Tru64 UNIX Best Practice

Backing Up LSM Volumes Using the Fast Plex Attach Feature

April 2003

Product Version: **Tru64 UNIX Version 5.1B and higher**

This Best Practice describes how to use the Fast Plex Attach feature of the Logical Storage Manager (LSM) to back up LSM volumes on the HP Tru64 UNIX operating system.

Hewlett-Packard Company
Palo Alto, California

Contents

Backing Up LSM Volumes Using the Fast Plex Attach Feature

Backing Up LSM Volumes Using the Fast Plex Attach Feature

This Best Practice describes how to use the Fast Plex Attach (FPA) feature of the Logical Storage Manager (LSM) to back up volumes.

FPA uses an existing complete plex of a mirrored LSM volume to create a separate volume for performing backups. This leaves the original volume available for use while the backup is occurring. When the backup is complete, you can return the "borrowed" plex to the original volume.

FPA reduces the time required to resynchronize the plex with the original volume through the use of dedicated FPA logs on both the original (primary) volume and the temporary (secondary) volume. The FPA logs track the changes to both volumes. When the plex returns to the original volume, LSM merges the two FPA logs and resynchronizes the returning plex to ensure that:

- Any changes that occurred to the primary volume are copied to the corresponding regions in the returning plex, bringing it up to date with the primary volume.
- Any changes that occurred to the secondary volume are overwritten by recopying the data from the corresponding regions in the primary volume to those regions in the returning plex.

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, your configuration must meet the requirements described in the following table:

Requirement	Description
Operating System	Tru64 UNIX Version 5.1B or higher on a standalone system or in a TruCluster Server environment

Requirement	Description
System Configuration	LSM must be installed and initialized. This Best Practice applies only to mirrored LSM volumes. The LSM disk group of the volume you want to back up must have enough free space to create two FPA logs (65 blocks/GB of volume size), and, if necessary, to add an additional data plex to the volume.
Impact on Availability	None. No reboot is required, no applications need to stop, and no file systems need to be unmounted.
Additional Requirements	You must be the superuser (root) or have root permissions.

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

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

Before You Begin

Before you apply the Best Practice for using the Fast Plex Attach (FPA) feature, you must understand some background information and, if necessary, perform some preliminary tasks.

FPA works only on mirrored LSM volumes with at least two complete, enabled, read-write plexes. LSM detaches one plex from the specified (primary) volume to create a temporary (secondary) volume that you can use to do your backups. Because LSM temporarily takes one plex away from the primary volume, the volume should have at least two plexes remaining to maintain data availability.

If the original volume does not have at least two suitable data plexes, you can do either one of the following:

- If the disk group has sufficient free space (equal to the size of the volume), add another data plex to the volume before using FPA.
- Use FPA with the force option, which leaves the primary volume with only one data plex (therefore, unmirrored). This introduces the risk of data loss if a disk in the remaining plex fails.

LSM uses available space in the disk group to create the FPA logs for both the primary and secondary volumes. Each FPA log needs approximately 65 blocks of disk space per gigabyte of volume size. LSM creates the FPA logs of the appropriate size for the volumes.

The free space can be on any LSM disks in the disk group including disks marked as hot-spares, if no other space is available.

Note

If LSM must use hot-spare disks to create the FPA logs, that space will no longer be available to perform hot-sparing operations.

Applying the Best Practice

Before you use the Fast Plex Attach feature to create a backup volume, be sure to follow the recommendations in *Before You Begin*.

To use FPA to create a secondary volume from an existing plex of a mirrored LSM volume:

1. Identify the volume you want to back up:

   # **volprint -g *disk_group* -vt**

2. Display detailed information about the volume to determine whether it meets the requirements specified in *Before You Begin*:

   # **volprint -g *disk_group* -ht *volume***

3. If necessary, add another data plex to the volume:

   a. Display the available disk space in the disk group containing the volume:

      # **voldg -g *disk_group* free**

      There must be free space equal to the size of the volume. If not, add one or more disks to the disk group.

   b. Add a data plex to the volume. In this example, the `snapstart` keyword is used expressly to create a temporary data plex for backups. The FPA feature will automatically use the new plex.

      # **volassist -g *disk_group* snapstart *volume***

4. Create the secondary volume, assigning it a name of your choice (which must be unique in the disk group) and add FPA support to both the primary and secondary volumes:

   ```
   # volassist -g disk_group snapfast \
   primary_volume secondary_volume
   ```

   For example:

   ```
   # volassist -g dg1 snapfast vol-01 vol-01BK
   ```

   If you added a new plex to the volume with `volassist snapstart`, LSM uses that plex to create the secondary volume. Otherwise, LSM uses any qualifying data plex.

5. Back up the secondary volume using your accustomed backup method or procedure.

6. When the backup is complete, run the following command to merge the FPA logs and return and resynchronize the migrant plex to the primary volume:

   ```
   # volassist -g disk_group snapback \
   secondary_volume primary_volume
   ```

   For example:

   ```
   # volassist -g dg1 snapback vol-01BK vol-01
   ```

   If the secondary volume has no other plexes, it is removed as part of this operation.

After the migrant plex is resynchronized to the primary volume, the FPA log on the primary volume is cleared (zeroed out) and remains attached to the volume for future use.

Example of FPA in Action

This section illustrates the process of using the FPA feature on a volume.

Volume Before Using FPA Feature

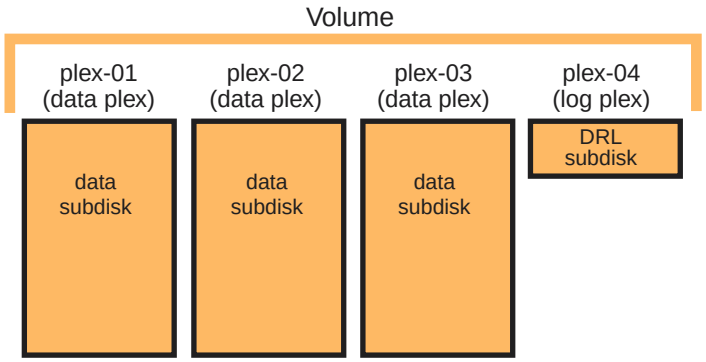

ZK-1875U-AI

The `volassist snapfast` command:

- Creates an FPA log subdisk and an FPA log plex and associates it with the primary volume as shown in the following figure:

Creation of Primary Volume's FPA Subdisk and Plex

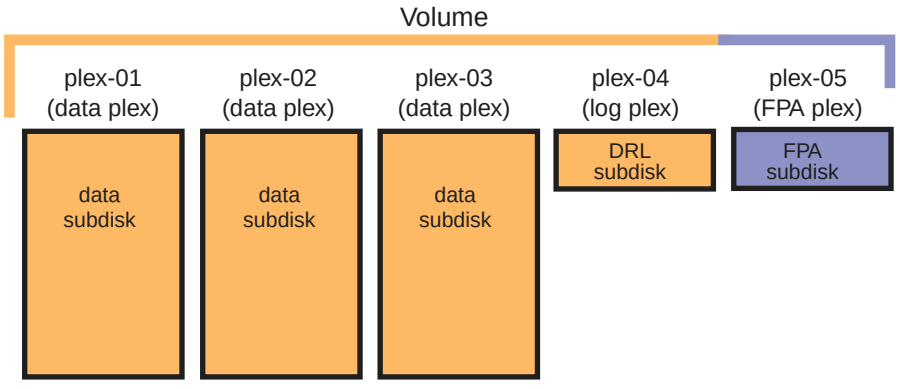

ZK-1876U-AI

- Selects a complete data plex to create the secondary volume. (If you added a new plex with the `volassist snapstart` command, LSM uses this plex for the secondary volume.)
- Creates another FPA log subdisk and associates it with the selected plex, now called the migrant plex, detaches the migrant plex and its

FPA log subdisk from the primary volume, and creates a secondary volume as shown in the following figure:

Creation of Secondary Volume

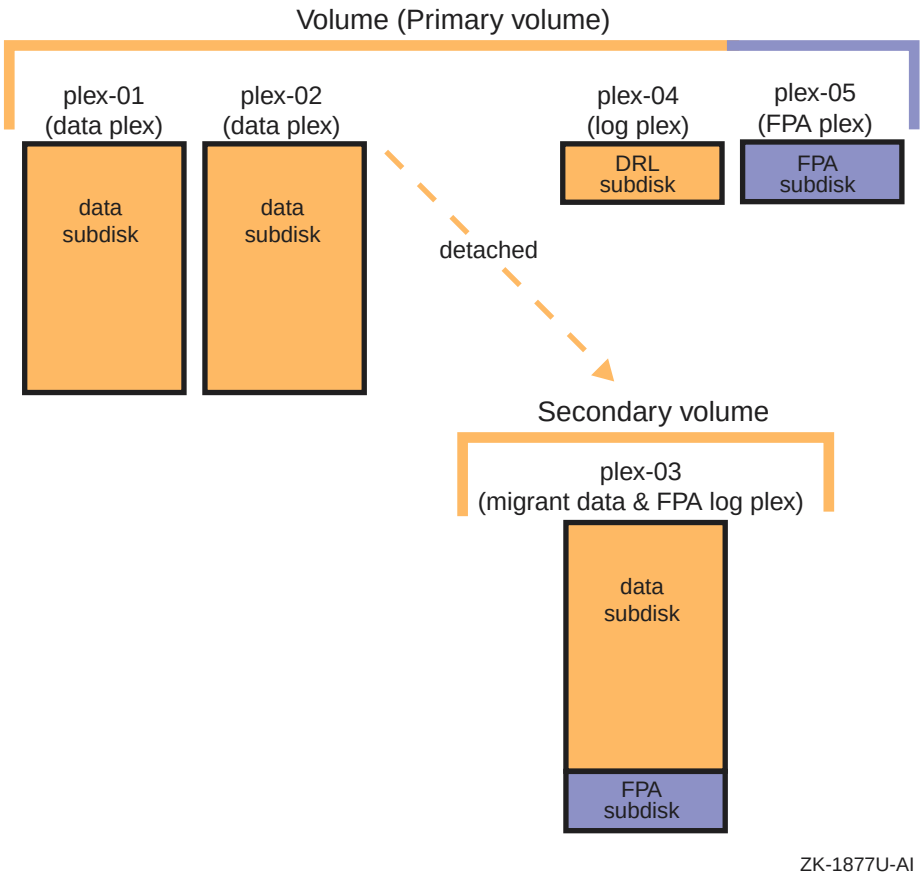

As I/O occurs to both the primary and secondary volumes, any changes are recorded in their respective FPA logs as shown in the following figure:

Writes Being Recorded in Both FPA Logs

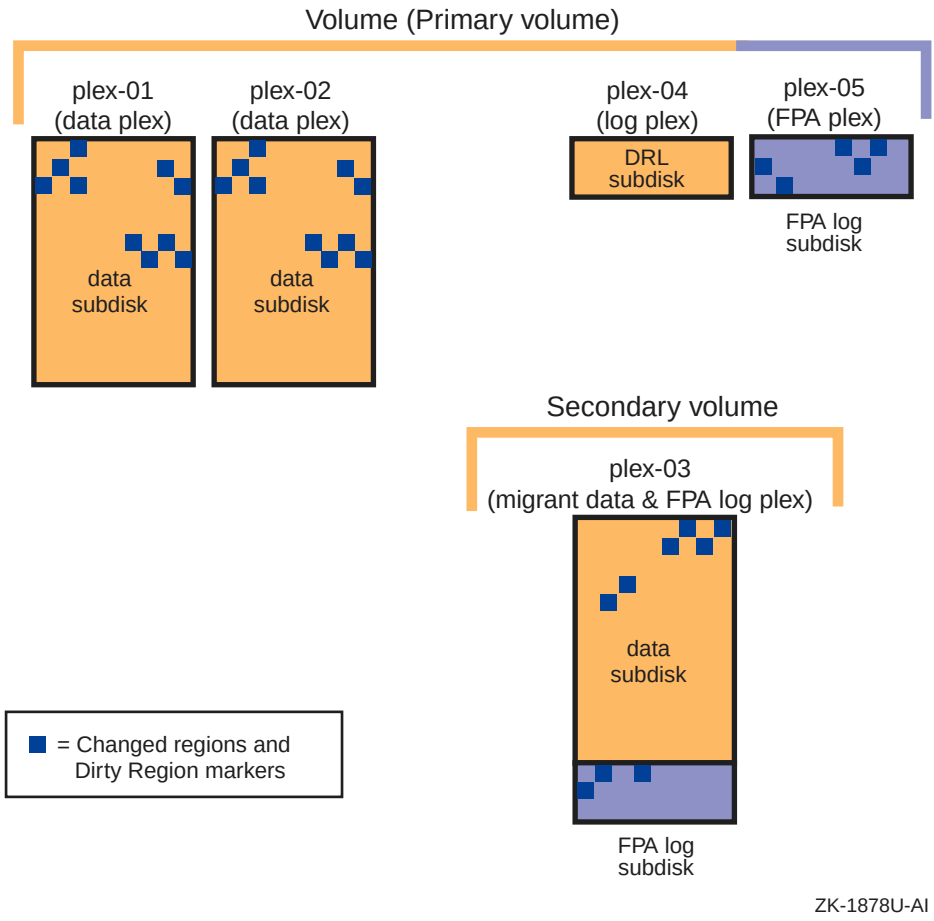

When you use the `volassist snapback` command to return the migrant plex to the primary volume, LSM merges the FPA logs of both volumes as shown in the following figure:

Merging of FPA Logs

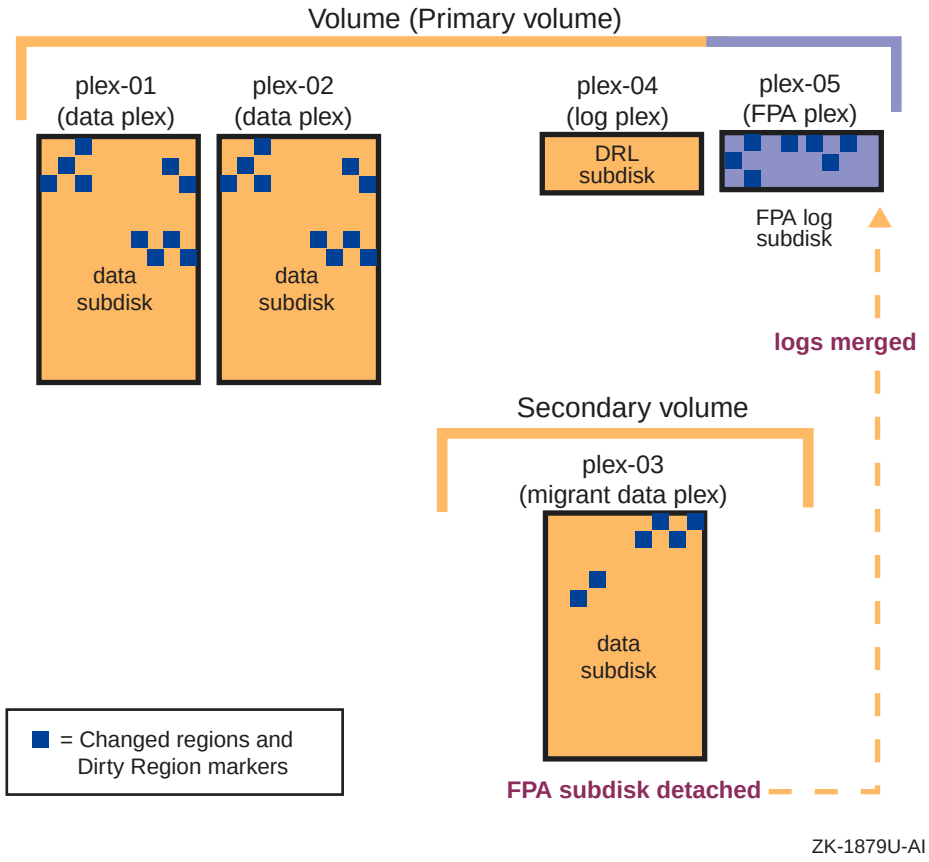

LSM then detaches the migrant plex from the secondary volume and reattaches it to the primary volume, using the information in the merged logs to resynchronize only the regions that differ between the migrant plex and the primary volume as shown in the following figure:

Resynchronization of Migrant Plex to Primary Volume

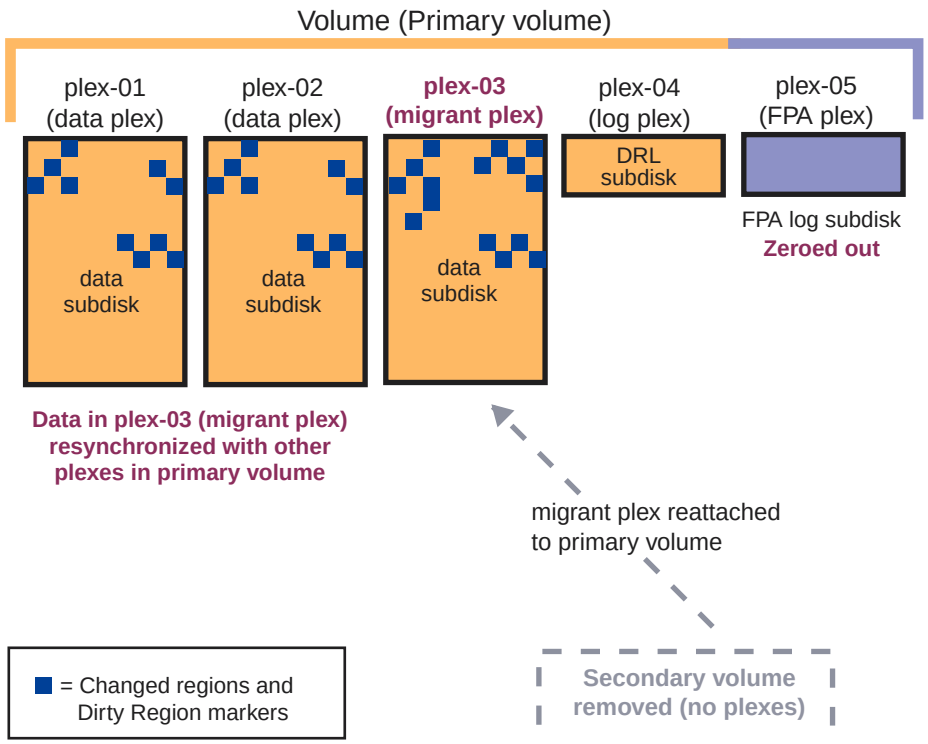

ZK-1880U-AI

Verifying Success

Most of the steps in this Best Practice are self-verifying, or verification is unnecessary. If problems occur, LSM displays error messages about the problem and, usually, indicates the steps needed to correct it.

See *Troubleshooting* for information about identifying and solving problems.

You can use the LSM `volprint` command at any time to display the properties of the primary and secondary LSM volumes.

For example, the following command displays the properties of the primary volume `vol-01` and the secondary volume `vol-01BK` in the `rootdg` disk group:

```
# volprint -h vol-01 vol-01BK
Disk group: rootdg

V  NAME       USETYPE    KSTATE   STATE    LENGTH   READPOL   PREFPLEX
PL NAME       VOLUME     KSTATE   STATE    LENGTH   LAYOUT    NCOL/WID MODE
SD NAME       PLEX       DISK     DISKOFFS LENGTH   [COL/]OFF DEVICE   MODE

v  vol-01     fsgen      ENABLED  ACTIVE   409600   SELECT    -
pl plex-01    vol-01     ENABLED  ACTIVE   409600   CONCAT    -        RW
sd dsk0-01    plex-01    dsk0     0        409600   0         dsk0     ENA
pl plex-02    vol-01     ENABLED  ACTIVE   409600   CONCAT    -        RW
sd dsk2-01    plex-02    dsk2     0        409600   0         dsk2     ENA
pl plex-04    vol-01     ENABLED  ACTIVE   LOGONLY  CONCAT    -        RW
sd dsk4-01    plex-04    dsk4     0        65       LOG       dsk4     ENA
pl plex-05    vol-01     ENABLED  ACTIVE   FPAONLY  CONCAT    -        RW    1
sd dsk0-02    plex-05    dsk0     409600   65       FPA       dsk0     ENA

v  vol-01BK   fsgen      ENABLED  ACTIVE   409600   ROUND     -
pl plex-03    vol-01BK   ENABLED  ACTIVE   409600   CONCAT    -        RW    2
sd dsk0-03    plex-03    dsk0     409665   65       FPA       dsk0     ENA   3
sd dsk3-01    plex-03    dsk3     0        409600   0         dsk3     ENA   4
```

1 Plex `plex-05` is the FPA log plex for volume `vol-01`.

2 Plex `plex-03` is the migrant plex that was detached from volume `vol-01`.

3 Subdisk `dsk0-03` contains the FPA log for volume `vol-01BK`, and is associated to plex `plex-03`.

4 Subdisk `dsk3-01` contains the data in the migrant data plex `plex-03`.

Troubleshooting

Use the following list to identify and solve problems.

Problem

Not enough free disk space in the disk group to create the FPA logs:

```
# volassist snapfast vol-01 vol-01BK
lsm:volassist: ERROR: Cannot allocate space for 195 block log
lsm:volassist: ERROR: Could not attach log plex(es) to volume vol-01
```

FPA logs are the same size as dirty region logs (DRLs): approximately 65 blocks per GB of volume size. The disk group must have enough free space to create FPA logs for both the primary and secondary volumes.

Solution

Initialize a new LSM disk (if necessary) and add the disk to the disk group:

```
# voldisksetup -i dskn
# voldg -g disk_group adddisk dskn
```

Problem

Error message when trying to create a secondary volume:

```
# volassist snapfast vol_01 vol_01BK
lsm:volassist: ERROR: vol_01: Not in any imported disk group
```

Possible causes:

- The primary volume name was misspelled. (In this case, `vol_01` and `vol_01BK` should be `vol-01` and `vol-01BK`; with dashes, not underscores.)
- The correct disk group was not specified (`rootdg` is assumed).
- The secondary volume was listed before the primary volume (syntax error).

Solutions

- Verify the name of the primary volume and the disk group containing it.
- Specify the disk group if other than `rootdg`.

  For example:

  ```
  # volassist -g dg1 snapfast vol-01 vol-01BK
  ```

- Make sure the primary volume precedes the secondary volume on the command line.

  There must not already be a volume in the disk group with the name you assign to the secondary volume.

Problem

Error message from naming a volume other than the primary volume in a `volassist snapback` operation:

```
# volassist snapback vol-01BK vol-02
lsm:volassist: ERROR: Reattaching FPA plex plex-03 to non-primary \
volume is not allowed
```

You can return the migrant plex only to the volume from which it came.

Solution

Specify the primary volume for the `snapback` operation:

```
# volassist snapback vol-01BK vol-01
```

Problem

Error message when trying to use FPA on a volume with only two plexes:

```
# volassist snapfast vol-02 vol-02BK
lsm:volassist: ERROR: Operation would cause volume vol-02 to be degraded \
to a single plex, use '-f' force.
```

Solutions

- Add another data plex to the volume before using FPA (see *Is This Best Practice Right for You?*)
- Retry the command with the force (`-f`) option:

  ```
  # volassist -f snapfast vol-02 vol-02BK
  ```

Caution

The primary volume is no longer mirrored and is vulnerable to data loss if a disk fails.

Problem

Error message due to stale FPA log:

```
#  volassist snapback vol-01BK vol-01
ERROR: The FPA log is STALE use -f to force an attach without FPA
```

Possible causes:

- A disk failure affected the FPA log for either the primary or secondary volume.
- The FPA log plex on the primary volume was explicitly detached or dissociated.

Solution

Retry the command with the force (`-f`) option:

```
# volassist -f snapback vol-01BK vol-01
```

Note

If an FPA log has been detached or dissociated from its volume, due to disk failure or by an explicit command, FPA logging is stopped and cannot be resumed. Because the

FPA logs are no longer in effect, LSM initiates a complete resynchronization of all volume data to the returning plex.

Alternative Practices

This Best Practice applies only to mirrored LSM volumes, and therefore, only to volumes with concatenated or striped data plexes, and is the recommended method for using the Fast Plex Attach feature to back up these volumes.

If your system or LSM configuration does not meet the requirements described in *Is This Best Practice Right for You?* — for example, the volume you want to back up has a RAID5 layout or cannot be mirrored — see the *Logical Storage Manager* manual.

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