



SnapDrive for UNIX options, keywords, and arguments

Snapdrive for Unix

NetApp
February 12, 2024

This PDF was generated from https://docs.netapp.com/us-en/snapdrive-unix/aix/reference_command_line_options.html on February 12, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- SnapDrive for UNIX options, keywords, and arguments 1
 - Command-line options 1
 - Rules for keywords. 5
 - Command-line keywords 6
 - Command-line arguments 8

SnapDrive for UNIX options, keywords, and arguments

SnapDrive for UNIX commands has options, keywords, and arguments associated with them.

Command-line options

There are various options that are used with SnapDrive for UNIX commands.

SnapDrive for UNIX enables you to include the following options as appropriate with its commands. In certain cases, you can abbreviate these options. For example, you can use `-h` instead of `-help`.

| Option | Description |
|--|--|
| <code>-addlun</code> | Tells SnapDrive for UNIX to add a new, internally-generated LUN to a storage entity in order to increase its size. |
| <code>-all</code> | Used with the <code>snapdrive storage {show list}</code> command to display all devices and LVM entities known to the host. |
| <code>-autoexpand</code> | Used with the <code>snapdrive snap connect</code> command to enable you to request that a disk group be connected when you supply a subset of the logical volumes or file systems in the disk group. |
| <code>-autorename</code> | Used with the <code>snapdrive snap connect</code> command to enable the command to rename any newly-connected LVM entities for which the default name is already in use. |
| <code>-clone type</code> | Clone-method to be used during <code>snap connect</code> operation. Here type refers to <code>lunclone</code> (creates lun clone(s)), <code>Optimal</code> (SnapDrive automatically chooses between <code>Restricted FlexClone(s)</code> and <code>LUN clone</code> depending on the storage configuration) and <code>Unrestricted</code> (creates <code>FlexClone(s)</code> which can be used for provisioning and Snapshot operations, just like normal flexible volumes). |
| <code>-capabilities</code> | used with <code>snapdrive storage show</code> command to know the allowed operations on the host file specs. |
| <code>-devices</code> or <code>-dev</code> | Used with the <code>storage {show list}</code> command to display all devices known to the host. |

| Option | Description |
|--|---|
| <code>-dgsiz</code> or <code>-vgsiz</code> | Used with the <code>snapdrive storage create</code> command to specify the size in bytes of the disk group you want to create. |
| <code>-force</code> (or <code>-f</code>) | Causes operations to be attempted that SnapDrive for UNIX would not undertake ordinarily. SnapDrive for UNIX prompts you to ask for confirmation before it executes the operation. |
| <code>-fsops</code> | <p>The options you want passed to the host operation that creates the new file system. Depending on your host operating system, this host operation might be a command such as the <code>mkfs</code> command.</p> <p>The argument you supply with this option usually needs to be specified as a quoted string and must contain the exact text to be passed to the command.</p> <p>For example, you might enter <code>-o largefiles</code> as the option you want passed to the host operation.</p> |
| <code>-fstype</code> | <p>The type of file system you want to use for the SnapDrive for UNIX operations. The file system must be a type that SnapDrive for UNIX supports for your operating system. Current values that you can set for this variable is “jfs”, “jfs2” or “vxfs”</p> <p>You can also specify the type of file system that you want to use by using the <code>-fstype</code> configuration variable.</p> |
| <code>-full</code> | Allows operations on a specified host-side entity to be performed even if the entity is not empty (for example, the entity might be a volume group containing one or more logical volumes). |
| <code>-growby</code> | The number of bytes you want to add to a LUN or disk group in order to increase its size. |
| <code>-growto</code> | The target size in bytes for a LUN, disk group, or volume group. SnapDrive for UNIX automatically calculates the number of bytes necessary to reach the target size and increases the size of the object by that number of bytes. |
| <code>-help</code> | Prints out the usage message for the command and operation. Enter this option by itself without other options. Following are the examples of possible command lines. |

| Option | Description |
|----------------------------|--|
| <code>-lunsize</code> | The size of the LUN in bytes to be created by a given command. |
| <code>-mntopts</code> | <p>Specifies options that you want passed to the host mount command (for example, to specify file system logging behavior). Options are also stored in the host file system table file. The options allowed depend on the host file system type.</p> <p>The <code>-mntopts</code> argument that you supply is a file system-type option that is specified using the mount command “- o” flag. Do not include the “- o” flag in the <code>-mntopts</code> argument. For example, the sequence <code>-mntopts tmplo_g</code> passes the string <code>-o tmplog_</code> to the mount command line, and inserts the text “tmplog” on a new command line.</p> |
| <code>-nofilerfence</code> | <p>Suppresses the use of the Data ONTAP consistency group feature in creating Snapshot copies that span multiple filer volumes.</p> <p>In Data ONTAP 7.2 or above, you can suspend access to an entire filer volume. By using the <code>-nofilerfence</code> option, you can freeze access to an individual LUN.</p> |
| <code>-nolv</code> | <p>Connects or creates a file system directly on a LUN without involving the host LVM.</p> <p>All commands that take this option for connecting or creating a file system directly on a LUN will not accept it for host cluster or shared resources. This option is allowed only for local resources.</p> |
| <code>-nopersist</code> | Connects or creates a file system, or a Snapshot copy that has a file system, without adding an entry in the host's persistent mount entry file. |
| <code>-prefixfv</code> | prefix to be used while generating cloned volume name. The format of the name of the new volume would be <code><pre-fix>_<original_volume_name></code> . |

| Option | Description |
|---|--|
| <code>-reserve - noreserve</code> | Used with the <code>snapdrive storage create</code> , <code>snapdrive snap connect</code> or <code>snapdrive snap restore</code> commands to specify whether or not SnapDrive for UNIX creates a space reservation. By default, SnapDrive for UNIX creates reservation for storage create, resize, and Snapshot create operations, and does not create reservation for Snapshot connect operation. |
| <code>-noprompt</code> | Suppresses prompting during command execution. By default, any operation that might have dangerous or non-intuitive side effects prompts you to confirm that SnapDrive for UNIX should be attempted. This option overrides that prompt; when combined with the <code>-force</code> option, SnapDrive for UNIX performs the operation without asking for confirmation. |
| <code>-quiet</code> (or <code>-q</code>) | <p>Suppresses the reporting of errors and warnings, regardless of whether they are normal or diagnostic. It returns zero (success) or non-zero status. The <code>-quiet</code> option overrides the <code>-verbose</code> option.</p> <p>This option will be ignored for <code>snapdrive storage show</code>, <code>snapdrive snap show</code>, and <code>snapdrive config show</code> commands.</p> |
| <code>-readonly</code> | <p>Required for configurations with Data ONTAP 7.1 or any configuration that uses traditional volumes. Connects the NFS file or directory with read-only access.</p> <p>Optional for configurations with Data ONTAP 7.0 that use FlexVol volumes. Connects the NFS file or directory tree with read-only access. (Default is read/write).</p> |
| <code>-split</code> | <p>Enables to split the cloned volumes or LUNs during Snapshot connect and Snapshot disconnect operations.</p> <p>You can also split the cloned volumes or LUNs by using the <code>enable-split-clone</code> configuration variable.</p> |
| <code>-status</code> | Used with the <code>snapdrive storage show</code> command to know if the volume or LUN is cloned. |

| Option | Description |
|---|--|
| <code>-unrelated</code> | Creates a Snapshot copy of <code>file_spec</code> entities that have no dependent writes when the Snapshot copy is taken. Because the entities have no dependent writes, SnapDrive for UNIX creates a crash-consistent Snapshot copy of the individual storage entities, but does not take steps to make the entities consistent with each other. |
| <code>-verbose</code> (or <code>-v</code>) | Displays detailed output, wherever appropriate. All commands and operations accept this option, although some might ignore it. |
| <code>-vgsize</code> or <code>-dgsiz</code> | Used with the <code>storage create</code> command to specify the size in bytes of the volume group you want to create. |
| <code>-vmtype</code> | <p>The type of volume manager you want to use for the SnapDrive for UNIX operations.</p> <p>If the user specifies the <code>-vmtype</code> option in the command line explicitly, SnapDrive for UNIX uses the value specified in the option irrespective of the value specified in the <code>vmtype</code> configuration variable. If the <code>-vmtype</code> option is not specified in the command-line option, SnapDrive for UNIX uses the volume manager that is in the configuration file.</p> <p>The volume manager must be a type that SnapDrive for UNIX supports for your operating system. Current values that you can set for this variable as <code>vxvm</code> or <code>lvm</code>.</p> <p>You can also specify the type of volume manager that you want to use by using the <code>vmtype</code> configuration variable.</p> |
| <code>-vbsr {preview execute}</code> | The <code>preview</code> option initiates a volume based SnapRestore preview mechanism for the given host filespec. With the <code>execute</code> option, SnapDrive for UNIX proceeds with volume based SnapRestore for the specified filespec. |

Rules for keywords

SnapDrive for UNIX uses keywords to specify the target of the SnapDrive for UNIX operations.

SnapDrive for UNIX uses keywords to specify sequences of strings corresponding to the host and storage

system objects with which you are working. The following rules apply to SnapDrive for UNIX keywords:

- Precede each keyword with a hyphen (-).
- Do not concatenate keywords.
- Enter the entire keyword and hyphen, not an abbreviation.


Command-line keywords

The SnapDrive for UNIX uses keywords to specify the name of host disk group, the destination group, volume, or FlexClone, the NFS file, the storage system, and so on.


Here are the keywords you can use with the SnapDrive for UNIX commands. You use them to specify the targets of the SnapDrive for UNIX operations. These keywords can take one or more arguments.



Some LVMs refer to disk groups and some refer to volume groups. In SnapDrive for UNIX, these terms are treated as synonyms. Moreover, some LVMs refer to logical volumes and some refer to volumes. SnapDrive for UNIX treats the term host volume (which was created to avoid confusing host logical volumes with storage system volumes) and the term logical volume as synonymous.

| Keyword | Argument used with this keyword |
|--|--|
| -dg (synonymous with -vg) | The name of the host disk group. You can enter the name of either a disk group or a volume group with this option. |
| -destdg -desthv -destlv -destvg | The destination group or volume. |
| -destfv | <div>The name of the FlexClone volume specified on the command line for volume clones created by SnapDrive for UNIX during the NFS Snapshot connect operation.</div> <div> This argument supports NFS volumes only and not NFS directories.</div> |
| -file | The name of a NFS file. |
| -filer | The name of a storage system. |
| -filervol | The name of the storage system and a volume on it. |

| Keyword | Argument used with this keyword |
|-------------------|--|
| -fs | The name of a file system on the host. The name used is the directory where the file system is currently mounted or is to be mounted (the mountpoint). |
| -hostvol or -lvol | The host volume name, including the disk group that contains it. For example, you might enter <code>large_vg/accounting_lvol</code> . |
| -igroup | <p>The name of an initiator group (igroup).</p> <p>NetApp strongly recommends that you use the default igroup that SnapDrive for UNIX creates instead of specifying an igroup on the target storage system. The default igroup is <code>hostname_protocol_SdIlg</code>.</p> <ul style="list-style-type: none"> • <i>hostname</i> is the local (non-domain qualified) name of the current host. • <i>protocol</i> is either FCP or iSCSI, depending on the protocol the host is using. <p>If the igroup <code>hostname_protocol_SdIlg</code> does not exist, SnapDrive for UNIX creates it and places all the initiators for the host in it.</p> <p>If it exists and has the correct initiators, SnapDrive for UNIX uses the existing igroup.</p> <p>If the igroup exists, but does not contain the initiators for this host, SnapDrive for UNIX creates a new igroup with a different name and uses that igroup in the current operation. To avoid using the same name, SnapDrive for UNIX includes a unique number when it creates the new name. In this case, the name format is <code>hostname-number_protocol_SdIlg</code>.</p> <p>If you supply your own igroup name, SnapDrive for UNIX does not validate the contents of the igroup. This is because it cannot always determine which igroups corresponding to the host are present on the storage system.</p> <p>All commands that take this option for specifying initiator groups cannot accept it with shared disk groups and file systems. This option is allowed only for local resources.</p> <p>The SnapDrive for UNIX command fails if any foreign igroups are involved in the command line. Ensure that all the igroups specified in the command line contain initiators from the local host.</p> |

| Keyword | Argument used with this keyword |
|--------------------|--|
| -lun | <p>The name of a LUN on a storage system. For the first LUN name you supply with this keyword, you must supply the full path name (storage system name, volume name, and LUN name). For additional LUN names, you can specify either only the names within their volume (if the volume stays unchanged) or a path to indicate a new storage system name or a new volume name (if you just want to switch volumes).</p> <div>  <p>In a <code>snapdrive snap connect</code> command, the <code>lun_name</code> should be in the <code>lun_name</code> or <code>tree_name/lun_name</code> format.</p> </div> |
| -lvol or - hostvol | The logical volume name, including the volume group that contains it. For example, you might enter <code>large_vg/accounting_lvol</code> as the logical volume name. |
| -snapname | The name of a Snapshot copy. |
| -vg or -dg | The name of the volume group. You can enter the name of either a disk group or a volume group with this option. |

Command-line arguments

The SnapDrive for UNIX takes arguments in a specific format.

The following table describes the arguments you can specify with the keywords.

Use the format `snapdrive type_name operation_name [<keyword/option> <arguments>]`; for example, if you wanted to create a Snapshot copy called `snap_hr` from the host file system `/mnt/dir`, you would enter the following command line:


```
snapdrive snap create -fs /mnt/dir -snapname snap_hr.
```

| Argument | Description |
|-----------|---|
| dest_spec | The name by which the target entity will be accessible after its disk groups or LUNs are connected. |
| dgname | The name of a disk group or volume group. |

| Argument | Description |
|------------|--|
| d_lun_name | Allows you to specify a destination name that SnapDrive for UNIX uses to make the LUN available in the newly-connected copy of the Snapshot copy. |
| filename | The name of a storage system. |
| filer_path | <p>A path name to a storage system object. This name can contain the storage system name and volume, but it does not have to if SnapDrive for UNIX can use default values for the missing components based on values supplied in the previous arguments. The following are examples of path names:</p> <ul style="list-style-type: none"> • test_filer:/vol/vol3/qtrees_2 • /vol/vol3/qtrees_2 • qtrees_2 |

| Argument | Description |
|-----------|---|
| file_spec | <p>The name of a storage entity, such as a host volume, LUN, disk or volume group, file system, or NFS directory tree.</p> <p>In general, you use the file_spec argument as one of the following:</p> <ul style="list-style-type: none"> • An object you want SnapDrive for UNIX to make a Snapshot copy of or to restore from a Snapshot copy • An object that you want to either create or use when provisioning storage <p>The objects do not have to be all of the same type. If you supply multiple host volumes, they must all belong to the same volume manager.</p> <p>If you supply values for this argument that resolve to redundant disk groups or host volumes, the command fails.</p> <p>Example of incorrect usage: This example assumes dg1 has host volumes hv1 and hv2, with file systems fs1 and fs2. As a result, the following arguments would fail because they involve redundant disk groups or host volumes.</p> <pre>-dg dg1 -hostvol dg1/hv1</pre> <pre>-dg dg1</pre> <pre>-fs/fs1</pre> <pre>-hostvol dg1/hv1 -fs/fs1</pre> <p>Example of correct usage: This example shows the correct usage for this argument.</p> <pre>-hostvoldg1/hv1dg1/hv2</pre> <pre>-fs/fs1/fs2</pre> <pre>-hostvoldg1/hv1 -fs/fs2</pre> |

| Argument | Description |
|-----------|--|
| fspec_set | <p>Used with the snap connect command to identify:</p> <ul style="list-style-type: none"> • A host LVM entity • A file system contained on a LUN <p>The argument also lets you specify a set of destination names that SnapDrive for UNIX uses when it makes the entity available in the newly connected copy of the Snapshot copy.</p> <p>The format for fspec_set is: { -vg -dg -fs -lvol -hostvol } <i>src_fspect</i> [<i>dest_fspect</i>] [{ -destdg -destvg } <i>dg_name</i>] [{ - destlv -desthv } <i>lv_name</i>]</p> |

| Argument | Description |
|-----------------|---|
| host_lvm_fspect | <p>Lets you specify whether you want to create a file system, logical volume, or disk group when you are executing the storage create command. This argument might have any of the three formats as described in the following. The format you use depends on the entity you want to create.</p> <div>  <p>The <code>-dg</code> and <code>-vg</code> options are synonyms that reflect the fact that some operating systems refer to disk groups and others refer to volume groups. In addition, <code>-lvol</code> and <code>-hostvol</code> are also synonyms. This guide uses <code>-dg</code> to refer to both disk groups and volume groups and <code>-hostvol</code> to refer to both logical volumes and host volumes.</p> </div> <p>To create a file system, use this format: <code>-fs file_spec [-fstype type] [-fsopts options] [-hostvol file_spec] [-dg dg_name]</code> To create a logical or host volume, use this format: <code>[-hostvol file_spec] [-dg dg_name] -hostvol</code> To create a disk or volume group, use this format: <code>file_spec [-dg dg_name] -dg dg_name</code></p> <p>You must name the top-level entity that you are creating. You do not need to supply names for any underlying entities. If you do not supply names for the underlying entities, SnapDrive for UNIX creates them with internally generated names.</p> <p>If you specify that SnapDrive for UNIX create a file system, you must specify a type that SnapDrive for UNIX supports with the host LVM. These types include JFS2 or VxFS.</p> <p>The option <code>-fsopts</code> is used to specify options to be passed to the host operation that creates the new file system; for example, <code>mkfs</code>.</p> |
| ig_name | The name of an initiator group. |

| Argument | Description |
|-----------------|---|
| long_filer_path | <p>A path name that includes the storage system name, volume name, and possibly other directory and file elements within that volume. The following are examples of long path names:</p> <pre>test_filer:/vol/vol3/qtrees_2</pre> <pre>10.10.10.1:/vol/vol4/lun_21</pre> |
| long_lun_name | <p>A name that includes the storage system name, volume, and LUN name. The following is an example of a long LUN name:</p> <pre>test_filer:/vol/vol1/lunA</pre> |
| long_snap_name | <p>A name that includes the storage system name, volume, and Snapshot copy name. The following is an example of a long Snapshot copy name:</p> <pre>test_filer:/vol/account_vol:snap_20040202</pre> <p>With the <code>snapdrive snap show</code> and <code>snapdrive snap delete</code> commands, you can use the asterisk (*) character as a wildcard to match any part of a Snapshot copy name. If you use a wildcard character, you must place it at the end of the Snapshot copy name. SnapDrive for UNIX displays an error message if you use a wildcard at any other point in a name.</p> <p>Example: This example uses wildcards with both the <code>snap show</code> command and the <code>snap delete</code> command:</p> <pre>myfiler:/vol/vol2:mynaps*</pre> <pre>myfiler:/vol/vol2:/yoursnap* snap show</pre> <pre>myfiler:/vol/vol1/qtrees1:qtrees_snap*</pre> <pre>snap delete</pre> <pre>10.10.10.10:/vol/vol2:mynaps*</pre> <pre>10.10.10.11:/vol/vol3:yoursnap* hersnap</pre> <p>Limitation for wildcards: You cannot enter a wildcard in the middle of a Snapshot copy name. For example, the following command line produces an error message because the wildcard is in the middle of the Snapshot copy name:</p> <pre>banana:/vol/vol1:my*snap</pre> |

| Argument | Description |
|---------------|---|
| lun_name | The name of a LUN. This name does not include the storage system and volume where the LUN is located. The following is an example of a LUN name: lunA |
| path | Any path name. |
| prefix_string | prefix used in the volume clone's name generation |
| s_lun_name | Indicates a LUN entity that is captured in the Snapshot copy specified by <i>long_snap_name</i> . |

Related information

[Storage provisioning command lines](#)

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.