



ALICE Operations

GDB

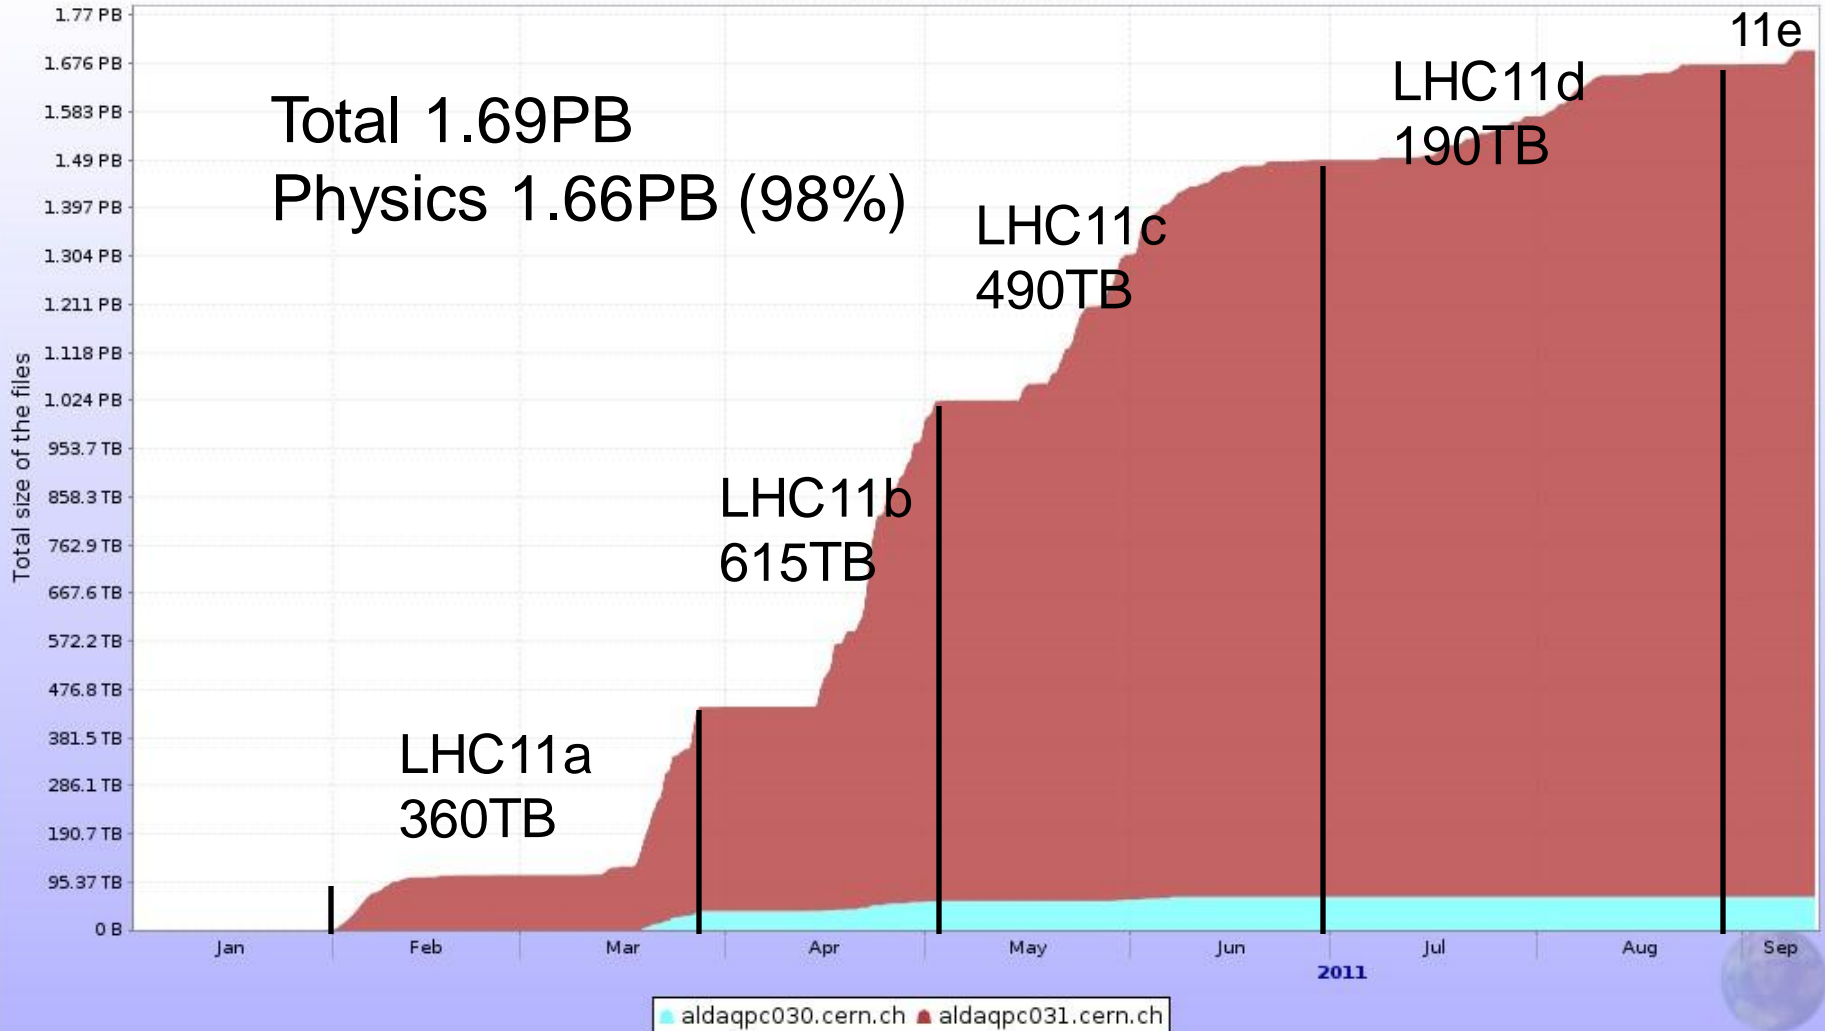
September 14, 2011

News

- Data taking and productions
- User activity – Grid analysis for a major physics conference
- Job efficiency – problems and remedies
- Sites operations, briefly
- Short term plans

RAW data collection

Total size of the files



RAW data processing

- LHC11a – Pass1

TOTAL	122,347/129,527	94.5%	<u>381,799,429</u>	216 jobs
-------	-----------------	-------	--------------------	----------
- 2.76TeV Pass3

TOTAL	32,760/32,997	99.3%	<u>95,398,974</u>	24 jobs
-------	---------------	-------	-------------------	---------
- LHC11b – Pass1

TOTAL	53,519/59,901	89.3%	<u>320,585,734</u>	168 jobs
-------	---------------	-------	--------------------	----------
- LHC11c – Pass1


TOTAL	97,683/116,621	83.8%	<u>147,595,998</u>	222 jobs
-------	----------------	-------	--------------------	----------
- LHC11d - CPass

TOTAL	30,240/81,791	36.10%	<u>24,684,987</u>	264 jobs
-------	---------------	--------	-------------------	----------

User analysis

- 2011 has been the first year of massive user analysis on the Grid
- In general a successful endeavor, with the following points in need of improvement
 - Move chaotic user analysis to smaller containers: ESD analysis (10% of RAW size) is still majority, large effort to migrate to AODs (2% of RAW size), not yet complete
 - High memory footprint: considerable reduction, more 'analysis trains' running

User activity (2)




- 6000K user jobs in average, ~30% of CPU resources

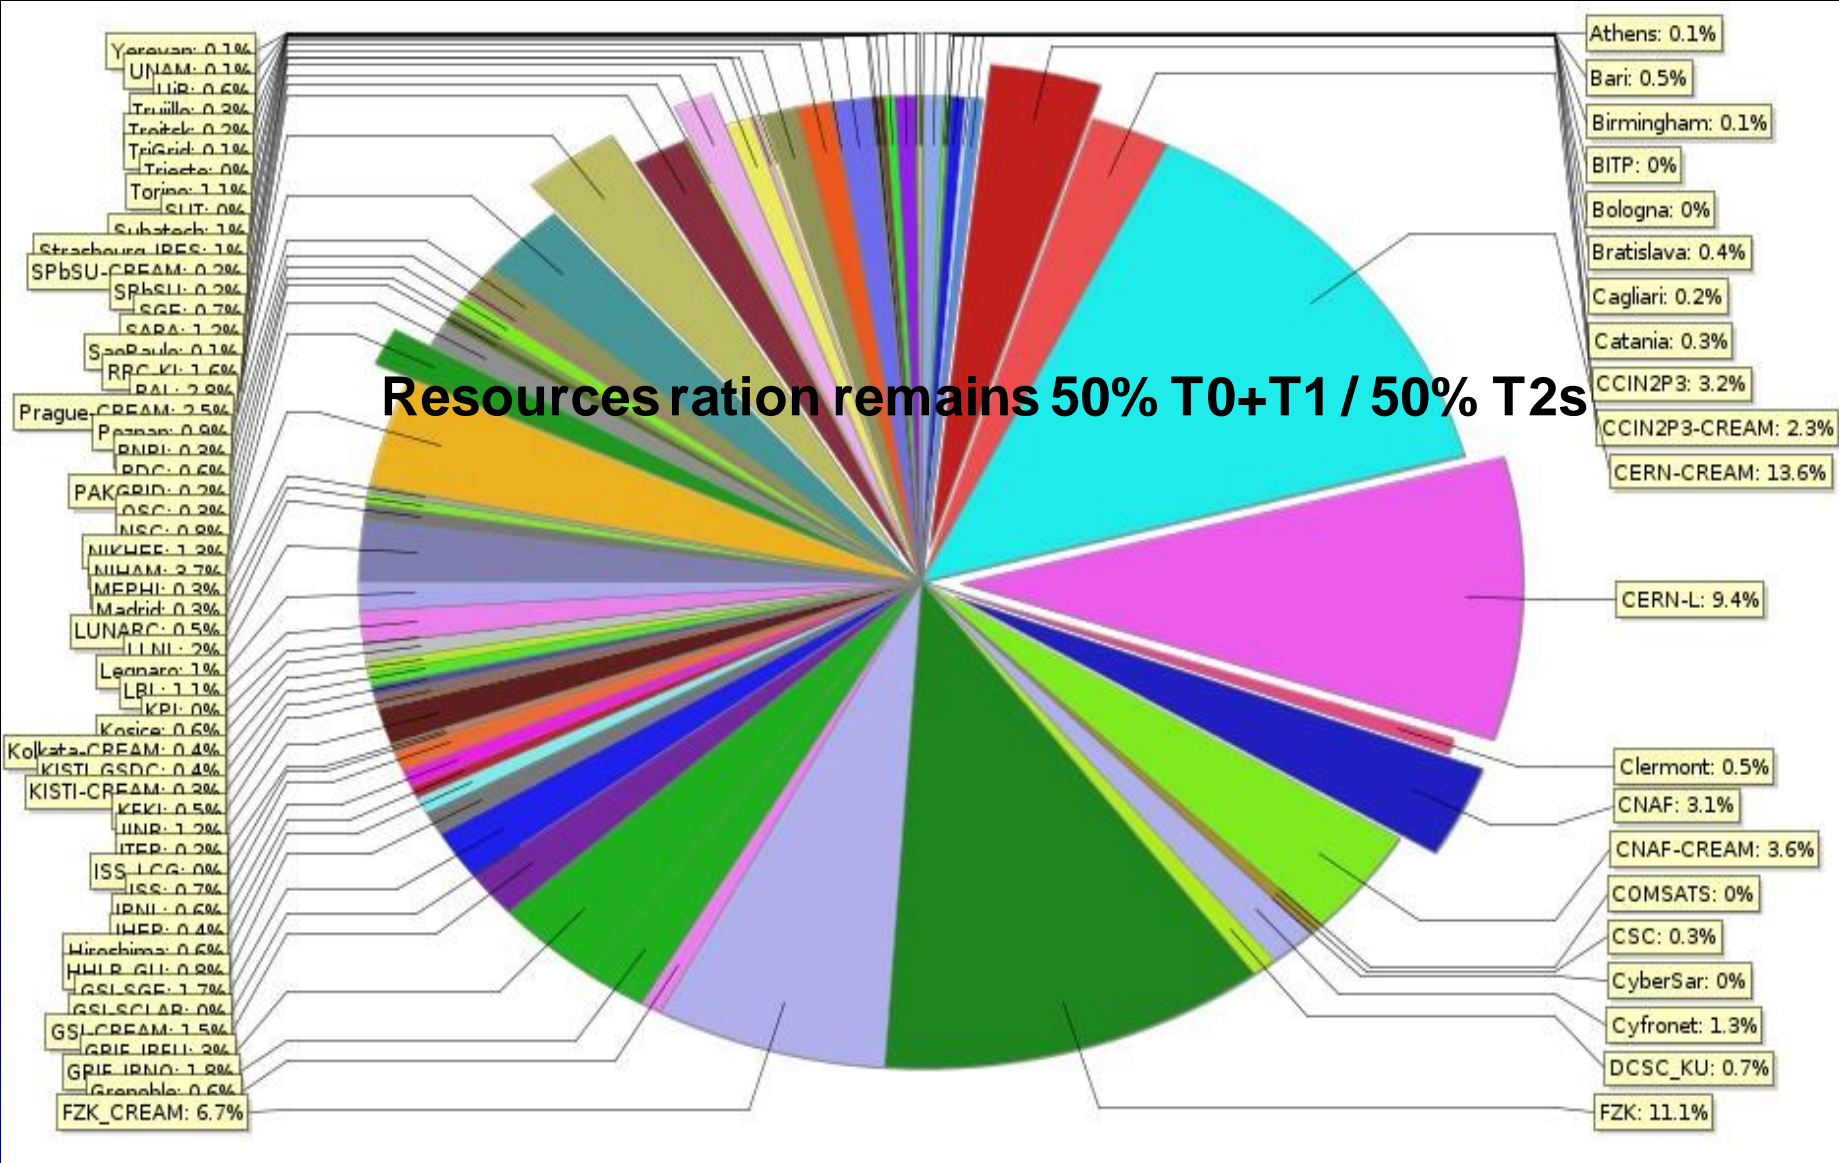
All activity

Running Jobs

21K jobs average over 6 months
69 active sites




- Athens — Bari — Birmingham — BITP — Bologna — Bratislava — Cagliari — Catania — CCIN2P3 — CCIN2P3-CREAM — CERN-CREAM — CERN-L — Clermont — CNAF — CNAF-CREAM — COMSATS
- CSC — CyberSar — Cyfronet — DCSC_KU — FZK — FZK_CREAM — Grenoble — GRIF_IPNO — GRIF_IRFU — GSI-CREAM — GSI-SCLAB — GSI-SGE — HHLR_GU — Hiroshima — IHEP — IPNL — ISS
- ISS_LCG — ITEP — JINR — KFKI — KISTI-CREAM — KISTI_GSDC — Kolkata-CREAM — Kosice — KPI — LBL — Legnaro — LLNL — Loewe_CSC — LUNARC — Madrid — MEPHI — NIHAM — NIKHEF — NSC
- OSC — PAKGRID — PDC — PNPI — Poznan — Prague-CREAM — RAL — RRC-KI — SaoPaulo — SARA — SGE — SPbSU — SPbSU-CREAM — Strasbourg_IRES — Subatech — SUT — Torino — Trieste
- TriGrid — Troitsk — Trujillo — UiB — UNAM — Yerevan — SUM



Grid operation - efficiency

- Efficiency (CPU/Wall) time for ALICE jobs has been 'sliding' over the year



Grid operation – efficiency (2)

- Causes
 - Special Pb+Pb MC cycle(s) with extensive ad-hoc use of RAW OCDB
 - Introduction of possibility to do a 'on the fly' corrections for already processed data (tenders) for ESD->AOD production
 - Large amount of chaotic user analysis (also with tenders), especially in the period April-May 2011 (up to 60% of the total ALICE Grid capacity)

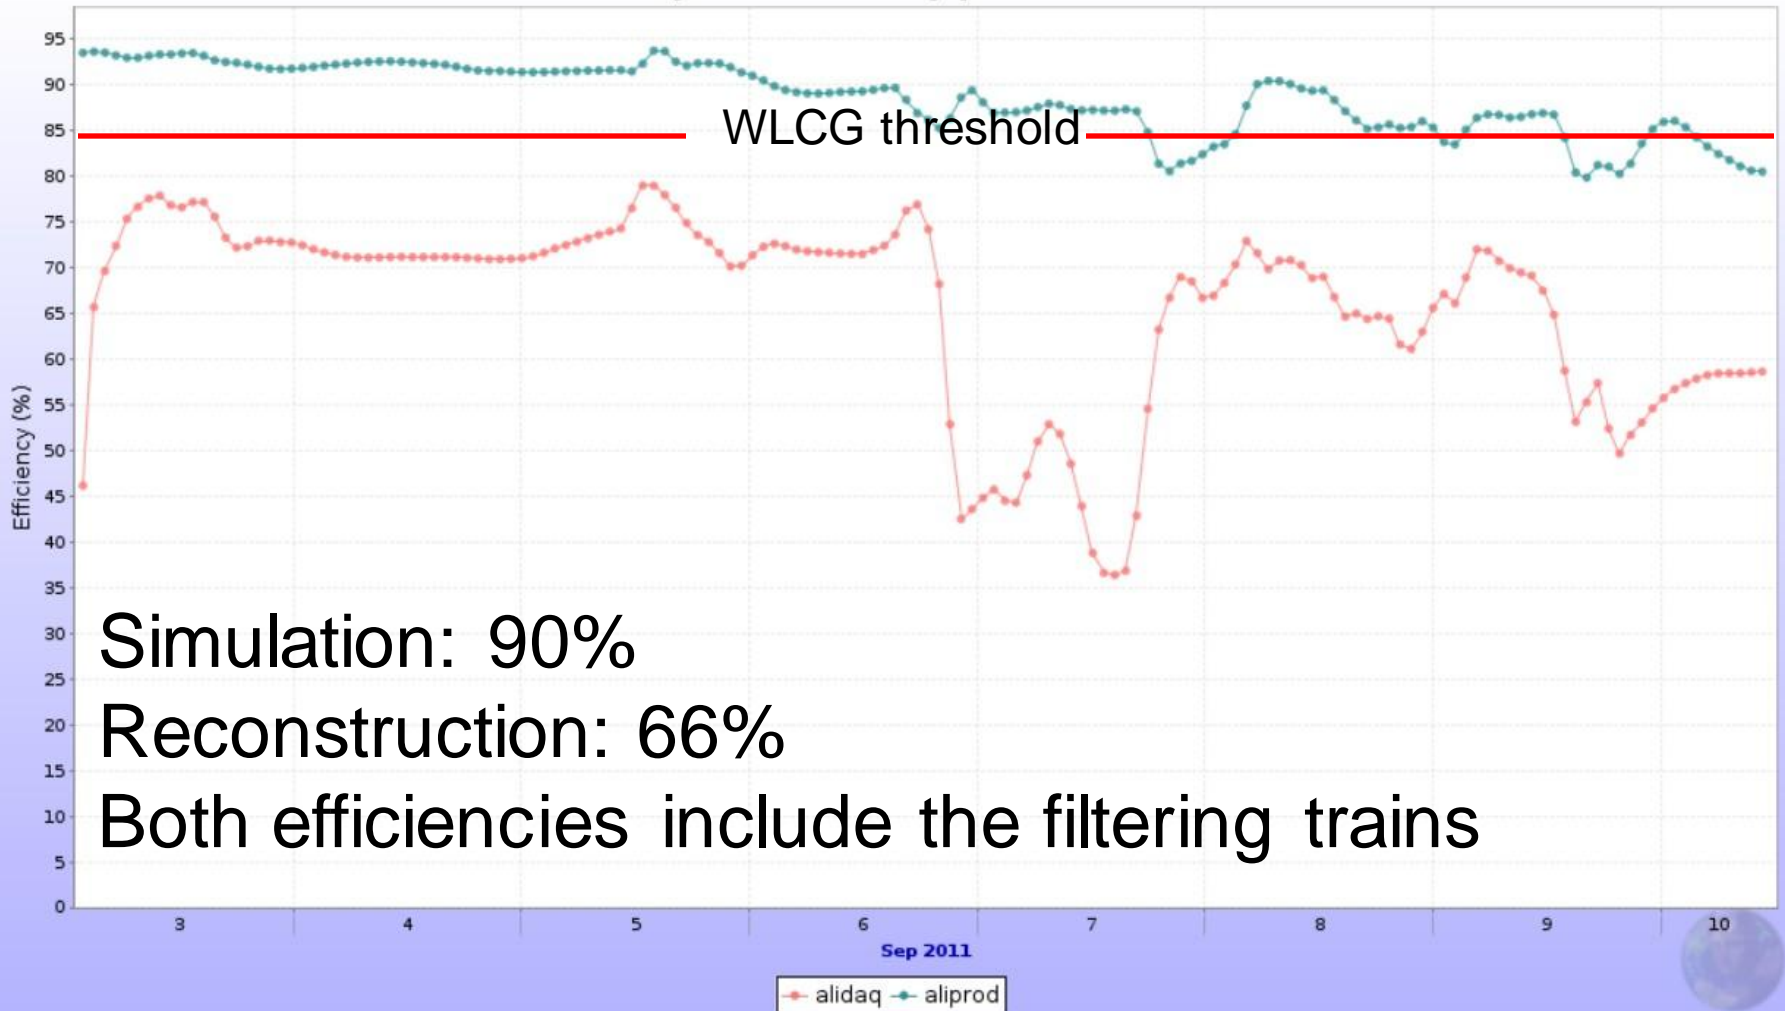
Grid operation – efficiency (3)

- Remedies

- Modification to the MC code to use RAW OCDB in a streamlined regime - single call to the conditions DB/storage - **Advanced**
- Upgrade of the server infrastructure which holds the conditions data (storage) and the metadata (AliEn catalogue) - **Done**
- Introduction of catalogue caches and OCDB snapshots – **Done/Ongoing**
- Modification of the tender code to optimize the access to OCDB, only for the necessary data – **Ongoing**
- Optimization of the data access from the user analysis code – **Ongoing/Consolidation**

Efficiency – production users

Jobs' efficiency per user



Simulation: 90%

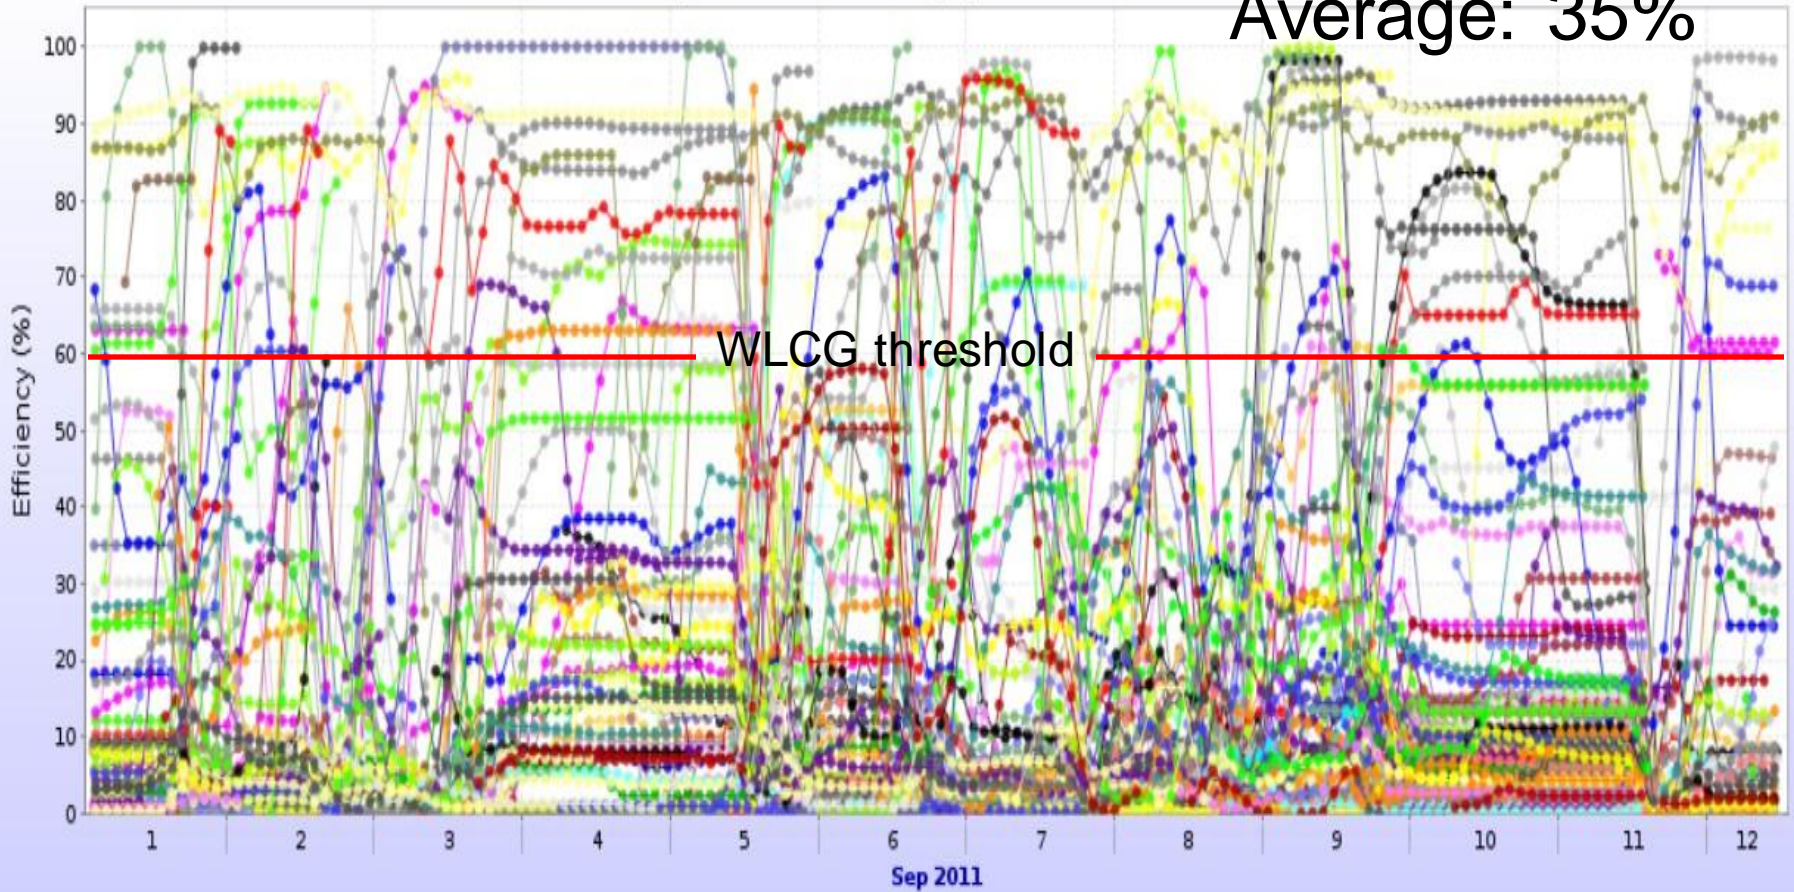
Reconstruction: 66%

Both efficiencies include the filtering trains

Efficiency – chaotic users

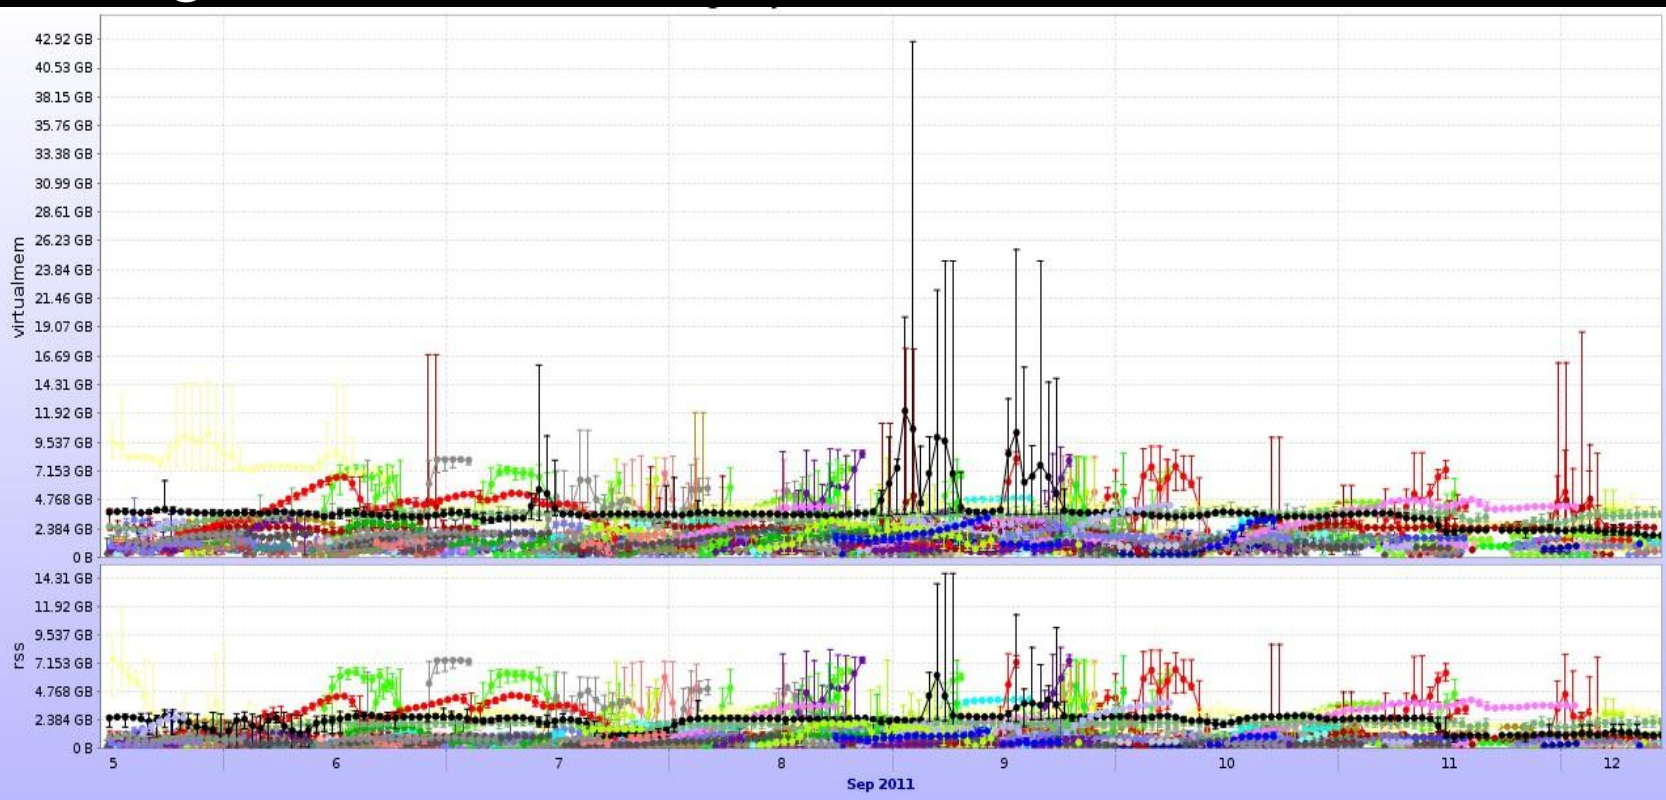
Jobs' efficiency per user

Average: 35%



Memory profiles

- Simulation/reconstruction/analysis trains – under control
- User analysis – good progress, but still quite high and uneven



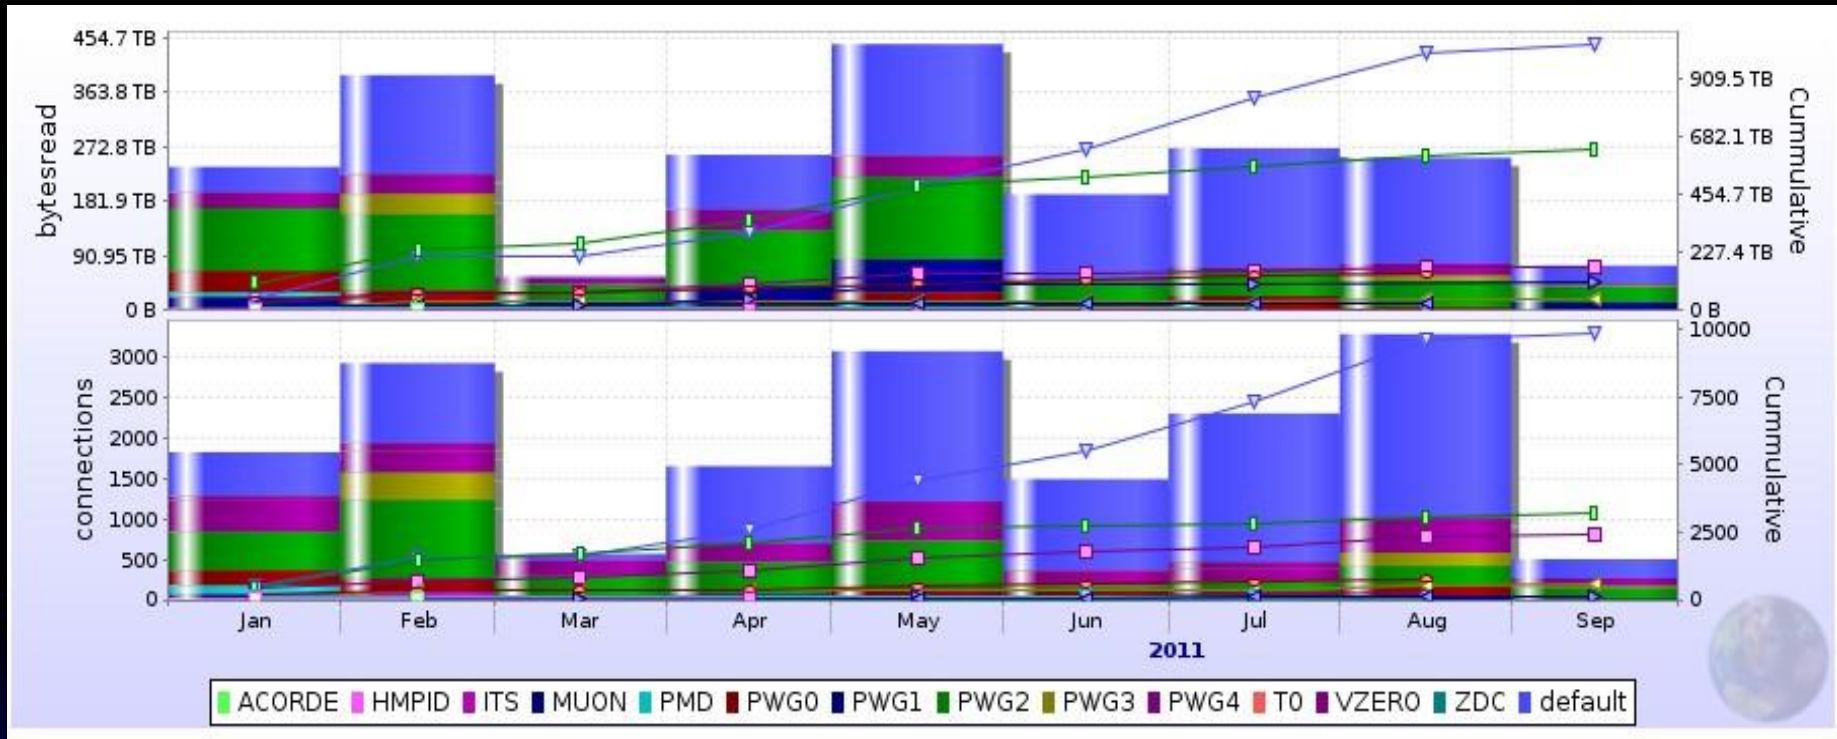
Memory profiles (2)

- Fast memory allocations still cause troubles at sites – kills the WNs
 - Generally available batch system process monitoring tools are not always effective
- Soon in AliEn – memory management and process killer, especially for fast allocations
 - End September

PROOF Analysis facilities

- Two common clusters CAF and SKAF
- KIAF (KISTI, South Korea) in preparation
- Three private clusters (LAF, SAF, JRAF)
- Stable operation throughout the year
- Aggressive ROOT updates – always running with the newest version
 - This was instrumental for the successful operation
- Grid/AF software packages synchronized – no exceptions, important for results reproducibility

Analysis facilities - profile



- 1.8 PB of data through CAF, 550TB through SKAF
- For comparison – on the Grid, we have written 15PB, read 37PB

Sites operation

- Stable throughout the year
- The software and support is mature
 - Any (minor) problems are solved quickly
 - Few updates are in the works – the most notable is the possibility to handle independent clusters at the same site through the same VO-box
- The fabric can handle the most demanding task – user analysis
 - The ‘storage would be saturated’ concern did not materialize
 - The network situation is excellent!
- In general – the tiered structure is blurred, the Grid is effectively a cloud

Short-term plans

- The infrastructure is ready for the Pb+Pb
 - This period will be longer than 2010, more data
- New data calibration strategy introduced in July
 - Shortens considerably the time to Pass2 of RAW data
- Post Pb+Pb – New AliEn version in January
 - Code consolidation
 - New central services and user client (possibly before that)

Summary

- Successful operation since beginning of 2011
 - 2 major physics conferences were 'covered' without too many problems and no showstoppers
- The Grid and AF resources are adequate to serve the RAW/MC and user analysis tasks
 - More resources would be better...
 - We are still 50% 'in the hole' with respect to what is needed
 - Some of the 'savings' from 2010 came in handy (storage)

Summary (2)

- The sites operation is very stable
 - The gLite (now EMI) software is mature and few changes are necessary
- No major perturbation of any activity due to sites updates/unavailability
 - Sufficient redundancy
- Many thanks to the 70 computing centres for the continuous and efficient support of the ALICE operations