CMS Experiment Status, Run II Plans, & Federated Requirements

Maria Girone, CERN

XrootD Workshop, January 27, 2015

Outline

- Computing Status and Preparation for Run II
- CMS Computing Improvements during Ls1
  - Data Management
  - Data Placement
  - Data Access
- Federation Progress and Requirements
- Validation and Commissioning during the Computing Software and Analysis Challenge (CSA14) and current state of the federation
- Outlook

Size of the Problem in Run II

- In Run II, CMS expects to promptly reconstruct 1000Hz of data (a factor 2 more) at an average of twice the reconstruction time (another factor of 2 more)
  - Sample of events to collect, promptly reconstruct and eventually reprocess is 3B events in 2015
- CMS expects to generate 1.3 times simulated events for every event collected (4B GEN-SIM by end of 2015)
  - Budget for processing is ~1.5 times events collected
- Analysis facilities were specified to primarily provide resources to analyze Run II data only
  - The resource request assumes that Run I data analysis will ramp down with the start of the run

Resource Balancing

- Organized processing for data and simulation reconstruction and reprocessing capacity will nearly double wrt 2012, while processing needs increase by more than a factor of 4
  - We need to do more with less, use resources more efficiently and in a more flexible way
- Disk resources double at Tier-0, but increase only 15% at Tier-1s and Tier-2s, despite the doubling of the trigger rate and the increase in event size
- The expectation is that we must use the storage much more effectively
  - The data federation is a crucial element in this

Facilities and Resources in 2015

- CERN
  - 12k cores for Tier-0 and additional 15k when the HLT farm is available
  - 15PB of disk and 31PB of tape
  - Network between 10-100Gb/s to Tier1s
- Tier-1
  - 7 facilities primarily at national labs or large computing centers
  - ~40k cores 27PB of disk, and 74PB of tape
  - Network between 1Gb/s – 100Gb/s to Tier2s
- Tier-2
  - 50 facilities primarily at university centers
  - 80k cores and 31PB of disk

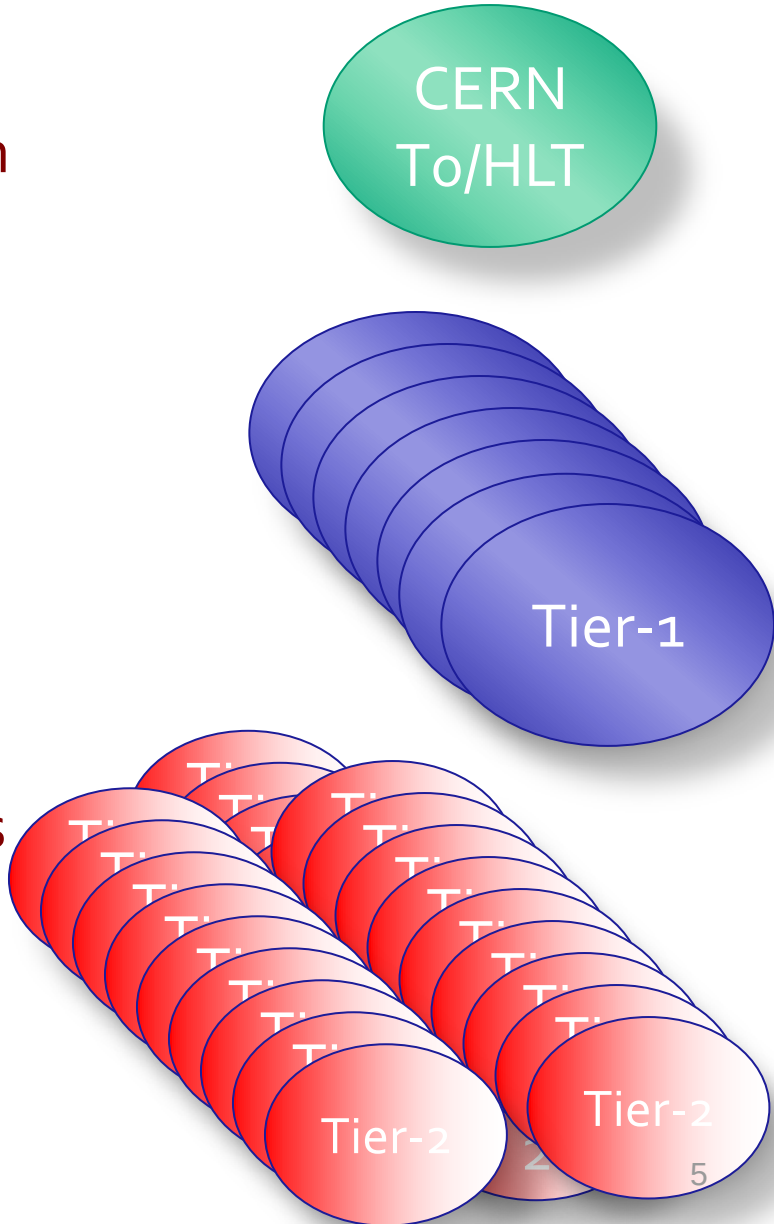

Removing Site Boundaries

- In LS1 we have made 2 significant changes
  - Logically the disk and tape storage systems have been split
    - A file written to disk on a Tier-1 site needs to be explicitly subscribed to a tape endpoint
  - A data federation has been deployed across CERN, all Tier-1s, and 90% of the Tier-2 disk space
    - Hierarchical redirectors have been installed

- Boundaries between sites have been removed
  - The difference between capability of all the Tiers has been reduced
  - Improved networks are key to this

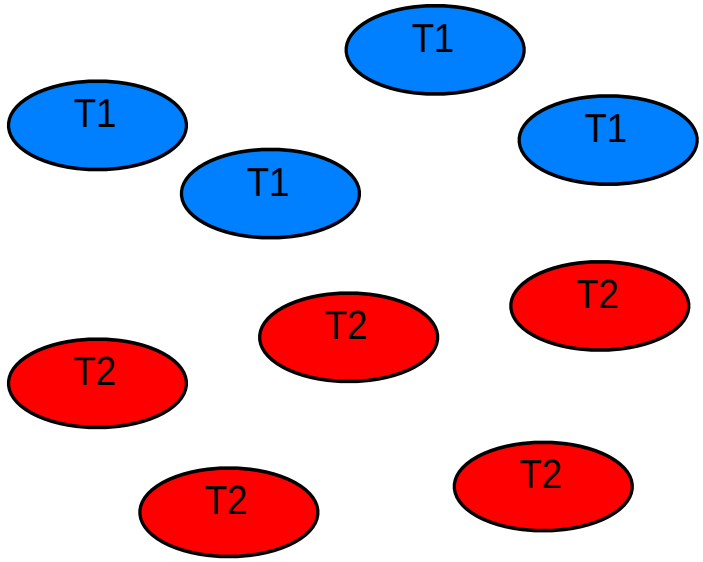

Resources and Workflows

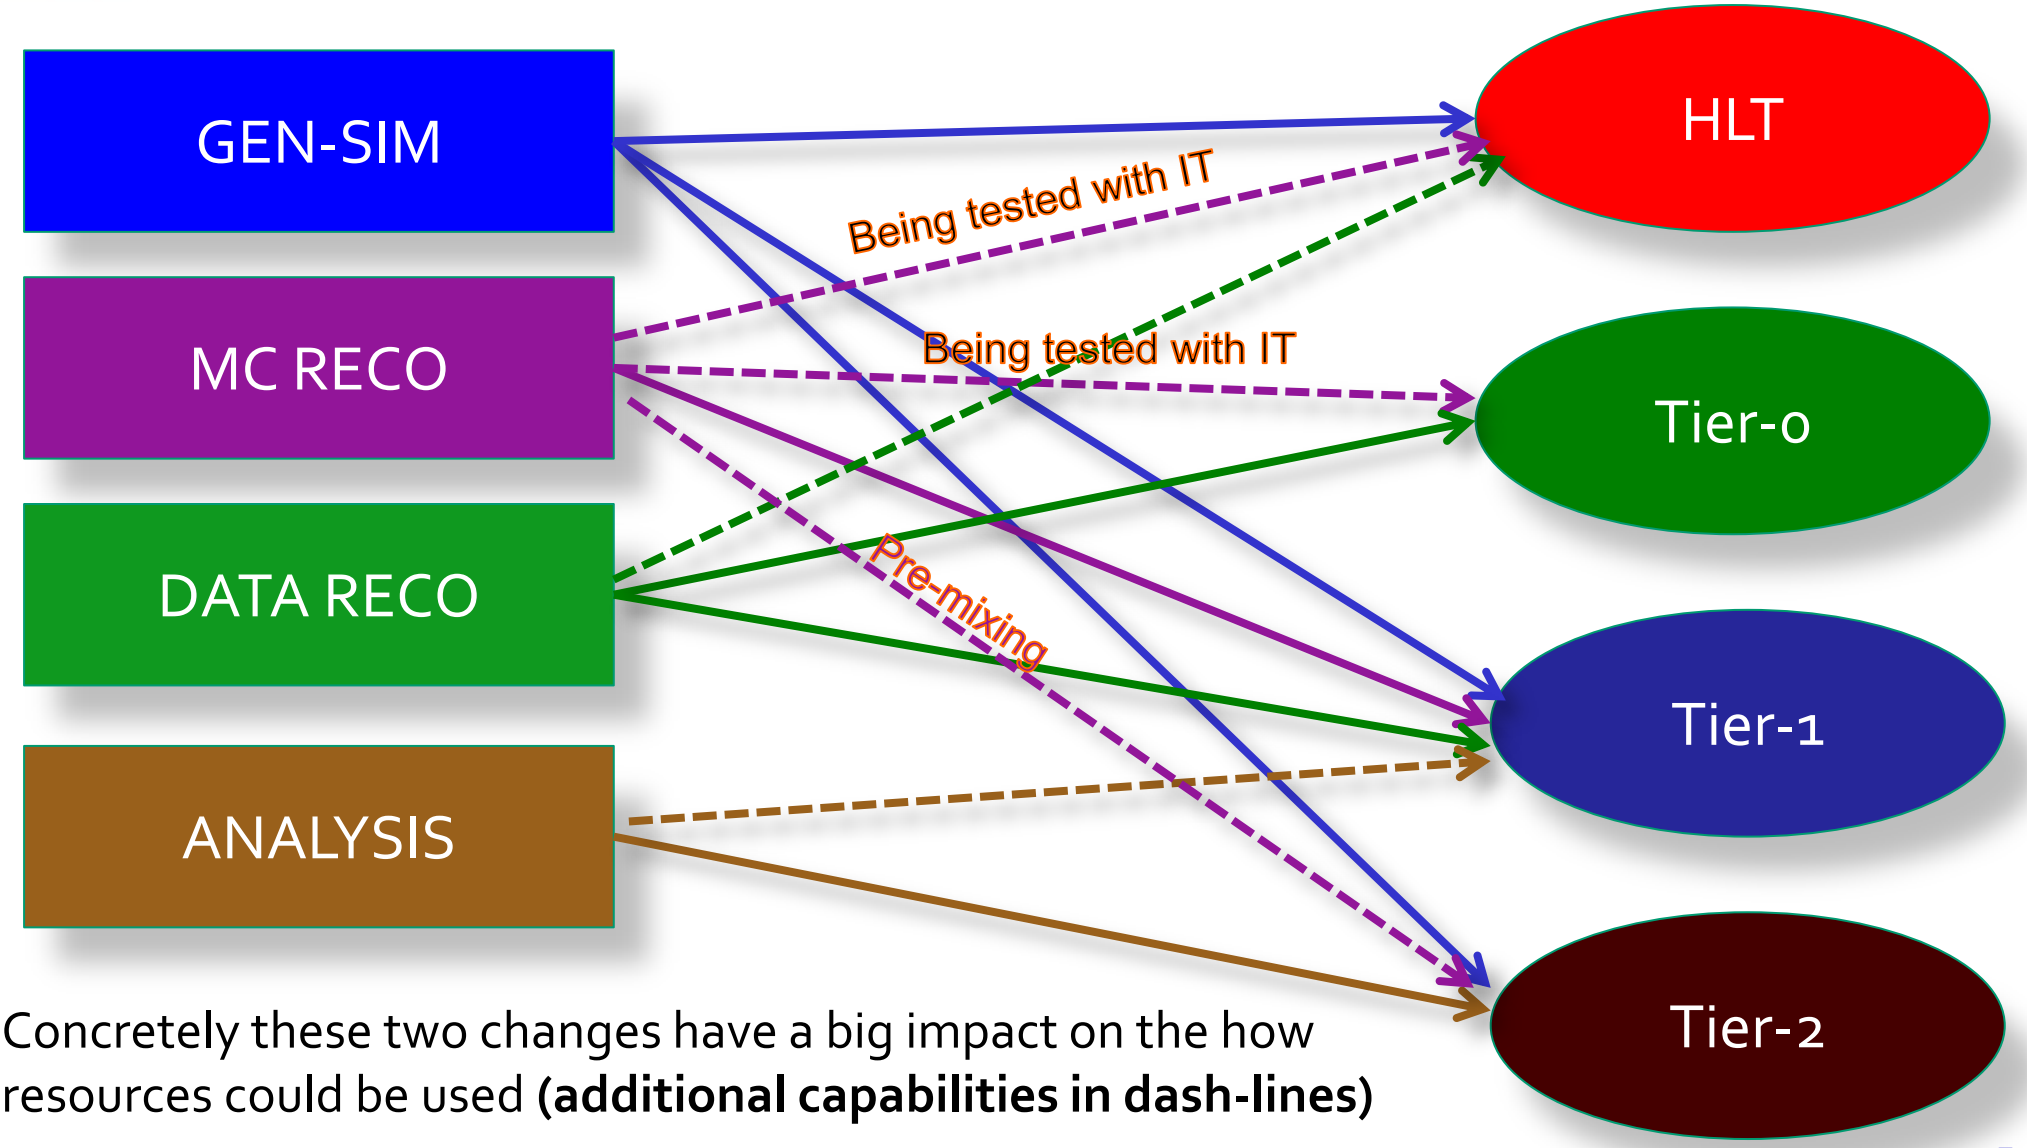

Concretely these two changes have a big impact on the how resources could be used **(additional capabilities in dash-lines)**

The HLT Farm

- An addition for Run II is the use of the High Level Trigger (HLT) farm for offline processing
  - It is a large computing resource (15k cores) that is similar in size to the Tier-0 in terms of number of cores, but we cannot reach this scale until March
  - Successfully interfaced using cloud computing tools. It is similar to the Tier-0 AI
- In 2014 the network link P5 to the computing center was upgraded from 20 to 60Gb/s
  - Far larger than needed for data taking but necessary to access the storage in the computing center for simulation reconstruction
  - **Will be upgraded to 120Gb/s before the 2015 run starts**
- Production workflows have been commissioned including the HI reprocessing, Gen-Sim, and Simulation reconstruction
  - All access to data is through the data federation and primarily served from CERN

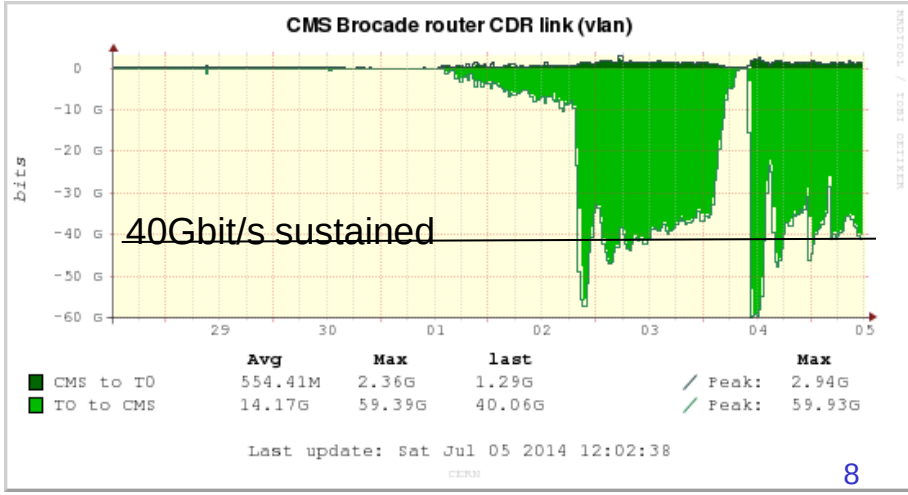

Data Management

- In addition to data federation, we have improved our traditional data placement and access
  - The use of samples is continuously monitored through the data popularity
- A new miniAOD format has been added for 80% of the analysis in 10% of the size
- Samples will be replicated dynamically if the load is high and replicas removed if they are not accessed for a period of time
  - Data and MC samples are distributed to Tier-2 (Tier-1) sites automatically
  - Number of replicas depends on the popularity of the datasets
  - Samples are only deleted when there is new data to replicate, disks are kept full
  - We are working to improve and automate the monitoring of the disk usage
- This is needed in part because the storage is inefficiently used
- Even at a 6M window, 27% of the disk space for AODSIM is used by un-accessed samples (10% of the total space)
- The zero bin includes un-accessed replicas and datasets that have only **one** copy on disk

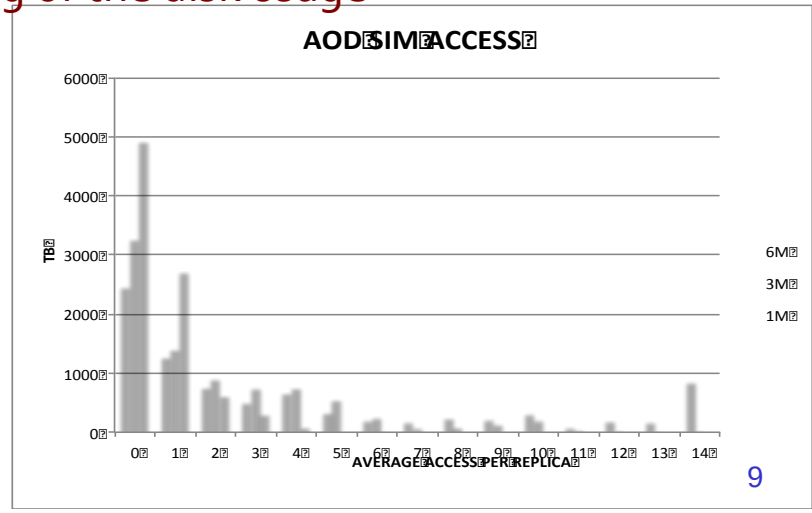

Data Access

- Access to data is CMS is done in 2 ways
  - At a local site with an interactive application by specifying the Logical file name (LFN), which is resolved into a file to open
    - This is done by users and debuggers
  - Or remote submission, where a job is specified by dataset (a large group of LFNs) and a large number of jobs is launched to grid enabled resources.
    - Distributed analysis and production
  - There have been significant improvements in both with the commissioning of a data federation

Federation Requirements

- **System Requirements**
  - Provide remote access to CERN, Tier-1, and nearly all Tier-2 disk storage through XrootD
    - More than 90% of all current CMS data should be accessible
    - Sufficient IO capacity to provide 20% of the total access (~100TB/day)
    - Enable to system of hierarchical redirectors to maintain access within geographic regions when possible

- **Functionality Requirements**
  - Interactive access
    - User opens a file from anywhere. Performance and stability are secondary to convenience
  - Failover protection
    - A cluster local storage fails to deliver a file and the federation serves as a backup. These jobs were going to die and the federation can save some
  - Overflow
    - A site is busy and jobs are intentionally redirected to a site nearby. Analysis Operations chooses close reliable sites
    - A user intentionally directs jobs to any site and accesses samples over the wide area at her/his own risk
  - Production
    - Production jobs share workflows across sites and the federation is used to serve the data. In order to have good operational efficiency the federation has to have similar performance and reliability to local storage

Data Federation

- Any Data, Anytime, Anywhere (AAA) has been a primary focus area in 2014
  - CERN, all Tier-1s, most of the Tier-2 sites serve data in the federation
  - Nearly all sites are configured to use the federation to access samples, if they aren't available locally

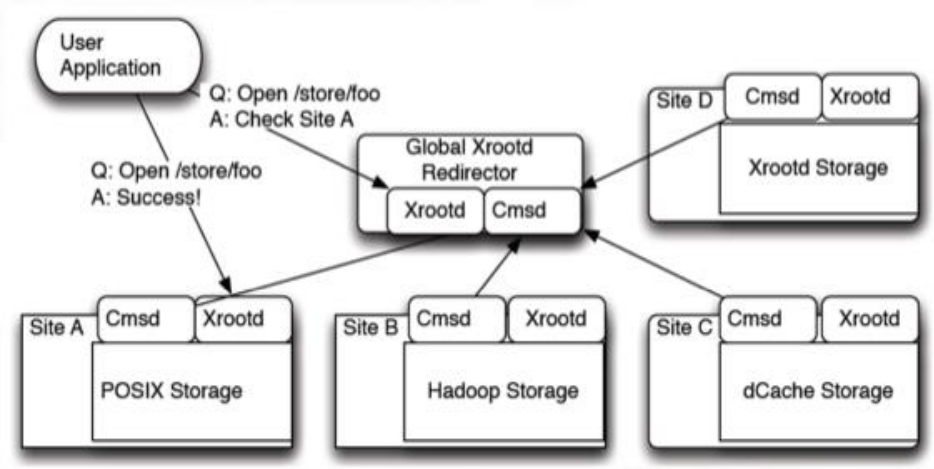

- Optimization of the IO has been an ongoing activity for several years, which has paid off in high CPU efficiency over the wide area
- Big push in 2014 to commission sites to measure IO and file open rates and to understand the sustainable load and to deploy and use advanced monitoring
- Target for Run II is that 20% of all access is through the federation

Commissioning Work

- We validated small scale use of non-local data access in CSA14
  - Fall-back when CRAB3 jobs don't find input data locally and in "ignore locality" mode
  - Very good feedback by users
- After CSA14 scale tests were performed in Europe and the US
  - 20% of jobs were able to access data over the wide area (60k files/day, O(100TB)/day)
    - Tests showed that the scale could be reached, but that the job success rates were sensitive to the health of all the sites

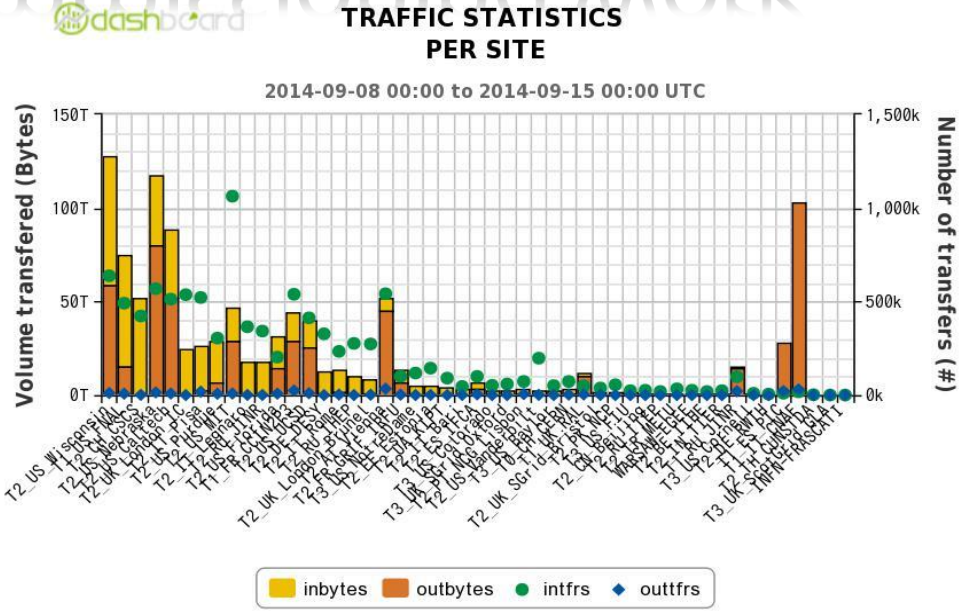

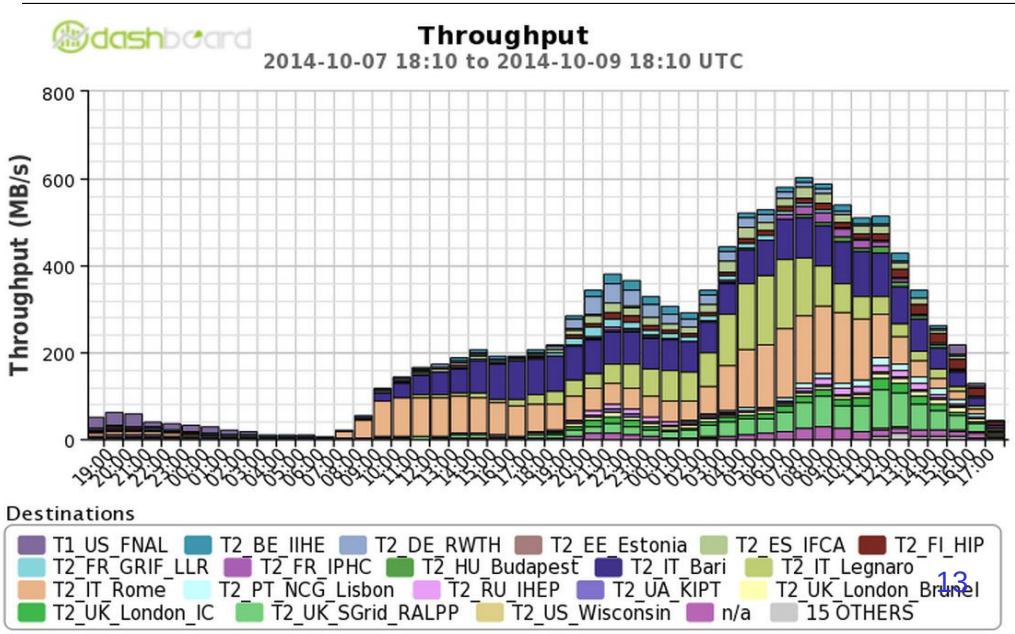

Creation of the Transitional Federation

- The addition of the production workflow puts additional constraints on the required reliability of the Federation

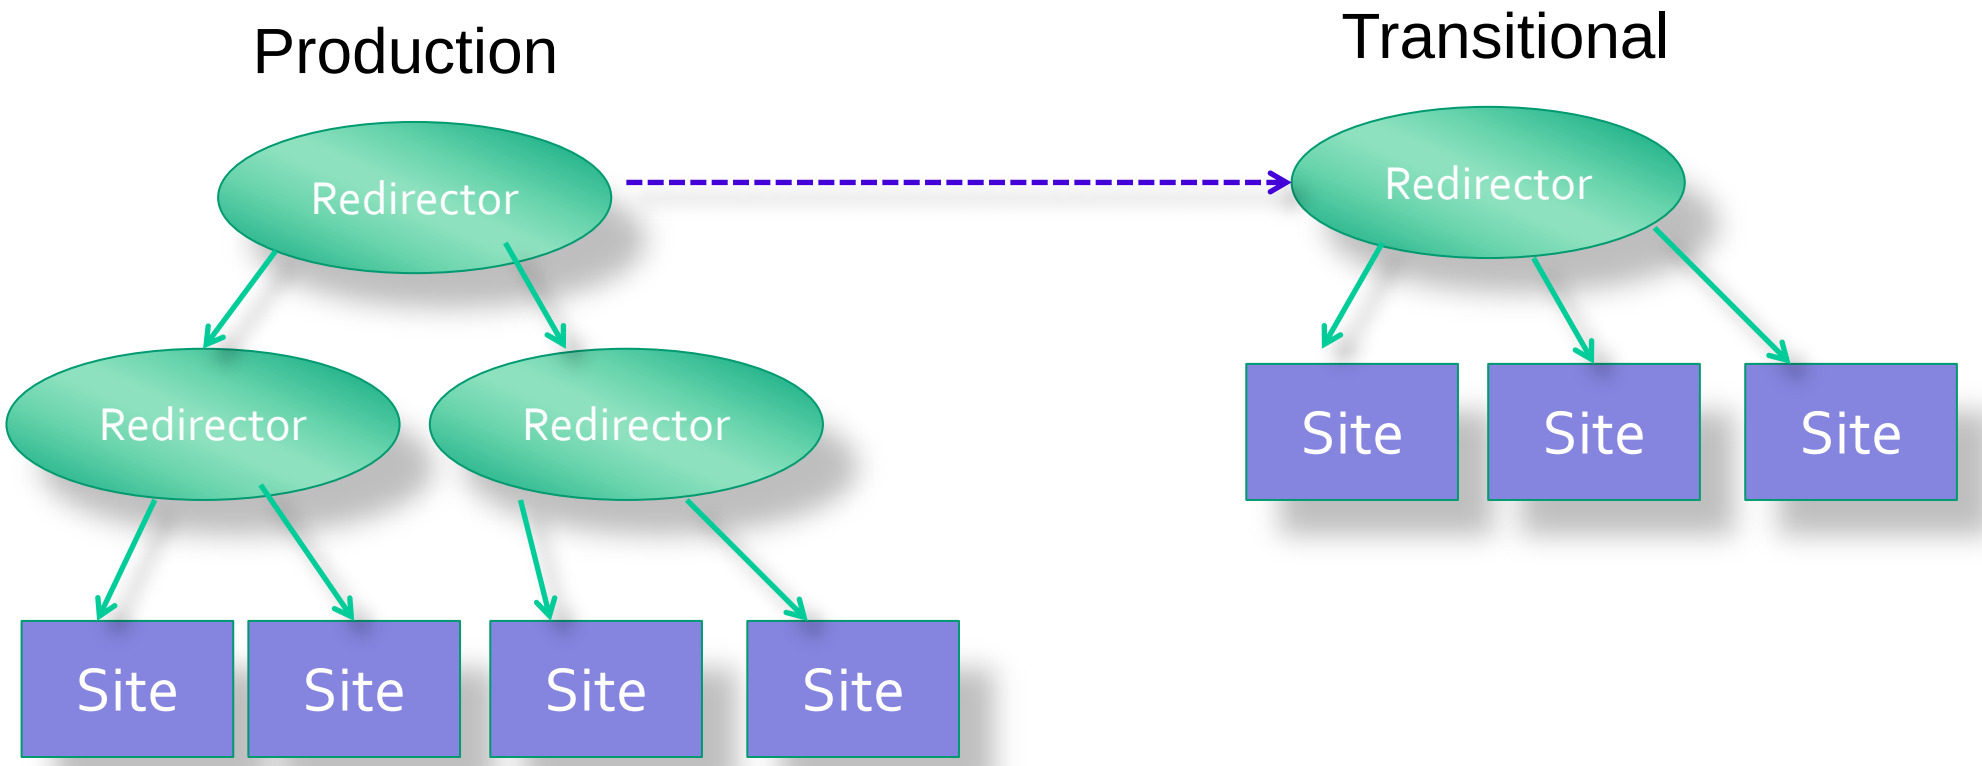

- Validated sites are in the production federation, sites being commissioned are in an independent federation and only when a sample cannot be found in production are they used

Interactive Access

- In Run I most of the low latency commissioning and analysis activities had to be done on the CAF (CERN Analysis Facility)
  - The express data is processed and delivered in an hour, but that doesn't allow time for merging so the data could not be efficiently replicated to remote sites
  - Access to the storage and CPU on the CAF was limited to about 10% of the collaboration
- In 2015 CERN will have 12PB of data serving local users and the federation
  - The express data, in addition to all samples, will be accessible to everyone as soon as it's produced
  - The federation equalizes access to samples

Failover and Overflow

- Currently all sites in CMS have the failover protection enabled
  - A small site configuration change protects against site storage failures
    - It's only a few percent of jobs, but good to have protection
- Overflow is enabled for US sites; we are preparing a European deployment
  - A matrix of well connected sites is created and jobs are directed to alternative locations when the queuing time becomes too large
    - Normal operations is ~10% of the jobs, and this improves the average waiting time

Production

- The use of Data Federation in production is still under commissioning

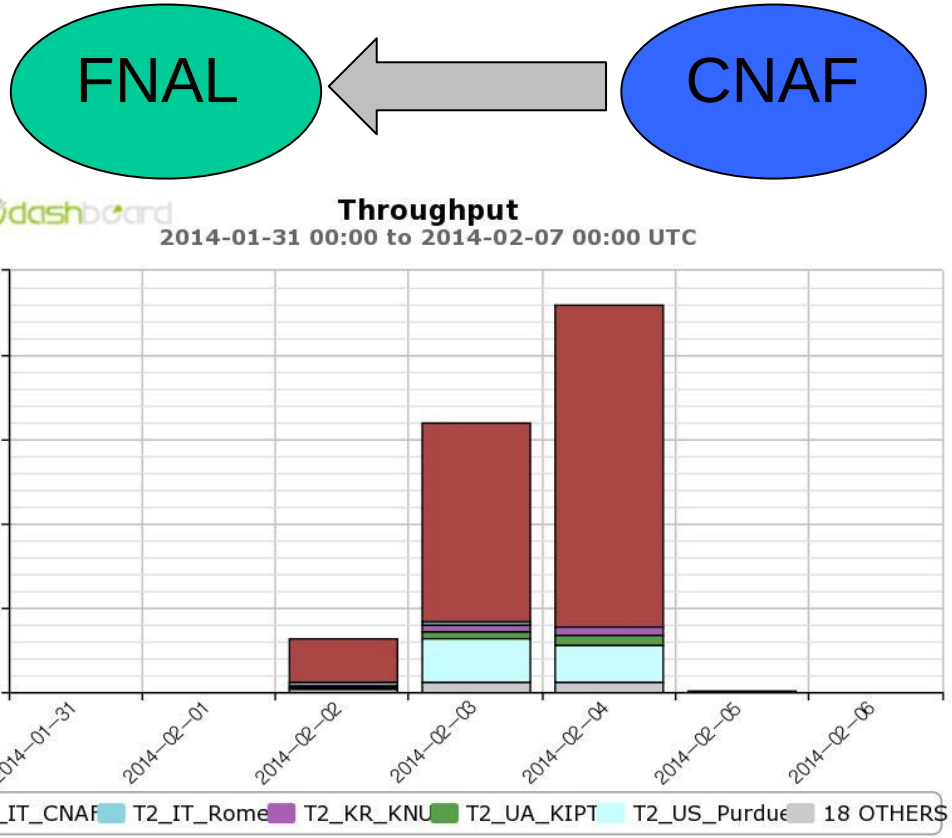

Proof of Concept

- Last January during a storage update, FNAL accidently pointed to an obsolete DB and production traffic failed over to Xrootd
  - Data was read from CNAF disk at 400MB/s for 2 days
  - Production jobs continued to succeed at the same rate but the network monitors noticed traffic with >3Gb/s on the OPN

This was a good scale test for AAA for shared production

- Benefited from WN on OPN (case of FNAL)

Production Changes on the Software Side

- After significant software development, CMS now has a production ready multi-core framework
  - More than 99% of the running core can run in parallel. This allows us to take better advantage of the growth in cores per system, decrease the memory footprint, reduced the number of running processes to track, improved merging, and simplified the IO paths
    - Big memory savings: 0.35 GB per additional thread instead of 1.8GB/job

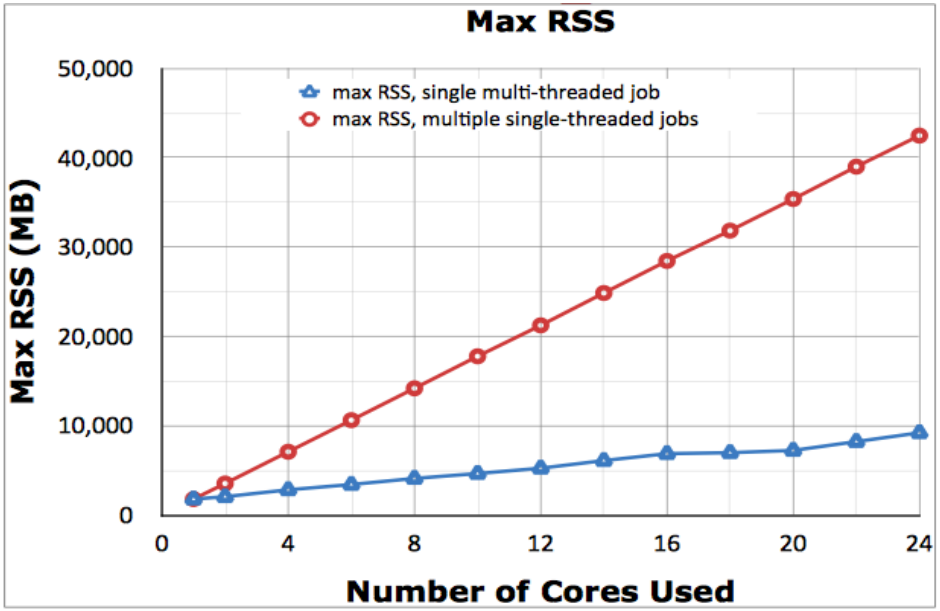

- An interesting side effect of this is that the IO of the application increases by a factor of 4-8 depending on the number of cores used
  - Instead of many small IO processes (50-100kB/s for Reco) we have fewer processes with higher IO (200kB/s-1MB/s)
    - For DIGI 3-5MB/s becomes 12-40MB/s

Production Use of Opportunistic Computing

- One of the challenges of opportunistic resources is access to data
  - Batch queues are dynamic and access is given. Storage usage is more static and resources are not given temporarily, because people are much more willing to kill jobs than to delete samples
    - Setting up a storage element an a PhEDEx instance is labor intensive and doesn't lend itself to dynamic deployment
  - Up to now this has limited to kind of processing we did opportunistically to activities that didn't require much input samples
- The data federation allows data to be streamed to opportunistic CPUs and opens up a class of activities on opportunistic computing

Outlook

- We believe that access to the samples will be dramatically improved in Run II
  - Use of data federation and dynamic data management should ensure that users have access to everything from wherever they choose to work
- Run II will be a challenge and improvements in data access were necessary to make the best use of the resources
  - We will have less computing and storage per event collected and simulated than we did in Run I
  - We need to improve efficiency
- Data Federation improves data access for everyone and improves the flexibility for using resources for organized processing
  - Feedback from users has been extremely positive