



Snapshot copy operations

Snapdrive for Unix

NetApp

June 20, 2025

This PDF was generated from https://docs.netapp.com/us-en/snapdrive-unix/aix/concept_crash_consistency_with_data_ontap_7_2_and_later.html on June 20, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Snapshot copy operations	1
Crash-consistent Snapshot copies	1
Crash consistency with Data ONTAP 7.2 and later	1
Consistency group Snapshot copies in SnapDrive for UNIX	2
Creating a consistency group Snapshot copy	2
Disabling consistency group Snapshots copies	2
Application-consistent Snapshot copies	3
Snapshot copies that span storage systems or volumes	3
Creating Snapshot copies of unrelated entities	4
Guidelines for Snapshot copy creation	5
Information required for using the snapdrive snap create command	5
Creating a Snapshot copy	7

Snapshot copy operations

You can create Snapshot copies by using the `snapdrive snap create` command.

Crash-consistent Snapshot copies

You might have to create a crash-consistent Snapshot copies of your file system or disk groups. SnapDrive for UNIX creates Snapshot copies that contain the image of all the storage system volumes specified in the entity.

When you create a Snapshot copy of a storage entity, such as a file system or disk group, SnapDrive for UNIX creates a Snapshot copy that contains the image of all the storage system volumes that comprise the entity you specified using a `file_spec` argument. The `file_spec` argument specifies the storage entity, such as the file system, LUN, or NFS directory tree that SnapDrive for UNIX uses to create the Snapshot copy.

SnapDrive for UNIX makes consistent storage components that compose the entity you requested in the Snapshot copy. This means that LUNs or directories being used outside those specified by the `snapdrive snap create` command `file_spec` argument might not have consistent images in the Snapshot copy. SnapDrive for UNIX enables you to restore only the entities specified by the `file_spec` argument that are consistent in the Snapshot copy.

Snapshot copies of entities contained on a single storage system volume are always crash-consistent. SnapDrive for UNIX takes special steps to ensure that Snapshot copies that span multiple storage systems or storage system volumes are also crash-consistent. The method that SnapDrive for UNIX uses to ensure crash consistency depends on the Data ONTAP version used where the storage entities in your Snapshot copy resides.

Crash consistency with Data ONTAP 7.2 and later

SnapDrive for UNIX uses the support for consistency groups provided by Data ONTAP 7.2 and later versions, such that all Snapshot copies that span multiple volumes are crash consistent.

Data ONTAP versions 7.2 and greater provides support for consistency groups and storage system fencing. SnapDrive for UNIX uses these features to ensure that all Snapshot copies that span multiple volumes are crash consistent.

To create a crash consistent Snapshot copy across multiple volumes, SnapDrive for UNIX does the following:

- Fences (freezes) I/O to every volume that contains a storage entity.
- Takes a Snapshot copy of each volume.

The time it takes to fence the volume and create the Snapshot copy is limited, and is controlled by Data ONTAP.

The `snapcreate-cg-timeout` parameter in the `snapdrive.conf` file specifies the amount of time, within Data ONTAP limitations, that you wish to allow for storage system fencing. You can specify an interval that is urgent, medium, or relaxed. If the storage system requires more time than allowed to complete the fencing operation, SnapDrive for UNIX creates the Snapshot copy using the consistency methodology for previous Data ONTAP 7.2 versions. You can also specify this methodology by using the `-nofilerfence` option when you create the Snapshot copy.

If you request a Snapshot copy for a storage entity that spans storage systems with both Data ONTAP 7.2 and previous Data ONTAP versions, SnapDrive for UNIX also creates the Snapshot copy using the consistency method for Data ONTAP versions before 7.2.

Consistency group Snapshot copies in SnapDrive for UNIX

Consistency Group Snapshot is a Snapshot copy of a set of volumes that span different Vservers or nodes of a cluster, which is managed as a single entity.

SnapDrive for UNIX captures crash-consistent Snapshot copies across all volumes spanning different Vservers or nodes of a cluster. You can also configure the time period within which the Snapshot copy is to be captured.

SnapDrive for UNIX captures consistency group Snapshot copies by default. You can disable this feature and revert to capturing Snapshot copies in best-effort mode.



SnapDrive for UNIX 5.2 supports consistency group Snapshot copies for clustered Data ONTAP only in Data ONTAP 8.2 or later versions.

Related information

[Creating a consistency group Snapshot copy](#)

[Disabling consistency group Snapshots copies](#)

Creating a consistency group Snapshot copy

You can configure SnapDrive for UNIX to create a consistency group Snapshot copy.

Steps

1. Enter the following command on the host:

```
snapdrive snap create -fs /mnt/test -snapname snapshotname -f -noprompt.
```

snapshotname is the name specified for the consistency group Snapshot copy.

Example

The following is an example of the command:

```
snapdrive snap create -fs /mnt/test -snapname snap_123 -f -noprompt
```

The consistency group Snapshot copy is successfully created.

Disabling consistency group Snapshots copies

You can configure SnapDrive for UNIX to disable a consistency group Snapshot copy.

Steps

1. Enter:

```
snapdrive snap create -fs /mnt/fs3 -snapname nfs_snap -nofilerfence
```

The consistency group Snapshot copy is successfully disabled.

Application-consistent Snapshot copies

To make an application-consistent Snapshot copy, you should halt the application before the Snapshot operation.

To ensure that a Snapshot copy is application-consistent, you might need to stop or perform the necessary steps to quiesce the application before taking the Snapshot copy. Note that database hot backup facilities depend on the methods used by the DBMS, and do not always quiesce I/O to database files.

If the application has not completed its transactions and written data to the storage system, the resulting Snapshot copy might not be application-consistent.



If your application can recover from a crash-consistent Snapshot copy, you do not need to stop it. Consult the documentation for your application. For more information about taking application-consistent Snapshot copies.

You should take a new Snapshot copy whenever you add or remove a host volume, LUN, or NFS directory tree, or resize host volumes or file systems. This ensures that you have a consistent copy of the newly configured disk group that you can use if you need to restore the disk group.

Snapshot copies that span storage systems or volumes

SnapDrive for UNIX enables you to take Snapshot copies that reside in multiple storage system volumes on the same or different storage systems.

SnapDrive for UNIX allows you to take Snapshot copies that span across multiple storage system volumes or multiple storage systems. These volumes can reside on the same storage system or different storage systems. Although the `snapdrive snap create` command creates a Snapshot copy of all the volumes that comprise the entity you request, SnapDrive for UNIX restores only the entities that you specify in the `snapdrive snap create` command.

When you use the `snapdrive snap create` command to make a Snapshot copy that spans multiple volumes, you do not need to name the volumes on the command prompt. SnapDrive for UNIX gets this information from the `file_spec` argument that you specify.

- If the `file_spec` you enter requests a disk group, or a file system or host volume that resides on a disk group, SnapDrive for UNIX automatically creates a Snapshot copy that includes all the storage system volumes for the disk group, volume, or file system you specified.
- If the `file_spec` you enter requests a LUN, SnapDrive for UNIX takes a Snapshot copy of the storage system volume that contains the LUN.
- If the `file_spec` you enter requests a file system that resides directly on a LUN, SnapDrive for UNIX takes a Snapshot copy of the storage system volume that contains the LUN and file system that you specified.
- If the `file_spec` you enter requests an NFS directory, SnapDrive for UNIX creates a Snapshot copy of the volume that contains the NFS directory tree.

In addition to using a `file_spec` argument that is built on entities from multiple storage systems and storage system volumes, you can also use a combination of `file_spec` arguments where each value is based on single storage system or storage system volume. For example, suppose you have a setup where the disk group `dg1` spans the storage systems `storage system2` and `storage system3`, `dg2` is on `storage system2`, and `dg3` is on `storage system3`. In this case, any of the following command lines would be correct:

```
snapdrive snap create -dg dg1 -snapname snapdg1  
snapdrive snap create -dg dg2 dg3 -snapname snapdg23  
snapdrive snap create -dg dg1 dg2 dg3 -snapname snapdg123
```

Something to keep in mind when creating Snapshot copies that span storage systems and volumes is that SnapDrive for UNIX creates the Snapshot copy on each storage systems volume using a short name. It does not include the full path name in the name, even if the volumes are on different storage system. This means that if you later rename the Snapshot copy, you must go to each storage system and volume and rename it there as well.

Creating Snapshot copies of unrelated entities

SnapDrive for UNIX creates Snapshot copies of unrelated entities by maintaining individual crash-consistent Snapshot copies.

Unless you specify otherwise, SnapDrive for UNIX assumes that all entities that you specify on a particular `snapdrive snap create` command line are related; in other words the validity of updates to one entity can depend on updates to the other entities specified. When storage entities have dependent writes in this way, SnapDrive for UNIX takes steps to create a Snapshot copy that is crash consistent for all storage entities as a group.

The following example shows how SnapDrive for UNIX creates a Snapshot copy of storage entities that may have dependent writes. In the following example, the `snapdrive snap create` command specifies a file system on a LUN and also a disk group. The disk group consists of LUNs residing on a single storage system (see [Creating a Snapshot copy](#)). The file system on a LUN resides on a different storage system and storage system volume. As a group, the file system and the disk group span multiple storage system volumes; individually they do not.

The following command specifies a Snapshot copy that contains both the file system `/mnt/fs1` and the disk group `dg1`:

```
snapdrive snap create -fs /mnt/fs1 -dg dg1 -snapname fs1_dg1
```

Because these storage entities can have dependent writes, SnapDrive for UNIX attempts to create a crash-consistent Snapshot copy, and treats the file system `/mnt/fs1` and the disk group `dg1` as a group. This means SnapDrive for UNIX is required to freeze I/O operations to the storage system volumes before creating the Snapshot copy.

Creating crash-consistent Snapshot copies for multiple storage entities across volumes takes extra time, and is not always possible if SnapDrive for UNIX cannot freeze I/O operations. Because this is so, SnapDrive for UNIX allows you to create Snapshot copies of unrelated storage entities. Unrelated storage entities are entities that you can specify that have no dependent writes when the Snapshot copy is taken. Because the entities have no dependent writes, SnapDrive for UNIX does not take steps to make the entities consistent as a group. Instead, SnapDrive for UNIX creates a Snapshot copy in which each of the individual storage entities is crash-consistent.

The following command specifies a Snapshot copy of the file system on a LUN and the disk group described

previously. Because the `-unrelated` option is specified, SnapDrive for UNIX creates a Snapshot copy in which the file system `/mnt/fs1` and the disk group `dg1` are crash-consistent as individual storage entities, but are not treated as a group. The following command does not require SnapDrive for UNIX to freeze I/O operations on the storage system volumes: `snapdrive snap create -fs /mnt/fs1 -dg dg1 -unrelated -snapname fs1_dg1`

Related information

[Creating a Snapshot copy](#)

Guidelines for Snapshot copy creation

Consider the following while creating Snapshot copies using SnapDrive for UNIX: you can keep maximum 255 Snapshot copies per volume, SnapDrive for UNIX supports only the Snapshot copies that it creates, you cannot create Snapshot copies of root disk groups, and boot device or swap device, and SnapDrive for UNIX requires a freeze operation to maintain crash-consistency.

Follow these guidelines when you enter commands that create Snapshot copies:

- You can keep a maximum of 255 Snapshot copies per storage system volume. This limit is set by the storage system. The total number can vary depending on whether other tools use these Snapshot copies.

When the number of Snapshot copies has reached the maximum limit, the Snapshot create operation fails. You must delete some of the old Snapshot copies before you can use SnapDrive for UNIX to take anymore.

- SnapDrive for UNIX does not support Snapshot copies that it does not create. For example, it does not support Snapshot copies that are created from the storage system console, because such a practice can lead to inconsistencies within the file system.
- You cannot use SnapDrive for UNIX to create Snapshot copies of the following:
 - Root disk groups

The Snapshot create operation fails when you try to take a Snapshot copy of a root disk group for an LVM.

- Boot device or swap device

SnapDrive for UNIX does not take a Snapshot copy of a system boot device or a system swap device.

- When a Snapshot copy spans multiple storage systems or storage system volumes, SnapDrive for UNIX requires a freeze operation to guarantee crash-consistency. For information about creating Snapshot copies on configurations for which a freeze operation is not provided.

Information required for using the `snapdrive snap create` command

When you create a Snapshot copy, you should determine the storage entity that you want to capture and specify a name of the Snapshot copy.

The following table provides the information you need when you use the `snapdrive snap create`

command:

Requirement/Options	Argument
Determine the type of storage entity you want to capture in the Snapshot copy. You can specify NFS entities, LUNs, file systems created directly on LUNs, and LVM entities on a single command line.	
Supply that entity's name with the appropriate argument. This is the value for the <code>file_spec</code> argument.	
<ul style="list-style-type: none">• If you specify a disk group that has a host volume or file specification, the argument translates into a set of disk groups on the storage system. SnapDrive for UNIX creates the entire disk group containing the entity, even if the entity is a host volume or file system.• If you specify a file specification that is an NFS mount point, the argument translates to the directory tree on the storage system volume.• If you specify a LUN, or a LUN that has a file system, the argument translates to the LUN, or to the LUN that has the file system.	
 You cannot specify special characters with the storage entities such as <code>-vg</code> , <code>-dg</code> , <code>-fs</code> , <code>-lvol</code> , and <code>-hostvol</code> . However, slash (/) is allowed for storage entities such as <code>-fs</code> , <code>-lvol</code> , and <code>-hostvol</code> .	
LUN (<code>-lun file_spec</code>)	<p>Name of the LUN</p> <p>You must include the name of the storage system, volume, and LUN.</p>
Disk group (<code>-dg file_spec</code>) or volume group (<code>-vg file_spec</code>)	Name of the disk or volume group
File system (<code>-fs file_spec</code>)	<code>filesystem_name</code>
Host volume (<code>-hostvol file_spec</code>) or logical volume (<code>-lvol file_spec</code>)	<p>Name of the host or logical volume</p>  You must supply both the requested volume and the disk group containing it; for example, <code>-hostvol dg3/acct_volume</code> .
Snapshot copy name (<code>-snapname snap_name</code>)	<code>Snapshot copy_name</code>
Specify the name for the Snapshot copy. This can be either the long version of the name that includes the storage system and volume with the Snapshot copy name or the short version that is the Snapshot copy name.	
 You must ensure that the value specified for <code>snapname</code> does not exceed 200 characters.	
<code>-unrelated</code>	<code>~</code>

Requirement/Options	Argument
Optional: To create a Snapshot copy of storage entities that have no dependent writes when the Snapshot copy is created, SnapDrive for UNIX creates a crash-consistent Snapshot copy of the individual storage entities, but does not try to make the entities consistent with each other.	
-force	~
-noprompt	~
Optional: Decide if you want to overwrite an existing Snapshot copy. Without this option, this operation halts if you supply the name of an existing Snapshot copy. When you supply this option and specify the name of an existing Snapshot copy, the command prompts you to confirm that you want to overwrite the Snapshot copy. To prevent SnapDrive for UNIX from displaying the prompt, include the -noprompt option also. (You must always include the -force option if you want to use the -noprompt option.)	
-fstype	<i>type</i>
-vmttype	<i>type</i>
Optional: Specify the type of file system and volume manager to be used for SnapDrive for UNIX operations.	

Creating a Snapshot copy

To create a Snapshot copy, run the `snapdrive snap create` command.

Before you execute this syntax, you must understand the options, keywords, and arguments mentioned in this command.

Steps

1. Enter the following command syntax to create a Snapshot copy:

```
snapdrive snap create {-lun | -dg | -fs | -hostvol } file_spec[file_spec ...] [ -lun | -dg | -fs | -hostvol ] file_spec [file_spec...]] -snapname snap_name [-force [-noprompt]] [-unrelated] [-nofilerfence] [-fstype type] [-vmttype type]
```

The *file_spec* arguments represent a set of storage entities on one or more storage systems. The Snapshot create operation takes a Snapshot copy of the storage system volume containing those entities and gives it the name specified in the *snap_name* argument.

Related information

[Creating Snapshot copies of unrelated entities](#)

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—with 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.