Computing at CMS - Operations

PHYSICS AT THE TERA SCALE
Helmholtz Alliance

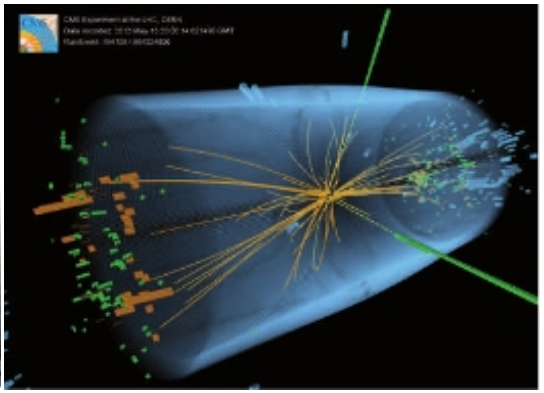

9th Annual Terascale Workshop
18th November 2015

Christoph Wissing (DESY)

HELMHOLTZ | ASSOCIATION

DESY

Introduction: CMS Detector on one Page

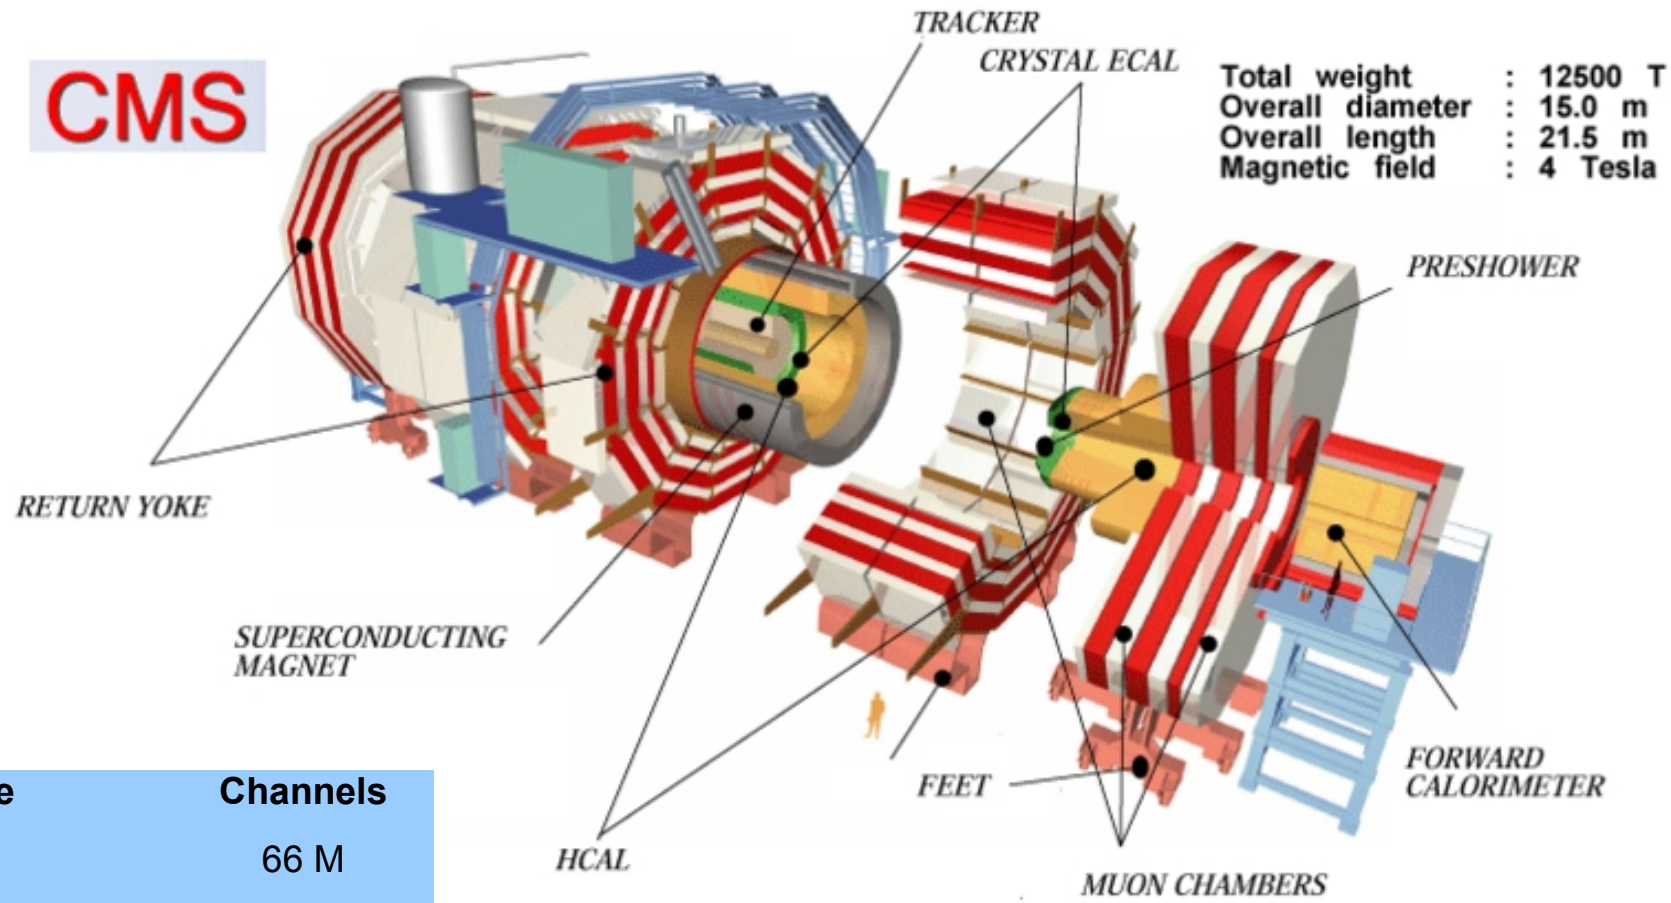

Name	Channels
Pixel	66 M
SiStrip	9.3 M
Muons	890 k
Preshower	140 k
ECAL	76 k
HCAL	9 k

Data from: 2008 JINST 3 S08004, Table 9.1

RAW data*
0.8 MB/event
1.5 MB/ev (MC)

→ RECO data*
2 MB/event

→ AOD data
400 kB/event

→ miniAOD data
50 kB/event

2 stage Trigger
~1000 ev/s logging &
PromptRECO

* Events sizes depent on pile-up

Distributed Computing Infrastructure

More than 50 CMS centers, in more than 20 countries

Flags taken from Wikipedia:
http://de.wikipedia.org/wiki/Liste_der_Nationalflaggen

During Run1:
Rather strict coupling of workflow types to tiers

- Tier 0
  - Main task in Run1:
    - Prompt reconstruction
    - Store RAW data and export to T1s
- Tier 1
  - Main tasks in Run1:
    - Re-recontruction & MC production
    - Long term storage of RAW and MC files
- Tier 2
  - Main tasks in Run1:
    - MC production
    - User analysis
- Tier 3
  - Main tasks in Run1:
    - User analysis

Evolution toward more Flexibility

- More Dynamic Data Management
  - Aim: Optimized usage of disk space at Tier-1 and Tier-2 sites
  - Automatic creation of replicas for popular data, deletion of unpopular replicas
- Disk-Tap Separation
  - Decouple processing and archiving location
  - Control content of Tier-1 disk more efficiently
- Data Federation Any data Anytime Anywhere (AAA)
  - Relaxation of constraints on datasets location and workflow execution
  - Improve capabilities of CMS applications for remote data reading
- Global Pool via GlideinWMS
  - Central control of job priorities, simplified infrastructure
- Ability to provision Cloud resources via GlideinWMS
  - Integrate resources like HLT, academic or commercial Clouds

Dynamic Data Management at Tier-2 Sites

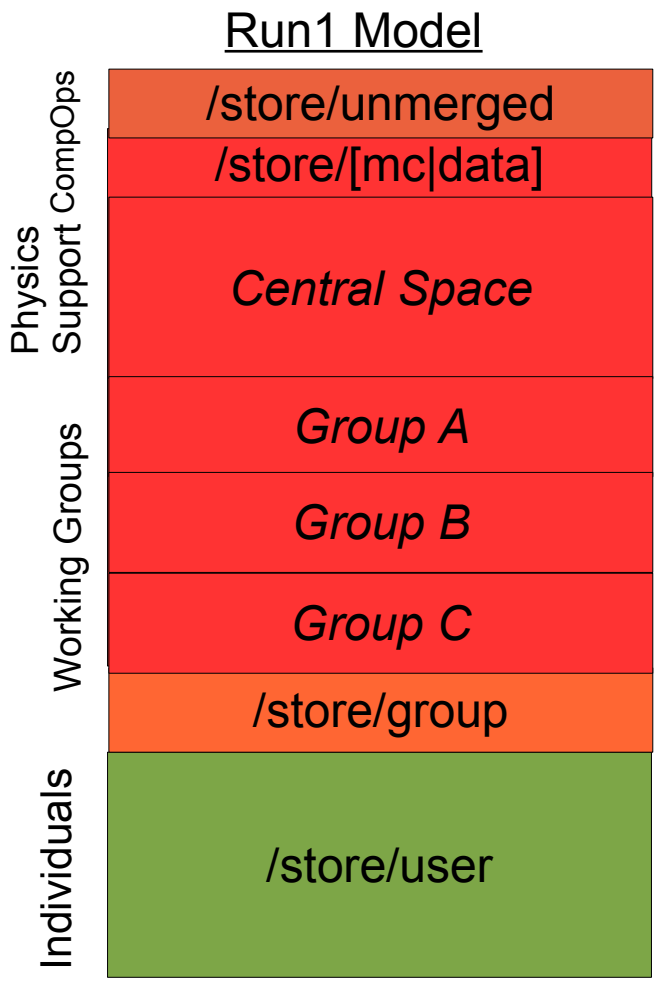

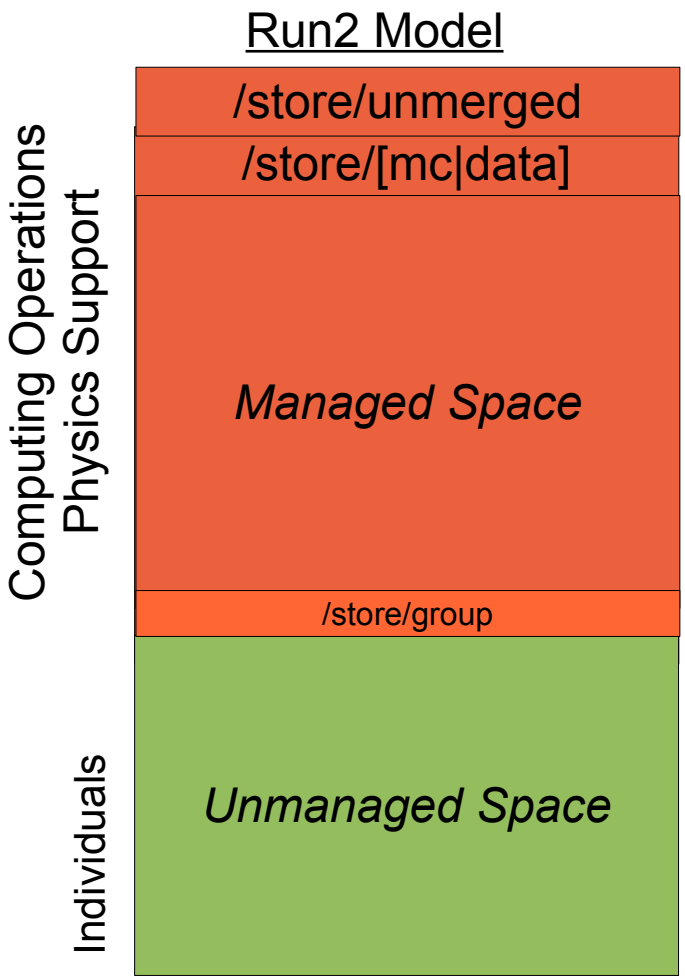

- All spaces formally managed by groups transferred into centrally managed space
- Central space is ~60% of pledged disk space
- Least popular data gets cleared when site reaches a quota value
- Create replicas for often accessed data
- Overall system is deployed including all Tier-1 sites

Disk Tape Separation

Run1: Disk-Tape coupling at Tier-1

- Uncontrolled access to files on tape forbids user analysis at Tier-1
- Produced files go to tape “immediately”
  - Cannot use site without writing to tape → inefficient and inflexible

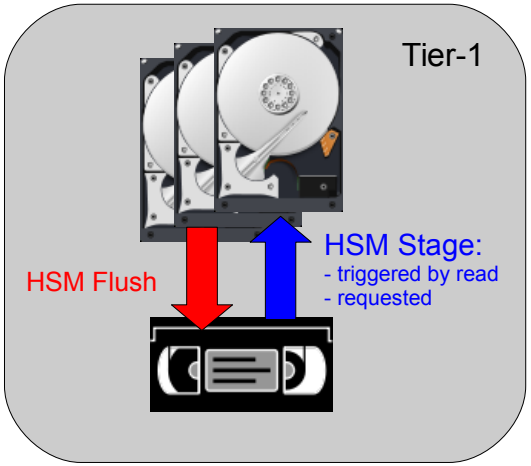

Run2: Separate disk and tape

- Large disk pool and small tape read/write pool
- Staging and flushing from/to tape is a <u>subscription</u> in CMS data management tool
- Files on disk can be read via WAN xrootd (never trigger tape operations)
- Tier-1s can be open for 'chaotic' analysis jobs
- Production and (tape) archiving location now independent
- Completed at all Tier-1 sites

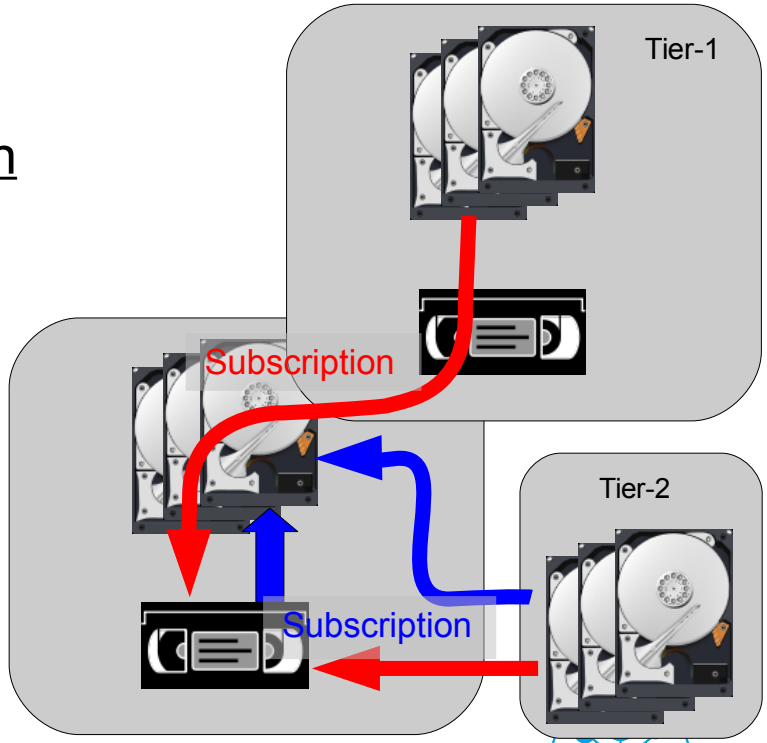

Multi-core Processing

- Need for multi-core jobs
  - With Run2 conditions RECO for one “Lumi Section” will not fit into usual 48-hour batch queue
    - Cannot split easily beyond Lumi section boundaries
  - Large potential to save memory pro CPU core
  - Less pilots to be sent – should increase scaling capability of infrastructure
- Efficiency of multi-core applications
  - Big fraction of code needs to be thread-safe [Amdahl's law]
  - Achieved good CPU efficiency in recent software releases
- Mixture of multi-core & single-core jobs
  - Partitionable pilots
    - 8 core pilot can execute e.g. 1x4+1x2+2x1 core-jobs
  - Ongoing discussion about pilot utilization
    - There are inefficiencies towards the end of the lifetime of a pilot

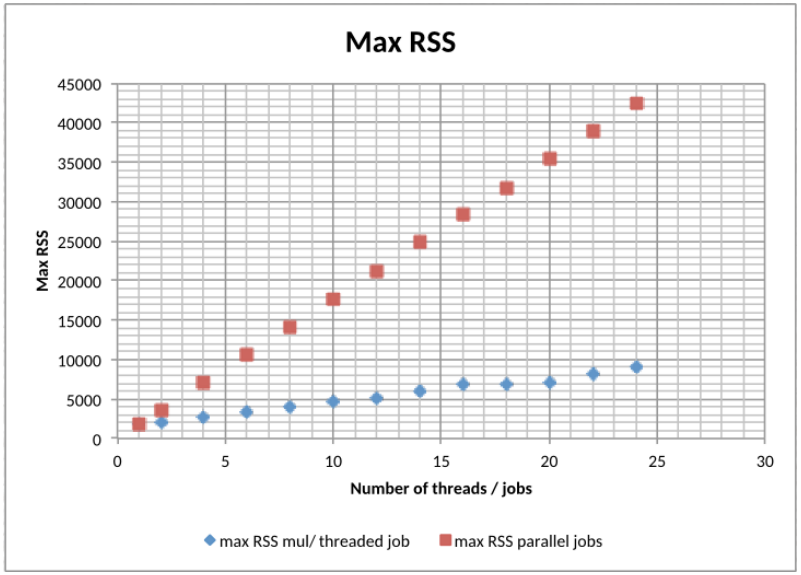

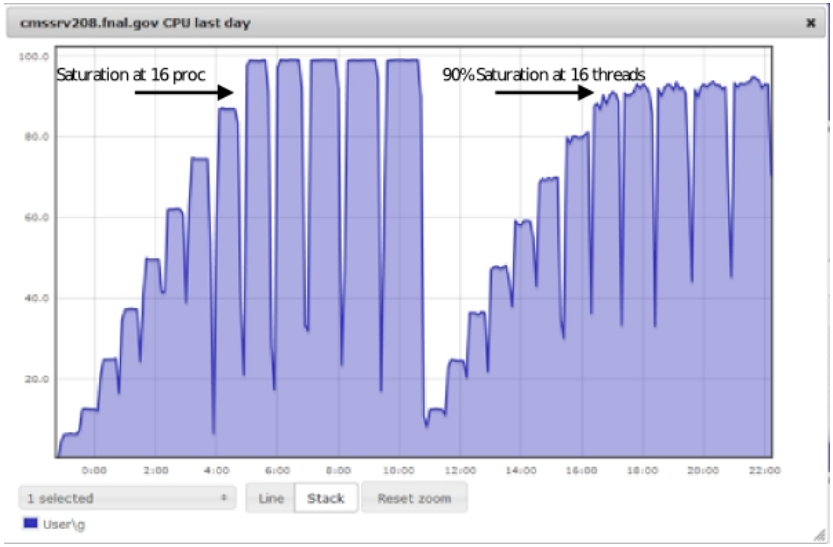

Resource Provisioning

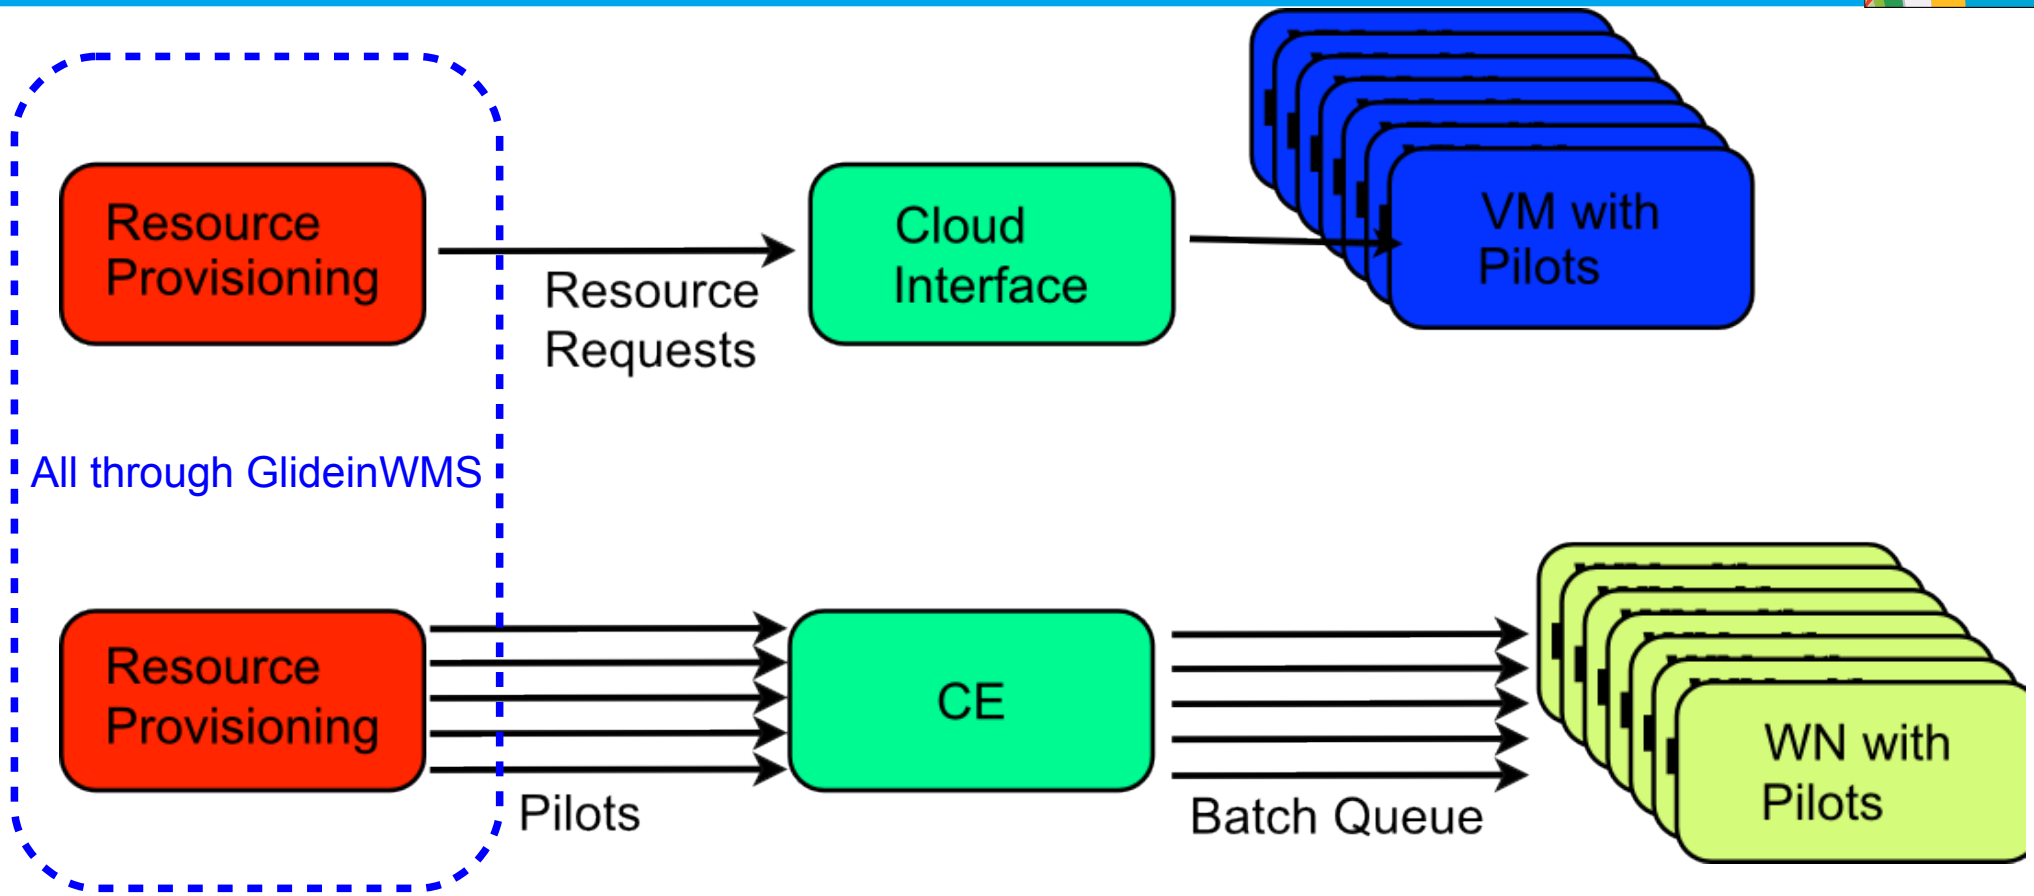

- All resource allocation through GlideinWMS
- In production
  - Grid sites with various CE flavours
  - OpenStack Cloud interface
- Expanding to opportunistic resources

Tier-0 Application: Resource View

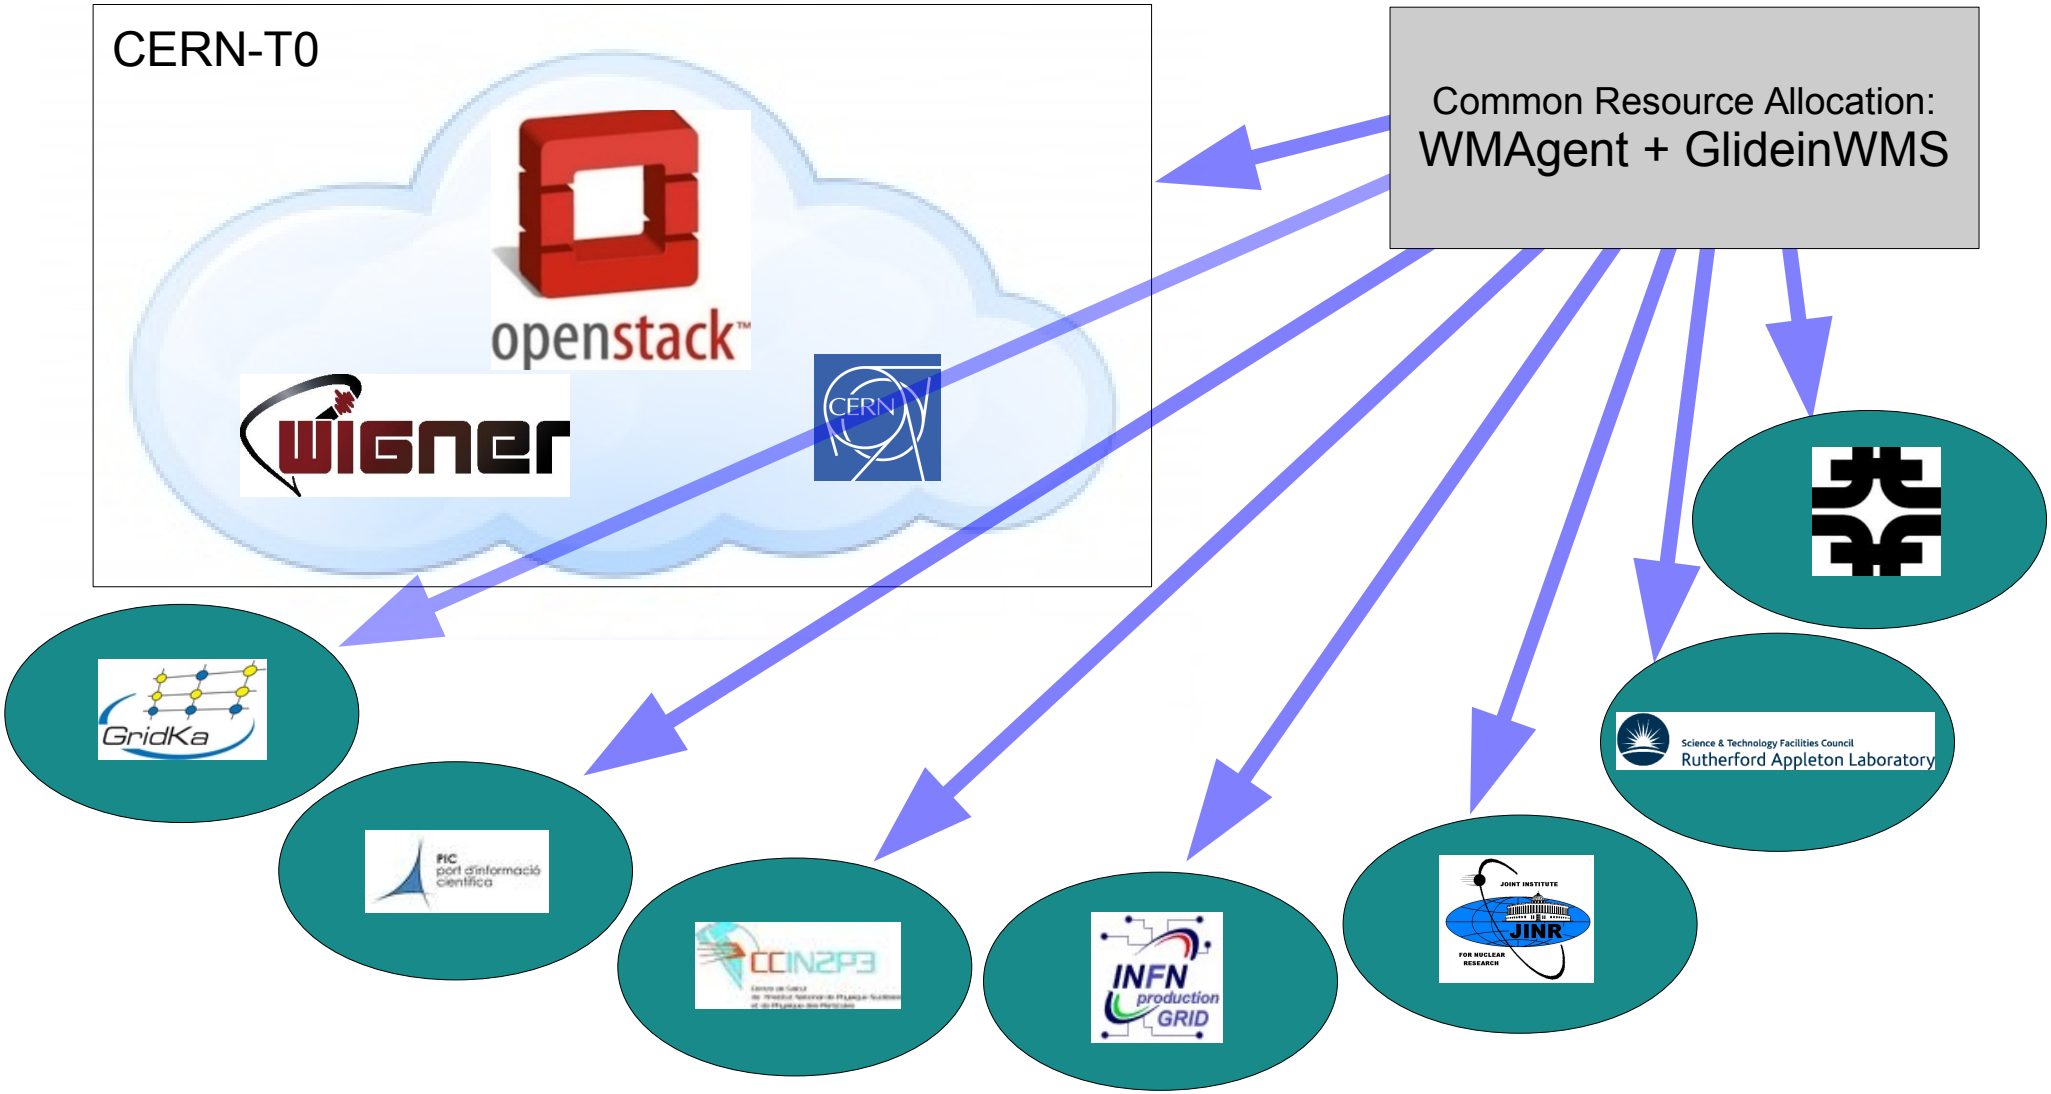

2015: CERN CPU resources sufficient for PromptRECO

Global Pool

- Global Pool
  - One single pool for production and analysis
  - All priorities handled inside the pool (not at the sites)
- Status
  - Migration done “on the fly”
  - Reached ~150k concurrent jobs!
- Pilots at Sites
  - Tier-2 only VOMS role “pilot”
  - Tier-1 for the time being still mixture of roles “pilot” (for analysis) and “production”
    - Limits scaling
    - Switch soon to one type of pilots for Tier-1

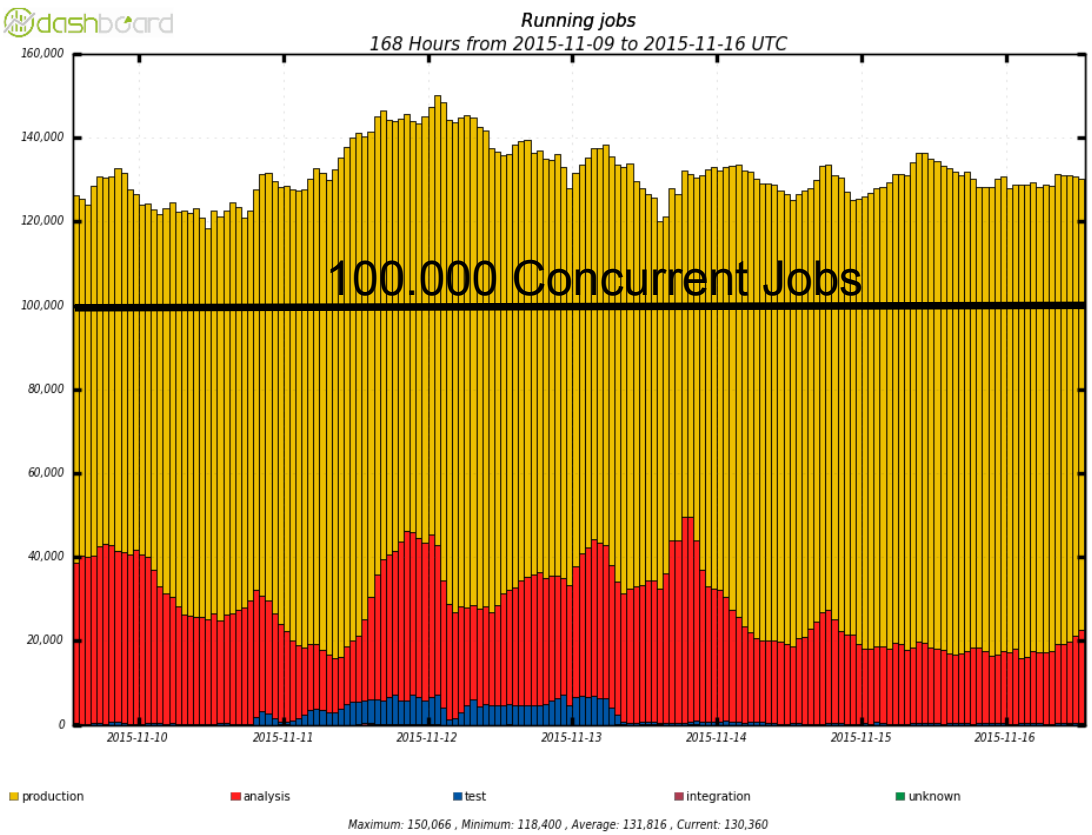

Lowering Site Boundaries

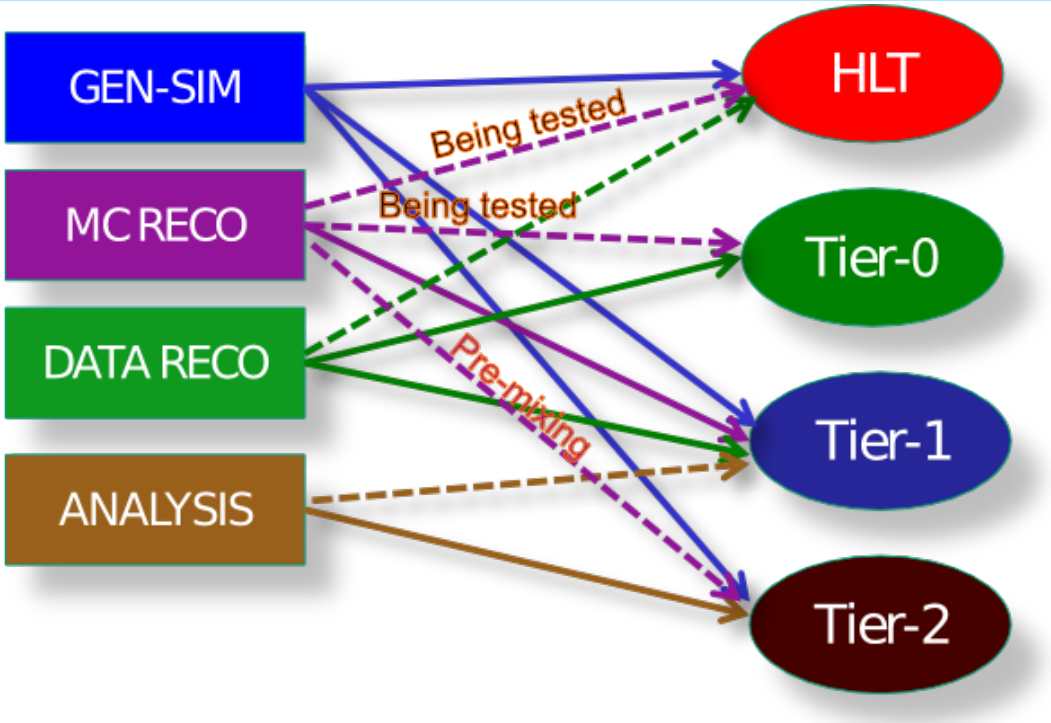

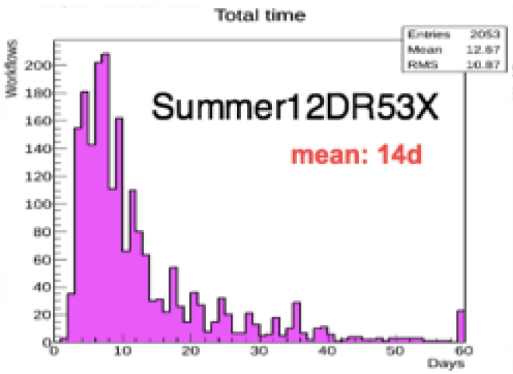

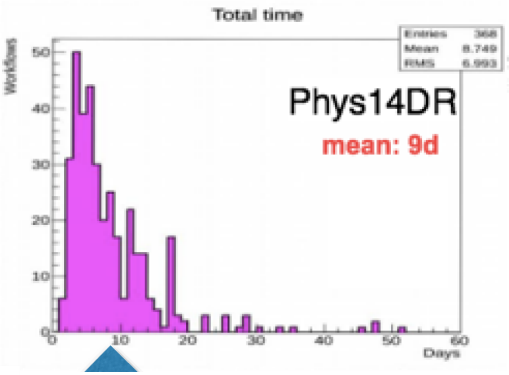

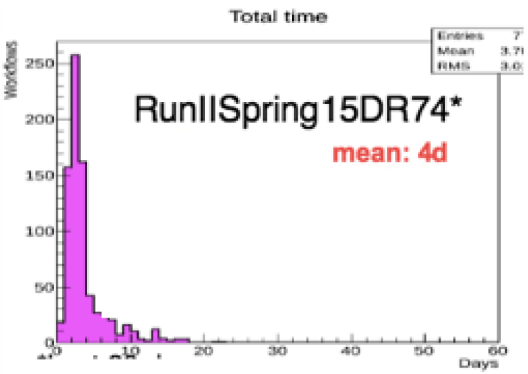

- Gaining flexibility where to run which kind of workflow
  - Use of ~20 Tier-2 sites for DIGI-RECO routinely (even without pre-mixing)
- Improved stability in WMAgent
- Largely automatized assignment and close-out of workflows
  - Presently available for MC workflows
  - Extension to data (re)-reprocessing in development

Lowering Site Boundaries

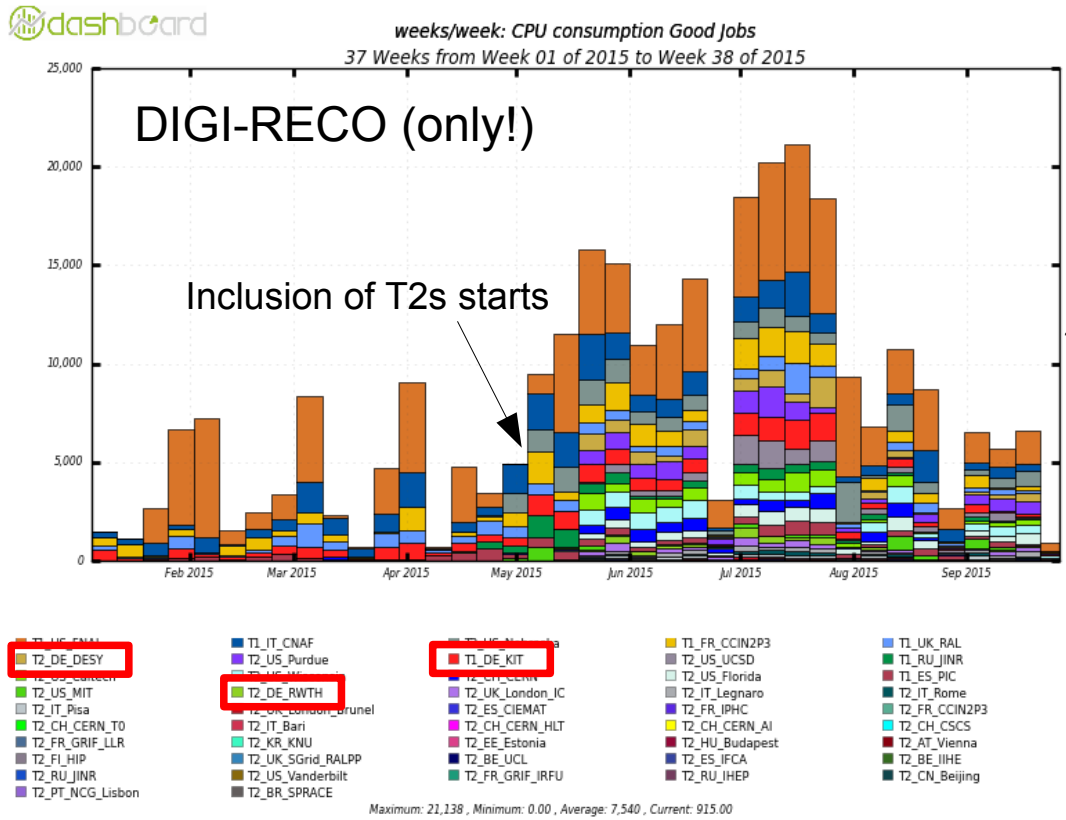

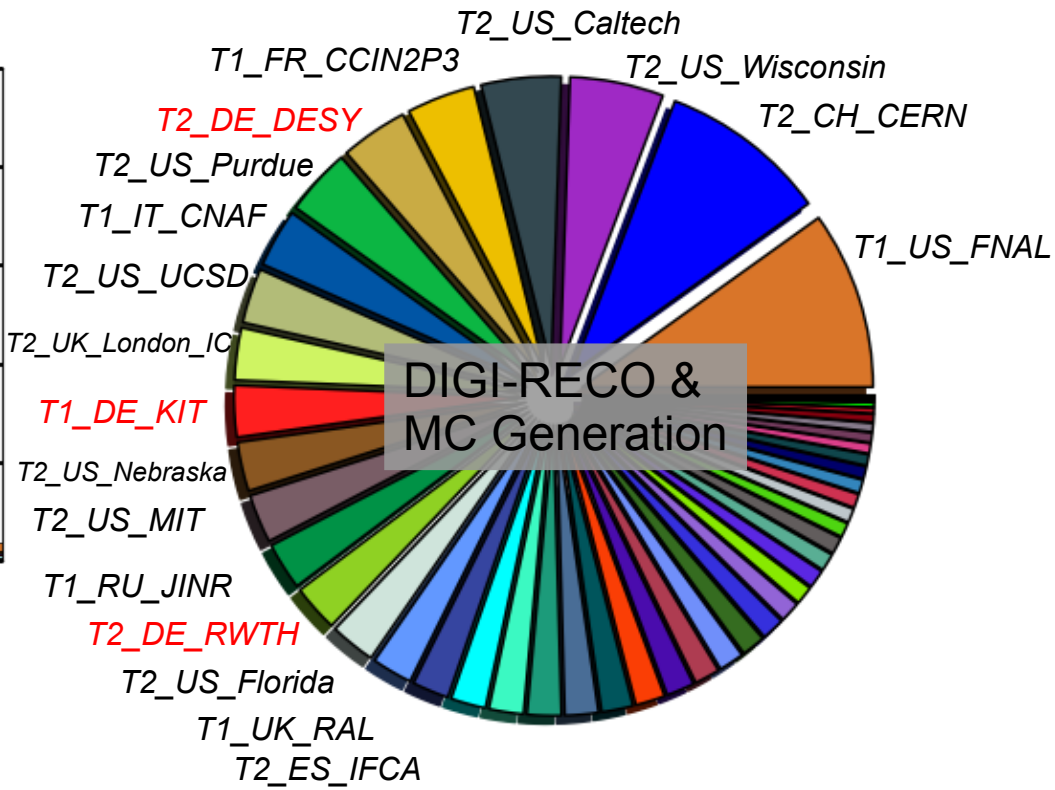

- DIGI-RECO workflows traditionally executed at Tier-1
- From May 2015 half of the Tier-2 sites added
  - Huge gain in CPU capacity for that workflow type
- Strong contribution from German sites to central processing

Opportunistic Resources

- Resources are tight for Run2
  - Have to restrict physics program to the 'necessary'
  - Resources beyond WLCG pledges (=opportunistic) offer some opportunities
- Various kind of resources and access possibilities
  - Grid resources beyond WLCG pledges – we used to have Tier-3s also in the past
    - CMS sites and non-CMS sites
    - Access is straight forward and supported
  - Access to academic clouds
    - Base experiences from HLT cloud
    - Likely to find other places
    - GlideinWMS provides EC2 connection
  - Temporary access to High performance clusters
    - Usually do no have Grid interfaces like Compute Elements
    - BOSCO interface GlideinWMS can start pilots through SSH connection
    - Won't have root access to install e.g. CVMFS for software distribution
    - Parrot allows “installation” of CVMFS as non-root user at run-time
    - Can be complex, e.g. Linux different flavour different from SL
  - Commercial clouds (might be cheap when weakly utilized)

Exploring Amazon

- CMS negotiated a “matching grant” with Amazon
  - Successful proposal – similar to ATLAS
  - CMS pays ~10% of the real cost, 90% sponsored by vendor
  - Started mid of this here and extends until March 2017
- Implementation plan
  - Via “dynamic extension” of FNAL resources
  - Buy “cheap” on the spot market

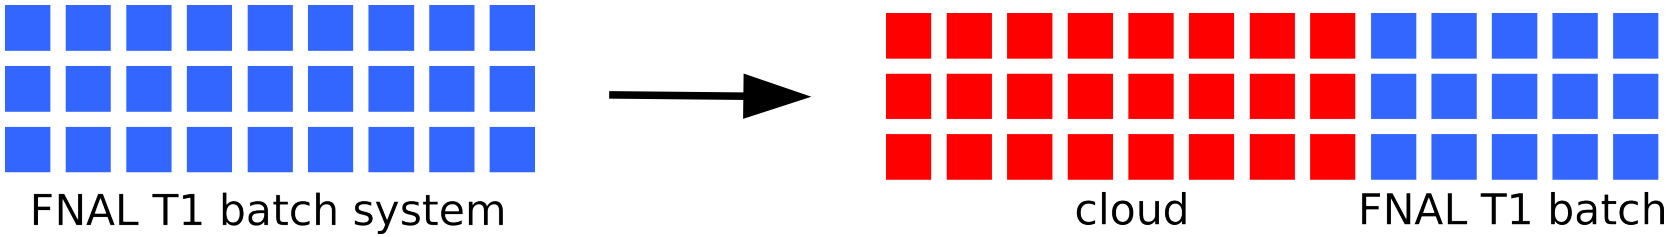

- Challenges are in the details
  - Avoid traffic into/out of the Amazon cloud (costs money)
  - Make selected workflows use the extension
  - Transparent integration into existing systems

A new Data-Tier: miniAOD

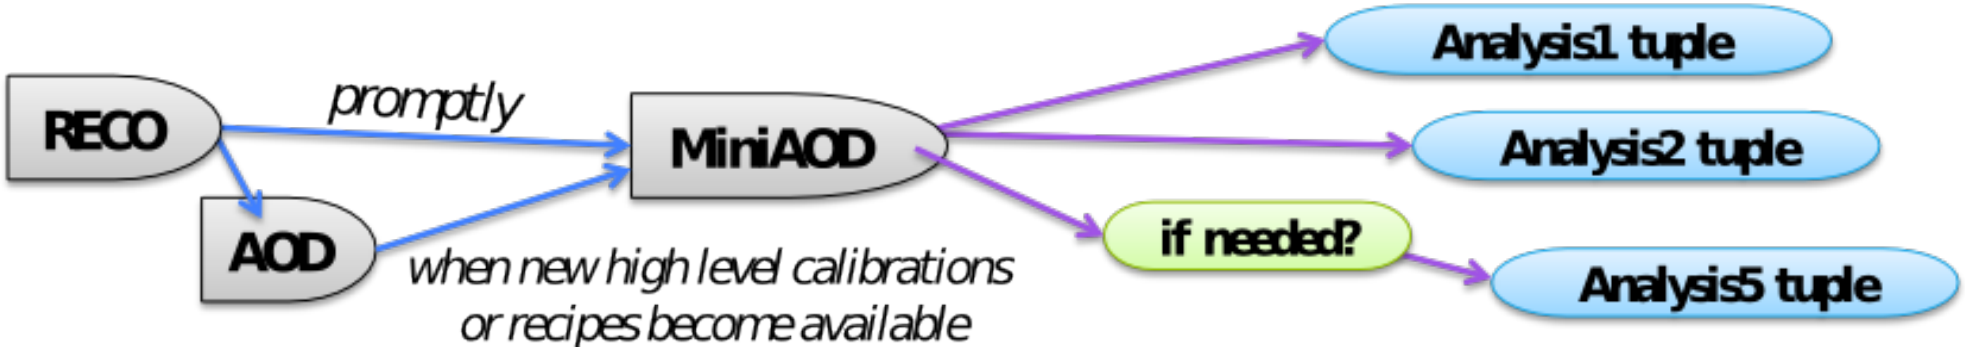

- Replace the dozens of “Group ntuples/trees”
- Size ~50kB/events
  - High level physics objects (Leptons, Jets...)
  - Particle Flow candidates in packed format (to allow for re-clustering, jet-structures etc)
- Should satisfy ~80% of all analysis cases
- Large exercises during CSA14
  - Actively used for many Run2 analysis
- MiniAOD production
  - Additional output of PromptRECO
  - Can be re-done a few times per year

User Analysis Job Submission: CRAB3

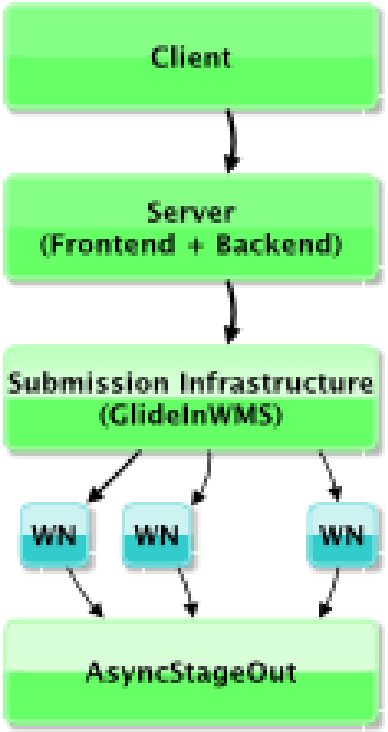

From M.Mascheroni

Submission Use Case

- Client sends a submit request to the server
- The server frontend checks auth/authz and store the request. The backend "prepares the jobs" and send them to the submission infrastructure
- The jobs are sent to the Grid. They contact the AsyncStageOut (ASO) server to start the transfer
- ASO transfers the files (using FTS) and it publishes them

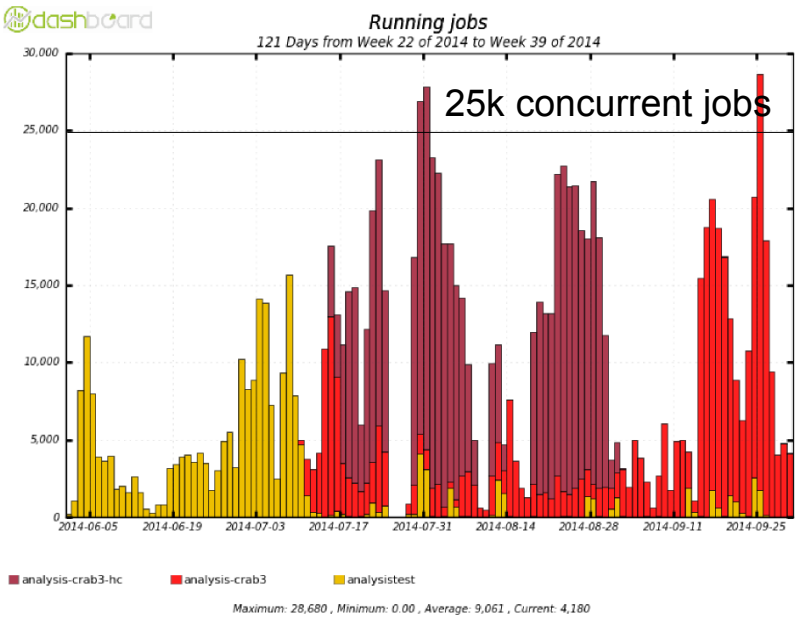

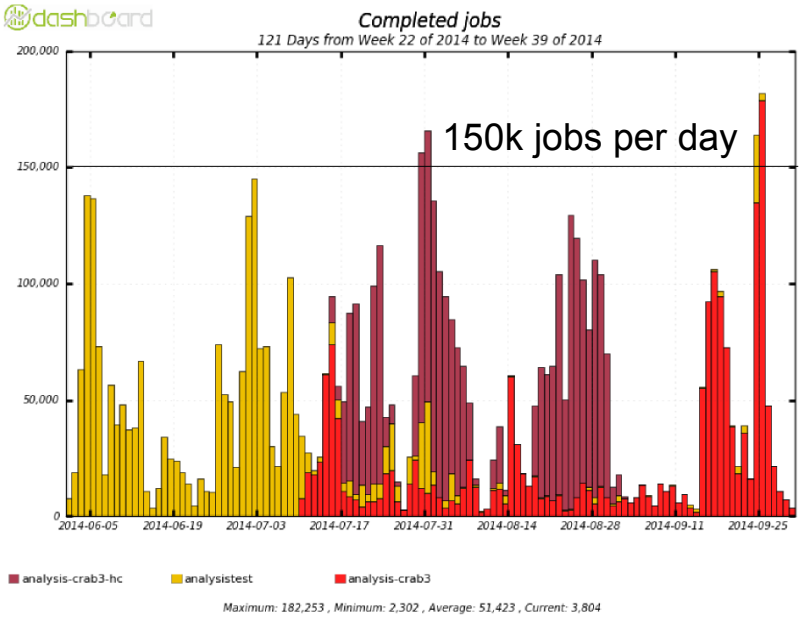

- Default tool for user analysis in Run2
- ~30k parallel jobs reached
- Processed up to 150k jobs per day

Summary

- Coming Run2 could become resource constrained in 2016
  - Improve CPU performance of applications
  - Increase the flexibility of resource usage
- A number of efforts to prepare for Run2
  - Completely Cloud/Openstack based Tier-0
  - Expanding PromptRECO to Tier-1 sites
  - Commissioning of the HLT for Production and Processing
  - Data federation to ease remote data access
  - Dynamic data management to improve usage of disk space
  - A new tool for user analysis (CRAB3)
- Experiments want to discover new physics in Run2
  - Computing relies on good planning
  - Difficult to plan for the unknown

Will have an exiting to time to deal with this “conflict”