



Computing Sector

Six Month Progress Report

“Run II Data Preservation Project”

Version 1.0

8/30/2013

CD-DocDB #5186

PREPARED BY:

Robert D. Kennedy, Joseph B. Boyd, Kenneth Richard Herner, Bodhitha Jayatilaka, Stephen P. White

CONCURRENCES:

Rob Roser

Date

Project Sponsor

<Official copy is maintained electronically – printed copy may be obsolete>

Table of Contents

- 1. PROJECT OVERVIEW 1**
- 2. PROJECT STATUS..... 1**
 - 2.1. PROJECT SCOPE.....1
 - 2.2. PROJECT SCHEDULE.....1
 - 2.3. PROJECT RESOURCES.....2
 - 2.4. PROJECT COST2
- 3. PROGRESS ON 6 MONTH MILESTONES 3**
 - 3.1. DATA MIGRATION3
 - 3.2. SOFTWARE INTO CENTRAL REPOSITORIES4
 - 3.3. ESTABLISH VALIDATION PROCEDURE4
 - 3.4. DEVELOP A PLAN TO RESOLVE OPEN ISSUES.....5
- 4. PROGRESS SUMMARY 5**
- APPENDIX 1. OPEN ISSUES LIST 6**

1. Project Overview

The purpose of this project is to implement data preservation for the Tevatron Run II experiments. The Run II data sets are very unique data sample of 1 TeV protons colliding with 1 TeV anti-protons, and represent a considerable investment. It is unlikely another data sample like this will ever be collected again. This sample is in many ways complimentary to the LHC data sample of multi-TeV protons on multi-TeV protons. A number of important “legacy” Tevatron physics measurements have been identified. While the majority of existing analysis tools will be available and supported for five years after end of data taking, the datasets will have value in generating and checking physics results throughout another decade.

Data preservation means different things to different people. The approach for this project will be to:

- curate the data from CDF and D0 Run II (both simulated and actual) by continuously migrating the data to modern storage media,
- maintain the full infrastructure ability to generate new Monte Carlo samples, simulate them, reconstruct them, and process them, and
- maintain the ability to perform physics analysis on both the simulated and actual data.

With this, a complete physics analysis can be carried out on Run II data well into the future.

2. Project Status

2.1. Project Scope

The Project Charter v1.0 (CD-DocDB #5072, 5/7/2013) was approved by the Project Sponsor and defines the project scope. There have been no scope changes since that charter version. Note that the project scope for each Run II experiment varies according to the needs of each individual experiment, so some project scope items refer only to one experiment or the other.

2.2. Project Schedule

The Project Charter v1.0 (CD-DocDB #5072, 5/7/2013) defines the project time frame. The project has chosen to include an additional set of milestones at 18 months (ending June 30,2014) in order to press the early delivery of the Run 2 Data Preservation system, allowing sufficient time for testing, adjustment, and hand-off to operations of the system before the project closes out. The timeline that the project is working towards is a little more aggressive than the charter:

- Jan – Jun 2013: Evaluate current state and plan the R2DP system
 - July 2013: Assess 6 Month milestones
- July – Dec 2013: Implement the R2DP pilot system, address some open issues
 - First significant changes/deployments permitted by experiments in Aug/Sep 2013
 - Jan 2014: Assess 12 Month milestones
- Jan – Jun 2014: Gap-fit the pilot system, address all open issues, finalize R2DP system plan
 - July 2014: Assess Internal 18 Month milestones
- July – Dec 2014: Deploy the R2DP system for both experiments and close-out the project
 - Oct – Nov 2014: Assess 24 Month milestones and begin project close-out

Because of the distributed decision-making, broad scope, and under-defined deliverables of the project, we have chosen not to maintain a detailed MS Project schedule. Rather, we are applying a rolling wave approach to plan the goals for each 6-month period, and allowing each experiment team to implement those goals in the manner appropriate for that experiment, including the work in-scope accepted by that experiment. As the project proceeds, we will plan the upcoming 6-month period in more detail as stakeholders provide feedback on the evolving design of the Run II Data Preservation system and deliverables.

While we achieved the milestones for the current 6-month period (see Section 3), we have also been working with the experiments towards the goal for the next 6-month period, to deliver a pilot Run II Data Preservation system for early testing. Both experiments are building SLF6-compatible offline releases that the project will use as the foundation of the R2DP pilot system, along the timeline given in Figure 1. We are also staging work on Open Issues to begin after the summer 2013 conference season is over.

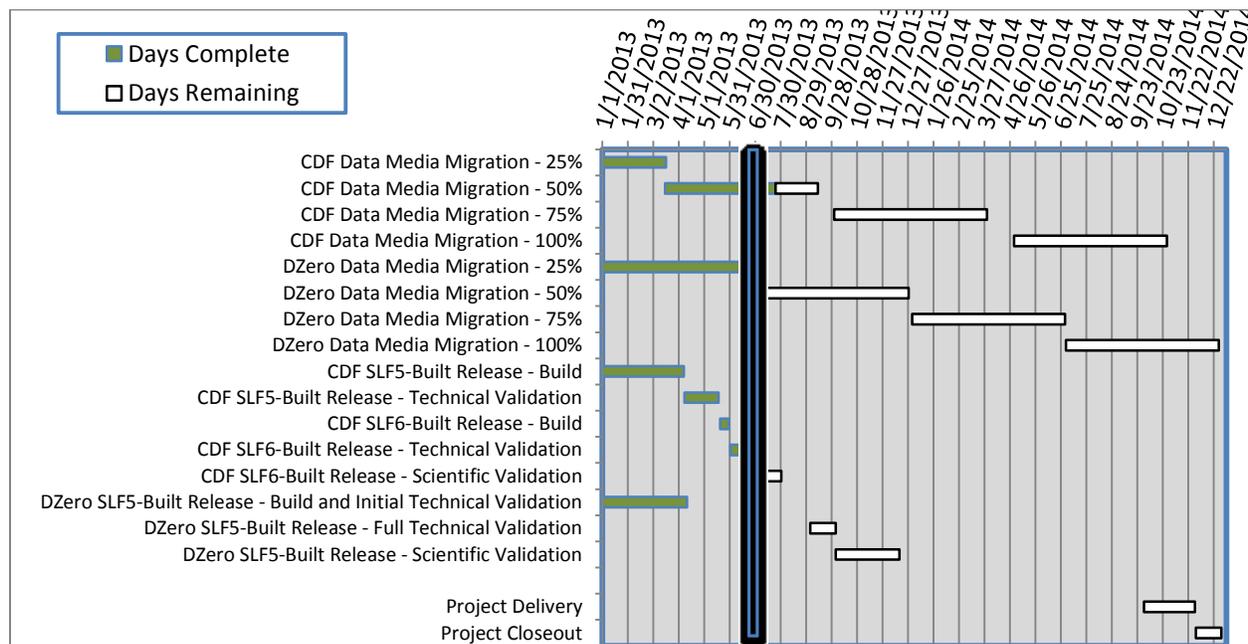


Figure 1: Early Timeline for Identified Run II Data Preservation Work towards the 12-Month Pilot

2.3. Project Resources

The Project Charter v1.0 (CD-DocDB #5072, 5/7/2013) defines the project resources and organization at a high level. There have been no major changes to the resources and organization since that charter version.

2.4. Project Cost

The Project Charter v1.0 (CD-DocDB #5072, 5/7/2013) defines the project budget. There have been no changes to the budget since that charter version.

The budget for FY13 is 3.965 FTEs in SWF and zero for M&S. As of July 1, 2.3643 FTEs of effort have been charged to the project. With 59% of the SWF budget has been spent after 75% of the fiscal year has passed, the project is mildly underspent. The cumulative project effort to date for FY13 is show in the chart below. This is due to our lower level of effort required at present to assess the requirements and needs of the Run 2 data analysis systems. This has led to proposals for modifications to the data analysis

systems which have either been declined or deferred until after the summer 2013 conference season by the Run 2 experiments. We expect to use a greater level of effort per month beginning in late August or early September as we begin execution of some planned changes as well as testing of the Run 2 Data Preservation pilot deliverables.

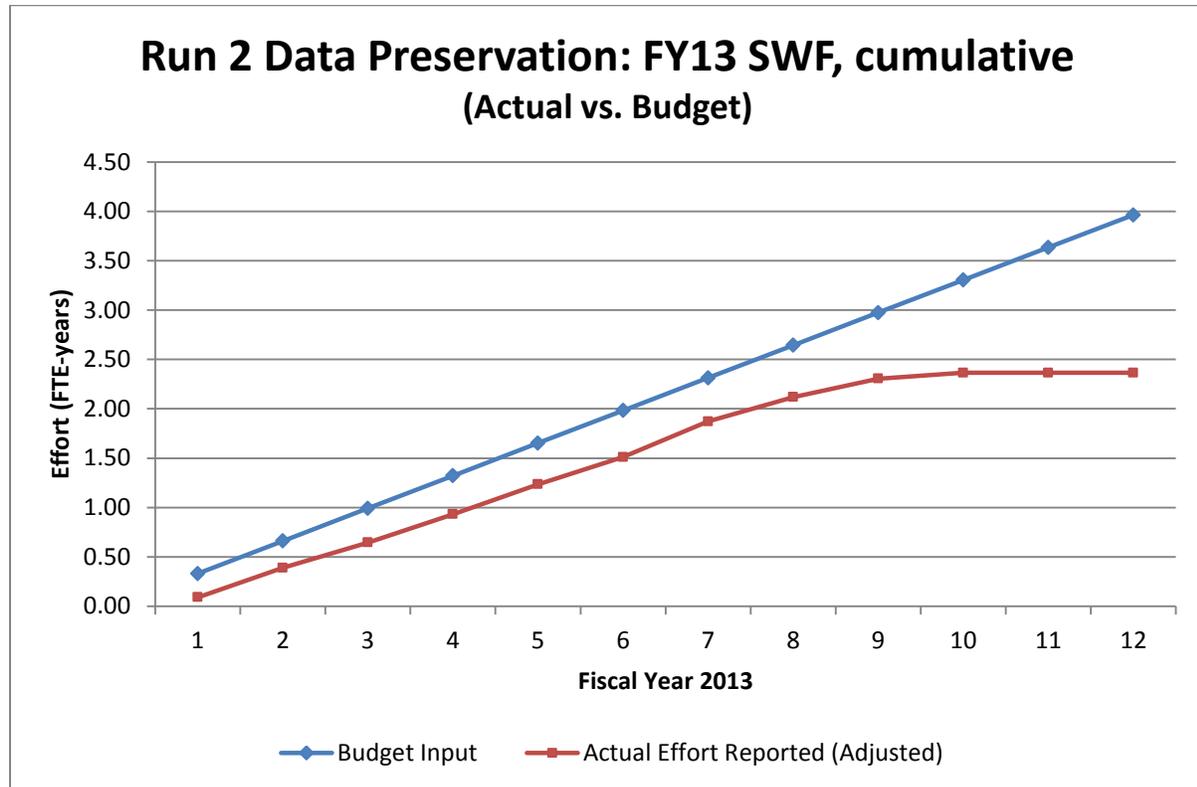


Figure 2: Cumulative Effort Reported to the Run II Data Preservation Project to Date in FY13

Note that months 10-12 (July-September 2013) in the figure above lack data at this time, as expected.

3. Progress on 6 Month Milestones

3.1. Data Migration

Goal: Complete 25% of data migration to modern T10K tape technology

Description: The project oversees this request to the Mass Storage Services (MSS) group, who will actually perform the actual work. The project and the MSS group work with the experiments to assess priorities in data migration, to insure that on-going data analysis takes priority over this data migration. In order to achieve the 25% data migration goal per 6 months of the project, each experiment data migration effort should sustain on average about 1% of the data migration per week.

Results: as of 7/1/2013

- CDF: 44.0% of data migrated
 - The CDF data migration began in January 2012, long before the project began. It has sustained an average data migration rate recently of 1.2% per week. If this rate continues, the CDF data migration will be completed in June 2014.
- DZero: 24.3% of data migrated

- The DZero data migration began in July 2012. It has sustained a rate of 0.68% migrated per week recently, but in June 2013 suffered a significant slowdown which is being investigated. In addition, we will discuss with the MSS group how this sustained rate can be increased to be closer to the goal rate of 1%. If necessary, we will work with the experiment to prioritize the data migration to ensure the relevant data at least is migrated before the end of the project.

Assessment: Goal achieved.

Next Steps: Project should continue to track the DZero data migration rate closely.

3.2. Software Into Central Repositories

Goal: Collect software in tape-backed central software repositories (eg. CVS), and work with physics groups to collect analysis-specific code.

Description: The project encourages the Run 2 experiments to review the status of the complete chain of analysis software in tape-backed central software repositories. While it has been a long-standing policy in both Run 2 experiments to reposit all software in official Offline releases, there have been varying degrees of repository use by physics groups and end-users. Ultimately, the experiments perform the work of repositing code or taking the risk of code loss.

Results:

- CDF Offline releases in CVS. Author lists in CVS. Project has emphasized to physics group conveners the ideal of having individual analyses in CVS, and there are some examples of this.
- DZero central code is in CVS. Several analyses are in CVS in various life stages. Not all end-user analyses are in CVS, but that is to be expected. Everything needed to run validation, except last specific piece, is in CVS.

Assessment: Goal achieved.

Next Steps: We should revisit this goal in later in the project as the example analyses are tested and the validation suite defined in the Run II Data Preservation pilot. These should all be in CVS as well.

3.3. Establish Validation Procedure

Goal: Establish validation procedure

Description: The project will work with the experiments to define a validation procedure for the Run II Data Preservation system. This will function not only as the “User Acceptance Test” for the project, but also be used by R2DP operations to validate the system remains unchanged by changes to the operating environment such as infrastructure hardware upgrades or software updates.

Results:

- CDF: The CDF Offline validation procedure uses 10k events in production, standard monte carlo generation. There is also a B ntuple-based validation. CDF is comfortable with the impact of the changes to the computing infrastructure on this validation procedure.
- DZero: DZero will use their established validation procedure for a new offline release. Need set of scripts and docs to teach a new person how to do it. Some pieces of doc do exist. DZero does

need to collect the existing validation documentation together, and interview or delegate to experienced people to packaging the procedure. Time frame for doing this can be up to six months from now, or until the experienced personnel have moved on.

Assessment: Goal achieved.

Next Steps: We should revisit this goal in later in the project as the validation suite is used to test the Run II Data Preservation pilot system. We wish to see the validation procedures sufficiently documented to allow new personnel to perform the validation for both experiments.

3.4. Develop a Plan to Resolve Open Issues

Goal: Develop a plan to resolve Open Issues

Description: The project will identify Open Issues and develop a plan to resolve them. Open Issues are high-level unresolved issues worthy of reporting to the Project Sponsor. Some Open Issues are in fact risk mitigation believed to be necessary to ensure that the Run II Data Preservation system is sustainable by the Fermilab Scientific Computing Division with minimal effort from the Run II experiments.

Results:

- Of the identified 11 high-level Open Issues:
 - 3 are CLOSED by virtue of the R2DP design or project scope definition
 - 8 are OPEN and have at least a brief plan to address each component of the issue.
- Appendix 1 contains the Open Issues list with brief plans for each.

Assessment: Goal achieved.

Next Steps: The project will execute the Open Issues plans after the summer 2013 conference season draws to a close to avoid disrupting the on-going analyses of the Run II experiments. For Open Issues that are determined to not be actionable, the project will develop risk plans and support agreements that all parties (CDF, DZero, SCD) can accept.

4. Progress Summary

Overall the progress in the project to date has been very good. Both experiments and SCD have responded to the project in a positive fashion, allowing a collaborative approach to achieve these goals.

While high-level plans are in place, the challenging work of the project is still ahead: to determine a proper balance of risk for all parties while enabling analysis of Run II data to continue beyond 2015.

Appendix 1. Open Issues List

We have grouped the Open Issues into two main categories, issues from the users' point of view and issues from the service providers' point, to reflect our goal of delivering a usable and sustainable system. Documentation references are to the R2DP Project site and require Fermilab SERVICES login.

To date, 3 of 11 Open Issues are considered closed.

How will data analysis be performed in the Run II Data Preservation system? – users' point of view

1. How are we going to perform bookkeeping?
 - a. Plan: Both Run II experiments will move from using SAM to using SAMWeb for a sustainable data bookkeeping solution.
 - b. CDF Documentation: [SAMWeb for CDF \(Illingworth\).pdf](#)
 - c. DZero Documentation: [SAMWeb for DZero.pdf](#)
 - d. Status: OPEN

2. How will a user be authorized to access data?
 - a. Answer: The R2DP system will continue with the existing experiment management plan for conceptual authorization even after the Run II experiments end.
 - i. This depends on some informal Run II collaboration existence that can define who is a member of the "experiment" and thus qualified to access the data.
 - ii. The technical means of authorization will remain the same as well.
 - b. Status: CLOSED

3. How will jobs be run? How will the user start his/her job?
 - a. Plan: Submit jobs using the Intensity Frontier (IF) job submission system.
 - i. DZero Plan: We plan to integrate existing DZero job submission scripts with the IF system to hide this change from the end-users.
 - ii. CDF Plan: We plan to use the same CAFsubmit. Nothing changes for end-users.
 - b. Plan: Jobs will run on a grid or cloud node. Details of this are under discussion.
 - c. Status: OPEN

4. Is our MC framework flexible enough to interface with new generators?
 - a. Answer: Yes, because the input format to the MC framework can always be flat ASCII text file.
 - b. Status: CLOSED

5. Can we use new data analysis programs as they become available?
 - a. Answer: This is up to the user to test and implement, not the project. Past experience with ROOT encourages us to believe that users will be able to use new versions of ROOT on existing data either directly or after applying a ROOT-supplied conversion tool.
Note:
 - i. New versions of ROOT will not be centrally managed or deployed by project.
 - ii. Experiment code will not be built and validated against new versions of ROOT by the project.
 - b. Status: CLOSED

6. How will users be supported for data analysis questions?
 - a. Answer: While not strictly in project scope, perceived success of system depends on this. The elements the project may pursue to encourage this outcome include the following:
 - i. Identify and address gaps in support process due to fewer active users and analyses. Experiment task. Project is a stakeholder, but not responsible for execution.
 - ii. Update user documentation known to be out-of-date. Experiment task. Project is a stakeholder, but not responsible for execution.
 - iii. User analysis support system – there should be reduced expectations for support response time, and greater reliance on documentation. This is an Experiment task. Project is a stakeholder, but not responsible for execution.
 - b. Plan: The Project will help identify areas where data analysis support may be weak to the experiment as we enable example analyses to run in the Run II Data Preservation system, to ensure the perceived quality of the project deliverable is not reduced by lack of analysis support.
 - c. Status: OPEN

How will the infrastructure be maintained for a Post-2015 system? – service provider's point of view

7. What needs to be upgraded to current versions to re-establish version currency?
 - a. Data migration to modern T10k media
 - i. Plan: track data migration. Work with the Mass Storage group and the experiments to address bottlenecks in the migration.
 - ii. Timeframe: This migration will continue until at least July 1, 2014.
 - iii. Migration Progress Tracking:
<https://sharepoint.fnal.gov/project/DataPreservation/SitePages/TapeMigration.aspx>
 - b. CDF: Offline Oracle RDBMS to 11
 - i. Plan: upgrade the Offline development RDBMS first and test operation. Then, after the summer 2013 conferences, upgrade the Offline production RDBMS.
 - ii. Timeframe: September 2013 to October 2013.
 - c. CDF: Frontier/Squid
 - i. Plan: upgrade to the most modern version of Frontier and its underlying toolset that are usable by CDF.
 - ii. Timeframe: September 2013
 - iii. CDF Documentation: [CDF Frontier.docx](#)
 - d. Overall Status: OPEN
8. What should be changed/upgraded to reduce the risk of obsolescence in existing infrastructure?
 - a. Convert to using new SAM server
 - i. Plan: Both Run II experiments will move from using SAM to using SAMWeb for a sustainable data bookkeeping solution. DZero has basically already completed this process, though SAMWeb is not yet the default.
 - ii. CDF Documentation: [SAMWeb for CDF \(Illingworth\).pdf](#)
 - iii. DZero Documentation: [SAMWeb for DZero.pdf](#)
 - iv. Status: Open

- b. DZero: Calib DB Servers
 - i. Plan: Deploy a virtual node with the new SLF6-based releases of all Calibration DB servers. This will replace the Calibration DB servers currently on d0dbsrv20. We plan to execute this move after the Summer 2013 conference season.
 - ii. Note: This only partially overlaps Issue (9a) since the service software is also being updated and ported to work on a newer operating system.
 - iii. Status: Open.
 - c. Migrate logbooks to supported tools
 - i. CDF Status: already using the supported tool.
 - ii. DZero Status: Migration has been converted to ECL.
 - 1. See <http://dbweb0.fnal.gov/ECL/dzero>
 - iii. Status: Closed
 - d. Migrate some auxiliary databases, database applications to supported tools
 - i. CDF Plan: People DB does not have a plan yet. Internal Notes going to Spires.
 - ii. DZero Plan: Whod0, Speakers Bureau will be dumped into a simple archive format. Internal Notes migrated to Spires.
 - iii. Status: Open
 - e. DZero: Transition from CORBA to HTML
 - i. Plan: Transition plan documented, but experiment chose not to proceed
 - ii. Documentation: DZero Documentation: [D0 Calibration DbServers](#), [D0 CORBA Emails](#), [CORBAtoHTTP Effort Summary](#)
 - iii. Agreement: We have met with the DB group and the D0 algorithms group and agreed that the html transition will occur if and when CORBA stops working. CORBA infrastructure will not be fixed. D0 understands that the CORBA-to-HTTP transition may take longer to complete if the work is done later rather than now owing to experts potentially moving on to other projects. D0 also understands that MC production will be unavailable for the entire CORBA-to-HTML transition period. D0 accepts the risks of not performing the transition now in order to save valuable manpower for other projects during the 2013-14 timeframe.
 - iv. Status: Closed
 - f. How long are certain production servers needed, before being retired?
 - i. Plan: We are collecting a comprehensive list of candidate production servers. Many have already been migrated as driven by the refresh cycle or risk issues.
 - ii. Documentation: [DPPlan.docx](#)
 - iii. Status: Open
 - g. Overall Status: OPEN
9. How/where will we adapt to enterprise architecture guidance to reduce operations costs?
- a. Move services on “physical nodes” to virtual nodes
 - i. Plan: These transitions are on-going, and will continue until all services for the Run II Data Preservation system that can be migrated to virtual nodes are.
 - ii. Status: Some work has already been done driven by refresh cycle or risk issues.
 - iii. Documentation: [DPPlan.docx](#)
 - iv. Note: This is closely tied to Issue (8f).
 - b. Move services on experiment-specific nodes/systems to shared nodes/systems where possible
 - i. Plan: List which services can be moved/shared and transition accordingly.

- ii. Status: Not yet done. Plans are being developed for larger such systems such as Mass Storage which may or may not be moved to shared systems/nodes.
 - c. Overall Status: OPEN

- 10. How is this system going to be technically sustainable from 2015 through 2020? Roadmap?
 - a. SLF6 versus SLF7
 - i. Plan: Focus on an all-SLF6 system that operates in a virtual environment. The host CPUs may run a different OS, but the analysis environment will be SLF6.
 - ii. Documentation: [D0 s/w beyond 2020.pdf](#), [CDF task force report](#), [DPPlan.docx](#)
 - b. Oracle RDBMS
 - i. Plan: Investigate the potential roadmap for Oracle. Develop a risk plan that takes into account the cost of upgrading versus not-upgrading especially given recent positive experience operating on unsupported Oracle RDBMS versions at CDF.
 - ii. Documentation: [DPPlan.docx](#)
 - c. DZero: SAM Cache and dCache (for local data caching)
 - i. Plan: Transition to using dCache for this role, a product with long-term support.
 - ii. Documentation: [Dzero Job Submission Talk](#)
 - d. Overall Status: OPEN

- 11. How is this system going to be financially sustainable from 2015 through 2020? Support plan?
 - a. Determine how to prove the R2DP system is operational after changes.
 - i. Plan: Identify and document official validation test procedures, standard analyses
 - ii. Plan: Determine if automating some or all validation testing is feasible.
 - b. CDF: Frontier/Squid Support
 - i. Plan: Develop a plan to migrate support responsibility from CDF to SCD.
 - ii. Status: Squid support has not yet been formally worked out with SCD.
 - c. How will costs for a possible future media migration be addressed?
 - i. Plan: engage SCD management on this question. This is heavily dependent on future CMS choices of media, equipment, and robotics.
 - d. Overall Status: OPEN