

Status progress since last meeting

- ➔ **Fairly stable operations of both clusters** (`ce01.lhep.unibe.ch`; `ce02.lhep.unibe.ch`)
- ➔ **ARC upgraded to 4.1.0-1.el6** (smooth upgrade on live system)
 Mandatory since ATLAS ops moved to Rucio as DDM tool in May, making the previous version of ARC obsolete.
- ➔ **ANALY_UNIBE-LHEP, UNIBE-LHEP, UNIBE-LHEP-UBELIX PanDA queues created in AGIS, added to HammerCloud**, following the commissioning of the new ATLAS Control Tower developed by Andrej, running in CERN, which has many new functionalities, e.g.:
 - Now PanDA queues at each ARC site rather than have all sites grouped under ARC and ARC_T2 PanDA queues
 - Can be tested by HammerCloud functional tests (automated auto-exclusion, whitelisting)
 - Automated exclusion from production and analysis in case of GOCDDB downtime
 - Can use the DATADISK and SCRATCHDISK at the local site, rather than only use the ND T1 SE
- ➔ **Both clusters commissioned for ATLAS analysis payloads and Multi-core ATHENA workloads**
- ➔ **DPM SE re-configured for xrootd data access for ATLAS, joined the German Federation (FAX)**
 - <http://dashb-atlas-ssb.cern.ch/dashboard/request.py/siteview#currentView=FAX+endpoints&fullscreen=true&highlight=false>
 - DPM versions from epel: 1.8.8-4.el5 (head node), 1.8.8-4.el6.x86_64 (disk servers)
- ➔ **New ARC CE (ce03) setup, SLURM master and 1 WN (all VMs) as testbed for HPC submission**
- ➔ **GIIS (`giis.lhep.unibe.ch`) and VOMS (`voms.lhep.unibe.ch`) services commissioned and in production**
 - Clients and ARC CEs configurations needed change to point to them (was `giis.smscg.ch` – `voms.smscg.ch`)



HammerCloud Gangarobot

History Legend



Historic view for "panda_queues_all" from 00:00 01.04.2014 to 00:00 17.08.2014

Show entries

Search:

PANDA queue	SITE Name	TIER	CLOUD	History plot time bin = 276 hours	offline		brokeroff		online		NoQueue		test	
					%	count	%	count	%	count	%	count	%	count
ANALY_LUNARC	SE-SNIC-T2	T2	ND		0.02	1	0	0	52.7	63	0	0	9.04	80
ANALY_SiNET	SINET	T2D	ND		0.05	1	0	0	63.86	166	0	0	12.02	6
ANALY_UNIBE-LHEP	UNIBE-LHEP	T2	ND		6.19	2	0	0	54.57	49	0	0	18.32	131
ANALY_UPPMAX	SE-SNIC-T2	T2	ND		0.09	2	59.1	35	0.47	28	0	0	0.52	79
LUNARC	SE-SNIC-T2	T2	ND		0.01	1	0	0	47.45	33	0	0	8.32	8
SINET	SINET	T2D	ND		0.02	1	0	0	56.49	71	0	0	13.54	56
UNIBE-LHEP	UNIBE-LHEP	T2	ND		0.03	1	0	0	51.87	27	0	0	2.41	16
UNIBE-LHEP-UBELIX	UNIBE-LHEP	T2	ND		0.03	1	0	0	53.46	35	0	0	0.81	2


Showing 1 to 8 of 8 entries

First Previous 1 Next Last

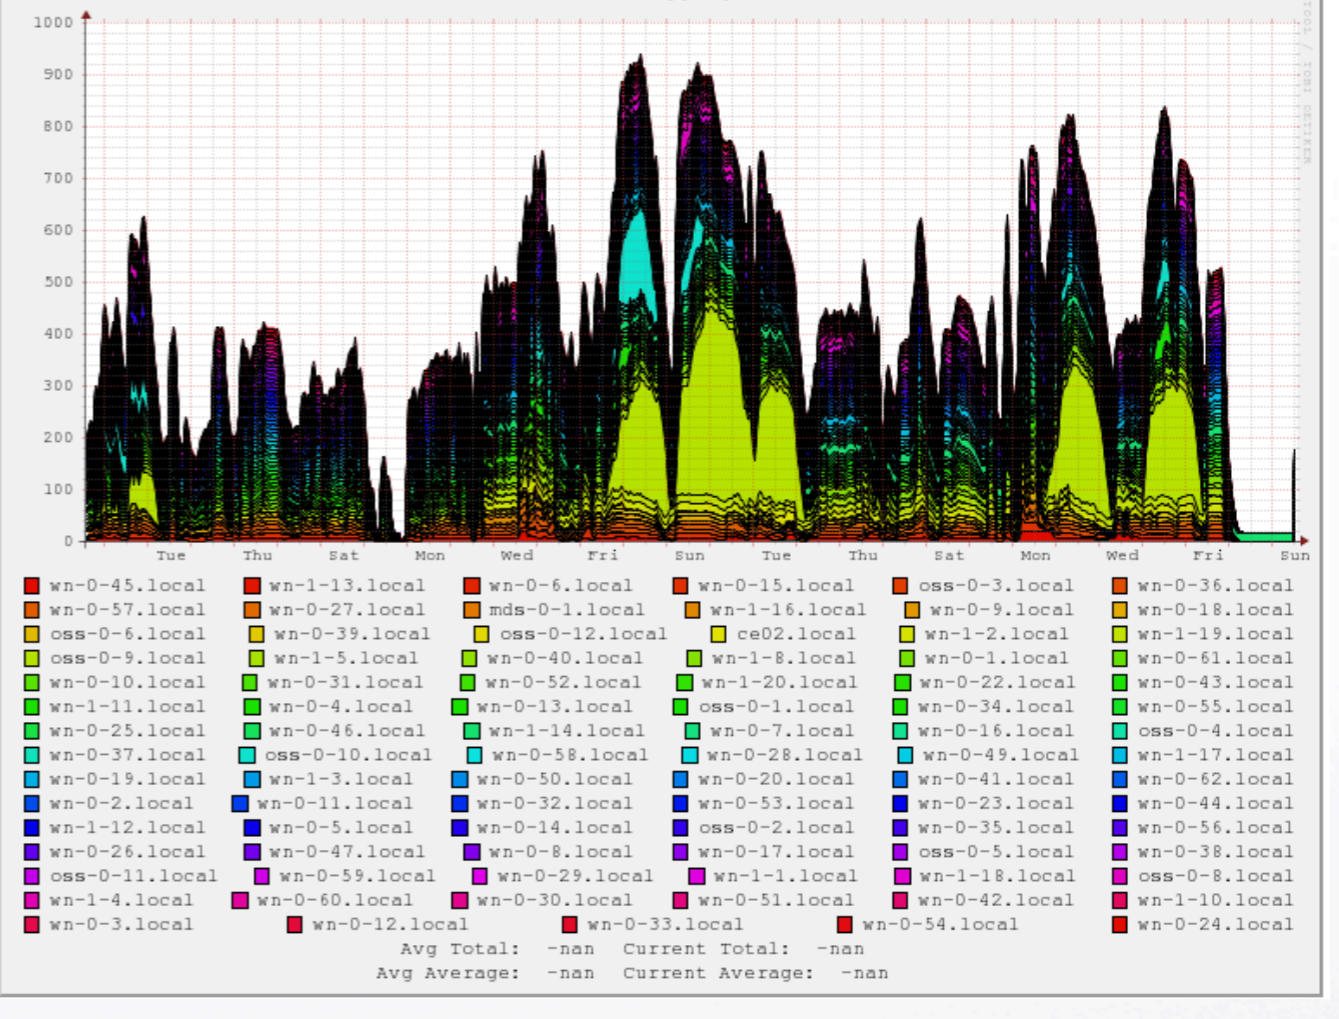
ce01

ce02

UNIBE-LHEP WLCG Cluster aggregated load_one last month



UNIBE-LHEP WLCG Cluster aggregated load_one last month



Issues and mitigations (1/3)

➔ **Issues with stale files in ARC sessiondir.** These are files left over by failed jobs, which end up clogging up the directory => the Infosys becomes very slow/unresponsive, clusters not visible in GIIS

Mitigation: Added a weekly cron to perform a cleanup

S

➔ **Issues with large amount of files in ARC sessiondir.** These are files related to jobs not retrieved by the users => the Infosys becomes very slow/unresponsive, clusters not visible in GIIS

Solution: Asked t2k.org users to retrieve/clean up their job outputs

➔ **Lustre MDS glitch on Fri 16th May on the ce01 cluster** caused the cluster to hang for some hours.

Solution: kill all jobs, unmount all clients, stop all OSTs. Power-cycle the MDS (stopping the MDT (unmount) would not work). Recovery took ~1.5h

➔ **EGI ops availability/reliability 33% in April**

probe org.sam.SRM-GetURLs-/ops/NGI/Germany failing, causing all the other SRM probes to go to Unknown state.

Solution: The problem self-resolved on 2nd May (no changes on the SE).

Could not request availability/reliability re-calculation

Issues and mitigations (2/3)

➔ **Middleware issues:**

- ▶ **GLUE2 Validator Warnings due to a minor ARC infosys bug** (also in latest version 4.1.0-1.el6)
(https://xgus.ggus.eu/ngi_ch/?mode=ticket_info&ticket_id=314)

Mitigation: **patched SGEmod.pm**

- ▶ **bdii service crashed on the site-bdii** (one-time occurrence)

left behind the slapd process running, preventing the restart cron from fixing it within the 15min window.

Solution: **following Nagios alert, restarted services manually**

- ▶ **xrootd broken (segfaults) in May**

Solution: **upgrade DPM to latest version 1.8.8**

- ▶ **a-rex crashes regularly on the ce02 cluster** (on average twice a month)

Solution: **issue not investigated/resolved. Following Nagios alerts, restart service manually**

Issues and mitigations (3/3)

➔ Accounting issues:

▶ Migration from ur-logger/SGAS to Jura/APEL nightmarish

Migration scheme: ur-logger/SGAS => Jura/SGAS => Jura/APEL

Timeline: started migration in March 2014, completed by end of June 2014

Issues: complex mechanism, poor documentation, poor support (ARC), obscure operational changes (APEL), etc

Outcome: migration accomplished, but loss of accounted job records for some periods.

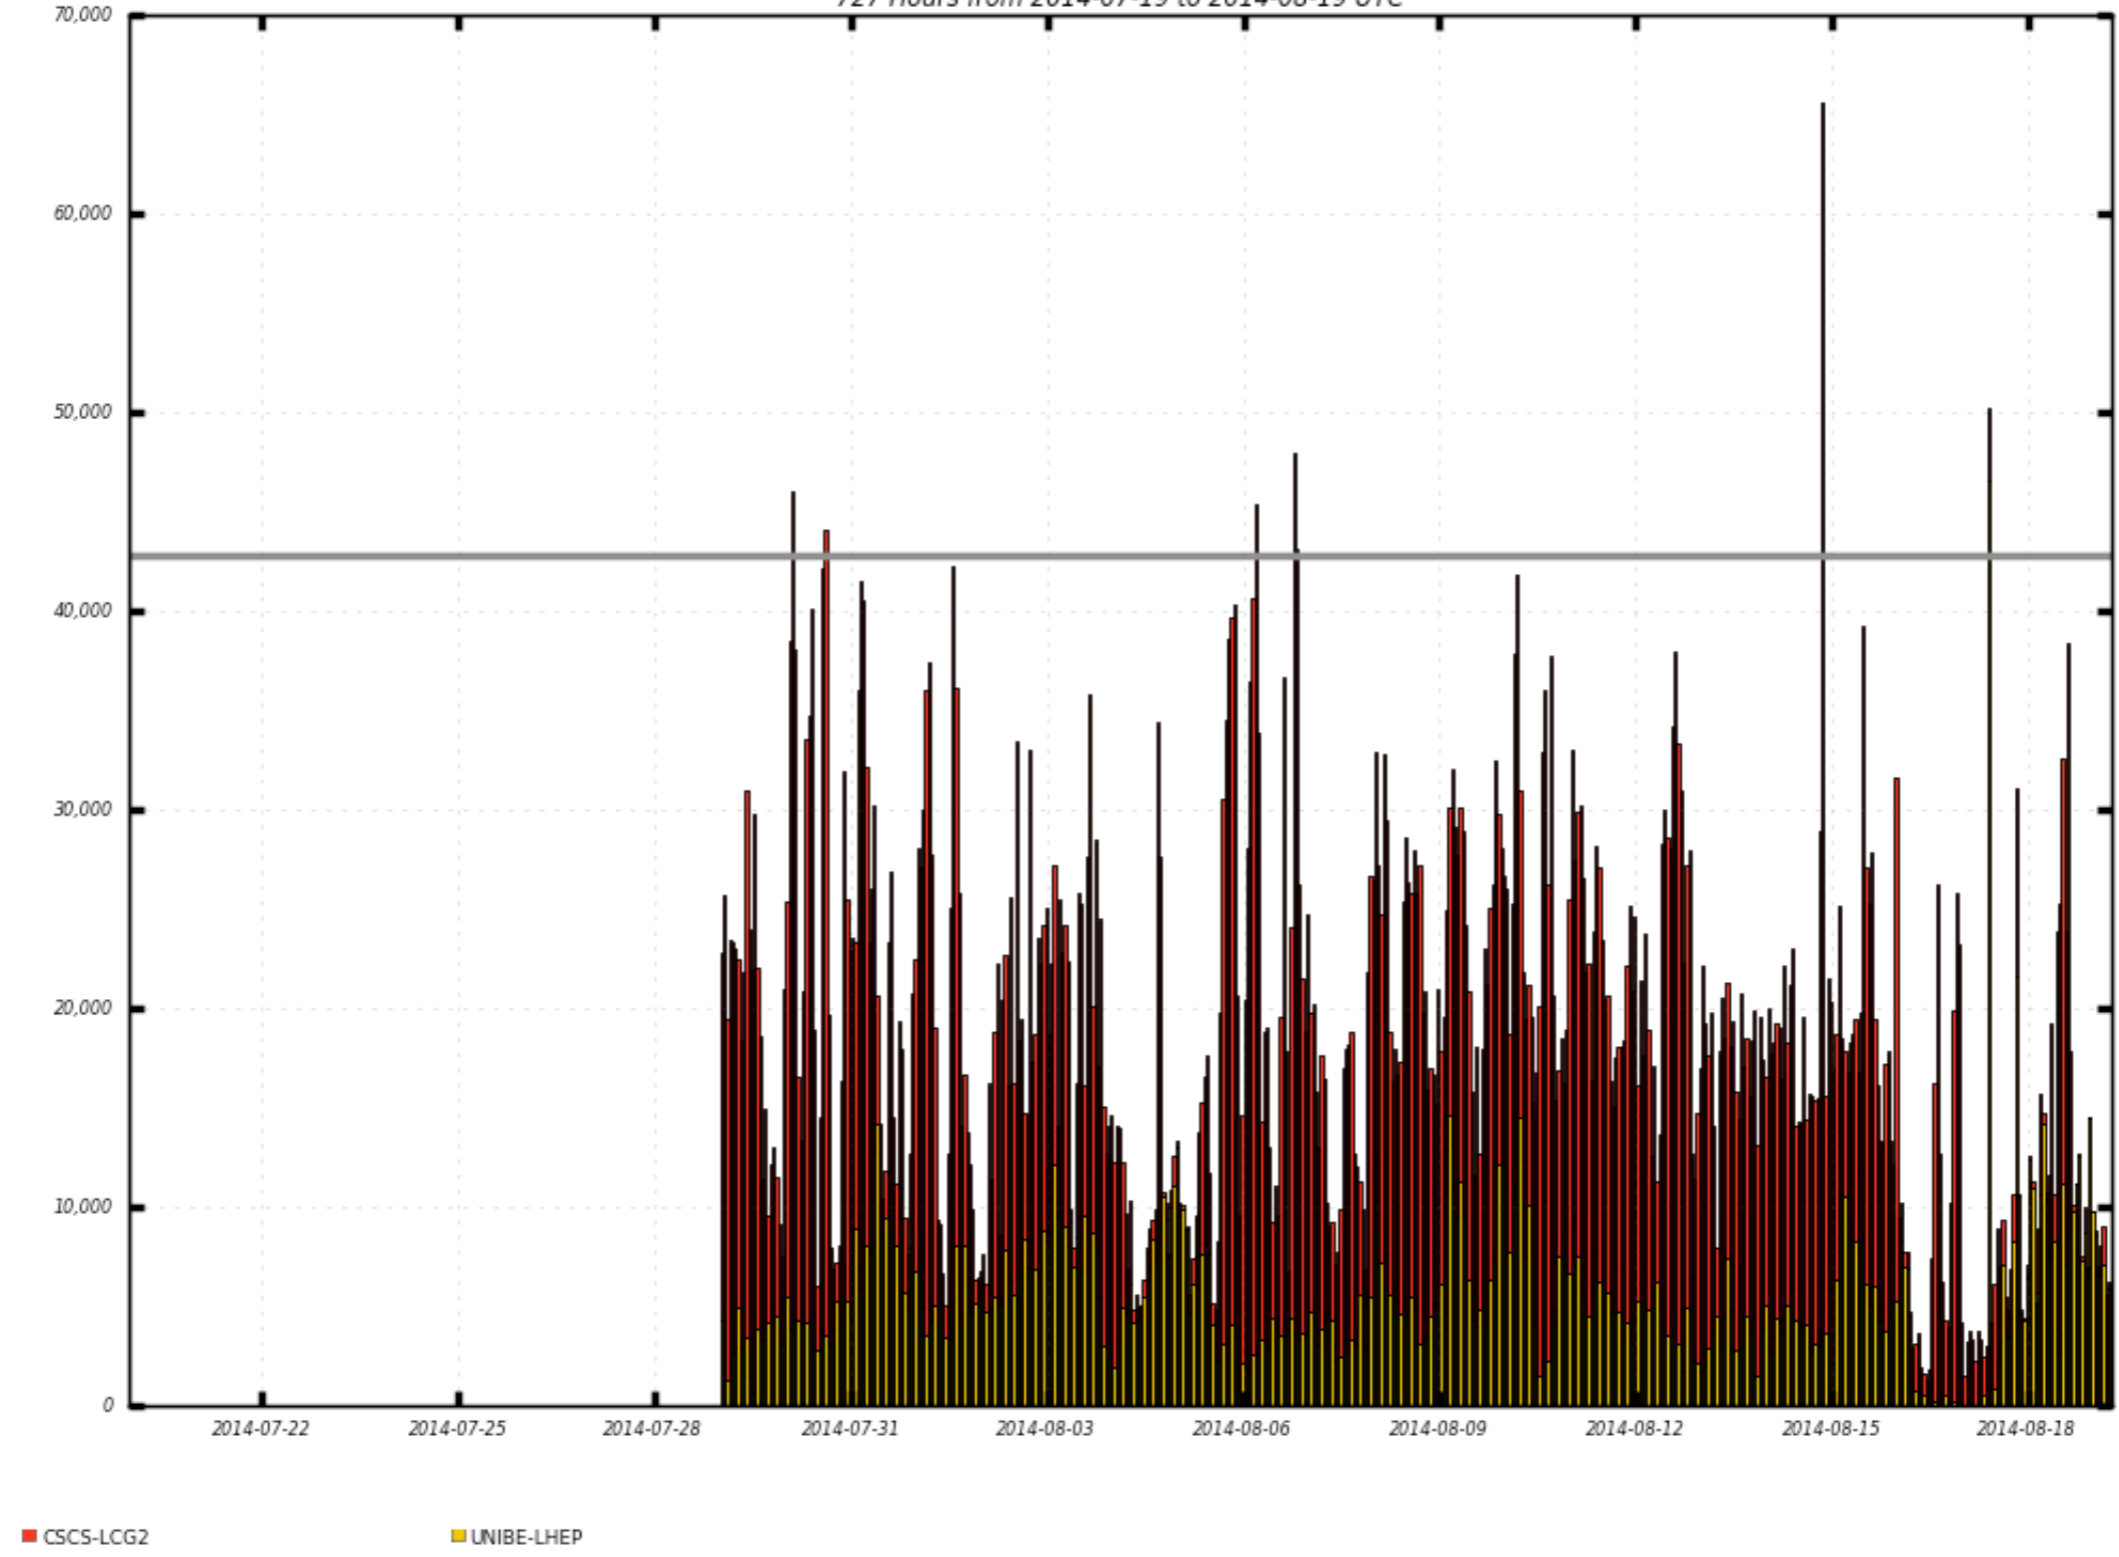
- ▶ Cross-check against batch server records not possible; gridengine does not account properly for multi-core jobs
- ▶ Attempted a cross-check against ATLAS own accounting: **found little sensible correlation**
- ▶ Estimated a ~10% loss over about 4 months

▶ **Recovery plans (?)**: in principle some of the lost records have an archive copy.

- ▶ A procedure exists to re-create appropriately formatted job records from archived records.
- ▶ Complex/clumsy. Will attempt only if time allows.



CPU HEPSEC06 Hours
727 Hours from 2014-07-19 to 2014-08-19 UTC

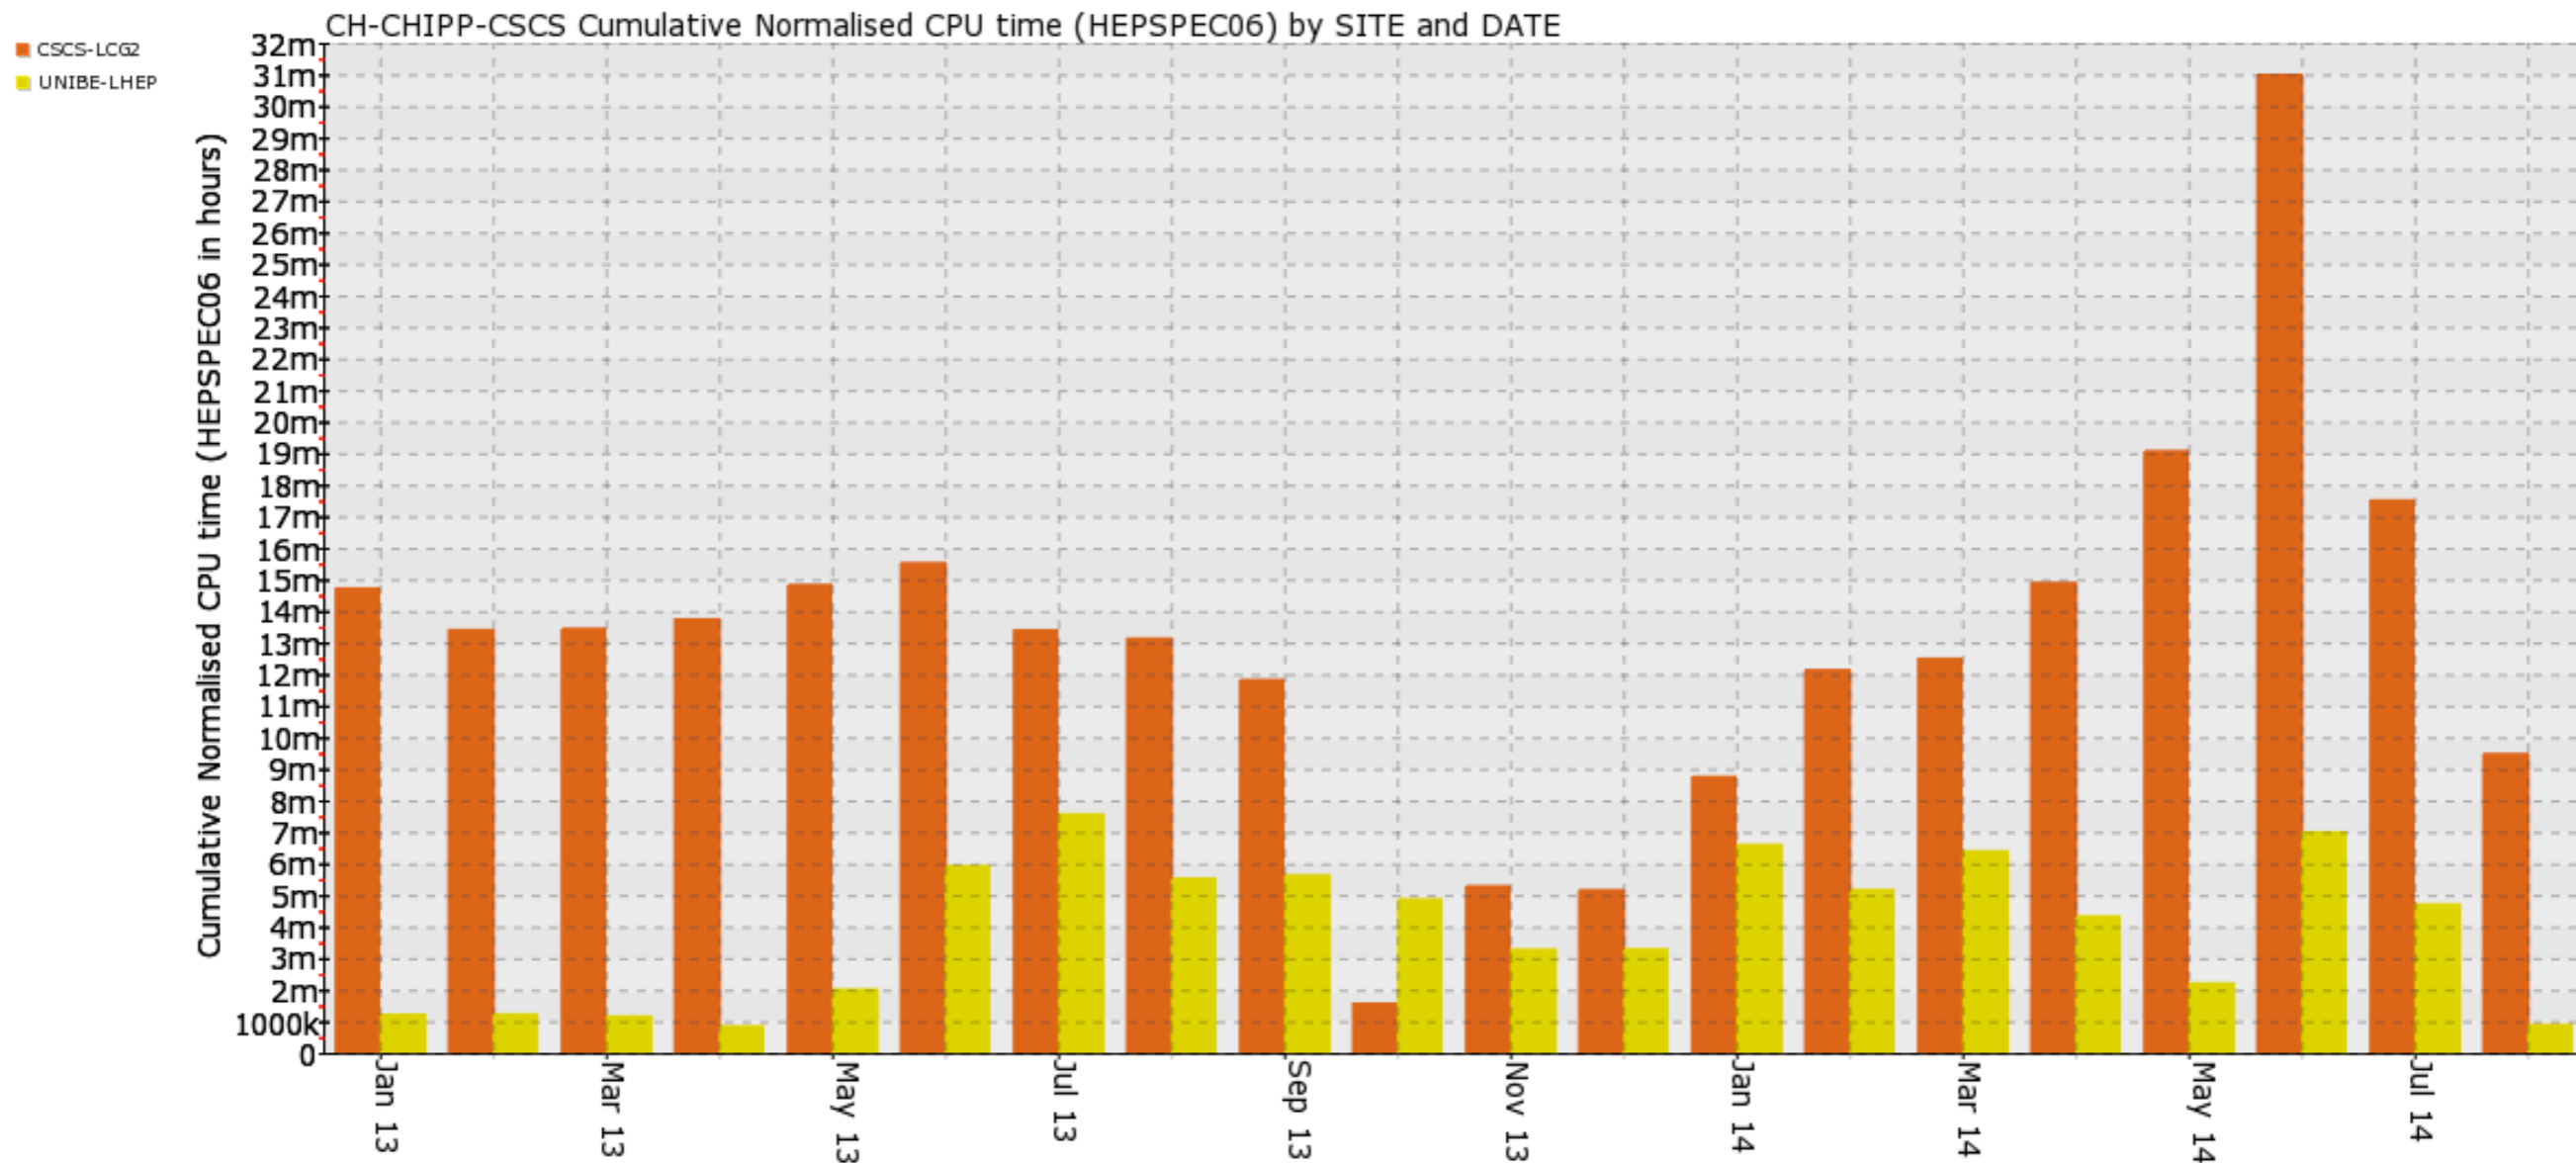


Maximum: 65,541 , Minimum: 0.00 , Average: 13,182 , Current: 6,234

Chart showing the Cumulative Normalised CPU time (HEPSPEC06) grouped by SITE and DATE (only information about LHC VOs is returned).

Developed by CESGA EGI View: / normcpu+HEPSPEC06 / 2013:1-2014:8 / SITE-DATE / lhc (x) / GRBAR-LIN / i

2014-08-19 00:01



UNIBE-LHEP-UBELIX

➔ Report from Michael:

- ▶ – buildup a proper tree-based IB topology and move GPFS to IB
- ▶- move from RHEL to CentOS
- ▶- usage Puppet (and the Foreman?) for life cycle and configuration management (replaces Red Hat Satellite)
- ▶- buy new compute nodes to be able to abandon really old, inefficient sun nodes
- ▶- new scratch storage (~400 TB)