

“SAM for Users” Review

Reviewers: Mine Altunay, Ray Culbertson, Oliver Gutsche, Burt Holzman (chair), Michael Kirby, Burt Holzman, Jim Kowalkowski, Adam Lyon

Review material: <http://indico.fnal.gov/event/s4u>

Introduction

Scientific Computing Division senior management requested a review of the “SAM for Users” program. The charge:

sam4users consists of a set of utility scripts for the SAM data management package that allows analysis users to readily store their datasets and output files in SAM and the subsequently read them back as requested. Enabling of sam4users is one of strategies that SCD will use to reduce the need for NFS mounts on the batch system worker nodes as many of the analysis files are stored on Bluarc.

I am requesting a review of this capability in terms of operational readiness. In particular, we would like the committee to address the following questions:

- Is the underlying infrastructure adequate to support the needs of the experiments, particularly of the current heavy bluearc users NOvA, MINERvA, Minos and of the up and coming data taking experiments - microBoone, DUNE 35t, and Iariat.
- Does the architecture allow for easy expansion as demand grows? I.e., is it simply a matter of purchasing additional commodity hardware or would a redesign be necessary?
- Is the architecture consistent with the architecture for storage?
- Is the end user documentation adequate?
- Are we sufficiently staffed to both onboard users and support the additional operational load?

Overall observations

The set of tools that has been referred to as “SAM for Users” is better described as a set of tools that encapsulate common operation sequences that greatly simplify the use of SAM with regards to dataset management. These tools also handle errors that can occur during their use. It is likely that this “SAM for Users” release will cause changes to the current support structure. The release of these tools, as presented, does not necessitate deploying a new service, and therefore new SLAs are not required.

The current installation procedure places these new SAM tools into an existing package from FIFE called *fife_utils*. The definition, scope, and purpose of *fife_utils* was not made clear at the

SAM for Users Review - October 26, 2015

review session. What we do know is that it is a package from a larger structure, and that a release of SAM should not be dependent on some other package that inserts something into `fife_utils`. This set of tools should be installed into a place that is directly under the SAM project and released with SAM (including testing and other validation activities). If it is appropriate, `fife_utils` can cause the SAM dataset utilities to become visible and active.

Charge 1: Adequacy of the underlying infrastructure

We have observed the following:

- These tools ought to help with the issue of overloaded BlueArc¹,
- The tools do not directly address the issue of eliminating the need for BlueArc.
- The tools provide a way to more easily use the dataset management system, which ought to help reduce the need for BlueArc.
- There are use cases where BlueArc is used that are not addressed by these tools.

It is difficult to estimate how the current release of these dataset tools will change the habits of the experiments' use of BlueArc. Handling of data from interactive nodes is not directly covered by these tools. It is unclear how these kinds of use cases impact BlueArc use and how these new SAM tools will change that behavior. For the more common use cases involving data produced at worker nodes, it is more likely to have a direct impact because it is more straightforward (and perhaps permits more flexibility) to use the data handling system than write data directly into BlueArc.

Given that an iterative development methodology was employed to prepare the "SAM for Users" release, it is likely that one or more development cycles will be needed to better answer the question of impact and coverage. The current set of command line tools was determined by observation of how the system is being used, and backed up by a first-round of requirements gathering that was presented at a Computing Sector experimental liaison meeting. From the review, we expect the initial set to contain the most important set of tools. We also expect the list of grow or change over the next year.

We recommend that SCD proactively promote the use of this release, attempting to target users that rely heavily on BlueArc. We recommend to follow up on the question of adequate infrastructure (focused on heavy BlueArc usage) after this recommendation is carried out. An assessment of impact can perhaps then be made.

Charge 2: Expansion of the architecture with demand

The SAM for Users architecture should be able to expand with increased usage and stay ahead of the increasing demand expected from additional usage. There are several components upon

¹ BlueArc is an NFS-mounted filesystem tailored to perform large amounts of I/O operations but is a poor fit for large-scale distributed access to experiment data.

which the SAM for Users architecture relies that are potential bottlenecks or points of failure: SAM station DB, samweb HTTP server, and the underlying storage elements distributing the data files. The last element of this architecture is not within the domain of SAM for Users, but it should be recognized that increasing the ease with which data can be automatically accessed will increase the load on storage elements. There should be some effort to coordinate with the support team of each storage element that is registered with an experiment's SAM instance to monitor load on the SE and access point.

The SAM station for each experiment is recognized as a limited resource, but to mitigate that there is the possibility of SAM station duplication. As well, the fact that each experiment has a separate SAM station allows for scalability and isolation of any loading effects. The SAM for Users team says that duplication is possible for any SAM station but not without considerable effort. The increased utilization is not expected to be drastic but should be well measured as experiments are on-boarded. With this understanding, there should be time to react and implement new servers without creating a bottleneck. It should be noted that the increased load should not create downtimes or outages due to the samweb HTTP interface between the user and the SAM station. No direct connection to the database is generated, and overall with this proxy interface in place, the team is not concerned about database stability or write performance. The database is run on a centrally maintained postgres server by CCD.

The samweb HTTP server has not yet shown problems with connections, with the extreme example being Mu2e not yet able to produce back pressure with more than 1 million writes during their current campaign. There is currently a bottleneck in the HTTP server which is the HTTP-to-CORBA bridge, but the SAM team is actively testing the removal of CORBA from the architecture. They plan to deploy and test native HTTP support on SAM stations in the next month which should remove this difficulty. The samweb servers are run on FEF-managed VMs and scaling up this aspect of the architecture should be possible. There is little chance of outages or downtime due to the nature of the HTTP server, and monitoring should provide ample metrics for when scaling up should be needed.

Overall, the SAM system has shown stability under heavy usage, there are non-trivial but possible methods for scaling up, and modifications in testing to make the system more robust under heavy load. The scalability of SAM for Users does not seem to be a challenge going forward.

Charge 3: Consistency of the “SAM for Users” architecture with the Storage Access Architecture document

Sam4users is a toolkit and script set to use the SAM and IFDH data management tools efficiently and make them accessible for the average user. In that, it fulfills one key guideline of the storage access architecture that all batch storage interaction should best go through supported data management systems. It allows users to register files in datasets from all categories of storage. Sam4users does not prevent users from registering files on BlueArc in SAM and using these files through the data management system from batch worker nodes; this specific use case is contrary to one of the goals of the SCD storage access architecture (“removing POSIX access from worker nodes” and “removing mounted file systems on the worker nodes”). Because it is made very easy for the users to register and also move around datasets between storage categories (from BlueArc to non-tape-backed MSS for example), it might be advisable to design safeguards into the SAM for Users architectures that actively prevent users from accessing SAM datasets from BlueArc storage locations.

Charge 4: Sufficiency of end-user documentation

During the review, proponents highlighted the existing documentation for SAM:

<https://cdcvs.fnal.gov/redmine/projects/sam/wiki>

This SAM documentation has not been modified to include a SAM for Users overview nor specific commands. The project has provided man pages for all commands. The man pages should come from the product, so will be available at remote sites. The commands themselves have help messages which are the synopsis and usage parts of the man page. The source code is easy to find and read. In an informal and incomplete survey of SCD liaisons to the targeted experiments, we found none which had put up internal documentation. Proponents stated that an important part of the rollout would be through SCD experiment liaisons, experiment production groups, and individual power users. All FIFE team members should be counted as power users and able to train users, and respond to tickets. There was some discussion of possible training sessions, perhaps in experiment-based meetings, but no specific plan was proposed.

Users of the project will come with a variety of experiences. Advanced users of SCD data handling products will immediately recognize the purposes and methods of the commands, as well as the likely failure modes and shortcomings. In fact, advanced users are likely to have personal scripts which accomplish the same results. For these users, the existing documentation, primarily the man pages and source code, is more than adequate. Documentation by man page is not typical in SCD products. SAM for Users, samweb and IFDH commands all respond to the “-h” switch, so this important facet is coordinated. On the other end of the user scale is the naive user who is not sure what SAM is, and simply follows recipes

when requested. For this user, the documentation is adequate to run commands and get through the recipe. In general, this user should work towards becoming an advanced user, and therefore should also have some basic overview information they can easily find and understand.

We find

- documentation is adequate to proceed with release to all experiments
- dedicated training sessions are not needed given the current limited understanding of the users
- the documentation is not completely integrated into SCD DH documentation

We recommend

- error messages be made more meaningful. In particular, the “oops” string should be replaced with messages concerning the relevant action or state
- proponents consider adding examples to man pages
- an overview of SAM for Users be added to “Scientific Data Management” section of the wiki <https://cdcvs.fnal.gov/redmine/projects>. In particular, the relationship to IFDH and sam, the setup procedure, some examples of intended usage, any overall guidelines (such as how many commands can be run in parallel, when large datasets should be broken into smaller steps, how fast commands run), how to (not) use on the grid, how users can find summaries of their datasets created by SAM for Users, what the validate command does, and how to get help
- an SCD wiki page for data handling overview, pointing to the component (elements - bluearc,dCaches and tools - SAM, IFDH, SAM for Users) documentation
- SCD should explicitly implement training for FIFE team members and experimental liaisons, as needed, to make sure team members can act as power users
- SCD monitor feedback from the rollout and follow up on providing further presentations or providing slides, additional web pages or other materials to experiments as needed

Charge 5: Sufficiency of existing on-program effort to on-board users and handle additional operational load

It is implied by the charge that we should be reviewing the effort profile specifically with respect to operations and support. As the SAM for Users tools are mainly wrappers and scripts on the client side, the relevant program to inspect is the SAM service itself. The current effort load is 1.15 FTE (“best effort”, spread among 6 people who serve shifts, with 8x5 coverage). The SAM service is designed to be fault-tolerant to reduce the need for 24x7 coverage. On average, the shifters respond to one incident a day, and one service request per week. The turnaround time on tickets has a wide distribution - minutes/hours for trivial incidents (web server needs to rebooted) to weeks/months for solving performance issues that may require feedback from data transfer and database experts.

MINERvA, NOvA, mu2e, holometer are currently using SAM; LArIAT, DUNE 35ton prototype, Seaquest, g-2, and MicroBooNE (offline) are planned to be on-boarded. For SAM, this activity has historically proceeded via experiment's production group. There are no clear plans (or process for generating plans) for on-boarding nor any timescale for completion. The onboarding plan for the SAM for Users toolset is also not clearly defined. They would like to "find one person that everyone seems to follow - the same way [Andrew] got NOvA to use SAM to begin with - the hot-shot postdoc [...] and convert him, and then he's the one that everyone comes to and gives them the magic incantation." We recommend that if this strategy is followed, that the team develop a set of training materials to hand to these power users - the materials could be developed during the process of teaching the FIFE team about Sam for Users.

The SAM for Users developers propose that the support/operations load will be reduced with the on-boarding of the SAM for Users tools - since the tools present production-quality interfaces to the users for common tasks that many users script themselves, it should decrease questions related to SAM primitives and common tasks. (Given that the tools represent best practices, it should not significantly increase the development load.) For tools that are well-documented and for user base that is well-trained, this is true in the long-term. In the short and medium terms, our shared experience is that the support load increases - effort for engagement and outreach and support for early adopters must ramp up while the support for the rest of the community slowly ramps down.

We recommend that the program create a plan for on-boarding communities over the next 6-12 months with well-defined metrics, required effort, and a clearly defined schedule. This is particularly critical as a significant fraction of the current effort will likely be reprogrammed to work on a collaborative project with the Simons Foundation. It is clear even in absence of these plans that the current operations/support effort is insufficient. However, SCD should have the capability to slow the pace of on-boarding to better match available effort.

Summary and Conclusions

The new tools contained within this "SAM4users" release of SAM provide an excellent start at making SAM dataset management visible and usable. Many of the concerns contained within the charge questions are addressed positively by this SAM release. It needs to be recognized that this release does not exist to eliminate the need for BlueArc, although through time and training, will likely help to alleviate the inclination for people to use it.

The underlying infrastructure for SAM for Users appears to be adequate and capable of scaling to demand. We recommend that that SCD proactively promote Sam for Users, attempting to target users that rely heavily on BlueArc. The SAM for Users toolset is compatible with the current SCD Storage Access Architecture; however, it may be advisable to add additional constraints into the system to disallow accessing experimental data from BlueArc.

The end-user documentation is adequate for the current state of the toolset, despite not being fully integrated into SCD documentation. We recommend that current documentation (and error messages) be more fully fleshed out; that a wiki and additional documentation be added to Redmine to link together this with other data handling tools; FIFE team members and liaisons be trained as “power users”; and that the rollout is monitored and feedback used to iterate on documentation and training materials.

There is no clearly defined plan for moving forward with on-boarding, support, and operations. We expect the amount of SCD effort on-project to increase during this activity. We recommend that a clear plan be developed for the next 6-12 months including with measurable goals, scheduling, and resource (effort) constraints.