

QMUL: Site Report

Christopher J. Walker

Queen Mary, University of London

Overview

- Hardware
 - WN
 - Storage
- Software
- STEP09 experiences
- Conclusions

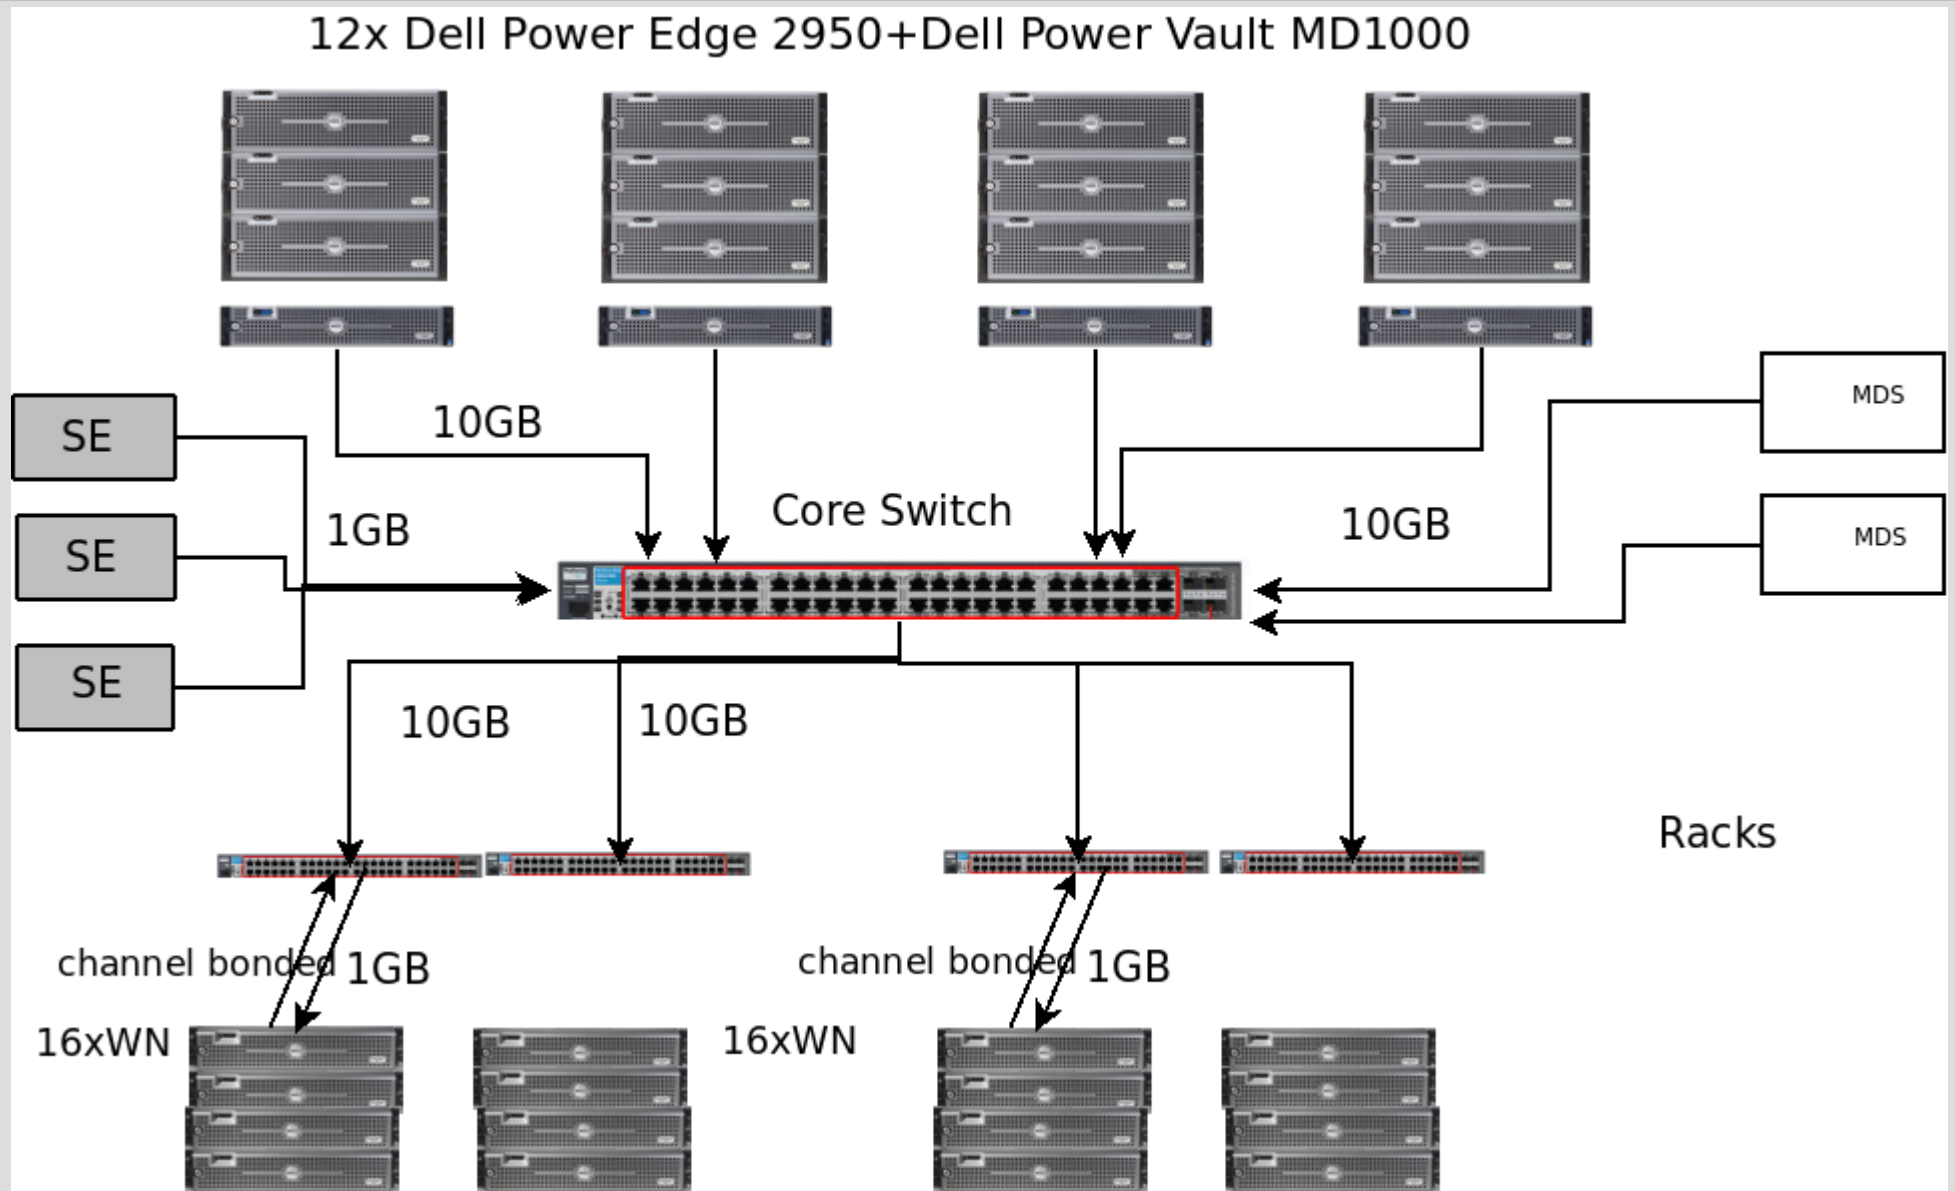
Hardware

- Worker nodes
 - 64* 2*4core Xeon E5420 (2.5GHz) 8 GB RAM
 - 2 * 2*4core Xeon E5420 (2.5GHz) 32 GB RAM
 - 160 *4core Opteron 270 4GB RAM
 - 128*2core(2 thread) Xeon (2.8GHz) 2GB RAM
- Total 1424 cores

Storage

- Lustre
 - 2*MDS
 - Failover
- Lustre_0 (290TiB)
 - 12 OSS
 - 30*1TB
 - Raid6 + 1 hot spare
 - 10GigE
- Lustre_1 (30TiB)
 - 4*OSS
 - GigE

Network



CE/SE

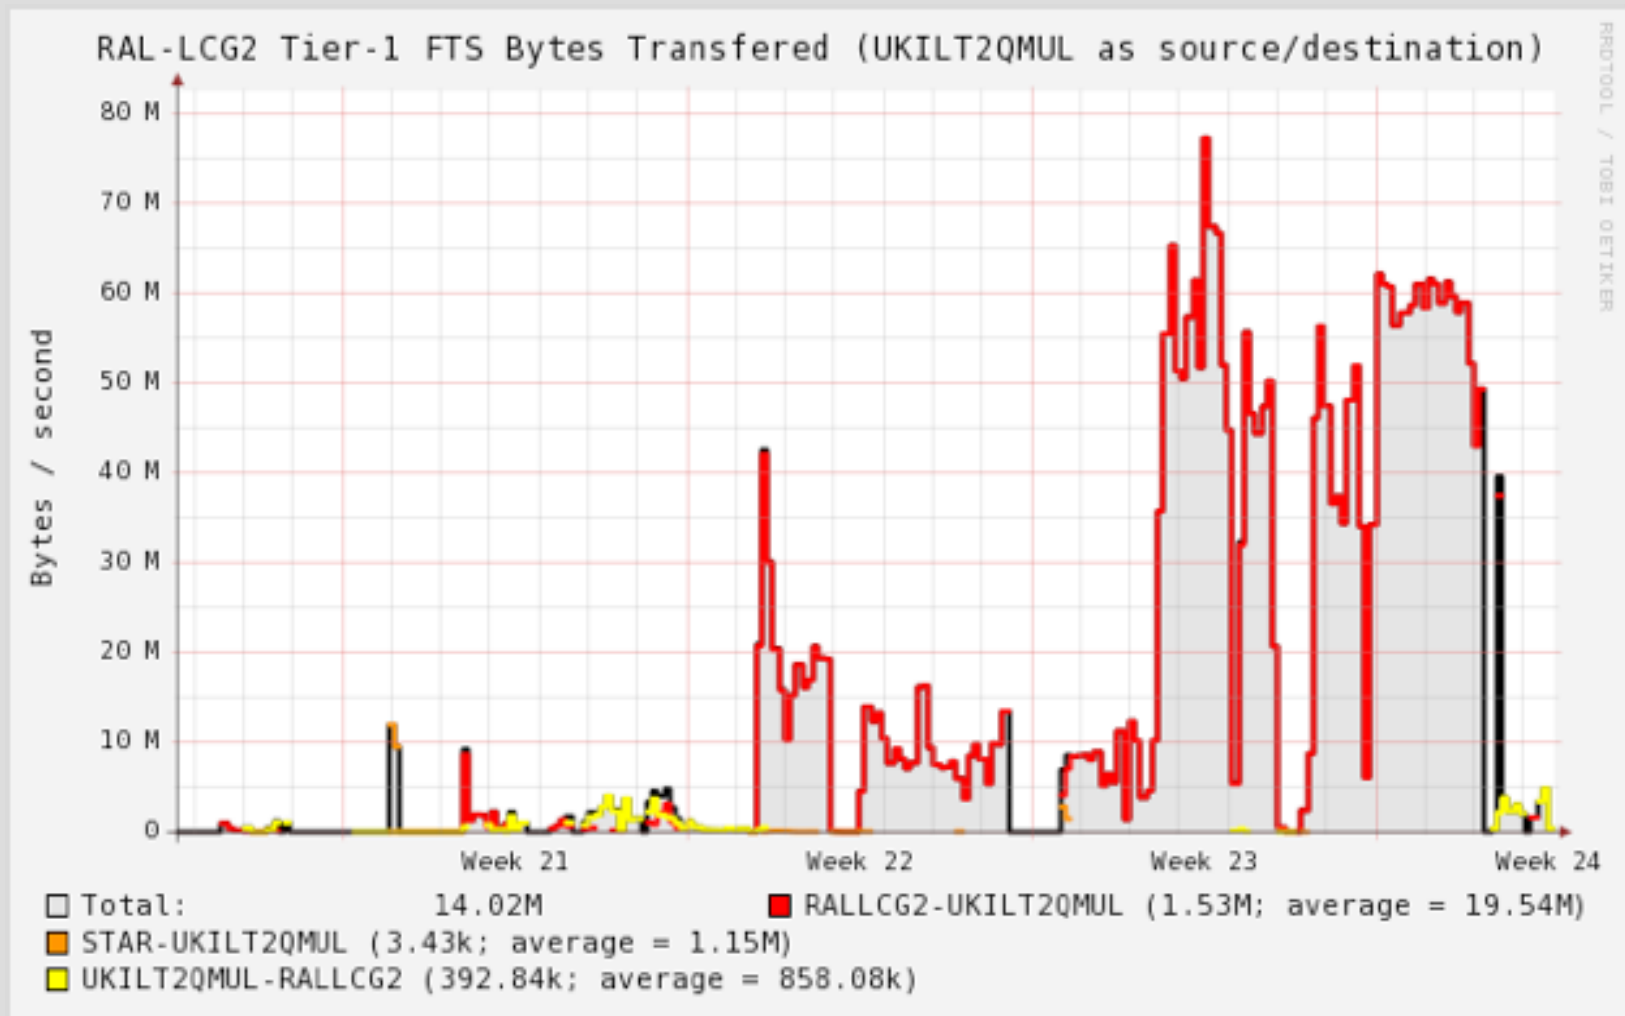
- The Past
 - SRM
 - se01 Storm (1.3.20)/Lustre
 - se02 DPM/Lustre
 - CE
 - ce01
 - ce02
- The Future
 - se03 Storm (1.4) / Lustre
 - ce03

Step 09 - Leadup

- New CE
 - Ready Mid May 2009
- New SE
 - Installed 27 May
 - Consistent user mapping 29 May
 - Announced 30 May
 - FTS 1 Jun
- Decide to try new SE for Step09

Step09 - Start

- Identify and eliminate network pinch points



Step 09

- Major Problem
 - Analysis jobs not using `file://` protocol
 - 1GigE limit
- Other Issues
 - Data transfers taking priority over MC
 - Lustre filesystem crash
- Monte Carlo
 - CMS 2.7×10^8 CPU seconds (89% success)
 - Atlas
 - 5 TB data
 - 28 TB MC

Conclusions

- STEP09 very worthwhile
 - Lack of [file://](#) Protocol support a major issue
 - Network pinch points identified and resolved
 - 33TiB of data received
 - Improved monitoring
- QMUL in good shape

Step 09 Outcome

- 5TiB ATLASDATA
- 28TiB MCDATA
-