

SAMGrid Status Report

Adam Lyon, 30 August 2005 GDM

1 Project Drivers, Scope and Milestones

1.1 *Inter-VO Grid Access to Fermilab Resources [COMMON]*

Parag, in conjunction with Tanya, Gabriele, Steve White, and Alan Jonckherre from DØ, is working to implement the VO integration with the SAM database. He is writing test cases for the administration interface (to be implemented within the existing SAM pyapi framework).

1.2 *Fermilab Computing and Strategic Storage Resources Accessible to the Open Science Grid [COMMON]*

DØ is using the Fermilab CMS cluster with its OSG infrastructure. Gabriele is investigating OSG services through gLite for resource selection.

1.3 *Minimize Operational and Deployment Loads [COMMON]*

See other sections for details on the SAMGrid DH and JS deployment projects.

1.4 *DØ Reprocessing [DØ]*

Gabriele, Andrew, and Parag have been supporting the DØ Reprocessing effort's use of SAMGrid for job and data management (under "*SAMGrid: Thick Job Manager Development*" and "*SAMGrid JS Deployment to Production [DØ]*" projects). Westgrid and in2p3 are still the main contributors of reprocessed events (over 200 million events processed at each facility). As of August 30, over 760 million events have been reprocessed (out of 1 billion desired); 690 million done remotely; 25 million on the OSG Fermilab CMS cluster.

The LCG interoperability project is going very well. Gabriele and Parag spent two weeks in Lyon, France to interact directly with LCG people. The trip was a success as DØ can now submit jobs to LCG via SAMGrid with a prototype job forwarding system. A report from the trip may be found at <http://www-d0.fnal.gov/computing/grid/doc/SAMGrid-LCG-integration-Lyon-report.pdf>. The DØ experiment has decided to use this resource for MC generation instead of reprocessing.

1.5 *Large Pick Events [DØ]*

As mentioned in previous reports, DØ would like to protect the SAM cache from large pick events projects. The current implementation of groups and

quotas does not work in the d0mino-less station configuration (DØ now has many cache disks on many nodes). Igor Mandrichenko (on loan to the SAMGrid team) has been working to implement changes to the SAM station to address the deficiencies. Up until recently, however, his time has been spent on a different non-SAM project.

1.6 Current Stable Operation

- Members of the SAMGrid team rotate a weekly expert shift. Issues are first filtered by the DØ SAM Shifter.
- The issue tracker has been in use for the past month at DØ. Overall, it works quite well but it is not without its annoyances..
- CDF is now using the issue tracker, and are much less tolerant of its annoyances than DØ.
- The "MIS" project is ongoing. Steve Sherwood is nearly complete with a final implementation. Sinisa is working on a prototype Information Service (meant to reduce SAM station reliance on the central database).

1.7 User Analysis using Grid Resources [DØ]

Resources cannot be devoted to this driver at this time due to the DØ Reprocessing. Though of course work and investigations for the Reprocessing may be applied to learning how to handle user analysis jobs submitted via the JS part of SAMGrid.

1.8 Production use of SAM for MC [CDF]

Gabriele, Valeria and Andrew have a system ready for testing, including merging as per CDF's merge specifications. They are awaiting CDF personnel to do the testing. The main risk here is again CDF's commitment to using SAMGrid for MC production (there apparently is competition from another system). Note that this task has been in this state for the past three months and we are still awaiting a decision from CDF to go or not go with this project.

1.9 Production Use of SAM on CAF & DCAF [CDF]

This driver is still one of the major sources of activity in the SAMGrid group.

- The v7 frozen client and CAF submission script conflicts have been solved (thanks to Elliot and the CAF team)
- SAM is still not in production for the CAF. See below for details.

In June, we deployed v7 of SAM db servers and client to the CAF at CDF. A few days later, a skimming group attempted ~ 20,000 simultaneous file declares, which caused the DB server to fail and created a long SAM service outage at CDF. CDF soon thereafter backed off from v7 and restored the old v6 (unfrozen) DB server and SAM client. Although the enormous number of

file declares would seem to be an abuse of the system, in fact this episode has pointed out several deficiencies in the SAM system and the deployment...

- The DB server poorly handles large numbers of connections and does not recover when they are attempted
- There is little documented large scale testing of SAM client and servers in a CDF production-like environment. We didn't know what would work, and we didn't know what it would take to break it
- There are no service limits attached to the deployment, so there is a "SAM must handle everything we can throw at it" expectation

In light of these realizations, the following actions were initiated...

- 1) The DB server connection framework was reviewed by Jim Kowalkowski and Sasha Moibenko. With their advice, a connection cleanup mechanism was implemented in the server and a smart retry system was written into the client. One other major problem was that old clients connecting to the v6/v7 DB server leave connections that are no longer in use. We will implement a client version detection mechanism in the DB server that can tell the DB server to either reject the old clients (forcing all remote sites to upgrade) or at least put a message in the log file so that problems can be traced more easily.
- 2) We have initiated a testing program at CDF with the CDF test CAF (many thanks to Doug, Krzysztof, and Elliot for making this facility available and for their assistance) in the integration SAM environment (slow machines). We are first testing basic SAM services (delivering files, declaring meta-data, setting locations). Then CDF will test experiment specific activities (farm scripts, data stripping). The test CAF can bring nearly 1000 virtual machines to simultaneously hammer the SAM systems. These tests are documented at <https://plone3.fnal.gov/SAMGrid/Wiki/CDFTesting>. We have already learned about station limitations and Andrew has found one easy fix buying an increase of 50% in speed. We are now focusing on the speed of file declares. Right now, however, the test CAF is out of service as machines are being moved around. We eagerly await its return.
- 3) Clearly the "SAM must handle everything" expectation cannot be realized. With the testing above, we are developing service limits that will apply to a deployment. For example, on September 15 (a CDF down day), CDF will deploy the v7 client/DB servers for CAF with a service limit of 100k delivered files per day. This service limit means that the SAM team is confident of smooth operations up to the service limit. If CDF exceeds the limit, then there will be slow downs in

deliveries. If CDF needs to exceed this limit for normal operations, then such a request needs to be made and development work will need to be performed. The 100k limit far exceeds CDF's current nominal DFC+SAM delivery of 50k files per day.

The testing program is still in somewhat in a state of flux and I am working on a formal document to detail the testing the SAM team will perform for a release.

I had stopped deployment of v7 until such documented testing could be performed. Given that we now have a much better picture of SAM limits, we are prepared to have CDF deploy v7 for the CAF for file read activities with a service limit of 100k files per day (it could be raised to 150k files per day if we can test a faster version of the station on the test CAF). We are now investigating service limits of file declares (necessary for stripping).

The testing is also in response to new extremely aggressive CDF requirements (some services levels many thousands of times greater than the load that DØ and CDF currently generate). We cannot respond to the new requirements until we know the current capabilities of SAM. See <https://plone3.fnal.gov/SAMGrid/Wiki/CDFRequirements> for the requirements and a preliminary response.

The Italian group had been having problems using SAM for their stripping efforts (Donatella Lucchesi was quoted in email as writing "Sam is a nightmare"). Andrew and I met with her and Armando Fella to investigate the problem. We discovered that they are using a non-official version of the DB server (one that was hack-fixed) and they were attempting to perform 600 simultaneous file declares. The file declare problem was in fact the only problem they were having, but it was holding up their entire production activity. Andrew and I sat with Armando to look at his code and offered some suggestions to reduce their load on the DB server. They seem to be running much better now, though we do not have hard numbers to make an improvement clear.

1.10 Production Use of SAM on Reconstruction Farm [CDF]

The farm is running fairly smoothly in "phase 2" with the v6 DB server and can process > 2.5 TB/day. Due to processing mistakes, they need to add a reprocessing activity that will bring the load to 4 TB/day. We have seen evidence of connection problems with their v6 DB server that are fixed in v7, but CDF wants to do full scale farm script testing with v7 before they will upgrade the farm.

1.11 Deployment of SAM [MINOS]

MINOS is using SAM for data delivery in a pre-production phase. Art has performed some service limit tests. There is anecdotal evidence that the MINOS configuration can deliver 500K files per day; far above the MINOS load.

1.12 V6/V7 Deployment at DØ

As of early July, DØ is running with v7 in production on the analysis farms. Reprocessing activities, however, remain at v5.

1.13 Enth Distributed Dimensions [SBIR II]

The SBIR has officially concluded. Matt is writing up the final report and is packaging the Enth dimension server so we can do testing here.

1.14 Packaging and Configuration

Now much activity as the v7 frozen client is working well. Efforts were diverted to shoring up the test harness in v7.

1.15 Request System

The request system has reached the point that it requires coordination from DØ to deploy. We must also await upgrades in MC runjob.

1.16 Other projects not started

Station-SRM integration (on hold - prototype completed); SQL builder; RCP to SAM parameters (CDF request); Process bookkeeping (CDF/DØ design request); JIM Job Brokering.

2 Main Project Milestones

- SAM on CDF CAF [initial deployment in September]
- SAM on CDF Reco Farm [running with v6, will need to test with v7]
- SAMGrid v6/v7 (frozen) on DØ systems [DONE]
- SAMGrid LCG/OSG interoperability [OSG/CMS cluster in use, LCG integration in proto-production state]
- SAMGrid for CDF MC (ready for testing by CDF) [awaiting decision by CDF]
- Packaging/configuration Freezing (deployed - in use at CDF and DØ)
- MC Request system (awaiting DØ coordination)
- MIS server used for SAMGrid specific monitoring with SAM HDTV (will enter proto-production in a month)

3 Effort

Fermilab CD effort is 9.5 FTE's as of August 2005

- 100%: Andrew, Gabriele (effective), Lauri, Sinisa, Parag, Krzysztof, Valeria, Steve Sherwood
- 50%: Randolph, Adam, Steve White, Robert, Igor Mandrichenko

Operations support

- Art, Dehong, Liz Buckley-Geer

As of mid-August, Lauri has departed the SAM team. As of 9/1, Sinisa will depart the SAM team. We have reassigned responsibilities, but we know that we will greatly miss their deep knowledge of the system. We will need to retain them as consultants for the short term future.

4 Risks

The main risks, as I see them, are

- The CDF requirements and unreasonable expectations (e.g. 20,000 simultaneous file declares). At least we are now starting to get a much better feel for the capabilities of SAM.
- Scope/Feature creep. The diversion of deterministic file delivery has been difficult to manage.
- For SAMGrid, the risk is delivery of an interoperable solution for OSG/LCG. The reliance on external middleware that has not been production-battle tested is a risk.
- CDF's use of SAMGrid for MC production or another system.
- The loss of Sinisa and Lauri means that the remaining members of the SAM team will spend much more time on operations than previously.