




FedoraCommons™

Fedora UK&I 13 December 2010

Fedora Commons Repository Update


Stephen Bayliss, Acuity Unlimited

Fedora 3.4 Release Manager




Agenda

- Recent releases
- Roadmap
- Upcoming release
- Community participation



Recent releases

- Release 3.4 – August 2010
 - DC, RELS-EXT, RELS-INT as Managed Content
 - REST API relationships methods
 - Enhanced Content Models
 - Optimistic Locking
 - FeSL Authentication as default (JAAS)
 - FeSL policies as Fedora Object
 - Logging reconfiguration without restart
 - Akubra low-level storage
 - REST API improvements and bug fixes
 - Deprecation of "LITE" APIs
- Release 3.4.1 – October 2010
 - Security update – DoS vulnerability




Roadmap – major themes

- London committer meeting February 2010 and Madrid committer meeting July 2010 (OR2010)
- Improving the build and release process
- Enhanced Content Models
- Fedora Command-Line interface and tools
- High Level Storage
- Module Architecture Development
- Improve service deployment
- Improve support for semantic web and linked data
- Consolidation of APIs (API-A/API-M)




Enhanced Content Models

- Extend DS_COMPOSITE_MODEL
 - include XSD references
- ONTOLOGY datastream
 - OWL LITE to declare allowed relationships, cardinality and target “classes”
- Validation (after ingest) against XSD and OWL
- Cloning objects via Fedora API, Template objects
- Repository “views” (for atomistic models)



High-Level Storage

- Pluggable storage interface in Fedora is currently “low level” – storing objects and datastreams as anonymous blobs
- Introduce a high-level storage interface
 - storage layer aware of the structural semantics of the object
 - multiplexing, transactions, locking strategies, indexing, managing
 - greater flexibility in adapting to the needs of preservation, performance, cloud and analysis-oriented use-cases




Module Architecture Development

- Long-term goal is to move to OSGi
 - but, steep learning curve, some way to go yet
- Shorter term – move to the Spring dependency injection framework
 - replacing Fedora’s “home grown” module architecture
- Enabler for many future improvements




Improve Service Deployment

- Reworking of service deployment in the current 3.x CMA
- Current implementation leaves much to be desired, eg current WSDL implementation
- Simpler description language for services
- Support for SOAP bindings
- Support for additional HTTP verbs – PUT, POST, DELETE
- Tooling and documentation improvements




Semantic web / Linked Data

- Extending Fedora's current support for RDF
- Persistent http URIs for resources
 - REST API and Resource Index
- Support for less-restricted RDF
 - current limitations: objects and datastreams as triple subjects
 - eg through named graphs
- REST API support for relationships




API Consolidation

- Two APIs – API-M and API-A, two sets of code
- SOAP splits these, REST does not
- Merge these two APIs
 - current split is cumbersome, particularly for AuthZ



Release 3.5


- Currently at the planning stage
- Tentatively planned for end March 2011
- Release Manager – Aaron Birkland, NSDL, Cornell University



Release 3.5 candidate items

- improve build and release process
- improve repository scalability and performance
- move to a space-per-feature-release model for documentation
- move from SourceForge to GitHub
- move to a modern dependency injection framework (Spring)
- allow Fedora to be deployed as a WAR (without an installer)
- rework client command-line tools
- improve HTTP 1.1 support (content-length and content-range)
- out-of-box SWORD interface
- improve authentication support for externally-managed content
- remove dependency on Berkeley DB XML for FeSL

Not a full or finalised list, this release is still at the planning stage




Other items

- New Committer – Scott Prater, University of Wisconsin - Madison
- “Nightly” builds
 - maintenance 3.4 branch (bug fixes)
 - release artefacts available after each commit
- Mailing list configuration changes
 - reply-to default to the list



Over to you...

- Top ten enhancements, improvements, fixes that you'd like to see
- Areas for improvement – What could we do better?
- Questions




Community participation

- Fedora is a community project, your participation is vital
- tell us what you want, give us feedback
 - mailing lists
 - voting and commenting on JIRA issues
- code contributions
- documentation updates
- testing



Links

- Fedora repository project issue tracker
 - <https://jira.duraspace.org/browse/FCREPO>
- Fedora repository development space
 - <https://wiki.duraspace.org/display/FCREPO>
- Mailing lists
 - <http://www.fedora-commons.org/community/userlist>
 - <http://www.fedora-commons.org/community/developers/devlist>



The Committers

- Benjamin Armintor - Columbia University
- Asger Askov Blekinge - State and University Library, Denmark
- Steve Bayliss - Acuity Unlimited
- Aaron Birkland - NSDL, Cornell University
- Bill Branan - DuraSpace
- Daniel Davis - Cornell University
- Gert Schmeltz Pedersen - Technical University of Denmark
- Scott Prater - University of Wisconsin - Madison
- Edwin Shin - Media Shelf
- Kai Strnad - FIZ Karlsruhe
- Chris Wilper - DuraSpace
- Andrew Woods - DuraSpace