


LHCb

DIRAC

COMMUNITY GRID SOLUTION™

LHCb session @ WLCG
workshop






- **Extended physics scope**
 - Charm physics in addition to beauty
 - Trigger rate increased: 3000 evt/s to 4500 evts/s
- **Higher energy**
 - 15% higher multiplicity
- **Pileup**
 - Slightly higher than 2012 ($4 \cdot 10^{32} \text{ cm}^{-2}\text{s}^{-1}$ instead of 3.5 for same number of bunches)
 - ☆ ... but much better than early 2011
- **Reconstruction SW**
 - Better tracking than 2011
- **Consequences:**
 - More tracks, therefore more CPU time for reconstruction
 - More combinatorics, therefore more CPU time for stripping
- **Higher compression**
 - Less space used but slightly more CPU for writing
- **Stripping also MC data (space reduction, not CPU!)**



Resources for real data in 2012

○ Reconstruction:

- average 25 HS06.s (expected 12), some tails
- Memory: some spikes due to the heap for some events
 - ☆ Is it possible to allow temporary spikes in memory?
 - ☆ It is released at the end of the event



○ Stripping:

- average 3 HS06.s (expected 1.2)
- High memory consumption:
 - ☆ mostly due to ROOT buffers (100 MB per output stream, 14 streams...), but also complexity of the application (800 algos!)
 - ☆ Needs around 3.5 GB
 - ☆ Working on reducing the memory footprint
 - * Hopefully this is going to go down (but code from physicists! :-)

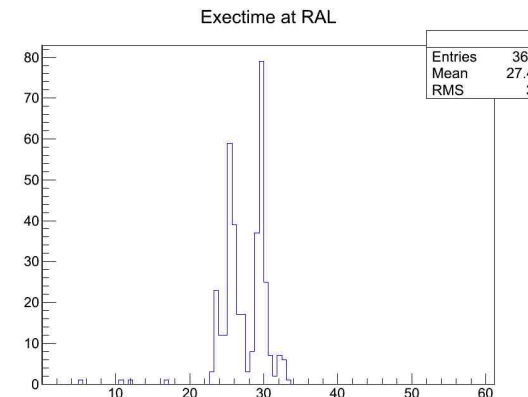
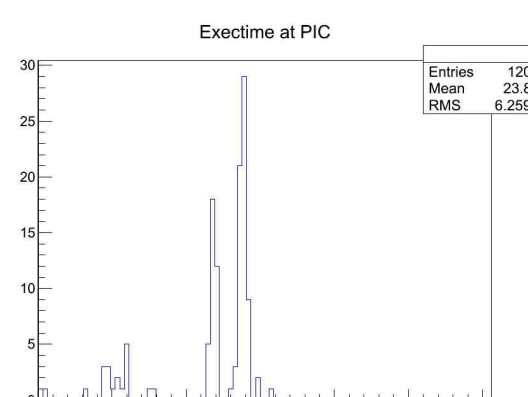
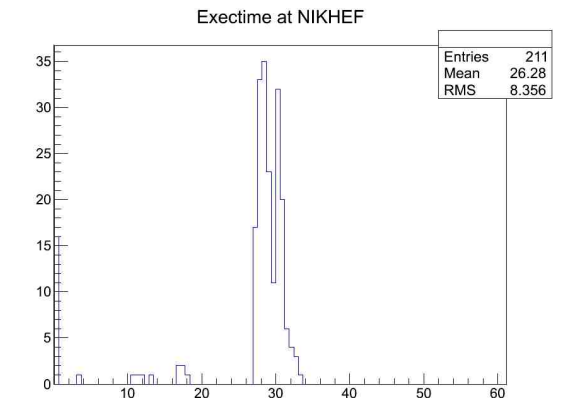
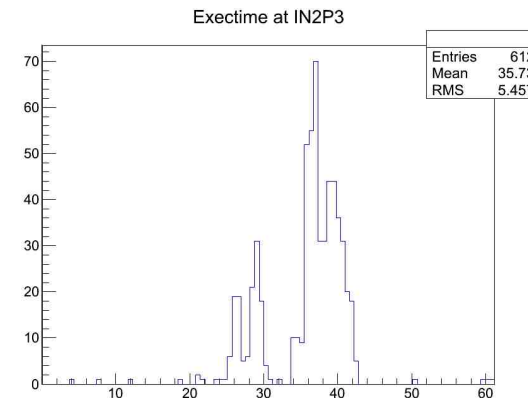
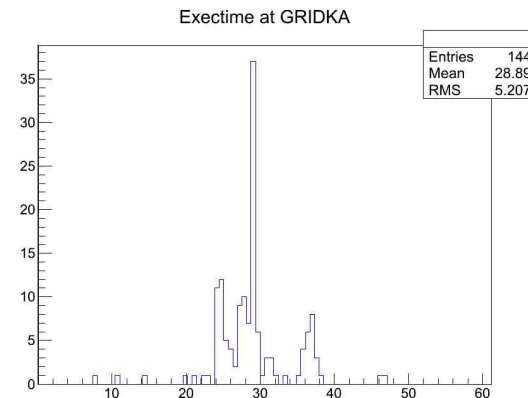
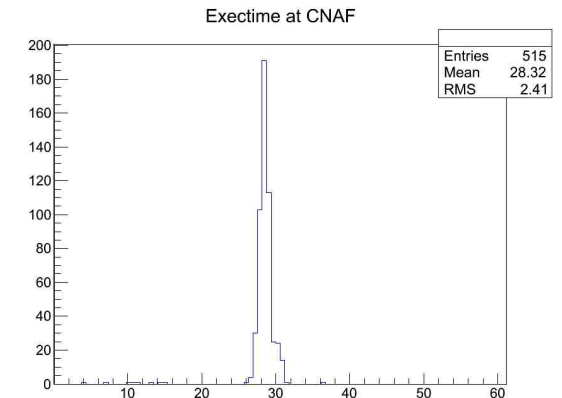
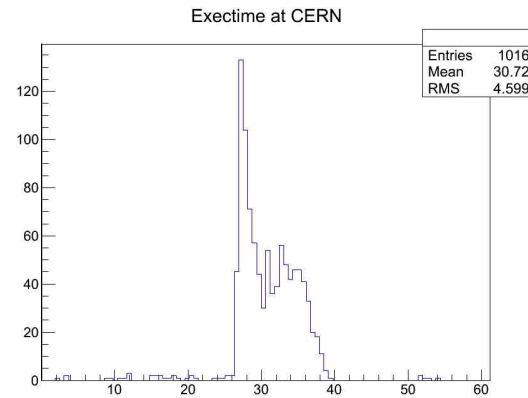
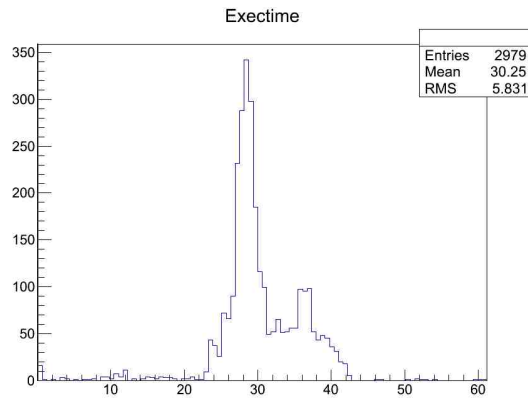


CPU & Memory requirements

- Reconstruction is the most demanding
 - After a lot of effort on reconstruction SW
 - 50,000 events per file (3GB): 1.5 MHS06.s
 - ☆ For a 10 HS06 core, this means 42 hours
- WN normalisation
 - We need to know how much time is left in the queue
 - WN normalisation is required
 - ☆ Estimated by DIRAC
 - ☆ Why isn't it published?
 - ☆ Each batch system is different, each site is different
 - ☆ We also need it for our internal accounting
 - * Should be identical to the normalisation used for accounting
- Memory
 - What counts is whether the process is swapping
 - Instantaneous VM usage is not necessarily a good metrics
 - In any case only the offending process should be killed
 - ☆ Catchable signal, allowing recovering information



CPU for reconstruction

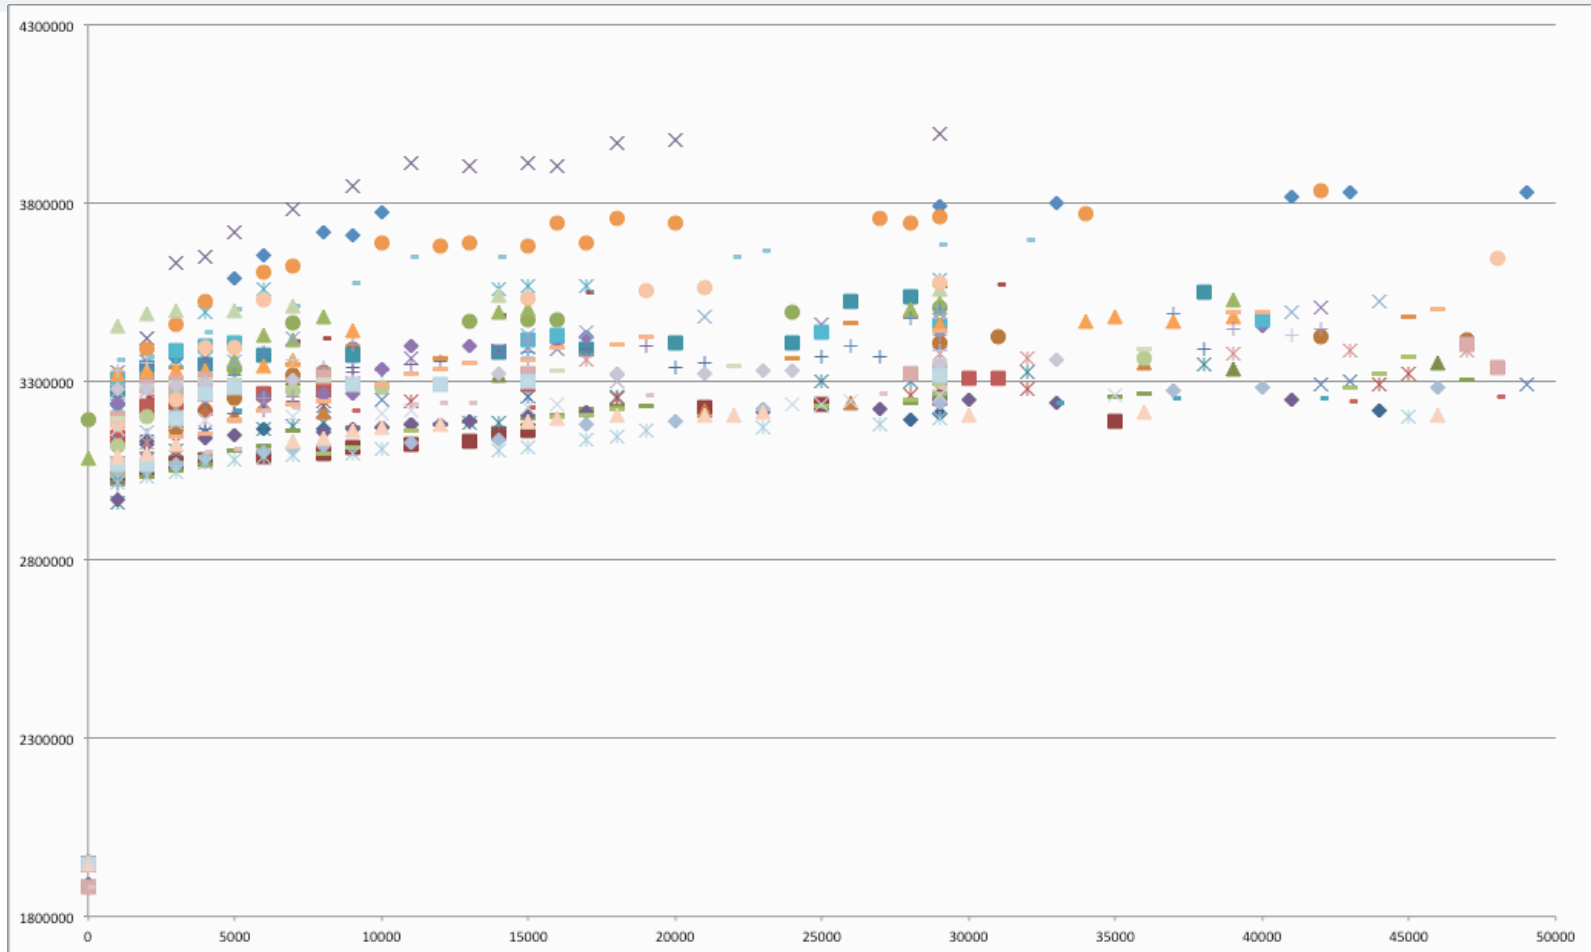


Maximum 45 hours

- CERN & IN2P3
- 10% to be added for monitoring appl.



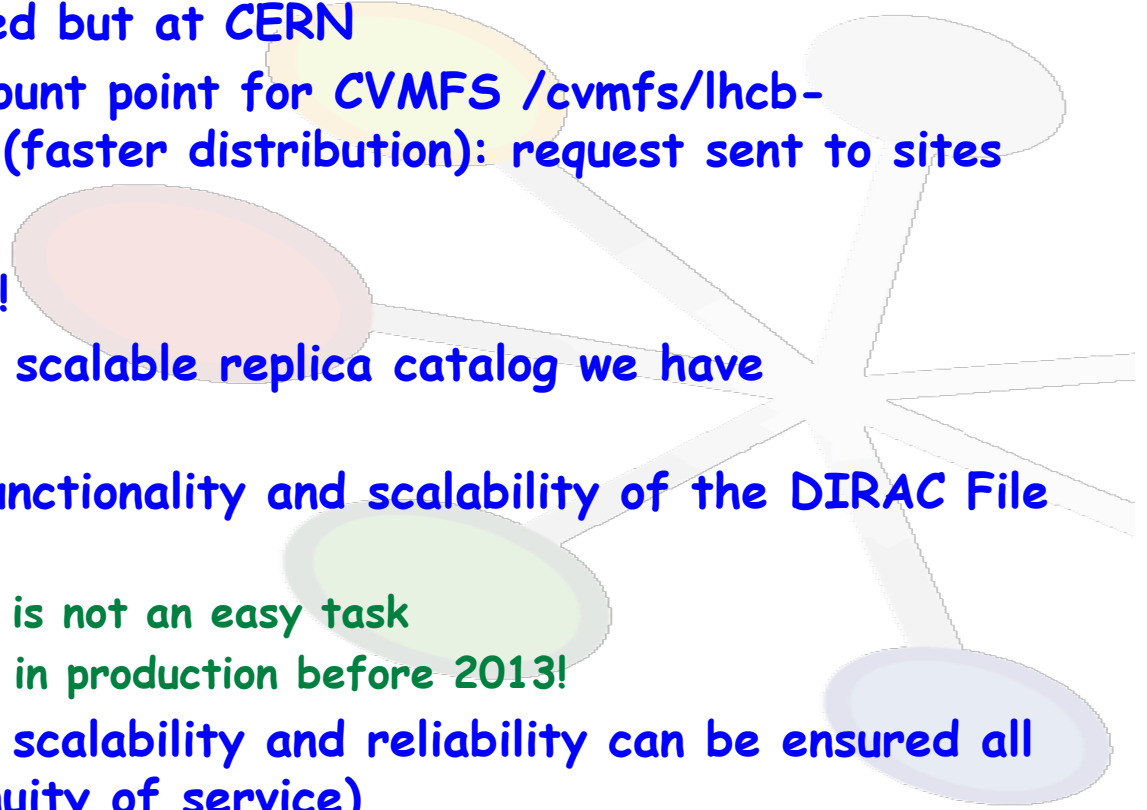
Memory for stripping jobs



- Application: < 2 GB
- Jump by ~1.5 GB due to ROOT buffers (14 output streams)
- Slow increase related to event buffer usage (large events)



- **Condition DB**
 - Currently at each Tier1, replication using 3D (Oracle streams)
 - Distribution moved to snapshots on CVMFS and web portal
 - No longer needed but at CERN
 - However new mount point for CVMFS `/cvmfs/lhcb-conddb.cern.ch` (faster distribution): request sent to sites
- **LFC**
 - We still need it!
 - This is the only scalable replica catalog we have commissioned
 - Going to test functionality and scalability of the DIRAC File Catalog
 - ☆ Migration plan is not an easy task
 - ☆ Therefore not in production before 2013!
 - Should check if scalability and reliability can be ensured all at CERN (continuity of service)
 - ☆ If successful, possibly decommission Tier1 LFCs





- **CVMFS is a very successful experience!**
 - Brilliant example of simplification (that works!)
 - Last problem fixed and deployed at CERN within less than one week (cache corruption pb!)
- Thanks to developers and to sites who have adopted it already!
 - Remember you can also decommission our NFS deployment area :-)
- We would like that **ALL sites move to CVMFS!**
 - Simple to put in place and configure
 - Reduces considerably the local support
 - How can we achieve this? Does it require coordination/help between sites?

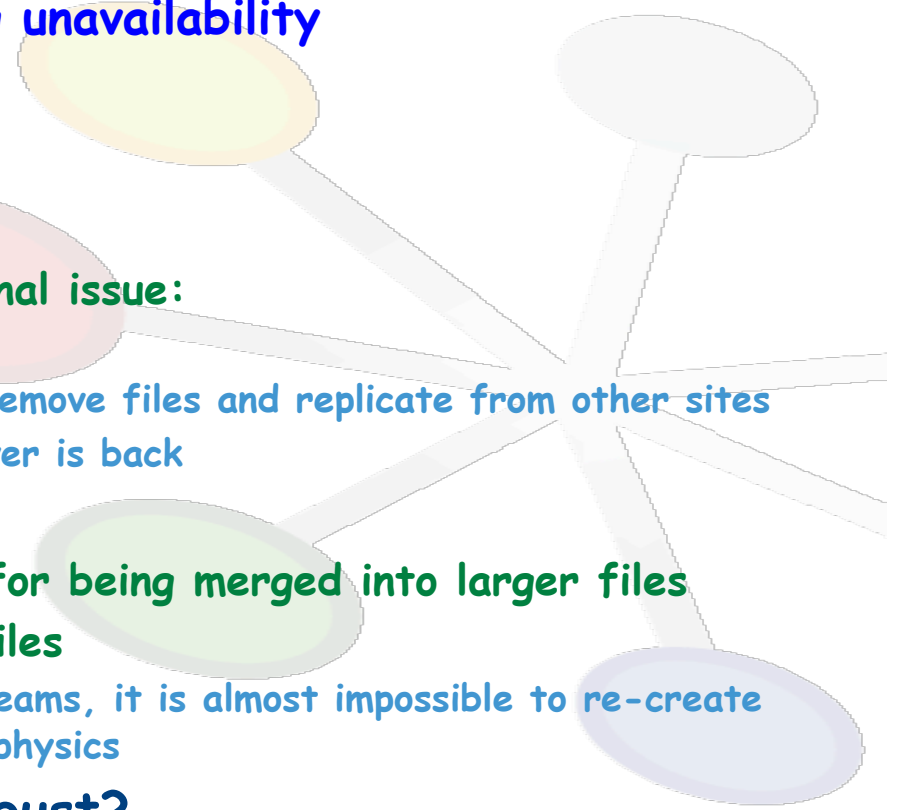




- Currently use: Castor (CERN, RAL), dCache (KIT, IN2P3, NL-T1, PIC), StoRM (CNAF)
- DIRAC heavily relies on SRM
 - As does the current FTS implementation
- Simplification in Space tokens
 - LHCb-Tape, LHCb-Disk, LHCb-USER
- Issues with custodial storage usage
 - DIRAC stager is "job-driven"
 - ☆ Files are staged when jobs are submitted
 - ☆ Jobs can be matched by pilot only when file is staged
 - Currently pinning for a fixed time (24 hours)
 - ☆ Implement internal "cache management"
 - Depends heavily on the cache configuration and job rate
 - ☆ Not trivial to tune
 - Can we simplify the cache architecture (up to 3 sets of pools at some sites: write, recall and read pools)?
 - StoRM (CNAF) by far the most flexible:
 - ☆ All spaces located on the same set of disk servers (across)



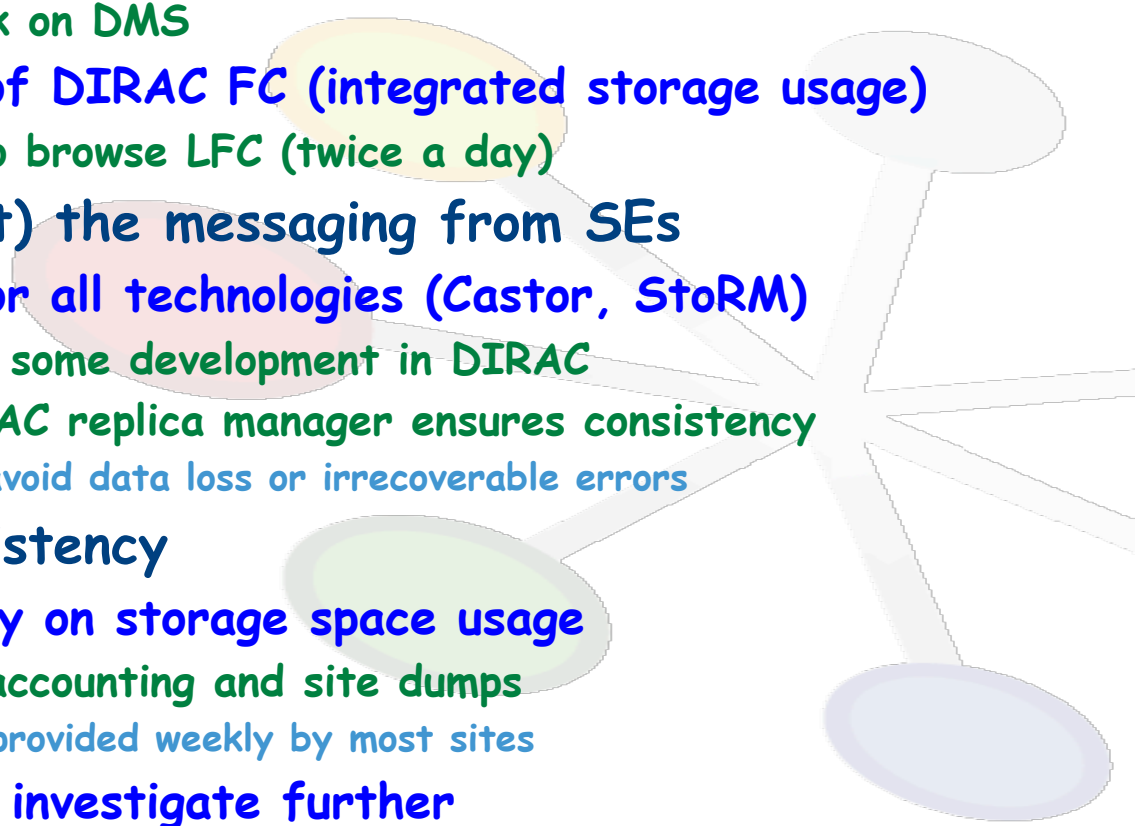
- Disk server reliability is an issue
- Larger and larger disk space per server
 - Many files unavailable when a server does offline
 - Several cases of very long unavailability
 - ☆ Up to irrecoverable data
- Two cases:
 - Replicated data
 - ☆ This is mostly an operational issue:
 - * Flag files in LFC
 - * If server irrecoverable: remove files and replicate from other sites
 - * Else: re-enable when server is back
 - Temporary files
 - ☆ Created by jobs, waiting for being merged into larger files
 - ☆ Single instance of these files
 - * As jobs produce many streams, it is almost impossible to re-create these files: just lost for physics
- Should servers be more robust?
 - Or TOD2 spaces for scratch space? (investigated on EOS)





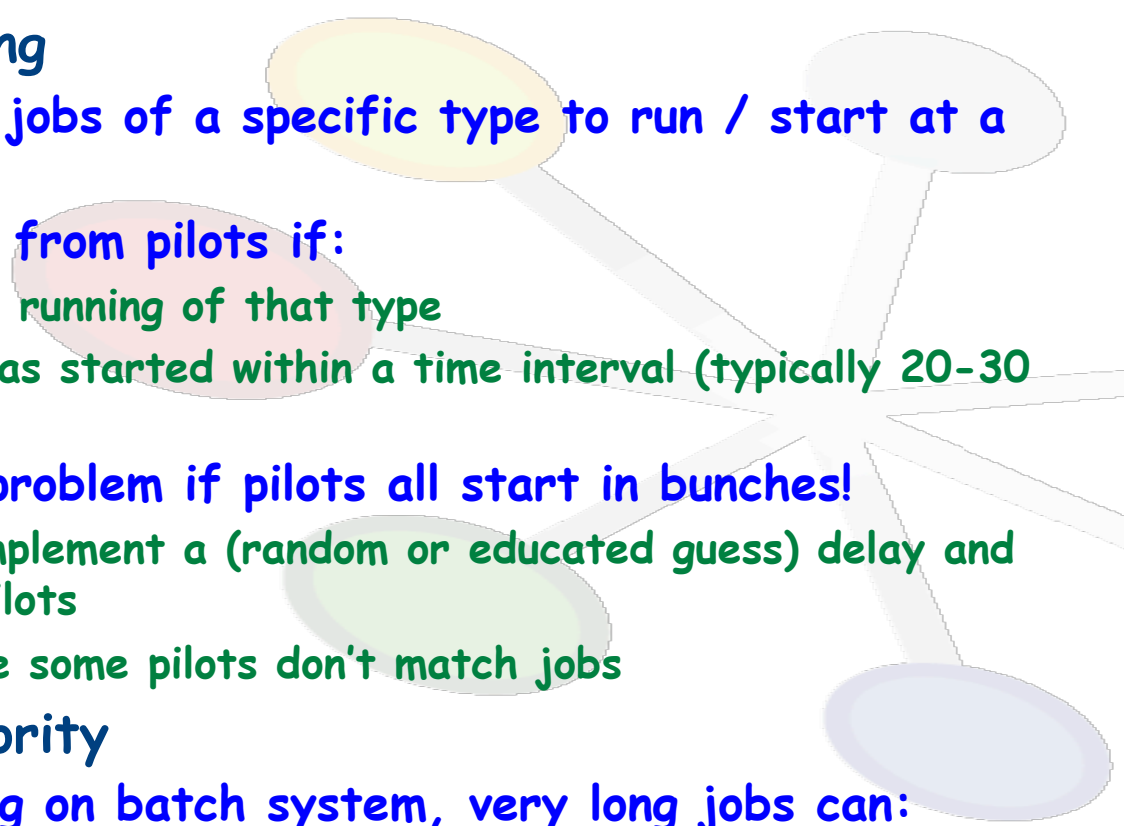
Storage and Catalog synchronization

- Storage monitoring
 - Keep storage usage per directory, and link to datasets
 - Accounting as a function of time (per site, dataset...)
 - ☆ C.f. CHEP talk on DMS
 - Study the use of DIRAC FC (integrated storage usage)
 - ☆ Avoids need to browse LFC (twice a day)
- We don't use (yet) the messaging from SEs
 - Not available for all technologies (Castor, StoRM)
 - ☆ Requires quite some development in DIRAC
 - ☆ Currently DIRAC replica manager ensures consistency
 - * But cannot avoid data loss or irrecoverable errors
- SE \leftrightarrow LFC consistency
 - First consistency on storage space usage
 - ☆ Between LFC accounting and site dumps
 - * Site dumps provided weekly by most sites
 - If discrepancy, investigate further
 - ☆ No way to get an exact match (no instantaneous picture)





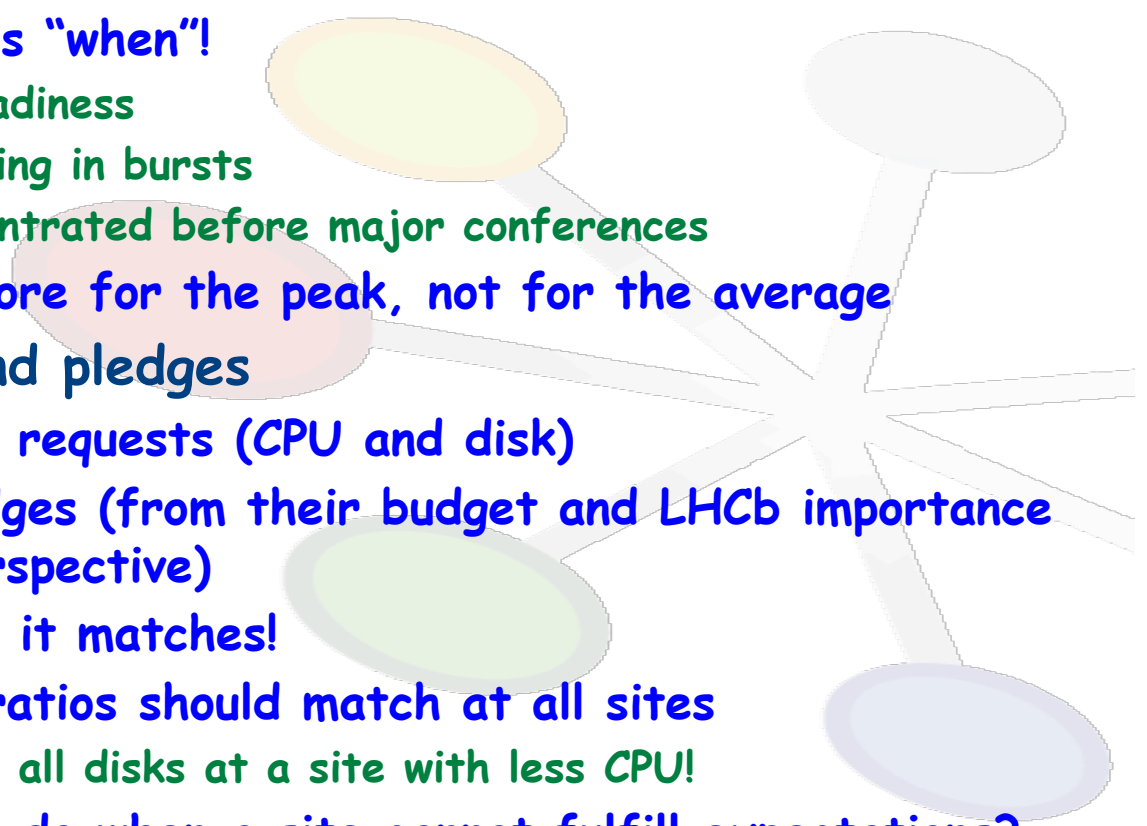
- Moving to direct CREAM submission
 - When done: no need for gLite WMS any longer
 - Mostly internal operational issue: tune pilot submission parameters
- Matching throttling
 - Avoid too many jobs of a specific type to run / start at a site
 - Not match jobs from pilots if:
 - ☆ Too many jobs running of that type
 - ☆ Another job was started within a time interval (typically 20-30 seconds)
 - This may be a problem if pilots all start in bunches!
 - ☆ We need to implement a (random or educated guess) delay and retry in the pilots
 - ☆ Currently quite some pilots don't match jobs
- Batch system priority
 - Issues depending on batch system, very long jobs can:
 - ☆ Reduce priority (LSF @ CERN, CNAF)
 - ☆ Increase priority (SGE at IN2P3)





Resource requirements and pledges

- CPU work and power:
 - What we know best is the “work” to be achieved
 - Some knowledge on how long it should take (reprocessing, restripping)
 - Less well know is “when”!
 - ☆ Application readiness
 - ☆ MC needs coming in bursts
 - ☆ Analysis concentrated before major conferences
 - Power is therefore for the peak, not for the average
- Tier1 requests and pledges
 - We make global requests (CPU and disk)
 - Sites make pledges (from their budget and LHCb importance from a local perspective)
 - Then we expect it matches!
 - CPU/disk/tape ratios should match at all sites
 - ☆ Useless to get all disks at a site with less CPU!
 - What should we do when a site cannot fulfill expectations?
 - ☆ Or delay installation of pledges





- LHCb is quite happy with all sites' responsiveness and stability

- **Warm thanks to everybody!**

- However...

- The system is fragile!

- Small hiccoughs can have severe consequences

- Quite manpower intensive to follow all sites

- ☆ LHCb has a centralised operations team

- ☆ Only LHCb people at Tier1s

- The more we can simplify, the better, however...

- ☆ Don't forget any change requires software changes and should be well prepared and scheduled (SRM, LFC,...)

