




MONASH University

Monash Site Update

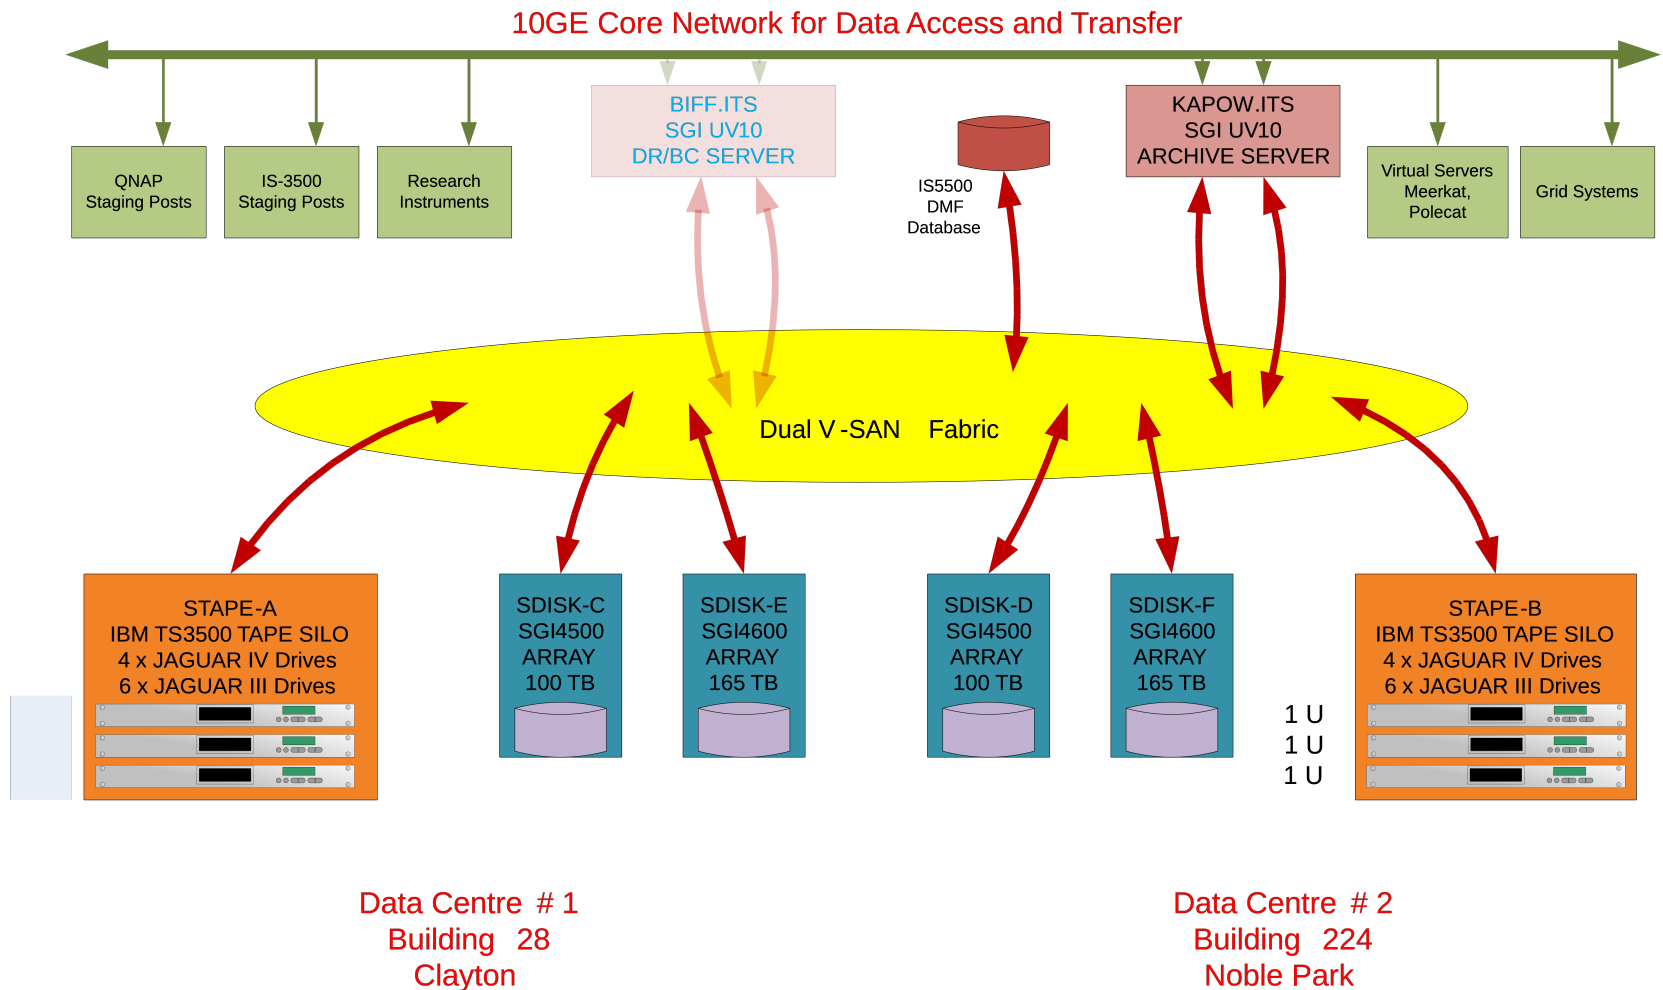
Growth in LaRDS/DMF Capacity

Stephen Dart
Research Storage Manager
Monash eResearch Centre
eSolutions




4th December 2013

Monash DMF Configuration




Monash Quarterly Growth Since 2009




Predicted Growth Characteristics to 2015

- Multiples of (step + linear) per group/project



Tape Library Expansion – Midyear 2012




Capacity Uplift and Growth Limit Extended

- More frontend file system caches allocated
 - Dedicated to MASSIVE and Synchrotron
- Reduce tape copies from 4 to 2.
- 180 JC Tapes (4TB per cartridge)= 720TB
- Migrate volume sets to higher density media
- Still Retiring JB's in favour of JCs
- Improved system support and monitoring
 - <http://meraem.its.monash.edu.au/munin/LaRDS/sirius.its.monash.edu.au.html>

Filesystem usage (in %) - by day

ROPTOOL / TOBI OETIKER

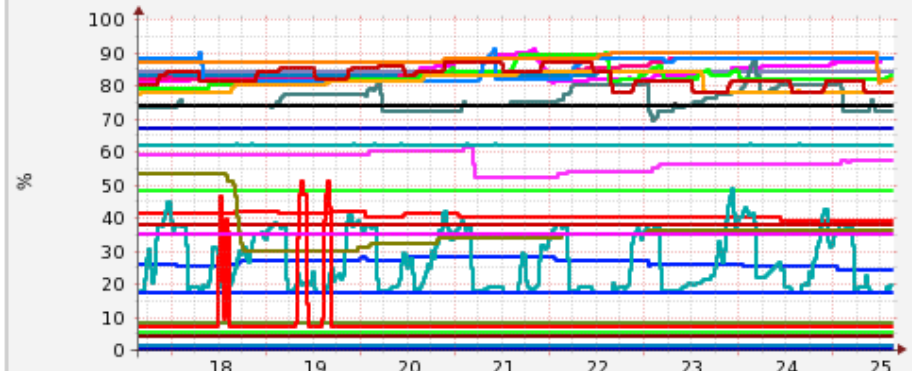


	Cur:	Min:	Avg:	Max:
/dmf-home	48.00	48.00	48.00	48.00
/dmf-journals	24.00	24.00	24.53	25.00
/dmf-spool	39.00	39.00	39.04	40.00
/dmf-tmp	19.07	17.00	24.76	43.00
/mnt/hpc	88.00	86.00	86.62	88.00
/mnt/a1	5.00	5.00	5.00	5.00
/mnt/r0	38.00	38.00	38.00	38.00
/mnt/r1	67.00	67.00	67.00	67.00
/mnt/r2	84.00	84.00	84.00	84.00
/mnt/r21	84.00	84.00	84.00	84.00
/mnt/r22	88.00	87.00	87.98	88.00
/mnt/r23	1.00	1.00	1.00	1.00
/mnt/r3	8.00	8.00	8.00	8.00
/mnt/r31	72.00	71.25	76.68	82.00
/mnt/r32	36.00	36.00	36.00	36.00
/mnt/r33	74.00	74.00	74.00	74.00
/mnt/r34	83.00	82.00	82.04	83.00
/mnt/r35	88.00	88.00	88.00	88.00
/mnt/r36	82.00	80.00	88.78	91.00
/mnt/r37	4.00	4.00	4.00	4.00
/mnt/r38	57.00	56.00	56.44	57.00
/mnt/rdb	5.00	5.00	5.00	5.00
/home	17.00	17.00	17.00	17.00
/	7.00	7.00	7.00	7.00
/var	62.00	62.00	62.00	62.00
/boot	35.00	35.00	35.00	35.00
/mnt/backup2	78.00	78.00	78.01	78.94
/mnt/backup1	78.00	78.00	79.38	81.00
/dev/shm	0.00	0.00	0.00	0.00
/dev	1.00	1.00	1.00	1.00

Last update: Mon Nov 25 15:20:19 2013

Filesystem usage (in %) - by week

ROPTOOL / TOBI OETIKER



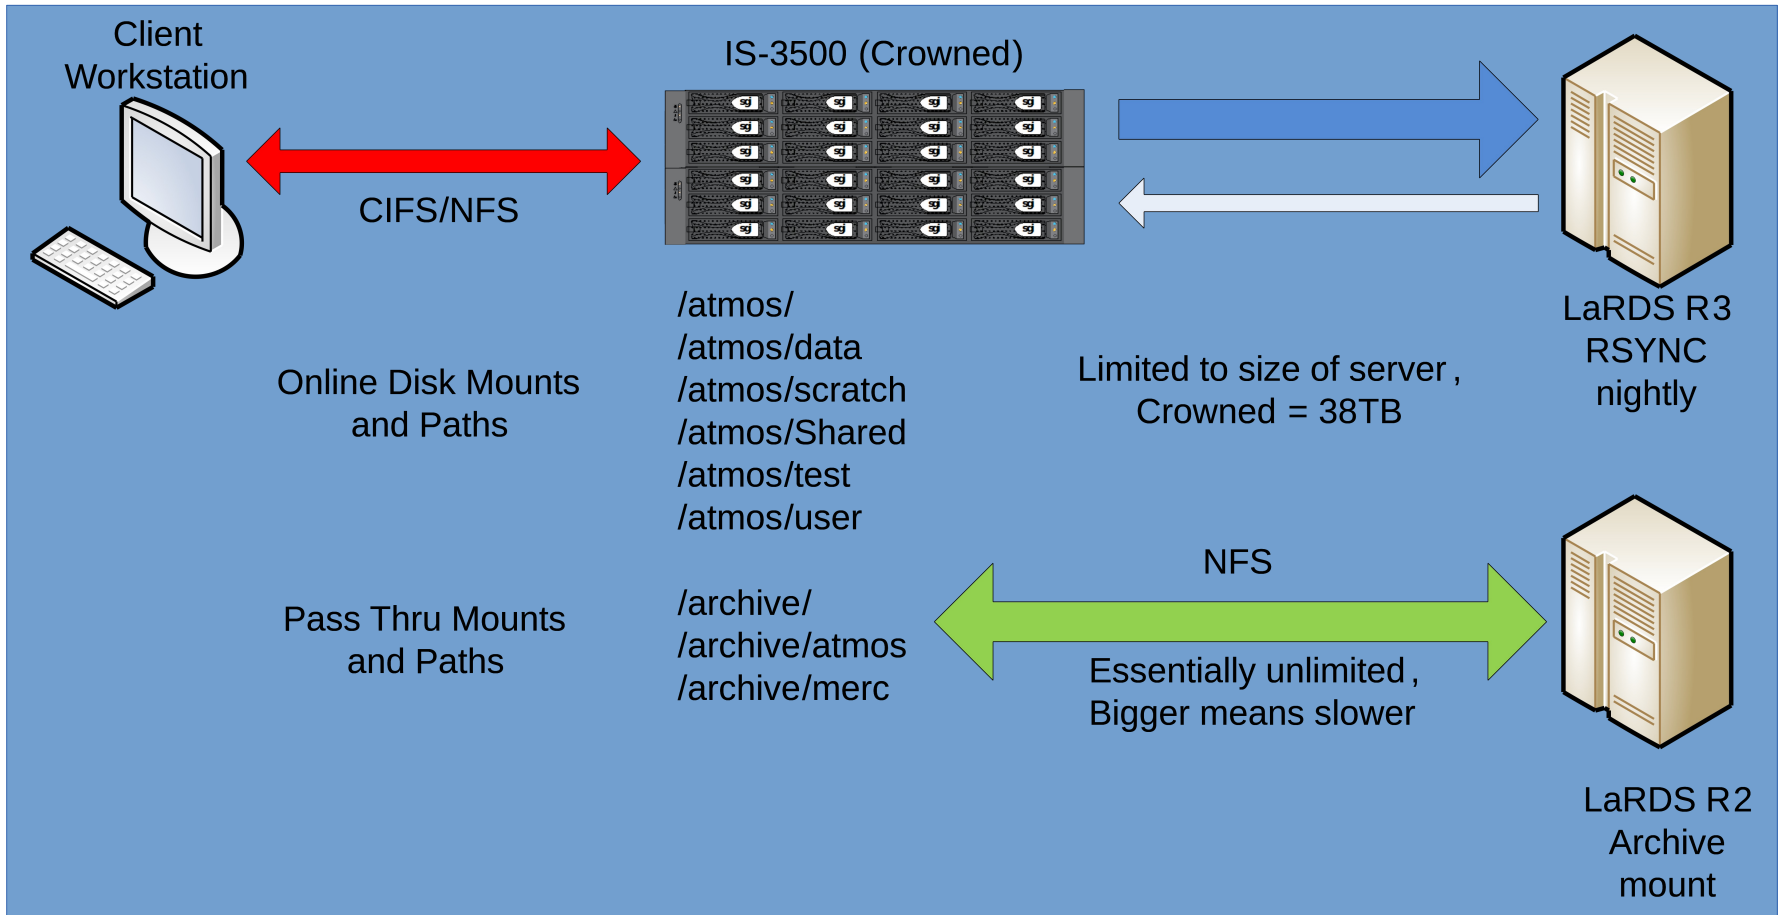
	Cur:	Min:	Avg:	Max:
/dmf-home	48.00	47.83	48.00	48.00
/dmf-journals	24.00	24.00	26.40	28.00
/dmf-spool	39.00	39.00	40.33	42.00
/dmf-tmp	19.33	16.98	26.15	50.00
/mnt/hpc	88.00	80.07	84.21	91.00
/mnt/a1	5.00	4.99	5.00	5.00
/mnt/r0	38.00	37.94	38.00	38.00
/mnt/r1	67.00	66.90	67.00	67.00
/mnt/r2	84.00	82.87	83.39	84.00
/mnt/r21	84.00	83.87	84.00	84.00
/mnt/r22	88.00	82.00	84.62	88.00
/mnt/r23	1.00	998.48m	1000.00m	1.00
/mnt/r3	8.00	7.99	8.00	8.00
/mnt/r31	72.00	69.00	75.71	88.00
/mnt/r32	36.00	30.00	36.42	53.00
/mnt/r33	74.00	73.87	74.00	74.00
/mnt/r34	83.00	79.00	82.66	91.00
/mnt/r35	88.00	81.00	84.78	91.00
/mnt/r36	81.82	80.00	88.18	91.00
/mnt/r37	4.00	3.99	4.00	4.00
/mnt/r38	57.00	51.89	56.85	62.00
/mnt/rdb	5.00	4.99	5.00	5.00
/home	17.00	16.96	17.00	17.00
/	7.00	6.99	8.17	57.50
/var	62.00	61.82	62.00	62.00
/boot	35.00	34.92	35.00	35.00
/mnt/backup2	78.00	77.00	80.73	84.93
/mnt/backup1	78.00	78.00	82.52	87.00
/dev/shm	0.00	0.00	0.00	0.00
/dev	1.00	998.08m	999.99m	1.00

Last update: Mon Nov 25 15:20:19 2013

LaRDS Cache Stabilizing Strategies

- Grid is largest LaRDS usage and still growing
 - Grid scratch moving to CEPH/Swift with NeCTAR
 - LaRDS share used for grid user home directories
- Staging posts used to move load
 - IS-3500 and QNAP used to implement effective always online
 - IS-3500 translating NFS/DMF to end user CIFS to give access to “Archive” share

Archive share has easier growth path



Issues:

- NFSv3 only supports 16 groups per user
 - Majority of end users unaffected
 - Admin class user first to be affected, but mainly due to share testing and config time
 - Perhaps move to NFSv4 or just give up on dual NFS/CIFS access model
- Path for DMF status blocked through NFS to CIFS translation
 - Samba can present on-line/off-line icon but only for local disk, not a NFS backend reshared
 - Linux STAT command reveals what is offline by blocks vs bytes (even through CIFS)
- Vendor has update patches and reconfiguration recommendation
 - Not all IS-3500s are at SP2
 - Patches have to be tested, then can be scheduled for all like machines
 - Testing reconfiguration for SMB2 is more tricky, but must be done.
 - Expansion of disk needed, but systems to be retired 2015.