



DMF UGM February 2015 at UNSW Site report: CSIRO SC Data Store

Robert C. Bell | HPC Partner Services Manager
18 February 2015

CSIRO IM&T SCIENTIFIC COMPUTING SERVICES
www.csiro.au



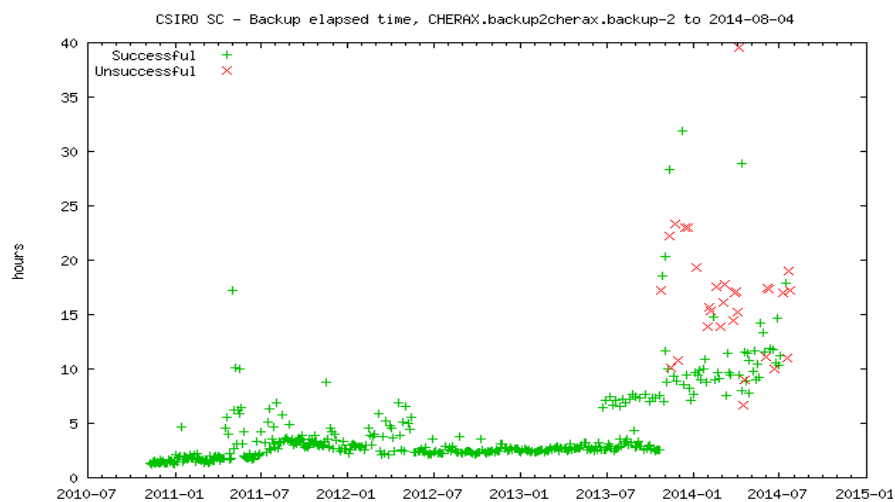
CSIRO Docklands Site report

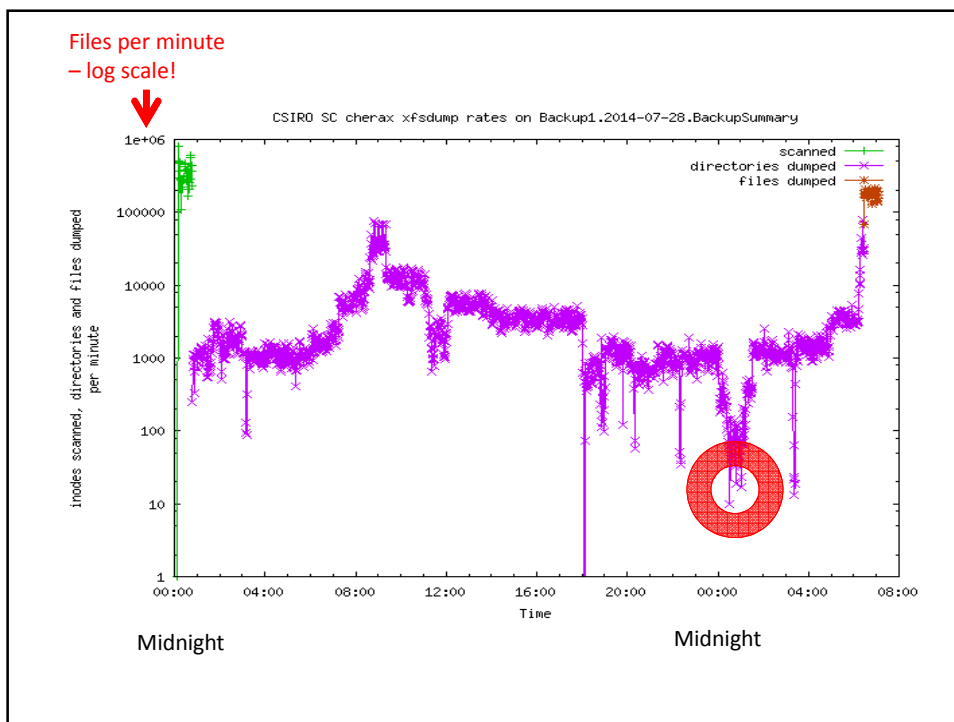
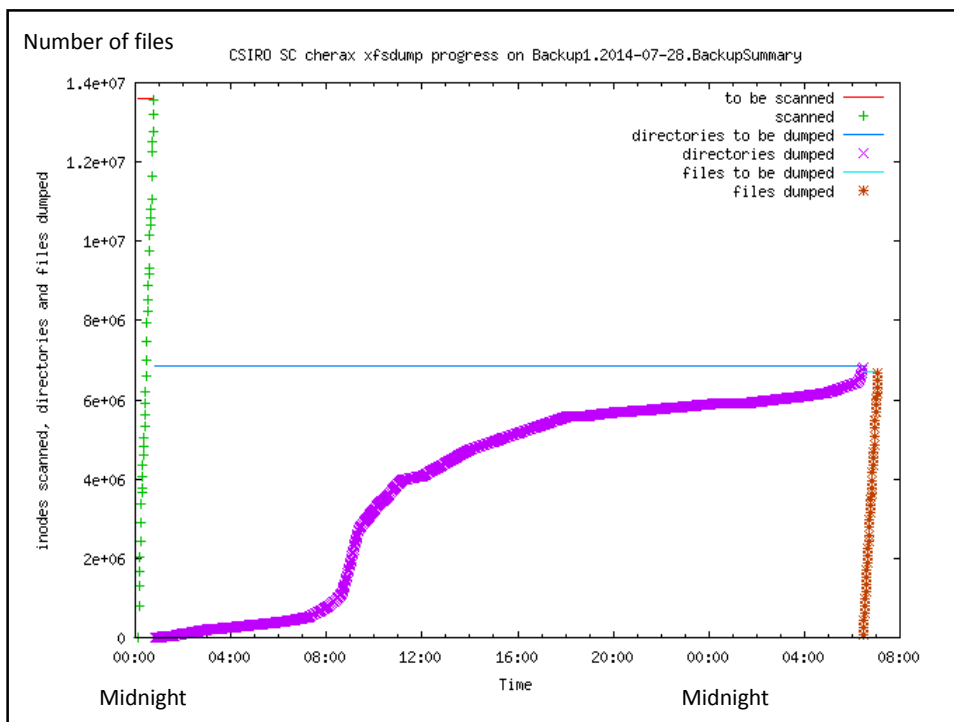
- A problem 'solved'
- A new service (see Peter's talk)
- History
- Current
- Future



The problem

- xfsdumps very slow after upgrade to SLES11 Service Pack 2
- Particularly slow for xfsdumps of filesystems that were the targets for backups using rsync

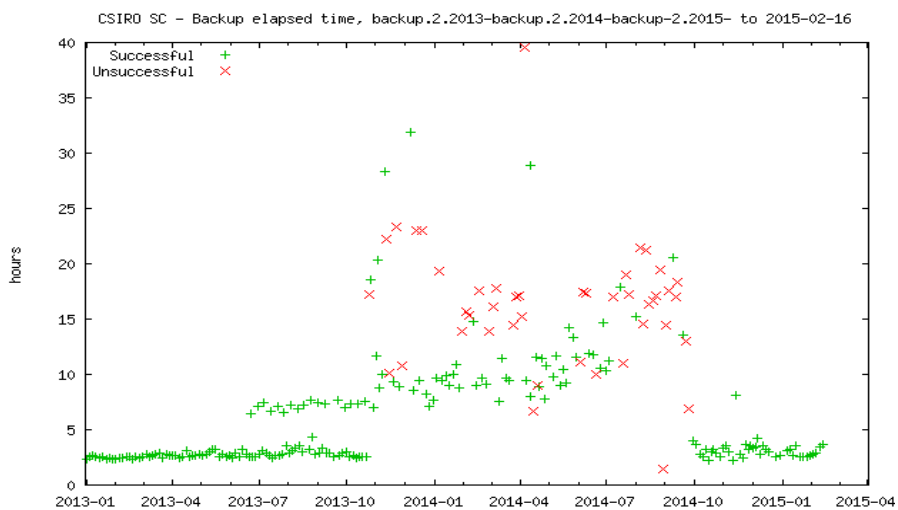




The solution:

- xfsdump or filesystem performance issued not solved!
- Problem circumvented: replaced IS220 (cheap and not cheerful!) with IS5500 disc with IS4600 SSD
- xfsdump performance now acceptable

7 | Off-site storage | Robert C. Bell



8 | Off-site storage | Robert C. Bell



The solution:

- rsync backups to the targets now much faster:
 - can do backup of filesystems with 30 million files and 10 Tbyte of data in $\frac{3}{4}$ hour
 - typically 0.5% of files, 1% of data change in a day

History

- Next month, it will be 25 years since the filesystem under the current Data Store was established

CSIRO ASC Data Store: History


- 1990: JSF: Cray Y-MP2/216 – “big”, expensive, fast disc
 - Turned a compute problem into a storage problem!



11 | Off-site storage | Robert C. Bell

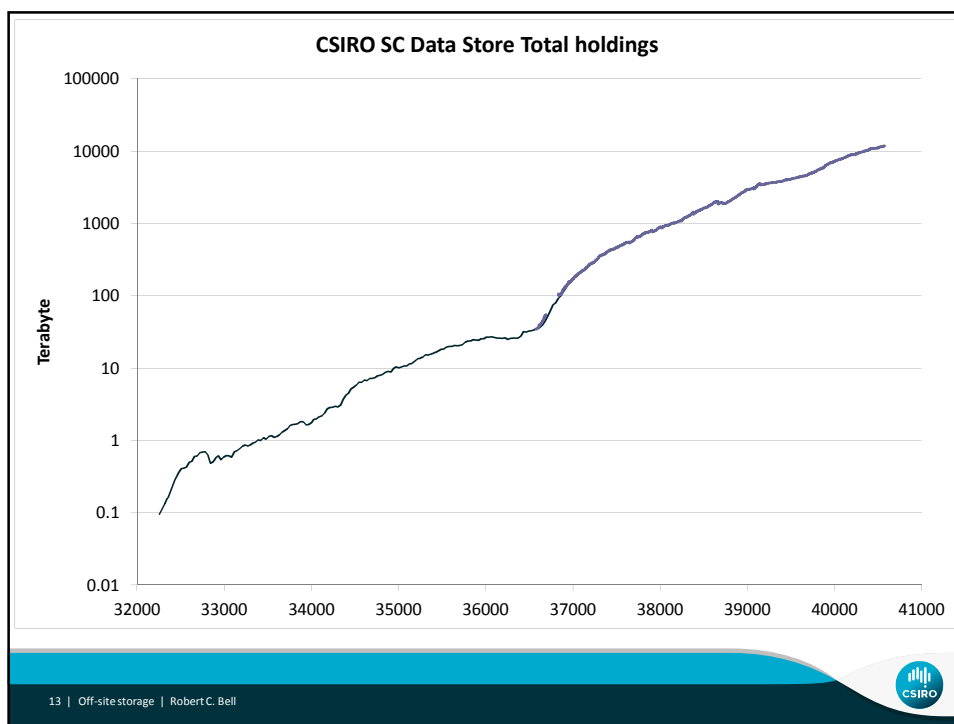


CSIRO Peak Computing Systems



12 | Off-site storage | Robert C. Bell





History

- 4 sites, 6 hosts, 6 disc systems, 11 tape technologies
- Oldest content identified dates back to mid-1960s
- About $\frac{3}{4}$ of files accessed since created
- Never had to put size limits on users' total holdings
- Limits on inodes (numbers of files) have been critical
- Direct user access to migrating file system – \$HOME

Current

- Datastore and 4 backup target areas
- Current holdings under DMF management:
 - 116M inodes on discs
 - Threshold for migration:
 - 28 kbyte (for /datastore), 4 kbyte (for /backup*)
 - Data holdings (Tbyte):
 - First copy: 5445 Total active: 11782
 - Primary disc: 29
 - DCM disc: 48
 - MAID: 849 active, 736 hard-deleted
 - Tape: 10890 active, 502 hard-deleted

Future

- Need to vacate Bureau of Meteorology's Central Computing Facility at Docklands this year
- Plan to move cherax services to a new host in Canberra, to sit beside major CSIRO clusters, and near NCI NF
- Tape libraries in Black Mountain and Clayton – dual site at last!

Thank you

CSIRO IM&T Scientific Computing Services
Robert C. Bell

t +61 3 9669 8102
e Robert.Bell@csiro.au
w www.hpsc.csiro.au

CSIRO IM&T SCIENTIFIC COMPUTING SERVICES
www.csiro.au

