

# Experiment to Investigate the Scalability of a DSpace-based Archive

Dharitri Misra, Ph. D.

National Library of Medicine  
April 30, 2008



THE LISTER HILL NATIONAL CENTER  
FOR BIOMEDICAL COMMUNICATIONS

*A Research Division of the U.S. National Library of Medicine*

# Outline

- SPER Archival System
- Testing SPER's Scalability
- Test Details
- Results
- Conclusion



THE LISTER HILL NATIONAL CENTER  
FOR BIOMEDICAL COMMUNICATIONS

*A Research Division of the U.S. National Library of Medicine*

# Description of SPER

**SPER (System for the Preservation of Electronic Resources) is an R&D-based archival system**

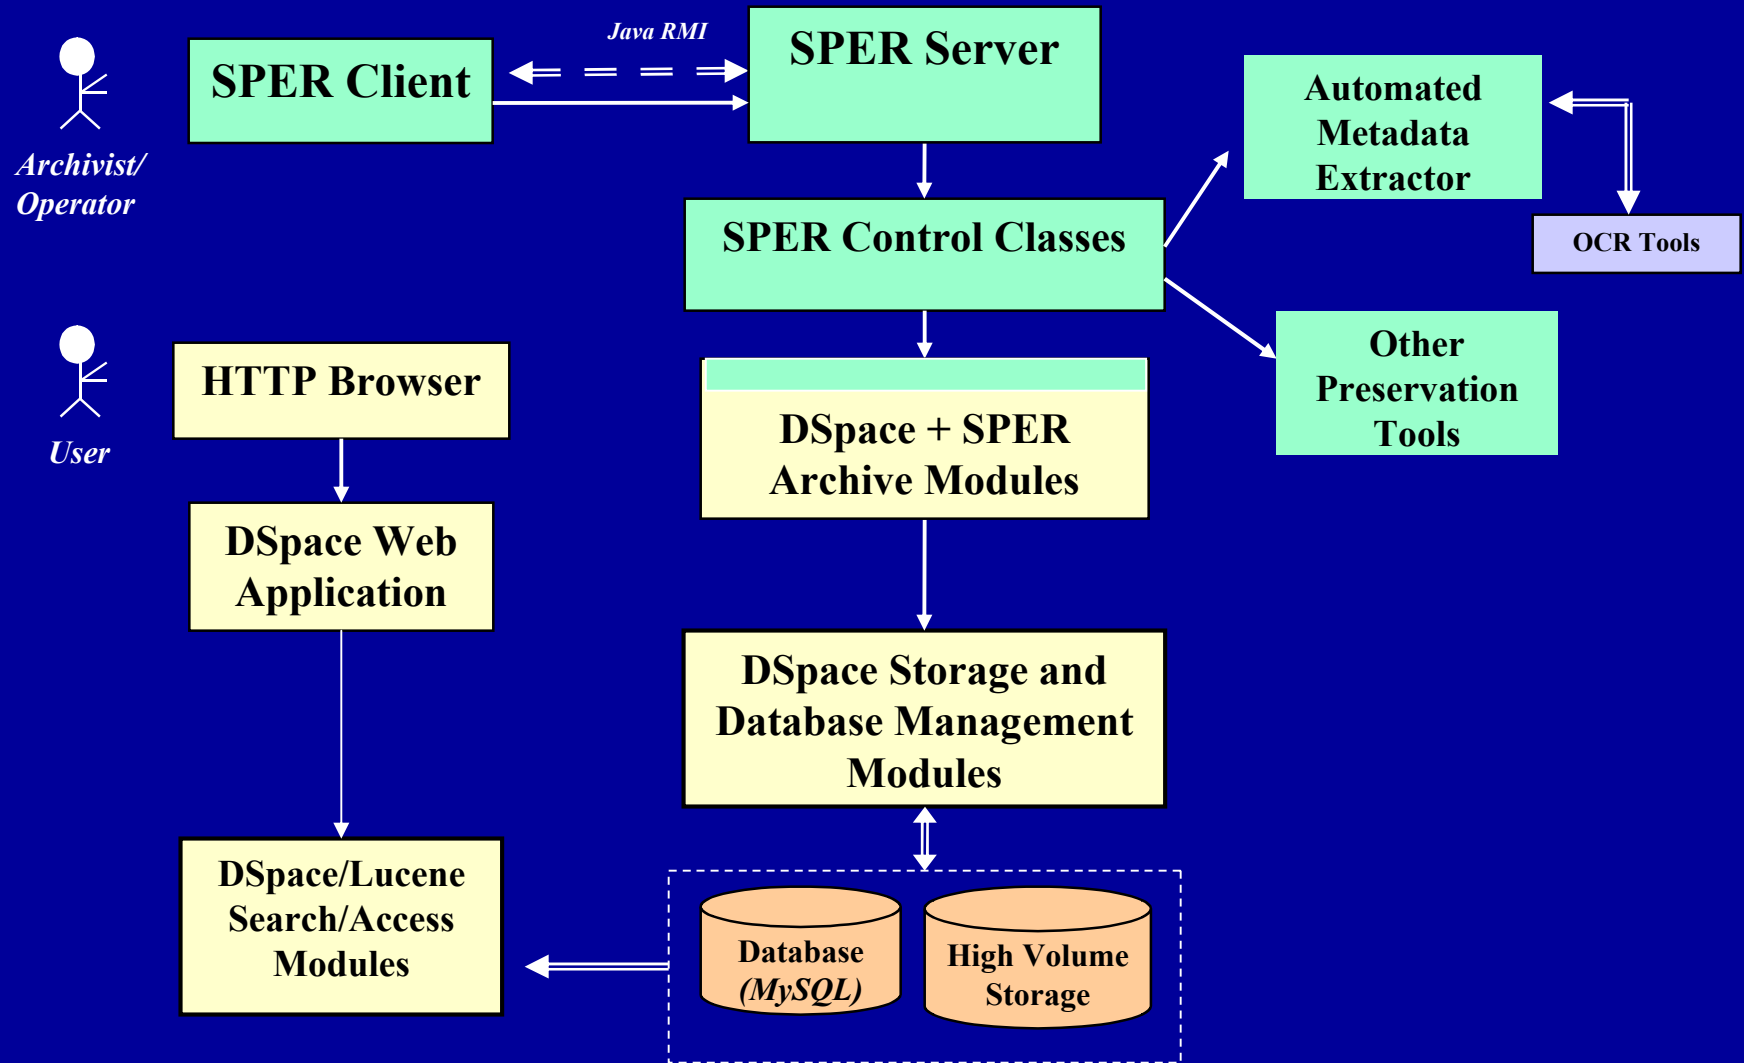
- Developed at NLM to study important aspects of digital preservation
- **Ingest and access capability provided by DSpace**
- Integrates other functional capabilities (e.g., automated metadata extraction)
- Customized batch-based operations
- Java RMI client-server model and Swing-based user interface
- Currently used to archive a collection of historic medico-legal documents from FDA



THE LISTER HILL NATIONAL CENTER  
FOR BIOMEDICAL COMMUNICATIONS

*A Research Division of the U.S. National Library of Medicine*

# SPER System Architecture



# Why Test Scalability

**Scalability of an Archive: Ability of the system to accommodate large number of items without compromising performance**

- Unsatisfactory performance leads to redesign or rejection of the archival system later
- Ingest, access and indexing performance must be evaluated in building a large archive
- Benchmarking difficult due to lack of large amount of realistic data and test software
- **Often ignored in evaluating/choosing an archival system at the installation site**



THE LISTER HILL NATIONAL CENTER  
FOR BIOMEDICAL COMMUNICATIONS

*A Research Division of the U.S. National Library of Medicine*

# Testing the Scalability of SPER

- SPER uses DSpace for building the archive  
*Scalability of SPER => Scalability of DSpace*
- No performance-based benchmarks found for large DSpace archives (or other well-known archival systems)
- **Task:**  
*Measure ingest performance of the SPER/DSpace archive, as a function of archive size, up to a million items*

# Test Components

## ■ Ingest data

Created through replication of two medico-legal FDA documents collections archived by SPER

- Monochrome scanned TIFF images
- Automatically extracted metadata files
- OCR Textline files

## ■ Test procedures and software

Developed in-house using Java, Perl and SQL Query

## ■ Test environment

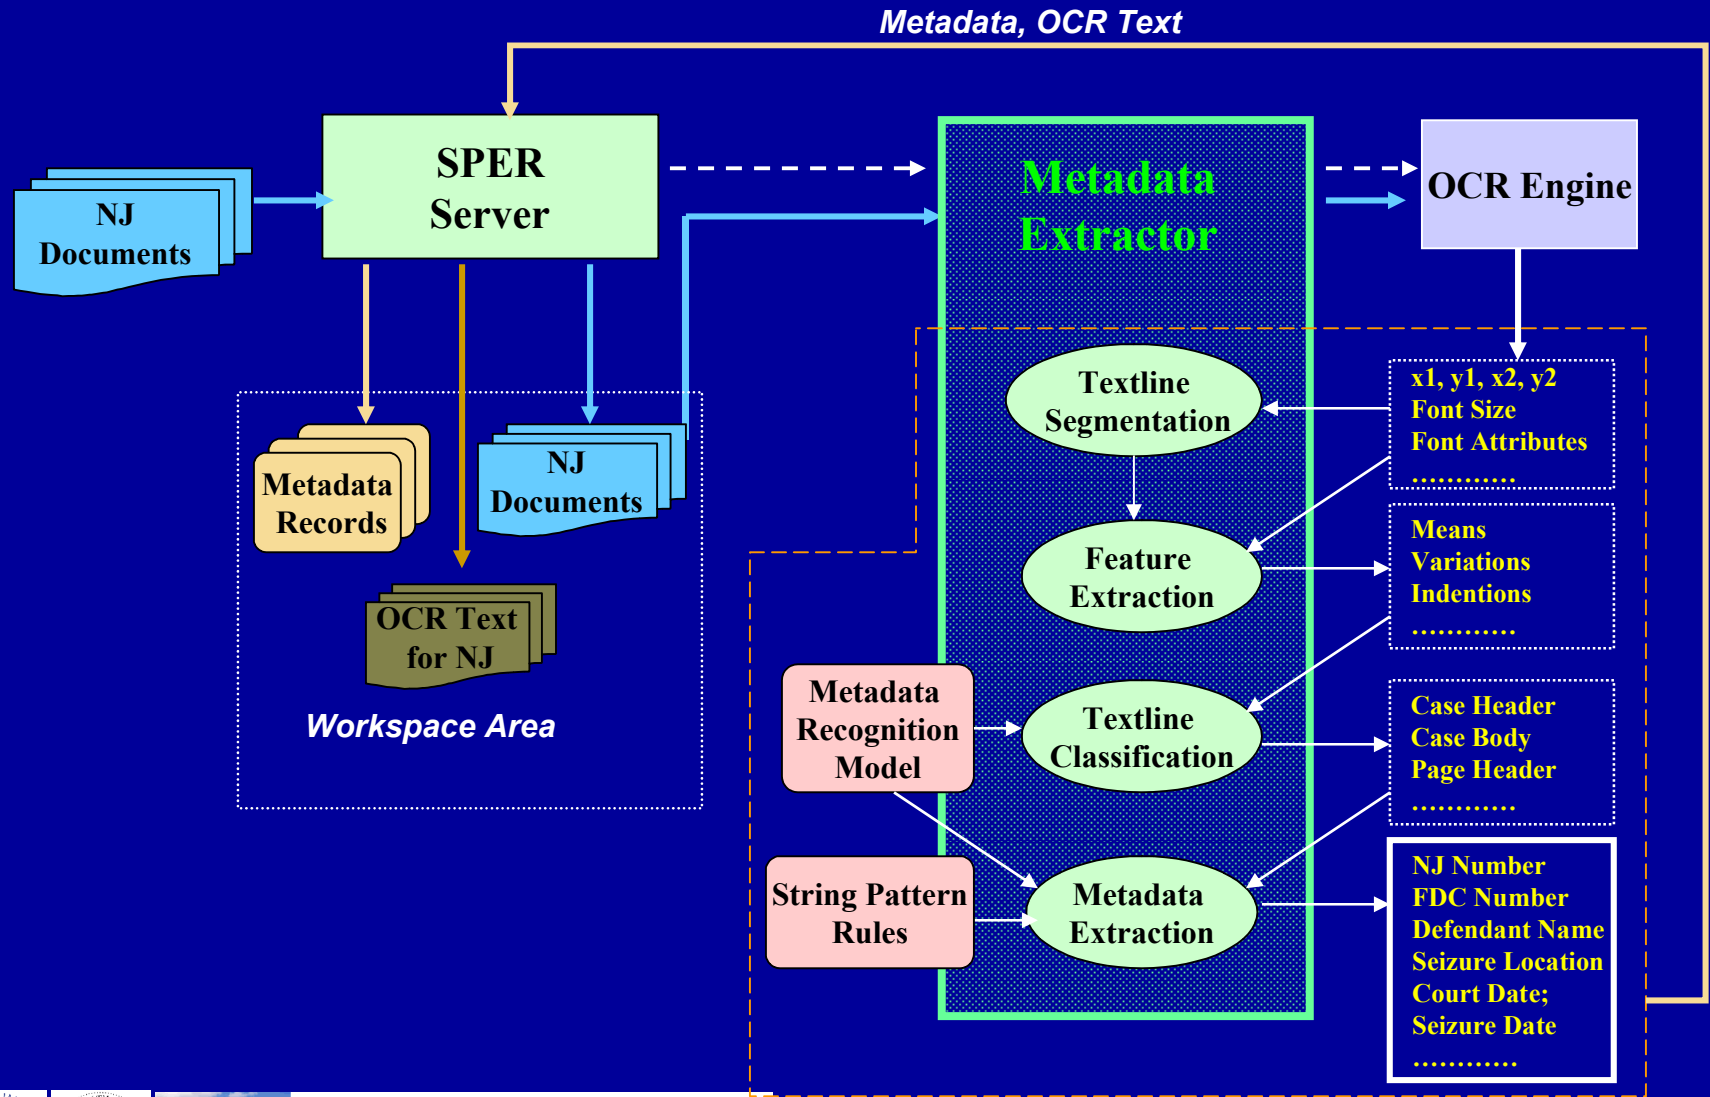
- Dedicated virtual machine on a Sun X4500 server 4.5 TB disk space and with 3 GB pre-allocated memory for test
- Sun Solaris 10, Java 1.4.2
- DSpace: V1.4.1 with *MySQL V5.01* database



THE LISTER HILL NATIONAL CENTER  
FOR BIOMEDICAL COMMUNICATIONS

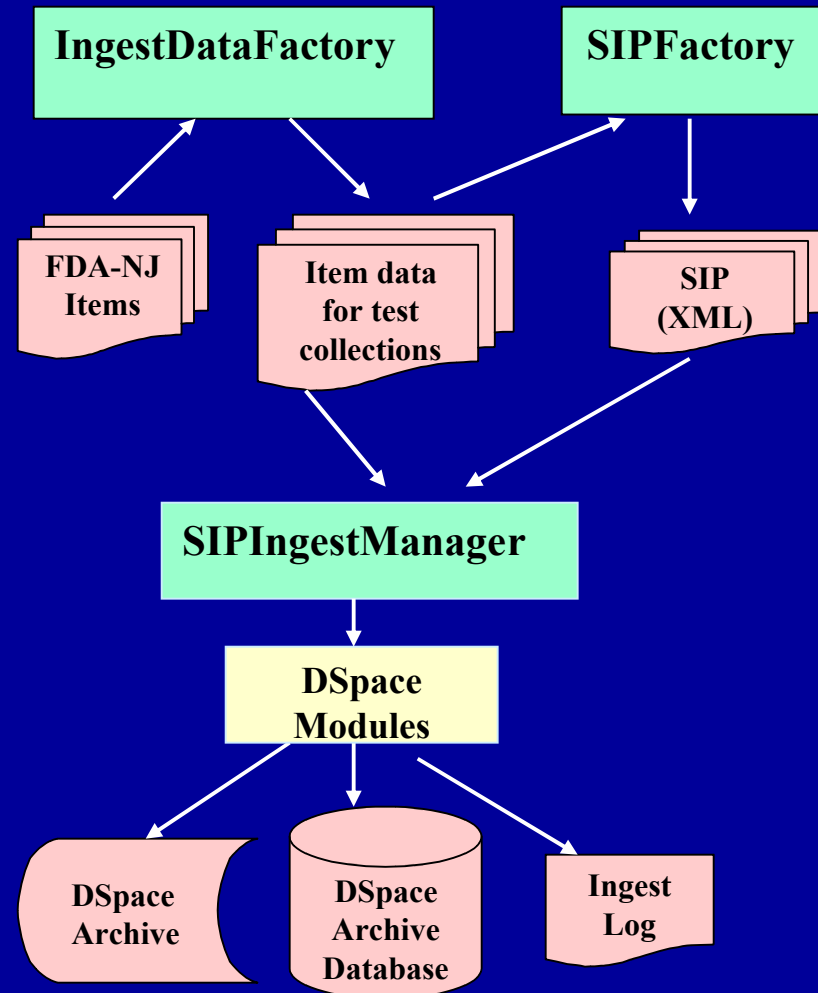
*A Research Division of the U.S. National Library of Medicine*

# Generation of FDA-NJ Collection Data



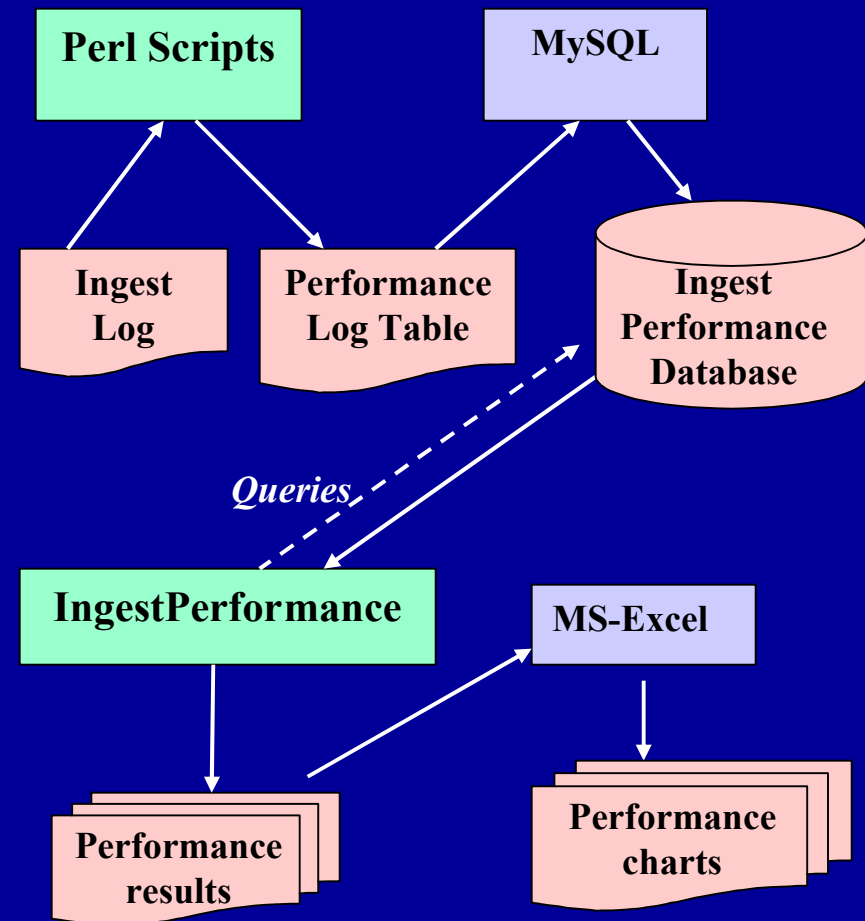
# Generation of Ingest Data

- Two FDA collections with 11000 and 220 items respectively submitted to IngestDataFactory for replication
- Data from one or more test collections batched into a single SIP (XML file), by SIPFactory
- Size of a SIP limited to 36,000 items
- One or more SIPs submitted to the SIPIngestManager in each invocation to be batch ingested to DSpace



# Computation of Performance Results

- Time of occurrence of various events in the ingest pipeline from Log file extracted and formatted into Performance Log tables
- Output loaded into an Ingest Performance database through MySQL scripts
- IngestPerformance module issues various queries to this database and outputs performance results to text files
- Finally, Performance result text files are plotted through MS-Excel into result charts



# Structure of the Test Archive

<b>No. of DSpace Communities</b>	<b>32</b>
<b>No. of Collections</b>	<b>109</b>
<b>No. of Collections per Community</b>	<b>varied between 1 to 5</b>
<b>No. of indexed fields in the Dublin Core metadata file</b>	<b>12</b>
<b>No. of Ingested Items</b>	<b>1,041,790</b>
<b>No. of Items with 4, 5 and 10 associated bitstreams</b>	<b>50%, 25% and 1% respectively</b>
<b>No. of Items with 50, 60 and 103 bitstreams</b>	<b>58 each</b>
<b>Average amount of data associated with an item</b>	<b>140 KB</b>



**THE LISTER HILL NATIONAL CENTER  
FOR BIOMEDICAL COMMUNICATIONS**

*A Research Division of the U.S. National Library of Medicine*

# Measured Parameters

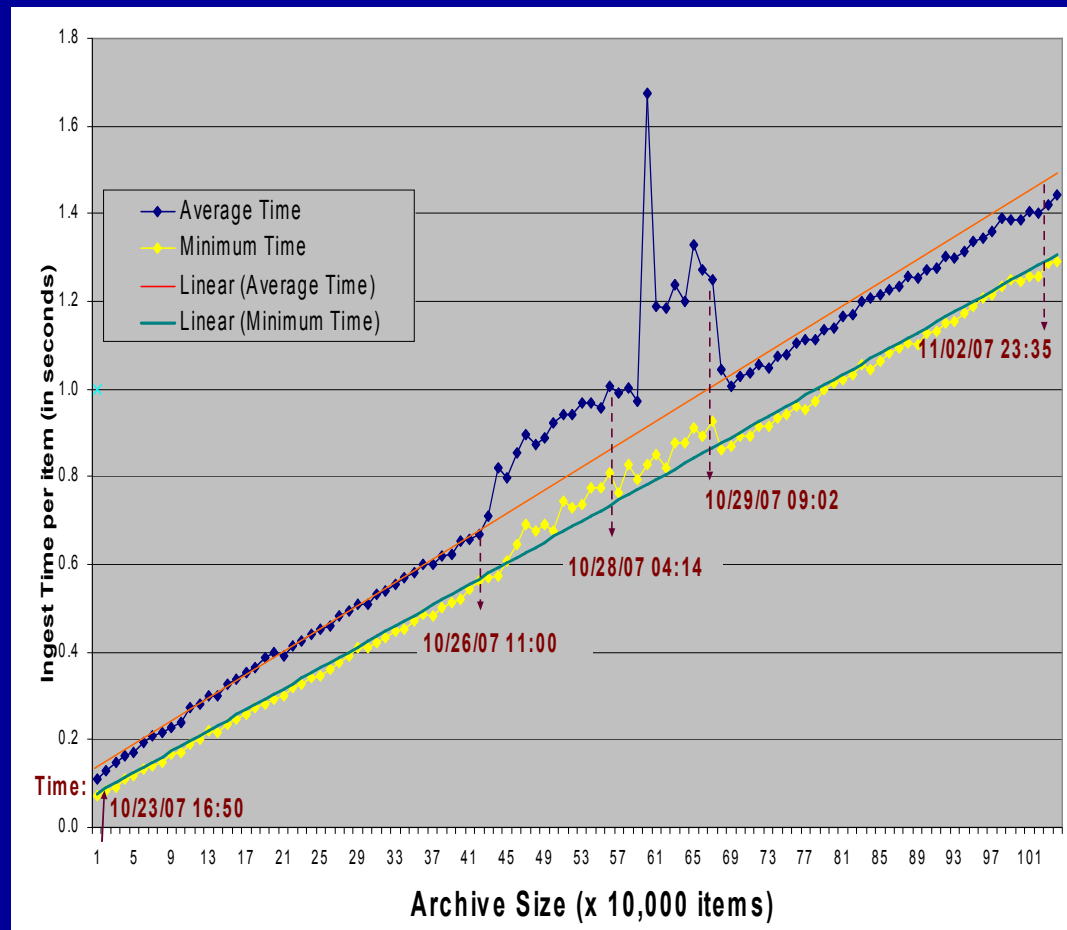
- Ingested items grouped into sets of 10,000 items
  - Each set comprises one point of measurement
  - Over 100 measurement points for “Archive Size”
- Parameters measured/computed for each set:
  - Average time to ingest an item
  - Minimum time to ingest an item
  - Total number of items ingested within +/- 25 percent of average ingest time
- Average time taken to ingest items with 4, 6, 8 and more bitstreams as a function of archive size.
- Test was conducted twice, using the same data but under different operational conditions  
*(Second one was a stress test for the system)*



THE LISTER HILL NATIONAL CENTER  
FOR BIOMEDICAL COMMUNICATIONS

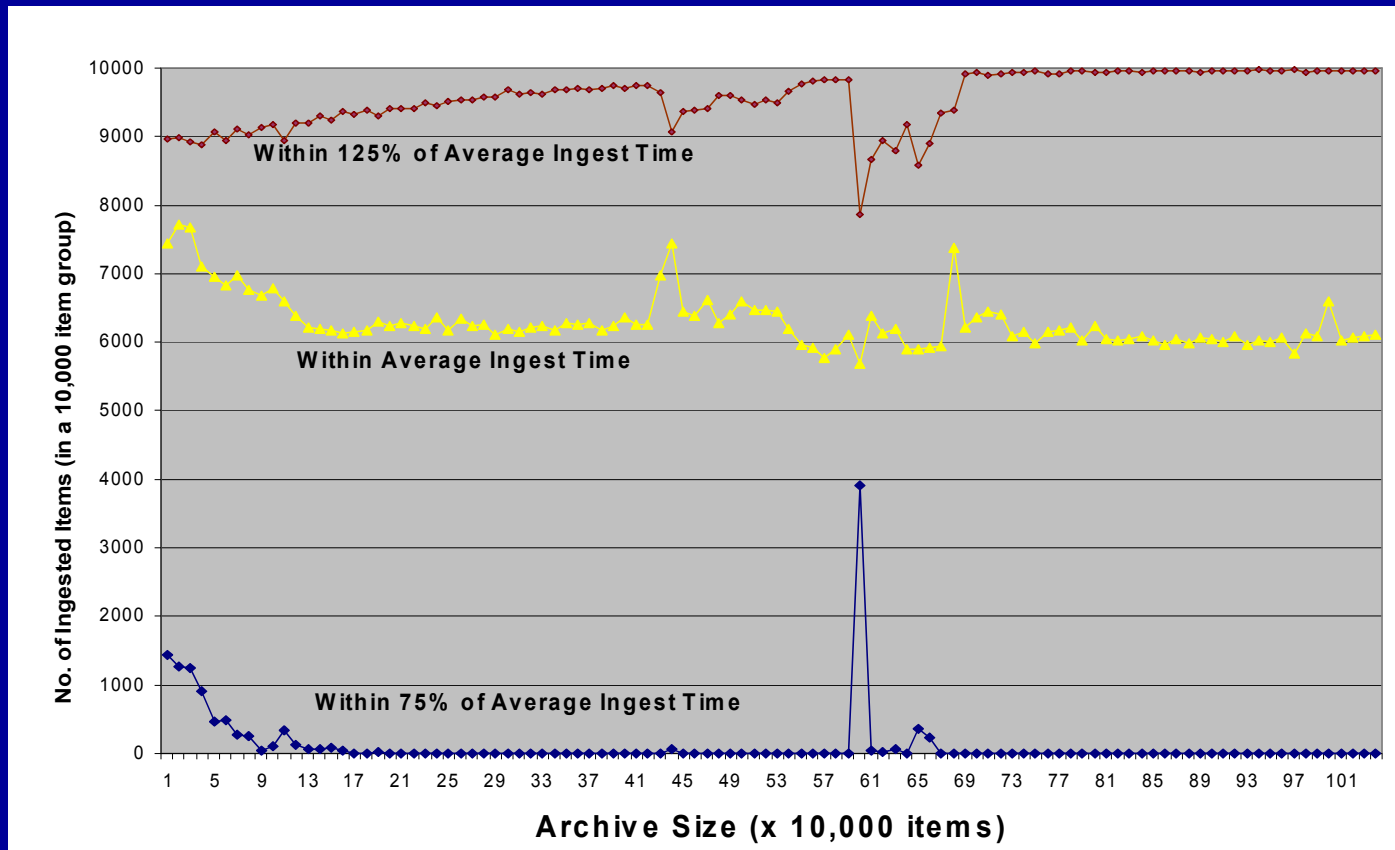
*A Research Division of the U.S. National Library of Medicine*

# Average and Minimum Ingest Times



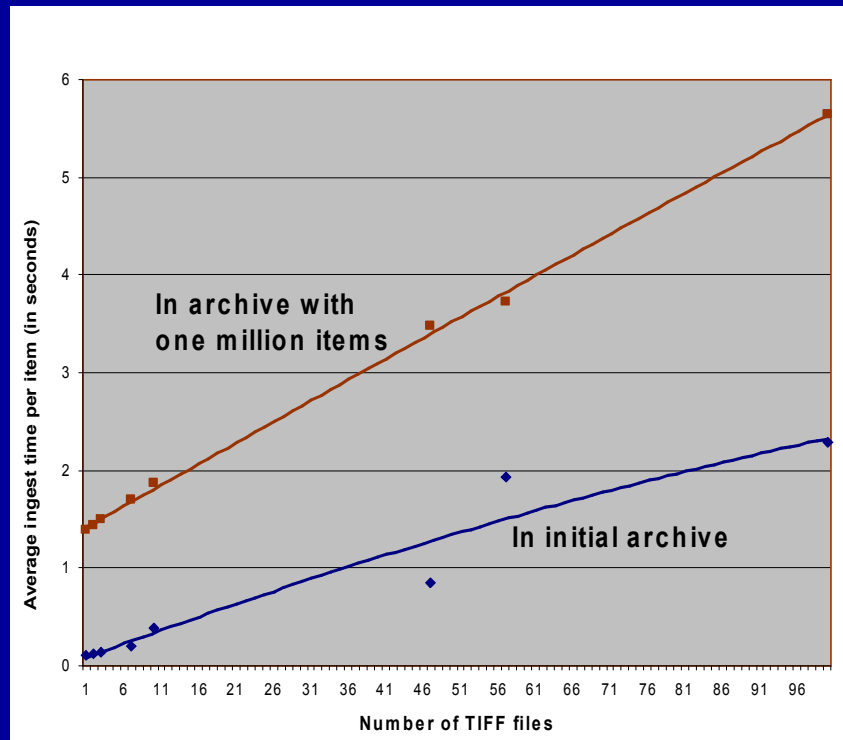
- Total time to ingest one million items ~ 240 hours
- Average initial ingest time 0.11 sec
- Average final ingest time: 1.4 sec
- Steady slopes for both average and minimum ingest times
- Deviation occurred due to network activities external to SPER

# Number of Items Ingested within Different Time spreads vs. Archive size



- **Most items ingested within +/- 25% of average ingest time, irrespective of archive size**

# Ingest Time vs. Number of TIFF files



# of Bitstreams	# of TIFF files	Av. time at start (sec)	Av. time at one million (sec)	Factor w.r.t. (X)	% increase per TIFF
4	1	.106	1.39 (X)	1.0	-
5	2	.125	1.435	1.032	3.2
6	3	.144	1.5	1.079	3.95
10	7	.205	1.698	1.22	3.66
13	10	.389	1.874	1.35	3.88
50	47	.845	3.475	2.5	3.26
60	57	1.93	3.723	2.68	3.0
103	100	2.28	5.651	4.06	3.10

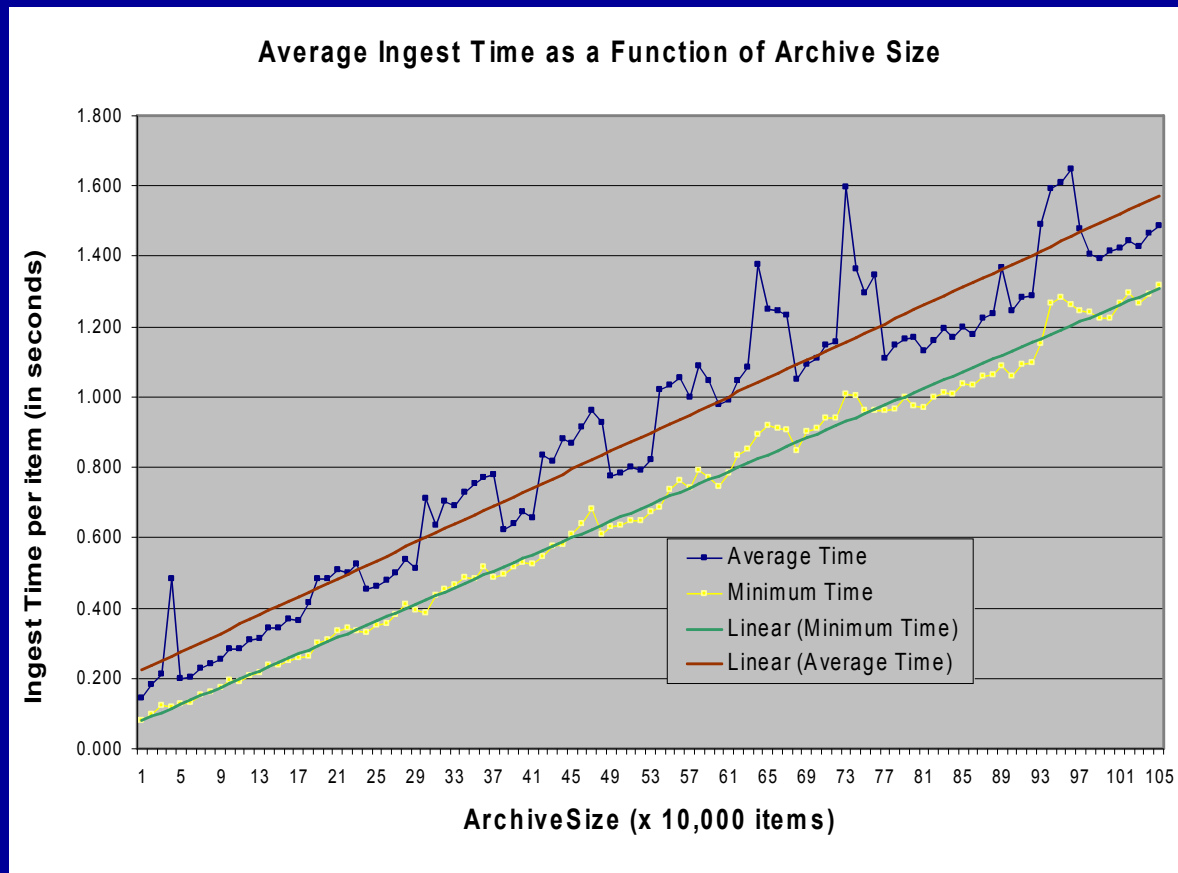
➤ Increase in Ingest time for each additional bitstream: 3 to 4 %



THE LISTER HILL NATIONAL CENTER  
FOR BIOMEDICAL COMMUNICATIONS

A Research Division of the U.S. National Library of Medicine

# Ingest Times under a Realistic Scenario



- Ingest operation not contiguous, done with long intermediate time gaps (over one month)
- Several peaks in ingest time, around new ingest starts
- Peaks most likely caused by system, cache, database initialization, etc.
- Averages and trends similar to that under more controlled environment

# Summary

- No hidden problems or flaws in system architecture detected in building a one-million item archive
- Ingesting monochrome TIFF images on a Sun X4500 Server increases almost linearly from ~0.1 to 1.4 sec as the DSpace archive increases in size
- The tests are repeatable (yielding similar results)
- File upload and other overheads negligible (approx. 10%) compared to actual archiving and indexing of data



THE LISTER HILL NATIONAL CENTER  
FOR BIOMEDICAL COMMUNICATIONS

*A Research Division of the U.S. National Library of Medicine*

# Conclusion

- Our archive, built on DSpace, shows acceptable performance in ingesting up to a million items
- Larger file sizes will not significantly affect performance
- Additional bitstreams would cause some increase (3% to 4% per bitstream) in ingest time
- Additional benchmarks needed for building DSpace archives much larger than 1 million items
- **Our benchmarks should be useful to other DSpace installation sites concerned with performance**



THE LISTER HILL NATIONAL CENTER  
FOR BIOMEDICAL COMMUNICATIONS

*A Research Division of the U.S. National Library of Medicine*