

Glossary

Fedora 5 is a [Linked Data Platform server](#) as defined by the [Linked Data Platform 1.0 specification](#). Many of the terms in this glossary are drawn directly from this specification.

- [Access Role](#)
- [Binary](#)
- [Checksum](#)
- [Children](#)
- [Container](#)
- [Datastream](#)
- [Dissemination](#)
- [Federation \(aka Projection\)](#)
- [Fixity](#)
- [Indexer](#)
- [Linked Data Platform \(LDP\)](#)
- [Managed External Content](#)
- [Namespace](#)
- [NonRdfSourceDescription \(formerly the properties of a Fedora 3 Datastream\)](#)
- [Policy Enforcement Point \(PEP\)](#)
- [Predicate](#)
- [Prefix](#)
- [Projection \(aka: Federation\)](#)
- [Properties](#)
- [rdf:type](#)
- [Resource](#)
- [Tombstone](#)
- [Transactions](#)
- [Triple](#)
- [Uniform resource identifier \(URI\)](#)
- [Universally Unique Identifier \(UUID\)](#)
- [Version](#)

Access Role

A named role, such as "writer", assigned to a user, a group, or some other identifying agent within the repository. Roles may be used by the [policy enforcement point \(PEP\)](#) to authorize actions taken in the repository.

Binary

A binary in a Fedora repository is a [nonRdfSource resource](#). In other contexts, binaries would sometimes be described as bitstreams or files. Binaries are always accompanied by a [nonRdfSourceDescription](#).

Checksum

A computed fingerprint for binary content, used to ensure a complete transfer or the [fixity](#) of stored information. Fedora supports the SHA-1 checksum algorithm by default, but can also support SHA-256 and MD5.

Children

The resources that are immediate children of a given [container](#). [Properties](#) are not children.

Container

A container is a [resource](#) that represents intellectual entities and can also be used to aggregate other [resources](#) in a Fedora repository. Containers may contain other containers or [binaries](#) and their [nonRdfSourceDescriptions](#).

Datastream

Datastream is principally a Fedora 3 concept. Within the Fedora 4 model, the closest equivalent to a datastream is a [binary](#) and its [NonRdfSourceDescription](#). Treat both together as a datastream; they only ever occur in a pair.

Dissemination

An approach available in Fedora 3 that allowed for binding a service from outside the repository to a content model or specific content type. Note: *disseminators* as such do not exist in Fedora 4 itself, but the [API-X](#) effort provides similar functionality by extending the core with specifications and optional machinery.

Federation (aka Projection)

See: [Projection](#)

Fixity

Fixity is a measurement or characterization of stored information at a given time, a reference to the stored information being 'fixed' or otherwise unchanged. Comparing two values of fixity generated at different times, such as comparing two checksums for a given binary file with one having been generated before ingest to a repository and the second after ingest, provides a means to determine whether the information has changed.

Indexer

Creating, modifying or deleting [resources](#) in the repository generates events. An (optional) indexer monitors and processes these events, such as loading RDF from Fedora to an external triplestore, for example.

Linked Data Platform (LDP)

The W3C [Linked Data Platform](#) (LDP) specification describes a set of best practices and a simple approach for a read-write Linked Data architecture, based on HTTP access to web resources that describe their state using RDF as the data model. Fedora 5 implements the LDP specification for create, read, update and delete (CRUD), allowing HTTP, REST, and linked data clients to interact with Fedora 5.

Managed External Content

Managed External Content is content that resides outside of the Fedora 5 repository, but is managed by Fedora through a [projection](#). Managed external content can be accessible via the Fedora API as well as have Fedora-managed audits.

Namespace

A namespace is a container for a set of identifiers (also known as symbols, names). In Fedora 5, [resource](#) properties may belong to any namespace. Namespaces are restricted to being an empty string or to a [URI](#) as defined in [section 3](#) of RFC3986. An example of a namespace would be "<http://purl.org/dc/elements/1.1/>".

NonRdfSourceDescription (formerly the properties of a Fedora 3 [Datastream](#))

A `nonRdfSourceDescription` is a [resource](#) that describes a [binary](#) resource within the Fedora repository. A `nonRdfSourceDescription` is always accompanied by a [binary](#), and vice versa.

Policy Enforcement Point (PEP)

This is a pluggable component in the Fedora framework that is responsible for authorizing all actions taken on [resources](#).

Predicate

A predicate expresses a relationship between the subject and the object of a [triple](#).

Prefix

Typically a short string representation of an associated [namespace](#). For example the prefix "dc" is frequently used to represent the namespace "<http://purl.org/dc/elements/1.1/>". When using defined namespace prefixes, it is possible, for example, to use "dc:contributor" versus the full form "<http://purl.org/dc/elements/1.1/contributor>" to reference the Dublin Core "contributor" element.

Prefixes are limited to the following grammar:

```
LocalName ::= ValidString - SelfOrParent
/* Any ValidString except SelfOrParent */
SelfOrParent ::= '.' | '..'
ValidString ::= ValidChar {ValidChar}
ValidChar ::= XmlChar - InvalidChar
/* Any XmlChar except InvalidChar */
InvalidChar ::= '/' | ':' | '[' | ']' | '|' | '*'
XmlChar ::= /* Any character that matches the Char production at http://www.w3.org/TR/xml/#NT-Char */
```

Projection (aka: Federation)

The process by which a repository may present [resources](#) through the API that are actually stored in a different system, such as a file system or database.

Properties

Properties are name/value pairs that belong to [resources](#). The name of a property can be any term from a [namespaced](#) vocabulary. When [RDF](#) is generated in response to a request for a [resource](#) that contains properties, the [RDF](#) will contain [triples](#) for each property where the subject of the [triple](#) is the resource itself, the [predicate](#) of the [triple](#) is the property name, and the object of the triple is the value of the property. Property values can be of any valid [rdf:type](#).

[rdf:type](#)

[Containers](#) are defined by one or more [rdf:types](#) that describe the nature of the [container](#). Similarly, every [property](#) will be an instance of a single [rdf:type](#). The official definition can be found [here](#).

Resource

Resources are the primary organizational structure in the repository. A resource is any web-addressable entity, such as a [container](#), a [nonRdfSourceDescription](#), or a [binary](#). Every resource is uniquely identified, its identifier representing a repository path. Resources are comprised of zero or more properties and/or child resources.

Tombstone

A tombstone is a repository marker indicating that a [container](#), a [nonRdfSourceDescription](#), or a [binary](#) used to exist at a given URL. A tombstone is created when a resource is deleted or moved.

Transactions

A transaction represents a series of changes to the repository that must execute successfully and completely or not at all. Transactions are client initiated and should be used to ensure consistency. Each transaction succeeds or fails as a complete unit; it cannot remain in an intermediate state.

Triple

A triple is a fundamental building block of RDF. It consists of: a subject, [predicate](#), and an object. In this way, a triple can describe a relationship (via the [predicate](#)) of the subject to the object. The official definition can be found [here](#).

Uniform resource identifier (URI)

A string of characters used to uniquely identify a [resource](#). It is defined in [RFC3986, section 3](#). An example of a URI would be "<http://id.loc.gov/authorities/subjects/sh2002000192>".

Universally Unique Identifier (UUID)

A universally unique identifier is a 36 hexadecimal number that is, for all practical purposes, unique (though there is an extraordinarily slight chance of a duplicate identifier being created in the future or already existing). It is used to identify a [resource](#).

Version

A snapshot of a [resource](#) that is saved in version history for later access.