

# **SAMGrid Status Report (and homework assignment c)**

**Adam Lyon, 9 May 2006 GDM**

## **1 Project Drivers, Scope and Milestones**

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

### **1.1 Continuing Activities**

#### **1.1.1 Continue Smooth Operations**

No major show-stopping SAM bugs have been found in the past month.

#### **1.1.2 Complete full deployment of SAM DH at CDF**

- 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 the CAF restart (we are still waiting on CDF testing).

#### **1.1.3 DØ MC & SAMGrid**

Parag and Joel Snow are working on improved SAMGrid integration with the DØ Monte Carlo tools and framework.

#### **1.1.4 LCG & OSG Integration**

The LCG integration project is essentially finished (Andrew wrote a document describing the end of the project). OSG integration is ongoing. Gabriele will start working with Parag (and student?) on integrating the resource broker.

#### **1.1.5 SQLBuilder**

No progress – Randolph is working on DB server.

#### **1.1.6 DØ upgrade from v5 to v7**

Progress has been made, though slowly.

- The conversion of the DØ online system to SAM v7 is nearly complete. Two large-scale tests were successfully completed. There is a remaining problem that for some files the CRC information is not

propagated from enstore back into SAM. This is under investigation.

- The main JIM code has been upgraded from using SAM v5 to v7 (and it uses Python 2.4 too!). The remaining part is to integrate JIM with Runjob, the latter also requires some work to make it compatible with SAM v7 (that work is not part of SAMGrid, but we are dependent on it).
- The MC request system is all new in v7. Testing has been very slow. Firstly we are testing in the integration DB to ensure we do not corrupt production (even though the new MC request tables are different than the old ones). We discovered that the last export of production to integration was over two years ago, so many new requests were not in the system. Integration was since refreshed. We have also had confusion as to how to search for requests and the meaning of dates (the date the request was entered? processed? completed?). These are being resolved now. The DØ MC group is continuing with testing and updating their scripts to work with the new system.

### **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 first attempt to try MIS on clued0 failed (the MIS server did not receive MIS events from the clued0 station). This is under investigation and awaiting time from Steve Sherwood, who is busy with other items.

### **1.1.9 Testing**

While we have been testing DB server functionality for quite awhile with the test harness, we would like to expand the test harness for more SAM tests. We have a testing requirements document for Station tests, and are implementing these tests into the test harness now.

### **1.1.10 Studying efficiency and scalability**

DØ tracks the ratio CPU time to Wall clock time for its batch systems. For cabsrv1 and cabsrv2, we have seen a steady decline in this ratio since last October to about 0.2. We discovered a bottleneck in the SAM station where a file removal blocks other activities until the removal is complete. In the past, this operation was very fast, but now seems to take tens of seconds. We have made the file removals asynchronous in the station – this increased the CPU/Wall time ratio to 0.4. We are investigating why this number is still low.

First indications show that SAM file deliveries are running efficiently and are scaling. There is anecdotal evidence that disk contention (two CPUs share the same cache disk) and excessive swapping contribute to the low CPU usage. Under investigation.

## ***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. See 1.1.6 above.

### **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 "tunnel" may need to be investigated in the future.

### **1.3.3 SAM usage of SRM capable storage elements**

Andrew is in discussions with Donetella and others for design on a SAM interface to SRM.

### **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.

## 2 Effort

Sorry - didn't have time to extract the information from the Effort Reporting.

## 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. UPDATE – the upgrade has been done for CDF production. Extensive testing in integration showed that many queries sped up in 10g, but many also slowed down, some of them substantially (many factors). We determined that the queries that became significantly slower were only run occasionally and so are not a large effect on SAM. 10g effectively runs at the same speed as 9i.
- 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.

## 4 Homework assignment c

c) Make a table of stakeholder data caching current and future performance metrics.

Right now, SAM has little cache performance metrics. For a recent day, we can determine cache turnover and hit rate upon request by analyzing log files. The hit rate used to be tracked for central-analysis, but these plots were never made for cabsrv1 and cabsrv2.

MIS will make such plots much easier to produce. I would like to track turnover, hit rate, and file lifetime for all of our cache systems for DØ.

For CDF, we use dCache and so we would not do our own metrics.