2017-04 Fedora Newsletter

This is the April 2017 edition of the Fedora Newsletter. This newsletter summarizes the most significant activities within the Fedora community over the last month.

Call for Action

Fedora is designed, built, used, and supported by the community. An easy and important way that you can contribute to the effort is by helping resolve outstanding bugs. If you have an interest in gaining a better understanding of the Fedora code base, or a specific interest in any of these bugs, please add a comment to a ticket and we can work together to move your interest forward.

Software development

4.6.2 and 4.7.2 Releases

Earlier this month, two common Java library vulnerabilities surfaced:

- https://issues.apache.org/jira/browse/VELTOOLS-163
- https://web.nvd.nist.gov/view/vuln/detail?vulnId=CVE-2014-0114

Although related libraries have historically been included in Fedora's deployable artifacts, investigation confirmed that Fedora is not subject to the vulnerability; in fact, Fedora is not actually making use of the affected libraries. Because of this, and to eliminate any doubt of a potential vulnerability, the libraries have been removed from the codebase and a patch release has been issued to the 4.7 line and backported to the 4.6 line.

Release notes for Fedora 4.6.2 and 4.7.2 are available on the wiki.

Standards

Fedora API Specification

The Fedora community is working to establish a clearly defined specification for the core Fedora services. This specification details the exact services and interactions required for a server implementation to be verified as "doing Fedora".

Community members raised issues around implementing Batch Atomic Operations in a distributed architecture, leading to a decision to separate this element into its own specification document. We are also seeking further input on supporting different kinds of external content.

The draft specification has been published and is looking for community comments. Following this round of input, we plan on an initial release of the specification in the spring of 2017.

Fedora API Adopters Guide

As the formalization of the Fedora API Specification matures, it will be increasingly important for existing applications and frameworks over Fedora to adjust client/server interactions to the specification. The services defined in the Fedora API Specification are the same ones that are currently provided by Fedora 4, but the interaction models in some cases are changing to be more in line with broader standards. The effort of this group is to facilitate the adjustment of client-side tooling by detailing the "deltas" between the current Fedora 4 implementation and the emerging specification.

Please stay tuned for updates and opportunities for engagement.

Community-driven Activity

Fedora Contributor Profiles

The Fedora project relies on many individuals and institutions to ensure the project's success. We are grateful for their commitment and will showcase their contributions in a series of community profiles aimed at recognizing our contributors' achievements, and introducing them to the rest of the community. The first article in the Fedora Repository Contributor Profile series features Danny Bernstein, who has a long history with DuraSpace technologies and is a new contributor to the Fedora repository team. His focus is on core Fedora performance issues particularly as they relate to development of the new Hyku repository. The full article can be found on the DuraSpace website.

Import/Export Tooling

One of the design goals of Fedora is to simplify the process of both getting your resources into and out of Fedora in a standardized way. This enables the reuse of Fedora resources in other contexts, such as exporting to a separate preservation system, as well as provides a pathway for migrating across Fedora installations. This design goal is being met by an external import/export tool that supports transacting in RDF and (optionally) BagIt bags.

We will be organizing another code sprint to address phase 3 requirements in the near future. If you would like to participate as developer, tester, or to write documentation, please contact Andrew Woods.

Performance and Scale

The Performance and Scale group met on March 20 and discussed progress on the "many members" performance issue, which affects retrieval times for resources with a large number of members. Danny Bernstein pursued a number of optimization strategies and found that a combination of parallel processing streams and increasing the size of the ModeShape cache produced significant performance gains of 5-8x.

Conferences and events

In an attempt to simplify the task of keeping up with Fedora-related meetings and events, a Fedora calendar is available to the community as HTML and iCal .

CNI Spring Meeting and DuraSpace Summit

The CNI Spring membership meeting will take place April 3-4 in Albuquerque, NM. The meeting will feature project briefings and panel discussions, and will be immediately followed by the annual DuraSpace Summit on April 4-5. The state of DuraSpace will be presented highlighting 2016 successes, ideas for new strategic directions, and open source project updates. In addition to the opportunity for small and large group discussions, each project will hold their own meeting, as organized by the respective Project Steering Groups, to discuss projects' goals, objectives, and resources for 2017.

ILIDE

The annual Innovative Library in Digital Era conference will take place April 3-5 in Jasná, Slovakia. The aim of the conference is to present visionary and original ideas based on the extensive experience of the participating experts and institutions. This year's agenda will feature a number of sessions on research data management, including a presentation on stewarding research data with Fedora by David Wilcox.

RDA 9th Plenary Meeting

The 9th RDA Plenary Meeting will take place April 5-7 in Barcelona, Spain. The plenary will feature interest and working group sessions on a number of topics related to repositories and Fedora, including the Research Data Repository Interoperability working group, and the Libraries for Research Data, Long Tail of Research Data, and Repository Platforms for Research Data interest groups.

Hydra Camp

The next Hydra Camp will take place April 17-20 at Emory University in Atlanta, GA. Led by Data Curation Experts, Hydra Camp is open to all developers interested in building skills working with the Hydra technology framework. Please register in advance to attend.

Museums and the Web

Museums and the Web will be hosted in Cleveland, OH from April 19-22. The conference features advanced research and exemplary applications of digital practice for cultural, natural and scientific heritage. This year's conference includes a workshop on Hydra-in-a-box for museums led by David Wilcox and Stefano Cossu. Workshop seats are limited so please register in advance.

IFLA International News Media Conference

From printed newspapers to born-digital news, libraries and other cultural heritage institutions have a central role in ensuring future access to news content. The IFLA News Media Conference, which takes place April 27-28 in Reykjavík, Iceland, will examine issues and challenges in collecting and preserving the news and making it available to users. The conference will include a Fedora workshop delivered by David Wilcox.

Advanced Hydra Camp

Data Curation Experts is offering an Avanced Hydra Camp on May 8-10 in Minneapolis, MN. This camp is for developers with a year or more hands-on Hydra development experience who have already attended Hydra Camp or have equivalent knowledge. Please register in advance to attend.

Islandora Conference

The second Islandoracon will be held in Hamilton, Ontario from May 15 - 19, 2017. The conference schedule will take place over five days, including a day of post-conference sessions and a full-day Hackfest. Please register in advance to attend.