


# Operational Experience from the Experiments

July 11, 2011  
Ian Fisk

# 2011

- ➔ Second year of running is a lot more luminosity
  - Each fill is now similar to the data taken in 2010
  - New computing challenges and activity levels

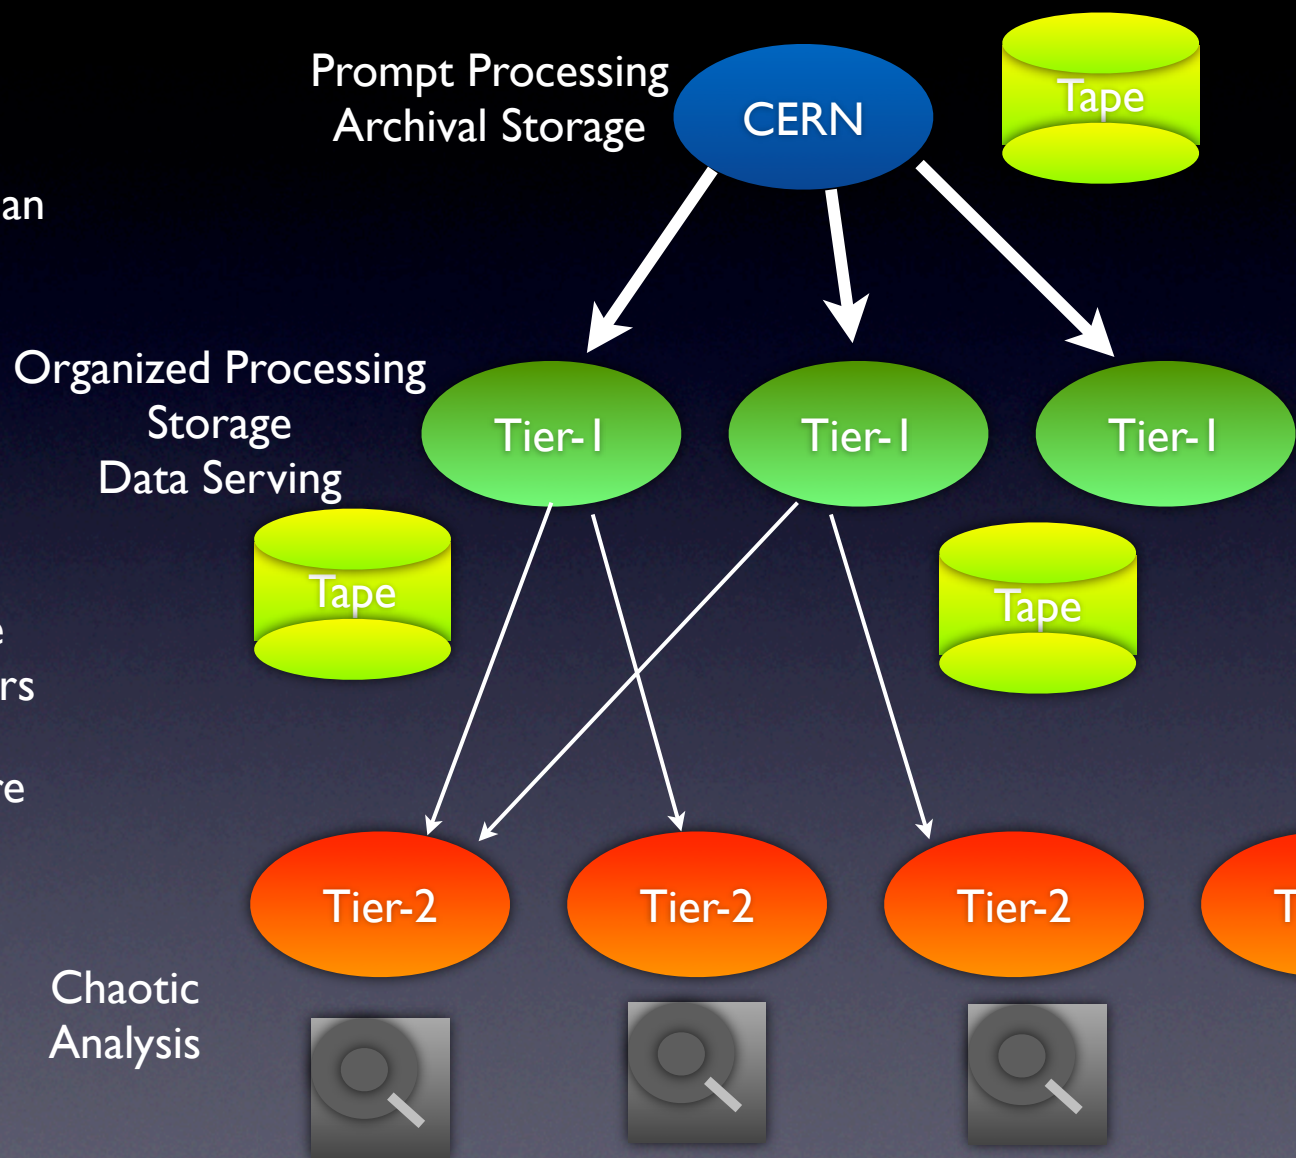




# Refinements

➔ Refinements rather than revolutions

- Indicates the experiment computing models are holding up
- ATLAS using more inter-cloud transfers
- LHCb Making more use of Tier-2s
- CMS moving data Tier-2 to Tier-2
- Etc.



# Generally



- ➔ Generally the WLCG has been a stable and reliable service
  - Components, services and sites are performing well
  - Distributed analysis and processing is a big success
  - We have survived our first year in operations and systems are scaling, including another resource increase in 2011
  - Component migrations in progress
    - ✦ LCG-CE to CREAM CE, changes in BDII architecture, roll-out of glExec



# Tier-0




- ➔ CERN usage is higher than 2010 with more complex events that take longer to process and a harder triggering process
- Still lower than expected
- Built for peak, well utilized for bursts

Usage of CERN




# Tier-0 Data

- LHC Experiments now have PB datasets
- Bigger simulation samples



# Tier-0 to Tier-I

- ➔ Replication of data from CERN has been lower in 2011
- ➔ Event sizes and data volumes are larger
  - Fewer redundant copies





- More stable utilization



# Tier-1 Utilization

- ➔ Average Tier-1 utilization is higher than 2010
  - Good usage by all VOs
  - At the end of the year we will be resource constrained
  - Big increase in resources in April and already well used

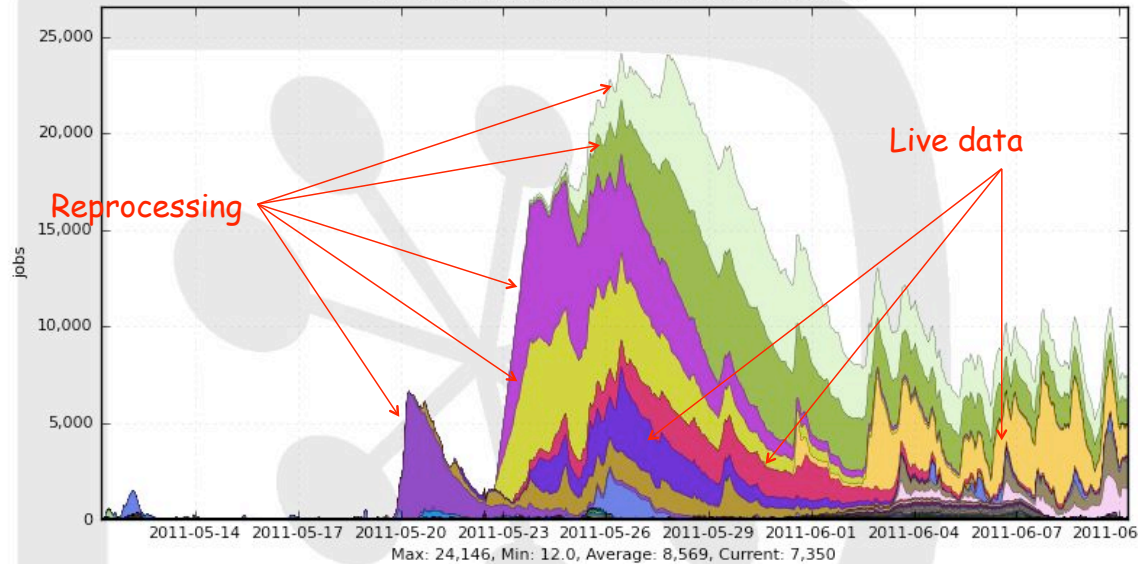




# Tier-1 Reprocessing


Waiting jobs at Tier 1

30 Days from 2011-05-11 to 2011-06-10




Running jobs at Tier 1

30 Days from 2011-05-11 to 2011-06-10



LHCb CPU @ T1

21 Weeks from Week 00 of 2011 to Week 22 of 2011




➔ LHCb


- Heavy use of distributed Tier-1s

# ATLAS Data Processing

Completed Jobs (last 6 months, daily)



CPU time (last 6 months, daily)



Grid global downtime to scale up DB capacity for  $\geq 2011$

- CPU  $\sim 4 \cdot 10^9$  sec daily
- $> 600K$  jobs /day


- Reprocessing of 0.5 PB of select data: HI, Muon, Fast, Tags
- $\sim$  one cycle/month of group reprocessing (NTUPs)
- MC production 2011
- new G4 production: 8400 tasks: 800M events produced
- digi + reco
  - MC10a campaign (2011 data analysis), start beginning of April ( $\sim 2$  weeks)
  - MC10b campaign (2011 data analysis), start end of April ( $\sim 5$  weeks for the bulk, new requests still coming in)
- MC11 new event generation and G4 started: progressing as expected

5 July 2011



# ALICE


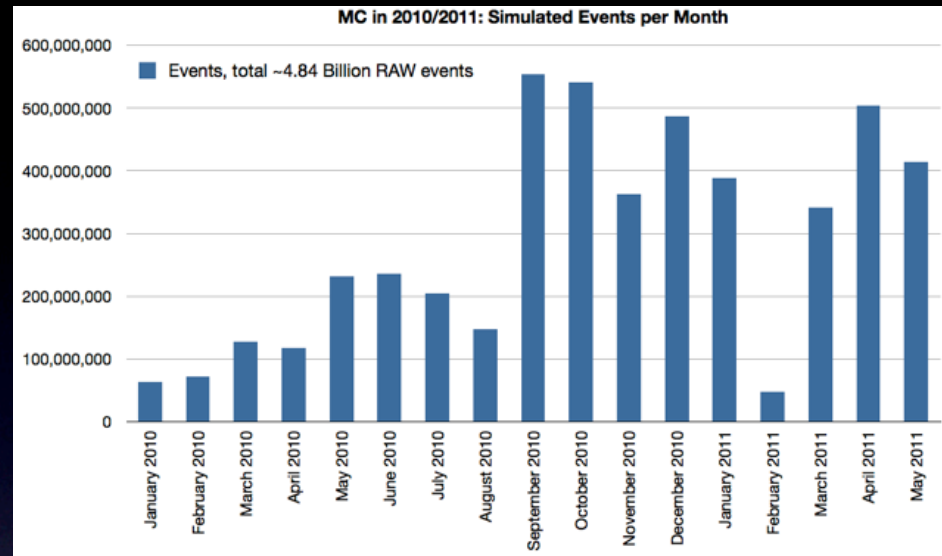
➔ Average of 17k jobs running (2/3rds are Production Jobs)



# CMS MC Production

➔ Planning was for 1.2B Simulated Events to be Completed by the end of May


- Simulation Production was 1.6B, but first 0.4M were 8TeV from before the planning change and were deleted
- Total Re-reconstructed is 1.5B. Roughly half with full out of time pile-up
- ✦ RECO Event size is larger than estimates
  - 0.97MB for GEN-SIM RECO (Compared to 0.63)
- ✦ AOD is very close at 0.25






# CERN Analysis

- ➔ Both ATLAS and CMS has moved to EOS for data access at CERN
  - EOS: high performance highly scalable redundant disk storage, based on xrootd framework
    - ✦ ATLAS tested and stress tested this new solution at the end of 2010 - start of 2011
    - ✦ to improve the CERN reliability for analysis, ATLAS is now migrating the CASTOR disk-only pools to the new EOS disk only storage. CMS has moved the CAF



# Distributed Analysis


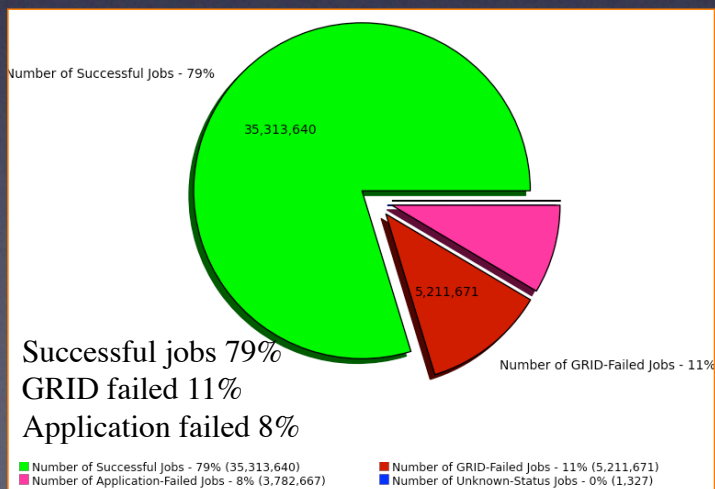

- ➔ Distributed Analysis on the Tier-2s remains an extremely successful application






- ➔ Generally 100% utilization of the pledge on average despite a big increase in T2 resources in April



# ALICE and ATLAS



# CMS





# Popularity

Common Project in IT-ES supporting ATLAS and CMS

Show  entries


Data Tier	Number of access	CPU Time	Number of Users
ALCARECO	37	0	1
AOD	890268	2171413937	830
AODSIM	756939	1038795500	677
DQM	154	0	5
FEVT	20680	0	9
GEN-SIM	20	0	2
GEN-SIM-DIGI-RAW-HLTDEBUG	2023	0	8
GEN-SIM-DIGI-RECO	220280	5475706	17
GEN-SIM-RAW	5050	800378	5
GEN-SIM-RAW-RECO	163	6087	5
GEN-SIM-RECO	291753	84403131	162
GEN-SIM-RECODEBUG	16273	16535210	15
RAW	12857	2894465	25
RAW-RECO	22937	64814428	51
RECO	418386	246586849	272
unknown	1	12	1
USER	197866	88028660	236

Data Tier      Number of access      CPU Time      Number of Users

Can now see what data is being accessed most and how they are accessed, what tiers and what users are looking at things

Allows us to see the transition to AOD

**DataTier Time Evolution**




# In Progress

- ➔ A number of interesting optimizations some of which will be more thoroughly covered
  - Xrootd for remote data access beyond ALICE
  - Network improvements and more site connections with LHCOne
  - Better disk management and cleaning separation of disk and tape.
  - Looking at cloud components and broader use of virtualization






# Distribution

Change from



To



# Outlook

- ➔ First year of running on LHC went well
  - There are some interesting adjustments to improve the efficiency, but the biggest change in 2011 will be resources constraints
    - ✦ Already we see high utilization of Tier-2s and growing Tier-1s
    - ✦ Resources will be a premium and optimization and choices will be needed.