

Data and Database Access at Remote Sites: Current and Future

Lee Lueking

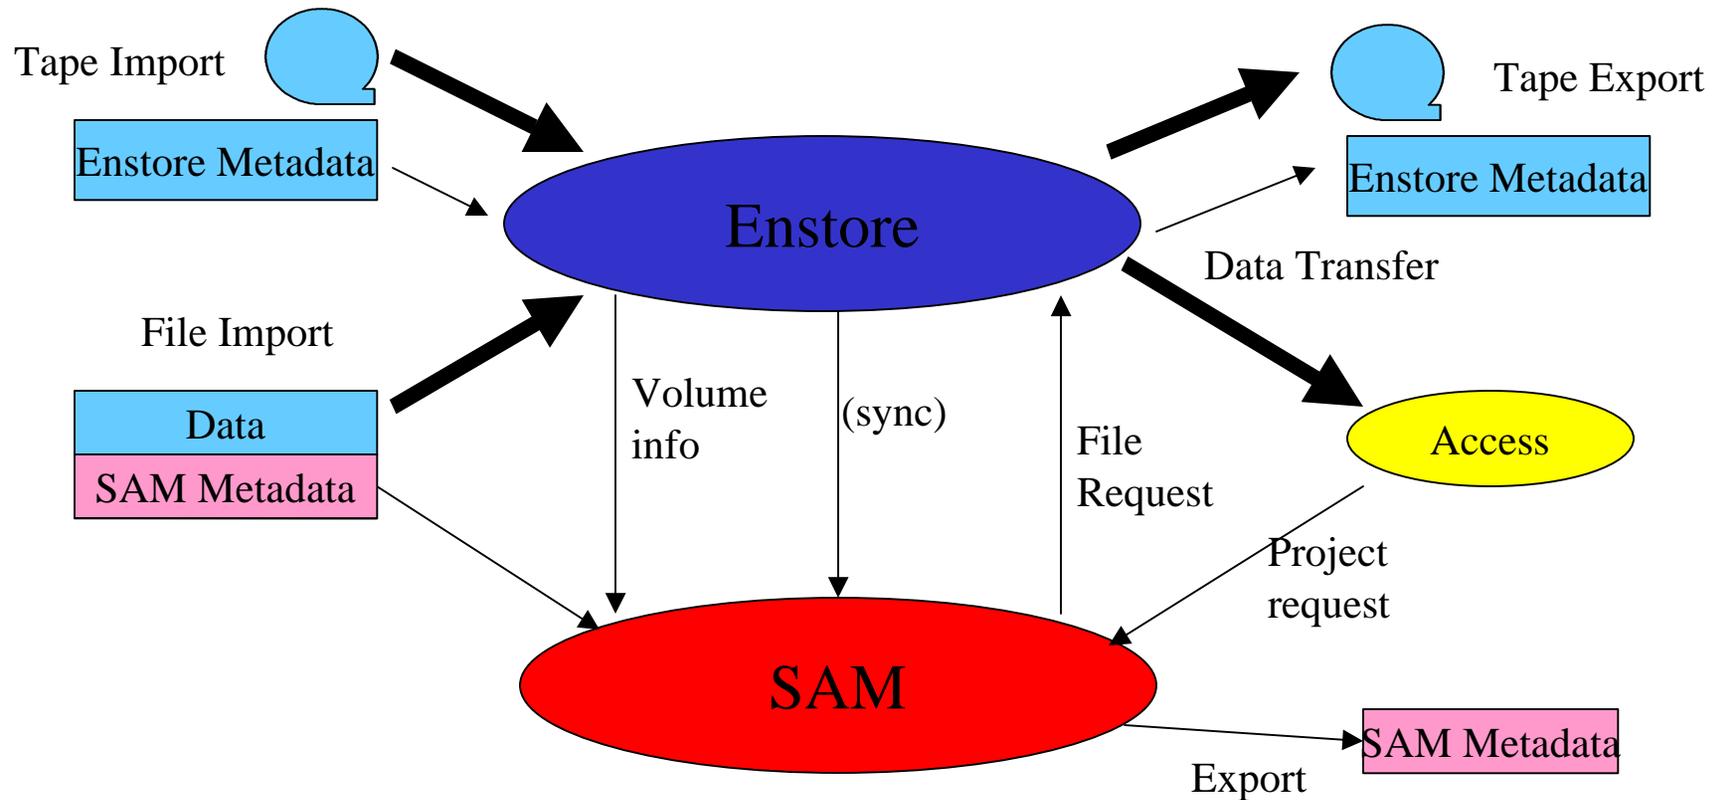
The D0 Off-site Computing Workshop

April 10-11, 2000

The Current Functional System

- Data Import
 - Network
 - Tape
 - Online and Farm
- Data Access: SAM “stations”
 - D0mino: central-analysis
 - D0bbin: D0 Farm
- Databases
 - SAM, miscellaneous
 - D0 FNAL Oracle Production Server

Data Import/Access/Export



Monte Carlo Network Import

- Cache areas are reserved on D0mino for buffering data: /sam/cache1/import and /sam/cache2/import.
- Three parts needed: 1) a Description file, 2) Data file(s), and 3) parameter file(s). Copied to the cache area.
- Clean up jobs are run to store metadata/data to sam/enstore about once per day.
- Sometimes the areas fill up due to problems with description files, or sam/enstore delays. Generally OK.
- Kors and others have worked on scripts to run at remote sites which deal with the error conditions and retry transfers when they fail. Parallel ftp transfers can be very fast.

Monte Carlo Tape Import

- Tapes are written at remote sites using `enstore_tape` utility.
- Tapes are shipped to Fermilab and inserted in to the D0 robot.
- Enstore metadata, sam description files, and parameter files are ftp'd to `/sam/cache1/import_tape` area.
- Data is declared to enstore and sam, then is available for use.
- This process has worked very well.
 - UTA is only site so far.
 - Data shipped on Mammoth I tapes.
 - Have DLT, Sony AIT I drives in the robot also.

Enstore Tape Import/Export

- Remote users can create cpio formatted tapes which can be added directly to the FNAL Enstore system.
- Interface

```
enstore_tape init[--tape-device=devname] [--tape-db=dbdir] [--verbose] --  
volume-label=label
```

```
enstore_tape write [--tape-device=devname] [--tape-db=dbdir] [--verbose] --  
volume-label=label file_list [file_list ...]
```

```
enstore_tape flatten_db [dbname] | output_file
```

perscription for reading tapes using generally available GNU cpio

- More details at: rip8.fnal.gov/enstore/volume_import.html

Enstore/D0 Tape Format

- Beginning:
 - ANSI Volume Header (VOL1, label, padding)
 - FILEMARK
- Data Files
 - Data in cpio wrapper
 - 2 GB filesize limit
 - FILEMARK between files
- End:
 - EOT Volume header (if written by Enstore)
 - FILEMARK
 - FILEMARK
 - EOT Volume header (if written by `enstore_tape`)

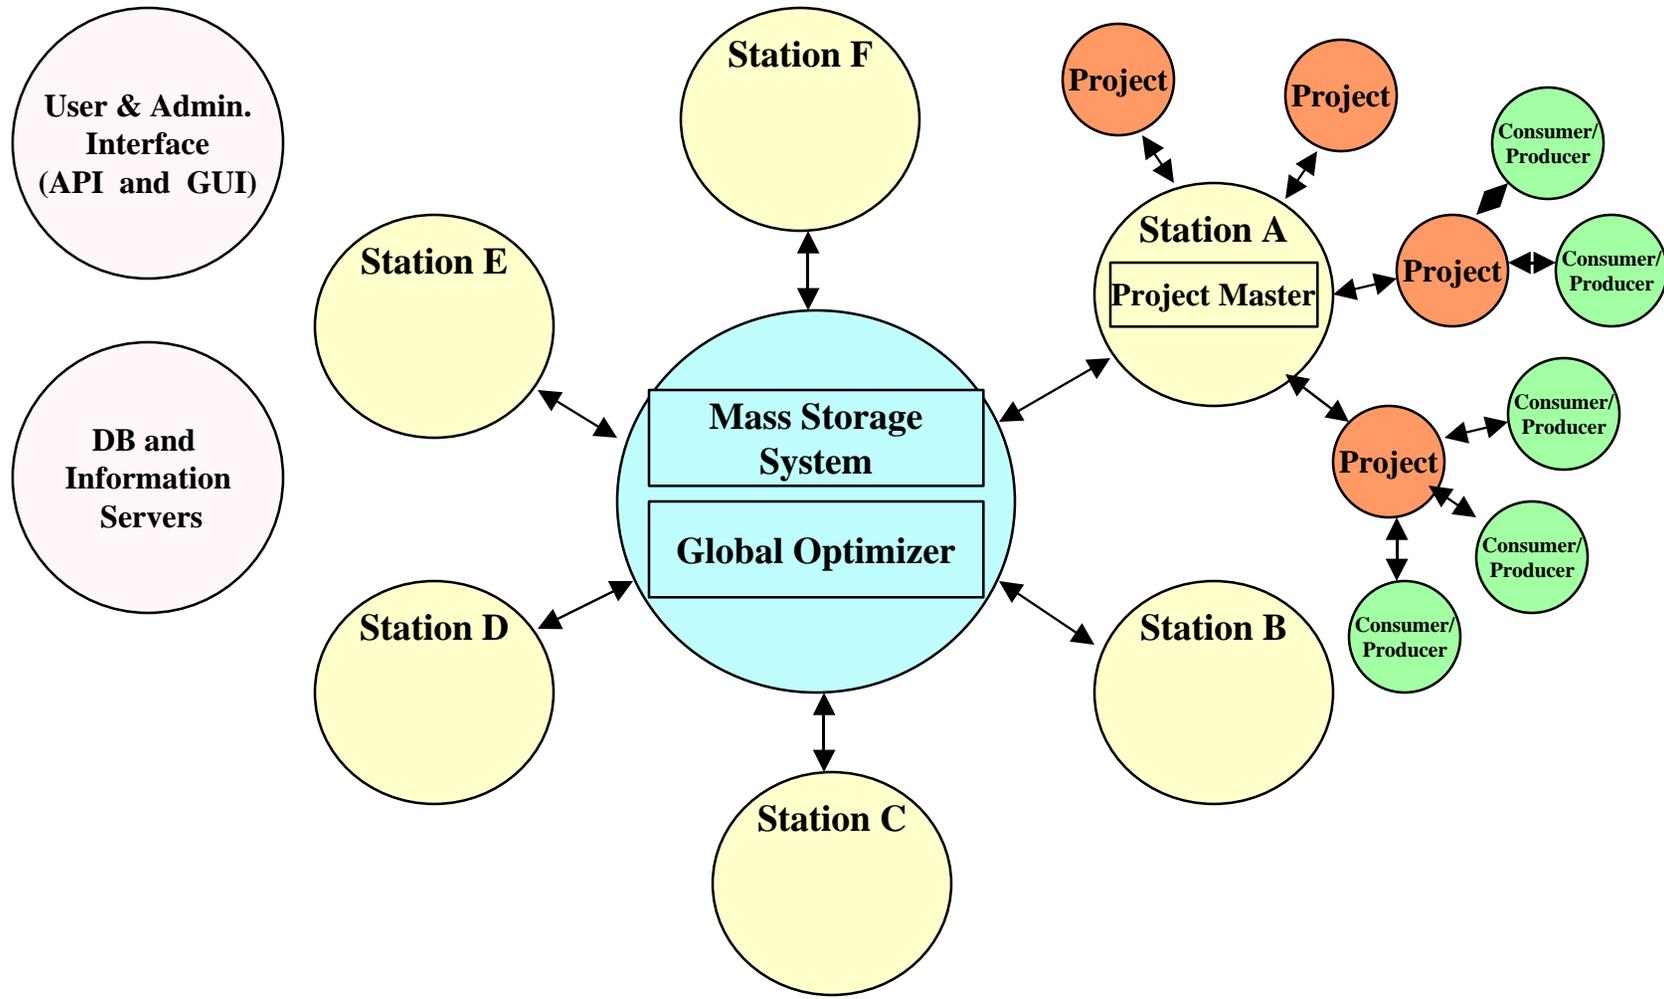
Online and Farm

- Use sam utility to store data and metadata
- `sam store --description=description_file \`
`--source=/source/dir --dest=/destination/dir`

Data Access via Stations

- Sam stations provide file stagers, and disk cache management.
- Projects are run under a station using allocated resources.
- The “central-analysis” station runs on D0mino
- Work is being done to test stations running outside FNAL. For example, in France, the “ccin2p3-analysis” station has been run under development.

The Sam Hierarchy



Data Access and Import In the Future

- Sam is designed to be highly scalable.
- Station support at remote sites is anticipated to be routine in the future.
- Heavily used data can be cached in the stations at remote sites.
- Features like parallel ftp will be built into the transfer mechanism and moving data to and from remote sites will be as simple as “sam run project” and “sam store”.
- A “shadow” stager/buffer at FNAL may be needed to match tape rate to slow network rates.

Data Access and Import In the Future (continued)

- Data produced at remote sites can be stored in their local tape management systems and the locations recorded in sam. Interfaces to their local sam stations will make the data accessible anywhere a sam station exists.

Oracle Database Servers: Present and Future

- D0 Oracle servers 1 & 2
 - D0Ora1 - development SUN 4500, 4 x 336 MHz Ultra SPARC CPU, 200 GB RAID disk.
 - D0Ora2 - production SUN 3500, 4 x 336 MHz Ultra SPARC, 100 GB disk.
- Plan to add CPU and RAID disk storage to 4500 and make it the production machine later in the spring. (May-June)
- We are looking into options to make this pair of servers a High Availability (HA) system, with fail-over capability so if one server dies, the other will take over. (July-August)



Future Database

- We do not have a clear picture of what the needs of remote database users are.
- Solutions beyond our current Fermi-central database servers could turn into a lot of support effort. As examples, consider distributed Oracle databases, or porting code and data to free databases like mySQL.
- For now, the tools will be DSPACK for calibration, flat files for trigger, and/or Oracle central db server access over the network.
- Serious planning and requirements are needed before we go beyond what we now have.

That's All Folks!