

NCI Storage Infrastructure

Ahmed El Zein

Manager, Data Storage Services, NCI



- Management thought it would be a good idea if the ANU IT Services and the NCI had a joint tender and tried to merge the data storage teams into a single team.
- The end result has a lot of complexity that could be avoided.
- However, we did get some good hardware and are in the process of simplifying our overall infrastructure.

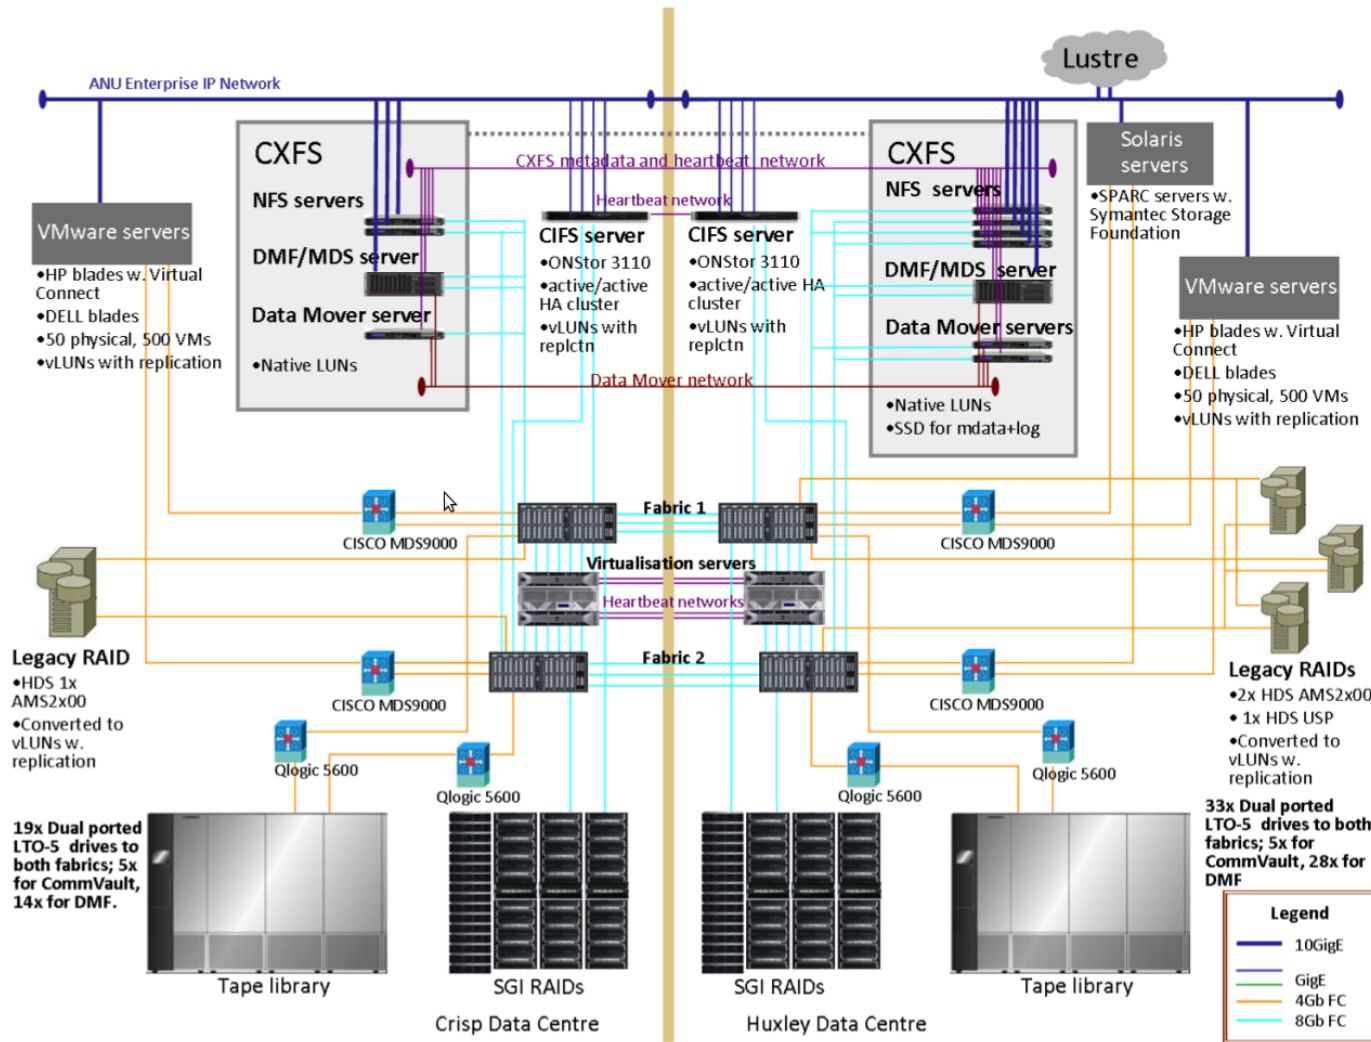


Figure 5 ANU-NCI Data Cloud

sgi[®]DMF

BAD



4x Qlogic 9200: 416 8Gbps FC ports

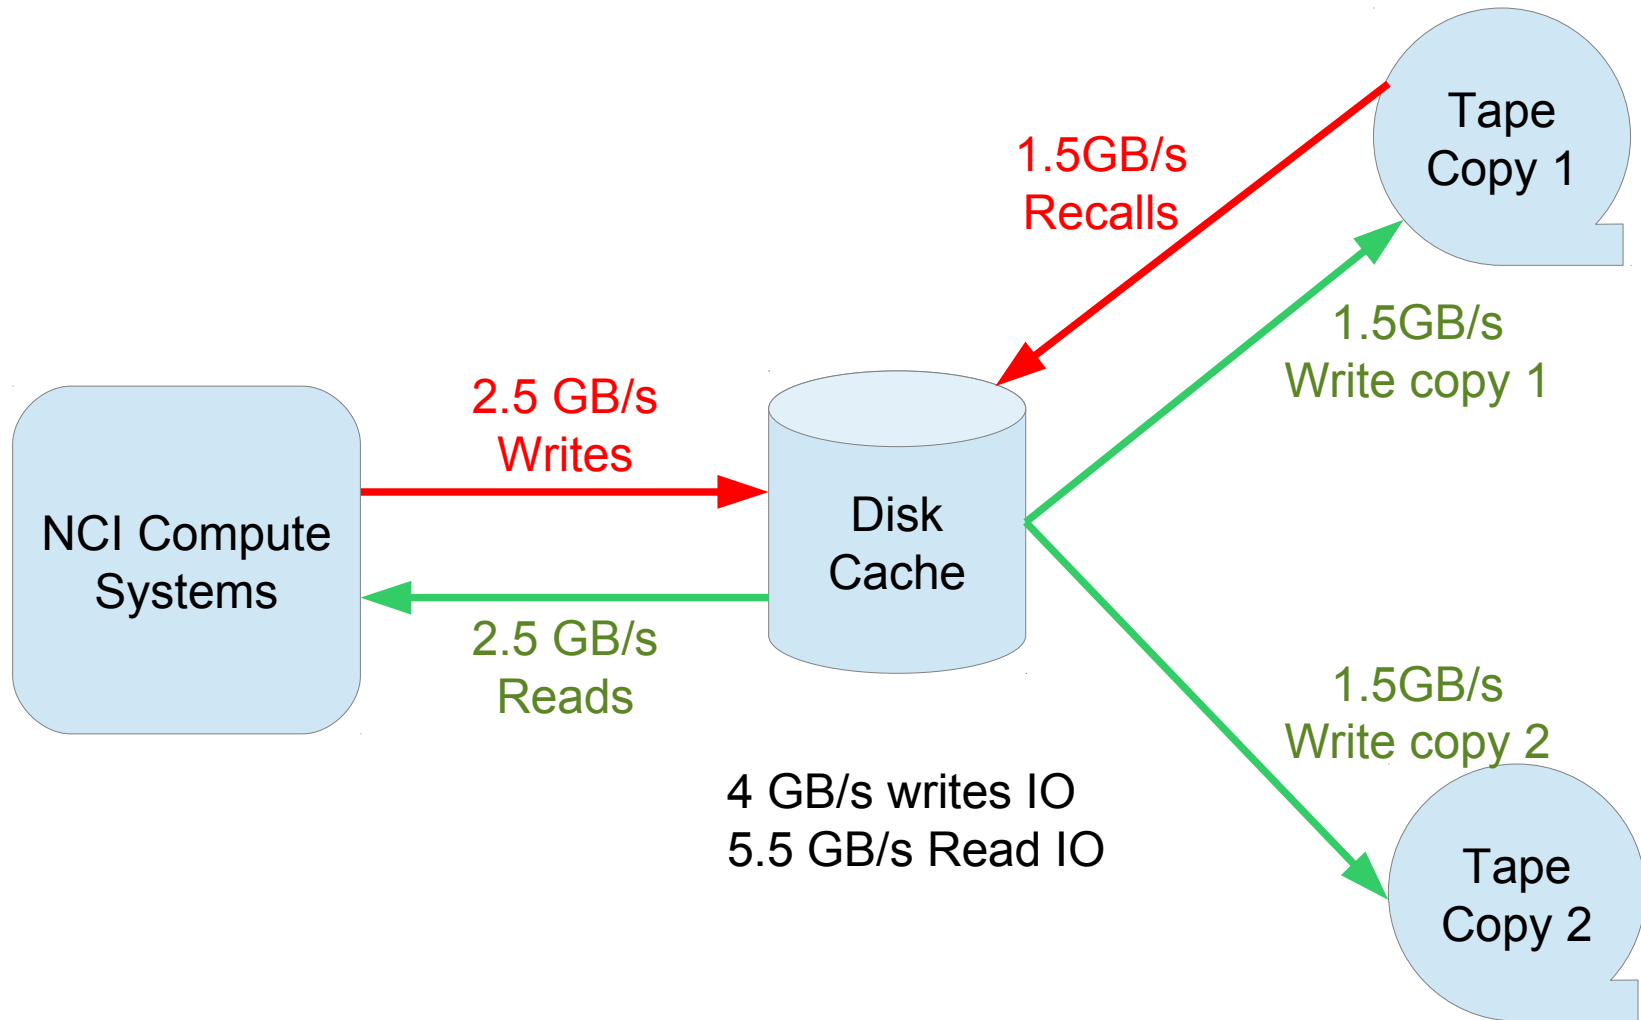


2x Spectra Logic T950: 42 Tape drives



11x SGI IS4600: 5280, 2TB SATA Disks

- The NCI was looking for a system that would be able to sustain 1.5GB/s duplex streams to tape.
- Should also be able to sustain 2.5GB/s Bursts to the HSM.



- Primary Site
 - DMF Server
 - 3 data movers
 - 4 edge nodes
 - 28 tape drives
- Secondary Site
 - DMF Server
 - 2 data movers
 - 2 edge nodes
 - 14 tape drives

- We have 3 main file systems
 - We had another 3 smaller ones that got merged into the main 3.

File system	/massdata	/projects	/rsync
Purpose	Archival	Online	Backups
# files	47.7 Million	42.8 Million	90 Million
Disk Cache	350 TB	1.6 PB	59 TB
Stored Data	1.8 PB	1.5 PB	73.7 TB
Average file size	38.8 MB	36.3 MB	N/A
Median file size	192 KB	77.3 KB	

Move to Lustre

Migration In progress

BAD

OK

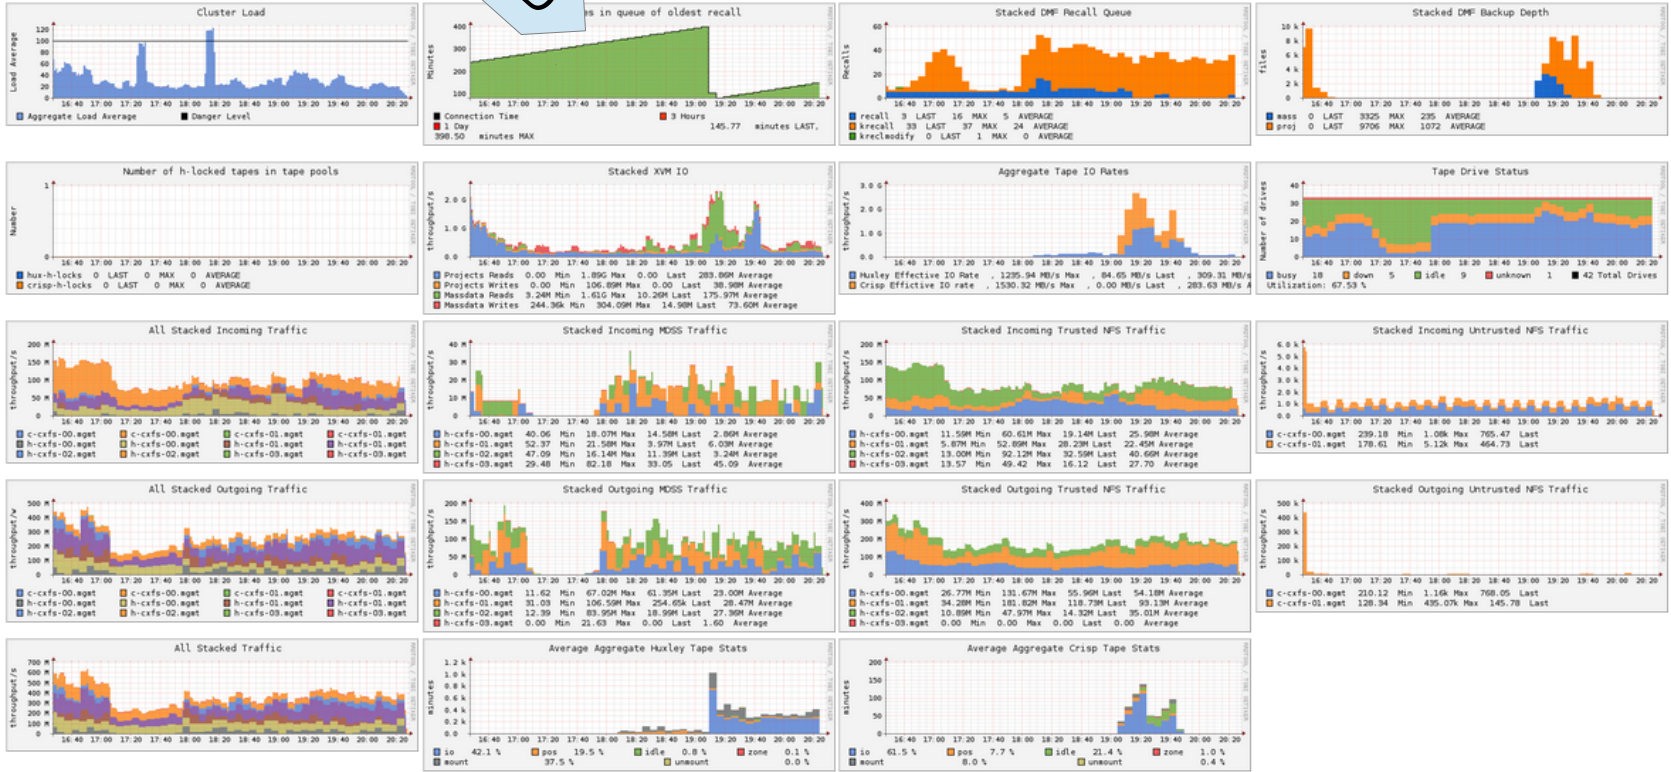
- DMF FTP back-end is quite a good solution.
- We had a lot of slowdowns from all the small files.
- Jason Ozolins will be talking about this in detail (as well as some other lessons learned)

- Not enough time to familiarise ourselves with DMF. System delivered on the 22nd December and started production use on 18th January.
- Large number of Spectra Logic robotic failures.
- DMF HA false positives... turned off Linux HA.
- Large amount of instability in CXFS edge servers... fixed now!
- I hate SLES!
- MDSS still not DMF aware... we are testing this now. Talk on this later.
- Not easy to “see” the status of the DMF environment


- We wanted a dashboard to quickly evaluate the status of the DMF environment.
- Nagios keeps history and creates graphs
- Still figuring out what how to turn all the data into information and what data should be combined to create useful graphs.

BAD

4 hour resolution [Nasion Tactical Overview](#) [Nasion](#)



Capacity Graphs (one month resolution)



We are very interested in any advice
you might have.