

QMUL – Lustre Storage

Christopher J. Walker

Queen Mary, University of London

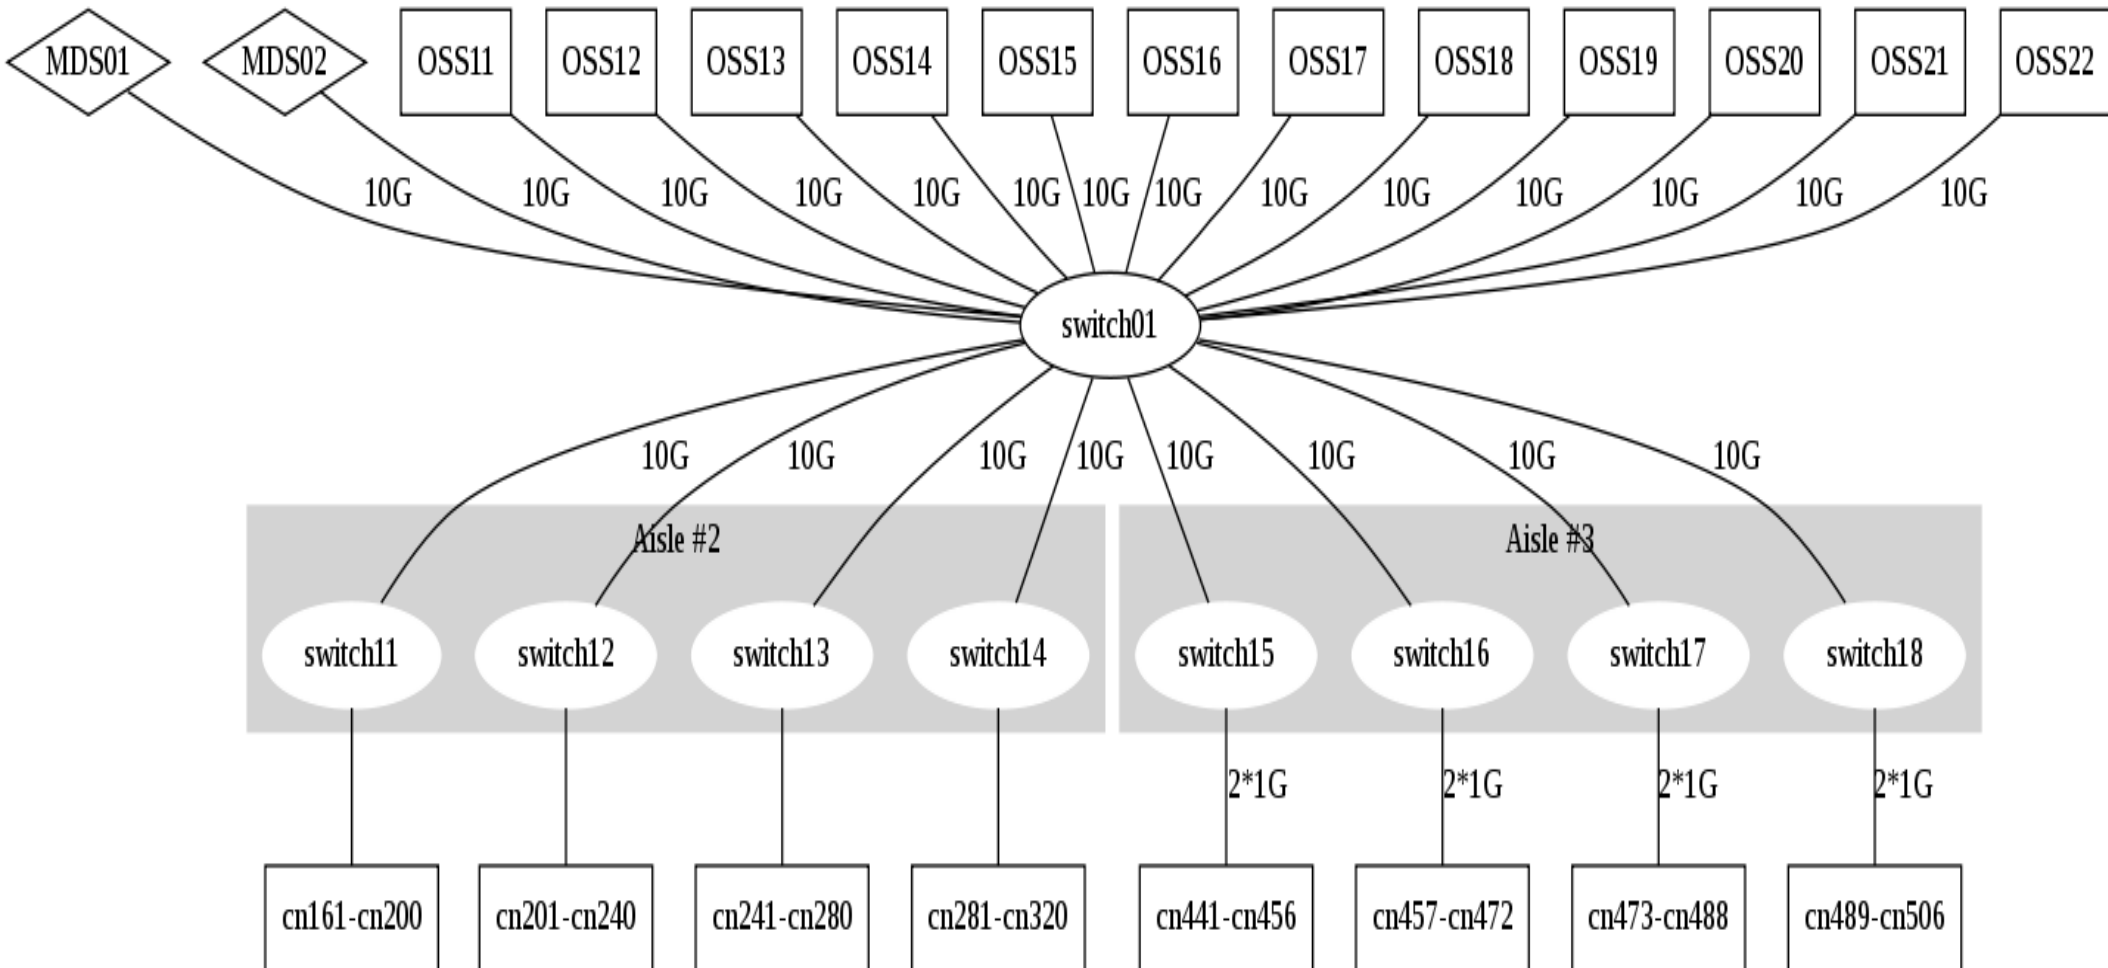
Overview

- Design
- Network
- Hardware
- Performance
- StoRM
- Conclusions

Design criteria

- Storage provided
- Enough performance
 - Storage bandwidth matches compute bandwidth.
- Cost!!!
 - Lots of Storage
 - OK performance

Old QMUL Network

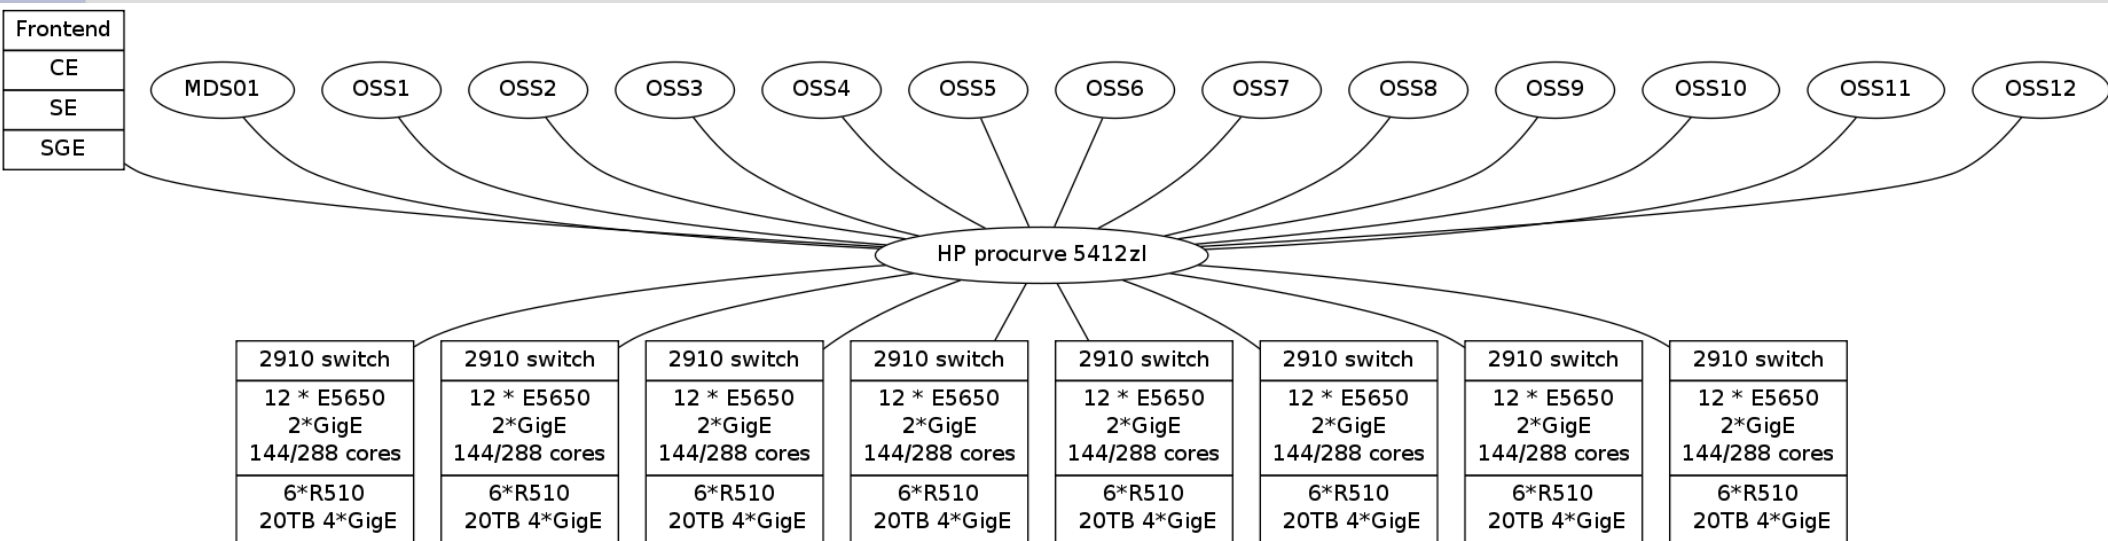


Lustre “Brick”

- 2900 Switch
 - 48 ports (24 storage, 24 compute)
 - 10Gig uplink
- 6 * Storage nodes
 - R510
 - 4*GigE
 - 12*2TB disk (20TB)
- 12 * Compute node
 - C6100 (contains 4 compute nodes)
 - 2*GigE
 - 24Gig RAM

New QMUL Network

- 48 *1Gig per switch
- 24 – storage
- 24 – CPU



Hardware

- 56* R510
 - 12*2TB SATA disk
 - H700 RAID controller
 - 12 Gig RAM
 - 4*GigE
- 4 *R510
 - 12*2TB SATA disk
 - 2*10GigE



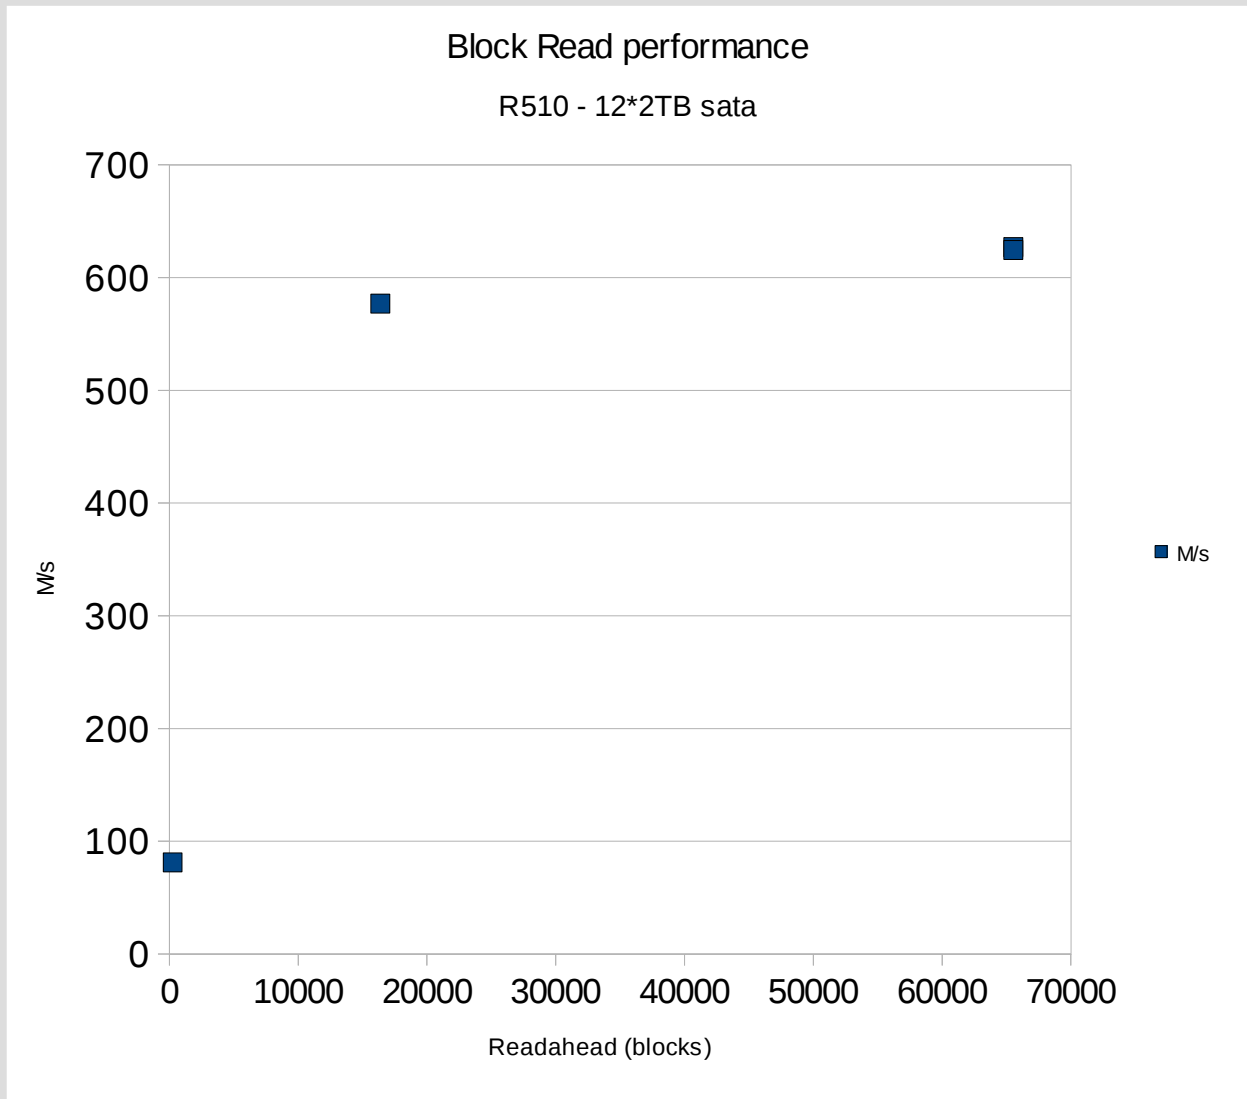
Storage Burn in

- Hepspec06
 - 2 machines low
 - (power saving mode)
- RAL disk thrashing scripts
 - 1 Backplane failure
 - 2 disk failures
- 10Gig cards in x8 slots

Storage Performance

R510

- Disk
 - 600M/s
- Network
 - 4Gbit/s
-



StoRM

- 1.4 (SL4) QMUL's "production SE"
 - Se03.esc.qmul.ac.uk
 - No checksums
 - Limited permissions
 - Publishing format problems
- 1.5 (SL4) WLCG baseline
 - Se01.esc.qmul.ac.uk
 - Reporting of space used incorrect on Lustre
- 1.6.2 (SL5) Staged Rollout
 - Se04.esc.qmul.ac.uk
 - Reporting of space used incorrect on Lustre
 - But I can run du once a day

Further Work

- StoRM 1.6.2
 - Next week?
- New Lustre
 - Benchmarks
 - MDS failover
 - Transition
- Reliability
 - Lcg-cp using cp
 - Grid as tape backend
 - Storm syncat dumper
 - Automate?