



Configure policies

SnapManager for Hyper-V

NetApp
June 18, 2025

This PDF was generated from <https://docs.netapp.com/us-en/snapmanager-hyper-v/reference-requirements-for-adding-policies-to-a-dataset.html> on June 18, 2025. Always check docs.netapp.com for the latest.

Table of Contents

configure policies	1
Requirements for adding policies to a dataset	1
Policy name and description	1
Backup retention limits	1
Scheduled backup job name	1
Permission to schedule backups	1
Number of possible datasets scheduled for backup simultaneously	1
Type of scheduled backup	2
Backup options	2
Secondary storage in a SnapMirror backup	2
Backup scripts	2
Add policies	3
Modify policies	4
View policies	4
Delete policies	4

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](#)

[ONTAP 9 Volume Backup Using SnapVault Express Guide](#)

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.

Copyright information

Copyright © 2025 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.