



**UNIVERSITÉ
DE GENÈVE**

ATLAS Tier-3 cluster @ UniGe

Luis March and Yann Meunier (Université de Genève)

**CHIPP + CSCS GRID: Face To Face meeting
CERN, September 1st 2016**

Description of ATLAS Tier-3 cluster at UniGe

The ATLAS Tier-3 cluster at UniGe-DPNC is physically located at [Uni Dufour](#) (~ 500 m away from UniGe-DPNC building)

Grid services (NordUGrid):
ARC-CE, BDII, proxy, DPM SE

Batch system (62+12 nodes):
Worker Nodes ~ 656+96 cores
Memory/process = 2 – 10 GB
(option defined by user)

Storage system (DPM):
ATLAS pool
Reserved = 450.0 TB
File Servers = 15

New space for DAMPE: 17 TB
Inside DPM Grid Storage



ATLAS Space Tokens	Capacity (TB)	Used (TB)	Free (TB)
ATLASGROUPDISK	25.0	9.85	15.15 (~ 60.6%)
ATLASLOCALGROUPDISK	420.0	380.14	39.86 (~ 9.5%)
ATLASSCRATCHDISK	5.0	0.22	4.78 (~ 95.7%)

Fine for
now

Description of extra Tier-3 cluster at UniGe

Some extra Tier-3 cluster resources at UniGe-DPNC, for different experiments (not only ATLAS), which are also physically located at [Uni Dufour](#)

User Interfaces (login machines for users):

SLC6 (3 nodes) = 48 cores

SLC5 (3 nodes) = 48 cores → They will be used for other services

In addition to DPM SE, we have NFS disk servers for local storage:

<code>/atlas/users</code>	→ Intended for software development (3 TB)	
<code>/atlas/software</code>	→ Intended for common ATLAS software (local users) (2 TB)	
<code>/cvmfs/*.cern.ch</code>	→ Official software tools for (some) experiments (mounted)	
<code>/atlas/data</code>	→ Data storage for UniGe ATLAS users	108.0 TB
<code>/neutrino/data</code>	→ Data storage for UniGe neutrino users	82.0 TB
<code>/ams/data</code>	→ Data storage for UniGe AMS users	103.0 TB
<code>/icecube/data</code>	→ Data storage for UniGe IceCube users	2.0 TB
<code>/dampe/data</code>	→ Data storage for UniGe DAMPE users	58.6 TB

Total NFS disk space for local storage = ~5 TB + 353.6 TB = ~ 358.6 TB

Operations

Grid services (Nordugrid):

ARC-CE → “nordugrid-arc-ce-5.0.5”

DPM SE → “glite-yaim-dpm 4.2.20-1” (we should upgrade it with Puppet)

GGUS ticket/s:

Ticket-ID 117900 → About ATLAS storage (monthly) consistency checks
Data management: **Status = closed**

```
srmls -l srm://grid05.unige.ch:8446/srm/managerv2?SFN=/dpm/unige.ch/home/atlas/atlaslocalgroupdisk/dumps/
```

```
srmls -l srm://grid05.unige.ch:8446/srm/managerv2?SFN=/dpm/unige.ch/home/atlas/atlasscratchdisk/dumps/
```

```
srmls -l srm://grid05.unige.ch:8446/srm/managerv2?SFN=/dpm/unige.ch/home/atlas/atlasgroupdisk/trig-daq/dumps/
```

Ticket-ID 120979 → About ATLAS storage deletions: **Status = closed**

Deletion failures at UNIGE-DPNC_SCRATCHDISK, due to permissions


Once the permissions were changed, no failures observed yet

In general, running smoothly:

ATLAS Production jobs → UNIGE-DPNC came back to production on July 23rd

UniGe local users → Increased activity (job submission) for last months

Accounting along time: UniGe cluster



Stats from October 2015 (4th quarter 2015) to June 2016 (2nd quarter 2016)



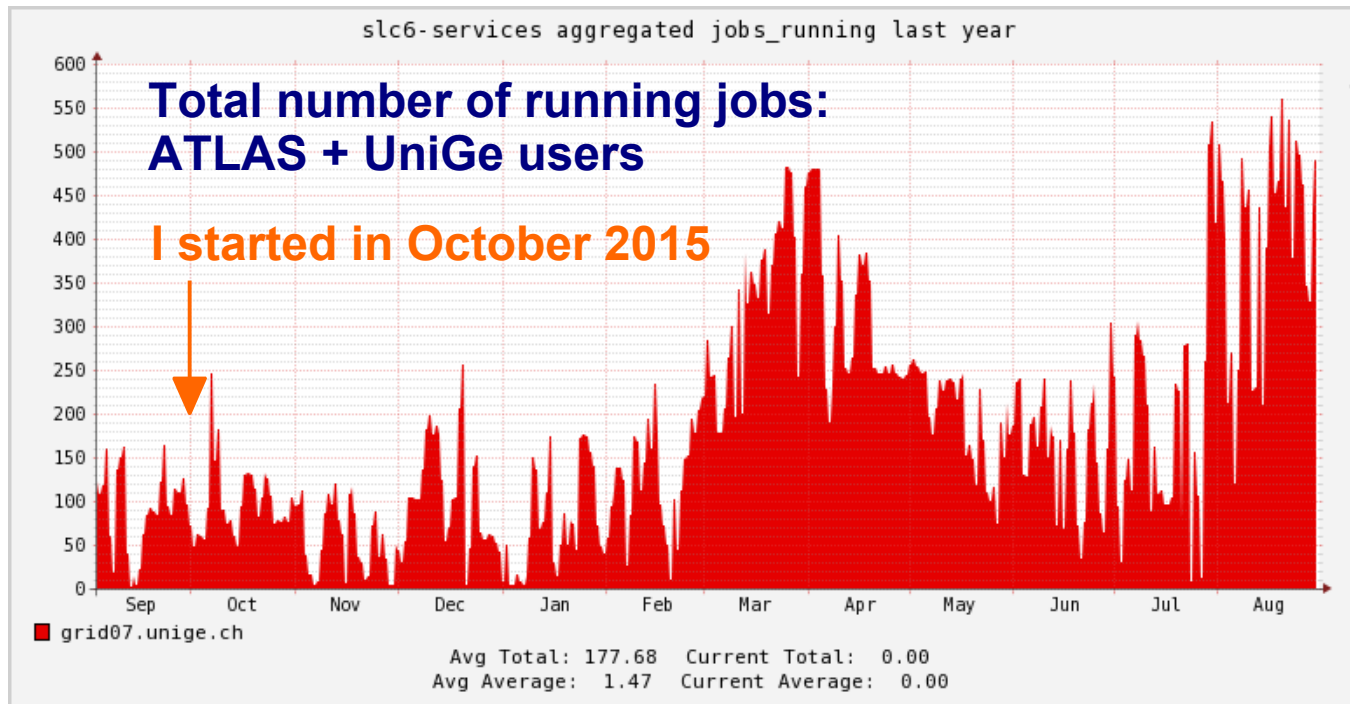
1st & 2nd quarter 2016: Highest local user activity (main users: ATLAS and DAMPE)

Lower ATLAS production activity since June 2015 (checked: related to memory)

ATLAS production activity re-started since ~ July 23rd

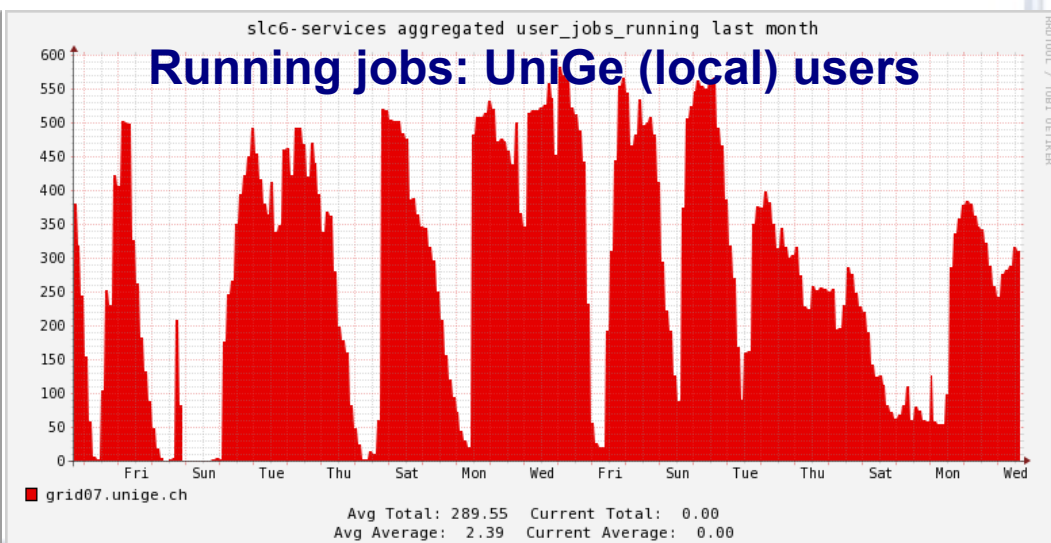
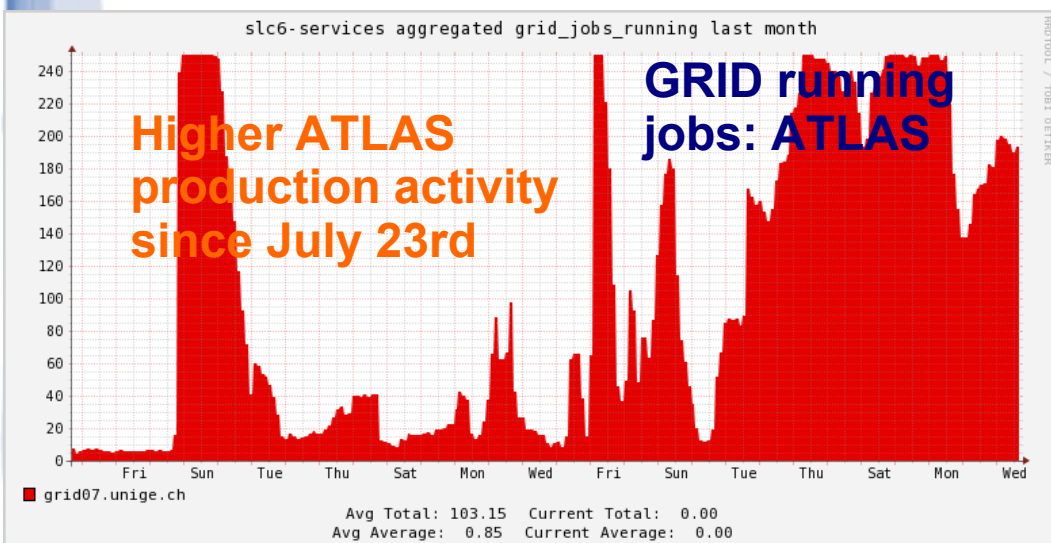

Accounting: Monitoring at UniGe (1)

Stats:
Last year




Accounting: Monitoring at UniGe (2)

Stats:
Last month



Monthly accounting: UniGe cluster


WallTime days - Batch system - Monthly 2016




Stats until August 31st at ~ 14 h

Two remarks: ATLAS group (local) and ATLAS production (grid) increased their usage

July 2016 - WallTime days per DPNC group




August 2016 - WallTime days per DPNC group



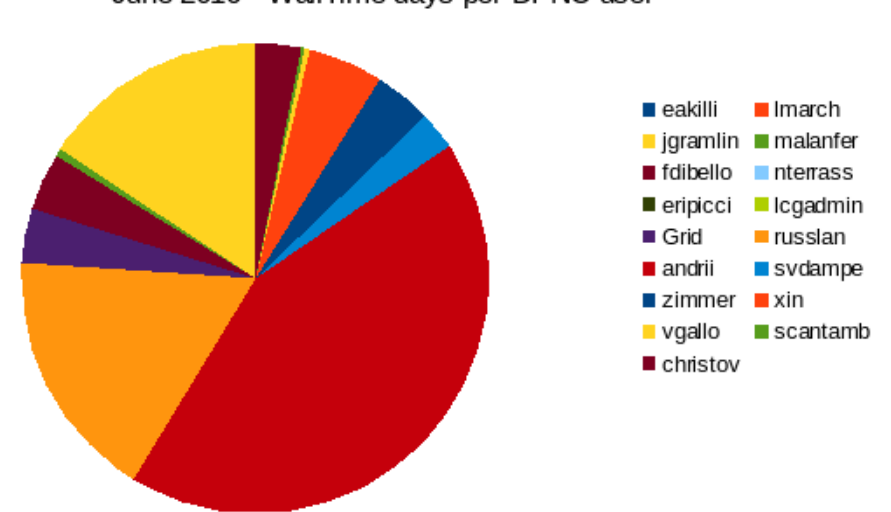
DAMPE group: Decreased activity for July and August

Monthly accounting: Local users


May 2016 - WallTime days per DPNC user



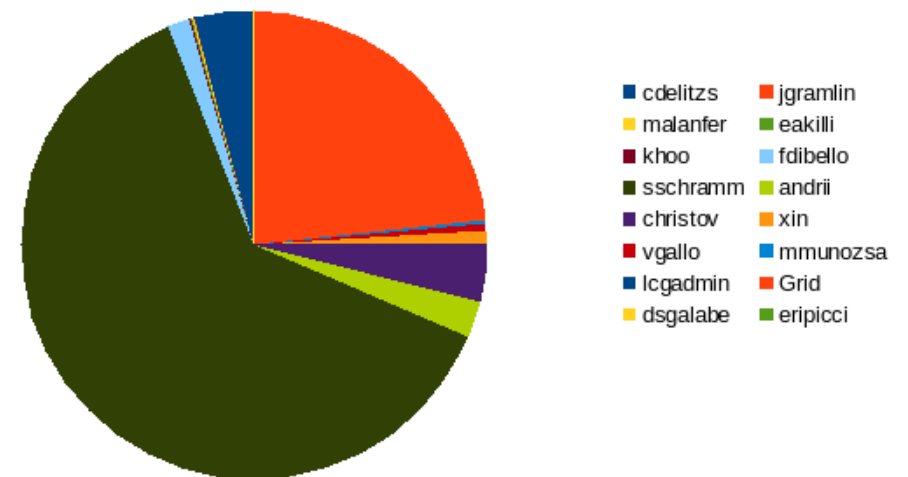
June 2016 - WallTime days per DPNC user



July 2016 - WallTime days per DPNC user




August 2016 - WallTime days per DPNC user




Monthly accounting: Local groups


May 2016 - WallTime days per DPNC group




June 2016 - WallTime days per DPNC group



July 2016 - WallTime days per DPNC group



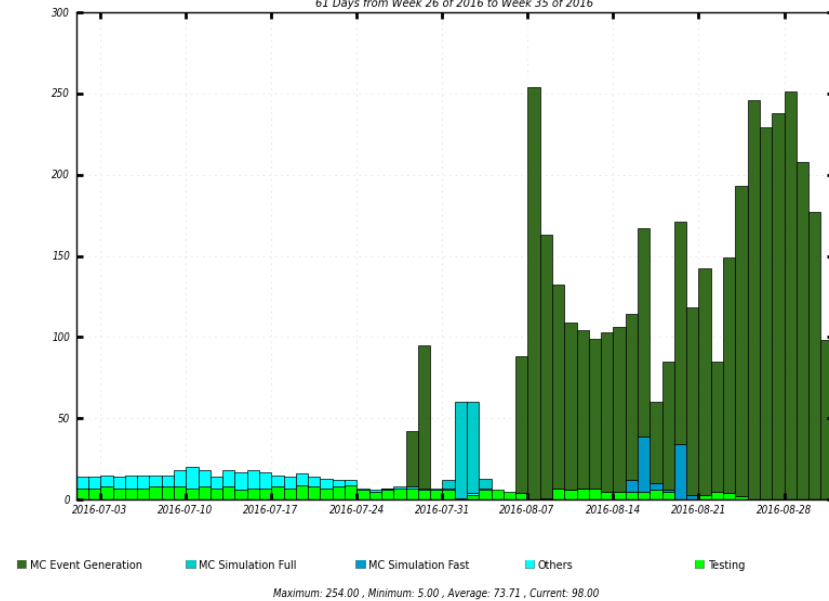
August 2016 - WallTime days per DPNC group



Accounting: ATLAS production (1)




Running jobs
61 Days from Week 26 of 2016 to Week 35 of 2016

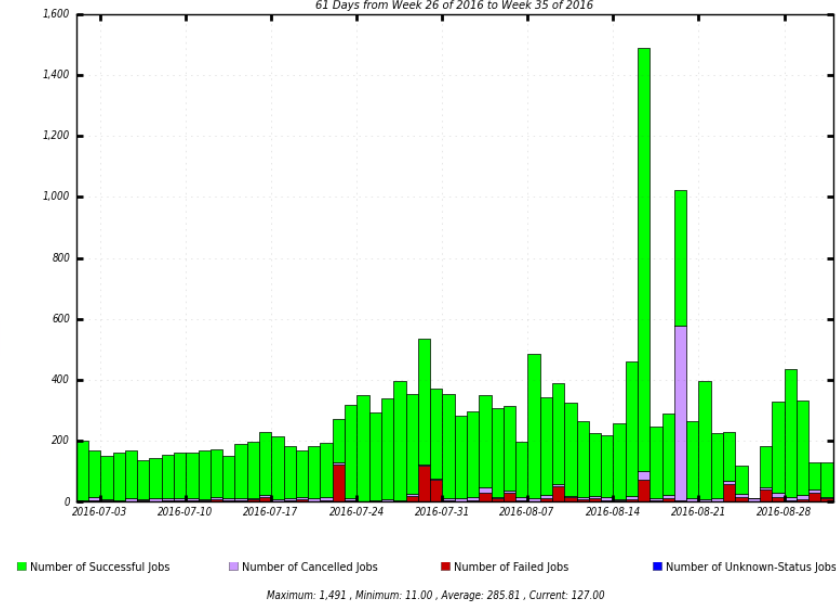


Number of Successful and Failed Jobs (Pie Graph) (Sum: 17,720)

Percentage of Successful Jobs - 89.11%




Number of Successful and Failed Jobs (Time Stacked Bar Graph)
61 Days from Week 26 of 2016 to Week 35 of 2016



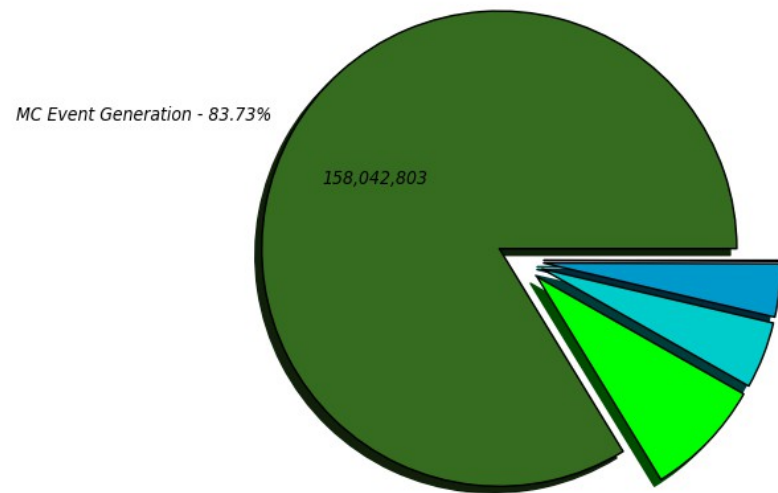
Accounting: ATLAS production (2)



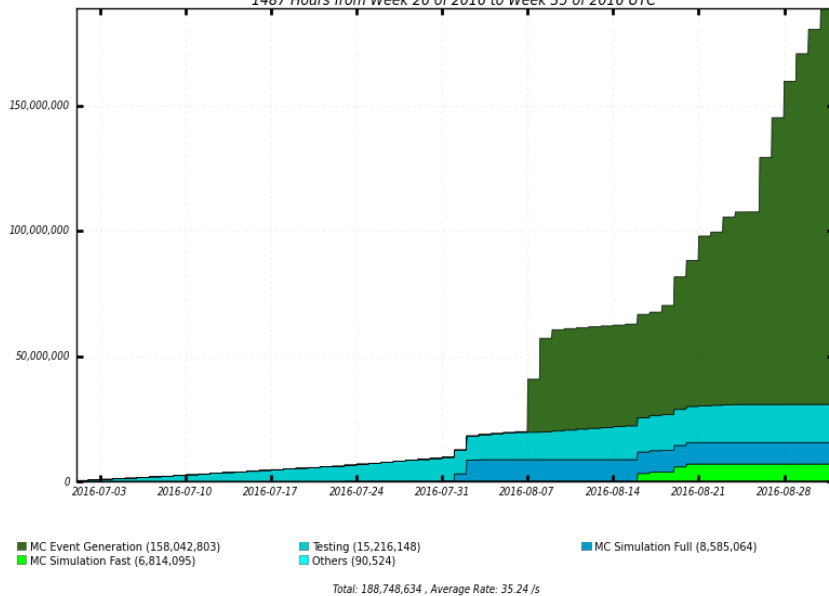
CPU consumption All Jobs in seconds
61 Days from Week 26 of 2016 to Week 35 of 2016



CPU consumption All Jobs in seconds (Sum: 188,748,634)



CPU consumption All Jobs in seconds
1487 Hours from Week 26 of 2016 to Week 35 of 2016 UTC



MC Event Generation - 83.73% (158,042,803)
 Testing - 8.06% (15,216,148)
 MC Simulation Full - 4.55% (8,585,064)
 MC Simulation Fast - 3.61% (6,814,095)
 Others - 0.05% (90,524)

Outlook

Operations:

CPU/cores

→ We have up to ~136 cores more to be added at the cluster

Operating system

→ Currently SLC6, but moving to CentOS at some point

Batch system

→ We would like to move to **SLURM** (currently Torque/PBS)

Testbed

→ We could use some of these CPUs to be tested with SLURM

GPU machines

→ **Funding request submitted to UniGe: Approved**
Finally, GPUs would be added into **Baobab HPC cluster**

ATLAS Production

→ ATLAS production re-started at UniGe since July 23rd
We should cross-check/review our accounting + Multi-Core

Storage:

Disk servers

→ We are going to add 1 (70 TB) disk server (maybe 2) to DPM
Created a DAMPE pool of ~ 20 TB on DPM

11 File Disk Servers with SCL5 (upgrade to SLC6 only if necessary)

DPM SE

→ We would like to move to **Puppet** (currently YAIM)

Testbed

→ We would like to make a small testbed:

1 service machine: Puppet

1 Head Node: DPM (newer version than current one)

1 File Disk Server: Data to be managed by DPM

Network:

Upgrade to 10 Gb/s → **Funding request submitted to UniGe: Approved**

Data transfers from/to NFS disk servers: Performance tests scheduled soon

Back-up Slides

Disk server intervention

atlasfs15, atlasfs18 & atlasfs20: RAID controller card damaged

RAID cards replaced and put them back into service mode

