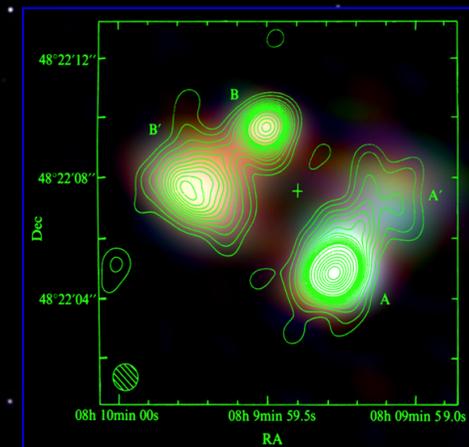
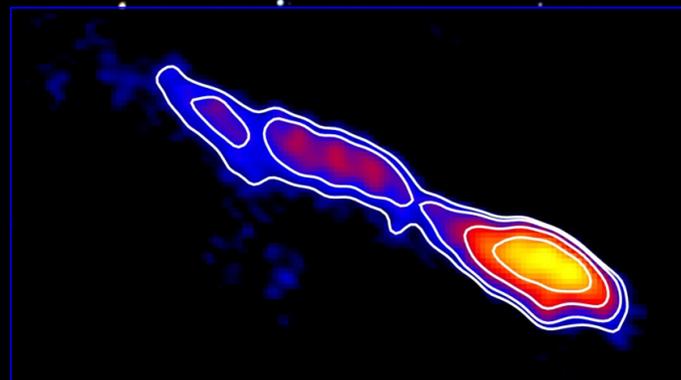
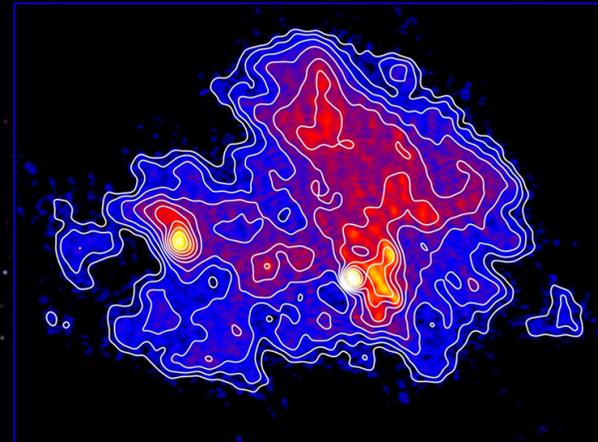
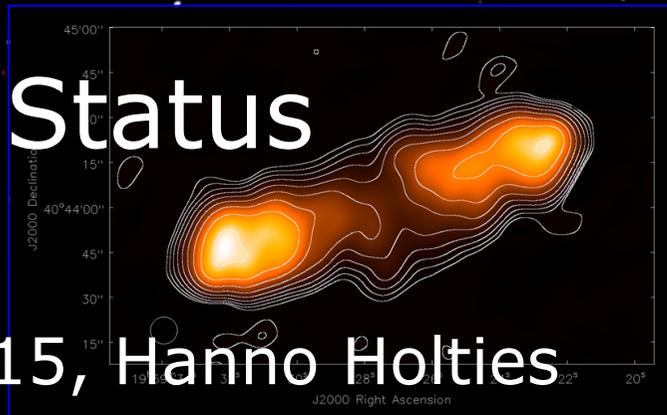


CEP4 Status

LSM

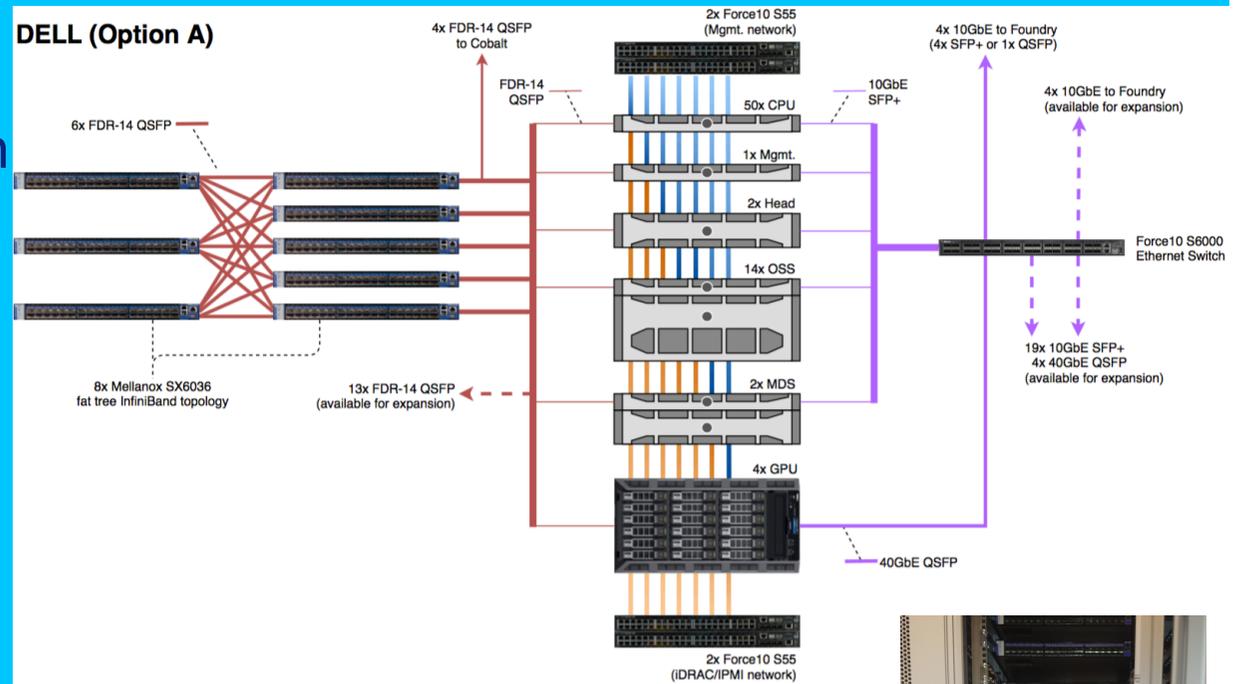
9/12/2015, Hanno Holties



CEP4 Tender



- Tender concluded
- Selected DELL option
- Separate compute and storage
- 50 Compute nodes
- 4 GPU Nodes
- 2+ PB Lustre PDFS



Compute cluster

ASTRON

Compute nodes

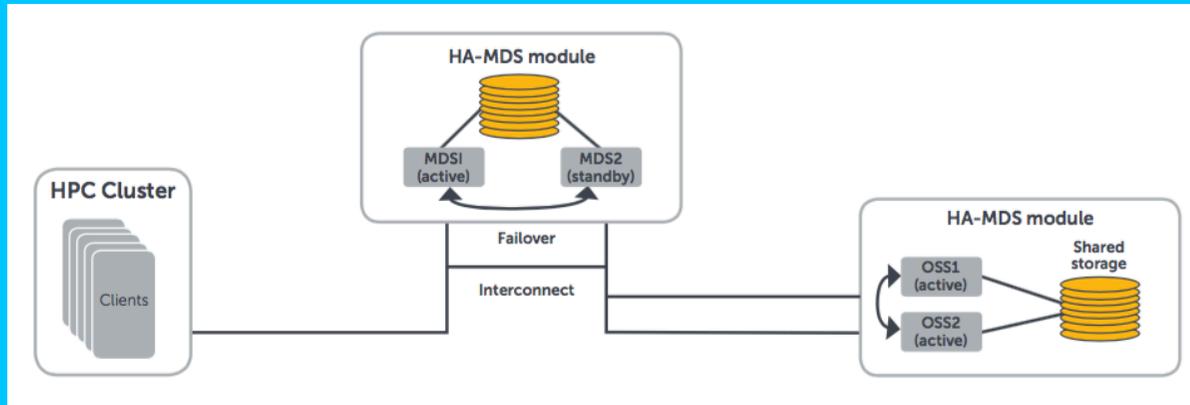
- 50x DELL R630
- Dual Intel Xeon E5-2680v3 (per server 24 cores @ 2.5 GHz)
- 256 GB memory
- FDR 56 Gbps infiniband interconnect

GPU nodes

- 4x DELL T630
- Dual Intel Xeon E5-2603v3 (per server 12 cores @ 1.6 GHz)
- 320 GB memory
- 4x Tesla K40C GPU
(per GPU 2880 cores, 4.29 Tflops SP, 1.43 Tflops DP)
- 2x 6TB internal storage

Lustre File System

ASTRON



<http://i.dell.com/sites/content/business/solutions/hpcc/en/Documents/Lustre-HPC-Whitepaper-10082011.pdf>

DELL Lustre HPC storage

- Architecture: white paper Cambridge/DELL
- 18x DELL R630 Storage server
- 18x DELL MD3460 Storage Array
- 60x 4TB disks each
- Total usable capacity up to 3PB



Impact on LOFAR

Lustre: Shared file system

- Improved scheduling efficiency
- Main challenge: Read/write IO throughput to many clients
 - Working with DELL to optimize for LOFAR
- Impact of filesystem failure
 - Lustre proven solution
 - No single point of failure
 - Looking into impact cache settings; relation performance

Slurm: Batch scheduling

- Improved scheduling efficiency
- First step: run pipelines as single job
- Implementing scheduler as a service

Planning

Current focus

- Functional cluster with most systems are in place
- Impact analysis proposed solution DELL underway
- Additional hardware for 80 GBps balanced read & write TBD

Next steps

- Installation new Lustre version by DELL
- Finish acceptance
- Handover December/early January
- Commissioning February/early March
 - Assuming no further issues from acceptance

Questions?