Data Preservation Activities at D0

Michael Diesburg, Fermilab
for the D0 Collaboration
DPHEP Meeting
Nov 20, 2012

Data Preservation at D0

I will cover data preservation activities at D0:

- Scope of preservation effort
- Investigations into what needs to be preserved
- Principles guiding preservation effort
- Assumptions about future environment
- Current preservation investigations
- Summary

Scope of D0 Data Preservation

- We divide preservation into two categories:
  - Gold Standard
    - Ability to return to data in N years and do a complete, publication quality analysis that can pass normal editorial review for publication.
  - Outreach Preservation:
    - Preserve some subset of data which would be useful for educational type activities, but would not be adequate for full publication quality analysis.
  - We are concerned here primarily with the first type.

- Time Scope of Preservation:
  - Actual time extent is arbitrary (No requirement for minimum period)
  - ~10 years seems reasonable (How long is our data still interesting?)
  - Actual cutoff date will be determined by practical considerations, not clearly defined need.

Preservation Investigations

- A Task Force was convened over the summer to examine preservation issues at D0.
- First part of charge to the group was to identify what should be preserved.
  - Easy answer: EVERYTHING
  - Really attempting to preserve complete intellectual infrastructure of D0
  - Collider data itself is a trivial and almost insignificant part of this
- Real job of the Task Force is not to choose what should be preserved, but to make sure nothing is forgotten
- Second charge to Task Force was to identify necessary tasks and make plans to accomplish the preservation.

What Do We Preserve?

- Collider data:
  - Exists in many forms. Dominated by raw, reconstructed, and skimmed.
  - Little to be gained by cherry picking what to keep (7.5PB vs 9.0PB).
  - Effort needed to identify and separate out what is unnecessary far exceeds effort and resources required to simply keep it.

- Computing Infrastructure:
  - Data exists in many forms: raw data, metadata, luminosity data, calibration data, data quality info, etc.
  - Need preserve all of this and consequently whatever utilities and systems are needed to access the data.
  - Must preserve the functionality of the entire offline computing infrastructure.

What Do We Preserve?

- Documentation:
  - Spread across many locations
    - Logbooks Email Talks
    - Internal Notes Agenda Servers Meeting Minutes
    - Must preserve access methods or convert to common format
  - Some institutional knowledge only exists as oral legend. This has to be captured.
- Management Infrastructure:
  - Current bylaws needed to be clarified/extended to function in an environment of sparse participation
  - Can we still implement the current analysis approval process?
    - Do we still have Spokesmen in 10 years?
    - How do we empanel editorial board in absence of other experts?
    - Whose names go on the papers?

Guiding Principles

- A few principles have guided our thoughts about how to achieve preservation goals:
  - Do no harm
  - Do not disrupt any currently active analysis effort
  - For those things which must change in the future, minimize the differences that users will see
  - Become as independent of experiment specific hardware installations as possible
  - Never step into territory where the security of the OS becomes an issue
  - Leverage common solutions wherever possible

Assumptions About Environment

- We can't expect to continue using exact same tools and systems for the 10+ years. Some things will have to change.
- We will make assumptions about longevity of some parts of the computing environment.
- Some things will remain the same either because they are stable and easily supported, or because they are impossible to get away from.
- Stable:
  - Code management: CVS
  - Environment tools: UPS/UPD (Fermilab product)
  - Build/Release tools: SRT (originally from BaBar)
- Minor interface changes, but improved fucntionality:
  - SAM data cataloging and file delivery tools

Assumptions About Environment

- Too much effort to abandon:
  - Data storage: ENSTORE (Works fine. No need to replace)
  - Databases: Oracle
    - Verification of any replacement would be difficult.

- In all the above there may be changes in implementation, but we should be able to protect users from any major interface changes.

- We also have to assume that supporting entities will continue to exist
  - RedHat, Oracle, etc
  - 10 years is a long time in the computing world

Current Preservation Activities

- We will maintain existing facilities, but at diminishing capacity until 2015.
  - Normal 3-4 year replacement cycle will become retirement cycle.
  - Should provide adequate resources that any current analysis can complete without disruption.
- Beyond 2015 we must be able to use alternate facilities.
- We are currently investigating on three areas relevant to 2015 cutover:
  - OS version. What will carry us for ~10 years?
  - Conversion to new version of SAM (platform independent data delivery)
  - Distribution of D0 runtime and development environment

Current Activities

- OS Version:
  - Which OS version we use can have significant implications for how preservation is approached. This is primarily due to security concerns
  - Easiest way to stay out of trouble is to simply ensure we are always running under a currently supported OS
  - Can we claim today that we can do this 8-10 years from now?
    - End of life for SLF 5 is 2017. Not far enough out.
    - End of life for SLF 6 is 2020. That's 8 years out and maybe quite as long as one would like, but it would be acceptable.
      - RedHat will extend support to 2023 for $$
      - Exact cost and some issues of distribution rights need investigation
    - There is always SLF 7
      - Won't be able to test until next year

Current Activities

- OS Version:
  - Most batch and interactive work today runs on SLF 5.
  - Library builds and some aspects of production processing are running under SLF 6.
    - Still need to do some further verification, but expect no problems.
    - Any problems that do arise will be corrected.
  - Movement to new OS versions has been relatively painless due to compatibility mechanism used in D0 environment.
    - Standard versions shared libraries and compilers are part of D0 environment.
    - Setting up D0 environment puts the standard libraries into users LD_LIBRARY_PATH so we are always using same runtime environment.
    - Change of OS version does not imply change of compiler version
  - Same mechanism should carry us to SLF 7

Current Activities

- Possible Conversion to new version of SAM
  - Current version of data handling system requires a stager process to be running on any node accessing data.
  - Not acceptable for Grid operation
  - New version delivers data via http. Only requires web access.
  - New version will also be used by most new FNAL experiments so long term support is guaranteed.
  - New and old version can co-exist so we can implement new version without disrupting any ongoing analysis.
    - Users can choose to use either at least until ~2015 when new system will become mandatory.

Current Activities

- Distribution of D0 runtime and development environment
  - Existing D0 interactive and batch facilities assume the presence of D0 code libraries and product libraries on the local node. Currently accomplished with NFS mounts
  - Not viable in a Grid environment. Need a better way to distribute this.
  - Investigating CVMFS as distribution tool
    - Initial tests look good. Load from D0 applications should not be an issue.
    - We expect this to be available on FermiGrid nodes in near future.
    - Long term support is less clear than other products mentioned earlier simply because FNAL is not primary source of support.
      - Use by most LHC experiments and some Intensity Frontier experiments at FNAL makes it a good bet.
    - Does not remove need for interactive login facilities, but should significantly decrease dependence on them.

Summary

- D0 has convened a preservation group to identify major elements of the experiment to preserve and to chart a course to accomplish that preservation out to ~10 years.
- We have some principles to guide how we approach preservation and we are narrowing our options down to a preferred path.
- Efforts are underway to investigate and implement newer computing infrastructure which will have long-term viability and help remove dependence on D0 specific hardware installations.
- Just getting started. There is much left to be done.