



Progress on Data Handling and Central Analysis Facility

Robert M. Harris
CDF DH Review
Dec. 19, 2001



Outline



- Management
 - ➔ New project managers
- Legacy System Operations
 - ➔ Additional personnel for operations
 - ➔ Solutions to operational problems
- Hardware & CD
 - ➔ Progress on Enstore + STK prototype
 - ➔ Progress on new Central Analysis Facility
- Collaboration Help
 - ➔ UK/GRID



Management



- Jeff Tseng
 - ➔ Co-head of Data Handling concentrating on development.
 - ➔ I will remain co-head for operations temporarily

- Frank Wuerthwein
 - ➔ Deputy Head of Offline Operations in charge of Central Analysis Facility Upgrade.
 - ➔ Frank has been offered a two-year guest scientist position in CDF department of computing division.
 - ➔ Frank is spearheading our effort to create a user analysis farm of PCs reading from lots of network attached disk.



Help for Legacy Operations



- Rich Glosson

- ➔ Has been trained by Mark Leininger for operations.
- ➔ Formerly worked on framework and banks.
 - Contract worker who was going to be terminated Jan. 31.
- ➔ I've extended Rich's contract until March 31.
 - Will help Mark until Enstore begins to replace legacy system.
 - Costs a lot but worth it to keep system running and Mark happy.

- CD team of tape experts to analyze operations.

- ➔ Margaret Votava and Luciano Piccoli from ODS
- ➔ Marc Mengel from OSS.
- ➔ Analysis was completed and solutions implemented.



Analysis & Repair of Operations



- Analysis: Votava report of Nov. 29
 - ➔ The system does not handle error conditions well.
 - ➔ Making it robust for all run 2A is a huge job.
 - ➔ Recommended test stand, simple patches & debug output.
- Repair by CD experts and DH group
 - ➔ The data handling test stand was made operational.
 - ➔ Operational problems were analyzed on test stand.
 - ➔ Patches and diagnostic output were implemented on stand.
 - ➔ Software patches were installed on fcdfsgi1/2 on Dec. 4.
 - ➔ Margaret gave a closeout on Dec. 5.



Operations



- Dismount problem patched
 - ➔ Most significant operational problem when reading tapes.
- Minor problems remain
 - ➔ Paul, Mark and Rich working on fixing problems.
 - ➔ Hope to have time over the holidays to catch up.
 - ➔ CD experts will serve as consultants as necessary.
- Operations are significantly smoother
 - ➔ Expect to be able to accommodate 8 streams / 20 datasets



Enstore + STK Project



- Huge Effort involving
 - ➔ CDF
 - ➔ Integrated Systems Department (ISD) of CD
 - ➔ Data Communications Department (DCD) of CD
 - ➔ Computing Services Department (CSD) of CD
 - ➔ D0 Department of CD
- Coordinating and staying on schedule is a lot of work.



Staged Plan



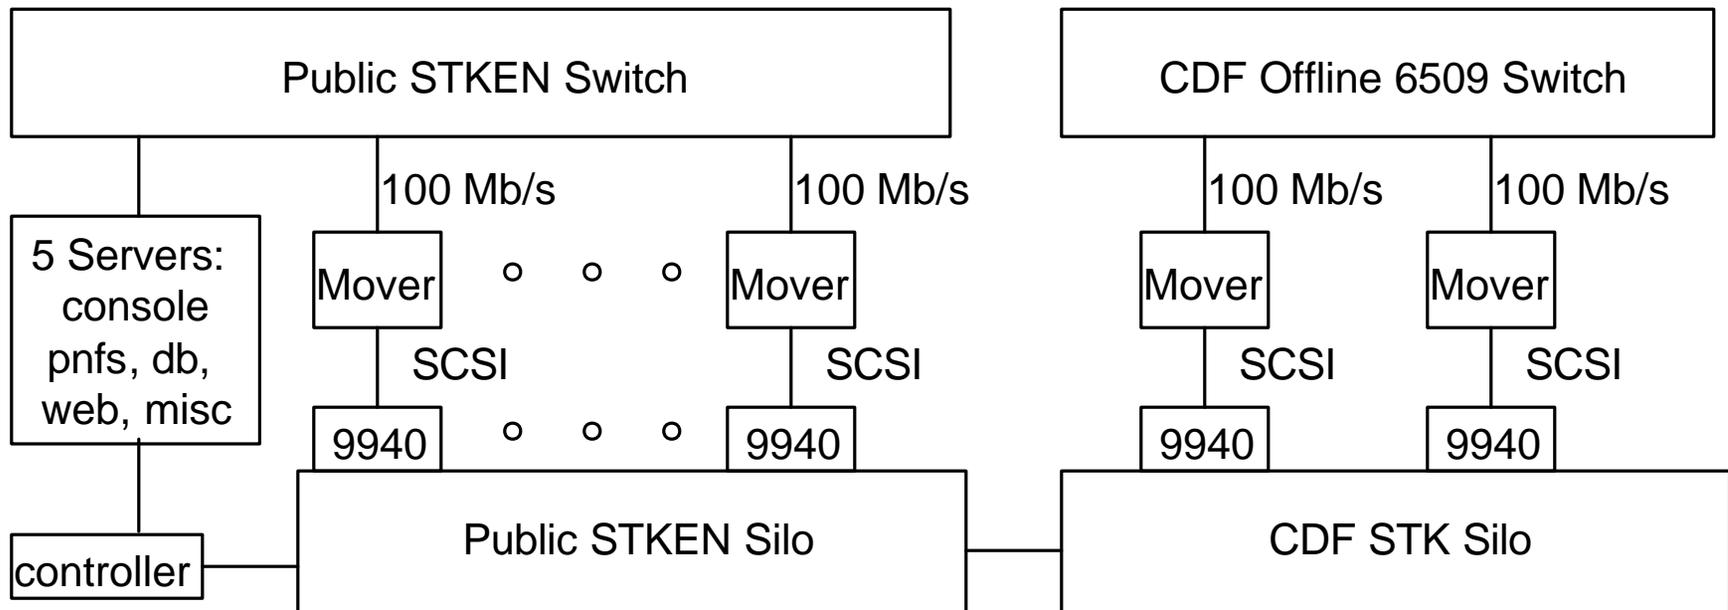
Stage	Arrival Date	Purpose	Drives & Rate	Robot
Prototype	Jan 1, 2002	Testing	2 Drives 20 MB/s	Run 1 STK Silo becomes CDFEN. Share STKEN servers for admin.
Stage 1	As early as Late Jan 2002	DC Data rate: 30 MB/s raw & farms. 70 MB/s users	10 Drives 100 MB/s	Standalone CDFEN. 330 TB capacity.
Stage 2	As early as fall 2002 ?	Full Run 2A: 60 MB/s raw & farms. 150 MB/s user	5 fast drives in 1 + 7 fast drives in 2	1 st CDFEN + 2 nd CDFEN with faster drives. 1.4 PB capacity.



Quick Prototype for Tape Handling



- Re-use the existing CDF Run 1 Silo and STKEN servers.
 - ➔ Two STK 9940 nodes will be installed in the CDF STK silo.
 - ➔ CDFEN robot will use the 5 STKEN servers for robot administration.
 - ➔ Two mover nodes will pump the data onto the offline 6509 switch.
 - ➔ One GBit NIC on fcdfsi2 with single CPU dedicated to robot traffic.
 - ➔ Prototype will be capable of 20 MB/s.





Prototype Installation Schedule



- Prototype by Jan. 2
 - ➔ Emptying Silo
 - 1.3 TB of Run 1 secondary data being emptied onto cdfsga by CDF.
 - >90% finished. Two more days to complete job working round the clock.
 - CSD will remove tapes from robot Dec. 22 – 23. Blanks on Dec. 20.
 - ➔ New Drives
 - Purchase order released Dec. 11 and first drives should be on site soon.
 - All 10 STK 9940 drives scheduled to be installed Dec. 27 by STK.
 - ➔ Mover Nodes
 - Two nodes will be attached to drives by ISD beginning Dec. 28.
 - Work scheduled to begin on Dec. 28 system will be handed to CDF Jan. 2.
 - ➔ Networking
 - Networking for mover nodes put in place by DCD.
 - GBit NIC for fcdfsig2 dedicated Enstore traffic should arrive Dec. 28.
 - Will be installed by CDF and configured to dedicated processor by D0.



Stage 1 Schedule



- Stage 1 hardware in as little as 6 more weeks.
 - ➔ ISD installs server nodes dedicated to CDFEN robot.
 - ➔ In parallel with these 6 weeks
 - ISD installs 8 additional mover nodes.
 - Obtain and install additional gigabit NICS for fcdffsgi1 / 2.
 - Purchase of 1000 tapes. Directorate has signed requisition but no purchase order yet.

- Stage 1 planning meeting today at 2:30 p.m.

- MOU drafted for support of Enstore.
 - ➔ ISD provides support in collaboration with CDF & DCD.



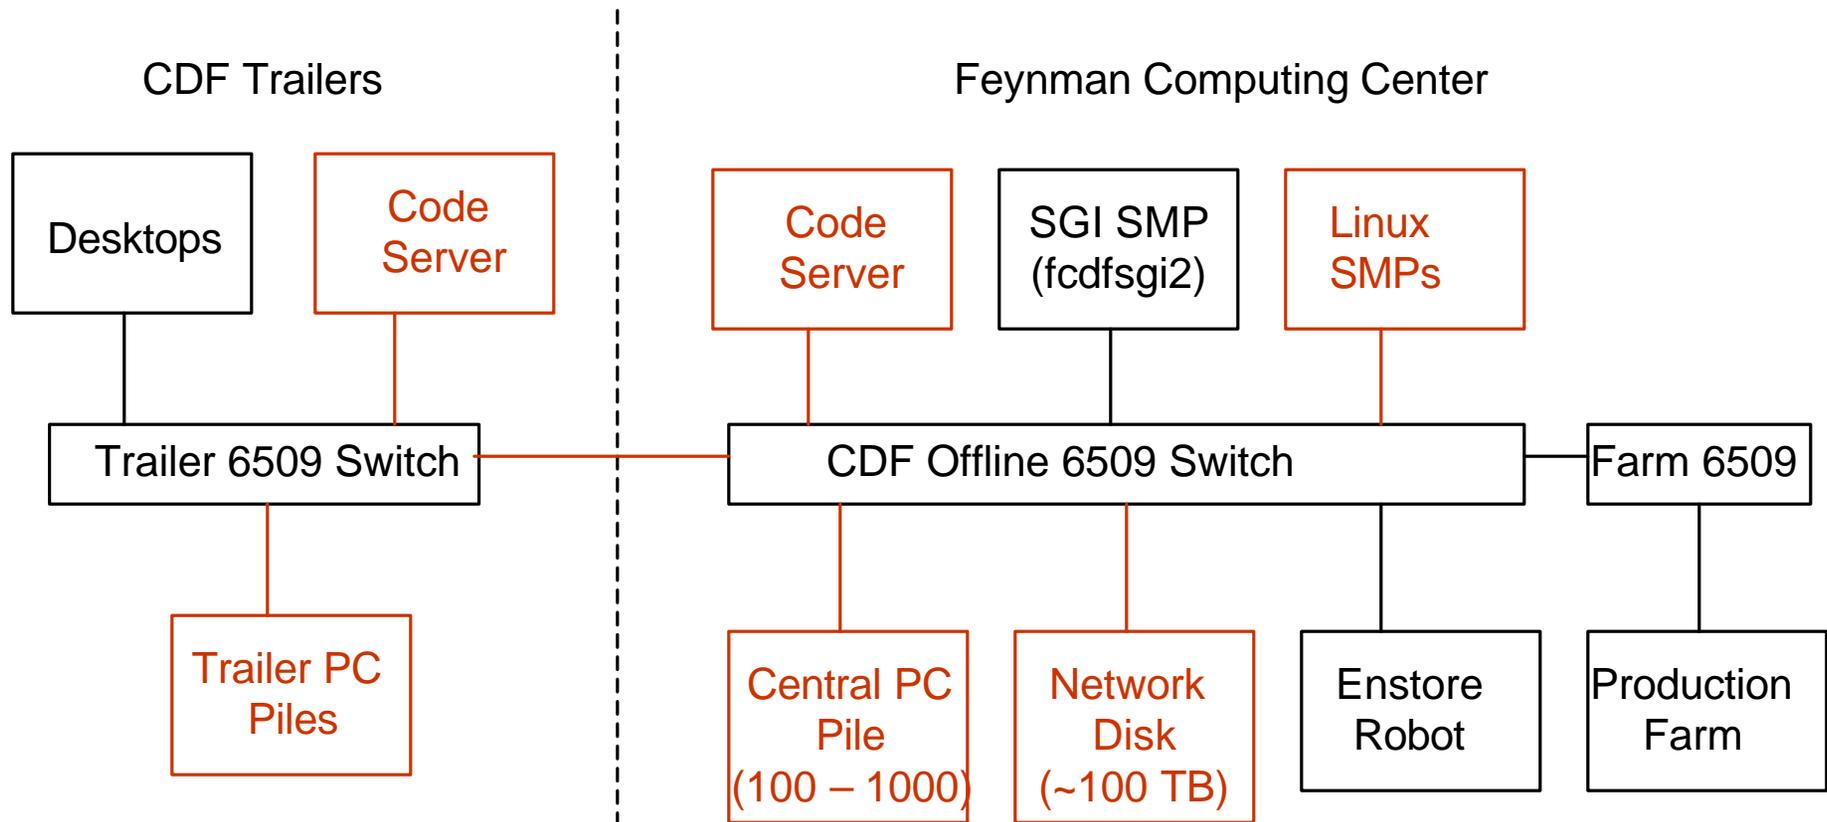
Farms and Online



- Meeting on Dec. 12 with farms group agreed that
 - ➔ Mid-Jan: farms complete concatenation on worker nodes.
 - ➔ End-Jan: farms implements writing to Enstore from fcdfsi1
 - Use's existing gigabit connection to farms on fcdfsi1.
 - Uses FSTD software that Dmitri has already implemented.
 - ➔ End-Feb: farms implements direct writing to Enstore from farms worker nodes.
- Online will wait until fcdfsi1 logs raw data to Enstore
 - ➔ Wants to see stability of logging raw data before B0 machines tries to log raw data to Enstore directly.



Central Analysis Facility Sketch



— Exists or already planned

— Proposed to be added or upgraded



Progress on Central Analysis Facility



● User Analysis Farms

- ➔ Frank Wuerthwein has been actively planning with CD.
- ➔ Meeting between CDF and CD farms group in OSS.
 - Agreed on procurement of 48 duals to be ready by April 1.
 - ➔ **Requisition, installation and burn in by OSS.**
 - ➔ **Joint project on administration by OSS and CDF.**
 - ➔ **OSS provides training and expertise for a CDF new hire.**
- ➔ Meeting between CDF and DCD on networking.
 - Agreed to initial network topology for new CAF.
 - Started requisitions for user analysis farms.
- ➔ Meeting with CDF task force and collaboration on prototype
 - Farm prototype will be contributed by universities in the trailers.
 - Network attached disk prototypes to be setup in trailers too.



Longer Term Enstore Requirements



- How long will it take to fill our 1st CDFEN robot?
 - ➔ We have 70 TB of data now we will transfer to robot.
 - ➔ If the accelerator & detector do very well we could
 - Write 0.5 TB/day of raw data
 - ➔ **20 MB/s capability and 30% accelerator*detector efficiency**
 - 0.5 TB/day of reconstructed data and secondary datasets
 - 70 TB + 1 TB / Day
 - ➔ **Robot fills up in 260 days, which is September 2002.**

- We should plan for another robot by September 1, 2002.
 - ➔ Using the next generation drives due out by June 2002.
 - 30 MB/s rate and 200 GB/cartridge capability.
 - One more silo would provide an additional 1.1 PB of storage.
 - 5 upgraded drives in 1st CDFEN robot would handle 150 MB/s user rate.
 - 7 upgraded drives in 2nd CDFEN robot would handle 210 MB/s total rate.



CDF/UK GRID: Chronology



- CDF/UK proposed a contribution to GridPP.
 - ➔ Oct. 5: Submitted original proposal.
 - ➔ Oct. 25: Update with more specifics on deliverables.
 - ➔ Nov. 1: Short plan on how CDF/D0 could work together.
 - ➔ Nov. 28: Offline management endorses CDF/UK effort.
 - ➔ Nov. 30: Submitted bids for allocation of 4 positions in UK.
 - 1 CDF, 1 D0, and 2 CDF/D0 (SAM).



CDF/UK Grid Proposal



- Three month evaluation of SAM and DataGrid tools
 - ➔ Afterwards CDF/UK will report and make recommendation.
 - ➔ CDF offline will review recommendation.
 - ➔ Decision will be made on which tools to use, if any.

- Remaining two years if SAM is chosen.
 - ➔ CDF/UK will collaborate closely with CDF, D0 and CD on
 - Incorporating the CDF Data File Catalog into SAM at CDF.
 - Interfacing SAM to a disk cache at CDF.
 - Any other changes that are necessary to mesh the two systems for data distribution to and among the remote institutions.



Offline Management's Opinion



- Enthusiastic about common CDF/D0 solutions.
 - ➔ We have no desire to reinvent the wheel.
- We have endorsed the CDF/UK exploratory effort.
 - ➔ We are waiting to see what they learn.
- Hope this will help with CDF data distribution.
- CDF may be interested in use of SAM for FNAL DH.
 - ➔ But our primary interest now is in data distribution.



Status as of early Decemer



- CDF/UK have made significant progress
 - ➔ Setup SAM stations at all four UK institutions.
 - ➔ Moved a file from Glasgow to UCL via the SAM DB
 - Used SAM bbftp.
 - ➔ Entered 3 files into the SAM DB with the metadata from CDF Data File Catalog.
 - Three tiers: raw file, reconstructed file, top dataset file.
 - ➔ Declared the top dataset file storage location to SAM.
 - ➔ Defined a project with that file and started it at UCL.
 - The file appeared in the UCL cache.
 - ➔ CDF now has it's own instance of SAM DB
 - Imported D0 development instance into CDF development DB.



Conclusions



- CDF Data Handling has made rapid progress.
 - ➔ New management has been put in place
 - ➔ Legacy system has been patched and staffed.
 - ➔ New tape handling system is on schedule
 - ➔ CD is making huge contributions.
 - Interdepartmental effort: CDF, CSD, D0, DCD, ISD, ODS, OSS.
 - ➔ Collaborators are coming on board.
 - MIT group contributing heavily to project management.
 - CDF / UK group moving rapidly on GRID.
 - Universities pitching in for CAF prototype work.

- We need the continued support of the committee.