

NetApp® OnCommand™ Plug-in for VMware® **Backup and Recovery Administration Guide** For Use with Host Package 1.0

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089 USA
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 4-NETAPP
Documentation comments: doccomments@netapp.com
Information Web: <http://www.netapp.com>

Part number: 215-06026_A0
August 2011

Contents

Overview of the OnCommand Plug-in for VMware	5
About Monitoring and Host Configuration	5
About Provisioning and Cloning	6
About Backup and Recovery	8
Working with alarms	9
About Legacy Backup and Recovery	9
Changing the port number of the OnCommand Plug-in for VMware	9
Creating backups using datasets	11
Dataset concepts	11
Virtual objects that a dataset can include	12
Types of backups	12
Creating a dataset	13
Editing a dataset	14
Deleting a dataset	15
Suspending protection of a dataset	15
Resuming protection on a suspended dataset	16
Selecting a virtual machine or datastore to create a new dataset	16
Selecting a virtual machine or datastore to add to an existing dataset	17
Scheduling backups	18
Local policies supplied by the OnCommand Plug-in for VMware	18
Guidelines for creating a local policy	18
Preconfigured storage services supplied by the OnCommand Plug-in for VMware	19
Adding a local policy	20
Copying a local policy	21
Editing a local policy	21
Deleting a local policy	22
Assigning a storage service for remote protection of virtual objects	22
Creating on-demand backups	23
Guidelines for performing an on-demand backup	23
Performing an on-demand backup of a dataset	24
Backing up unprotected resources	25

Scheduling a backup with unprotected resources	26
Descriptions of dataset protection status	27
Restoring data from backups	29
Where to restore a backup	29
Restoring data from backups created by the OnCommand consoleOnCommand	
Plug-in for VMware	29
Searching for backups	30
Restoring a datastore	30
Restoring a virtual machine or its disk files	30
Mounting a backup	31
Unmounting a backup	31
Restoring data from legacy backups	32
Searching for legacy backups	32
Restoring a datastore	33
Restoring a virtual machine or its disk files	33
Managing jobs	35
Understanding jobs	35
Canceling jobs	35
Monitoring jobs	36
Copyright information	37
Trademark information	39
How to send your comments	41
Index	43

Overview of the OnCommand Plug-in for VMware

The OnCommand Plug-in for VMware is a single vCenter Server plug-in that provides end-to-end virtual machine life cycle management for VMware environments running NetApp storage. The plug-in provides storage configuration and monitoring using the Monitoring and Host Configuration capability, datastore provisioning and virtual machine cloning using the Provisioning and Cloning capability, and backup and recovery of virtual machines and datastores using Backup and Recovery capability.

Because the OnCommand Plug-in for VMware is a vCenter Server plug-in, it is available to all vSphere Clients that connect to the vCenter Server. This is different than a client-side plug-in that must be installed on every vSphere Client. You can install the OnCommand Plug-in for VMware on either the vCenter Server or on another Windows server in your data center; do not install it on a client computer.

The OnCommand Plug-in for VMware adds a NetApp icon to the Solutions and Applications panel of the vSphere Client home page.

When you select the About panel, the OnCommand Plug-in for VMware displays version information for each of the installed capabilities.

Note: You must log in to the DataFabric Manager server the first time you click on any of the panels in the Backup and Recovery navigation pane. You are then prompted for the user credentials for the DataFabric Manager server. If you inadvertently click on any of the panels in the Backup and Recovery navigation pane, you can ignore the log in prompt and navigate to a different capability.

About Monitoring and Host Configuration

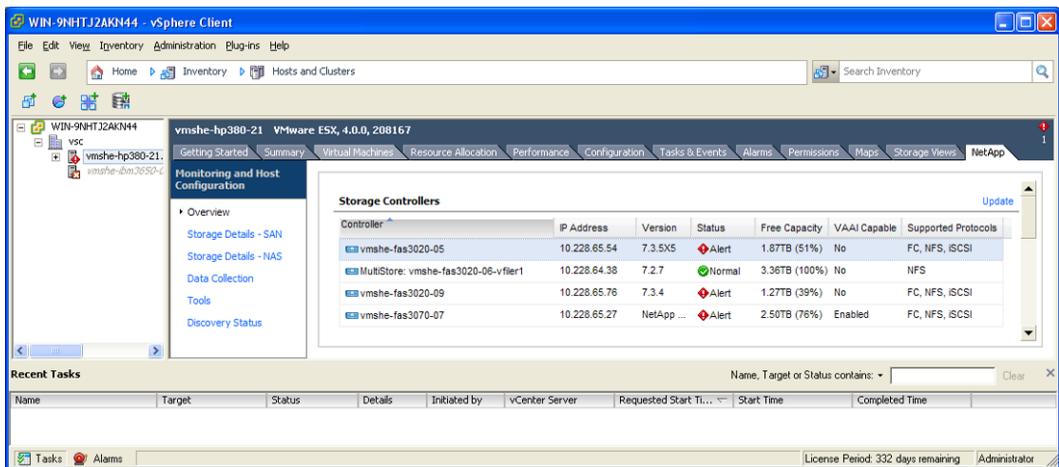
The Monitoring and Host Configuration capability enables you to manage ESX and ESXi servers connected to NetApp storage systems. You can set host timeout, NAS, and multipathing values, view storage details, and collect diagnostic data.

You can use this capability to:

- View the status of storage controllers from a SAN (FC, FCoE, and iSCSI) perspective
- View the status of storage controllers from a NAS (NFS) perspective
- View SAN and NAS datastore capacity utilization
- View the status of VMware vStorage APIs for Array Integration (VAAI) support in the storage controller
- View the status of ESX hosts, including ESX version and overall status
- Check at a glance whether the following are configured correctly, and if not, automatically set the correct values

- Storage adapter timeouts
- Multipathing settings
- NFS settings
- Set credentials to access storage controllers
- Launch the FilerView GUI to create LUNs and manage storage controllers
- Collect diagnostic information from the ESX hosts, storage controllers, and Fibre Channel switches
- Access tools to set guest operating system timeouts and to identify and correct misaligned disk partitions

When you click the NetApp icon in the vCenter Server and click Monitoring and Host Configuration in the navigation pane, the Overview panel appears. It is similar to the following:



About Provisioning and Cloning

The Provisioning and Cloning capability of Virtual Storage Console helps you to provision datastores and quickly create multiple clones of virtual machines in the VMware environment. Using FlexClone technology, the Provisioning and Cloning capability allows you to efficiently create, deploy and manage the lifecycle of virtual machines from an easy-to-use interface integrated into the VMware environment. It is ideal for virtual server, desktop, and cloud environments.

You can use this capability to:

- Clone individual virtual machines and place in new or existing datastores
- Create, resize, or delete datastores
- Apply guest customization specifications and power up new virtual machines
- Run deduplication operations
- Monitor storage savings
- Redeploy virtual machines from a baseline image

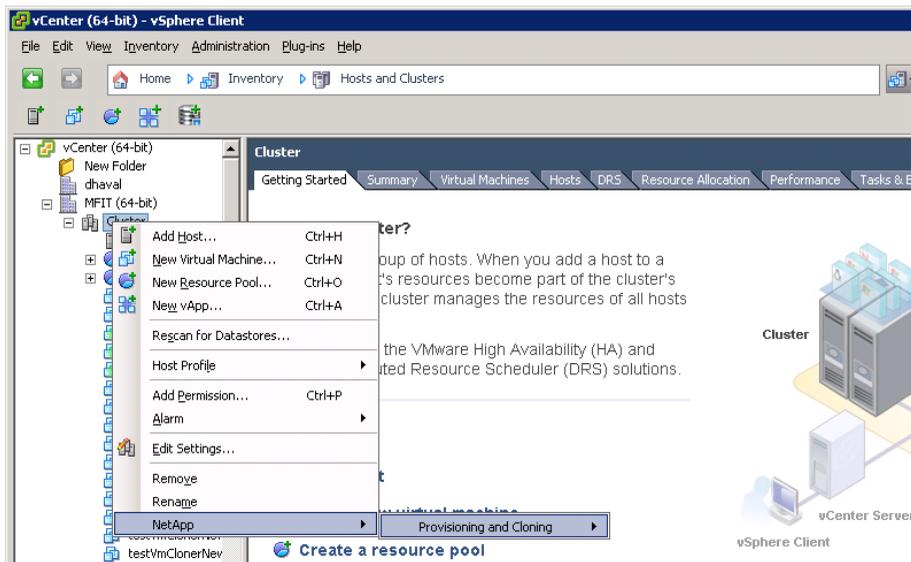
- Replicate NFS datastores across sites
- Import virtual machines into virtual desktop infrastructure connection brokers and management tools

You can access the various Provisioning and Cloning features using the following methods.

Managing datastores and cloning virtual machines

To manage datastores and clone virtual machines, right-click an object in the Inventory panel of the vSphere Client and select **NetApp > Provisioning and Cloning**.

- Right-click a powered-down virtual machine or template to create clones
- Right-click a datacenter, cluster, or host to provision datastores



Managing controllers and connection brokers, replicating datastores, and redeploying clones

Click the **Inventory** button in the navigation bar, and then select **Solutions and Applications > NetApp**.

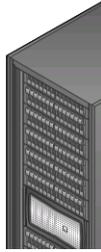
- Select **Storage controllers** to add, remove, or modify properties of storage controllers
- Select **Connection brokers** to add and remove connection broker definitions
- Select **DS Remote Replication** to clone NFS datastore templates to multiple target sites
- Select **Redeploy** to redeploy virtual machines.



What is the Provisioning and Cloning Capability?

The Provisioning and Cloning Capability combined with NetApp FlexClone allows users to quickly and efficiently create, deploy and manage the lifecycle of virtual machines from an easy-to-use user interface integrated into VMware vCenter. It is ideal for virtual server, desktop and dynamic cloud environments. The tool can be used to:

- Clone entire datastores or individual virtual machines
- Create, resize, and delete datastores
- Apply guest customization specifications to virtual machines
- Power up virtual machines
- Run deduplication operations
- Monitor storage savings
- Redeploy virtual machines from a baseline image
- Import virtual machines into Virtual Desktop Infrastructure connection brokers and management tools



License requirements

This task	Requires this license
Provision datastores	NFS, iSCSI, or FCP
Use vFilers in provisioning and cloning operations	MultiStore
Clone virtual machines	FlexClone
Configure deduplication settings	A_SIS
Distribute templates to remote vCenters	SnapMirror

About Backup and Recovery

The Backup and Recovery capability of Virtual Storage Console OnCommand Plug-in for VMware provides rapid backup and recovery of multihost configurations running on NetApp storage systems.

You can use this capability to:

- Perform policy-based backups of individual virtual machines, datastores, or a datacenter of a local or remote system by grouping these objects into a dataset, applying a local policy to the dataset, and selecting a storage service
- Support virtual machines and datastores that are located on either NFS directories or VMFS file systems
- Mount a backup to verify its content prior to restoration, or to restore one or more files to a guest virtual machine disk (VMDK) without having to restore the entire virtual machine
- Restore datastores or virtual machines to the original location
- Restore VMDKs to the original or an alternate location
- Restore one or more files to a guest VMDK without having to restore the entire virtual machine or VMDK using single file restore feature

Working with alarms

Alarms are configured notifications that are sent whenever a specific event or an event of a specific severity type occurs. You can create alarms for any defined resource group or the global group for which you want automatic notification of events.

You can access the Alarms tab in the OnCommand console to configure alarms for immediate notification of events. You can also configure alarms even before a particular event occurs. You can access this tab by clicking **Administration > NetApp > Alarms**.

About Legacy Backup and Recovery

The Legacy Backup and Recovery capability that is included in this release, enables you to restore data from legacy backups. You should begin to use the Backup and Recovery capability of the OnCommand Plug-in for VMware and shift from using the Legacy Backup and Recovery capability.

You must use the new version of the Backup and Recovery capability that is included with the OnCommand Plug-in for VMware to create a local or remote backup. The Legacy Backup and Recovery is accessible from the navigation bar. Legacy backup jobs continue to run until you delete them.

OnCommand Plug-in for VMware supports restoring data from backups taken with one of the following applications:

- or earlier
- SnapManager 2.0 for Virtual Infrastructure or earlier

Legacy backup jobs continue to run and provide data protection while you migrate your backup data using the OnCommand Plug-in for VMware.

Changing the port number of the OnCommand Plug-in for VMware

By default, the OnCommand Plug-in for VMware connects to the vCenter Server and the vSphere Client via port 8143 and the Backup and Recovery capability connects via port 8043. You can change the default port number if these port assignments pose a conflict for your other applications.

Steps

1. Create a file called `nvpf.override` for the OnCommand Plug-in for VMware or a file called `smvi.override` for the Backup and Recovery capability if the file does not exist.
2. For the OnCommand Plug-in for VMware, add this file to the directory `%PROGRAM FILES%\netapp\virtual storage console\etc`.

3. For the Backup and Recovery capability, add this file to the directory %PROGRAM FILES%\netapp\virtual storage console\smvi\server\etc.
4. Add the line http.port=8043 or http.port=8143 in the new file.
5. Restart the NetApp vSphere Plugin Framework service or the SnapManager for Virtual Infrastructure service.

Creating backups using datasets

You can create backup jobs that include which virtual machines or datastores to back up, where to store the backups, and for how long. After you specify the dataset that you want to back up, you can define the retention type and the backup settings for your backup job and the OnCommand Plug-in for VMware uses this information to create new backups and the retention policy to remove specific older ones.

When backups reach their expiration, they are deleted from the storage system. You can choose to keep the backup until you manually delete it, or you can assign a retention duration. By specifying a length of time to keep the on-demand local backup, you can override the retention duration of the local policy you assigned to the dataset for this backup. The on-demand local and remote backups default to a retention type of Monthly. Remote backups inherit the retention duration specified in the storage service assigned to the Dataset.

Dataset concepts

Associating data protection, a local policy, or 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 virtual objects	A collection of supported 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.
Data protection policy	Defines 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. Protection policies can be assigned indirectly to datasets of virtual objects through a storage service.
Provisioning policy	Defines 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. Provisioning policies can be assigned indirectly to datasets of virtual objects through a storage service.
Storage service	A single dataset configuration package that consists of a protection policy, provisioning policies, 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 setting up similar vFiler unit

attachments to each of them. 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.

Virtual objects that a dataset can include

Datasets that you configure to include virtual objects are constrained by VMware-specific restrictions.

- A dataset can include datacenter, virtual machine, and datastore objects.
- VMware datacenter objects that you include in a dataset cannot be empty. They must contain datastore or virtual machine objects for successful backup.
- VMware objects and storage system container objects (such as aggregates, volumes, and qtrees) cannot coexist in the same dataset.
- If you add a datastore object to a dataset, all the virtual machine objects that are contained in that datastore are protected by the dataset's assigned local backup policy or storage service.
- 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.
- If a virtual machine is added to a dataset, all of its VMDKs are protected by default unless one of the VMDKs is on a datastore that is in the "exclusion list" of that dataset.
- 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 could fail.

Types of backups

You can perform scheduled or on-demand local backups, or remote backups of datasets. Depending on what you need, the different types of backups offer variation in how you protect your data.

Scheduled local backups You can create scheduled local backups by adding or editing datasets and their application policies. The host service runs local backups so even when the OnCommand console is down, your backups continue to run.

On-demand local backups You can create on-demand local backups as you need them. On-demand backups apply to datasets. You can add specific virtual machines or datastores to existing or new datasets for backup. You can also select specific settings for on-demand backups that might differ from the local policy that might be attached to a dataset, including starting a remote backup after the local backup operation.

Remote backups You can create remote backups by assigning a storage service to the selected dataset. The storage service you assign to the dataset determines when and how the remote backup occurs. To create a remote backup, you must first create or edit a local dataset backup, or perform an on-demand dataset backup. During dataset creation, you can add the storage service to the dataset you want to back up.

Creating a dataset

You can create datasets to manage the protection of virtual objects in a VMware environment. A dataset is a set of virtual data storage containers that you can configure as a unit for the purpose of group protection operations, like having multiple backup schedules for members in a dataset.

Before you begin

Have the following protection information available:

- The name you want to give this dataset
- Whether you want to specify dataset-level custom naming formats for the Snapshot copy, volume, and qtree objects that are generated by local policy or storage service protection jobs on the virtual objects in this dataset.

About this task

The Create Dataset dialog box allows you to create an empty dataset, a populated but unprotected dataset, a populated, remotely protected dataset, a populated locally protected dataset, or a populated partially-configured or fully-configured dataset. A minimally configured dataset is an empty dataset. Virtual objects are unprotected if they do not belong to a dataset that has been assigned a local policy.

Steps

1. Click the **NetApp** tab, and then click **Datasets** under Backup and Recovery in the navigation pane.
2. In the **Datasets** panel, click **Create**.
3. In the **Create Dataset** dialog box, select the **Name** option and enter the requested information in the sub-tabs of the associated content area.
 - a. Enter the dataset name and administrative contact information in the **General Properties** tab.
 - b. If you want to specify dataset-level naming formats to apply to the Snapshot copy, volume, and qtree objects that are generated by local policy or storage service protection jobs that are run on this dataset, select the appropriate options in the **Naming Properties** tab.
4. If you want to specify, at this time, the virtual objects to include in this dataset, select the **Data** option and make your selections in the associated content area.

You can also add or change this information for this dataset at a later time.

5. If you want to specify, at this time, a storage service that will execute remote protection for the objects in this dataset, select the **Storage Service** option and make your selection.

You can also add or change this information for this dataset at a later time.

6. If you want to specify or create and configure, at this time, a local policy that will execute local protection for the objects in this dataset, select the **Local Policy** option and make your local policy selection or configuration.

You can also add or change this information for this dataset at a later time.

7. After you specify your desired amount of information about this dataset, click **OK**.

Editing a dataset

You can make modifications to existing datasets that were created to contain virtual objects.

About this task

You can modify the following in an existing dataset:

- Add virtual objects.
- Modify the remote protection of a dataset's virtual objects by assigning a storage service to that dataset.
- Modify local policies to automate local protection of datasets containing virtual objects.
- Modify the schedule of local backup jobs that are configured in the local policy assigned to a dataset containing virtual objects.
- Permanently remove dataset-based local or remote protection from a virtual object by removing that object from the protected dataset.

Steps

1. Click the **NetApp** tab, and then click **Datasets** under Backup and Recovery in the navigation pane.
2. In the **Datasets** panel, select the dataset that you want to edit and click **Edit**.
3. In the **Edit Dataset** dialog box and complete any of the following actions:
 - a. Select the **Data** option and select the virtual objects to add to the dataset.
 - b. Select the **Storage Service** option to view and assign a storage service that will execute remote protection for the objects in this dataset.
 - c. Select the local policy from the Local Policy drop-down list to assign a new or existing local policy to this dataset.
 - d. Select the **Local Policy** option to change the schedule for the local policy for this dataset.

- e. Select the **Data** option and select the virtual objects that you want to remove as dataset members.
4. After you specify your desired amount of information about this dataset, click **OK**.

Deleting a dataset

You can delete a dataset if you want to stop protection for all of its members and stop conformance checking against its assigned protection policies.

About this task

When you delete a dataset, the physical resources that compose the dataset are not deleted.

Steps

1. Click the **NetApp** tab, and then click **Datasets** under Backup and Recovery in the navigation pane.
2. In the **Datasets** panel, select the dataset that you want to delete and click **Delete**.
3. Click **Yes** to confirm the delete request or **No** to cancel the request and close the dialog box.

The deleted dataset is no longer listed in the Datasets panel.

Suspending protection of a dataset

You can temporarily stop data protection (if a protection policy is assigned) of a dataset. This task suspends all policies that are assigned to the dataset.

Steps

1. Click the **NetApp** tab, and then click **Datasets** under Backup and Recovery in the navigation pane.
2. In the **Datasets** panel, select the dataset that you want to suspend and click **Suspend**.
3. In the confirmation dialog box, click **Yes**.

The Details area of the Datasets panel shows a Protection Suspended status if the dataset has a protection policy assigned.

Result

You can resume data protection from the Datasets panel.

Resuming protection on a suspended dataset

You can resume data protection after it was suspended on a dataset. If you suspended protection on a dataset in order to perform maintenance on its volumes used as destinations for backups or mirror copies, you can resume operations after maintenance is complete.

Steps

1. Click the **NetApp** tab, and then click **Datasets** under Backup and Recovery in the navigation pane.
2. In the **Datasets** panel, select the dataset on which you want to resume protection and click **Resume**.
3. In the confirmation dialog box, click **Yes**.

The Details area of the Datasets panel shows a Protection enabled status if the dataset has a protection policy assigned.

Result

All scheduled backup operations are resumed.

Selecting a virtual machine or datastore to create a new dataset

You can select a virtual machine or datastore that you want to manage as one group and combine them into a new dataset. 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.

Before you begin

- Have the protection information available that you need to complete this task:
 - The name you want to give this dataset
 - The user group that you want this dataset visible to.
 - Whether you want to use the dataset name or substitute name as a prefix to names assigned to backup copies that are generated for this dataset.

Steps

1. Select a virtual machine or datastore in the **Inventory** panel, and then click the **NetApp** tab.
2. Select a virtual machine or datastore in the **Inventory** panel, and then right-click **NetApp > Backup and Recovery > Add to new Dataset**.

3. In the **Add to New Dataset** dialog box, select the **Name** option and enter the requested information in the associated content area.
4. If you want to specify, at this time, a storage service that will execute remote protection for the objects in this dataset, select the **Storage Service** option and make your selection.
You can also add or change this information for this dataset at a later time.
5. If you want to specify or create and configure, at this time, a local policy that will execute local protection for the objects in this dataset, select the **Local Policy** option and make your local policy selection or configuration.
You can also add or change this information for this dataset at a later time.
6. After you specify your desired amount of information about this dataset, click **OK**.

Selecting a virtual machine or datastore to add to an existing dataset

You can select a virtual machine or datastore and add it directly into an existing dataset. You can later edit a partially configured dataset to fully configure its membership or data protection.

Before you begin

- Have the protection information available that you need to complete this task:
 - The type of virtual objects that you want to add to an existing dataset.
 - The names of the virtual objects that you want to add.
 - The name of the dataset to which you want to add your selected objects.

Steps

1. Select a virtual machine or datastore in the **Inventory** panel, and then click the **NetApp** tab.
2. Select a virtual machine or datastore in the **Inventory** panel, and then right-click **NetApp > Backup and Recovery > Add to existing Dataset**.

The Add to Existing Dataset dialog box appears listing the selected objects and the existing datasets that can accept objects of their type.

3. Select the dataset to which you want to add your selected objects and click **Add**.

Scheduling backups

When you create a local policy for a dataset, it enables you to specify the start times, stop times, frequency, retention time, and warning and error event thresholds for local backups of virtual objects.

A local policy is a configured combination of Snapshot copy schedules, retention times, warning threshold, and error threshold levels that you can assign to a dataset. After you assign a local policy to a dataset, that policy applies to all the virtual objects that are included in that dataset.

Local policies supplied by the OnCommand Plug-in for VMware

To simplify the task of providing local protection to virtual objects in a dataset the OnCommand Plug-in for VMware provides a set of local policies (preconfigured combinations of local backup schedules, local backup retention settings, lag warning and lag error thresholds, and optional backup scripts) specifically to support local backup of certain types of data.

The preconfigured local policies include the following:

VMware local backup policy template	<p>This default policy enforces the following VMware environment-optimized settings related to local backup scheduling and retention. This policy can also be renamed and modified, to accommodate different circumstances.</p> <ul style="list-style-type: none"> • Hourly backups without VMware snapshot (crash consistent) every hour between 7 am and 7 pm every day including weekends • Daily backups with VMware snapshot at 10 PM every night every day including weekends • Weekly backup with VMware snapshot every Sunday midnight • Retention settings: Hourly backups for 2 days, Daily backups for 1 week, Weekly backups for 2 weeks • Issue a warning if there are no backups for: 1.5 days • Issue an error if there are no backups for: 2 days • Backup script path: empty
--	--

Guidelines for creating a local policy

Before you use the Create Local Policy dialog box or Edit Local Policy dialog box to create or edit a local policy, you must make some decisions about the policy to input in your command selection and in the configuration page.

Name

Name Your company might have a naming convention for policies. When specifying a name for a new policy, make sure you follow that convention.

Description Use a description that helps someone unfamiliar with the policy to understand its purpose.

Schedule and Retention

Schedule You can set up multiple schedules of multiple types (Hourly, Daily, Weekly) for your local backups. Each schedule has a start time, a stop time and a frequency with which its backups are executed.

Retention You can specify a retention period to be associated with each type of backup schedule (Hourly, Daily, or Weekly). A retention period specifies the minimum length of time that a backup copy is maintained, before it is eligible to be purged. A retention period assigned to one type of backup schedule applies to all backup copies of that type.

Backup Options Depending on the type of virtual objects the dataset contains, you can enable additional operations to be performed on those objects during backup.

- **Create VMware Snapshots**
Enables creation of VMware formatted snapshots (in addition to the storage system Snapshot copies) during local backup operations (displayed for datasets of VMware objects).
- **Include independent disks**
Applies local backups associated with this schedule to VMDKs that belong to VMware virtual machines in the current dataset, but reside on datastores that are not part of the current dataset (displayed for datasets of VMware objects).
- **Start remote backup after local backup**
Starts a remote backup of data to secondary storage after the local backup is finished if a storage service that specifies a remote backup is also assigned to the dataset.

Preconfigured storage services supplied by the OnCommand Plug-in for VMware

To simplify the task of providing provisioning storage for virtual objects and remote protection to virtual objects in a dataset, the OnCommand Plug-in for VMware provides a set of storage services, preconfigured combinations of protection and provisioning policies specifically designed to facilitate remote protection of virtual objects.

The preconfigured storage services include the following provisioning and protection policies:

Reserved Data Space for NFS Datastores This storage service is configured to provide thick provisioning for NFS datastores.

- Provisioning policy: Reserved Data Space for NFS Datastores
- Protection policy: None

Reserved Data Space for VMFS Datastores	<p>This storage service is configured to provide provisioning for write guaranteed VMFS datastores.</p> <ul style="list-style-type: none"> • Provisioning policy: Reserved Data Space for VMFS Datastores • Protection policy: None
Thin Provisioned Space for NFS Datastores	<p>This storage service is configured to provide thin provisioning for NFS datastore objects.</p> <ul style="list-style-type: none"> • Provisioning policy: Thin Provisioned Space for NFS Datastores • Protection policy: None
Thin Provisioned Space for VMFS Datastores	<p>This storage service is configured to provide thin provisioning for VMFS datastore objects.</p> <ul style="list-style-type: none"> • Provisioning policy: Thin Provisioned Space for VMFS Datastores • Protection policy: None

Adding a local policy

You can create local policies to schedule local backup jobs and designate retention periods for the local backup copies for datasets of virtual objects. A dataset's local policy in OnCommand enables you to specify the start times, stop times, frequency, retention time, and warning levels for local backups of its VMware objects.

About this task

The Create Local Policy dialog box enables you to create a new local policy, or add a preconfigured local policy by selecting the type of backup: partial, independent, or VMware snapshot and specifying the schedule/retention options.

Steps

1. Click the **Inventory** button in the navigation bar, and then select **Solutions and Applications > NetApp**.
2. Click **Local Policies** under Backup and Recovery in the navigation pane.
3. Click **Add** in the **Local Policies** panel.
4. In the **Create Local Policy** dialog box, select the **Name** option and enter the requested information in the associated content area.
5. Select the **Schedule & Retention** option and click **Add**.
 - Select the schedule type (Hourly, Daily, Weekly, or Monthly) that specifies the duration that you want the backup copy retained.
 - Click **View** in the Backup Options column and specify the following additional options that you can enable for the selected local policy.

- If you want a remote backup to begin after the local backup has finished, click the **Start remote backup after local backup** check box.
 - If you want VMware formatted snapshots created in addition to the storage system Snapshot copies during local backup operations, click the **Create VMware snapshot** check box.
 - If you want the local backups associated with this schedule applied to VMDKs that belong to virtual machines in the current dataset, but reside on datastores that are not part of the current dataset, click the **Include independent disks** check box.
6. Select the **Backup Settings** option and make your selections in the associated content area.
You can also change this information for this policy at a later time.
 7. After you enter your desired amount of information about this policy, click **OK**.

Copying a local policy

You can make a new local policy by copying an existing local policy and modifying it to meet your particular requirements.

Steps

1. Click the **Inventory** button in the navigation bar, and then select **Solutions and Applications > NetApp**.
2. Click **Local Policies** under Backup and Recovery in the navigation pane.
3. In the **Local Policies** panel, select the local policy that you want to copy and click **Copy**.
4. Select the copied policy and click **Edit**.
5. Make any required name changes and schedule modifications in the **Edit Local Policy** dialog box and click **OK**.

Editing a local policy

You can edit an existing local policy if you want to modify the name or description, schedule and retention times, or the backup options assigned to the policy.

About this task

The Edit Local Policy dialog box enables you to modify an existing local policy.

Steps

1. Click the **Inventory** button in the navigation bar, and then select **Solutions and Applications > NetApp**.
2. Click **Local Policies** under Backup and Recovery in the navigation pane.
3. In the **Local Policies** panel, select the local policy that you want to edit and click **Edit**.

4. In the **Edit Local Policy** dialog box, click the selection button on one or more of the following options to display and modify their related settings:
 - Name
 - Schedule & Retention
 - Backup Settings
5. After you have changed the desired information, click **OK**.

Deleting a local policy

You can delete local policies that have not yet been assigned to datasets of virtual objects.

Steps

1. Click the **Inventory** button in the navigation bar, and then select **Solutions and Applications > NetApp**.
2. Click **Local Policies** under Backup and Recovery in the navigation pane.
3. In the **Local Policies** panel, select the local policy that you want to delete and click **Delete**.
4. Click the **Dependencies** tab to determine whether this local policy is currently assigned to one or more datasets.

If one or more assigned datasets are listed on the Dependencies tab, either edit those datasets to unassign this local policy or do not delete this local policy.

5. Click **Delete** if no assigned datasets are listed In the **Dependencies** tab
6. Click **Yes** to confirm the delete request or **No** to cancel the request and close the dialog box.

Assigning a storage service for remote protection of virtual objects

After Snapshot copies of a dataset's virtual objects are generated as local backup, remote protection operations that are specified in an assigned storage service configuration can save these backup copies to secondary storage. You can configure remote protection of a dataset's virtual objects by assigning a storage service to that dataset.

Before you begin

Secondary or tertiary protection of virtual objects cannot be accomplished unless Snapshot copies have been generated on the primary node by backup jobs carried out on demand or through local policy.

Steps

1. Click the **NetApp** tab, and then click **Datasets** under Backup and Recovery in the navigation pane.

2. In the **Datasets** panel, select the dataset for whose virtual objects you want to configure remote protection, and click **Edit**.
3. In the **Edit Dataset** dialog box use the **Storage Service** option to view and assign a storage service that will execute remote protection for the virtual objects in this dataset.

As you make your selection, the associated content area displays the administrative, topological, protection schedule, and backup retention information for the storage service that you select.

4. After you make your selection, if you want to test whether your dataset's settings and resources conform to the requirements of its newly assigned storage service before you apply that storage service, click **Test Conformance** to display the Dataset Conformance Report.
 - If the Dataset Conformance Report displays no warning or error information, click **Close** and continue.
 - If the Dataset Conformance Report displays warning or error information, read the Action and Suggestion information resolve the conformance issue, then click **Close** and continue.
5. When you are ready to confirm the storage service assignment, click **OK**.

Creating on-demand backups

You can protect your virtual machines or datastores by adding them to an existing or new dataset and performing an on-demand backup.

Guidelines for performing an on-demand backup

Before performing an on-demand backup of a dataset, you must decide how you want to assign resources and assign protection settings.

Local protection settings	You can define the retention duration and the backup settings for your on-demand backup, as needed.
Retention	<p>You can choose to keep a backup until you manually delete it, or you can assign a retention duration. By specifying a length of time to keep the on-demand local backup, you can override the retention duration in the local policy you assigned to the dataset for this backup. The retention duration of a local backup defaults to a retention type for the remote backup.</p> <p>A combination of both the remote backup retention type and storage service is used to determine the remote backup retention duration.</p> <p>For example, if you specify a local backup retention duration of two days, the retention type of the remote backup is Daily. The dataset storage service then verifies how long daily remote</p>

backups are kept and applies this to the backup. This is the retention duration of the remote backup.

The following table lists the local backup retention durations and the equivalent remote backup retention type:

Local retention duration	Remote retention type
Less than 24 hours	Hourly
1 day up to, but not including, 7 days	Daily
1 week up to, but not including, 31 days	Weekly
More than 31 days	Monthly

Backup Settings

You can choose your on-demand backup settings based on the type of virtual objects you want to back up.

Create VMware snapshot

You can choose to create a VMware formatted in addition to the storage system Snapshot copies created during local backup operations.

Include independent disks

You can include independent disks. VMDKs belong to VMware virtual machines in the current dataset, but reside on datastores that are not part of the current dataset.

Backup Script Path

You can specify a script that is invoked before and after the local backup. The script is invoked on the host service and the path is local to the host service. If you use a PowerShell script, you should use the drive letter convention. For other types of scripts, you can use either the drive letter convention or the Universal Naming Convention.

Backup Description

You can provide a description for the on-demand backup so you can easily find it when you need it.

If you want a remote backup to begin after the local backup has finished, click the **Start remote backup after local backup** check box.

Performing an on-demand backup of a dataset

You can perform a local on-demand dataset backup to protect your virtual objects. You can choose any virtual machine or datastore to back up, but it must first belong to a dataset. If the virtual

machine or datastore does not belong to a dataset, you can choose to add it to an existing dataset. You can also add a new dataset and add the virtual machine or datastore to the dataset.

Before you begin

You must have the following information available:

- Dataset name
- Retention type
- Backup settings
- Backup script location
- Backup description

About this task

You must have added the virtual objects to an existing dataset or have added a dataset that contains the virtual objects you want to back up.

Steps

1. Click the **NetApp** tab, and then click **Datasets** under Backup and Recovery in the navigation pane.
2. In the **Datasets** panel, choose the dataset that you want to back up.
3. Click **Back Up Now**.
4. In the **Back Up Now** dialog box, specify the local protection settings, backup script path, and backup description for the on-demand backup.

If you have already established local policies for the dataset, that information automatically appears for the local protection settings for the on-demand backup. If you change the local protection settings, the new settings override any existing local policies for the dataset.

5. If you want a remote backup to begin after the local backup has finished, click the **Start remote backup after local backup** check box.
6. Click **Back Up Now**.

Backing up unprotected resources

You can define local and remote backups for the unprotected virtual machines or datastores.

You can complete the following tasks:

- Add unprotected virtual machines or datastores to an existing dataset.
- Select unprotected virtual machines or datastores from the drop-down list.
- Create an application VMware-specific local policy consisting of backup schedules and settings.

- View the list of available storage services and select one for remote backups.
- Create a dataset with a local policy for local backups and a storage service for the remote backups.

Scheduling a backup with unprotected resources

You can schedule backups of unprotected resources by first selecting and adding the unprotected resources to a new dataset or an existing dataset.

About this task

Virtual objects are unprotected if they do not belong to a dataset that has been assigned a local policy.

Steps

1. Click the **Inventory** button in the navigation bar, and then select **Solutions and Applications > NetApp**.
2. Click **Unprotected Resources** under Backup and Recovery in the navigation pane.
3. Select **Virtual Machines** or **Datastores** from the drop-down list in the **Unprotected Resources** panel.
4. Select a virtual machine or datastore and click **Backup**.
5. Select one of the following options:
 - Schedule a backup for the selected virtual machine or datastore by adding it to an existing dataset.
Note: Both the add to existing dataset and create new dataset options appear when the unprotected resource is not part of any dataset.
 - Schedule a backup for the selected virtual machine or datastore by creating a new dataset.
Note: This is the only option that appears when there are no existing datasets to which you can add this unprotected resource.
 - Back up now.
Note: The create new dataset, add to existing dataset, and back up now options all appear when the unprotected resource is already a part of a dataset.
6. Click **Schedule Backup** after making your selection.
The Add to Existing Dataset dialog box, Create Dataset dialog box, or the Back Up Now dialog box appear.
7. After you select an existing dataset or specify your desired amount of information about this dataset, click **Add**, **OK**, or **Back Up Now**.

Descriptions of dataset protection status

A dataset is remotely protected only if the secondary storage system specified in the protection relationship is successfully backing up data, and if the copies of the data can be restored. You can monitor dataset status using Datasets panel.

You should regularly monitor a dataset's protection, because the OnCommand Plug-in for VMware cannot sufficiently protect the dataset under the following conditions:

- If a secondary storage system runs out of storage space necessary to meet the retention duration required by the protection policy
- If the lag thresholds specified by the policy are exceeded

The following list describes protection status values and their descriptions:

Baseline Failed No initial baseline data transfers have registered a backup version.

Initializing The dataset is conforming to the protection policy and the initial baseline data transfer is in process.

Job Failure The most recent protection job did not succeed.

This status might result for any of the following reasons:

- A backup from a SnapVault or Qtree SnapMirror relationship failed or could not be registered.
- A mirror copy from a SnapMirror relationship failed or could not be registered.
- Local backups (Snapshot copies) failed on the primary node.

Lag Error The dataset has reached or exceeded the lag error threshold specified in the assigned protection policy. This value indicates that there has been no successful backup or mirror copy of a node's data within a specified period of time.

This status might result for any of the following reasons:

- The most recent local backup (Snapshot copy) on the primary node is older than the threshold setting permits.
- The most recent backup (SnapVault or Qtree SnapMirror) is older than the lag threshold setting or no backup jobs have completed since the dataset was created.
- The most recent mirror (SnapMirror) copy is older than the lag threshold setting or no mirror jobs have completed since the dataset was created.

Lag Warning The dataset has reached or exceeded the lag warning threshold specified in the assigned protection policy. This value indicates that there has been no successful backup or mirror copy of a node's data within a specified period of time.

This status might result for any of the following reasons:

- The most recent local backup (Snapshot copy) on the primary node is older than the threshold setting permits.
- The most recent backup (SnapVault or Qtree SnapMirror) is older than the lag threshold setting or no backup jobs have completed since the dataset was created.
- The most recent mirror (SnapMirror) copy is older than the lag threshold setting or no mirror jobs have completed since the dataset was created.

No Protection Policy

The dataset is managed by the OnCommand Plug-in for VMware but no protection policy has been assigned to the dataset.

Protected

The dataset has an assigned policy and it has conformed to that policy at least once.

Protection Suspended

An administrator has requested that scheduled backups be temporarily halted until the administrator requests that they be resumed.

Uninitialized

This status might result for any of the following reasons:

- The dataset has a protection policy that does not have any protection operations scheduled.
- The dataset does not contain any data to be protected.
- The dataset does not contain storage for one or more destination nodes.
- The single node dataset does not have any backup versions.
- An application dataset requires at least one backup version associated with it.
- The dataset does not contain any backup or mirror relationships.

No Local Policy A dataset of virtual objects has no local policy assigned.

Attention: Importing an external relationship into a dataset temporarily changes the dataset's protection status to Uninitialized. When the next scheduled backup or mirror backup job runs or when you run an on-demand backup the protection status changes to reflect the results of the protection job.

Restoring data from backups

The OnCommand console OnCommand Plug-in for VMware allows you to restore your virtual machines and datastores from legacy backups and from backups taken of newly created datasets that contain your virtual machines and datastores. The OnCommand console OnCommand Plug-in for VMware supports restore from local and remote backups and from backups that contain VMware-based snapshots.

Where to restore a backup

OnCommand console OnCommand Plug-in for VMware allows you to select a restore destination from the Restore Wizard.

- Original location
The backup of an entire datastore, a single virtual machine, or a virtual machine's disk files are restored to the original location. You set this location by choosing **The entire virtual machine** option.
- Different location
The backup of the virtual machine disk files are restored to a different location. You select the destination location by setting the **Particular virtual disks** option.

Note: If you are recovering data from a remote backup, data access from the ESX server to a Hyper-V controller using FC, iSCSI, or NFS protocols must exist to ensure that an NFS mount can be done from a secondary node to the ESX server.

Restoring data from backups created by the OnCommand console OnCommand Plug-in for VMware

You can restore a datastore, virtual machine, or its disk files to its original location or an alternate location. From the Backup Management panel, you can sort the backup listings by vendor type to help you find your backups.

From the Backup Management panel, you can do the following:

- Restore a datastore, virtual machine, or its disk files from a local or remote backup to an original location
- Restore virtual machine disk files from a local or remote backup to a different location.
- Restore from a backup that has a VMware snapshot.

Searching for backups

You can locate a specific backup by searching the backup table. After you locate a backup, you can then restore from it, view its status, or delete it.

Steps

1. From the navigation panel, click **Restore** panel to display the available backups.
2. In the Search text box, type one or more search terms and press Enter.

Available criteria are the backup ID, dataset name, dataset description, node name (local or remote), VMware snapshot, mount state or resource name.

A filtered list displays in the Backup Management panel.

Restoring a datastore

You can use the OnCommand Plug-in for VMware to restore a datastore. By doing so, you overwrite the existing content with the backup you select.

Steps

1. In the **Restore** panel, select a backup of the datastore.
2. Click **Restore**.

Restoring a datastore powers off all virtual machines.

3. In the **Restore dialog box**, click **Restore**.

Restoring a virtual machine or its disk files

You can use the Backup and Recovery capability to restore an entire virtual machine or particular virtual disks of a given virtual machine. By doing so, you overwrite the existing content with the backup you select.

Before you begin

You must have already backed up a virtual machine or its disk files using OnCommand Plug-in for VMware or Legacy Backup and Recovery before you can recover it.

Steps

1. Select a virtual machine in the **Inventory** panel, and then right-click the node and select the **NetApp > Backup and Recovery > Restore** from the context menu.

The Restore wizard opens and lists all backups which include the virtual machine.

2. From the Backups table, select a backup of the virtual machine and click **Next** to display the Virtual Machine Component Selection page.

3. Select one of the following recovery options:

Option	Description
The entire virtual machine	Restores the contents of your virtual machine from a Snapshot copy. The Start VM after restore checkbox is enabled if you select this option and the virtual machine is registered.
Particular virtual disks	Restores the contents of the virtual disks on a virtual machine. This option is enabled if you uncheck the entire virtual machine option.

4. In the **ESX Host Name field**, select the name of the ESX host.

This option is available if you want to restore virtual disk files or the virtual machine is on a VMFS datastore.

5. Click **Next** to display the Summary page, which displays information about the restore that is about to be performed.
6. From this page, view the settings and click **Restore** to begin the restoration.

After you finish

After you end the Restore wizard and start the restore job, you can track the progress of the restore job from the Task tab of the Status panel, and monitor the job for possible errors.

Mounting a backup

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.

Steps

1. Select a virtual machine or datastore in the **Inventory** panel, and then click the **NetApp** tab.
2. Click **Restore** under Backup and Recovery in the navigation pane.
3. In the **Restore** panel, select an unmounted backup that you want to mount.
4. Click **Mount**.
5. 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.

6. Click **Mount**.

Unmounting a backup

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 that it was mounted to.

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.

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. Select a virtual machine or datastore in the **Inventory** panel, and then click the **NetApp** tab.
2. Click **Restore** under Backup and Recovery in the navigation pane.
3. In the **Restore** panel, select a mounted backup to unmount.
4. Click **Unmount**.
5. At the confirmation prompt, click **Yes**.

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.

Restoring data from legacy backups

OnCommand Plug-in for VMware supports legacy backups created using the Legacy Backup and Recovery capability or SnapManager for Virtual Infrastructure application. You can restore from legacy backups until the legacy backup expires and you have replaced the job with one that you created by using a dataset, application policy and storage service.

Searching for legacy backups

You can search for Snapshot-based backups and recover your entire datastore or virtual machine, or virtual disk files in the virtual machine from the Restore panel.

If you recover data from a virtual machine or its virtual disk files to its original location on an NFS datastore, they are restored using single file snap restore. For all other data recovery, the backup is mounted and corresponding files are copied to the restore location.

Restoring a datastore

You can use the legacy Backup and Recovery to recover a datastore. By doing so, you overwrite the existing content with the legacy backup you select.

Before you begin

You must have already backed up a datastore using Virtual Storage Console before you can recover it.

Steps

1. Select a datastore in the **Inventory** panel, and then right-click the node and select the **NetApp > Backup and Recovery > Restore** from the context menu.

The Restore wizard opens and lists all backups which include the datastore, or any virtual machine in that datastore.

2. From the Backups table, select a backup of the datastore.
3. From the Backed-up entities table, select the datastore you want to restore.
4. Click **Next** to display the Summary page, which displays information about the restore that is about to be performed.
5. From this page, view settings and click **Restore** to begin the restoration.

After you finish

After you end the Restore wizard and start the restore job, you can track the progress of the restore job from the Task tab of the Status panel, and monitor the job for possible errors.

Restoring a virtual machine or its disk files

You can use the Backup and Recovery capability to restore an entire virtual machine or particular virtual disks of a given virtual machine. By doing so, you overwrite the existing content with the backup you select.

Before you begin

You must have already backed up a virtual machine or its disk files using OnCommand Plug-in for VMware or Legacy Backup and Recovery before you can recover it.

Steps

1. Select a virtual machine in the **Inventory** panel, and then right-click the node and select the **NetApp > Backup and Recovery > Restore** from the context menu.

The Restore wizard opens and lists all backups which include the virtual machine.

- From the Backups table, select a backup of the virtual machine and click **Next** to display the Virtual Machine Component Selection page.
- Select one of the following recovery options:

Option	Description
The entire virtual machine	Restores the contents of your virtual machine from a Snapshot copy. The Start VM after restore checkbox is enabled if you select this option and the virtual machine is registered.
Particular virtual disks	Restores the contents of the virtual disks on a virtual machine. This option is enabled if you uncheck the entire virtual machine option.

- In the **ESX Host Name field**, select the name of the ESX host.
This option is available if you want to restore virtual disk files or the virtual machine is on a VMFS datastore.
- Click **Next** to display the Summary page, which displays information about the restore that is about to be performed.
- From this page, view the settings and click **Restore** to begin the restoration.

After you finish

After you end the Restore wizard and start the restore job, you can track the progress of the restore job from the Task tab of the Status panel, and monitor the job for possible errors.

Managing jobs

You can view the current status and other information about all jobs that are currently running as well as jobs that have completed in the Jobs panel.

Understanding jobs

A job is typically a long-running operation. OnCommand enables you to create, manage, and monitor jobs. From the Jobs panel, you can view all jobs that are currently running as well as jobs that have completed.

Following are three examples of a job:

- A scheduled local backup of a dataset
- A mirror transfer job
- A mount or unmount of a VMware snapshot copy

Canceling jobs

You can use the Jobs panel to cancel a job if it is taking too long to complete, encountering too many errors, or is no longer needed. You can cancel a job only if its status and type allow it. You can cancel any job that has the status Running or Running with failures.

Steps

1. Click the **Inventory** button in the navigation bar, and then select **Solutions and Applications > NetApp**.
2. Click **Jobs** under Backup and Recovery in the navigation pane.
3. From the list of jobs, select one or more jobs.
4. Click **Cancel**.

Note: The Cancel button is enabled only for jobs that are either Running or Running with failures. If the Cancel button is not enabled, that job type cannot be canceled.

5. At the confirmation prompt, click **Yes** to cancel the selected job.

Monitoring jobs

You can monitor for job status and other job details using the Jobs panel. For example, you can view the progress of an on-demand backup job and see whether there are any errors.

Steps

1. Click the **Inventory** button in the navigation bar, and then select **Solutions and Applications > NetApp**.
2. Click **Jobs** under Backup and Recovery in the navigation pane.
3. Select a job in the jobs list to see information about that job.

Copyright information

Copyright © 1994–2011 NetApp, Inc. All rights reserved. Printed in the U.S.A.

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

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.

IBM, the IBM logo, and 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 www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the U.S.A. and/or other countries.

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.

How to send your comments

You can help us to improve the quality of our documentation by sending us your feedback.

Your feedback is important in helping us to provide the most accurate and high-quality information. If you have suggestions for improving this document, send us your comments by e-mail to doccomments@netapp.com. To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support Telephone: +1 (888) 4-NETAPP

Index

B

- Backup and Recovery
 - about 8
- Backup Management panel 29
- backups
 - creating 11
 - finding 30
 - mounting 31
 - on-demand 12
 - remote 12
 - restoring data from 29
 - restoring from legacy backups 32
 - scheduled local 12
 - scheduling 18
 - scheduling with unprotected resources 26
 - selecting the restore destination 29
 - unmounting 31

D

- datasets
 - creating 13
 - deleting 15
 - editing 14
 - general concepts 11
 - performing an on-demand backup 24
 - resuming protection 16
 - suspending protection 15
 - types of protection status, defined 27
 - virtual objects that they can include 12

datastores

- recovering 33
- restoring a 30
- searching for backups of 30

E

- ESX hosts
 - recovering virtual disks on VMFS datastores 30, 33

I

- Inventory panel
 - selecting a virtual entity 33

J

- jobs
 - canceling 35
 - managing 35
 - monitoring 36
 - understanding 35

- jobs, defined 35

L

- Legacy Backup and Recovery
 - about 9
- local policies
 - adding 20
 - copying 21
 - defined 18
 - deleting 22
 - editing 21
 - guidelines for adding or editing 18, 19

M

- Monitoring and Host Configuration
 - about 5

N

- NFS datastores 32

O

- on-demand backups
 - guidelines 23
 - prerequisites 23
- OnCommand Plug-in for VMware
 - Legacy Backup and Recovery capability 9

P

- protected data
 - dataset protection status, described 27
- protection policies
 - overview 11
- Provisioning and Cloning
 - about 6
- provisioning policies
 - overview 11

R

- recover virtual data
 - from a datastore 33
 - from a virtual machine 30, 33
 - from virtual disk files 30, 33
- remote backups 12
- remote protection
 - assigning to virtual objects 22
- restore
 - a datastore 30
 - where to restore a backup 29
- Restore panel
 - searching for backups 30
 - using 32
- Restore wizard
 - recovering a datastore 33
 - recovering a virtual machine or its disk files 30, 33
- restoring
 - data from backups 29
 - legacy backups 32

S

- scheduled backups 12
- Search field
 - using the 30
- search for a backup 30

- status definitions, dataset
 - protection 27
- Status panel
 - viewing job status 33
- storage services
 - assigning to a dataset of virtual objects 22
 - overview 11
 - supplied with the product 19

V

- virtual disk files
 - recovering 30, 33
- virtual machines
 - recovering 30, 33
 - restarting 30, 33
 - searching for backups of 30
- virtual objects
 - configuring remote protection for 22
- Virtual Storage Console
 - Backup and Recovery capability 8
 - Monitoring and Host Configuration capability 5
 - Provisioning and Cloning capability 6
- VMFS datastores 30, 33
- VMware
 - objects that a dataset can include 12