



Upgrade to ISSP 3.0 – our experience

Lindsay Brebber
DMFUG – December 2013



History

- Using DMF since 1996
- Multiple data copies, primary in tape library, offsite in commercial storage facility
- Primary data storage
 - ❑ 0 to 1 Petabyte in 15 years
 - ❑ 1 Petabyte to 2 Petabytes in 10 months
- Relocated entire facility in February 2011

Computing facilities

- SGI Altix ICE - 96 nodes – installed 2009
- SGI cluster – 20 node – Windows 2008 R2 HPC

Major use

- Scientific research and operations
- Remote sensing research, vegetation management, soil and water
- Climate modelling, global circulation models and regional climate models (joint research with CSIRO)





Facilities

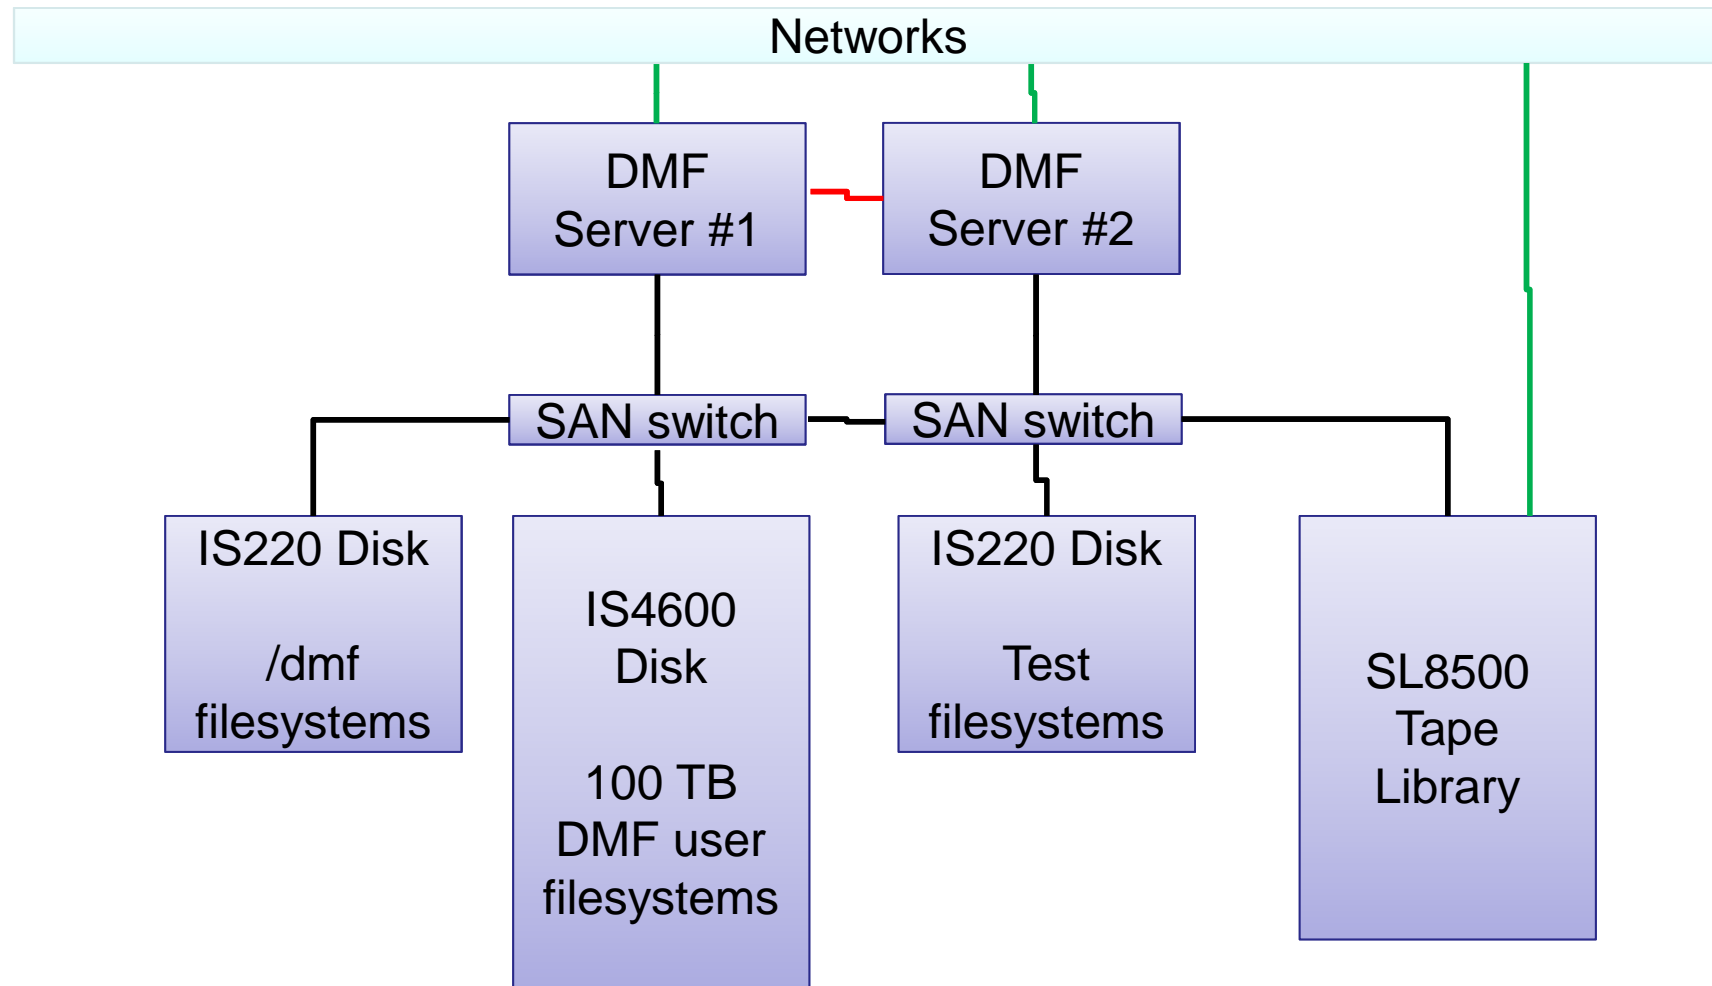
Tape

- Oracle SL8500
- 10,000 slots, licenced for 7,000 slots
- 6 x HP LTO4 – 800GB
- 2 x IBM LTO5 – 1500GB
- 2 x Oracle T10000C – 5000GB





High Availability configuration - outline





Upgrade Path

- Old configuration
 - Altix 450 (IA64)
 - 8 Gbit/s fibre channel
 - 10 gigabit Ethernet
 - DDR Infiniband
 - SLES11SP1
 - ISSP 2.5 (DMF 5.5)
- New configuration
 - C2108-RP2 (x86_64)
 - 8 Gbit/s fibre channel
 - 10 gigabit Ethernet
 - QDR Infiniband
 - SLES11SP2
 - ISSP 3.0 (DMF 6.0)



Preparation phase

- Setup two new x86_64 servers
 - T9840D tape drives
 - IS220 disk array (/dmf/home, /dmf/spool, /dmf/journals, /test1)
 - SLES11SP2
 - ISSP 3.0
 - SuSE High availability extension (HAE)
- Configuration
 - Configure HAE from scratch based on steps in DMF HA guide
 - High Availability Guide for SGI InfiniteStorage (007-5617-006)
 - Replica of production environment, tape drives, disk, networks on a smaller scale
- Testing
 - Running mix of continual recalls and migrates with network load



Migration phase

- Live HAE migration from IA64 to x86_64 servers
 - Setup both new systems with HAE with only stonith defined
 - Shutdown inactive HAE node on IA64 systems
 - Moved all SAN/network cables from IA64 server to new server
 - Booted new server and confirmed disk/tape/networks ok
 - Exported/Imported HAE config (CIB) from production (IA64) system to new (x86_64) system after removing stonith definitions
 - Stopped all HAE services on active HAE node on IA64 server
 - Started all HAE services on new server
 - Tested services from client systems Altix ICE, Windows PCs
 - Shutdown other IA64 server, moved cables, tested, booted, tested HAE failover



What didn't work

- Existing features
 - dmaudit from command line - LC_ALL sort problem (SGI patch)
 - tape/disk I/O performance is not the same (under investigation)
 - stonith deathmatches using bonded (backup) heartbeat network
- New features
 - dmarchive and Panasas (under investigation)



Lessons

- Good documentation: High availability guide for SGI Infinitestorage
- Apart from some issues with IPMI and stonith, it all worked



Plans and Issues

Plans

- Start using some of new features in DMF 6.0/ISSP 3.0
 - dmarchive
 - alternate recall options if primary drives are busy
- Integrate dmarchive use with our Panasas disk array on Linux cluster
- Increase size of disk array on DMF server

Issues

- Change in data storage patterns due to change in business/research
- Budget



Questions

Remote Sensing Centre,
Department of Science Information Technology Innovation and the Arts
<http://www.qld.gov.au/environment/land/vegetation/mapping/remote-sensing/>

The screenshot shows the Queensland Government website. At the top left is the Queensland Government logo. To its right are links for 'Contact us' and a search box labeled 'Search website'. Below this is a navigation bar with 'For Queenslanders' and 'Business and industry'. A breadcrumb trail reads: 'Queensland Government home > For Queenslanders > Environment, land and water > Land, housing and property > Land and vegetation management > Satellite mapping and monitoring > Remote sensing centre'. A left-hand menu titled 'Satellite mapping and monitoring' includes 'Remote sensing centre' (highlighted), 'Statewide monitoring programs', 'Assessing land clearing', 'Partnerships', and 'Maps, data and imagery'. The main content area is titled 'Remote sensing centre' and contains the following text: 'Scientists in our remote sensing centre use data from satellites, airborne sensors and ground-based platforms to map and monitor Queensland's landscape. They look at imagery of groundcover, land use, vegetation change, fire scars and water bodies across the state.' To the right of this text is a satellite image of a landscape with a river and a channel. Below the image is the caption: 'Satellite image of the Herbert River, Ingham on the left and the Hinchinbrook Channel off to the right in northern Queensland.'