Optimise resource usage and operation of lightweight Grid sites

S. Jézequel (LAPP)

ATLAS Jamboree – March 2019

Introduction

- Updated version from ADC weekly presentation (Jan 2019)
- Overview : Recommendation (ICB-2018) to redirect funding from storage to CPUs for **lightweight Grid site (2018 limit : 460 TB, 2019 (+15%): 520 TB)**
  - Already implemented on volontary basis on 5 sites also had small amount of CPUs (<200 cores) (slide 3)
    - 3 T2s with pledge CPU resources (HEP-UIBK, RO-14,RO-16)
    - Support is now close to zero (and no worry about future storage migrations)
- Question : Can this setup be extended to larger Grid sites (200-1000 cores)
  - Grid site : Host client commands to copy/read files on remote Grid SE
- Volunteer computing (Boinc) is another option
  - Already 10k cores managed
  - Reinforced by recent implementation of correct report (at least for sites without pledge resources)

Challenges

- Identified during the SE decomissioning in UKI-SOUTHGRID-BHAM (2018) (1k cores)
  - Network saturation : Direct remote access requires stable connectivity
  - Too much load on remote SE : Blow up with network instability/saturation

    → Should adapt the global IO to effective connectivity

  → Would have required a tool to limit the nb of jobs with large IO : not existing yet

  → Proposition for 2019 : Run only low IO jobs

- Questions :

  1) Will it reduce significantly the Grid CPU resources for large IO activity ?

  → Would not accept reduction by more than few % (total Grid cores : 300k)

  2) Will it threaten the available resources to run Grid anlysis jobs

  3) Will low IO jobs keep the resources busy ?

Isolated T2/T3 sites

Site	DATA DISK (TB)	SCRATCH DISK (TB)
HK-LCG2	420	22
ZA-CHPC	-	-
BEIJING-LCG2	310	60
EELA-UTFSM	360	30

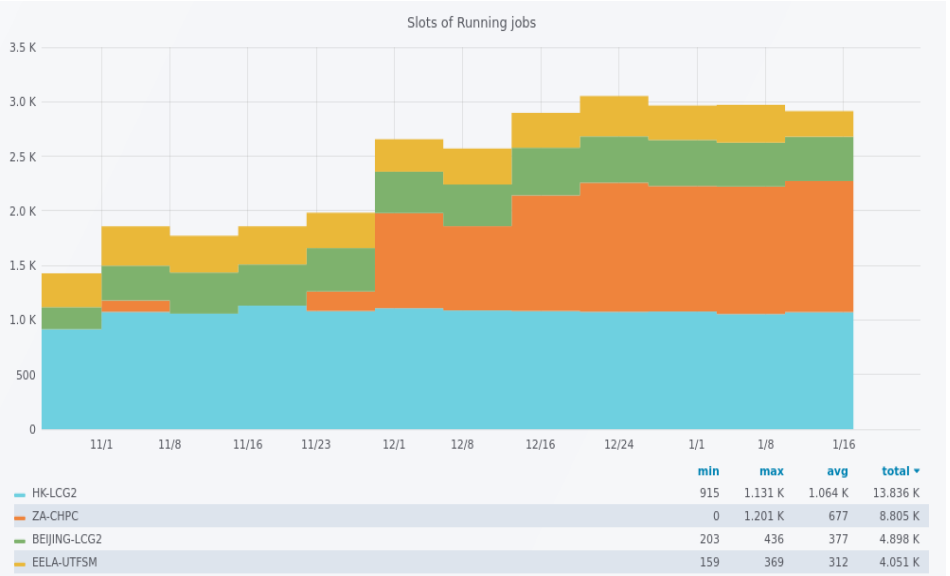

3000 cores → 1 % of Grid capacity

Storage capacity : ~1.2 PB

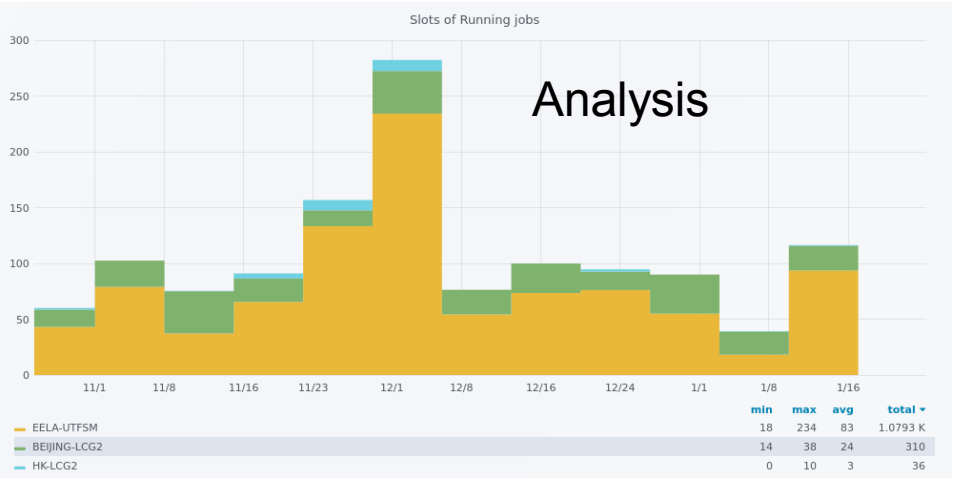

- ZA-CHPC (South Africa) (associated to INFN-T1) experience demonstrates no need for local DATADISK if low IO activity (some inefficiency due to very limited network bandwidth)

Sites within T2 federation

Site	DATA DISK (TB)	SCRATCH DISK (TB)
DURHAM	192	-
SHEF	450	11
SUSX	10	-
Brunel	33	-
BHAM	-	-
CAM	190	8

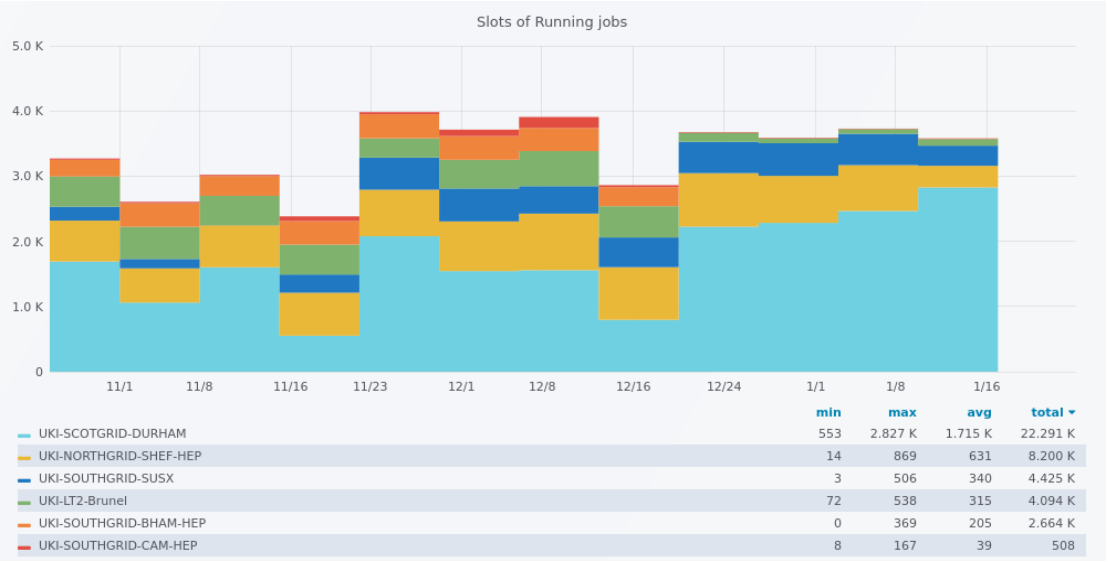

Site	DATA DISK (TB)	SCRATCH DISK (TB)
WEIZMANN*	192	40
NCG-INGRID-PT*	177	11
PSNC	341	22
RO-02-NIPNE*	319	16

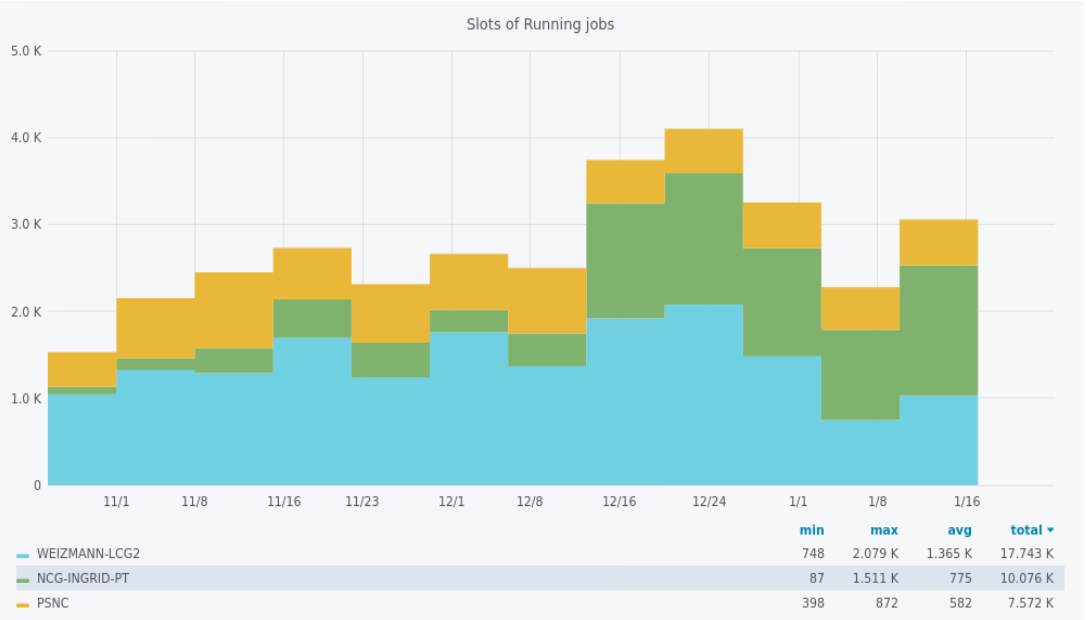

8000 cores → 2.5 % of Grid capacity

Storage capacity : ~2 PB

* : Site broken for analysis > 2 months in 2018

Sites within T2 federation (2)

Site	DATA DISK (TB)	SCRATCH DISK (TB)
COSENZA	330	55
GENOVA	1.1	4.4
LECCE	1.00	5
ROMA3	3.3	7.7

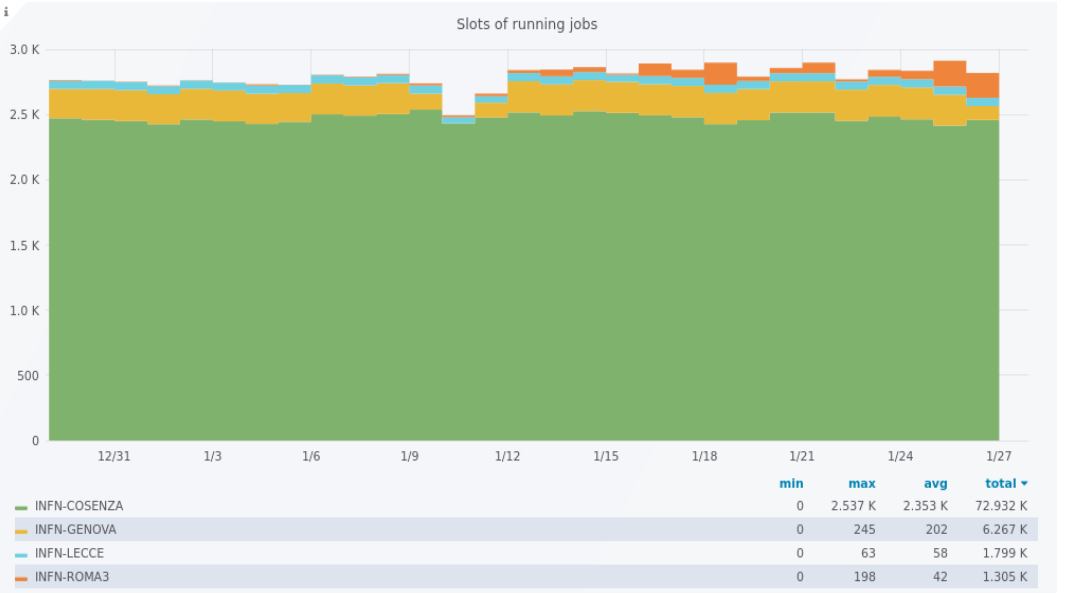

Site	DATA DISK (TB)	SCRATCH DISK (TB)
NERSC	192	-
UTA-SWT2	325	-
OU_OSCER_ATLAS	400	50

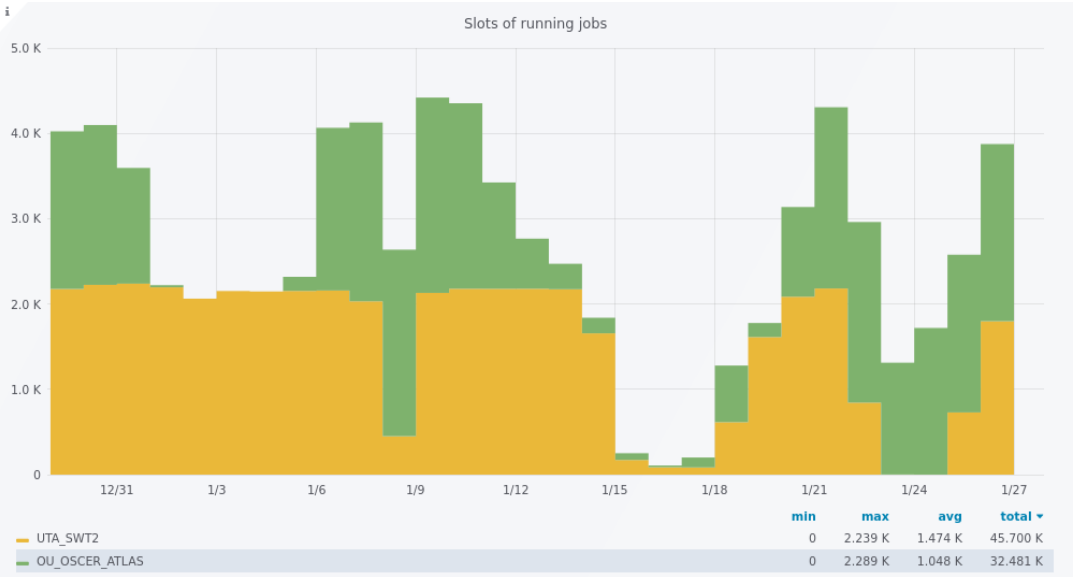

6000 cores → 2 % of Grid capacity

Storage capacity : ~1.4 PB

Loose CPU capacity for analysis ?

No local Grid storage → Stronger request on remote connectivity (bandwith and reliability)

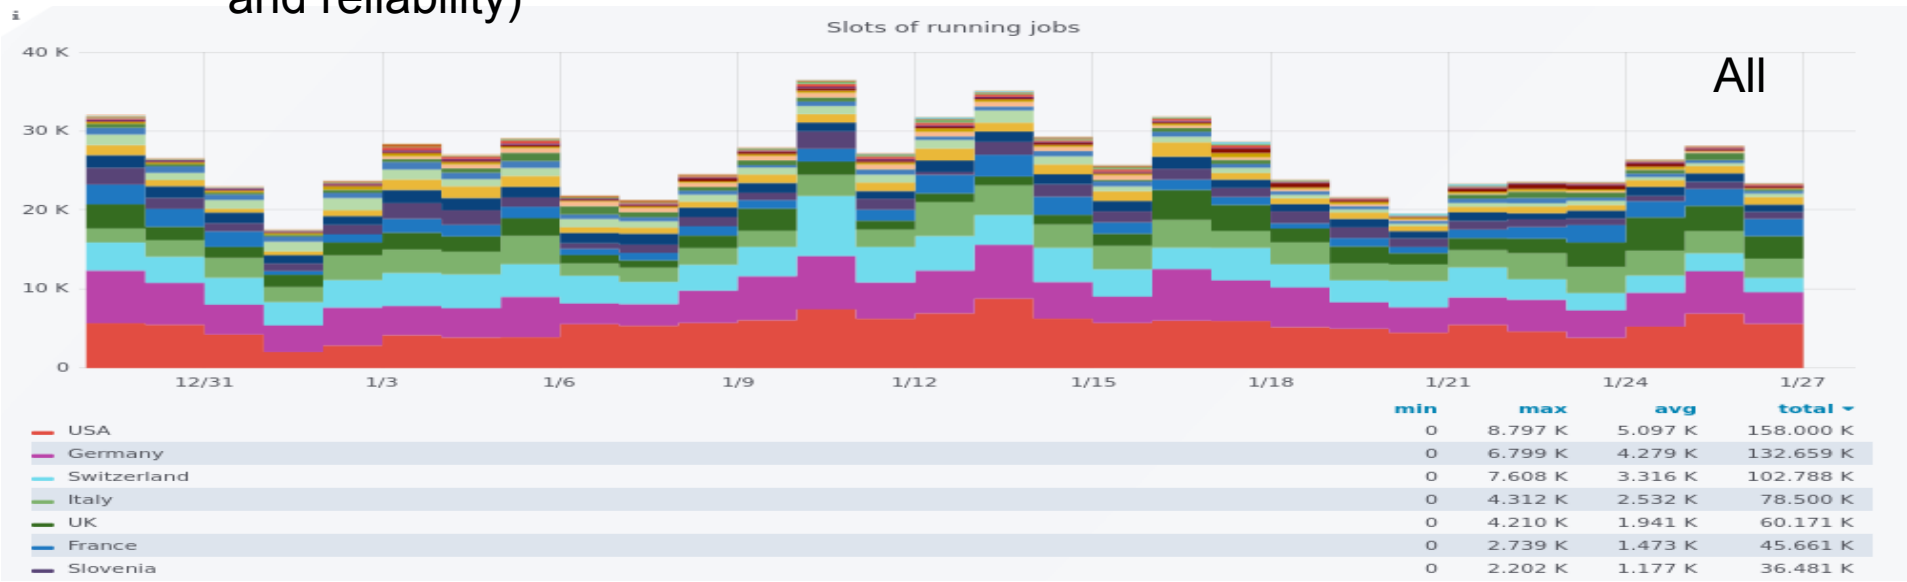

	min	max	avg	total
USA	0	8.797 K	5.097 K	158.000 K
Germany	0	6.799 K	4.279 K	132.659 K
Switzerland	0	7.608 K	3.316 K	102.788 K
Italy	0	4.312 K	2.532 K	78.500 K
UK	0	4.210 K	1.941 K	60.171 K
France	0	2.739 K	1.473 K	45.661 K
Slovenia	0	2.202 K	1.177 K	36.481 K

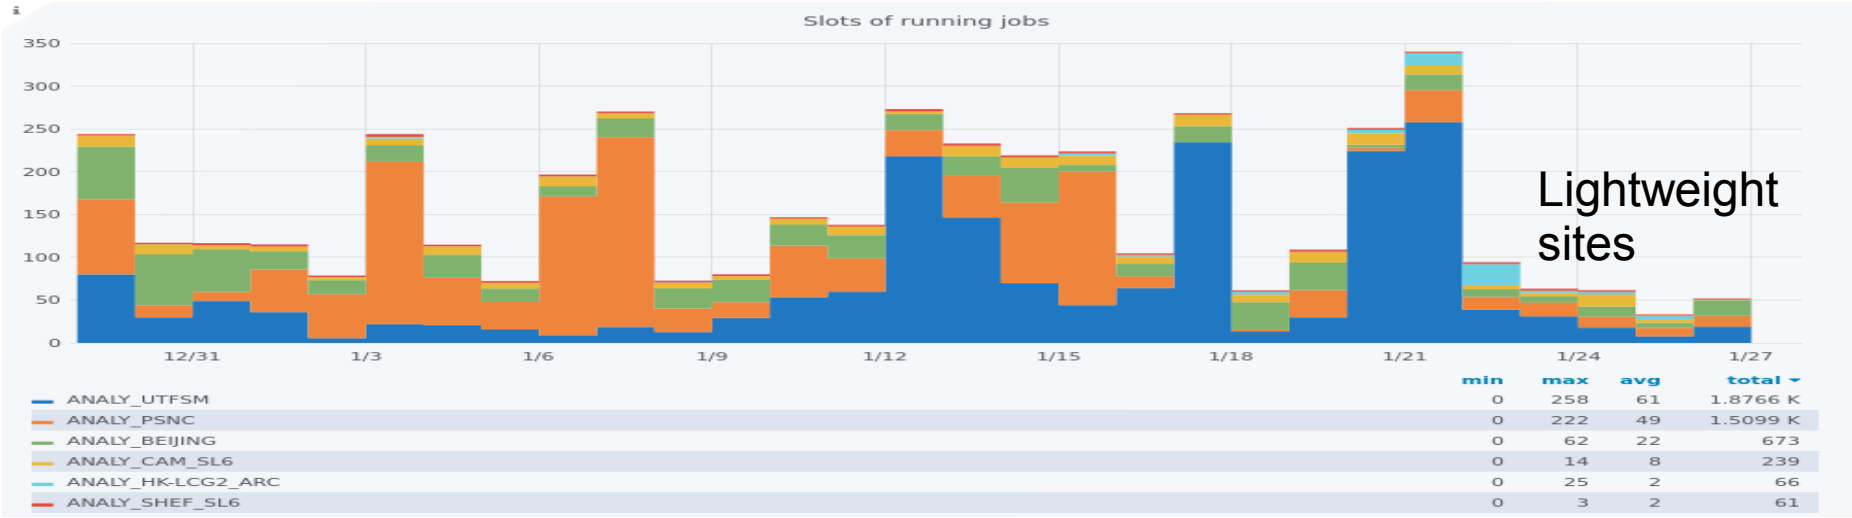

	min	max	avg	total
ANALY_UTFSM	0	258	61	1.8766 K
ANALY_PSNC	0	222	49	1.5099 K
ANALY_BEIJING	0	62	22	673
ANALY_CAM_SL6	0	14	8	239
ANALY_HK-LCG2_ARC	0	25	2	66
ANALY_SHEF_SL6	0	3	2	61

→ Would loose O (1%) of analysis capacity

Enough activity with simul/evgen ?

- Simul and evgen processed by
  - HPC
  - Cloud computing
  - Volunteer computing
- Processing request so huge that all sites can be kept busy with simulation

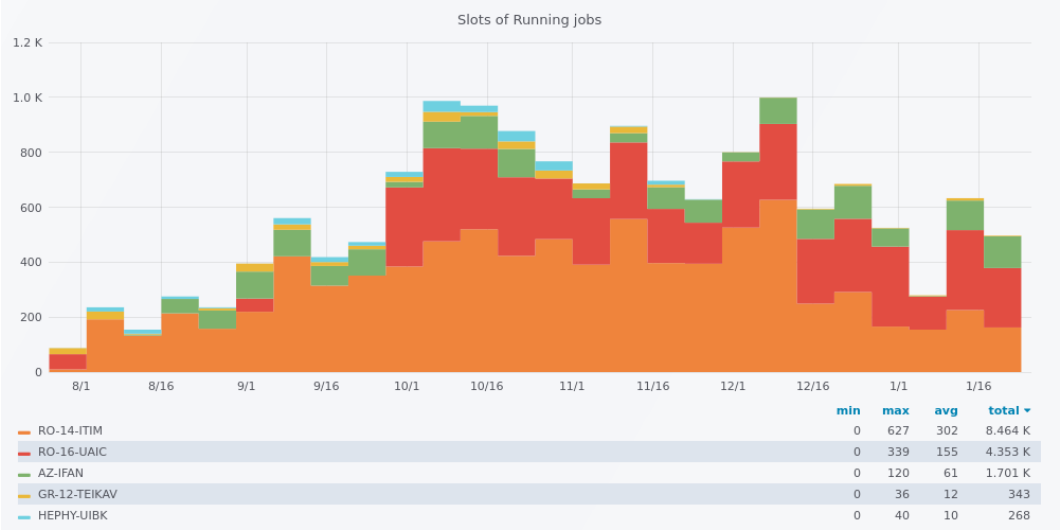

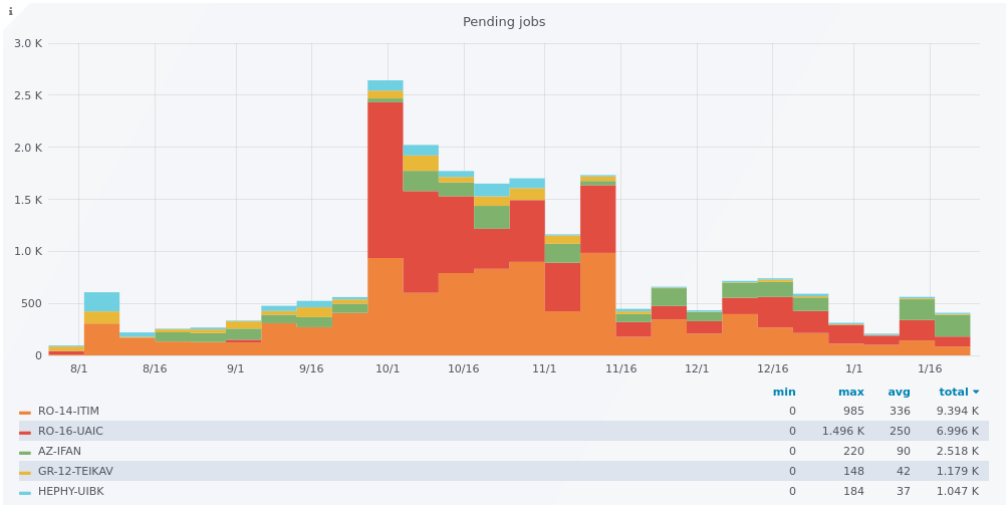

→ Current diskless sites kept busy in last 6 months with simul only

600 cores → 0.2 % of Grid capacity

Downtime management

- Switcher updated recently to handle downtime for diskless site
  - Destination SE downtime → stop PQ at the diskless site

    (J. Caballero this week)

Conclusion

- Diskless site : Should optimise the
- Focus activity to low IO production jobs
  - Lightweight Grid site represents ~15k (5%) of Grid CPU capacity

    → Transforming to diskless would not affect significantly high IO jobs processing
  - Low IO → Almost not affected by network occupancy → Lower operational burden
  - No more analysis queue (currently~2%)→No interest to keep DATADISK
  - **Migrating to ATLAS@Home/BOINC would be even simpler solution (monitoring report issue solved recently)**
- **Up to the cloud or country coordination to drive reorganisation**
  - **Convince FA that it is optimal usage of funding**

Recommendation to diskless sites

- **Up to the cloud or country coordination to organise migration in liaison with ADC**
- Actions from site admins (minimising changes requested to sites) :
  - Close local analysis queue and reallocate slots to Production
  - Close Grid storage
  - If site wants to keep storage for some time, reallocate all Grid storage to SCRATCHDISK and contribute to storage of second replicas of user outputs (lifetime of 2 weeks → storage decomissioning could be fast)

Beyond 2019

- Since more sites are expected to become diskless, necessary to setup ATLAS site configuration to enable reasonable level of high I/O jobs according to network connectivity
  - Stay diskless : All informations are available (network connectivity per experiment, data flow rate per job type) but requires implementation in the workflow
  - Use cache if necessary to smooth network usage (ARC-CE, xcache,...)
  - Usefull configuration for commercial cloud or HPC
- Other options
  - Federate storage :Reduction of manpower but responsibility of site admins to make it transparent to application (network, downtime)
  - Dynafed : One entry point for the experiment but sites keep their own management

DOMA R&D : Place to discuss/test/evaluate different options