

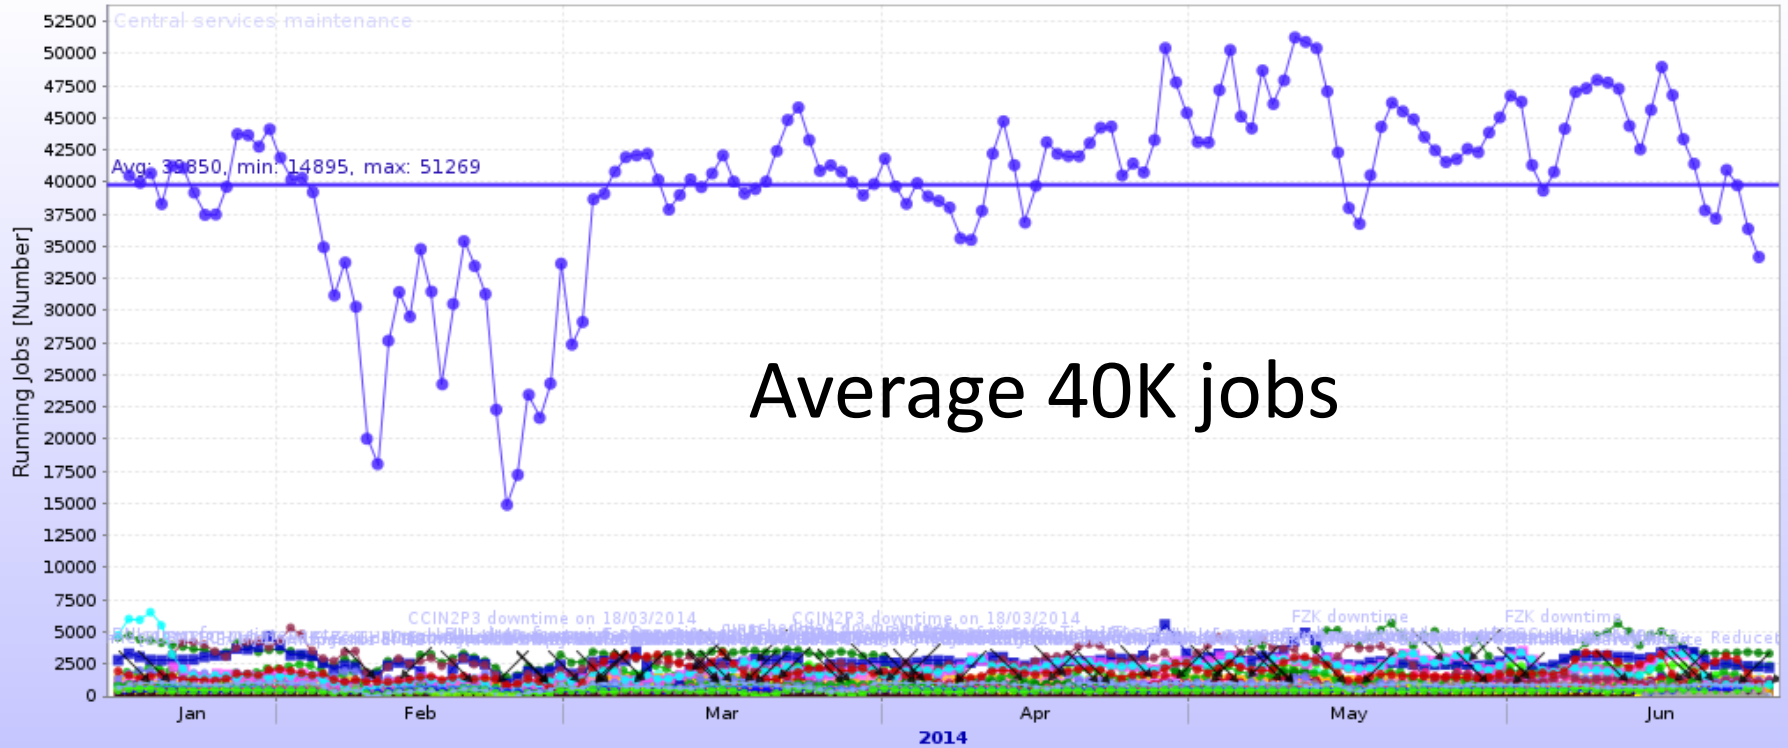
ALICE Grid operations

Offline week

26/06/2014


The GRID job profile in 2014

Running Jobs



- Athens — Bandung — Bari — Birmingham — BITP — BITP_ARC — Bologna — Bratislava — Cagliari — Catania — CCIN2P3 — CERN — CERN-CREAM
- CERN (Wigner) — CERN-HW6 — CERN-L — CERN (Meyrin) — CERN-VM6 — CERN-W5 — Clermont — CNAF — COMSATS — CSC — Cyfronet — DCSC_KU — FZK
- Grenoble — GRIF_IPNO — GRIF_IRFU — GRIF_IRFU_SHARED — GSI_2 — HHLR_GU — Hiroshima — IHEP — IPNL — ISMA — ISS — ISS_LCG — ITEP — JINR
- KFKI — KISTI-CREAM — KISTI_GSDC — KNU — Kolkata-CREAM — Kosice — LBL — Legnaro — LLNL — LUNARC — Madrid — MEPHI — NECTEC — NERSC
- NIHAM — NIKHEF — NIPNE — ORNL — OSC — Oxford — PAKGRID — pcalice92.cern.ch — PNPI — Poznan — Prague — RAL — RRC-KI — RRC_KI_T1
- SaoPaulo — SARA — SNIC — SPbSU — Strasbourg_IRES — Subatech — SUT — TACC — Torino — Torino-Torrent — Trieste — TriGrid_Catania — Troitsk
- Trujillo — UIB — UNAM — UNAM_T1 — Yerevan — ZA_CHPC — SUM

Resources delivery distribution



50/50 T0+T1/T2
 ~110M Wall hours
 Equivalent to
 same period in 2013
 (flat resources)


Status of infrastructure

- Operating system on all sites SL6 (or compatible)
 - Exception – CERN – about 50% still on SLC5
- Storage
 - xrootd updated to 3.2.x+
 - 5/60 sites have installed EOS
 - 22.6 PB available disk storage, 14.8 used (65%)
 - All major cleanups are done

Jobs by user since January 2014

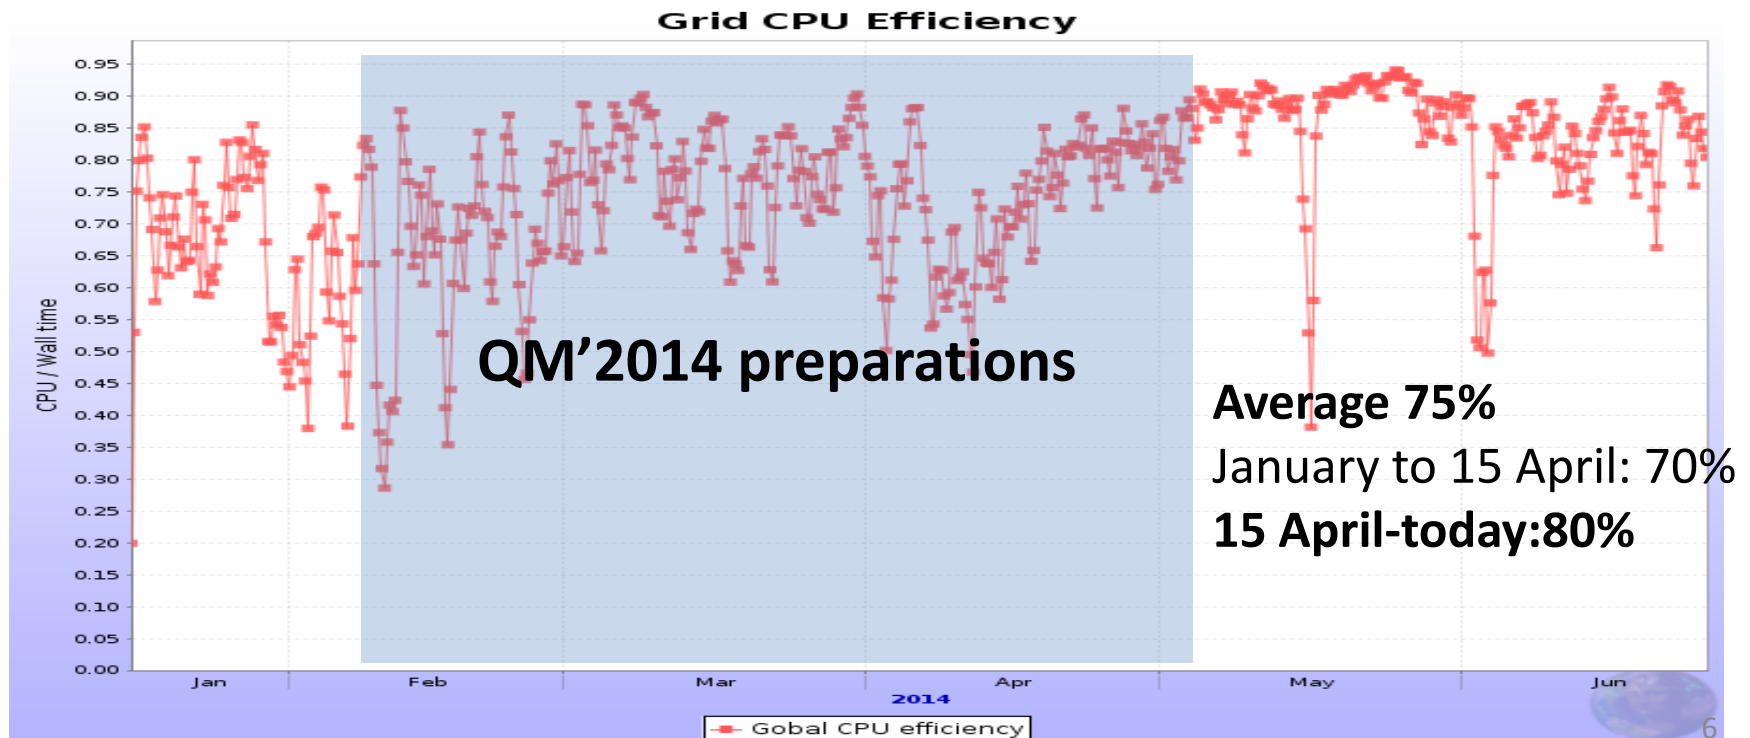
- 76% MonteCarlo (unchanged)
- 16% Organised analysis in trains (+6%)
- 2% RAW data reconstruction (-8%, software upgrades)
- 6% Individual user analysis (-6%)

Running jobs per user




Efficiency since January 2014

- Small effect due to high volume user activities (QM 2014)
- A fix in replica access algorithm (bug discovered in April, fixed in days by Costin, Andrei and Miguel) further increases the overall analysis efficiency




Replica access fix effect on LEGO

Jobs' efficiency per user



Data volumes since January 2014

- 111.4 PB read, 10.5 PB written
- Cleanup of storage done – close to zero ‘dark data’



Summary on efficiency

- Small negative effect due to QM analysis this year!
 - There was *no contention* for resources
- Replica access fix – overall efficiency gain of 10%
- T2s efficiencies are not much below T0/T1s
- Storage stability has increased since beginning of 2014

Activities

- AliEn code consolidation and cleanup – Miguel's presentation
- More detailed monitoring on data access – Costin (not shown)
- Investigation of lower efficiency in some of the T0 queues – ongoing, not shown
- CVMFS squid caches monitoring and alarms – ongoing, not shown

Site activities

- Migration to EMI3 – completed
- Deployment of additional storage resources
- Retirement of obsolete storage resources
 - Requires evacuation of data – better automatic tools needed
- **The infrastructure is stable – ‘flat resources’ year**
 - Higher level of activities is expected in 2015

Productions

- Steady rate of MC production
 - Number of production cycles is lower than in 2013
- RAW data production is delayed
 - So far we have LHC11a CPass0/CPass1
- Priorities may change (2011->2012)... the general workflow will be the same

MC as service tasks

- Very successful
 - Vladimir Kovalenko and Catalin-Lucian Ristea are handling all MC productions
- The number of productions is below expectations
 - i.e. they are still to be overloaded 😊

Summary

- The Grid is performing normally in 2014
- Software updates successfully completed
- MC productions completed according to requests
- Some backlog of RAW data (and corresponding MC) productions, hopefully we will be ready with these before start of 2015 data taking
- Analysis trains share continues to grow – very beneficial for stable resources utilization