

# Features

## Authorization

Fedora 4 authorization is designed to be fine grained, while at the same time manageable by administrators and end users. Authentication is managed by the servlet container or externally, but authorization mechanisms are open to extension and many reference implementations are included. Roles-based access control is an included feature that makes permissions more manageable and at the same time easier for external applications to retrieve, index and enforce. Finer grained security checks have no

## Backup and Restore

The Fedora 4 Backup capability allows a user, such as the repository manager, make a REST call to have the repository binaries and metadata exported to the local file system. Inversely, the Restore capability allows a user to make a REST call to have the repository restored from the contents of a previous Backup operation. In addition, with the default configuration, files are stored on disk named according to their SHA1 digest, so a filesystem backup approach can also be used.

## External Search

To support the differing needs for sophisticated, rich searching, Fedora 4 comes with a standard mechanism and integration point for indexing content in an external service. This could be a general search service such as Apache Solr or a standalone triplestore such as Sesame or Fuseki.

## External Triplestore

RDF support is a core feature of Fedora 4, used as the primary data format for the REST API. A triplestore is not bundled into the repository itself. Instead, Fedora 4 sends events when the repository is updated, and the Indexer copies RDF from the repository to an external triplestore to keep it in sync with the repository. This pattern, which is also used for search functionality for the same reasons, allows maximum flexibility about what triplestore to use, and removes the overhead of keep

## RESTful HTTP API

The Fedora 4 HTTP API is generally a RESTful API. HTTP methods like GET, PUT, POST and DELETE are implemented on most resource paths. The API also relies heavily on content negotiation to deliver context-appropriate responses, and a HATEOAS-driven text/html response (providing a decent GUI experience on top of the repository).

## Transactions

Fedora 4 supports the ability to wrap multiple REST API calls into a single transaction that can be committed or rolled back as an atomic operation.

## Versioning

Within Fedora 4, snapshots of the current state of a resource may be saved into the version history. The RDF for historic version snapshots may be browsed and old non-RDF content may be downloaded. Furthermore, an object or subgraph may be reverted to the state that it existed in a historic version.

## Clustering

To support horizontal scalability use cases as well as geographic distribution, Fedora 4 can be configured as a cluster of application servers.

## Identifiers

Identifiers can be specified in REST API calls and generated either automatically using the internal PID minter or via an external REST service.

## Linked Data Platform

## External Content