



An Oracle White Paper
April 2010

Working with Oracle[®] Solaris ZFS Snapshots

Introduction.....	1
Oracle Solaris ZFS Snapshots: Overview	2
Setting Up the File System	2
Taking a Snapshot.....	3
Rolling Back a Snapshot	3
Copying Individual Files From a Snapshot.....	4
Storing a Snapshot on Your System	5
Sending a Snapshot to Another System.....	5
For More Information	6

Introduction

This guide is intended to show a new user the capabilities of the snapshot feature of Oracle® Solaris ZFS. It describes the steps necessary to set up an Oracle Solaris ZFS file system, as well as how to create snapshots, use them for backup and restore purposes, and migrate them between systems. After reading this guide, the user will have a basic understanding of how snapshots can be integrated into system administration procedures.

Oracle Solaris ZFS Snapshots: Overview

An Oracle Solaris ZFS snapshot is a read-only copy of an Oracle Solaris ZFS file system or volume. Snapshots can be created almost instantly and initially consume no additional disk space within the pool. Snapshots are a valuable tool for system administrators needing to perform backups and other users needing to save the state of a file system at a particular point in time and possibly recreate it later on the same or another machine. It is also possible to extract individual files from a snapshot. These tasks can be performed with Oracle Solaris ZFS without the need for additional software. In this short guide, we take a look at the simple command syntax necessary to achieve these tasks.

Setting Up the File System

The following steps explain how to set up the file system.

1. First, we create a pool (which we call *pool*) and display it.

```
# zpool create -f pool c0d0s5
# zpool list
```

NAME	SIZE	USED	AVAIL	CAP	HEALTH	ALTROOT
pool	3.11G	75K	3.11G	0%	ONLINE	-

2. Then we create a file system (called *filesystem*) in our pool and confirm that we have done so.

```
# zfs create pool/filesystem
# zfs list
```

NAME	USED	AVAIL	REFER	MOUNTPOINT
pool	97.5K	3.06G	18K	/pool
pool/filesystem	18K	3.06G	18K	/pool/filesystem

3. Now to illustrate our example we fill the file system with some data.

```
# cd /platform
# du -h -s .
261M .
# find . -print | cpio -pd /pool/filesystem
536032 blocks
# zfs list
```

NAME	USED	AVAIL	REFER	MOUNTPOINT
pool	206M	2.86G	19K	/pool
pool/filesystem	206M	2.86G	206M	/pool/filesystem

We are now ready to start working with snapshots.

Taking a Snapshot

1. Snapshots are named with the syntax *pool/fs@something*, where *something* can be a fairly arbitrary name, but ideally one that means something to the creator.

```
# zfs snapshot pool/filesystem@thursday
```

2. The snapshot is then visible using the *zfs list* command.

```
# zfs list
# zfs list
NAME                                USED    AVAIL    REFER    MOUNTPOINT
pool                                262M    2.81G    19K      /pool
pool/filesystem                     262M    2.81G    262M    /pool/filesystem
pool/filesystem@thursday            0       -        262M    -
```

3. However, the snapshot does not appear as a file system when using the *df* command. The reason it is hidden from normal Oracle Solaris utilities, such as *ls*, *tar*, *cpio*, and others, is to prevent the snapshot from appearing in backups.

```
# df -h
Filesystem      SIZE    USED    AVAIL    CAP    MOUNTED ON
pool            3.1G    19K     2.8G    1%     /pool
pool/filesystem 3.1G    262M    2.8G    9%     /pool/filesystem
```

Rolling Back a Snapshot

1. Our snapshot can now be used as a recovery mechanism. First, we “accidentally” delete all the files in our file system. We see that the files have been removed and the size of the data reported for our file system has decreased appropriately.

```
# cd /pool/filesystem
# ls
i86hvm i86pc i86xpv
# rm -rf *
# ls
# df -h /pool/filesystem
Filesystem      SIZE    USED    AVAIL    CAP    MOUNTED ON
pool/filesystem 3.1G    18K     2.8G    1%     /pool/filesystem
```

- Rolling back the snapshot to restore all our missing files is trivial. We can see that the files have been returned and the space consumed again.

```
# zfs list
NAME                                USED    AVAIL    REFER    MOUNTPOINT
pool                                262M    2.81G    19K      /pool
pool/filesystem                     262M    2.81G    18K      /pool/filesystem
pool/filesystem@thursday            262M    -        262M    -

# zfs rollback pool/filesystem@thursday
# cd /pool/filesystem
# ls
i86hvm i86pc i86xp

# df -h /pool/filesystem
Filesystem      SIZE    USED    AVAIL    CAP    MOUNTED ON
pool/filesystem 3.1G    262M    2.8G     9%    /pool/filesystem
```

Copying Individual Files From a Snapshot

It is possible to copy individual files from a snapshot by changing into the hidden `.zfs` directory of the pool that has been snapped.

```
# cd /pool
# ls -la
total 8
drwxr-xr-x  3 root root   3  Sep 11 15:33 .
drwxr-xr-x 23 root root 512  Sep 11 15:30 ..
drwxr-xr-x  2 root root   2  Sep 11 17:23 filesystem
# cd filesystem
# ls -la
total 6
drwxr-xr-x  2 root root   2  Sep 11 17:23 .
drwxr-xr-x  3 root root   3  Sep 11 15:33 ..
# cd .zfs
# ls
snapshot
# cd snapshot
# ls
thursday
# cd thursday
# ls
i86hvm i86pc i86xp
```

Storing a Snapshot on Your System

Storing snapshots on your system is a good practice for short-term storage as long as the snapshots are recreated regularly as data changes or the Oracle Solaris operating system is upgraded. Consider using an enterprise backup solution to save important data for long-term storage.

In the following sequence of commands, we send the snapshot into a file and then compress it. It can then be retrieved from the file when required. This is also shown.

```
# zfs send pool/filesystem@thursday > /var/tmp/thursday.snap
# gzip -9 -v /var/tmp/thursday.snap
# zfs create pool/thursday
# gzip -d -c /var/tmp/thursday.snap.gz | zfs receive -F pool/thursday
```

Sending a Snapshot to Another System

You can send the snapshot to another system and install it there as a usable file system.

1. First, create a pool to receive the snapshot on the target system.

```
otherhost# zpool create -f otherpool c0d0s7
otherhost# zpool list
NAME          SIZE  USED  AVAIL  CAP  HEALTH  ALTROOT
otherpool    6.22G  75K   6.22G  0%   ONLINE  -
```

2. Then *send* the snapshot over the network and receive it into the pool using a combination of the *zfs send* and *zfs receive* commands and a network pipe.

```
# zfs send pool/filesystem@thursday | ssh otherhost "/usr/sbin/zfs
receive otherpool/myfs"
```

3. The received snapshot is then visible in the pool on the other host.

```
otherhost# zfs list
NAME          USED  AVAIL  REFER  MOUNTPOINT
otherpool          262M  5.87G  19K    /otherpool
otherpool/myfs    262M  5.87G  262M   /otherpool/myfs
otherpool/myfs@thursday  0      -      262M   -
```

For More Information

There is more to the use of Oracle Solaris ZFS snapshots that we have not covered in this brief treatment. More comprehensive coverage of Oracle Solaris ZFS snapshots and clones, a related concept, are covered in the *Solaris Oracle Solaris ZFS Administration Guide* at *Chapter 7: Working With Oracle Solaris ZFS Snapshots and Clones* at <http://docs.sun.com/app/docs/doc/819-5461>. The Oracle Solaris ZFS manual is at <http://docs.sun.com>. Other sources of information are collected in Table 1.

TABLE 1. WEB RESOURCES

The man pages for Oracle Solaris ZFS	http://docs.sun.com/app/docs/doc/819-2240/zfs-1m http://docs.sun.com/app/docs/doc/819-2240/zpool-1m
Oracle Solaris ZFS	http://www.oracle.com/us/products/servers-storage/storage/storage-software/031857.htm
Oracle Solaris ZFS Wiki	http://www.solarisinternals.com/wiki/index.php?title=Category:Oracle Solaris ZFS
OpenSolaris Oracle Solaris ZFS Community (The OpenSolaris Oracle Solaris ZFS manual is found here.)	http://www.opensolaris.org/os/community/zfs/
Opensolaris Mail Alias Archive	http://www.opensolaris.org/jive/forum.jspa?forumID=80
OpenSolaris advocacy group presentations	http://www.opensolaris.org/os/community/advocacy/os-presentations/
Search for Oracle Solaris ZFS	http://blogs.sun.com



Working with Oracle Solaris ZFS Snapshots
April 2010
Author: Dominic Kay
Contributing Authors: Paul Eggleton, Cindy Swearingen

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2010, Oracle and/or its affiliates. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0310

SOFTWARE. HARDWARE. COMPLETE.