

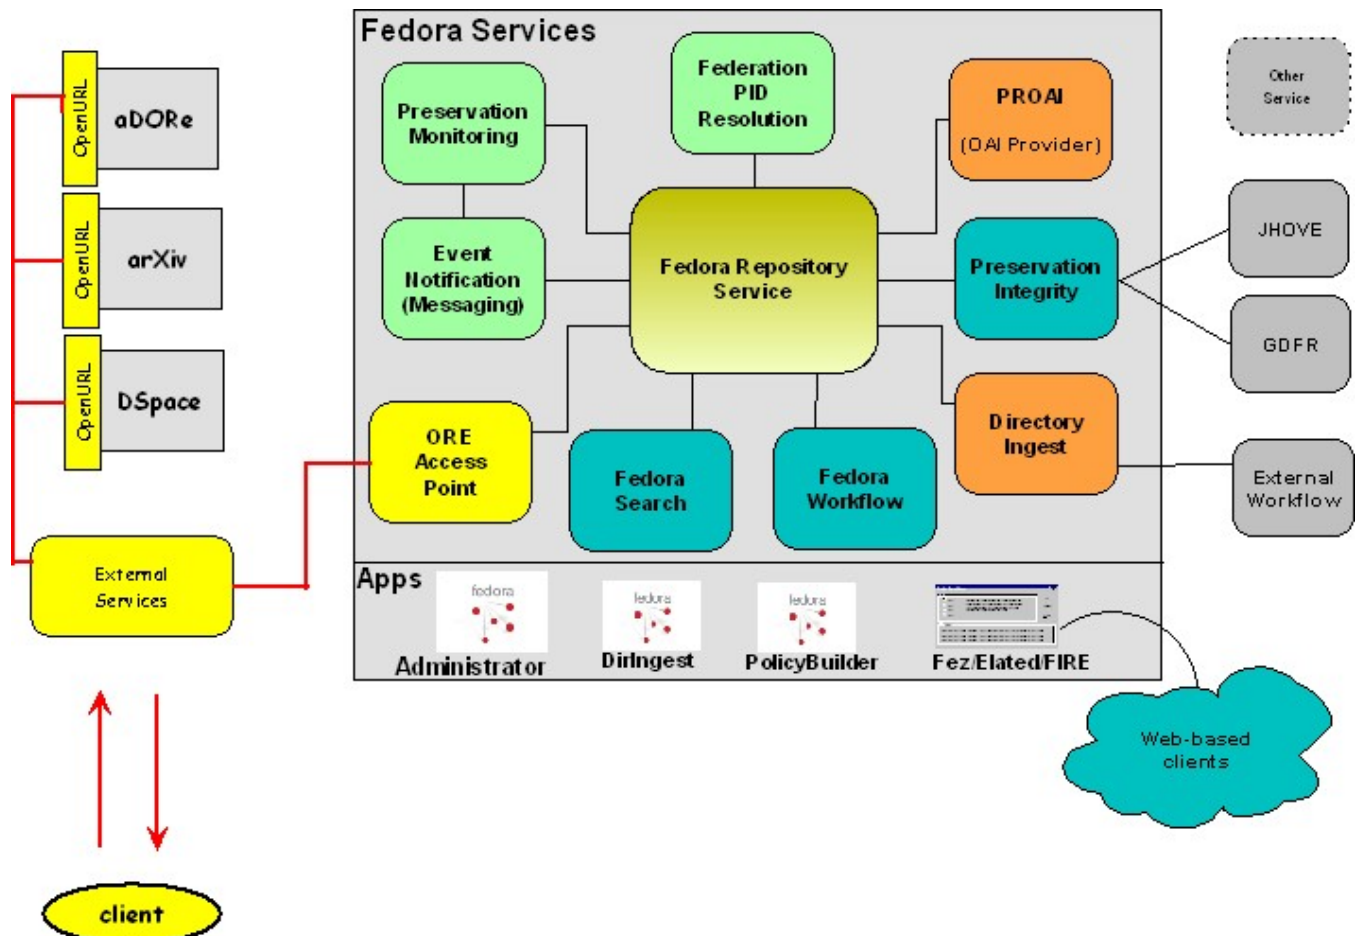
Service Framework

Introduction

As of Fedora 2.1, the Fedora Service Framework was introduced to facilitate the integration of new services with the Fedora repository. The framework takes a service-oriented architecture approach to building new functionality around a Fedora repository. While the Fedora repository, itself, exposes its functionality as a set of web service interfaces, all of these interfaces belong to the Fedora web application that runs in its own Tomcat. The new Fedora Service Framework allows new services to be built around the core repository - as stand-alone web applications that run independently of the Fedora repository. While Fedora repository functionality can still be extended with new modules, the intent is to keep the repository service focused on the core functions of a repository. Yet, there are many other services that are beneficial companions to a repository, such as specialized ingest services, workflow services, preservation services, and many others. These are the kinds of services that the framework is intended to support. There are two main benefits to the service framework approach: (1) it allows new functionality to be added as atomic, modular services that can interact with Fedora repositories, yet not be part of the repository, (2) it makes co-development of new services for Fedora easier since each service can be independently developed and plugged into the framework. As of Fedora 2.1, the Fedora development team has released an initial set of services ([Directory Ingest](#) and [OAI Provider](#) described below), and will continue to develop new services over the course of Fedora Phase 2 (2005-2007) and beyond, especially services for workflow and preservation. Services that are part of the framework will be packaged as part of the Fedora open-source software distribution and will be kept up to date with new versions of the core Fedora repository service. Members of the Fedora community will be collaborating on the development of services and will contributed back to the Fedora Project. Further documentation will be provided to establish guidelines on how services should be designed to effectively plug into the framework. In the mean time, developers of new services can follow the design patterns of the Directory Ingest and OAI Provider services.

Framework and Services

The Fedora Service Framework establishes a means for coupling new services with the core Fedora repository service. The framework allows for the creation of atomic, modular services that can interact with the Fedora repository or each other. The diagram below depicts the Fedora Service Framework as it is envisioned to evolve during 2005-2007. The repository service was the first deliverable of the Fedora Project (2002-present). In 2006, the next two services were introduced with Fedora 2.1: the OAI Provider Service and the Directory Ingest services. In January 2007, the Fedora Search Service is being deployed as part of the framework. The Search Service (known as GSearch) was contributed to the Fedora Project by community member Gert Schmeltz Pedersen of the Danish Technical University. During the rest of Phase 2 of the Fedora Project, both the Fedora development team, and the Fedora community will develop the other services to fit into the framework. In 2007, development will begin on the Fedora Preservation, Event/Messaging, and Workflow services. During this phase of development we will move Fedora towards an "enterprise" architectural pattern.



The Fedora Service Framework can evolve to include new services conceived of by the Fedora community. Listed below is a brief description of each service, links to specifications when available, and status.

- **Repository Service:** at the heart of Fedora is the repository service that enables the creation, management, storage, access, and reuse of digital objects.
 - **OAI Provider Service:** a configurable OAI Provider service for harvesting metadata out of a Fedora repository via OAI-PMH. The service can be configured to harvest any type Datastream or dissemination from objects in the repository. It also supports OAI sets.
 - **Directory Ingest Service:** a service to ingest a hierarchical directory of files into a Fedora repository. The service will accept a Submission Information Package (SIP) in the form of a .zip archive that contains directories of files along with a METS-based manifest file that describes the directory hierarchy. The default parent-child relationships that characterize a directory hierarchy can be overridden and refined to have other semantic meaning (e.g., collection-member, folder-document). Upon receipt of the SIP, the Directory Ingest service will process the .zip archive and create a Fedora digital object in the repository for every file and every directory, plus it will record the relationships among them in Fedora's RDF-based relationships datastream. *A new web-based client for creating the SIP is now available (see [SIP Creator](#)). This client will enable a user select files from a file system, add metadata about files, and assert semantically-meaningful relationships, and ultimately submit the SIP to a Fedora repository.*
 - **Search Service:** a configurable search service that can index any datastream or dissemination of Fedora digital objects. The service is pluggable, and will provide adapters for Lucene and Zebra as the default backend search engines. Other search adapters can be developed to plug into other engines.
 - **Workflow and Orchestration Service:** (planned 2007, under specification by Fedora Workflow Working Group and Fedora Development Team)
 - **Preservation Integrity Service:** (planned 2007, currently under specification by Fedora Preservation Working Group)
 - **Preservation Monitoring and Alerting Services:** planned 2007, currently under specification by Fedora Preservation Working Group)
 - **Event Notification Service (Messaging):** (planned 2007)
 - **Persistent Identifier Resolution Service**
 - **Object Reuse and Exchange (ORE) Access Point:** an interface to a Fedora repository to facilitate cross-repository interoperability (2007-2008)
- The Fedora Service Framework provides building blocks for higher-level customized services and user applications.

Core Repository Service

Fedora Core Repository Service can be run as a stand-alone service, or it can be situated within the Fedora Service Framework, where a suite of companion services can be loosely coupled with the repository to provide additional functionality for integrating the repository into a broader service-oriented architecture (SOA) pattern. The core repository can be accessed via web service interfaces to its core functionality. The core repository service actually has several web service APIs: an interface for repository management (API-M); an interface for repository access (API-A); interface for basic repository search; and an interface for RDF-based search of the Resource Index. All of these web service interfaces are available on the Fedora repository server web application that runs in Tomcat. The repository service is built in a modular manner, so that each inner function is implemented as a java-based module. The inner modules are configurable, and they can be replaced with alternate implementations.

The Fedora repository service is the core service in the Fedora Service Framework, and was depicted in the above diagram in the center of all other services. Below, the Fedora repository service is depicted in more detail, with its inner modules exposed, and all repository interfaces. The diagram depicts the repository service from the perspective of how it maps to the [Open Archival Information System \(OAIS\)](#) reference model which has been approved as an ISO standard.

