Online data handling with Lustre at the CMS experiment

Lavinia Darlea, on behalf of CMS DAQ Group

MIT/DAQ CMS

September 17, 2015

CERN

LHCb
ATLAS
CERN Meyrin
CERN Prévessin
SPS 7 km
SUISSE
FRANCE
CMS
ALICE
LHC 27 km
CERN

CERN was founded 1954: 12 European States

"Science for Peace"

Today: 21 Member States

~ 2300 staff
~ 1050 other paid personnel
> 11000 users

Member States: Austria, Belgium, Bulgaria, the Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Italy, Israel, the Netherlands, Norway, Poland, Portugal, Slovakia, Spain, Sweden, Switzerland and the United Kingdom

Candidate for Accession: Romania

Associate Members in the Pre-Stage to Membership: Serbia

Applicant States: Cyprus, Slovenia, Turkey

Observers to Council: India, Japan, the Russian Federation, the United States of America, Turkey, the European Commission and UNESCO

CERN

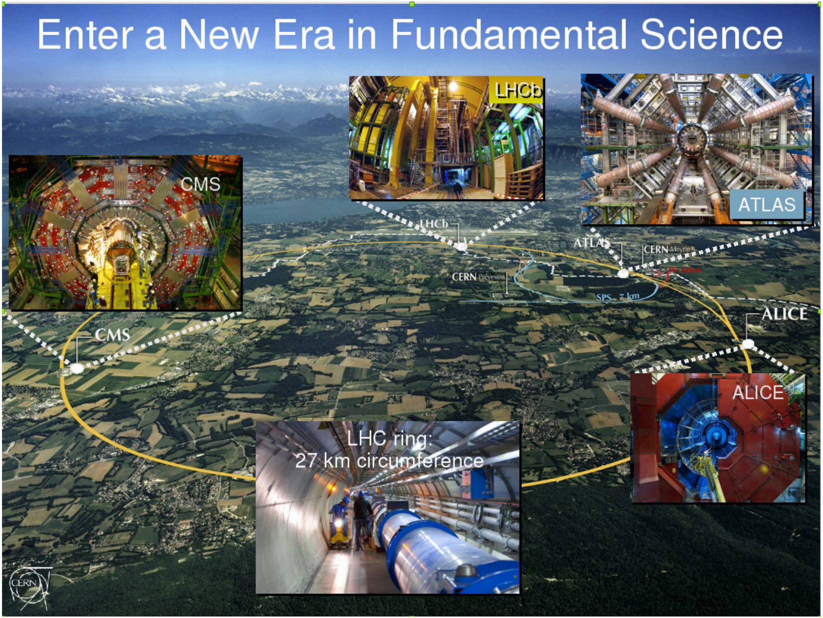
Enter a New Era in Fundamental Science
LHCb
ATLAS
CMS
LHCb
ATLAS
CERN Meyrin
CERN Prévessin
SPS 7 km
ALICE
CMS
ALICE
LHC ring:
27 km circumference
CERN

Protons collide in the CMS detector

- Took ~2000 scientists and engineers more than 20 years to design and build
- Is about 15 metres wide and 21.5 metres long
- Weighs twice as much as the Eiffel Tower – about 14000t
- Uses the largest, most powerful magnet of its kind ever made

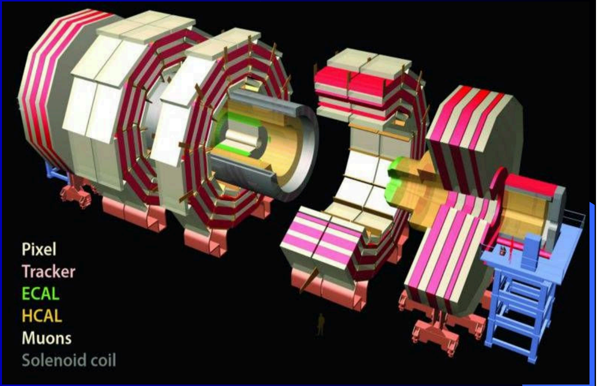

CERN

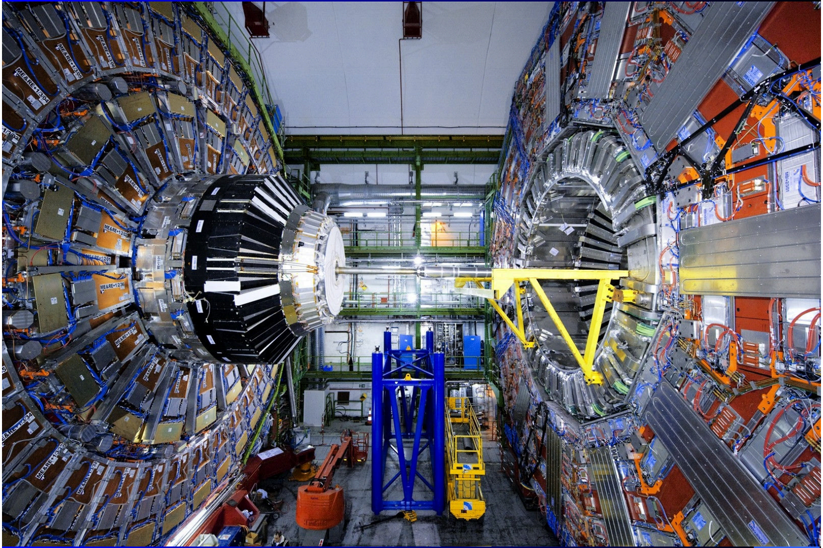

Introduction

CMS data flow, rates and computing model at LHC

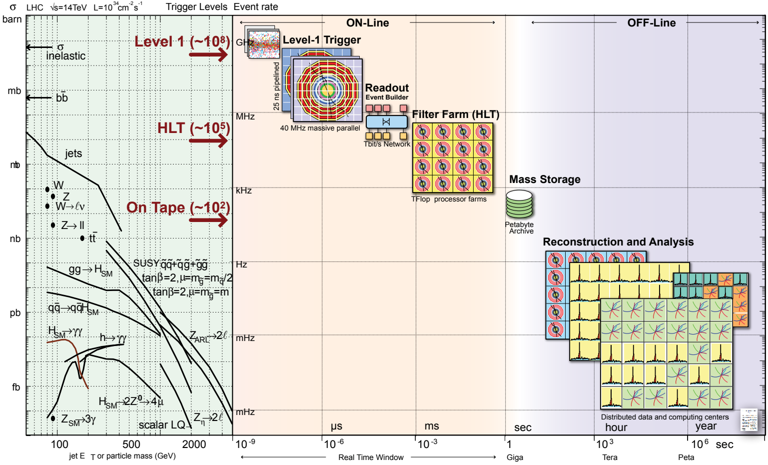

CMS DAQ2 System

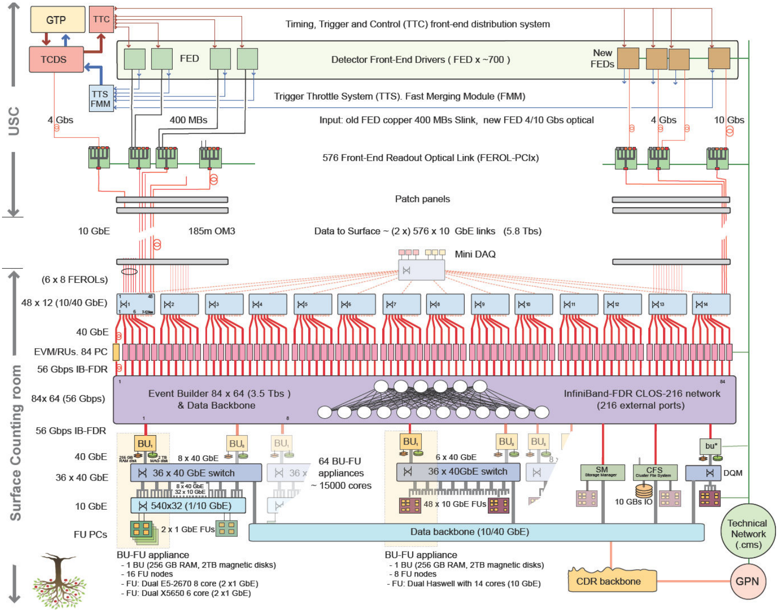

CMS DAQ2 System

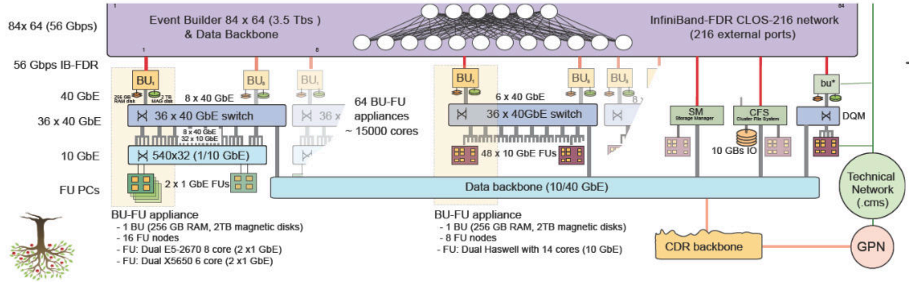

Storage Manager and Transfer System (SMTS) in the DAQ chain

SMTS and DAQ

- input: end of the Data AcQuisition chain*
- last part of the data flow: ensure safe storage and transfer to Tier0

**E. Meschi, File-based data flow in the CMS Filter Farm, CHEP 2015*

Storage Manager and Transfer System Role

Data flow pattern to and from Lustre

- Storage and Transfer Management System (STMS) is the final component of the DAQ system
- Data flow status before the STMS:
  - Builder Units (BU) nodes receive event fragments that they reconstruct
  - They distribute these events to the Filter Units (FU) for selection
  - The FUs send back the meaningful events to the BUs
- The Storage Manager merges the data into a smaller number of files
- The resulting files are temporarily stored in the Lustre File System (LFS)
- The Transfer System picks the files up from LFS and ships them for analysis processing and permanent storage into the CERN Tier0 (EOS/Castor)

Storage Manager and Transfer System Role

Implementation Stages

- merge the filter units output as to obtain 1 data and 1 metadata file/LS/stream
- buffer the data until it is safely stored on tape in Tier0/Castor
- copy the final files according to their intended destination:
  - Tier0 for the main data streams
  - various sub–detectors for online consumption: DQM, EventDisplay
  - store locally for local calibration of various sub–detectors
- ensure hand–shake with Tier0 for proper accountability

Simplified glossary

- LS: 23s
- stream: grouping of similar datasets

Storage and Transfer System Requirements

Merger System

- “merge” data at the BU level such as to obtain 1 file/BU/LS/Stream (mini–merger)
- centralize and merge all the BU outputs such as to obtain 1 file/LS/Stream (macro–merger)
- latency: a maximum of 2LS (1LS = 23s) delay in the macro–merger is considered acceptable
- provide input for the online monitoring system – 1 additional metadata file per data file*
- not only “concatenate”, but deal with special files, such as histograms and jsn files

**S. Morovic, A scalable monitoring for the CMS Filter Farm based on elasticsearch, CHEP 2015*

Storage and Transfer System Requirements

Storage and Transfer

- buffer a minimum of 3 days of continuous running (estimated 250TB)
- aggregated SM input from the 62 BUs is expected to reach a maximum of 2GB/s – mini merger write to LFS
- rate of files can be as high as 20 streams x 3files/LS, ∼120 files/minute; additionally: bookkeeping/locking files, up to 60 BUs x 20 streams, ∼1200 files/minute
- the macro–merger needs to consume this data online (2GB/s read the fragments, 2GB/s write the final merged file): 4GB/s(!)
- the transfer system is expected to transfer most of the data to Tier0 at 1GB/s*
- overall: LFS needs to guarantee a total of sustained 7GB/s parallel r/w

*Recently this requirement increased to 3GB/s

Mergers

2 available options

- "A"dditive
  - mini–mergers write a file/BU/LS/Stream, macro–merger merges them and makes them available for the TS
  - easy debugging, reliable, "standard" logic
- "C"opyless
  - mini–mergers write in parallel in the final file, macro–merger checks for completion and makes it available for the transfer system
  - reduce the required bandwidth with 4GB/s
  - reduce the number of temporary files by a factor of 60 (number of BUs)
  - fast due to parallel writing in the same file
  - more sensitive to corruption

Lustre File System – Implementation

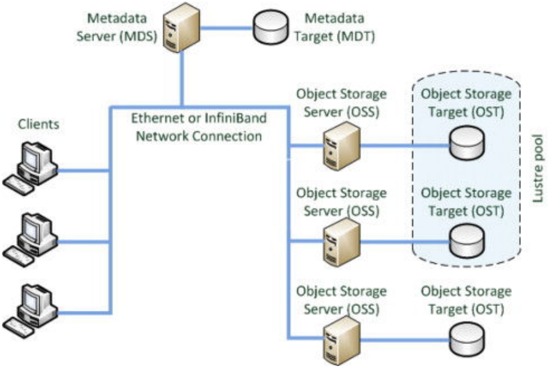

Lustre FS architecture

- current Intel Enterprise Edition for Lustre version: 2.2.0.2
- servers: 6 DELL R720
  - 2 MDS nodes, one active at a time
  - 4 OSS nodes, each controls 6 OSTs

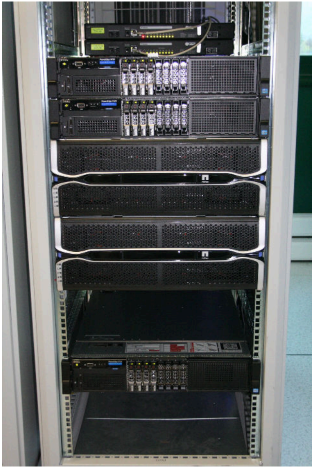

Rack view – MDT (low), 1 OST controller and 1 disk shelves expansion enclosure

Lustre File System – Implementation

Meta–Data Configuration

- 16 drives of 1TB in 1 volume group, 8 hot spares
- only 10% of the disks capacity is used in order to increase performance
- partitions: 10GB for MGT, 1TB for MDT
- connection to servers: Mini-SAS HD to Mini-SAS
- redundancy: RAID6

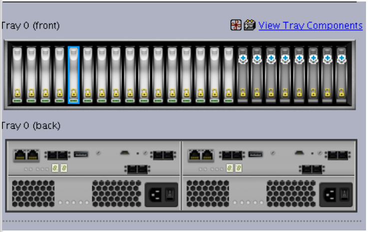

MDT: NetApp E2724 front and rear view

Lustre File System – Implementation

Object Storage Configuration

- 2 OST controllers: NetApp E5560
- each controller manages one disk expansion enclosure DE6600
- each disk shelf enclosure contains 60 disks of 2TB each
- total raw disk space: 240 disks x 2TB = 480 TB
- physical installation: 2 racks, 1 controller and its expansion enclosure per rack
- connection to servers: Mini-SAS to Mini-SAS

Front OST

Disk shelves

Lustre File System – Implementation

OST: Volume Configuration

- each controller/expansion shelf is organized in 6 RAID6 volume groups
- the volume groups are physically allocated vertically to ensure resilience to single shelf damage
- total usable space: 249TB

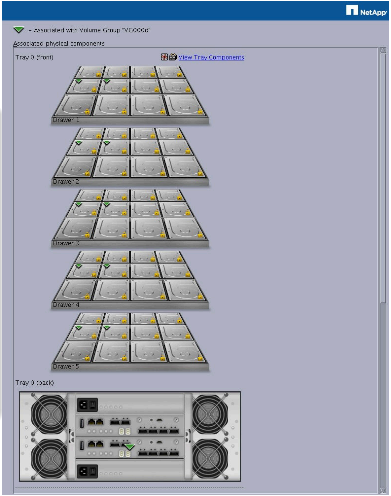

Volumes configuration

Lustre File System – Implementation

High Availability

- volumes repartition to provide full shelf failure redundancy
- all volumes are RAID6
- all devices (controllers, shelves, servers) are dual powered (normal and UPS)
- all servers configured in active/passive failover mode via corosync/pacemaker: MDS in neighbouring racks, OSS within the same rack
- LFS nominal availability: 40GE and InfiniBand (56Gb) data networks*

Lustre File System – Control and Monitoring

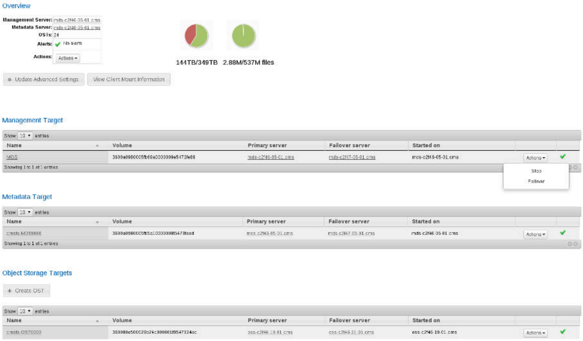

IML: Lustre FS control and monitoring interface

- mostly used for control and base FS operations
- the dashboard provides useful information for debugging an overloaded system
- very demading installation requirements
- not fully reliable: fake BMC monitoring warnings, false status reports upon major FS failures

Lustre File System – Control and Monitoring

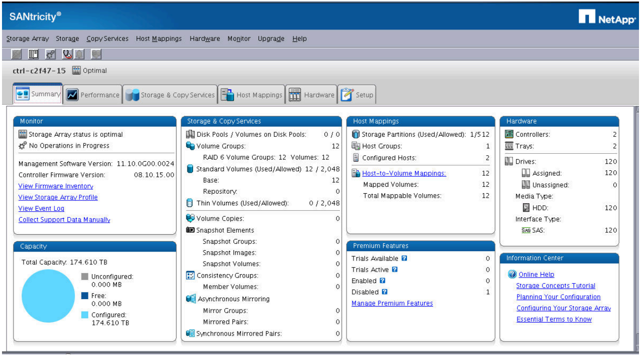

SANtricity

- mostly used for monitoring bandwidth usage per controller
- reports detailed text bandwidth usage per volume
- provides useful information and alerts on hardware status

Bandwidth Validation

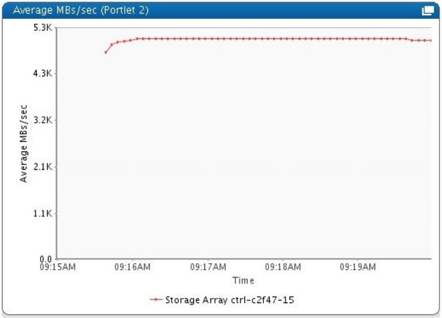

Commissioning Acceptance

Proven steady 10GB/s rate in r/w mode

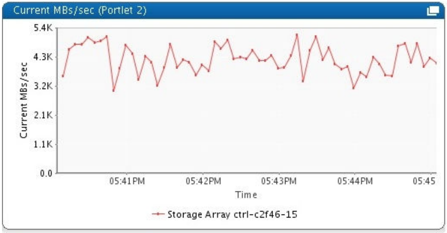

Merger emulation

Proven steady 7.5GB/s rate

Validation

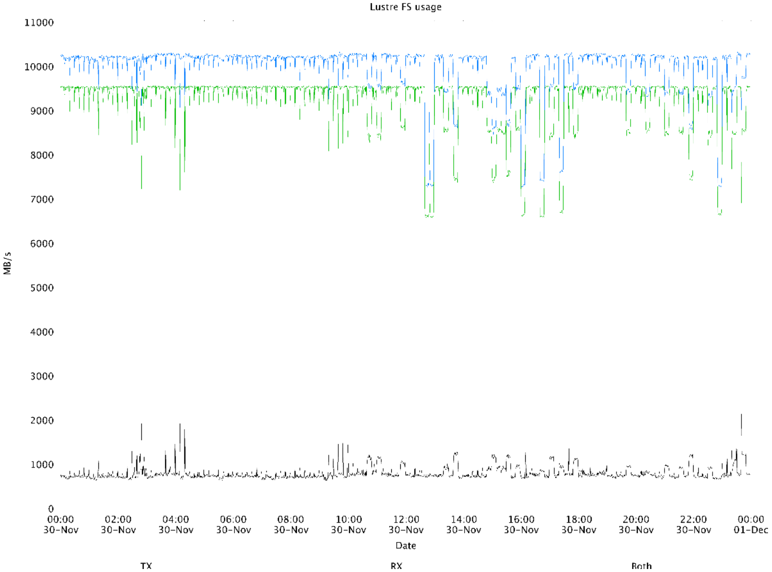

Read vs Write

Validation

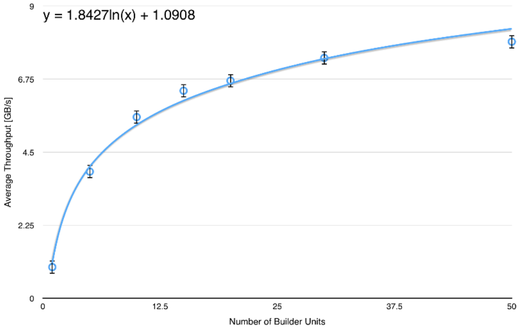

LFS bandwidth benchmarking

Emulation tests using the production computing cluster

- tests performed using different fractions of the available computing farm
- obvious non–linear behaviour with the number of BUs
- transfer system (read operations) were not considered during the tests
- saturation is expected around 8.5GB/s

Real Life Usage

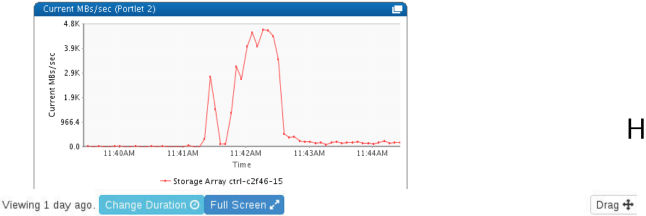

High trigger rate

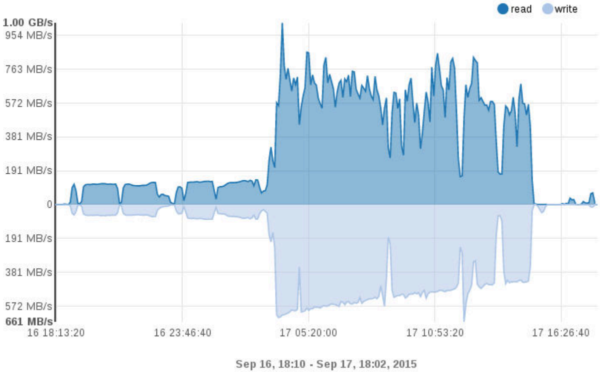

Stable beams runs

Conclusion

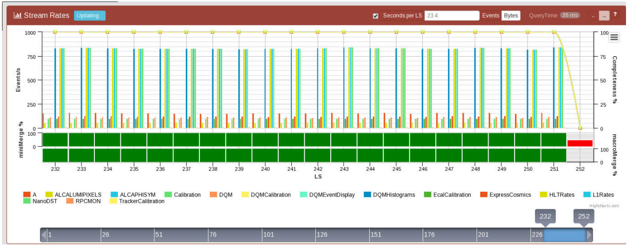

Mergers monitoring sample

Mini Merger

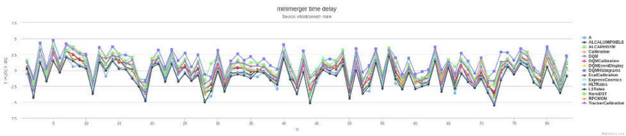

Macro Merger

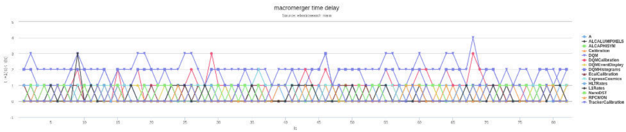

Mergers delays sample

SMTS Validation

- stable behaviour in 5 months of production running mode
- general latencies within the requirements
- proven reliability and availability
- a few glitches, have been followed up and mostly solved

Conclusion

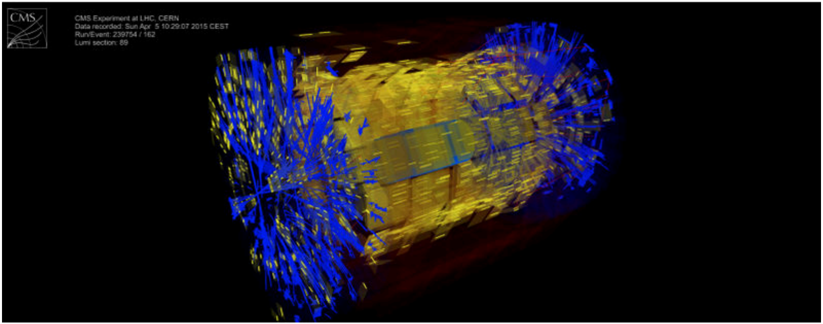

Event display of one of the first particle splashes seen in CMS during Run2

Conclusion

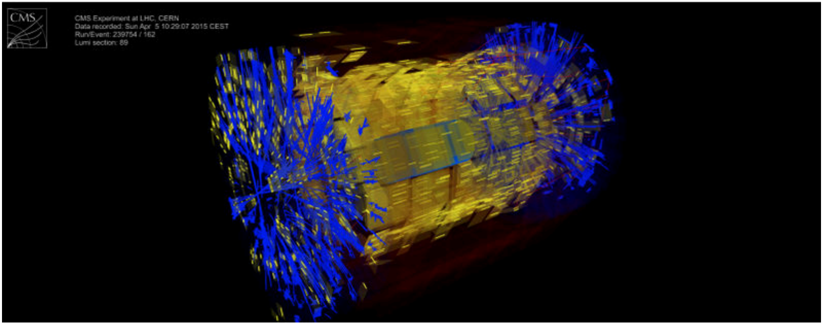

Event display of one of the first particle splashes seen in CMS during Run2

... only a few minutes before one of the OSS servers crashed...

Conclusion

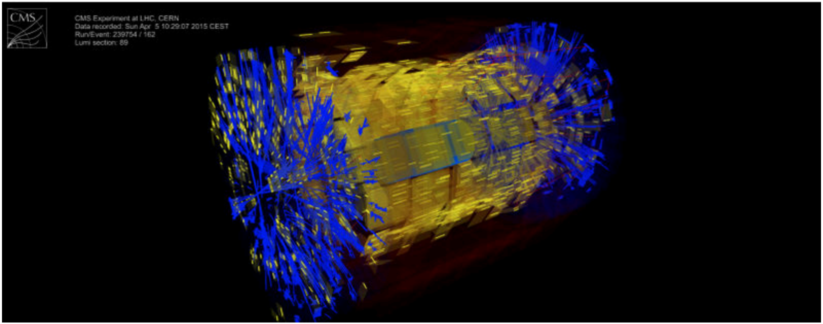

Event display of one of the first particle splashes seen in CMS during Run2

... only a few minutes before one of the OSS servers crashed...
... and the failover mechanism failed ...

Conclusion

SMTS team interaction with Lustre

- Lustre appeared to be very sensitive to network glitches
- IML can be misleading, but provides very intuitive ways of controlling the FS
- sub–optimal application architecture artificially increased the load on the FS (fixed)
- a few FS issues have been identified, but they have been/are being fixed
- clients recover pretty fast and painlessly after FS unavailability
- Intel's Lustre and NetApp's E-Series seem to play nicely together and they deliver the required bandwidth performance

Questions?