

Repository Hardware

Hardware requirements

- Any files that are stored on disk should be backed up, either to HPSS or our Tivoli tape system.
- It's worthwhile to keep using HPSS, because it takes care of our disaster-recovery needs. However, since it is slow to access (compounded by the aggregation issue), it may be worthwhile to store masters locally as well.
- We may want to replace Gavotte with a faster machine, because it takes a long time to process derivatives when we get a batch from somewhere. (Rifias took three months.)

Current service arrangement

Thalia

- Production Fedora server
- Storage for Fedora objects and datastreams
- Production Fedora services
- Temporary storage for other uses

Rhyme:

- Development Fedora server
- Fedora ingest system
- Storage for development Fedora objects and datastreams
- Development Fedora services

Feta:

- Production Tomcat, new
- Purl Resolver

Erato:

- Production Tomcat
- Production database
- Storage for existing collections that have not been moved to Fedora
- HPSS archiving

Brie:

- Development Tomcat, new
- subversion (SVN)

Urania:

- Development database
- CVS
- Development Tomcat

Euterpe:

- Storage for newly-digitized content that cannot be processed yet

Clio:

- Storage for existing collections that have not been moved to Fedora
- Serving (relatively) static web pages

Gigue:

- Persistent name mapping (PURLs, handles, etc.)
- Backups

Gavotte:

- Image processing

Algernon:

- Confluence/Jira Tomcat
- (DLXS) Storage for existing collections that have not been moved to Fedora

- Text indexing
- Text searching
- XML database

Melpomene:

- OAI

Not currently provided:

- JPEG2000 decoding
- Preservation integrity checking

Proposed future arrangement

This arrangement re-organizes the current layout a bit, centralizing functionality onto fewer machines. We would have to move to this arrangement gradually.

Thalia, production Fedora:

- Production Fedora server
- Tomcat for Fedora-specific utilities
- Storage for production Fedora objects and datastreams
- (Fedora) Text indexing
- (Fedora) Text searching
- (Fedora) OAI provider

Rhyme, development Fedora:

- Fedora ingest system
- Development Fedora server
- Fedora services
- Staging/test Fedora server

Feta, production webapps:

- Production Tomcat
- Serving (relatively) static web pages
- Storage for existing collections that have not been moved to Fedora
- XML database

Brie, development webapp/database:

- Development Tomcat
- Development database for Fedora
- Development database for other apps
- CVS

Bleu, off-line processing:

- Xsubmit
- Image processing (including JPEG2000 encoding)
- Storage for newly-digitized content that cannot be processed yet
- Preservation integrity checking
- HPSS archiving

????, database:

- Production database

Algernon:

- Confluence/Jira Tomcat
- (DLXS) Text indexing
- (DLXS) Text searching
- (DLXS) Storage for existing collections that have not been moved to Fedora

Gigue

- Persistent name mapping (PURLs, handles, etc.)
- Backups

Melpomene, Clio, and Euterpe will disappear. Euterpe's file serving capabilities may move to LIT.

The most pressing needs are to purchase replacements for Gavotte, Erato, and Urania. We need to develop short-term backup solution, while we wait for the UITS system-wide backup to be implemented.

Storage needs (as of 2006)

Rough estimates of throughput (master files only):

- Sheet music is 20MB/page. Maxed out at 20 pages/hr, it could produce 3.2 GB per day
- With the new digital camera, we could produce a maximum of 54GB/day
- Eviada is currently about 5TB/yr. If they expand to lossless compression, it could be 25TB/yr.

Selected stats for current collections:

- Hohenberger: 2143 images, each with master, thumbnail, screen, and large JPG. Masters take 13GB. Derivatives take 1.4GB. Full Fedora storage (derivatives, metadata, and resource index) takes 2.5GB.
- DIDO: 40,000 images, each with master, thumbnail, screen, and large images. Derivatives take 15GB.
- US Steel: 2200 images with master, thumbnail, and screen. Masters take 20GB.
- Cushman: 14,500 images. Derivatives (including notebooks) 3.2GB.
- Victorian Women Writers: (text only)
- Wright American Fiction: 400,000 pages. Derivatives are generated on the fly. Master images and text files are 111GB.
- Sheet Music (DeVincent): 50,000 pages. Screen and full/large for all pages; thumbnails for some. Derivatives 11GB.
- Sheet Music (Starr): Derivatives 600MB.
- Jane Johnson: Derivatives 300MB.
- IU Archives: 600 images. Does not have large size yet. Derivatives 113MB.
- The master files in euterpe\digitize currently take up 540GB. 200GB of this is under Indiana, which mostly includes Indiana Authors.

Rough guideline: If we create three derivative sizes for each image, 1GB per thousand images/pages.

An incredibly rough estimate: 100GB to store derivatives for all currently deployed collections (excluding audio/video), 1.5TB if we store all the masters as well. But this could be under-estimating by a bit, since we have many collections in the pipeline. Including the files in euterpe\digitize, and assuming roughly a 10:1 master to derivative size, we may have more like 200GB for derivatives and 4TB total.

We should definitely err on the high side when estimating storage needs. For image collections delivered through DLXS (like Wright American Fiction), the master files must be stored locally, and derivatives are generated on the fly.

Note that if we move our derivatives to JPEG2000, storage needs may change slightly, but not by very much.

We also need a "drop box" for bulk content acquisitions, like the 500GB of content we're getting for Newton.

Cost estimates (as of 2006)

We can get 5TB of disk in an Apple XServ Raid for \$20k.

A powerful 4-processor (dual core) machine from Sun with 32GB of RAM is available for \$50k.

We could completely replace the current tape system with a newer, bigger one for \$54k.

Virtualization

The UIS group is moving toward a virtualized system, where they have a big pool of server power that can be dynamically allocated. We may consider using it for some of our needs instead of purchasing new machines for everything.

Schedule:

- May 2007 - All UIS applications will be on the virtual system
- Fall 2007 - A pricing schedule for other departments will be available

Open questions

1. Are we relying on HPSS for storage of masters, or is it feasible to do something else ourselves?
 - a. We can probably handle the raw cost outlay of storing them ourselves. It's the extra workload that would be a problem. But keeping HPSS introduces its own workload.
 - b. Keeping local copies of the masters would allow us to salvage lost/broken items, making our preservation system more robust. The preservation system could periodically download a single aggregate from HPSS, and compare all of its contents to the local copies, alerting someone if there are any differences.
2. Do we need specific replacements for Algernon and Gigue? Or can their functions be moved onto other machines?