

Tru64 UNIX

Using StorageWorks HSG80 Controller-Based Cloning and Snapshotting

September 2003

This Best Practice describes how to use the controller-based cloning and snapshotting capability of HP StorageWorks products to safely create consistent point-in-time copies of Tru64 UNIX file systems.

Contents

Using StorageWorks HSG80 Controller-Based Cloning and Snapshotting

Is This Best Practice Right for You?	2
Before You Begin	2
Applying the Best Practice	4
Example 1: Safely Cloning a Single-LUN, Single-Volume AdvFS Domain Using HSG80 Commands (Tru64 UNIX Version 5.0 or Higher)	5
Example 2: Safely Cloning a Multi-LUN, Multivolume AdvFS Domain Managed by a Single Pair of HSG80 Controllers Using HSG80 Commands (Tru64 UNIX Version 5.0 or Higher)	8
Example 3: Safely Snapshotting a Single-LUN, UFS File System Using HSG80 Commands (Tru64 UNIX Version 4.0F or Version 4.0G)	12
Alternative Practices	15
Comments and Questions	16
Legal Notice	16

Using StorageWorks HSG80 Controller-Based Cloning and Snapshotting

This Best Practice describes how to use the HSG80 controller-based cloning and snapshotting capabilities of HP StorageWorks™ products to safely create consistent point-in-time copies of Tru64 UNIX file systems.

- Controller-based cloning creates a full duplicate data set on different physical storage media from the original data set.
- Controller-based snapshotting copies only the critical RAID metadata for a data set to other media and employs copy-on-write technology to preserve a point-in-time copy of the original data set.

Controller-based clones and snapshots can be mounted as new file systems. These new file systems can be used as the source of backup operations, either on the original host or on a different host. By mounting the new file systems on a different host, backups can be performed without introducing a performance drain on the source system during the backup operation. Controller-based clones and snapshots mounted on the original host or on a different host can also serve as the data source for a decision-support application.

This document is not a complete tutorial on controller-based cloning and snapshotting capabilities. It points out specific issues related to these capabilities in a Tru64 UNIX environment to help you avoid file system data and metadata corruption. For more information on HSG80 controller-based cloning and snapshotting, see the HSG80 documentation:

<http://h18004.www1.hp.com/products/storageworks/acs/g80tru64.html>

The scope of this Best Practice is limited to the following environments:

- UFS and/or AdvFS file systems.
- Logical Storage Manager (LSM) is not used on disks containing file system data to be copied.

See the StorageWorks products Web page for more information about StorageWorks products:

<http://h18006.www1.hp.com/storage/>

See the Tru64 UNIX Best Practices Web page for more information about Best Practices documentation:

http://tru64unix.compaq.com/docs/best_practices/

Is This Best Practice Right for You?

Not all Best Practices apply to all configurations, so you must be sure that it is appropriate for your system and circumstances. To use this Best Practice, you must meet the requirements described in the following table:

Requirement	Description
Operating System	Tru64 UNIX Version 4.0F or higher.
Local File Systems	UFS and/or AdvFS.
Do not use Logical Storage Manager	No LSM use on any disk containing a partition used by a UFS file system or an AdvFS domain being copied through controller-based cloning or snapshotting.
RAID hardware	HSG80-based RAID array controller. All logical unit numbers (LUNs) affected must be on line.
HSG80 firmware	Controller-based snapshotting is supported only in Version 8.5P, Version 8.5S, or higher of the HSG80 firmware. Controller-based cloning is supported in all versions of the HSG80 firmware.

If you do not meet the previous requirements, see *Alternative Practices* for information.

Before You Begin

Before you apply this Best Practice for Using StorageWorks Controller-Based Cloning and Snapshotting, you must understand some background information and perform some preliminary tasks.

To preserve data and file system metadata integrity in a clone or snapshot, the clone or snapshot must be created atomically for all LUNs that constitute the AdvFS domain or UFS file system. The HSG80-based RAID array controller supports the atomic creation of clones and snapshots for a single LUN managed by the HSG80 controller. Therefore, this Best Practice describes two kinds of controller-based clone and snapshot scenarios:

- Online scenarios in which an AdvFS domain or a UFS file system to be cloned or snapshot uses only one HSG80 LUN. In these scenarios, the domain or file system being cloned or snapshot does not have to be unmounted before the clone and/or snapshot operation.
- Offline scenarios in which an AdvFS domain to be cloned or snapshot spans more than one LUN. In these scenarios, the domain being cloned or snapshot must be unmounted prior to the clone and/or snapshot operation.

Some AdvFS and UFS commands work on unmounted file systems or domains. Be careful not to run any of the commands that work on unmounted file systems or domains while any controller-based cloning or snapshotting activities are occurring. The file system commands that modify file system metadata on unmounted file systems include, but might not be limited to, the following:

- `chfsets` (AdvFS)
- `extendfs` (UFS)
- `fsck` (UFS)

Caution

Data or file system metadata corruption can occur if this Best Practice is not followed. The corruption might or might not be evident. In any case, the result of such corruption might include one or more of the following:

- Unmountable file system
 - User data corruption
 - AdvFS domain panic
 - Tru64 UNIX kernel panic
-

This Best Practice describes how to safely use the controller-based cloning and snapshotting capabilities to create crash-consistent data sets. That is, the data in the clone or snapshot will be as consistent — from a Tru64 UNIX perspective — as if the system had crashed. The use of this Best Practice does not guarantee data consistency from an application perspective; controller-based cloning and snapshotting works only at the storage level. Application data that is cached in the application's memory or in the Tru64 UNIX buffer cache will not be captured in the controller-based

clone or snapshot. Therefore, before applying this Best Practice, you must take steps to flush and quiesce application data. For example, if you are using this Best Practice to create a controller-based snapshot or clone for your Oracle database, you might want to place the Oracle tablespace into online backup mode using the Oracle SQL command before creating the controller-based snapshot or clone:

```
ALTER TABLESPACE name BEGIN BACKUP
```

Then, once the controller-based snapshot or clone has been created, you can enter the Oracle SQL command to take the tablespace out of online backup mode:

```
ALTER TABLESPACE name END BACKUP
```

If the applications using the file systems that are to be the subject of a controller-based cloning or snapshotting operation do not have a quiesce capability, Tru64 UNIX provides several ways to force any pending updates stored in the Tru64 UNIX buffer cache to be flushed to storage:

- An application can issue an `fsync()` system call to flush the updates to a particular file.
- An application can open a file with the `O_SYNC` or `O_DSYNC` flag. This will cause all updates to the file to be synchronously written to storage.
- A file system can be mounted with the `-o sync` option. This will force updates to files in the file system to be synchronously written to storage. Use of this option can cause a significant loss in performance to the applications using that file system. It is preferable to have applications use one of the preceding methods to flush only the updates that need to be flushed to storage.
- The `chfile -l on` command can be used on a file residing in an AdvFS file system to cause updates to that file to be synchronously flushed to storage, regardless of the I/O mode requested by the applications using the file.

Applying the Best Practice

Before you use controller-based cloning and snapshotting with StorageWorks, be sure to follow the recommendations in *Before You Begin*.

The following examples show how to safely use controller-based cloning and snapshotting with Tru64 UNIX:

- *Example 1: Safely Cloning a Single-LUN, Single-Volume AdvFS Domain Using HSG80 Commands (Tru64 UNIX Version 5.0 or Higher).* Because only one LUN is involved, this is an online example.
- *Example 2: Safely Cloning a Multi-LUN, Multivolume AdvFS Domain Managed by a Single Pair of HSG80 Controllers Using HSG80 Commands (Tru64 UNIX Version 5.0 or Higher).* Because more than one LUN is involved, this is an offline example.
- *Example 3: Safely Snapshotting a Single-LUN, UFS File System Using HSG80 Commands (Tru64 UNIX Version 4.0F or Version 4.0G).* Because only one LUN is involved, this is an online example.

Note: These examples work on one AdvFS domain or UFS file system. If you need to clone or snapshot more than one domain or file system, repeat the recommended practice for other domains or file systems.

Example 1: Safely Cloning a Single-LUN, Single-Volume AdvFS Domain Using HSG80 Commands (Tru64 UNIX Version 5.0 or Higher)

The configuration for this example:

- Tru64 UNIX Version 5.0 or higher.
- AdvFS domain to be cloned, `oracle_dmn`, has one fileset, `oracle`, and uses storage on one volume, `/dev/disk/dsk23c`.
- `/dev/disk/dsk23c` is a single LUN, known to the HSG80 as `D1`, consisting of a striped mirrorset, `STRIPEM1`.
- `oracle_dmn#oracle` is mounted on `/oracle` and stays mounted there for the duration of the example.
- `STRIPEM1` is a 3-way stripeset consisting of three 3-way mirrorsets, `MIRR1`, `MIRR2`, and `MIRR3`.
- `MIRR1` is made up of disks `DISK10000`, `DISK10100`, and `DISK10200`.
- `MIRR2` is made up of disks `DISK20000`, `DISK20100`, and `DISK20200`.
- `MIRR3` is made up of disks `DISK30000`, `DISK30100`, and `DISK30200`.

The procedure for this example is as follows:

1. From the HSG80 console:

- a. Atomically remove a disk from each mirror making up the striped mirrorset:

```
REDUCE DISK10200 DISK20200 DISK30200
```

Caution

- It is critical that all disks to be used in the clone are removed from the original striped mirrorset by using one REDUCE command. Using multiple REDUCE commands to remove the disks can cause data or file system metadata corruption in the clone LUN.
 - Because the example involves only one LUN, there is no need to unmount the filesets in the domain being cloned.
-

b. Create temporary mirrorsets using the removed disks:

```
ADD MIRRORSET TEMPM1 DISK10200
ADD MIRRORSET TEMPM2 DISK20200
ADD MIRRORSET TEMPM3 DISK30200
```

c. Create a stripeset from the three temporary mirrorsets:

```
ADD STRIPESET CLONE1 TEMPM1 TEMPM2 TEMPM3
```

d. Initialize the clone stripeset without destroying its contents:

```
INIT CLONE1 NODESTROY
```

e. Create a LUN for the clone stripeset:

- If the controller is connected to only one host or cluster, use:

```
ADD UNIT D2 CLONE1
```

- If the controller is shared by multiple hosts or clusters, use:

```
ADD UNIT D2 CLONE1 DISABLE_ACCESS_PATH=ALL
```

f. For easier identification at the Tru64 UNIX level, set an identifier on the new LUN:

```
SET D2 IDENTIFIER=2
```

Note: If cloning or snapshotting is done often, it is helpful to establish a naming convention for these identifiers.

- g. If the controller is shared by multiple hosts or clusters, you must enable access to the newly created LUN.

- i. Find the connections that should be granted access to this LUN:

```
SHOW CONNECTIONS
```

- ii. Enable access to the LUN for the desired connections:

```
SET D2 ENABLE_ACCESS_PATH=connectionname...
```

Refer to the Hardware Configuration Technical Update for Fibre Channel for more information:

http://www.tru64unix.compaq.com/docs/updates/TCR51_FC/TITLE.HTM

2. From Tru64 UNIX:

- a. Create device special files for the new LUN:

- On a standalone system, use:

```
/sbin/hwmgr -scan scsi
```

- On a cluster, use:

```
/sbin/hwmgr scan component -category scsi_bus  
-cluster
```

This command is asynchronous, so it could complete before the device special files are created.

Note: In Tru64 UNIX Version 5.0 or higher, every time a new LUN is created, a new set of device special files is created. To see a list of device special files that are available for use:

```
/sbin/hwmgr -view devices
```

- b. Assuming that `/dev/disk/dsk24c` was created in step 2a, create the directory and symbolic link necessary to access the AdvFS domain on `/dev/disk/dsk24c`:

```
mkdir /etc/fdmns/clone_domain  
ln -fs /dev/disk/dsk24c /etc/fdmns/clone_domain
```

Caution

Do not use `/sbin/mkfdmn` on `/dev/disk/dsk24c`.
Doing so will cause the cloned data to become
inaccessible.

c. Mount the `oracle` fileset in `clone_domain`:

```
mount [-o dual] clone_domain#oracle /backup
```

- If the clone domain and fileset are being mounted on the same host as the original, you must use the `-o dual` option.
- If the domain and fileset are being mounted on a different host, the `-o dual` option is not necessary.

The `/backup` mount point can now be used to access the clone copy of the data.

Example 2: Safely Cloning a Multi-LUN, Multivolume AdvFS Domain Managed by a Single Pair of HSG80 Controllers Using HSG80 Commands (Tru64 UNIX Version 5.0 or Higher)

The configuration for this example:

- Tru64 UNIX Version 5.0 or higher.
- AdvFS domain to be cloned, `oracle_dmn`, has one fileset, `oracle`, and uses storage on two volumes, `/dev/disk/dsk23c` and `/dev/disk/dsk24c`.
- `/dev/disk/dsk23c` is a single LUN, known to the HSG80 as D1, consisting of a striped mirrorset, `STRIPEM1`.
- `/dev/disk/dsk24c` is a single LUN, known to the HSG80 as D2, consisting of a striped mirrorset, `STRIPEM2`.
- Because more than one LUN makes up the domain, `oracle_dmn#oracle` must be unmounted until step 1a has been completed.
- Alternatively, for Tru64 UNIX Version 5.1B or higher, the filesystem can be "frozen" by using the `freezefs` command. This command prevents modifications to the filesystem until the filesystem is "thawed" using the `thawfs` command (see the following caution for more information). This allows the filesystem to remain on line and accessible (for example, for reading), but prevents modifications that would ordinarily render the clone invalid.

The basic method to freeze the filesystem is to use the following command:

```
freezefs -t -1 <mount-point>
```

The `-t -1` specifies that no timeout will be followed; that is, the filesystem will have to be explicitly thawed.

After cloning, the filesystem is thawed using the following command:

```
thawfs <mount-point>
```

Caution

Freezing a filesystem prevents metadata updates, but does not always prevent data updates. For example, on a single-node (nonclustered) system, freezing an AdvFS filesystem will prevent writes that would require additional storage to be allocated to the file (for example, appending to the file), but will *not* prevent writes on storage that has already been allocated to a file (for example, modifying an already existing page). It is therefore possible for the HSG80 clone to be consistent as far as filesystem metadata are concerned, but nevertheless to be inconsistent as far as application data are concerned. If the application can be placed in a quiescent state (for example, see the previous discussion of the Oracle database), this danger can be avoided, but if the application cannot be placed in such a state, it is safer to take the filesystem off line by unmounting it.

On a cluster, CFS filesystems are not subject to this danger: freezing the filesystem will prevent *all* writes, not just metadata writes. However, there is a different danger: certain events will thaw a filesystem automatically, even if the timeout has been explicitly disabled. For example, node failure or explicit shutdown of *any* node will cause an automatic thaw of *all* frozen filesystems . Therefore, it is important to check the return value of the `thawfs` command: if the return value indicates that the filesystem was not frozen, an automatic thaw must have taken place. In such a case, assume that the clone is invalid and try the operation again. Checking return values should be done in all cases, but it is particularly important here.

There are also restrictions on what filesystems can be frozen, for example, the root, the `/usr`, and the `/var` filesystems cannot be frozen. For more information, see the `freezefs(8)` and `thawfs(8)` reference pages.

- STRIPEM1 is a three-way stripeset consisting of three 3-way mirrorsets, MIRR1, MIRR2, and MIRR3.

- STRIPEM2 is a three-way stripeset consisting of three 3-way mirrorsets, MIRR4, MIRR5, and MIRR6.
- MIRR1 is made up of disks DISK10000, DISK10100, and DISK10200.
- MIRR2 is made up of disks DISK20000, DISK20100, and DISK20200.
- MIRR3 is made up of disks DISK30000, DISK30100, and DISK30200.
- MIRR4 is made up of disks DISK40000, DISK40100, and DISK40200.
- MIRR5 is made up of disks DISK50000, DISK50100, and DISK50200.
- MIRR6 is made up of disks DISK60000, DISK60100, and DISK60200.

The procedure for this example is as follows:

1. From the HSG80 console:

- a. Remove a disk from each mirror making up the two striped mirrorsets:

```
REDUCE DISK10200 DISK20200 DISK30200
REDUCE DISK40200 DISK50200 DISK60200
```

Caution

- Because the HSG80 REDUCE command can act on only one LUN at a time and because there are multiple LUNs making up the domain, multiple REDUCE commands must be used. Therefore, it is critical that the domain be unmounted before issuing the REDUCE commands.
 - Without LSM, it is not possible to have a UFS file system with its storage on multiple LUNs, so this multi-LUN example is specific to AdvFS.
-

- b. Create temporary mirrorsets using the removed disks:

```
ADD MIRRORSET TEMPM1 DISK10200
ADD MIRRORSET TEMPM2 DISK20200
ADD MIRRORSET TEMPM3 DISK30200
ADD MIRRORSET TEMPM4 DISK40200
ADD MIRRORSET TEMPM5 DISK50200
ADD MIRRORSET TEMPM6 DISK60200
```

- c. Create two clone stripesets from the six temporary mirrorsets:

```
ADD STRIPESET CLONE1 TEMPM1 TEMPM2 TEMPM3
ADD STRIPESET CLONE2 TEMPM4 TEMPM5 TEMPM6
```

- d. Initialize the clone stripesets without destroying their contents:

```
INIT CLONE1 NODESTROY
INIT CLONE2 NODESTROY
```

- e. Create LUNs for the clone stripeset.

- If the controller is connected to only one host or cluster, use:

```
ADD UNIT D3 CLONE1
ADD UNIT D4 CLONE2
```

- If the controller is shared by multiple hosts or clusters, use:

```
ADD UNIT D3 CLONE1 DISABLE_ACCESS=ALL
ADD UNIT D4 CLONE2 DISABLE_ACCESS=ALL
```

- f. For ease of identification at the Tru64 UNIX level, set an identifier on each of the new LUNs:

```
SET D3 IDENTIFIER=3
SET D4 IDENTIFIER=4
```

Note: If cloning or snapshotting is done often, it is helpful to establish a naming convention for these identifiers.

- g. If the controller is shared by multiple hosts or clusters, further work must be done to enable access to the newly created LUNs.

- i. Find the connections that should be granted access to this LUN:

```
SHOW CONNECTIONS
```

- ii. Enable access to the LUNs for the desired connections:

```
SET D3 ENABLE_ACCESS=connectionname...
SET D4 ENABLE_ACCESS=connectionname...
```

See the Hardware Configuration Technical Update for Fibre Channel for more information:

http://www.tru64unix.compaq.com/docs/updates/TCR51_FC/TITLE.HTM

2. From Tru64 UNIX:

- a. Create device special files for the new LUNs.

- On a standalone system, use:

```
/sbin/hwmgr -scan scsi
```

- On a cluster, use:

```
/sbin/hwmgr scan component -category scsi_bus
-cluster
```

This command is asynchronous, so it could complete before the device special files are created.

Note: In Tru64 UNIX Version 5.0 or higher, every time a new LUN is created, a new set of device special files is created. To see a list of device special files that are available for use:

```
/sbin/hwmgr -view devices
```

- b. Assuming that `/dev/disk/dsk25c` and `/dev/disk/dsk26c` were created in step 2a, create the directory and symbolic links necessary to access the AdvFS domain on `/dev/disk/dsk25c` and `/dev/disk/dsk26c`.

```
ln -fs /dev/disk/dsk25c /etc/fdmns/clone_domain
ln -fs /dev/disk/dsk26c /etc/fdmns/clone_domain
```

Caution

Do not use `/sbin/mkfdmn` on `/dev/disk/dsk25c` or `/dev/disk/dsk26c`. Doing so will cause the cloned data to become inaccessible.

- c. Mount the oracle fileset in `clone_domain`:

```
mount [-o dual] clone_domain#oracle /backup
```

- If the clone domain and fileset are being mounted on the same host as the original, you must use the `-o dual` option.
- If the domain and fileset are being mounted on a different host, the `-o dual` option is not necessary.

The `/backup` mount point can now be used to access the clone copy of the data.

Example 3: Safely Snapshotting a Single-LUN, UFS File System Using HSG80 Commands (Tru64 UNIX Version 4.0F or Version 4.0G)

The configuration for this example:

- Tru64 UNIX Version 4.0F or Version 4.0G.
- UFS file system to be snapshotted has storage on `/dev/rz23c`.

- /dev/rz23c is a single LUN, known to the HSG80 as D1, consisting of a striped mirrorset, STRIPEM1.
- /dev/rz23c is mounted on /oracle and stays mounted there for the duration of the example.
- STRIPEM1 is a 3-way stripeset consisting of three 3-way mirrorsets, MIRR1, MIRR2, and MIRR3.
- MIRR1 is made up of disks DISK10000, DISK10100, and DISK10200.
- MIRR2 is made up of disks DISK20000, DISK20100, and DISK20200.
- MIRR3 is made up of disks DISK30000, DISK30100, and DISK30200.
- SNAP1 is a 3-way stripeset consisting of three 3-way mirrorsets, MIRR4, MIRR5, and MIRR6.
- MIRR4 is made up of disks DISK40000, DISK40100, and DISK40200.
- MIRR5 is made up of disks DISK50000, DISK50100, and DISK50200.
- MIRR6 is made up of disks DISK60000, DISK60100, and DISK60200.

The procedure for this example:

1. From the HSG80 console:
 - a. If you have not set up preferred paths for the LUN to be snapshotted, you must configure a preferred path before adding the snapshot unit. For example:


```
SET D1 PREFERRED_PATH=THIS
```
 - b. Create a new LUN, D2, which is a snapshot of D1:


```
ADD SNAPSHOT_UNITS D2 SNAP1 D1
```

Note that you can use the `USE_PARENT_WWID` option, which allows the Tru64 UNIX host to assign the same `/dev/disk/diskxx` to each instance of a snapshot of the same parent.
 - c. For easier identification at the Tru64 UNIX level, set an identifier on the new LUN:


```
SET D2 IDENTIFIER=2
```

Note: If cloning or snapshotting is done often, it is helpful to establish a naming convention for these identifiers.
 - d. Access to snapshot units is initially disabled. You must re-enable access before the snapshot LUN can be used:

- If the controller is connected to only one host or cluster, use:
- If the controller is shared by multiple hosts or clusters, you must enable access to the newly created LUN:

```
SET D2 ENABLE_ACCESS_PATH=ALL
```

- Find the connections that should be granted access to this LUN:

```
SHOW CONNECTIONS
```

- Enable access to the snapshot LUN for the desired connections:

```
SET D2 ENABLE_ACCESS_PATH=connectionname...
```

Refer to the Hardware Configuration Technical Update for Fibre Channel for more information:

http://www.tru64unix.compaq.com/docs/updates/TCR51_FC/TITLE.HTM

2. From Tru64 UNIX:

- Create device special files for the new LUN:

```
/sbin/scsimgr -scan_all
```

- Assuming that `/dev/rz24c` was created in step 2a, run `fsck` on it:

```
/sbin/fsck /dev/rz24c
```

- Mount the snapshotted UFS file system:

```
mount /dev/rz24c /backup
```

The `/backup` mount point can now be used to access the snapshotted copy of the data.

Caution

- It is not possible to atomically create snapshots of multiple LUNs with HSG80 commands. Therefore, it is not possible to safely create a multi-LUN snapshot of a multivolume AdvFS domain using the HSG80 commands, unless all filesets in the domain are first unmounted. Entering multiple `ADD SNAPSHOT_UNITS` commands to create snapshots of multiple LUNs that make up a mounted multivolume AdvFS domain might cause data or file system metadata corruption.

- Only HSG80 firmware Version 8.5P, Version 8.5S, or higher support snapshotting.
 - Because this example involves only one LUN, there is no need to unmount the file system being snapshotted.
-

Using Snapshots

Consider the following scenario to safely use snapshots. First, create a snapshot and then take a backup of that snapshot. The next day you delete the old snapshot, create a new snapshot, back that one up, and so on. There are two safe ways to do this:

- Method 1: After you enter the HSG80 controller command to delete the old snapshot, enter `hwmgr -delete` on the host to delete the host's old view of the old snapshot. Omitting this step may cause unexpected behavior on the new snapshot. If you use this method, your new snapshot will be given a new name (`/dev/disk/dskxx`).
 - Method 2: After you enter the HSG80 controller command to delete the old snapshot, enter a host command to try to access the deleted unit. For example, enter `file /dev/rdisk/diskxx` or `disklabel -r /dev/rdisk/diskxx`. Either command will fail. You can then proceed to create and use a new snapshot. Omitting the host command to attempt access to the old name can result in unexpected behavior on the new snapshot. If you use this method, each snapshot may get a new name or reuse an existing name if the WWID matches by using the `USE_PARENT_WWID` option on the controller command.
-

Alternative Practices

Depending on the solution being sought, an alternative to using controller-based cloning and snapshotting is AdvFS clone filesets, which require the AdvFS Utilities license. At a logical level, AdvFS fileset cloning

works much like controller-based snapshotting. When a fileset clone is created, critical file system metadata is copied. The point-in-time view of the fileset is maintained using copy-on-write technology.

AdvFS clone filesets can be a good alternative to controller-based cloning and snapshotting, if your configuration and needs meet the following requirements:

- An application's data does not span AdvFS filesets. AdvFS cloning guarantees atomicity at the fileset level only.
- The clone fileset will be mounted on the same host as the original fileset.
- Only one clone fileset is needed at a time.
- The cloned data will only be read. AdvFS clone filesets are read only.

AdvFS clone filesets do not require additional hardware, only available storage within the AdvFS domain. For more information on AdvFS clone filesets, see the AdvFS Administration guide for your version of the operating system:

http://tru64unix.compaq.com/docs/pub_page/doc_list.html

Comments and Questions

We value your comments and questions on the information in this document. Please mail your comments to us at this address:

best_practices@zk3.dec.com

Legal Notice

UNIX is a trademark of The Open Group in the United States and other countries.

All other product names mentioned herein may be trademarks or registered trademarks of their respective companies.

Proprietary computer software. Valid license from HP and/or its subsidiaries required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Neither HP nor any of its subsidiaries shall be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for HP products are set forth in the express limited

warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.