

# UniGE Tier 3, an update

Aug 19<sup>th</sup>, 2014

Szymon Gadomski

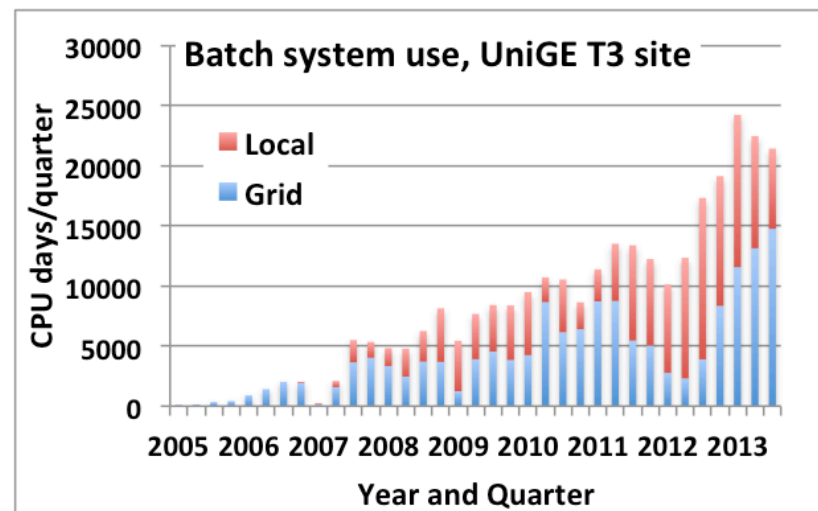
DPNC, Université de Genève

1. Hardware numbers and usage
2. Recent developments
3. Near future plans

# The Geneva T3, August 2014



- **784 CPU cores (5990 HS06)**
  - 656 in batch, 96 login, 32 Windows
- **614 TB net**
  - 474 in a grid Storage Element (DPM)
  - 140 NFS
- **10 Gb/s direct to CERN IT and to the SWITCH network**
- **used mainly for ATLAS, but not only**
- **ATLAS grid jobs using spare CPU**
- **gradual build up since 2005**



# Improvements of cooling in the machine room



# The 2014 upgrade is in progress

- **Replacing the oldest storage, including user home and software**
  - dedicated hardware for /home and /software
  - separate from large data volumes
- **New dedicated hardware to run all the virtualized critical services**
  - head node of the SE, NorduGrid front-end, site BDII and the batch server
- **Double CPU and memory in our latest computing nodes**
  - two machines will have 64 cores and 192 GB RAM

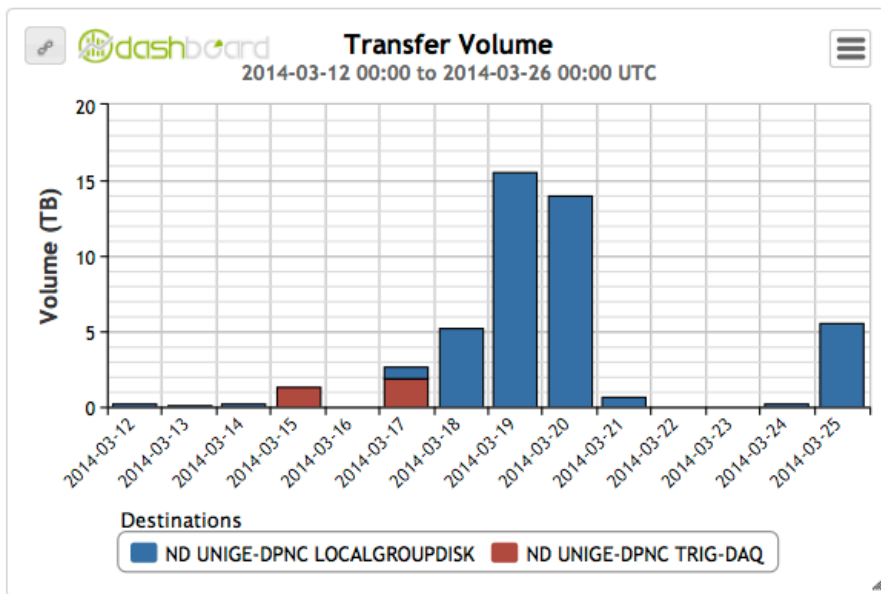

| <b>Machine type</b>         | <b>Purpose</b>                | <b>Units</b> |
|-----------------------------|-------------------------------|--------------|
| Sun Server X4-2             | user home and software        | 2            |
| IBM 3630 M4                 | data storage                  | 1            |
| IBM 3550 M4                 | virtualized critical services | 2            |
| CPU and memory for IBM 3755 |                               | 2            |

# Hardware repairs

- On 12 machines of a certain type (IBM x3630 M3, 2011-12) we had four cases of overheating hardware RAID
- The remaining machines opened and checked in June
  - two were repaired



# Data transfers to Geneva, March 2014




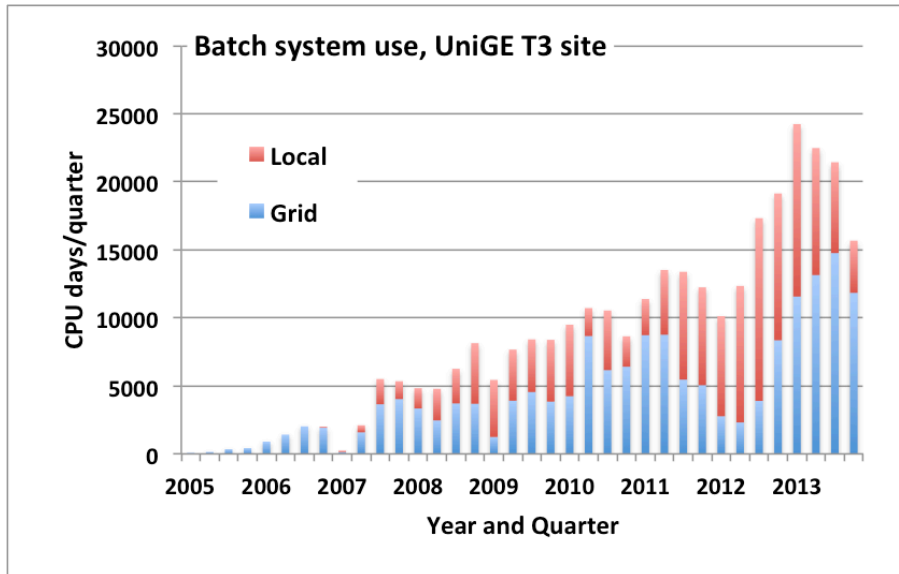
- Networks improve everywhere
- The ATLAS DDM software improves as well
- Data transfers to Geneva have reached 500 MB/s this Spring, even from Nordic T1
  - good news, data transfers of 15 TB/day have become possible
- Concern for the Uni network experts
  - we now have 10 Gbps to CERN and to the main switch of the Uni
  - the Uni has 10 Gbps to SWITCH and a firewall (~6 Gbps?)
- New protocol (SRM3) allowing protections set up for us by ATLAS DDM experts

# Access to data at other sites (FAX)

- The XrootD protocol allows data access from remote sites
- Our colleagues are working on “Federated ATLAS Data Access using XrootD”
  - automatic redirections can even find data for you
- Our site is already “federated”
- Redirection to CERN needed a fix, fixed now
- Performance tests
  - 1.5 to 2.5 MB/s from the US
  - 2.5 to 50 MB/s from CERN
  - redirection can take 15 sec/file, but the result is cached, much faster the 2<sup>nd</sup> time
- Can be useful, especially if the data are at CERN

<https://twiki.cern.ch/twiki/bin/view/AtlasComputing/UsingFAXforEndUsers>


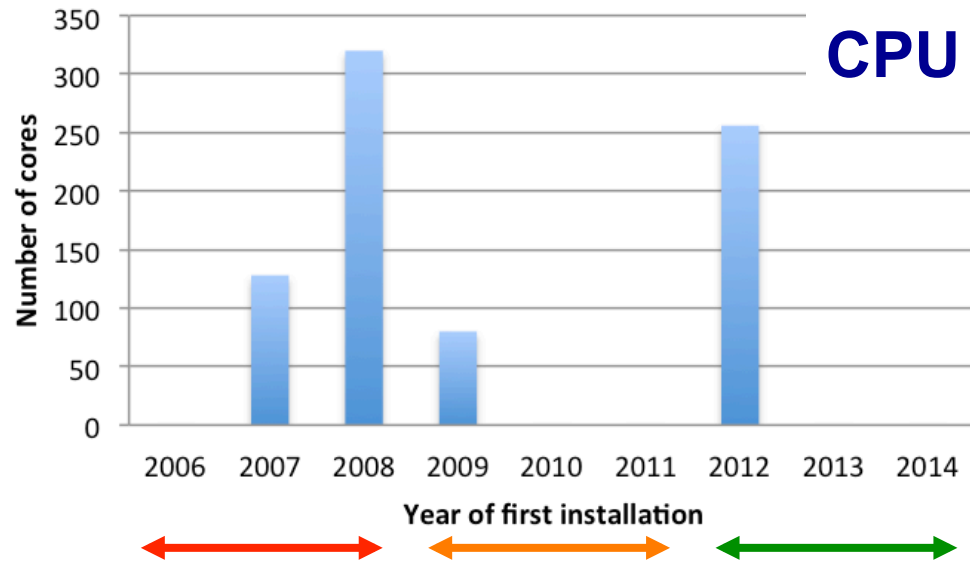
# Recent CPU usage, ATLAS plans



- Plan for next 2 years is due soon
- All numbers helping to quantify our needs are helpful



# Age distribution of our CPU and disk



- The disk servers from 2006 are being replaced now
- Old CPU servers, no longer reliable, do very well as batch worker nodes
- Need to be careful about ageing disk servers
- **Some renewal of both CPU and storage will be necessary in 2015-2016**

# Summary and plans

- **It is going well. Several little improvements**
  - hardware upgrade
  - hardware repairs
  - machine room
  - remote data access
- **Next**
  - finish the 2014 upgrade
  - plan hardware renewal for 2015-2016

# Other recent work

- **Migration to SLC6**
  - done for all batch machines
  - three login machines still run SLC5, plan to keep as long as necessary
  - **make another alias, like the lxplus?**
- **Upgrade of the Storage Element**
  - DPM 1.8.8 (latest)
  - head node and 15 disk servers
  - fix of xrootd crashes
- **NorduGrid upgrade**
  - new Computing Element in a virtual machine
  - issue with time drift in the VM
- **SE Cleanup**
  - 88 TB (18%) liberated
  - execution taking time and getting stuck, but finished yesterday

