Run II Database Applications

Run II Computing Review

September 11, 2003

Lee Lueking
CEPA Department, DBS Group

Outline

- Players
  - CDF, DØ, DBS
- CDF Apps overview
- DØ Apps Overview
- Common projects
  - Monitoring, N-Tier architecture, SAM (defer to Wyatt)

Manpower (FY 03 → FY 04)

Underline means "left", **bold** means "joined"

- CDF (~3 FTE)
  - Dmitri Litnintsev (Head, former Deputy)
  - Alan Sill (former Head)
  - **Patar Maksimovic**(New Deputy)
  - Other Majors: Jack Cranshaw, **Matthew Marten**, Bill Badgett, Tom Wright, Donatella Torretta, Larry Kirsch, Andrew Hamilton, Art Kreymer, Randy Herber, Finnish Group.
- DØ (~3 → 1.5 FTE)
  - Taka Yasuda (Head)
  - Other Majors: Elizabeth Gallas, Robert Illingworth, Jeremy Simons, Herb Greenlee, and Sub-detector contributors
- CEPA/DBS (~4 excl. SAM)
  - Lee Lueking (Head, Monitoring)
  - Dennis Box (Mis. CDF)
  - Yuyi Guo (Monitoring)
  - Chih-Hao Huang (Enstore)
  - Carmenita Moore (SAM)
  - Margherita Vittone (Monitoring, Slow controls)
  - John Weigand (SAM)
  - Steve White (DAN,SAM,Mon)
  - Eric Wicklund (Mon)
- **CEPA/APS (Consulting and Design)**
  - Jim Kowalkowski
  - Mark Paterno

CDF: Apps

- Hardware
- Run
- Calibration
- Trigger
- Slow Controls
- SAM, DFC

CDF: Projects

Green means "done", Blue is "ongoing"

- CodeGen rewrite: Rewrite CodeGen to make it more maintainable and robust is completed.
- Connection code review: done, recommendations were implemented. The changed allowed to increase robustness and maintainability of the code. We no longer have dangling connection problems
- DbMetering: done, this code allows to protect Oracle server from overload (client queries the load in terms of connections and sleeps if the number of connections exceeds limits, useful for CAF operations)
- dbMonitoring: (common project) CDF will need certain functionalities expanded from time to time like the addition of new monitored events, query durations, name of the module that executes query, et cetera. Summary tables and history plots need to be completed.
- DbSchema Review : reviewed. TriggerDB API and agreed to get rid of cumbersome views. Hardware DB (HWDB) - huge tables are being split into many small ones (tables approach Oracle limits on # of columns).
- Freeware Database: (in Nelly's talk, see appendix B slide of this talk)

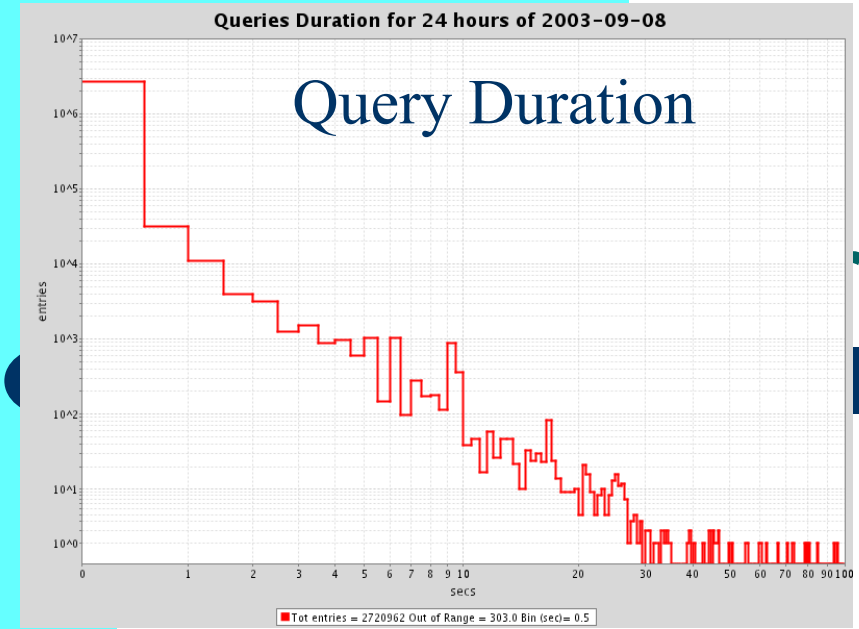
Queries Duration for 24 hours of 2003-09-08
Query Duration
entries
secs
Tot entries = 2720962 Out of Range = 303.0 Bin (sec)= 0.5

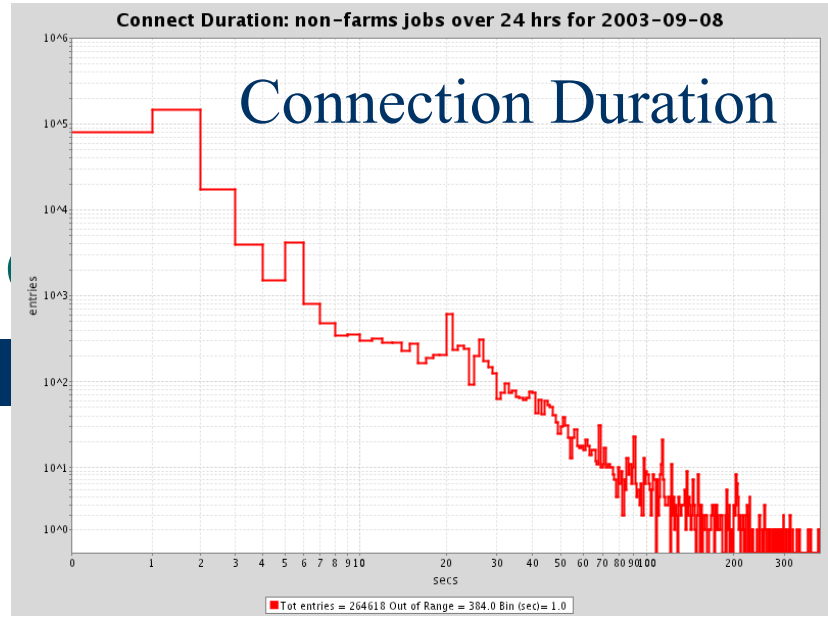
Connect Duration: non-farms jobs over 24 hrs for 2003-09-08
Connection Duration
entries
secs
Tot entries = 264618 Out of Range = 384.0 Bin (sec)= 1.0

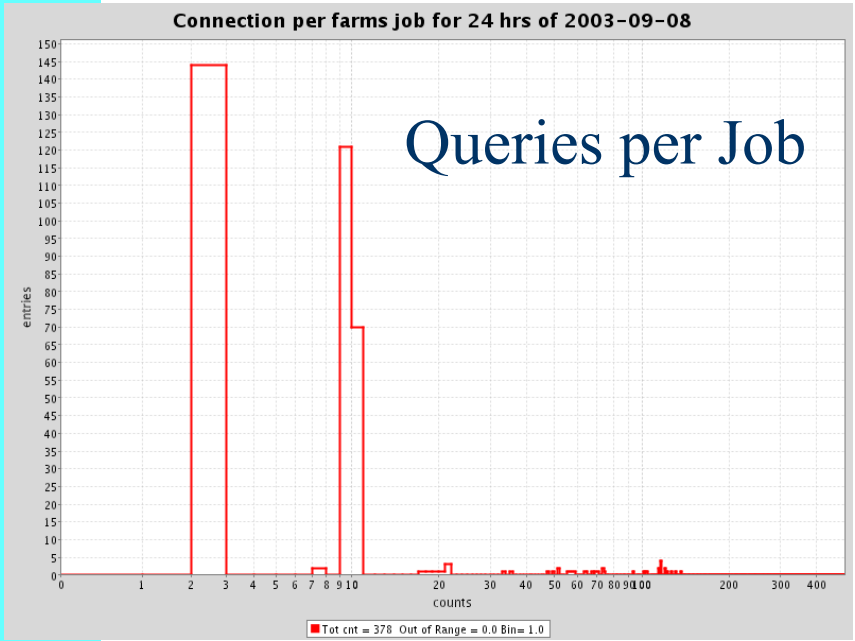
Connection per farms job for 24 hrs of 2003-09-08
Queries per Job
entries
counts
Tot cnt = 378 Out of Range = 0.0 Bin= 1.0

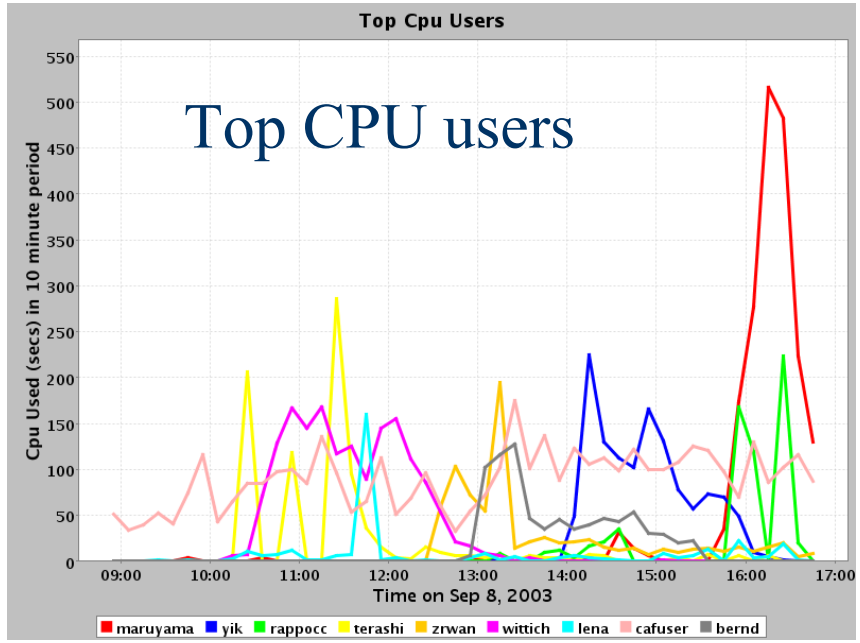
Top Cpu Users
Top CPU users
Cpu Used (secs) in 10 minute period
Time on Sep 8, 2003
maruyama
yik
rappocc
terashi
zrwan
wittich
lena
cafuser
bernd

DØ Applications:

- Trigger DB
- Luminosity (still in development)
- Calibration (8 sub-detector apps + 1 “top level”)
- Runs and Run quality
- SAM
- Miscellaneous: Speakers Bureau, Release

DØ: Projects

Green means "done", Blue is "ongoing"

- **DAN:** Development of the Middle Tier server for caching and database connection management (more next slide).
- **DAN Client Debugging**: Making all D0 client applications work with DAN, CPS is last but will be completed in Reco v16 release.
- **DAN Deployment**: Additional work needed to deploy DAN to remote sites.
- **Ongoing Trigger DB** maintenance and improvements.
- **Luminosity DB development**: Using flat files but this will be replaced with database in coming months.
- **Monitoring**: (common project) Working to include SAM DB server, and to satisfy additional monitoring needs, in addition to calibration DB apps. Ongoing improvements. Summary tables and history plots need to be completed.
- **Calibration DB Browser**: Need to develop web based browser to plot calibration data. Student worked on JSP application last summer, need upgraded tomcat server. Possible PHP options too.

DØ Applications: DAN (Database Access Network)

- CORBA interface to Client apps
- Memory (L1) and Disk (L2) caching
- Connection management to Database
- Server has common code base with SAM DB server

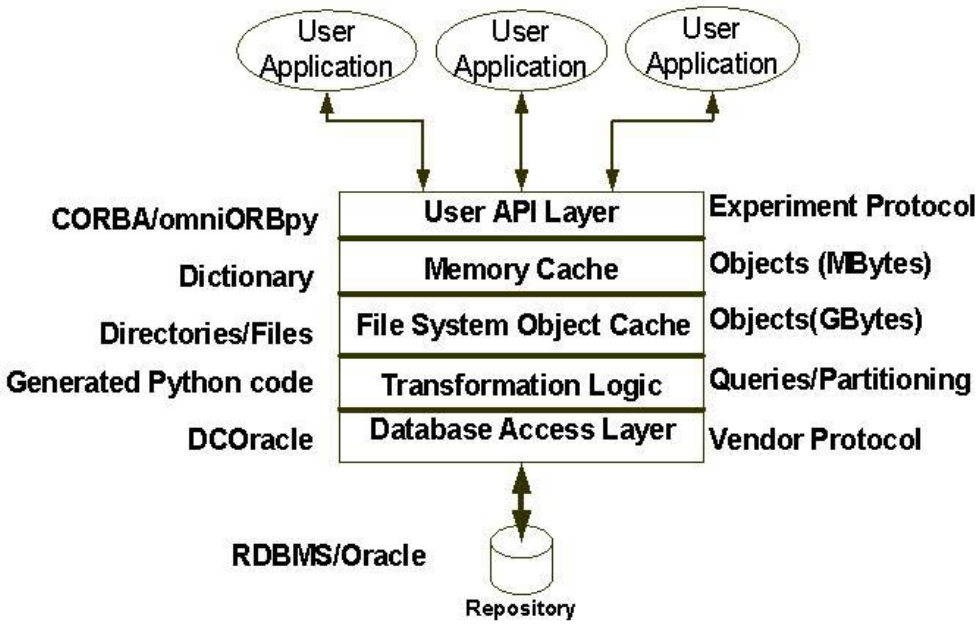

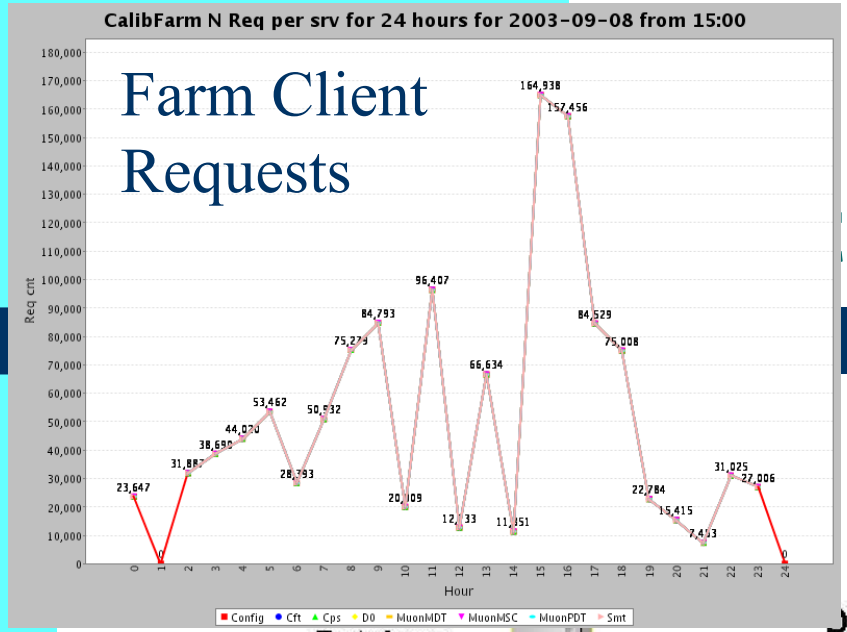

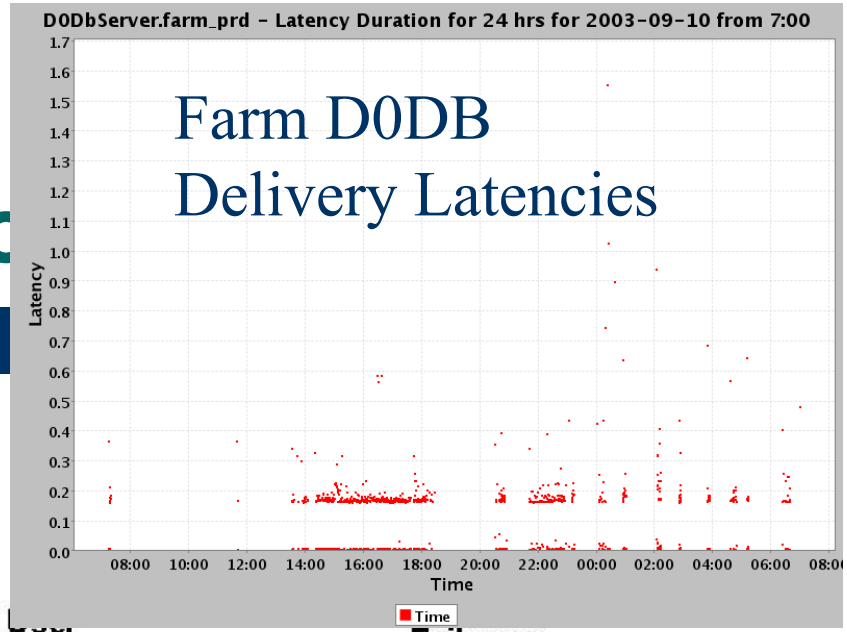

Database

Fail over

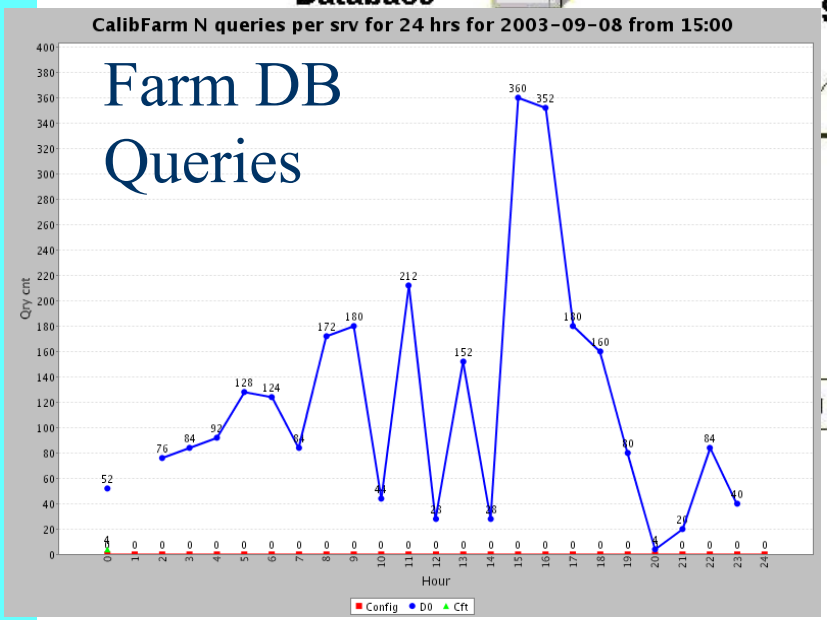

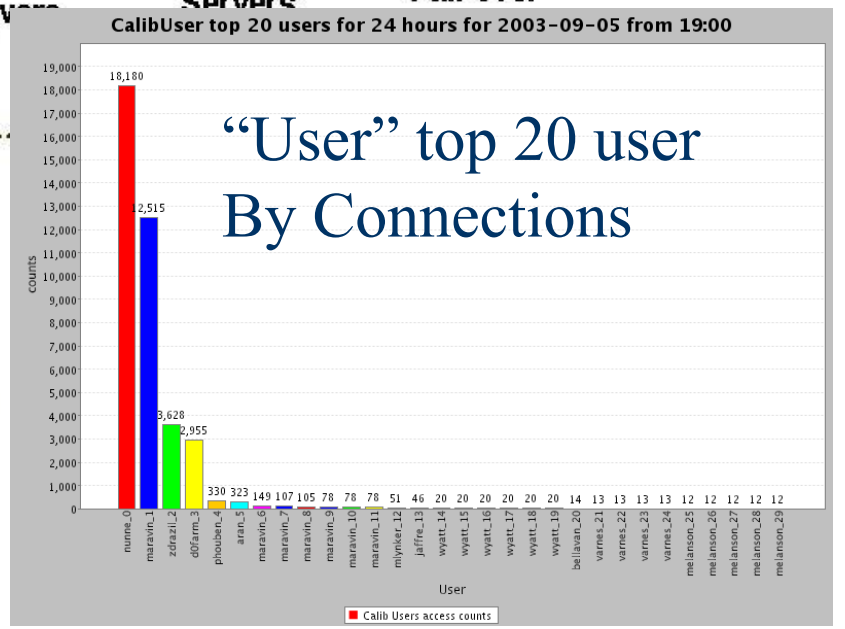

Common Applications: DBSMon

- Project Goal: Common tools for Application Monitoring
- Information Generation (InfoGen) is Exp. Specific.
- Collector/Parser
- Archiver using MySQL Repository
- Plotting tools using JavaFreeChart
- Histogramming part uses JAIDA
- Admin and automation scripts

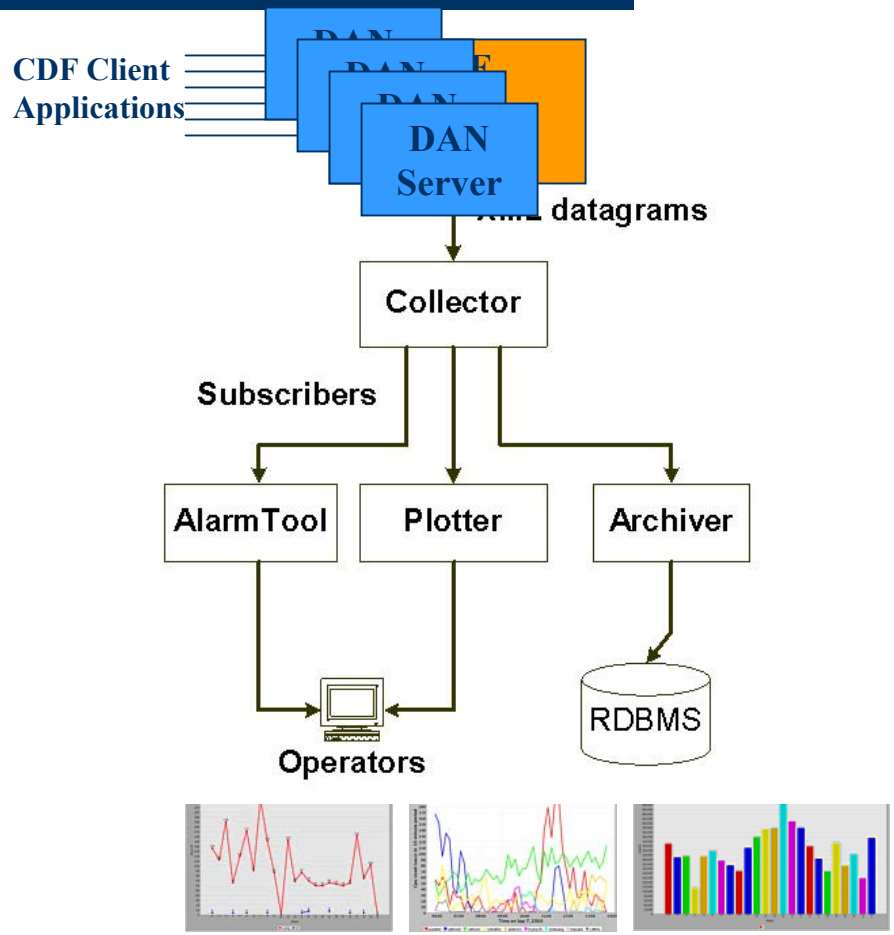

http://dbsmon.fnal.gov

Common Applications: DBSMon

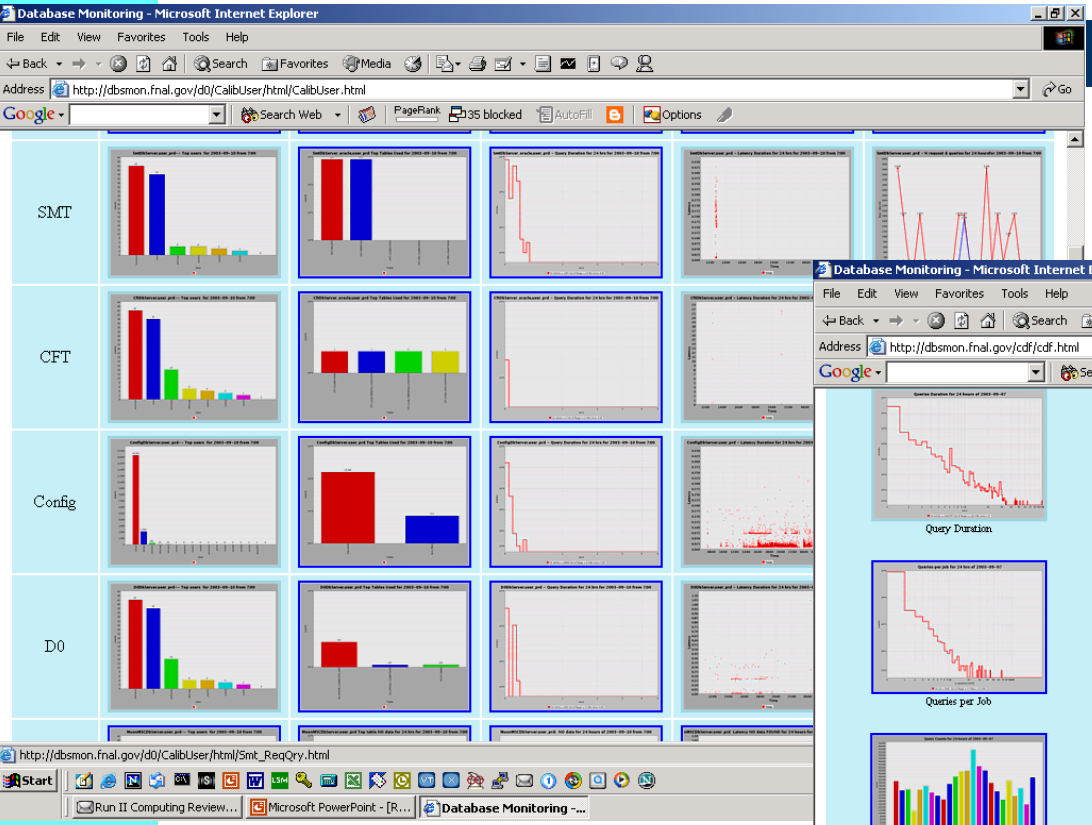

CDF

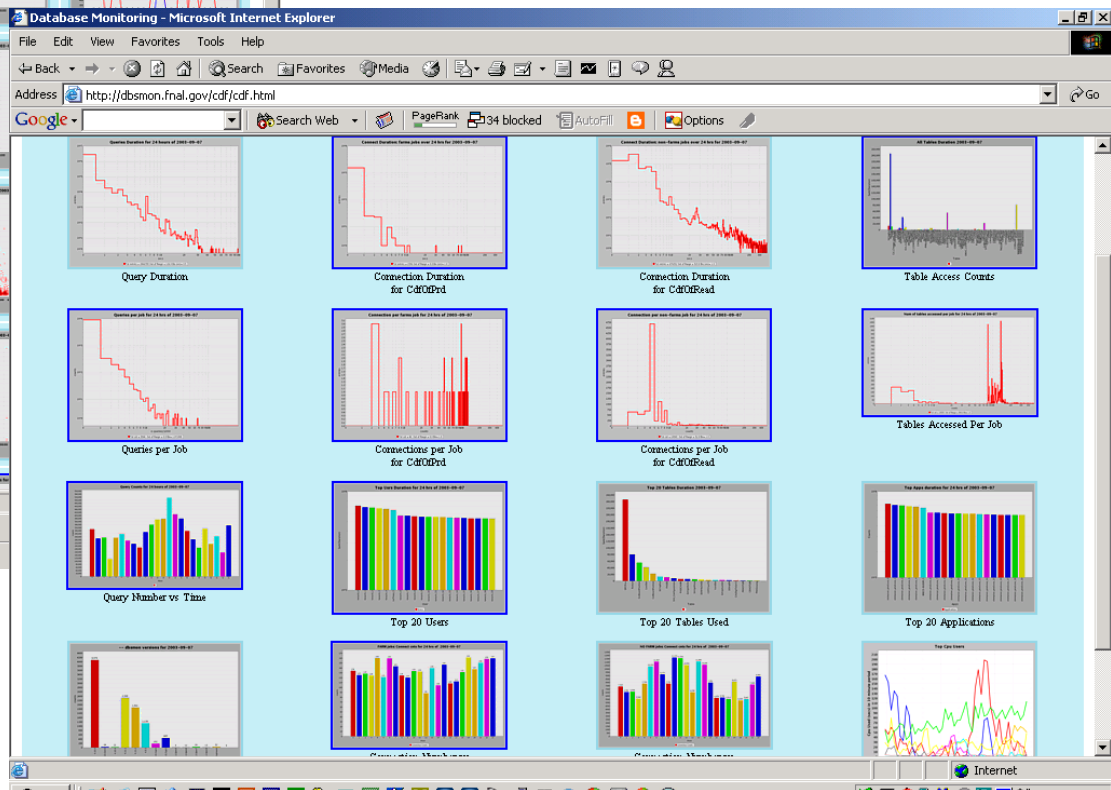

DØ

Common Applications: N-Tier Pilot

- DØ is successfully using a middle tier approach for DB access in DAN. BUT, it has not all been clear sailing:
  - The current implementation in Python has required the development of a “tandem” C++ caching server in some cases (SMT, CPS).
  - CORBA is a “heavy” solution for the functionality needed in the client interfaces.
- CDF is interested in using an N-Tier approach for read only database access. Their current replication scheme is working well, but the low administrative overhead and other advantages of a middle tier approach are attractive. (see more details of CDF view in appendix slide)
- We are exploring using more standard technologies, such as Apache, Tomcat, and HTTP to build a modular, high performance system. The first steps are to understand throughput and scalability via prototyping. Much of the DAN design and experience is directly applicable.

Summary

- Both Experiments have their database applications in full operation. Some of the code has been, or is being, revised to make it more robust and easier to maintain.
- CDF is successfully using replicated databases with clients connecting to the database directly. Connection metering is used.
- DØ's use of the DAN server is quite successful, full deployment to remote sites still in progress.
- Monitoring of Database Applications is now using tools common to both experiments.
- Serving DB information to clients through an N-Tier approach has many advantages. We are exploring a common project to develop an architecture similar to DAN with updated technology.

Common Applications: N-Tier Pilot

(Appendix A: CDF DB Group view of N-Tier)

- The scope of this project is being defined now. CDF finds n-tier approach especially attractive for remote institutions as it allows to fully utilize local CPU resources w/o administrative overhead and in a seamless fashion. Also, there are many other advantages such as 1) DB Connection management (Potential reduction in Oracle licensing), 2) decoupling of client code from DB, 3) decoupling of client code from DB schema, and 4) local caching.
- CDF finds this solution attractive for read-only DB access, characteristic for Calibration and calibration-like API access.
- The general idea is the following: CDF is successfully analyzing data and DB access at CDF is in good shape, especially for on-site work. To accommodate the remote access needs, w/o major disturbance to existing system we are taking two prong approach:
  - use available solution of replicating the data that is necessary for an analysis on a given dataset to the remote site.
  - in longer term we are looking forward toward developing (together with CD) a n-tier approach as it seems more convenient.

CDF: Projects

(Appendix B: CDF Freeware DB project)

- Freeware Database: initial work was done.
  - Were able to set up and populate on demand mySQL and PostSQL servers. The CDF code was proven to work against mySQL server.
  - Karlsruhe physicists were able to run on large statistics samples off Fermilab mySQL replica. Performance in terms of speed is adequate.
  - There were unresolved issues with data rounding off and memory leaks in DB access API. The CDF DB group will continue to pursue this, and ideally a remote institution will be able to set up a local mySQL server and populate it on demand from Fermilab Oracle server, and in the future from the 'main mySQL' server at Fermilab.