

Improving Fedora 4 to Work with Web-scale Storage and Services

The Digital Curation Innovation Center (DCIC) at the University of Maryland's College of Information Studies (iSchool) requests \$240,138 to research ways to improve the performance and scalability of Fedora 4 for the Fedora user community.

1. Statement of National Need

This project addresses high priority challenges that face the Fedora user community. The National Digital Platform is interested in expanding the scalability and interoperability of major open source software used by libraries and archives. Our project aims at identifying current storage limitations in Fedora 4 and researching and prototyping ways of improving its use when managing web-scale collections.

Fedora has consistently provided an object storage framework for repositories of content and data for decades. Fedora developers and users have created hundreds of applications, a strong community, and a rich software ecosystem of shareable code. The Fedora 4 project worked to redefine Fedora's conventions, the application programming interface (API) and the data model, in the age of linked data. The storage demands of Fedora users for objects and metadata are extreme when compared to typical content management applications. Fedora users store large data files in their repositories, unlike the typical media objects used in publishing. They may also store millions of objects with extensive metadata attached to each. Despite this, the Fedora 4 software implementation is limited by an underlying software component that cannot meet these demands. Importantly, Fedora 4's storage layer is limited to a single large storage pool and this is a performance bottleneck. The single storage pool also limits the total storage capacity of any repository. Institutions adjust their budgets to purchase expensive high-capacity storage arrays, so that their single storage pool can grow as large as possible. Fedora 4's storage pool is a single point of failure that impacts reliability and remains the limiting factor on repository growth for the Fedora community.

2. Project Design

The Fedora user community's needs for storage are unique among content management applications, but not unique among web-scale companies. Storing millions of objects and their metadata is not unlike storing the data and preferences for millions of users. Large companies like Netflix, Facebook, and Google have consistently succeeded in meeting this challenge by adopting a different strategy for storage. The Fedora 4 strategy for growth is one where you plan, budget, build, and finally migrate into a single storage system that is large enough to meet expected demand. In contrast, the web-scale companies meet demand by adding more, separate storage systems to a group of similar systems, a cluster, that work together to store the data. Each system in the cluster is relatively small and cheap, so you no longer have to plan and build expensive storage arrays. A storage cluster can grow very large while maintaining a smooth, predictable cost for increases in capacity. The performance of a storage cluster is not limited by a single bottleneck system.

Most Fedora institutions have trouble estimating costs and planning capacity for a large new collection. When current storage is nearly full, adding new collections may be postponed for years, until the next, higher capacity, storage system can be designed, budgeted, and built. These planning cycles make it very difficult to estimate costs for collections. In contrast, using a cluster approach allows us to simply count the number of new servers required to store the collection. The costs are simple and servers may be purchased and added to the existing

cluster as needed. The current Fedora strategy can have a serious impact on collection development, budgeting, and planning long before it has a technical impact on performance and total capacity.

The DCIC is in a unique position to deliver a cluster strategy for the Fedora community. We have been involved in developing a new cluster-based archival repository, called DRAS-TIC (Digital Repository At Scale - That Invites Computation). DRAS-TIC implements an open-source industry API standard called the Cloud Data Management Interface (CDMI) that defines the way in which applications can create, retrieve, update, and delete datasets and metadata in the cloud. It demonstrates dramatic improvements in performance and capacity without the liabilities of expensive storage, capacity planning, and data loss that have plagued non-cluster storage strategies. We currently manage 100 Million files with metadata. This project will work with the Fedora community to develop an alternate implementation of the Fedora API that addresses a very real shortcoming in the existing implementation, towards meeting the most demanding and extreme storage requirements.

Richard Marciano, professor, will act as project director. Marciano has been the Principal Investigator for grants from IMLS, NSF, the Library of Congress, and the National. Greg Jansen is principal research software architect, a contributor to the Fedora 4 Technical Team, and creator of the Curator's Workbench.

We propose the following work plan:

1. Continue to work with Fedora developers and the steering group to ensure that this project aligns with Fedora conventions and the needs outlined above (*we have conferred with the Fedora Steering Group*).
2. Create a prototype of the Fedora API running on DRAS-TIC (*DRAS-TIC-as-Fedora*) by implementing create, read, update and delete operations.
3. Create a virtual machine for development and evaluation by the Fedora community at large.
4. Run Fedora's API-based test suites to measure API compliance and performance. Produce metrics that can be compared with the Fedora 4 implementation.
5. Implement additional operations from the Fedora API, suitable for a full repository.
6. Develop a how-to guide for running a DRAS-TIC-as-Fedora repository, including:
 - a. How to create and manage the storage cluster.
 - b. How to manage backup and recovery for Fedora data.
7. Reach out to the broader Fedora community for adoption and support and present at DLF, CNI, and Open Repositories conferences.

3. National Impact

The proposed project will generate open-source software, documentation, and best-practice guides on how to manage Fedora 4 with large-scale collections and thus contribute to big-data ready national software infrastructure. Jansen and Marciano will also participate in the IMLS-funded "Always Already Computational: Library Collections as Data" project.

Budget

The team requests from IMLS a total budget of \$240,138 over the 2-year period (Oct. 1, 2017 to Sep. 30, 2019) without cost sharing. This includes \$157,986 in direct costs and \$82,153 in indirect costs, calculated at UMD's negotiated rate. The direct cost is further broken down into \$120,118 for salaries and wages, \$27,868 for fringe benefits, \$10,000 for travel.