ASTRON

Netherlands Institute for Radio Astronomy

CEP4 Update

October 26, 2016

The road we've traveled: Timeline

CEP4
Timeline

ASTRON

- **August 2015**: Target system delivery
- **October 2015**: Realized hardware delivery; Start acceptance
- **January 2016**: Agreement on performance results: 4 extra storage servers needed; Start system configuration by CIT
- **April 2016**: Additional storage servers delivered; CEP4 support in specification & control systems
- **May/June 2016**: CEP4 support in pipelines, Cobalt, new scheduler; Internal CEP4 training workshops; CIT role changed from installation to support; Start commissioning
- **July 2016**: CEP4 clean-up tool delivered; Problematic data server repaired
- **August 2016**: Last activity DELL (configuration)

The bare facts: System description

Hardware

ASTRON

- 2 head nodes
  - 3.6 TB local storage
  - Failover configuration
- 3.1 PB global storage
  - Available on all nodes + lexar003/4 (LTA ingest)

System	Compute	Memory
head01.cep4 head02.cep4	12 cores @ 1.6 GHz	128 GB
50x cpuXX.cep4	24 cores @ 2.5 GHz	256 GB
4x gpuXX.cep4	16 cores @ 2.4 GHz 4x Tesla K40C	320 GB

Performance

ASTRON

- Observational throughput Cobalt -> CEP4:

Mode	CEP2	CEP4
BF superterp	<37 Gbps	>110 Gbps
Correlator	~40 Gbps	~69 Gbps

- Pipeline throughput TBD. CEP2 -> CEP4:
  - Batch scheduling -> better utilization
  - 4.8x more FLOPs
  - 0.8x memory bandwidth
  - 1.5x disk speed (but higher latency!)


A bumpy ride:
Challenges encountered

CEP4
Challenges

ASTRON

- CEP4 is not a drop-in replacement for CEP2
  - New technologies introduced
  - Development & testing require representative environment
- CEP4 is not a standard HPC cluster
  - Combination of real-time & offline processes
  - Part of an operational instrument

CEP4
Challenges encountered & addressed

ASTRON

- IO performance acceptance
  - Vendor initial design based on reference implementation
  - Next time: representative LOFAR benchmark?
- Lustre behavior different from local disks
  - Higher throughput but also higher latency: random access significantly slower (e.g. generation inspection plots)
  - Large scale filesystem operations expensive: introduced “Robinhood” file system monitoring tool but synchronization challenging although it is an accepted tool. LOFAR file generation mechanisms atypical?
- Technology maturity
  - Bug in Slurm/Docker/Kernel cgroup handling

CEP4
Challenge remaining

ASTRON

- Cobalt data loss
  - Data loss observed during start-up of imaging observations and throughout high throughput observation (LOTAAS)
  - Related to simultaneous offline process activity
  - First update OS, docker & configuration (CIT)


A bright future:
New features

CEP4
Brings us many new things

ASTRON

- Lustre: global filesystem
  - Easier data management & improved reliability/flexibility
- SLURM: batch scheduling
  - Saves significant operator effort (automation & maintenance)
  - Collects job statistics: characterization & optimization
- Docker: software configuration management
  - Improves version handling/availability (continuity, rollback)
  - Easier deployment of user-provided tools
- C++11: compatibility with latest Linux & compilers
- Supervisord: process management
- Ganglia: cluster monitoring

Monitoring CEP4 Utilization (Ganglia)

- Many metrics: CPU/network/memory/SLURM/…
- Extensive querying: per day/week/month/…, graph/cvs/json

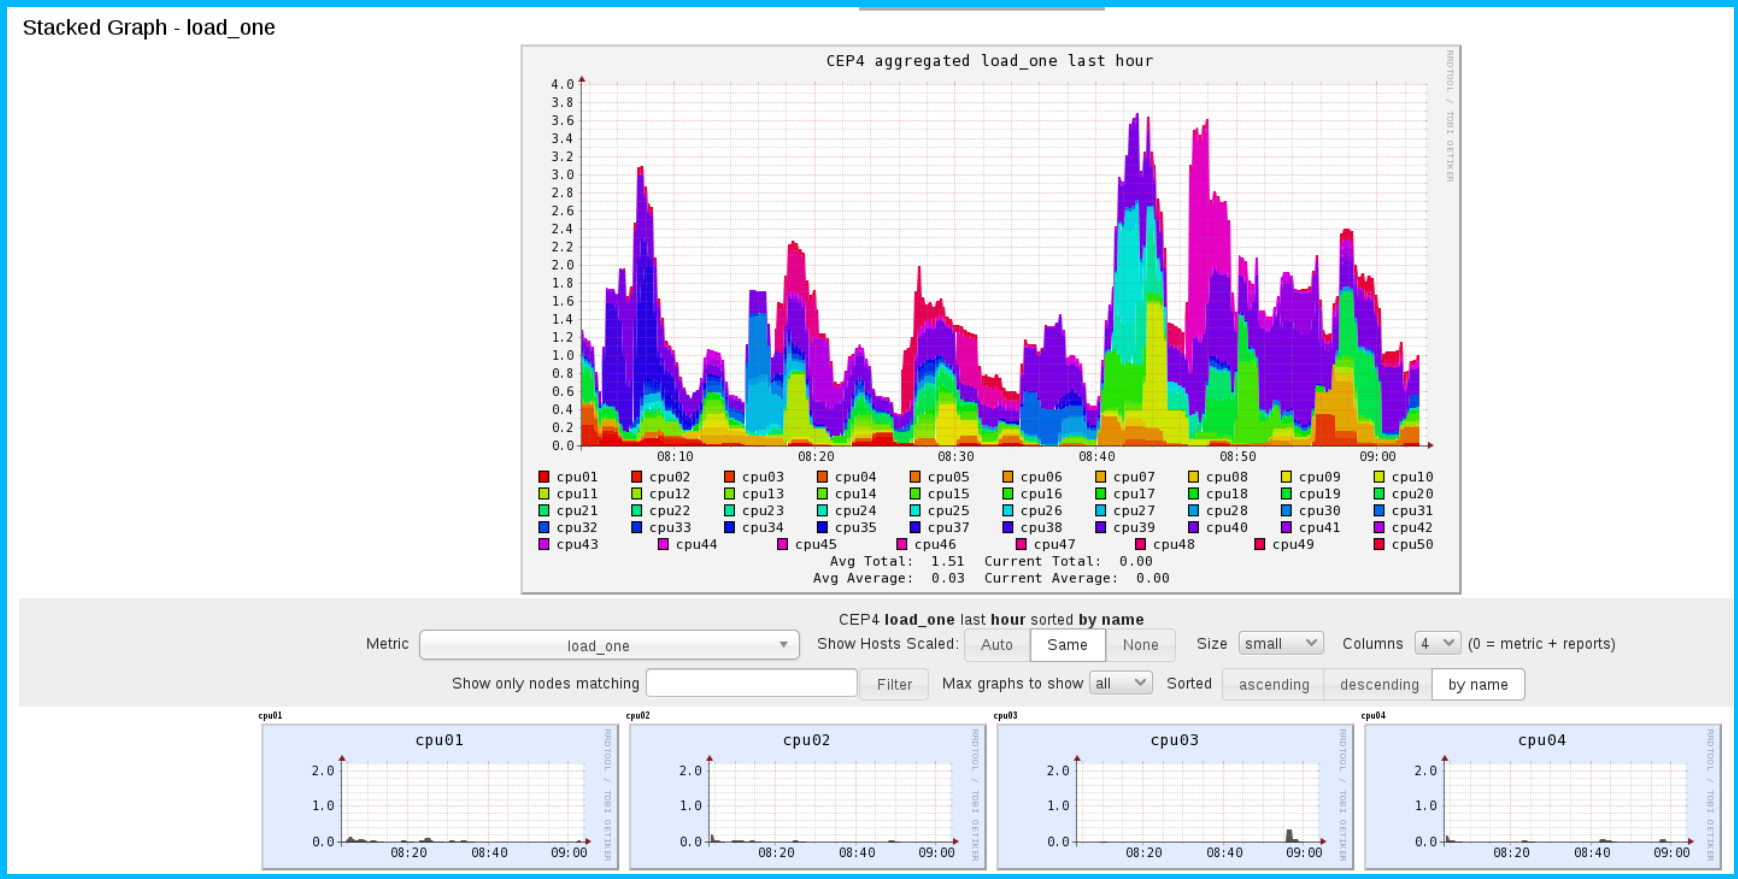

...and already set in motion before, but realized on CEP4

ASTRON

- Services Framework: more flexible interaction between components
  - Modular: easier maintenance & development
  - Open interfaces: easier access to information

...and already set in motion before, but realized on CEP4

- ResourceAssigner: a new Scheduler
  - Real-time scheduling: ready for responsive telescope
  - Web-based
  - Faster & interactive

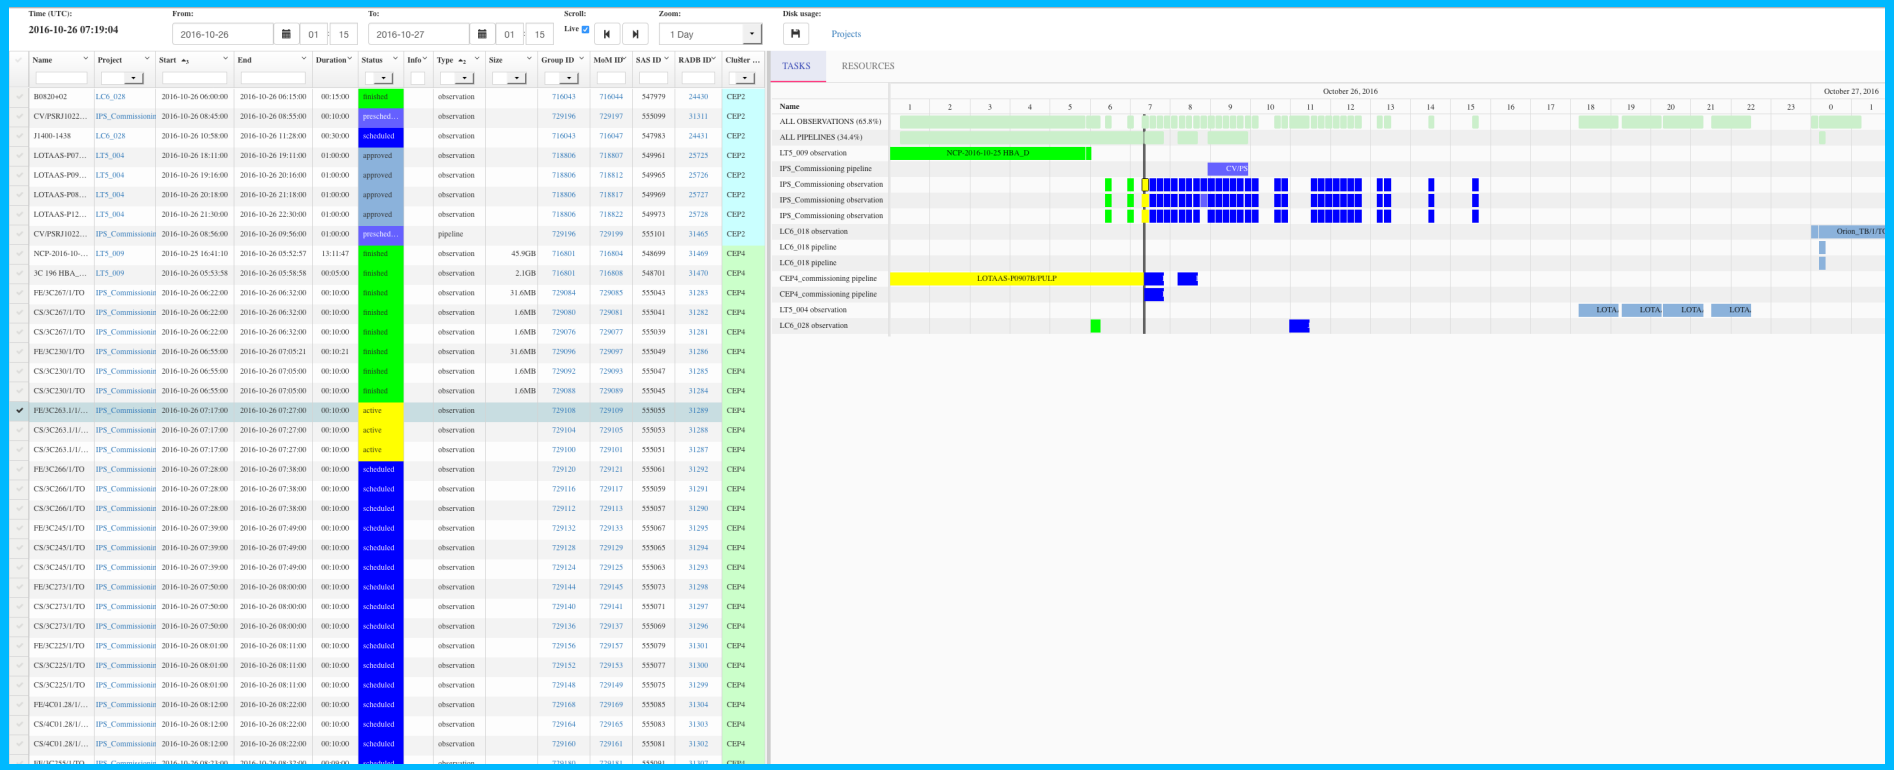


Fin
Questions?