

Repository Redux:

The Past, Present & Future of Fedora

Open Repositories 2013
Charlottetown, Prince Edward Island
Thursday, 11 July 2013

Tom Cramer
@tcramer

A (Funding) History of Fedora

- Original software created 2000–08 with \$2.4M from the Mellon Foundation
 - ⇒ +\$500,000 from UVA Library
- Moore Grant: 2007–11, \$4.9M
- Committers from 10 institutions
- 2009–present: DuraSpace's sponsorship program has provided funding for a tech lead

The Success of Fedora

- Architecture is demonstrably Flexible & Extensible
- Support for Durability
- One foot in the linked data world
- A decade of maturity & proven use
- Substantial community of adopters, contributors, vendors





Looming Storm Clouds

Technical

- Performance
- Fault tolerance
- Scalability
- Complexity
- Aging code base

Organizational

- Small, fixed pool of developers
- Demands outstrip applied resources
- Lack of vibrancy

New Opportunities

危機

- Front-ends: eSciDoc, Hydra, Islandora
 - Attracting new energy and adopters
 - Creating new technical demands
- Evolved technical environment
 - Web architectures & horizontal scaling
 - Linked data
- Data management mandates

Fedora Futures Takes Shape

- OR12: Ad hoc meeting
- September '12:
 - Meet to compare needs and notes
 - Charter a 3 month investigation
- December '12:
 - Commit to a three year project
 - Announce, Invite and Launch



Fedora Futures Objectives

- Preserve the strengths of the architecture and community
- Address the needs for robust and full-featured repository services (that we now understand very well)
- Provide a platform in the repository ecosystem for the next 5-10 years

Technical Requirements

- Highly scalable
- High availability
- Higher performance
- Flexible storage
- Robust auditing, reporting and metrics
- Enhanced fixity and versioning

Requirements, continued

- Work for small, medium and large institutions
 - Easy to deploy, administer
- Support breadth of needs
 - Traditional IR
 - Heterogeneous content (e.g., media)
 - Emerging data management needs
- Interoperate with other systems
 - Lean core, modular, APIs

Organizational Requirements

- Revitalized corps of developers
- Robust community investment and governance
- Bigger community base
 - Geographic
 - Commercial & Non-Profit
 - Additional domains

Fedora Futures Delivers

- January – Feb `13
 - Evaluate platforms for Alpha
- March – June `13
 - Alpha v1 development
- July `13
 - Alpha v1 released!
- 2nd half of `13: Beta development

Fedora 4 Alpha 1 Highlights

- Roughly 80% of Fedora 3.x functionality
 - in 7% of the lines of code
 - with 72% test coverage (vs. 10% for Fedora 3.x)
- Clustering
- Batch operations
- Transaction support
- Policy-driven & projected storage
- Self-healing
- One step install...

Who Has Contributed So Far...

- Brown University
- Center for Disease Control (CDC)
- Columbia University
- Discovery Garden
- DuraSpace
- FIZ Karlsruhe
- Indiana University
- MediaShelf
- Northeastern University
- Oxford University
- Penn State
- Smithsonian Inst.
- Stanford
- UCSD
- UNC – Chapel Hill
- University of Virginia
- University of New South Wales



We Need You...

- Financial contributions
- Developer contributions
- Use cases & priorities
- Integration & testing
- Advocacy & evangelism

Next Steps

1. Donate money
<http://duraspace.org/sponsors>
2. Add a developer
contact awoods@duraspace.org
3. Join the email list
ff-tech@googlegroups.com
4. Install the alpha
github.com/futures/fcrepo4/
5. Chime in!
Give use cases, feedback

