



CDF Computing and Analysis

Robert M. Harris
CD Briefing & Status
Nov 7, 2002



Outline



- CDF Department
 - Goals
 - CDF Organization & CDF Department Members
 - Collaboration with other CD Departments
 - Computing & Physics Activities of each Person
- CDF Computing Operations
 - Overview given at September 5 CDF Collaboration meet
 - Included for completeness. I won't show most of these slides.
- Monthly Status Report
 - An update on selected activities



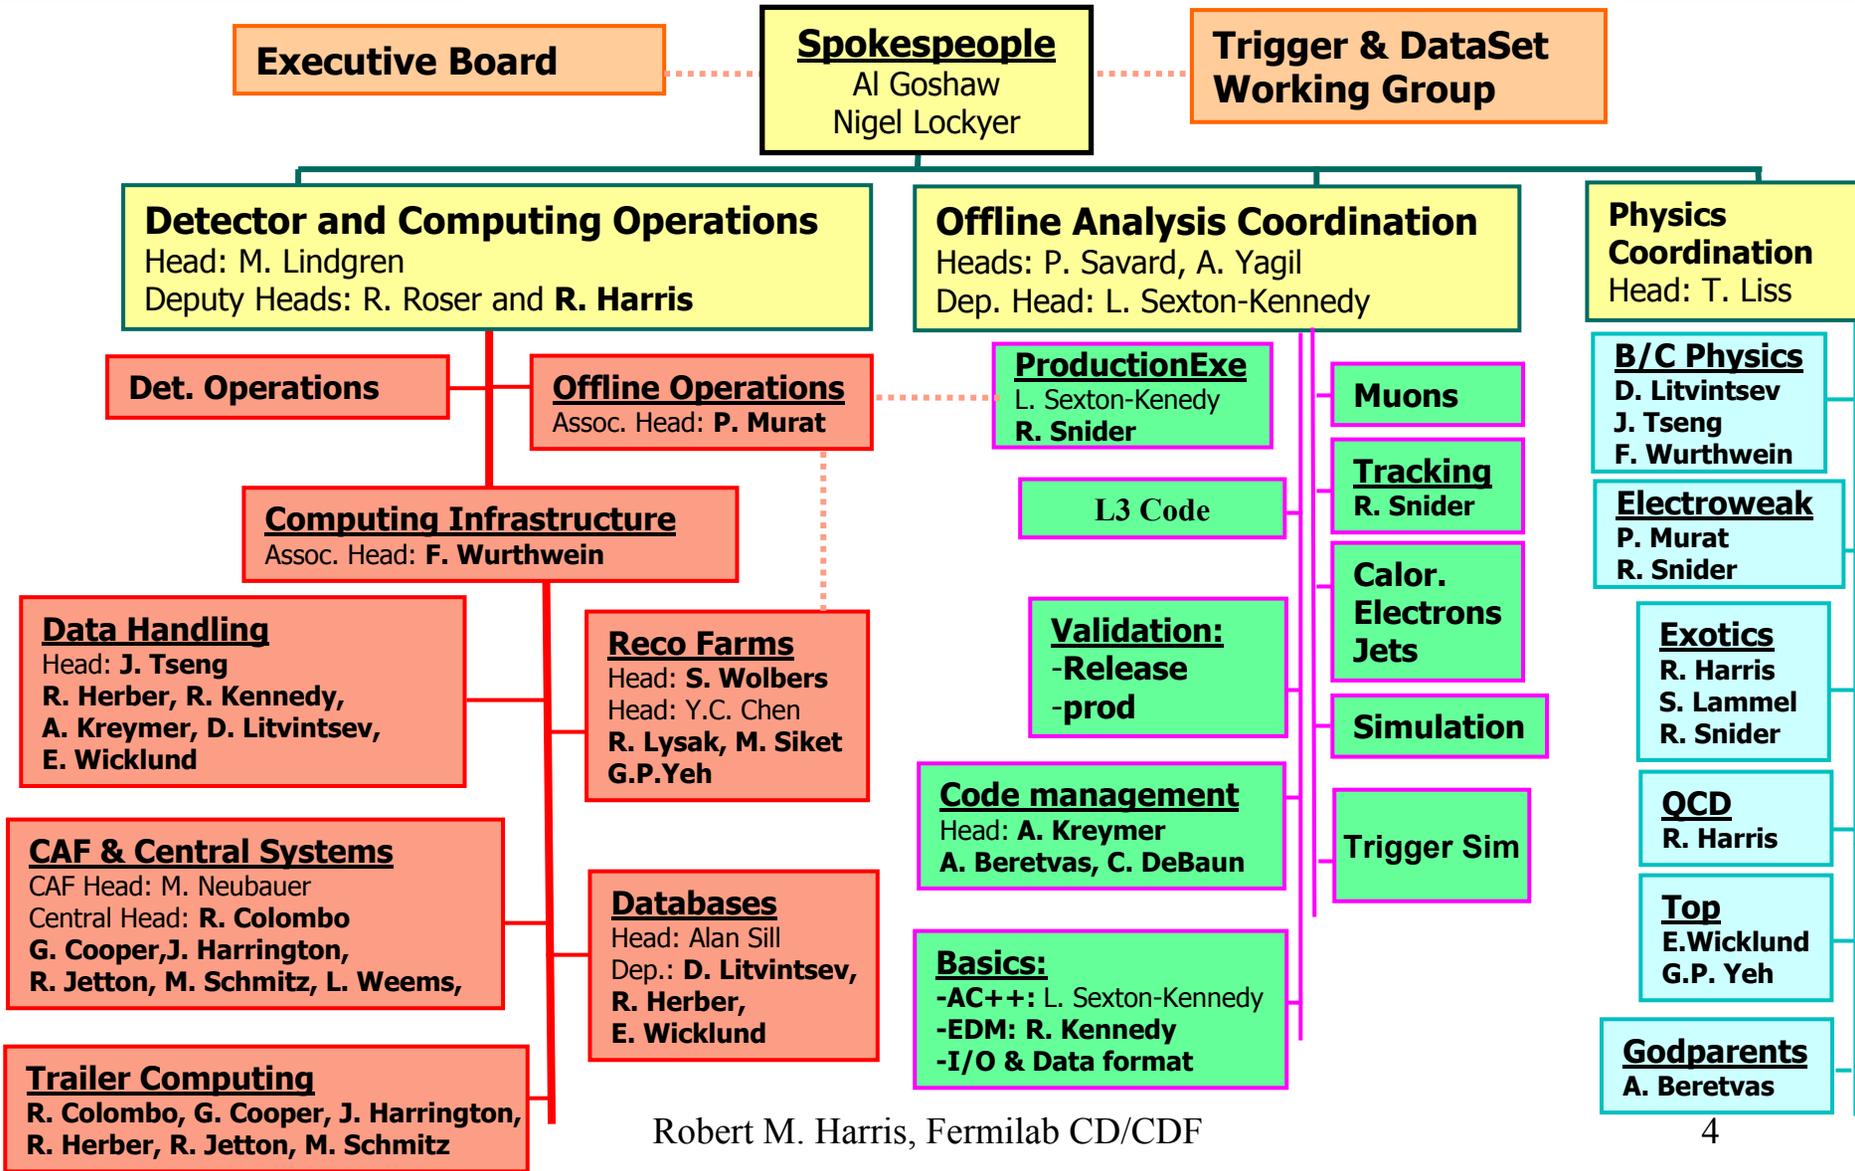
- Goals

- In coordination and cooperation with the rest of the division, provide & operate the computing infrastructure needed for the CDF collaboration to record, process and analyze data.
 - Data Handling
 - Reconstruction Farms
 - Central Computing Systems & Central Analysis Farms (CAF)
 - Databases
 - Desktops & Computer Security
 - Code Management & Software Infrastructure
- Participate in the analysis of CDF data, both to provide a home for active scientists in the department and to help insure our computing infrastructure stays focused on what is really needed for physics.



CDF Organization & CDF Department

(CDF Dept Members in bold font)





Collaboration with other CD Departments

Critical for success of CDF !



- CPD
 - Framework, C++ consulting, ROOT, physics generators.
- CSD
 - Helpdesk, paging & tape ops.
- CMS
 - Linux system monitoring.
- CST
 - Computer security.
- DCD
 - Networking, especially for CAF.
- DO
 - SAM, Enstore & SGIs.
- ESD
 - PREP, contracts, hardware install.
- ESE
 - Electronics for detector (PPD).
- FESS
 - Space, power, budget & admin.
- ISD
 - Data Handling (Enstore, dCache, dfarm) and GRID.
- OSS
 - Farms, CAF, and Linux.
- ODS
 - DB administration & software, licenses, SoftRelTools and GRID.



General Directions



- Technology
 - Network intensive systems.
 - Data handling via networked attached tape drives & disk.
 - Commodity computing
 - Central Analysis Farms (CAF) from inexpensive PCs and IDE disk.
 - Distributed computing
 - SAM, GRID, and decentralized CAF to integrate remote resources.

- Organization
 - Close collaboration with CDF institutions and CD departments.

- Priorities
 - Concentrate on computing systems & software infrastructure.
 - Reconstruction & analysis software is property of all CDF & lower CD priority.

- External Hardware Contributions
 - Incentives for CDF institutions to contribute to the CAF for everyones benefit.



Department Computing Activities



- Andy Beretvas (Application Physicist)
 - Code Management
 - Validation of releases and secondary release manager (C. DeBaun is primary).

- Chuck DeBaun (Computing Professional)
 - Code Management
 - Primary release manager: building and distribution of releases.
 - Special code management scripts as needed.

- Art Kreymer (Application Physicist)
 - Code Management
 - Leadership and problem solving in all areas. Effective supervisor of Chuck and Andy.
 - Data Handling
 - Learning to do configuration management and distribution for SAM.
 - CAF & Trailers
 - Participates in code server mgmt and areas of overlap with code mgmt.
 - Databases
 - Assists in problem solving associated with connections and poor code performance.



Department Computing Activities (cont.)



- Randy Herber (Computing Professional)
 - Database and DH development: DB Browser, SAM developer.
 - User and department consulting on difficult UNIX problems.
- Rob Kennedy (Computing Professional)
 - Infrastructure Software
 - Event Data Model (EDM), ROOT, loaders, compilers, debuggers, Linux, C++ working group.
 - Data Handling
 - CDF dCache project, liaison to ISD on dCache, integration of dCache with SAM.
- Dmitri Litvintsev (Post Doc)
 - Data Handling
 - Raw and reconstructed data logging and dataset management.
 - Data File Catalog design, operations and maintenance.
 - Deputy Head of CDF Databases
- Jeff Tseng (Guest Scientist, Senior Post Doc)
 - Head of CDF Data Handling, co-head of SAM, co-chair of R2D2.
 - Leads CDF data handling development and operations including GRID.
- Eric Wicklund (Application Physicist)
 - Data Handling and Databases: monitors and tests Data File Catalog and Databases.
 - Identifies, reports and solves problems.



Department Computing Activities (cont).



- Roman Lysak (Guest Scientist, Grad Student)
 - Being trained as a Reconstruction Farms expert and operator.
- Pasha Murat (Application Physicist)
 - Assoc. Head of Detector & Computing Operations in charge of Offline Operations.
 - Training & supervising of offline shift personnel, and operational aspects of producing & testing new production executables.
 - Coordinates weekly offline operations meeting and presents to CDF weekly meeting on issues of Farms processing, calibrations, Code Management, CAF, DB, DH, etc.
- Rick Snider (Associate Scientist)
 - Reconstruction software and the production executable.
 - Tracking software consulting, run 2b tracking committees, DB software committee.
- Miroslav Siket (Guest Scientist, Grad Student)
 - Reconstruction Farms expert, designer, operator and maintainer of CDF farms.
- Steve Wolbers (Scientist)
 - Head of CDF Farms group.
 - Division Office Duties.



Department Computing Activities (cont.)



- Robert Harris (Scientist)
 - Head of CDF Department
 - Deputy head of CDF Detector and Computing Operations.

- Stephan Lammel (Scientist)
 - Design, development and management of ATOM cluster: physics analysis cluster for CDF physicists associated with Fermilab. Supervising Rich Krull, sys admin. in PPD.

- Frank Wurthwein (Guest Scientist, Assistant Professor)
 - Associate Head of Detector and Computing Operations for Computing Infrastructure
 - Line management responsibility for all computing infrastructure operations.
 - Leader of Central Analysis Farms (CAF) and GRID related CAF development.

- G.P. Yeh (Scientist)
 - International Relations & Public Outreach
 - WWW event displays, development of Okinawa graduate school, Local organizer of support for Taiwan Farms personnel, Advisor on Science and Technology to leaders in Taiwan and Japan, DuPage Research Park.
 - Boards and Committees
 - CDF Executive Board, CDF International Coordination Board, CDF Speakers, ACFA and world wide Linear Collider studies and workshops, Lepton Photon 2003.



CDF Task Force Activities



- Computing professionals on 24/7 pager rotation for support of CDF computing.
 - All are being cross-trained to support the rapidly expanding CDF CAF.
- Rick Colombo
 - Associate Head of CDF Department and group leader of CDF task force.
 - Coordination of requisition and delivery of all computing equipment. Laptop support.
- Jason Harrington
 - System administration of Linux Desktops, CDF workgroup. Linux system development.
- Mark Schmitz
 - System administration of IRIX and Linux Desktops.
- Glenn Cooper
 - System administrator for IRIX SMPs: cdfsga, fcdfsi1, fcdfsi2. Systems development.
- Richard Jetton
 - System administrator for CDF database machines. Systems development.
 - Deputy CDF General Computer Security Coordinator & FCIRT member.
- Lance Weems
 - Primary system administrator for CAF. Testing and development of CAF systems.



Physics Activities of CDF Department



- Dmitri Litvintsev (Bottom/Charm)
 - L3/offline tracking efficiency for J/ψ cross section paper.
 - Reconstruction of the decay $\Lambda_B \rightarrow J/\psi \Lambda$.
 - Goal of approved CDF result by Lepton-Photon, August 2003.
 - Interested in reconstructing other bottom baryons Σ_B & Λ_B^* after that.
- Pasha Murat (Electroweak)
 - Test of lepton universality in W decays: $BR(W \rightarrow \tau \nu) / BR(W \rightarrow e \nu)$.
 - Goal of approved CDF result by LaThuille, March 2003.
 - Interested in τ polarization, and charged higgs decays to $\tau \nu$ after that.
- Rick Snider (Exotics & Electroweak)
 - Search for New Massive Stable Charged Particles: uses new TOF.
 - Goal of approved CDF result by LaThuille, March 2003.
 - Improved lepton identification in forward regions for W asymmetry.



Physics Activities of CDF Dept (cont.)



- Andy Beretvas
 - Godparent for roughly 2 CDF papers per year.
- Robert Harris (QCD & Exotics)
 - Searching for new particles decaying to dijets. Also thinking about $t\bar{t}$ resonances.
- Stephan Lammel (Exotics)
 - Co-head of CDF Supersymmetry working group.
 - Run 2 SUSY dilepton triggers, datasets and analysis.
 - Guiding students & postdocs, and producing papers on run 1 data.
- Jeff Tseng (Bottom/Charm)
 - Searching for $\eta_c \rightarrow J/\psi J/\psi$ in run 1 & run 2.
- Eric Wicklund (Top)
 - Analysis of Top decays to dileptons in run 1 and run 2.
- G. P. Yeh (Top)
 - Leading Taiwan-Slovakia group in top analysis for run 2.



CDF Computing Operations

Robert M. Harris

September 5, 2002

CDF Collaboration Meeting



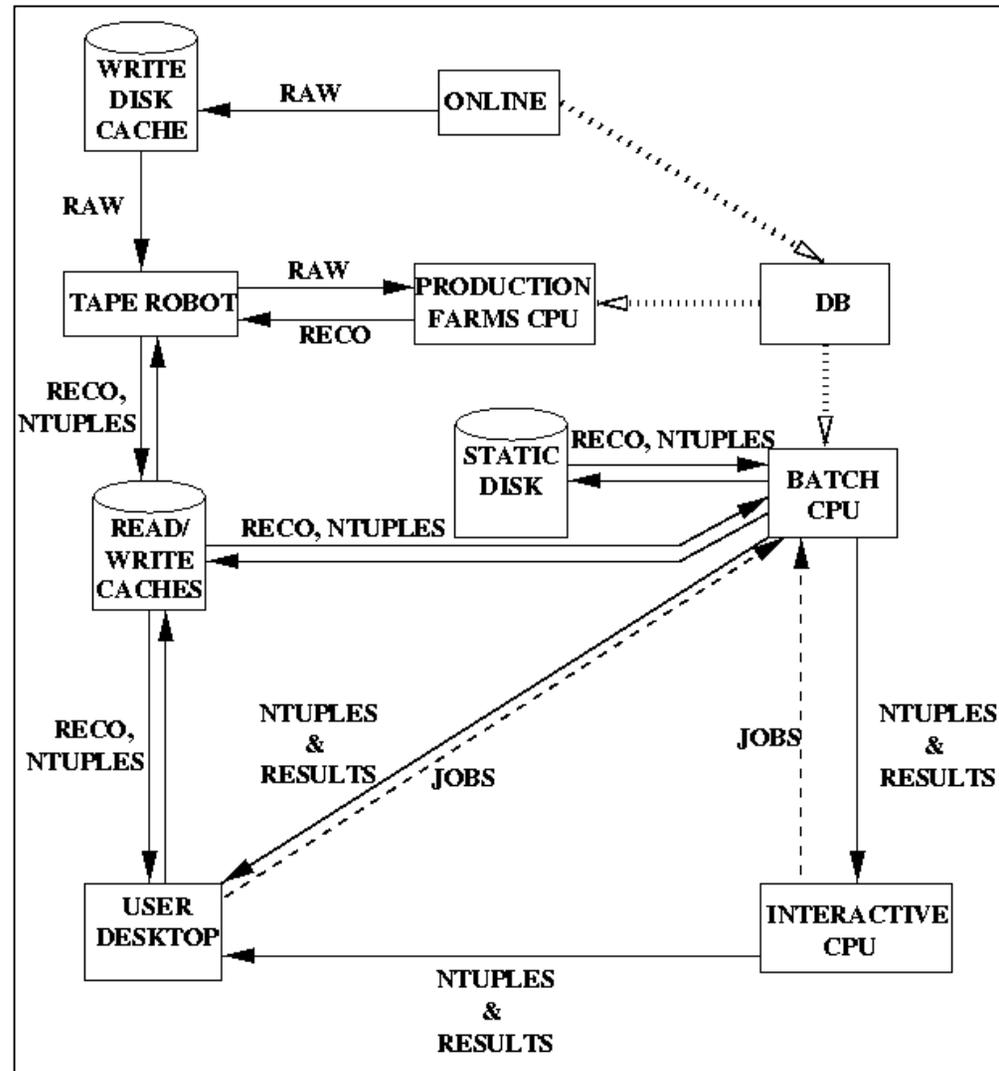
Outline



- Data Handling
- Central Analysis Farms (CAF)
- Other Central Computing Systems
- Databases
- Reconstruction Farms
- Software Infrastructure



- Raw Data
 - ➔ Written to write cache before being archived in tape robot.
 - ➔ Reconstructed by production farms.
- Reconstructed Data
 - ➔ Written by farms to tape robot.
 - ➔ Read by batch CPU via read cache.
 - ➔ Stripped and stored on static disk.
- Batch CPU (CAF).
 - ➔ Produces secondary datasets and root ntuples for static disk.
 - ➔ Analyzes secondary datasets and ntuples.
- Interactive CPU and desktops
 - ➔ Debug, link and send jobs to CAF.
 - ➔ Access data from cache and CAF.
 - ➔ Write data to robot via cache.
- Database and replicas provide
 - ➔ Constants for farms, CAF, users.





Data Handling



- People
CD/CDF: J. Tseng, R. Kennedy, D. Litvintsev, E. Wicklund, R. Herber, A. Kreymer.
CD/ISD: D. Petravick, J. Bakken, B. Alcorn, R. Wellner, Enstore-Admin.
Rutgers: F. Ratnikov Glasgow: R. St. Denis.
- Systems
 - CDF deploys a wide variety of tools for data handling.
 - Disk Inventory Manager (DIM): in use on fcdfsi2 only.
 - dCache: DIM replacement in beta use on most central + trailer systems.
 - rootd: server run on central systems for trailer + remote clients.
 - SAM: the future CDF DH system now in beta use in trailers + remotely.
- CDF/D0/CD Joint Project
 - Run 2 Data Handling and Distributed Computing (R2D2)
 - Coordinate SAM and GRID efforts.
 - Jeff Tseng (CDF) and Lee Leuking (D0) are in charge of SAM.
 - Steering Committee with CDF, D0 and CD representatives.
 - Explicit non-Fermilab experiment representatives.
- See next talk by Jeff Tseng for a more detailed discussion of DH & GRID.



Tape Robot & Drives



- CDFEN: STK Robot using Enstore for data transfer.
 - ➔ Installed in February, used for all DH since May, operated by CD/ISD.
 - ➔ 10 T9940A drives: write at 10 MB/s and 60 GB/cartridge.
 - ➔ Robot capacity of 330 TB has 170 TB of data and is filling at 0.3 TB/day.
- CDF has purchased a 2nd STK robot in FY02.
 - ➔ Will be installed in september on the FCC mezzanine.
 - ➔ We will receive ISD's STKEN robot in exchange.
 - On 2nd floor of FCC next to CDFEN, "pass-through" allows two robots to act as one.
- CDF is purchasing 10 STK T9940B drives in FY02.
 - ➔ Triple I/O and capacity of T9940A drives: 30 MB/s and 200 GB/cartridge.
 - ➔ The drives have been tested by CD/ISD on 25 TB of data and passed.
 - Currently achieve 16 – 22 MB/s but working with STK on improving rate.
- Two robots and 10 T9940B drives upper capacity: 300 MB/s and 2.2 PB.
 - ➔ Meets our needs for next 2 years. Thanks to CD / ISD for hard work on R&D!
 - By avoiding FY03 tape drive costs we can redirect \$300K to analysis CPU & Disk.



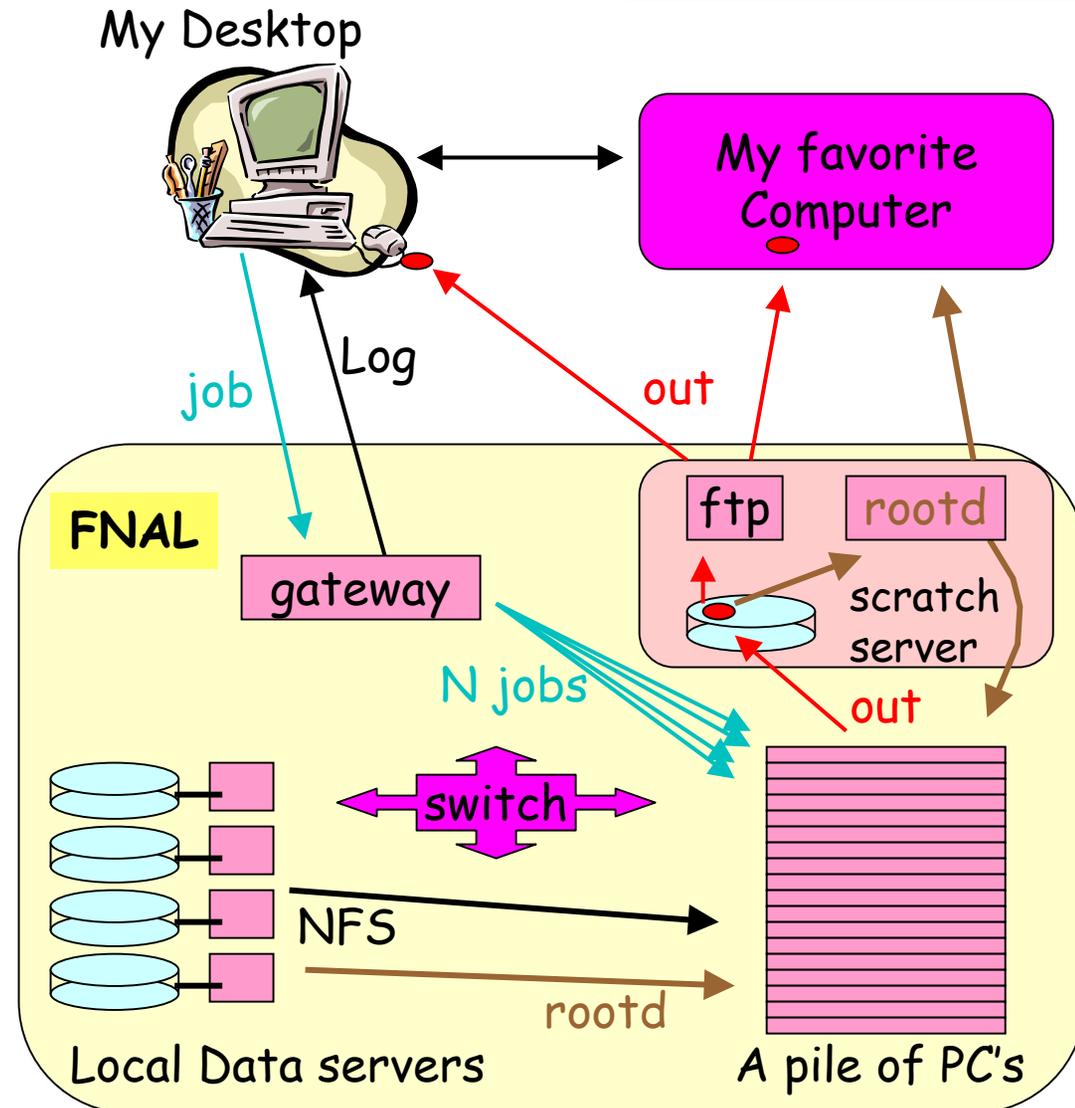
Central Analysis Farms (CAF)



- The CAF is a batch analysis engine utilizing plenty of cheap PCs and disk.
- Broad CDF participation in development of the CAF.
 - **MIT:** T.Kim, M. Neubauer, F. Wuerthwein
 - **CD:** R. Colombo, G. Cooper, R. Harris, R. Jetton, A.Kreymer, I.Mandrighenko, L. Weems
 - **INFN Italy:** S. Belforte, M. Casara, S. Giagu, O. Pinazza, F. Semaria, I. Sfligoi, A. Sidoti
 - **Pittsburgh:** J. Boudreau, Y. Gotra
 - **Rutgers:** F. Ratnikov
 - **Carnegie Mellon:** M. Paulini
 - **Rochester:** K. McFarland
- Current configuration: Stage 1
 - 63 dual worker nodes contributed by CD, Carnegie Mellon, Pittsburgh, INFN.
 - Roughly 160 GHz total, compared to 38 GHz for fcdfsgi2, and 8 GHz for cdfsga.
 - 16 file servers with 2 TB each from CD and MIT.
 - 7 file servers for physics (~50% full), 5 for DH, 2 for development, 2 for user scratch.

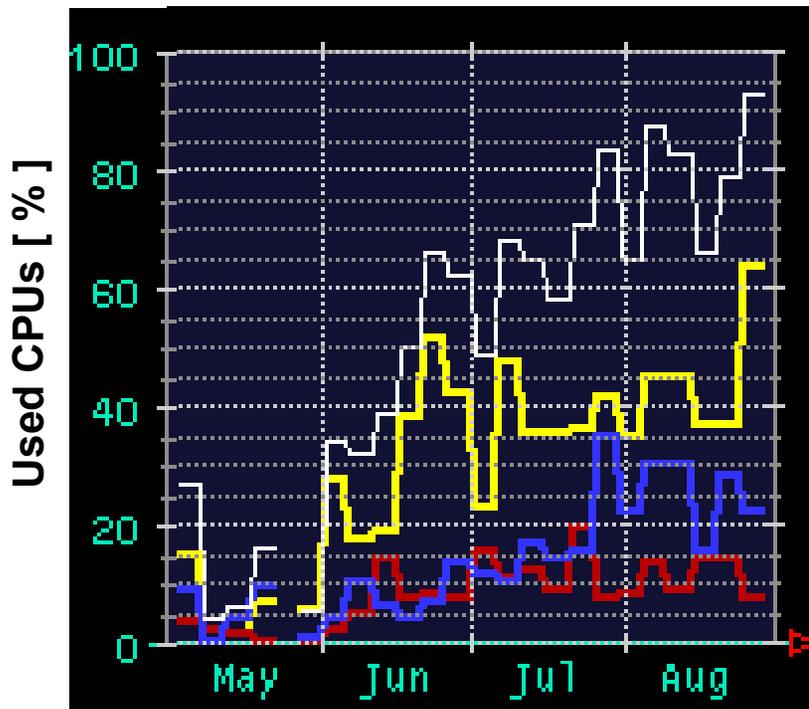


- Compile/link/debug job anywhere (Linux).
- Submit job from anywhere to CAF
 - ➔ Submission of N parallel jobs (sections) with single command.
 - ➔ Jobs run on pile of PC's and access data on network attached filesystems.
- Get output anywhere
- Store output on local scratch disk for reanalysis.
 - ➔ Access to scratch disk from anywhere.



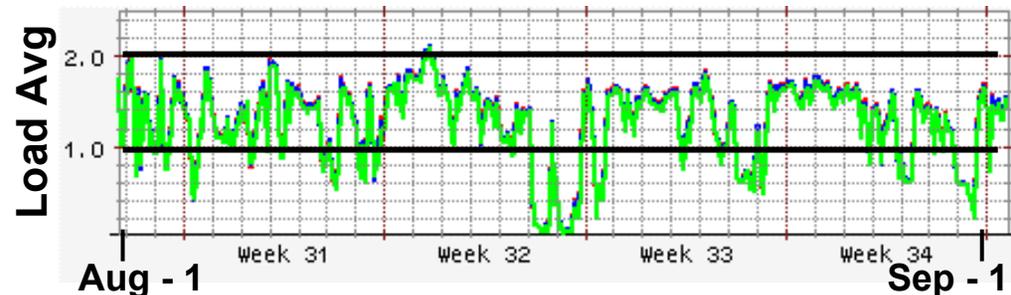


- The usage of the CAF has been ramping up steadily since May.
 - ➔ Supplied critical CPU needed for ICHEP and other summer conferences.



ptype	Average %	Current %
■ short	30.6	63.4
■ Medium	8.1	7.3
■ Long	13.7	21.7
■ All processes	52.4	92.4

Updated: Sep 1 10:55:01 2002



- Clear need for more CPU.
 - ➔ Current usage of 92% and load averages per dual usually between 1 & 2.



CAF CPU Usage (continued)



• Users

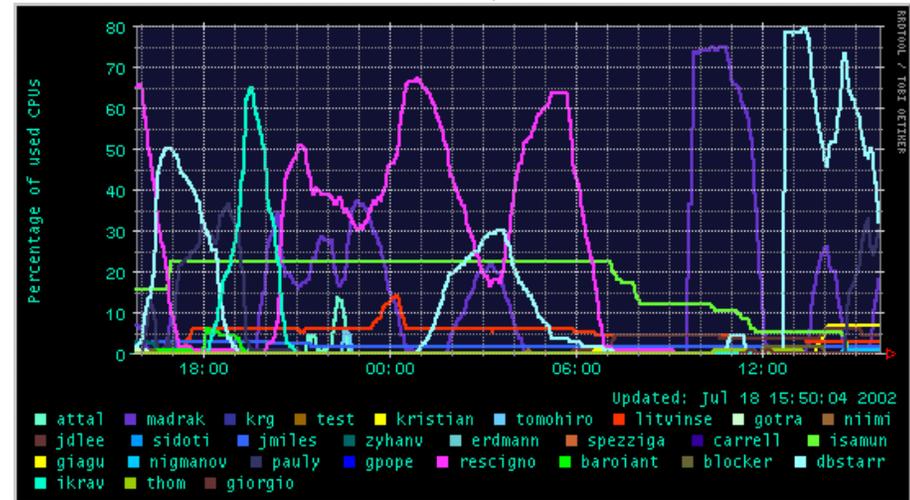
- 235 users signed up
- 30 users per typical day.
- Top 15 used 85% of CAF last month.
- Used by remote collaborators.
- Competition for cycles increasing.

- See plots of two days month apart.

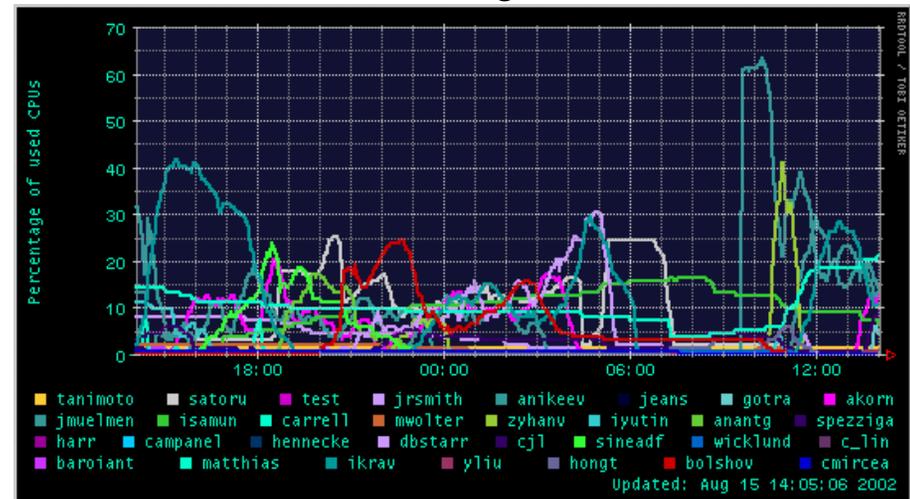
- User support has covered more than just CAF.

- Problems with user software on CAF often sent to CAF team first.
- Users should test short jobs on fcdflnx2 before asking for help from CAF support.

Users July 18



Users Aug 15



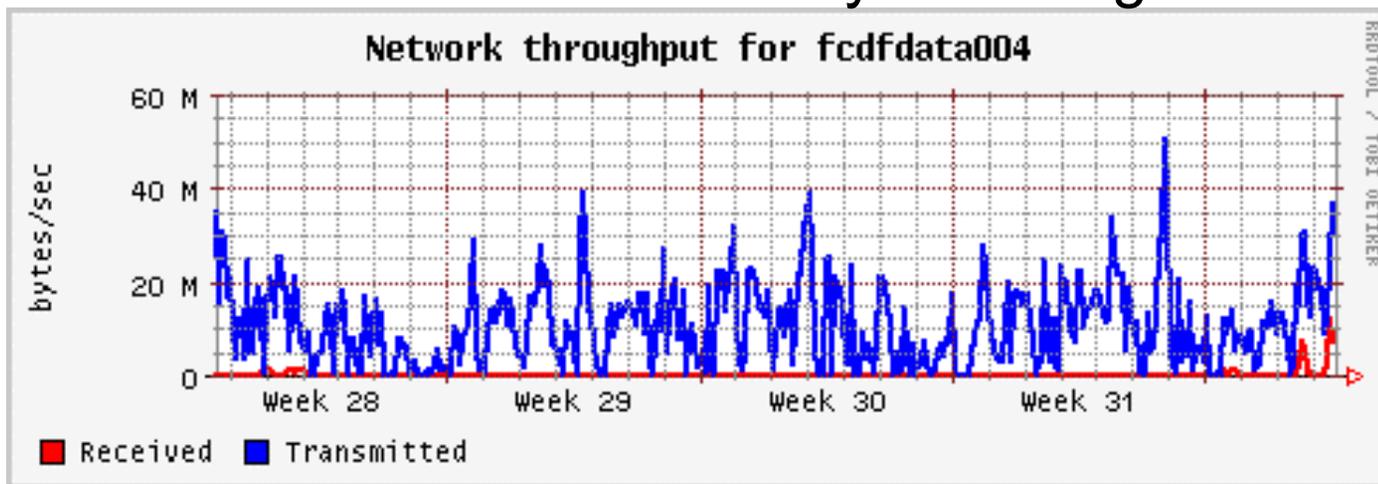


CAF Fileserver I/O Usage



Static Dataset Fileserver July 15 – Aug 15

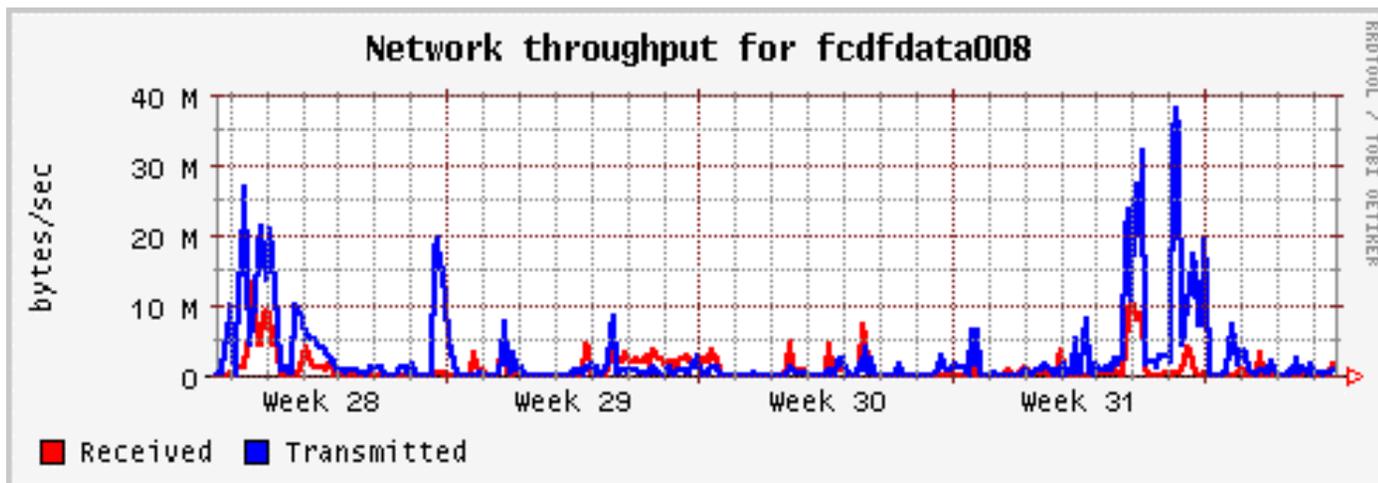
Write once
read many
usage of Static
dataset fileserver



dCache DH Fileserver July 15 – Aug 15

Write many
read many
usage of dCache
fileserver.

dCache
still in beta.

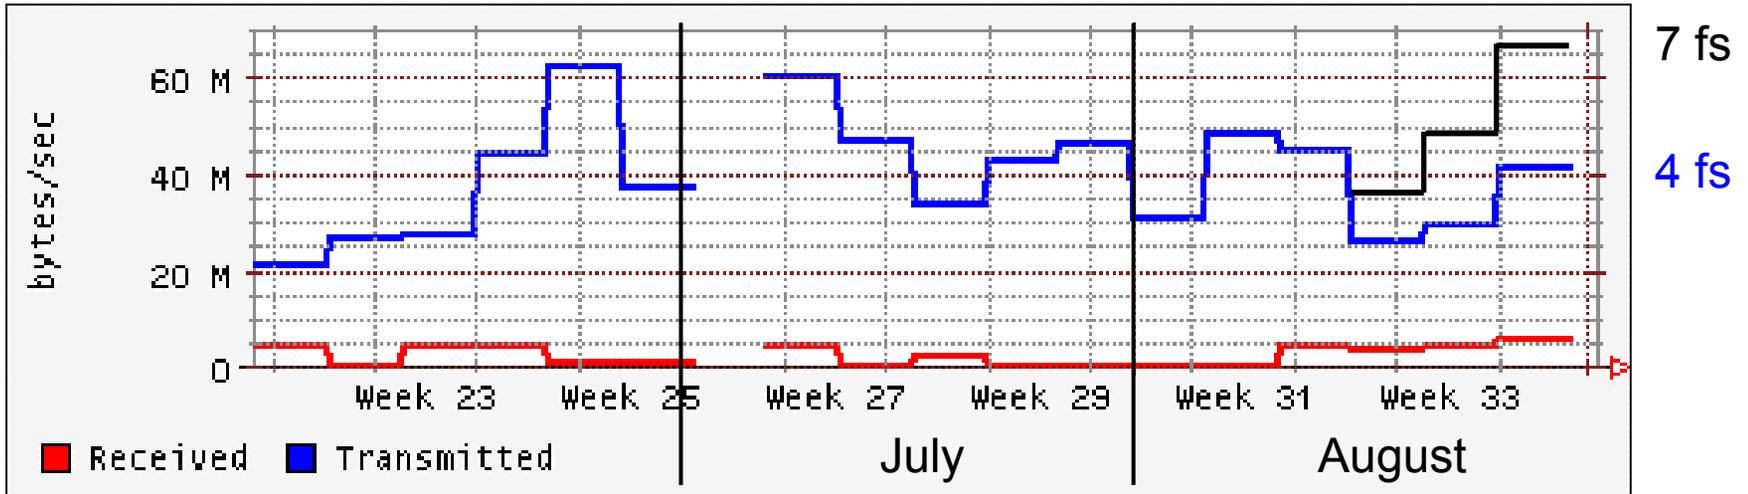




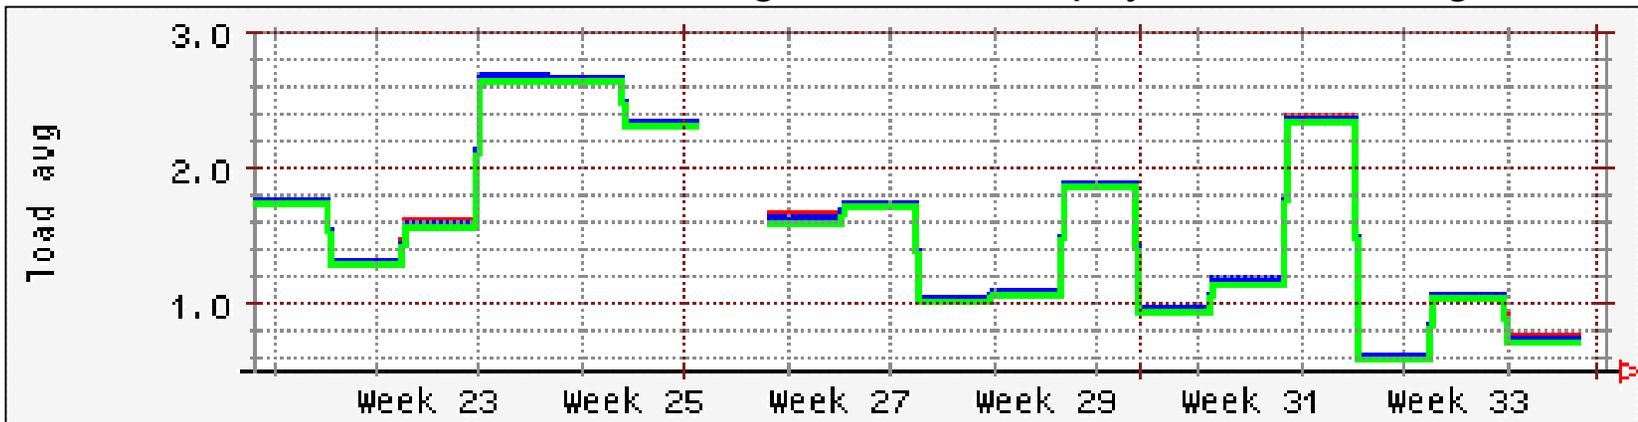
CAF Fileservers for Physics Datasets



- 350 TB served in 3 months at weekly average rates from 22 to 66 MB/s.



- Load average on first 4 fileservers filled was high (2.0 is 100%).
 - Load is more tolerable after using 3 more fs for physics in mid-August.





CAF Operational Issues



- Utilization
 - Wait times for CPU are increasing and winter conferences are approaching.
 - Stage 2 of the CAF should resolve this beginning in November
- Hardware
 - Disk drive "glitches" @ 4/week/kilodrive (contacting 3Ware).
 - Disk drive failures @ 1/week/kilodrive as expected.
 - Translates to 1.6 drives per week for 100 servers with 16 drives.
 - Three instances of power failure for a single fileserver is under investigation.
 - Jobs getting hung on worker nodes
 - Sleeping and consuming no CPU until they run out of time and are killed.
 - ServerWorks IDE chipset problems on Intel worker nodes.
 - Causing file system corruption which hangs one node per week now.
 - Buying Athlons in the future would resolve this.
- Software
 - Learning how to operate the CAF for best possible user experience.
- DB
 - CAF is stressing the DB in terms of both number of connections and CPU.



CAF Stage 2



- A five fold increase of the CAF is underway.
 - CPU is being bid on by the vendors.
 - 184 Dual Worker Nodes with roughly 1.7 GHz processors.
 - **108 FNAL, 38 Japan, 32 Italy, 3 JHU, 2 Yale, 1 UCD.**
 - Disk server bids are back from vendors and being evaluated
 - 66 file servers with 2 TB each.
 - **35 FNAL, 15 UK, 3 Japan, 2 Italy, 11 from universities**
 - **Purdue, UIUC, JHU, UCD, Penn, Cantabria, Toronto, Florida, Yale, Korea, Duke and Harvard.**
 - Networking
 - CISCO 6513 switch with GigE and FE modules was delivered on Aug-15.
 - Facilities
 - Space and power has been made available on 1st floor of FCC.
- Institutions contributing to CAF will receive proportional benefits.
 - Disk paid for by institution will be theirs to allocate.
 - Half for arbitrary usage, half for physics datasets they host accessible to all CDF.
 - CPU purchased by institution will be compensated for with special queues.
 - Gives institution high priority on their CPU.
 - CPU cycles unused by institution will be used by entire collaboration.

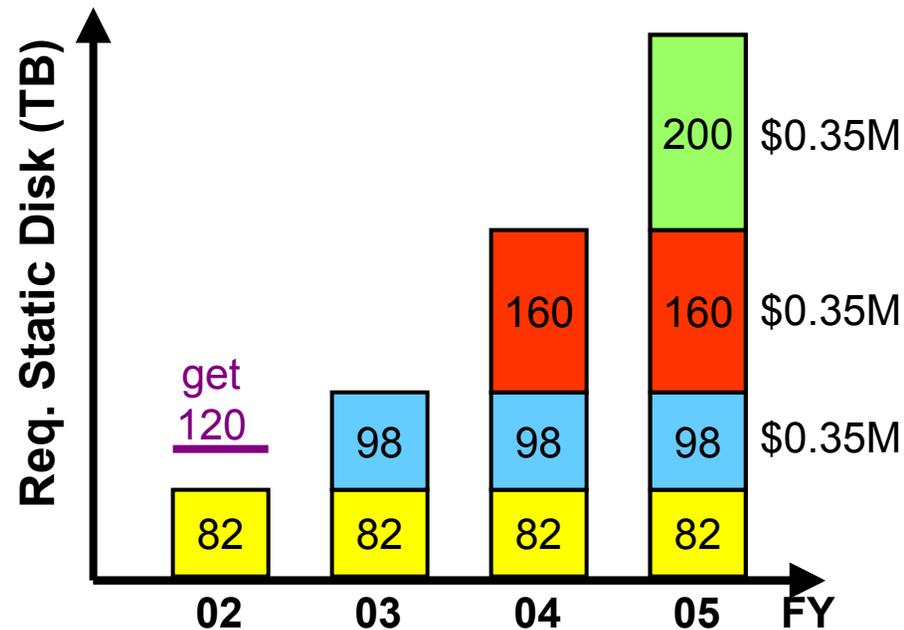
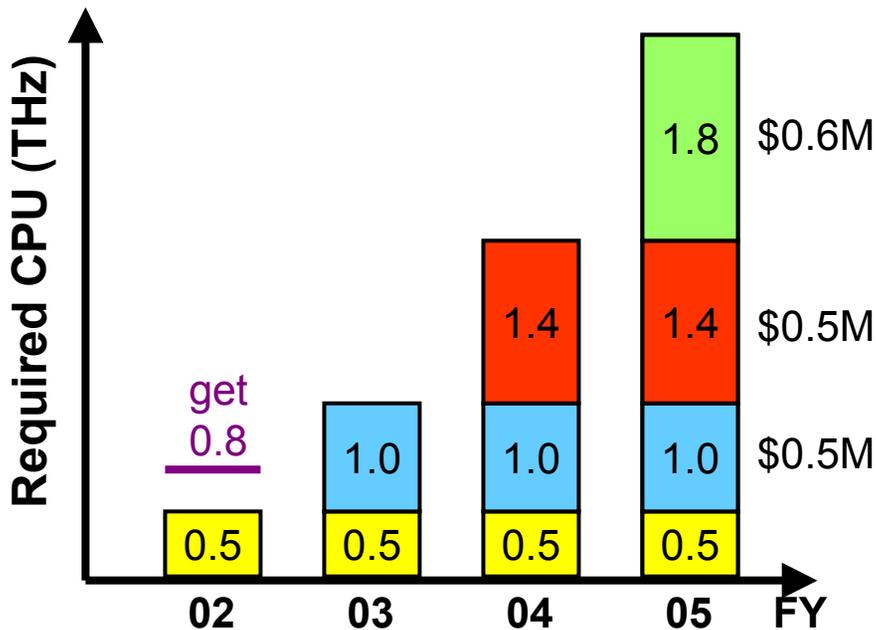


CAF over Longer Term



- Issues

- LAN challenge
 - Two switches with potential bottleneck between them. Then three, then four, ...
- SAM on CAF.
- DB replication to support CAF load.
- Interactive “login pool”.
- Meeting requirements from CDF 5914 in face of potentially dwindling budget.
 - FY02 exceeds requirements, but FY03 budget guidance reduced from \$2M to \$1.5M.

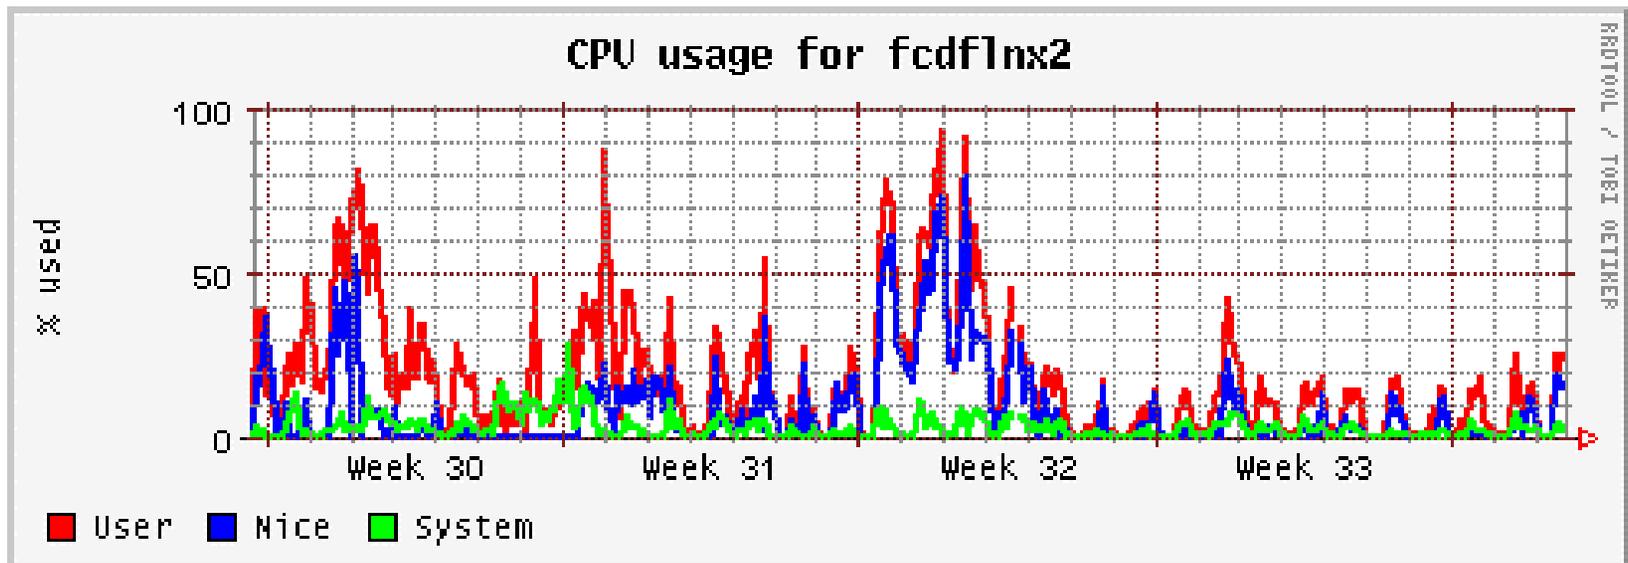




Central Interactive Linux Systems



- fcdflnx2 and 3
 - Supported by Lance Weems and CDF Task Force.
 - Heavily loaded machines for CAF job submission job and Linux development.
 - Each an 8-processor 700 MHz Intel/Linux box.
 - 2 TB of scratch space for users.
 - 1 TB of production farm space on fcdflnx3
 - Plan to expand to 2 TB
 - 500 accounts, 70 logons per machine, 30 active users.





Central Interactive IRIX Systems



- Fcdfsgi2
 - Supported by Glenn Cooper and CDF Task Force.
 - Run 2 SGI now operating fairly stably.
 - Moderately loaded at 40-80% of full load average (128).
 - 128 processors 300 MHz with up to 39 TB of disk.
 - 27 TB assigned
 - 12 TB data handling cache disk in disk inventory manager
 - 10 TB physics groups static datasets
 - 2 TB raw and production data look areas
 - 1 TB physics stripping areas
 - 1 TB detector and offline groups
 - 1 TB miscellaneous
 - 7 TB available but not yet assigned
 - **Most of this will be assigned for winter conference use.**
 - 5 TB in various stages of hardware repairs
 - There will be no more disk purchased for fcdfsgi2 in the future.



Central and Desktop Systems



- cdfsga
 - Supported by Glenn Cooper and CDF Task Force.
 - 28 processors 300 MHz with 3 TB of disk (all used).
 - Heavily loaded at 75 – 100% of full load average.
 - Run 1 students need cdfsga cycles. Please use cdfsga for run 1 analysis only.
- cdfsga and fcdfsgi2 will not be supported forever.
 - cdfsga needs to be available until larger run 2 datasets exist.
 - Migration of users off cdfsga could begin soon after that (late 2003?).
 - fcdfsgi2 could remain until we have a large interactive Linux system.
 - Probably interactive login pools with large amounts of shared disk.
 - After that we need to start migrating users off of fcdfsgi2.
 - We may reduce the number of processors as early as FY2003.
- Desktop Computing
 - Supported by Jason Harrington, Mark Schmitz and CDF Task Force.
 - 295 Linux nodes and 40 IRIX nodes.
 - Code Server: replacing 300 MHz & 100 Mbit with Dual 1.4 GHz & Gbit server.
 - Hopefully will reduce linking times in trailers.



Databases



- People

Texas Tech: A. Sill, J. Cranshaw
CD/ODS: J. Trumbo, N. Stanfield,
A. Kumar, D. Box, Y. Gao,
L. Lebedova, M. Vittone.
CD/CPD: L. Sexton-Kennedy,
J. Kowalkowski
CD/CDF: D. Litvintsev, E. Wicklund,
A. Kreymer, R. Herber,
R. Jetton.
U.C. London: R. Snihur, D. Waters.
OSU: R. Hughes.
PPD/CDF: B. Badgett, F. Chlebana
INFN: Rodolfo Pellizoni

- Most swamped with operations
 - Need help in many areas.

- Hardware

- Production and development DB.
 - Online Suns b0dau35 & 36.
 - Offline Suns fcdfora1 & 2.
- Online/Offline replication working.
- Online DB stable.
- Offline production DB overloaded.

- Offline DB load problems.

- CAF jobs with many simultaneous sections contribute to critical DB load
 - See example on next slide.
 - Lead CAF to modify how they started many jobs that access DB.
 - Introduced a 10 s startup delay.
 - Situation has improved, but load on fcdfora1 is still unacceptable.
- Bugs & poor DB usage in the CDF software overloads the DB.
 - DB group has made progress here.

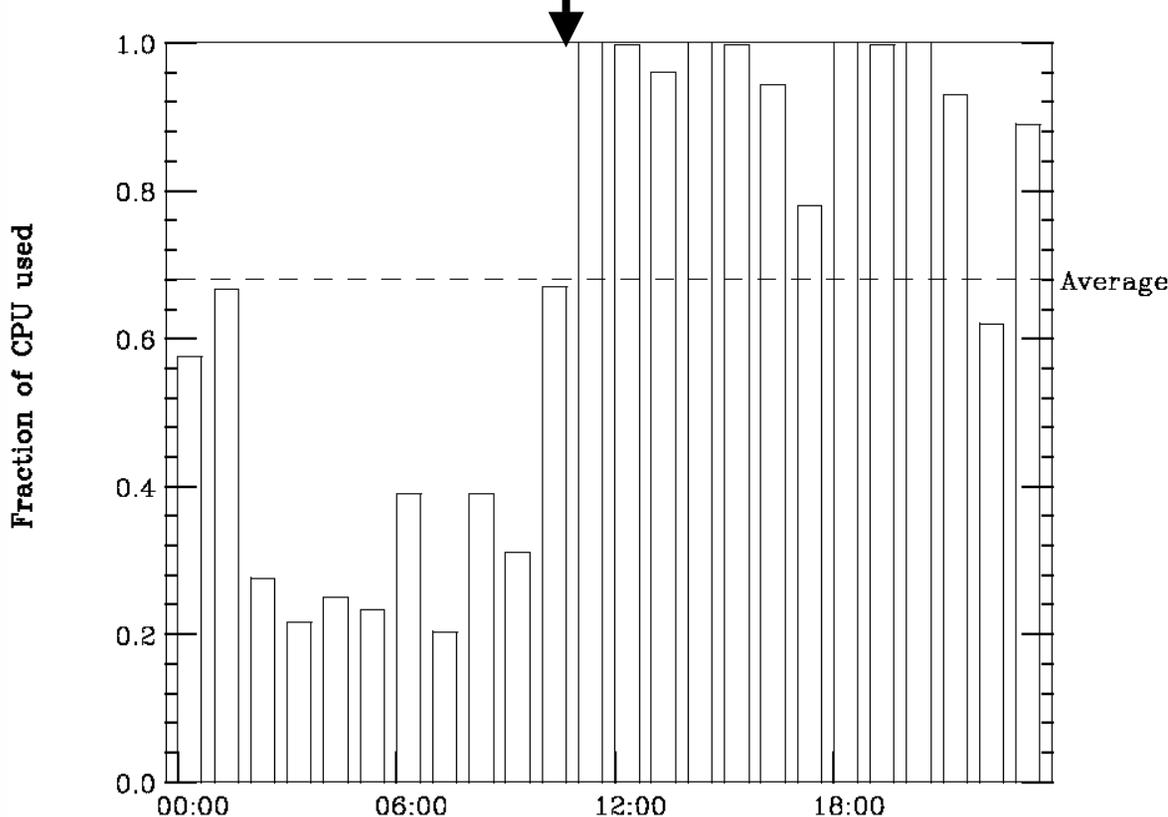
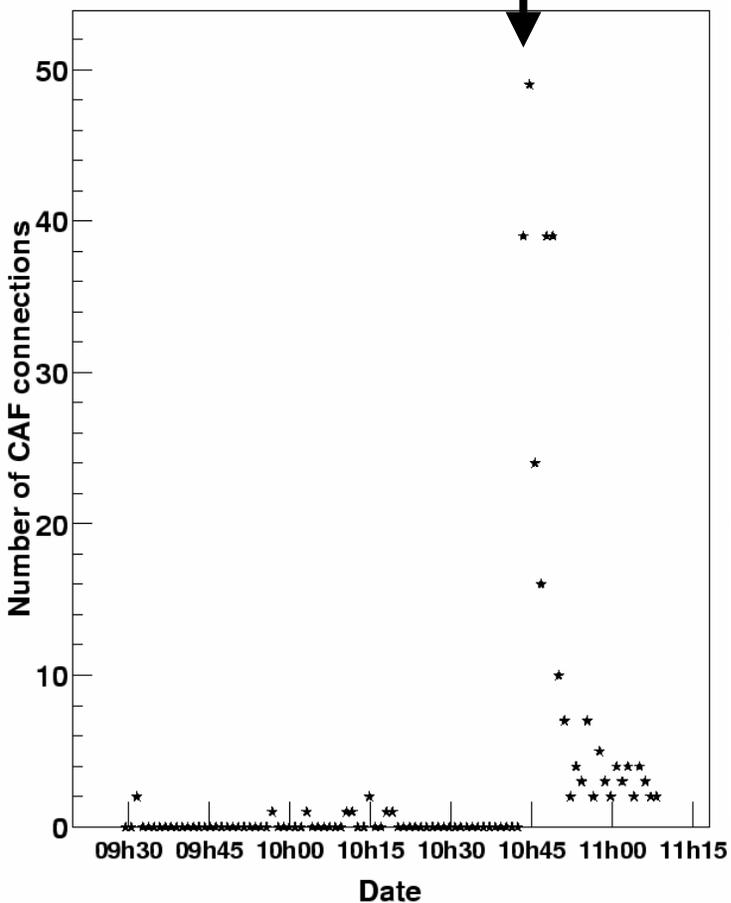


Example: DB overload initiated by CAF?

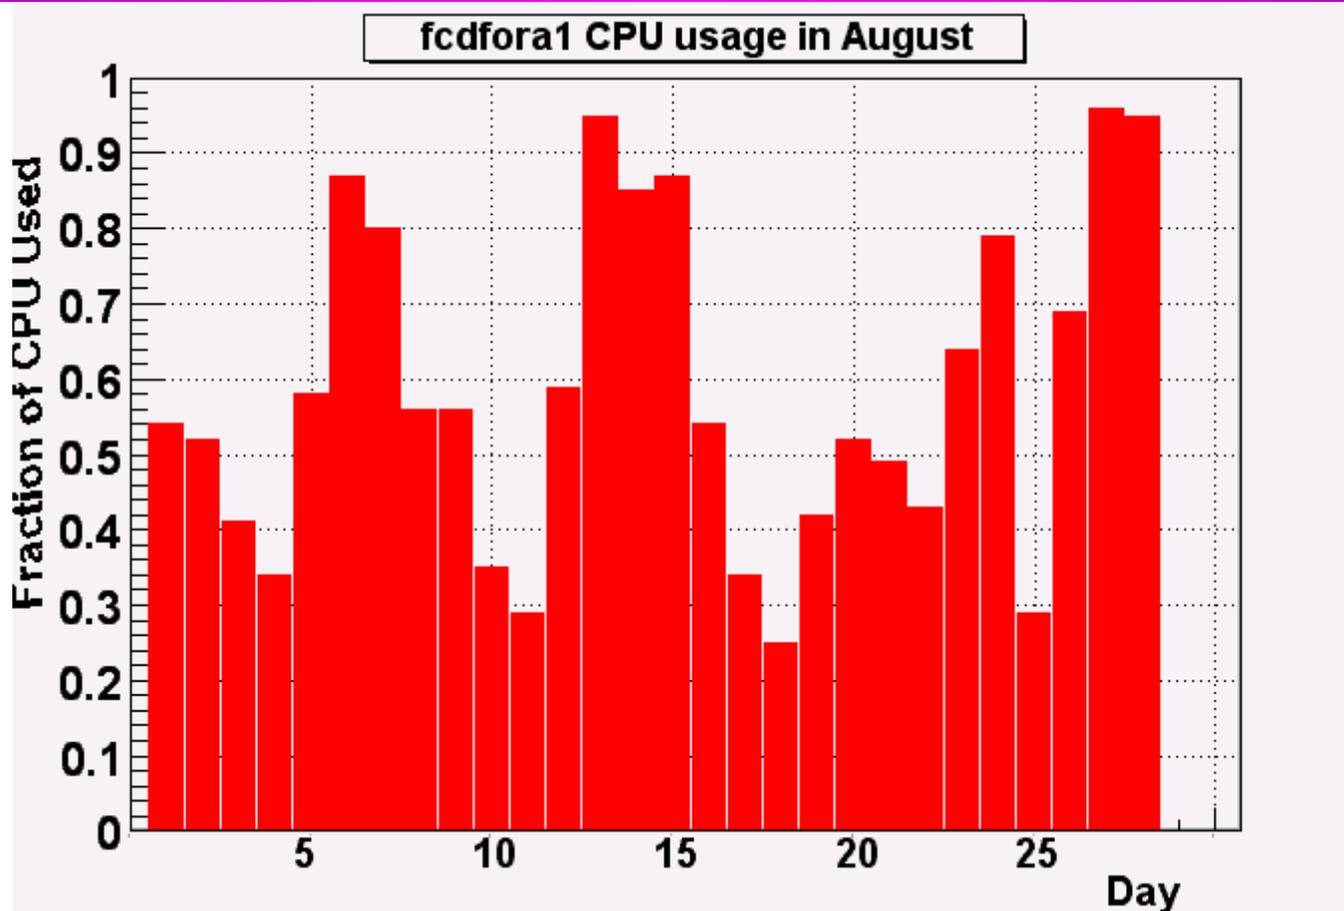


CAF connections to database

fcdfora1 Fri Jul 12 2002



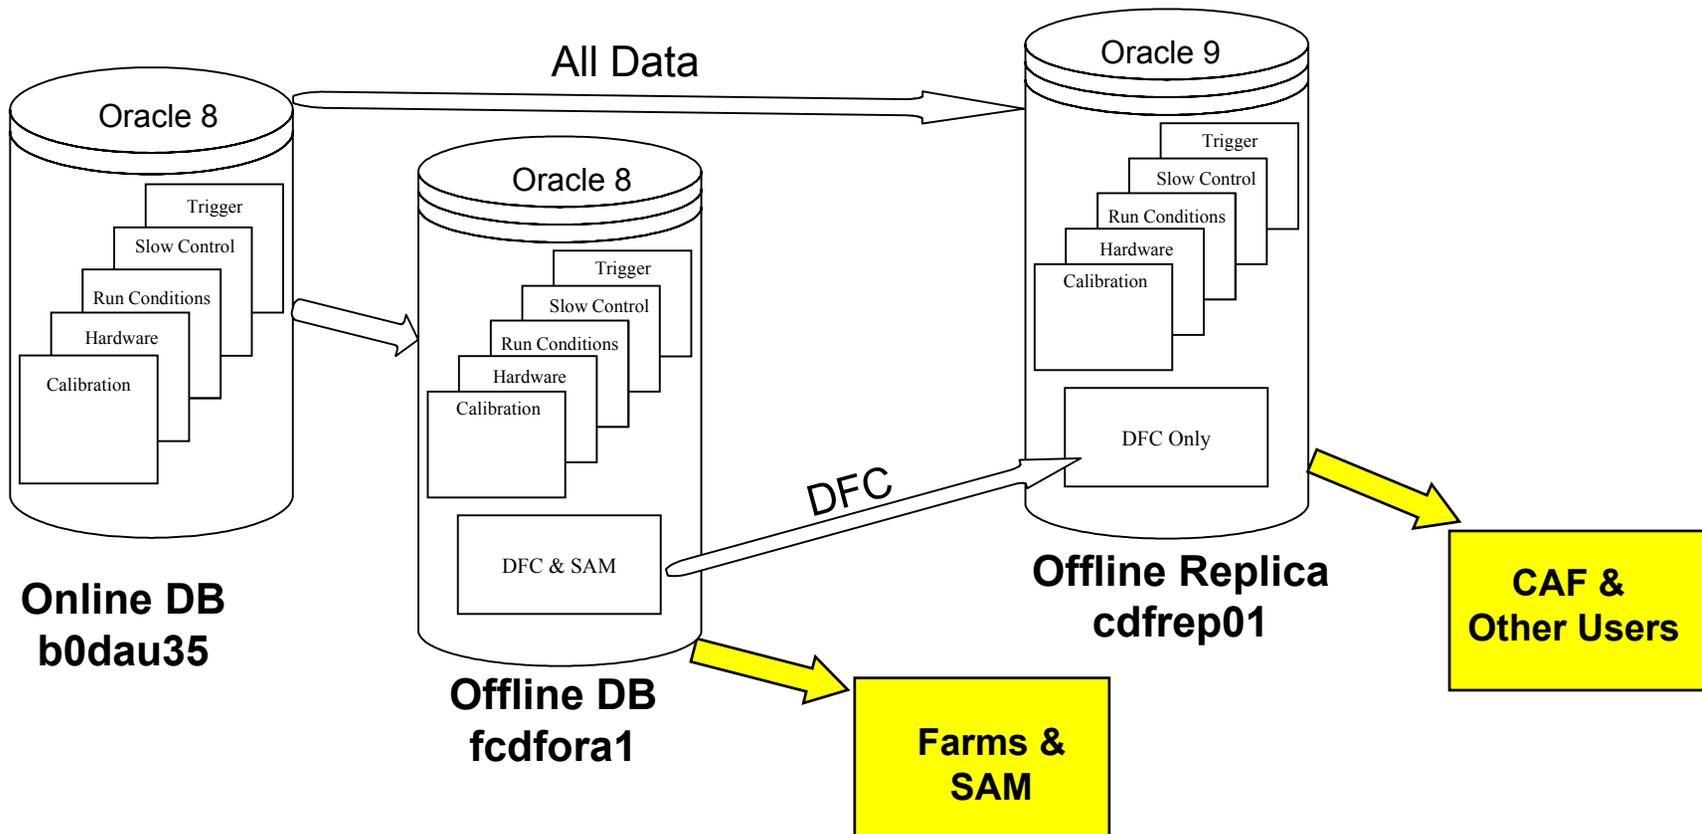
Spikes of 100% usage in 7 separate hours after CAF job startup.



- Unacceptable loads of up to 95% CPU use over day, 100% for hours.
 - ➔ Causes long delays, connection timeouts, and interferes with farms operations.

Plan to handle offline DB load

- Replicate DB to more powerful machine
 - ➔ Original DB for farms and SAM, replicas for CAF / Users.
 - cdfrep01: 4-way 700 MHz Linux replica, 6 times fcdfora1, now operating in beta.
 - ➔ Introduce additional replicas as necessary.





DB Plans



- Near term plans:
 - Proceed with timing, performance, replication, and other operational tests
 - Implement at least one load-sharing offline replica
 - Finnish statistics code in the calibration API
 - Important to understand usage patterns and to help in diagnosing problems.
- Longer term plans through rest of Run 11a:
 - Develop capability to field and deploy multiple replicas (Oracle 9i v2)
 - Development of connection broker by CD/ODS.
 - Controls number of connections and coordinates which replica gets a connection.
 - Begin to prepare for grid-like deployment
- If time allows
 - Implement freeware replicas of parts of the database.
 - CD/ODS is willing to contribute personnel here. Option for long term CD strategy.
 - Making all applications work under freeware is a major effort that requires people.



Some areas where people are needed.



- Database monitoring
 - Analyze connection usage patterns using the statistics package being developed.
- API design, coding and testing
 - Need APIs to read the Hardware, Run Configuration, Trigger & Slow Controls DB.
- Freeware port of database and replication
 - More CDF people needed for databases other than calibration.
- Connection broker / negotiator design & testing
 - Help develop and implement an architecture that will protect the database from overload.
- SAM/Database/Grid test stand
 - Setup small-scale array to test distributing & connecting to database in various ways
- Study of slow controls and monitoring system
 - Guide redesign of tables and construction of accessors with eye toward analysis.
- Replication procedures and design
 - Possibility of secondary sourcing, prefetching of data, and replicating only portion of DB.



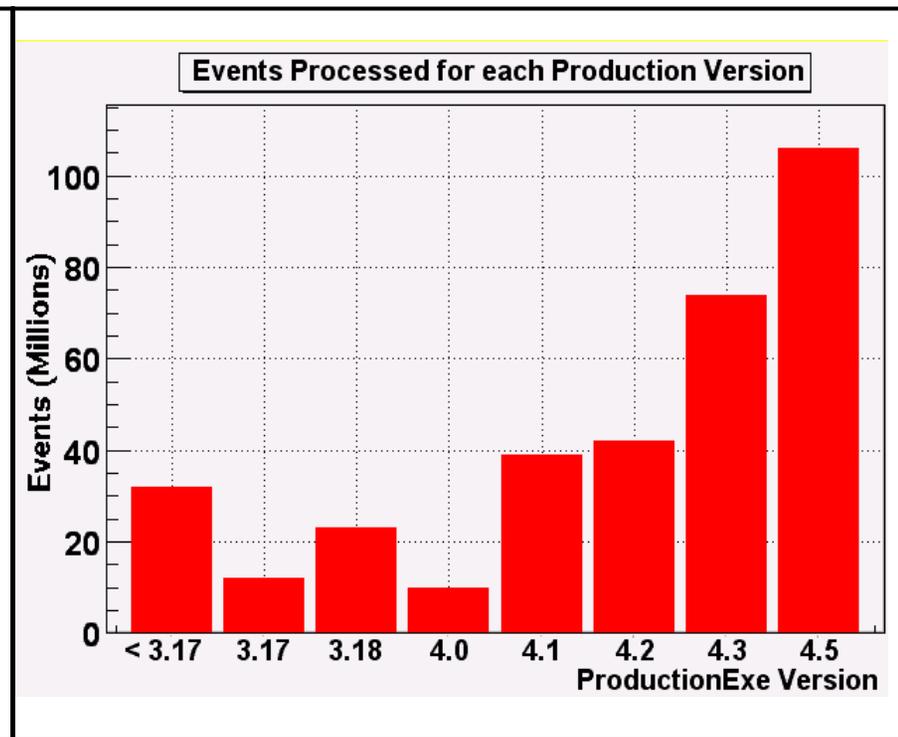
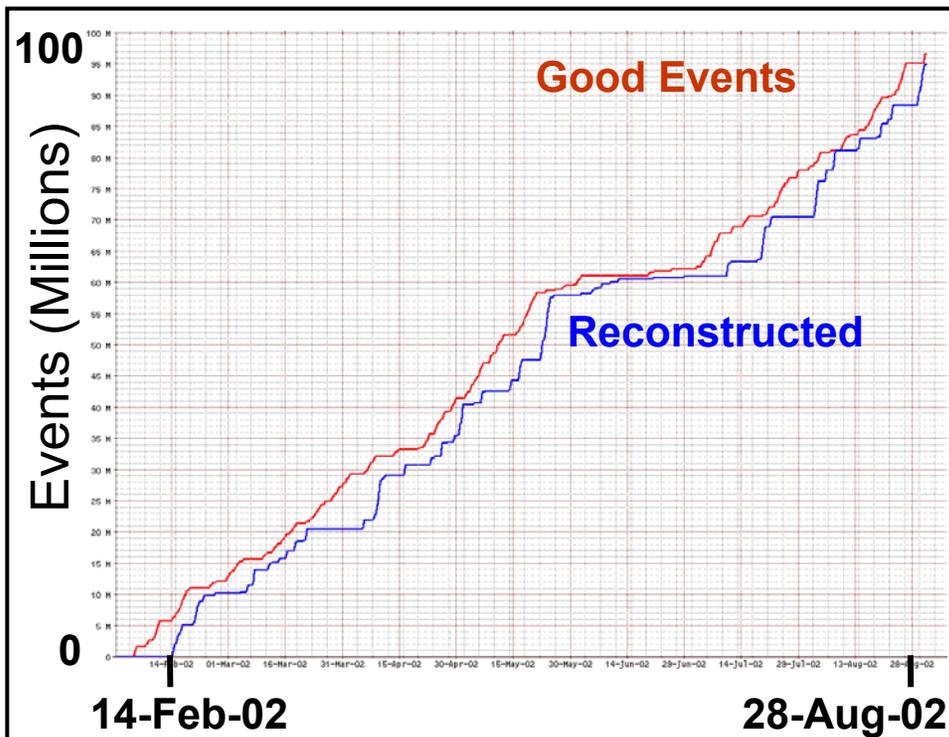
Reconstruction Farms



- Contributors: M. Siket, M. Babik, Y. Chen, S. Wolbers, G. P. Yeh.
 - R. Lysak joining in October.
 - S. Timm leads CD/OSS support.
- Hardware
 - 169 worker nodes
 - 50 x 500 MHz duals
 - 23 x 800 MHz duals
 - 64 x 1 GHz duals
 - 32 x 1.26 GHz duals
 - Equal to 600 x 500 MHz
 - 32 duals more in FY02
 - Networking
 - CISCO 6509 switch
 - 165 duals on 100 Mbit
 - 4 duals on Gbit
 - Input and Web Servers
- Data Flow
 - Data staged from tape robot and copied to worker nodes.
 - Output stored on worker disk.
 - Output copied to 4 concatenation nodes (Gbit).
 - Concatenated filesets are copied to tape.
- Capacity
 - 5 million ev/day assuming 8 sec/event on 500 MHz and 80% efficiency.
- Calibrations
 - Farms often wait 2-4 days for calibrations to be ready.



- Farms keeping up with data and have lots of reprocessing capacity.
 - ➔ 95 million events reconstructed at least once since 14-Feb-02.
 - ➔ Large datasets from stable executables used for physics: 4.3.2 & 4.5.2.
 - ➔ See Pierre Savard's talk on Friday for future farms reconstruction plans.
 - Rapid progress has been made on a ProductionExe for winter conferences.





MC Production on Farms



- 6.1 million events requested by physics groups / users processed in last 3 months.

<u>Physics Group</u>	<u>Requester</u>	<u>Process</u>	<u>Generator</u>	<u>Executable</u>	<u>Events</u>
Top/EWK	Vaiciulis	t tbar	Herwig	4.5.0	246 K
QCD	Field	Dijets	"	"	300 K
Exotics	Culbertson	Diphoton	Pythia	4.5.3	90 K
"	"	Photon	"	"	90 K
"	Pratt	Z -> mu mu	"	"	1040 K
"	Pratt	W -> mu nu	"	"	900 K
Top/EWK	Holker	Z -> e e	"	"	300 K
"	Goldstein	W -> e nu	Herwig	"	300 K
"	Cabrera	W W	Pythia	"	200 K
"	Coca	Z -> tau tau	Pythia & Herwig	"	720 K
QCD	Wyatt	Dijets	Herwig	"	600 K
?	Frisch	?	CompHEP	"	1000 K
Top/EWK	Kuznestova	b bbar	Pythia	4.6.2	1K
"	Lys	Photon	"	"	128 K
"	Vaiciulis	t tbar	Herwig	"	201 K

- Farms can do ~200K ev/day with 20% of CPU, assuming 40 s/ev on 500 MHz PIII.
 - If significantly more events are needed we can use remote resources or the CAF.
 - IPP is considering producing ~300 Million events/year on their 224 node Beowulf cluster.
 - The MC could be written to tape at Fermilab using SAM at an average rate of ~1 MB/s.
 - Efforts like this are welcome and the run 2 computing review encouraged central coordination.



Software Infrastructure: Compilers



- People: R. Kennedy, L. Kennedy, A. Kreymer, A. Beretvas, C. Debaun
- KAI is going away
 - New licenses outside of FNAL.GOV domain will not be available after Dec. 31.
- gcc v3.1
 - Port has been essentially completed.
 - CDF code compiles, links and runs.
 - gcc v3.1.x patched to eliminate unused debugger symbols. Fed back to gcc support.
 - Reduced library sizes so that library and executable sizes comparable to KAI.
 - Compilation and execution times and size comparable to KAI.
 - Link times much larger with g++ than KAI.
 - COT Tracking gives essentially the same results with both KAI and gcc.
 - Soon we will turn on automatic notifications to developers of gcc build failures.
- Issues
 - Transition plans for making gcc a supported compiler by end of year.
 - How long can we support two compilers?
 - Desirable but unclear if we have enough manpower.



Software Infrastructure



- FRHL 7.3
 - Out in beta. Final version intended for use in trailers, farms, CAF when ready.
 - Native support of kernel 2.4.18 needed for CDF software.
 - Device drivers that work with recent hardware.
 - Issue: Linker LD is consuming 2.5 times more memory than 7.1
 - Unknown currently whether simple fix exists.
- Debuggers
 - Have fermilab supported patched version of gdb to work with KAI.
 - Developed by Rob Kennedy for gdb version 5.0.
 - When we migrate to gcc then gdb will work without any patches.
 - Recommendation: Gdb + ddd for linux, totalview for IRIX.
- Memory leak checkers.
 - Purify on IRIX is state of the art.
 - Under Linux none of the half dozen works well enough for us.
 - Deficiencies with large multi-threaded executables like our production executable.
 - Temporary patch: wrote our own tool for productionExe memory monitoring.
 - Have a memory leak checker that works under Linux? Please, help us test it.



Summary



- CDF is working with D0 and CD on DH, SAM and GRID.
 - Cost to develop and support solutions for CDF alone is too high.
- CAF stage 1 delivered the computing we needed for ICHEP.
 - Stage 2 for the winter conferences will be 5 times more powerful.
- Databases need more attention from CDF.
 - A critical part of our analysis infrastructure that must not fail.
 - Proceeding with a reasonable plan but need more people.
- Farms are keeping up with the raw data and producing requested MC.
 - Plenty of reprocessing capacity, and used only 1/3 of MC budget (20%).
- The port to gcc is essentially complete.
 - Significant work remains in the KAI to gcc transition and support.



CDF Department Monthly Status Report

Nov. 7, 2002



Personnel



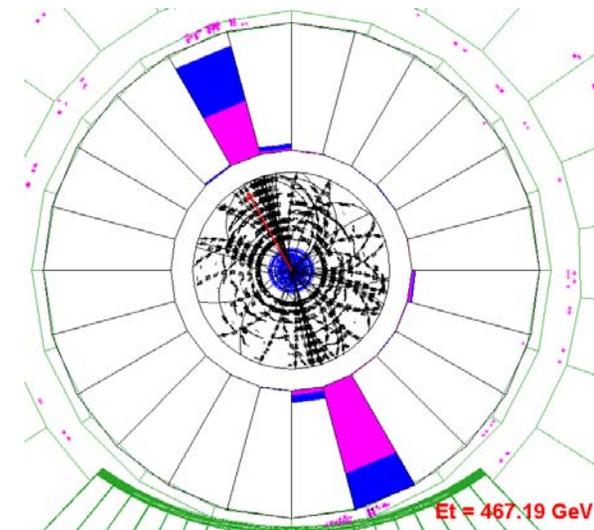
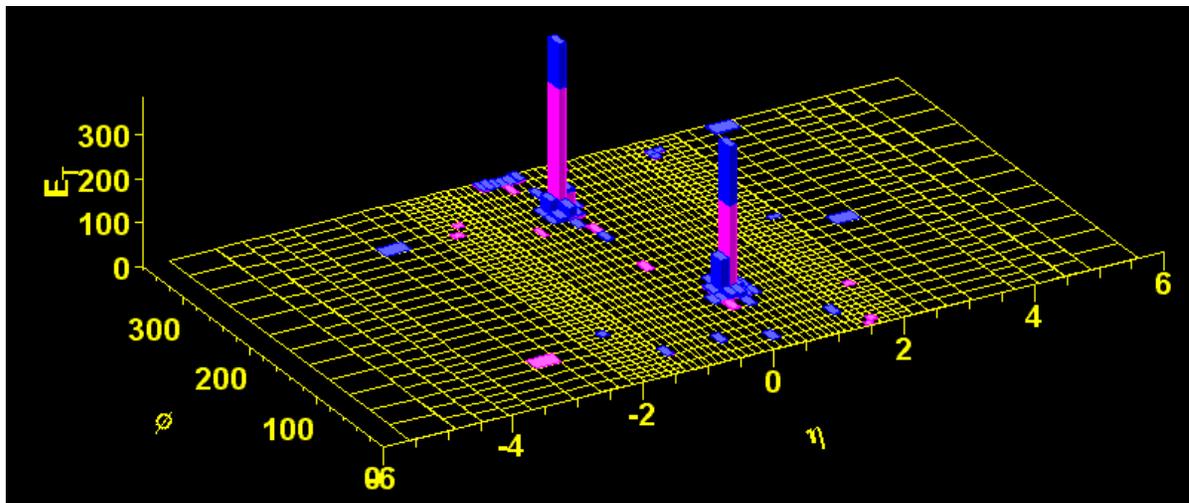
- Reviews
 - Completed all FY02 performances reviews and FY03 goals.

- Collaboration
 - CDF executive board on Nov. 14 will consider UCSD proposal to join CDF
 - Frank Wurthwein will leave MIT to be an Associate Professor at UCSD.
 - First proposal based solely on computing.
 - CAF: Frank Wuerthwein (80%), 2 post-docs, 1 student.
 - GRID: Frank Wuerthwein (20%), Jim Branson (50%), 1 Comp. Prof., 1 post-doc.
 - MC EvtGen: 1 student.
 - Develop, operate and maintain CAF infrastructure software.
 - Hardware maintenance will continue to be done by CDF task force.
 - Develop CDF GRID with emphasis on using remote computing farms.
 - Wurthwein will extend CAF to remote computing farms.
 - Jim Branson, leader of CMS distributed computing at UCSD, will join.
 - CD should give strong support to this proposal.
 - UCSD is volunteering for operational responsibilities that historically were CDs.
 - Trend that I believe we should strongly encourage.

- Luminosity

- Currently have about 60 pb^{-1} on tape. Run 1 was 100 pb^{-1} .
 - Last week was a record 6.9 pb^{-1} delivered, 5.8 pb^{-1} to tape (84% efficiency).
- Projections presented at DOE Review of Tevatron Run 2 Plans
 - FY03: base projection 200 pb^{-1} , stretch projection 320 pb^{-1}
 - Consistent with CDF plan of 200 pb^{-1} for Lepton-Photon at Fermilab.
 - Support emerging for moving 6 week shutdown to summer
 - FY03-FY08: Base 6.5 fb^{-1} stretch 11.0 fb^{-1} . 2-3 fb^{-1}/yr after FY08

Dijet Event with High Mass: $\sim 1.1 \text{ TeV}$





Central Analysis Farms (CAF)

F. Wurthwein & CDF Task Force



- Stage 1 use last 3 months

- CPU usage

- All 120 CPUs being used.
- Short queue increasingly popular.

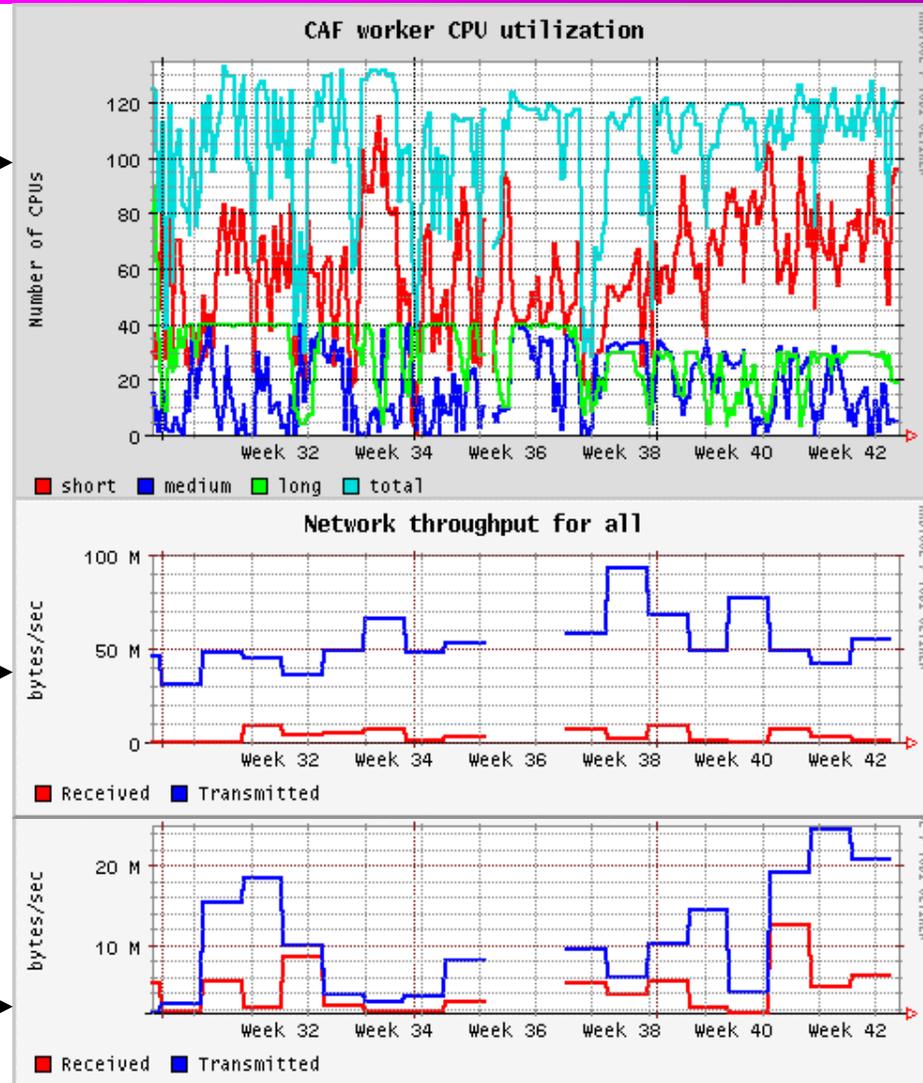


- I/O usage

- 7 fileservers (13 TB) holding static datasets being read at total rate of 50 -100 MB/s.



- 4 fileservers (7 TB) for dCache beginning to be heavily used at 5-10 MB/s total writing and 20 MB/s total reading.





CAF Hardware & Development

F. Wurthwein & CDF Task Force



- Stage 2 Hardware
 - Fileservers
 - 25 fs (50 TB) out of 75 total (150 TB) have arrived.
 - 8 have been assigned to dCache for immediate use.
 - Worker Nodes
 - 238 Athlon MP2000+ (1.67 GHz) duals expected to arrive November 12.
 - Networking
 - 6513 in place with most of the needed gigabit and FE modules.
 - Infrastructure
 - FESS is working on getting necessary power.
 - **Early retirement of 50 PII 500 MHz farm nodes from CDF should help.**
- Development
 - Decentralized CAF (DCAF)
 - Installed on development CAF in B0 trailers.
 - Working out issues regarding packaging of CAF for offsite users.
 - SAM on CAF: working out issues.
 - Monitoring: starting to think about long term strategy.



Data Handling Operations

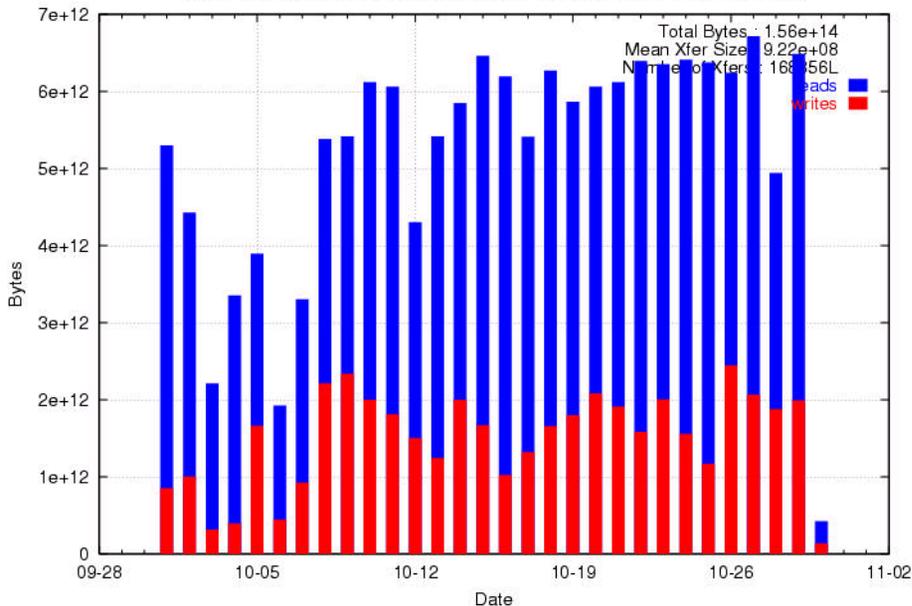
J. Tseng, D. Litvintsev, R. Harris, R. Kennedy, E. Wicklund, ISD



- Saturating capability of CDFEN robot
 - ➔ Transferring about 6 TB/day: 2 TB/day writes and 4 TB/day reads.
 - Eric has been chasing down inefficiencies and helping to improve performance.
 - ➔ All 10 T9940A drives are being used all the time.
 - Have limited the farms to roughly 5 drives to avoid locking out users.
- Pass through between CDFEN and new robot being installed today. Thanks ISD!
 - ➔ Thanks D0 for loaning 5 mover nodes for new robot while our movers are being bought.

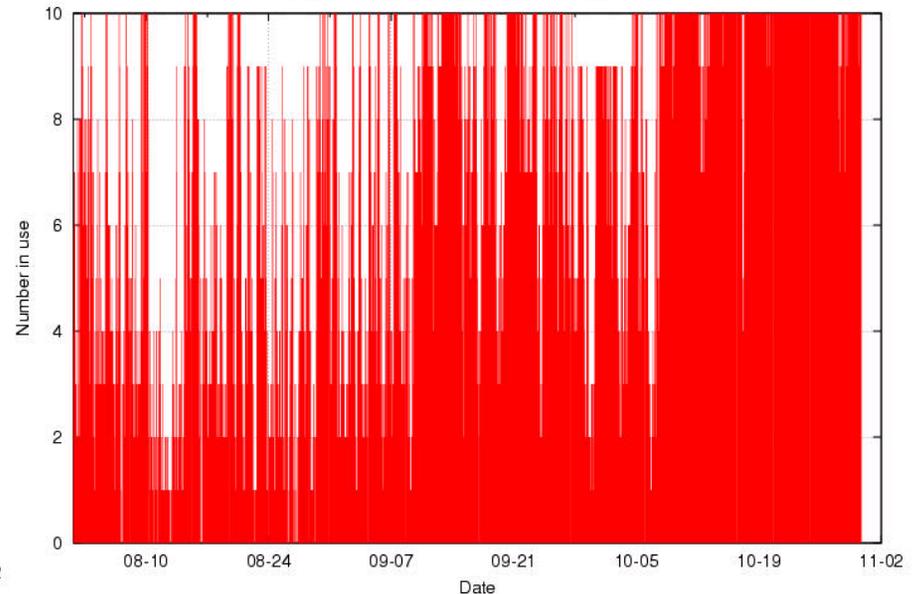
CDFEN Bytes/Day in October

Total Bytes Transferred Per Day (no null mvms) (Plotted: 2002-Oct-30 01:34:06)



CDFEN Drive Utilization August-October

STK 9940 Drive Utilization 2002-10-30:17:00:09





Tapes

R. Harris, D. Litvintsev, ISD



- Tapes for CDFEN are in short supply.
 - As of Nov. 5 we had 395 tapes and using ~30 tapes/day.
 - Requisition 159603 was cut in half from 1700 to 850 tapes.
 - Purchasing says PO went out Nov. 6 with delivery 3-5 days after that.
 - Requisition 160371 for 850 more tapes approved but no buyer as of Nov. 5
- History
 - September 22 noticed that rate of usage had tripled.
 - From about 160 per month to 500 per month, and only 320 tapes were left.
 - Due to Tevatron performance, reprocessing, larger event size (silicon).
 - Dmitri & ISD recycled 500 tapes and pushed through FY02 order for 500 more.
 - Difficult to recycle more tapes. Start cutting into datasets used for physics & PhD's.
- Future
 - Need to make T9940B drives in new robot available by early 2003.
 - CDFEN Robot can hold only 1500 more tapes.
 - Robot fills up by January 1 at current rate, but reprocessing should stop soon.
 - T9940B increases tape capacity from 60 GB to 200 GB.



Data Handling: dCache

R. Kennedy, ISD



- Operations
 - pretty stable running over the past few weeks
 - upgrading to incorporate new file servers: 7 TB -> 21 TB
 - This week make dCache default caching mechanism in development.
 - DIM remains default for older offline versions and can still be selected by users.
 - DIM currently has 11 TB, and I/O rate of 9 MB/s and a cache lifetime of two weeks.
- Development
 - write pools
 - file family affinity
 - helps accommodate different caching behaviors and refine static datasets implementation
 - SRM (Storage Resource Manager)
 - Will go into dCache interface to facilitate load balancing.
 - Also found useful to deal with firewalls aborting idle control connections.
 - An "srmcp" will be implemented for use by offsite users
 - **Rutgers DIM as well as SAM**



Data Handling Development

J. Tseng, D. Litvintsev, R. Herber, A. Kreymer, ISD



- SAM

- Sam Shifts

- Art Kreymer completed training and 1st week of shifts.
- Loaded Enstore CRC values into SAM tables for compatibility with newer software.

- DFC migration to SAM

- Finished redesign of SAM metadata necessary to preserve CDF DFC functionality.
- Consulting with the CDF luminosity group, D0 and ODS/DBA on planned changes.

- New SAM stations

- Karlsruhe Technical University is now up with a cluster.
- Working with the Karlsruhe Research Center to set up a joint cluster with D0
 - **Also has to play well with CMS and BaBar. Issues of accounts, security, etc.**

- GRID

- gridftp access to dCache tested at Rutgers (and later srmcp).

- No evidence of file corruption. Uses the gridftp door/server implementation from ISD.

- gridftp tested by Stefan Stonjek (Oxford) for SAM.

- Out of 40 files, no file corruptions.

- GRID Demo for SC2002

- F. Wurthwein & I. Terekhov plan to use JIM, SAM and CAF to analyze CDF data.
- Hope to jumpstart SAM on CAF at Fermilab, RAL, Toronto, Ruthers, Texas Tech.



Reconstruction Farms

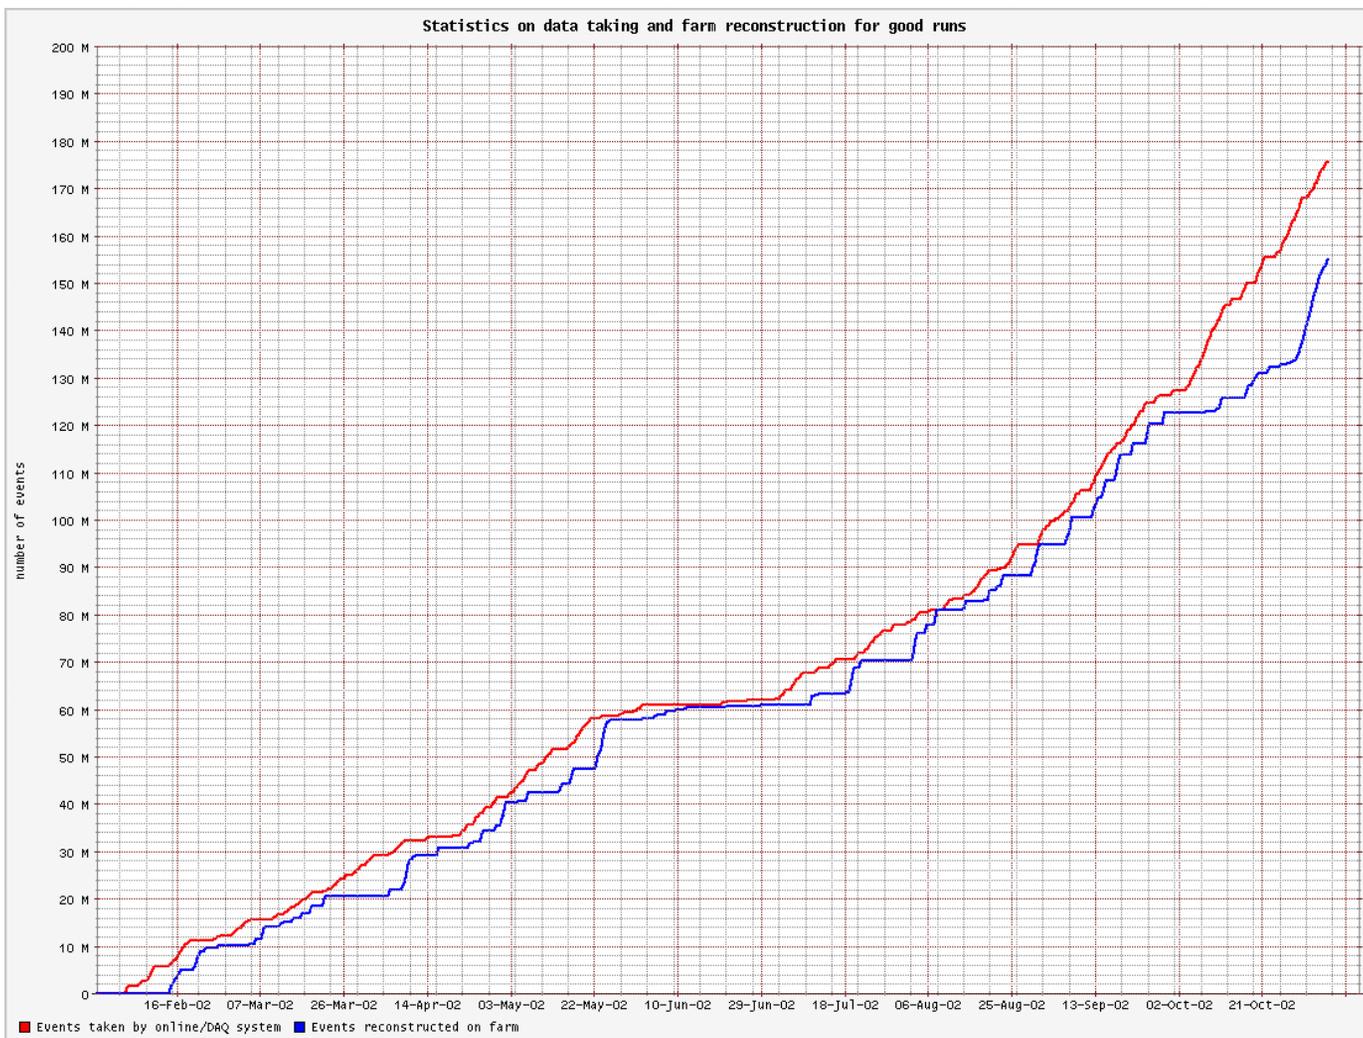
S. Wolbers, M. Siket, R. Lysak, OSS



- Processing: 155 out of 176 Million events complete for all 8 streams.
 - Slightly behind due to reprocessing effort and “staggered” operation.
 - Process recent data from 1st four streams while reprocessing old data from 2nd four streams, then switch.
 - Some streams have all 176 million events processed, but not all 8 streams.
- Reprocessing for Winter Conferences
 - All data needs to be reprocessed with 4.8.4a
 - Began reprocessing on Oct 11.
 - 130 out of 176 million events reprocessed as of Nov. 5
 - Average reprocessing rate was 6 million events per day.
 - Now doing 5.5 million events/day to save drives. 1.2 TB/day output.
- Hardware
 - Expecting 32 Athlon MP2000+ duals to arrive Nov. 12 (Thanks OSS!)
 - Should be available for use a month later.
 - Roughly twice the CPU of the 50 PIII 500MHz duals that we plan to turn off Nov. 15.
 - Additional disk and Gbit connectivity will come as part of upgrades to farm for
 - faster tapedrives, additional output datasets (35), and higher rate.



Reconstruction Farms Progress





Databases

D. Litvintsev, R. Herber, E. Wicklund



- DB Replica
 - Stays within 30% of its CPU capacity
 - Crashed last week, failover mechanism to production DB worked.
- Work on merging DFC into SAM.
 - Randy writing a paper on “predator” which copies DFC data to SAM.
- Dmitri chaired the database connection code committee
 - Made recommendations for improvement of operational efficiency and maintenance.
- Randy made a number changes to the Data Base Browser
 - Improve efficiency, error reporting and made a few reports more useful.



Offline Operations & Code Management

P. Murat, A. Kreymer, C. DeBaun, A. Beretvas



- Production Release for Winter Conferences
 - Production executable crashes once every 100K events. Acceptable for operations
 - Supported release 4.8.4
 - Trained collaborators M. Herndon & P. Savard to tag and build releases.
- Database
 - Deployed new Oracle configuration files to support the new DB replica server.
- Compilers
 - Started sending nightly gcc build error messages to developers.
 - Continued general support of development, integration, distribution, KAI licensing, etc.
- OS
 - Continued tests of patched ld to solve Fermi Linux 7.3 performance issues.
- Code Development
 - Code integration 4.9.0int1 built
 - 4.9.1 MC consolidation release built successfully, being tested.
 - 10K production test runs successfully.

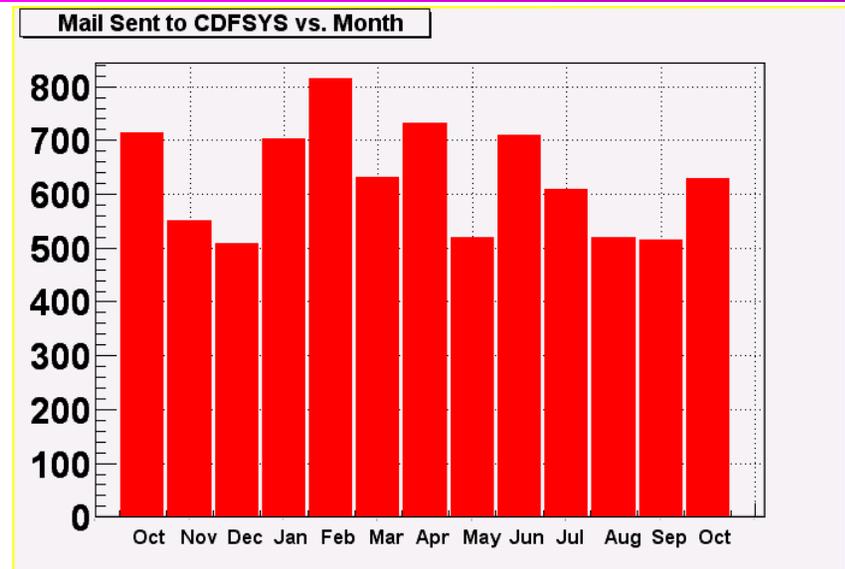


CDF Task Force

R. Colombo, G. Cooper, J. Harrington,
R. Jetton, M. Schmitz, L. Weems



- Mail to CDFSYS:
 - A monthly measure of questions & requests to CDF Task Force
- fcdfsgi2
 - Pretty stable over the last month.
 - One RAID controller failed. Replaced it and sent it off for repair.
- fcdfora1/2
 - No system outages for either node during the past month.
- fcdfsun1
 - Network logins were disabled on October 31.
 - Plans in place to reinstall and configure as the CDF SAM server.
 - Work to begin on or about November 8.





CDF Task Force: Central Systems



- FcdfInx1 (DB Replica)
 - System crashed on October 18 and it did not auto-reboot.
 - A new kernel was installed on October 30.
- CAF and related projects
 - Slowly getting Lance some help by training Glenn and Richard.
 - Upgraded memory on 18 worker nodes so that all have 2 GB.
 - 17 file servers received and installed. 16 are running burn-in tests.
 - Three new console servers installed, configured, now in service.
- CDF Offline Web Server
 - Specifying a new web server so that this functionality can be moved from cdfsga.
 - Due to our recent experiences with Dell duals, it is likely that a similar platform will be chosen.
- CDF Offline DB Replication Server, Stage 2
 - Have provided purchase order details to the appropriate parties.
 - This new replication DB server will have dual P4 CPUs, and about 2.8TB of raw storage.
- Monitoring for CAF
 - Investigating ganglia in collaboration with Hans Wenzel of CMS.
 - Would also be useful with the trailer desktops
 - Hope to integrate the monitoring daemon and local configuration into the cdf workgroup, so all machines are monitored by default..



CDF Task Force: Desktop Systems



- Administer 306 linux nodes and roughly 40 IRIX nodes.
- NCDF209, new server for CDF trailers operating
 - Nis services for linux nodes
 - Will takeover from NCDF09 in serving cdf code to desktops without local install.
 - Currently 16 clients on NCDF209 and growing
 - Being used as the fermi linux distribution server for the trailers
 - Keep changes from being made to the base installation.
 - Provide uniformity between caf stage ii and the desktops.
 - fcdcode1 will clone the distribution from ncdf209 to spread the nfs load...
 - Jason Harrington will manage workgroup changes on the main servers, ncdf209 and fcdcode1.
- Fermi Linux 7.3
 - 3 machines running fermi linux 7.3.1 and there have been no complaints yet.
- CDF Online
 - Jeff Schmidt (CDF/PPD) and Jason Harrington are preparing the CDF Online workgroup for Fermil Linux 7.3.1
 - Implementing a slash server for informal discussion/troubleshooting.