Your Data, Our Responsibility CSIRO IM&T Advanced Scientific Computing



Robert Bell*, Jeroen van den Muyzenberg, Peter Edwards

Introduction

High Performance Computing (HPC) centres have become holders of high volumes of data. Users would like unlimited storage that is high performance, visible from everywhere, secure from loss, and at zero cost.

Storage Dilemma

- HPC centres should meet the users' needs, but are constrained by budgets and technologies. Truism: "good, cheap or quick – choose any two".
- With storage, there are even more considerations: capacity, performance, visibility, reliability, resiliency, recoverability, recovery time, latency and backup.
- Centres cannot provide storage with all the desirable characteristics in the one system They usually provide a range of file systems, with different policies.
- The facilities and policies for users' data affect the productivity of the users and their perception of the service.
- Managing data in many locations, and managing transfers between the locations is error-prone, and highly unproductive for users.
- Centres must make backups of files to provide some protection against loss.
 - · Need to do a risk assessment
 - What if the data was lost? How could the data be lost?

CSIRO Data Store Total holdings (log scale)

1991 1994 1997 1999 2002 2005 2008 2010

10000

1000

100

10

1

0.1

0.01

Need current snapshot, plus history

Hierarchical Storage Management (HSM) – Why?

- · Virtually infinite storage for users
- Easy to add capacity
- Offline data is no longer dead data
- Easy transitions
 - CSIRO at 4th site, 5th host, 9th type of tape since 1991
- · Less energy usage than disc alone
- Reduces costs
- Gets around backup problem
 trickles data from disc
- · Backup target for other systems

rsync and the Tower of Hanoi for backups of systems into an HSM provide the following features:

F1. centralised tape management

- F2. full backups every time, for the cost of incrementals
- F3. simple visibility of the backup holdings
- F4. simple recovery for individual files and complete file systems
- F5. no vendor lock-in
- F6. space saving of about a factor of five compared with conventional backups
- F7. coverage back in time automatically adjusting to the likelihood of recovery being needed



* Contact: 0428 108 333

rsync with --link-dest=DIR

Backup from disc to disc, then tape (using HSM) *F1* Sequence of backups – any frequency.

To build a new backup directory, have two sources for active files:

- 1. the source file system (for all files)
- copy just the new files using rsync
 2. the previous backup directory (for old files)
 - just hard-link old files from the previous directory
 - rsync with --link-dest=DIR does this



Each directory is a full backup, but only changes need to be transferred *F2*

Can use standard UNIX/Linux to look at backups, with a view back in time *F3*

Can use standard UNIX/Linux cp/scp to recover individual files, or rsync to recover the lot *F4* All standard UNIX/Linux commands plus rsync *F5* Allows multiple snapshots, without large growth

(no duplication in backups – de-duplication) **F6**

Management of backup sets Backup rotation scheme - Tower of Hanoi

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

Number the rings from number 1 (the smallest). The moves are: 1, 2, 1, 3, 1, 2, 1, 4, 1, 2, 1, 3, 1,

2, 1, 5, 1, 2, etc.

Use these ring numbers as labels for backup sets. Keep one of each set (directories or tapes).

20090722.seq.0 to set 0 20090823.seq.32 to set 6 20090909.seq.48 to set 5 20090917.seq.56 to set 4 20090921.seq.60 to set 3 20090922.seq.61.recycle to set 1 20090923.seq.62 to set 2 20090924.seq.63 to set 1 **F6 and F7**



urther information contact: ASC Help Desk whone: (03) 9669 8103 or (03) 8601 3800 mmail: hpchelp@csiro.au

ununu esiro au