Usage of CC-IN2P3

Frédéric Derue, LPNHE Paris

Calcul ATLAS France (CAF) meeting
CC-IN2P3 Lyon, 1st April 2019

Usage of sps (1/3)

cctools view [*link*]

- **sps (under gpfs)**
  - 291 TB allocated, 160 TB used (at previous CAF was 156 TB)
  - ‘daily’ recovery of GPFS disk from groups moving to new platform.
  - expected 360 TB by end of March

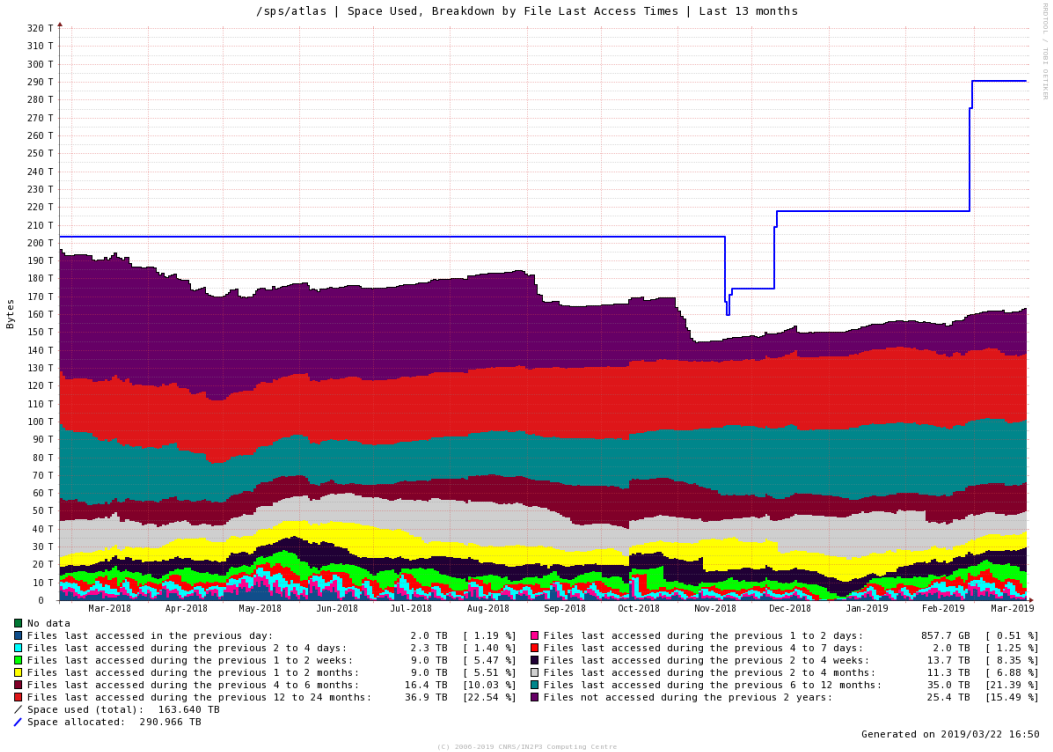

- Usage by « users »
  list of users by decreasing order : [*link*]

- **cleaning procedure**
  data not accessed since a year moved to ATLASLOCALGROUPTAPE
  ex : recently 31 TB not accessed since 01/01/2018 were moved
  technical procedure is described *here* - thanks to Manoulis (also on wiki?)

Usage of sps (2/3)

- **sps (under isilon)**
  isilon mountpoint with 5TB is ready (since >1 month)
  - thanks to Manoulis et al.
    /sps/atlastest the users just have to do (from a cca) :
    cd /sps/atlastest ; mkdir $USER
    In order to create a personal directory on top of this mountpoint

- Three ATLAS users voluntaries to run their usual analysis on sps gpfs and isilon : K. Al Khoury (LAL), M. Escalier (LAL), E. Sauvan (LAPP)

- **Tests of Konie / Emmanuel**
  → no feedback of Konie

  → from Emmanuel : no details on tests (some issues as tests were done on batch at the moment of some batch system perturbation)
  No differences seen in performance between jobs on sps (under gpfs) and sps (under isilon)

Usage of sps (3/3)

● **Test of Marc** : interactive jobs on cca010

program that read MxAOD produced by the team H→gamgam
sample of ZH with H→gamgam. Then it applies a selection dedicated to a tempory version for the analysis H(bb)H(gamgam).

READ=the location where is the MxAOD input file.
WRITE=location where I put the destination file (write only at the end of the program)
Hz (frequency of events treated per second)
time : the total duration of the program (in minutes =' and seconds='')

The duration of the program before the loop on the events itself is around 45 s.
It doesn't vary with the architecture of the destination.
I did the exercice in the ordering of the number in parenthesis (1=first, 2=second)

READ	gpfs		isilon
WRITE	gpfs	(1) ~16 Hz, 10'32"	isilon (2)~16 Hz, 10'58"

After the test, I do (3) = again the configuration "(1)" : it makes 10'10" instead of o10'32"<=>very stable<=>no influence in case there would be a variation of activity from other users of the computer.

**Outputs : /sps/atlastest/escalier
**Conclusion :There is no significative gain observed with my program that reads rather intensily the events, when using isilon, as compared to gpfs.

Usage of LOCALGROUPDISK

dashboard SRM view [link]

- **LOCALGROUPDISK**
  525 TB, among which 200 TB left (at previous CAF was 150 TB)

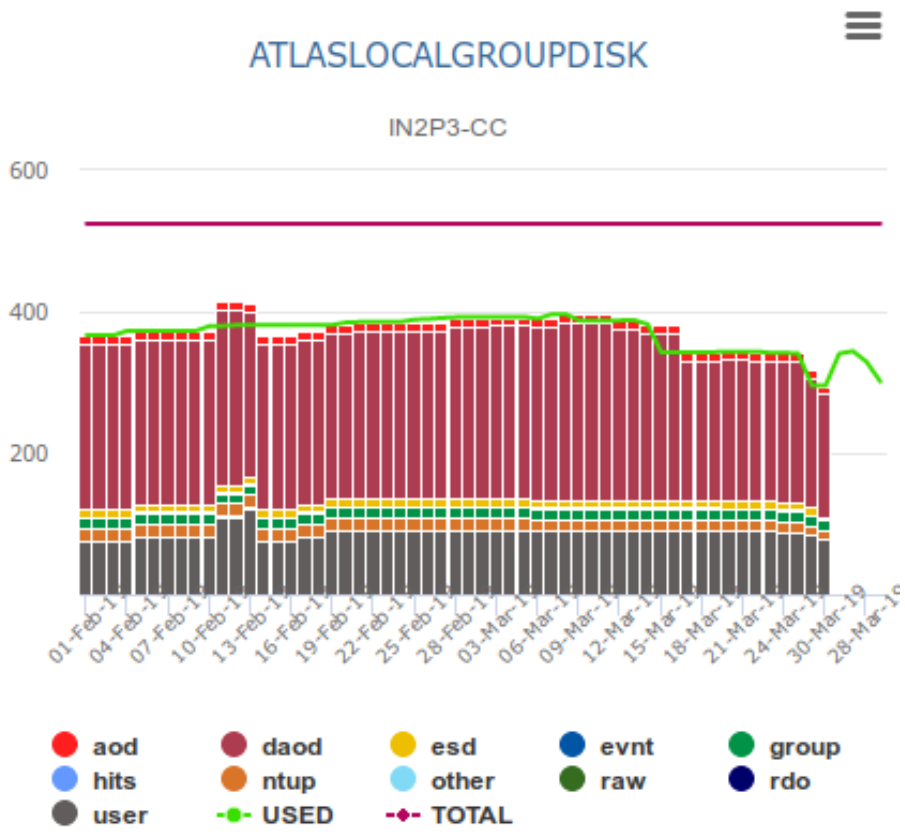

Usage of LOCALGROUPTAPE

dashboard SRM view [link]

- **LOCALGROUPTAPE :**
  for long tetrm storage only, panda queues have no acces to this RSE
  ~240 TB used
  recent request from J. Stark (LPSC)
  to store some analysis ntuples
  (some already on grid,
  some other not yet)
  → not so straightforward to do it
  (still ongoing)

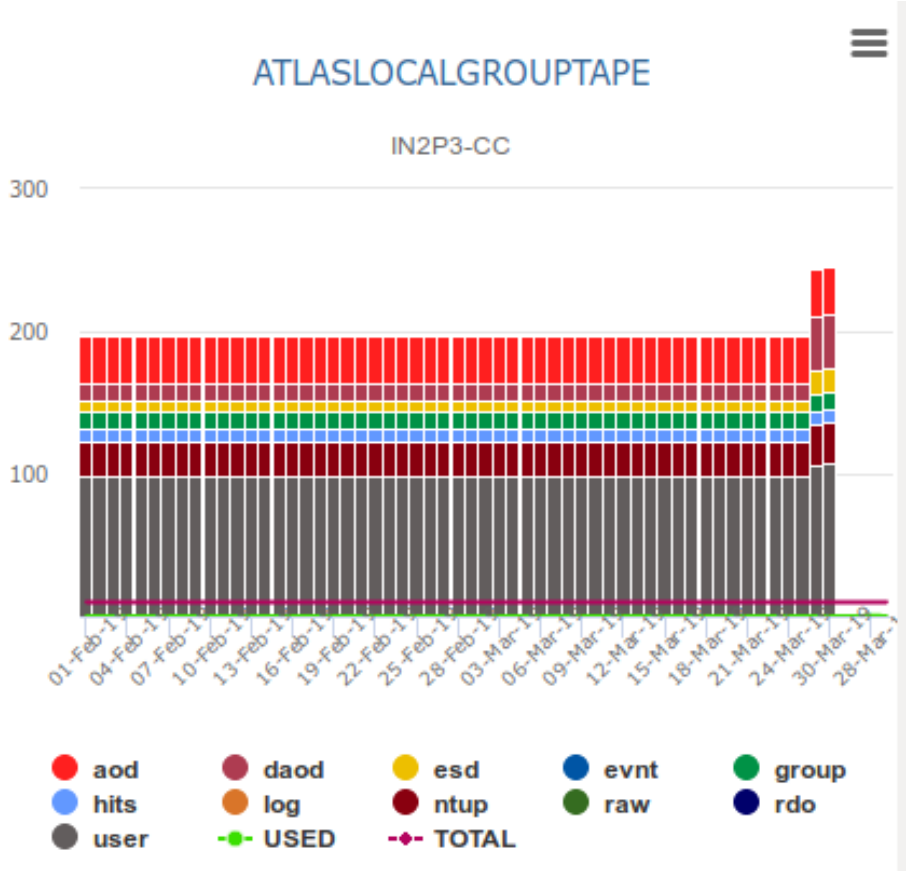

Usage of LOCALGROUPDISK-MW

dashboard SRM view [link]

- **LOCALGROUPDISK-MW** for SM (IRFU)
  75 TB used, on disks which are no more under warranty
  no recent feedback

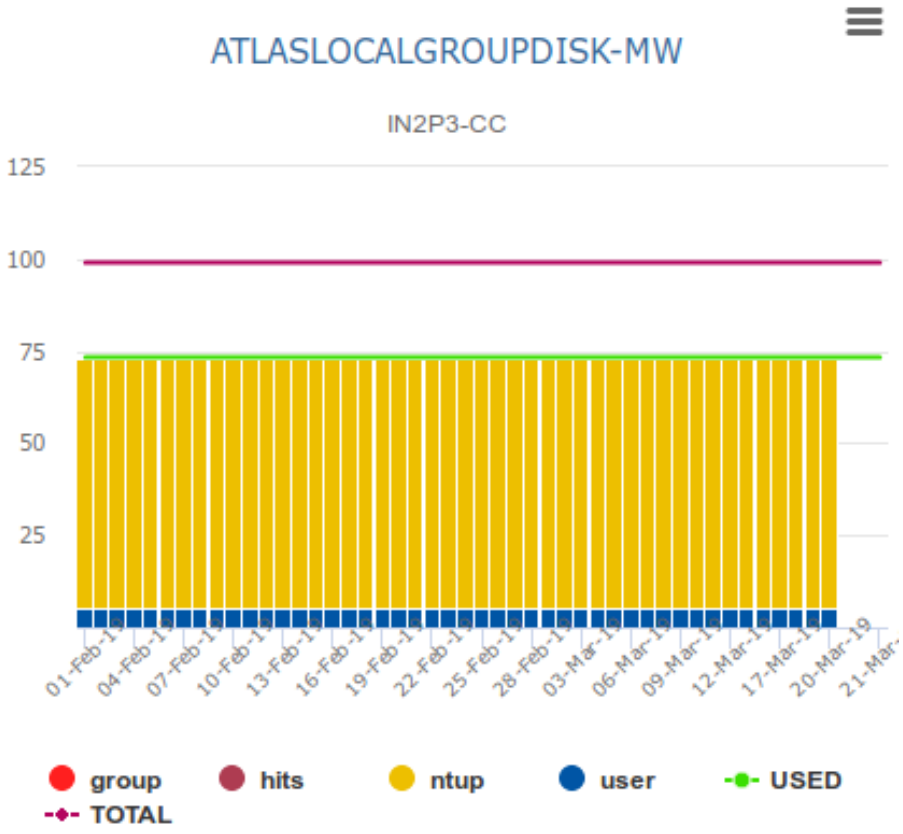

Usage of local batch system

cctools view [*link*]

- Running and requested jobs accessing sps (last 8 weeks)

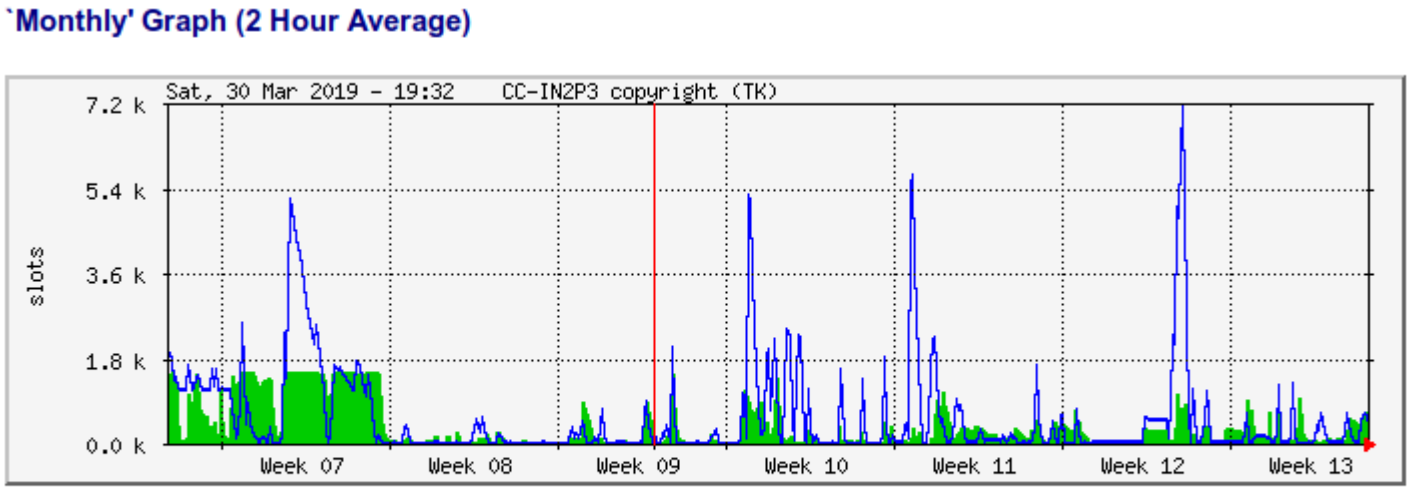

	Max	Average	Current
used:	1499.0	339.0	627.0
requested slots:	7083.0	515.0	406.0

→ at previous CAF <slots used>=612, <slots requested>=232
max slots requested>=4016

→ drop in usage since a few weeks

→ next slides for other « atlas group sub-projects »

→ see last slide of Manoulis on the agenda, to get the correspondance between the projects and the panda queues

Usage of local batch system

- Statistics of CC-IN2P3 (by Manoulis)
  Reported period from 01-12-2018 to 07-02 2019
  T3 atlas job on CC-IN2P3 BATCH farm
  Only queue long (max ~1500 current job, FIFO policy per user)

  - Qtime = start_time - submition_time
  - WallClock = end_time – start_time
  - Aratio = Qtime / WallClock
  - « A percentile (or a centile) is a measure used in statistics indicating the value below which a given percentage of observations in a group of observations falls»
  - « For example, the 50th percentile is the value (or score) below which 50% of the observations may be found (e.g. the median of the distribution) »

Distribution of the Wallclock time of all jobs

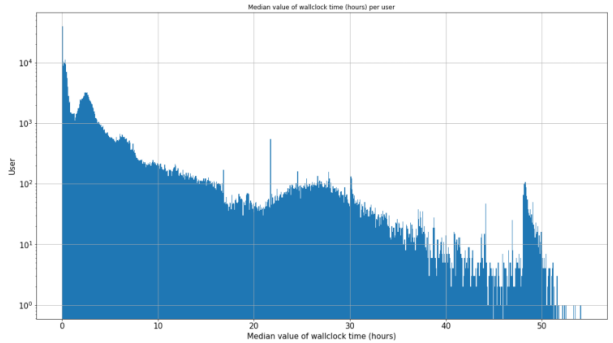

- 50% of the job exhibits WallClock time less tan 1.82 Hours
- Where a 5% of the job exhibits WallClock time greater than 23.42 Hours

Percentiles

qname	count	mean	std	min	5%	10%	25%	50%	75%	90%	95%	max
long	265955.0	4.32	7.53	0.0	0.0	0.02	0.27	1.82	4.16	11.99	23.42	54.08

Usage of local batch system

- Statistics of CC-IN2P3 (by Manoulis)
  - Median wallclock and qtime of the jobs are not to high ~O(1h)
    this is not a worry about the average behavior of the system.
    But wallclock and qtime distributions exhibit tails (95% percentile ~ O(20H))
    and this could block some user on particular dates

  - High submition rate (particular one user) cause the saturation of the resources (up to the limit of the slots).

  - The user with high submition rate, submit the 48% of the total job and consumed the 68% of the total wall clock time for the given period (01/12/2018 to 07/02/2019)

Usage of altas ressources cctools view [*link*]

`Monthly' Graph (2 Hour Average) atlas project

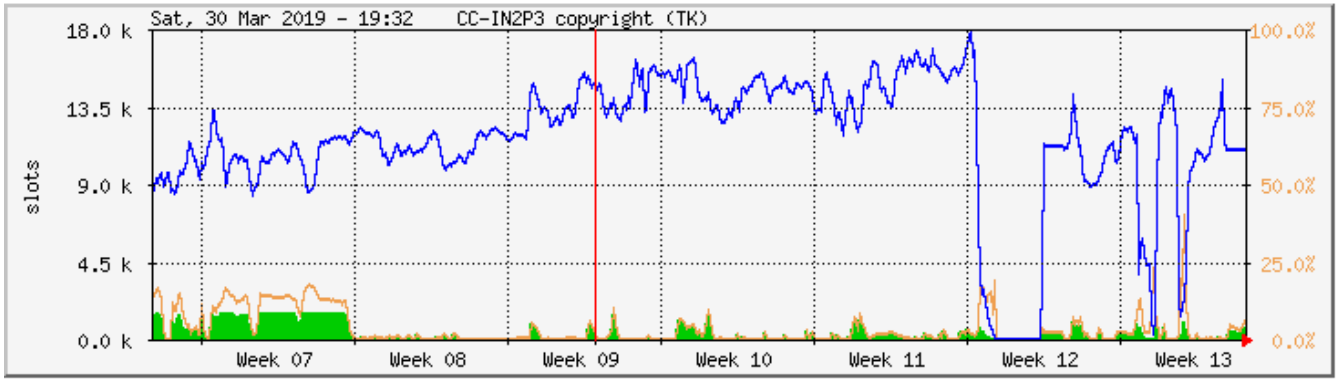

	Max	Average	Current
atlas project used slots:	1499.0	346.0	628.0
atlas group used slots:	17.8 k	11.6 k	11.0 k
Percentage	40.0 %	3.0 %	6.0 %

→atlas project : same as previous slide (batch queues using sps)

`Monthly' Graph (2 Hour Average) atlas T1ana (ANALY_IN2P3_CL7)

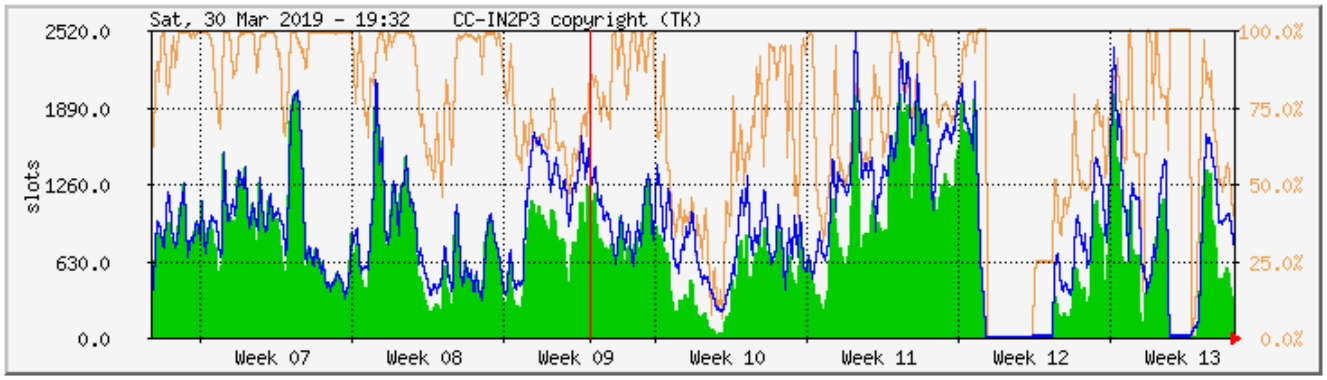

	Max	Average	Current
atlas_T1ana project used slots:	1997.0	716.0	327.0
atlas_T1ana project used & requested slots:	2497.0	906.0	733.0
Percentage	100.0 %	79.0 %	45.0 %

Usage of altas ressources cctools view [*link*]

atlas T1pmc « multicore » (IN2P3-CC_CL7_MCORE, IN2P3-CC_CL7_MCORE_HIMEM,)

`Monthly' Graph (2 Hour Average)

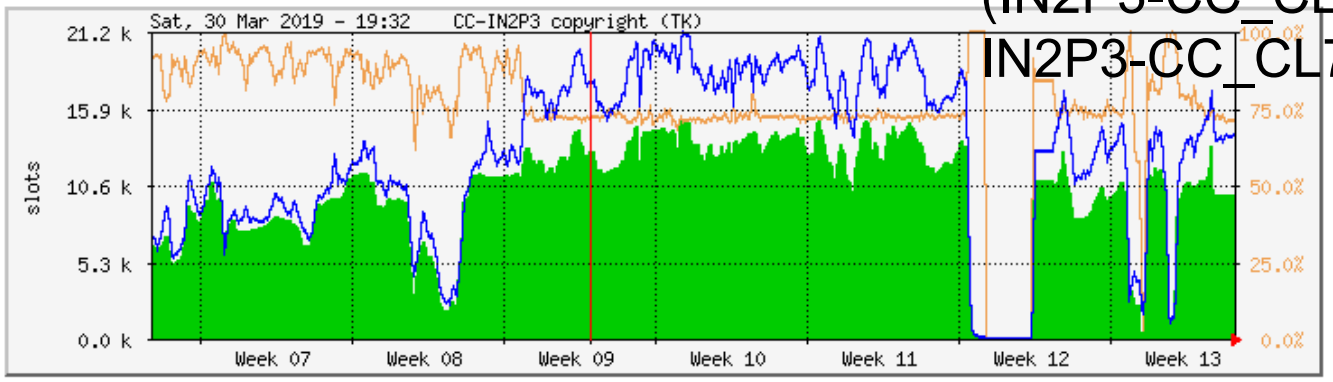

	Max	Average	Current
atlas_T1pmc project used slots:	15.0 k	9913.0	9957.0
atlas_T1pmc project used & requested slots:	21.2 k	12.8 k	14.1 k
Percentage	100.0 %	77.0 %	71.0 %

atlas T1prod (IN23P-CC_CL7, IN23P-CC_CL7_HIMEM, IN2P3_CC_CL7_VVL)

`Monthly' Graph (2 Hour Average)

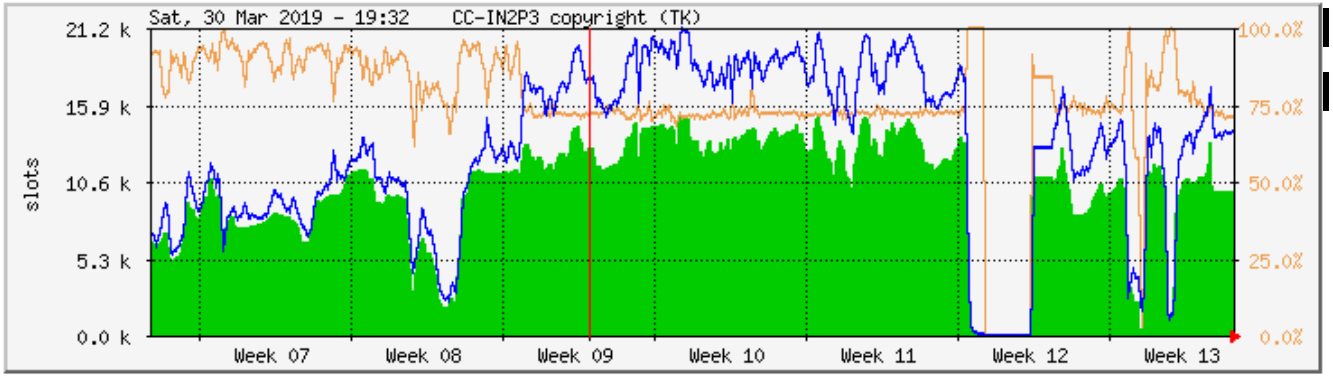

	Max	Average	Current
atlas_T1pmc project used slots:	15.0 k	9913.0	9957.0
atlas_T1pmc project used & requested slots:	21.2 k	12.8 k	14.1 k
Percentage	100.0 %	77.0 %	71.0 %

Usage of altas ressources cctools view [*link*]

atlas T1ufd « unified queues » - see presentation of Manoulis on the agenda IN2P3_CC_CL7_UCORE

`Monthly' Graph (2 Hour Average)

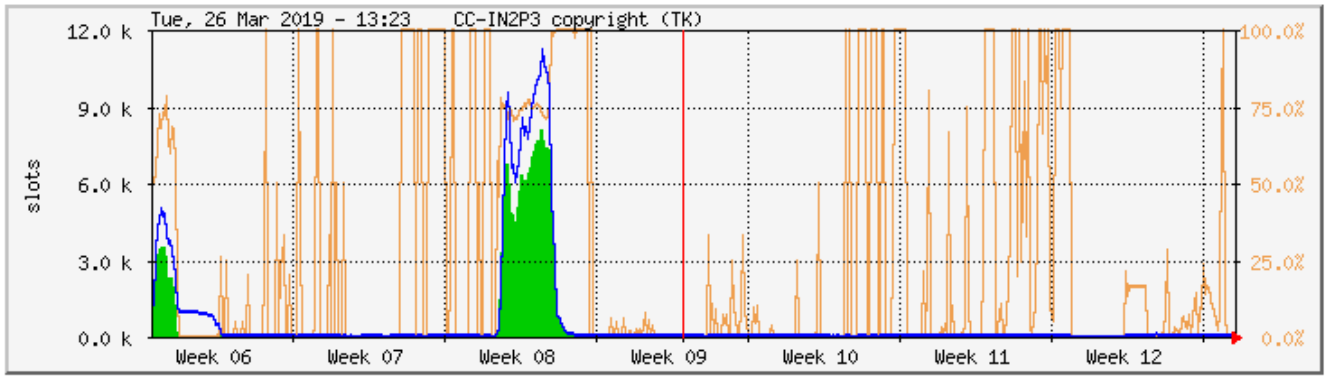

	Max	Average	Current
atlas_T1ufd project used slots:	8082.0	365.0	0.0
atlas_T1ufd project used & requested slots:	11.2 k	534.0	73.0
Percentage	100.0 %	68.0 %	0.0 %