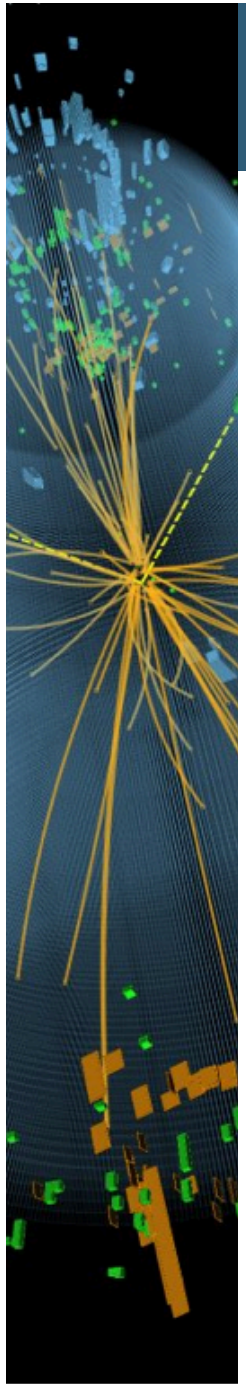

CMS Plans

OCT Planning Meeting

Nov 29

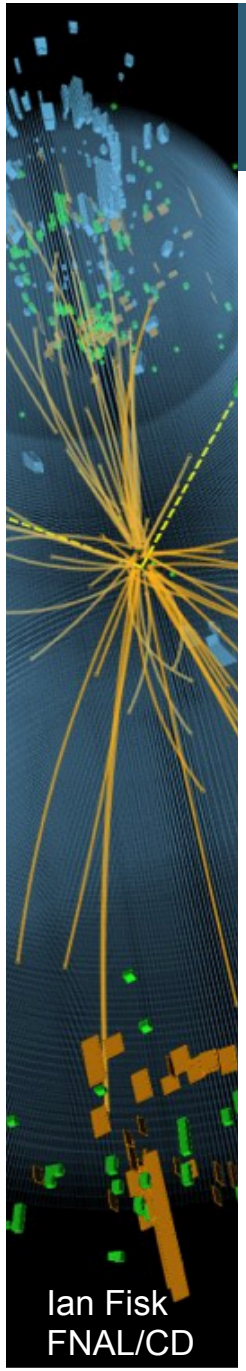

Plans

- CMS Medium Term Plans
  - Next 6-12 months involve
    - expanding the data federation
    - Changing the archiving model
    - Changing the storage management model
    - Deploying the HLT for use in organized processing
    - Retiring the use of gLite WMS
    - Sustaining HC and SAM Tests
- Longer Term
  - Multicore
  - Opportunistic Computing
  - Clouds

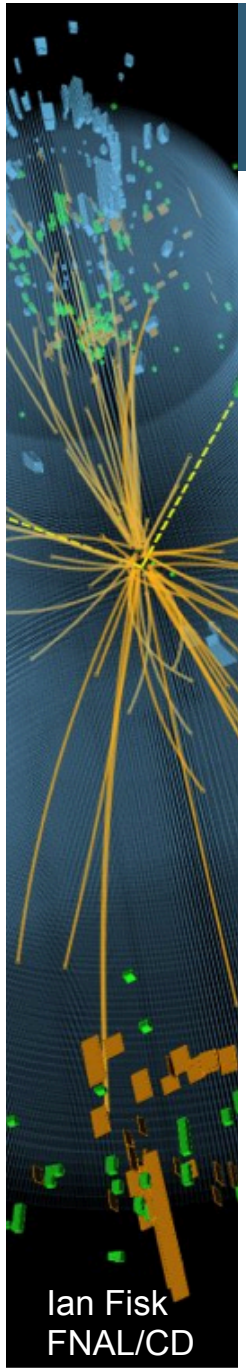

Xrootd

- We plan to get xrootd rolled out on all sites
  - Start with simply adding the Xrootd fall back to sites so that failures to open are retried on Xrootd
  - Next step is to push more facilities to enable xrootd access to storage
- We want to make SAM tests for xroot to monitor the health
- Expanded use of the xrootd popularity monitor to make sure the system is sustainable

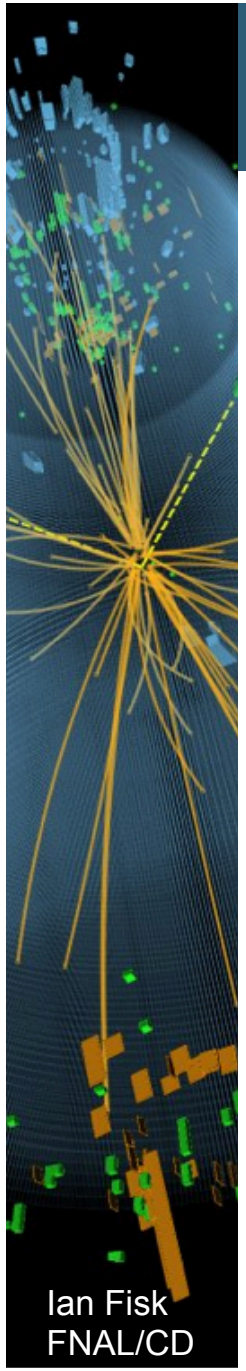

The Tier-1 Storage "Cloud"

- CMS is proposing to work on transparent access to data at Tier-1s using the OPN
  - Negotiations with sites for moving worker resources inside the OPN domain
  - Calculate and test the amount of access that could be sustained with our share
- Short term would eliminate individual workflow problems. In the long term could evolve CMS to much more dynamic use of the resources
- Instead of failing back to archive, we would fall over to xrootd if the data was accessible on another Tier-1
  - All items discussed at the Amsterdam workshop

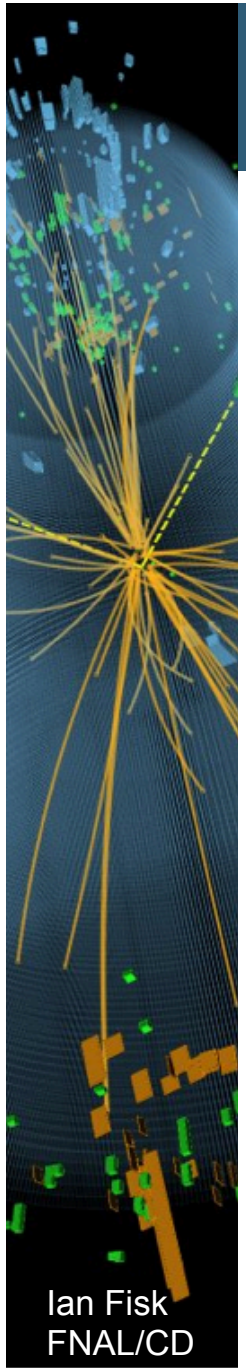

Impact

- Changing the WN access to the OPN will require a fair bit of negotiation
  - Requires reconfiguration at some sites
- Has the potential to benefit all the LHC experiments and could change to thinking of the Tier-1s as an aggregated resource instead of individual sites

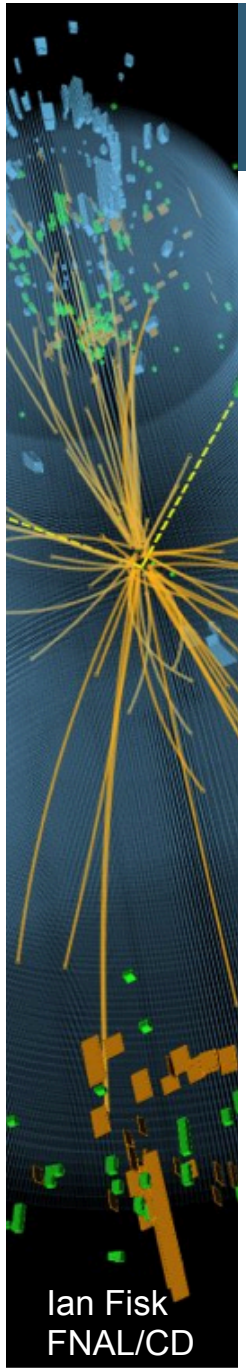

Disk/Tape Separation

- We will ask that all Tier-1 sites introduce two PhEDEx end points one disk and one tape
  - The disk end point should write only to disk
  - The tape end point should write to tape, but could also be hosted on disk
- Like any other endpoint data subscribed should be resident until deleted
  - Subscribing to disk will be the equivalent of the prestage and pin
  - Subscribing to tape will trigger and transfer to archive
  - Deleting from disk will release the cached space

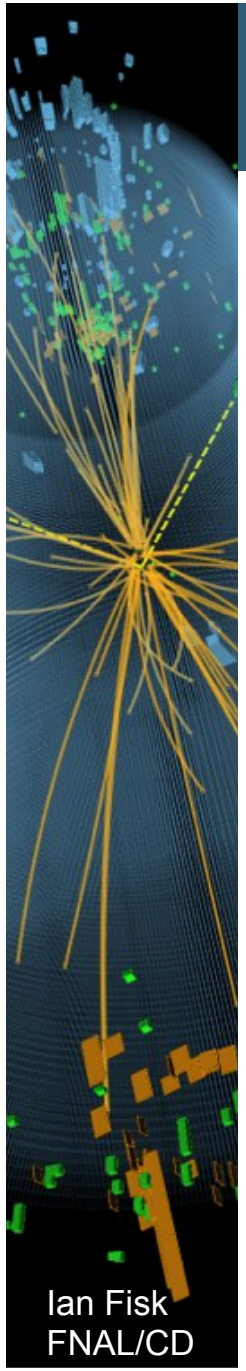

Disk/Tape Separation 2

- We choose this to have better disk management and to ensure we can enable xrootd without triggering unexpected staging
- We want to host and process data on Tier-1 disk without using tape, and we want to decide when and where to archive samples explicitly
- How sites choose to do the implementation is up to them
  - CERN has introduced a separate technology
  - RAL has done a split in the existing technology
  - Other technical solutions may exist

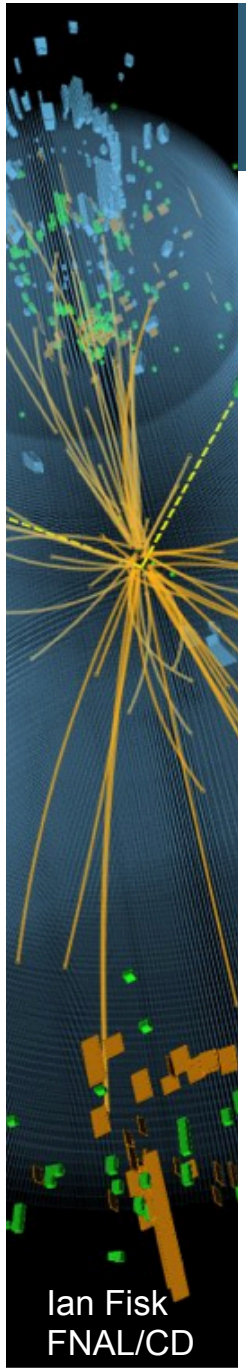

Data Management Changes

- We will propose a new split of the disk space at Tier-1 and Tier-2 centers
  - Managed vs Unmanaged
- Tier-1s will be entirely managed disks
- Tier-2s will be a percentage
  - Unmanaged User and group files
  - Official data managed automatically by us
- All data would be subscribed to Tier-2s as it was produced
  - We clean up the caches automatically with Victor
  - 2 copies, goes to 1 copy, goes to 1 xroot distributed copy from Tier-1, goes to an error condition

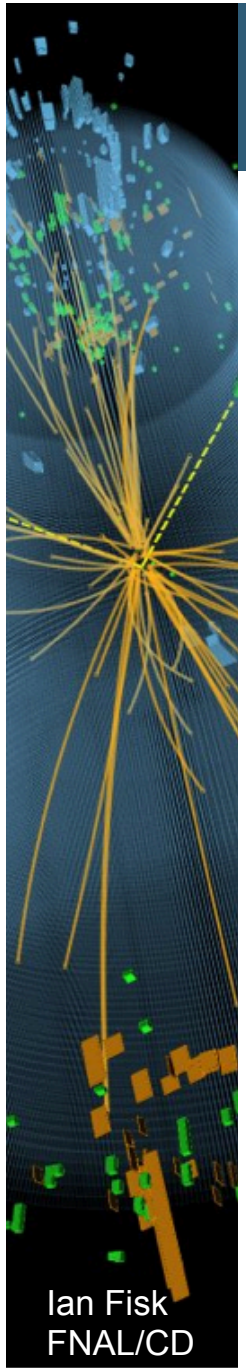

Disk Management Impact

- Will reduce the coupling of groups and sites for processing
  - Unmanaged data should be accessible via xroot
  - Managed data will be allocated automatically
- Requires improvements in priority for pilots, so we can give additional computing for individuals, groups, and roles.
- Should make better use of the disk space and change the slope of how disk needs increase

Text

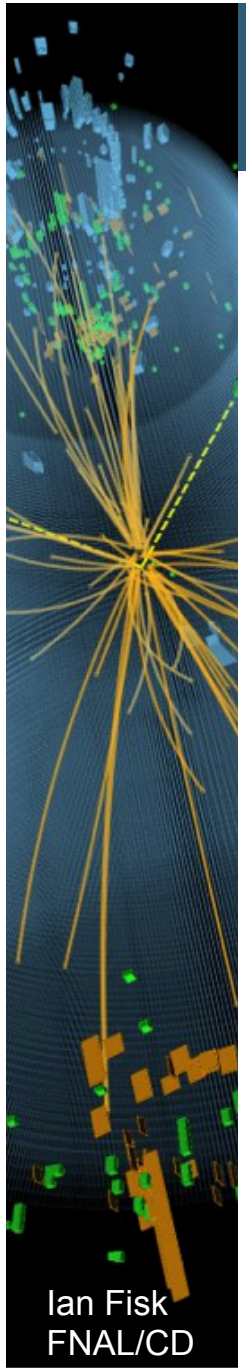

HLT Use

- CMS is testing this week the HLT farm in "Cloud-mode" for organized data reprocessing
  - Openstack with EC2 interface to start resources, driven by a glide-in WMS3 resource control
  - The CMS Production system through the glide-ins, accessing data from EOS over xrootd
  - Goal is running at scale over the Christmas stop and then smaller scale through the HI run
  - Most of the farm accessible after April

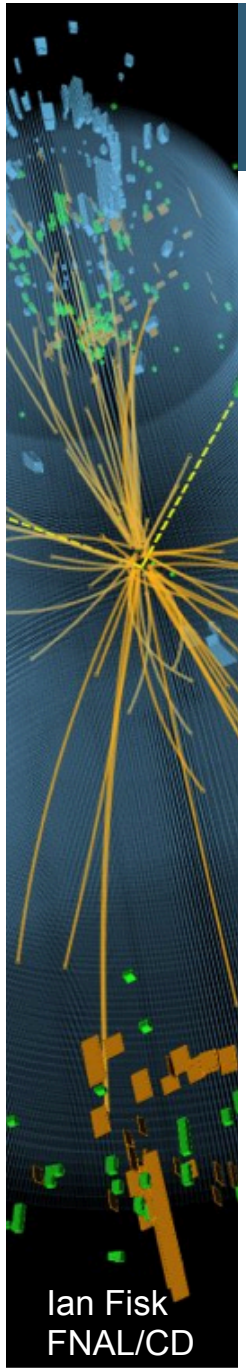

Retiring gLite WMS

- We would like to eliminate the last of the workflows that rely on the gLite WMS
  - SAM tests are in the process of moving to glide-in based tests
  - Analysis users through CRAB still submit both with client and server
  - Next week we will develop a release plan for the next generation of analysis tools that will support a pilot implementation by default

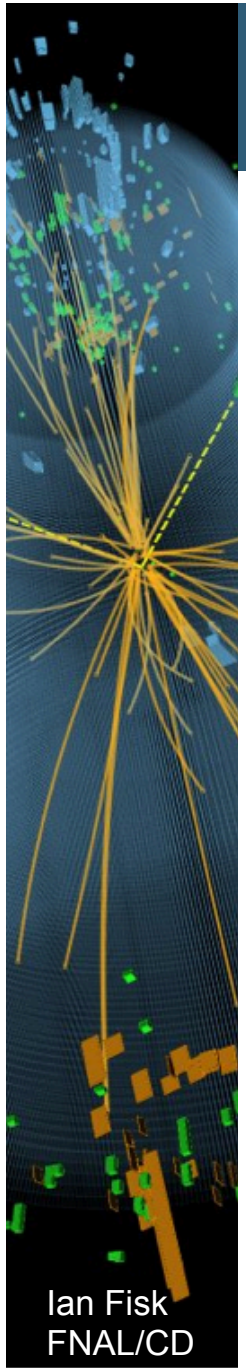

HC and SAM Tests

- Essential tools for the operation of the distributed infrastructure
  - Used effectively by CMS since several years to achieve and maintain a high site reliability level
  - Usage is increasing and closely following the evolution of the computing model
  - gLExec, storage federations, pilots, etc.
- Continued support of the frameworks is very important
  - New features in SAM: addition of site-level metrics, of new service types, glidein submission, etc.
  - New features in HammerCloud: full glidein support, better error reporting tools for operations, cloud submission, dynamic load management, etc.
  - Development of probes in CMS but the many commonalities with other experiments should be used to save effort
  - Already the case for SRM, gLExec tests

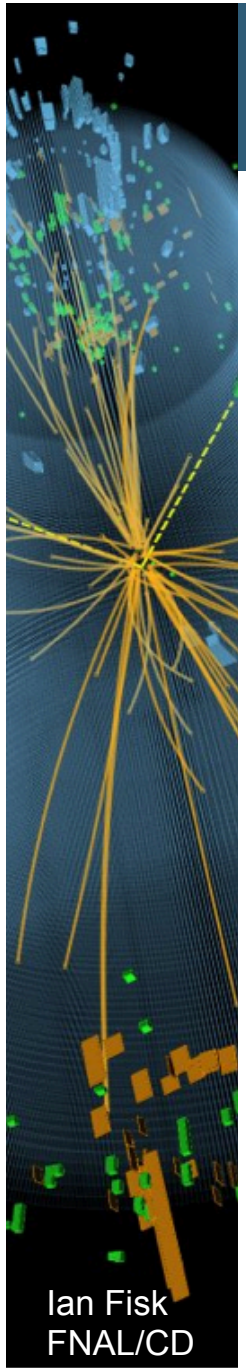

12-24 Months

- By the Fall of 2013 CMS expects to have a multi-threaded reconstruction release, and we will begin the transition of organized processing to multi-core
- There is a big push for better utilization of opportunistic computing
  - Leadership class systems (BlueGene etc.)
  - Opportunistic University Clusters
  - Clouds (Commercial and academic)
  - Requires some changes in the workflow
- Clouds: CMS would like to arrive at the end of LS1 with the demonstrated ability to utilize could resources at a similar scale to current simulation production