

SAMGrid

Adam Lyon
(Fermilab/SAMGrid Project Manager)

Run II Computing Review
September 2005

Outline

- ◆ SAMGrid cliff notes
- ◆ New Management
- ◆ Status of Experiments
- ◆ Future Ideas

Basic SAMGrid Requirements

- ◆ Transfer enormous amounts of data needed for different activities (**scalable**)
 - ◆ ... sometimes over large distances and with non-fault tolerant hardware (**robust**)
 - ◆ Maintain knowledge of what we are doing and what we did (**monitoring and bookkeeping**)
 - ◆ Use our limited resources effectively both at home and away (**efficient**)
 - ◆ Don't want to know the details [where files sit, where jobs run] (**helpful**)
-
- Solution...
- ◆ A data handling and job management system
 - ◆ **SAM + JIM = SAMGrid**

SAMGrid “Cliff Notes”

- ◆ **SAMGrid manages file storage (replica catalogs)**
 - ❖ Data files are stored in tape systems at Fermilab and elsewhere. Files are cached around the world for fast access
- ◆ **SAMGrid manages file delivery**
 - ❖ Users at Fermilab and remote sites retrieve files out of file storage. SAMGrid can handle caching for efficiency
 - ❖ You don't care about file locations
- ◆ **SAMGrid manages file metadata cataloging**
 - ❖ SAMGrid DB holds metadata for each file. You don't need to know the file names to get data
- ◆ **SAMGrid manages analysis bookkeeping**
 - ❖ SAMGrid remembers what files you ran over, what files you processed successfully, what applications you ran, when you ran them and where
- ◆ **SAMGrid manages jobs**
 - ❖ Choose execution site, deliver job and its needed dependencies and data, store output

New Management

- ◆ I became project manager in March 2005
 - ❖ Define project goals, set tasks, schedules, track work, interact with customers and stakeholders, interact with line managers (more than just Run II DH), held accountable

- ◆ Krzysztof Genser (CD) is Run II Data Handling Group Leader
 - ❖ Manages the people (all of whom are on the SAMGrid team)
 - ❖ Beyond this role
 - is a member of the SAMGrid team
 - leads CDF SAM deployment for CDF

Management Changes

- ◆ My science is DØ, but my SAMGrid is CD
 - ❖ Aside from being a DØ sam shifter, I have no DH responsibilities at DØ (nor anywhere else)
- ◆ Institute a more formal relationship with customers/stakeholders
 - ❖ SAM liaisons
 - ❖ Improve effectiveness of communication; Get it in writing
- ◆ Improve cohesiveness, focus, planning, documentation (both internal and external), and process
 - ❖ Requests to team are made formally
 - ❖ Team meets to respond to requests, discuss accomplishments, design
- ◆ Slowly changing the status quo

Issue Tracking Prehistory

◆ Before:

- ❖ Issues typically sent by users to an experiment SAM admin mailing list
- ❖ SAM shifter elevates issues by sending to the SAM expert mailing list

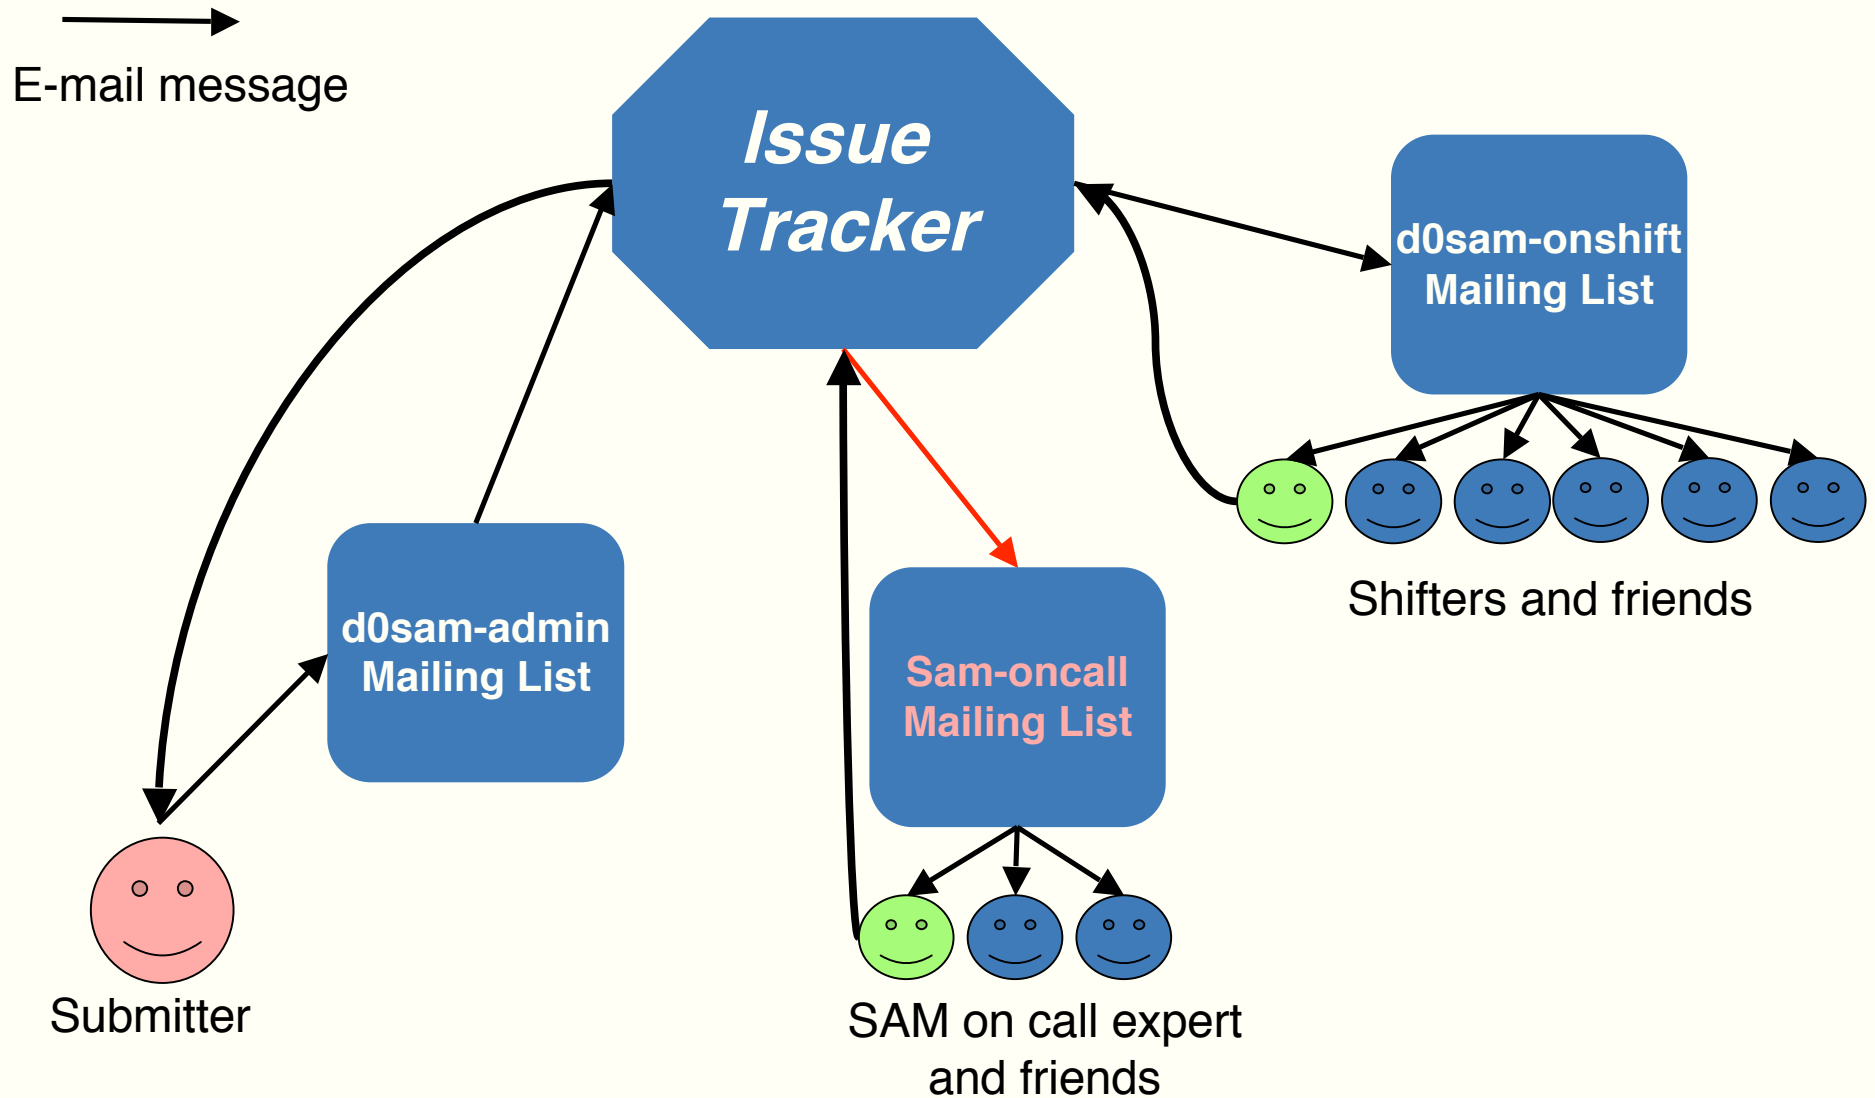
◆ Problems:

- ❖ Issues are buried in the mailing list archives
- ❖ Hard to determine disposition or even the full picture of an issue
- ❖ Original submitter is often out of the loop

Issue Tracking Requirements

- ◆ Want an issue tracker. Requirements:
 - ❖ Web based presentation of activity for an issue
 - ❖ Set state of issues
 - ❖ Assign issues to experts
 - ❖ Add to issues via e-mail
 - ❖ Search/count issues (e.g. by state, experiment)
 - ❖ **Automatic capture of issues sent to standard mailing lists**
 - ❖ Automatic alerts to mailing lists of issues
 - ❖ Original submitter sees all activity on issue
- ◆ Judy Nicholls (CD) was charged with finding a suitable program (I inherited her decision)
 - ❖ The *Plone* issue tracker fit the bill - supported/backed up by CD
 - ❖ The automatic e-mail rejects helpdesk “Remedy” system

Issue Tracker linked to e-mail



Issue Tracking Status

◆ Some annoyances

- ❖ Many separate e-mails and users may get multiple copies
- ❖ For e-mail reply, users must not quote original message, must not cc: mailing lists
- ❖ Formatting web input is not ideal
- ❖ Working with CD plone support

◆ DØ

- ❖ Transition began in April
- ❖ Issue Tracker is used for all issues and tied to standard sam “help” mailing list
- ❖ Modest early resistance quickly disappeared

◆ CDF

- ❖ Needed - issues were unknown, request timelines were confused
- ❖ Transition began in August
- ❖ Change and annoyances generated strong resistance
- ❖ In response CDF ties IT to new e-mail portal. Will only be activated by shifters/CDF experts
- ❖ Good: CDF manag. uses IT

◆ General

- ❖ All requests for SAM development go through IT
- ❖ Tied to sam-design list

DØ Status

- ◆ Originator of SAM & JIM
- ◆ DØ relies on SAM for all Data Handling
- ◆ Analysis Farm cache disks managed by SAM (> 130TB of “server cache”)
 - ❖ Increase of “server” cache and routing makes efficient use of tape system
 - ❖ Very smooth running
 - ❖ No manual prestaging
 - ❖ Problem: Groups and quotas
 - ❖ Need more monitoring
- ◆ 50 stations, 50B events consumed (> 2 PB) in past year
- ◆ Recent upgrade to v7
 - ❖ In conjunction with experiment management and schedule [change to status quo]
 - ❖ Successful transition over three months
 - ❖ JIM still uses old v5 for reprocessing
- ◆ v7 = frozen v6
 - ❖ Python and all libraries in one executable
 - ❖ Much faster over NFS
 - ❖ Easier to deploy
 - ❖ Less interference with experiment

DØ Status

- ◆ SAMGrid
 - ❖ Job submission and monitoring via standard Grid tools
 - ❖ Reprocessing effort (IB events, mostly offsite)
 - ❖ MC production (all offsite)
 - ❖ Well suited for production:
 - One use case
 - Known exes and dependencies
 - ❖ Analysis via SAMGrid is under investigation (varied exe & dep, varied use cases)
- ◆ SAMGrid interoperability
 - ❖ Forwarding to LCG
 - ❖ OSG interop
 - ❖ Use SAM for data handling
- ◆ SAMGrid installations often require expert help
 - ❖ Each cluster has different set of challenges
 - Different linux flavors
 - Firewalls
 - Sysadmins
 - ❖ Plans to make SAM less vulnerable to firewalls
 - ❖ Perhaps some automation is needed
- ◆ VO Management
 - ❖ Program to replace SAM's authentication/authorization mechanism with standard VO tools

MINOS Status

- ◆ Scale of data handling far less than DØ, CDF
- ◆ Some minor changes to DB schema needed (e.g. MINOS has two detectors!)
- ◆ Using SAM for physics analysis (dCache managed cache)
- ◆ All raw and reconstructed file metadata are declared to SAM database
- ◆ Reconstruction farm does not use SAM
 - ❖ Comfortable with current system

CDF Deployment Status

- ◆ Goal: Replace DFC with SAM; benefits
 - ❖ Rich metadata catalog
 - ❖ Bookkeeping
 - ❖ dCache throttling
 - ❖ Cache management (e.g. no manual population of cache)
- ◆ Online
 - ❖ Raw data logged into SAM & DFC
 - ❖ Path to SAM is not robust. DFC must populate SAM ~1 time/month
- ◆ Production farm (big success)
 - ❖ Running with SAM since July; new data going only into SAM catalog
 - ❖ Produces 2.5 TB/day (18M events). Meets CDF's needs
 - ❖ Running with older, less robust version of SAM code
 - ❖ Uses general dCache pool – competes with analysis farm
 - ❖ Efficiency issues

CDF Deployment Status

◆ MC Production

- ❖ Files produced offsite and metadata declared to DFC
- ❖ DFC metadata uploaded to SAM catalog

◆ User education (big success)

- ❖ CDF management has taken user education to heart
- ❖ Physics groups have SAM “power users”
- ❖ Lots of excellent documentation in CDF internal Wiki (incompatible with SAM’s Plone Wiki)
- ❖ Their Wiki also supports forums and issue discussions
- ❖ Most SAM developers cannot access that Wiki
- ❖ SAM Issue tracker used only by management

CDF Deployment Status - Analysis Farm

- ◆ Different use cases impact SAM differently
- ◆ Root tuple creation – substantial processing time per input file
 - ❖ Latest version of SAM software deployed for this use case. Service limit of 150,000 files per day at less than two file requests on average per second. CDF typically requests 50,000 files per day
- ◆ Data stripping (skimming) – 2 steps
 - ❖ Skim file creation: substantial processing time per input file; small output files declared to SAM durable storage
 - ❖ Merge small skim files to large output files for tape: negligible processing time; files declared to SAM for tape storage
 - ❖ The merge step can create many (perhaps 1000s) simultaneous file metadata declares. Service limit looks to be ~ 100 declares per minute. Still under study.

Reacting to past problems in depl.

- ◆ In late June, SAM deployed new versions of services
- ◆ Skimming group attempted > 10,000 simultaneous file metadata declares
- ◆ SAM service was not robust under high load. Exhausted computing resources. Melt down.
- ◆ CDF reverted to old version of SAM code
- ◆ Skim group reduced load to several hundred simultaneous declares – better, but hack was required to keep SAM running
- ◆ Outside task force examined SAM server code and made recommendations for improved reliability under high load. Most have been implemented
- ◆ How to we redeploy new version with confidence?

Lessons Learned (Change Status Quo)

- ◆ Only minimal testing of SAM code before release. Problems are found in production.
 - ❖ Since we controlled the deployment at DØ, we could easily fix versions before problem impacted too many people
 - ❖ Model cannot work at CDF where we do not control the deployment
 - ❖ We are developing a testing model and implementing some tests (see later slide)
- ◆ SAM services released without service limits
 - ❖ Expectation that SAM could operate in extreme load
 - ❖ Yes, SAM should handle extreme load gracefully instead of melting down
 - ❖ But experimenters should know boundaries of service that they should not exceed
 - ❖ Even we did not know the service limits

Lessons Learned (Change Status Quo)

- ◆ Requirements on SAM were unknown
 - ❖ CDF did release a DRAFT requirements document a few weeks after the problem
 - ❖ Requirements are extremely ambitious and far exceed any current load on SAM (rates are purported to be DC)

Quantity	CDF Requirement/s	DØ Mean/s	CDF is x times DØ mean	DØ max/s	CDF is x times DØ max
Project starts	1	0.0033	306	0.011	93
Process starts	1	0.029	34	0.09	11
File delivery	8	0.15	52	0.44	18
File storage	50	0.029	1748	0.098	513

- ❖ SAM is scalable – deploy more servers, but management gets difficult

SAM Testing Program

- ◆ Implement a mature testing program
 - ❖ Unit tests, Integration tests/Release verification, Production scale testing/Service level verification, Experiment specific tests
 - ❖ Tests and results must be documented!
 - ❖ Easy to write a plan - harder to implement! Started with the big tests and going backwards.
 - ❖ CDF has graciously given us use of their test CAF (~500 virtual machines) and people to help – an invaluable and crucial resource for the deployment's success!
- ◆ Using test CAF to determine service limits of SAM services and verify service
 - ❖ Survival under high load (slow down, don't die)
 - ❖ File deliveries, file metadata declares, location declares
 - ❖ We've already certified a release for a particular use case

Challenges to deployment

Why is the deployment still not done?

- ◆ Shifts in priority appear when we are near deployment
- ◆ Prefetching
 - ❖ SAM managed cache has prefetching built in
 - ❖ dCache SAM station would need a hack to do it – would rather talk to an SRM around dCache
- ◆ Deterministic file delivery
 - ❖ CAF job restart issue
- ◆ Suggestion: Develop a road map of functionality needed for initial deployment and later
- ◆ Suggestion: If new priorities appear, have a week cooling off period for evaluation
- ◆ Unknown thresholds for deployment
 - ❖ CDF will test latest version of SAM code on the farm. CDF is comparing new more robust version of SAM to less robust version. What is the criteria for success?
 - ❖ If we put out a version to fix a bug, what is the threshold for its deployment?
 - ❖ Suggestion: Road map to spell out requirements for deploying new versions
- ◆ Unknown requirements persist
 - ❖ Suggestion (both CDF & DØ): Revisit current and future requirements. Produce a timeline for required service load increases. Prioritize.

More challenges to deployment

- ◆ Many SAM functions have been wrapped by CDF code
 - ❖ Non-optimal use of SAM
 - ❖ Problems generated by these codes are difficult to debug
- ◆ Suggestion: CDF and SAM team work to eradicate wrapped SAM functions.
 - ❖ If SAM functionality is needed, we can put it into the SAM client
 - ❖ If CDF specific functionality is needed, then CDF specific code should produce output that can be fed into SAM commands
 - ❖ Remove explicit SQL queries
 - ❖ **PRIORITIZE**
- ◆ Outlook
 - ❖ **Enormous** real progress in the past five months. CDF is using SAM for production!
 - ❖ Jerry Guglielmo (special assignment) has been instrumental in keeping up the focus and pressure
 - ❖ The SAM team and CDF management are working better together
 - ❖ We've learned that the Status Quo has contributed to the inability to deploy. Change it instead of giving in!
- ◆ What is the plan to decommission the DFC?

Other activities

◆ SAMGrid at CDF

- ❖ We would like to see SAMGrid used for production (and perhaps analysis later) at CDF along with DØ. An installation of SAMGrid at CDF for MC production was made for evaluation. We have yet to receive useful feedback. We do not understand the disposition of SAMGrid at CDF, and if there is an experiment commitment.

◆ Improved monitoring of SAM

- ❖ Discussed last year - delayed due to personnel and priority changes. Implementation is active now, test deployment will occur shortly

◆ Improvements for CDF & DØ

- ❖ Redesign of MC request system for DØ
- ❖ Efficiency and usability improvements

Short term future

- ◆ Mitigate vulnerability to single point of failure database
 - ❖ New service that “caches” station and project state in the DB
 - ❖ Currently running projects can continue even if DB goes down
 - ❖ Have a basic working prototype
- ◆ Introduce web services
 - ❖ Prototype interface to SAM via web services (SOAP)
 - ❖ Much of user SAM client functionality is there (creating new datasets, querying metadata)
 - ❖ Files can also be delivered as SOAP attachments
 - ❖ Advantage: Clients no longer needs SAM specific code!
 - ❖ Needs to be tested for scalability – MINOS is interested
- ◆ SAMGrid - continue path to LCG and OSG interoperability and evolution to use common Grid services

Longer term future (DH)

- ◆ Direction from the division is to break monolithic SAM into distinct services for use by Grid systems
 - ❖ Data storage services: SRM interfaces
 - ❖ Metadata catalog services: SBIR
 - ❖ Bookkeeping services
 - ❖ Job management services
- ◆ Details of the future need to be worked out – maybe when the CDF deployment is complete

Manpower

- ◆ 9.5 FTEs on SAM team
 - ❖ 6 at 100% (including one guest scientist)
 - ❖ 7 at 50% (including one post-doc)
- ◆ Two main developers have left the project
 - ❖ I have arrangements with their new departments to allow for them to consult us on critical problems
 - ❖ Spreading their duties among the team. Reduce exclusive knowledge
- ◆ There has been no critical issue left unaddressed due to lack of SAM team manpower

Summary

- ◆ New management, new directions
- ◆ DØ and MINOS are running well
- ◆ CDF deployment is ongoing
 - ❖ We feel good about the recent important accomplishments
 - ❖ We feel disappointed that the deployment is still incomplete, and that we are just recently addressing the status quo
 - ❖ While the end is in sight, the path to get there is not crystal clear
 - ❖ There is strong commitment from all sides to see this deployment through
- ◆ Our monitoring will vastly improve soon
- ◆ SAMGrid evolving to use of common Grid services
- ◆ Our knowledge in data handling and job management is unique and valuable. We would like to share it with other systems and the Grid.