



Configuring SnapDrive for UNIX

Snapdrive for Unix

NetApp
June 20, 2025

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

Table of Contents

Configuring SnapDrive for UNIX	1
What the snapdrive.conf file is	1
Configuration options and their default values	1
What the SnapDrive configuration wizard does	28
SnapDrive configuration wizard	28
Some configuration commands	29
Using the SnapDrive configuration wizard	30
Steps to configure in NFS environment	30
Steps to configure in SAN environment	31
Steps to configure in Mixed SAN and NFS environment	32
Setting values in the snapdrive.conf file	33
Checking the version of SnapDrive for UNIX	34
Example	34

Configuring SnapDrive for UNIX

You can edit the `snapdrive.conf` file, which contains all the configuration variables and options used in SnapDrive for UNIX, to enable or disable options according to your requirements. You can also add variables to create host-specific usage.

What the `snapdrive.conf` file is

The `snapdrive.conf` file contains a name-value pair for each configurable variable that SnapDrive for UNIX use to function. SnapDrive for UNIX automatically checks the information in this file each time it starts. You can use a text editor to modify this file.

The `snapdrive.conf` file is located in the SnapDrive installation directory. The `snapdrive config show` command displays the current and active contents of the `snapdrive.conf` file.

Configuration options and their default values

You can determine the current configurable variables and their settings by running the `snapdrive config show` command.

The supported configurable items and their default settings can vary across host operating systems and the different versions of SnapDrive for UNIX. For example, on AIX the default path is `/var/log/...`

The following table describes the parameters in the `snapdrive.conf` file:

Variable	Description
<code>lunpath-monitor-frequency</code>	Enables you to specify how often SnapDrive for UNIX automatically fixes LUN paths. The default value is 24 hours.
<code>blacklist-interfaces</code>	Enables you to specify, when there are multiple Ethernet interfaces, the interfaces that you do not want to use, to reduce operation time. If the configuration has multiple Ethernet interfaces, SnapDrive for UNIX at times searches through the list of interfaces to determine if the interface can ping. If the interface fails to ping, it tries for five times before checking the next interface. Thus, the operation takes additional time to execute. If you want SnapDrive to ignore some of the interfaces, you can specify those interfaces in the <code>blacklist-interfaces</code> parameter. This reduces the operation time.

Variable	Description
all-access-if-rbac-unspecified=on	<p>Specifies the access-control permissions for each host on which SnapDrive for UNIX runs by entering the permission string in an access-control file. The string that you specify controls which SnapDrive for UNIX Snapshot copy and other storage operations a host might perform on a storage system. (These access permissions do not affect the show or list operations.)</p> <p>Set this value to either <code>on</code> or <code>off</code> where:</p> <ul style="list-style-type: none"> • <code>on</code> specifies that SnapDrive for UNIX enables all access permissions if no access-control permissions file exists on the storage system. The default value is <code>on</code>. • <code>off</code> specifies that the storage system allows the host only the permissions that are mentioned in the access-control permissions file. <p>If you provide an access-control file, this option has no effect.</p>
allow-partial-clone-connect=on	<p>SnapDrive for UNIX enables you to connect to a subset of file systems or only to the host volume of the cloned disk group.</p> <p>Set this value to <code>on</code> or <code>off</code>:</p> <ul style="list-style-type: none"> • <code>on</code> specifies that SnapDrive for UNIX enables you to connect to a subset of file systems or only to the host volume of the cloned disk group. • <code>off</code> determines that SnapDrive for UNIX cannot connect to a subset of file systems or only to the host volume of the cloned disk group.
audit-log-file="/var/log/sd-audit.log"	<p>Specifies the location where SnapDrive for UNIX writes the audit log file.</p> <p>The default value depends on your host operating system. The path shown in the example is the default path for a AIX host.</p>

Variable	Description
audit-log-max-size=20480	<p>Specifies the maximum size, in bytes, of the audit log file. When the file reaches this size, SnapDrive for UNIX renames it and starts a new audit log. The default value is 20480 bytes. Because SnapDrive for UNIX never starts a new log file in the middle of an operation, the correct size of the file could vary slightly from the value specified here.</p> <p> You should use the default value. If you decide to change the default value, remember that too many log files can take up space on your disk and might eventually affect performance.</p>
audit-log-save=2	<p>Determines how many old audit log files SnapDrive for UNIX should save. After this limit is reached, SnapDrive for UNIX discards the oldest file and creates a new one.</p> <p>SnapDrive for UNIX rotates this file based on the value you specify in the audit-log-save variable. The default value is 2.</p> <p> You should use the default value. If you decide to change the default value, remember that too many log files can take up space on your disk and might eventually affect performance.</p>
autosupport-enabled	<p>Determines that the option autosupport-enabled is on by default.</p> <p>This option is enabled by default to store the AutoSupport information in the Event Management System (EMS) log of the storage system.</p> <p> SnapDrive 4.2 for UNIX and later versions do not have the option autosupport-filer.</p>

Variable	Description
available-lun-reserve=8	<p>Specifies the number of LUNs that the host must be prepared to create when the current SnapDrive for UNIX operation completes. If few operating system resources are available to create the number of LUNs specified, SnapDrive for UNIX requests additional resources, based on the value supplied in the <i>enable-implicit-host-preparation</i> variable.</p> <p>The default value is 8.</p> <p> This variable applies only to systems that require host preparation before you can create LUNs. Hosts require this preparation.</p> <p>This variable is used on configurations that include LUNs.</p>
bypass-snapdrive-clone-generated-check	<p>Specifies that the deletion of the SnapDrive generated or Non-snapdrive generated FlexClone.</p> <p>Set this value to either <code>on</code> or <code>off</code> where:</p> <ul style="list-style-type: none"> • <code>on</code> - Specifies that SnapDrive for UNIX allows to delete the FlexClone volume of the snapdrive-generated and non-snapdrive generated FlexClone. • <code>off</code> - Specifies that SnapDrive for UNIX allows to delete only the FlexClone volume of the snapdrive-generated. The default value is <code>off</code>.

Variable	Description
check-export-permission-nfs-clone	<p>Determines that setting the NFS export permission allows/disables to create cloning in the secondary host (host which does not have export permissions on the parent volume) or storage system.</p> <ul style="list-style-type: none"> • on - SnapDrive for UNIX checks for appropriate export permission on the volume for the secondary host. The default value is on. • off - SnapDrive for UNIX does not check appropriate export permission on the volume for the secondary host. <p>SnapDrive for UNIX does not allow cloning if there is no export permission for a volume in an NFS entity. To overcome this situation, disable this variable in the <code>snapdrive.conf</code> file. As a result of the cloning operation, SnapDrive provides appropriate access permissions on the cloned volume.</p> <p>Setting the value to off enables secondary protection to work in clustered Data ONTAP.</p>

Variable	Description
cluster-operation-timeout-secs=600	<p>Specifies the host cluster operation timeout, in seconds. You should set this value when working with remote nodes and HA pair operations to determine when the SnapDrive for UNIX operation should time out. The default value is 600 seconds.</p> <p>Other than the non-master node, the host cluster master node can also be the remote node, if the SnapDrive for UNIX operation is initiated from a non-master node.</p> <p>If SnapDrive for UNIX operations on any node in the host cluster exceed the value you set, or the default of 600 seconds (if you set no value), the operation times out with the following message:</p> <div data-bbox="829 713 1496 1262" style="border: 1px solid black; padding: 10px;"> <p>Remote Execution of command on slave node sfrac-57 timed out. Possible reason could be that timeout is too less for that system. You can increase the cluster connect timeout in snapdrive.conf file. Please do the necessary cleanup manually. Also, please check the operation can be restricted to lesser jobs to be done so that time required is reduced.</p> </div>
contact-http-port=80	Specifies the HTTP port to use for communicating with a storage system. The default value is 80.
contact-ssl-port=443	Specifies the SSL port to use for communicating with a storage system. The default value is 443.
contact-http-port-sdu-daemon=4094	Specifies the HTTP port to use for communicating with the SnapDrive for UNIX daemon. The default value is 4094.
contact-http-dfm-port=8088	Specifies the HTTP port to use for communicating with an Operations Manager server. The default value is 8088.

Variable	Description
contact-ssl-dfm-port=8488	Specifies the SSL port to use for communicating with an Operations Manager server. The default value is 8488.
datamotion-cutover-wait=120	Specifies the number of seconds SnapDrive for UNIX waits for the DataMotion for vFiler (cutover phase) operations to complete and then retries the SnapDrive for UNIX commands. The default value is 120 seconds.
dfm-api-timeout=180	Specifies the number of seconds SnapDrive for UNIX waits for the DFM API to return. The default value is 180 seconds.
dfm-rbac-retries=12	Specifies the number of times SnapDrive for UNIX checks access retries for an Operations Manager refresh. The default value is 12.
dfm-rbac-retry-sleep-secs=15	Specifies the number of seconds SnapDrive for UNIX waits before retrying an access check for an Operations Manager refresh. The default value is 15.
default-noprompt=off	<p>Specify if you want the <code>-noprompt</code> option to be available. The default value is <code>off</code> (not available).</p> <p>If you change this option to <code>on</code> SnapDrive for UNIX does not prompt you to confirm an action requested by <code>-force</code>.</p>

Variable	Description
device-retries=3	<p>Specifies the number of inquiries that the SnapDrive for UNIX can make about the device where the LUN resides. The default value is 3.</p> <p>In normal circumstances, the default value should be adequate. In other circumstances, LUN queries for a snap create operation could fail because the storage system is exceptionally busy.</p> <p>If the LUN queries keep failing even though the LUNs are online and correctly configured, you might want to increase the number of retries.</p> <p>This variable is used on configurations that include LUNs.</p> <div data-bbox="856 813 905 865" style="border: 1px solid #ccc; border-radius: 50%; padding: 5px; text-align: center; width: 20px; height: 20px;">  </div> <p>You should configure the same value for the <code>device-retries</code> variable across all the nodes in the host cluster. Otherwise, the device discovery involving multiple host cluster nodes can fail on some nodes and succeed on others.</p>
device-retry-sleep-secs=1	<p>Specifies the number of seconds SnapDrive for UNIX waits between inquiries about the device where the LUN resides. The default value is 1 second.</p> <p>In normal circumstances, the default value should be adequate. In other circumstances, LUN queries for a snap create operation could fail because the storage system is exceptionally busy.</p> <p>If the LUN queries keep failing even though the LUNs are online and correctly configured, you might want to increase the number of seconds between retries.</p> <p>This variable is used on configurations that include LUNs.</p> <div data-bbox="856 1657 905 1710" style="border: 1px solid #ccc; border-radius: 50%; padding: 5px; text-align: center; width: 20px; height: 20px;">  </div> <p>You should configure the same value for the <code>device-retry-sleep-secs</code> option across all the nodes in the host cluster. Otherwise, the device discovery involving multiple host cluster nodes can fail on some nodes and succeed on others.</p>

Variable	Description
default-transport=iscsi	<p>Specifies the protocol that SnapDrive for UNIX uses as the transport type when creating storage, if a decision is required. The acceptable values are <code>iscsi</code> or <code>FCP</code>.</p> <p> If a host is configured for only one type of transport and that type is supported by SnapDrive for UNIX, SnapDrive for UNIX uses that transport type, irrespective of the type specified in the <code>snapdrive.conf</code> file.</p> <p>On AIX hosts, ensure the <code>multipathing-type</code> option is set correctly. If you specify <code>FCP</code>, you must set <code>multipathing-type</code> to one of the following values:</p> <ul style="list-style-type: none"> • <code>NativeMPIO</code> • <code>DMP</code>
enable-alua=on	<p>Determines that the ALUA is supported for multipathing on the igrup. The storage systems must be HA pair and the HA pair failover state in <i>single-image</i> mode.</p> <ul style="list-style-type: none"> • The default value is <code>on</code> to support ALUA for igrup • You can disable the ALUA support by setting the option <code>off</code>
enable-fcp-cache=on	<p>Specifies whether to enable or disable the cache. SnapDrive maintains a cache of available active ports and the port names (WWPNs) information to send the response faster.</p> <p>This variable is useful in few scenario where there is no FC cables connected to the port or wrap plug is used in the port, SnapDrive for UNIX may experience long delays to fetch the information about FC interface and their corresponding WWPNs. The caching helps to resolve/improve the performance of SnapDrive operations in such environments.</p> <p>The default value is <code>on</code>.</p>

Variable	Description
enable-implicit-host-preparation=on	<p>Determines whether SnapDrive for UNIX implicitly requests host preparation for LUNs or notifies you that it is required and exits.</p> <ul style="list-style-type: none"> • on - SnapDrive for UNIX implicitly requests the host to create more resources, if there is inadequate amount of resources available to create the required number of LUNs. The number of LUNs created is specified in the <i>available-lun-reserve</i> variable. The default value is on. • off - SnapDrive for UNIX informs you if additional host preparation is necessary for LUN creation and SnapDrive exits the operation. You can then perform the operations necessary to free up resources needed for LUN creation. For example, you can execute the <code>snapdrive config prepare luns</code> command. After the preparation is complete, you can reenter the current SnapDrive for UNIX command. <p> This variable applies only to systems where host preparation is needed before you can create LUNs for the hosts that require the preparation. This variable is used only on configurations that include LUNs.</p>

Variable	Description
enable-migrate-nfs-version	<p>Allows to clone/restore by using the higher version of NFS.</p> <p>In a pure NFSv4 environment, when snap management operations such as clone and restore are attempted with a Snapshot copy created on NFSv3, snap management operation fails.</p> <p>The default value is <code>off</code>. During this migration, only the protocol version is considered and other options such as <code>rw</code> and <code>largefiles</code> are not taken into account by SnapDrive for UNIX.</p> <p>Therefore, only the NFS version for the corresponding NFS filespec is added in the <code>/etc/fstab</code> file. Ensure that the appropriate NFS version is used to mount the file specification by using <code>-o vers=3</code> for NFSv3 and <code>-o vers=4</code> for NFSv4. If you want to migrate the NFS file specification with all the mount options, it is recommended to use <code>-mntopts</code> in the snap management operations. It is mandatory to use <code>nfs</code> in the attribute value of the Access Protocol in the export policy rules of the parent volume during migration in clustered Data ONTAP .</p> <p> Ensure that you use only the <code>nfsvers</code> or <code>vers</code> commands as the mount options, to check the NFS version.</p>
enable-mountguard-support	<p>Enables SnapDrive for UNIX support for the Mount Guard feature of AIX, which prevents simultaneous or concurrent mounts. If a file system is mounted on one node and the variable is enabled, AIX prevents the same file system from being mounted on another node. By default the <code>enable-mountguard-support</code> variable is set to <code>off</code>.</p>
enable-ping-to-check-filer-reachability	<p>If the ICMP protocol access is disabled or ICMP packets are dropped between the host and storage system network where SnapDrive for UNIX is deployed, this variable must be set to <code>off</code>, so that SnapDrive for UNIX does not ping to check if the storage system is reachable or not. If this variable is set to <code>on</code> only SnapDrive snap connect operation does not work due to the ping failure. By default, this variable is set to <code>on</code>.</p>

Variable	Description
enable-split-clone=off	<p>Enables splitting the cloned volumes or LUNs during Snapshot connect and Snapshot disconnect operations, if this variable is set to <code>on</code> or <code>sync</code>. You can set the following values for this variable:</p> <ul style="list-style-type: none"> • <code>on</code> - enables an asynchronous split of cloned volumes or LUNs. • <code>sync</code> - enables a synchronous split of cloned volumes or LUNs. • <code>off</code> - disables the split of cloned volumes or LUNs. The default value is <code>off</code>. <p>If you set this value to <code>on</code> or <code>sync</code> during the Snapshot connect operation and <code>off</code> during the Snapshot disconnect operation, SnapDrive for UNIX does not delete the original volume or LUN that is present in the Snapshot copy.</p> <p>You can also split the cloned volumes or LUNs by using the <code>-split</code> option.</p>
enforce-strong-ciphers=off	<p>Set this variable to <code>on</code> for the SnapDrive daemon to enforce TLSv1 to communicate with the client.</p> <p>It enhances the security of communication between the client and the SnapDrive daemon using better encryption.</p> <p>By default, this option is set to <code>off</code>.</p>
filer-restore-retries=140	<p>Specifies the number of times SnapDrive for UNIX attempts to restore a Snapshot copy on a storage system if a failure occurs during the restore. The default value is 140.</p> <p>In normal circumstances, the default value should be adequate. Under other circumstances, this operation could fail because the storage system is exceptionally busy. If it keeps failing even though the LUNs are online and correctly configured, you might want to increase the number of retries.</p>

Variable	Description
filer-restore-retry-sleep-secs=15	<p>Specifies the number of seconds SnapDrive for UNIX waits between attempts to restore a Snapshot copy. The default value is 15 seconds.</p> <p>In normal circumstances, the default value should be adequate. Under other circumstances, this operation could fail because the storage system is exceptionally busy. If it keeps failing even though the LUNs are online and correctly configured, you might want to increase the number of seconds between retries.</p>
filesystem-freeze-timeout-secs=300	<p>Specifies the number of seconds that SnapDrive for UNIX waits between attempts to access the file system. The default value is 300 seconds.</p> <p>This variable is used only on configurations that include LUNs.</p>
flexclone-writereserve-enabled=on	<p>It can take any one of the following values:</p> <ul style="list-style-type: none"> • on • off <p>Determines the space reservation of the FlexClone volume created. Acceptable values are <code>on</code> and <code>off</code>, based on the following rules.</p> <ul style="list-style-type: none"> • Reservation: on • Optimal: file • Unrestricted: volume • Reservation: off • Optimal: file • Unrestricted: none

Variable	Description
fstype=jfs2	<p>Specifies the type of file system that you want to use for SnapDrive for UNIX operations. The file system must be a type that SnapDrive for UNIX supports for your operating system.</p> <p>AIX: <code>jfs</code>, <code>jfs3</code> or <code>vxfs</code></p> <p>The default value is <code>jfs2</code>.</p> <p> The JFS file system type is supported only for Snapshot operations and not for storage operations.</p> <p>You can also specify the type of file system that you want to use by using the <code>-fstype</code> option through CLI.</p>
lun-onlining-in-progress-sleep-secs=3	<p>Specifies the number of seconds between retries during attempts to bring back online a LUN after a volume-based SnapRestore operation. The default value is 3.</p>
lun-on-onlining-in-progress-retries=40	<p>Specifies the number of retries during attempts to bring back online a LUN after a volume-based SnapRestore operation. The default value is 40.</p>
mgmt-retry-sleep-secs=2	<p>Specifies the number of seconds SnapDrive for UNIX waits before retrying an operation on the Manage ONTAP control channel. The default value is 2 seconds.</p>
mgmt-retry-sleep-long-secs=90	<p>Specifies the number of seconds SnapDrive for UNIX waits before retrying an operation on the Manage ONTAP control channel after a failover error message occurs. The default value is 90 seconds.</p>

Variable	Description
multipathing-type=NativeMPIO	<p>Specifies the multipathing software to use. The default value depends on the host operating system. This variable applies only if one of the following statements is true:</p> <ul style="list-style-type: none"> • More than one multipathing solution is available. • The configurations include LUNs. <p>The acceptable values are <code>none</code> or <code>nativempio</code>.</p> <p>You can set the following values for this variable:</p> <p>AIX: The value you set for AIX depends on which protocol you are using.</p> <ul style="list-style-type: none"> • If you are using FCP, set this to any one of the following values: <ul style="list-style-type: none"> ◦ NativeMPIO The default value is <code>none</code>. • In addition, set the <code>default-transport</code> option to FCP. • If you are using iSCSI, set this value to <code>none</code>. In addition, set the <code>default-transport</code> option to <code>iscsi</code>.
override-vbsr-snapmirror-check	<p>You can set the value of the <code>override-vbsr-snapmirror-check</code> variable to <code>on</code> to override the SnapMirror relationship, when a Snapshot copy to be restored is older than the SnapMirror baseline Snapshot copy, during volume-based SnapRestore (VBSR). You can use this variable only if the OnCommand Data Fabric Manager (DFM) is not configured.</p> <p>By default, the value is set to <code>off</code>. This variable is not applicable for clustered Data ONTAP version 8.2 or later.</p>
PATH="/sbin:/usr/sbin:/bin:/usr/lib/vxvm/bin:/usr/bin:/opt/NTAPontap/SANToolkit/bin:/opt/NTAPsanlun/bin:/opt/VRTS/bin:/etc/vx/bin"	<p>Specifies the search path the system uses to look for tools.</p> <p>You should verify that this is correct for your system. If it is incorrect, change it to the correct path.</p> <p>The default value might vary depending on your operating system. This path is the default for AIX host does not use this variable because they process the commands differently.</p>

Variable	Description
/opt/NetApp/snapdrive/.pwfile	<p>Specifies the location of the password file for the user login for the storage systems.</p> <p>The default value might vary depending on your operating system.</p> <p>The default path for Linux is <code>/opt/NetApp/snapdrive/.pwfile/opt/ontap/snapdrive/.pwfile</code></p>
ping-interfaces-with-same-octet	<p>Avoids unnecessary pings through all the available interfaces in the host that might have different subnet IPs configured. If this variable is set to <code>on</code>, SnapDrive for UNIX considers only the same subnet IPs of the storage system and pings the storage system to verify address response. If this variable is set to <code>off</code>, SnapDrive takes all the available IPs in the host system and pings the storage system to verify address resolution through each subnet, which may be locally detected as a ping attack.</p>
prefix-filer-lun	<p>Specifies the prefix that SnapDrive for UNIX applies to all LUN names it generates internally. The default value for this prefix is an empty string.</p> <p>This variable allows the names of all LUNs created from the current host, but not explicitly named on a SnapDrive for UNIX command line, to share an initial string.</p> <div data-bbox="850 1227 910 1290" style="border: 1px solid #ccc; border-radius: 50%; width: 37px; height: 37px; display: flex; align-items: center; justify-content: center; margin-right: 10px;"></div> <p>This variable is used only on configurations that include LUNs.</p>
prefix-clone-name	<p>The string given is appended with the original storage system volume name, to create a name for the FlexClone volume.</p>

Variable	Description
prepare-lun-count=16	<p>Specifies how many LUNs SnapDrive for UNIX should prepare to create. SnapDrive for UNIX checks this value when it receives a request to prepare the host to create additional LUNs.</p> <p>The default value is 16, which means the system is able to create 16 additional LUNs after the preparation is complete.</p> <p> This variable applies only to systems where host preparation is needed before you can create LUNs. This variable is used only on configurations that include LUNs. hosts require that preparation.</p>
rbac-method=dfm	<p>Specifies the access control methods. The possible values are <code>native</code> and <code>dfm</code>.</p> <p>If the variable is set to <code>native</code>, the access-control file that is stored in <code>/vol/vol0/sdprbac/sdhost-name.prbac</code> or <code>/vol/vol0/sdprbac/sdgeneric-name.prbac</code> is used for access checks.</p> <p>If the variable is set to <code>dfm</code>, Operations Manager is a prerequisite. In such a case, SnapDrive for UNIX issues access checks to Operations Manager.</p>
rbac-cache=off	<p>Specifies whether to enable or disable cache. SnapDrive for UNIX maintains a cache of access check queries and the corresponding results. SnapDrive for UNIX uses this cache only when all the configured Operations Manager servers are down.</p> <p>You can set the value of the variable to either <code>on</code> to enable cache, or to <code>off</code> to disable it. The default value is <code>off</code>, which configures SnapDrive for UNIX to use Operations Manager and the set <code>rbac-method</code> configuration variable to <code>dfm</code>.</p>
rbac-cache-timeout	<p>Specifies the rbac cache timeout period and is applicable only when <code>rbac-cache</code> is enabled. The default value is 24 hrs. SnapDrive for UNIX uses this cache only when all the configured Operations Manager servers are down.</p>

Variable	Description
recovery-log-file=/var/log/sdrecovery.log	<p>Specifies where SnapDrive for UNIX writes the recovery log file.</p> <p>The default value depends on your host operating system. The path shown in this example is the default path for a AIX host.</p>
recovery-log-save=20	<p>Specifies how many old recovery log files SnapDrive for UNIX should save. After this limit is reached, SnapDrive for UNIX discards the oldest file when it creates a new one.</p> <p>SnapDrive for UNIX rotates this log file each time it starts a new operation. The default value is 20.</p> <p> You should use the default value. If you decide to change the default, remember that having too many large log files can take up space on your disk and might eventually affect performance.</p>
san-clone-method	<p>Specifies the type of clone that you can create.</p> <p>It can take the following values:</p> <ul style="list-style-type: none"> • lunclone <ul style="list-style-type: none"> Allows a connection by creating a clone of the LUN in the same storage system volume. The default value is lunclone. • optimal <ul style="list-style-type: none"> Allows a connection by creating a restricted FlexClone volume of the storage system volume. • unrestricted <ul style="list-style-type: none"> Allows a connection by creating an unrestricted FlexClone volume of the storage system volume.

Variable	Description
secure-communication-among-clusternodes=on	<p>Specifies a secure communication within the host cluster nodes for remote execution of SnapDrive for UNIX commands.</p> <p>You can direct SnapDrive for UNIX to use RSH or SSH by changing the value of this configuration variable. The RSH or SSH methodology adopted by SnapDrive for UNIX for remote execution is determined only by the value set in the installation directory of the <code>snapdrive.conf</code> file of the following two components:</p> <ul style="list-style-type: none"> • The host on which the SnapDrive for UNIX operation is executed, to get the host WWPN information and device path information of remote nodes. <p>For example, <code>snapdrive storage create</code> executed on master host cluster node uses the RSH or SSH configuration variable only in the local <code>snapdrive.conf</code> file to do either of the following:</p> <ul style="list-style-type: none"> ◦ Determine the remote communication channel. ◦ Execute the <code>devfsadm</code> command on remote nodes. <ul style="list-style-type: none"> • The non-master host cluster node, if the SnapDrive for UNIX command is to be executed remotely on the master host cluster node. <p>To send the SnapDrive for UNIX command to the master host cluster node, the RSH or SSH configuration variable in the local <code>snapdrive.conf</code> file is consulted to determine the RSH or SSH mechanism for remote command execution.</p> <p>The default value of <code>on</code> means that SSH is used for remote command execution. The value <code>off</code> means that RSH is used for execution.</p>

Variable	Description
snapcreate-cg-timeout=relaxed	<p>Specifies the interval that the <code>snapdrive snap create</code> command allows for a storage system to complete fencing. Values for this variable are as follows:</p> <ul style="list-style-type: none"> • <code>urgent</code> - specifies a short interval. • <code>medium</code> - specifies an interval between <code>urgent</code> and <code>relaxed</code>. • <code>relaxed</code> - specifies the longest interval. This value is the default. <p>If a storage system does not complete fencing within the time allowed, SnapDrive for UNIX creates a Snapshot copy using the methodology for Data ONTAP versions before 7.2.</p>
snapcreate-check-nonpersistent-nfs=on	<p>Enables and disables the Snapshot create operation to work with a non-persistent NFS file system. Values for this variable are as follows:</p> <ul style="list-style-type: none"> • <code>on</code> - SnapDrive for UNIX checks whether NFS entities specified in the <code>snapdrive snap create</code> command are present in the file system mount table. The Snapshot create operation fails if the NFS entities are not persistently mounted through the file system mount table. This is the default value. • <code>off</code> - SnapDrive for UNIX creates a Snapshot copy of NFS entities that do not have a mount entry in the file system mount table. <p>The Snapshot restore operation automatically restores and mounts the NFS file or directory tree that you specify.</p> <p>You can use the <code>-nopersist</code> option in the <code>snapdrive snap connect</code> command to prevent NFS file systems from adding mount entries in the file system mount table.</p>
snapcreate-consistency-retry-sleep=1	<p>Specifies the number of seconds between best-effort Snapshot copy consistency retries. The default value is 1 second.</p>

Variable	Description
<code>snapconnect-nfs-removedirectories=off</code>	<p>Determines whether SnapDrive for UNIX deletes or retains the unwanted NFS directories from the FlexClone volume during the Snapshot connect operation.</p> <ul style="list-style-type: none"> • <code>on</code> - Deletes the unwanted NFS directories (storage system directories not mentioned in the <code>snapdrive snap connect</code> command) from the FlexClone volume during the Snapshot connect operation. • <code>off</code> - Retains the unwanted NFS storage system directories during the Snapshot connect operation. The default value is <code>off</code>. <p>The FlexClone volume is destroyed if it is empty during the Snapshot disconnect operation.</p> <p>During the Snapshot disconnect operation, only the specified storage system directories are unmounted from the host. If nothing is mounted from the FlexClone volume on the host, the FlexClone volume is destroyed during the Snapshot disconnect operation.</p> <p>If you set this variable to <code>off</code> during the connect operation or <code>on</code> during the disconnect operation, the FlexClone volume is not to be destroyed, even if it has unwanted storage system directories and is not empty.</p>
<code>snapcreate-must-make-snapinfo-on-qtree=off</code>	<p>Set this variable to <code>on</code> to enable the Snapshot create operation to create Snapshot copy information about a qtree. The default value is <code>off</code> (disabled).</p> <p>SnapDrive for UNIX always attempts to write <code>snapinfo</code> at the root of a qtree if the LUNs are still snapped and are at the qtree. When you set this variable to <code>on</code>, SnapDrive for UNIX fails the Snapshot create operation if it cannot write this data. You should set this variable only to <code>on</code> if you are replicating Snapshot copies using qtree SnapMirror.</p> <div data-bbox="850 1706 910 1769"> </div> <p>Snapshot copies of qtrees work the same way Snapshot copies of volumes do.</p>

Variable	Description
snapcreate-consistency-retries=3	<p>Specifies the number of times SnapDrive for UNIX attempts a consistency check on a Snapshot copy after it receives a message that a consistency check failed.</p> <p>This variable is particularly useful on host platforms that do not include a freeze function. This variable is used only on configurations that include LUNs.</p> <p>The default value is 3.</p>
snapdelete-delete-rollback-withsnap=off	<p>Set this value to <code>on</code> to delete all rollback Snapshot copies related to a Snapshot copy. Set it to <code>off</code> to disable this feature. The default value is <code>off</code>.</p> <p>This variable takes effect only during a Snapshot delete operation and is used by the recovery log file if you encounter a problem with an operation.</p> <p>It is best to accept the default setting.</p>
snapmirror-dest-multiple-filervolumesenabled=off	<p>Set this variable to <code>on</code> to restore Snapshot copies that span multiple storage systems or volumes on (mirrored) destination storage systems. Set it to <code>off</code> to disable this feature. The default value is <code>off</code>.</p>
snaprestore-delete-rollback-afterrestore=off	<p>Set this variable to <code>on</code> to delete all rollback Snapshot copies after a successful Snapshot restore operation. Set it to <code>off</code> to disable this feature. The default value is <code>off</code> (enabled).</p> <p>This option is used by the recovery log file if you encounter a problem with an operation.</p> <p>It is best to accept the default value.</p>

Variable	Description
snaprestore-make-rollback=on	<p>Set this value to either <code>on</code> to create a rollback Snapshot copy or <code>off</code> to disable this feature. The default value is <code>on</code>.</p> <p>A rollback is a copy of the data that SnapDrive makes on the storage system before it begins a Snapshot restore operation. If a problem occurs during the Snapshot restore operation, you can use the rollback Snapshot copy to restore the data to the state it was in before the operation began.</p> <p>If you do not want the extra security of a rollback Snapshot copy at restore time, set this option to <code>off</code>. If you want the rollback, but not enough for your Snapshot restore operation to fail if you cannot make one, set the variable <code>snaprestore-must-makero</code>llback to <code>off</code>.</p> <p>This variable is used by the recovery log file, which you send to NetApp technical support if you encounter a problem.</p> <p>It is best to accept the default value.</p>
snaprestore-must-make-rollback=on	<p>Set this variable to <code>on</code> to cause a Snapshot restore operation to fail if the rollback creation fails. Set it to <code>off</code> to disable this feature. The default value is <code>on</code>.</p> <ul style="list-style-type: none"> • <code>on</code> - SnapDrive for UNIX attempts to make a rollback copy of the data on the storage system before it begins the Snapshot restore operation. If it cannot make a rollback copy of the data, SnapDrive for UNIX halts the Snapshot restore operation. • <code>off</code> - Use this value if you want the extra security of a rollback Snapshot copy at restore time, but not enough for the Snapshot restore operation to fail if you cannot make one. <p>This variable is used by the recovery log file if you encounter a problem with an operation.</p> <p>It is best to accept the default value.</p>

Variable	Description
snaprestore-snapmirror-check=on	<p>Set this variable to <code>on</code> to enable the <code>snap restore</code> command to check the SnapMirror destination volume. If it is set to <code>off</code>, the <code>snap restore</code> command is unable to check the destination volume. The default value is <code>on</code>.</p> <p>If the value of this configuration variable is <code>on</code> and the SnapMirror relationship state is <code>broken-off</code>, the restore can still proceed.</p>
space-reservations-enabled=on	<p>Enables space reservation when creating LUNs. By default, this variable is set to <code>on</code>; therefore, the LUNs created by SnapDrive for UNIX have space reservation.</p> <p>You can use this variable to disable the space reservation for LUNs created by the <code>snap connect</code> command and <code>snapdrive storage create</code> command. It is best to use the <code>-reserve</code> and <code>-noreserve</code> command-line options to enable or disable LUN space reservation in the <code>snapdrive storage create</code>, <code>snapdrive snap connect</code>, and <code>snapdrive snap restore</code> commands.</p> <p>SnapDrive for UNIX creates LUNs, resizes storage, makes Snapshot copies, and connects or restores the Snapshot copies based on the space reservation permission that is specified in this variable or by the <code>-reserve</code> or <code>-noreserve</code> command-line options. It does not consider the storage system-side thin provisioning options before performing the preceding tasks.</p>
trace-enabled=on	<p>Set this variable to <code>on</code> to enable the trace log file, or to <code>off</code> to disable it. The default value is <code>on</code>. Enabling this file does not affect performance.</p>

Variable	Description
trace-level=7	<p>Specifies the types of messages SnapDrive for UNIX writes to the trace log file. This variable accepts the following values:</p> <ul style="list-style-type: none"> • 1 - Record fatal errors • 2 - Record admin errors • 3 - Record command errors • 4 - Record warnings • 5 - Record information messages • 6 - Record in verbose mode • 7 - Full diagnostic output <p>The default value is 7.</p> <p> It is best not to change the default value. Setting the value to something other than 7 does not gather adequate information for a successful diagnosis.</p>
trace-log-file=/var/log/sd-trace.log	<p>Specifies where SnapDrive for UNIX writes the trace log file.</p> <p>The default value varies depending on your host operating system.</p> <p>The path shown in this example is the default path for an AIX host.</p>

Variable	Description
<pre>trace-log-max-size=0</pre>	<p>Specifies the maximum size of the log file in bytes. When the log file reaches this size, SnapDrive for UNIX renames it and starts a new log file.</p> <p> However, no new trace log file is created when the trace log file reaches the maximum size. For the daemon trace log file, new log file is created when the log file reaches the maximum size.</p> <p>The default value is 0. SnapDrive for UNIX never starts a new log file in the middle of an operation. The actual size of the file could vary slightly from the value specified here.</p> <p> It is best to use the default value. If you change the default, remember that too many large log files can take up space on your disk and might eventually affect performance.</p>
<pre>trace-log-save=100</pre>	<p>Specifies how many old trace log files SnapDrive for UNIX should save. After this limit is reached, SnapDrive for UNIX discards the oldest file when it creates a new one. This variable works with the <i>tracelog-max-size</i> variable. By default, <i>trace-logmax-size=0</i> saves one command in each file, and <i>trace-log-save=100</i> retains the last 100 log files.</p>
<pre>use-https-to-dfm=on</pre>	<p>Specifies whether you want SnapDrive for UNIX to use SSL encryption (HTTPS) to communicate with Operations Manager.</p> <p>The default value is on.</p>

Variable	Description
use-https-to-filer=on	<p>Specifies whether you want SnapDrive for UNIX to use SSL encryption (HTTPS) when it communicates with the storage system.</p> <p>The default value is <code>on</code>.</p> <p> If you are using a version of Data ONTAP earlier to 7.0, you might see slower performance with HTTPS enabled. Slow performance is not an issue if you are running Data ONTAP 7.0 or later.</p>
vmttype=lvm	<p>Specify the type of volume manager you want to use for SnapDrive for UNIX operations. The volume manager must be a type that SnapDrive for UNIX supports for your operating system. Following are the values that you can set for this variable, and the default value varies depending on the host operating systems:</p> <ul style="list-style-type: none"> • AIX: <code>vxvm</code> or <code>lvm</code> <p>The default value is <code>lvm</code></p> <p>You can also specify the type of volume manager that you want to use by using the <code>-vmttype</code> option.</p>
vol-restore	<p>Determines whether SnapDrive for UNIX should perform volume-based snap restore (<code>vbsr</code>) or single-file snap restore (<code>sfsr</code>).</p> <p>The following are the possible values.</p> <ul style="list-style-type: none"> • <code>preview</code> - Specifies that SnapDrive for UNIX initiates a volume-based SnapRestore preview mechanism for the given host file specification. • <code>execute</code> - Specifies that SnapDrive for UNIX proceeds with volume based SnapRestore for the specified filespec. • <code>off</code> - Disables the <code>vbsr</code> option and enables the <code>sfsr</code> option. The default value is <code>off</code>. <p> If the variable is set to <code>preview/execute</code>, then you cannot override this setting by using CLI to perform SFSR operations.</p>

Variable	Description
volmove-cutover-retry=3	<p>Specifies the number of times SnapDrive for UNIX retries the operation during the volume migration cut-over phase.</p> <p>The default value is 3.</p>
volmove-cutover-retry-sleep=3	<p>Specifies the number of seconds SnapDrive for UNIX waits between the volume-move-cutover-retry operation.</p> <p>The default value is 3.</p>
volume-clone-retry=3	<p>Specifies the number of times, SnapDrive for UNIX retries the operation during FlexClone creation.</p> <p>The default value is 3.</p>
volume-clone-retry-sleep=3	<p>Specifies the number of seconds, SnapDrive for UNIX waits between the retries during FlexClone creation.</p> <p>The default value is 3.</p>

What the SnapDrive configuration wizard does

The SnapDrive configuration wizard enables you to configure SnapDrive for UNIX, and the NFS or SAN settings, depending on the environment. Alternatively, you can also open the `snapdrive.conf` file and edit the configuration variables.

SnapDrive configuration wizard

Use the SnapDrive configuration wizard to update the configuration variables in the wizard. You can run the configuration wizard at any time to modify your configuration changes to SnapDrive.

You can run the configuration wizard from `/opt/NetApp/snapdrive/setup/config_wizard`.

Enter `exit` to close the configuration wizard, without saving the changes.

Alternatively, you can modify the configuration variables value directly in the `snapdrive.conf` file.

Related information

[Stack requirements](#)

Some configuration commands

There are few commands that help you in configuring SnapDrive for UNIX.

The following table summarizes additional commands that are helpful in configuring SnapDrive for UNIX:

Command or action	Description
Configuring and verifying your version of SnapDrive for UNIX software	
snapshot config show	Check the values in the <code>snapshot.conf</code> file.
snapshot version	Check the version of SnapDrive for UNIX.
snapshotd start	Start the SnapDrive for UNIX daemon.
snapshot config prepare luns -count count_value	Prepare the host for creating a specific number of LUNs as well as to determine how many LUNs you can create.
snapshot clone split	Estimate, start, stop, and query the status of split for a volume clone or a LUN clone.
Edit the variables in the <code>snapshot.conf</code> file.	Change the path name and options for the log files. SnapDrive for UNIX logs information to three files: <ul style="list-style-type: none">• An audit file• A recovery file• A trace file
Setting and displaying access control permissions between a host and a storage system.	

Command or action	Description
Edit the access control permissions file (sd <i>hostname</i> .prbac) on the storage system associated with that host.	<p>Specify the access control permissions a host has on a storage system. You can set the following access levels for a host on a storage system:</p> <ul style="list-style-type: none"> • NONE—No access. • SNAP CREATE—Create Snapshot copies. • SNAP USE—Delete and rename Snapshot copies. • SNAP ALL—Create, restore, delete, and rename Snapshot copies. • STORAGE CREATE DELETE—Create, resize, and delete storage. • STORAGE USE—Connect and disconnect storage. • STORAGE ALL—Create, delete, connect, and disconnect storage. • ALL ACCESS—All operations.
<code>snapdrive config access show <filer_name></code>	Display information about the access control permissions a host has to a storage system.
<code>snapdrive config delete <filename> [<filename> . . .]</code>	Remove the specified user name-password pair from SnapDrive for UNIX.

Using the SnapDrive configuration wizard

The configuration wizard allows you to configure in NFS, SAN or Mixed environment.

Steps to configure in NFS environment

The following are the steps to configure in NFS environment.

Steps

1. Select the **NFS** profile.
2. Enable the Protection Manager Integration.
 - Select Yes to enable the access permission checks by using the DataFabric Manager.
 - Enter the DataFabric Manager server name or IP address followed by user name and password.
 - Enter the http/https port to communicate with the DataFabric Manager. The default value is 8088.
 - Enter the SSL server port to access the DataFabric Manager. The default value is 8488.
 - Enable the HTTPs enabled to communicate with the DataFabric Manager.
 - Select No to enable the access permission checks by using the rbac.

3. Specify the role-based access control methods. The possible values are `native` and `dfm`.
 - Select `native` to check the access permission for the host using the control file stored in `/vol/vol0/sdprbac/sdhost-name.prbac` or `/vol/vol0/sdprbac/sdgenericname.prbac`.
 - Select `dfm` to check the access permission using the Operations Manager console.

 If you select `dfm` as `rbac-method` without configuring DataFabric Manager, a warning message specifying that the RBAC method is selected as `dfm` without enabling Protection Manager Integration is displayed.
4. Specify `https` or `http` to communicate with the storage system.
5. The final step is to save the configuration changes in the `snapdrive.conf` file, and restart the daemon.
 - If you select `Yes`, the SnapDrive daemon is restarted and the configuration changes are reflected.
 - If you select `No`, the variable values are changed in `snapdrive.conf` file, but the changes are not reflected.

Steps to configure in SAN environment

The following are the steps to configure in SAN environment.

Steps

1. Select the SAN profile.
2. Select the required transport protocol.
 - Select `fcp` to set the default-transport.
 - Select `iscsi` to set the default-transport.
3. Select the SAN Storage Stack (combination of MPIO Solution, volume manager, and file system). The options are `native`, `veritas`, and `none`.

SnapDrive does not support `veritas` for iSCSI transport protocol.
4. Enable the Protection Manager Integration.
 - Select `Yes` to enable the access permission checks by using the DataFabric Manager.
 - Enter the DataFabric Manager server name or IP address followed by user name and password.
 - Enter the `http/https` port to communicate with the DataFabric Manager. The default value is `8088`.
 - Enter the SSL server port to access the DataFabric Manager. The default value is `8488`.
 - Enable the `HTTPs` enabled to communicate with the DataFabric Manager
 - Select `No` to enable the access permission checks by using the `rbac`.
5. Specify the role-based access control methods. The possible values are `native` and `dfm`.
 - Select `native` to check the access permission for the host using the control file stored in `/vol/vol0/sdprbac/sdhost-name.prbac` or `/vol/vol0/sdprbac/sdgenericname.prbac`.
 - Select `dfm` to check the access permission using the Operations Manager.



If you select `dfm` as `rbac-method` without configuring DataFabric Manager, a warning message specifying that the RBAC method is selected as `dfm` without enabling Protection Manager Integration is displayed.

6. Specify `https` or `http` to communicate with the storage system.
7. The final step is to save the configuration changes in the `snapdrive.conf` file, and restart the daemon.
 - If you select `Yes`, the SnapDrive daemon is restarted and the configuration changes are reflected.
 - If you select `No`, the variable values are changed in `snapdrive.conf` file, but the changes are not reflected.

Steps to configure in Mixed SAN and NFS environment

The following are the steps to configure in Mixed SAN and NFS environment.

Steps

1. Select the Mixed profile.
2. Select the required transport protocol.
 - Select `fcp` to set the default-transport.
 - Select `iscsi` to set the default-transport.
3. Select the SAN Storage Stack (combination of MPIO Solution, volume manager, file system). The options are `native`, `veritas`, and `none`.

SnapDrive does not support `veritas` for iSCSI transport protocol.
4. Enable the Protection Manager Integration.
 - Select `Yes` to enable the access permission checks by using the DataFabric Manager
 - Enter the DataFabric Manager server name or IP address followed by user name and password.
 - Enter the `http/https` port to communicate with the DataFabric Manager. The default value is `8088`.
 - Enter the SSL server port to access the DataFabric Manager. The default value is `8488`.
 - Enable the `HTTPs` enabled to communicate with the DataFabric Manager.
 - Select `No` to enable the access permission checks by using the `rbac`.
5. Specify the role-based access control methods. The possible values are `native` and `dfm`.
 - Select `native` to check the access permission for the host using the control file stored in `/vol/vol0/sdprbac/sdhost-name.prbac` or `/vol/vol0/sdprbac/sdgenericname.prbac`
 - Select `dfm` to check the access permission using the Operations Manager console.



If you select `dfm` as `rbac-method` without configuring DataFabric Manager, a warning message specifying that the RBAC method is selected as `dfm` without enabling Protection Manager Integration is displayed.

6. Specify `https` or `http` to communicate with the storage system.

7. The final step is to save the configuration changes in the `snapdrive.conf` file, and restart the daemon.
 - If you select Yes, the SnapDrive daemon is restarted and the configuration changes are reflected.
 - If you select No, the variable values are changed in `snapdrive.conf` file, but the changes are not reflected.

SnapDrive modifies the following variables in the `snapdrive.conf` file.

- `contact-http-dfm-port`
- `contact-ssl-dfm-port`
- `use-https-to-dfm`
- `default-transport`
- `use-https-to-filer`
- `fstype`
- `multipathing-type`
- `vmtype`
- `rbac-method`
- `rbac-cache`

Setting values in the `snapdrive.conf` file

You can change the values in the `snapdrive.conf` file or add new name-value pairs.

You must be logged in as a root user.

Steps

1. Back up the `snapdrive.conf` file.
2. Open the `snapdrive.conf` file in a text editor.
3. To add a name-value pair, use the following format:

```
config-option-name=value value # optional comment
```

`config-option-name` is the name of the variable you want to configure; for example, `audit-log-file`.
`value` is the value you want to assign to this option.

If you want to include a comment with the name-value pair, precede the comment with a number sign (#).

You should enter only one name-value pair, per line.

If the name or the value uses a string, enclose the string in either single (`) or double ("") quotation marks. You can place the quotation marks either around the entire name-value pair or around only the value. The following examples show how you can use quotation marks and comments with name-value pairs:

```
"config-option-one=string with white space" # double quotes around the pair
```

```
config-option-two="string with white space" # double quotes around the value
```

```
config-option-2B='string with white space' # single quotes around the value
```

4. To modify a name-value pair, replace the current value with the new value.

Follow the steps below to ensure that the default values are recorded in the file.

- Add the sign (#) to the line that you want to modify.
- Copy the line.
- Activate the copied text by removing the number sign (#).
- Modify the value.

If you want to specify a blank value (for example, to disable the audit log file), enter a pair of double quotation marks ("").

5. Save the file after you make your changes.

SnapDrive for UNIX automatically checks this file each time it starts. Your changes take effect the next time it starts.

6. Restart the SnapDrive for UNIX daemon by using the `snapdrived restart` command.

Restarting the daemon is necessary for `snapdrive.conf` file changes to take effect.

Checking the version of SnapDrive for UNIX

You can verify the version of SnapDrive for UNIX by entering the `snapdrive version` command.

Steps

1. At the CLI prompt, enter the following command:

```
snapdrive version
```

Example

```
# snapdrive version
snapdrive Version 5.2
snapdrive Daemon Version 5.2
```



The only argument this command accepts is `-v`, which displays additional version details. If you include additional arguments, SnapDrive for UNIX displays a warning and then the version number.

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.