ASTRON

LOFAR

Update LOFAR Long Term Archive

May 18th, 2009

Hanno Holties

LOFAR Long Term Archive (LTA) Update

ASTRON

- Status
- Architecture
- Data Management
- Integration LOFAR, Target, EGEE
- Data Challenges
- Computational Facilities
- Planning

The LOFAR "CS1" archive

- **Archived data (~20TB) stored on Dutch GRID facilities**
  - For access GRID certificate required
  - Most data migrated to tape
    - (Very) high latency
    - Staging required for efficient downloading
- **Some data available from the Astro-Wise prototype**
  - Visibility data only
  - Incomplete: not yet automatically updated
- **10 Gbps connection CIT – SARA operational**
  - Used with 10 file servers in data challenges

LOFAR LTA
Target/AstroWise Prototype

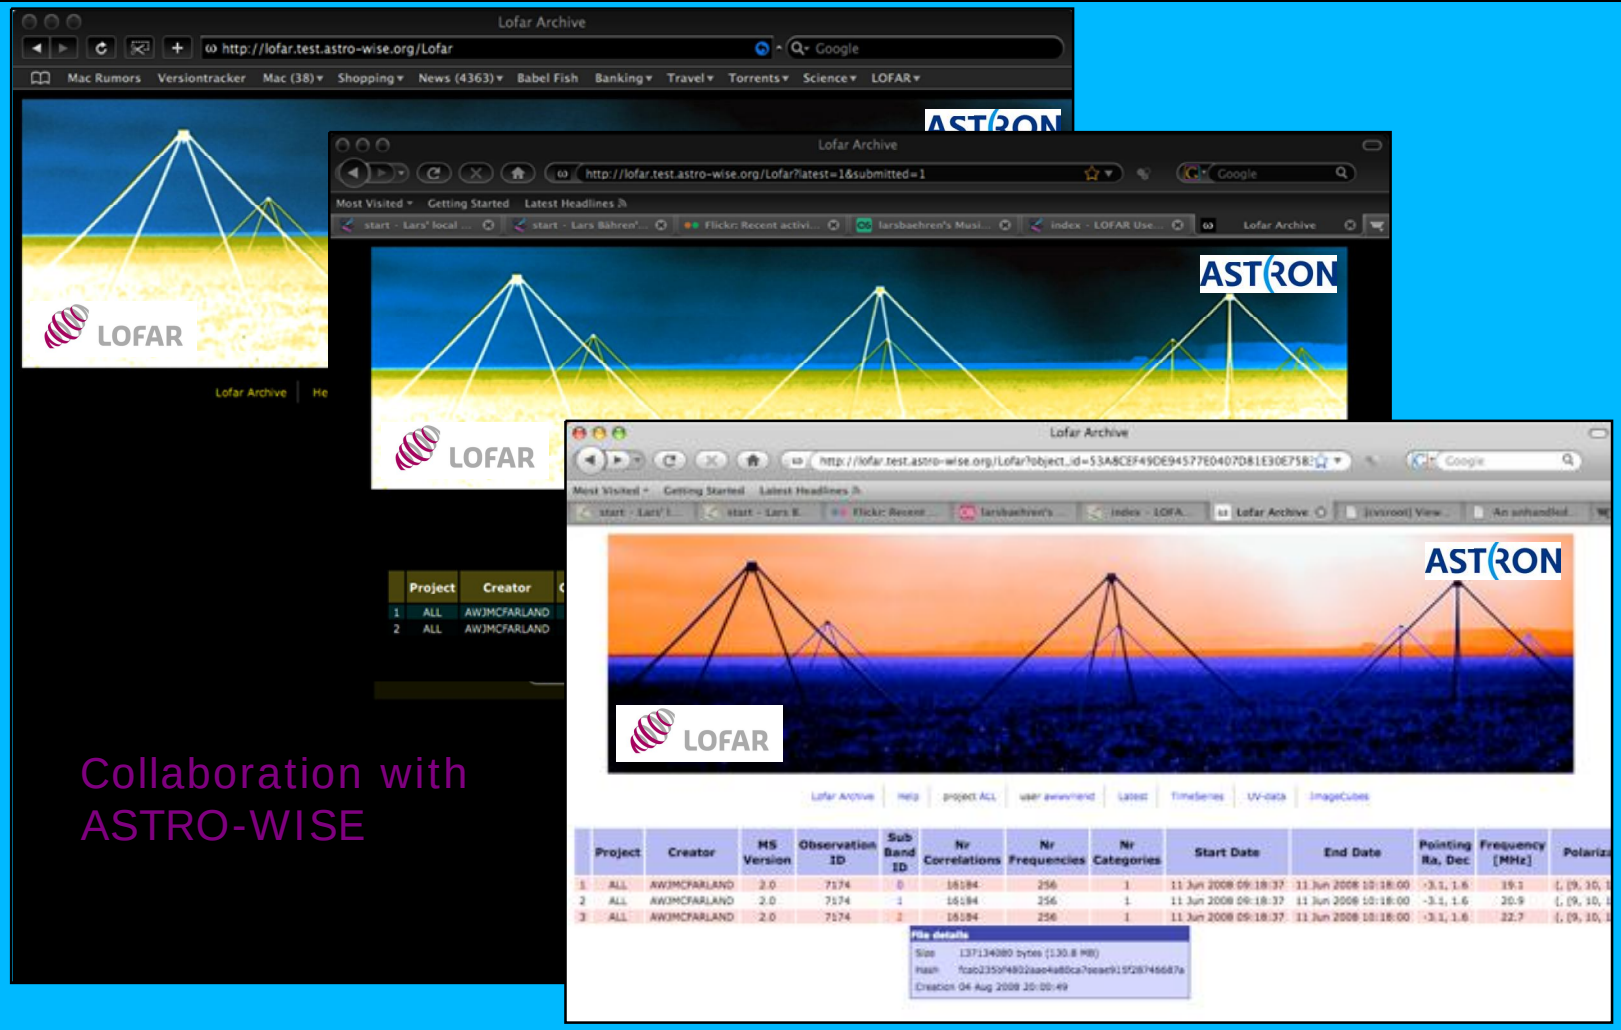

Collaboration with ASTRO-WISE

(courtesy W.-J. Friend and J. McFarland)

The LOFAR LTA Characteristics

- **Large** (Estimated yearly growth: 2.5 Petabyte)
  - Nevertheless storage (& computing) scarce resource
    - Allocated by Program Committee
  - Mix of technologies
    - Tape: cheap(er) & slow(er)
    - Disk: expensive & fast(er)
- **Distributed**
  - Groningen (CIT)
  - Amsterdam (SARA)
  - Jülich
  - ...
- **Integrated processing facilities**

The LOFAR LTA Activities

ASTRON

- **LOFAR / BiG Grid** (ASTRON, CIT, SARA, Nikhef)
  - Blue Print defined
  - Working on detailed design
  - Acquisition systems later this year
  - Plan to provide 200 – 300 TB storage this year
- **Target** (CIT, RUG, ASTRON)
  - Budget awarded
  - Working on detailed design
  - Plan to provide central database & File Server this year
- **GLOW** (Jülich)
  - Being discussed
  - Plan to provide 1 PB storage

The LOFAR LTA
The Central Processing facility

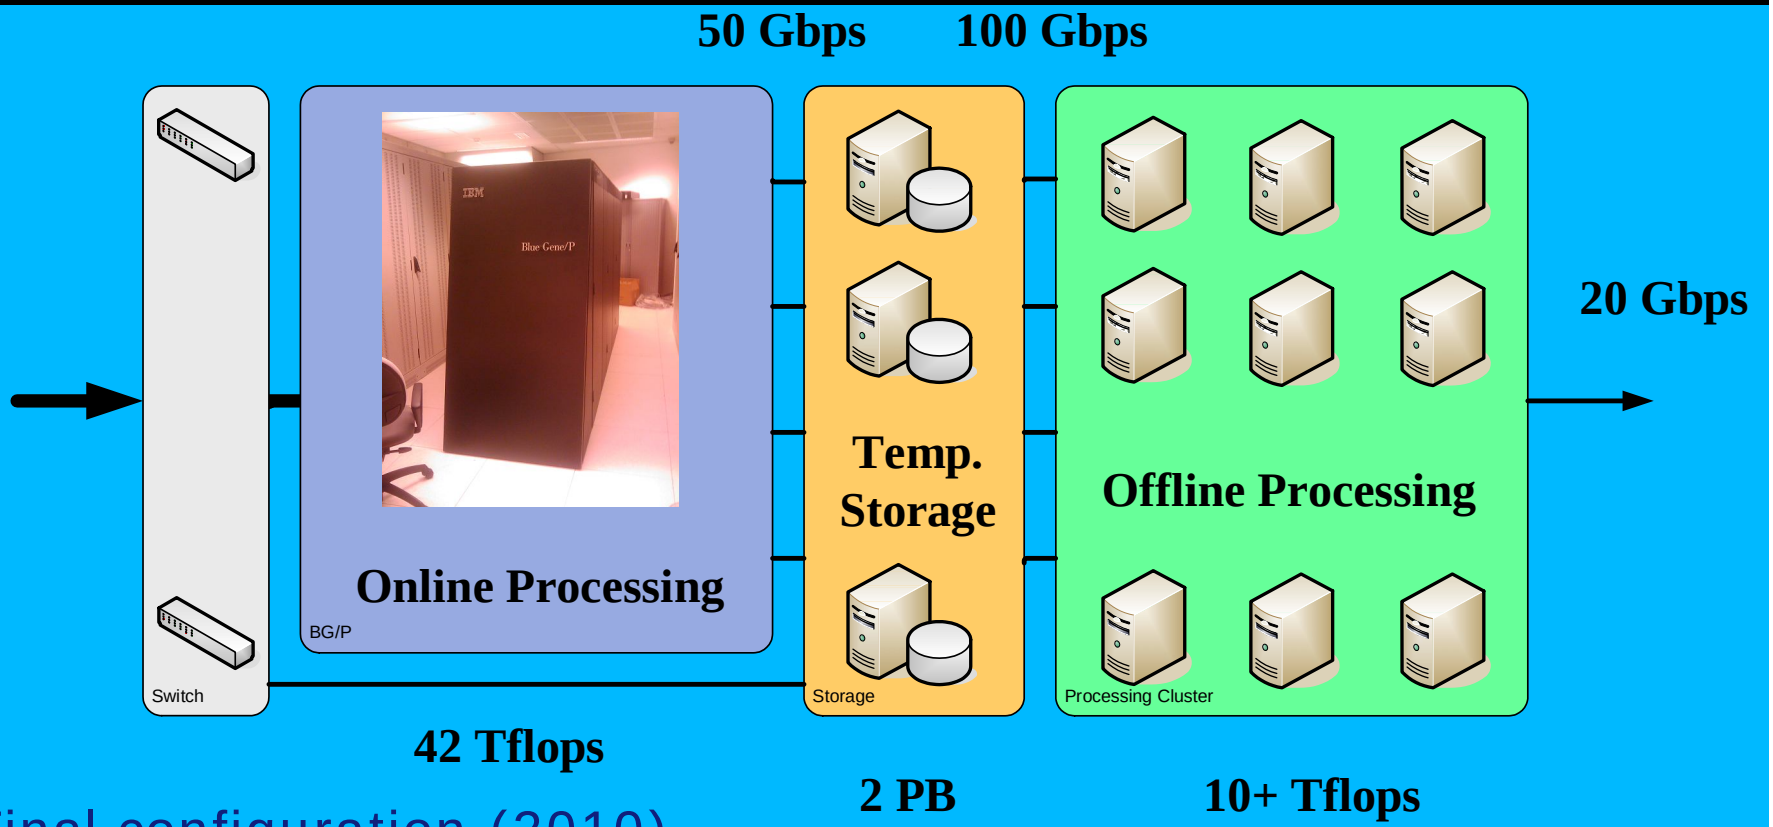

- Final configuration (2010)
- Output online processing ~50 gbps (0.5 PB per day!)
- Expect typically 1-2 weeks retention time

The LOFAR LTA Architecture

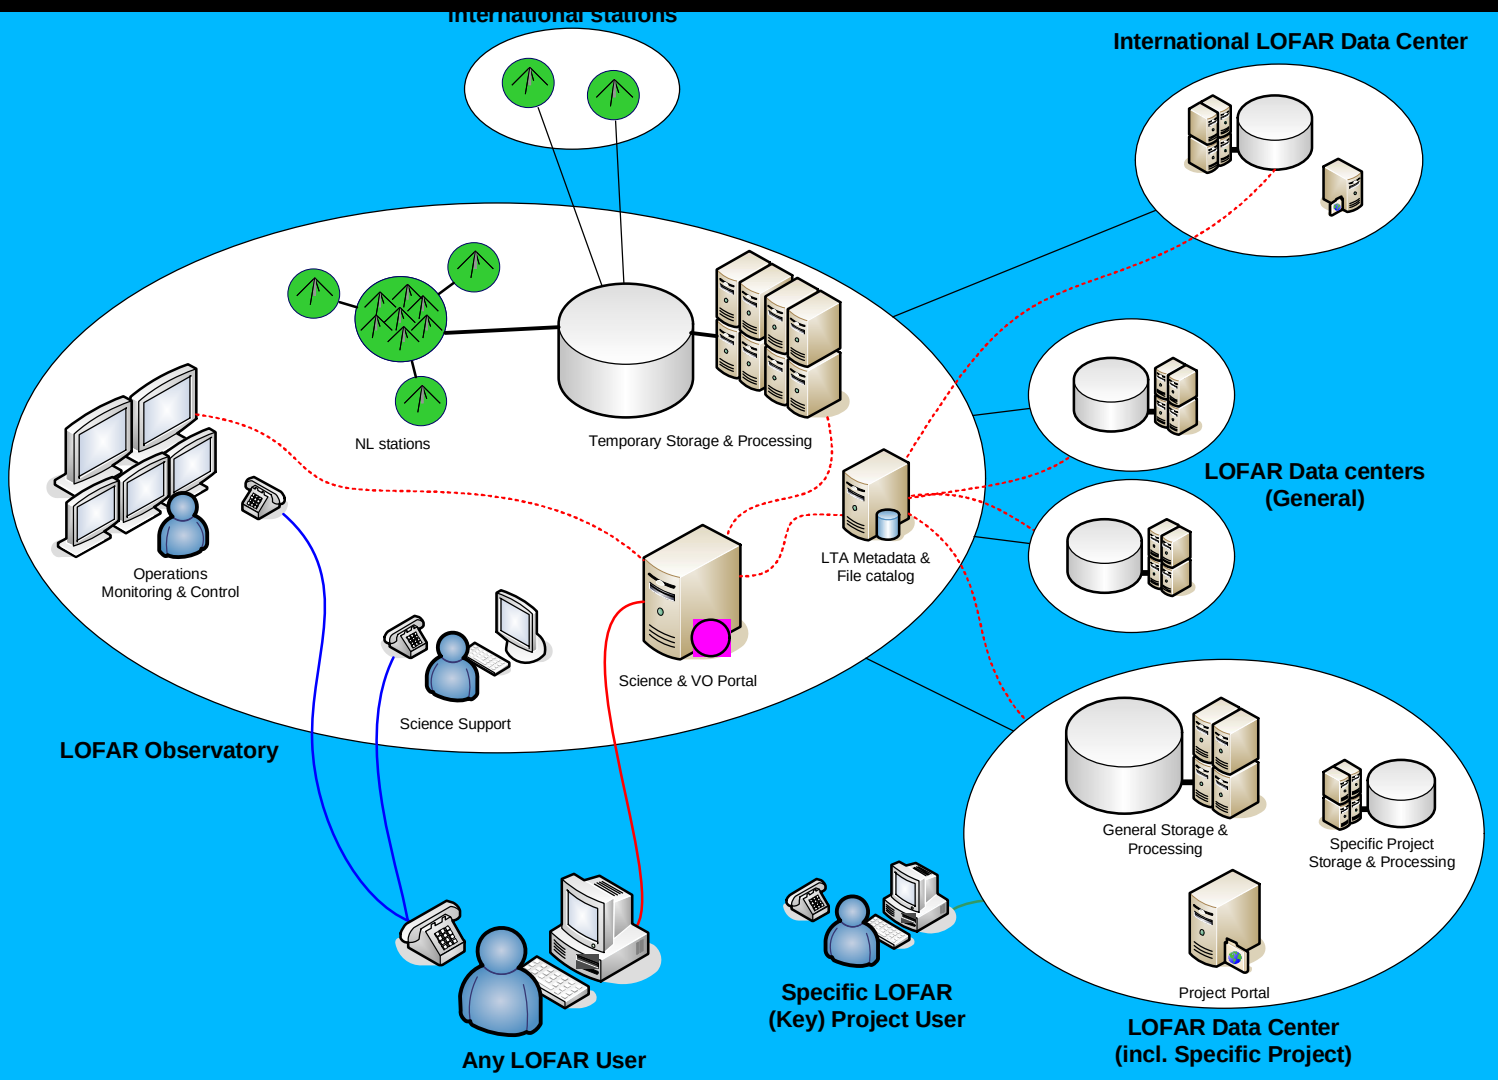

LOFAR LTA
Tier 1 Integration

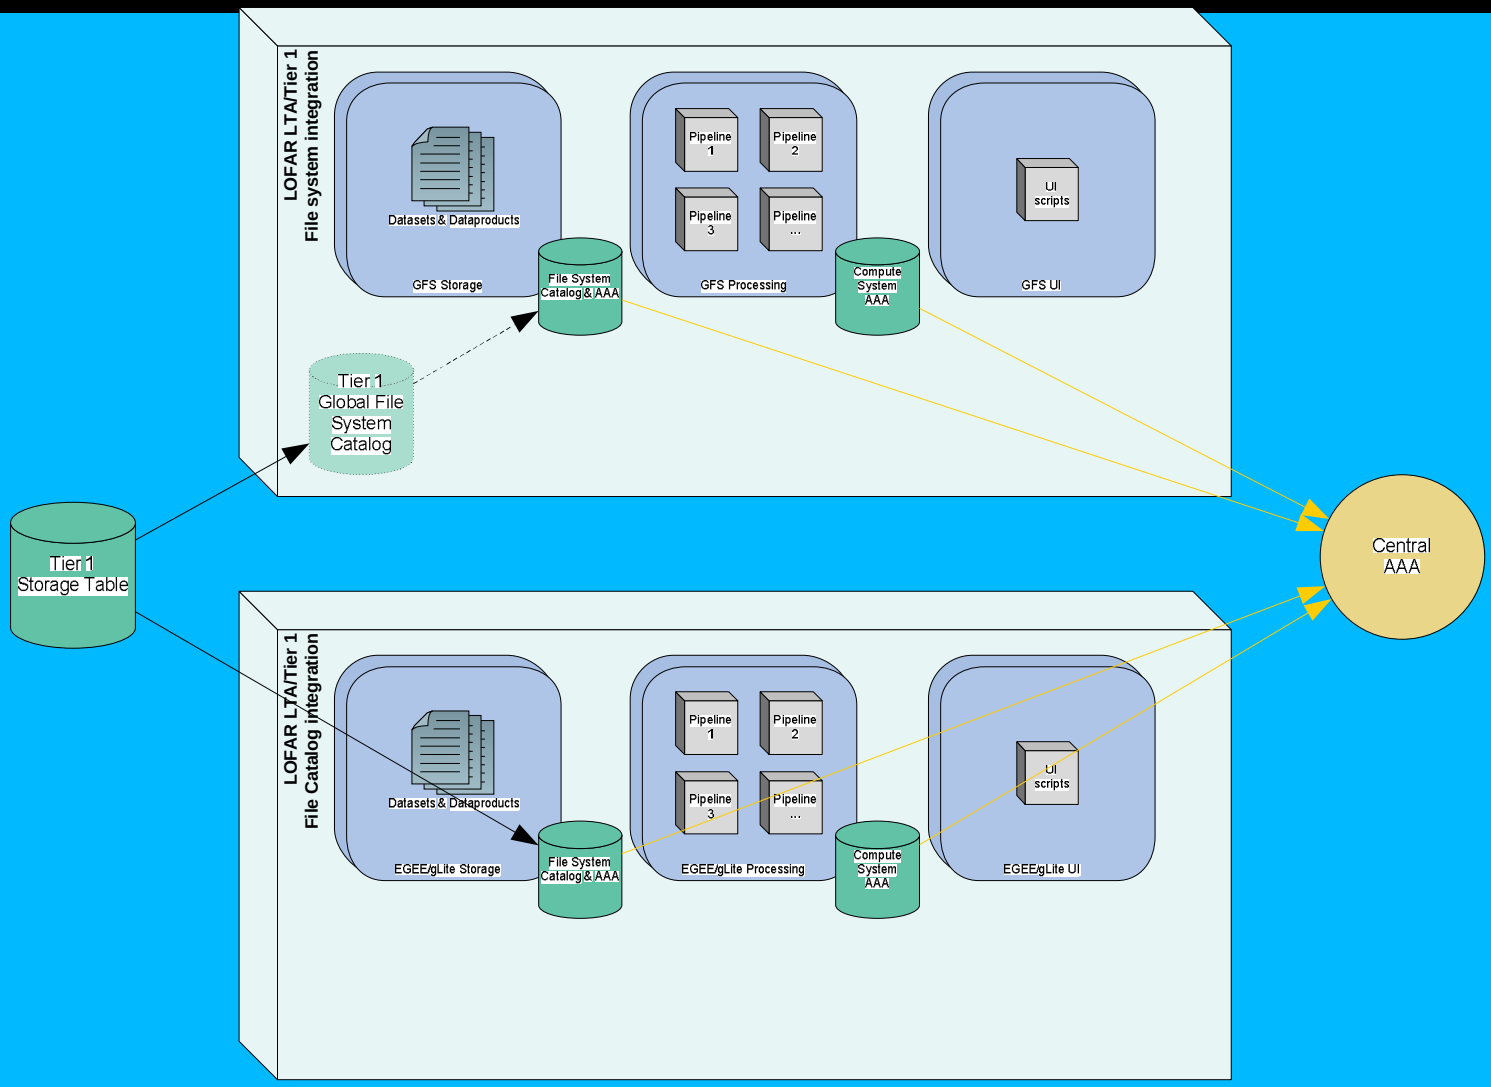

LTA/Tier 1 – Site Integration

The LOFAR LTA
Data management

- **Data retrieved from archive (not CEP)**
- **Data owned by LOFAR**
  - Proprietary period may apply
  - Most data will become publicly available
- **Security applied throughout the LTA**
  - Uses LOFAR user accounts
  - Use of personal (GSI) certificates is option
  - Authorization based on project administration
  - Resource allocation & usage accounted
- **Central file catalog**
  - Keeps track of location(s) of data stored in the LTA
  - Provides references to physical file locations for retrieval
- **Derived data products**
  - can be added to the archive (by user)

LOFAR LTA
Data streams

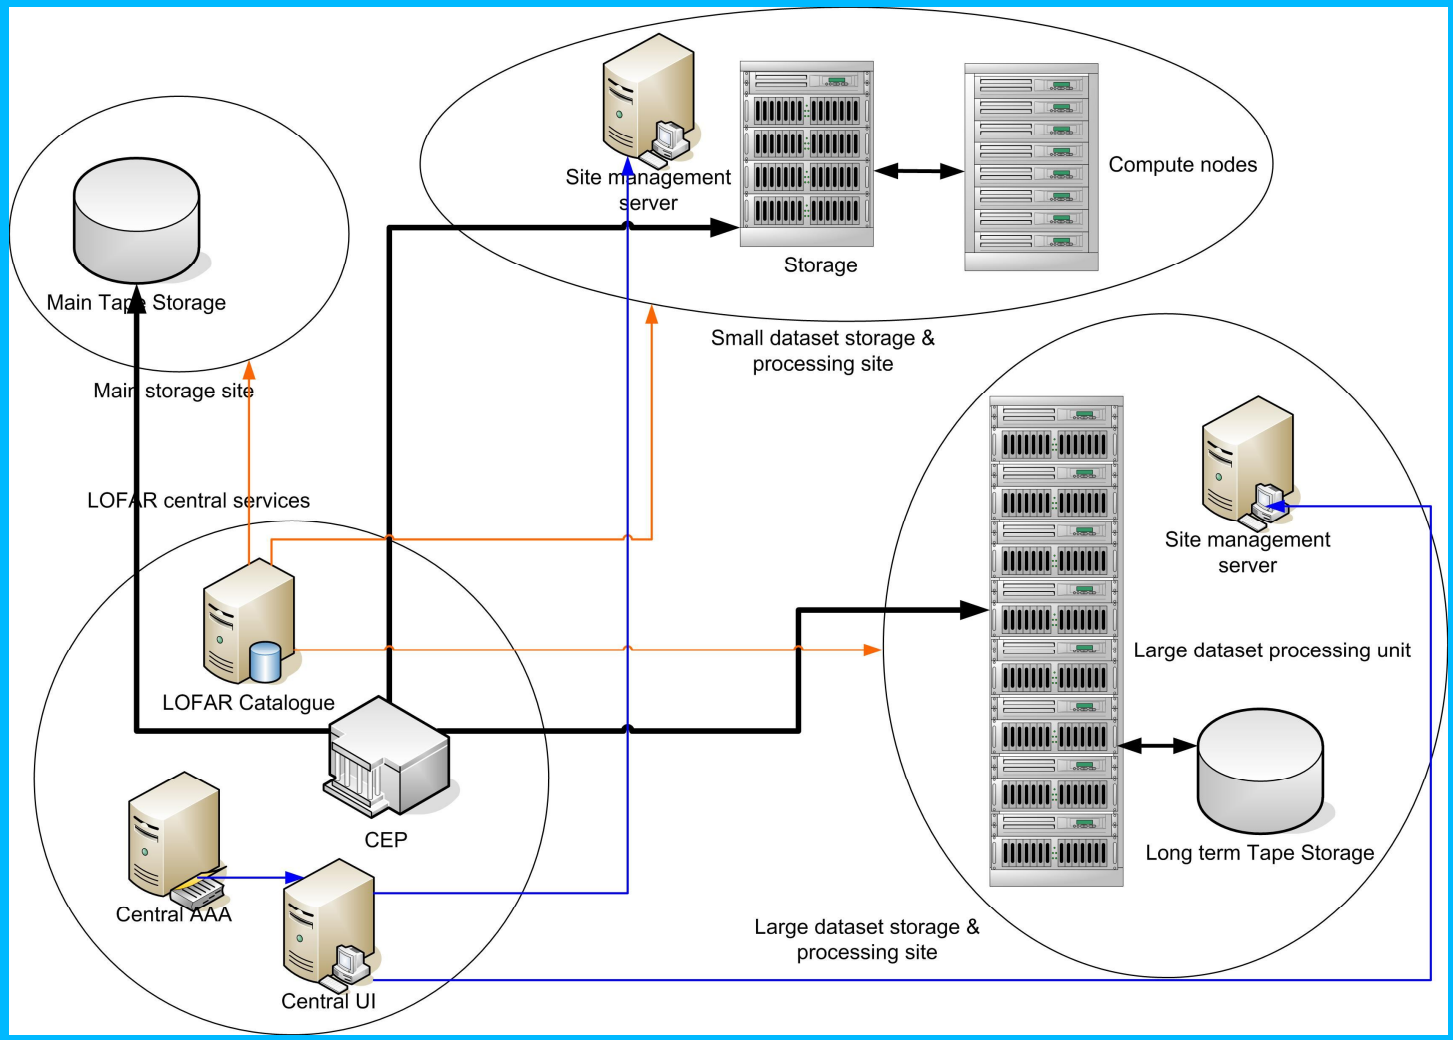

LOFAR LTA
Data challenges

1) Available effective bandwidth CEP – Tier 1

- Memory – Memory; aim: > 9 Gbps

2) Basic file transfer

- Disk – Disk; aim: > 9 Gbps

3) Stress test file transfers to online systems

- 24 hour; Disk – Disk; aim: 2 Gbps sustained

4) Basic beginning to end file transfer

- Disk – Tape; aim : 2 Gbps sustained

5) Operational use case

- 24 hour; Disk – Tape; aim: > 30 TB transferred

LOFAR LTA
Data challenges CEP - SARA

NB 10 available “old” servers using spare Gbps ports

Results:

1) Memory – Memory: 8.4 Gbps

- Probable bottle neck network cards used at CEP

2) Disk – Disk: 4.8 Gbps

- Bottleneck File servers at SARA side

3) Stress test Disk – Disk: 2.5 Gbps (not optimized)

- > 5Gbps in “burst-mode”

4) Disk – Tape: 0.8 Gbps

- Limited by available SARA tape systems

5) On hold

- Waiting for new CEP systems & SARA LHC tape units

LOFAR LTA
Data challenges CEP - SARA

3) 24 hour Disk – Disk

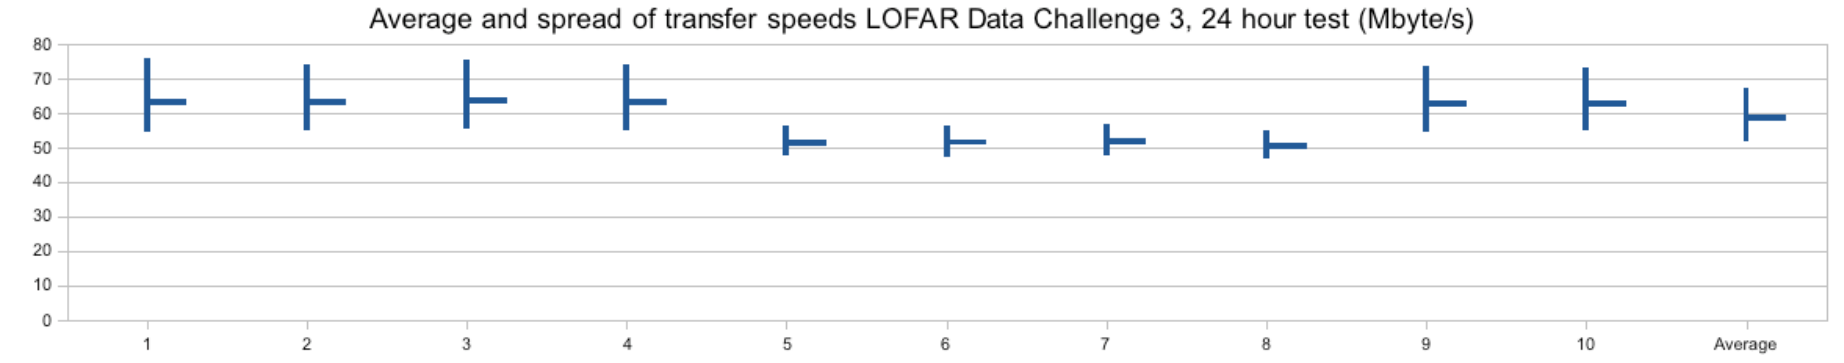

4 SARA file servers also provide shared disk storage

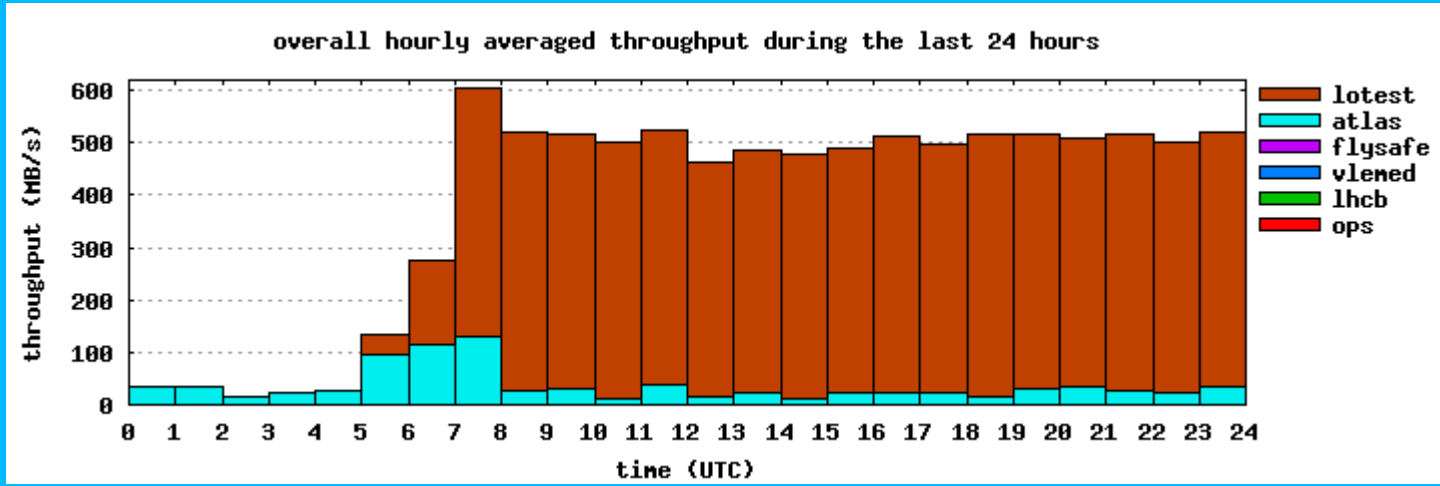

LOFAR LTA
Data challenges CEP - SARA

4) Disk – Tape

3 Stage process:

i) Storage on disk based
file servers

ii) Automatic migration to
tape staging area (disk)

iii) Migration to tape
2x 60 MBps unit

Room for improvement!

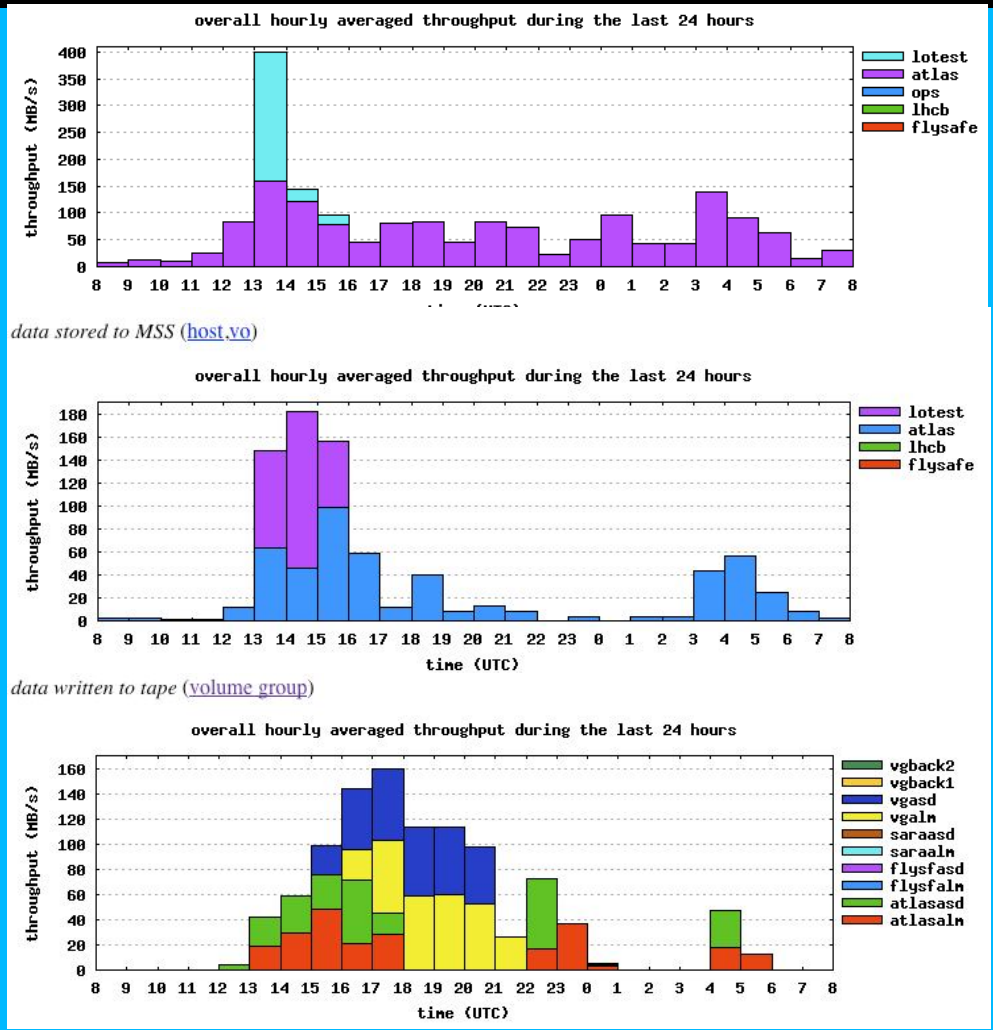

LOFAR LTA
Processing in the LTA

ASTRON

Still many uncertainties (what is exactly needed)

Some typical (?) processing steps have been considered:

- Calibration: BBS estimates (mostly embarrassingly parallel)
- Projection: 5 TB visibility dataset (can be parallelized)
- FFT: 10,000 × 10,000

Preliminary conclusions:

- Likely to fit on standard 8 core (dual quad core) nodes with 16 or 32 Gbytes of RAM
- Bandwidth limited: requires fast network to keep compute nodes busy
- MPI & SMP needed to use cores efficiently with 10GB+ datasets

LOFAR LTA Planning

ASTRON

- **Connection CEP – Tier 1**
  - SARA (update), TARGET/CIT, Jülich
  - Data challenges
- **Security implementation**
  - Central LOFAR Identity Management
  - Provisioning through proposal application
  - Integration with Tier 1 facilities
- **Ingest procedure**
  - Bulk data transfers to distributed Tier 1 sites
  - Metadata to central LTA database
  - Storage table (LTA central file catalog) being designed

LOFAR LTA Conclusion

- **First implementation ready Q3 2009**
- **Data challenges progressing**
  - Preliminary results promising
  - Restrictions understood
- **In particular processing requirements uncertain**
  - Need substantiation through performing realistic tests!
- **BiG Grid facilities are available**
  - Storage
  - Processing (...)
- **Target prototype is available**
- **Jülich offers to provide a Petabyte**