The LOFAR phase II cluster
update

Chris Broekema

On behalf of
Harm Paas (RUG), Teun Grit, Arjen Koers (RUG) and
André Gunst

ASTRON LOFAR NWO

ASTRON JIVE Daily Image-Mozilla Firefox

File Edit View History Bookmarks Split Tools Help

http://www.astron.nl/dailyimage/

Wikipedia (en)

documents:minut... vliegengordijn.nl ... HomePage - Syste... Binsearch -- Usen... Motortransport.nl, ... Vergelijk energiet... Linux Assembly NZBIndex.nl - We i... net: Introduce rec... Astrophysics Genetic Algorithm ...

Digitenne: Ontvangst kwaliteit ver... | French Alps seen from the beak - ... | ASTRON | ASTRON | ASTRON JIVE Daily Image | nu.nl - De Jager | SpitsNieuws : Met Pasen massaal ...

ASTRON Netherlands Institute for Radio Astronomy

JIVE JOINT INSTITUTE FOR VLBI IN EUROPE

Daily Image

20-04-2011

Previous Next

Click here or on the picture for a full size image.

CEP 2.0 cluster

Submitter: Chris Broekema

Description: Today we proudly announce that the LOFAR phase II cluster is operational. Over the next couple of days we will show the first observations stored and reduced on new hardware recently installed in Groningen. We've added some 100 hybrid compute and storage nodes, containing 2448 cores, 6.4 terabyte of main memory and just over 2 petabyte of disk storage to our already impressive list of hardware. The entire cluster is in theory capable of performing just over 20.5×10^{12} floating point operations very second, four times that of the previous cluster.

This new cluster is designed to handle far greater bandwidths than the previous cluster. With this hardware we should (eventually) be able to handle up to 80 Gbps of data streaming from the Blue Gene/P supercomputer, allowing us to observe with more beams, at higher time resolution and with more frequency bands than previously possible. According to one enthusiastic pilot user we can now do with one node what until now required the entire cluster.

The picture above shows the evolution of the cluster, from empty racks to a full-fledged (and rather noisy) cluster of 100+ nodes. What can't be shown in pictures is the tremendous amount of work that was done, by personnel from both the University of Groningen and ASTRON, to get the operating system and software working in record time.

Find: broodje Previous Next Highlight all Match case

Done

ASTRON LOFAR NWO

Current status

- All nodes have been installed
- First observations have been done and reduced
  - Both correlated and beamformed data
  - Dutch TBBs can reach new cluster, Intl TBBs Groningen
- Software stack almost ready
  - Pulsar pipeline mostly operational
  - Imaging still needs a bit of work
- Reboot not always successful
  - Routing issues – solution identified
- Stability issues

Stability issues

- Last weekend ~20% of nodes failed
  - Spontaneous reboots: 14 nodes
  - Unreachable IB device: 2 nodes
  - Crashed: 5 nodes
- No obvious correlation found
  - Nodes in use / idle
  - Location in rack
  - Power Phase

Stability issues

- TARGET cluster has similar problems
- Both cluster share
  - Power (UPS)
  - Environment (Heat, airflow)
- Power stability or contamination prime suspect
  - Large scale investigation started by RUG FB
- 51 nodes identified as “Good”
  - Can be used until a solution found

ASTRON

NWO

Cluster HOWTO

- No separate compute and storage nodes
- In general: process where the data is
  - Every node has ~20 TiB storage and 24 cores
  - Although interconnected by IB – avoid moving data
- ~Two nodes dedicated as “hoover nodes”
  - Cluster nodes with only 2 TB scratch space
  - Further identical to other nodes
  - (can) mount all cluster nodes

Building - running

- Two head nodes
  - Fewer but faster cores
  - 5+ terabyte of RAID-5
- 500 GB of fast SSD based storage
  - /data/scratch

Optimal datastreams

- All traffic to and from the cluster goes through a **router node**
- There are four such nodes
- Currently two in use:
  - locus001 – locus050: router2
  - locus051 – locus100: router3
- Data from BG/P routed on a per I/O node basis
- Configuration of router nodes WIP
  - Currently at least 18.4 Gbps/router (total 73.6 Gbps)

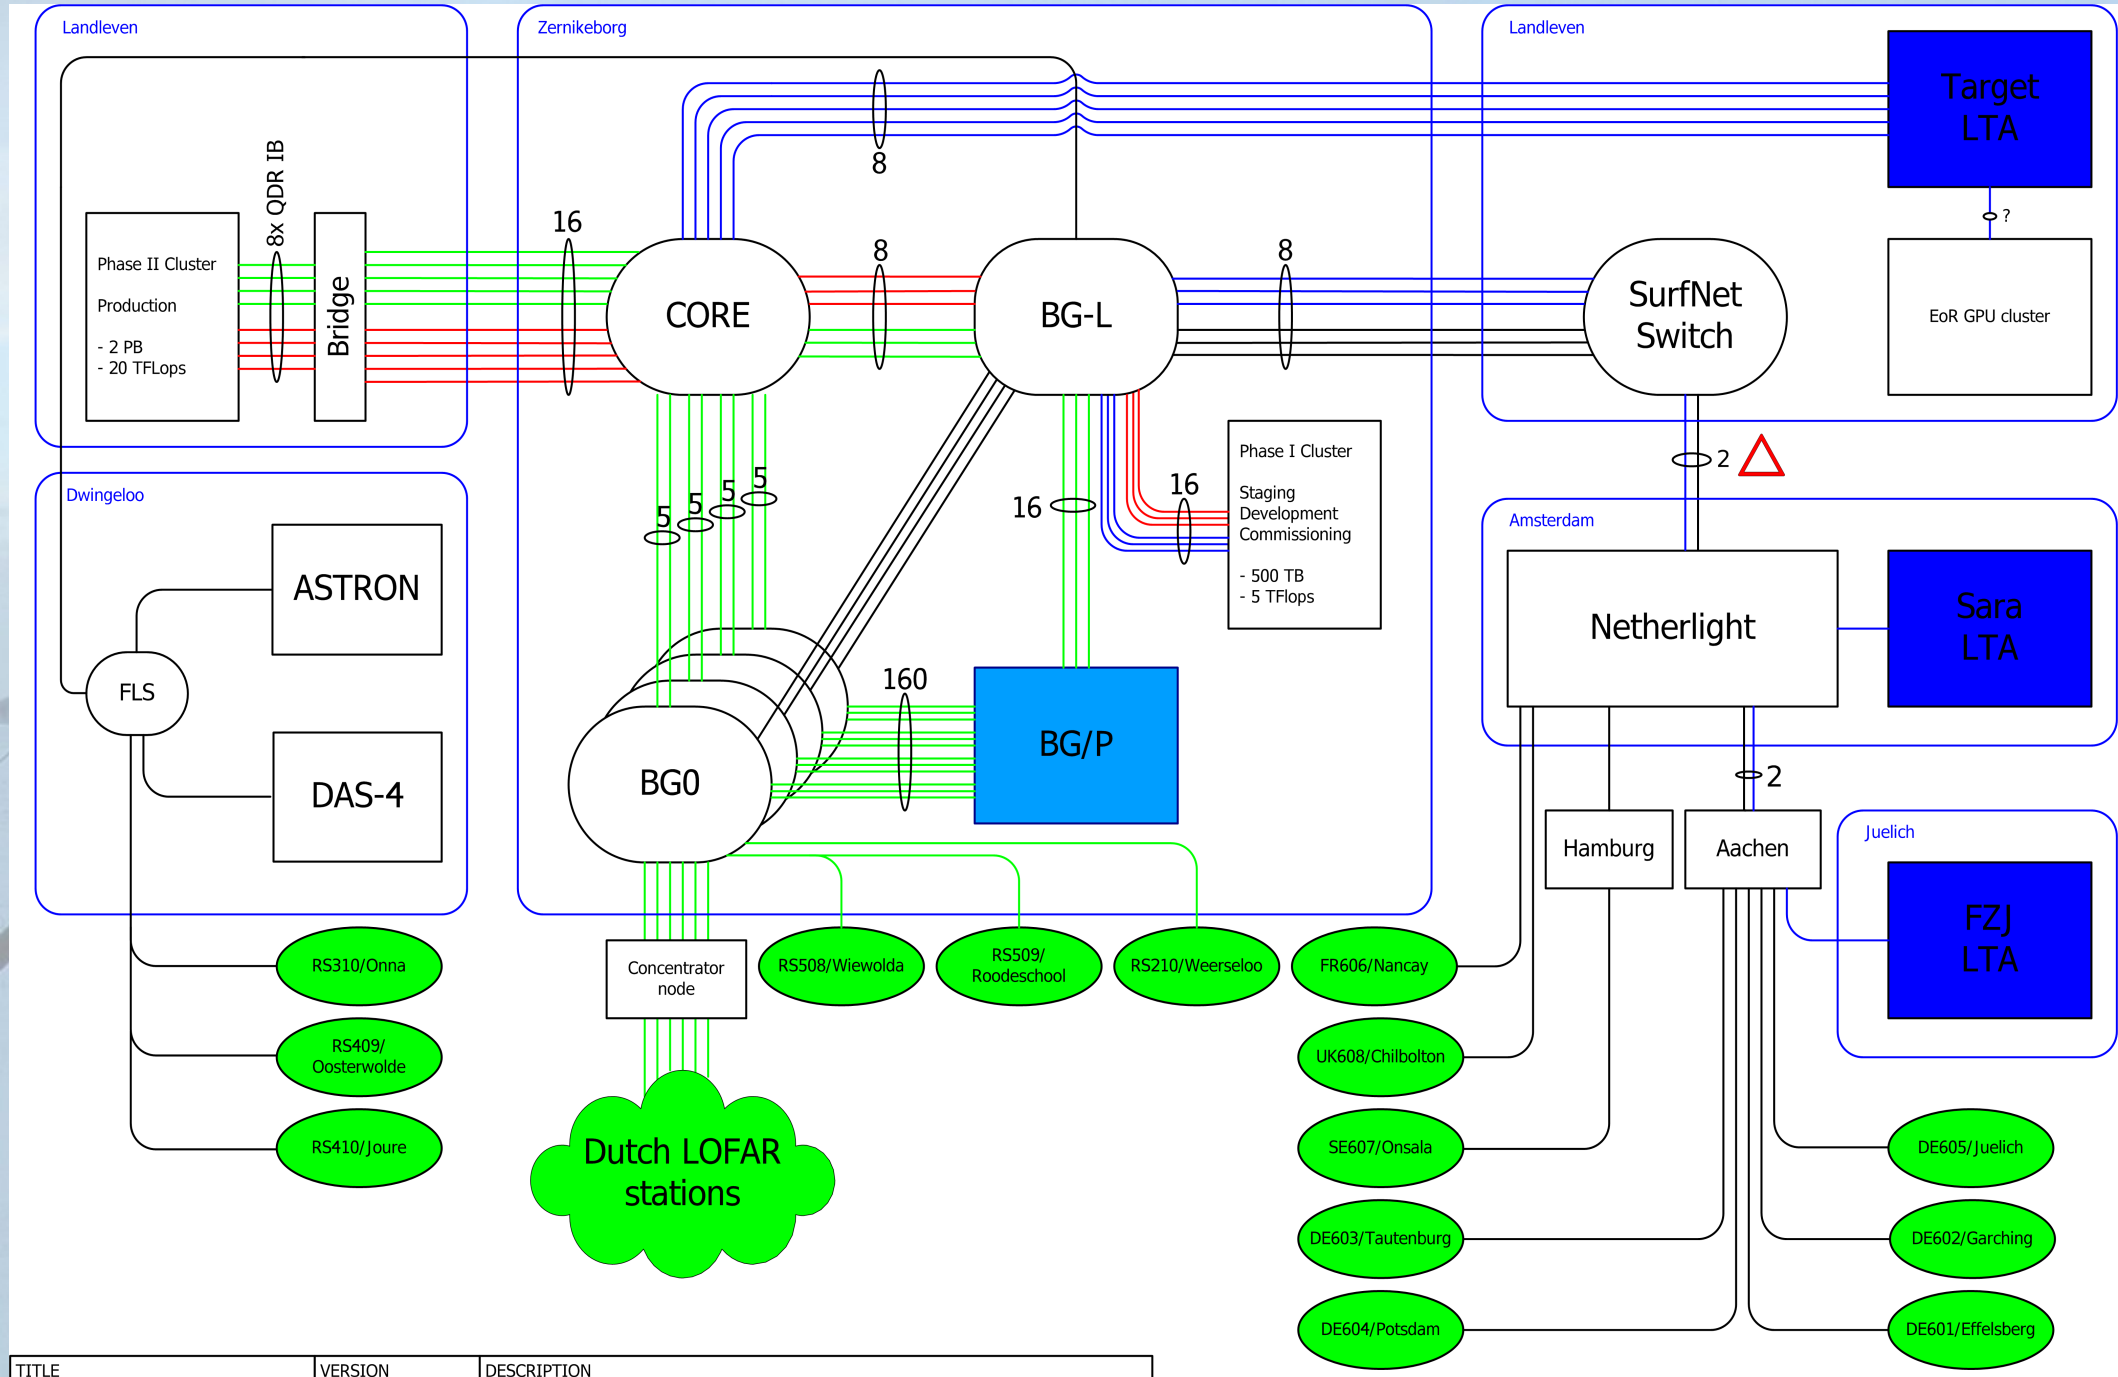

TITLE	VERSION	DESCRIPTION
LOFAR Phase 2 network	1.0	Projected LOFAR phase 2 network – work in progress
DRAWN BY	DATE	FULL FILENAME
P.C. Broekema	7/1/2011	Y:\PROJECTS\LOFAR\NETWORK\LOFAR NETWORK PHASE 2.VSD

LOFAR New Year's Resolutions

- More taart (i.e. nice results).
- More short LSM presentations about observational and reduction results.
- Communicate more between groups.

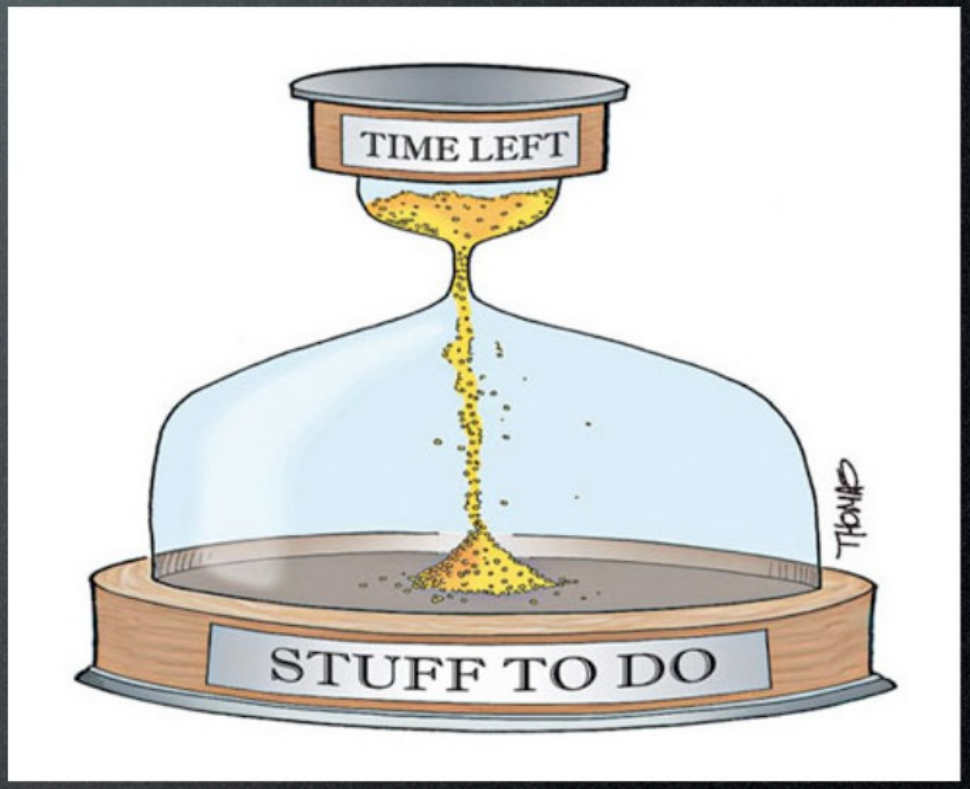

√ - Bigger basket

√ - More cake

√ - Better communication

This bunch of hardware should be able to handle 80Gbps

Proof me wrong!

ASTRON LOFAR NWO