

IBM System Storage N series



OnCommand Unified Manager Guide to Common Workflows for 7-Mode for Use with Core Package 5.1 and Host Package 1.2

Contents

Preface	6
Supported features	6
Websites	6
Getting information, help, and service	6
Before you call	7
Using the documentation	7
Hardware service and support	7
Firmware updates	8
How to send your comments	8
Introduction to OnCommand Unified Manager Guide to Common Workflows	9
Backing up physical storage objects	10
Adding a dataset of physical storage objects (7-Mode only)	11
Assigning or changing a protection policy	12
Adding a weekly protection schedule	13
Backing up unprotected physical storage objects	14
Identifying unprotected physical storage objects	15
Adding unprotected data to an existing dataset	15
Restoring a single file from a backup	17
Locating specific backups (7-Mode only)	18
Mounting backups in a VMware environment from the Backups tab (7-Mode only)	19
Using VMware to restore a single file	20
Unmounting backups in a VMware environment from the Backups tab (7-Mode only)	20
Supporting delegated management for virtual infrastructure administration	22
Launching the Operations Manager console	23
Creating an administrative role	23
Adding administrative users	24
Creating groups	25
Assigning permissions to an administrative role	25

Access permissions for the Virtual Infrastructure Administrator role	26
Editing group membership	27
Resolving issues with virtual objects	28
Monitoring dashboard panels	28
Monitoring dataset status (7-Mode only)	29
Editing a dataset containing virtual objects to reschedule or modify local backup jobs (7-Mode only)	30
Resolving issues with physical objects	31
Monitoring dashboard panels	32
Starting secondary space management	32
Changing the storage service for a virtual dataset	34
Changing a storage service for datasets of storage objects (7-Mode only)	34
Resolving dataset conformance issues	36
Listing nonconformant datasets and viewing details (7-Mode only)	37
Evaluating conformance error text	38
Resolving conformance issues automatically without a baseline transfer of data (7-Mode only)	40
Resolving conformance issues manually without a baseline transfer of data (7- Mode only)	41
Resolving conformance issues manually when a baseline transfer of data might be necessary (7-Mode only)	42
Appendix of additional information	44
Administrator roles and capabilities	44
Dataset concepts (7-Mode only)	45
Datasets of virtual objects (7-Mode only)	47
Decisions to make before adding a schedule	47
Decisions to make before adding datasets of physical storage objects (for protection) (7-Mode only)	49
Decisions to make before assigning or changing policies	52
Descriptions of dataset conformance status (7-Mode only)	54
Guidelines for adding a dataset of virtual objects (7-Mode only)	54
Guidelines for mounting or unmounting backups in a VMware environment (7- Mode only)	55
How the Secondary Space Management wizard works	56
Role of protection policies in dataset management (7-Mode only)	57

Requirements and restrictions when adding a dataset of virtual objects (7-Mode only)	57
What a dataset is (7-Mode only)	58
Why datasets fail to conform to policy (7-Mode only)	58
Copyright information	60
Trademark information	61
Index	64

Preface

Supported features

IBM System Storage N series storage systems are driven by NetApp Data ONTAP software. Some features described in the product software documentation are neither offered nor supported by IBM. Please contact your local IBM representative or reseller for further details.

Information about supported features can also be found on the N series support website (accessed and navigated as described in [Websites](#) on page 6).

Websites

IBM maintains pages on the World Wide Web where you can get the latest technical information and download device drivers and updates. The following web pages provide N series information:

- A listing of currently available N series products and features can be found at the following web page:
www.ibm.com/storage/nas/
- The IBM System Storage N series support website requires users to register in order to obtain access to N series support content on the web. To understand how the N series support web content is organized and navigated, and to access the N series support website, refer to the following publicly accessible web page:
www.ibm.com/storage/support/nseries/
This web page also provides links to AutoSupport information as well as other important N series product resources.
- IBM System Storage N series products attach to a variety of servers and operating systems. To determine the latest supported attachments, go to the IBM N series interoperability matrix at the following web page:
www.ibm.com/systems/storage/network/interophome.html
- For the latest N series hardware product documentation, including planning, installation and setup, and hardware monitoring, service and diagnostics, see the IBM N series Information Center at the following web page:
publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you. This section contains

information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your IBM N series product, and whom to call for service, if it is necessary.

Before you call

Before you call, make sure you have taken these steps to try to solve the problem yourself:

- Check all cables to make sure they are connected.
- Check the power switches to make sure the system is turned on.
- Use the troubleshooting information in your system documentation and use the diagnostic tools that come with your system.
- Refer to the N series support website (accessed and navigated as described in [Websites](#) on page 6) for information on known problems and limitations.

Using the documentation

The latest versions of N series software documentation, including Data ONTAP and other software products, are available on the N series support website (accessed and navigated as described in [Websites](#) on page 6).

Current N series hardware product documentation is shipped with your hardware product in printed documents or as PDF files on a documentation CD. For the latest N series hardware product documentation PDFs, go to the N series support website.

Hardware documentation, including planning, installation and setup, and hardware monitoring, service, and diagnostics, is also provided in an IBM N series Information Center at the following web page:

publib.boulder.ibm.com/infocenter/nasinfo/nseries/index.jsp

Hardware service and support

You can receive hardware service through IBM Integrated Technology Services. Visit the following web page for support telephone numbers:

www.ibm.com/planetwide/

Firmware updates

IBM N series product firmware is embedded in Data ONTAP. As with all devices, ensure that you run the latest level of firmware. Any firmware updates are posted to the N series support website (accessed and navigated as described in [Websites](#) on page 6).

Note: If you do not see new firmware updates on the N series support website, you are running the latest level of firmware.

Verify that the latest level of firmware is installed on your machine before contacting IBM for technical support.

How to send your comments

Your feedback helps us to provide the most accurate and high-quality information. If you have comments or suggestions for improving this document, please send them by email to starpubs@us.ibm.com.

Be sure to include the following:

- Exact publication title
- Publication form number (for example, GC26-1234-02)
- Page, table, or illustration numbers
- A detailed description of any information that should be changed

Introduction to OnCommand Unified Manager Guide to Common Workflows

This guide provides workflows for tasks that enable you to achieve the common administrative goals that form the basis for OnCommand Unified Manager. This guide is not comprehensive, but it does provide examples of the most common tasks that administrators perform.

This guide is intended for experienced storage and virtual infrastructure administrators. It includes administrative tasks that span multiple user interfaces that launch from a single user interface: the OnCommand console.

Completing tasks in this guide requires that you have both the OnCommand console and all host services installed and configured on your system. See the *OnCommand Unified Manager Installation and Setup Guide* for more information.

You can also use the appendix included in this guide as a source for background information about the guidelines and decisions that precede each workflow.

Related information

IBM N series support website: www.ibm.com/storage/support/nseries

Backing up physical storage objects

This workflow shows you how to create a backup of your physical storage objects by creating a new dataset and then assigning an existing protection policy and weekly backup schedule, using both the OnCommand console and the N series Management Console.

Before you begin

You must have installed the OnCommand console and N series Management Console.

You must be familiar with the protection concepts of a dataset, protection policy, and backup schedule. For more information about these concepts, see the OnCommand console Help.

About this task

This workflow is based on the assumption that you start in the OnCommand console and that you are creating an empty dataset and assigning an existing protection policy and weekly backup schedule as separate steps. You can also assign policies and backup schedules during the creation of the dataset.

During this task, the OnCommand console launches N series Management Console. Depending on your browser configuration, you can return to the OnCommand console by using the Alt-Tab key combination or clicking the OnCommand console browser tab. After the completion of this task, you can leave N series Management Console open, or you can close it to conserve bandwidth.

After you finish

After you have completed the tasks involved in creating a dataset, including assigning a local policy and storage service, you can perform an on-demand backup to ensure that no problems exist with the dataset configuration. For more information about how to perform an on-demand backup, see the OnCommand console Help.

Steps

1. [Add a new dataset of physical storage objects](#) on page 11
Creating a new dataset enables you to group and manage physical storage objects that share the same protection requirements. After creating a dataset, you must assign an existing protection policy.
2. [Assign a protection policy to the new dataset](#) on page 12
Assigning a protection policy to your new dataset enables you to specify how you want your data to be protected. This includes how you want to protect the dataset on primary, secondary, or tertiary storage; schedule backups; and retain backups. After assigning a protection policy, you can add a weekly backup schedule that is applied to your protection policy.
3. [Add a weekly protection schedule](#) on page 13

Adding a weekly protection schedule enables you to specify times for backup and mirror operations and lets you run existing daily schedules on specific days of the week. You can also designate backups to retain as weekly copies based on the protection policy settings.

Adding a dataset of physical storage objects (7-Mode only)

You can include multiple physical storage objects in a dataset for managing the protection requirements or provisioning the storage space and hardware requirements of those objects as a group.

Before you begin

- You must already be familiar with *Decisions to make before adding datasets of physical storage objects (for protection)* on page 49.
- You must have N series Management Console installed.
- You must have the protection information that you require to complete this task.
- You must have the provisioning information that you require to complete this task.
- You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

About this task

During this task, the OnCommand console launches N series Management Console. Depending on your browser configuration, you can return to the OnCommand console by using the Alt-Tab key combination or clicking the OnCommand console browser tab. After the completion of this task, you can leave N series Management Console open, or you can close it to conserve bandwidth.

Steps

1. Click the **View** menu and select the **Datasets** option.
2. In the **Datasets** tab, click **Create** and select the **Dataset with storage objects** option to start the N series Management Console **Add Dataset** wizard.
3. Complete the steps in the **Add Dataset** wizard to create a dataset of physical storage objects.
After you complete the wizard, the new dataset is listed in the Datasets tab.
4. Optional: To provide data protection or disaster recovery protection for the new dataset, complete the following steps in N series Management Console:
 - a) In the **Datasets** tab, select the dataset that you just created and click **Protection Policy** to start the **Dataset Policy Change** wizard.
 - b) Complete the steps in the **Dataset Policy Change** wizard to assign a protection policy to the new dataset.
5. When finished, press **Alt-Tab** or click the OnCommand console browser tab to return to the OnCommand console.

Related references

[Administrator roles and capabilities](#) on page 44

Assigning or changing a protection policy

You can assign a policy to a dataset or change the policy assigned to it. The policy specifies how the data is to be protected.

Before you begin

- You must have gathered the protection information that you need to complete this task:
 - Dataset properties
 - Group membership
 - Protection policy
- Determine which policy you want to assign to the dataset. You can review available protection policies on the Protection Policies window.
If no policy meets the requirements of your new dataset, you can create a new policy or modify a copy of an existing policy.
- You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

About this task

- You can use this procedure after you have created a new dataset and want to assign a policy to it, or when you want to change the protection policy assigned to a dataset. You can also use this procedure to protect a dataset that is listed on the Datasets window
- The N series Management Console operations described in this topic do not apply to virtual object type datasets.
Any dataset that you directly assign a protection policy or provisioning policies and resource pools directly through N series Management Console is displayed as a dataset of physical storage objects.

Steps

1. From the menu bar, click **Data > Datasets > Overview**.
2. Select a dataset and click **Protection Policy** to start the **Dataset Policy Change** wizard.
If you want to provision your nodes by resource pool, click the **Provision and attach resources using a policy** option when it is displayed.
3. Complete the steps in the wizard.

Related concepts

[Decisions to make before assigning or changing policies](#) on page 52

Related references

[Administrator roles and capabilities](#) on page 44

Adding a weekly protection schedule

You can use the Add Schedule wizard to create new weekly protection schedules. After you create a weekly schedule, you can apply it to protection policies to determine when hourly, daily, and weekly backup or mirror operations are executed.

Before you begin

Have the information available that you need to complete this task:

- Days and time to perform the mirror or backup operation and how often (required)
- Retention time of mirror or backup copies (optional)
- Throttle schedule (optional)

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

Steps

1. From the menu bar, click **Policies > Protection > Schedules**.
2. Click **Add** to start the **Add Schedule** wizard.
3. In the wizard, select the **Weekly** schedule option and complete the wizard to create the new Weekly schedule.

If you do not want to schedule weekly backups, continue past the Weekly Events property sheet without specifying a day or time.

Result

Your new schedule is listed on the Schedules tab.

Related concepts

[Decisions to make before adding a schedule](#) on page 47

Related references

[Administrator roles and capabilities](#) on page 44

Backing up unprotected physical storage objects

This workflow shows you how to create a backup for unprotected data. You can identify unprotected volumes, add them to an existing dataset, and schedule backups using both the OnCommand console and the N series Management Console.

Before you begin

You must have installed the OnCommand console and N series Management Console.

You must be familiar with the protection concepts of a dataset, protection policy, and backup schedule. For more information about these concepts, see the OnCommand console Help.

About this task

This workflow is based on the assumption that you start in OnCommand console.

During this task, the OnCommand console launches N series Management Console. Depending on your browser configuration, you can return to the OnCommand console by using the Alt-Tab key combination or clicking the OnCommand console browser tab. After the completion of this task, you can leave N series Management Console open, or you can close it to conserve bandwidth.

After you finish

After you have completed the tasks involved in creating a dataset, including assigning a local policy and storage service, you can perform an on-demand backup to ensure that no problems exist with the dataset configuration. For more information about how to perform an on-demand backup, see the OnCommand console Help.

Steps

1. [Identify unprotected physical storage objects](#) on page 15
Identifying unprotected physical storage objects in the OnCommand console enables you to add them to an existing dataset in N series Management Console.
2. [Add unprotected data to an existing dataset](#) on page 15
Adding unprotected data to an existing dataset in N series Management Console enables you to protect that data and change protection policies based on your requirements.

Identifying unprotected physical storage objects

You can locate your unprotected physical storage objects so that you can add them to existing datasets or to create a new dataset for the purpose of protection.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

About this task

During this task, the OnCommand console launches N series Management Console. Depending on your browser configuration, you can return to the OnCommand console by using the Alt-Tab key combination or clicking the OnCommand console browser tab. After the completion of this task, you can leave N series Management Console open, or you can close it to conserve bandwidth.

Steps

1. If the Dashboard is not visible, click the **Dashboard** tab.
2. Locate the **Unprotected Data** dashboard panel.
3. Click **Volumes**.

OnCommand console takes you to the N series Management Console Unprotected Data window.

Related references

[Administrator roles and capabilities](#) on page 44

Adding unprotected data to an existing dataset

You can add unprotected data to an existing dataset, even if the dataset is currently unprotected. You can browse and add an entire host, or individual aggregates, volumes, qtrees, virtual machines, or Open Systems SnapVault directories on the host.

Before you begin

- Determine the name of the dataset to which you want to add the resource.
- You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

About this task

- When you add an unprotected qtree or volume to a dataset that has an assigned protection policy, the N series Management Console data protection capability starts an initial baseline transfer of the data from the primary to the secondary node of the dataset.

Note: An initial baseline transfer of existing qtree or volume data might require many times more bandwidth and time than subsequent incremental transfers of that data will require.

- If you add an unprotected resource to a dataset that does not have an assigned protection policy, the data is unprotected until you assign a protection policy to the dataset and the dataset conforms to that policy.
- The N series Management Console operations described in this topic do not apply to virtual object type datasets. To carry out similar operations on virtual object type datasets, use the OnCommand console.

Steps

1. From the menu bar, click **Data > Unprotected Data > Resources**.

2. Select the resource that you want to add to a dataset:

- a) Select one or more hosts that contain the data that you want to protect.

The aggregates, volumes, and qtrees contained in that host are displayed in the bottom left pane, in a hierarchical list.

- b) (Optional) Use the **View Resources** buttons to select which resource type to display, then select the specific aggregates, volumes, qtrees, hosts, or Open Systems SnapVault directories that contain the unprotected data that you want to add to a dataset.

3. Click **Add to existing Dataset**.

A dialog box opens, from which you can select an existing dataset. This list includes both protected and unprotected datasets.

4. Verify that the existing dataset contains the new resource by viewing the dataset member list in the **Datasets** window **Overview** tab.

Result

The qtrees, volumes, aggregates, or hosts that you added to the dataset are displayed in the Datasets window and are no longer displayed in the hierarchy or host lists in the Unprotected Data window Resources tab.

After you finish

To protect the data in the dataset, you must assign a protection policy to the dataset.

Related references

[Administrator roles and capabilities](#) on page 44

Restoring a single file from a backup

This workflow shows you how to restore a single file from a backup made in a VMware environment. You might want to do this if a file, virtual machine, or database has been corrupted. You can locate the backup, mount, and unmount the backup using OnCommand console, but you must use the vSphere Client to restore the file.

Before you begin

You must have installed and logged into the following software:

- OnCommand console
- vSphere Client

You must be familiar with the protection concepts of a dataset, protection policy, and backup schedule. For more information about these concepts, see the OnCommand console help.

About this task

VMware ESX 4.x currently does not support hot-removal of LUNs. Hot-removal of LUNs can also affect restore unmounting processes within backup and recovery, as well as destroy datastores within provisioning and cloning. For more information, see VMware KB article 1015084.

Steps

1. [Locate a backup](#) on page 18
You must first locate the backup in OnCommand console that contains the file that you want to restore. This enables you to use the backup version containing the information you want to restore.
2. [Mount a backup](#) on page 19
You can mount a backup using OnCommand console. After the backup is mounted on an ESX server, you can locate the file that you want to restore. Using the VMware application, you can restore the file you have selected.
3. [Restore a single file using VMware](#) on page 20
After you have mounted a backup using OnCommand console, you can use the VMware application to restore the single file. After the restore is complete, you must unmount the backup.
4. [Unmount a backup](#) on page 20
After you have located the file and restored it, you can use OnCommand console to unmount the backup. Unmounted backups are not deleted.

Locating specific backups (7-Mode only)

You can locate a specific backup copy by searching for different criteria such as names and descriptions. After you locate a backup, you can then restore it, view the status of it, or delete it.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

About this task

You can locate a specific backup copy by searching one of the following criteria:

- Whole or partial backup description
- Partial current name of a primary object in the backup
- Partial name of a virtual object when the backup was taken
- UUID of a virtual object within the backup

Steps

1. Click the **View** menu, then click the **Backups** option.
2. From the **Backups** tab, locate the **Search** field, and type in all or part of the backup description or version.

You can locate multiple backup versions by inserting a comma between search terms and you can clear the search field to view all backups.

3. Click **Find**.

Related references

[Administrator roles and capabilities](#) on page 44

Mounting backups in a VMware environment from the Backups tab (7-Mode only)

You can mount existing backups onto an ESX server for backup verification prior to completing a restore operation or to restore a virtual machine to an alternate location. All the datastores and the virtual machines within the backup are mounted to the ESX server that you specify.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

About this task

In the OnCommand console Backups tab, deleting a mirror source backup copy prevents you from mounting its partner mirror destination backup copy. For a Mirror-generated destination backup copy to be mountable, its associated mirror source backup copy must still exist on the source node.

Steps

1. Click the **View** menu, then click the **Backups** option.
2. In the **Backups** tab, select an unmounted backup that you want to mount.
3. Click **Mount**.
4. In the **Mount Backup** dialog box, select from the drop-down list the name of the ESX server to which you want to mount the backup.

You can only mount one backup each time and you cannot mount a mounted backup.

5. Click **Mount**.

A dialog box appears with a link to the mount job and when you click the link, the Jobs tab appears.

After you finish

You can monitor the status of your mount and unmount jobs in the Jobs tab.

Related concepts

[Guidelines for mounting or unmounting backups in a VMware environment \(7-Mode only\)](#)
on page 55

Related references

[Administrator roles and capabilities](#) on page 44

Using VMware to restore a single file

After you have mounted a backup, you can use VMware to restore a single file.

About this task

For more detailed information about the steps performed in this task, see Knowledge Base article 1011968.

Steps

1. Find the VMDK that you need from within the mounted backup datastore.
2. Attach that VMDK to the virtual machine to which you need to restore the file.
3. If the partitions on the newly mounted VMDK do not appear as drive letters in Windows, bring the disks online in the Windows Disk Management Utility.
4. Restore the file by dragging it to the appropriate location, using Windows Explorer or other file browser.
5. Detach the backup VMDK from the virtual machine.

After you finish

You must now unmount the backup in the OnCommand console.

Related information

IBM N series support website: www.ibm.com/storage/support/nseries

Unmounting backups in a VMware environment from the Backups tab (7-Mode only)

After you are done using a mounted backup for verification or to restore a virtual machine to an alternate location, you can unmount the mounted backup from the ESX server.

When you unmount a backup, all the datastores in that backup are unmounted and can no longer be seen from the ESX server that you specify.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

About this task

If there are virtual objects in use from the previously mounted datastores of a backup, the unmount operation fails. You must manually clean up the backup prior to mounting the backup again because its state reverts to not mounted.

If all the datastores of the backup are in use, the unmount operation fails but this backup's state changes to mounted. You can unmount the backup after determining the datastores are not in use.

Steps

1. Click the **View** menu, then click the **Backups** option.
2. In the **Backups** tab, select a mounted backup to unmount.
3. Click **Unmount**.
4. At the confirmation prompt, click **Yes**.

A dialog box opens with a link to the unmount job and when you click the link, the Jobs tab appears.

After you finish

If the ESX server becomes inactive or restarts during an unmount operation, the job is terminated and the mount state remains mounted and the backup stays mounted on the ESX server.

You can monitor the status of your mount and unmount jobs in the Jobs tab.

Related concepts

[*Guidelines for mounting or unmounting backups in a VMware environment \(7-Mode only\)*](#)
on page 55

Related references

[*Administrator roles and capabilities*](#) on page 44

Supporting delegated management for virtual infrastructure administration

This workflow shows you how to delegate protection of virtual objects to a virtual infrastructure administrator. You delegate control to the virtual infrastructure administrator when there are two or more administrators in your environment and you must separate system responsibilities.

Before you begin

You must have installed the OnCommand console and Operations Manager console.

You must have already created the storage services that you want to allow the virtual infrastructure administrator to use.

You must be familiar with the concepts of a virtual dataset and role based access control management (RBAC). For more information about dataset concepts, see the OnCommand console help. For more information about RBAC concepts, see the Operations Manager console help.

About this task

Typically, the storage administrator administers the system and has full control. The virtual administrator administers only the virtual portion of the environment, including creating, backing up, and restoring virtual machines. You can give a virtual infrastructure administrator focused, appropriate control without full access to the entire system.

Steps

1. [Launch the Operations Manager console](#) on page 23
Creating an account for the virtual infrastructure administrator is performed in the Operations Manager console. You can launch the Operations Manager console from the OnCommand console.
2. [Create an administrative role](#) on page 23
This administrative role contains the virtual objects and permissions that you will later assign to a user account for the virtual administrator. Before adding the virtual objects and permissions to the role, you must create an empty group, protection policies, and storage services for the virtual administrator.
3. [Create an administrative user](#) on page 24
Creating a user account for a virtual infrastructure administrator allows you to assign a role containing virtual objects and permissions to the virtual administrator.
4. [Create a group](#) on page 25
This empty group contains the operations and permissions specific to the virtual administrator, along with datasets and local policies that the virtual administrator will create later.
5. [Assign permissions to an administrative role](#) on page 25
After you have created an administrative role, group, application and protection policies, and storage services for the virtual administrator, you can assign the operations and permissions for

the administrative user. For information about which access permissions to assign to an administrative role for a virtual administrator, see [Access permissions for Virtual Infrastructure Administrator role](#).

6. [Edit the group membership to include the vCenter member](#) on page 27

You must edit the group membership for a virtual infrastructure administrator to include the vCenter member.

Related concepts

[What a dataset is \(7-Mode only\)](#) on page 58

Related references

[Access permissions for the Virtual Infrastructure Administrator role](#) on page 26

Launching the Operations Manager console

You can launch the Operations Manager console from the OnCommand console to perform many of your physical storage tasks.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

About this task

During this task, the OnCommand console launches the Operations Manager console. Depending on your browser configuration, you can return to the OnCommand console by using the Alt-Tab key combination or clicking the OnCommand console browser tab. After the completion of this task, you can leave the Operations Manager console open, or you can close it to conserve bandwidth.

Step

1. Click the **File** menu, then click **Operations Manager**.

The Operations Manager console opens in a separate browser tab or window.

Creating an administrative role

You can create a role for an administrator from the Setup menu in Operations Manager.

Steps

1. Select **Roles** from the **Setup** menu.

2. Enter the name of the role and the description of the role that you want to add.
3. Select any of the pre-created roles that you want to assign to this role; otherwise, select **None**.
4. Click **Add Role**.

Adding administrative users

You can create administrator accounts from the Operations Manager console. Administrator accounts are either an individual administrator or a group of administrators.

Before you begin

The DataFabric Manager server user must be a local operating system user, or a domain user reachable by LDAP.

Steps

1. Log in to the Administrator account.
2. In the Operations Manager console, click **Setup > Administrative users**.
3. Type the name for the administrative user or domain name for the group of administrators.
When you add the user, they must be available locally.
4. If you have already created a role that you want to assign to this user or group of users, select the role in the left column of the displayed table and use the arrow button to move the role to the column on the right.
Roles in the column on the right are assigned to the user that you are creating.
5. Type the email address for the administrator or administrator group.
6. Enter the pager number for the administrator or administrator group.
7. Click **Add**.

In Windows, when you add a user to the Administrators group, the user gets added as a local admin user.

After you finish

If you have not created a role for the user you created, you must create a role.

Creating groups

You can create a new group from the Edit Groups page. You can group objects based on storage systems at a location, or all file systems that belong to a specific project or group in your organization.

Before you begin

To create a group, you must be logged in as an administrator with a role having database write capability on the parent group. To create a group directly under the Global group, the administrator must have a role with Database Write capability on the Global group.

Steps

1. From the **Control Center**, click the **Edit Groups**.
2. In the **Group Name** field, type the name of the group you want to create.
See “Naming conventions” for groups.
3. From the list of groups, select the parent group for the group you are creating.
You might need to expand the list to display the parent group you want.
4. Click **Add**.

Result

The new group is created. The Current Groups list in the left-pane area is updated with the new group. You might need to expand the Current Groups list to display the new group.

Assigning permissions to an administrative role

You can assign to a role permission specific to a user's task.

Before you begin

You must have created the group, storage services, and application policies, if appropriate, for a virtual infrastructure administrative role.

Steps

1. From the **Setup** menu, select **Roles**.
2. Select the role to which you want to add capabilities
3. Click **Add Capabilities**.

4. Select each resource for which you want to assign administrative permissions, then check the appropriate permissions for each corresponding operation.
See [Access permissions for the Virtual Infrastructure Administrator](#) on page 26 for a list of permissions you need to assign to the virtual administrative role.
5. Click **OK**.
6. Optionally, to copy capabilities from an existing role, select that role from the **Inherit Capabilities** list and click \>> to move the role to the list at the right.
7. Click **Update**.

Access permissions for the Virtual Infrastructure Administrator role

When you create a virtual infrastructure administrator, you must assign specific permissions to ensure that the administrator can view, back up, and recover the appropriate virtual objects.

A virtual infrastructure administrator role must have the following permissions for the resources:

Groups	The VI administrator will need the following operation permissions for the group created for the VI administrator role:										
	<table border="0" style="width: 100%;"> <tr> <td style="padding-left: 20px;">DFM.Database</td> <td style="text-align: right;">All</td> </tr> <tr> <td style="padding-left: 20px;">DFM.BackManager</td> <td style="text-align: right;">All</td> </tr> <tr> <td style="padding-left: 20px;">DFM.ApplicationPolicy</td> <td style="text-align: right;">All</td> </tr> <tr> <td style="padding-left: 20px;">DFM.Dataset</td> <td style="text-align: right;">All</td> </tr> <tr> <td style="padding-left: 20px;">DFM.Resource</td> <td style="text-align: right;">Control</td> </tr> </table>	DFM.Database	All	DFM.BackManager	All	DFM.ApplicationPolicy	All	DFM.Dataset	All	DFM.Resource	Control
DFM.Database	All										
DFM.BackManager	All										
DFM.ApplicationPolicy	All										
DFM.Dataset	All										
DFM.Resource	Control										
Policies	The VI administrator will need the following operation permissions for each policy template, located under Local Policies that you want the virtual administrator to be able to copy:										
	<table border="0" style="width: 100%;"> <tr> <td style="padding-left: 20px;">DFM.ApplicationPolicy</td> <td style="text-align: right;">Read</td> </tr> </table>	DFM.ApplicationPolicy	Read								
DFM.ApplicationPolicy	Read										
Storage services	The VI administrator will need the following operation permissions for each of the storage services that you want to allow the VI administrator to use:										
	<table border="0" style="width: 100%;"> <tr> <td style="padding-left: 20px;">DFM.StorageService</td> <td style="text-align: right;">Attach, read, detach, and clear</td> </tr> </table>	DFM.StorageService	Attach, read, detach, and clear								
DFM.StorageService	Attach, read, detach, and clear										
Protection Policies	These are the policies contained within the storage services that you selected above:										
	<table border="0" style="width: 100%;"> <tr> <td style="padding-left: 20px;">DFM.Policy</td> <td style="text-align: right;">All</td> </tr> </table>	DFM.Policy	All								
DFM.Policy	All										

Editing group membership

You can add members or delete members from the Edit Group dialog box.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

If you have created a new group for a virtual infrastructure administrator, you must have added vCenter to the group membership using the OnCommand console GUI.

Steps

1. Click the **Administration** menu, then click the **Groups** option.
2. From the **Groups** tab, select the group you want to modify.
3. Click **Edit**.
4. In the **Edit Group** dialog box, click the **Group Member** tab.
5. Select the type of member you want to add or delete from the Member Type field.
6. Perform one of the following steps to add or delete a group member:

To do this ...	Take this action ...
Add a member to the group ...	Move a member type to the Selected Members list.
Remove a member from the group ...	Move a member type to the Available Members list.

7. Click **OK**.

Resolving issues with virtual objects

This workflow shows you how to resolve an issue for a virtual object by monitoring and troubleshooting the virtual object, and correcting errors in the virtual object policy. You need to troubleshoot and resolve issues for your virtual objects if you discover problems while monitoring the Dataset Status dashboard panel.

Before you begin

You must have installed the OnCommand console.

About this task

This workflow is based on the assumption that a virtual object is being affected by two jobs that have been scheduled so that they overlap. This workflow assumes that you are familiar with the concepts of a virtual dataset and a protection policy. For more information about these concepts, see the OnCommand console Help.

Steps

1. [Monitor dashboard panels](#) on page 28
You can view the dashboard panel to monitor virtual datasets and identify datasets that are experiencing problems.
2. [Monitoring dataset status \(7-Mode only\)](#) on page 29
You can monitor the status of your datasets for possible errors.
3. [Edit a dataset containing virtual objects to reschedule or modify local backup jobs](#) on page 30
You can correct a problem with a dataset policy by editing the local settings of the dataset.

Related concepts

[What a dataset is \(7-Mode only\)](#) on page 58

[Role of protection policies in dataset management \(7-Mode only\)](#) on page 57

Monitoring dashboard panels

You can use the dashboard panels to monitor your physical storage, logical storage, virtual storage and non-storage objects.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

Steps

1. Log in to the OnCommand console. By default, the Dashboard tab is displayed.
2. To view details about any of the information displayed in the dashboard panels, click the panel heading to display the relevant OnCommand console tab.

You can also click any hypertext links in the individual dashboard panels to view detailed information.

Related references

[Administrator roles and capabilities](#) on page 44

Monitoring dataset status (7-Mode only)

You can monitor the status of your datasets for possible errors.

Before you begin


You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

Steps

1. Click the **View** menu and click the **Datasets** option.
2. In the **Datasets** tab, select the dataset whose status you want to monitor.

The Overview area of the Datasets tab displays the protection, conformance, resources, and space statuses of the selected dataset.

3. To view status details, click the  button, if displayed.

If the  button is displayed for resources, clicking it displays a dialog box that lists events related to warning-level or critical-level resource issues. You can use the Acknowledge button to mark an event as acknowledged. If you take actions outside of the dialog box that resolves an event issue, you can use the Resolve button to mark that event as resolved.

4. To view details about secondary or tertiary nodes, click the corresponding tabs for these nodes.

Related references

[Administrator roles and capabilities](#) on page 44

Editing a dataset containing virtual objects to reschedule or modify local backup jobs (7-Mode only)

You can modify the schedule of local backup jobs that are configured in the local policy assigned to a dataset containing VMware objects.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

About this task

To reschedule or modify the local backup jobs associated with the local policy of a dataset of virtual objects, you can edit the Local Policy settings.

Steps

1. Click the **View** menu, and then click the **Datasets** option.
2. In the **Datasets** tab, select the dataset on which you want to schedule and configure local backups and click **Edit**.
3. In the **Edit Dataset** dialog box, locate the **Local settings** option and click >.
4. Modify the local backup jobs as needed.
5. After you finish changing the schedule for the local policy to this dataset, click **OK**.

Result

Any local policy modification that you completed is applied to the local protection of the virtual objects in all datasets that use that local policy.

Related references

[Administrator roles and capabilities](#) on page 44

Resolving issues with physical objects

This workflow shows you how to resolve an issue for a physical object. You can monitor your physical objects by using the dashboard, and manually launch N series Management Console to resolve the issue.

Before you begin

You must have installed OnCommand console and N series Management Console.

You must be familiar with the protection concepts of a dataset and secondary space management. For more information about these concepts, see the OnCommand console Help.

About this task

This workflow is based on the assumption that you are resolving an issue with a volume that is going to run out of space soon. This is important to solve because applications that are run on this volume could run out of space and stop operating.

This workflow begins in the OnCommand console, but you must also manually launch N series Management Console to resolve the problem.

This workflow assists with one type of issue resolution. Depending on what your issue is, there are additional ways to resolve it.

During the performance of this workflow, you must manually launch N series Management Console, with the same server as OnCommand Unified Manager. The N series Management Console does not open automatically.

Steps

1. [Monitor dashboard panels](#) on page 28
You can use the Full Soon Storage dashboard panel to find volumes that are running out of space. After identifying the volume, you must manually start N series Management Console.
2. [Start Secondary Space Management wizard](#) on page 32
Using the Secondary Space Management wizard enables you to resize a volume that is running out of space. In addition to resizing, you can also delete backups, perform deduplication, or migrate volumes to another aggregate. Before using this wizard, you should understand how the Secondary Space Management wizard works.

Monitoring dashboard panels

You can use the dashboard panels to monitor your physical storage, logical storage, virtual storage and non-storage objects.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

Steps

1. Log in to the OnCommand console. By default, the Dashboard tab is displayed.
2. To view details about any of the information displayed in the dashboard panels, click the panel heading to display the relevant OnCommand console tab.

You can also click any hypertext links in the individual dashboard panels to view detailed information.

Related references

[Administrator roles and capabilities](#) on page 44

Starting secondary space management

You can identify volumes that have space issues, plan actions to resolve those issues, and implement the selected actions by using the Secondary Space Management wizard.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

Steps

1. From the menu bar, click **Hosts > Aggregates**.
2. Select an aggregate and click **Manage space** to start the **Secondary Space Management** wizard.
3. Complete the steps in the wizard to plan and implement tasks to reclaim space on the selected aggregate.

Related concepts

[How the Secondary Space Management wizard works](#) on page 56

Related references

Administrator roles and capabilities on page 44

Changing the storage service for a virtual dataset

This workflow shows you how to change the storage service for a virtual dataset by using the Change Storage Service wizard.

Before you begin

You must have installed the OnCommand console and N series Management Console.

About this task

This workflow is started from the OnCommand console, but it is performed in the N series Management Console and assumes that you are familiar with the concepts of a virtual dataset and a protection policy. For more information about these concepts, see the OnCommand console Help.

Changing a storage service for a virtual dataset is different from changing it for a physical storage dataset. This task can be started from the OnCommand console, but the system launches N series Management Console to perform the task.

You must change the storage service for your virtual dataset if you change your requirements for your dataset protection. For example, you need to change your storage service if you want to add mirroring to SnapVault backups or upgrade your service from bronze to gold. If you change the underlying topology of your system, you also need to change the provisioning, and so forth. For example, if you change from backup to mirror and then back to backup, you must reprovision and possibly rebaseline. Reprovisioning, rebaselining, and so forth can impact performance, and you should consider this impact prior changing your storage service.

Changing a storage service for datasets of storage objects (7-Mode only)

You can use the Change Storage Service wizard to change a storage service for datasets of storage objects that already have another storage service attached.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

The following conditions must exist:

- The new storage service that you want to attach is available in the group in which you want to locate that dataset.
- All datasets to which you want to attach the new storage service are currently attached to the same current storage service.

About this task

The Change Storage Service wizard allows you to select an alternative storage service, presents you with possible node remapping alternatives along with rebaselining requirements for each alternative, carries out a dry run of your request, and then implements your request upon your approval.

During this task, the OnCommand console launches N series Management Console. Depending on your browser configuration, you can return to the OnCommand console by using the Alt-Tab key combination or clicking the OnCommand console browser tab. After the completion of this task, you can leave N series Management Console open, or you can close it to conserve bandwidth.

Steps

1. Click the **View** menu and click the **Datasets** option.
2. In the **Datasets** tab, select the dataset whose storage service you want to change, click **More** then select **Storage Service** and **Change** to start the **Change Storage Service** wizard.
3. After you complete each property sheet in the wizard, click **Next**.
4. Confirm the details of the storage service and click **Finish**.
5. Press **Alt-Tab** or click the OnCommand console browser tab to return to the OnCommand console.
6. Refresh your browser to update the OnCommand console with the changes you made.

Result

The selected datasets are listed in the datasets table with their newly attached storage service named in the storage service column.

Related references

[Administrator roles and capabilities](#) on page 44

Resolving dataset conformance issues

This workflow shows you how to resolve dataset nonconformance. Resolving a dataset's conformance issues helps ensure that its protection jobs are fully successful.

Before you begin

You must have installed the OnCommand console.

You must be familiar with dataset virtual objects, storage services, protection policies, and conformance. For more information about these concepts, see the OnCommand console Help.

About this task

When the OnCommand console conformance checker encounters a dataset condition that is nonconformant with the provisioning and protection policies of the dataset's assigned storage service, it first attempts to correct that condition automatically. However, if the condition requires user consent to correct, or if the condition must be corrected manually, then the conformance checker assigns nonconformant status to the dataset. If the dataset is tagged with nonconformant status, operations associated with the assigned storage service's protection and provisioning policies cannot be fully executed until the nonconformant condition is resolved.

If the conformance checker encounters the nonconformant condition during initial storage service assignment or in an existing dataset later on, it flags that condition for the administrator in the Datasets tab. The administrator can select the nonconformant dataset and display the Conformance Details dialog box to view that dataset's conformance issues.

The conformance checker actively monitors a dataset's resources during dataset storage service assignment and at intervals thereafter to make sure that the resource configuration remains in conformance with (is able to support) the protection policies and provisioning policies that are included in the storage service assigned to it.

Steps

1. [List the nonconformant dataset and view the details](#) on page 37
OnCommand console routinely scans your datasets to ensure that they are in conformance with their assigned policies. You can list a dataset that has fallen out of conformance with its protection and provisioning policies, then view the nonconformance details.
2. [Evaluate the conformance error messages](#) on page 38
Use the Conformance Details dialog box to evaluate a dataset's conformance issues. You can evaluate the error and warning text that the dialog displays and attempt to resolve the conformance issues they describe.
3. [Resolve the conformance issues automatically without a baseline transfer](#) on page 40

If the text displayed in the Conformance Details dialog box identifies conformance issues that you can automatically resolve without a rebaseline of data, you can use the dialog box controls to do so.

4. [Resolve the conformance issues manually without a baseline transfer of data](#) on page 41
If the text displayed in the Conformance Details dialog box identifies conformance issues that you cannot or choose not to resolve automatically, you can resolve the conformance issues manually.
5. [Resolve conformance issues manually when a baseline transfer of data might be necessary](#) on page 42
If the text displayed in the Conformance Details dialog box identifies conformance issues whose automated resolution might require a baseline transfer, attempt a manual resolution before using automated resolution.

Related concepts

[What a dataset is \(7-Mode only\)](#) on page 58

[Role of protection policies in dataset management \(7-Mode only\)](#) on page 57

[Dataset concepts \(7-Mode only\)](#) on page 45

[Datasets of virtual objects \(7-Mode only\)](#) on page 47

[Why datasets fail to conform to policy \(7-Mode only\)](#) on page 58

Related references

[Descriptions of dataset conformance status \(7-Mode only\)](#) on page 54

Listing nonconformant datasets and viewing details (7-Mode only)


You can list existing datasets that have become nonconformant with their protection and provisioning policies. Then, you can view and attempt to resolve their nonconformance. Conformance issues can prevent completion of dataset protection and provisioning operations.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

Steps

1. Click the **View** menu and click the **Datasets** option to display the **Datasets** tab.
2. In the **Datasets** tab click the column header labeled **Conformance Status** and select **Nonconformant**.

3. If the **Datasets** tab lists a dataset with nonconformant status, select that dataset to display its Details area.
4. In the selected dataset's Details area, click the  button.
The Conformance Details dialog box displays the results of the most recent conformance check and suggestions for resolving the issues encountered.

After you finish

After you display the Conformance Details dialog box, you must address and resolve the issues that are indicated by its warning and error messages.

Related references

[Administrator roles and capabilities](#) on page 44


Evaluating conformance error text

You can use the Conformance Details dialog box to evaluate a dataset's conformance issues. You can evaluate the error and warning text that the dialog box displays and attempt to resolve the conformance issues that it describes.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

About this task

- This procedure assumes you are viewing the Conformance Details dialog box that you displayed by selecting a nonconformant dataset in the Datasets tab and clicking  after its Conformance status display.
- Because of the probable time and bandwidth required for a baseline transfer completion, a resolution that avoids a baseline transfer of data is preferable to a resolution that triggers one.
- The Conformance Details dialog box describes the nonconformance condition under the headings Information, Error, Action, Reason, and Suggestion.

Steps

1. Evaluate any labeled text that the dialog box displays:

Information Information text indicates configuration operations that the OnCommand console successfully completed without conformance issues.

Error	Error text indicates the configuration operations that the OnCommand console is unable to perform on this dataset due to conformance issues.
Action	Action text indicates what the OnCommand Unified Manager conformance engine did to discover the conformance issue.
Reason	Reason text indicates the probable cause of the conformance issue.
Suggestion	Suggestion text indicates a possible way of resolving the conformance issue. If the text is highlighted in yellow, a baseline transfer of data is involved.

2. Based on the dialog box text, decide the best way to resolve the conformance issue:
 - If the dialog box text indicates that the OnCommand Unified Manager conformance monitor cannot automatically resolve the conformance issue, resolve this issue manually.
For details, see "Resolving conformance issues manually without a baseline transfer of data."
 - If yellow Suggestion text indicates that automatically resolving the conformance issues requires a baseline transfer of data, first attempt to resolve this issue manually.
If unsuccessful, consider resolving the issue automatically, even if doing so requires a baseline transfer of data.
For details, see, "Resolving conformance issues manually when a baseline transfer of data might be necessary."
 - If the Suggestion text indicates that the OnCommand Unified Manager conformance monitor can resolve the conformance issue automatically without reinitiating a baseline transfer of data, first consider resolving the issue manually.
If unsuccessful, resolve the issue automatically.
For details, first see, "Resolving conformance issues manually without a baseline transfer of data."
Then, if necessary, see "Resolving conformance issues automatically without a baseline transfer of data."

After you finish

Proceed to the appropriate task to resolve the dataset conformance issues.

Related references

[Administrator roles and capabilities](#) on page 44


Resolving conformance issues automatically without a baseline transfer of data (7-Mode only)

If the text displayed in the Conformance Details dialog box identifies conformance issues that you can automatically resolve without reinitializing a baseline transfer of data, you can use the dialog box controls to do so.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

About this task

- Because of the probable time and bandwidth required for a baseline transfer completion, a resolution that avoids a baseline transfer of data is preferable to a resolution that triggers one.
- This procedure assumes that you are viewing the Conformance Details dialog box that you displayed by selecting a nonconformant dataset Datasets tab and clicking  after its nonconformant status display.

Steps

1. In the **Conformance Details** dialog box, read the text to determine the ability of the conformance engine to automatically resolve the nonconformant condition without reinitializing a baseline transfer of data.
2. If the text suggests that a simple automatic resolution is possible, click **Conform**.

The OnCommand console conformance engine closes the Conformance Details dialog box and attempts to reconfigure storage resources to resolve storage service protection and provisioning policy conformance issues automatically.

3. Monitor the conformance status on the **Datasets** tab for one of the following values:
 - Conformant: The conformance issue is resolved.
 - Nonconformant: The conformance issue is not resolved. Consider manual resolution of the issue.

After you finish

After you achieve dataset conformant status, continue with the operation that required the dataset to be conformant.

Related references

[Administrator roles and capabilities](#) on page 44


Resolving conformance issues manually without a baseline transfer of data (7-Mode only)

If the text displayed in the Conformance Details dialog box identifies conformance issues that you cannot resolve automatically, you can resolve the conformance issue manually.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

About this task

- Because of the probable time and bandwidth required for a baseline transfer completion, a resolution that avoids a baseline transfer of data is preferable to a resolution that triggers one.
- This procedure assumes you are viewing the Conformance Details dialog box that you displayed by selecting a nonconformant dataset Datasets tab and clicking  after its nonconformant status display.

Steps

1. In the **Conformance Details** dialog box, confirm that the messages indicate that the conformance issues cannot be resolved automatically.
2. Using the conformance messages, determine what is causing the nonconformance problem and attempt to correct the condition manually.

You might need to log in to another GUI or CLI console to resolve the issues.

3. After you have attempted to correct the condition, wait at least one hour for the conformance monitor to update the dataset's conformance status.
4. Return to the **Conformance Details** dialog box and click **Test Conformance** to determine if the conformance issue is resolved.

If the conformance issue is resolved, the Conformance Details dialog box does not display the **Conform** button.

5. If the conformance issue is resolved, click **Cancel**.
6. If the conformance issue is not resolved, repeat Steps 2, 3, and 4.

After you finish

After you achieve dataset conformant status, continue with the operation that required the dataset to be conformant.

Related references

[Administrator roles and capabilities](#) on page 44


Resolving conformance issues manually when a baseline transfer of data might be necessary (7-Mode only)

If the text displayed in the Conformance Details dialog box identifies conformance issues whose automated resolution might require a baseline transfer of data, attempt a manual resolution before using automated resolution.

Before you begin

You must be authorized to perform all the steps of this task; your RBAC administrator can confirm your authorization in advance.

About this task

- Because of the probable time and bandwidth required for baseline transfer completion, a resolution that avoids a baseline transfer of data is preferable to a resolution that triggers one.
- This procedure assumes you are viewing the Conformance Details dialog box that you displayed by selecting a nonconformant dataset in the Datasets tab and clicking  after its nonconformant status display.

Steps

1. In the **Conformance Details** dialog box, confirm that warning text is displayed that indicates that a reinitialized baseline transfer of data might be required.

You should try to resolve the conformance issues manually before initializing a time-consuming baseline transfer of your data.

2. Using the conformance messages, determine what is causing the conformance problem and attempt to correct the condition manually.

You might need to log in to another GUI or CLI console to resolve the issues.

3. After you have attempted to correct the condition, wait at least one hour for the conformance monitor to update the dataset's conformance status.
4. Return to the **Conformance Details** dialog box and click **Test Conformance** to determine if the conformance issue is resolved.

If the conformance issue is resolved, the Conformance Details dialog box does not display the **Conform** button.

5. If the conformance issue is not resolved, click **Conform** to attempt automated resolution and initiate a rebaseline of your data.

After you finish

After you achieve dataset conformant status, continue with the operation that required the dataset to be conformant.

Related references

[*Administrator roles and capabilities*](#) on page 44

Appendix of additional information

You can use the additional information provided to assist you in completing the workflows in this guide.

Administrator roles and capabilities

The RBAC administrator roles determine the tasks you can perform in the OnCommand console.

One or more capabilities must be specified for every role, and you can assign multiple capabilities if you want the administrator to have more control than a specific role provides. For example, if you want an administrator to perform both the backup and restore operations, you can create and assign to the administrator a single role that has both of these capabilities.

You can use the Operations Manager console to create new roles and to customize the default global roles provided by the DataFabric Manager server and the client applications. For more information about configuring RBAC, see the *OnCommand Unified Manager Operations Manager Administration Guide*.

Note: If you want a role with global host service management capability, create a role with the following properties:

- The role inherits from the GlobalHostService role.
- The role includes the DFM.Database.Read operation on a global level.

Note: A user who is part of the local administrators group is treated as a super-user and automatically granted full control.

Default global roles

GlobalApplicationProtection	Enables you to create and manage application policies, create datasets with application policies for local backups, use storage services for remote backups, perform scheduled and on-demand backups, perform restore operations, and generate reports.
GlobalBackup	Enables you to initiate a backup to any secondary volume and ignore discovered hosts.
GlobalDataProtection	Enables you to initiate a backup to any secondary volume; view backup configurations, events and alerts, and replication or failover policies; and import relationships into datasets.
GlobalDataset	Enables you to create, modify, and delete datasets.

GlobalDelete	Enables you to delete information in the DataFabric Manager server database, including groups and members of a group, monitored objects, custom views, primary and secondary storage systems, and backup relationships, schedules, and retention policies.
GlobalHostService	Enables you to authorize, configure, and unregister a host service.
GlobalEvent	Enables you to view, acknowledge, and delete events and alerts.
GlobalFullControl	Enables you to view and perform any operation on any object in the DataFabric Manager server database and configure administrator accounts. You cannot apply this role to accounts with group access control.
GlobalMirror	Enables you to create, destroy, and update replication or failover policies.
GlobalRead	Enables you to view the DataFabric Manager server database, backup and provisioning configurations, events and alerts, performance data, and policies.
GlobalRestore	Enables you to restore the primary data to a point in time or to a new location.
GlobalWrite	Enables you to view or write both primary and secondary data to the DataFabric Manager server database.
GlobalResourceControl	Enables you to add members to dataset nodes that are configured with provisioning policies.
GlobalProvisioning	Enables you to provision primary dataset nodes and can attach resource pools to secondary or tertiary dataset nodes. The GlobalProvisioning role also includes all the capabilities of the GlobalResourceControl, GlobalRead, and GlobalDataset roles for dataset nodes that are configured with provisioning and protection policies.
GlobalPerfManagement	Enables you to manage views, event thresholds, and alarms apart from viewing performance information in Performance Advisor.

Dataset concepts (7-Mode only)

Associating data protection, disaster recovery, a provisioning policy, or a storage service with a dataset enables storage administrators to automate tasks, such as assigning consistent policies to primary data, propagating policy changes, and provisioning new volumes, qtrees, or LUNs on primary and secondary dataset nodes.

Configuring a dataset combines the following objects:

Dataset of physical storage objects	<p>For protection purposes, a collection of physical resources on a primary node, such as volumes, flexible volumes, and qtrees, and the physical resources for copies of backed-up data on secondary and tertiary nodes.</p> <p>For provisioning purposes, a collection of physical resources, such as volumes, flexible volumes, qtrees, and LUNs are assigned to a dataset node. If the protection policy establishes a primary and one or more non-primary nodes, each node of the dataset is a collection of physical resources that might or might not be provisioned from the same resource pool.</p>
Dataset of virtual objects	<p>A collection of supported VMware virtual objects that reside on storage systems. These virtual objects can also be backed up locally and backed up or mirrored on secondary and tertiary nodes.</p>
Resource pool	<p>A collection of physical resources from which storage is provisioned. Resource pools can be used to group storage systems and aggregates by attributes, such as performance, cost, physical location, or availability. Resource pools can be assigned directly to the primary, secondary, or tertiary nodes of datasets of physical storage objects.</p> <p>They can be assigned indirectly both to datasets of virtual objects and to datasets of physical storage objects through a storage service.</p>
Data protection policy	<p>A set of rules that define how to protect primary data on primary, secondary or tertiary storage, as well as when to create copies of data and how many copies to keep.</p> <p>Protection policies can be assigned directly to datasets of physical storage objects. They can be assigned indirectly to both datasets of virtual objects and to datasets of physical storage objects through a storage service.</p>
Provisioning policy	<p>A set of rules that define how to provision storage for the primary or secondary dataset nodes, and provides rules for monitoring and managing storage space and for allocating storage space from available resource pools.</p> <p>Provisioning policies can be assigned directly to the primary, secondary, or tertiary nodes of datasets of physical storage objects. They can be assigned indirectly to both datasets of virtual objects and datasets of physical storage objects through a storage service.</p>
Storage service	<p>A single dataset configuration package that consists of a protection policy, provisioning policies, resource pools, and an optional vFiler template (for vFiler unit creation). You can assign a single uniform storage service to datasets with common configuration requirements as an alternative to separately assigning the same protection policy, provisioning policies, and resource pools, and to set up similar vFiler unit attachments to each of them.</p> <p>The only way to configure a dataset of virtual objects with secondary or tertiary backup and mirror protection and provisioning is by assignment of a storage</p>

service. You cannot configure secondary storage vFiler attachments for datasets of virtual objects.

Local policy	A policy that schedules local backup jobs and designates retention periods for the local backup copies for datasets of virtual objects.
Protection or provisioning related objects	Snapshot copies, primary volumes, secondary volumes, or secondary qtrees that are generated as a result of local policy or storage service protection jobs or provisioning jobs. The OnCommand console lists related objects for each dataset on the Datasets tab.
Naming settings	Character strings and naming formats that are applied when naming related objects that are generated as a result of local policy or storage service protection jobs or provisioning jobs.

Datasets of virtual objects (7-Mode only)

The OnCommand console enables you to group VMware virtual objects that reside as data on your storage systems into datasets for purposes of data protection.

You can set up and enhance automated protection and provisioning for a dataset of virtual objects by configuring the following types of protection:

- You can assign a local policy to configure local backup job scheduling and local backup copy retention of your virtual object data.
- You can assign a storage service to configure secondary storage and tertiary storage backup and mirroring of your virtual object data.

Decisions to make before adding a schedule

Before you use the Add Schedule wizard to add a new protection or throttle schedule to your list of existing schedules, you must evaluate whether you really need a new schedule and, if so, what type it should be.

Preliminary schedule decisions

- Can the protection policies that you want to apply use their currently assigned schedules? If not, do you want to apply different daily, weekly, monthly, or throttle schedules to the primary node, backup connection, or mirror connection components of those protection policies?
- If a policy requires a different daily, weekly, monthly, or throttle schedule, can you assign an existing schedule of that type, or do you need to create a new schedule?
- If you need to create a new schedule, what is its name and description?

Daily schedule creation decisions

If you need to add a daily schedule, make the following decisions before starting the Add Schedule wizard:

- Do you want to use this daily schedule for specifying hourly backup and mirror operations and, if so, over which hours and at which time intervals?
A common practice is to schedule frequent hourly backup and mirror operations at the Primary Data node during working hours or periods of heavy data input.
- Do you want to schedule hourly backup operations at intervals shorter than an hour?
Protection configurations on which you plan to schedule backup operations at intervals of less than an hour require that both the source and destination nodes are preconfigured with the SnapMirror license.
- Do you want daily backups for retention purposes and, if so, at what times of day?
When you apply this schedule to a protection policy node, you can assign different retention durations to your hourly and daily backups. A common practice is to schedule one or two daily backup operations per day, at least one of which is during nonworking hours.
- Do you want to schedule daily mirror operations at unique times in addition to the regular hourly mirror operations that you set up in this schedule?

Weekly schedule creation decisions

If you need to add a weekly schedule, make the following decisions before starting the Add Schedule wizard:

- Do you want to apply existing daily schedules to this weekly schedule?
If so, for which days of the week (for example, Daily schedule A: Saturday through Sunday, Daily schedule B: Monday through Friday)?
The hourly and daily backups in the applied daily schedules are automatically included in the new weekly schedule.
- Do you want weekly backups for retention purposes?
When you apply this schedule to a protection policy node, you can assign different retention durations to your hourly, daily, and weekly backups.
- Do you want to schedule weekly mirror operations at unique times in addition to the times already specified for the hourly or daily mirror operations of the applied daily schedules?
- Which days and times during the week do you want to schedule weekly backup or mirror operations?
A common practice is to schedule one or two weekly backups per week.

Monthly schedule creation decisions

If you need to add a monthly schedule, make the following decisions before starting the Add Schedule wizard:

- Do you want to apply an existing daily or weekly schedule to this monthly schedule?

The hourly, daily, and weekly backup and mirror operations in the applied daily or weekly schedule are automatically included in the new monthly schedule.

If you specify a daily schedule, that schedule applies to every day of the month; if you specify a weekly schedule, that schedule applies to every week of the month.

- Do you want monthly backups for retention purposes?
When you apply this schedule to a protection policy node, you can assign different retention durations to your hourly, daily, weekly, and monthly backups.
- Do you want to schedule monthly mirror operations at unique times outside the times already specified for the hourly, daily, or weekly mirror operations of the applied daily or weekly schedules?
- Which days and times during the month do you want to schedule monthly backup or mirror operations?
A common practice is to schedule one or two monthly backups per month.

Throttle schedule creation decisions

Data protection operations can consume large amounts of network bandwidth. If you need to create a schedule to throttle the bandwidth availability used by these operations, make the following decisions before starting the Add Schedule wizard:

- Do you want to restrict the network bandwidth available for backup or mirror operations at certain times of the day?
- Do you want to prevent new backup or mirror operations completely, at certain times of the day?
Zero network bandwidth allotment prevents all new backup or mirror operations during the period and on the connection in which it is in effect.

Note: When the N series Management Console data protection capability executes a mirror operation that consists of multiple simultaneous data transfers, N series Management Console divides the total bandwidth allotted to this operation and distributes it equally to each transfer.

Decisions to make before adding datasets of physical storage objects (for protection) (7-Mode only)

Before you create a new dataset of physical storage objects, considering some preconfiguration questions can help you choose a dataset name, how related objects generated for this dataset are named, and what kind of protection and resources are assigned to the dataset's members.

Dataset properties

- Is there a dataset naming convention that you can use to help administrators easily locate and identify datasets?
Dataset names can include the following characters but cannot be only numeric:

a to z

A to Z
 0 to 9
 . (period)
 _ (underscore)
 - (hyphen)
 space

If you use any other characters when naming the dataset, they do not appear in the name.

- What is a good description of the dataset membership?
Use a description that helps someone unfamiliar with the dataset to understand its purpose.
- Who is the owner of the dataset?
- If an event on the dataset triggers an alarm, who should be contacted?
You can specify one or more individual e-mail addresses or a distribution list of people to be contacted.
- If dataset members exist in multiple time zones, which time zone do you want to use to schedule operations on the dataset?
You can specify a time zone in the wizard or use the default time zone, which is the system time zone used by the DataFabric Manager server.

Dataset naming properties

- Do you want to use the actual dataset name or a custom label in your dataset-level Snapshot copy, primary volume, secondary volume, or secondary qtree naming?
- For customizing the naming settings of object types, do you want the current default naming format to apply to one or more object types that are generated in this dataset?
If you want to customize the dataset-level naming formats for one or more object types, in what order do you want to enter the naming attributes for Snapshot copy, primary volume, secondary volume, or secondary qtree?

Group membership

- Do you need to create a group of datasets and resource pools based on common characteristics, such as location, project, or owning organization?
- Is there an existing group to which you want to add this dataset?

Resources for primary storage

Will you assign a resource pool or individual physical resources as destinations for your primary storage?

If using a resource pool, consider the following details:

- For the primary node in the dataset, which resource pool meets its provisioning requirements?
- If no resource pool meets the requirements of the primary node, you can create a new resource pool for each node at the Resource Pools window.

- Verify that you have the appropriate software licenses on the storage you intend to use.

If using individual resources, consider the following details:

- If you prefer not to use resource pools for automatic provisioning, you can select individual physical resources as members of your dataset.
- Verify that you have the appropriate software licenses on the storage you intend to use.

Protection policy

After you create a new dataset of physical objects, you protect it by running the N series Management Console Dataset Policy Change wizard to assign a protection policy.

- Which protection policy meets the requirements of the dataset?
Review the policies listed on the Protection Policies window to see if any are suitable for your new dataset.
- If no protection policy meets the requirements of your new dataset, is there a protection policy that would be suitable with minor modifications?
If so, you can copy that protection policy to create a new policy you can modify as needed for the new dataset. If not, you can run the Add Protection Policy wizard to create a new policy for the dataset.

Resources for secondary or tertiary storage

When you assign a protection policy, will you assign a resource pool or individual physical resources as destinations for your backups and mirror copies?

You do not have to assign a resource pool or physical resources to a node to create a new dataset. However, the dataset will be nonconformant with its policy until resources are assigned to each node, because the N series Management Console data protection capability cannot perform the protection specified by the policy.

If using a resource pool:

- For the secondary or tertiary nodes in the dataset, which resource pool meets their provisioning requirements?
For example, the resource pool you assign to a mirror node should contain only physical resources that would be acceptable destinations for mirror copies created of the dataset members.
- If no resource pool meets the requirements of a node, you can create a new resource pool for each node at the Resource Pools window.
- Verify that you have the appropriate software licenses on the storage you intend to use.

If using individual resources:

- If you prefer not to use resource pools for automatic provisioning, you can select individual physical resources as destinations for backups and mirror copies of your dataset.

- Verify that you have the appropriate software licenses on the storage you intend to use.

Decisions to make before assigning or changing policies

Before you assign or change a policy, you need to gather information about the dataset and policies that you want the dataset to have.

You need to gather the following information:

Protection policy

- Which protection policy meets the requirements of the dataset?
Review the policies listed on the Protection Policies window to see if any is suitable for your new dataset.
- If no protection policy meets the requirements of your new dataset, is there a protection policy that would be suitable with minor modifications?
If so, you can copy that protection policy to create a new policy you can modify as needed for the new dataset. If no suitable protection policy exists, you can run the Add Protection Policy wizard to create a new policy for the dataset.

Disaster recovery policy

- What type of disaster recovery capable protection policy do you need?
The N series Management Console data protection capability provides disaster recovery capable protection policies that function similarly to the backup policies.
Note: When you change a policy from backup to mirror or mirror to backup, the Dataset Policy Change wizard prompts you to establish a new baseline for the relationship. If you do, old data is retained, and the N series Management Console data protection capability makes a new copy of the entire dataset and transfers the active file system on the secondary. After reinitialization, you can manually delete the Snapshot copy.
- If you are changing the protection policy to a disaster recovery policy, do you want to map the settings from a node in the old dataset to a node in the new dataset?
You should copy the settings only if the path from the primary node is the same in the new policy as it was in the old policy.
- Do you plan to use a failover script to shut down processes before the N series Management Console data protection capability invokes failover?
If so, you need to define the path to a failover script.

Backup and mirror node resources

- If you are changing the protection policy for a dataset with backup and mirror nodes, do you want to use the same resource assignments that were used in the previous policy?

For example, if you have a dataset using the Mirror policy and you want to change to the Chain of two mirrors policy, you can choose to copy resources used for the single mirror node in the current policy to one of the two mirror nodes in the new policy. After you have copied resources from a node in the current policy, you cannot copy resources from that same node to any other node in the new policy.

- If you are assigning a policy for the first time or if you do not want to copy resources used in the current policy, is there a resource pool that meets the provisioning requirements of the dataset?

For example, the resource pool you assign to a mirror node should contain physical resources that would all be acceptable destinations for mirror copies created of the dataset members. If no resource pool meets the requirements of a non-primary node, you can create a new resource pool for each backup and mirror node using the Add Resource Pool wizard.

- If you prefer to not use resource pools for automatic provisioning, which physical resources would be suitable as destinations for backups and mirror copies of the dataset?

You do not have to assign a resource pool or physical resources to a destination node to assign or change a policy. However, the dataset will be nonconformant with its new policy until resources are assigned to each destination node, because the N series Management Console data protection capability cannot carry out the protection specified by the policy.

Provisioning policy

- On which node do you want to assign or change the provisioning policy?
Your dataset might have a primary and one or more non-primary nodes. You can assign the same provisioning policy to every node in the dataset, or you can assign a different provisioning policy to each node.
- Which provisioning policy meets the requirements of the dataset node?
Review the policies listed on the Provisioning Policies window to see if any is suitable for the dataset node.
If you are changing the provisioning policy, the policy type (NAS or SAN) on the primary dataset node must match the policy type assigned to the non-primary node. If you want to assign a NAS or SAN type provisioning policy instead of a secondary type policy to a non-primary node, the dataset must be disaster recovery capable. This means that it must also have a protection policy assigned that supports disaster recovery, and the node must be the disaster recovery capable node.
- If no provisioning policy meets the requirements of the dataset node, is there a provisioning policy that would be suitable with minor modifications?
If so, you can copy that policy to create a new policy that you can modify as needed. If not, you can run the Add Provisioning Policy wizard to create a new policy for the dataset node.

Descriptions of dataset conformance status (7-Mode only)

The dataset conformance status indicates whether a dataset is configured according to its local policy or storage service's protection policy. To be in conformance, all secondary and tertiary storage that is part of the backup relationship must be successfully provisioned and the provisioned objects must match the requirements of the primary data. You can monitor dataset status by using the Datasets tab.

The OnCommand console regularly checks a dataset for conformance. If it detects changes in the dataset's membership or policy definition, the console does one of three things:

- Automatically performs corrective steps to bring a dataset back into conformance
- Presents you with a list of actions for your approval prior to correction
- Lists conditions that it cannot resolve

You can view these actions and approve them in the Conformance Details dialog box.

A dataset might be nonconformant because there are no available resources from which to provision the storage or because the N series Management Console data protection capability does not have the necessary credentials to provision the storage resources.

The following list describes dataset conformance values:

- | | |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Conformant | The dataset is conformant with all associated policies. |
| Conforming | The dataset is not in conformance with all associated policies. The OnCommand console is performing actions to bring the dataset into conformance. |
| Nonconformant | The OnCommand console cannot bring the dataset into conformance with all associated policies and might require your approval or intervention to complete this task. |

Guidelines for adding a dataset of virtual objects (7-Mode only)

When you create or edit a dataset of virtual objects, following guidelines specific to datasets of virtual objects helps you to avoid some performance and space usage problems after configuration.

The following configuration guidelines apply to datasets of VMware objects:

- To avoid conformance and local backup issues caused by primary volumes reaching their Snapshot copy maximum of 255, best practice is to limit the number of virtual objects included in a primary volume, and limit the number of datasets to which each primary volume is directly or indirectly included as a member.

A primary volume that hosts virtual objects that are included in multiple datasets has an additional Snapshot copy of itself retained for every local backup on any dataset of which any of the primary volume's virtual object children are members.

- To avoid backup schedule inconsistencies, best practice is to include only virtual objects that are located in the same time zone in one dataset.

The schedules for the local protection jobs and remote protection jobs specified in the local policies and storage services that are assigned a dataset of virtual objects are performed according to the times on the host systems that are associated with the dataset's virtual objects.

- If a virtual machine resides on more than one datastore, you can exclude one or more of those datastores from the dataset. No local or remote protection is configured for the excluded datastores.

You might want to exclude datastores that contain swap files that you want to exclude from backup.

- To avoid an excessive amount of secondary space provisioned for backup, best practice when creating volumes to host the VMware datastores whose virtual machines will be protected by the OnCommand console backup is to size those volumes so that they are not much larger than the datastores they host.

The reason for this practice is that when provisioning secondary storage space to back up virtual machines that are members of datastores, the OnCommand console allocates secondary space that is equal to the total space of the volume or volumes in which those datastores are located. If the host volumes are much larger than the datastores they hold, an excessive amount of provisioned secondary space can result.

Guidelines for mounting or unmounting backups in a VMware environment (7-Mode only)

After you create a full or partial backup of your virtual machine or datastore, you can mount the backup onto an ESX server to verify what the virtual machine or datastore contains. After you verify the newly created backup to compare its content to the original and to look for mismatches, you can unmount the mounted backup.

When you mount or unmount backups, follow these guidelines:

- You cannot mount a backup on the same or a different ESX server if that backup is already mounted.

You must unmount this backup from the first ESX server prior to mounting a backup to a different ESX server.

- You can mount a local backup and a remote backup on any ESX host that is managed by the same host service that was used when the backup was created.
- If you include the same datastore in multiple backups and those backups are mounted, that datastore is mounted multiple times.

These datastores can be differentiated because the name includes a mounted timestamp and contains the dataset name.

- Backup and restore of mounted objects is not supported.
- If there is some data written in the mounted datastore, that data is lost when you unmount the backup.

- If a backup is mounted, you cannot delete it, even if it has expired, until you unmount the backup.
- While mounting a remote mirror backup, if the corresponding primary mirror backup has already been deleted, the mount request fails with a backup not found error.
- After you mount a backup, the time it takes to copy data from the datastore system depends on your network bandwidth and whether this datastore is on a secondary storage system.

How the Secondary Space Management wizard works

Before you use the Secondary Space Management wizard to plan and resolve your space management issues, it is helpful to understand how the wizard works.

Summary of the wizard workflow

The wizard is iterative, which means that instead of moving through all the wizard pages in a linear path, you move through some of the pages to plan a task for a selected volume. Then the wizard returns you to the starting page to plan another task.

At any time in the iterative planning process, you can review your planned tasks to see the estimated space savings in the aggregate. When you have planned all the tasks you want, you can have the wizard implement all your planned tasks or choose to copy the plan to use later.

Summary of how to create a space management plan

1. Select a volume and task

The **Volume and Tasks** page displays information about the aggregate space usage, lists the volumes in the aggregate and space usage information about each volume, and lists the tasks that are allowed for the selected volume. You select a volume and the task that you want to plan.

2. Provide details for the task and view estimated space reclaimed

The wizard moves you through a series of pages that help you to enter specific information that is needed to implement the task that you selected. On the **Task Summary** page, the wizard displays the estimated amount of space in the aggregate that the task can reclaim.

3. View list of tasks planned and plan another task

When you click **Next** at the end of a task summary, the task is added to the space management plan and the wizard returns to the **Volume and Tasks** page. The aggregate space usage information is updated to reflect the projected results of the task you just planned. You can plan as many tasks as you want by adding each additional task to the space management plan.

Summary of how to implement your space management plan

1. Review all tasks (and optionally revert a task)

When the estimated aggregate space usage resulting from your planned tasks reaches your goal, you can click **Review and Commit** on the **Volume and Tasks** page. This option displays your entire space management plan and the estimated total amount of space that will be reclaimed. You can revert any of the planned tasks.

2. Implement all tasks

You can click **Finish** to have the wizard implement all the tasks in the plan. If you do not want to implement the plan immediately, you can copy the list of tasks to your management station clipboard and save it to re-create the space management plan at another time.

Role of protection policies in dataset management (7-Mode only)

A protection policy specifies a dataset's primary node, secondary node, and tertiary node data protection topology, backup or mirror schedules, backup copy retention times, backup bandwidth consumption, and other aspects related to backing up physical storage objects and virtual objects in a dataset.

Direct or indirect assignment of a protection policy

Protection policies can be assigned to datasets directly or indirectly, through storage services.

- You can assign protection policies directly to datasets that are configured to include and manage physical storage objects as members.
- You can also assign protection policies to storage services, which are preconfigured combinations of protection policies, provisioning policies, and resource pools.
You can then assign storage services directly both to datasets configured for physical storage objects and datasets configured for virtual objects.

Where you can create, modify, and assign protection policies

You can create and modify protection policies, and assign them to storage datasets or storage services by using the associated program, N series Management Console. Information about any of those tasks is in the N series Management Console help.

Requirements and restrictions when adding a dataset of virtual objects (7-Mode only)

You must be aware of the requirements and restrictions when creating or editing a dataset of virtual objects.

- OnCommand console administrators attempting to assign a storage service to a dataset require RBAC read permission (DFM.Policy.Read) to access any protection policy included in that storage service.
This is in addition to having RBAC permissions DFM.Policy.Read, DFM.StorageService.Read, and DFM.StorageService.Attach to the storage service itself.
- Administrators of datasets of virtual objects who are attempting to attach storage services to those datasets require RBAC DFM.Resource.Control permission if the storage service they are attempting to attach assigns a provisioning policy to the dataset primary node.

Even though provisioning policy assignment does not actually apply to datasets of virtual objects, the DFM.Resource.Control permission is necessary to allow access to the underlying storage on which the virtual objects are located.

- Virtual objects (VMware objects) cannot coexist in the same dataset with storage system container objects (such as aggregates, volumes, and qtrees).
- To provide an application dataset with application-consistent backup protection, the OnCommand console administrator must assign to that application dataset a storage service that is configured with a protection policy that uses a "Mirror then backup" protection topology.
- VMware datacenter objects that you include in a dataset must not be empty. They must contain datastore or virtual machine objects for successful backup.
- VMDKs on a datastore object in a dataset must be contained within folders in that datastore. If VMDKs exist outside of folders on the datastore, and that data is backed up, restoring the backup might fail.

What a dataset is (7-Mode only)

A dataset is a set of virtual or physical containers that you can configure as a unit for the purpose of group protection or group provisioning operations.

- You can use the OnCommand console to configure datasets that contain virtual VMware objects.
- You can also use the OnCommand console and the associated N series Management Console to configure datasets that contain physical storage systems with aggregates, volumes, qtrees, and LUNs

During dataset configuration, you can additionally configure or assign local protection or remote protection arrangements and schedules that apply to all objects in that dataset. You can also start on-demand protection operations for all objects in that dataset with one command.

Why datasets fail to conform to policy (7-Mode only)

A dataset must meet several conditions to be conformant to its assigned policy.

A dataset is conformant with its policy when it meets the following conditions:

- Its member storage systems are properly configured.
- Its assigned secondary storage system is provisioned and has enough backup space.
- Its protection policy includes all necessary relationships to enforce data backups or mirror copies.

Following are some of the common reasons datasets fail to conform to their protection policies:

- Dataset storage service and protection policy definitions changed.
- Dataset membership changed.
- Volumes or qtrees were created or deleted at the storage system (external to the OnCommand console).

- The configuration of a dataset of virtual objects has been updated, but not yet communicated from the DataFabric Manager server to the host service.
This nonconformance normally lasts only a few minutes, until the DataFabric Manager server updates the host service with the latest configuration.

Copyright and trademark information

Copyright ©1994 - 2012 NetApp, Inc. All rights reserved. Printed in the U.S.A.

Portions copyright © 2012 IBM Corporation. All rights reserved.

US Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

No part of this document covered by copyright may be reproduced in any form or by any means— graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

References in this documentation to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's or NetApp's intellectual property rights may be used instead of the IBM or NetApp product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM and NetApp, are the user's responsibility.

No part of this document covered by copyright may be reproduced in any form or by any means— graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT

(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S.A. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service

Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, N.Y. 10504-1785
U.S.A.

For additional information, visit the web at:
<http://www.ibm.com/ibm/licensing/contact/>

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM web sites are provided for convenience only and do not in any manner serve as an endorsement of those web sites. The materials at those web sites are not part of the materials for this IBM product and use of those web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

If you are viewing this information in softcopy, the photographs and color illustrations may not appear.

Index

- ## A
- Add Dataset wizard
 - decisions to make for protection 49
 - Add Schedule wizard
 - decisions to make 47–49
 - task for weekly 13
 - administrative users
 - adding 24
 - administrator roles
 - list and descriptions 44
- ## B
- backups
 - creating for physical storage objects 10
 - creating for unprotected physical storage objects 14
 - guidelines for mounting or unmounting backups in a VMware environment 55
 - locating specific copies 18
 - mounting using the Backups tab 19
 - restoring a single file in a VMware environment 17, 20
 - searching for 18
 - unmounting using the Backups tab 20
- ## C
- conformance
 - dataset status 54
 - datasets, failure to conform 58
 - evaluating error text 38
 - resolving dataset issues manually with a baseline transfer of data 41
 - resolving issues automatically without a baseline transfer of data 40
 - resolving manually in datasets when a baseline transfer of data might be necessary 42
 - troubleshooting 37
 - workflow for resolving dataset conformance issues 36
 - custom label 49
- ## D
- dashboard panels
 - monitoring objects 28, 32
 - data
 - adding unprotected data to an existing dataset 15
 - Dataset Policy Change wizard
 - decisions to make 52
 - task 12
 - datasets
 - adding to manage physical storage objects 11
 - adding unprotected data to an existing 15
 - assigning a protection policy 12
 - changing storage services 34
 - conformance status values 54
 - decisions to make before adding to manage physical storage objects 49
 - defined 58
 - evaluating conformance issues 38
 - general concepts 45
 - guidelines when configuring datasets of virtual objects 54
 - listing nonconformant datasets 37
 - monitoring status 29
 - of virtual objects 47
 - policies, decisions to make before changing 52
 - reasons for failure to conform to policy 58
 - resolving conformance issues 36
 - resolving conformance issues automatically without a baseline transfer of data 40
 - resolving conformance issues manually when a baseline transfer of data might be necessary 42
 - resolving conformance issues manually without a baseline transfer of data 41
 - role of protection policies 57
 - troubleshooting 37
 - types of conformance status, defined 54
 - virtual objects, guidelines when configuring datasets of 54
 - datasets of physical storage objects
 - adding, decisions to make for protection 49
 - names, acceptable characters 49
 - properties of, for protection 49
 - delegated management
 - supporting for virtual infrastructure administration 22
 - disaster recovery policies
 - decisions to make before applying to datasets 52

N

- naming properties
 - custom label 49
- naming settings
 - definition of 45
- nonconformance
 - resolving dataset nonconformance 36

O

- Operations Manager Console
 - launching 23

P

- physical storage objects
 - adding datasets for 11
 - backing up 10
 - backing up unprotected 14
 - identifying unprotected 15
 - resolving issues 31
 - unprotected 14
- properties
 - of datasets of physical storage objects 49
- protected data
 - dataset conformance status, described 54
 - reasons for failure to conform to policy 58
- protection
 - adding unprotected data to a dataset 15
 - decisions before adding a dataset of physical storage objects 49
- protection or provisioning related objects
 - definition of 45
- protection policies
 - assigning to datasets 12
 - decisions to make before applying to datasets 52
 - overview 45
 - role in dataset management 57
- provisioning policies
 - datasets rebaselining after changing policies 52
 - decisions to make before applying to datasets 52
 - overview 45

R

- RBAC
 - capabilities 44
 - default roles 44
- rebaselining
 - about 52

- resolving issues
 - physical storage objects 31
 - virtual objects 28
- role-based access control
 - See* RBAC

S

- schedules
 - adding weekly 13
 - decisions to make before creating 47–49
- secondary space management
 - how the wizard works 56
 - starting the wizard 32
- Secondary Space Management wizard
 - how it works 56
- single file backups
 - restoring in a VMware environment 17, 20
- status definitions
 - dataset conformance 54
- storage objects
 - adding datasets for 11
 - unprotected physical objects 15
- storage services
 - changing for datasets 34
 - changing for virtual datasets 34
 - overview 45

T

- troubleshooting
 - dataset conformance issues 37
 - dataset failure to conform 58
 - listing nonconformant datasets 37

U

- unprotected data
 - assigning a protection policy 12
- users
 - adding, administrative 24

V

- virtual datasets
 - changing storage services 34
- virtual infrastructure
 - supporting delegated management for administration 22

- virtual objects
 - definition of 45
 - guidelines when configuring datasets of virtual objects 54
 - resolving issues 28
- VMware
 - guidelines when configuring datasets of VMware objects 54
- VMware environment
 - guidelines for mounting or unmounting backups 55

- mounting backups using the Backups tab 19
- restoring a single file 17, 20
- unmounting backups using the Backups tab 20

W

- weekly schedules
 - adding 13



210-05543_A0, Printed in USA

GA32-1021-01

