Centre de Calcul de l'Institut National de Physique Nucléaire et de Physique des Particules

Source: LSST

LSST data release processing at CC-IN2P3

status and perspectives

fabio hernandez

CCIN2P3

LSST-France, Paris, Nov 8th, 2018

cnrs

Highlights for 2018

Highlights for 2018

- DESC DC2 as a unique opportunity to exercise CC-IN2P3's platform for data release processing

  see details in talk about DC2 by Johann

- Stabilisation of cvmfs as main channel for delivering the LSST software pipelines to all relevant hosts at CC-IN2P3 as well as personal computers

  the `/cvmfs/sw.lsst.eu` *software repository is now also available at NERSC, OpenScienceGrid sites and replicated at CERN*

  details of how to use: **https://sw.lsst.eu**

Highlights for 2018 (cont.)

- Jupiter-based Python notebooks, equipped with LSST software and with access to datasets, considered stable and usable

  see documentation on how to use

  elasticity (i.e. capacity to resize the backend according to demand) not yet implemented

  working on adding visualisation component based on Firefly

- Qserv cluster

  all the components of the cluster now managed by Puppet ⇒ reproducibility of the software installation (thanks to F. Wernli and F. Jammes)

  prepared platform for exporting data from Qserv cluster to Google cloud, as part of the evaluation of Google platform for hosting the catalog database (see F. Jammes talk)

Highlights for 2018 (cont.)

- Shared group disk area under `/pbs/throng/lsst`

  for sharing small- to medium-sized files (e.g. software, code, scripts, documents, etc.) but not for large datasets

  see guidelines in the documentation

- Set of 10 worker nodes in the production batch farm configured with 10 Gbps network links to GPFS servers

  we are collecting accounting data produced by GridEngine for understanding impact

  dedicated experiment to measure benefits (if any) for LSST jobs yet to be organised in the framework of DESC DC2

- Tools for import/export of data being regularly exercised between CC-IN2P3 and NERSC with DESC DC2 data

  ongoing development work to make them easily usable by anybody

Highlights for 2018 (cont.)

Data flow: **NERSC** (GPFS) → **CC-IN2P3** (GPFS) [3 servers, 4 clients]

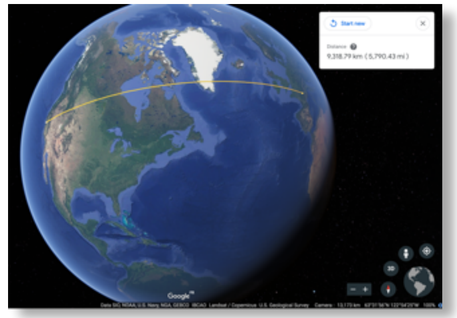

Aggregated application-level network throughput: **1.5 GB/s (12 Gbps)**

secure HTTP ⇒ integrity, confidentiality

pull model, disk-to-disk transfer, wide area network, 150ms RTT

Connectivity provided by

RENATER
CONNECTEUR DE SAVOIRS

Resource utilisation

CPU allocation

	2018 Q1	2018 Q2	2018 Q3	2018 Q4
Request [HS06 hours]	10M	10M	10M	20M
Allocation [HS06 hours]	10M	10M	10M	20M

10M HS06.hours per quarter is roughly equivalent to **430 recent CPU cores** permanently usable by LSST

CPU utilisation

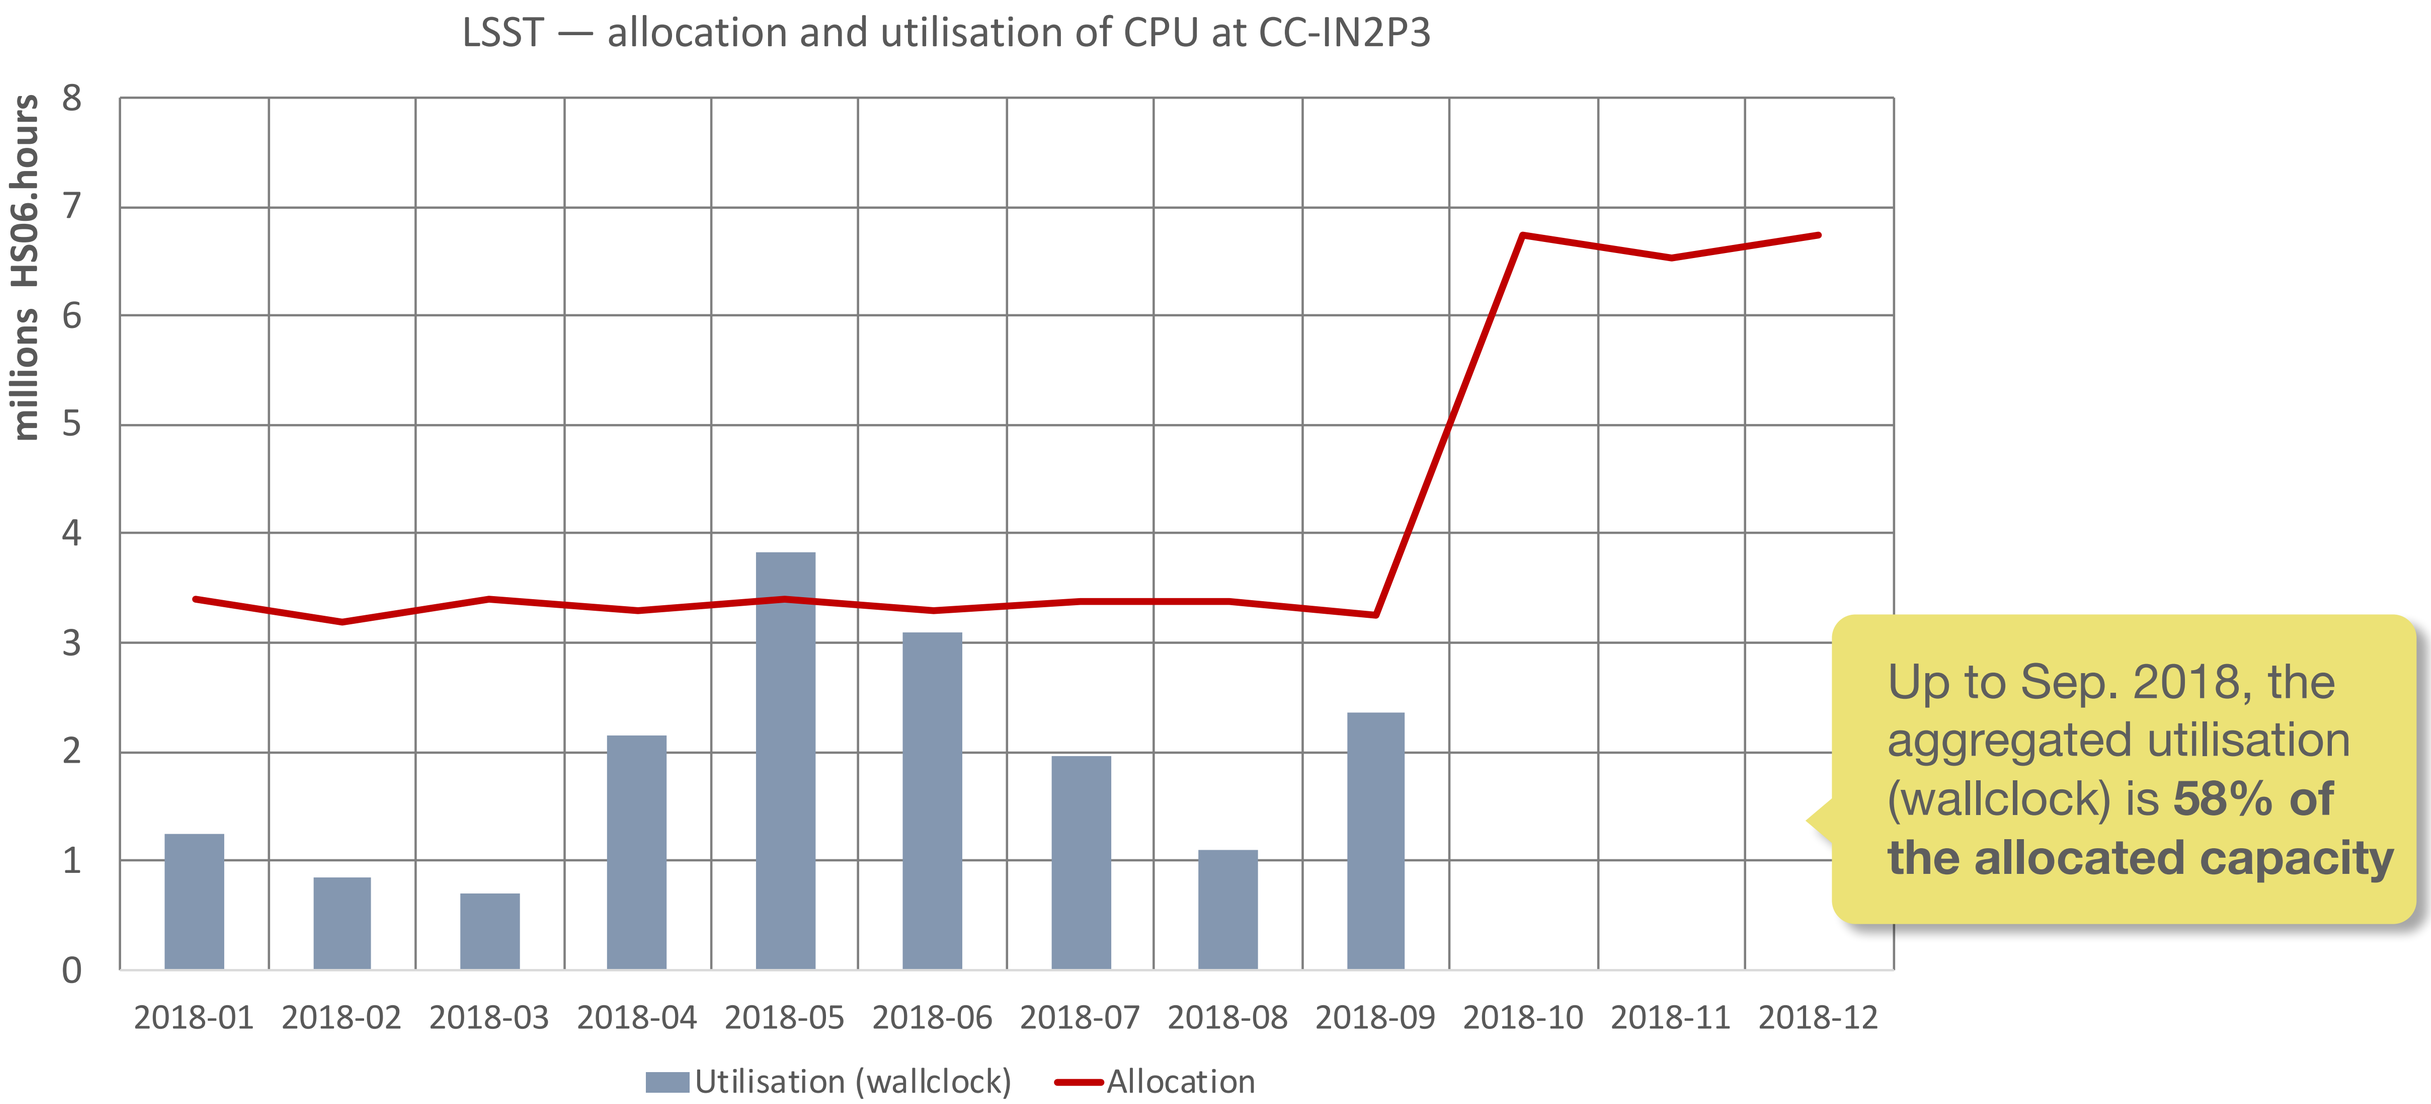

Up to Sep. 2018, the aggregated utilisation (wallclock) is **58% of the allocated capacity**

Disk allocation /sps/lsst

	2017 Q4	2018 Q1	2018 Q2	2018 Q3	2018 Q4
Request [TB]		+200	+300	+500	+500
Allocation [TB]	233		+303	+524	

- By the end of 2017, the disk allocation was 233 TB, 95% of that was used
- A reconfiguration of GPFS, including better hardware and LSST data redistribution, was implemented for improving performance.
  Example: SSD is now used for fast access to file system metadata and for storing small files (thanks to L. Tortay)
- No further allocation increase expected before end 2018

Disk storage utilisation /sps/lsst

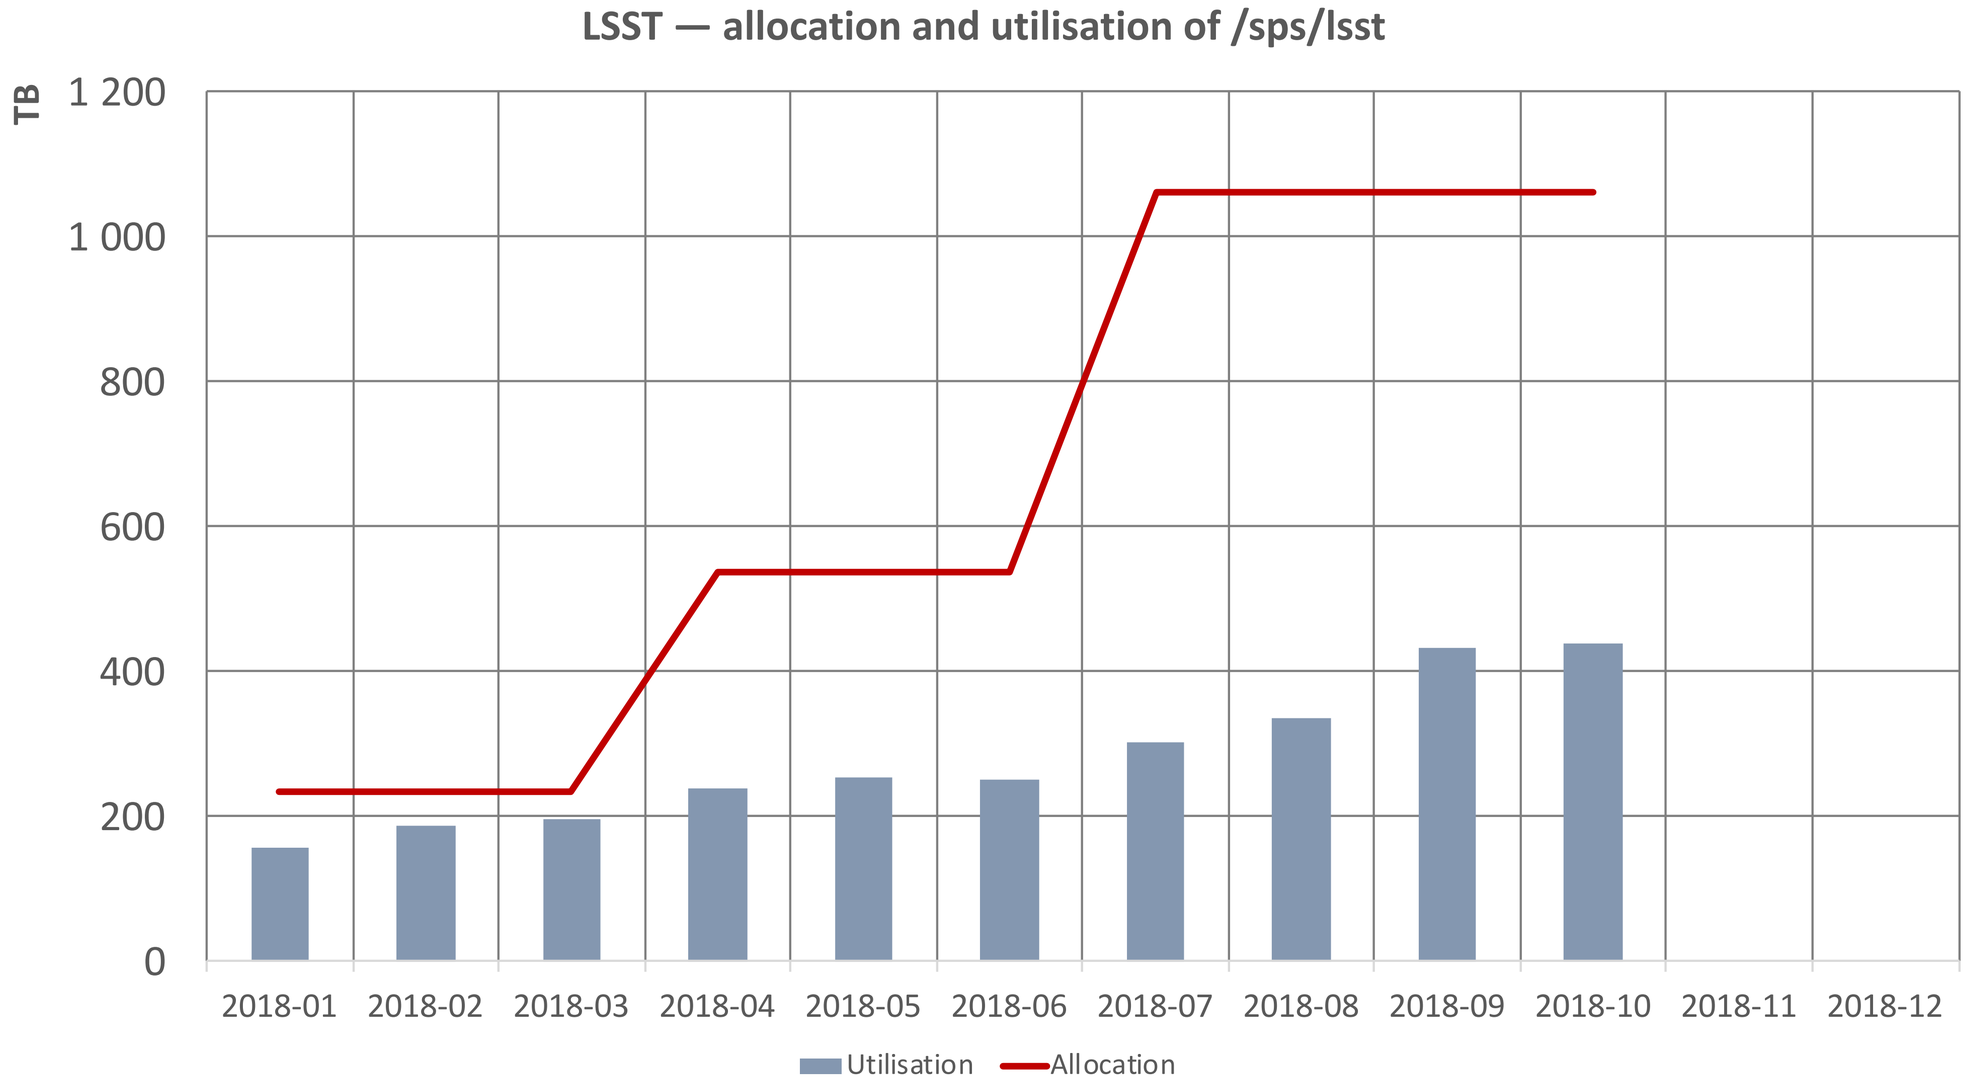

Training

Training

- Two upcoming training sessions, not devoted to but hopefully relevant for LSST members
- Python, Nov. 12th - 14th

  https://indico.in2p3.fr/event/17933
- SQL, Dec. 12th - 14th

  https://indico.in2p3.fr/event/17977

Perspectives

Perspectives

- Migration of your `$HOME` is imminent

  this is part of a campaign to decommission AFS: you will be notified by e-mail

  your new `$HOME` *will be under* `/pbs/home` *and you will have a default quota of* ***20 GB per account***

  already done for production accounts: `lsstprod` *and* `descprod`

  NFSv4 ***Access Control Lists*** *can be used with your future new* `$HOME`

  see `nfs4_getfacl(1)` *and* `nfs4_setfacl(1)`

- Allocation and reorganisation of namespace under `/sps/lsst` for groups yet to be performed

  beware that this is likely to be disruptive, but necessary

Perspectives (cont.)

- Hosts in the Qserv cluster are out of the 3 years-long warranty period

  currently investigating if we purchase an 2 years-long extension of warranty, or

  use that money to start a progressive replacement by new hardware better suited for database servers: unlikely to be possible with 2018 budget

- Work just started to build dashboards of LSST-related activity at CC-IN2P3

  e.g. batch farm, disk storage, login farm, data exchange, etc.

  uses the production data collection infrastructure based on ElasticSeach, plus Grafana and Kibana for visualisation

  the intention is that those dashboard will be accessible to authenticated LSST members

  work lead by B. Chambon with contribution by F. Wernli

Perspectives (cont.)

- Ongoing work to develop tools for detailed analysis of LSST batch activity

  based on accounting data emitted by GridEngine augmented with data extracted from the jobs' logs

- Network bandwidth between NCSA and CC-IN2P3 to be effectively 20 Gbps by the end of 2018

  currently, the bottleneck link is 10 Gbps

- New research engineer hired to work on LSST matters

  limited duration contract, effective start expected in December 2018

Questions & Comments