

Solaris™ UFS

Backup and Recovery in a SAN Environment

Toby Creek | 11/10/2003 | TR-3292

TECHNICAL REPORT

Network Appliance, a pioneer and industry leader in data storage technology, helps organizations understand and meet complex technical challenges with advanced storage solutions and global data management strategies.

Abstract

This paper details the implementation of file system freeze/thaw functionality with the Solaris Filesystem (UFS) with NetApp® filer Snapshot™ technology to perform backup and recovery using iSCSI and Fibre-Channel filers.

Table of Contents

1	Purpose and Scope	3
2	Requirements and Assumptions	3
3	Best Practices	3
3.1	Snapshot and Storage Design	3
3.2	Archive to Tape Considerations	4
3.3	VERITAS Volume Manager Considerations	4
4	Using Snapshot Technology with the UFS File System	4
4.1	Overview	4
4.2	Creating a Snapshot Record of a UFS File System	5
4.3	Restoring UFS File Systems Using SnapRestore®	6
4.4	Mounting Writable Snapshot LUNs	7
5	Conclusions	7
6	Caveats	7
7	References	8

1) Purpose and Scope

This paper covers the techniques for utilizing Snapshot technology available in Network Appliance™ storage systems with the Sun™ Solaris UFS file system. Specifically, this paper covers the following issues:

- Backing up a UFS file system using Snapshot
- Restoring a UFS file system using Snapshot
- Connecting writable Snapshot UFS LUNs to a host

2) Requirements and Assumptions

For the methods and procedures in this paper to be useful to the reader, several assumptions are made:

- The reader has at least basic UNIX® administration skills and has access to the administrative login for the server.
- The reader has at least basic Network Appliance administration skills and has administrative access to the filer via the command-line interface.
- The filer has the necessary licenses to perform the activities outlined in this paper.
- The target system has the required block-level and network protocol interconnects to perform the activities outlined in this paper.

In the examples in this paper, all administrative commands are performed at the server or filer console for clarity. Web-based management tools can also be used.

This paper and its code samples are written for the Sun Solaris operating environment. Minor modifications to the source code to fit a particular environment may be required.

3) Best Practices

3.1) Snapshot and Storage Design

To facilitate most effective use of NetApp Snapshot technology, recommendations for designing the storage configuration are presented in this section. These recommendations are designed to prevent configuration issues from impacting data integrity or the ability to take or restore Snapshot records.

All UFS-formatted LUNs in a given filer volume should be frozen during the Snapshot process.

The Snapshot process occurs at the volume level on the filer. All LUNs within the filer volume are included in the Snapshot and should be consistent at the time the Snapshot is taken. All file systems using those LUNs should be frozen for the Snapshot creation. If this procedure is followed, any LUNs that are restored from the Snapshot will be consistent.

If multiple hosts have LUNs in a single filer volume, prefix the Snapshot name with the hostname. Snapshot records can be synchronized with only one host at a time. Using the hostname as the prefix will allow the administrator to quickly identify which Snapshot records are consistent for a given host. In these environments, only the single-LUN SnapRestore feature should be used.

Additional recommendations are presented in the Network Appliance SAN System Administrator's Guide. A link to this document can be found in the references section of this paper.

3.2) Archive to Tape Considerations

When using Snapshot records as the source for archiving to tape, some conventional backup-to-tape solutions are no longer relevant. This section will present a few high-level recommendations for designing a three-tier backup solution using Snapshot records.

Integrate the freeze/snap/thaw process into the prebackup facilities of the backup package. All enterprise-level backup packages have the ability to call other scripts to prepare the system for backup. The scripts presented in this paper can easily be integrated with leading packages to handle the creation of a Snapshot record before the backup-to-tape operation begins.

Allow the backup software to be the scheduler for the Snapshot process. All enterprise-level backup software packages have error notification and reporting functions built in. If the backup software is used as the Snapshot scheduler, rather than another facility such as “cron,” the reporting function can be accessed from a single console. The reporting capabilities will often be much more robust as well, with features such as notification to a network management console or pager.

If using NDMP, make sure that a consistent Snapshot record is used as the source for the backup. During NDMP backup operations, the filer will create a Snapshot record to use as a static backup image unless a specific Snapshot path is specified in the backup. Since the filer automatically creates the Snapshot record, no file system synchronization is performed, and the file system being backed up may not be consistent. Specifying a Snapshot path, such as /vol/vol1/.snapshot/dbhost.0, will allow the administrator to select a consistent Snapshot record for backup.

Additional recommendations may be made by the backup software vendor and Network Appliance. Consult the relevant documentation when designing a backup solution.

3.3) VERITAS Volume Manager Considerations

Special care must be taken when UFS file systems reside in Volume Manager (VxVM) volumes to maintain the correct disk group configuration.

The LUNs that make up a VxVM disk group must be treated as a single unit. When a Snapshot process is triggered, all LUNs in that disk group must exist in the Snapshot record. This will require that all file systems on volumes in the disk group be frozen during the Snapshot process. When a volume within the disk group is restored, all LUNs (and consequently, all volumes and file systems) in the disk group must be restored from the same Snapshot record. Restoring the disk group will require the disk group to be deported for the restore, so all file systems within the disk group must be unmounted and unavailable during the restore. If new LUNs have been added to the disk group since the Snapshot record was created, they must be removed when the Snapshot record is restored to ensure that the disk group configuration is returned to the configuration at the time the Snapshot record was created.

4) Using Snapshot Technology with the UFS File System

4.1) Overview

The process of creating Snapshot records in the SAN environment differs from the process in the NAS environment in one very fundamental way: in the SAN environment, the filer does not control the state of the file system. For this reason, Snapshot processes must be initiated from the host after the appropriate operations have been performed to ensure that a consistent file system image is obtained in the Snapshot record. These operations are commonly referred to as “freeze” and “thaw.” The freeze

operation flushes any dirty buffers in the file system cache and then suspends new activity on the file system until the thaw operation is performed.

If freeze and thaw are not performed on the file system, there are several failure scenarios that are possible with file systems in general and UFS in particular. In the first, if logging is enabled, minimal log replay will be required, in which changes from the journal are applied to the file system. This is generally not a time-consuming operation. In the second, a full file system sanity check, performed by the “fsck” process, will be required to return the file system to a usable state. This check can take from minutes to hours depending on the size of the file system, during which time the file system cannot be mounted or otherwise accessed. Data may be lost during the fsck process. The final possibility is that the file system is completely unusable and all data is lost.

Sun provides a command-line utility called “lockfs” to flush and freeze a file system prior to creating a Snapshot record. This program takes a file system mount point and an operation as arguments. While the file system is frozen, processes that attempt to write to the file system will block until the thaw operation occurs.

The simplest way to automate the freeze/snap/thaw process is to use a scripting language available on the host to call “lockfs” to perform the freeze and “rsh” on the filer to manage the Snapshot records. The script presented uses the Korn shell. PERL, C-shell, and many other languages are also suitable to the task.

4.2) Creating a Snapshot Record of a UFS File System

Generally, it is preferable to script the file system freeze, flush, and thaw commands with the commands to the filer. A sample script to create a Snapshot record of the /u01 file system is presented below:

```
#!/bin/ksh

# Catch signals and unlock all file systems
trap "lockfs -ua" 1 2 3 15

# Rename and delete old snapshot records
rsh -l root data snap delete voll sunhost.3
rsh -l root data snap rename voll sunhost.2 sunhost.3
rsh -l root data snap rename voll sunhost.1 sunhost.2
rsh -l root data snap rename voll sunhost sunhost.1

# Flush and suspend write I/O
lockfs -w -c "locked for snapshot" /u01
lockfs -f /u01

# Take a new snapshot
rsh -l root data snap create voll sunhost

# Resume normal I/O
lockfs -u /u01
```

```
# Exit nicely
exit 0
```

To create a Snapshot record of a UFS file system, execute the script from the command line as in the following example:

```
# ksh nasnapufs.ksh
```

The script will output basic status messages as it performs operations on the filer and the file systems involved in the Snapshot process.

4.3) Restoring UFS File Systems Using SnapRestore

Snapshot records provide a very efficient and time-conserving way to restore file systems. Restoring a LUN that contains a UFS file system is easily accomplished. The steps to restore a pool are detailed below:

1. Unmount the file system to be restored. If a volume-level SnapRestore process is used, all file systems with LUNs in the filer volume being restored must be unmounted.

```
# umount /u01
```

2. Use the appropriate SnapRestore command on the filer. To restore a single LUN:

```
filer> lun offline /vol/vol1/ufs0.lun
filer> snap restore -t file -s snap.0 /vol/vol1/ufs0.lun
filer> lun online /vol/vol1/ufs0.lun
```

To restore a filer volume and all of its LUNs:

```
filer> snap restore -t vol -s snap.0 vol1
```

```
WARNING! This will revert the volume to a previous snapshot.
All modifications to the volume after the snapshot will be
irrevocably lost.
```

```
Volume vol1 will be made restricted briefly before coming back
online.
```

```
Are you sure you want to do this? y
```

```
You have selected volume vol1, snapshot snap.0
```

```
Proceed with revert? y
Volume vol1: revert successful.
```

3. Remount the restored file systems once the SnapRestore process has completed.

```
# mount /u01
```

For environments where VERITAS Volume Manager is used, the disk group consisting of filer LUNs must be deported before the SnapRestore process and imported after the SnapRestore process completes. As mentioned previously, all LUNs in the disk group must be restored at the same time.

4.4) Mounting Writable Snapshot LUNs

The mounting of writable Snapshot LUNs can be used to restore individual files or to allow the backup-to-tape process to occur on a second host to offload the process.

The procedure for mounting writable Snapshot LUNs is as follows:

1. Create the writable Snapshot LUN on the filer and map it to the host.

```
filer> lun create -b /vol/vol1/.snapshot/snap.0/vxfs0.lun
/vol/vol1/vxfssnap0.lun
filer> lun map /vol/vol1/vxfssnap0.lun vxfshost
lun map: auto-assigned vxfshost=1
```

2. Create the device files if none exist already.

```
# devfsadm
```

3. Mount the file system on the host.

```
# mount /dev/dsk/c1t2d1s6 /u01
```

4. When the LUN is no longer needed, it can be unmounted and destroyed.

```
# umount /mnt
filer> lun destroy -f /vol/vol1/vxfssnap0.lun
```

This example is specific to Solaris without any volume management software. This process does not work on a host using VERITAS Volume Manager. Because of the unique disk identifiers that VxVM assigns to the LUNs, it is difficult to mount writable Snapshot LUNs on the same host where the current incarnation of the disk group resides. Though there are unsupported procedures for changing the disk identifiers to make this possible, it is recommended that the disks be imported on a second host instead.

5) Conclusions

A Network Appliance filer offers the Solaris administrator using UFS compelling advantages in terms of backup and recovery. Use of Snapshot records, combined with conventional backup-to-tape techniques, can dramatically optimize the server backup operation. Retaining a number of online Snapshot records allows the system administrator to restore file systems without the necessity to restore from tape in many circumstances. Backup and recovery performance is dramatically improved over conventional local disk and SAN configurations, improving Mean-Time-to-Recovery (MTTR) intervals.

6) Caveats

This paper is not intended to be a definitive implementation guide. There are many factors that may not be addressed in this paper. Expertise may be required to solve logistical problems when the system is designed and built. NetApp has not tested this procedure with all of the combinations of hardware and

software options available on Solaris. There may be significant differences in your configuration that will alter the procedures necessary to accomplish the objectives outlined in this paper. If you find that any of these procedures do not work in your environment, please contact the [author](#) immediately.

7) References

Sun documentation on the Web:

<http://docs.sun.com/>

Network Appliance SAN System Administrator's Guide:

http://now.netapp.com/NOW/knowledge/docs/ontap/rel641/html/ontap/san_sag/index.htm

© 2003 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, the Network Appliance logo, and SnapRestore are registered trademarks and Network Appliance and Snapshot are trademarks of Network Appliance, Inc. in the U.S. and other countries. Sun and Solaris are trademarks of Sun Microsystems, Inc. UNIX is a registered trademark of The Open Group. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.



Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089
www.netapp.com

Network Appliance, Inc. Company Confidential and Proprietary.

© 2002 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, the Network Appliance logo, FAServer, FilerView, NetCache, SecureShare, SnapManager, SnapMirror, SnapRestore, and WAFL are registered trademarks and Network Appliance, ApplianceWatch, BareMetal, Camera-to-Viewer, Center-to-Edge, ContentDirector, ContentFabric, ContentReporter, DataFabric, Data ONTAP, EdgeFiler, HyperSAN, InfoFabric, MultiStore, NearStore, NetApp Availability Assurance, NetApp ProTech Expert, NOW, NOW (NetApp on the Web), RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, Smart SAN, SnapCache, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapMigrator, Snapshot, SnapSuite, SnapVault, SohoCache, SohoFiler, The evolution of storage, Vfiler, and Web Filer are trademarks of Network Appliance, Inc. in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.