



# Open Source Long Term Preservation Archives

Richard Matthews  
Sun Microsystems, Inc.



# Presentation

- Prepared by:
  - > Keith Rajecki
  - > Industry Solutions Architect
  - > Global Education & Research
- Presented by:
  - > Rick Matthews
  - > Sr. Staff Engineer
  - > Solaris Software, Archive Products Group

# Agenda

- Sun, Open Source, and Communities
- Preservation Archiving Trends
- Sun Archiving Storage Solutions
- References

# Sun Microsystems Today

Fortune  
**211**  
Company

Annual Revenues  
**\$13+ Billion**

Annual R&D  
**~\$2 Billion**

Worldwide  
Employees  
**35,000**

Java Devices  
**6 Billion**

Java Developers  
**5 Million**

Annual Storage  
Petabytes  
Shipped  
**410**

SPARC  
Embedded  
Processors  
**44+ Million**

U.S. Patents  
**5,000+**


Solaris 10  
Licenses  
**7 Million**

Cash  
**\$4.8 Billion**

Business  
Presence  
**100 Countries**

# Sun's Open Source Strategy

## Developer Preference



- More core developers
- More deploying developers
- More partners

## User Preference



- Free to use
- More platform choice
- More suppliers
- Larger user community

## Value Proposition



- Business Deployment
- Sun's target market
- Binary distribution
- Pay for value

# Sun and Open Source Software

- OSS has been part of Sun's DNA for awhile

“Every software asset we produce is open source. If it isn't today, it will be pretty damn quickly.”

**Jonathan Schwartz**

CEO, Sun Microsystems

January, 2007

- Sun's commitment to OSS Communities
  - > This includes Open Repository Communities
  - > More than just a Storage perspective
  - > Dedicated organization to support OSS Communities

# Sun's Open Stack

Flexible and Heterogeneous with Zero Barrier to Exit

<b>Database Platform</b>	
<b>Application Infrastructure</b>	
<b>Virtualization</b>	
<b>Operating System</b>	
<b>Partners</b>	
<b>Architecture</b>	



# OpenSPARC

The Most Open Platform on the Planet

- 4300 downloads to date
- 14 million lines of source code
- Community interest: 1 to 1000 core systems
- First derivative design: SimplyRISC S1 core



OpenSolaris + OpenSPARC = Only Truly Open Platform

[www.opensparc.net](http://www.opensparc.net)



# OpenSolaris



# opensolaris™

90,000 Members

64 Community Projects,  
BrandZ, DTrace, Solaris ZFS, Zones

53 User Groups Worldwide

260 Code Contributions

2006 Codie “Best Open Source Solution”

2005 Open Source World Editor's Choice

2005 InfoWorld Innovators Award

2005 MIT Young Innovator Bryan Cantrill

Innovation Happens Everywhere

www.opensolaris.org

# Sun is Committed to Developer Communities

Building  
**Open  
and  
Free**  
Communities



Java

Solaris

SPARC

Building a Vibrant Ecosystem: Sun is the Largest Commercial Contributor to Open Source Communities

Community

Infrastructure

*The Source for Java Technology Collaboration*

Ecosystem

# Sun Preservation Archiving Community

- PASIG Meeting May 27-30 San Francisco  
[www.sun-pasig.org](http://www.sun-pasig.org)
- Comparison of high-level OAIS architectures, services-oriented architecture, and use cases
- Sharing of best practices and software code
- Cooperation on standard, open, 'in-a-box' solutions around repository technologies
- Review of preservation and archiving storage architectures and eResearch data set management
- Discussion of the uses of commercial third-party and community-developed solutions

## Trends: Growth and Preservation

**Humans created 161 exabytes of data in 2006, approximately 3 million times the information in all the books ever written, according to IDC.**

*Source: "An Inconvenient IT Truth," By Michael Vizard, eWeek, 06/22/07*

**80% of all movies made before 1940 are gone.**

*Source: CLIR Report "The Future of the Past", 1999*

# Repository Archiving Projects

- Compliance
- Book and Image Digitization and Sharing
- National Heritage Content
- Newspapers
- Replicated, Tiered Repositories for Archived Materials
- Research Data, Applications, and Systems
- Science, Technology, and Medical Journals
- Born Digital Materials

# What's The Buzz?

- Problem:
  - > Exponential growth of digital content, now and in the future
  - > Powerful, flexible infrastructure required to archive:
    - Store **unstructured, fixed content**
    - Search that content
    - Preserve that content for the long-term
- One Proposed Solution:
  - > Fedora plus Sun StorageTek 5800 “*Honeycomb*”

# SAM-QFS Infinite Archive System

- SAM-QFS: World's best policy based multi-tiered archive manager
  - > Application Transparent dynamic data movement
  - > Four tiers, local and remote
  - > "Continuous Archive" = CDP
  - > WORM & Retention management
- Infinite Archive System
  - > Scalable multi-tiered SAM-QFS
  - > platform base
  - > 10-256TB systems.
  - > Data-In-Place upgrade





# Sun StorageTek 5800 “Honeycomb”

## Content-aware Open Storage



- What It Is
  - > 'Smart', network-attached, clustered, racked storage system
- What It Provides
  - > 64TB (raw) per rack: data objects + metadata
  - > WORM Objects
  - > Metadata awareness built into the design
  - > Reliability, persistency, currency assurances
- Open System, Open Source Software

# Sun StorageTek 5800 “Honeycomb”


## Content-aware Open Storage



- RAIN architecture
  - > Symmetric cluster CPU, memory, SATA Disks
- Each node
  - > Opteron-based SunFire server
  - > Solaris 10
  - > 3 GB RAM
  - > Dual Gig-E
  - > 4 x 500GB SATA
- L2 load-spreading switches
- Service processor

# Sun StorageTek 5800 “Honeycomb”

## Content-aware Open Storage



'Half Cell'  
8 servers  
16 TB (raw)

'Full Cell'  
16 servers  
32 TB (raw)

'Rack' (2 Cells)  
16 servers  
32 TB (raw)

'Hive' (N Cells)  
*[Future]*

'Hot Scaling'



# Why Honeycomb?

- Architecture optimized to store and retrieve **unstructured fixed content**
- Object storage, metadata aware
- Extreme data protection via RAID6, data self-healing, bit-rot detection
  - > Mean Time To Data Loss > 2M years
- A commitment to standards
  - > Dublin Core metadata
  - > Web DAV
  - > Future: XAM (metadata + query model)

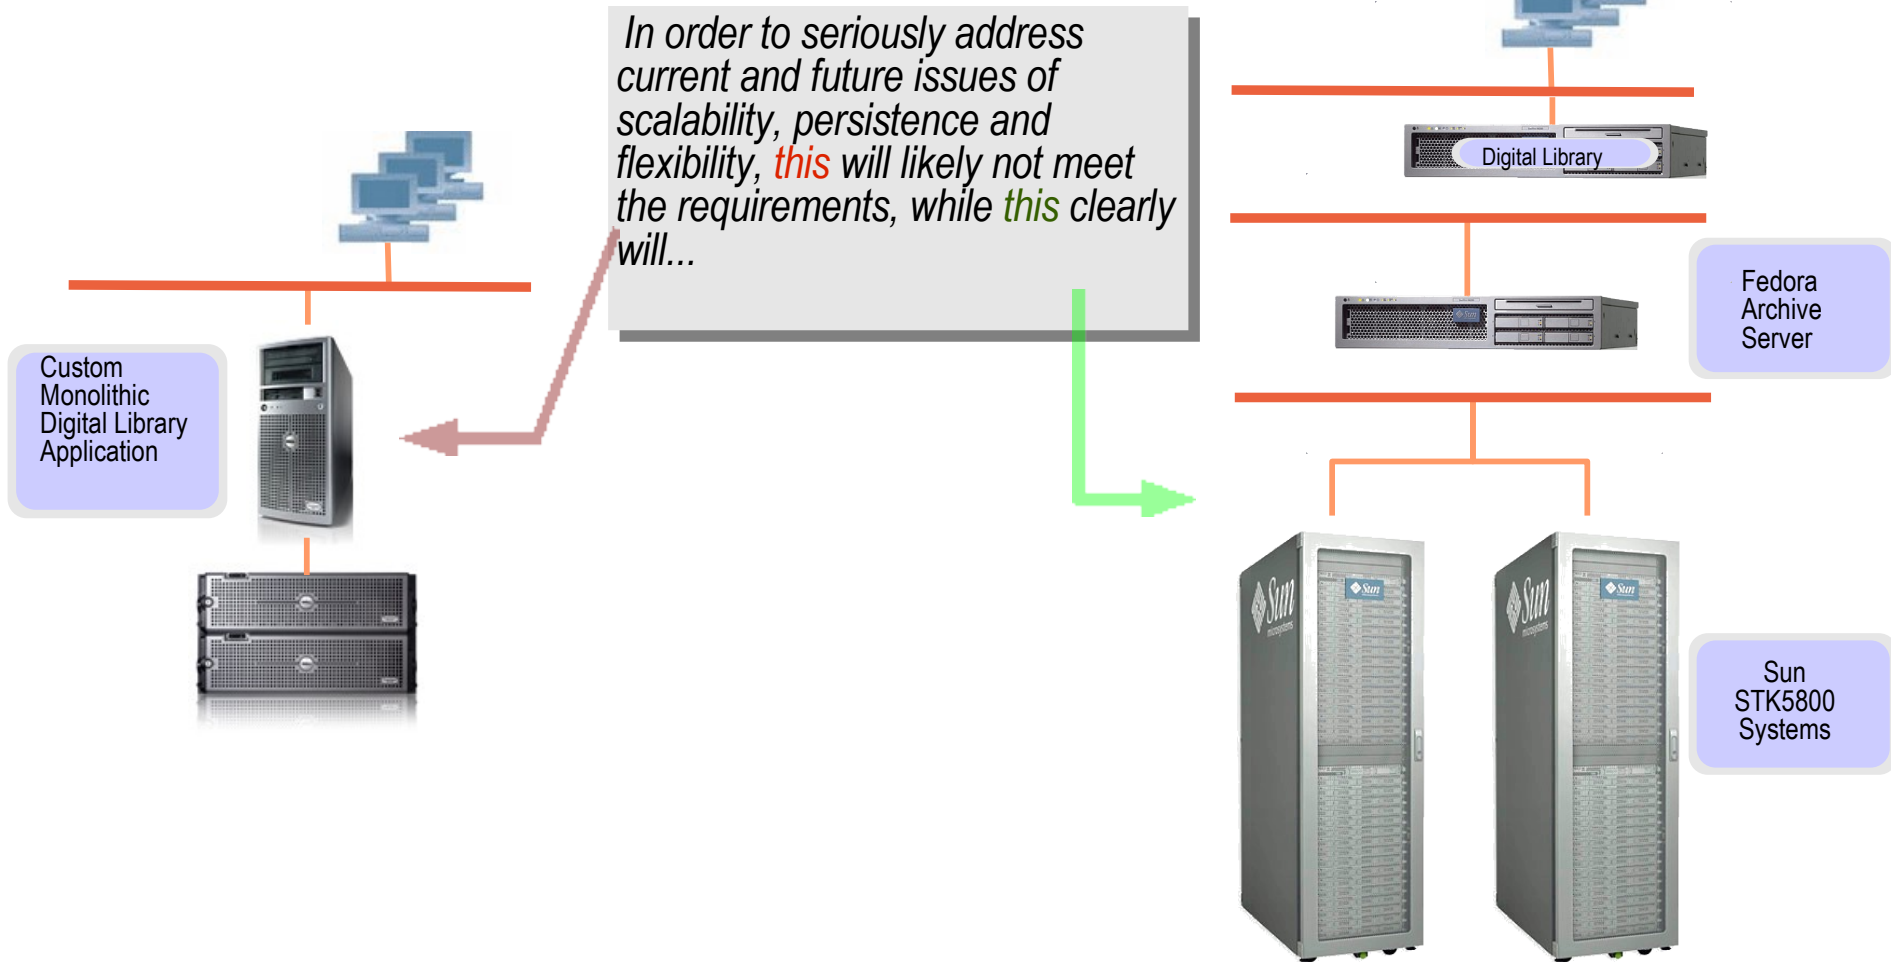
# Why Honeycomb?

- Open Source strategy fits with majority of repository/archive software efforts
- Standard Java and C APIs in SDK
- Horizontal Scaling as storage needs grow
- Dublin Core is only the beginning
- Platform-agnostic
- [Near Future] On-board local data services available (“Storage Beans”)

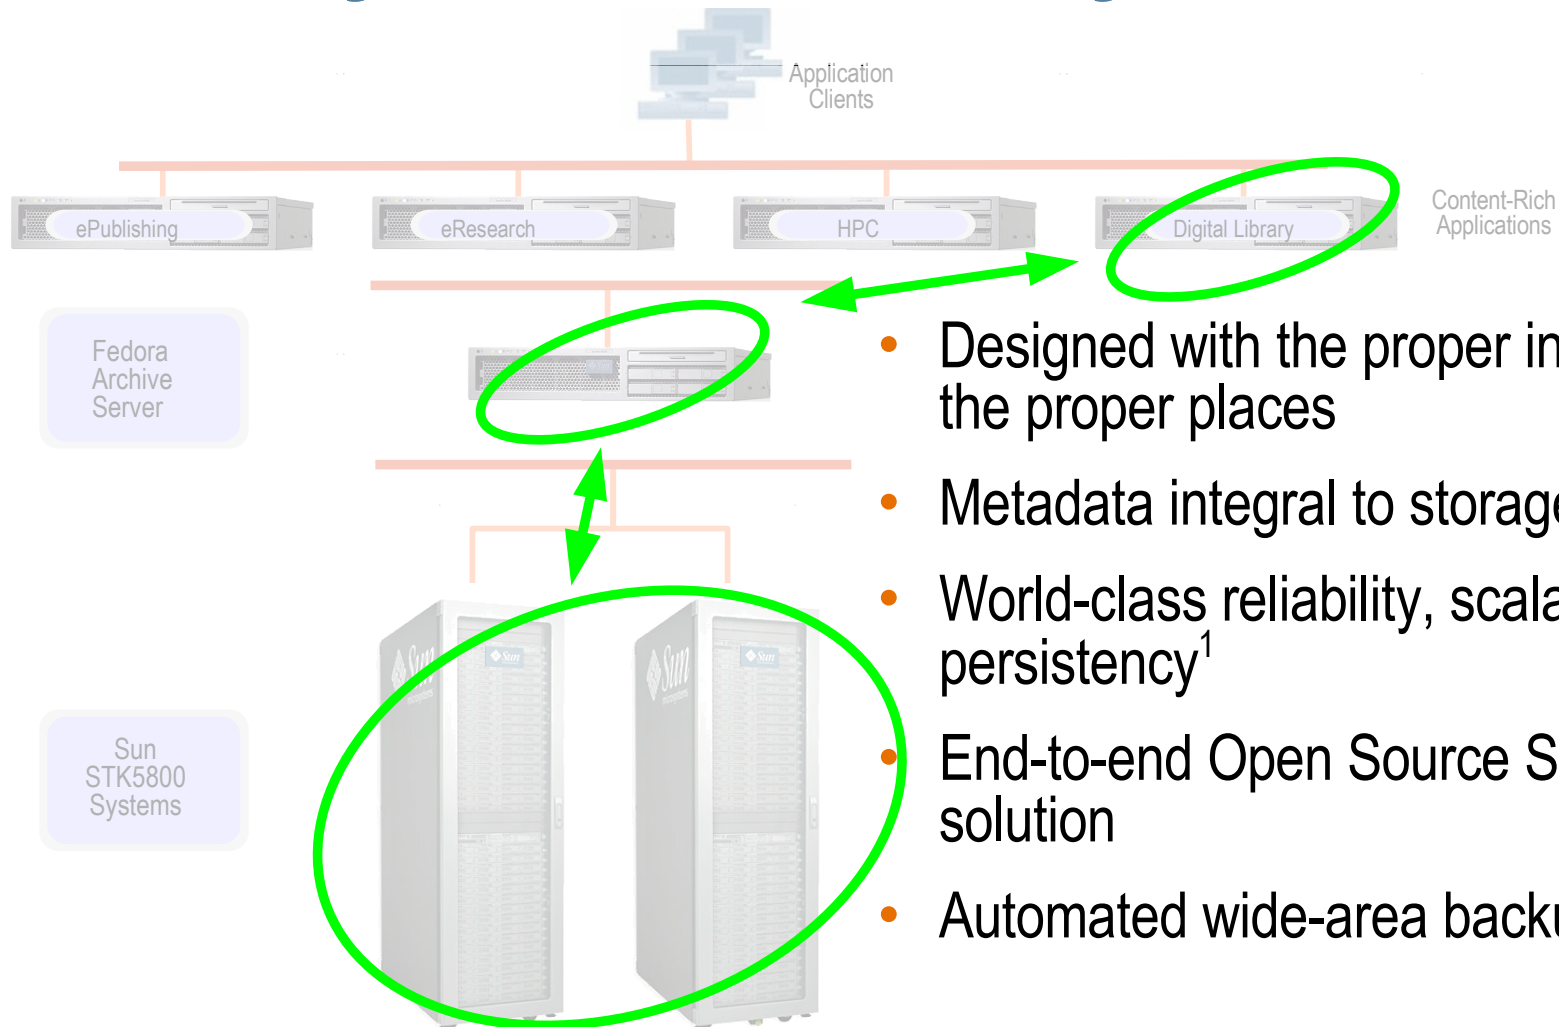
# Why Fedora + Honeycomb?

## Answer: Archival Storage

*In order to seriously address current and future issues of scalability, persistence and flexibility, **this** will likely not meet the requirements, while **this** clearly will...*



# Why Fedora + Honeycomb?



- Designed with the proper intelligence in the proper places
- Metadata integral to storage
- World-class reliability, scalability and persistency<sup>1</sup>
- End-to-end Open Source Software solution
- Automated wide-area backup option

<sup>1</sup> Yes, this really is a word



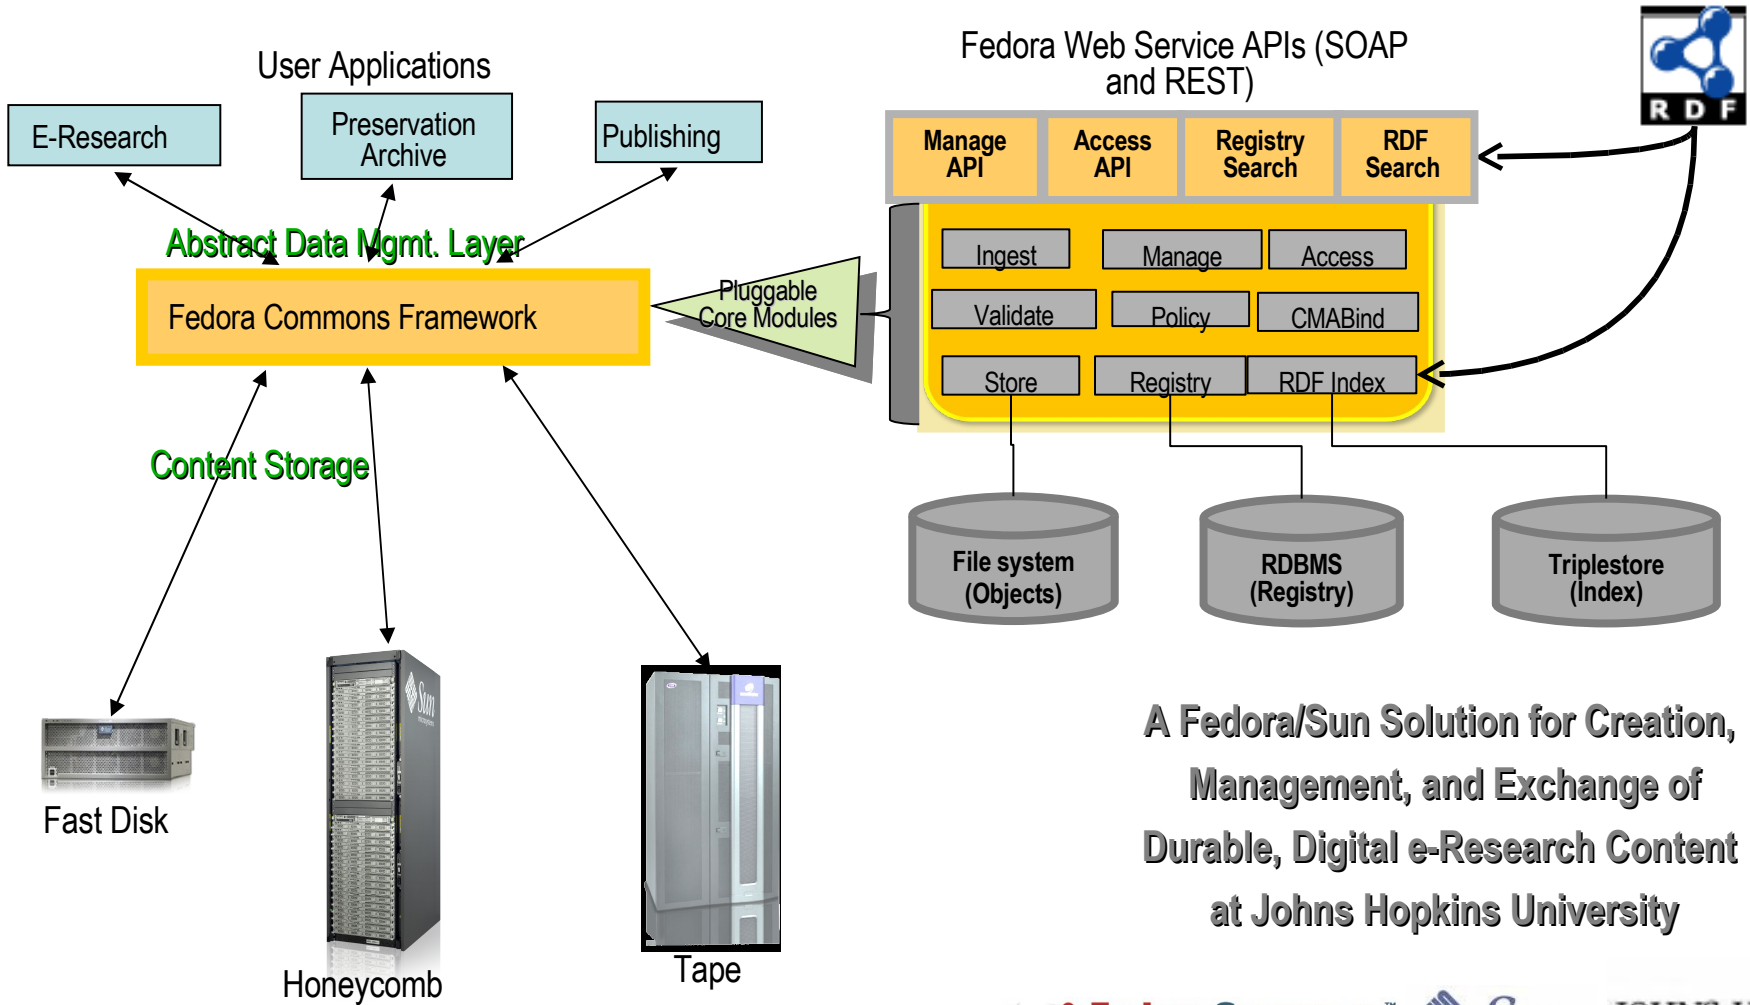
# Storage Beans

- Discrete Services inside Honeycomb
- What will they do? That's up to **you!**
- Asynchronous (Background)
  - > Transformations
  - > Periodic Data Scrubbing
  - > Duplicate Consolidation
- Synchronous (Real-time)
  - > Audit logs
  - > Watermarking
  - > Encryption

# Sun/Fedora Efforts

- Fedora runs now on Solaris/Open Solaris
  - > Server + Storage reference configurations
  - > Inclusion of Fedora 3 in the Open Solaris 'Indiana' Repository
  - > Fedora on Solaris
    - How does it perform?
    - How does it scale?
    - What are the advantages to running on Solaris?
    - Best Practices on Sun

# Proof Point: Fedora/Sun/JHU

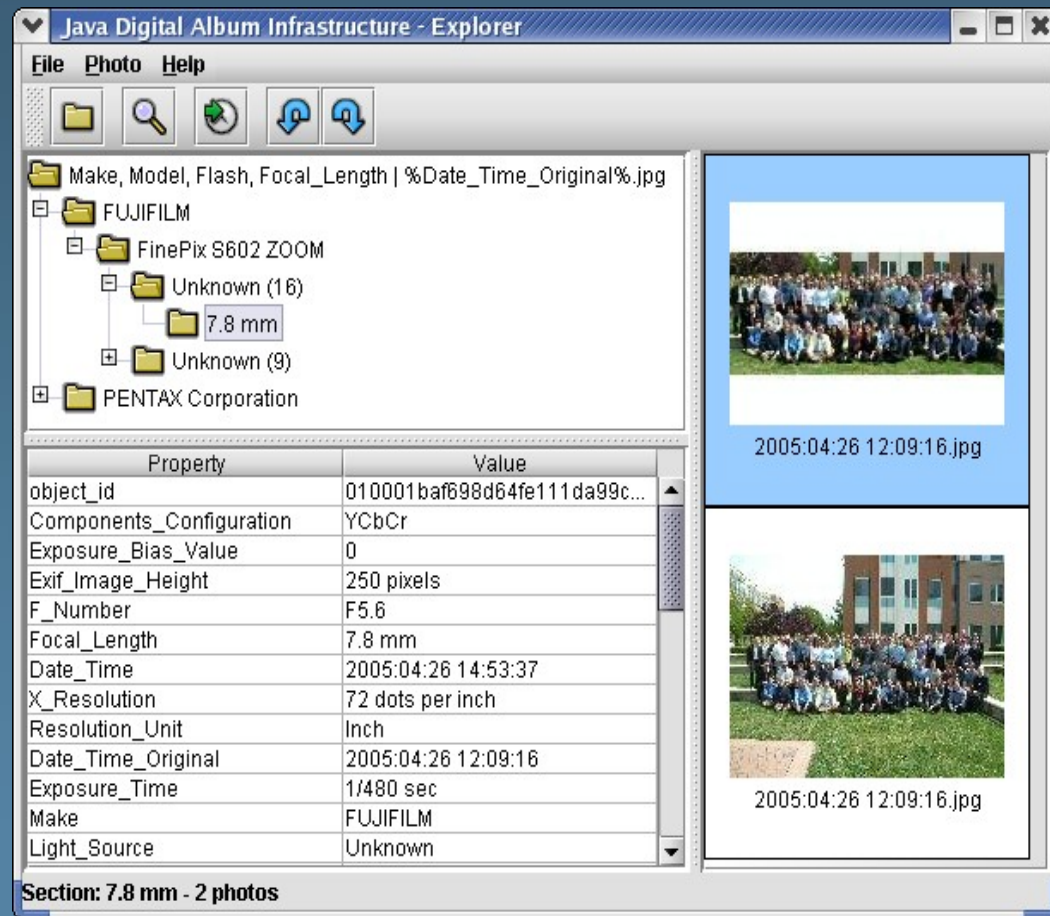


A Fedora/Sun Solution for Creation, Management, and Exchange of Durable, Digital e-Research Content at Johns Hopkins University

# Honeycomb Virtual Views Example

## Photo Application

- Photo demo application on top of StorageTek 5800 or ST5800 Emulator
- Leverages metadata to organize and present the content in logical views
- Photo app extracts, stores, and displays embedded EXIF jpeg metadata



The screenshot shows a Java application window titled "Java Digital Album Infrastructure - Explorer". The interface includes a menu bar (File, Photo, Help), a toolbar with icons for folder, search, refresh, and navigation, and a file tree on the left. The file tree shows a hierarchy: "Make, Model, Flash, Focal\_Length | %Date\_Time\_Original%.jpg" > "FUJIFILM" > "FinePix S602 ZOOM" > "Unknown (16)" > "7.8 mm". Below the tree is a table of EXIF metadata for the selected photo.

Property	Value
object_id	010001baf698d64fe111da99c...
Components_Configuration	YCbCr
Exposure_Bias_Value	0
Exif_Image_Height	250 pixels
F_Number	F5.6
Focal_Length	7.8 mm
Date_Time	2005:04:26 14:53:37
X_Resolution	72 dots per inch
Resolution_Unit	Inch
Date_Time_Original	2005:04:26 12:09:16
Exposure_Time	1/480 sec
Make	FUJIFILM
Light_Source	Unknown

At the bottom of the window, a status bar reads "Section: 7.8 mm - 2 photos". On the right side, two photo thumbnails are displayed, both showing a group of people in front of a building, with the filename "2005:04:26 12:09:16.jpg" below each.

# Digital Repository References

- New York U. Digital Media Management
- Stanford U. OAIS Digital Repository
- Johns Hopkins U. eResearch (Fedora)
- Purdue U. eResearch (Fedora)
- Oxford University Google Project (Fedora/VITAL)
- National Library of New Zealand Digital Preservation (Ex Libris)
- California Digital Library Large Scale Digitization
- Swedish Archive for Sound/Recording Digital Media Management
- Southampton U. EPrints Repository
- San Diego Supercomputer Large Dataset Storage
- U. Michigan D-Space Repository
- The Alberta Library OAIS Digital Repository

# For More Information

- Storage Archive Manager  
[http://www.sun.com/storagetek/management\\_software/data\\_management/sam/index.xml/](http://www.sun.com/storagetek/management_software/data_management/sam/index.xml/)
- Honeycomb  
[http://www.sun.com/storagetek/disk\\_systems/enterprise/5800/index.xml](http://www.sun.com/storagetek/disk_systems/enterprise/5800/index.xml)
- Honeycomb Architecture Document:  
[http://www.sun.com/storagetek/disk\\_systems/enterprise/5800/5800-Arch-final-LR.pdf](http://www.sun.com/storagetek/disk_systems/enterprise/5800/5800-Arch-final-LR.pdf)
- Sun Preservation and Archiving Community  
<http://www.sun-pasig.org>
- Open Source Honeycomb Software  
<http://www.opensolaris.org/os/project/honeycomb/>



# Thank You

# Q&A

