Computing Infrastructure for XENON1T

Alfredo Davide Ferella for the XENON Computing WG

INFN Computing Workshop, February 26th, 2014

The XENON Project

Detectors of increasing size and increasing computing technology needs

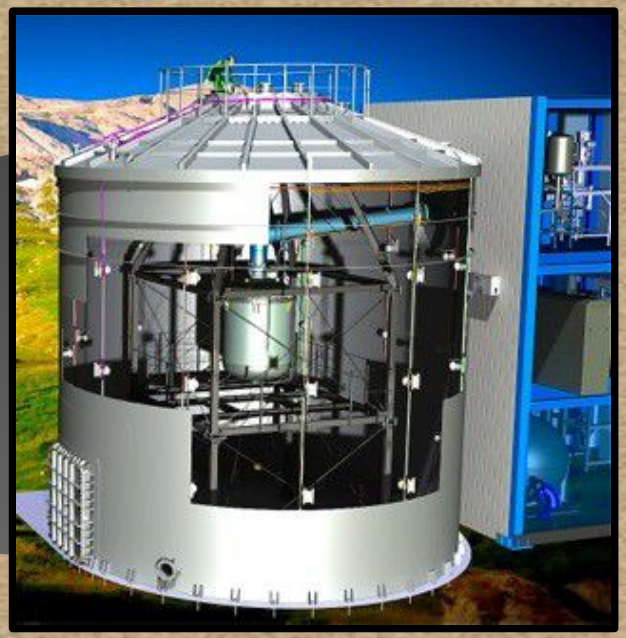

XENON10

Data taking in 2006-2008

Scientific goal achieved

Program terminated

XENON100

Data taking in 2009-today

Scientific goal achieved

Still in operation
to go beyond
its expected performances!

XENON1T

Under construction

Scientific data taking
foreseen in 2015

The XENON100 Detector activity

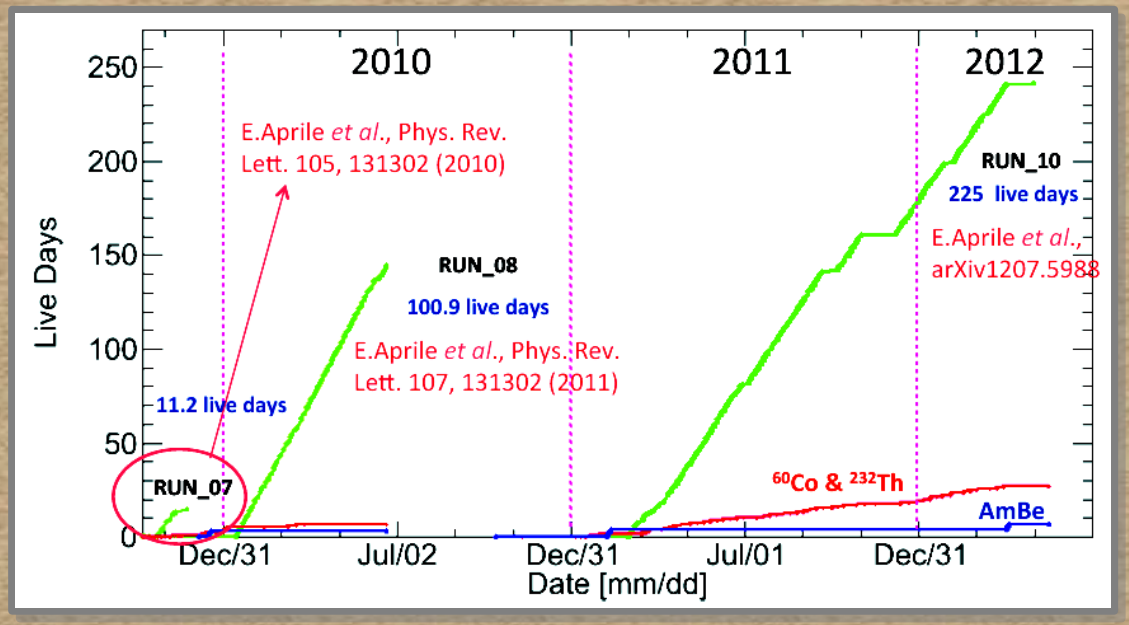

As an example, the latest published scientific run (2011-2012):

- 225 days of Dark Matter search in the cosmic silence (light-weight data)

- 41 days of gamma calibration (heavy-weight data!)

- 7 days of neutron calibration (light-weight but needs of prompt answers)

The XENON100 computing infrastructure

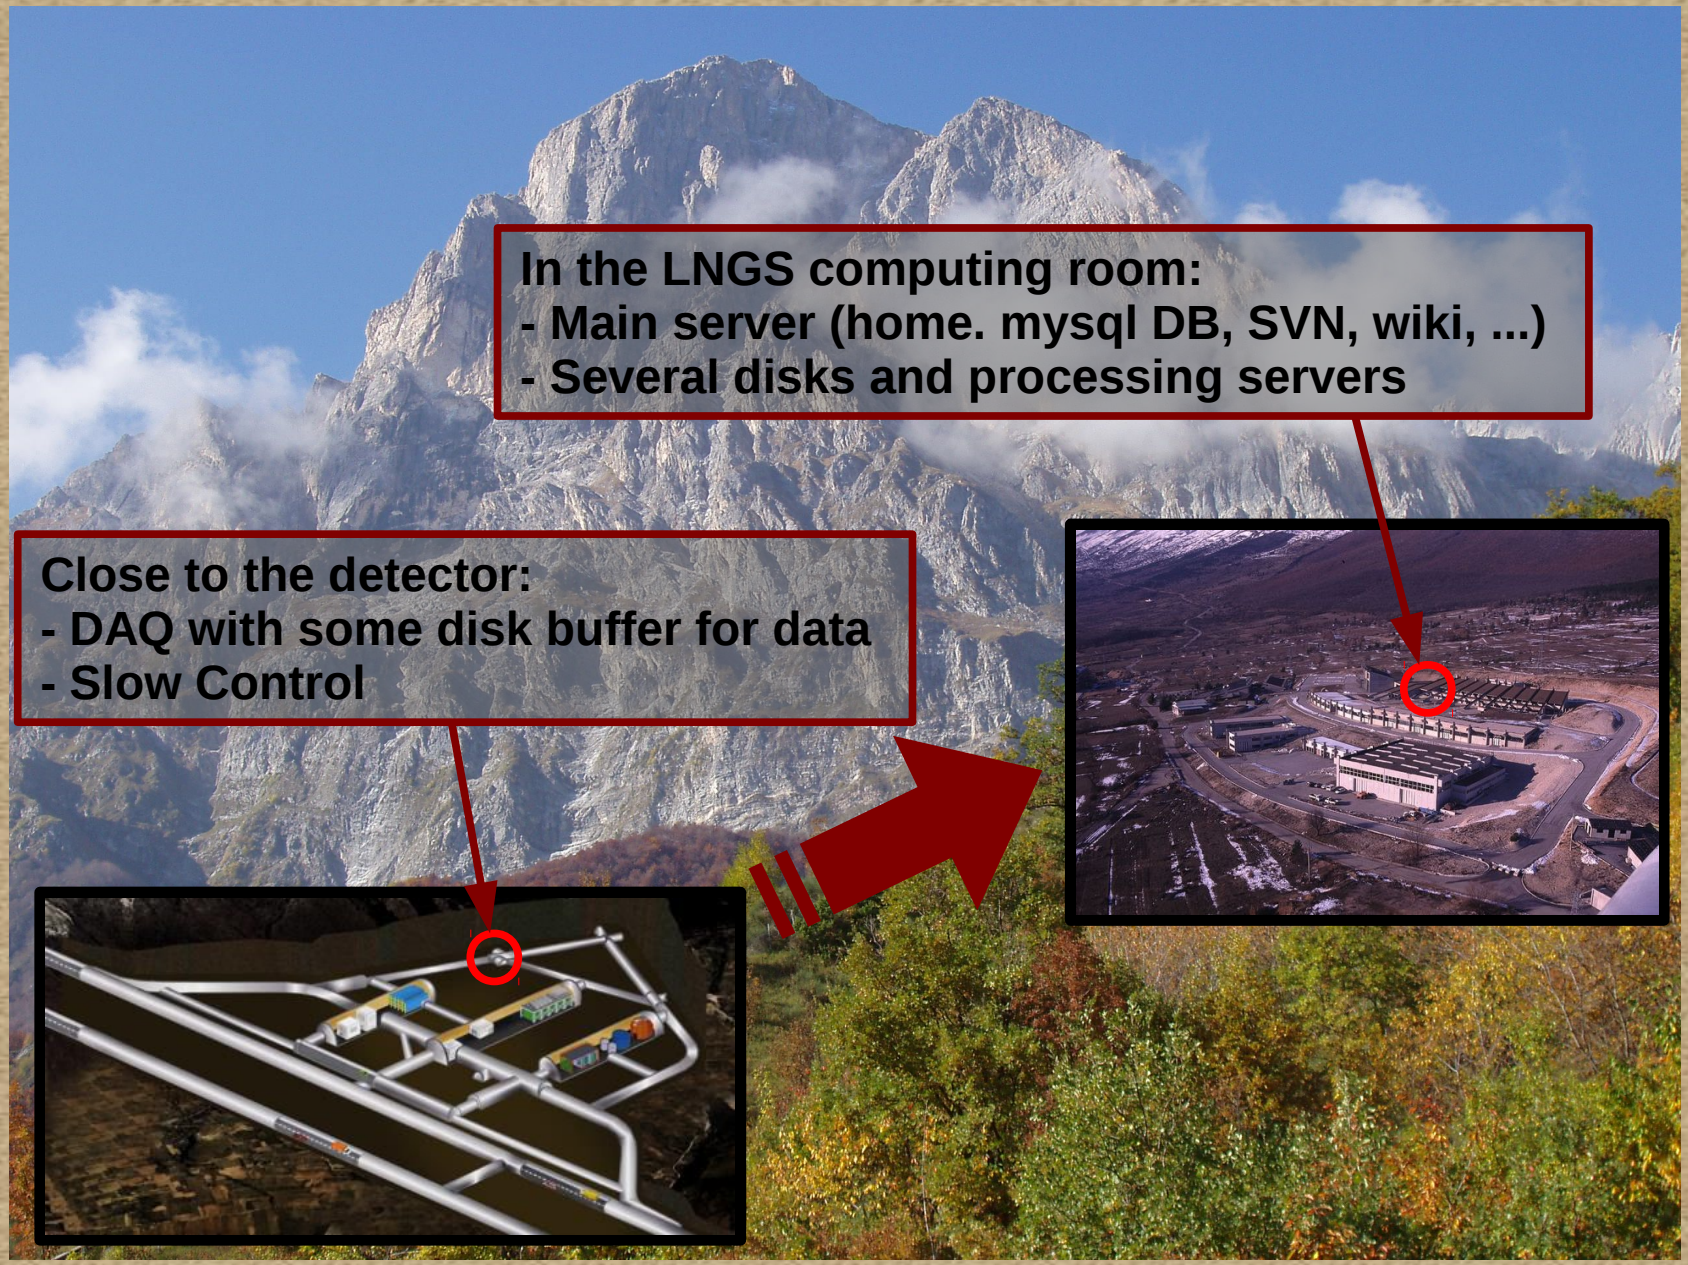

The XENON100 infrastructure: DAQ

Located underground

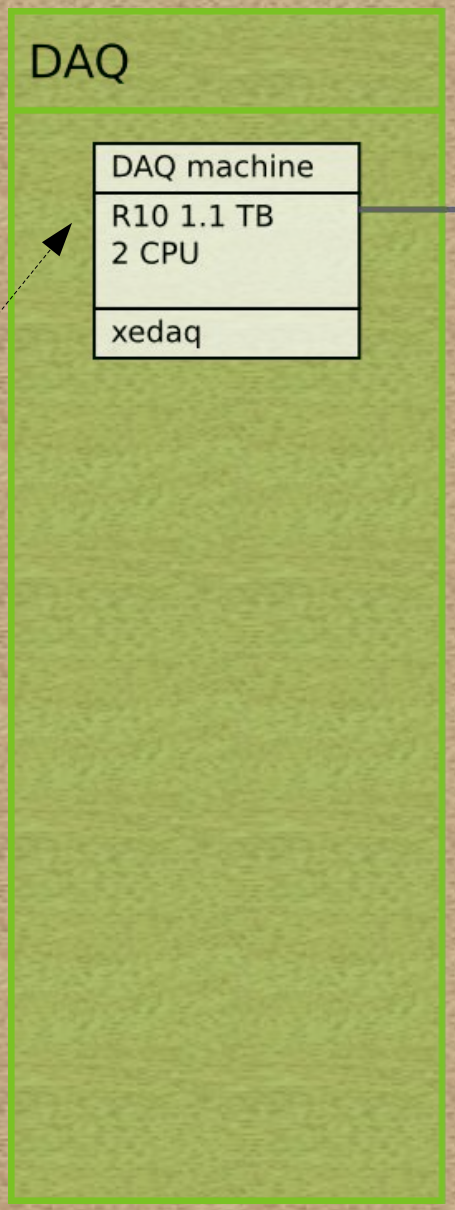

Storage buffer: 1.1 TB

To continue taking data
In case of network issues

Connection to the XENON computing facility in surface

The XENON100 infrastructure: Storage

Located in surface

Local Servers

Disk Server
R6 46 TB 24 CPU Raw and proc data
xecluster07

Disk Server
R6 13 + 9 = 22 TB 4 CPU Raw, proc, common
xecluster19

Disk Server
R5 15 TB 2 CPU Raw and proc data
xecluster20

Storage

Connection to the DAQ

Processing Server
4.6 TB 32 CPU Data processing
xecluster06

Home/DB/Web svr
R5 2.1 TB (home) 4 CPU apache+svn+db
xecluster01

Processing

Connection to external facilities

Analysis Server
8 CPU Analysis
xecluster02

Analysis Server
8 CPU Analysis
xecluster03

Analysis

The XENON100 infrastructure: DAQ

Data flow

Dark Matter search

Frequence: 1 Hz
Event occupancy: 0.9 MB/event

0.9 MB/s

78 GB/day

Gamma calibration

Frequence: 25 Hz
Event occupancy: 0.6 MB/event

15 MB/s

1.3 TB/day

Neutron calibration

Frequence: 25 Hz
Event occupancy: 0.07 MB/event

1.7 MB/s

150 GB/day

The XENON100 infrastructure: DAQ

	Data flow	Total data in one run
Dark Matter search Frequence: 1 Hz Event occupancy: 0.9 MB/event	0.9 MB/s	17 TB
Gamma calibration Frequence: 25 Hz Event occupancy: 0.6 MB/event	15 MB/s	53 TB
Neutron calibration Frequence: 25 Hz Event occupancy: 0.07 MB/event	1.7 MB/s	1 TB

Resources used (so far) in XENON100

Latest scientific run:

Dark Matter search

17 TB of raw data
0.4 TB of processed data
11k CPU-hours to process them

Gamma calibration

53 TB of raw data
4 TB of processed data
61k CPU-hours to process them

Neutron calibration

1 TB of raw data
0.1 TB of processed data
3.5k CPU-hours to process them

TOTAL USED RESOURCES

71 TB of raw data
4.5 TB of processed data
76k processing CPU-hours

Data are stored in disks and also in tapes as backup copy

The whole XENON100 infrastructure

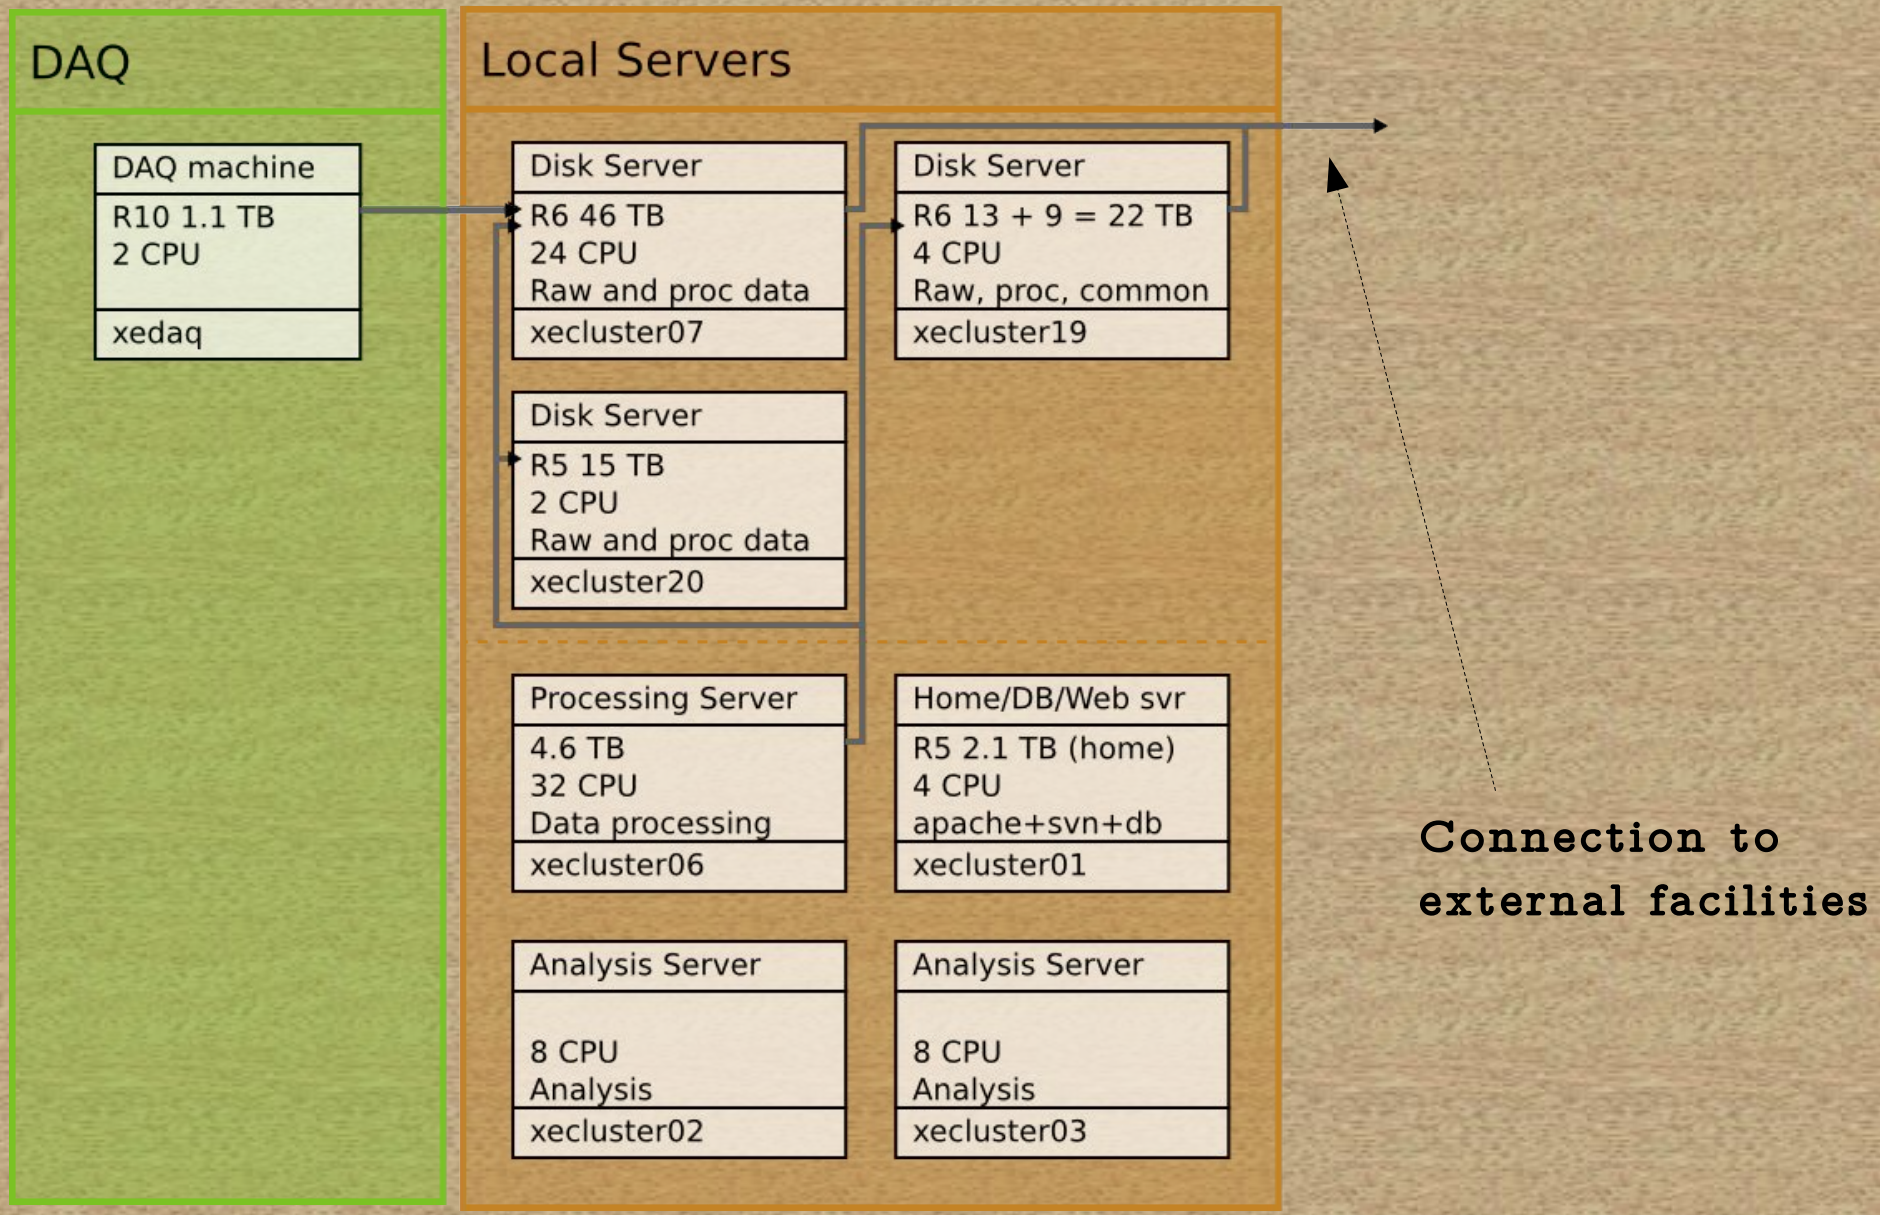

Using LNGS resources

In few years, the high amount of XENON100 data required the use of **increased storage and data processing resources** (CPU)

In 2012:
- Our **disk capacity** was 80 TB of data and it was almost saturated
- We had 48 cores. Enough for daily **data processing**, but not for data reprocessing, requiring at least a factor 2 more.
- Need of a more automated **backup** system and abandon the old LT04 tapes

Since 2012, we took advantage of three LNGS resources, so that today we can profit of:

- **Storage**: 131 TB more. Disks bought by us but maintained by LNGS
- **CPU**: intensive use of U-LITE : 6 nodes with 24 shared CPUs each, plus 16 CPUs always available for XENON. Used for raw **data processing and MC simulations**
- **Tapes**: data backup with LT04 (until we finish our stock), then with new LT06 tapes

Since 2012, all XENON100 data processing and XENON1T MC simulation could not have been done so efficiently without the contribution of the LNGS resources

Using GRID for MC simulations

Monte Carlo simulation is required for XENON1T for:
- background estimation
- detector performances (e.m. and neutron background, light collection efficiency, waveform generation, ...)

In 2012 we started investigating also the use of **GRID**.
Technology and software existing already for XENON1T MC.
Under test by few users.

Several resources are available: **CNAF**, **Nikhef** and **USA** (opensciencegrid.org)

VO: xenon.biggrid.nl

UI at CNAF, LNGS and Nikhef
(CNAF and Nikhef usable
by local users only)

GRID resources available so far

CNAF:
- Storage: 20 TB available
- CPU: 500 HS06 CPU power (~50 cores) allocated,
  allowed peaks of 3000 HS06 per day

They have also Tapes (never used, but we are aware of this possibility)

After a few months of "burn-up" (end of 2013)
we're ready to fully run our MC simulations on GRID

Nikhef:
- Storage: 17 TB available
- CPU: 5000 CPUh / year

USA (opensciencegrid.org):
- CPU: 1 MCPUh / year (shared with others)

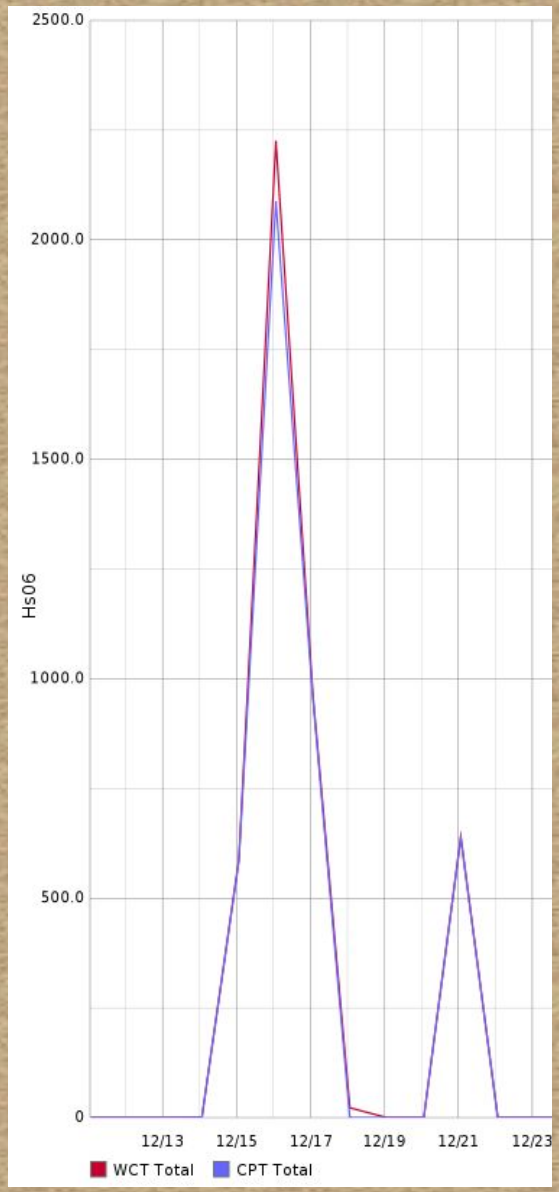

The whole XENON100 infrastructure

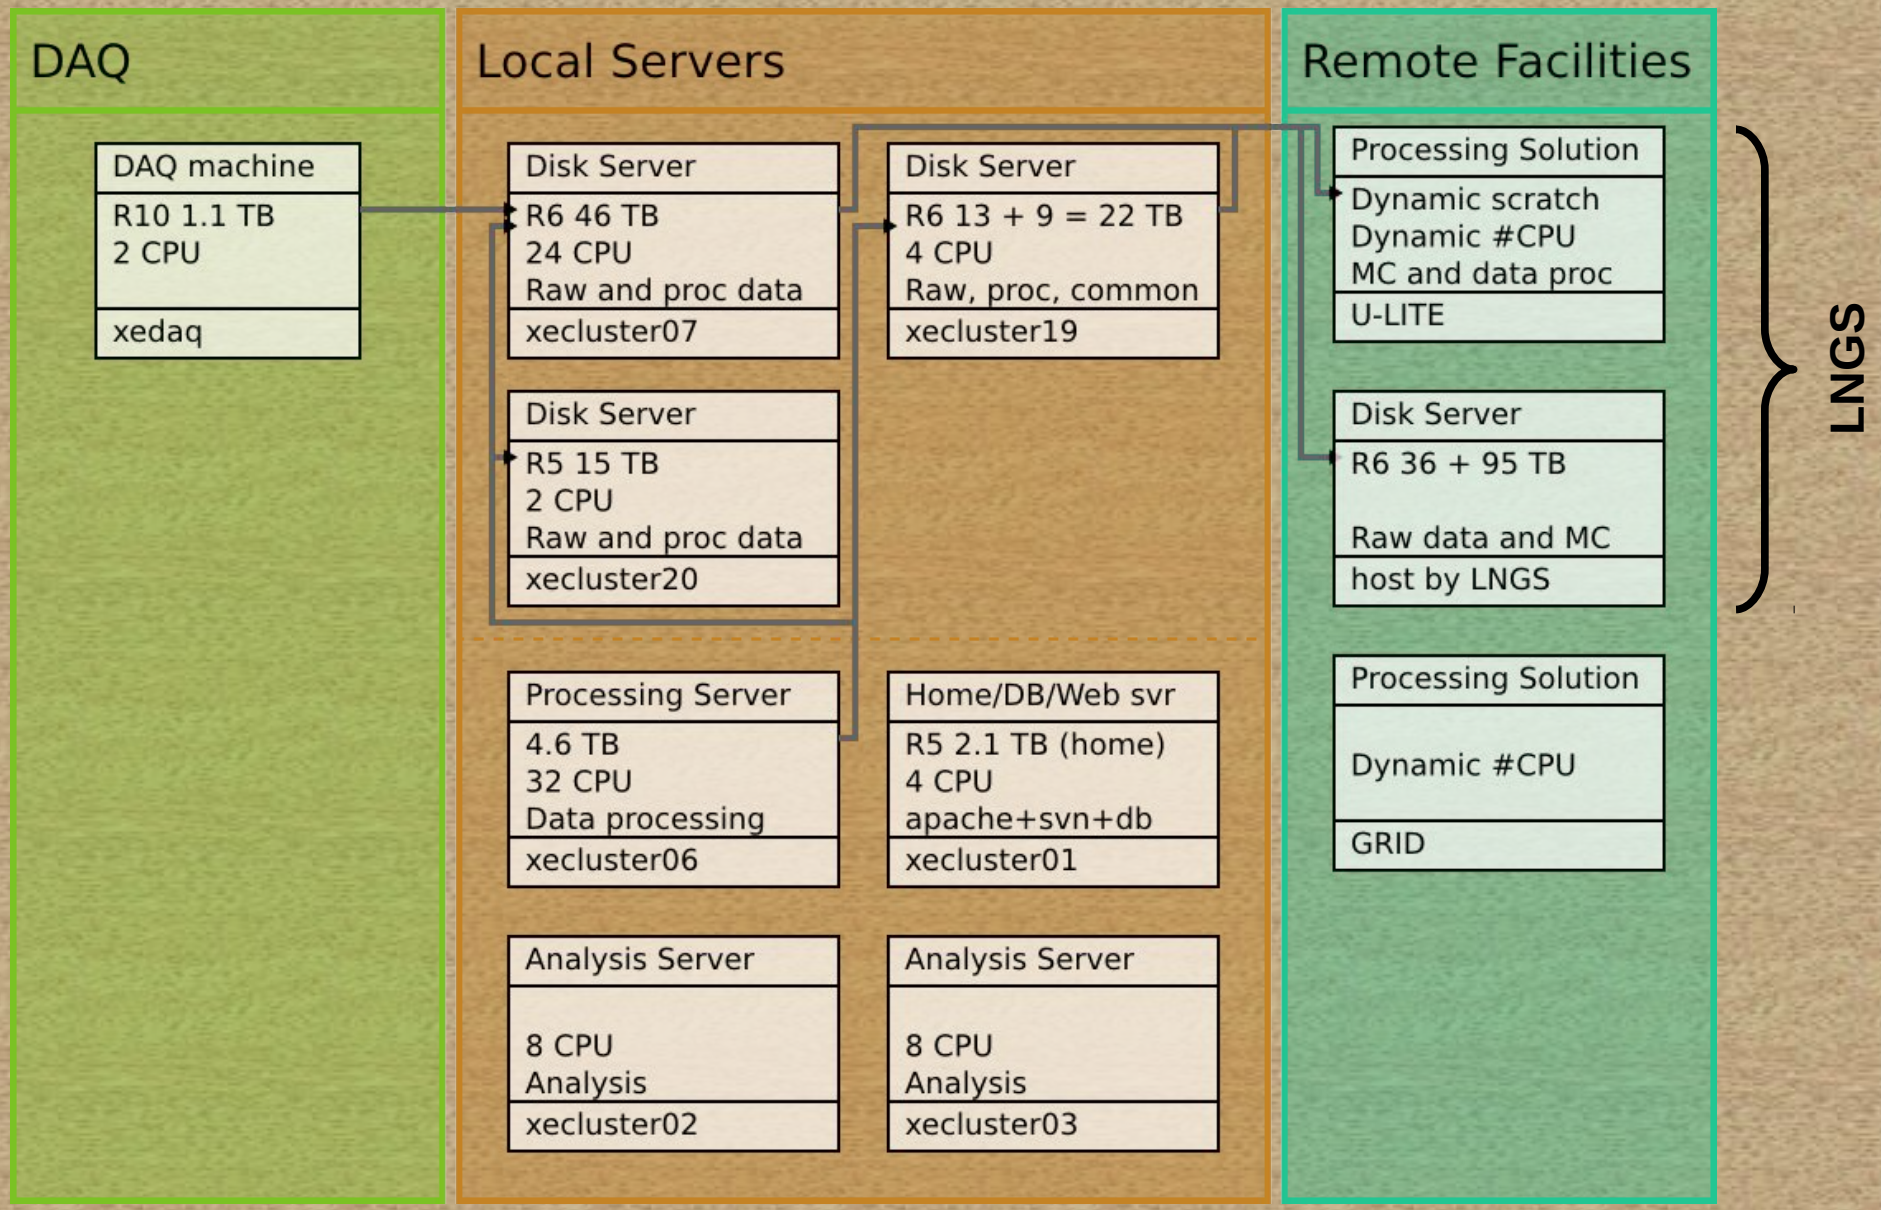

Scaling from XENON100 to XENON1T

Scaling factors are well known:

- **DAQ** : higher data flow (technology up to 300 MB/s)

- **Storage** : expected from few to several hundreds of TB

- **Processing** : expected processing resource for the single event as for XENON100, but we have to scale to much more events! Few hundreds of thousands of CPU-hours

Policy for the XENON1T data handling

General policy to build the XENON1T computing infrastructure:

- Build a **local facility** able to handle (store + process + analyze) the data flow during **limited times** (i.e., we cannot stop data taking because of computing issues!)

- Local facility must be able to handle **"Dark Matter search" data completely alone** (i.e., our DM data are blinded, so they must stay "at home")

- Extensive use of **extra** (internal or external) facilities for:
    - Monte Carlo simulation production
    - Calibration data storage
    - Calibration data processing
    - Data re-processing (boost of resources!)
    - Storage on tapes for backup

XENON1T expected resources

Baseline for numbers are well established:

- **Data flow:**
  - **DM search** : 10 MB/s
  - **Calibration** : Min 50 MB/s, Max limited by DAQ (300 MB/s)

- **Storage** : expected from few to several hundreds of TB
  - **DM search** : 50 TB / year
  - **Calibration** : 500 TB/ year

- **Processing** : Few hundreds of thousands of **average** CPU-hours
  - **DM search** : 20 kCPUh / year
  - **Calibration** : 600 kCPUh / year
  - **Monte Carlo simulations** : 80 kCPUh / year
  - **Analysis** : 44 kCPUh / year
  - **Total** : 0.7 MCPUh / year (baseline : 85 cores, but a factor x10 is required at least for **burst processing**)

Conclusion

XENON1T will start taking data in **2015**

XENON1T **computing facilities** must be ready and well tested before that date

General policy is:

- **Local facility** to handle data flow during limited times (network interruption, remote facility issues, …)

- **Remote facilities** extensively used for most of heavy work, but the Collaboration has not yet made a final decision (several alternatives are under study; using the LNGS facilities is one of the alternatives)

- **GRID** is a good solution for MC simulations (but not for raw data processing, unless of a reasonable bandwidth)