



www.csiro.au

DMF as a target for backup

Robert C Bell, Jeroen van den Muyzenberg, Peter Edwards

CSIRO Advanced Scientific Computing



Features

The techniques in use provide:

- 1. coverage back in time adjusting to the likelihood of recovery being needed**
- 2. full backups every time, for the cost of incrementals**
- 3. simple visibility of the backup holdings**
- 4. simple recovery for individual files and complete file systems**
- 5. no vendor dependency**
- 6. centralised tape management**
- 7. space saving of about a factor of five compared with conventional backups**

This is for non-DMF-managed file systems

Outline

- **The need for backups**
- **Unitary versus discrete**
- **Disc versus tape target (feature 6)**
- **Rsync with - -link-dest (features 2, 3, 4, 5, 7)**
- **Tower of Hanoi (features 1, 3)**

The need for backups

- **Need to do risk assessment**
 - What if the data was lost?
 - How could the data be lost?
 - User errors (oops, wrong directory)
 - Systems staff errors (even bigger oops)
 - System software failures (RAID software)
 - Hardware failures (disc drives)
 - External event
- **Backups are snapshots of data holdings, made to allow recovery in some events.**
 - Need latest, and to go back in time.
- **Backups are not archives!**

Unitary versus discrete

Dilemma: have many files. Make backup to:

1.a container (dump file or tape); or

2.a file system (equal numbers of files)

Option 1 is easy to manage, and good for recovery of the lot, but requires serial access (hence slow) for individual file recovery, and has no addressability. So, if you want to keep a current full dump, have to write the full information each time.

Option 2 is not normally feasible, because of the number of files. Good for individual file recovery.

Disc versus tape target (feature 6)

- **Backups to disc**
 - attractive – easy to manage, reliable
- **Backups to tape**
 - off-line, hence safer, but need drive on each system, or backup over network: less reliable
- **Backups to disc then to tape**
 - best of both worlds.
- **With an HSM already dealing with tapes, no extra tape management is needed for the backups – just target to a DMF filesystem.**
 - Even better!
- **6. Provides for centralised tape management**

Rsync with --link-dest

- **rsync: copies files, but only as necessary.**

source -->

destination directory

- **Colleague, Jeroen van den Muyzenberg discovered a feature of the rsync utility in 2007**

--link-dest= previous backup directory

source ->

destination directory

- **Hard-links formed between files in previous and current destination directories.**
- **Hard-links – same file, but aliased in multiple directories**

Backup directories and hard links

Backup															
seq 0	█	█	█	█	█										
seq 1	█	█	█		█	█	█	█	█						
seq 2	█	█	█			█	█	█		█	█	█			
seq 3	█		█			█	█			█	█		█	█	
seq 4	█					█				█			█		█
<i>link-count</i>	5	4	3	2	1	4	3	2	1	3	2	1	2	1	1

Rsync with --link-dest (features 2, 3)

- **Each directory is a full backup, but only changes need to be transferred.**
- **2. full backups every time, for the cost of incrementals**
- **Can use standard UNIX/Linux to look at backups (ls, etc).**
- **Can see a view back in time.**
- **3. simple visibility of the backup holdings**

Rsync with --link-dest (feature 3)

```
pwd ; dmls -al ....
```

```
/backup2/backup/test-trout-dumps
```

```
-rw-rw----1 2521929 2008-10-16 13:51 (OFL) 20081016T1626.seq.0/peter  
-rw----- 2 3271394 2009-03-25 12:09 (OFL) 20090325T1429.seq.64/peter  
-rw----- 1 3273354 2009-03-26 09:15 (OFL) 20090326T0946.seq.96/peter  
-rw----- 1 5683063 2009-07-13 10:34 (DUL) 20090713T1310.seq.144/peter  
-rw----- 3 6818519 2009-08-27 10:15 (DUL) 20090827T1113.seq.148/peter  
-rw----- 3 6818519 2009-08-27 10:15 (DUL) 20090827T1156.seq.151/peter  
-rw----- 3 6818519 2009-08-27 10:15 (DUL) 20090827T1203.seq.152/peter  
-rw----- 1 6827918 2009-08-27 17:03 (DUL) 20090827T1952.seq.153/peter  
-rw----- 1 6916875 2009-08-31 09:28 (DUL) 20090831T0942.seq.154/peter  
-rw----- 1 10853271 2009-09-24 16:00 (REG) 20090925T1047.seq.155/peter
```


Rsync with --link-dest (features 4, 5, 7)

- **Can use standard UNIX/Linux cp to recover individual files, or rsync to recover the lot.**
- **4. simple recovery for individual files and complete file systems**
- **All standard UNIX/Linux commands plus rsync**
- **5. no vendor dependency**
- **Allows multiple snapshots, without large growth**
- **7. space saving of about a factor of five compared with conventional backups**

Rsync with --link-dest

- **Lots of problems, nearly all solved**
 - Number of inodes grows – directories – tolerable
 - Purging and thinning
 - Recycling
 - rsync deficiencies when file exists on source and link-dest, but older file in destination – copies from source instead of linking from link-dest – heuristic to fix this, or `cp -l`.
 - Schemes for superseded backup series – culling
 - Consolidation – planned
- **Use DMF, and xfsdump 3 times per week to ensure recoverability – all eggs in one basket otherwise.**
- **Big issue – only one physical copy**

Tower of Hanoi



Aim: to move all the rings from one tower to another, one at a time, with no ring is placed on a smaller ring.

So, if we number the rings from number 1, the smallest, the moves are in the order: ring 1, ring 2, ring 1, 3, 1, 2, 1, 4, 1, 2, 1, 3, 1, 2, 1, 5, 1, 2, etc.

Tower of Hanoi

Tower of Hanoi move number	Tower of Hanoi ring to move	
1	1	00001
2	2	00010
3	1	00011
4	3	00100
5	1	00101
6	2	00110
7	1	00111
8	4	01000
Backup sequence number	Backup set number	Binary sequence number

Tower of Hanoi

- **Usage:**

- assign a sequence number, then a set number (from ToH)
- keep one (directory, tape) of each set number

Assigning slash.20090722.seq.0 to set 0

Assigning slash.20090823.seq.32 to set 6

Assigning slash.20090909.seq.48 to set 5

Assigning slash.20090917.seq.56 to set 4

Assigning slash.20090920.seq.59.recycle to set 1

Assigning slash.20090921.seq.60 to set 3

Assigning slash.20090922.seq.61 to set 1

Assigning slash.20090923.seq.62 to set 2

Assigning slash.20090924.seq.63 to set 1

Found 9 backups as expected up to sequence number 63

Tower of Hanoi

- **Advantages**

- Easily programmed to determine which backup set to use. No complication of days of the week, days/weeks per month, etc.
- Can use a sequence number rather than just a date. This allows arbitrary spacing of backups.
- The coverage of backups spaces out over time. For example, with daily backups, there will be better than two backups in the last two days, 3 backups in the last 4 days, 4 backups in the last 8 days, etc.

- **1. coverage back in time adjusting to the likelihood of recovery being needed**

- The storage grows as $\log_2 n$, where n is the number of backups taken.

- **7. space saving of about a factor of five compared with conventional backups**

- The system self-heals.

Tower of Hanoi

- **Improvements**

1. Start with sequence number zero, set 0, and never remove it – baseline backup. The cost is one extra backup set.
2. To provide better coverage of recent days (or times), we keep two set 1s. The cost is one extra backup set.
3. We make backups to a new area before discarding or recycling. The cost is one extra backup set.

These together provide even more coverage – we always keep at least the last 5 backups.

The second enhancement can be easily extended: for example, if you want 24 hours coverage of hourly backups, you would keep 12 set 1s, 6 set 2s, 3 set 3s, 2 set 4s, and every higher set number.

Features

The techniques in use provide:

- 1. coverage back in time adjusting to the likelihood of recovery being needed**
- 2. full backups every time, for the cost of incrementals**
- 3. simple visibility of the backup holdings**
- 4. simple recovery for individual files and complete file systems**
- 5. no vendor dependency**
- 6. centralised tape management**
- 7. space saving of about a factor of five compared with conventional backups**

Implementation – another time!



CSIRO

[Home](#) | [Contact Us](#) | [Facilities](#) | [Software](#) | [Support](#) | [News](#) | [Partnerships](#) | [User Information](#) |

CSIRO ADVANCED SCIENTIFIC COMPUTING

User Filesystem Backup Status

Last updated Tue 2009-09-29 11:00:01 +1000

Error codes 23 and 24 are soft errors (files disappeared or open) - these can be ignored.

Sorted by date:

Host	Filesystem	Last backup	Status	Files Transferred	Total Files	Bytes Transferred	Total Bytes
cherax	/cs/datastore_0	Sun 2009-09-27 11:47	0	15555342	16188904	82704722344	4283195858944
burnet-fs0	/system/cs/home	Mon 2009-09-28 18:30	0	8526	846308	4442311132	313213832298
CBF	/current	Mon 2009-09-28 19:03	0	397	7244997	103613919	905086125446
sx-6	/cs/home	Mon 2009-09-28 20:48	23	1618	862158	29749655382	712029664160
cherax	/cs/datastore_9	Tue 2009-09-29 01:03	0	501940	16393816	3647467608	4554001555456

Powderhorn tape library



SL8500 tape library and Gareth Williams



Your data, our responsibility



CSIRO ASC

Dr Robert Bell

Phone: +61 3 9669 8102

Email: Robert.Bell@csiro.au

Web:

http://hpsc.csiro.au/users/dmfug/Presentations_Oct09/DMF_as_a_target_for_backups/

Thank you

Contact Us

Phone: 1300 363 400 or +61 3 9545 2176

Email: Enquiries@csiro.au **Web:** www.csiro.au

