

Integrated Backup and Recovery in DMF

<Presenter's Name>



Outline

- Backup methods in DMF
- Integrated backup design goals
- Design
 - Dump
 - Restore
 - Disaster recovery
- Caveats
- Future directions

Backup methods in DMF

- Tape-based
 - A filesystem is directly backed on to tape
- Disk-based
 - A filesystem is backed on to disk

Tape-based vs. disk-based DMF backups

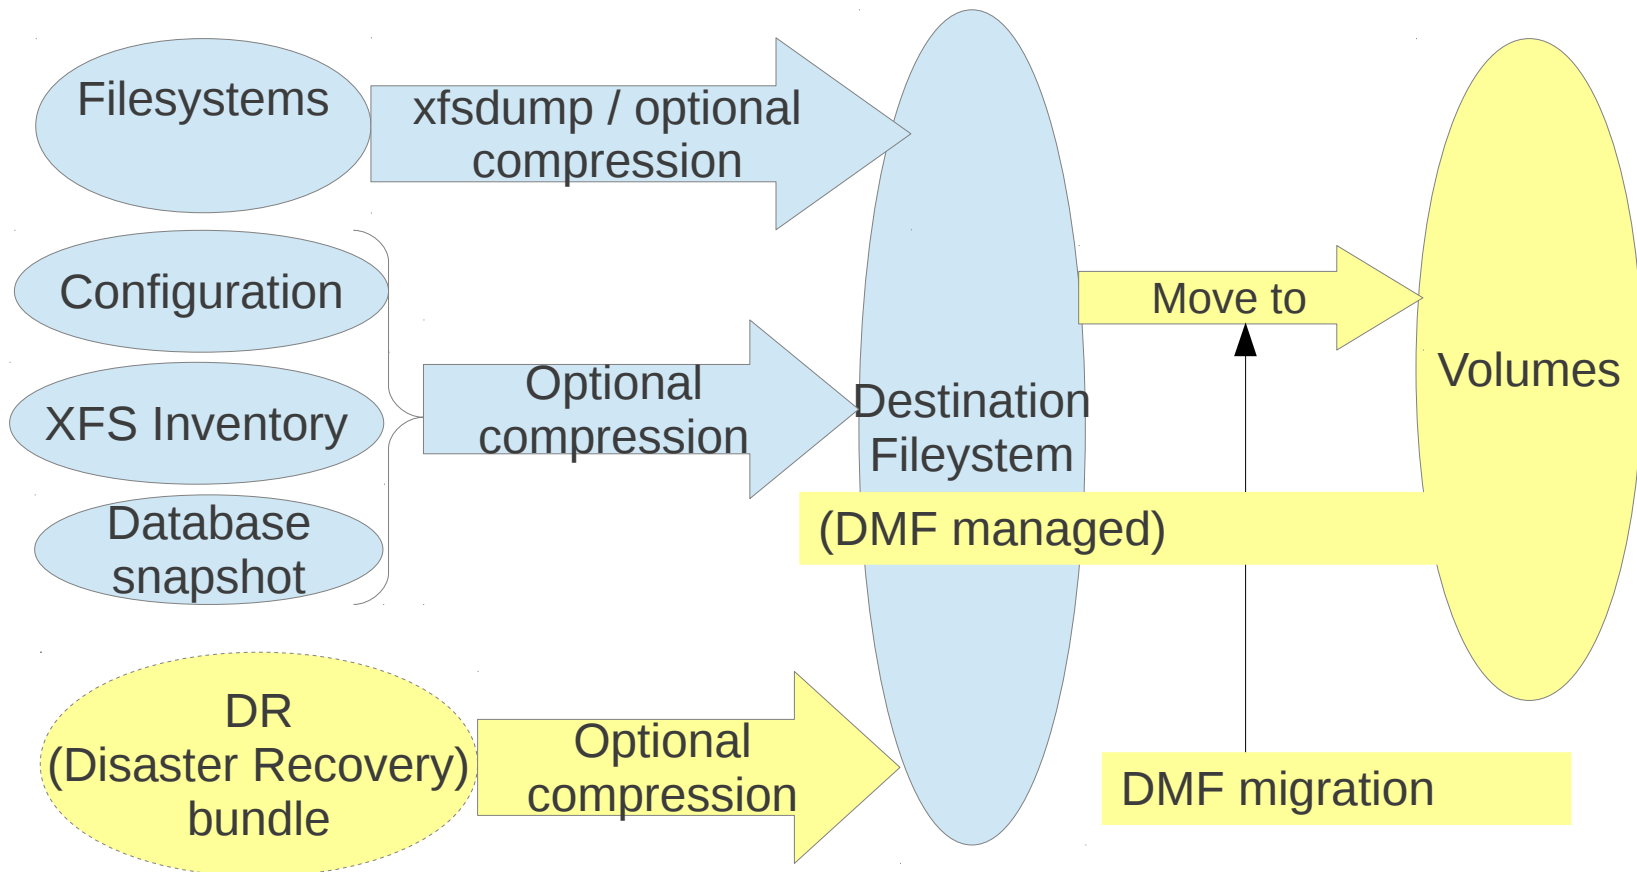
Property	Tape	Disk
Cost	Low	High
Dumps & restores	Slow	Fast
Usage of space	Sub-optimal (constrained by implementation)	Optimal
Copies	None (constrained by implementation)	Remote / local filesystems as dump mirrors
Incremental dumps	Levels 0-1 (constrained by implementation)	All levels (0-10)
Multi-stream support	None (requires multiple tapes at a time)	Up to 20 streams
Space management	Manual	Not required till filesystem runs out of space

Integrated backup design goal

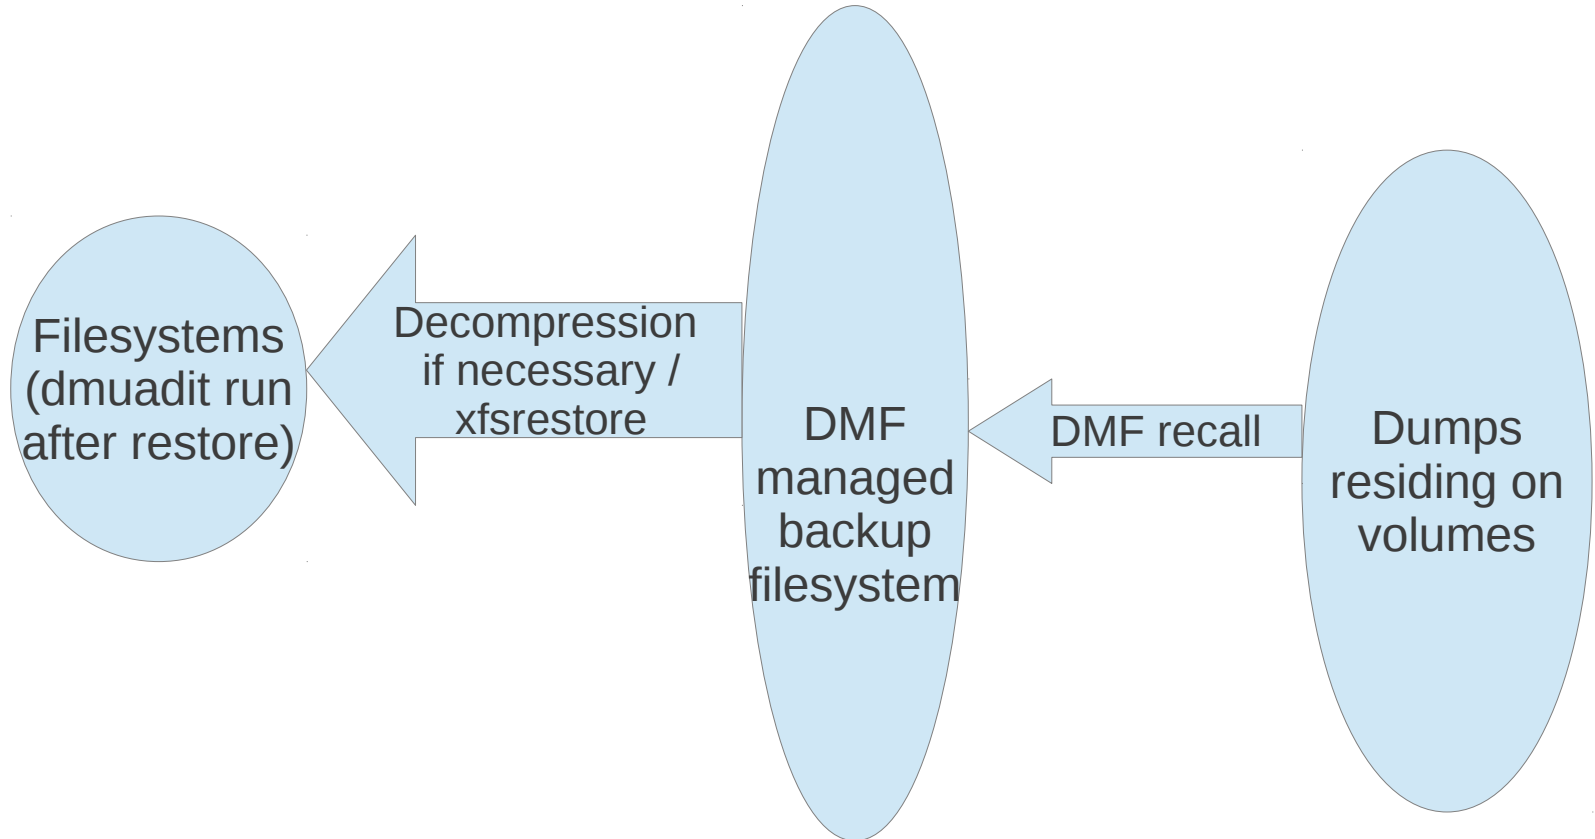
Retaining the advantages of a disk-based scheme while reducing the cost. In particular:

- ◆ Faster dumps by dumping first to a filesystem
- ◆ Safer dumps by moving dumps to multiple volumes
- ◆ Efficient storage of data on volumes
- ◆ Ease of use

An integrated backup scheme



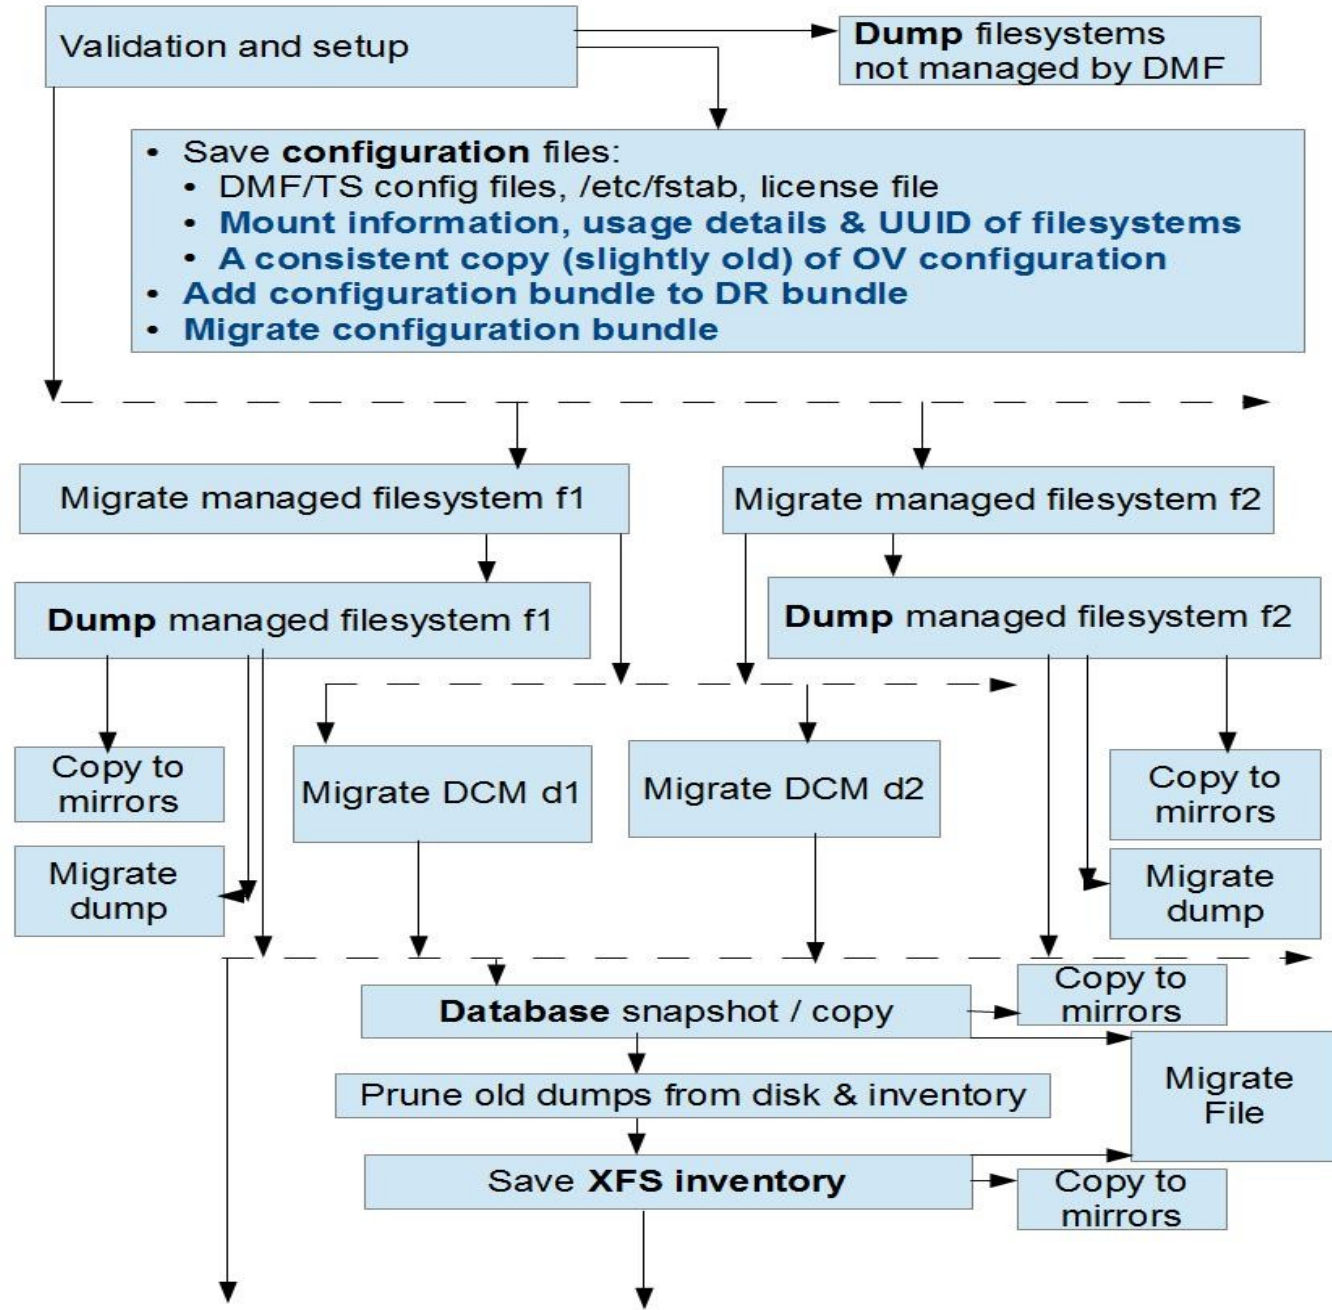
Restore



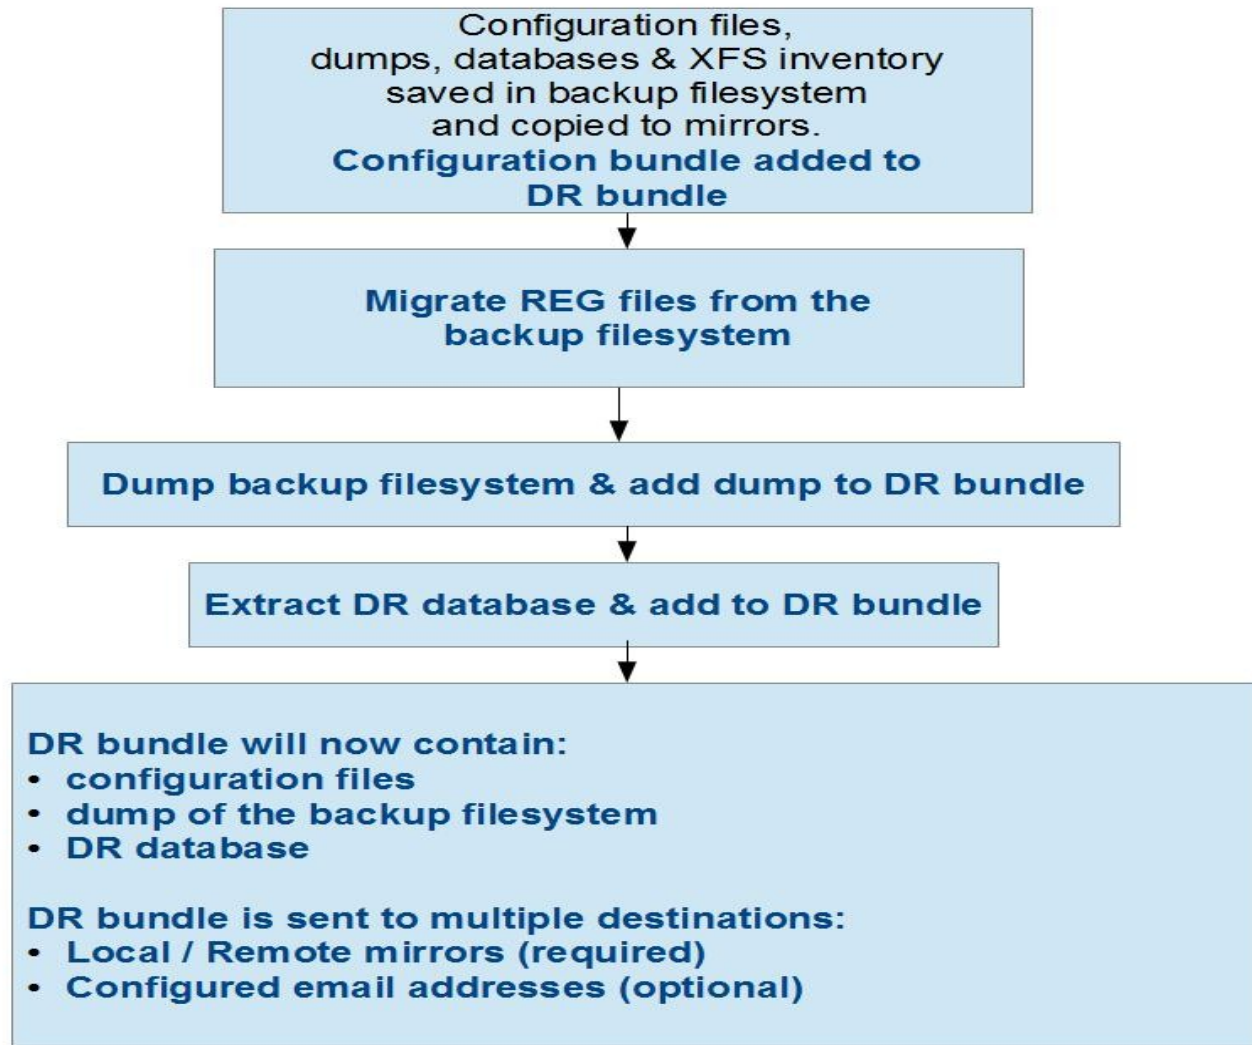
Configuration of backup filesystem

- The backup filesystem is a DMF managed filesystem marked exclusively for backups
- The backup filesystem is allocated at least 2 exclusive MSPs, one of which must be a library-server-based MSP.


DMF dump session (1/2)



DMF dump session (2/2)





Restore interface (1/2)

Configuration ▾ Storage ▾ Messages ▾ Statistics ▾ Backups ▾ Help ▾ [Change Password](#) [Log Out](#) 

Restore Filesystems x

Available dumps to restore from

1 - 10 of 37 items 10 | 25 | 50 | 100

	Date	Filesystem	Level
 No filter applied 			
<input type="radio"/>	2013-11-15T00:00:01	/dmfusr	1
<input type="radio"/>	2013-11-14T23:05:01	/dmfusr	0
<input type="radio"/>	2013-11-14T00:00:01	/dmfusr	1
<input type="radio"/>	2013-11-13T23:05:01	/dmfusr	0
<input type="radio"/>	2013-11-13T03:42:40	/dmfusr	1
<input type="radio"/>	2013-11-13T03:42:39	/dmfusr	1
<input type="radio"/>	2013-11-13T03:11:02	/dmfusr	0
<input type="radio"/>	2013-11-13T03:11:01	/dmfusr	0
<input type="radio"/>	2013-11-13T00:00:01	/dmfusr	1
<input type="radio"/>	2013-11-11T00:00:01	/dmfusr	0

Restore /dmfusr

Selected session:

Source filesystem: /dmfusr
Date: 2013-11-12T23:05:01
Dump level: 0
Session ID: bbcedb29-b262-459a-bd7e-d12f3837c21b
Method: Integrated
DMAPI: yes

Select restore options:


Source directory or file relative to the filesystem:
Leave empty for restoring the entire session:

Destination filesystem: /dmfusr

Destination directory relative to the destination filesystem:
Leave empty for restoring to the top level directory of the destination filesystem:

Do not overwrite existing files
Do not overwrite newer files
Overwrite existing files

Restore interface (2/2)

Configuration ▾ Storage ▾ Messages ▾ Statistics ▾ Backups ▾ Help ▾ [Change Password](#) [Log Out](#) 

Restore Filesystems x

```
Log for restore of /dmf/backups/test2
Restore command:
xfsrestore -D
-S b75625d1-d36f-4365-be41-90b14fc94007
-f fifo00
/dmf/backups/test2

xfsrestore: using file dump (drive_simple) strategy
xfsrestore: version 3.1.0 (dump format 3.0)
xfsrestore: using online session inventory
xfsrestore: searching media for directory dump
xfsrestore: examining media file 0
xfsrestore: reading directories
xfsrestore: 24 directories and 57 entries processed
xfsrestore: directory post-processing
xfsrestore: restoring non-directory files
xfsrestore: restore complete: 0 seconds elapsed
xfsrestore: Restore Summary:
xfsrestore:  stream 0 /dmf/tmp/dmxfrestore.14803/fifo00 OK (success)
xfsrestore: Restore Status: SUCCESS

Dump:  vajra:/dmfusr at 2013-11-11 00:00:01 (level 1)
Session: d4ba17ac-aaac-4123-a42d-1a43445fbf7
Media:  /dmf/backups/filesystems/12e40d9f-5abe-4331-b966-6ec49c998fcd/dmfusr-20131111_000000-
l1-s00.xfsdump.gz

Restore command:
xfsrestore -D
-S d4ba17ac-aaac-4123-a42d-1a43445fbf7
-f fifo00
/dmf/backups/test2

xfsrestore: using file dump (drive_simple) strategy
xfsrestore: version 3.1.0 (dump format 3.0)
xfsrestore: using online session inventory
xfsrestore: searching media for directory dump
xfsrestore: examining media file 0
```

Status of restore processes

1 - 6 of 6 items 10 | 25 | 50 | 100

<input type="checkbox"/>	Destination	Time Started	Time Ended	Status
<input type="checkbox"/>	No filter applied			
<input type="checkbox"/>	/dmf/backups/test1	2013-11-13T06:30:31	2013-11-13T06:30:33	Completed
<input type="checkbox"/>	/dmf/backups/test1	2013-11-13T06:28:11	2013-11-13T06:28:43	Error
<input type="checkbox"/>	/dmf/backups/test3	2013-11-13T06:20:03	2013-11-13T06:20:05	Completed
<input type="checkbox"/>	/dmf/backups/test3	2013-11-13T06:19:57	2013-11-13T06:19:59	Completed
<input type="checkbox"/>	/dmf/backups/test3	2013-11-13T03:43:44	2013-11-13T03:43:46	Completed
<input checked="" type="checkbox"/>	/dmf/backups/test2	2013-11-11T03:13:30	2013-11-11T03:13:30	Completed

Acknowledge
Kill...
Show Log

Disaster Recovery

- Done when one or more of all the following is lost:
 - Configuration files on the root filesystem
 - Backup filesystem
 - DMF database
- Restore is performed in the following order using the DR bundle:
 - Configuration files
 - Backup filesystem
 - DR database
 - DMF database

Pruning dumps & DR bundles

- Space on volumes recycled based on DUMP_RETENTION
- Space occupied by DR bundles recycled based on a fixed retention period

Future directions:

- Better retention policies & pruning methods.
- Records related to existing dumps must not be hard-deleted.

Assumptions

- HFREE_TIME and JOURNAL_RETENTION are set appropriately
- The user wants to go back only to the last sane state of the database
- **Either** of the following is true:
 - No merge/dmmove/dmemptytape operations have occurred between that sane state and recovery point.
 - Either of the HOME_DIR or JOURNAL_DIR filesystems is safe

Caveats

- Certain DMF operations (e.g. merge, move, etc.) could cause:
 - Invalidation of DR databases
 - Invalidation of metadata backups
- What if both HOME_DIR and JOURNAL_DIR filesystems are lost ?
 - Only partial database recovery is possible based on the last successful backup
 - All the more important to have HOME_DIR and JOURNAL_DIR filesystems mounted on different devices

Future directions

- OpenVault database snapshot to get a consistent current copy of the OpenVault database
- OpenVault configuration restore capability
- DMF journal snapshot to conserve backup media space
- Better retention policies while retaining validity of existing dumps

sgi