


ASTRON



LOFAR

The LOFAR Long Term Archive

February 12th, 2009
Hanno Holties




The LOFAR LTA The Central Processing facility



ASTRON

2009



3 Gbps/station

35 Tflops

0.25 PB

10+ Tflops

- Output online processing < 50 gbps (0.5 PB per day!)
- Expect typically 1-2 weeks retention time (~0.1 PB per day on average)

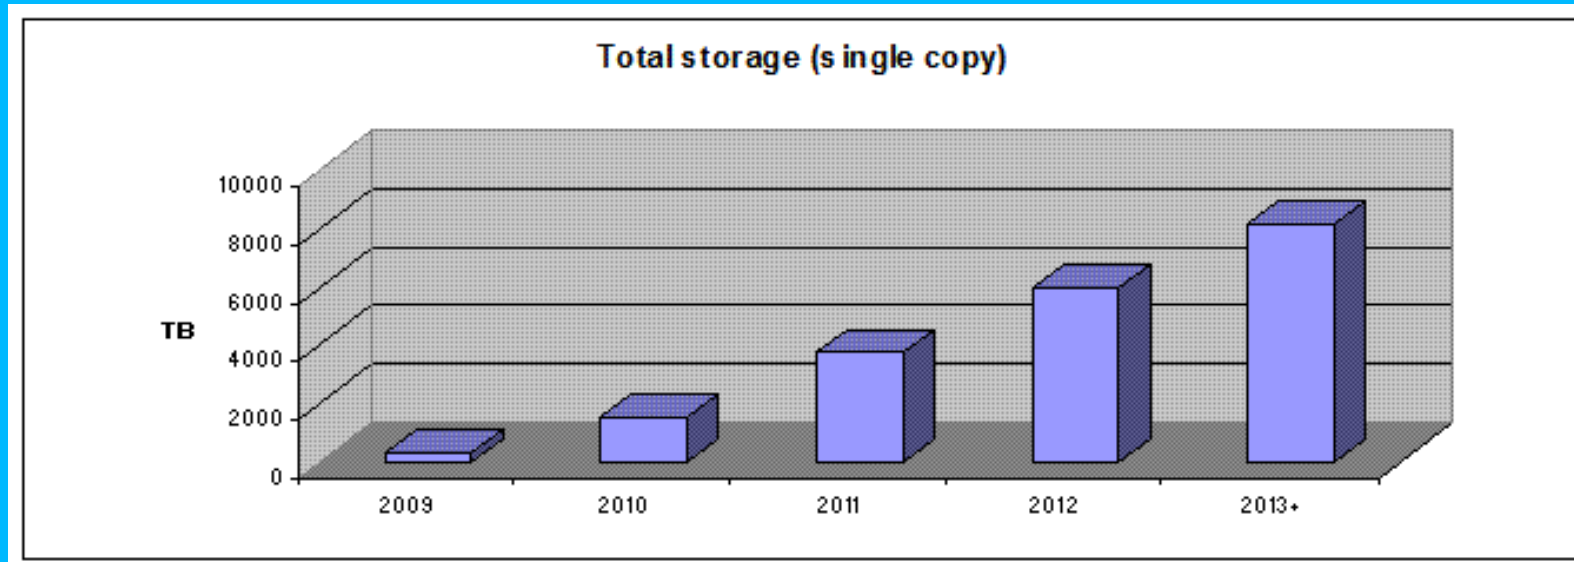
The LOFAR LTA

The Central Processing facility



- **At CEP**
 - **'Online' processing on BG/P**
 - **'Offline' processing on cluster**
 - Within 2 weeks
 - Limited capacity (2 runs?)
 - When in full operation: no direct user access
 - **Before archiving**
 - Quality analysis
 - Data sufficiently flagged & calibrated to allow significant reduction in size
 - Derived data products 'packaged'
 - **Data does not get read back from the archive**

The LOFAR LTA Estimated growth archived data



- Data reduction factor before archiving 10–30x
- Long term data rate: 2.5 PB/yr (single copy!)
times two to allow recovery operations: ~1.3 Gbps sustained
- Catalog(s) 10^7+ entries
- Data integrity mechanisms & redundancy required

The LOFAR LTA Characteristics



- **Large** (Multi-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 Data management



- **Data retrieved from archive (not CEP)**
- **Data owned by LOFAR**
 - Proprietary period may apply
 - Most data will become publicly available
- **Derived data products** can be added to the archive (by user)

- **NB Commissioning data not public!**

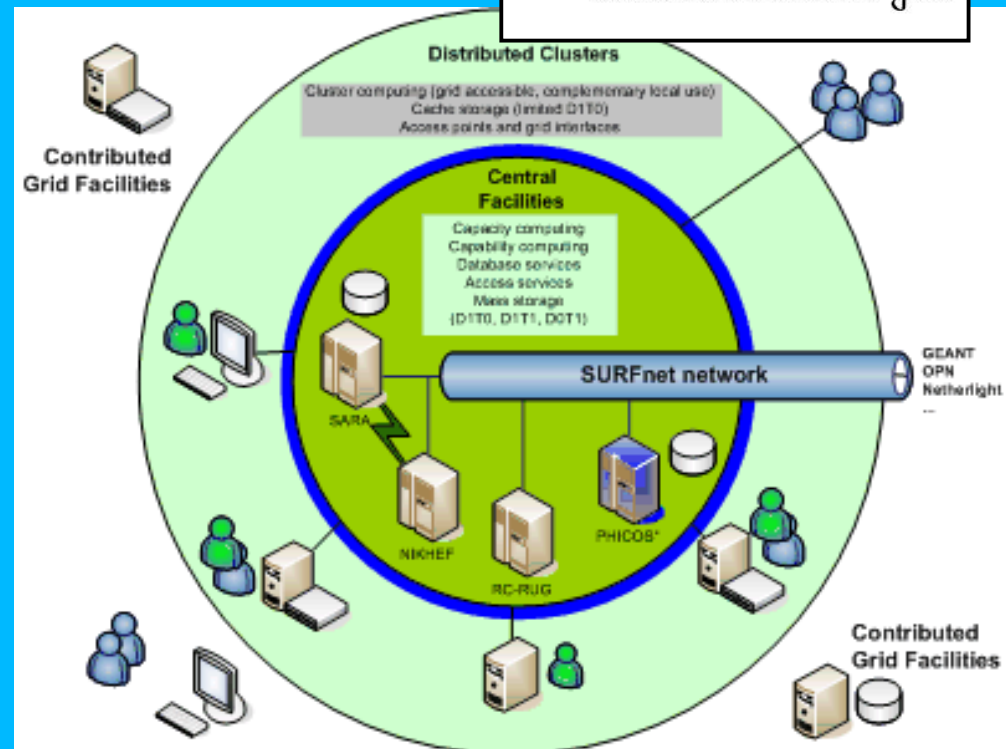
The LOFAR LTA BiG Grid



ASTRON

Dutch project for GRID computing

- NIKHEF, NCF, NBIC
 - Supported by many science & technology organizations
- Multi Petabyte storage
 - Disk
 - Tape
- Thousands of processing nodes
- Utilizing EGEE middleware (gLite)



The LOFAR LTA Target



Dutch Northern region project

- Intelligent information systems for sensor networks
- Centered around University of Groningen
- Includes
 - Astronomy
 - Lifelines
 - Artificial Intelligence
 - Commercial partners



- **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
 - Not complete

The CS1 archive



■ Finding observations

- The observation catalog (do not trust the location)

- Or: ask someone

■ Finding the location

- On LISTFEN:

- lofstorman
- observerui.csh

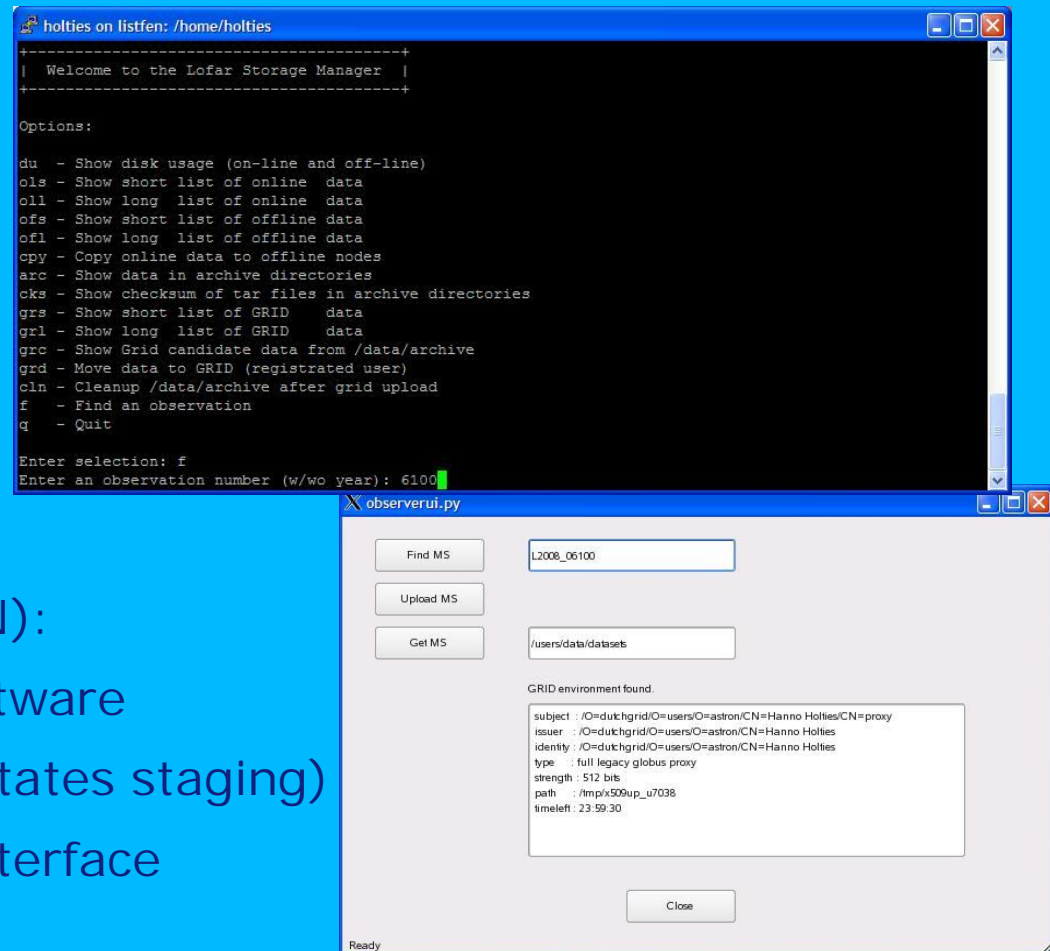
- Or: ask operator

■ Retrieving from the archive

- On GRID UI (e.g. LISTFEN):

- gLite GRID client software
- observerui.csh (facilitates staging)
- Via prototype web interface

- Or: ask operator



LOFAR LTA Prototype



Collaboration with
ASTRO-WISE

Project	Creator
1 ALL	AWJMCFARLAND
2 ALL	AWJMCFARLAND

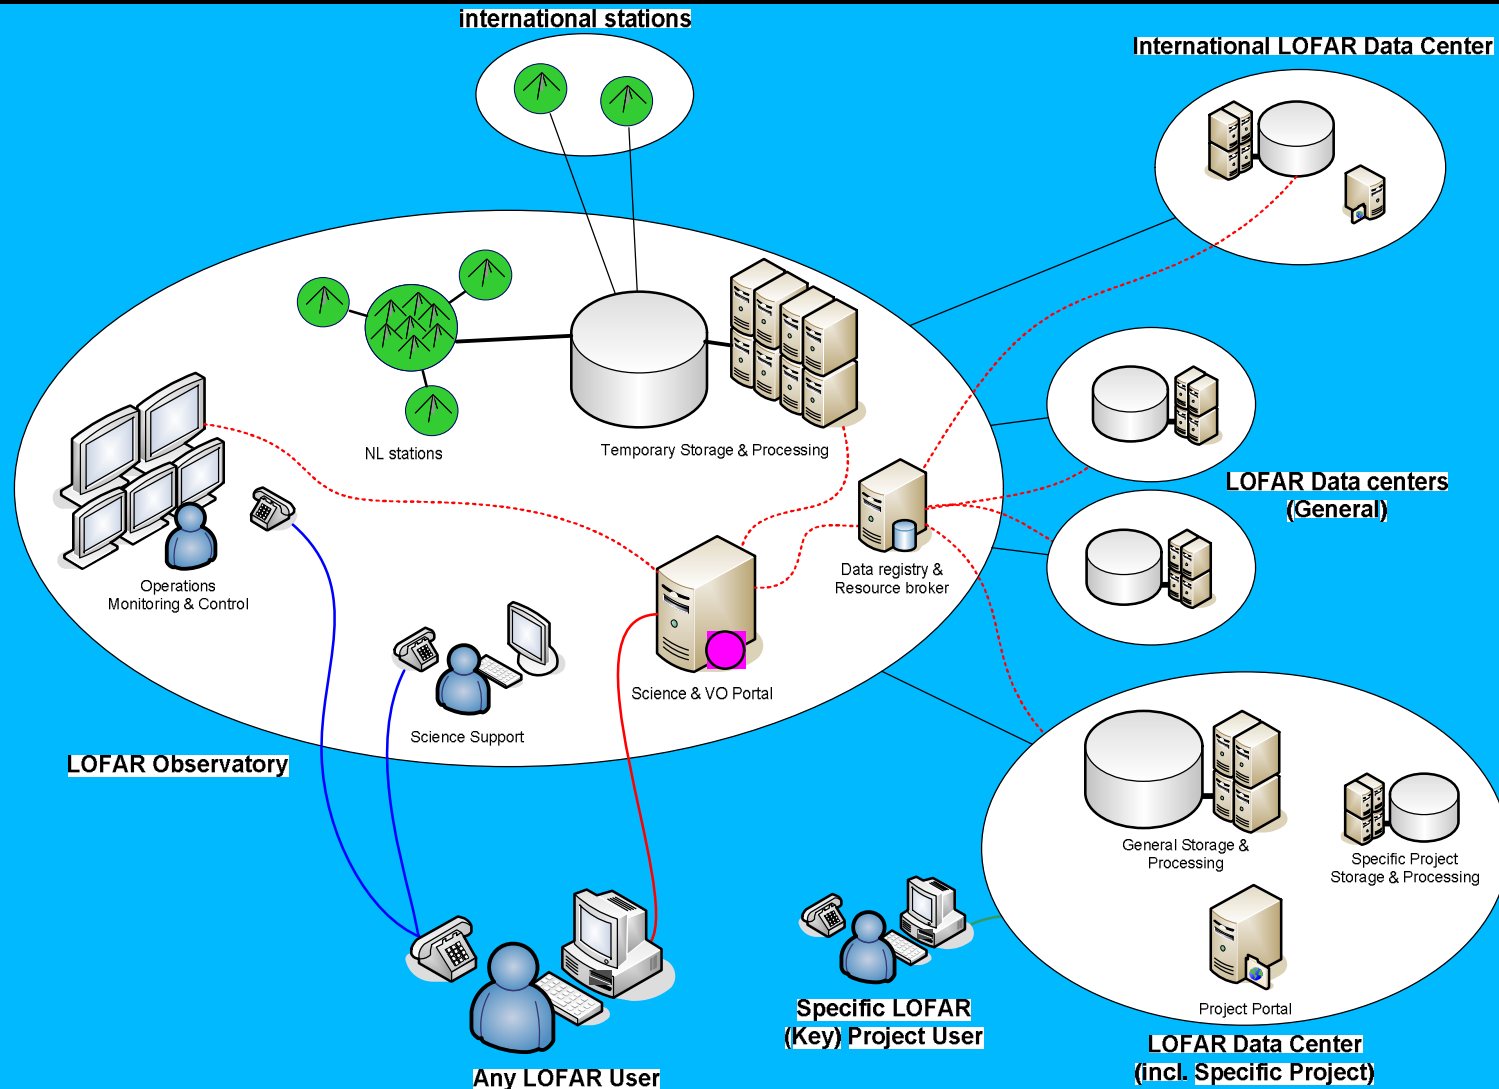
Project	Creator	MS Version	Observation ID	Sub-Band ID	Nr Correlations	Nr Frequencies	Nr Categories	Start Date	End Date	Pointing Ra, Dec	Frequency [MHz]	Polarization	
1	ALL	AWJMCFARLAND	2.0	7174	0	16194	256	1	11 Jun 2008 09:18:37	11 Jun 2008 10:18:00	-3.1, 1.6	19.1	[, (9, 10, 11)
2	ALL	AWJMCFARLAND	2.0	7174	1	16194	256	1	11 Jun 2008 09:18:37	11 Jun 2008 10:18:00	-3.1, 1.6	20.9	[, (9, 10, 11)
3	ALL	AWJMCFARLAND	2.0	7174	2	16194	256	1	11 Jun 2008 09:18:37	11 Jun 2008 10:18:00	-3.1, 1.6	22.7	[, (9, 10, 11)

File details

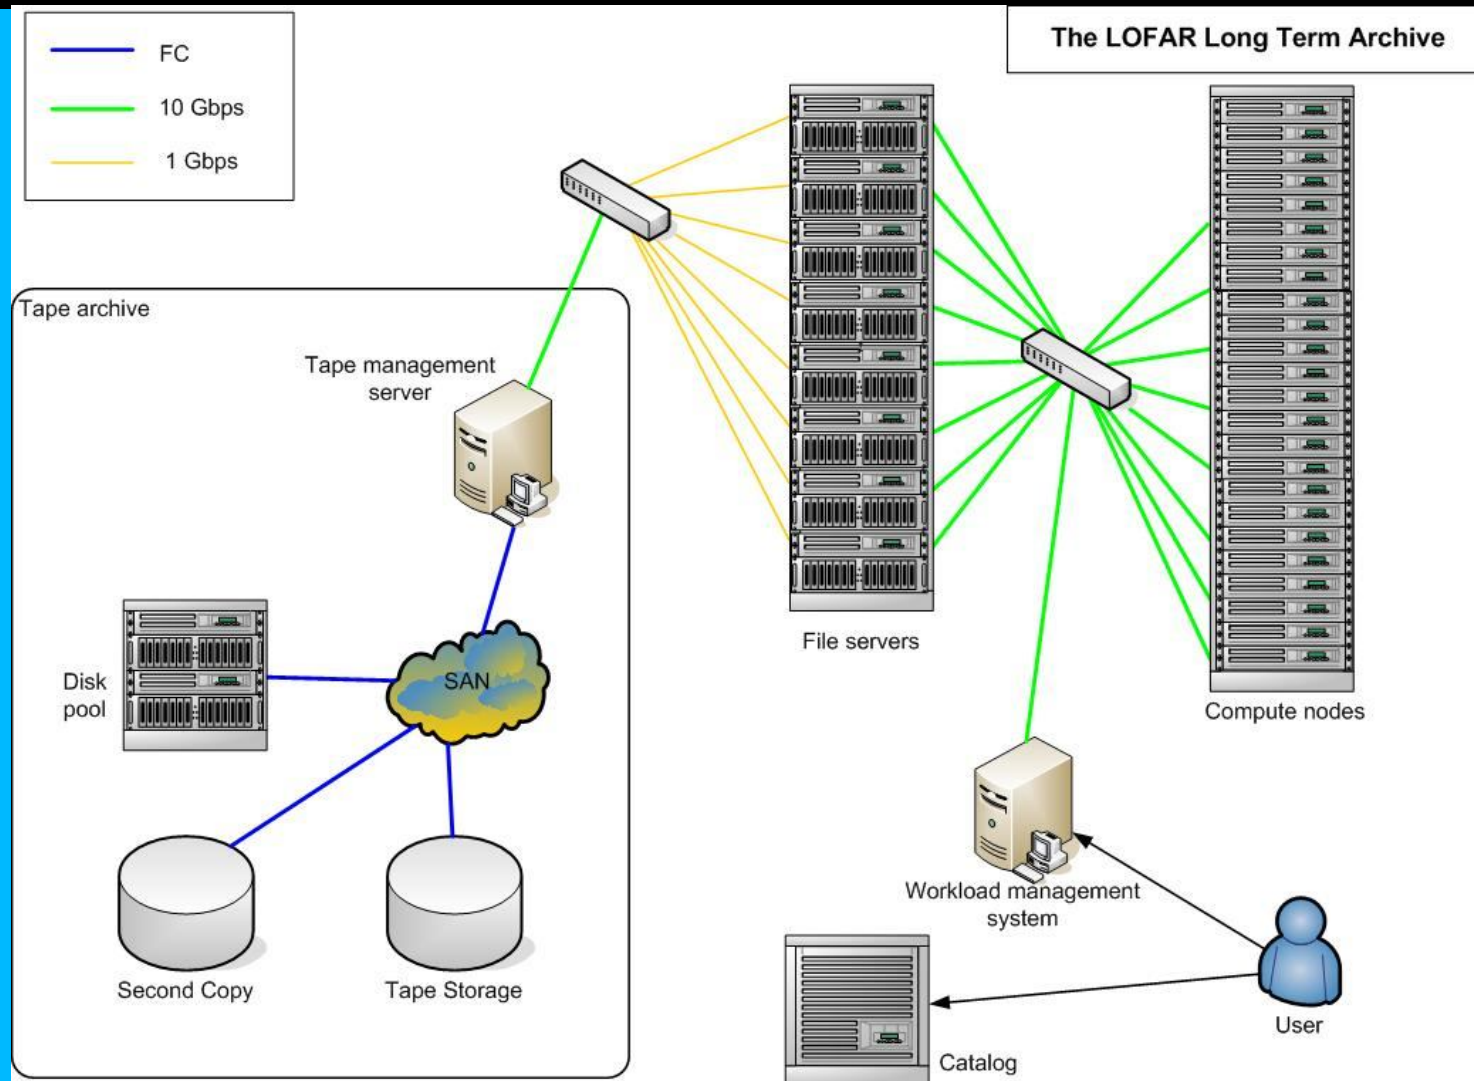
Size: 137134080 bytes (130.8 MB)
 Hash: f7a02350f4802aa04a80ca79ee915f28746687a
 Creation: 04 Aug 2008 10:00:49

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

The LOFAR LTA



The LOFAR LTA Initial configuration



Still many uncertainties!

Some typical (?) processing steps have been considered:

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

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

- **First implementation ready Q3 2009**
- **In particular processing requirements uncertain**
 - Need substantiation through performing realistic tests, soon!
- **BiG Grid facilities are available**
 - Storage
 - Processing
- **Target prototype is available**

- **LOFAR Wiki:** <http://www.lofar.org/operations>
 - System design documents
 - CS1 observation catalog
 - GRID User Interface tutorial
- **LOFAR LTA prototype:**
 - username : awworkshop2009
 - password : rafol
- **BiG Grid:** <http://www.biggrid.nl>
- **USG Wiki:** <http://usg.lofar.org>
- **LOFAR Operator:** observer@astron.nl
 - Coordination usage LOFAR resources