COBALT

ASTRON

Project Update

R.J. Nijboer

LOFAR

ASTRON

COBALT

COrrelator and Beamformer Application platform for the Lofar Telescope

http://www.lofar.org/wiki/doku.php?id=cobalt:start

LOFAR

IBM Blue Gene / P

ASTRON

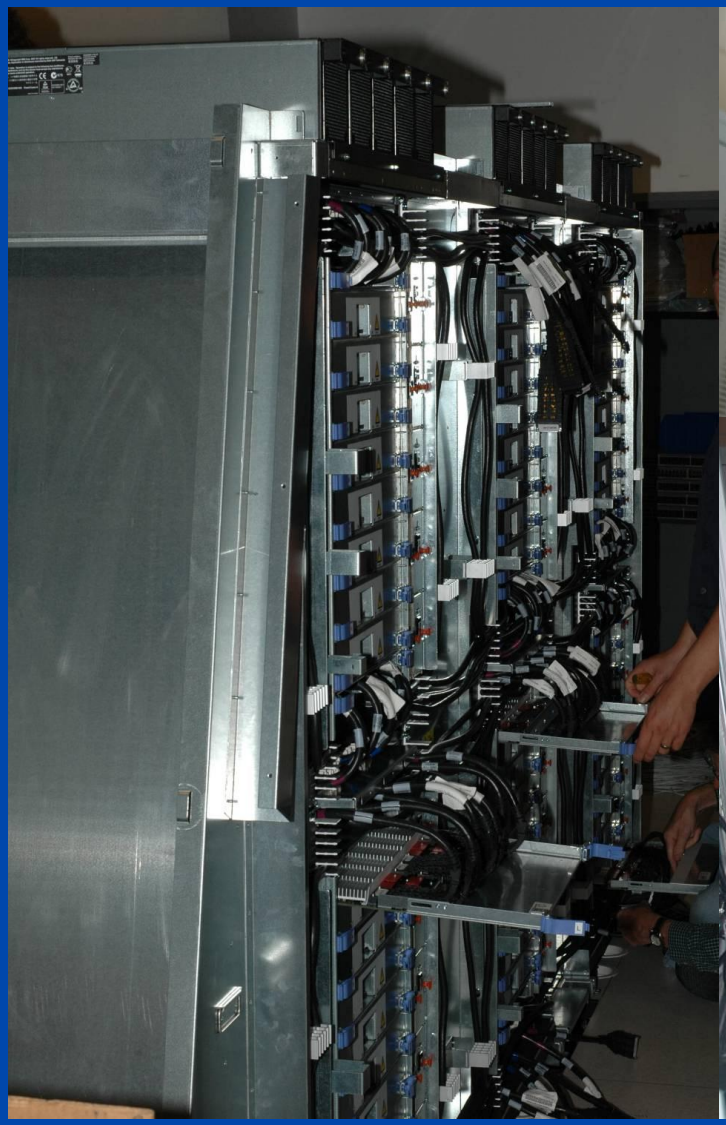

LOFAR

NVIDIA Tesla K10

ASTRON

- COBALT based on Tesla K10 GPU

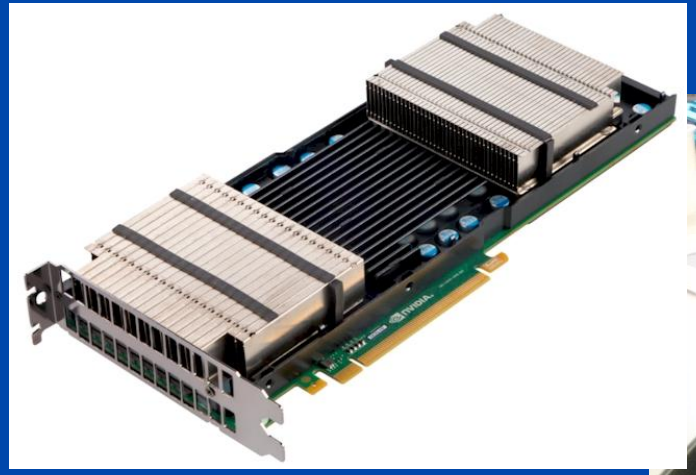

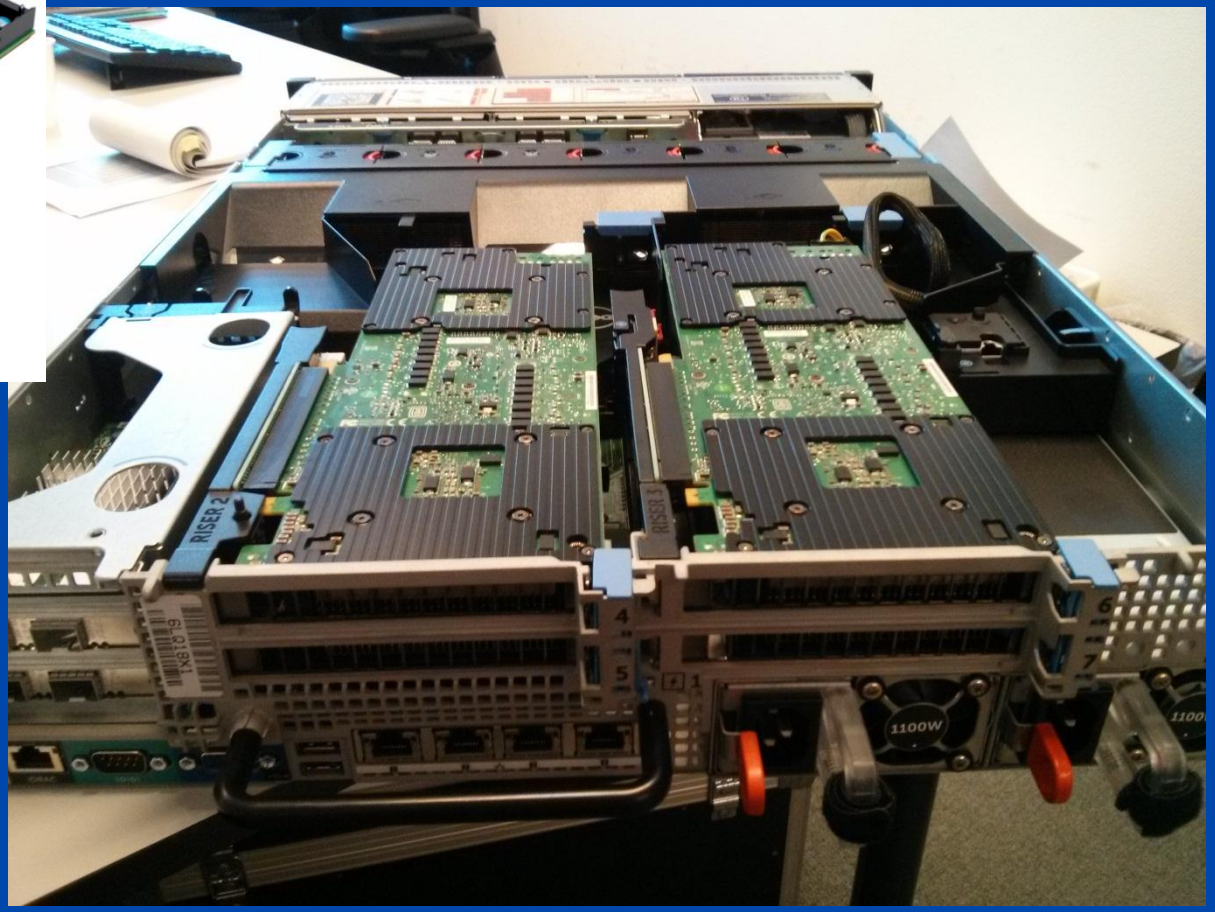

LOFAR

Requirements

ASTRON

1. Port all current operational functionality
2. Do this in a staged approach and be ready at the end of 2013
3. Do this within budget
4. Be ready to scale-up and / or extent
5. Maximize robustness, while minimizing operational and maintenance costs

Functionality

ASTRON

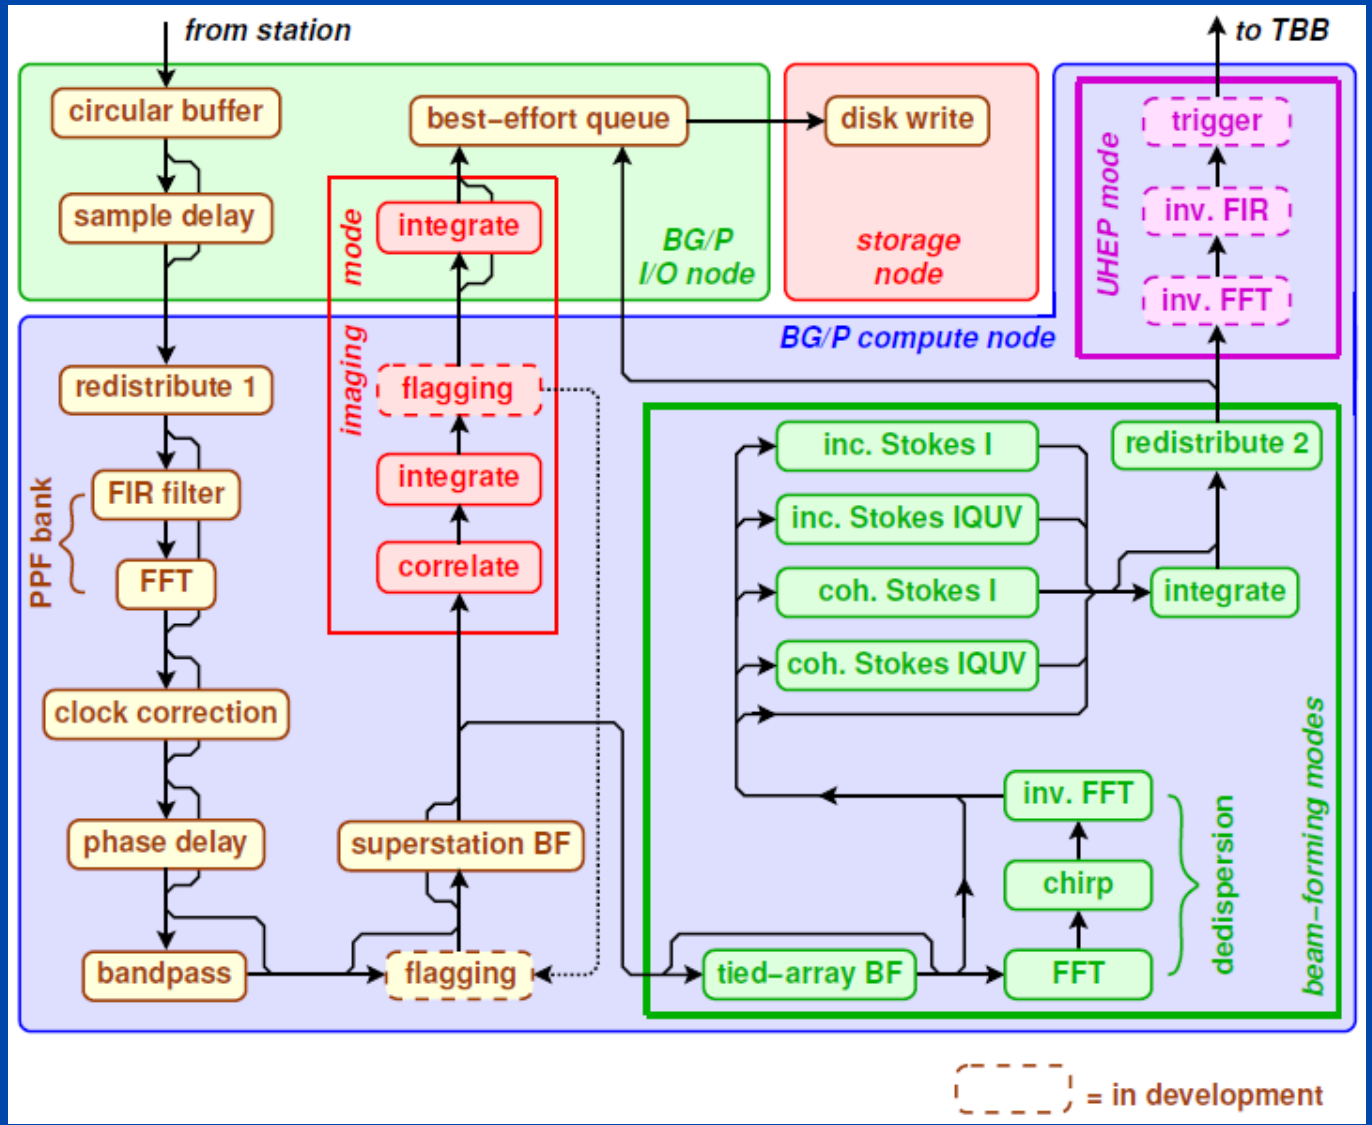

LOFAR

Specifications

ASTRON

- Be able to correlate 80 stations on 1 sec / 1 kHz resolution
  - Assumption: enough hw for beamforming 127 TABs
- An improved sw design
  - Processing pipelines will be independent to allow for efficient parallel and / or commensal observing modes.
  - allow for parallel observations with sub-arrays.
  - minimize latencies to allow for responsive telescope and UHEP mode.
  - allow for UHEP mode.
- A more efficient implementation of the beamforming pipeline
  - including an improved implementation of the on-line coherent de-dispersion part.
- Better integrate the development in the LOFAR continuous integration system and by using development cycles
  - facilitate external contributions by e.g. the eScience Center.
- Better integration with MAC software to prepare for responsive telescope

Preliminary HW System Design – Scenario 1

ASTRON

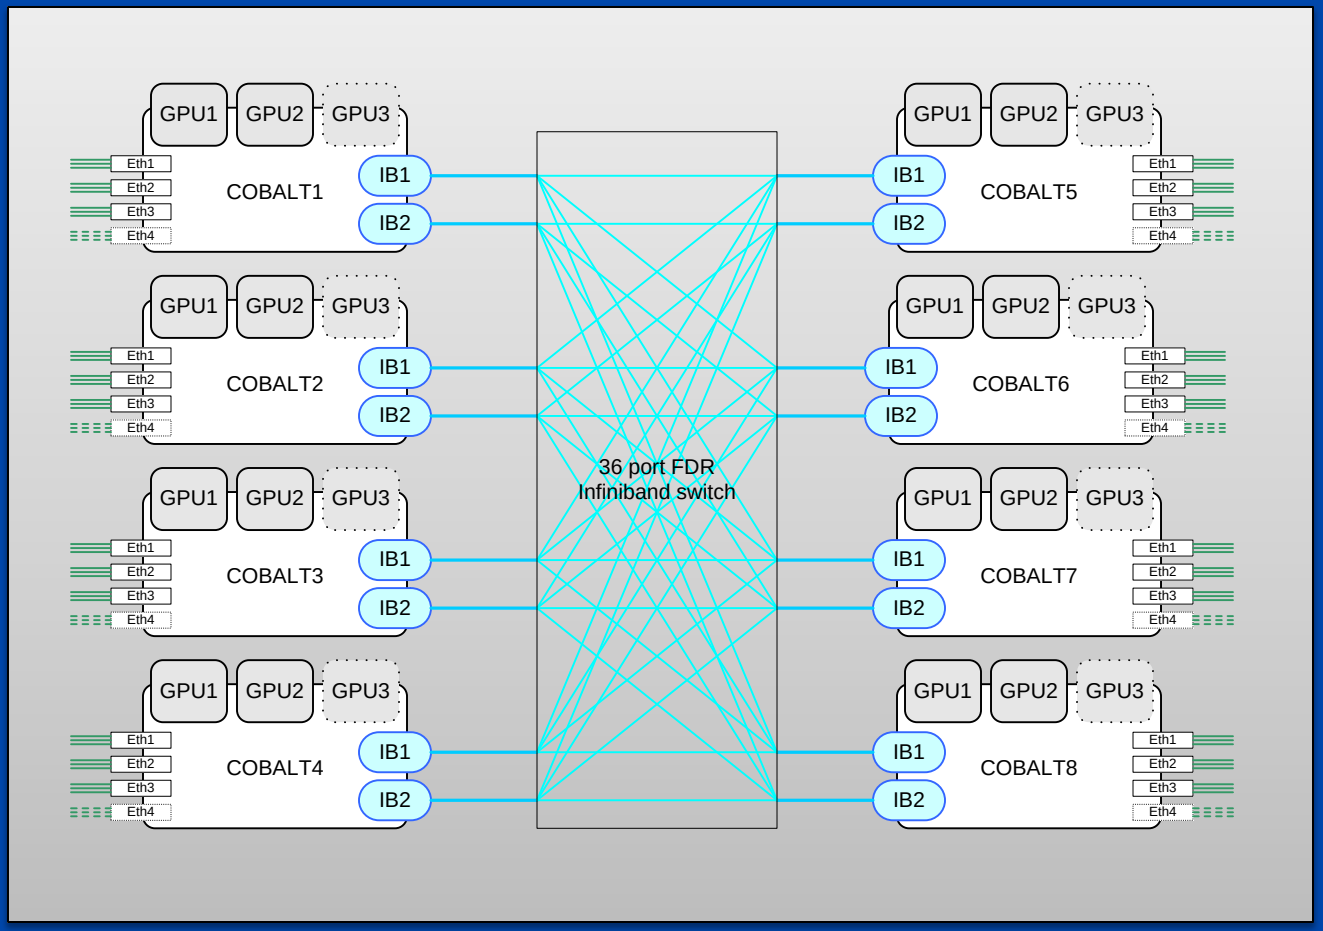

Minimum requirement: 6 nodes

Assumed size: 8 nodes

LOFAR

Preliminary HW System Design – Scenario 2

ASTRON

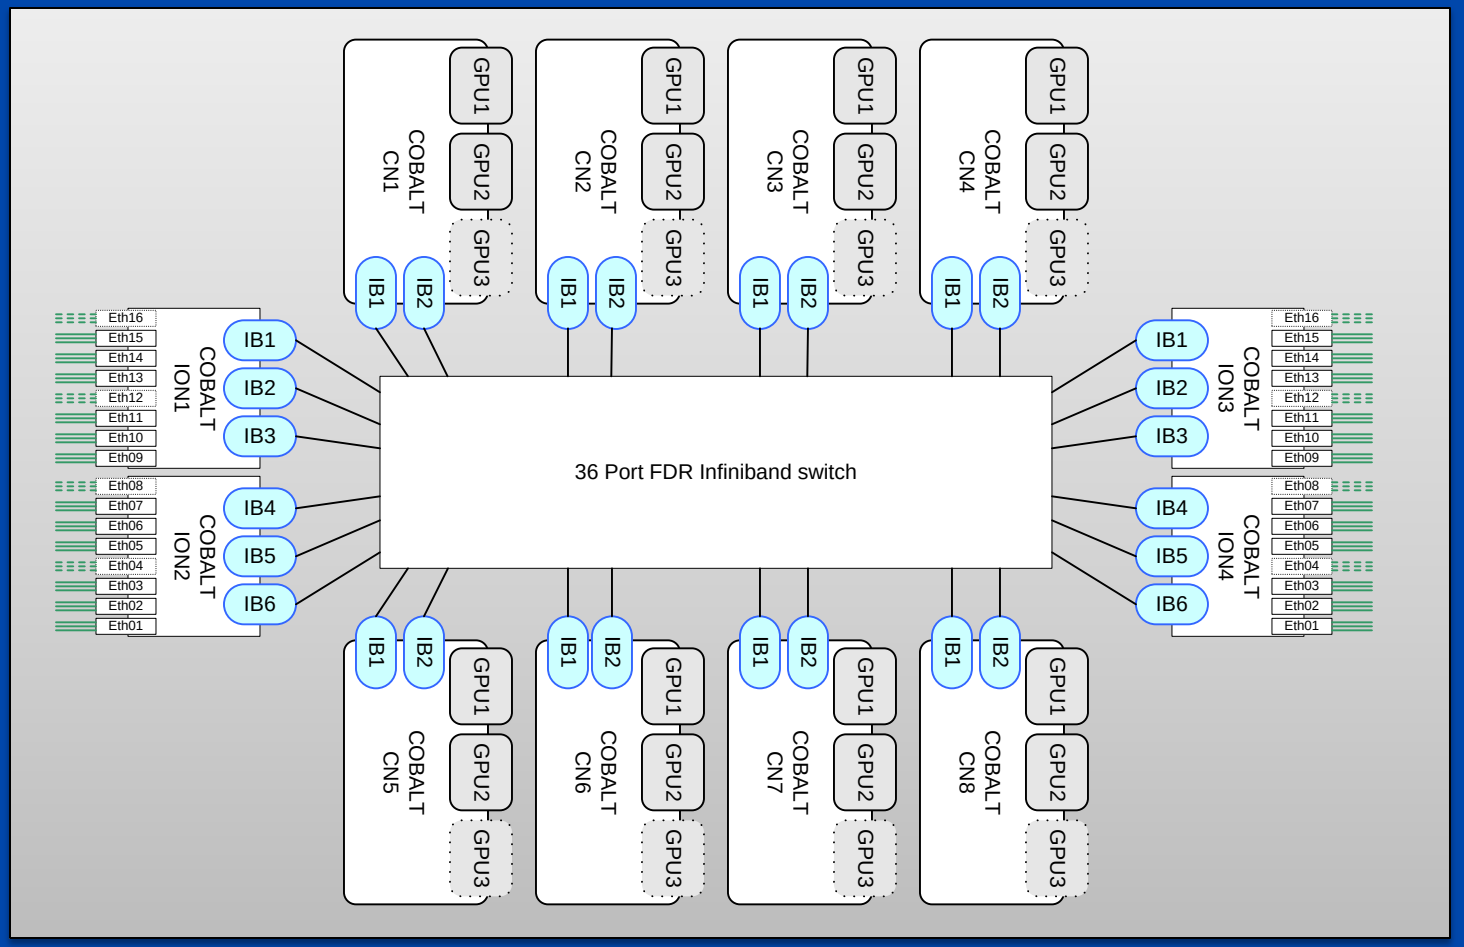

LOFAR

Current LOFAR network

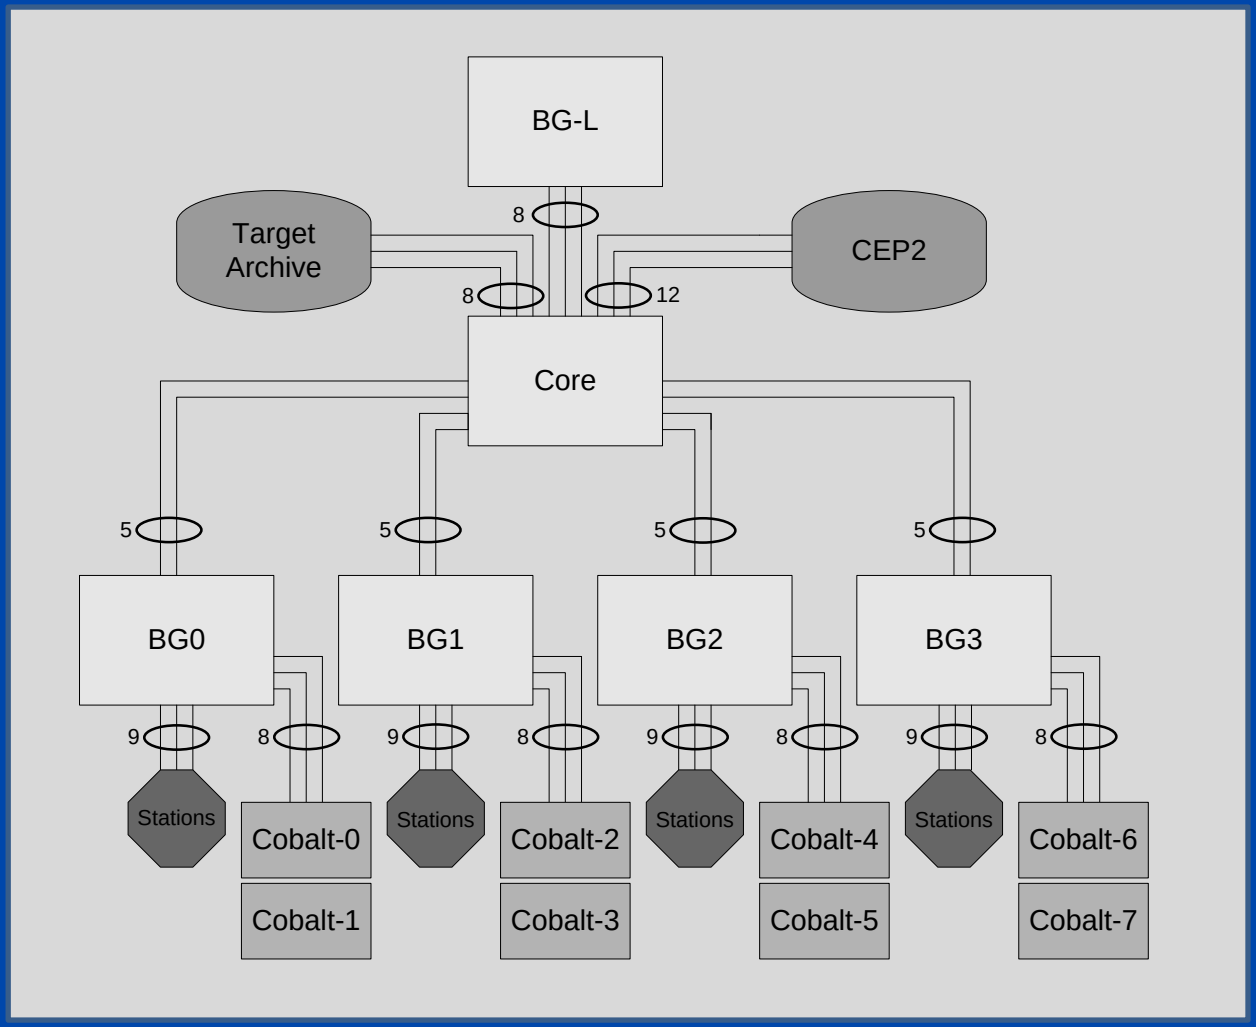

Option: update LOFAR network

ASTRON

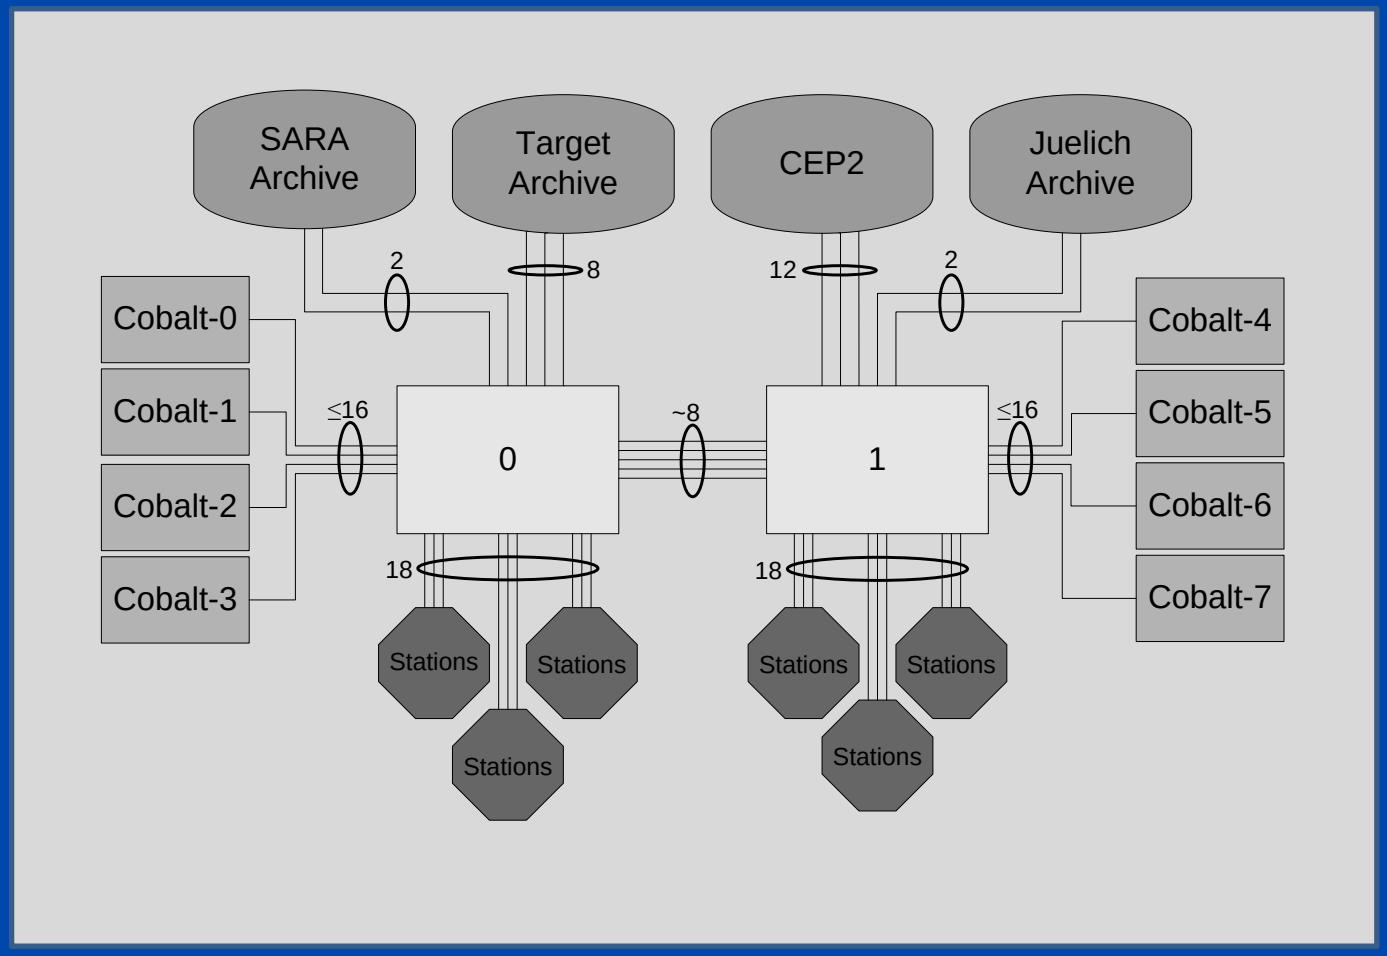

LOFAR

SW Design Overview

ASTRON

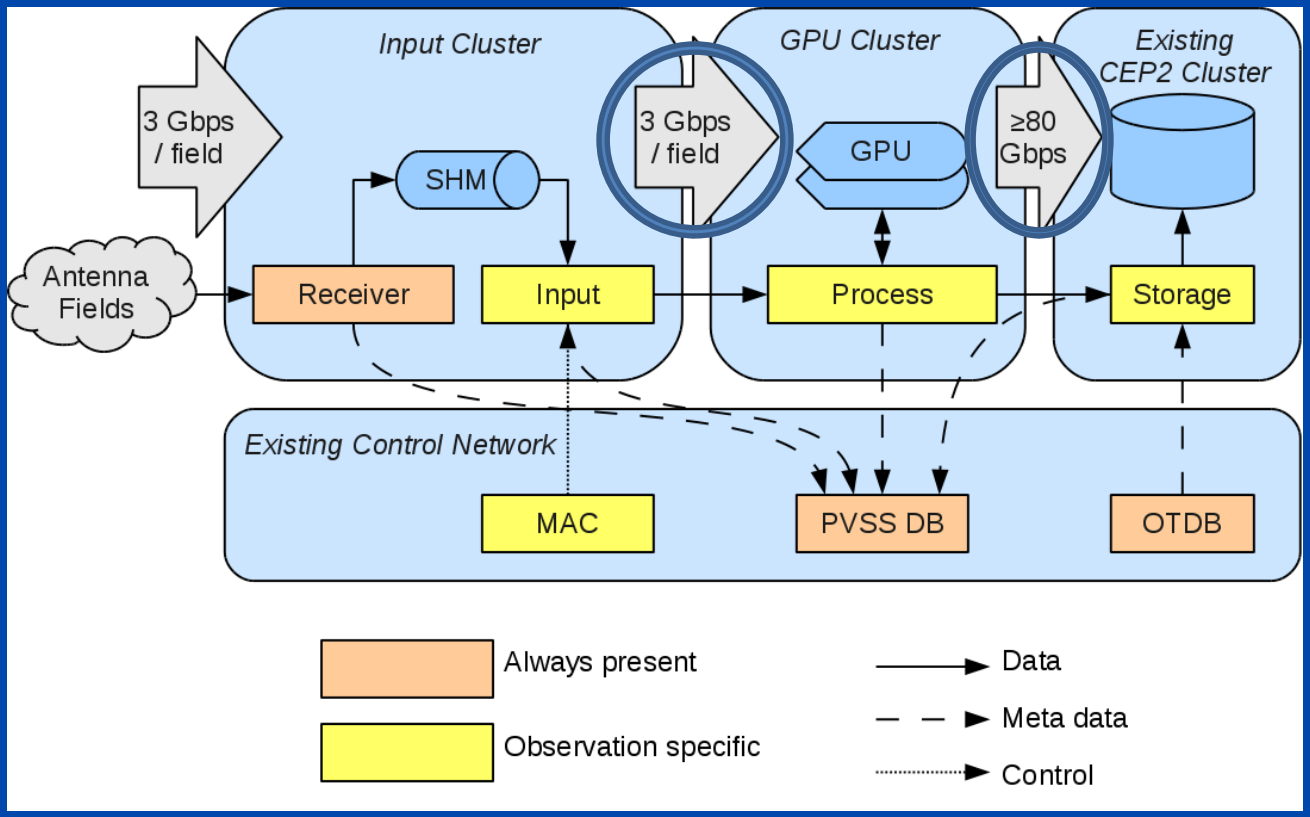

LOFAR

Latencies

ASTRON

Transport / Process	Latency
Antenna -> Cobalt	< 200 ms
Accumulate 5ms in circular buffer	5 ms
Read 5ms from circular buffer	< 5 ms
5 ms from 50 buffers -> GPU nodes	20 ms
Latency on 32 GPU processors	< 160 ms
Cobalt -> Stations	< 200 ms
Total	**< 590 ms**

- Total latency low enough for using TBBs

Resources

ASTRON

- Project leader: Ronald Nijboer
- Advisor: John Romein
- Project Scientist: 2 or 3; TBD
- Liason CIT Groningen: Hanno Holties
- Lead HW engineer: Chris Broekema
- Lead SW engineer: Jan David Mol
- Sw engineers: Alexander van Amesfoort, Marcel Loose, Wouter Klijn
- Sw engineers: Arthur Coolen, Arno Schoenmakers, Teun Grit, Alwin de Jong, Nico Vermaas

Status

ASTRON

- Design documents delivered
- Passed System Review
- Hw and sw choices made
  - Planning to be updated for this
- SW
  - Code refactoring
  - Functional / unit testing
  - CPU tasks

Impact on LOFAR

ASTRON

- Minimize impact; to be planned with Science Support
- Some downtime to be expected
  - Connecting GPU cluster: minimal
  - Network reconfiguration: 2 weeks in September
- Observing time
  - Initial test observations: June, July, August
  - Full system tests: October / November
    - LOFAR bandwidth split: 50% to BG/P, 50% to Cobalt
- Resources needed for porting the BG/P functionality remain available for LOFAR sw maintenance

LOFAR

Timeline 2013

ASTRON

Date	Activity	Deliverable
1 January	Start of Project	
Jan / Feb	Preparations	
12 February	Start sw implementation	Draft documents
18 February	Documentation send to CDR panel	Project Plan, Requirements, Top-level Design documents
28 February	System Review	
April / May	Procurement of Scenario 1 hw	
June	Scenario 1 hw available First Correlator pipeline (TBC)	Correlator commissioning plan (TBC)
July / August	Testing	
1 September	If needed: Scenario 2 hw available First Beamforming pipeline (TBC)	Beamformer commissioning plan (TBC)
September	Network reconfiguration	
Oct./ Nov.	Full system tests (incl. Observatory side)	
1 December	Cobalt replaces BG/P as production system	
31 December	End of Project	

Questions?