



WORKING WITH **ZFS SNAPSHOTS**

> **Solaris™ 10** How-To Guides



About This Solaris™ How-To Guide

This guide is intended to show a new user the capabilities of the ZFS snapshots feature. It describes the steps necessary to set up a ZFS filesystem and the use of snapshots including how to create them, use them for backup and restore purposes, and how to 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.

Contributors: Reviewed by Paul Eggleton & Cindy Swearingen

Contents

ZFS Snapshots: Overview	Page 1
Setting up the Filesystem	Page 1
Taking and Using a Snapshot	Page 1
Restoring From a Snapshot	Page 2
Restoring Individual Files	Page 3
Moving the Snapshot to Another System	Page 3
Archiving and Restoring Snapshots	Page 4
For More Information	Page 5

ZFS Snapshots How-To Guide

ZFS Snapshots: Overview

A *zfs snapshot* is a read-only copy of a Solaris™ ZFS file system or volume. Snapshots can be created almost instantly and initially consume no additional disk space within the pool. They are a valuable tool both for system administrators needing to perform backups and other users who need to save the state of their file system at a particular point in time and possibly recreate it later on the same machine or some other. It is also possible to extract individual files from a snapshot. These tasks can be performed with ZFS without the need for any 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

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 file system) 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 and Using 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
```

- The snapshot is then visible using `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     -
```

- However, the snapshot does not appear as a file system using `df`:

```
# 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
```

The reason it is hidden from normal Solaris utilities such as `ls`, `tar`, `cpio`, and others is to prevent the snapshot from appearing in backups.

Restoring From a Snapshot

- Our snapshot can now be used as a recovery mechanism. First, we “accidentally” delete all the files in our file system:

```
# cd /pool/filesystem
# ls
i86hvm i86pc i86xpvc
# rm -rf *
# ls
# df -h /pool/filesystem
Filesystem      SIZE      USED      AVAIL     CAP    MOUNTED ON
pool/filesystem 3.1G      18K       2.8G      1%     /pool/filesystem
```

We see that the files have been removed and the size of the data reported for our file system has decreased appropriately.

- Rolling back the snapshot to restore all our missing files is trivial:

```
# 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 i86xpvc
# df -h /pool/filesystem
Filesystem      SIZE      USED      AVAIL     CAP    MOUNTED ON
pool/filesystem 3.1G      262M      2.8G      9%     /pool/filesystem
```

We can see that the files have been returned and the space consumed again.

Restoring Individual Files

1. 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  i86xpv
```

Moving the Snapshot to Another System

1. You can move the snapshot to another system and install it there as a usable filesystem. 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/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       -
```

Archiving and Restoring Snapshots

Another use of snapshots is to create archives for long-term storage elsewhere.

1. 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
```

For More Information

There is more to the use of ZFS snapshots that we have not covered in this brief treatment. More comprehensive coverage of both ZFS snapshots and clones, a related concept, are covered in the ZFS Administration Guide at Chapter 6: Working With ZFS Snapshots and Clones at <http://docs.sun.com/app/docs/doc/819-5461>. The Solaris ZFS manual is at <http://docs.sun.com>.

Other sources of information are collected together below:

Web Resources

The “man” pages for ZFS	http://docs.sun.com/app/docs/doc/819-2240/zfs-1m http://docs.sun.com/app/docs/doc/819-2240/zpool-1m
ZFS Learning Centre	http://www.sun.com/software/solaris/zfs_learning_center.jsp
OpenSolaris ZFS Community (The OpenSolaris ZFS manual is found here.)	http://www.opensolaris.org/os/community/zfs/
ZFS Wiki	http://www.solarisinternals.com/wiki/index.php?title=Category:ZFS
OpenSolaris advocacy group presentations	http://www.opensolaris.org/os/community/advocacy/os-presentations/
Opensolaris mail alias archive	http://www.opensolaris.org/jive/forum.jspa?forumID=80
Search for ZFS	http://blogs.sun.com

sun.com/solaris

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA Phone 1-650-960-1300 or 1-800-555-9SUN Web sun.com



© 2009 Sun Microsystems, Inc. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD appliances, licensed from the University of California.

Sun, Sun Microsystems, Solaris, ZFS, and the Sun logo are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. or its subsidiaries in the U.S. and other countries.

The OPEN LOOK and Sun's Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(j)(2)(6/87) and FAR 52.227-1887, or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a). DOCUMENTATION IS PROVIDED AS IS AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS HELD TO BE LEGALLY INVALID. SunWIN # 554555 Lit # SWWP14764-0 01/09