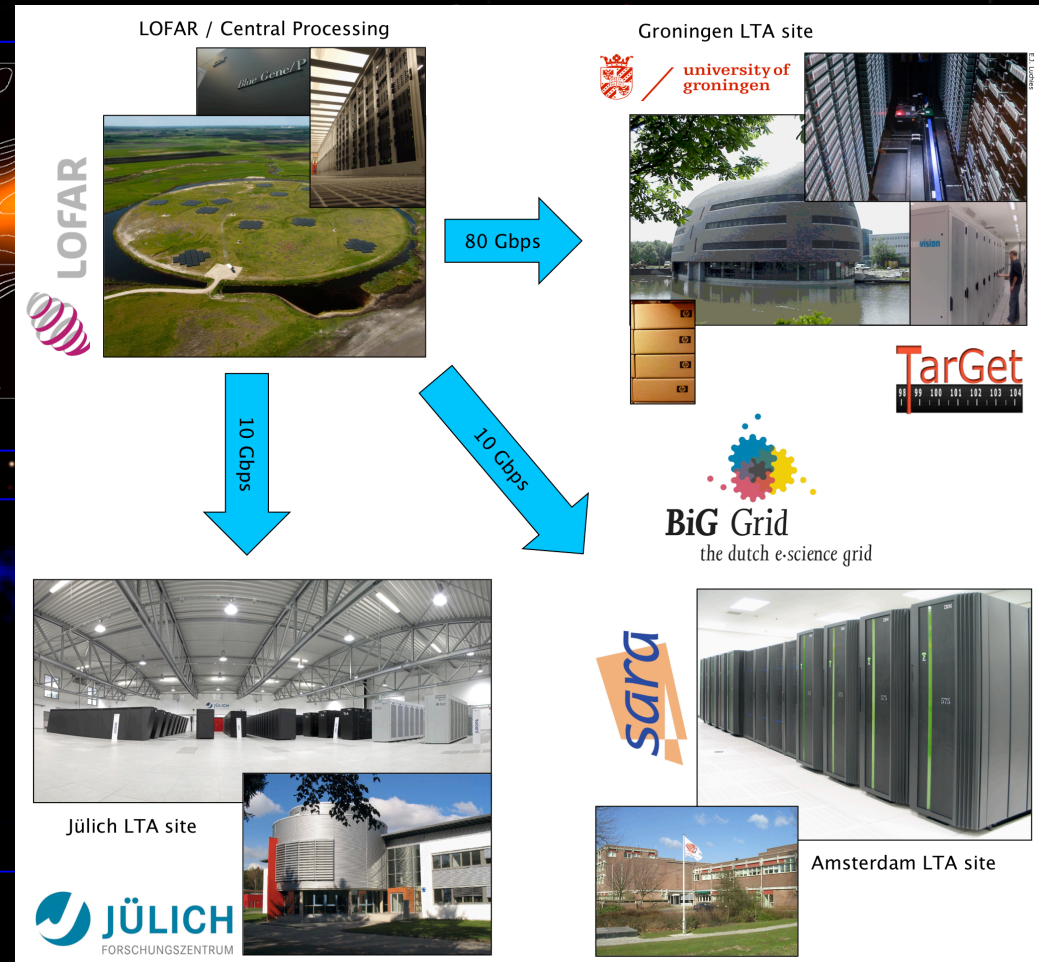

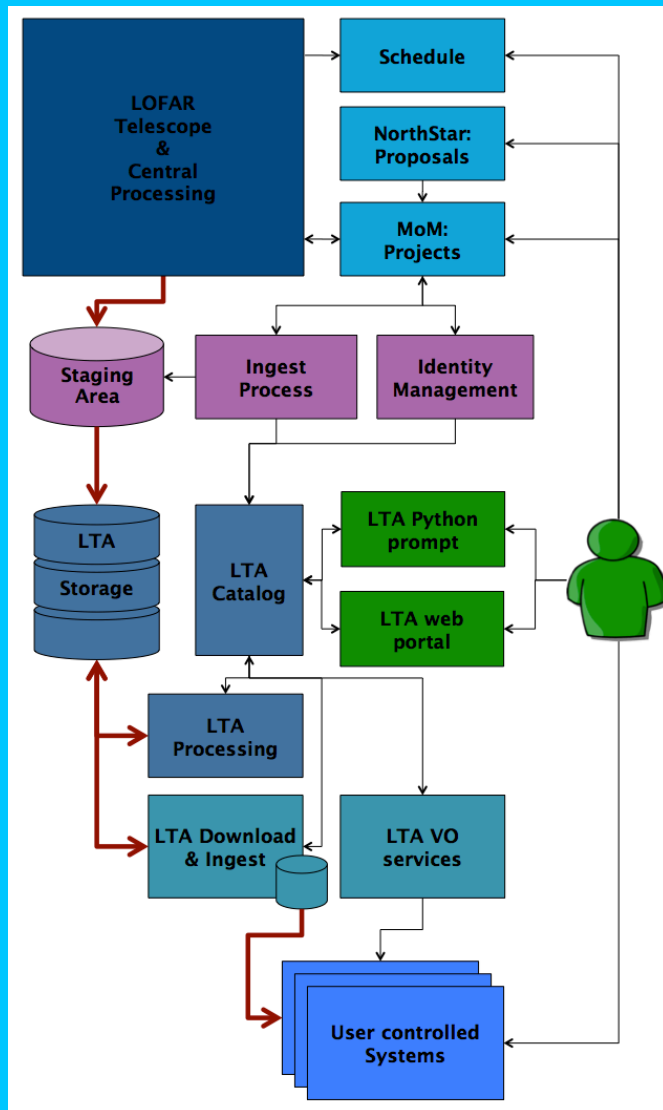


## LOFAR Long Term Archive Overview and Status



# LOFAR LTA System Architecture



## Tier 0 (a.k.a. CEP)

- ( Short term storage.
- ( First iteration processing.
- ( After quality inspection move data to LTA Storage and metadata to LTA Catalog.

## Tier 1 (a.k.a. LTA)

- ( Long term storage.
- ( Further processing.
- ( Catalog can be queried via prompt, web- and VO interfaces.
- ( Data made available via Download Service.

## Tier 2 (a.k.a. 'the rest of the world')

- ( User controlled.
- ( Includes university or faculty storage and processing systems.

# Regular transfers from CEP to LTA



## ( LTA copy

Move data to LTA storage;  
no registration in LTA catalog  
(manual administration)

Applies to everything stored  
in the LTA until now

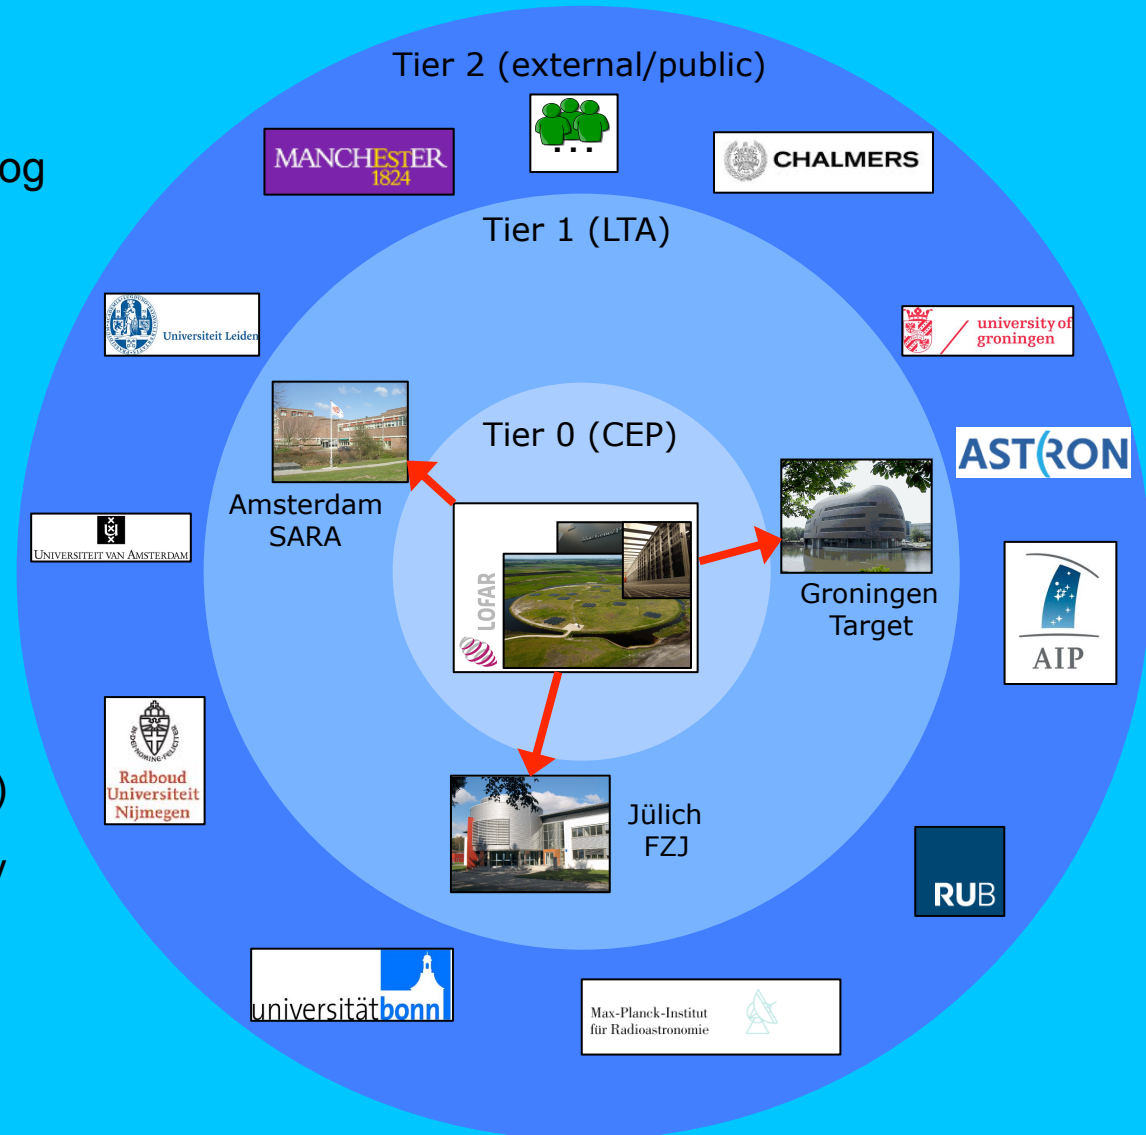
## ( Data ingest (LOFAR 1.0)

Move data to LTA storage  
& update LTA catalog  
in single transaction

From CEP:  
MoM provides required  
metadata (from SAS/MAC)  
Responsibility Observatory  
(timing, placement)

Functional/tested

low user initiated ingest  
from Tiers 1 & 2



# Projects and accounts



## ( Project synchronization

MoM provides LOFAR project administration

Provisioned to LTA catalog when an 'LTA Resource' (storage/processing) is allocated to project

## ( Account synchronization

NorthStar/MoM accounts synchronized with LTA when account is 'LTA enabled'

Automatic for members of project with LTA Resource

Can not distribute MoM/NorthStar password hash

User will receive request by mail to generate new password

Looking into using federated accounts (with SurfNet & ASTRON ICT)

Allow home institute accounts for logging in

A screenshot of the ASTRON login interface. At the top left is the ASTRON logo. To its right, the text reads 'Please Log In' and 'ASTRON Web Applications Password Self Service'. The main content area is a light grey box containing a login form. The form has two input fields: 'Username' with the text 'holties' and 'Current Password' with masked characters '\*\*\*\*\*'. Below the fields are three buttons: 'Login', 'Clear', and 'Cancel'. At the bottom of the form, there is a link for 'Forgotten Password' with the text 'Recover your forgotten password.' At the very bottom of the page, there is a small footer that reads 'Idle Timeout: 4 minutes' and '82.95.247.128 | English'.

# LTA Catalog



## ( LTA Catalog DB

Oracle RAC 11g

LOFAR Datamodel  
Complete data lineage

- Observations
- Pipelines
- Data products

## ( LTA Catalog UI

Builds on AstroWise

Several views:

- Web interface
- DB view
- Console

Request download

- Now: scripts;  
assumes data online
- Soon: staging request  
plus user notification

Home | Contact | Help | user awtier0 | project MSSS | Preferences | Tables

### ImagingPipeline

- Order by: <none>  descending  ascending
- Maximum number of rows: 100
- Show only data within project:  yes  no
- Show expanded attributes:  yes  no [much faster]
- Export options: HTML

Submit

expand all attributes

ImagingPipeline

- calibrationStrategy
- creationDate
- demixing
- duration
- endTime
- frequencyIntegrationStep
- imagerIntegrationTime
- imagerStrategy
- numberOfMajorCycles
- observationId
- observationIdSource
- pipelineName
- pipelineVersion
- processIdentifier
- processIdentifierName
- processIdentifierSource
- skyModelDatabase
- startTime
- strategyDescription
- strategyName
- timeIntegrationStep
- parset
- projectInformation

Submit



### Projects of db.lofar.target.rug.nl

- Number of projects : 2
- Number of Users : 4
- Current user : AWANONYMOUS

Click on a project name to set the project

ID	Project	Privileges	Instrument	Member of	Member count	Manager(s)
1	ALL	4	None	True	public	AWTIERG
401551	MSSS	2	LOFAR	False	2	AWTIERG

powered by



# LTA Catalog AstroWise web interface



**Query SkyImageDataProduct**

**ObservationId**   
**DataProductIdentifier**   
**Pointing** RA  DEC  Range  °  
**TargetName**   
**ObservationDescription**   
**StrategyDescription**  select  
**StationSelection**  
 All  
 Core  
 Custom  
 Dutch  
 Single  
**AntennaSet**  select  
**Stations** Select individual stations +/-

powered by



[Home](#) | 
 [Help](#) | 
 [login \(awtler0\)](#) | 
 [project \(MSSS\)](#) | 
 [Search](#) | 
 [Show Latest](#)

## CorrelatedDataProduct (total 80)

[edit columns](#) | 
 [download script](#) | 
 [http srm](#)

#	dataProductIdentifier	fileFormat	subArrayPointingIdentifier	subband	startTime	duration	centralFrequency [MHz]	rightAscension [degrees]	declination [degrees]	creationDate	pipeline	Observations	Stations	pipeline Results
1	2362470	AIPS++/CASA	109110	0	2012-03-31 03:45:01	120.0	30.078125	299.867916667	40.7339138889	2012-05-07 16:06:12		1	show	
2	2362493	AIPS++/CASA	109110	23	2012-03-31 03:45:01	120.0	42.578125	299.867916667	40.7339138889	2012-05-07 16:04:08		1	show	
3	2362494	AIPS++/CASA	109110	24	2012-03-31 03:45:01	120.0	42.7734375	299.867916667	40.7339138889	2012-05-07 16:04:08		1	show	
4	2362549	AIPS++/CASA	109110	79	2012-03-31 03:45:01	120.0	74.8046875	299.867916667	40.7339138889	2012-05-07 15:37:59		1	show	
5	2362547	<b>File Details</b>					74.4140625	299.867916667	40.7339138889	2012-05-07 15:37:58		1	show	
6	2362548						74.609375	299.867916667	40.7339138889	2012-05-07 15:37:57		1	show	
7	2362546	filename L55294_SAP000_SB077_uv_MS_de98df86.tar					74.21875	299.867916667	40.7339138889	2012-05-07 15:37:37		1	show	
8	2362545	filesize 202.8 MB					74.0234375	299.867916667	40.7339138889	2012-05-07 15:37:18		1	show	
9	2362543	creation_date 2012-05-07 15:37:58					73.6328125	299.867916667	40.7339138889	2012-05-07 15:37:17		1	show	
10	2362544	URI srm://srm.grid.sara.nl:8443/pnfs/grid.sara.nl/data/lofar/ops/test/msss/55294					73.828125	299.867916667	40.7339138889	2012-05-07 15:37:17		1	show	
11	2362542	/L55294_SAP000_SB077_uv_MS_de98df86.tar					73.4375	299.867916667	40.7339138889	2012-05-07 15:36:56		1	show	
12	2362541	protocol srm					73.2421875	299.867916667	40.7339138889	2012-05-07 15:36:37		1	show	
13	2362539	hash_md5 cab977f92000318d0ded1f147ec9980e					67.1875	299.867916667	40.7339138889	2012-05-07 15:36:32		1	show	
14	2362540	hash_adler32 71a95f46					73.046875	299.867916667	40.7339138889	2012-05-07 15:36:32		1	show	
15	2362538	AIPS++/CASA	109110	68	2012-03-31 03:45:01	120.0	66.9921875	299.867916667	40.7339138889	2012-05-07 15:36:20		1	show	
16	2362537	AIPS++/CASA	109110	67	2012-03-31 03:45:01	120.0	66.796875	299.867916667	40.7339138889	2012-05-07 15:35:55		1	show	
17	2362535	AIPS++/CASA	109110	65	2012-03-31 03:45:01	120.0	66.40625	299.867916667	40.7339138889	2012-05-07 15:35:47		1	show	
18	2362536	AIPS++/CASA	109110	66	2012-03-31 03:45:01	120.0	66.6015625	299.867916667	40.7339138889	2012-05-07 15:35:45		1	show	
19	2362534	AIPS++/CASA	109110	64	2012-03-31 03:45:01	120.0	66.2109375	299.867916667	40.7339138889	2012-05-07 15:35:35		1	show	
20	2362533	AIPS++/CASA	109110	63	2012-03-31 03:45:01	120.0	66.015625	299.867916667	40.7339138889	2012-05-07 15:35:15		1	show	
21	2362531	AIPS++/CASA	109110	61	2012-03-31 03:45:01	120.0	65.625	299.867916667	40.7339138889	2012-05-07 15:35:06		1	show	
22	2362532	AIPS++/CASA	109110	62	2012-03-31 03:45:01	120.0	65.8203125	299.867916667	40.7339138889	2012-05-07 15:35:05		1	show	
23	2362530	AIPS++/CASA	109110	60	2012-03-31 03:45:01	120.0	65.4296875	299.867916667	40.7339138889	2012-05-07 15:34:55		1	show	

# LTA data retrieval

## ( Web based download server

'LTA enabled' ASTRON/  
LOFAR account

Low threshold

Primarily for few files  
& smaller volumes

## ( GridFTP

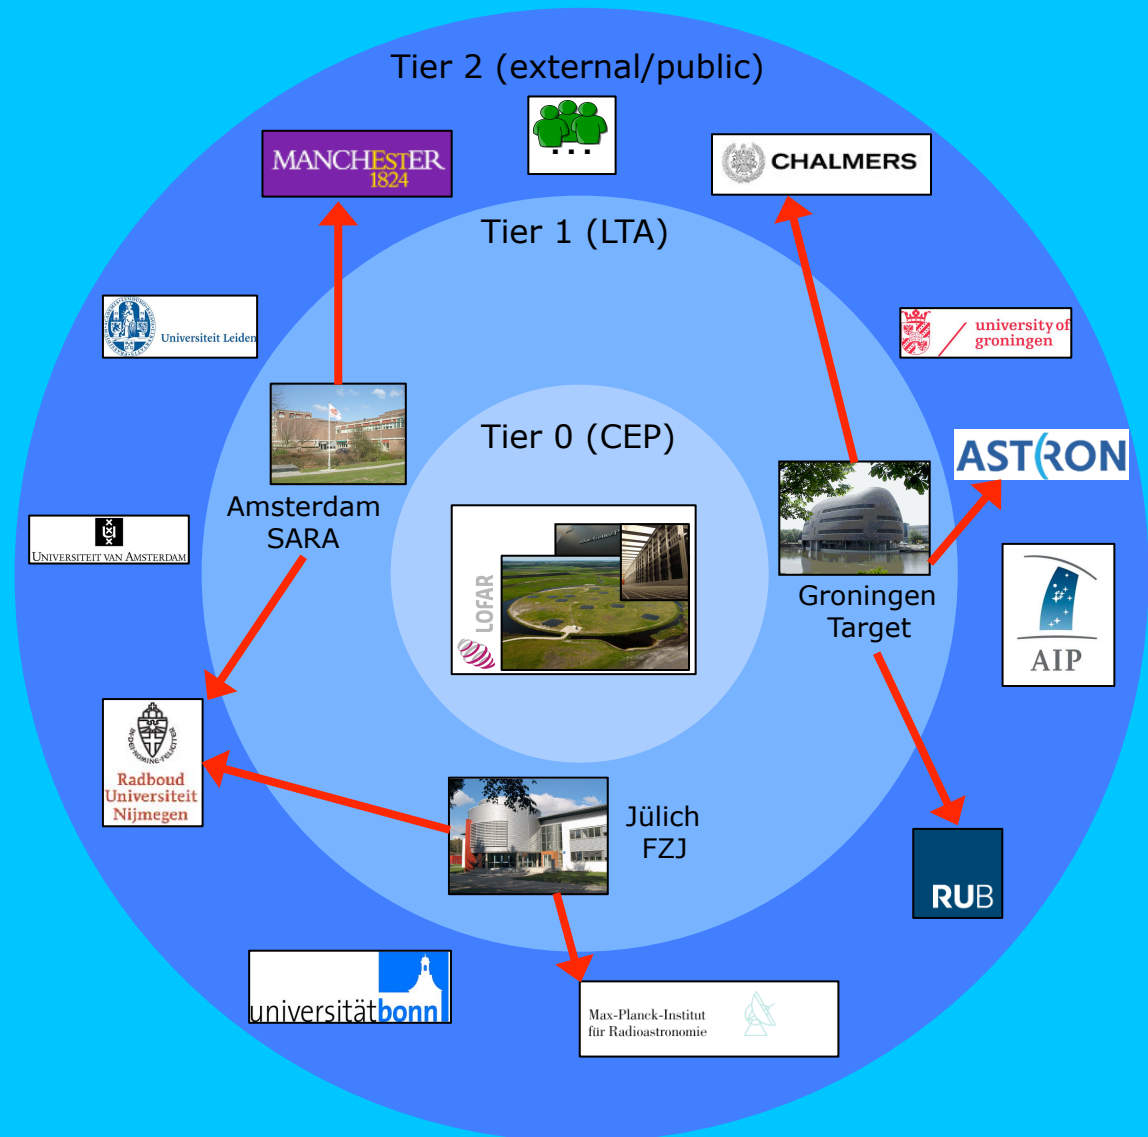
Requires grid  
user certificate

Requires grid  
client installation

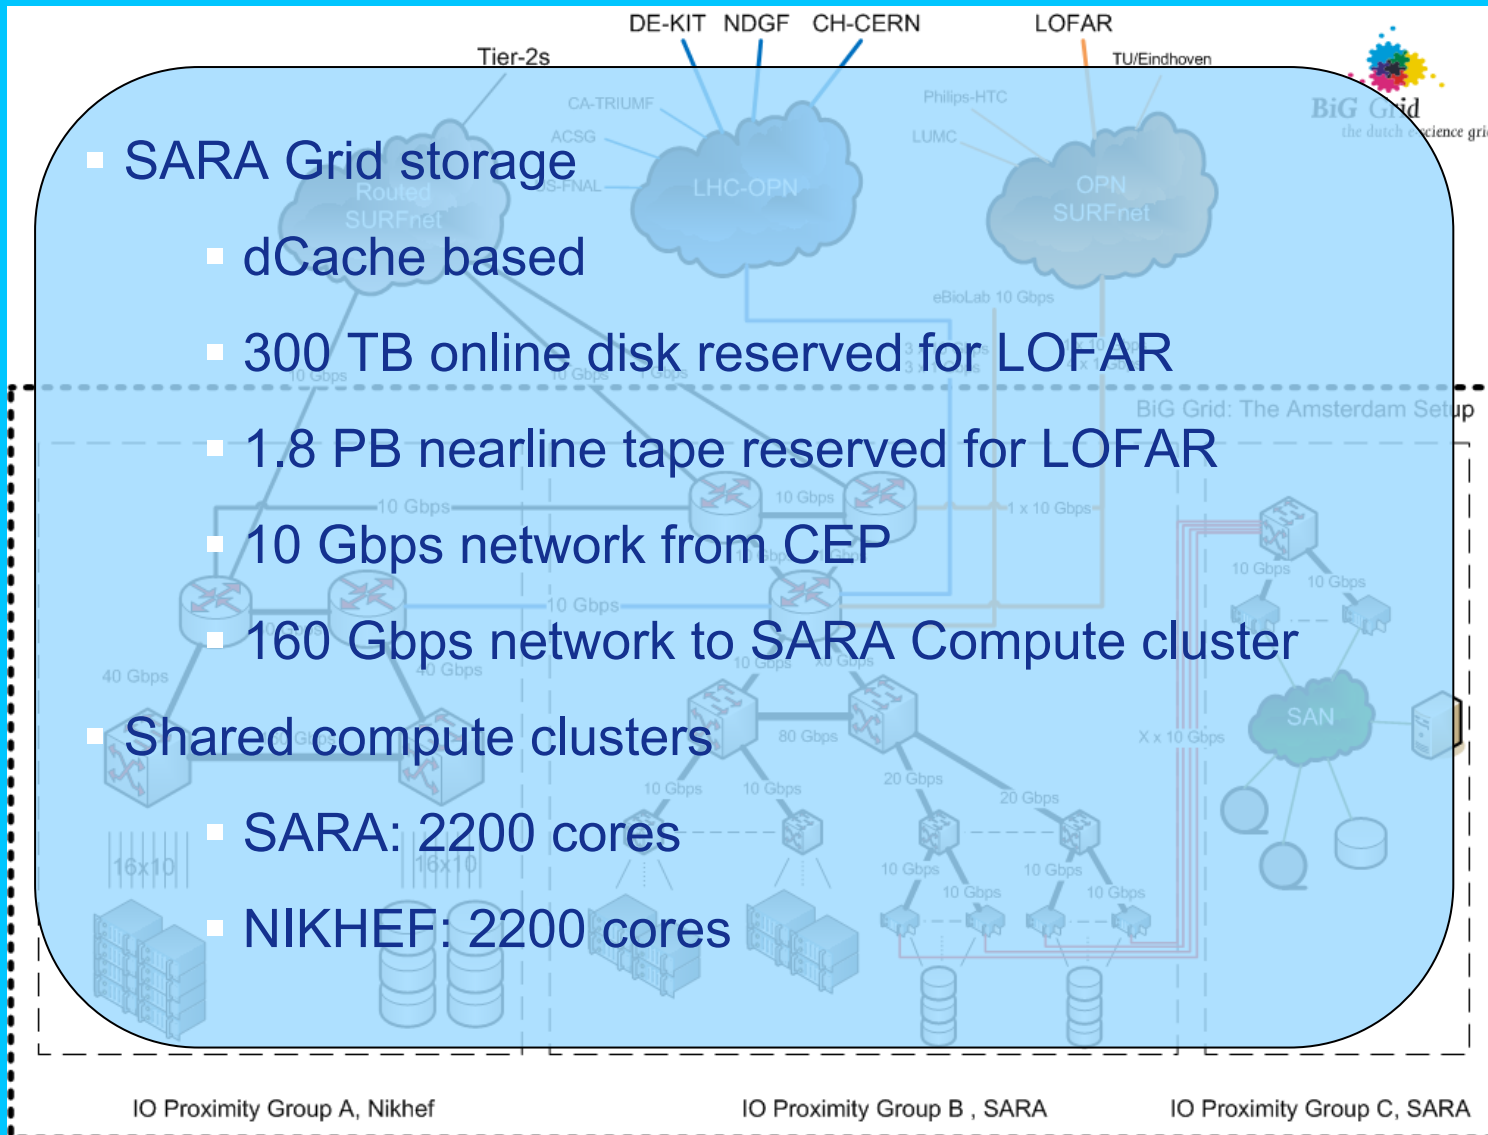
More robust;  
superior performance

## ( Future: client tool

Best of both worlds?



# Amsterdam/SARA






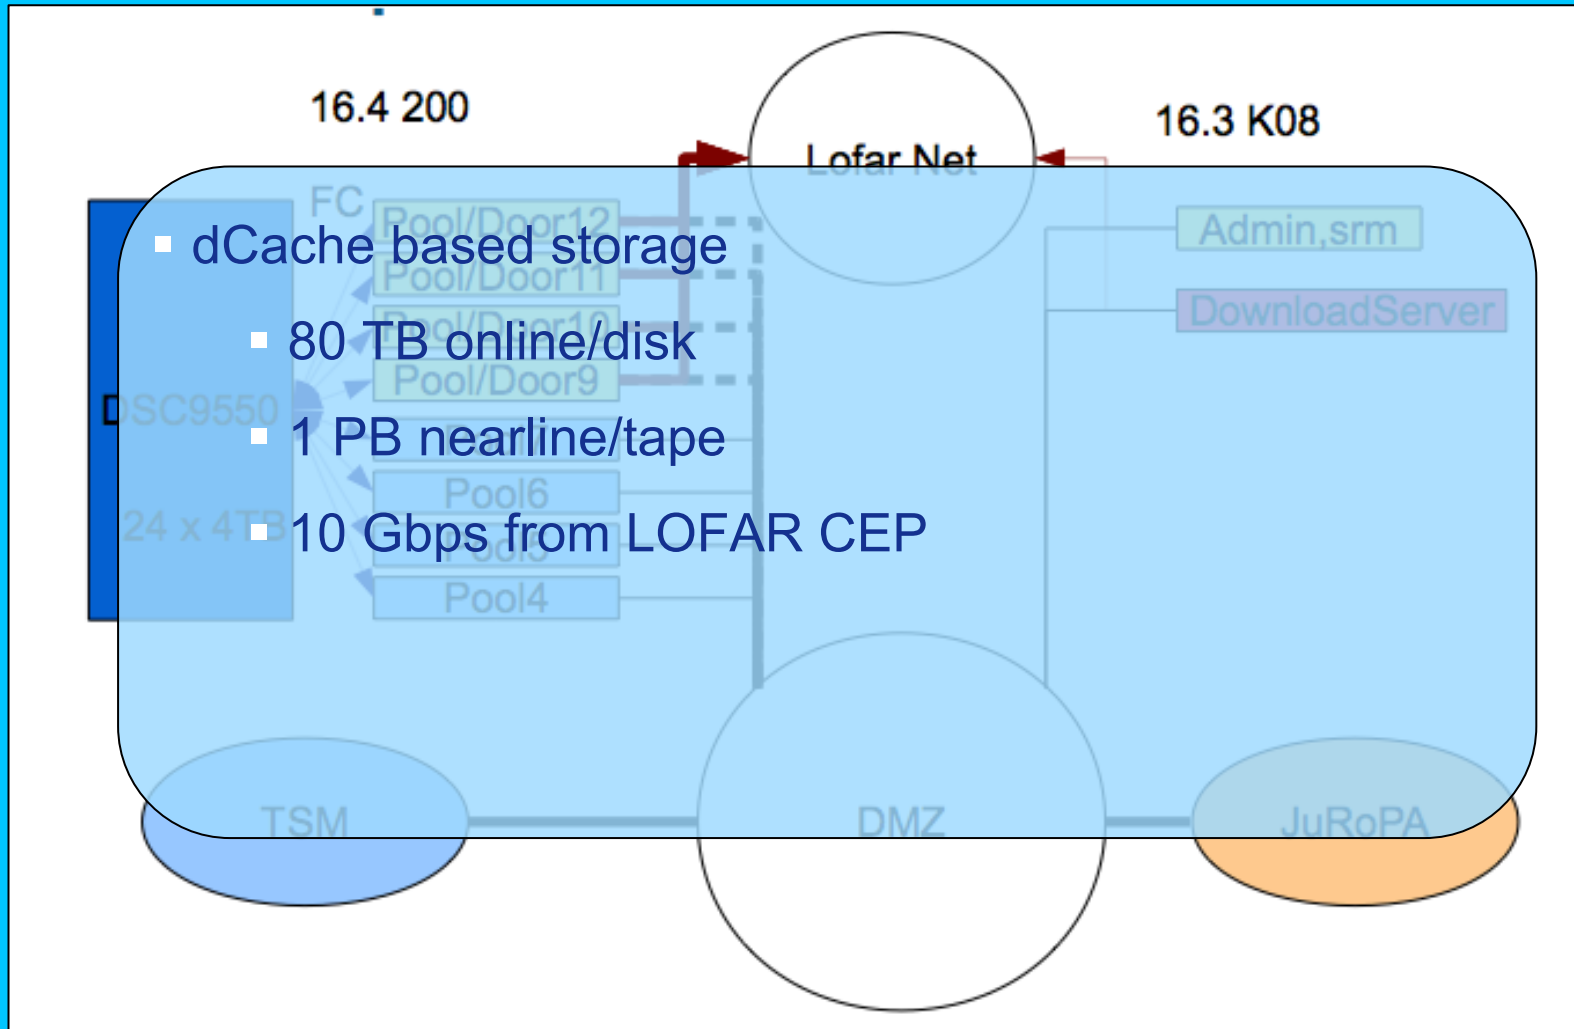


# Amsterdam/SARA Throughput



- ( ~ 800Mbps; drops when online disk area allocated to LOFAR is full
- Migration to tape becomes limiting
- Additional migration capacity can (temporarily) be allocated on request





# Jülich/FZJ Throughput




( ~ 600MBps first couple of hours; drops significantly after some time

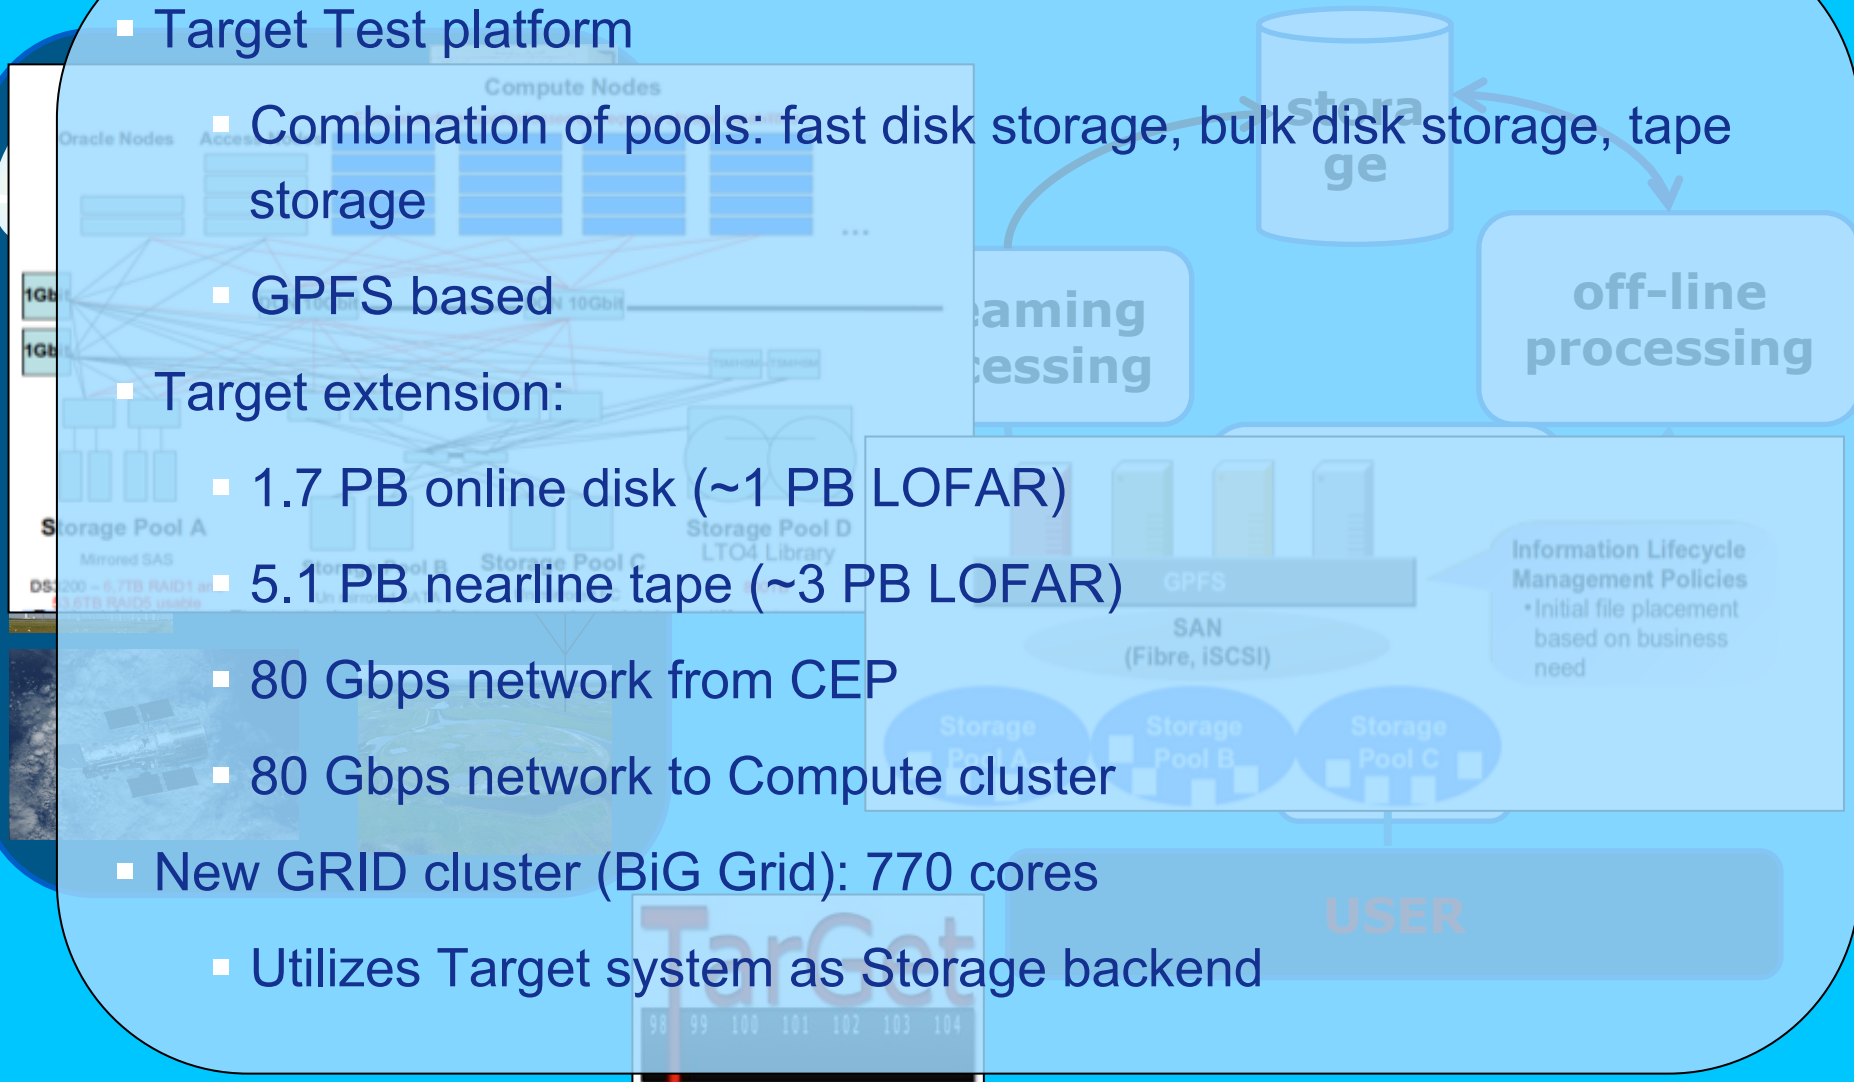
( Inefficient scheduling mechanism of dCache as installed at FZJ

Expect improvement after dCache upgrade (being tested)

Further improvement expected from disksreplacements this summer



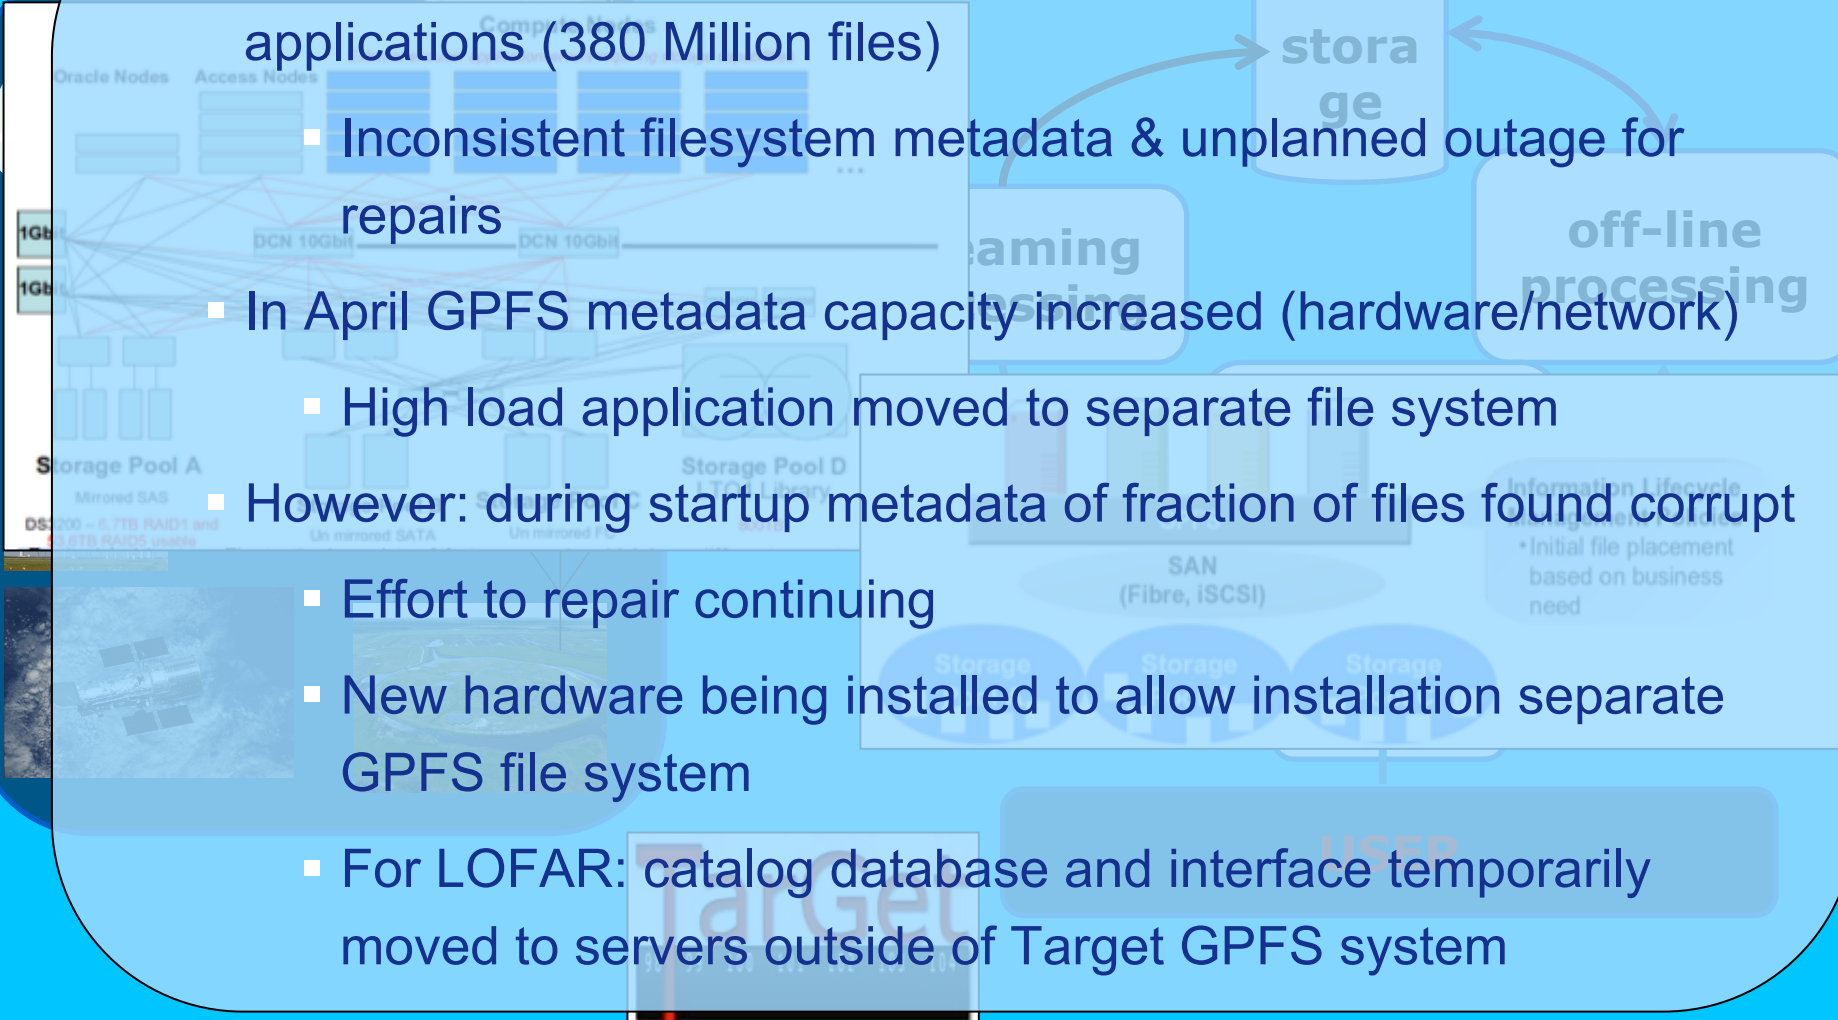
- Target Test platform
  - Combination of pools: fast disk storage, bulk disk storage, tape storage
  - GPFS based
- Target extension:
  - 1.7 PB online disk (~1 PB LOFAR)
  - 5.1 PB nearline tape (~3 PB LOFAR)
  - 80 Gbps network from CEP
  - 80 Gbps network to Compute cluster
- New GRID cluster (BiG Grid): 770 cores
  - Utilizes Target system as Storage backend



# Groningen/Target: recent issues



- GPFS filesystem unable to handle load generated by one of the applications (380 Million files)
  - Inconsistent filesystem metadata & unplanned outage for repairs
  - In April GPFS metadata capacity increased (hardware/network)
    - High load application moved to separate file system
  - However: during startup metadata of fraction of files found corrupt
    - Effort to repair continuing
    - New hardware being installed to allow installation separate GPFS file system
  - For LOFAR: catalog database and interface temporarily moved to servers outside of Target GPFS system



# Groningen/Target Throughput



## ( **Not yet tested**

Before upgrade (single Target access node): at least as good as SARA


## ( **Upgrade:**

Four Target access nodes with 2x 10 Gbps each

Objective is to achieve > 60 Gbps for offloading observation data from CEP2

# LTA Storage Capacity

All in TB	Disk 1/2012	1/2013	In use 5/2012	Tape 1/2012	1/2013	In use 5/2012
SARA	300	500	250	1800	3700	1000
TARGET	300	1000	70	0	1700	350
FZJ	80	?	80	1000	?	150



# 2012 Activities



## ( **Support for other dataproducts**

Beamformed

TBB

Single station

## ( **Enhancing the user interface**

Usability

Robustness

Functionality

## ( **Source catalog**

## ( **Extend LTA datamodel**

a.o. quality related metadata

## ( **Manually operated processing**

Pulsar processing (started)

Imaging processing

NDPPP/BBS/Imager components

MSSS preprocessing

LOFAR software available!

## ( **LTA pipeline integration**

AstroWise DPU pipeline framework

Direct access to LTA catalog

Ingest output

## ( **Bandwidth on Demand**

Data distribution

With SurfNet/JIVE (NEXPreS)



# Data distribution within LTA

## Requirements

### i) Duplication/repair

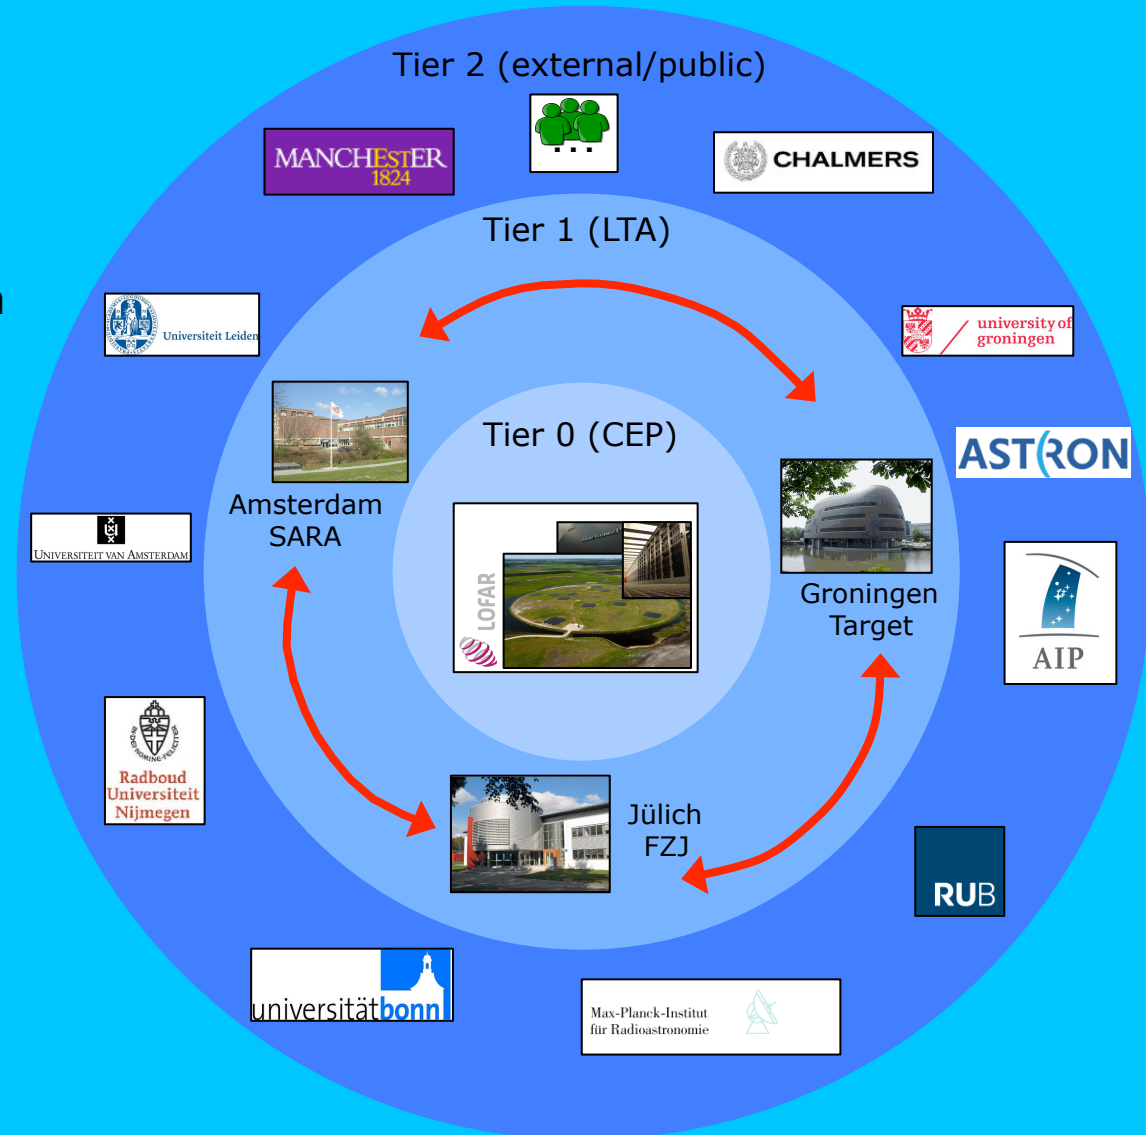
Flexible scheduling

QOS more important than exact bandwidth

### ii) Remote processing

Full chain scheduling

Specific bandwidth;  
- NEXPreS: ~10 Gbps  
- Ideally comparable to LAN speeds: ~80 Gbps



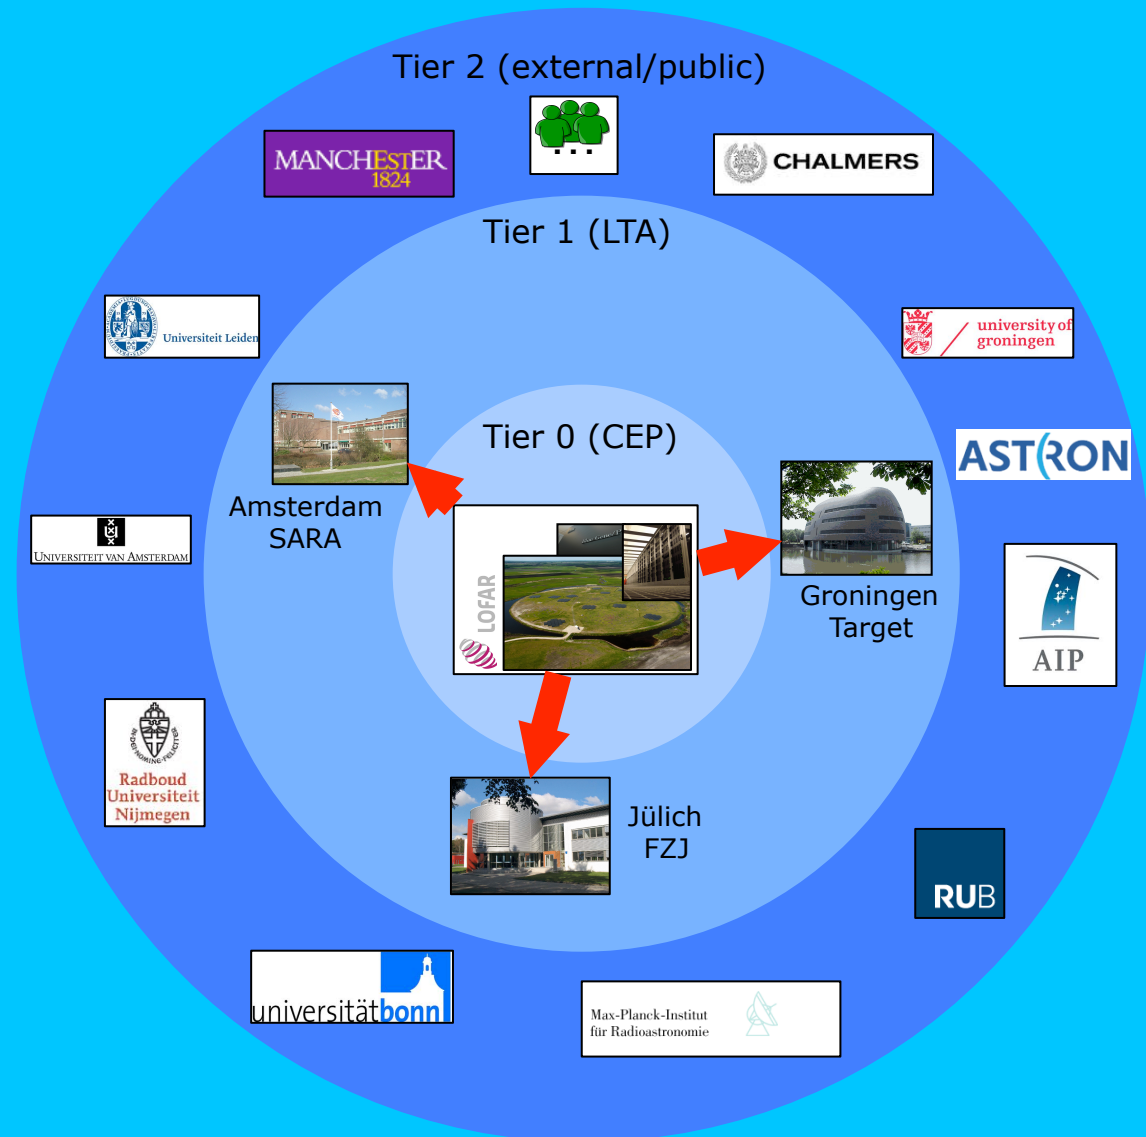
# Case IV: Raw data off-loading from Tier 0 to Tier 1

## Requirements

- ( Unidirectional (mostly)
- ( X PB per year
- ( 60 - 80 Gbps
- Keep up with peak observation data rate
- Only fraction of time
- ( Reliability important
- Not 'life threatening'
- But impacts observation program (which is the driver)

## Infrastructure

- ( Shared with station network (mostly); c.w. Case III

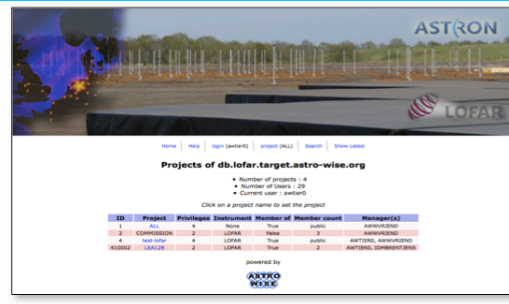


# LTA Science Interfaces




## Data Access

Open access to all LOFAR and APERTIF standard science data and derived products is provided via the archive.



## Data Processing

The archive offers users capabilities for custom processing of their data. By unlocking the power of international HPC infrastructure, potential science output is increased dramatically.



## Data Connections

Existing VO tools and protocols allow researchers to easily combine data from many telescopes simultaneously. By utilizing the VO infrastructure, the LTA will add LOFAR and APERTIF data to the astronomers toolbox.