ALICE report

Maarten Litmaath
CERN
IT-SDC

WLCG Collaboration Workshop
Barcelona
July 8, 2014

Part 1

- Overview
- Data management
- CVMFS monitoring
- Clouds
- Run 3

RUN 2 physics programme and rates

- Target - integrated luminosity of $1nb^{-1}$ of Pb-Pb collisions (combined RUN 1+RUN 2)
  - Consistent with the ALICE approved programme
  - 4-fold increase in instant luminosity for Pb-Pb
- Double event rate of TPC/TRD
- Increased capacity of HLT system and DAQ
  - Rate up to 8GB/sec to T0

Heavy Ion data taking

RUN 2 detector upgrades

- TPC, TRD readout electronics consolidation
- TRD full azimuthal coverage

(+5 modules)

- +1 PHOS calorimeter

module

- New DCAL calorimeter

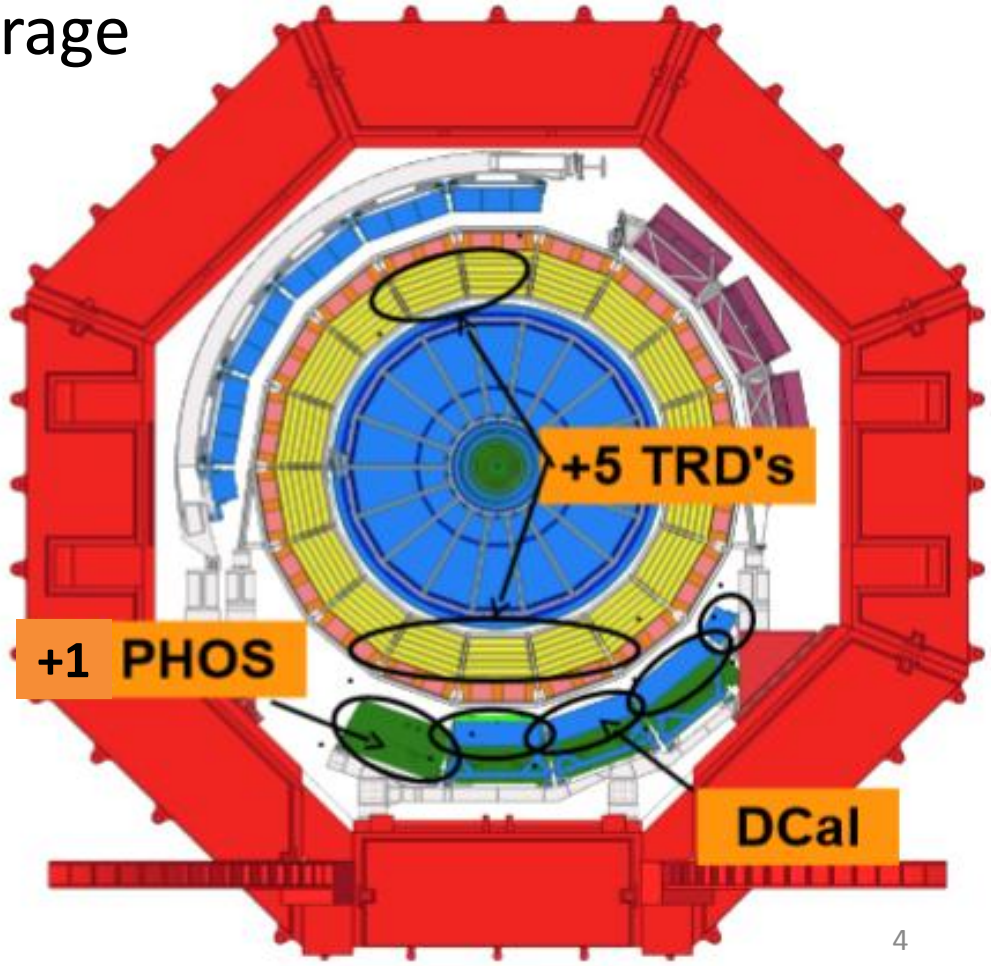

RUN 2 resources considerations

- Same CPU power needed for reconstruction
- 25% larger raw event size
  - Additional detectors
  - Higher track multiplicity with increased beam energy and event pileup
- ALICE requirements for RUN2 were approved by CRSG in April 2014
- The CPU request growth is compatible with 'flat' budget, i.e. depends purely on technology development
- Major demand on resources towards the end of 2015 (Pb-Pb data taking)

New T2 at Universidad Nacional Autónoma de México (UNAM)

Entered ***production*** 28 June

Initial capacity: 1000 cores, 570 TB

Plan:

- Wide area network tuning – 10 Gbit capacity (shared)
- Sign WLCG MoU for T2 by end of this year
- Work toward T1 status (next year)

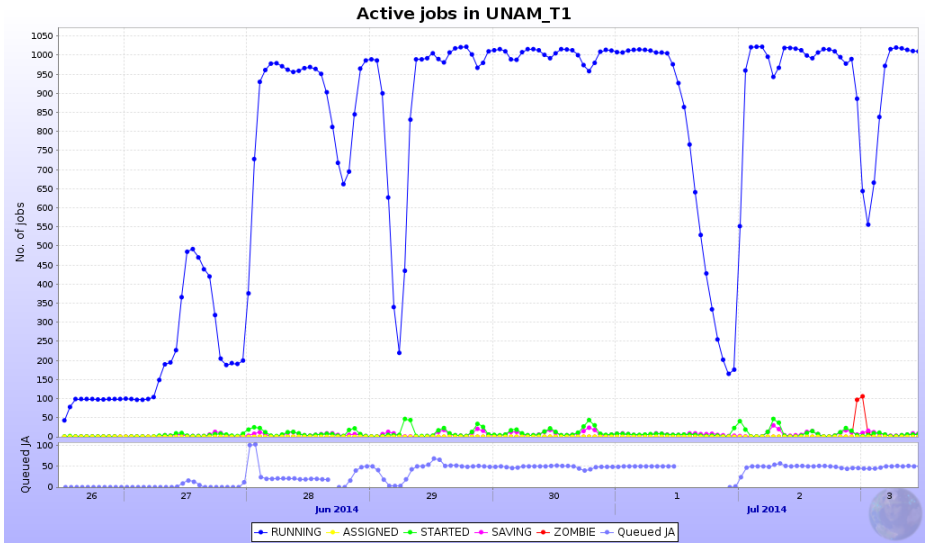

Many thanks to UNAM team!

New proto-T1 at RRC-KI-T1

Entered ***production*** 25 January

Current capacity: 700 slots, 150 TB

Joined OPN with 10 Gbit on Apr 25

Plan:

- Custodial storage
- Work toward T1 status

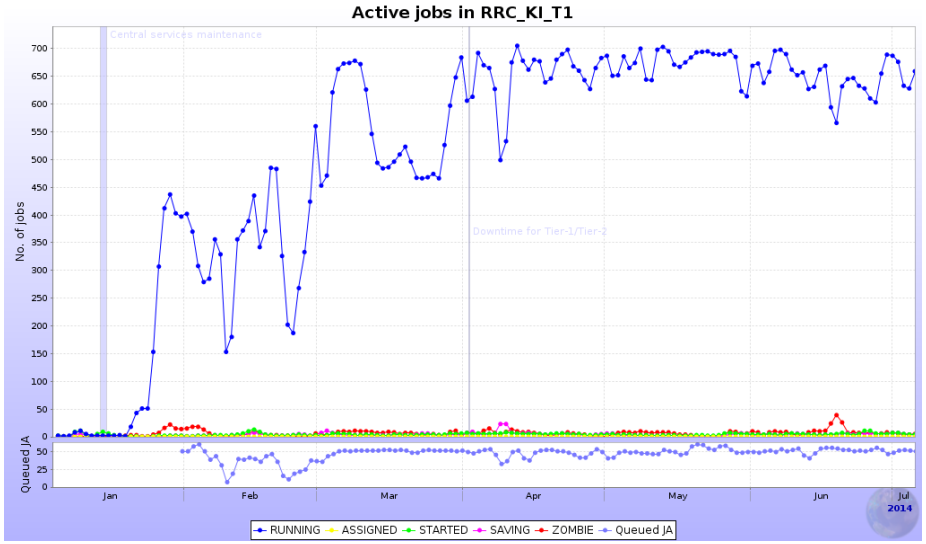

Many thanks to RRC-KI-T1 team!

Software and process improvements

- Moving one calibration iteration to online
  - Software under tests
- Using HLT track seeds for offline reconstruction
  - Comparison of methods ongoing
- Improving performance of GEANT4 simulation for ALICE + development of fast and parameterized simulation
- Collaborating with other experiments to explore contributed resources (i.e. spare CPU cycles on supercomputers)
  - Centres in US, leveraging existing PanDA development
    - New T2 being set up at ORNL, host of Titan supercomputer

Software and process improvements (2)

- HLT farm for offline processing
  - Additional 3% CPU resources
- Improving the performance of organized analysis trains
  - Faster turnaround
- Speeding up and improving the efficiency of the analysis activities by active data management
  - Consolidation of datasets where applicable, introduction of smaller containers (nanoAODs) for specific analysis types

Infrastructure improvements

- Focus on SE stability – major factor for successful analysis and high CPU efficiency

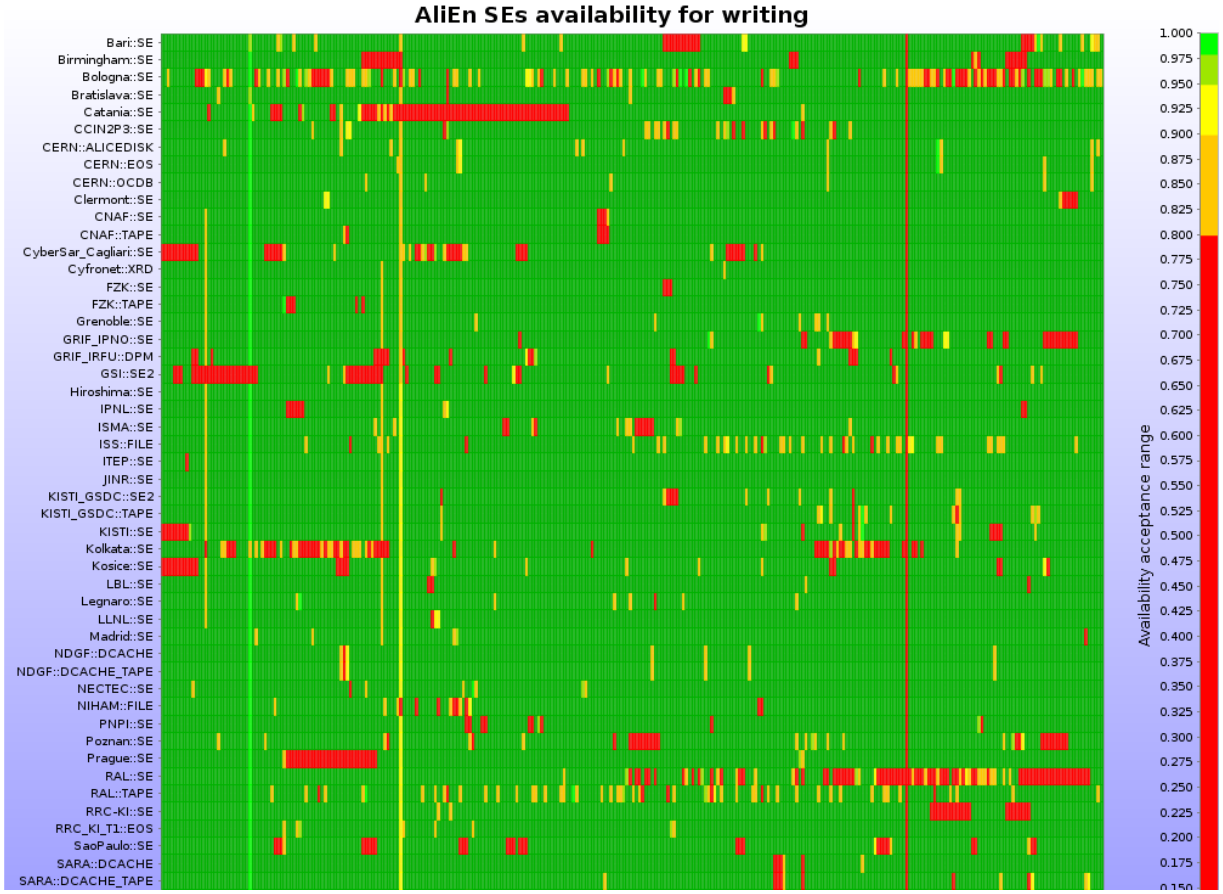

Goal for all SEs >98% availability

See data management section later in this presentation

Infrastructure improvements (2)

- LHCONE programme
  - Europe is largely covered, focus on South America and Asia
  - Larger data volumes, more to transfer between sites
  - Remote access to storage in certain analysis/reconstruction scenarios

Infrastructure improvements (3)

- IPv6 readiness
  - IPv4 address depletion is already a fact for new sites
  - ALICE services are largely IPv6 ready
  - xrootd v.4 should be IPv6 ready (released June 5)
  - Other services are being brought into compliance

Infrastructure improvements (4)

- Refurbishment of SAM/SUM tests
  - WLCG monitoring consolidation project in advanced status
- Site tests will reflect more and more the VO tests
  - In the ALICE case – provided by MonALISA

New critical tests will be introduced very carefully

Operations in 2014

- Steady running: 37K jobs

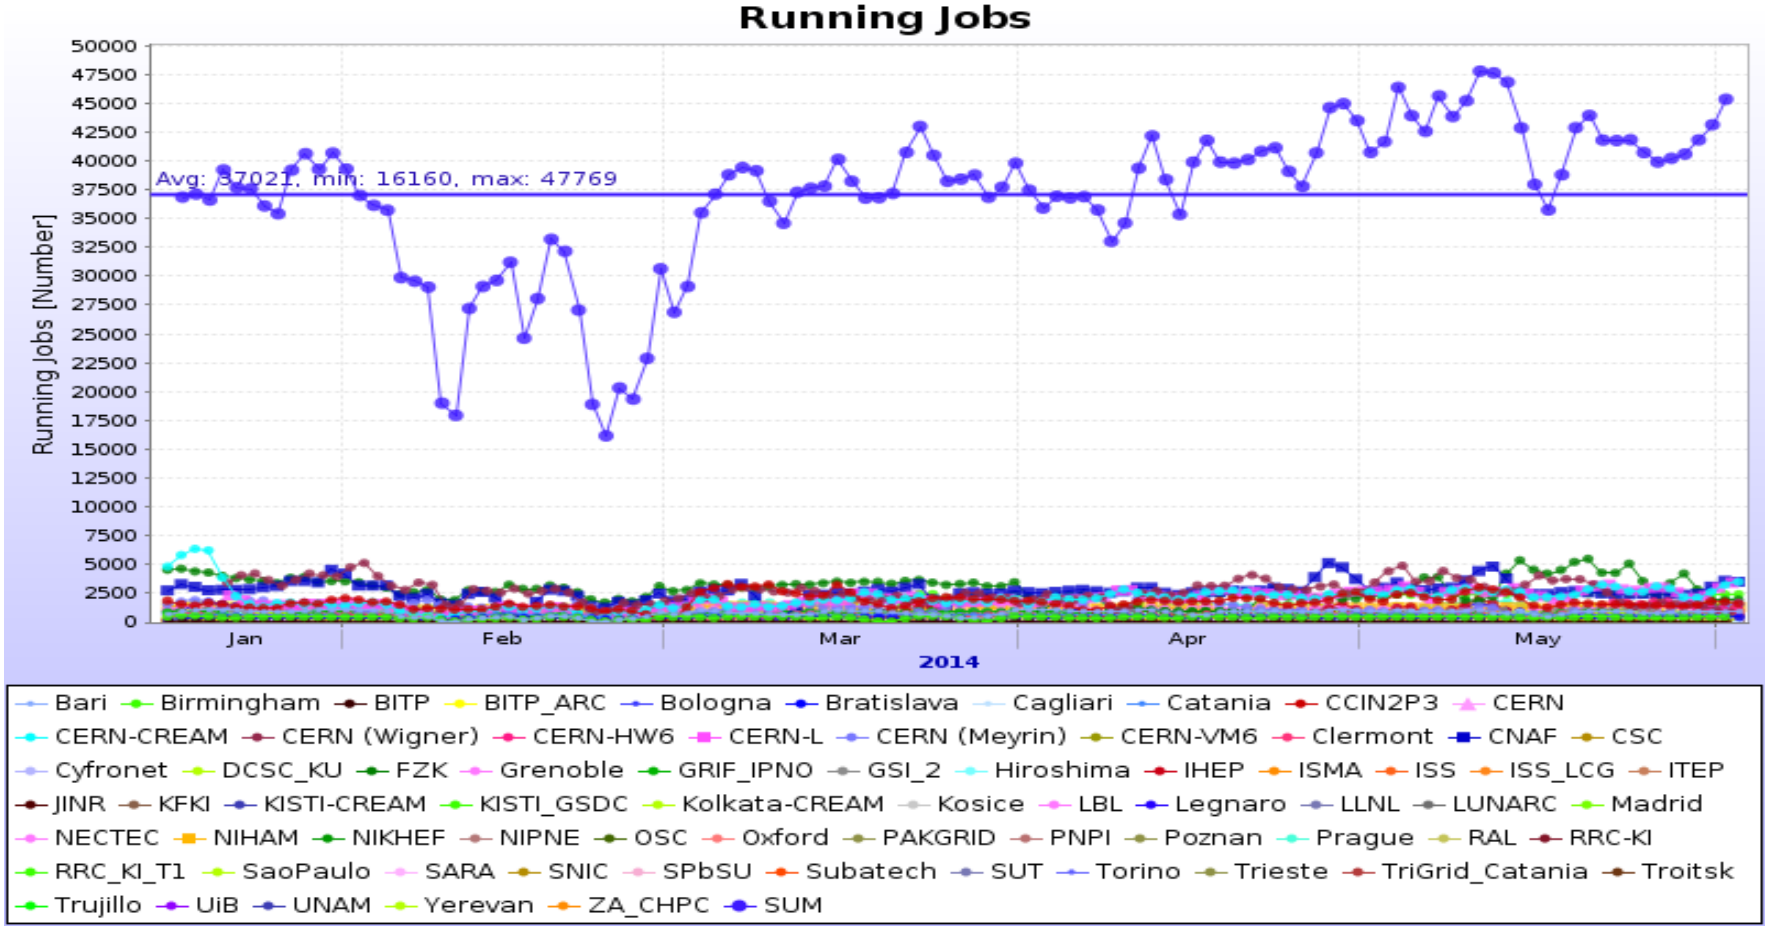

Jobs by user since January 2014

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

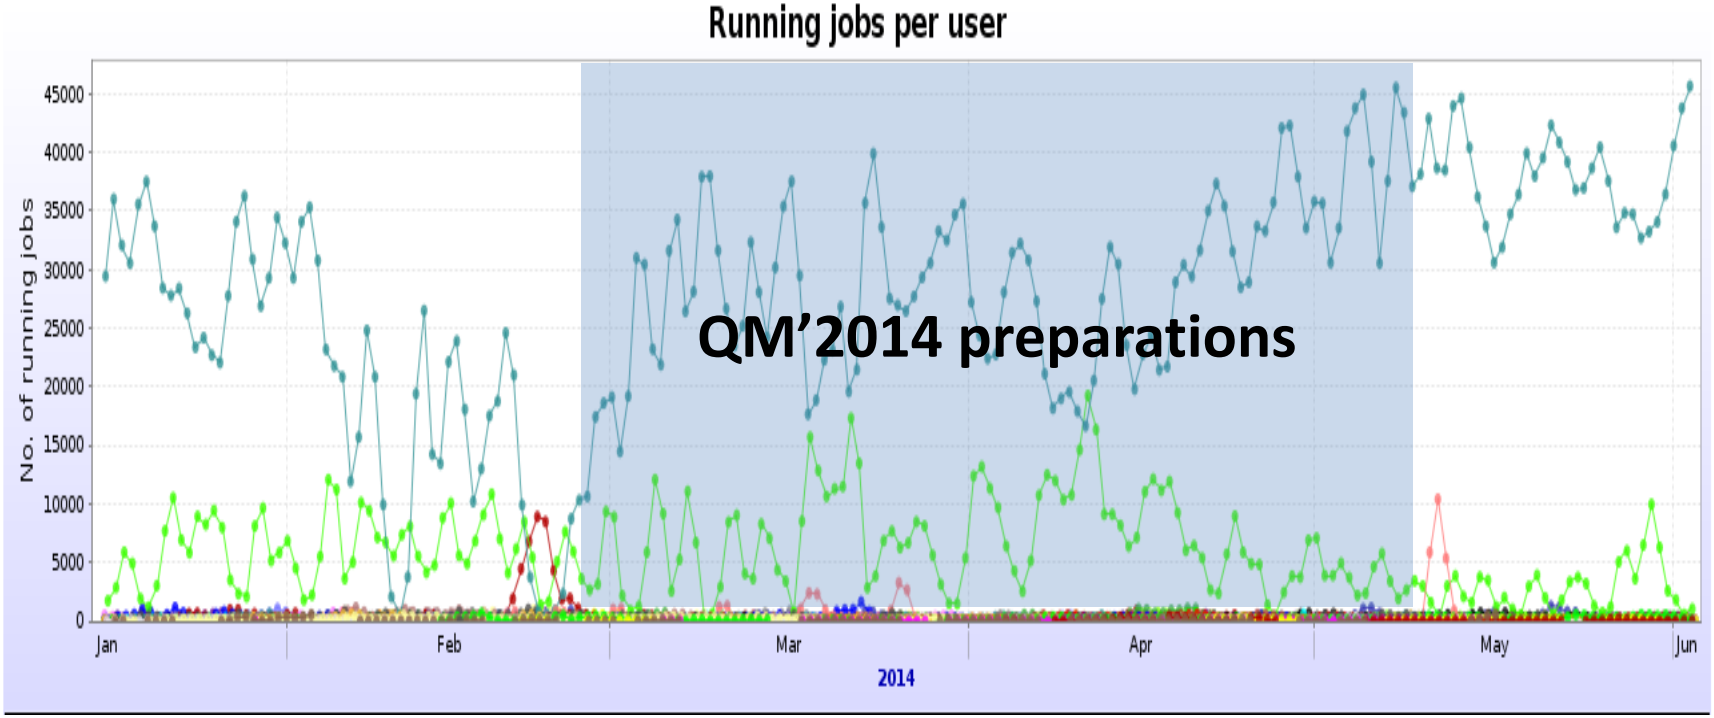

Efficiency since January 2014

- Small effect due to high volume user activities
- A fix in replica access algorithm (bug discovered in April) further increases the overall analysis efficiency
- Continue pushing for larger share of organized analysis (daily software AN tags)

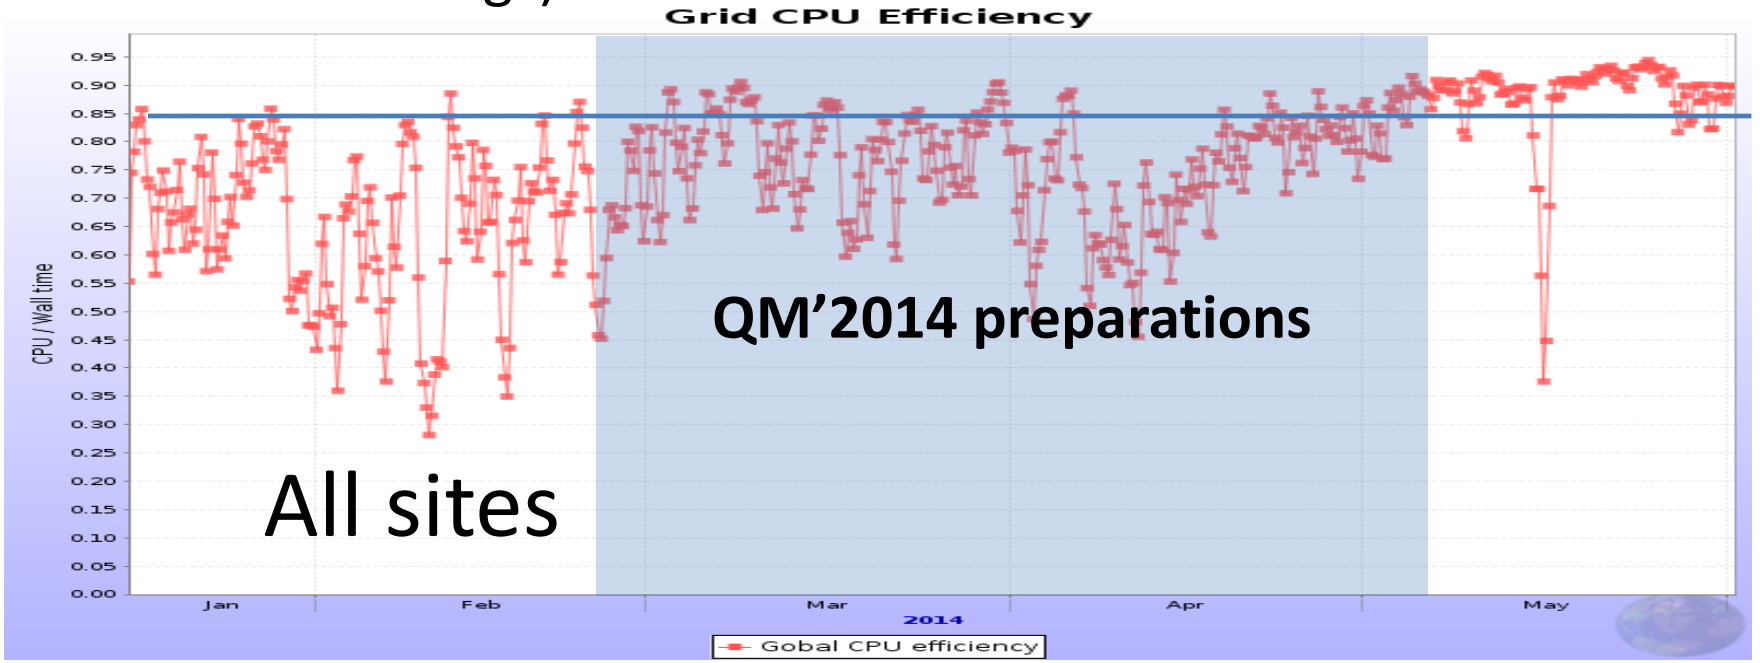

Data volumes since January 2014

- 111.4 PB read, 10.5 PB written
- Regular ‘inactive data sets’ cleanups, popularity service being put in production

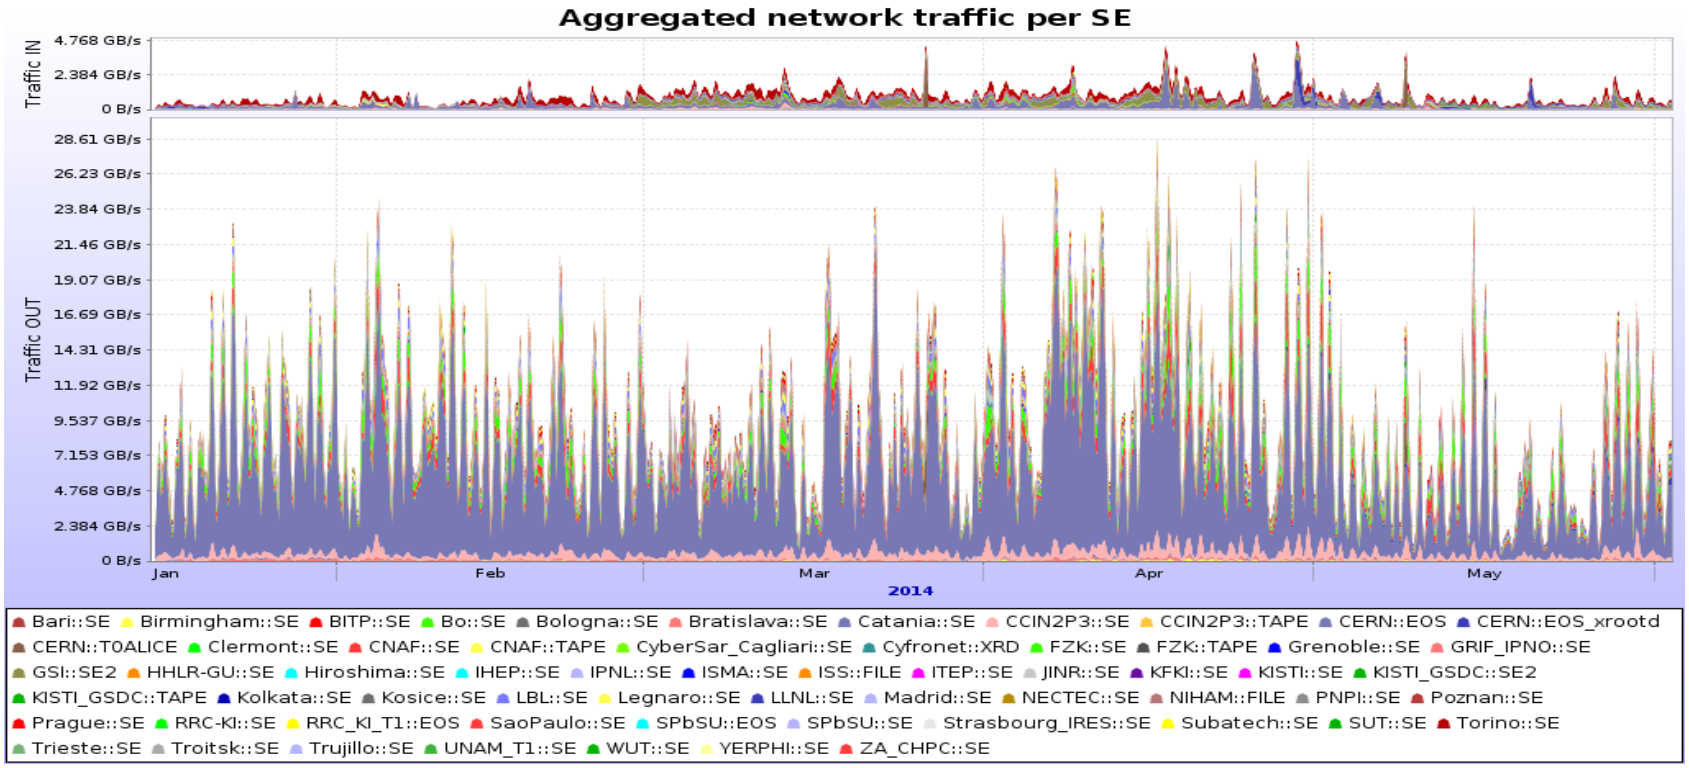

Plans for the next 6 months

- Pass 2/3 of 2011 p-p data and associated MC
  - Full detector recalibration, 2 years of software updates
- Pass 2 of LHC12 p-p, Pass3 of p-Pb data
- From August/September – start cosmics trigger data taking
  - Upgraded detectors readout, Trigger, DAQ, HLT
  - Data will be reconstructed Offline
- No special plans for Grid data challenges
  - All data processing aspects are covered by daily activities

ALICE upgrade

- 13 Computing Working Groups
- Particulars
  - Data flow simulation
  - Data transport model
  - FLP (first level processors)/EPN (event processing nodes) traffic shaping, buffers and system scalability
  - Computing platforms
  - Software framework development (ALFA)
- Computing TDR writing in progress
  - Detailed report: Status of O2 project and TDR

Architecture

O² Architecture and data flow

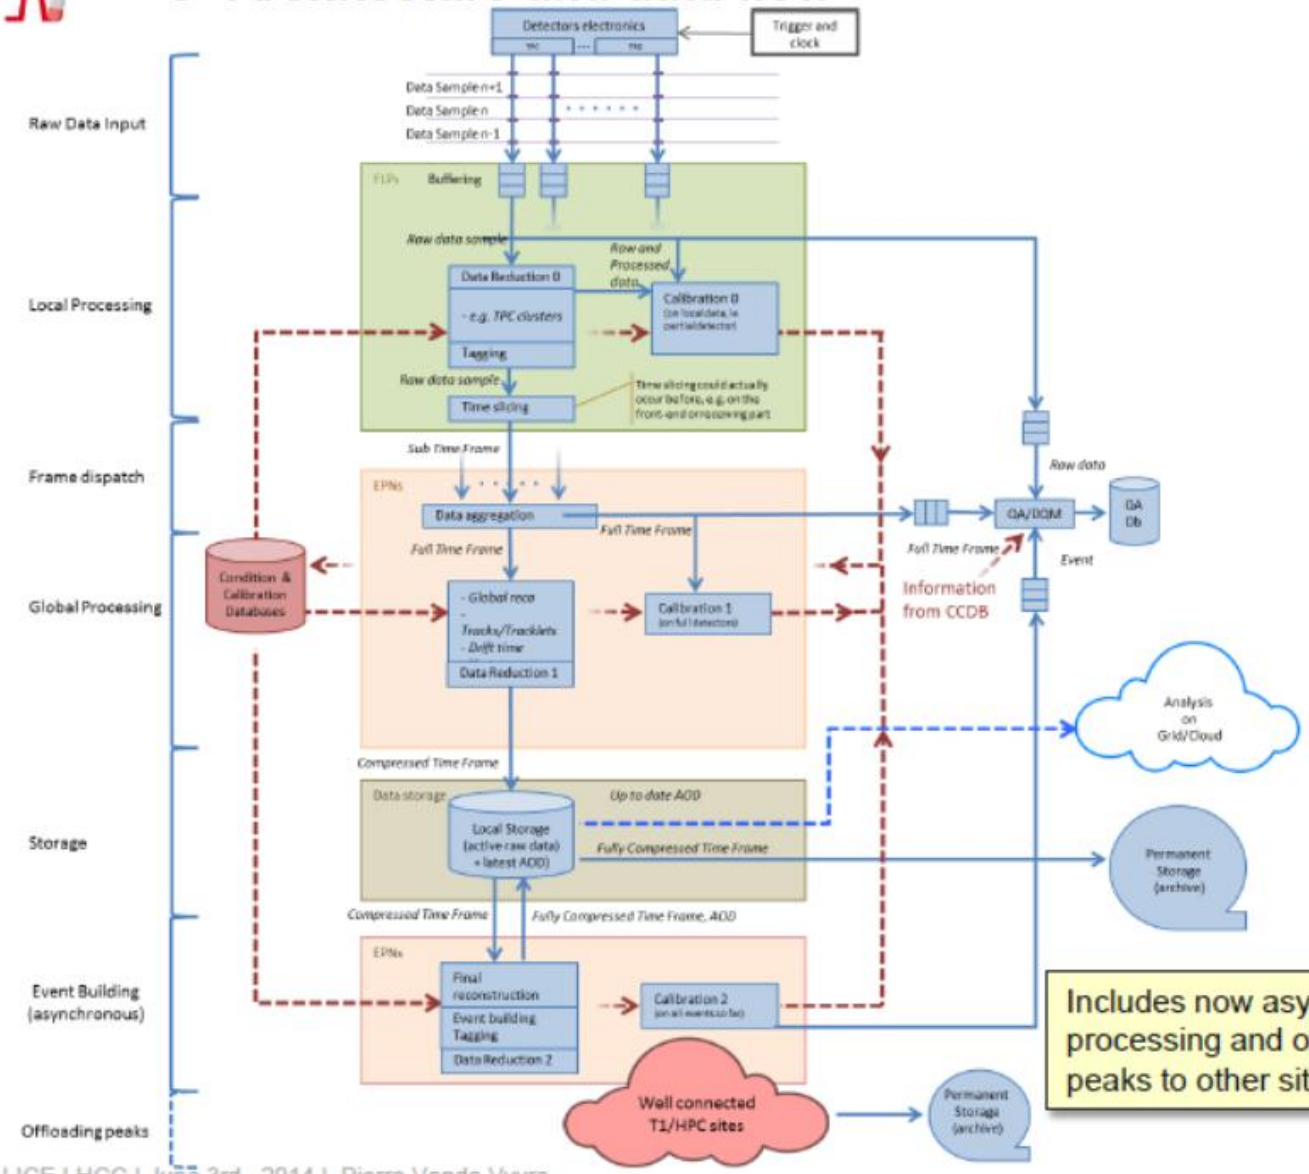

TPC Data
1 TB/s

200 GB/s

100 GB/s

50 GB/s

Includes now asynchronous processing and offloading peaks to other sites.

Summary of part 1

- Steady operations in the past 6 months
  - Emphasis on increasing the share of organized analysis and overall efficiency
- Gradual software and infrastructure upgrade plans leading up to Run2
  - No dramatic changes of computing and operations model
- Resources will be adequate to cover the ALICE physics programme
  - Resources request 2015-2017 endorsed by CRSG
- Preparations for the ALICE upgrade
  - Ongoing work on system design and simulations, software demonstrators and Computing TRD

Part 2

- Overview
- Data management
- CVMFS monitoring
- Clouds
- Run 3

ALICE data access model

- Central catalogue of LFNs and replica locations
- Data files are accessed directly from the storage
- Jobs go to where a copy of the data is
  - Other required files are read remotely (configuration, calibration, executing and validating scripts etc)
  - Urgent tasks (organized analysis) relax the locality constraint to get the job done quickly for the 'tail' (few last percent) of the jobs
- For all requests the client gets a sorted list of replica locations, function of storage availability and its location, i.e. closest (local) first
  - An omission in that code was fixed mid April

Organized analysis trains

- Running many user tasks over the same input data
- Users are strongly encouraged to join the trains instead of running their own tasks
- The most IO-demanding central processing
- The **average** analysis requires **2MB/s/core** to be 100% CPU efficient, but majority of the current infrastructure doesn't support that
  - New CPUs require much more…

Analysis trains activity

- Example week of organized analysis train read volume:
  - Local site storage: 1.56 PB
  - Remote storage: 64 TB
  - 4% remote data access (failover, lifted locality restrictions)
- Read throughput:
  - Local site storage avg: 1.35 MB/s
  - Remote site storage avg: 0.73 MB/s
  - Reading remotely introduces a large penalty!

SE monitoring

- xrootd and EOS data servers publish two monitoring streams
  - ApMon daemon reporting the data server host monitoring and external xrootd parameters
    - Node total traffic, load, IO, sockets, disk IO, memory …
    - Version, total and used space
  - xrootd monitoring configured as:
    - `xrootd.monitor all flush 60s window 30s dest files info user MONALISA_HOST:9930`
    - Client IP, read and written bytes

Infrastructure monitoring

- On each site VoBox a MonALISA service collects
  - Local SE monitoring data (network interface activity, load, sockets, client access stats etc)
  - Job resource consumption, WN host monitoring …
- Traffic data is aggregated in client IPv4 C-class, LAN/WAN, client site, server site
- ML services perform VoBox to VoBox measurements
  - *traceroute / tracepath*
  - 1 stream available bandwidth measurements (FDT)
    - This is what impacts the job efficiency
- All results are archived and we also infer the network topology and utilization from them

SE functional tests

- Performed centrally every 2h, targeting the declared redirector
  - *add/get/rm* suite using the entire AliEn stack
  - Or just *get* if the storage is full
- The dynamically discovered xrootd data servers are tested individually, with a simplified suite
- Monitor discrepancies between declared volume and total space currently seen by the redirector
- Site admins prompted to solve the above issues
  - And many other related tests, like insufficiently large TCP buffer sizes

Replica discovery mechanism

- Closest working SEs are used for both reading and writing
  - Sorting the SEs by the network distance to the client making the request
    - Combining network topology data with the geographical location
  - Leaving as last resort the SEs that fail the respective functional test
  - Weighted with their free space and recent reliability
- Writing is slightly randomized for more 'democratic' data distribution

Weight factors

- Free space modifies the distance with
  - f (ln(free space / 5TB))
- Recent history of *add*, resp. *get* contribute with
  - 75% * last day success ratio +
  - 25% * last week success ratio
- The result is a uniform federation with a fully automatic data placement procedure based on monitoring data

Plans

- In the near future ALICE will upgrade to xrootd 4.0 centrally
  - AliEn, using the xrootd the command line
  - ROOT, using xrootd as library
  - Eventually replacing *xrd3cp* with the new client that implements the same functionality
- Implement IPv6 network topology discovery and use it for SE discovery
  - We have already started getting requests on IPv6
- Retry using async IO in ROOT with the new releases

Site plans

- Long overdue xrootd upgrade
  - A few sites still run rather old versions
- Will ask all to upgrade to 4.0 as soon as it is stable
- For newly deployed storage we recommend EOS
  - Five sites already using it
  - Site admins have provided recipes, being tuned further
- Work closer with the sites to identify IO bottlenecks and solve them
  - Keeping in mind the target of 2MB/s/core

Part 3

- Overview
- Data management
- CVMFS monitoring
- Clouds
- Run 3

Motivation

- CVMFS is now a critical services for all VOs
- Currently missing information about the performance of the Stratum 0/1 and the local site proxies
  - Some bits of information in various places, like the availability of Stratum 0 services, awstats …
  - Not enough to assess whether the services performance is OK
- Some sites are alerted for failures by the users (tasks failing)

Proposed solution

- As presented to the WLCG Monitoring Consolidation WG on 2014/06/06
- Deploy a monitoring service on each server of the infrastructure
  - Full host monitoring (CPU, memory, network IO, disk IO performance, sockets and processes in each state)
  - CVMFS and Squid-specific probes
- Real time access to the parameters, plus
  - Alarms, history of all parameters, simple display options
  - Trivial to integrate in dashboard

Prototype

- http://cvmfsmon.cern.ch/
- Very simple package to deploy on the Squid proxy server
  - Either as RPM or as a script to run locally
- Based on MonALISA monitoring service
  - All host monitoring modules enabled
  - New Squid module querying the status 1/minute
  - Network topology discovery enabled for the group
- Installed on 35 servers already (6 T1 sites)
  - Many thanks for the help of Subatech, BITP and ISS site admins during the initial iterations and to all who have deployed and provided feedback!

Part 4

- Overview
- Data management
- CVMFS monitoring
- Clouds
- Run 3

CernVM Elastic Clusters

What is an Elastic Cluster?

- A cluster of CernVM virtual machines: one head node, many workers
- Running the HTCondor job scheduler
- Capable of growing and shrinking based on the load with elastiq

- Configured via a web interface: cernvm-online.cern.ch
- Entire cluster launched with a single command
- User interacts only by submitting jobs
- No external tools: embedded elasticity, ideal for opportunistic clouds

Release Validation cluster

Fully disposable Elastic Cluster for running the Release Validation

- No need to carry the software with the VMs
  - AliRoot versions to validate on CernVM-FS
- The cluster *(incl. the head node)* can be thrown away after use
  - Worker VMs automatically wiped out when validation completes
  - Output and log files stored on shared storage (EOS)
- Procedure fully repeatable
  - Cluster can be rebuilt using a configuration file
  - The very same environment can be restored as it was

A private cloud on the HLT farm

- HLT farm is a delicate real-time environment
  - Opportunistic exploitation can by no means interfere with standard HLT operations
- Hard separation of HLT environment and the opportunistic one
  - Best isolation technique: configure HLT nodes as a private cloud
- We start working on the current “devel” farm
  - Configuration will be moved to the forthcoming “production” farm
  - We are considering OpenStack → popular, lots of support
- Ideal type of opportunistic jobs: CPU-intensive → Monte Carlos
    - I/O uplink and gateway might be a bottleneck on HLT

Preliminary milestones

- August:
  - base OpenStack services configured on the devel cluster
  - network isolation operational
- September:
  - test the devel configuration on the production cluster
  - network hardware configured for traffic shaping
  - configure the special AliEn VOBOX
- October:
  - ready for running special AliEn jobs

Part 5

- Overview
- Data management
- CVMFS monitoring
- Clouds
- Run 3

Towards Run 3

- Detectors and running scenario:
  - ALICE upgrade aiming to high statistics sample (10 nb^{-1})
  - continuous readout TPC, upgraded ITS
  - 50 kHz PbPb interaction rate (current rate x100)
  - ~1.1 TB /s detector readout
- Data reduction strategy
  - Up to factor 20 by doing online reconstruction and compression
  - Store only the reconstruction results, discard raw data
    - demonstrated with TPC cluster finder running on HLT since PbPb 2011
  - Requires ~250k CPU cores
- Overall ALICE expects the WLCG Grid to grow by a factor 2.5 by 2019 and provide ~125k jobs slots to ALICE
  - Based on 20% (CPU) and 15% (Disk) yearly growth

O2 Facility and computing model for Run 3

- ALICE Online-Offline (O2) Facility at P2 will have to provide the remaining CPU capacity (at least 50% of total CPU budget) as well as a large disk buffer to allow all data to be stored after the first stage reconstruction and while waiting for asynchronous second stage calibration and reconstruction
  - Large disk buffer (25 PB)
- Custodial (archive) storage at T0 and T1s
  - Subsequent re-reconstruction passes at O2 facility, T0/T1
- Simulation and associated reconstruction will run mostly at T2s

New software framework for Run 3

- Should work in Offline and Online environment
- Based on new technologies
  - Root 6.x, C++11
- Optimized for I/O
  - New data model
- Capable of utilizing hardware accelerators
  - FPGA, GPU, MIC …

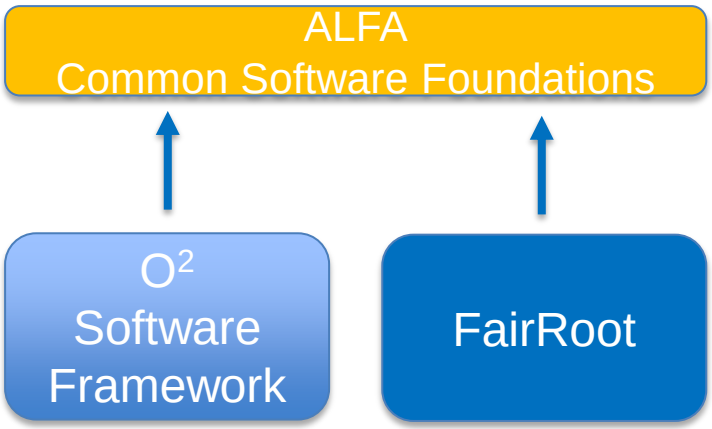

- AliRoot 6.x, our new software framework will be based on ALFA - common software foundation jointly developed between ALICE & GSI/FAIR
  - Support for concurrency in an heterogeneous and distributed environment
  - Based on 0MQ messaging framework

Reducing complexity is the key

for illustration only

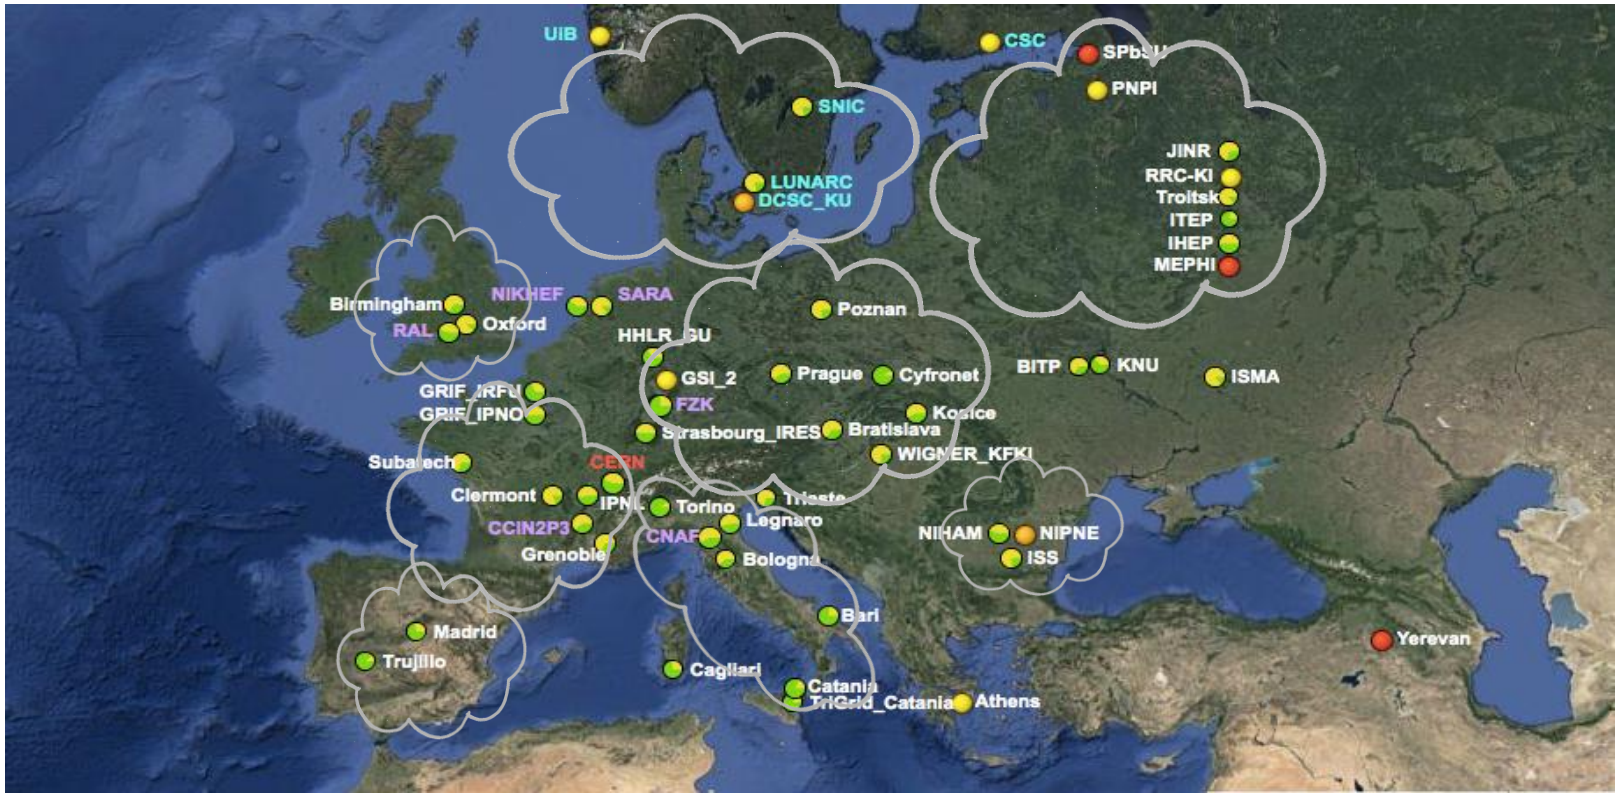

- Virtually joining together the sites based on proximity (latency) and network capacity into Regional Data Clouds
- Each cloud/region provides reliable data management and sufficient processing capability
  - Dealing with handful of clouds/regions instead of the individual sites

Time is short

- There is no time to reinvent the wheel
  - Run2 is about to start and Run3 is only a few years away
  - Data is the product of every experiment
  - Data management is complex and it takes a very long time to iron out all bugs and have confidence in data management software
- We must start from something we already trust
  - EOS is the product that most closely matches our needs and expectations

Why EOS?

- EOS is already tested to the scale required to manage ALICE internal disk buffer@P2
  - some extras might be needed
- Provides data access via
  - high performance private xroot protocol
  - standard HTTP protocol
- On the Grid side, we would need scalable global name space
  - Replacement for a global file catalog
- Federation of Regional Data Clouds