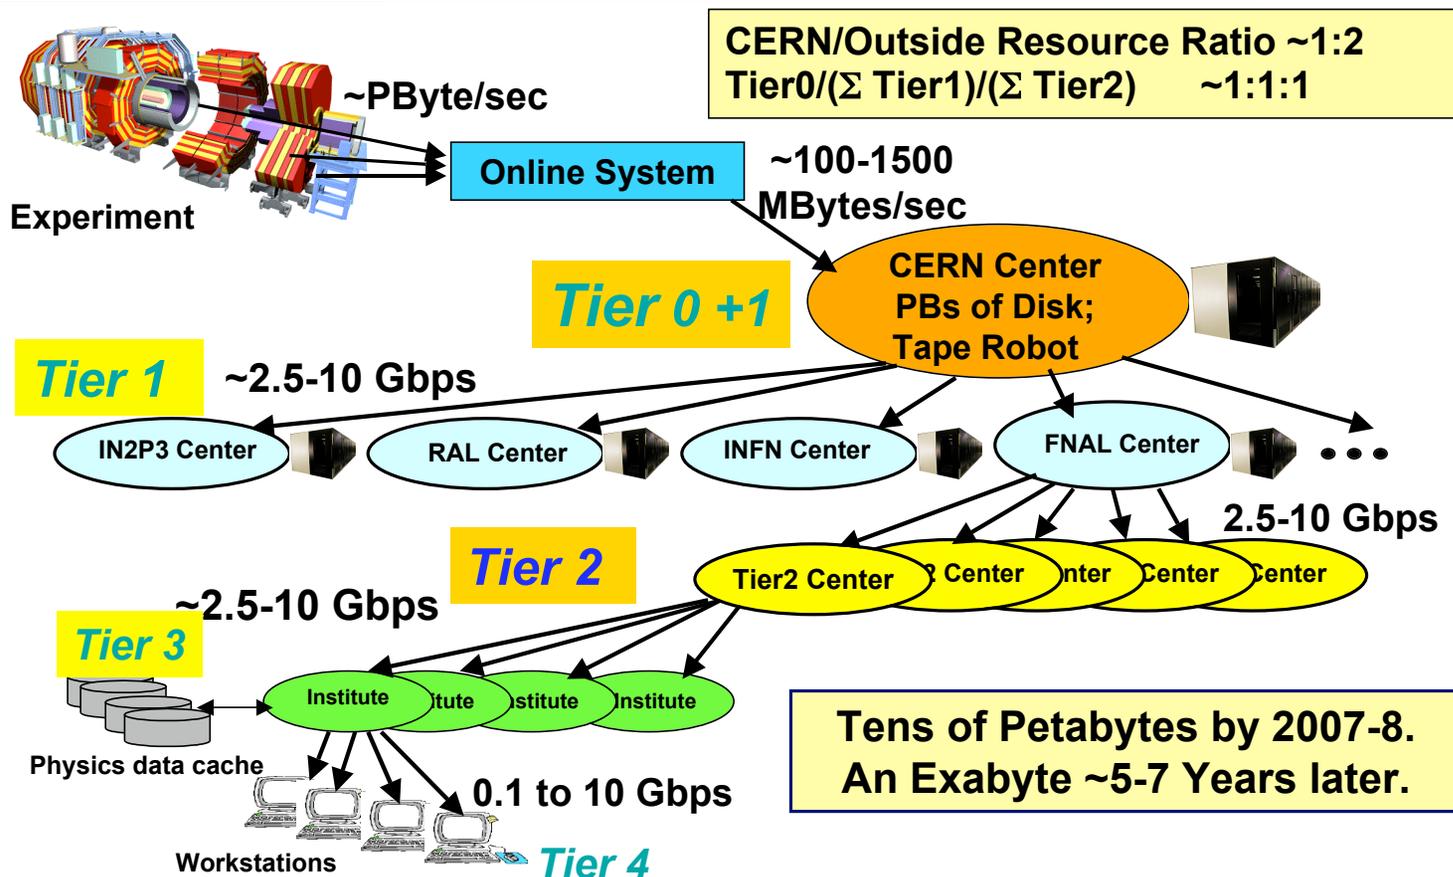


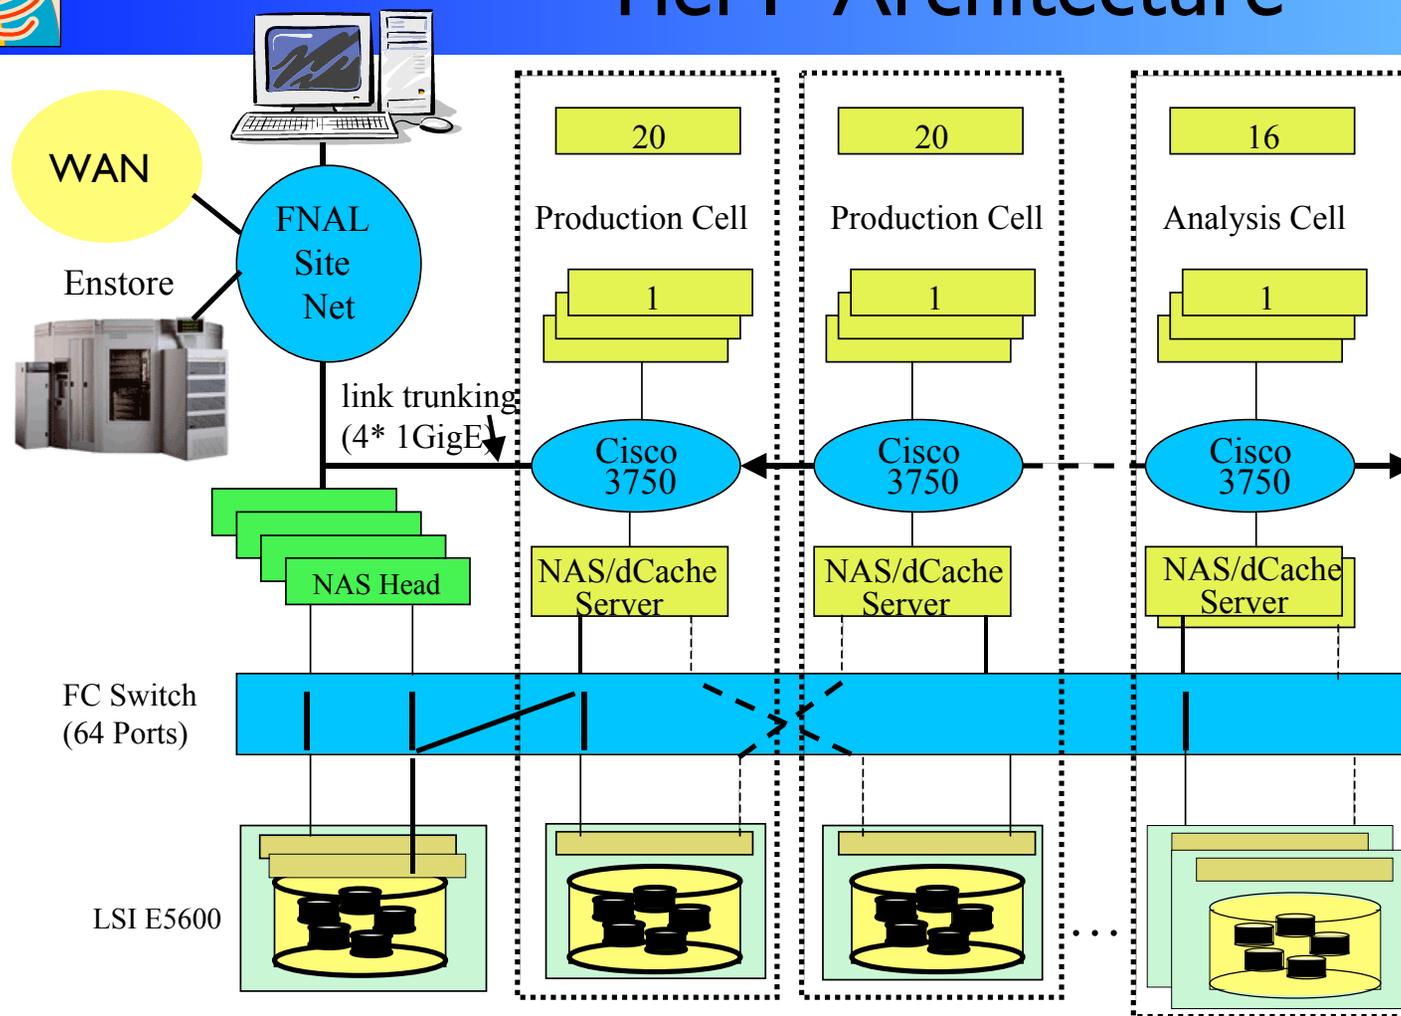
US-CMS Tier I and DC04

January 28, 2003



US-CMS Tier I Facility in Fermilab represents about 10% of the total CMS computing

- ➔ Offsite Tier I and Tier2 Centers represent the vast majority of CMS analysis resources



New Tier I Systems are Dual P4 Xeons connected over gigabit

➔ Lots of memory and significant local disk

4TB RAID arrays serve as dCache POOL nodes on the cells



User Analysis Facility (UAF)



Fermilab is building up the capability of the User Analysis Facility for CMS

- ➔ Up to now a fairly small scale facility with a small number of users
- ➔ Equal numbers of systems for Objectivity and POOL data

CERN is struggling to provide reasonable resources to the international CMS user community.

Fermilab is ramping up

- ➔ Procured new CPU for the user analysis farm
- ➔ Increased dCache space for serving datasets
- ➔ Implemented a new storage architecture

Beginning to pay off

- ➔ Head of CMS reconstruction software now sometimes works at Fermilab
 - He is working to get other PRS members to get accounts here

In the long term this will improve the quality of service and make sure FNAL is a productive analysis center for CMS



Hardware and Data Access Services

- ➔ The new storage network designed significantly improves the capability of the analysis cluster
- LSI currently has 3TB reserved primarily for user space
 - 4 clients connected over a fiber channel switch have sequentially read or wrote 400MB/s
- Chose a distributed file system from IBRIX
 - Allows data to be distributed over NFS from IBRIX servers connected to the LSI
 - Preliminary performance numbers are approximately 100MB/s
 - Nearly wire speed
- All UAF computing systems are connected over gigabit to the data
- We can provide each user with 50-100GB of updateable user space
- ➔ The Analysis cell has two 4TB RAID devices for dCache pools
 - fast and reliable

We expect to have a production level service by the end of February



UAF Hardware Services



The UAF processing capacity has increased by 16 new dual Xeon Systems

- ➔ Used for batch and interactive uses

If we began to see significant batch requests from the user community, we could send those requests to production resources at Fermilab

Access and load balancing is still handled by FBS batch queuing system

- ➔ Log in through central machine and then session open on available system
- ➔ Working on a load balancing system liked the one used for Ixplus as CERN.
 - Always log into Ixplus and the DNS server sends you to a free system



UAF Software Services



The US-CMS Software environment is linked directly to CERN using the database from CMS the configuration tool SCRAM

- ➔ All software releases and pre-releases are immediately available
- ➔ Uses afs
 - It can be a little slow the first time while the local caches are updated, but so far it has proved acceptable and reliable.
 - scram list on UAF systems will show the full list
- ➔ Some portion of the stable code will be installed locally
 - Hans Wenzel arranged for the code you are using tomorrow to be locally installed
 - Improves the performance and protects against network failure



Data Hosting and Cataloguing



The CMS Physics groups have requested about 70 million events so far for the preparation of the physics TDR

- ➔ The first 4M should hopefully be available next week
 - They will be replicated to all Tier 1s when available
- ➔ JetMet group is the only PRS group so far to request that their data is hosted at Fermilab
 - We have the tape space, the tape mounts, and the cache space to host all of them

There are about 50k files associated with the current request

- ➔ We are currently deploying local catalog services from the POOL persistency mechanism.
 - This simplifies but does not eliminate the problem. A user still needs to know the dataset names
 - Need better ways to publish dataset names



MC Generation



Currently physics groups make Monte Carlo requests

- ➔ Individuals submit to the coordinators who prioritize the list

The production team arranges to produce the MC

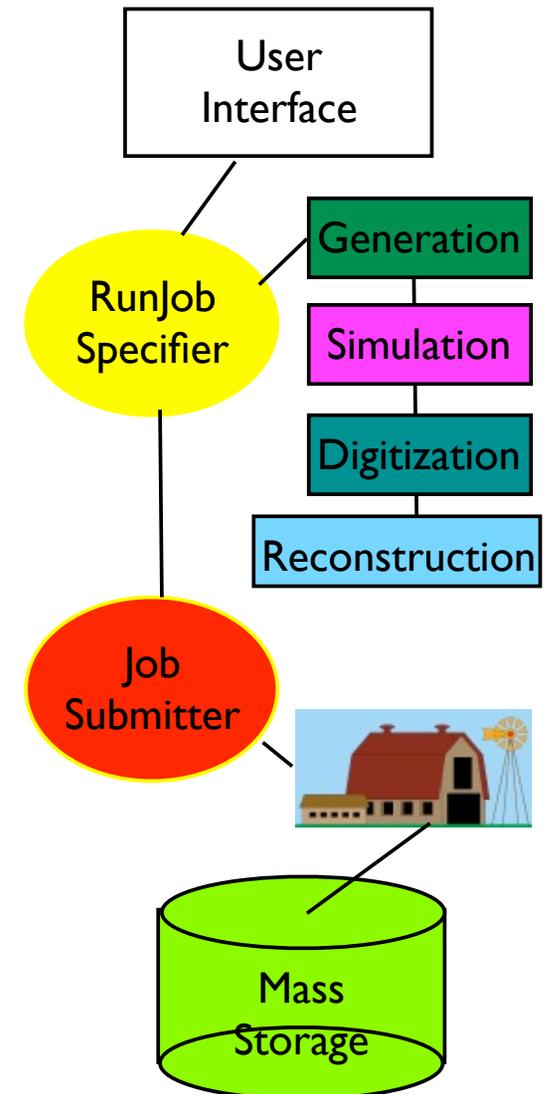
- ➔ Optimized to complete the entire sample
 - Each step is run in series
 - The intermediate steps are pretty worthless to the requestor
- ➔ Due to software delays this year, the time lag between request and delivery is about 6 months.
 - The DAQ TDR production was 4 months
 - Some of the specialized physics ntuple productions were faster, but still pretty slow

It would be good to reduce this time for small samples and to allow individuals to make custom requests



We are working to develop services for Monte Carlo Generation

- ➔ Build on the experience with production tools to be a system to allow individual users to create their own MC.
 - Starting with generator files or data cards
 - Progress through simulation and digitization
 - Make the events accessible to the user
- ➔ Current analysis farm could handle around around 10k events per day
 - Would be possible to submit the specified jobs to larger production farms if need arises
- ➔ Could dramatically reduce turn around time and allow users more freedom to choose what samples were run





Preparations for DC04



Data Challenge 04 is the first of the major CMS data challenge milestones. It is scheduled for March 1 and has three main areas

- ➔ Reconstruction of events at CERN at a continuous rate of 20Hz for a month
 - Software functionality and stability challenge
- ➔ Transfer the reconstructed data and some percentage of the raw data to Tier-I facilities for analysis, rereconstruction, and archiving
 - Challenges data management and transport tools
 - Data replication and cataloguing tools
- ➔ Once data is available, it is accessed at the Tier-I centers
 - Data must be spun over to demonstrate it is accessible and analysis is possible in a real time environment
 - Physics quantities can be analyzed and plotted

First area is primarily a Tier-0 responsibility, but areas 2 and 3 have a large Tier-I component



Analysis Components



Once data arrives at the Tier I

- ➔ It is archived in Enstore and made visible in dCache
- ➔ It is cataloged in the local mySQL Pool catalog

The physics groups need to spin over the data in real time at Tier I centers

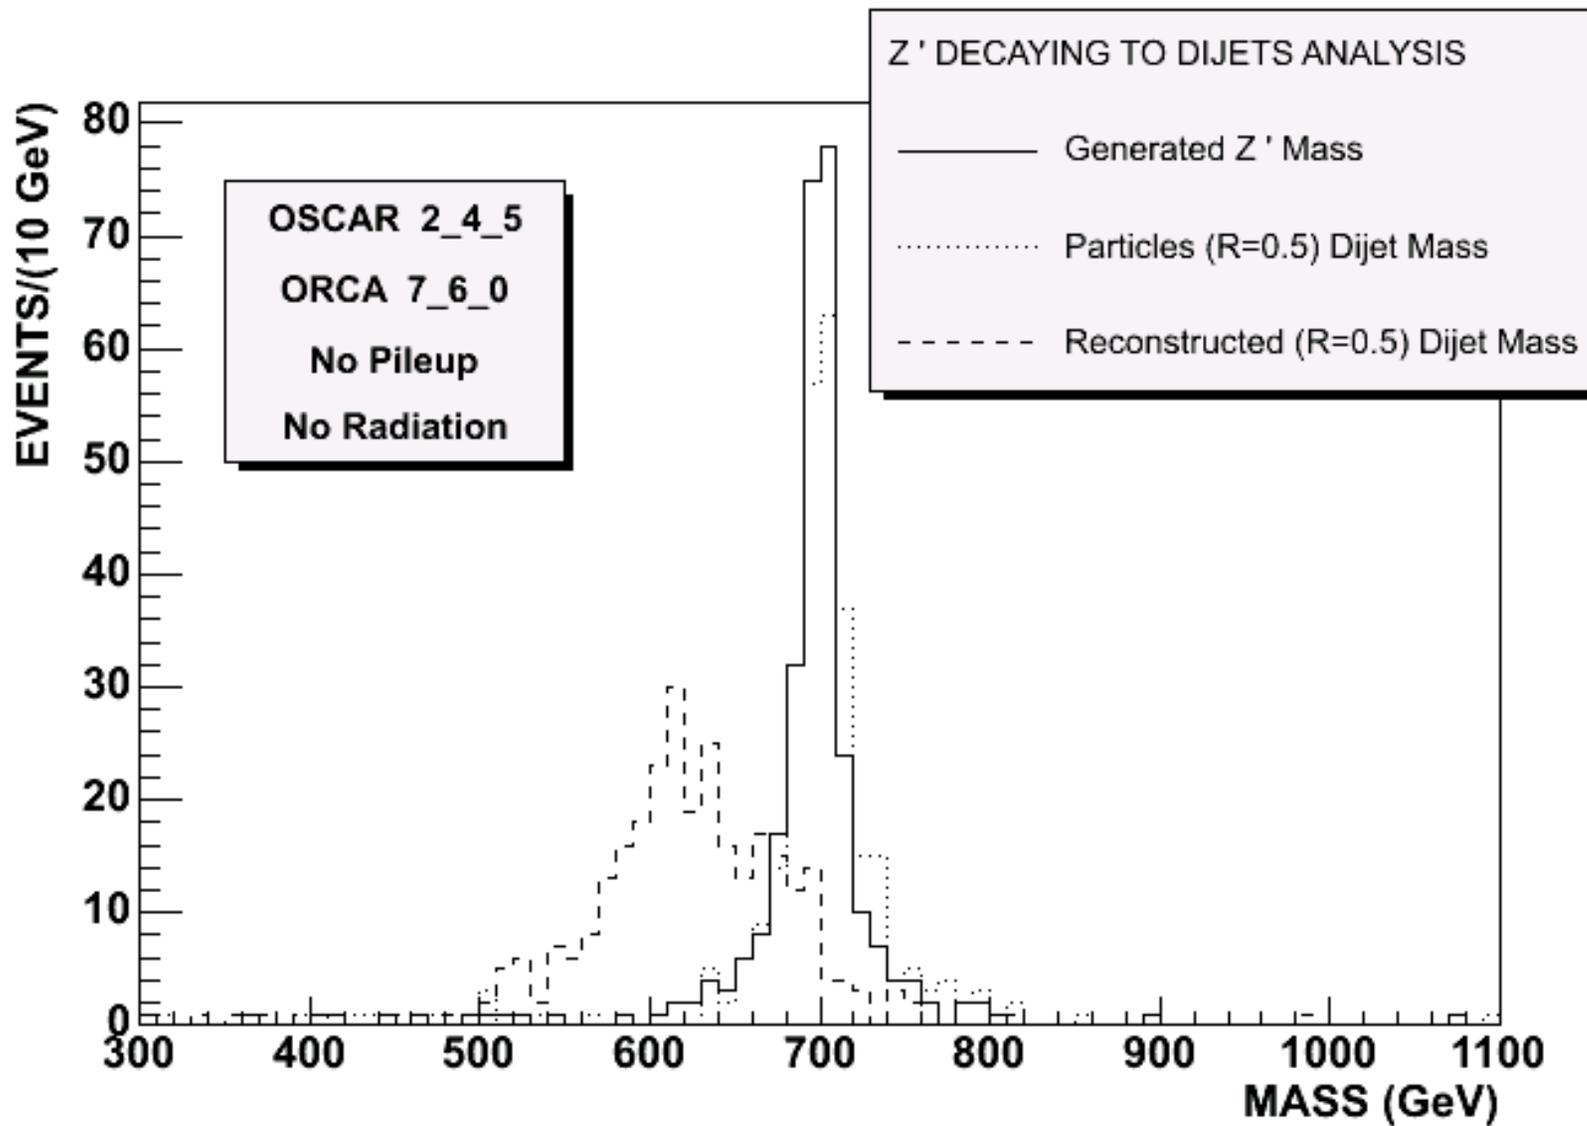
- ➔ Since this is a data challenge there was a decision early on not to transfer MC information
 - This has proved difficult because a lot of the PRS applications have expected it
- ➔ It was decided to send complete raw data to some sites and summarized data to all sites
 - The lack of examples of what was to be sent has hampered PRS preparation



Rob Harris and Hans Wenzel have been working on a physics monitor for DC04

- ➔ Starting with the JetMet PRS group analysis application, which is code which generates a Root tree from the raw data file
 - They've been able to reconstruct Z' masses from decays to dijets
 - Currently working on getting the JetMet code to run with the set of data files available for the challenge

- ➔ We need a few infrastructure components to run large numbers of jobs in an automated way
 - Modifications to the job submission tools used for production
 - Automated publishing of physics plots





Exciting time at the Tier I center

- ➔ Improvements in the infrastructure
 - Hopefully a corresponding improvement in activity
- ➔ We need feedback from the physics groups
 - Hans Wenzel is the US-CMS User Services Coordinator
 - There is a UAF requirements document for 2004
 - http://www.uscms.org/s&c/reviews/doe-nsf/2004-01/docs/UAF_requirements.pdf
 - Please provide feedback and suggestions

Data Challenge has sparked a lot of activity

- ➔ Should an interesting couple of months
- ➔ The analysis that follows in preparation for the Physics TDR should be fun