



What LHCb wants
from your storage.

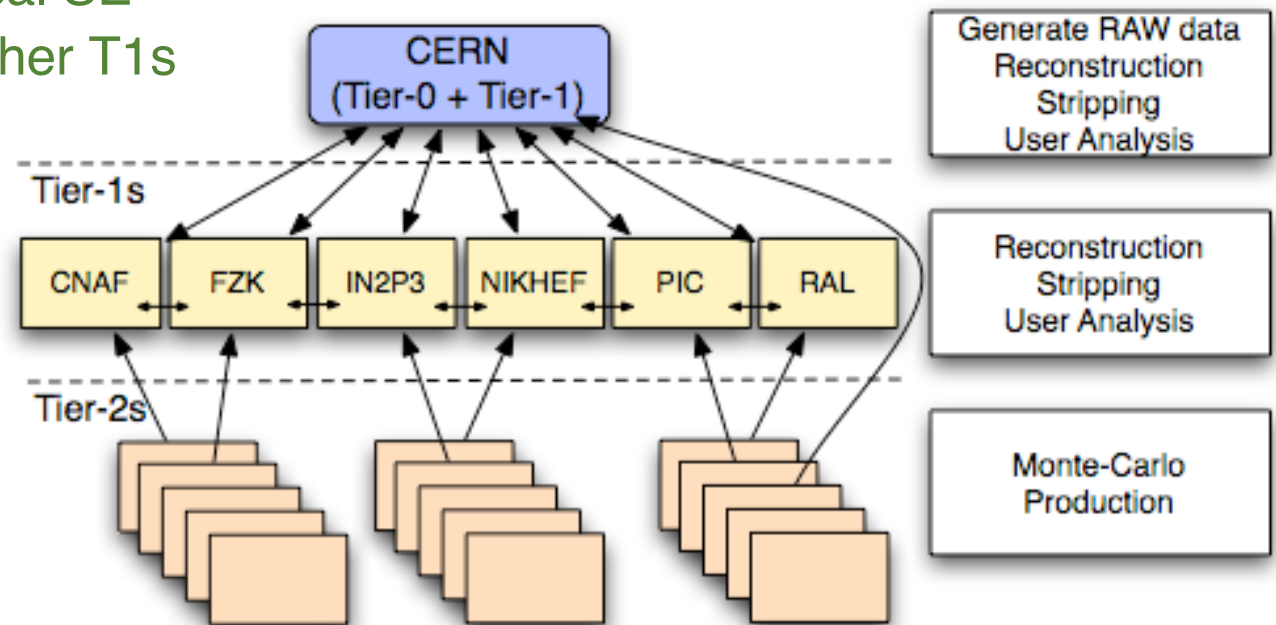
Andrew C. Smith - University of Edinburgh





Computing Model (real data)

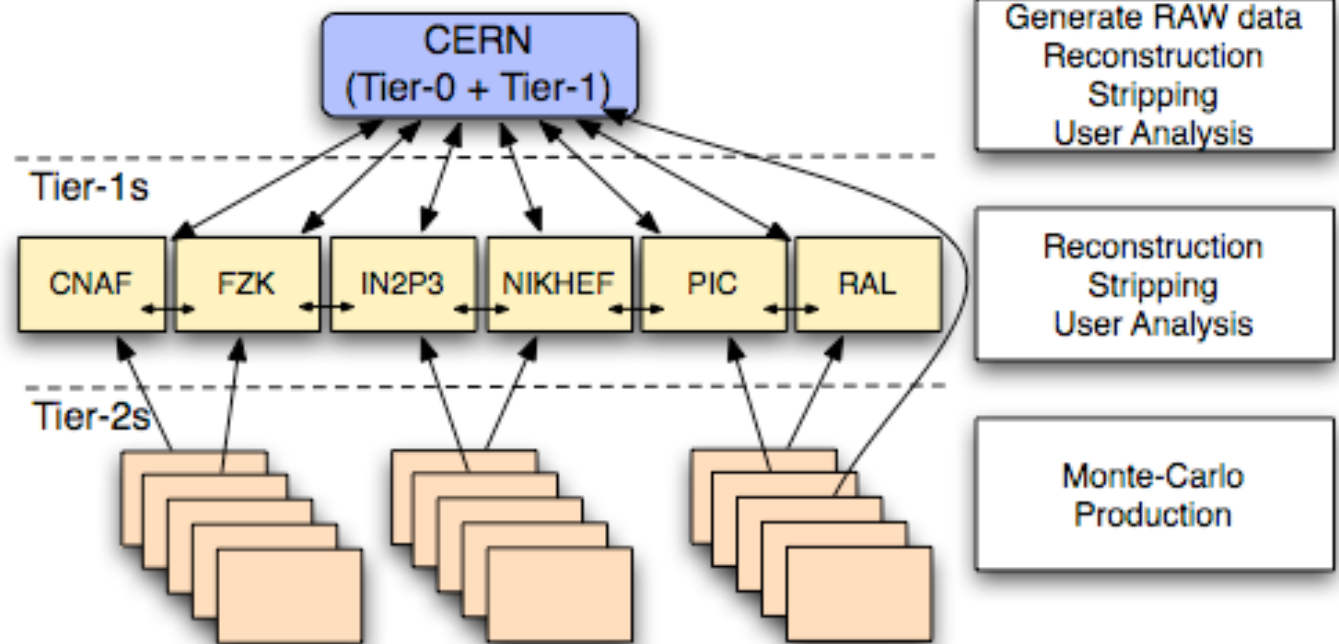
- 70MB/s RAW data from LHCb online system to Castor
- RAW data is replicated to 6 LHCb Tier1s
 - Split by computing share
 - Each RAW file has 1 distributed copy
- RAW reconstructed (per computing share)
 - This is done at Tier1s and CERN
 - Output (rDST) uploaded to local SE
- rDST stripped
 - Output DST to local SE
 - Replicated to 5 other T1s
 - and CERN
- DST for analysis





Computing Model (simulated data)

- Simulation performed at Tier2
 - Output (DST) to associated Tier1 SE
 - Replicated to another T1 and CERN
- DSTs input for analysis





- Real RAW transferred to 'LHCb_RAW' space
 - CERN
 - LAN access for writing data
 - WAN access for replicating
 - Tier1s
 - WAN access for writing data
 - Read from reconstruction job
 - Get TURL from SRM
 - Data read by application using supported local protocol

- rDST written from WN to 'LHCb_rDST' space
 - Using lcg-utils
 - Get TURL from SRM
 - gridFTP used for transport
 - Read from stripping job (along with ancestor RAW)
 - Get TURL from SRM
 - Data read by application using supported local protocol



- DST written from WN to 'LHCb_M_DST' space
 - Using lcg-utils
 - Get TURL from SRM
 - gridFTP used for transport
 - Replicated to 6 Tier1s using FTS
 - WAN access for reading
 - Read from analysis jobs
 - Get TURL from SRM
 - Data read by application using supported local protocol

- DSTs transferred from Tier1s to 'LHCb_DST' space
 - Replicated using FTS
 - WAN access for writing
 - Read from analysis jobs
 - Get TURL from SRM
 - Data read by application using supported local protocol



- MC DST written from Tier2 WN to Tier1 'LHCb_MC_M_DST' space
 - Using lcg-utils
 - Get TURL from SRM
 - gridFTP used for transport
 - Replicated to another Tier1 using FTS
 - WAN access for reading
 - Read from analysis jobs
 - Get TURL from SRM
 - Data read by application using supported local protocol

- DSTs transferred from Tier1s to 'LHCb_MC_DST' space
 - Replicated using FTS
 - WAN access for writing
 - Read from analysis jobs
 - Get TURL from SRM
 - Data read by application using supported local protocol



- LHCb_RAW
 - Tape1Disk0
 - Local access
 - CERN for writing from pit
 - Read access from WN
 - WAN access
 - Read at CERN
 - Write at Tier1s
- LHCb_rDST
 - Tape1Disk0
 - Only Local access
- LHCb_M_DST
 - Tape1Disk1
 - WAN read access
 - Local read access
- LHCb_DST
 - Tape0Disk1
 - WAN write access
 - Local read access





- LHCb_MC_M_DST
 - Tape1Disk1
 - WAN read access
 - Local access

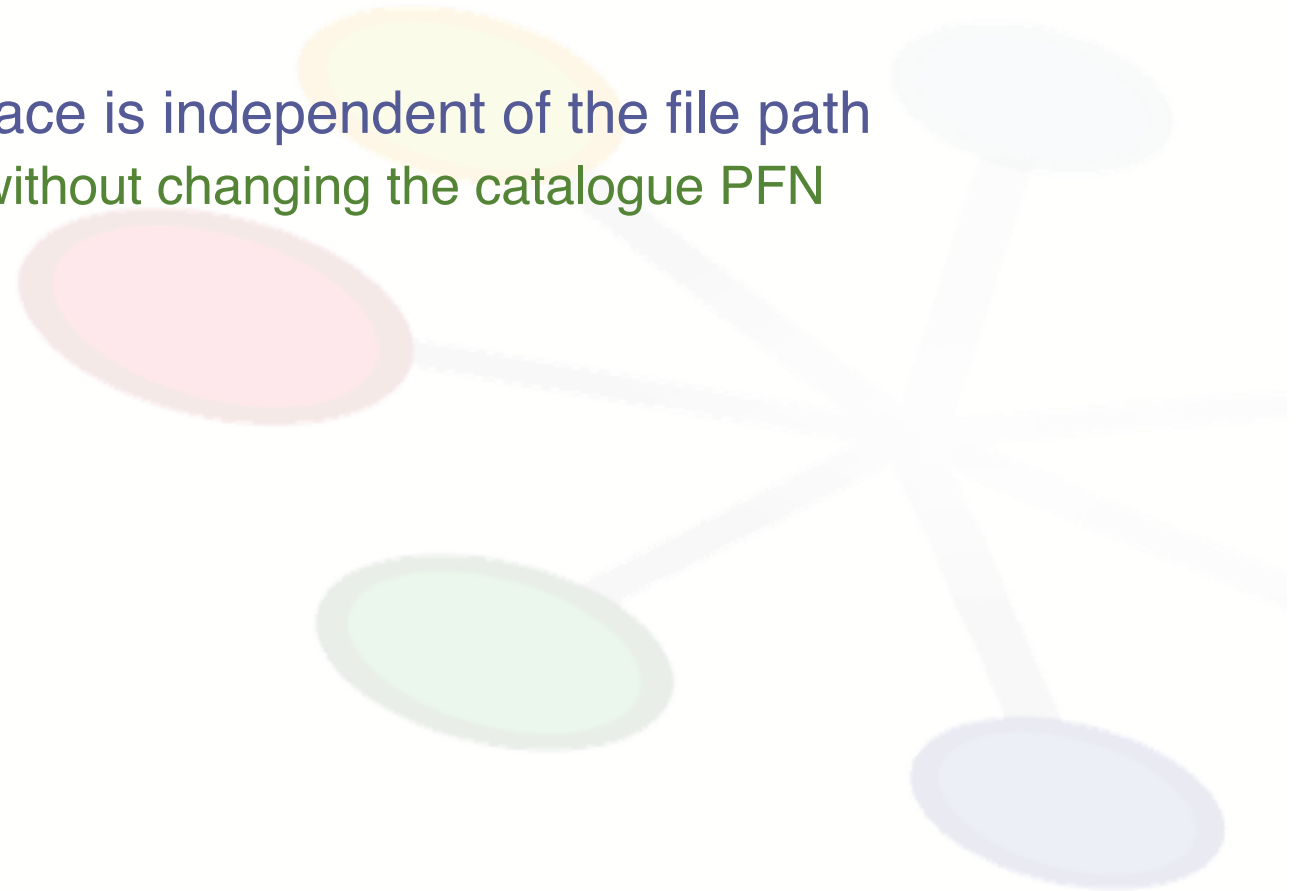
- LHCb_MC_DST
 - Tape0Disk1
 - WAN write access
 - Local access

- LHCb_USER
 - Tape1Disk1
 - Local access





- LHCb requires space is independent of the file path
 - Migrate space without changing the catalogue PFN






- 'LHCb_RAW' space
 - Files 'pinned' at write time
 - Accessed by reconstruction job
 - Pin released by reconstruction job

- 'LHCb_rDST' space
 - Files 'pinned' at write time
 - Accessed by stripping job
 - Pin released by stripping job

- 'LHCb_*_DST' space
 - Files on Tape*Disk1
 - No 'pin' required





- Reprocessing (1 re-reconstruction, 2 re-stripping)
 - Requires access to Tape1Disk0 files
 - Performed three times a year
 - O(month) after initial write
 - Probably cleared from cache

- Re-reconstruction
 - Files in 'LHCb_RAW' to be 'pre-staged'

- Re-stripping
 - Files in 'LHCb_RAW' and 'LHCb_rDST' to be 'pre-staged'
 - 'LHCb_RAW' files staged for re-reconstruction reused



- The DSTs become outdated after re-processing
 - The outdated DSTs are no longer needed on disk
 - Replica copies (LHCb_DST, LHCb_MC_DST) are removed
 - Master copies (LHCb_M_DST, LHCb_MC_M_DST) kept
 - Removed from disk
 - Need to perform Tape1Disk1 -> Tape1Disk0 migration



- LHCb_RAW
 - WAN access required for writing data into Tier1s
 - Also local access required for processing
 - Should there be distinct pools? Local access triggers disk-to-disk copy?

- LHCb_*_DST
 - WAN access required to replicating files
 - LHCb_M_DST/LHCb_MC_M_DST
 - Read access for replication
 - LHCb_DST/LHCb_MC_DST
 - Write access for replication
 - Local read access required for chaotic analysis activity
 - Should there be distinct pools? Local access triggers disk-to-disk copy?



LHCb GSSD twiki:

<https://twiki.cern.ch/twiki/bin/view/LCG/GSSDLHCB>

