

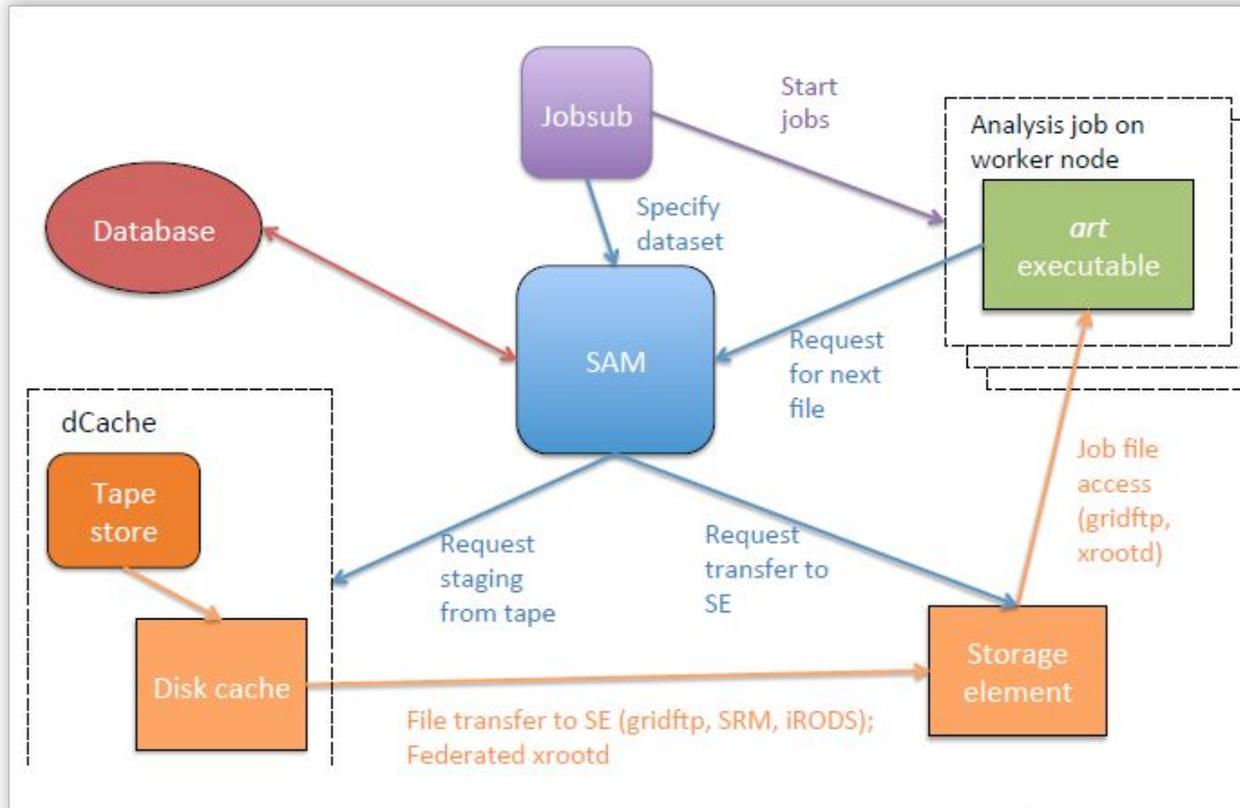
SAM data handling

Behind the scenes

Cristina Schlesier
UIUC

March 30, 2016

SAM (Sequential Access via Metadata)



Why use SAM?

Produce a large volume of data, both in number of files and total file size.

Impossible to maintain the paradigm where all files live in central NFS storage, and jobs are run on the local grid using explicit file lists.

It is necessary to move to using both local and off-site grids for MC production and to multiple storage modes (central NFS storage, large cache disks, and tape storage).

SAM Concepts

SAM is organized around

- Datasets
- Snapshots
- Projects
- Stations

SAM Datasets

Users often want to analyze a group of files that share characteristics.

This collection of files is called a dataset. It is convenient to give datasets symbolic and descriptive names that can be used by others when performing similar analysis.

Example: The user wants to analyze the set of all raw files which meet the following criteria:

1. Are from May 1 to June 1 of 2016
2. Produced by a particular gas gun
3. Are from runs with more than 10,000 events

Dataset name: MayJun16_BeamDynGG_greater10k

SAM Snapshots

Once a dataset definition is created, it is possible to search for the relevant files and store the result as a snapshot.

A SAM snapshot is the actual list of files that satisfy the metadata query at a particular point in time, i.e. when a particular analysis was run.

When this dataset definition is required in the future, the snapshot can be referred to instead of carrying out the search a second time, which could lead to a different set of files.

Using the same snapshot makes it possible to compare a current analysis with an earlier analysis and be assured of having the same files.

SAM Projects and Stations

A SAM project represents the pool of files available for delivery, and the collection of processes pulling files from that pool.

A SAM station is an application that coordinates all communication between projects, processes and the database, as well as all file delivery activities.

A station manages a collection of hardware resources. Projects are attached to stations.

http://samweb.fnal.gov:8480/station_monitor/

G-2 SAM station monitoring

Station	Projects	Last updated	Last activity
gm2	0	2016-03-30 11:30:00	No activity in last 6 hours
gm2-int	0	2016-03-30 11:30:00	No activity in last 6 hours

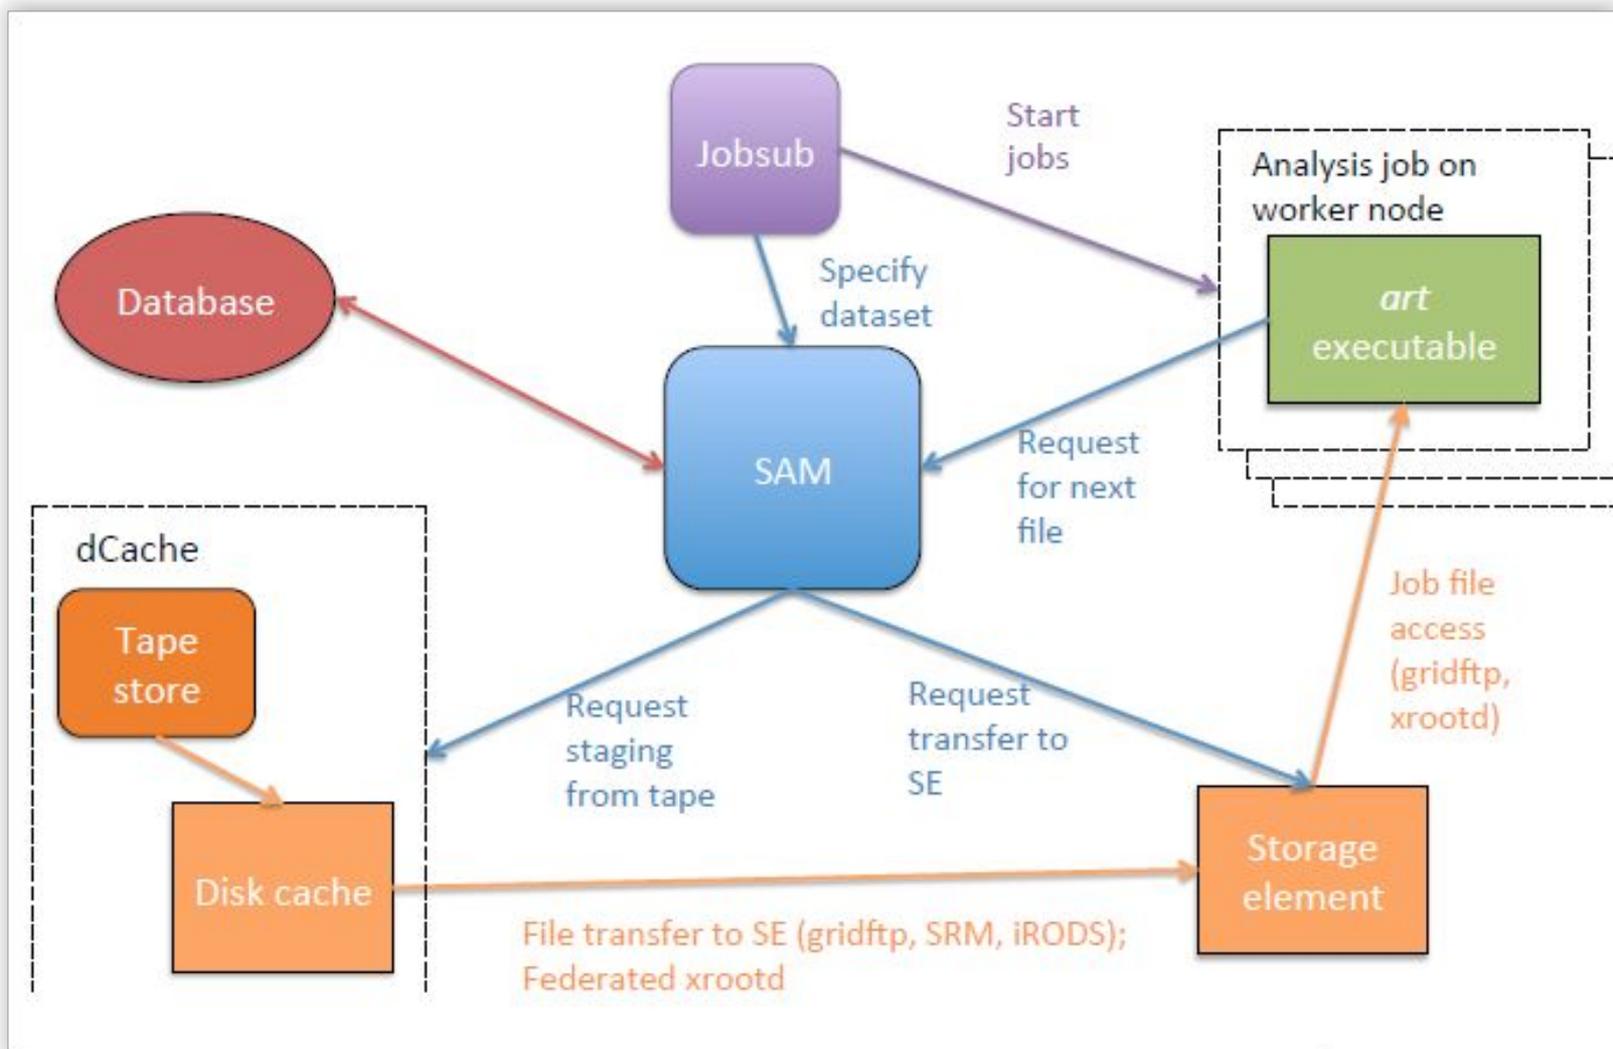
How does it work?

Files can be accessed by defining datasets, which are constructed in SAM through metadata queries.

At job run-time, SAM creates a project, defined as the list of all files satisfying the metadata query stored in the dataset from which the project was created.

Many worker nodes can connect to a project and request files, which are served by SAM as URL links.

Since each file can have multiple locations registered with SAM, SAM transparently chooses the appropriate one (for instance, choosing a file from central NFS storage rather than tape).



How does it work?

Once the worker node receives an access URL from SAM, it can use IFDH tools to fetch them.

IFDH (Intensity Frontier Data Handling) is a set of tools developed at Fermilab which facilitates the transfer of data from storage elements to compute nodes and back.

Allows the end user to run jobs on off-site grids without needing to know the details of where the data is coming from or what the correct transfer protocol is.

The way nova uses SAM

Metadata module (art module)

Manages the flow of metadata through files in various processing stages

The metadata module is in charge of collecting metadata from several sources, merging the metadata, and then forwarding it to the metadata service to be embedded in the output files.

Metadata is also collected from the input art files the job receives.

The metadata module uses the IFDH service to query SAM every time a new file is opened. This ensures that the metadata from the previous processing stage is preserved in each subsequent stage.

For g-2

We are making datasets and running jobs over them but not using all metadata capabilities

Current grid submission scripts written in bash

Want to have scripts modeled after Nova which are more flexible

- Command line options

- Stores a copy of the fcl file

- Metadata parameter set in the fcl file

Conclusion

Early on, it's possible to store all files on central NFS storage and run jobs on the local grid using explicit file lists

As the scale of raw data and Monte Carlo production grows, this paradigm became increasingly untenable

Use SAM to take care of data handling and focus on other things

G-2 production team in the process of incorporating SAM metadata options into grid submission scripts

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