



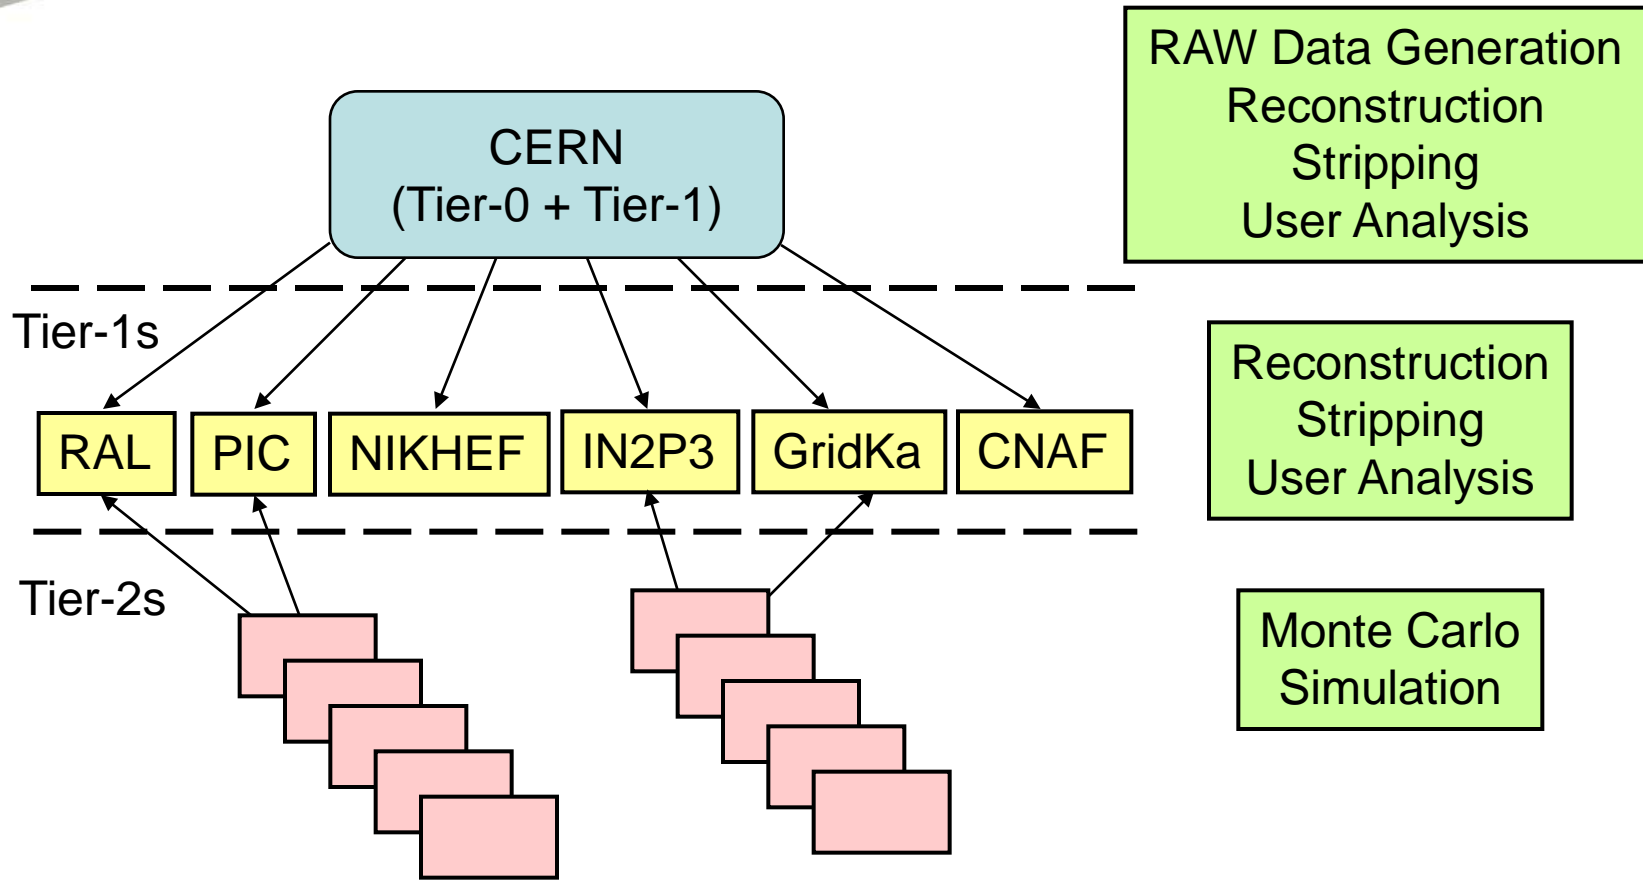
LHCb and Wile E. Coyote

Raja Nandakumar

- Computing model
 - Relevant features
- Highlights of the year
- Dave's questions
 - And LHCb's answers



LHCb Computing model



⇒ For LHCb : CERN = Tier-1++

- ⇒ Computing model enforced by DIRAC
- ⇒ Every file has at least one copy

- ≥2 Tier-1s

- CERN + Tier-1 for RAW data
- We hope so ...

- Files have different lifetimes

- ~2 years for rdst files
 - ⇒ Version of software + workflow used to create it
- RAW data lives "for ever"
- Can vary widely for user data

- ⇒ Failover mechanism for recoverable failures

- LFC, destination SE, ...

- Some problems are unrecoverable

- File availability, input sandbox, cpu / wall time, etc.
- Trap and report these problems if possible






Tier-1s in LHCb

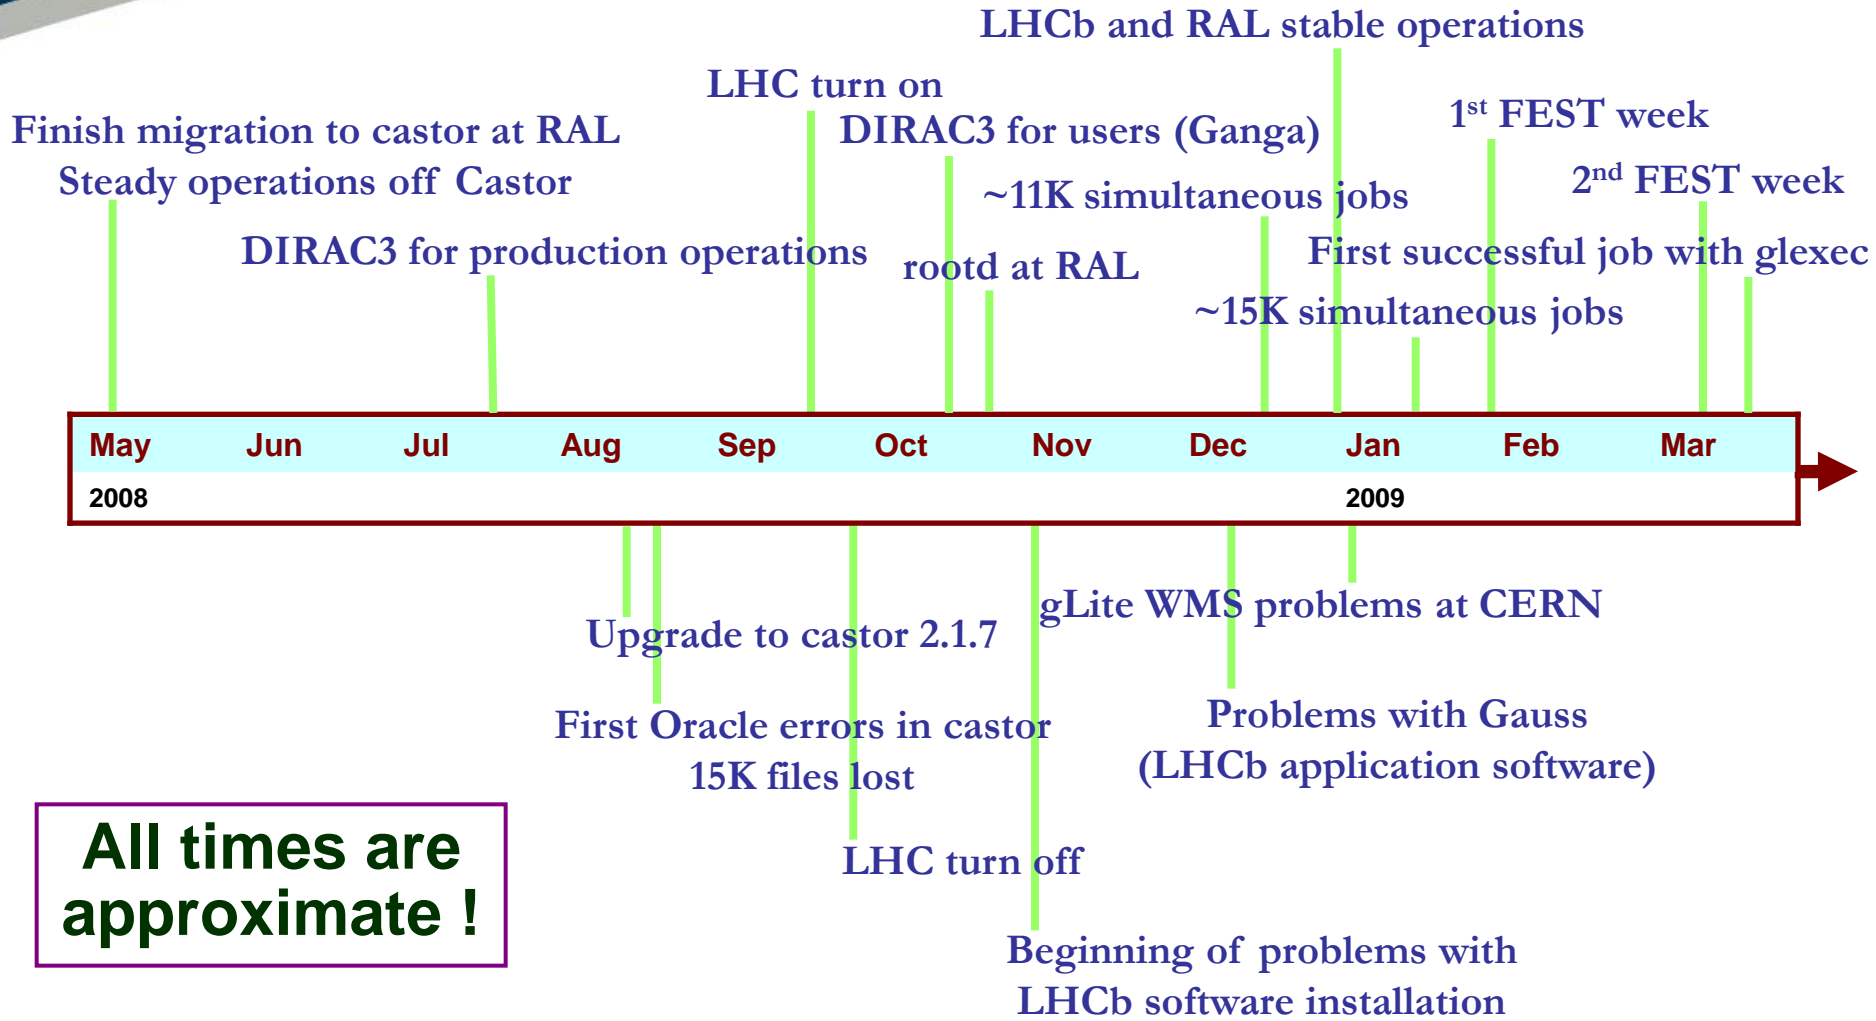
- ⇒ Tier-1s do most of the same work as CERN
 - Except producing RAW data
 - Will also store user data at some level
 - Currently all user data goes to CERN
 - Maybe ~15% of user data at each Tier-1
- ⇒ No MC simulation at Tier-1s
 - Unless there is no other LHCb activity
- ⇒ LFC (readonly), Conditions DB, VO-box at each Tier-1 site
 - VO-box to handle failover requests
 - At present from Production jobs, but from user jobs also in future
 - LFC and Conditions DB replicated from CERN



- ⇒ Tier-2s used for MC simulation (user, production)
 - Mapped to Tier-1s (for productions) for uploading data
 - Failover mechanism if Tier-1 is unavailable
 - Redistributed using FTS if desired
 - Minimal requirements on network, hardware
 - Less than ATLAS, CMS anyway
- ⇒ Some data analysis possible if enough storage is available on Tier-2
 - Storage support provided by local manpower
 - Job submission will be through Ganga/DIRAC as usual
 - ⇒ Will be given assistance to enable site for analysis within DIRAC
 - No global policy on Tier-2 analysis decided yet




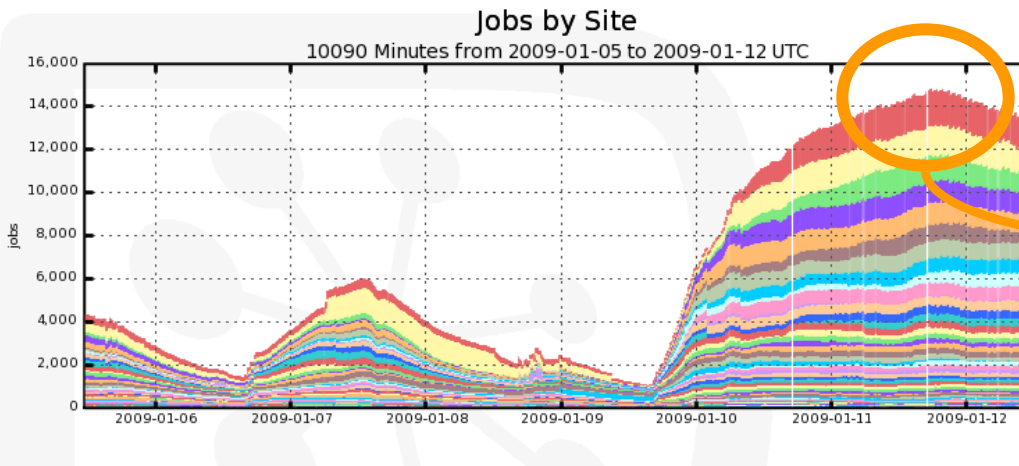
Highlights of last year



All times are approximate !

A few more details

- ⇒ DIRAC reviewed by group from GDB
 - Compliant with security requirements
 - Logs dn of all actions
- ⇒ Currently runs off a single server
 - Quad core, Intel Xeon, 3GHz cpu, 4GB RAM
 - No load problems seen so far



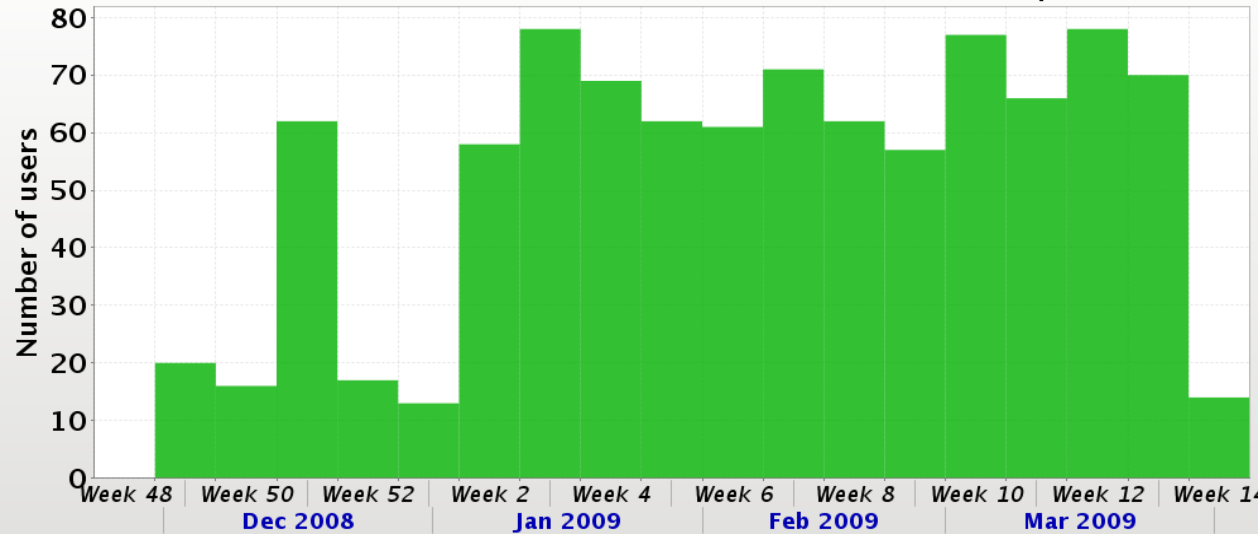
- Ganga is the standard LHCb user interface to the grid
 - Uses DIRAC as a backend for grid jobs
 - Running off DIRAC3 since September 2008
 - Users migrated fast
 - ⇒ Jobs ran faster under DIRAC3
 - ⇒ DIRAC3 is better supported anyway
 - DIRAC2 stopped in Jan 2009
 - ⇒ Services and machines (CERN, PIC) stopped
 - Many improvements in error messages
 - #of users & jobs has increased a lot (next slide)
 - Problems in middleware also found ...
- Ganga committed to implement all use cases for user analysis
 - Ongoing work
 - Increasing use brings more needs to light

User jobs on the grid

Last 4 months

Unique Users

Number of unique users: 220

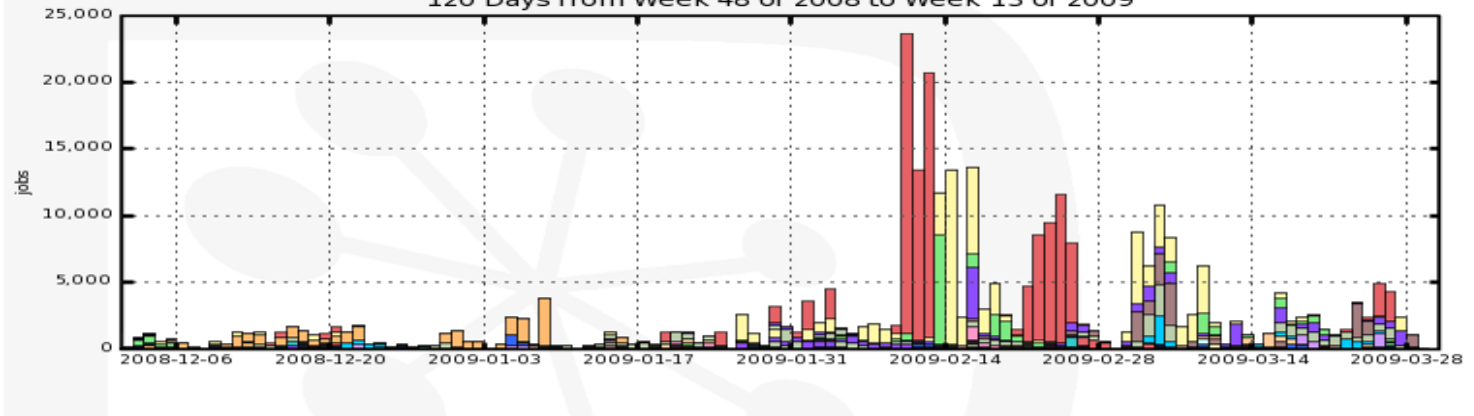


Ganga

Dirac

Jobs by User

120 Days from Week 48 of 2008 to Week 13 of 2009



⇒ Various problems seen

- Fire at CNAF
- Firewall at NIKHEF / SARA
- dCache / staging at IN2P3, NIKHEF
- Lost files at RAL, PIC, GridKa
- Quota / fairshare at CERN
- ...

⇒ Work around these problems

- Hope problems do not hit at the same time
 - At least 1 (out of 7) site up always
 - One occasion in 2006 with only 2 Tier-1s up
- If one / more Tier-1s are down, work with the usable Tier-1s
- If "we" know about problems, jobs should not reach the site
 - Banning within Dirac
 - Site stops accepting jobs and is unattractive in the bdi
- Much happier if the site does not have problems
 - And if so, it recovers asap
 - ⇒ RAL admins have been great in this



Other possibilities

⇒ Problems at Tier-2s

- Seen many of these over the years
 - nfs / sqlite problems most recently
- Does not affect us as much
 - GGUS ticket and ban the site and wait till the site recovers

⇒ Problems with DIRAC

- Been under load for quite some time now
- Some developer usually available at all times

⇒ Problems with LHCb software

- It is an important issue
 - Problems take time to fix
 - ⇒ Developers do not easily understand problems that are reported by a third party (not visible interactively)
 - Many external dependencies also
- Mainly affects production activities
 - User jobs install software locally if needed

Dave's questions - 1

- ➔ How would you (your experiment) handle a significant loss of data in the UK - either custodial T1 data or data at T2s. What is the recovery strategy (timescale and priorities) and what are the bandwidth or other implications?
 - There is no "official" data at the Tier-2s
 - At Tier-1s, we remove the "lost" replicas from the LFC and replicate the data as soon as possible
 - Custodial data : From our computing model
 - 2 replicas of RAW data - losing one is quite bad
 - ⇒ Will need to replicate immediately where ever possible
 - 1 replica of rdst - will be difficult to recreate from old workflows
 - Much prefer custodial data to be really safe
 - Recovery timescale & priority depends on the data
 - Some data more important than others
 - Bandwidth needed depends on the amount of data lost
 - Will start replicating most important data first
 - All such data movement through the FTS channels
 - ⇒ Should be able to handle the needed load
 - ⇒ Unlikely to need more than 1GB/s rates



Dave's questions - 2

→ How would you react to the loss of the compute or storage services? For example, the collapse of the job submission system (due to technical, security, or network issues for example)? What are the priorities? Who coordinates the response (central or local)? What are the implications?

- LHCb submission framework is centrally coordinated
 - Inputs from sites are always welcome
- The priorities are to get jobs running again
 - If needed, ban one or more users to get approval to run again
 - Copy critical data over if it is lost
- DIRAC currently submits jobs through WMS-es
 - Can submit jobs directly to the CE if needed
 - We will need open queues at the sites
- Network problems are recoverable (in a limited sense)
 - Heartbeats for the life time of the job
 - Use of failover mechanisms - jobs at Tier-1s should be able to save data on the same Tier-1 at least
 - Simulation jobs at Tier-2s are not critical and can be resubmitted if needed
 - Some job failures are inevitable



Dave's questions - 3

→ What response would your experiment suggest if we had a massive procurement failure (CPU or Disk is found not-fit-for purpose so that a 6 month delay ensued)? How would fair-shares and priorities be determined and communicated?

➤ Work around the problem once details are known

■ May need some changes to our computing model

⇒ Fewer replicas

■ Would like to be kept informed

⇒ WLCG & operations meetings, experiment contacts, ...

➤ Hope the site works to recover as fast as possible

➤ User analysis is the priority

■ Whatever is needed to make data available to users

⇒ Reconstruction, stripping

⇒ Maybe even move data to ensure this happens

➤ Writing up procedures for recovery (see last slide)



Dave's questions - 4

- How would your experiment plan to cope with an extended outage (>5 days) of the UK Tier-1?
 - This might be the whole site due to an extended power problem; fire; flood etc, or some significant fraction of the site (in terms of services)
 - Work around the problem (Slides 10 - 14)
- Please explain what short and longer term changes would be required to your experiment model and their implications.
 - Reduced number of file replicas
- If possible, please let us know the highest priority services to restore.
 - In order - **SE**, Batch system, Conditions DB, LFC
- Please also address how such an incident gets communicated to and through the experiment management and the users
 - WLCG and daily operations meetings
 - Experiment contacts (me, Glenn Patrick at RAL, Nick Brook at Bristol)

Summary

- ⇒ Aware that problems do happen
 - Have seen a few already and worked around them
- ⇒ Detailed procedures being collected for various items
 - <https://twiki.cern.ch/twiki/bin/view/LHCb/SiteFailureProcedures>
 - <https://twiki.cern.ch/twiki/bin/view/LHCb/ProductionProcedures>
 - <https://twiki.cern.ch/twiki/bin/view/LHCb/DiracProjectPage>
- ⇒ Manpower just as important as data
 - Both within LHCb and at the sites
 - Without effort, data cannot be recovered anyway