

SnapDrive® 5.0.1 for UNIX® Release Notes (AIX®, HP-UX®, Linux®, and Solaris®)

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089 USA
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: www.netapp.com
Feedback: doccomments@netapp.com
Part number: 215-07049_A0 ur001
August 2012

Contents

Release notes	6
New enhancements and features	7
Limitations	9
Fixed issues in SnapDrive 5.0.1 for UNIX	11
SnapDrive takes longer to start in storage systems with a large number of volumes	11
SnapDrive operations do not work in a pure NFSv4 environment in Data ONTAP operating in 7-Mode and Data ONTAP operating in Cluster-Mode	11
SnapDrive for UNIX prevents a clone volume from inheriting the exportfs option of the parent volume in NFS environments	12
SnapDrive does not allow you to create clones from the secondary host system in NFS environment	12
Volume-based SnapRestore fails for live filespecs with multiple host vols on dg in SFRAC environment	12
When you clone a file specification using -devicetype shared options and after reboot you cannot telnet or ssh the host in an SFRAC environment	13
Snap connect operation fails if the -destfv option is specified more than once in Data ONTAP 8.1 operating in Cluster-Mode in an NFS environment	13
SnapDrive daemon fails to start if there are mixed or upper case variable values in the snapdrive.conf file	14
SnapDrive may fail to start due to invalid format of the comment in the snapdrive.conf file	14
SnapDrive storage show -all command fails on Red Hat Enterprise Linux and Oracle Enterprise Linux	15
Known issues	16
Known issues: Related to all hosts	16
Snap management operation fails in RHEL/OEL 5.x due to an issue with umount command	16
Special characters in the AutoSupport messages shows garble values in the EMS log view of the storage system	16
SSL error occurs due to unsupported version	16

Snapshot copies created on the FlexClone volume are deleted after the clone split operation	17
SnapDrive clone split operation fails on the cloned volume, if the parent volume has been migrated in Data ONTAP operating in Cluster-Mode.	17
SnapDrive clone split estimate using Snapshot copy names along with the -volclone option does not display the underlying volume names	17
SnapDrive clone split estimate using the -fs or -dg option does not display the underlying multiple LUNs	17
Clone split cannot be performed for a LUN clone that is inside a FlexClone	18
Clone split cannot be performed for a LUN clone if an IP address is specified instead of a storage system hostname	18
Autodetect does not work for more than one LUN in a secondary volume	18
SnapDrive log files have read-write permissions only for the root users	19
Storage show for LUNs fail if volume group is unavailable in AIX and HP-UX	19
Snapdrive storage delete operation for a LUN hangs in DMP configuration.	19
If SnapDrive for UNIX daemon is started immediately after a stop, sometimes the start may fail	20
If the Operations Manager server fails to discover LUNs, SnapDrive for UNIX commands fail with an error message "Could not check"	20
Access control permissions for a file does not work as intended	20
Restoring the storage stack on the host after removing a LUN	21
Restoring qtrees in NFS-mounted file systems on SnapMirror destination	21
Restoring NFS directories with symbolic links	21
Restore operation fails with a message that the file system is mounted	21
Directory access and ownership rights change during the Snapshot restore operation	23
Known issues with Linux	23
Multiple values are not supported for anon and allow-suid options in SnapDrive for UNIX in Data ONTAP operating in Cluster-Mode	23

Attribute values of the inherited parent export policy rules are not same as the parent export policy rules in Data ONTAP operating in Cluster-Mode	23
Snap disconnect operation fails for the NFS file specification connected by using snap connect with the -readonly option	24
Known issues with AIX	24
SnapDrive storage provisioning commands fail when you specify both the options -nolvm and -nopersist during file system creation	24
Where to find product documentation and other information	25
Copyright information	26
Trademark information	27

Release notes

These release notes provide new features, limitations, and known issues for SnapDrive 5.0.1 for UNIX.

New enhancements and features

The SnapDrive 5.0.1 for UNIX release introduces new enhancements. It also lists the new features available in SnapDrive 5.0 for UNIX.

New enhancements in SnapDrive 5.0.1 for UNIX

- When snap management operations such as clone and restore are attempted to create a Snapshot copy by using NFSv3, the snap management operation fails in a pure NFSv4 environment. To overcome this issue, a new configuration variable *enable-migrate-nfs-version* is added in the `snapdrive.conf` file, that allows to clone/restore by using the higher version of NFS. By default, this variable is off. During this migration, only the protocol version is considered and other options such as `rw` and `largefiles` are not taken into account by SnapDrive for UNIX. Therefore, only the NFS version for the corresponding NFS filespec is added in the `/etc/fstab` file.

Ensure that the appropriate NFS version is used to mount the file specification by using `-o vers=3` for NFSv3 and `-o vers=4` for NFSv4.

If the user wants to migrate the NFS file specification with all the mount options, it is recommended to use `-mntopts` in the snap management operations. It is mandatory to use `nfs` in the attribute value of the `Access Protocol` in the export policy rules of the parent volume during migration in Data ONTAP operating in Cluster-Mode.

- A new variable *ping-interfaces-with-same-octet* is introduced in SnapDrive 5.0.1 for UNIX to avoid unnecessary pings through all the available interfaces in the host that might have different subnet IPs configured. If this variable is set to on, 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 off, 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.
- SnapDrive for UNIX does not allow to clone if there is no export permission for a volume in an NFS entity. To overcome this situation, disable the variable *check-export-permission-nfs-clone* in the `snapdrive.conf` file. As a result of the cloning operation, SnapDrive provides appropriate access permissions on the cloned volume.
- SnapDrive for UNIX supports Brocade FC HBA and CNAs.

New features in SnapDrive 5.0 for UNIX

- SnapDrive operations are supported after the storage system hostname is changed.
- `-mntopts` option is introduced for snap connect and snap restore operations.
- `-dump` and `-dump all` options are added to SnapDrive operations to automatically collect logs for a specific operation or for all available operations.
- Ext4 support is available from RHEL 6 and OEL 6 onwards.

- SnapDrive 5.0 for UNIX transparently handles the SnapDrive failures during the cutover phase in Data ONTAP features, such as volume migration and NetApp DataMotion for vFiler units. SnapDrive for UNIX uses the appropriate configuration variables for these failures.
- Support for Data ONTAP 8.1 operating either in 7-Mode or Cluster-Mode is available only with Data ONTAP 8.1 RC3 release and later. Data ONTAP operating in Cluster-Mode is supported on the Red Hat Enterprise Linux and SUSE Linux Enterprise Server platform.

Note: The Data ONTAP 8.1 7-Mode and Cluster-Mode features documented in the *SnapDrive 5.0 for UNIX Installation and Administration Guide* are supported from Data ONTAP 8.1 RC3 release and later.

Note: There is no updated SnapDrive for UNIX Installation and Administration Guide. For more information about the existing features, see the *SnapDrive 5.0 for UNIX Installation and Administration Guide*.

Limitations

The following is a list of the current known limitations in SnapDrive for UNIX.

Generic Limitations of SnapDrive for UNIX

- SnapDrive 5.0.1 for UNIX does not support snap operations on an NFS mount point when the volume is exported with security type as krb5, krb5i, or krb5p.
- SnapDrive for UNIX does not support symbolic links for any mount points.
- Snap operations may be inconsistent if you perform snap restore on a mount point where a different entity is mounted other than what was taken in the Snapshot copy.
- The `snapdrive config prepare luns` command cannot be used to prepare any RDM LUNs in the guest operating system.
- Due to the non-availability of the Consistency Group Snapshots (CGS) feature in Data ONTAP 8.1 operating in Cluster-Mode, SnapDrive 5.0.1 for UNIX does not support Data ONTAP 8.1 operating in Cluster-Mode in SAN environments that are running the UNIX or Linux operating systems. CGS is required in most SAN environments that are running the UNIX or Linux operating systems to take application-consistent Snapshot copies that can then be restored and recovered successfully.

However, SnapDrive 5.0.1 for UNIX continues to support Data ONTAP 8.1 operating in Cluster-Mode in NFS environments running Linux as mentioned in the Interoperability Matrix for SnapDrive for UNIX.

- SnapDrive clone split operation does not work when HTTPS is enabled for daemon communication.

Limitations of SnapDrive for UNIX on Linux platform

- SnapDrive for UNIX does not support operations on file specification/LUNs, which span across Data ONTAP operating in 7-Mode and Data ONTAP operating in Cluster-Mode storage system.
- In DATA ONTAP operating in Cluster-Mode, to configure Vserver with SnapDrive for UNIX, ensure that the IP address of the Vserver's management logical interface (LIF) is mapped to the Vserver name either in DNS or in the `/etc/hosts` file. You must also ensure that the Vserver name is configured in SnapDrive for UNIX by using the `snapdrive config set <vsadmin> <Vserver name>` command.
- Aggregates which contain Vserver volumes must be assigned to the Vserver's aggregate list to configure the Vserver and execute SnapDrive operations.
- SnapDrive 5.0.1 for UNIX does not support RDM LUN in the guest OS.

Limitations of SnapDrive for UNIX on Solaris platform

- SnapDrive 5.0.1 for UNIX supports Veritas 5.0 MP2 RP4.

Note: For more information, see bug number 392686 on Bugs Online.

- LUNs greater than 1 TB with SMI label is not supported
- SnapDrive for UNIX does not support the EFI and SMI LUNs in the same disk group.

Limitations of SnapDrive for UNIX on AIX platform

- SnapDrive 5.0.1 for UNIX does not support Live Partition Mobility (LPM) and Live Application Mobility (LAM).
- SnapDrive daemon takes a long time to restart if one or more FCP ports links are down or not available. However, SnapDrive daemon starts irrespective of the state of the FCP ports.

Fixed issues in SnapDrive 5.0.1 for UNIX

The following known issues are fixed in SnapDrive 5.0.1 for UNIX.

SnapDrive takes longer to start in storage systems with a large number of volumes

SnapDrive for UNIX takes longer to start in storage systems with a large number of volumes.

By default, AutoSupport is enabled in SnapDrive 4.2 for UNIX release and later versions. As part of the host side data collection, the number of LUNs, disk groups, host volumes and file systems, and all the volumes on the storage system are queried. In a large scale environment, where SnapDrive for UNIX is connected to multiple storage systems with a large number of volumes, the SnapDrive for UNIX daemon can take longer to start.

SnapDrive 5.0.1 for UNIX introduces a new configuration variable *autosupport-datacollection-enabled* in the `snapdrive.conf` file, which allows for data collection such as the number of LUNs, disk groups, host volumes and file systems, and all the storage system volumes. By default, this variable is set to `off`.

Workaround

It is recommended to set the *autosupport-enabled* variable to `on` in the `snapdrive.conf` file.

SnapDrive operations do not work in a pure NFSv4 environment in Data ONTAP operating in 7-Mode and Data ONTAP operating in Cluster-Mode

SnapDrive (snap restore) operations do not work in a pure NFSv4 environment where NFSv2 and/or NFSv3 are disabled in Data ONTAP operating in 7-Mode and Data ONTAP operating in Cluster-Mode.

However, NFSv4 is supported in the following scenarios:

- If `nfs.v2.enable` and `nfs.v3.enable` is set to `on` in Data ONTAP operating in 7-Mode.
- If the Vserver `nfs` service is created with NFSv3 and NFSv4, and the volume export policy rule attribute value of the `Access Protocol` is set to `nfs` or any in Data ONTAP operating in Cluster-Mode.

For more information, see bug number 619896 on Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

SnapDrive for UNIX prevents a clone volume from inheriting the `exportfs` option of the parent volume in NFS environments

A critical bug has been identified in SnapDrive 5.0 for UNIX, which prevents a clone volume from inheriting the `exportfs` option of the parent volume in NFS environments.

As an example, if the parent volume has the option `root=<hostname/IP>` specified, then the `root=<hostname/IP>` option gets deleted on the clone volume, which might cause any write operation on the clone volume to fail. If you are using any supported version of SnapManager for Oracle or SnapManager for SAP with SnapDrive 5.0 for UNIX in an NFS environment, then this bug might cause some SnapManager operations to fail.

This defect does not impact you if you have exported the parent volume with the `anon=0` option.

SnapDrive does not allow you to create clones from the secondary host system in NFS environment

SnapDrive does not allow you to create clones from the secondary host (host which does not have export permissions on the parent volume) system in NFS environment.

In the `snapdrive.conf` file, ensure that the `check-export-permission-nfs-clone` configuration variable is set to `off` to create a clone from the secondary host system in a NFS environment. By default, this variable is set to `on` to disable cloning in the secondary host system.

Volume-based SnapRestore fails for live filespecs with multiple host vols on dg in SFRAC environment

SnapDrive for UNIX works in volume-based SnapRestore in SFRAC environment. During this condition, the Veritas diskgroup import fails with the following error message

```
VxVM vxdg ERROR V-5-1-10978 Disk group 1245149592.131.hudson: import failed: Cannot find disk on slave node
```

SnapDrive for UNIX works for all live filespecs like `-fs`, `-dg`, `-hostvol`. However, the volume-based SnapRestore fails in diskgroups that has multiple host volumes. Thus, the single file SnapRestore must be used.

When you clone a file specification using `-devicetype shared` options and after reboot you cannot telnet or ssh the host in an SFRAC environment

In an SFRAC environment, when you clone a file specification using `-devicetype shared` options, and after reboot you cannot telnet or ssh the host system.

The cause of the failure is due to the `boot options` set to yes for shared mount point in the `/etc/vfstab` file.

In SnapDrive 5.0.1 for UNIX, the `boot options` set to no for shared mount point in the `/etc/vfstab` file.

Snap connect operation fails if the `-destfv` option is specified more than once in Data ONTAP 8.1 operating in Cluster-Mode in an NFS environment

When the snap connect operation is attempted with more than one `destfv` option for multiple NFS mountpoints, the `snapdrive snap connect` operation fails.

In Data ONTAP 8.1 operating in Cluster-Mode in an NFS environment, when you execute the `snapdrive snap connect` command with more than one `destfv` option to connect multiple NFS mountpoints, the snap connect operation fails.

Workaround

It is recommended not to use more than one `destfv` option in the `snapdrive snap connect` command while cloning multiple NFS mountpoints.

For more information, see bug number 594184 on Bugs Online.

Related information

[Bugs Online: support.netapp.com/NOW/cgi-bin/bol](http://support.netapp.com/NOW/cgi-bin/bol)

SnapDrive daemon fails to start if there are mixed or upper case variable values in the snapdrive.conf file

SnapDrive 5.0 for UNIX allowed only lower case entries in the `snapdrive.conf` file. In case, you enter any upper or mixed values in the `snapdrive.conf` file, snapdrive daemon fails to start while upgrade from SnapDrive 4.2 for UNIX to SnapDrive 5.0 for UNIX.

SnapDrive 5.0.1 for UNIX accepts all types of cases such as upper, lower, or mixed in the `snapdrive.conf` file, and thereby does not encounter any SnapDrive daemon failure.

Examples:

- **`use-https-to-filer=ON` #Communication with filer done via HTTPS instead of HTTP**
- **`use-https-to-dfm=On` #Communication with DFM done via HTTPS instead of HTTP**
- **`use-https-to-dfm=on` #Communication with DFM done via HTTPS instead of HTTP**

SnapDrive may fail to start due to invalid format of the comment in the snapdrive.conf file

In the `snapdrive.conf` file, a comment must be preceded by a hash mark (#).

The following is an example that displays the valid and invalid formats in the `snapdrive.conf` file.

Example 1:

- Valid format:
`enable-fcp-cache="on" #Enables FCP Cache in Assistants`
- Invalid format:
`enable-fcp-cache="on" Enables FCP Cache in Assistants`

You might encounter any one of the following error messages, if there is an invalid format in the `snapdrive.conf` file.

Error message 1:

```
# snapdrive version -v
Error:
-----
/opt/NetApp/snapdrive/snapdrive.conf failed validation checks.
Cause:
-----
Value of Configuration Variable is not recognized.
Details
-----
```

```
Variable Name      : enable-fcp-cache
Allowed  Regex Pattern: (^on$|^off$)
Value    Given      : on" Enable FCP Cache in Assistants
Corrective Action
-----
Set the variable value so that it matches the regex pattern.
```

Example 2:

- Valid format:
***trace-log-save=100* #Number of old copies of trace log file to save**
- Invalid format1:
***trace-log-save=100* Number of old copies of trace log file to save**
- Invalid format2:
***trace-log-save=100* #Number of old copies of trace log file to save
#test**

Error message 2:

```
# snapdrived start
terminate called after throwing an instance of
'boost::exception_detail::clone_impl<boost::exception_detail::error_in
fo_injector<boost::bad_lexical_cast> >'
what(): bad lexical cast: source type value could not be
interpreted as target
IOT/Abort trap (core dumped)
```

Workaround

Ensure that the values for any configuration variable is within double quotes and any succeeding characters after the double quotes are preceded with only one # mark.

SnapDrive storage show -all command fails on Red Hat Enterprise Linux and Oracle Enterprise Linux

If you are using SnapDrive for UNIX on Red Hat Enterprise Linux and Oracle Enterprise Linux, the `snapdrive storage show -all` command fails with the Status call to SDU daemon failed error message. To avoid this error message, set the `LVM_SUPPRESS_FD_WARNINGS` environment variable to 1 in the shell.

Known issues

To use the SnapDrive for UNIX product more effectively, you should be aware of the following known issues and workaround for some features.

Known issues related to all hosts

Some potentially undesirable behaviors that impact all supported host platforms have been identified in this release. Workaround is provided where available.

Snap management operation fails in RHEL/OEL 5.x due to an issue with umount command

Snap management operation fails in RHEL/OEL 5.x due to an issue with the `umount` command. To overcome this, ensure that the NFS entity is mounted using `-o vers=3` in the command line interface.

Special characters in the AutoSupport messages shows garble values in the EMS log view of the storage system

When AutoSupport is enabled in SnapDrive for UNIX, upon any operation failure, the error messages are logged in the EMS log view of the storage system. If the error message contains any special characters such as (`<`, `>`, `&`, `'`, `"`, and `\r`), garble values are displayed in the EMS log view of the storage system.

Workaround

None.

SSL error occurs due to unsupported version

An SSL error occurs, if you have installed an unsupported version of openssl software on the host.

Workaround

It is recommended to use the supported SSL software package that is bundled with the Native OS image.

Snapshot copies created on the FlexClone volume are deleted after the clone split operation

When a clone is split from the parent volume, the Snapshot copies created on the FlexClone volume are deleted as a result of the clone split operation.

Workaround

None.

SnapDrive clone split operation fails on the cloned volume, if the parent volume has been migrated in Data ONTAP operating in Cluster-Mode.

SnapDrive clone split operation fails on the cloned volume, if the parent volume has been migrated from one aggregate to another in Data ONTAP operating in Cluster-Mode.

Workaround

It is recommended to split the clone on the storage system.

SnapDrive clone split estimate using Snapshot copy names along with the -volclone option does not display the underlying volume names

SnapDrive clone split estimate using Snapshot copy with the `-volclone` option displays the resource name as Snapshot copy names, and if there are multiple volumes, it is difficult to identify the corresponding estimates for the underlying volumes.

Workaround

None.

SnapDrive clone split estimate using the -fs or -dg option does not display the underlying multiple LUNs

When you execute the `snapdrive clone split estimate` command by using the `-fs` or `-dg` option on a FlexClone containing multiple LUNs, it displays the file specifications instead of the corresponding underlying LUN names. Therefore, it is difficult to identify the clone split estimation for the corresponding underlying LUNs.

Workaround

None.

Clone split cannot be performed for a LUN clone that is inside a FlexClone

When you execute the `snapdrive clone split start` command, the clone split start operation fails for a LUN clone that resides inside the FlexClone.

Workaround

If you want to perform a clone split start for a LUN clone that is inside a FlexClone, first perform the clone split start operation on the storage system.

Clone split cannot be performed for a LUN clone if an IP address is specified instead of a storage system hostname

When you execute the `snapdrive clone split start` command and if an IP address is specified instead of a storage system hostname by using the `-lun` option, the clone split start operation fails.

Workaround

Provide the storage system path along with the IP address in the `-lun` option of the `snapdrive clone split start` command.

```
snapdrive clone split start -lun 10.7.122.7:/vol/vol1/lnx200-106_vol1_lun2 -splitmode
```

Autodetect does not work for more than one LUN in a secondary volume

Autodetect does not work for more than one LUN in a secondary volume.

If you delete more than one LUN in a secondary volume, and execute the `snapdrive storage delete` command, SnapDrive for UNIX displays an error message as "Could not detect type of the entity..." and deletes only the first LUN in the secondary volume.

```
# snapdrive storage delete storagesystem:/vol/user/test1 test2
/vol/user/test3 test4
Auto detection of file_spec(s) in progress ...
- 'storagesystem:/vol/user/test1' detected as LUN
- 'test2' detected as LUN
- '/vol/user/test3' detected as LUN
9000-305 Command error: Could not detect type of the entity test4. Provide
a specific option (-lun, -dg, -fs or -lvol) if you
know the type of the entity.
```

Workaround

None.

SnapDrive log files have read-write permissions only for the root users

The log files in SnapDrive for UNIX such as trace, audit, and recovery have read-write permissions only for the root users and read permission for all users.

Workaround

None.

Storage show for LUNs fail if volume group is unavailable in AIX and HP-UX

On AIX and HP-UX hosts, the `snapdrive storage show -lun lun_name` and the `snapdrive storage show -all` commands fail if any volume group on the system is inactive (HP-UX) or varied off (AIX).

SnapDrive for UNIX cannot determine the volume groups to which the LUNs belong. This might be because the command specifies one or more LUNs using the options `-lun` and `-all`, and the volume group might not be available, thus the command fails. If you use one of the other command options (`-dg`, `-hostvol`, or `-fs`), the command captures the information for that entity and returns it. Other forms of the `snapdrive storage show` command work without issues if you do not specify an inactive or varied off volume group.

Workaround

If possible, activate the inactive volume group using the `varyon` command in AIX and the `vgchange` command in HP-UX. If you have a host cluster, activating all the volumes might not be practical. In this case, use the `snapdrive storage show` command with the `-dg`, `-hostvol`, or `-fs` option rather than the `-lun` or `-all` option.

Snapdrive storage delete operation for a LUN hangs in DMP configuration.

On hosts with FC and DMP configurations, when you execute the `snapdrive storage delete -lun` command, the system hangs for more than 10 minutes.

This issue is caused by incorrect installation or configuration of the following components:

- NTAPasl
- Veritas stack (storage foundation)
- Multipathing licenses

Workaround

Refresh the DMP path information after any FC path is enabled, disabled, or added. To refresh the DMP path, execute the following commands in the sequential order:

- For an AIX host:
 1. `cfgmgr`

2. `vxdisk scandisk`
- For an HP-UX host:
 1. `ioscan -fnC disk`
 2. `ioinit -i`
 3. `vxdisk scandisks`
- For a Solaris host:
 1. `devfsadm -Cv`
 2. `vxdisk scandisk`

If SnapDrive for UNIX daemon is started immediately after a stop, sometimes the start may fail

If SnapDrive for UNIX daemon is started immediately after a stop, sometimes the start may fail due to child processes waiting on the port for a longer period. The waiting period usually depends on the TCP timeouts and the life of that child process. In such cases, the daemon cannot use that port and the daemon start might fail.

Workaround

Wait until processes in the Time Wait state are timed out. This wait period can be as long as five minutes.

Note: This issue might occur more frequently on Solaris platforms.

If the Operations Manager server fails to discover LUNs, SnapDrive for UNIX commands fail with an error message "Could not check"

SnapDrive for UNIX commands fail, if the operations manager server fails to discover the LUNs.

The following error message is displayed: 0002-332 Admin error: Could not check SD.Storage.Write access on lun storage_array1:/vol/vol1/lun1 for user unix-host\root on Operations Manager server(s) ops-mngr-server1

Workaround

None.

Access control permissions for a file does not work as intended

The local access control mechanism does not work as intended. This issue occurs when the storage system does not contain a `vol0`.

By default, SnapDrive for UNIX expects to find the permissions file for local access control at the path `/vol/vol0/sdprbac/sd<hostname>.prbac`, where `<hostname>` is the name of the local host. This problem does not occur if the `vol0` file is on the storage system.

Workaround

None.

Restoring the storage stack on the host after removing a LUN

You should always use the SnapDrive storage commands to remove LUNs. Using external commands to remove LUNs can have an unpredictable effect on the host system.

Workaround

If you remove a LUN without using a SnapDrive storage command, you should reboot the host to restore the storage system stack to a known state. You can also restore the host system state by entering the `snapdrive snap restore` command, and specifying a Snapshot copy that contains an instance of the removed LUN..

Restoring qtrees in NFS-mounted file systems on SnapMirror destination

SnapDrive for UNIX does not allow you to restore a Snapshot copy of a qtree, which is mounted on Network File (NFS-mounted) system on a SnapMirror destination. SnapDrive for UNIX allows you to restore LUNs on qtrees on a SnapMirror destination.

Workaround

None.

Restoring NFS directories with symbolic links

SnapDrive for UNIX fails if you use it to restore an NFS directory that contains symbolic links.

Workaround

None.

Restore operation fails with a message that the file system is mounted

If you unmount the file system and attempt to restore the Snapshot copy, the operation fails and you get a message indicating that the file system is mounted.

Using a Linux host, you can continuously mount the same file system. You can then make a Snapshot copy using SnapDrive for UNIX. The Snapshot copy, though, contains only the first entry in the mount table for the file system. When you issue an `umount` command, it unmounts the most recent mountpoint, and not the one shown in the Snapshot copy. When you attempt to restore the Snapshot copy, it fails with a message that the volume is already mounted. The following is an example to illustrate if a file system is mounted more than once.

```
# mount /dev/dg3/vol1 /var/tmp/mnts/dg3/vol1
# mount /dev/dg3/vol4 /mnt/um1
# mount /dev/dg3/vol4 /var/tmp/mnts/dg3/vol4
```

```
# mount |grep dg3
/dev/dg3/vol1 on /var/tmp/mnts/dg3/vol1 type ext3 (rw)
/dev/dg3/vol4 on /mnt/um1 type ext3 (rw)
/dev/dg3/vol4 on /var/tmp/mnts/dg3/vol4 type ext3 (rw)

# snapdrive snap create -dg dg3 -snapname dg3_vol2_snapshot
snap create: snapshot dg3_vol2_snapshot contains:
disk group dg3 containing host volumes
vol1 (filesystem: /var/tmp/mnts/dg3/vol1)
vol2
vol3
vol4 (filesystem: /mnt/um1)
snap create: created snapshot eccentric:/vol/vol1:dg3_vol2_snapshot

# snapdrive snap show -v -snapname eccentric:/vol/vol1:dg3_vol2_snapshot
snap name          host
date              snapped
-----
eccentric:/vol/vol1:dg3_vol2_snapshot childs Sep 1 14:14 dg3
host OS: Linux 2.4.21-9.ELsmp #1 SMP Thu Jan 8 17:08:56 EST 2004
snapshot name: dg3_vol2_snapshot
Volume Manager: linuxlvm 1.0.3
disk group:      dg3
host volume:    vol1
host volume:    vol2
host volume:    vol3
host volume:    vol4
file system: vol1 file system type: ext3      mountpoint:
/var/tmp/mnts/dg3/vol1
file system: vol4 file system type: ext3      mountpoint: /mnt/um1
lun
path
      dev paths
-----
eccentric:/vol/vol1/swl14          /dev/sds

# snapdrive snap restore -dg dg3 -snapname dg3_vol2_snapshot
Starting to restore LUNs in disk group dg3
WARNING: This can take several minutes.
DO NOT CONTROL-C!
If snap restore is interrupted, the disk group being restored may have
inconsistent or corrupted
data. For detailed progress information, see the log file /var/log/sd-
recovery.log
Importing dg3 Error (Command error 0001-034): mount failed:
mount: /dev/dg3/vol4 already mounted or
/mnt/um1 busy mount: according to mtab, /dev/dg3/vol4 is already mounted
on /mnt/um1
```

Workaround

Ensure that the file system is mounted only once.

Directory access and ownership rights change during the Snapshot restore operation

NFS directory access and ownership rights change during the Snapshot restore operation. For nonexistent NFS directories on Data ONTAP 7.0 and 7.1, and for existing and nonexistent NFS directories on Data ONTAP 7.2, the directory access and permission rights change from user to root user.

Workaround

To change the ownership rights, enter the following command at the command prompt: `chown username directory name`. To change the access rights, enter the following command at the command prompt: `chmod permission directory name`.

Known issues with Linux

Some potentially undesirable behaviors that impact Linux host platform have been identified in this release. Workaround is provided where available.

Multiple values are not supported for anon and allow-suid options in SnapDrive for UNIX in Data ONTAP operating in Cluster-Mode

In Data ONTAP operating in Cluster-Mode, NFS snap management operations fail, if the export policy with multiple rules have different attribute values for anon and allow-suid options in the export policy rules.

Workaround

It is recommended to have uniform attribute values for anon and allow-suid options for all the rules in the export policy.

Attribute values of the inherited parent export policy rules are not same as the parent export policy rules in Data ONTAP operating in Cluster-Mode

In Data ONTAP operating in Cluster-Mode, when you perform snap connect of NFS file specification from the secondary host (host which does not have export permissions on the parent volume), SnapDrive creates a new export policy, and inherits the parent export policy rules, and appends the rules in the secondary host, if the parent export policy rules does not have export permission.

The newly created export policy does not have same attribute values for Access Protocol and Superuser Security Flavors for parent (which are inherited) export policy rules.

Workaround

None.

Snap disconnect operation fails for the NFS file specification connected by using snap connect with the -readonly option

In Data ONTAP operating in Cluster-Mode, the snap disconnect operation fails for the NFS file specification connected by using `snap connect` with the `-readonly` option. This is due to the restriction in the junction path naming convention.

Workaround

The junction path should be same as a volume name in SnapDrive 5.0 for UNIX and SnapDrive 5.0.1 for UNIX versions.

Known issues with AIX

Some potentially undesirable behaviors that impact AIX host platform have been identified in this release. Workaround is provided where available.

SnapDrive storage provisioning commands fail when you specify both the options -nolvm and -nopersist during file system creation

When you create or connect a file system by using both the `-nolvm` and `-nopersist` options, `snapdrive storage create`, `snapdrive storage connect`, and `snapdrive host connect` commands fail in AIX.

Workaround

It is recommended to use any one of the options, `-nolvm` and `-nopersist`, when creating a file system on an AIX platform.

Where to find product documentation and other information

The NetApp Support site contains the most current documentation for SnapDrive for UNIX, FC and iSCSI Host Utilities, SAN storage system components, and various interconnects.

The following table describes some of the documents on the NetApp Support site that contain current information about host requirements, supported configurations, and procedures for installing and managing SnapDrive for UNIX.

If you need more information about...	Go to...
The latest supported configurations	The <i>NetApp Interoperability Matrix</i>
Supported Fibre Channel SAN topologies	The <i>SAN Configuration Guide</i>
How to download and install SnapDrive for UNIX	The <i>SnapDrive for UNIX Installation and Administration Guide</i>
Configuring your host for the Host Utilities	The Host Utilities documents are in the SAN/IPSAN Information Library

Related information

[Access and Order Product Documentation page - support.netapp.com/documentation/productsatoz/index.html](http://support.netapp.com/documentation/productsatoz/index.html)

[NetApp Interoperability Matrix -support.netapp.com/NOW/products/interoperability/](http://support.netapp.com/NOW/products/interoperability/)

[SAN Configuration Guide \(called Fibre Channel and iSCSI Configuration Guide in Data ONTAP 8.1 and earlier\)](#)

Copyright information

Copyright © 1994–2012 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the United States and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the United States and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the United States and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks.

NetApp, Inc. NetCache is certified RealSystem compatible.