SAMGrid Status Report

Adam Lyon, 28 February 2006 GDM

See the last section for implications of the DØ Grid Goals.

1 Project Drivers, Scope and Milestones

These are taken from the FY06 SAMGrid Budget Document. Updates are based on the January 2006 report.

1.1 Continuing Activities

1.1.1 Continue Smooth Operations

Robert Illingworth and Steve White have discovered the cause to seemingly random DB server inaccessibility observed at CDF. CDF has complained that after a DB server was up for ~ 24 hours, they would start to see jobs failing due to connection errors to the DB server. Restarting the DB server would fix the errors, but they would recur in about a day. We tracked the problem down to the DB server running out of file descriptors (each process is allowed 1024) due to client connections held open. Client connections were not closing due to a bug in Ominorb (our Corba ORB). A new version of Omniorb fixed the problem. We now see about 15 connections open to the DB server at both CDF and DØ.

1.1.2 Complete full deployment of SAM DH at CDF

- SAM on the farm is now at v7.
- SAM on the farm is using a private CDF file storage mechanism instead of SAM FSS. We had a few bugs in FSS and CDF did not have a robust system to assure storage (FSS was never meant to be a 100% robust solution). We have fixed the bugs in FSS and would like it to be reinstated for the reconstruction farm.
- SAM is being integrated in their CAF restart (we are waiting on CDF testing).
- CDF wants to be able to transfer a file out of SAM (e.g. a Root file) to any node running GridFTP. We have produced a v7 version of "sam get dataset" that performs this function. Awaiting minor change to dCache configuration.

1.1.3 DØ MC & SAMGrid & REFIXING

In late December, a calibration bug was discovered in the latest DØ fixing pass. Approximately 1.5B events needed to be "refixed". DØ has just recently completed an intense project to refix that data.

The first refixing pass was to process critical skims needed for analyses going to the winter conferences. The reconstruction farm, analysis farm, a large farm at IN2P3, the CMS farm and some other offsite locations participated in that first pass.

The second pass involved refixing all of the data. A completion deadline was set for March 1. This project was essentially all of the I/O of the recent yearlong p17 reprocessing effort squeezed into six weeks. Thanks to changes to make SAM data handling and SAMGrid more efficient, the refixing was finished over a week early.

1.4 billion events in 81K files were processed in the second pass. SAMGrid was used to process nearly 50% of the data sample offsite. Our gateways to LCG and OSG allowed us to use DØ and non-DØ owned resources for refixing. LCG resources processed ~8% (110M events) of the data and OSG resources processed ~1% (~10M events). The startup for OSG sites was slow due to firewall problems on the machine hosting the SAM station at Oklahoma University. Now that these gateway mechanisms have been thoroughly exercised, DØ is in a good position to utilize LCG and OSG for MC production.

Andrew and Parag were the main SAM participants in the refixing, and they did an enormous amount of work in a very short time. Their efforts were directly responsible for the success of the refixing project. It should be noted that Andrew realized the potential to heavily exercise LCG/OSG for this project and he was one of the main proponents of using those resources.

1.1.4 LCG & OSG Integration

See above for use of LCG and OSG for the refixing. LCG operations is now handled by Joe Steele. Parag still does OSG SAMGrid operations, but will soon turn it over to Joel Snow.

1.1.5 SQLBuilder

Randolph has been working on CDF DH issues, SAM autodest and DB server issues as well as his CHEP paper this month.

1.1.6 DØ upgrade from v5 to v7

Much of this effort has been put on hold due to the DØ refixing effort and is now resuming. Andrew and Parag are now nearly 100% tasked to the SAMGrid v7 upgrade. Dehong has really completed coding the online v7 upgrade (I thought this was nearing completion last month but a review of what he had done showed that he was going down an undesirable path). He is now testing and preparing for a full online test at DØ. The MC request system is under test by the DØ MC group with the integration database.

1.1.7 Rewrite broken groups and quotas for SAM managed cache

The human resources for this project have been retasked (Igor M.).

1.1.8 Deploy new SAM Data Handling Monitoring

A test version of "SamHDTV" is operational and working correctly. The MIS and mainline station codes were merged and final testing will begin before a test deployment. -- Still testing. -- A single station test is complete. Multiple station testing and DB impact study are ongoing.

1.1.9 Testing

All of the DB Server unit tests have been wrapped in our SAM Test Harness, making testing new releases much easier. We are now in the process of adding additional functionality, including specific client tests, SAM station tests and autodestination tests. We also have plans to add tests to confirm that bug fixes made in the past still function.

1.2 Moving forward with new technology (new activities)

1.2.1 Integrate SAMGrid with v6/7 compatible Run Job

RunJob is still being worked on by other people.

1.2.2 Investigate deployment of SAM redundant information services

This project is Sinisa's Information Service system. No work has been performed in the past month.

1.2.3 Investigate deployment of SAM web services

MINOS has been testing Sinisa's SAM web services prototype. They provided feedback and some bugs have been fixed. The SAM team right now does not have the resources to pursue large scale production testing at the moment.

1.2.4 Investigate use of Enth for data base queries (continuation of SBIR project)

No work has been done. Awaiting Matt's report. Still no report.

1.3 Providing new capabilities

1.3.1 SAM DH and Condor Glide in

Much of this work will depend on an SRM interface. See below.

1.3.2 SAM Edge Service prototype

The Wisconsin student has working scripts to do the on the fly deployment, but we have not seen a demonstration. One new wrinkle that has appeared is the need for the station machine to be on the "edge" of the site's network (needs to have ports open for CORBA). A CORBA "tunneler" may need to be investigated in the future.

1.3.3 SAM usage of SRM capable storage elements

We have begun planning and design of the SRM and SAM DH interface. We had a very fruitful meeting with Timur and now understand the current and future capabilities of the dCache SRM interface. Andrew is ready to start design discussions in earnest as time allows.

1.3.4 Implementation of SRM interface around SAM managed cache

No work has been performed yet.

1.3.5 Investigate breakup of SAM data handling services

No serious work has been performed yet.

1.3.6 Investigate SAMGrid for Analysis

No work has been performed yet.

1.4 Outreach

Several papers, posters, and talks (10 total) were displayed at the CHEP 2006 conference. A common theme of many submissions was lessons learned from running SAM data handling and SAMGrid at running experiments. Maybe LHC experiments took notice.

- Poster "The SAM-Grid / LCG interoperability system: a bridge between two Grids" (Gabriele, et. al.) [url]. This paper describes the SAMGrid to LCG forwarding and gateway system. It also nicely details some lessons that were learned including the "black hole" problem.
- Poster "A case for application-aware grid services" (Gabriele, et. al.) [url]. Describes the problem faced at DØ where a generic grid job service was not efficient for all types of job applications. The LHC experiments will surely face the same problem.
- Poster "Interface between data handling system (SAM) and CDF experiment software". (Valeria, et. al.) [url] Describes the interface layer between SAM and the CDF framework.

- Poster "SAMGrid Web Services" (Sinisa) [url] Describes the prototype access to SAM information via SOAP web services.
- Poster "SAMGrid Peer-to-Peer Information Service" (Sinisa) [url] Describes the redundant SAM information services.
- Poster "Lightweight deployment of the SAM grid data handling system to new experiments" (Art) [url] Describes the MINOS installation of SAM Data Handling
- Poster "SQLBuilder, a metadata language to SQL translator, an overview of its input language and internal structure (Randolph) [url] Describes lessons learned with the old dimensions parser that lead to features of the in development SQLBuilder
- Talk "Automated recovery of data-intensive jobs in D0 and CDF using SAM" [url] (Andrew, et. al.) Describes the interaction with SAM and CAF at CDF and general tools for automatic recovery of failed jobs.
- Poster "Experiences with operating SamGrid at the GermanGrid centre "GridKa" for CDF" [url] Describes the SAM configuration at GridKa for CDF.
- Talk "DØ Data Reprocessing with SAM-Grid" (Gavin, et. al.) [url]

2 Effort

Fermilab CD effort is ~6.0 FTE (as of February 2006 effort reporting)

- 100%: Andrew, Parag, Steve Sherwood
- 50%: Randolph, Adam, Steve White, Robert, Krzysztof, Dehong
- 20%: Gabriele

Breakdown of effort is below. Note that time off (vacation, sick, holiday) is not included, so the total effort will not match the available effort.

Effort	FTE
Core Development	2.5
Deployment to Production	1.3
Operational Support	1.0
Project Management	0.5

Outreach	0.5
Total	5.8

3 Risks

The risks are unchanged from November.

Some of the previous risks (unreasonable expectations and feature creep) are somewhat under control as we are now bringing related requests to the GDM instead of handling ourselves. A lesson that I'm learning is to always insist on use cases and requirements before any further consideration is made on a request.

Some new risks...

- Human resources: While I think we are in ok shape now to handle the projects we've started, we do not have the resources to start other important projects (Web services deployment, breaking up SAM into services). We are also undertaking some short term rapid projects to get some operational problems out of the way (speed up DB server). I am hoping that completing the CAF restart, monitoring, and DØ migration will free us up a bit.
 - Losing more time of Steve White. Steve is the only remaining DB server expert. Randolph is catching up, but Steve has been very important for rapid DB server fixes and debugging.
- Upgrade to Oracle 10g from 9i. It could be great or a disaster. The speed of many SAM queries worsened when we switched from 8 to 9. We need a good testing program to find problems before 10g goes into production. MINOS has done some preliminary testing and found no problems but their database is a small fraction of the size compared to CDF and DØ.
- Grid politics I understand that protocols and interfaces may be changed (e.g. Condor using their own protocols) from what we use currently. SAMGrid will need to keep up.