



Configure SnapManager for Hyper-V

SnapManager for Hyper-V

NetApp
February 12, 2024

Table of Contents


- Configure SnapManager for Hyper-V 1
 - Dashboard settings 1
 - Configure hosts 2
 - Configure SnapInfo directory settings 7
 - Configure datasets 9
 - configure policies 11
 - Configure SVMs or CIFS servers for Hyper-V over SMB 15

Configure SnapManager for Hyper-V

After you install SnapManager for Hyper-V, you can configure and manage your hosts and virtual machines by adding policies to protect and restore your data.

Dashboard settings

The SnapManager for Hyper-V dashboard displays an overview of resources that are currently being protected, as well as those that are not protected. You can select different segments of either the VM Protection Status pie chart or the Job History bar graph to view general information about the status of your jobs, resources, and history.



- **VM protection status**

When you select a segment in the VM Protection Status pie chart, you can view information about the protection status of the virtual machines in the Details pane. The descriptions for valid values are as follows:

- **OK**

- Displays the most recent successful backup of all virtual machines.

- **Failed**

- Displays the most recent failed backup of each virtual machine.

- **Unprotected**

- Displays the virtual machines that do not belong to any datasets and are therefore unprotected.

- **Job history**

When you select a segment in the Job History bar graph, you can view, in the Details pane, the history of completed, failed, and running jobs over a specified period of time. You can change the length of time that job details are displayed in the Job History bar graph. The default value is seven days.

- **Configure refresh**

You can change how often the dashboard refreshes the displayed information by using the **Configure refresh** button. The default value is 10 minutes.

Configure hosts

You can add, view, and remove Hyper-V parent hosts or clusters by using SnapManager for Hyper-V.

Requirements for adding a Hyper-V parent host or host cluster

You must have all necessary configuration information available before adding a parent host or host cluster to SnapManager for Hyper-V.

SnapManager for Hyper-V installation

SnapManager for Hyper-V must be installed on the Hyper-V host that you want to add.

If you do not have SnapManager for Hyper-V installed, you are prompted to run the Remote Host Install wizard. The same SnapManager for Hyper-V version must be installed on each cluster node.

Configuration settings

The Hyper-V parent host that you want to add must be configured for SnapManager for Hyper-V.

If the SnapInfo settings, report directory settings, and notification settings are not configured for SnapManager for Hyper-V, you can configure them after you add the host, using the Configuration wizard.

Initially, the **Manage Storage Connections** tab is empty. You can add the storage connections from the **Manage Storage Connections** tab, but the newly added connections are visible from SnapDrive for Windows (SDW) Transport Protocol Settings (TPS).

You must configure the backup repository and report directory settings to add and manage virtual machines by using SnapManager for Hyper-V. Notification settings are optional.

Virtual machines and ONTAP LUNs

All of the files associated with the virtual machines, including configuration files, Snapshot copy file location, and VHDs, must reside on ONTAP LUNs.

This is necessary to perform a successful backup.



If you change a virtual machine Snapshot file location to a different ONTAP LUN after creating the virtual machine, you should create at least one virtual machine Snapshot copy using Hyper-V Manager before making a backup by using SnapManager for Hyper-V. If you change the Snapshot copy file location and do not make a virtual machine Snapshot copy before making a backup, the backup operation could fail.

Dedicated and clustered virtual machines

Your virtual machines can be dedicated or part of a cluster.

If you add a single host, SnapManager for Hyper-V manages the dedicated virtual machines on that host. If you add a host cluster, SnapManager for Hyper-V manages the shared virtual machines on the host cluster. Virtual machines residing on SAN and NAS that belong to the same host cluster should not exist in the same dataset. Adding these types of resources to a single dataset can cause the dataset backup to fail.

For application-consistent backups, dataset backups of clustered virtual machines take longer to complete when the virtual machines run on different nodes of the cluster. When virtual machines run on different nodes, separate backup operations are required for each node in the cluster. If all virtual machines run on the same node, only one backup operation is required, resulting in a faster backup.

Number of virtual machines

If your Hyper-V host or host cluster has more than 1,000 virtual machines, you must increase the value of the maximum `Elements In Cache Before Scavenging` property in the `SnapMgrServiceHost.exe.config` file for Hyper-V Cache Manager. This value should be greater than or equal to the number of Hyper-V hosts running on a stand-alone host or cluster. The value should be changed on each node of the cluster, and the SnapManager for Hyper-V service must be restarted after changing this value. You must manually edit the `SnapMgrServiceHost.exe.config` file using a text editor.

```
<cacheManagers>
...
    <add name="HyperV Cache Manager"

type="Microsoft.Practices.EnterpriseLibrary.Caching.CacheManager,
        Microsoft.Practices.EnterpriseLibrary.Caching"
        expirationPollFrequencyInSeconds="60"
        maximumElementsInCacheBeforeScavenging="1200"
        numberToRemoveWhenScavenging="10"
        backingStoreName="inMemory" />
...
</cacheManagers>
```

SnapManager for Hyper-V service account requirements

When using SnapManager for Hyper-V to manage a Hyper-V host cluster, the SnapManager for Hyper-V and SnapDrive for Windows service accounts must be domain user accounts with local administrator rights on the server.

SnapManager for Hyper-V application-consistent backups run on the cluster node where the virtual machine is running. If Cluster Shared Volumes (CSVs) used by the virtual machine are not owned by the same node, virtual machine backups can fail when the SnapManager for Hyper-V service is using a local system account (even though the account has administrator privileges). In this case, SnapManager for Hyper-V cannot detect that the virtual machine files are residing on a CSV, causing the backup to fail.



For remote Volume Shadow Copy Service (VSS) operations with virtual machines stored on clustered Data ONTAP SMB 3.0 continuous availability (CA) shares to work properly, you must grant full control rights to the share for the SnapDrive for Windows service accounts and a minimum read-level access to the SnapManager for Hyper-V web service account.

Related information

[Microsoft TechNet: Hyper-V](#)

Add a Hyper-V parent host or host cluster

You can add a Hyper-V parent host or host cluster to back up and restore your virtual machines.

Steps

1. From the navigation pane, click **Protection**.
2. From the Actions pane, click **Add host**.
3. Run the **Add host** wizard.

After you finish

When you add a host to a cluster, the information about the new host is not automatically displayed in the GUI. Manually add the host information to the xml file in the installation directory.

SnapManager for Hyper-V must be installed on each cluster node. If you do not have SnapManager for Hyper-V installed, you are prompted to run the Remote Host Install wizard.

Manage storage connection settings

After you have added a host, you should enter all of the storage connections (SnapDrive for Windows and SnapManager for Hyper-V) for using Manage Storage Connection Settings in **Protection > Dataset Management**.

What you'll need

You must have at least one host added to SnapManager for Hyper-V before you can manage your storage connection settings.

Steps

1. From **Protection > Dataset Management**, select the **Manage Storage Connection** Settings.
2. Add the storage connections.

All storage connections can be viewed in SnapDrive for Windows TPS.

View a Hyper-V parent host or host cluster

You can view configuration information about a specific Hyper-V parent host or host cluster so that you can monitor its status.

Steps

1. From the navigation pane, click **Protection > Hosts**.
2. Select the host or host cluster that you want to view.

The Details pane displays the host or host cluster name, domain, cluster members (if applicable), and configuration messages. If you select a host that is not configured, the Details pane displays information about what is not configured.

View a virtual machine

From the Virtual Machine tab and VHD tab of the Details pane for a virtual machine, you can view information about and monitor the status of that machine.

Steps

1. From the navigation pane, click **Protection > Hosts > Protection > Datasets**.
2. Select the dataset or host to which the virtual machine belongs.
3. Select the appropriate virtual machine.

Results

The Virtual Machine tab displays the name, GUID, and state of the selected virtual machine.

The VHD tab displays system disk, mountpoint, VHD full path, LUN path, storage system name, serial number, and volume name associated with the selected virtual machine.

Migrate a Hyper-V virtual machine for SnapManager for Hyper-V operations

SnapManager for Hyper-V does not contain a migration wizard to help you migrate virtual machines (VMs) from non-ONTAP storage to ONTAP storage so that you can use them with SnapManager for Hyper-V. Instead, you must manually export and import the VM by using Server Manager.

Import or export host and dataset configuration information

Although you should manage a host from only one management console, if you need to do so from multiple consoles, you can import and export host and dataset configuration information from one remote management console to another to ensure data consistency.

About this task

You should not import or export configuration information to the directory on which SnapManager for Hyper-V is installed. If you uninstall SnapManager for Hyper-V, this file is lost.



You can use the Import and Export wizard to change host and dataset configuration settings to a previously exported setting. If you perform this operation in a clustered environment, you must import the settings on all nodes in the cluster so that all host and dataset configurations are the same.

Steps

1. From the navigation pane, click **Protection**.
2. From the **Actions** pane, click **Import and export**.

The Import and Export wizard appears.

3. Complete the steps in the wizard to export host and dataset configuration information from one management console to another.



The export file is static and current only at the time that export file was executed.

4. Complete the steps in the wizard to import host and dataset configuration information to the destination management console.

Remove a Hyper-V parent host or parent host cluster

You can remove a Hyper-V parent host or parent host cluster when you no longer want to manage it using SnapManager for Hyper-V.

Steps

1. From the navigation pane, click **Protection > Hosts**.
2. Select the host or host cluster you want to remove.
3. In the **Actions** pane, click **Remove**.

You can select **Delete all VM backups** to delete any virtual machine backups associated with the host.

The Hyper-V parent host or host cluster is removed from SnapManager for Hyper-V management but is not deleted permanently. The virtual machines belonging to that host or host cluster are also removed from any datasets to which they belonged.

Event notification settings

You can configure event notification settings to send e-mail, Syslog, and AutoSupport messages if an event occurs.

If event notification settings are not configured when a Hyper-V parent host is added to SnapManager for Hyper-V, you can configure those settings later, using the Configuration wizard.

You can change the event notification settings by using the Configuration wizard even after the host has been added to SnapManager for Hyper-V.

You can configure the event notification settings before you can add virtual machine resources to a dataset.

Configure email notifications

Multiple email recipients for notifications must be separated by commas.

When you configure multiple email recipients for email notifications in SnapManager for Hyper-V, separate each recipient with a comma. This requirement differs from SnapManager for SQL in which each email notification recipient must be separated by semicolons.

Report path settings

You can configure report path settings so that you can store reports for SnapManager for

Hyper-V operations. You must configure the report path settings before you can add virtual machine resources to a dataset.

If the report settings are not configured when a Hyper-V parent host is added to SnapManager for Hyper-V, you can configure (and even change) those settings later, using the Configuration wizard.

If you configure the report path settings for a parent host cluster, you must manually create the report directory on each cluster node. The report path should not reside on a cluster shared volume (CSV) or a shared LUN.

Related information

[Microsoft TechNet: Use Cluster Shared Volumes in a Failover Cluster](#)

Configure SnapInfo directory settings

You must configure the SnapInfo settings for a host before you can add the virtual machine resources within that host to a dataset. If the SnapInfo settings are not configured when a Hyper-V host is added to SnapManager for Hyper-V, you can configure those settings later, using the Configuration wizard or the **SnapInfo settings** action.

You can also change the SnapInfo settings after the host has been added to SnapManager for Hyper-V. However, if you change the SnapInfo settings, you must manually move all files to the new location; SnapManager for Hyper-V does not update them automatically. If you do not move the files, you cannot restore from or manage the backup copy, and SnapManager for Hyper-V does not list the backup copy.

Starting with SnapManager for Hyper-V, the SnapInfo path can reside on Cluster Shared Volumes (CSV) and it can also reside on SMB shares for Windows Server 2012.

Related information

[Microsoft TechNet: Use Cluster Shared Volumes in a Failover Cluster](#)

Set up a SnapInfo LUN

You must add a SnapInfo LUN in SnapManager for Hyper-V to store the dataset backup metadata. The SnapInfo path must reside on a ONTAP LUN, because SnapManager for Hyper-V makes a backup of the SnapInfo copy after a regular backup occurs.

What you'll need

The SnapInfo path can reside on a Cluster Shared Volume (CSV) if you are running a Windows Server 2012 cluster. If you manage dedicated virtual machines, the SnapInfo location must be to a dedicated ONTAP LUN. If you manage shared virtual machines, the SnapInfo location must be to a shared ONTAP LUN.

Steps

1. Create a new shared disk by using SnapDrive for Windows.
 - a. When given the option to choose a Microsoft Cluster Services Group, select the option **Create a new cluster group**.
 - b. Name the group `smhv_snapinfo` and complete the process.
2. Open Windows Failover Clustering (WFC) and verify that the new group is online.

3. Install SnapManager for Hyper-V on each node in the cluster.
4. Run the **Configuration** wizard and apply the SnapInfo configuration settings to all nodes in the cluster.
 - a. Select one of the hosts.
 - b. From the **Navigation** pane, click **Protection > Hosts**.
 - c. From the Actions pane, run the **Configuration** wizard.
 - d. Apply the SnapInfo settings to the newly created LUN.

Results

When the Configuration wizard is run, the SnapInfo configuration settings are replicated to all nodes in the cluster. **Related information**

[Error: Snapdrive SDDiscoveryFileSystemListInfo response is null](#)

Change the SnapInfo directory path

You can control SnapInfo directory path settings by using the Configuration wizard or the **SnapInfo settings** action.

About this task

SnapInfo directory settings are specified on the host level in SnapManager for Hyper-V. SnapManager for Hyper-V supports NAS (SMB) hosts and SAN hosts. For SAN hosts, the SnapInfo settings are applied on the volume level; for NAS hosts, the SnapInfo settings are applied on the SMB share level.

If you have added the IP address of the storage system to SnapDrive for Windows TPS, the storage settings from SnapDrive for Windows are automatically populated when you run the configuration wizard in SnapManager for Hyper-V. If you do not have SnapDrive for Windows TPS configured, you must specify the IP address of the storage system in the Manage Storage Connections tab in SnapManager for Hyper-V.

Steps

1. From the navigation pane, click **Protection > Hosts**.
2. Select the host for which you want to change the SnapInfo directory path.
3. From the **Actions** pane, select **SnapInfo settings**.

The **SnapInfo Settings** dialog box opens.

4. Select the storage type from the options shown:

Option	Description
SAN	This is the default storage type.
NAS	Use this option for SMB shares.

5. Click **Browse**.

The Browse for Folder window opens.

6. Select your SnapInfo storage system (SAN) or volume (NAS) and click **OK**.

The hosts that display are NAS shares corresponding to the storage systems that have been registered using the `Manage Storage Connections` option at the host level. If you do not see the shares you are looking for, ensure that `Manage Storage Connections` has been configured correctly.

7. From the **SnapInfo Settings** window, click **OK**.

Configure datasets

You can create, modify, view, and delete datasets based on your protection needs.

Requirements for creating a dataset

You must meet specific requirements when you want to create datasets to protect your data. You must first add the host or host cluster to SnapManager for Hyper-V and then add virtual machines to the host or host cluster.

Dataset name and description

When naming the dataset, you should use a naming convention at your site to help administrators locate and identify datasets, limited to these characters:

- a to z
- A to Z
- 0 to 9
- _ (underscore)
- - (hyphen)

Dataset resources

You must add the host or host cluster to SnapManager for Hyper-V before adding resources such as virtual machines to the dataset.

You can add hosts, dedicated virtual machines, or shared virtual machines to a dataset. If you add a host, you add all of the virtual machines that belong to the host. You can also add virtual machines belonging to different hosts to the dataset. Virtual machines can belong to multiple datasets.



Dedicated and shared virtual machines that belong to the same host cluster should not exist in the same dataset. Adding these types of resources to a single dataset can cause the dataset backup to fail.

Virtual machines and ONTAP LUNs

All of the files associated with the virtual machines, including configuration files, Snapshot copies, and VHDs, must reside on ONTAP LUNs.

Dataset resource consumption

Only one application-consistent backup operation can occur on a host at any given time. If the same virtual machines belong to different datasets, you should not schedule an application-consistent backup of the datasets at the same time. If this occurs, one of the backup operations fails.

When creating a dataset, you should select all virtual machines that reside on a particular ONTAP LUN. This enables you to get all backups in one Snapshot copy and to reduce the space consumption on the storage system.

Create a dataset

You can create datasets for virtual machine resources that share the same protection requirements. You can add virtual machines to multiple datasets, as necessary.

What you'll need

You must have the following information available:

- Dataset name and description
- Virtual machine resources you plan to add to the dataset

About this task

Dedicated and shared disks that belong to the same host cluster should not be placed in the same dataset. Adding these types of resources to a single dataset can cause the dataset backup to fail. You can have only one type of VM per dataset: NAS or SAN. You cannot have mixed-mode datasets.

The Validate Dataset check box is selected by default. SnapManager for Hyper-V checks for any configuration errors in all VMs during the creation or modification of a dataset. You must ensure that the check box is not selected if you do not want to enable validation of the dataset.

Steps

1. From the navigation pane, click **Protection > Datasets**.
2. From the Actions pane, click **Create dataset**.
3. Complete the pages of the wizard.

After you finish

You should next add protection policies to the dataset that you created.

Modify a dataset

After you have created a dataset, you can modify the dataset description and the resources associated with the dataset.

About this task

The Validate Dataset check box is selected by default. SnapManager for Hyper-V checks for any configuration errors in all VMs during the creation or modification of a dataset. You must ensure that the check box is not selected if you do not want to enable validation of the dataset.

Steps

1. From the navigation pane, click **Protection > Datasets**.
2. Select the dataset that you want to modify.
3. From the Actions pane, click **Modify dataset**.
4. Complete the steps in the wizard.

View a dataset

You can view the virtual machines associated with a dataset.

Steps

1. From the navigation pane, click **Protection > Datasets**.
2. Expand the tree view to look at the virtual machines that belong to the dataset.

Delete a dataset

You can delete a dataset as your protection needs change.

About this task

Deleting a dataset does not delete the virtual machines that belong to the dataset. After the dataset is deleted, the virtual machines that belonged to it become unprotected if they do not belong to another dataset.

Steps

1. From the navigation pane, click **Protection > Datasets**.
2. Select the dataset that you want to delete.
3. From the Actions pane, click **Delete** and click **OK** to delete the dataset.

Deleting the dataset also deletes the scheduled jobs from all hosts that are members of the dataset. SnapManager for Hyper-V no longer manages the retention of backups when you delete the dataset, even if it is re-created with the same name.

Results

SnapManager for Hyper-V no longer protects the resources associated with the deleted dataset unless those resources belong to another dataset.

configure policies

You can add, modify, or delete policies associated with datasets so that you can protect your data.

Requirements for adding policies to a dataset

You must meet specific requirements when you want to apply policies to datasets for backup or restore functionality. You can add multiple retention, scheduling, and replication policies to the same dataset.

Policy name and description

A policy name and description, limited to these characters:

- a to z
- A to Z
- 0 to 9

- _ (underscore)
- - (hyphen)

Backup retention limits

You must decide the minimum length of time that you want to keep your hourly, daily, weekly, or monthly backup copies before they are deleted.



Backups with a retention type of “Unlimited” are not deleted.

You can keep backups based on either time or a specified number. For example, you can keep the 10 most current backups, or you can delete backups older than 15 days.

If your system appears to retain old backups, you should check your retention policies. All objects being backed up that share a Snapshot copy must meet the backup deletion criteria for the retention policy to trigger the removal of a Snapshot copy.

Scheduled backup job name

You must assign a name to the scheduled backup job.

Permission to schedule backups

You must have appropriate credentials to schedule dataset backups.

Number of possible datasets scheduled for backup simultaneously

If the same virtual machines belong to different datasets, you should not schedule a backup of more than one dataset containing the same VM at the same time. When this occurs, one of the backup operations fails. Only one backup operation can occur on a host at any given time.

Type of scheduled backup

You can perform either an application-consistent or a crash-consistent backup.

Backup options

You must choose if you want to update the SnapMirror destination location after the backup is complete.

The update succeeds only if you already have SnapMirror configured, and if the LUNs containing the virtual machines in the dataset belong to the source SnapMirror volume.

The default behavior of SnapManager for Hyper-V is to fail a backup if one or more virtual machines cannot be backed up online. If a virtual machine is in the saved state or shut down, you cannot perform an online backup. In some cases, virtual machines are in the saved state or shut down for maintenance, but backups still need to proceed, even if an online backup is not possible. To do this, you can move the virtual machines that are in the saved state or shut down to a different dataset with a policy that allows saved state backups.

You can also select the Allow saved state VM backup check box to enable SnapManager for Hyper-V to back up the virtual machine in the saved state. If you select this option, SnapManager for Hyper-V does not fail the backup when the Hyper-V VSS writer backs up the virtual machine in the saved state or performs an offline backup of the virtual machine. Performing a saved-state or offline backup can cause downtime.

The distributed application-consistent backup feature enables multiple VMs running on the partner cluster nodes to be consistent in a single hardware Snapshot copy made from the backup node. This feature is supported for all the VMs running on a CSV 2.0 Windows volume across multiple nodes in a Windows Failover Cluster. To use this feature, select the Application-consistent backup type and select the **Enable Distributed Backup** check box.

Secondary storage in a SnapMirror backup


These options enable you to accept options applicable to a secondary storage defined in a SnapMirror relationship. Here, you can select **Update SnapMirror after backup**. Under the Vault label option pane, you can select **Update SnapVault after backup**. If you select **Update SnapVault after backup**, you must choose a vault label from the drop-down menu or enter a custom label.

Backup scripts

You must decide if you want the optional backup scripts to run either before or after the backup takes place.

These scripts run on all dataset member hosts unless you indicate a specific server.

Backup scripts run on each node in the dataset. You can set the dataset policy to specify the name of the host on which you want to run the scripts. The policy is processed on each node in the cluster where the VM to be backed up is running.



You can use the following environment variables in arguments for backup postscripts:

- **\$VMSnapshot**: Specifies the name of the Snapshot copy that is created on the storage system as a result of this backup. When you are performing application-consistent backups in ONTAP environments running in 7-Mode, this is the name of the second (backup) Snapshot copy. The first name is same as the second name but without the `_backup` suffix.
- **\$SnapInfoName**: Specifies the timestamp used in the SnapInfo directory name.
- **\$SnapInfoSnapshot**: Specifies the name of the SnapInfo Snapshot copy created on the storage system. SnapManager for Hyper-V makes a Snapshot copy of the SnapInfo LUN at the end of the dataset backup operation.



The **\$SnapInfoSnapshot** variable is supported for dedicated virtual machines only.

Related information

[Microsoft TechNet: Hyper-V](#)

Add policies

You can add retention, scheduling, and replication policies, as well as scripts, to your datasets so that you can protect your data.

What you'll need

You must have the following information available:

- Policy names and descriptions
- Retention information
- Scheduling information
- Backup options information
- Backup script information

Steps

1. From the navigation pane, click **Protection > Datasets**.
2. Select the dataset to which you want to add policies.
3. From the Actions pane, click **Add policy**.

The Create Policy wizard appears.

4. Complete the steps in the wizard to create protection policies for your dataset.

Modify policies

You can modify the policies that protect your datasets by using the Modify Policy wizard.

Steps

1. From the navigation pane, click **Protection > Datasets**.
2. Select the dataset that contains the policies that you want to modify.
3. Select the policy that you want to modify.
4. From the Actions pane, click **Modify policy**.

The Modify Policy wizard appears.

5. Complete the steps in the wizard to modify the protection policy for your dataset.

View policies

You can view policy details associated with a specific dataset.

Steps

1. From the navigation pane, click **Protection > Datasets**.
2. Select the dataset that contains the policies that you want to view.
3. From the Policies pane, select the specific policy for which you want to view details.

Information about the policy appears in the Details pane.

Delete policies

You can delete a policy from a dataset when it is no longer needed.

Steps

1. From the navigation pane, click **Protection > Datasets**.
2. Select the dataset that contains the policy that you want to delete.
3. From the Policies pane, select the specific policy that you want to delete.
4. From the Actions pane, click **Remove** and click **OK** to delete the policy.

Deleting the policy also deletes the scheduled jobs from all hosts that are members of the dataset.

Configure SVMs or CIFS servers for Hyper-V over SMB

Configuring a single storage virtual machine (SVM) or CIFS server for multiple applications can lead to resource sharing issues, which in turn impact the Hyper-V environment. You should configure dedicated SVMs or CIFS servers for Hyper-V over SMB, depending on your requirements.

Related information

[NetApp KB Article 1015099: How to set up SVM/CIFS for Hyper-V over SMB](#)

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. 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. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.