

AIX File Systems

Back Up and Restore Using a NetApp Filer in a SAN Environment

Toby Creek | 1/25/2004 | TR 3306

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 document details the implementation of file system backup and restore using AIX's journaling file systems (JFS and Enhanced JFS) 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	Storage Design Considerations	
3.2	Archive to Tape Considerations	
3.3	AIX Logical Volume Manager Considerations	
4	Using Snapshot Technology with the JFS File System	
4.1	Overview	
4.2	Creating a Snapshot Backup of a JFS File System	
4.3	Restoring JFS File Systems Using SnapRestore®	
4.4	Mounting Writable Snapshot LUNs	
5	Conclusions	
6	Caveats	
7	References	

1) Purpose and Scope

This document covers the techniques for utilizing Snapshot technology available in Network Appliance™ storage systems with IBM AIX journaling file system products. Specifically, this report covers the following issues:

- Backing up a JFS or Enhanced JFS (JFS2) file system using Snapshot technology
- Restoring a JFS or JFS2 file system from a Snapshot backup.
- Connecting writable Snapshot LUNs to a host

2) Requirements and Assumptions

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

The reader has at least basic AIX 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 and host have the necessary licenses to perform the activities outlined in this document

The target system has the required block-level and network protocol interconnects to perform the activities outlined in this document

The target file systems are using the JFS or Enhanced JFS file system format

In the examples in this report, all administrative commands are performed at the server or filer console for clarity. Web-based or other management tools such as Network Appliance FilerView® or the AIX System Management Interface Tool (SMIT) can also be used.

This report was written for the AIX operating environment, but may be applicable to other UNIX® variants. Minor modifications to the procedure provided may be required.

3) Best Practices

3.1) Storage Design Considerations

To help users utilize NetApp Snapshot technology most effectively, this section presents recommendations for designing the storage configuration. These recommendations are designed to prevent configuration issues from impacting data integrity or the ability to create or restore a Snapshot backup.

If SnapRestore is used to restore LVM LUNs, all LUNs in the volume group must be restored at the same time—when the Snapshot creation is triggered, all LUNs in the volume group must exist in the Snapshot image. The LVM volume group is the most granular object that can be restored from a Snapshot backup. When a volume group is restored, all LUNs (and consequently, all logical volumes and file systems) in the volume group are restored from the same Snapshot backup. The restore process will require the volume group to be deactivated for the restore, so all file systems on logical volumes within the volume group must be unmounted and unavailable during the restore. If new LUNs have been added to the volume group since the Snapshot image was created they must be removed

when the Snapshot image is restored to ensure that the volume group configuration is returned to the configuration at the time the Snapshot image was created.

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 a Snapshot instance 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 technology.

Integrate the Snapshot process into the prebackup facilities of the backup package. All enterprise-level backup packages have the ability to call 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 image 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 you utilize the backup software as the Snapshot backup 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 image is used as the source for the backup. During NDMP backup operations, the filer will create a Snapshot image to use as a static backup image unless a specific Snapshot path is specified in the backup. Since the filer automatically creates the Snapshot image, no file system synchronization is performed, and the file system being backed up may not be consistent. Specifying the Snapshot path, for example, `/vol/vol1/.snapshot/dbhost.0`, will allow the administrator to select a consistent Snapshot image 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) AIX Logical Volume Manager Considerations

Special care must be taken when JFS file systems reside in AIX Logical Volume Manager (LVM) volumes to maintain the correct volume group configuration.

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

4) Using Snapshot Technology with the JFS File System

4.1) Overview

The Snapshot process in the SAN environment differs from that of 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, the Snapshot process must be initiated from the host after the appropriate operations have been performed to ensure that a consistent file system state is obtained in the Snapshot image. 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, several failure scenarios are possible with journaling file systems in general. In the first, 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.

AIX JFS and JFS2 do not provide a user-controllable mechanism to quiesce the file system during Snapshot creation. Since the storage system does not have knowledge of the file system state during this operation, consistency can only be guaranteed by the effectiveness of the journaling implementation in the file system code. Log replay will be required when mounting the file system.

4.2) Creating a Snapshot Backup of a JFS File System

Snapshot image creation is triggered from the host using one of several methods available on the filer.

1. The remote shell (rsh or remsh) facility can be used to perform commands on the filer as if typed at the console. All commercial UNIX variants contain client remote shell programs. This method is generally easier to implement, but does not contain robust error reporting.
2. The ONTAPI™ management programming interfaces on the filer are accessible from many programming languages, including C and PERL. Though the code written to use these APIs is somewhat more complex than the remote shell commands, error reporting and functionality are much improved.

For simplicity, this paper will use the remote shell method to trigger Snapshot creation. A Snapshot image of a filer volume containing AIX LUNs can be created following the example below:

```
# rsh filer snap create voll snap.0
```

4.3) Restoring JFS File Systems Using SnapRestore

Snapshot technology provides a very efficient and time-conserving way to restore file systems. Restoring a LUN that contains a JFS file system is easily accomplished. The steps to restore a volume group are detailed below:

1. Unmount all file systems on logical volumes in the volume group to be restored.

```
# umount /u01
# umount /u02
```

2. Deactivate the LVM volume group. If a volume-level SnapRestore command is used, all volume groups with LUNs in the filer volume being restored must be deactivated.

```
# varyoffvg ntapvg
```

3. Use the appropriate SnapRestore command on the filer. To restore a volume group consisting of two LUNs:

```
filer> lun offline /vol/vol1/aixserv0.lun
filer> lun offline /vol/vol1/aixserv1.lun
filer> snap restore -t file -s snap.0 /vol/vol1/aixserv0.lun
filer> snap restore -t file -s snap.0 /vol/vol1/aixserv1.lun
filer> lun online /vol/vol1/aixserv0.lun
filer> lun online /vol/vol1/aixserv1.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.
```

4. Reactivate the LVM volume group.

```
# varyonvg ntapvg
```

5. Remount the restored file systems if they are not automatically remounted.

```
# mount /dev/ntapl01 /u01
# mount /dev/ntapl02 /u02
```

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/aixserv0.lun
/vol/vol1/aixservsnap0.lun
filer> lun create -b /vol/vol1/.snapshot/snap.0/aixserv1.lun
/vol/vol1/aixservsnap1.lun
filer> lun map /vol/vol1/aixservsnap0.lun aixserv
lun map: auto-assigned aixserv=2
filer> lun map /vol/vol1/oraservsnap1.lun aixserv
lun map: auto-assigned aixserv=3
```

2. Scan for new disk devices on the host.

```
# cfgmgr
```

3. Use the “sanlun” and “setsp” commands to determine the paths to the new disk devices.

```
# sanlun lun show
# setsp -a
```

4. Change the physical volume identifier (PVID) on the Snapshot LUNs.

```
# chdev -l hdisk4 -a pv=clear
# chdev -l hdisk8 -a pv=clear
```

5. Create a new volume group for the Snapshot LUNs.

```
# recreatevg -y snapvg -Y snap hdisk4 hdisk8
```

This command creates a volume group called “snapvg” and prefixes all of its logical volumes with “snap.” This helps the administrator easily identify the logical volumes provided from the Snapshot LUNs.

6. Activate the volume group.

```
# varyonvg snapvg
```

7. Mount the file systems from the logical volumes in the volume group.

```
# mount /dev/snaplv01 /u01.snap
# mount /dev/snaplv02 /u02.snap
```

8. When the LVM volume group is no longer needed, it can be unmounted and destroyed.

```
# vgchange -a n orasnapvg
# vgexport orasnapvg
# umount /u01.snap
# umount /u02.snap
```

```
filer> lun destroy -f /vol/voll/aixservsnap0.lun
filer> lun destroy -f /vol/voll/aixservsnap1.lun
```

5) Conclusions

A Network Appliance filer offers the UNIX administrator using AIX LVM and JFS compelling advantages in terms of backup and recovery. Use of Snapshot technology, combined with conventional backup-to-tape techniques, can dramatically optimize the server backup operation. Retaining a number of online Snapshot images allows the system administrator to restore file systems without the necessity of restoring from tape in many circumstances. Backup and recovery performance is dramatically improved over that of 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 document. 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 UNIX variants. 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

Network Appliance Product Documentation

<http://now.netapp.com/>

© 2004 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, the Network Appliance logo, FilerView, and SnapRestore are registered trademarks and Network Appliance, ONTAPI, and Snapshot are trademarks of Network Appliance, Inc. in the U.S. and other countries. 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.