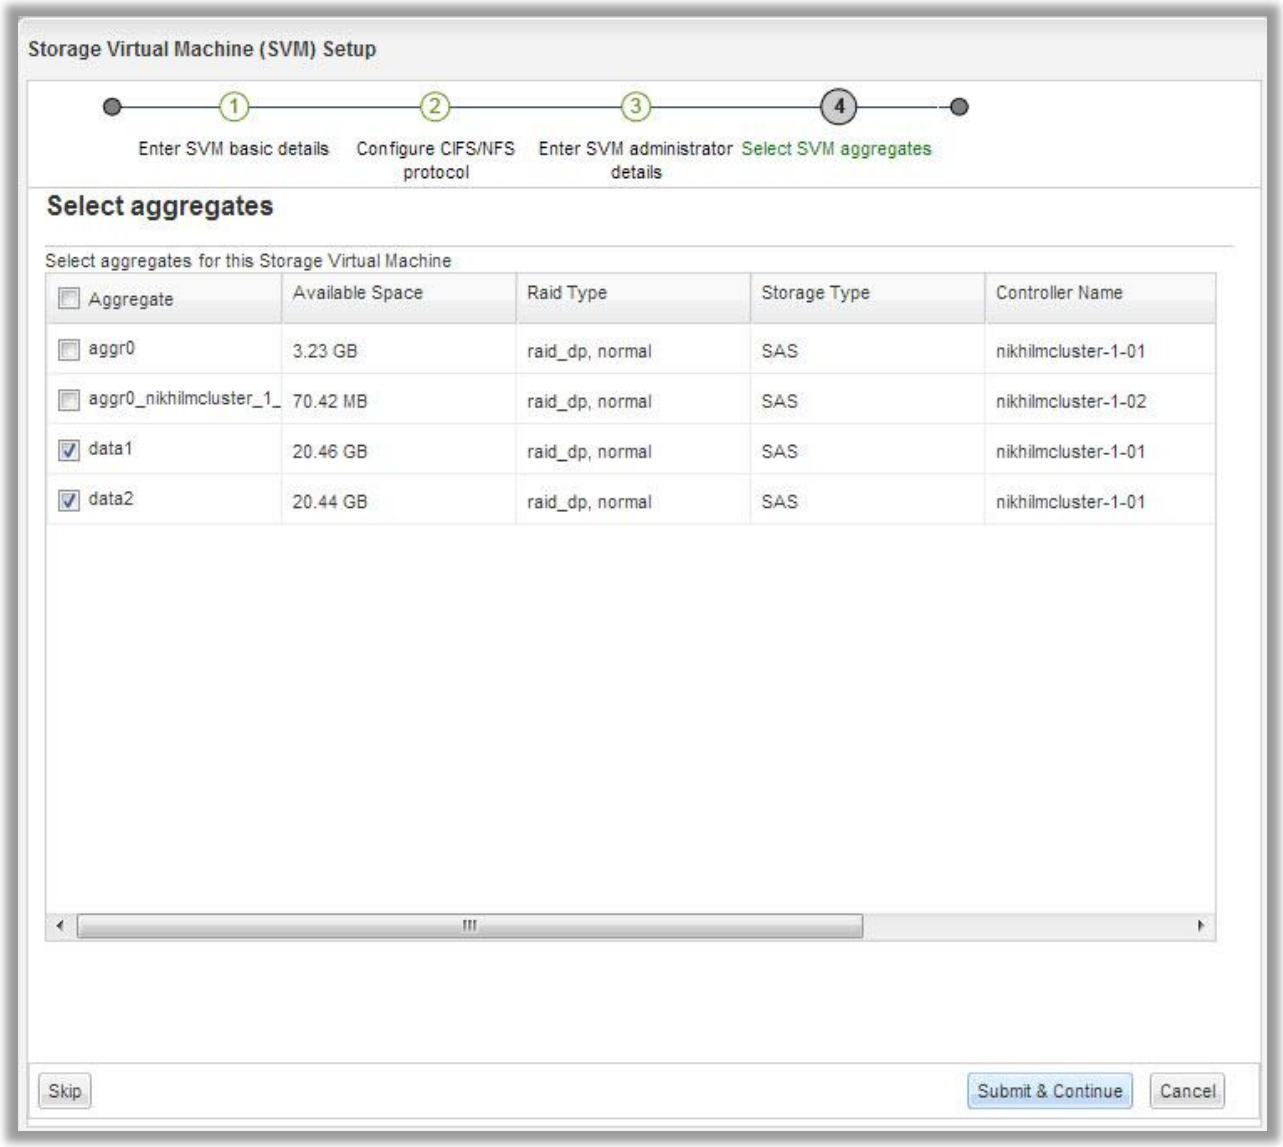

1. In the Select aggregates section, select the aggregates to assign to the SVM with Infinite Volume.
   Note: By default, the node root aggregates are not selected.
2. Click Submit & Continue.

Creating and Mounting an Infinite Volume

1. Expand the Storage Virtual Machines hierarchy in the left navigation pane.

2. In the navigation pane, select the SVM with Infinite Volume and click Storage > Infinite Volume. Click Create.

Figure 66) SVM inventory page.

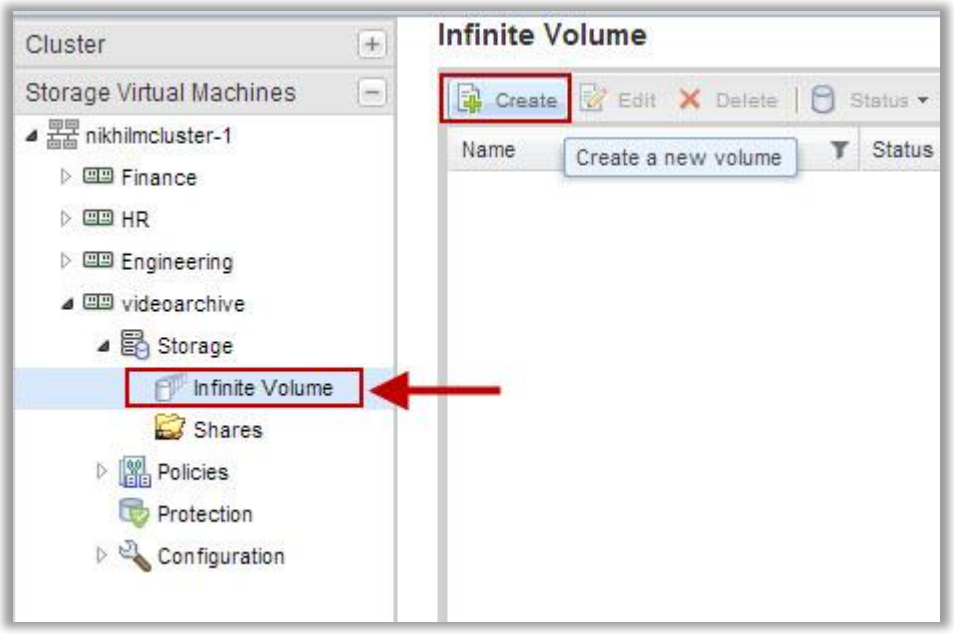

3. If you want to change the default name, specify a new name.
4. Specify a junction path to mount the volume.
5. Note: The junction path can only be a single element path, such as /NS or /InfiniteVol. More than one element (for example, "/NS/v1") is not allowed.
6. Optional: Select Data Protection if the Infinite Volume you are creating is a SnapMirror destination volume. You are provided read-only access to this volume.
7. The number of aggregates that the volume spans is displayed.
8. Click Edit to modify the list of aggregates that are available to the Infinite Volume.
9. Specify the size of the volume and the percentage of the total volume size that you want to reserve for Snapshot copies.
10. Note: The minimum size of the volume is 1.33TB for each node used. The default space reserved for Snapshot copies is 5 percent.
11. Optional: Select Enable SnapDiff to enable incremental tape backup of the volume. Incremental tape backup requires SnapDiff support from your backup application.
12. Optional: If you want to enable deduplication, compression, or both on the volume, make the necessary changes in the Storage Efficiency tab.
13. Click Create.
14. Verify that the volume you created is displayed in the Infinite Volume window.

Mount and unmount operations on an Infinite Volume can be performed from the Infinite Volume inventory page. You can use the SnapMirror feature available under the Protect by option to replicate data from a source Infinite Volume to a destination Infinite Volume.

Figure 67) Creating and mounting an Infinite Volume.

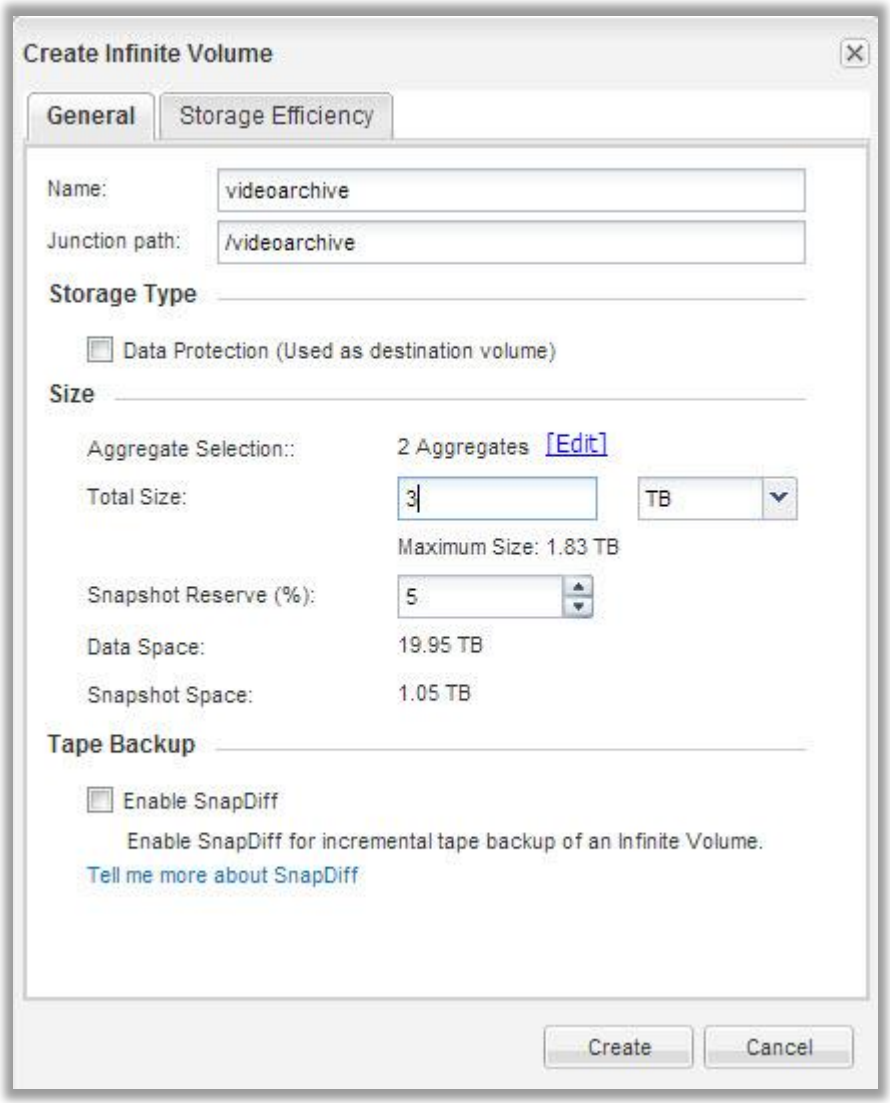

Figure 68) Infinite Volume.

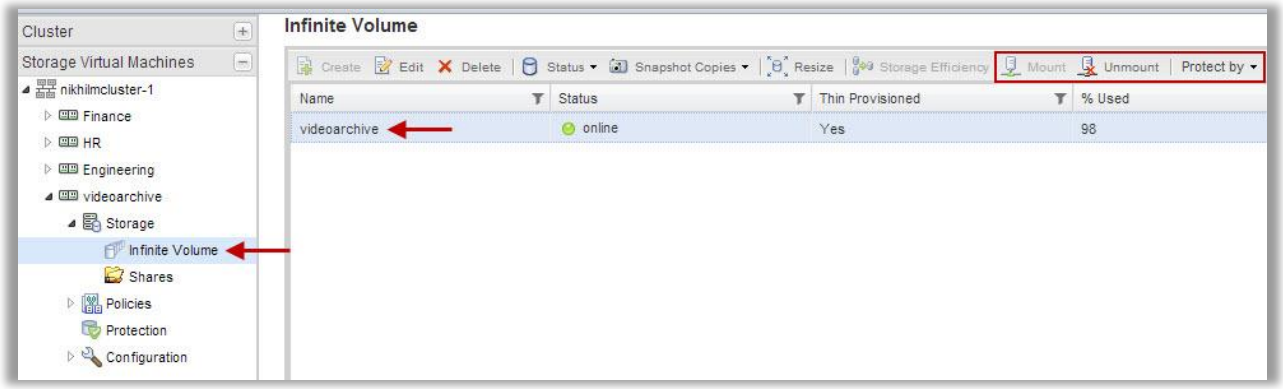

Creating a CIFS Share on Infinite Volume

1. In the navigation pane, select the SVM with Infinite Volume and click Storage > Shares.
2. Click Create Share.
3. Specify the junction path of the volume that should be shared.
4. Specify a name for the new CIFS share.
5. Provide a description for the share and click Create.

Figure 69) Creating a CIFS share.

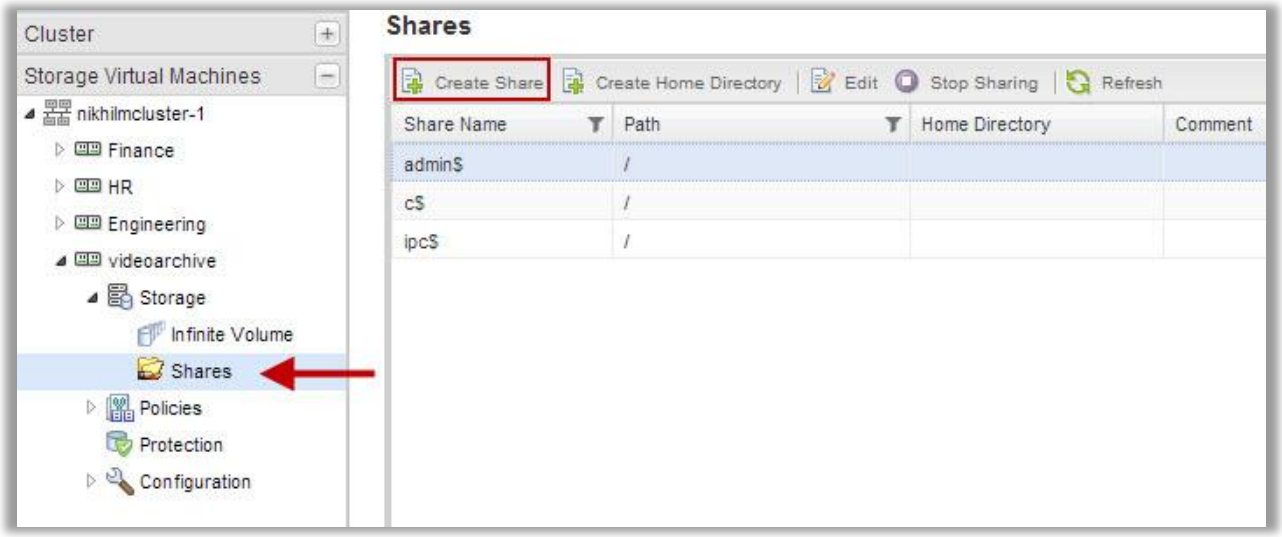

Figure 60) Create Share window.

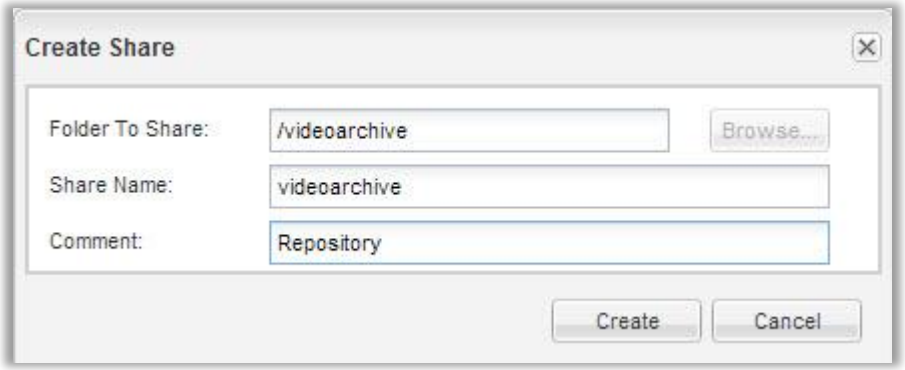

In this workflow, an SVM with Infinite volume was created followed by configuring CIFS protocol on the SVM. We also assigned specific aggregates to SVM with Infinite volume so that the Infinite Volume can only span across those aggregates. Then, we created an Infinite volume and configured a share on the Infinite Volume.

3.3 Initiating Manual Takeover/Giveback of Nodes in an HA Pair

With System Manager 3.1, you can monitor the state and interconnect statuses of the high-availability (HA) pairs in a cluster, as well as perform manual takeover and giveback operations on the nodes in an HA pair.

The ability to perform manual takeover and giveback operation is available for clustered Data ONTAP 8.2.1 clusters.

For demonstrations about performing HA takeover/giveback by using System Manager 3.1, see the demonstration videos on NetApp Community.

Use case: A storage administrator wants to perform maintenance activities on one of the nodes in a clustered Data ONTAP 8.2.1 cluster. The administrator wants to perform a takeover operation on the node, monitor the status of HA pairs in a cluster during takeover, and then initiate a giveback operation. After the maintenance operations on the node are complete, the administrator must monitor the same.

As a part of the above workflow, we perform the following steps:

- Ascertain the status of the nodes in the HA pair.
- Initiate a manual takeover of the node, which will undergo maintenance.
- Monitor takeover operation.
- Initiate a manual giveback to the node that underwent maintenance.
- Monitor giveback operation.

Workflow Steps

Ascertain the Status of the Nodes in HA Pair

You can monitor the state and interconnect status of all the HA pairs in a cluster running clustered Data ONTAP 8.2.x.

You can verify whether takeover or giveback operations are enabled or have occurred and view the reasons due to which the takeover or giveback operations are not currently possible.

1. From the homepage, double-click the appropriate storage system.
2. Expand the Cluster hierarchy in the left navigation pane.
3. In the navigation pane, click High Availability.

Figure 61) High Availability window.

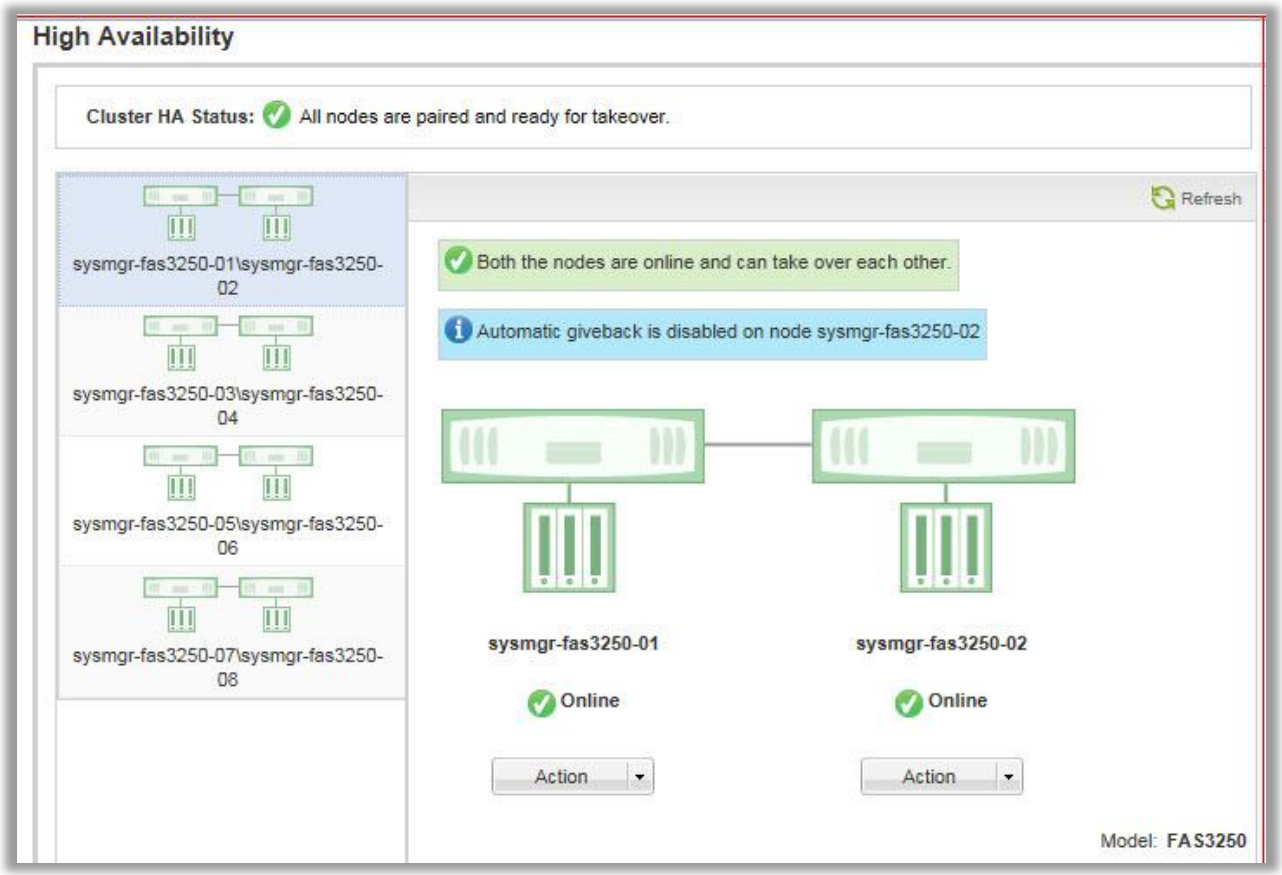

4. In the High Availability window, click the HA pair image to view details such as the cluster HA status, node status, interconnect status, and hardware model of each node.

Note: If the cluster management LIF or the data LIFs of a node are not in their home node, a warning message is displayed indicating that the node has some LIFs that are not in the home node.

The High Availability window provides a pictorial representation of the HA state, interconnect status, and takeover or giveback status of all the HA pairs in clustered Data ONTAP. You can view details such as the takeover or giveback status and interconnect status by clicking the HA pair image.

The following colors indicate the below HA pair statuses:

- Green: Indicates that the HA pair and the interconnect are optimally configured and available for takeover or giveback. It also indicates takeover is in progress, giveback is in progress, and waiting for giveback states.
- Red: Indicates a downgraded state such as a takeover failure.
- Yellow: Indicates that the interconnect status is down.

Initiate a Manual Takeover and Monitor Takeover Operation

Considering that node cluster-1-20 in the HA pair is the node that will undergo maintenance operation, let's initiate a takeover from node cluster-1-19.

1. Click Action under node cluster-1-19 and select Takeover of node cluster-1-20 to initiate takeover operation.

   A confirmation window pops up to confirm the takeover operation.

Figure 62) Initiate takeover.

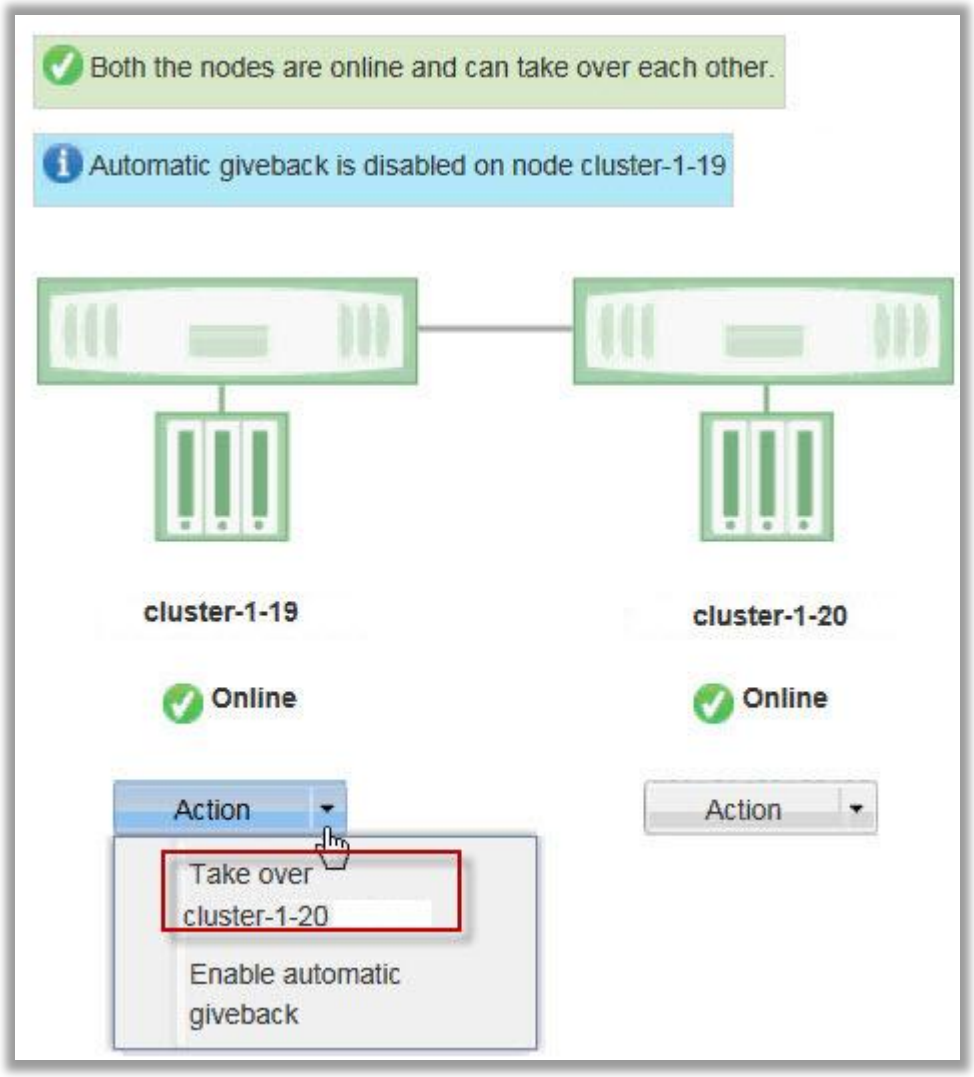

Figure 63) Takeover Confirmation window.

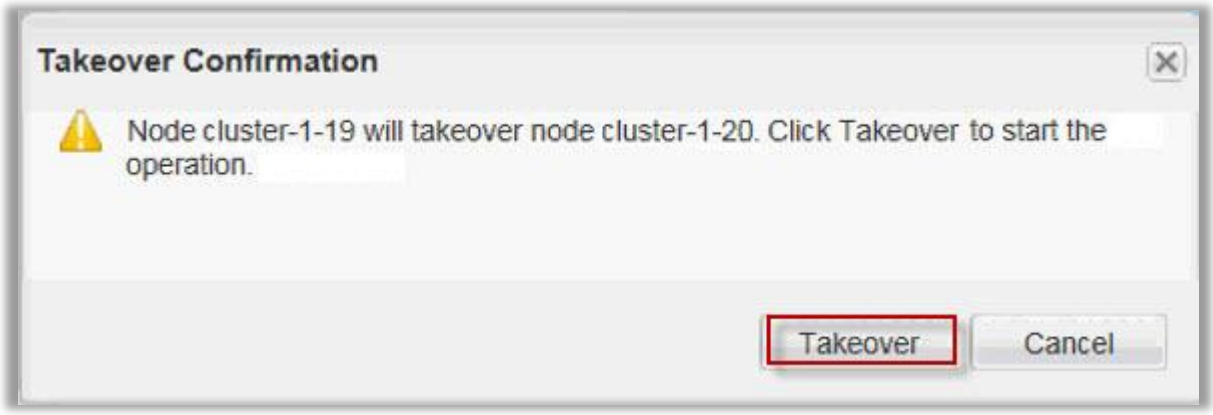

2. Click Takeover to start the operation.
   The High Availability window refreshes every minute to provide the status of the takeover operation.
   By refreshing the contents of the High Availability window (the graphical representation of the HA state, interconnect status and the node) periodically, System Manager enables you to know when takeover has completed successfully. The color red on the node cluster-1-20 indicates a downgraded state. In our scenario the node cluster-1-20 has been taken offline for maintenance purposes.

Figure 64) Takeover in progress.

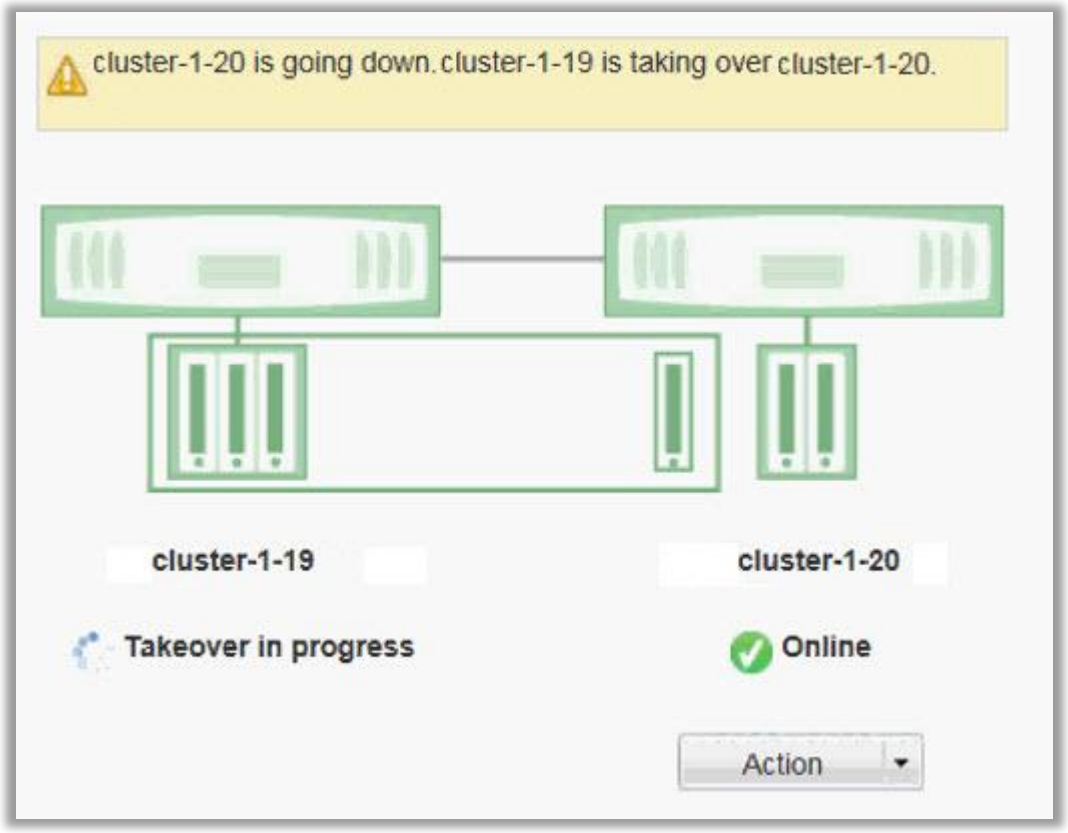

Figure 75) Takeover complete.

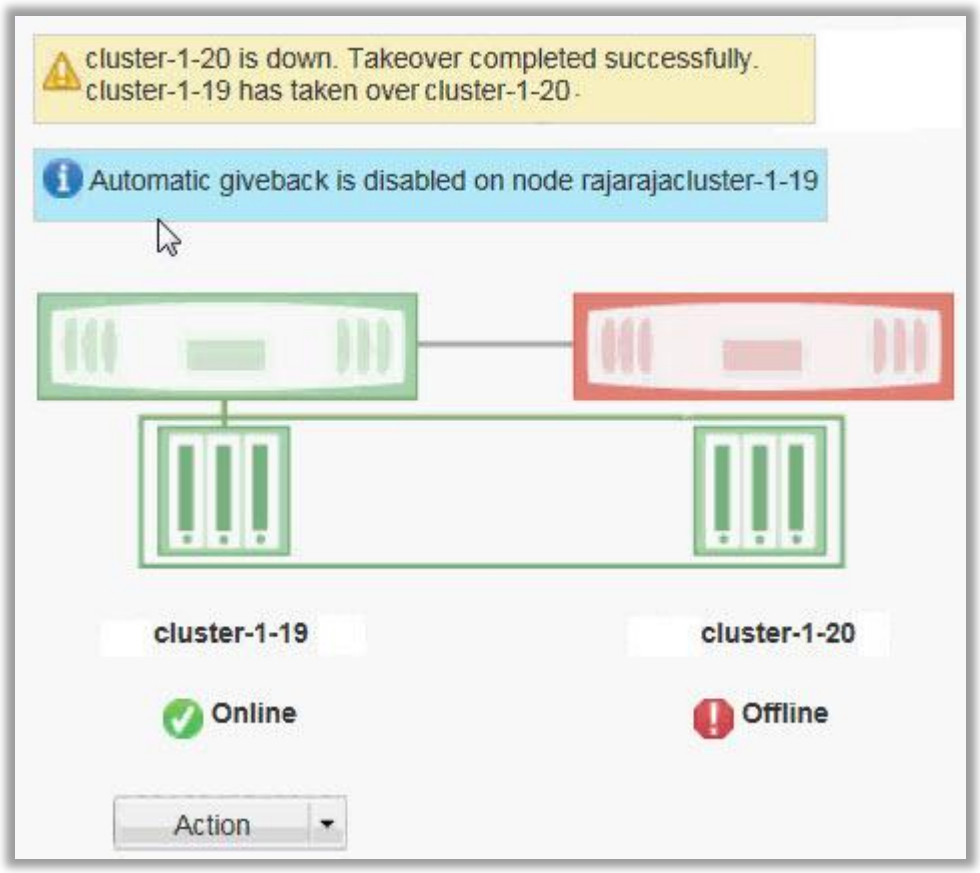

Initiate a Manual Giveback and Monitor Giveback Operation

While a maintenance operation is in progress, disable automatic giveback on the node cluster-1-19 which has taken over resources of the other node cluster-1-20 by using Action.

1. After the maintenance operation is complete, pictorial representation on node cluster-1-20 turns to green color with a status of waiting for giveback state.

Figure 65) Waiting for giveback.

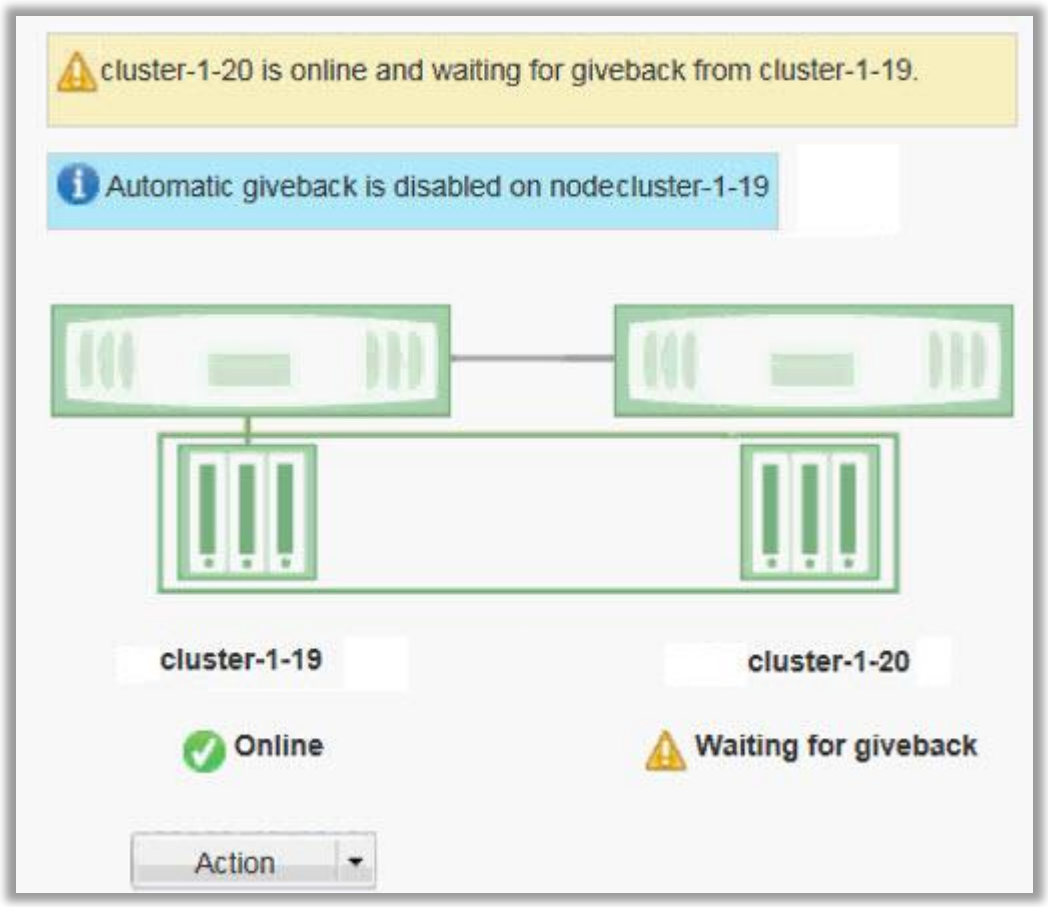

Figure 66) Initiating giveback.

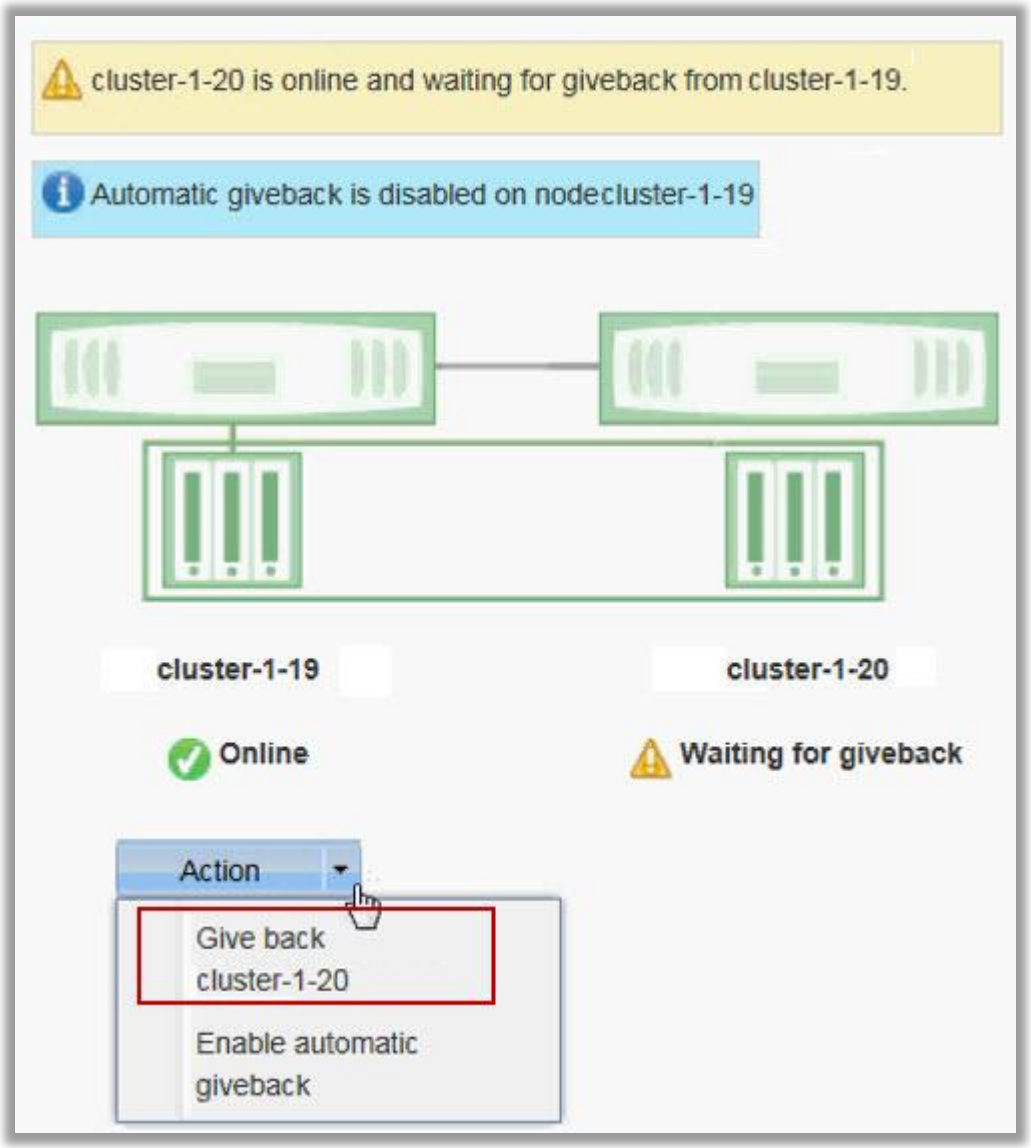

2. Because automatic giveback is disabled, click Action under node cluster-1-19 and select Give back to node cluster-1-20 to initiate giveback operation.
3. The High Availability window refreshes every minute to provide the status of the giveback operation.

Figure 67) Giveback in progress.

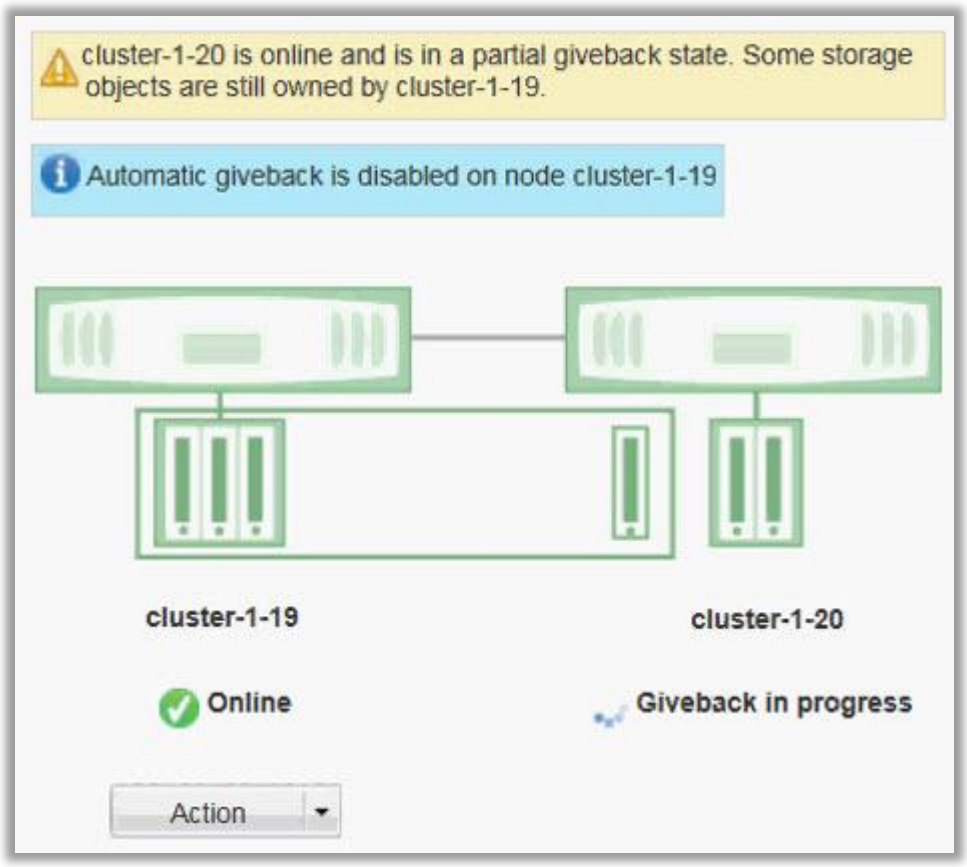

4. By refreshing periodically, the graphical representation of the HA state, interconnect status, and node lets you know when the giveback operation completed successfully. The color green on the node cluster-1-20 indicates that the HA pair and the interconnect are optimally configured.

Figure 68) Nodes online.

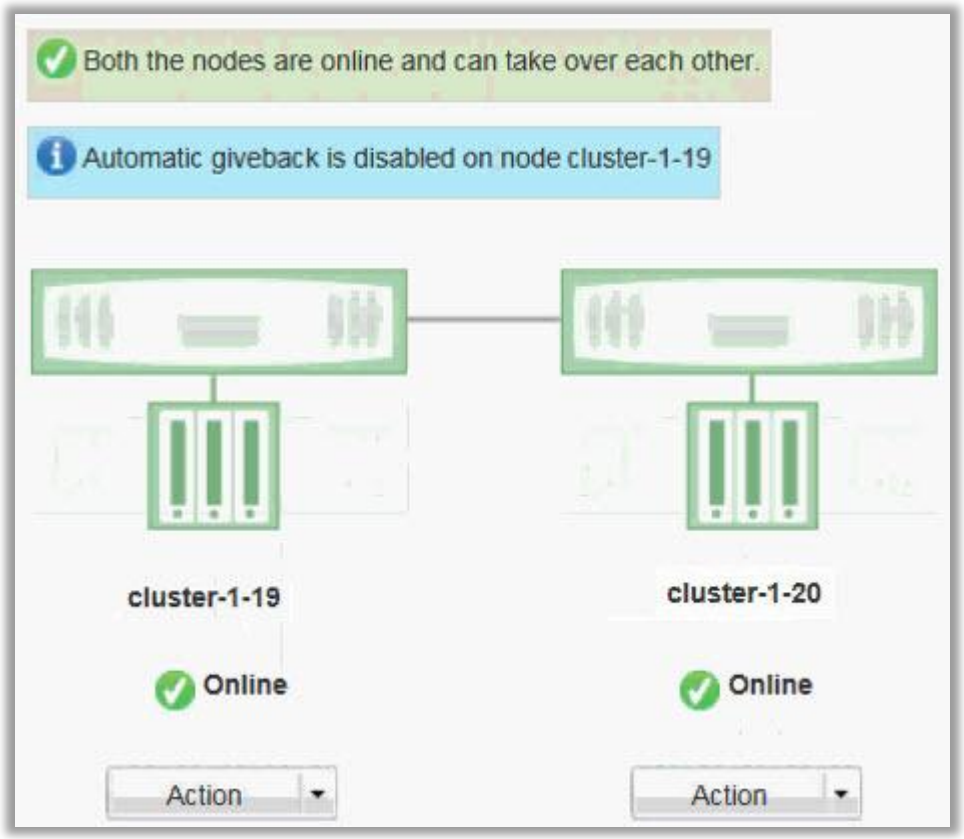

As a part of the preceding workflow, we performed the following steps:

- Ascertained the status of the nodes in the HA pair.
- Initiated a manual takeover of the node, which will undergo maintenance.
- Monitored takeover operation.
- Initiated a manual giveback to the node that underwent maintenance.
- Monitored giveback operation.

4 Summary

This document provides an overview of new and enhanced workflows introduced with OnCommand System Manager 3.0 and 3.1. With System Manager 3.0 and 3.1, new and enhanced workflows are available for clustered Data ONTAP 8.2.x.

For any specific workflow that is not listed in this document, refer to the “OnCommand System Manager 3.0 and 3.1 Help for Use with Clustered Data ONTAP” guide for support information.

The System Manager 3.0 release also includes existing workflows from the System Manager 2.2 release.

For information on System Manager 2.2 workflows, refer to TR-4031: “System Manager 2.2 Workflow Guide.” This guide lists the most common workflows used by storage administrators for the configuration and ongoing management of storage controllers for both 7-Mode and Cluster-Mode operations.

References

The following reference was used in this technical report.

- OnCommand System Manager 2.2 Guide to Common Workflows
  http://www.netapp.com/us/system/pdf-reader.aspx?pdfuri=tcm:10-107717-16&m=tr-4031.pdf

Version History

Version	Date	Document Version History
Version 1.0	August 2013	System Manager 3.0 Workflow Guide
Version 1.1	December 2013	System Manager 3.0 and 3.1 Workflow Guide

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®

www.netapp.com