


Internal Database Tables

Fedora uses several database tables and indexes to provide fast retrieval of some information. This page documents the various tables Fedora creates. Due to differences in databases, some datatypes might not match the implemented type in your database exactly (ie. datetime versus timestamp).

When a Fedora repository is started with an empty database it begins an indexing process to create these indexes from the content on disk. You can force the rebuilding of these internal indexes by setting the `fcrepo.rebuild.on.start` JVM property to `true`. For more information again visit the [Properties](#) page.

For the full list of database properties, please refer to the **Database** section of the [Properties](#) documentation.

 You can find additional information on this wiki about how Fedora [rebuilds these caches](#) and [configuring supported databases](#).

Technical Information

Containment Index

The containment index is a series of tables/indexes used to track the `ldp:contains` relationships of containers in Fedora.

containment

The table of committed containment relationships. Accessible to all requests.

Field	Type	Constraints	Description
fedora_id	varchar(503)	NOT NULL, PRIMARY KEY	The internal ID of the resource
parent	varchar(503)	NOT NULL	The internal ID of the containing resource
start_time	datetime	NOT NULL	The datetime when the resource was created
end_time	datetime		The datetime when the resource was deleted (but still has a tombstone)
updated	datetime		The datetime when any contained resources were updated (used for eTags)

containment_transactions

The table of containment relationship operations that not been committed to the **containment** table yet. These are only accessible to the indicated transaction.

Field	Type	Constraints	Description
fedora_id	varchar(503)	NOT NULL	The internal ID of the resource
parent	varchar(503)	NOT NULL	The internal ID of the containing resource
start_time	datetime		The datetime when the resource was created
end_time	datetime		The datetime when the resource was deleted (but still has a tombstone)
transaction_id	varchar(255)	NOT NULL	The transaction id this record belongs to
operation	varchar(10)	NOT NULL	The type of operation this is, one of 'add', 'delete' or 'purge'

Membership Index

The membership index contains relationships generated by `ldp:DirectContainer` and `ldp:IndirectContainer` resources.

membership

The main table of membership relationships. Accessible to all requests.

Field	Type	Constraints	Description
subject_id	varchar(503)	NOT NULL	The internal ID of the resource initiating the relationship (ie. the subject)
property	varchar(503)	NOT NULL	The predicate of the relationship
object_id	varchar(503)	NOT NULL	The internal ID of the resource which is the target of the relationship (ie. the object)
source_id	varchar(503)	NOT NULL	The internal ID of the direct or indirect container containing the resource that created the relationship.
proxy_id	varchar(503)	NOT NULL	The internal ID of the resource that created the relationship. This is contained within the <code>source_id</code>
start_time	datetime		The datetime the relationship started.

end_time	datetime		The datetime the relationship ended.
----------	----------	--	--------------------------------------

membership_tx_operations

The table of membership relationships operations not yet committed to the **membership** table. These are only accessible to the specified transaction.

Field	Type	Constraints	Description
subject_id	varchar(503)	NOT NULL	The internal ID of the resource initiating the relationship (ie. the subject)
property	varchar(503)	NOT NULL	The predicate of the relationship
object_id	varchar(503)	NOT NULL	The internal ID of the resource which is the target of the relationship (ie. the object)
source_id	varchar(503)	NOT NULL	The internal ID of the direct or indirect container containing the resource that created the relationship.
proxy_id	varchar(503)	NOT NULL	The internal ID of the resource that created the relationship. This is contained within the source_id
start_time	datetime		The datetime the relationship started.
end_time	datetime		The datetime the relationship ended.
last_updated	datetime		The datetime the relationship was last updated. (used for eTags)
tx_id	varchar(36)	NOT NULL	The transaction id this record belongs to
operation	varchar(10)	NOT NULL	The type of operation this is, one of 'add' or 'delete'
force_flag	varchar(10)		Indicates that a 'delete' operation should purge an entry from the membership index completely, rather than assigning it and end_time. Used for on-demand versioning when the membership relation existed because of uncommitted state in a mutable head.

Inbound Reference Index

The inbound reference tables hold relationships between two resources in the repository, due to there being no single subject restriction it is possible for a 3rd resource to generate an inbound reference from a 1st resource to a 2nd resource. This relationship will persist even if the 2nd resource was deleted.

reference

The main source of inbound/outbound referenced information. Accessible to all requests.

Field	Type	Constraints	Description
fedora_id	varchar(503)	NOT NULL	The internal ID of the resource creating the relationship
subject_id	varchar(503)	NOT NULL	The internal ID of the resource which is the subject of the relationship.
property	varchar(503)	NOT NULL	The predicate of the relationship.
object_id	varchar(503)	NOT NULL	The internal ID of the resource which is the object of the relationship.

reference_transaction_operations

The reference operations being performed in a transaction. Only accessible to the same transaction.

Field	Type	Constraints	Description
fedora_id	varchar(503)	NOT NULL	The internal ID of the resource creating the relationship
subject_id	varchar(503)	NOT NULL	The internal ID of the resource which is the subject of the relationship.
property	varchar(503)	NOT NULL	The predicate of the relationship.
object_id	varchar(503)	NOT NULL	The internal ID of the resource which is the object of the relationship.
transaction_id	varchar(255)	NOT NULL	The transaction ID this relationship operation is part of
operation	varchar(10)	NOT NULL	The operation being performed. One of 'add' or 'delete'

OCFL Index

The OCFL index provides a mapping between a Fedora URI its base URI (which is different in the case of an Archival Group) and its OCFL ID.

ocfl_id_map

The main source of OCFL mappings. Accessible to all requests.

Field	Type	Constraints	Description
fedora_id	<code>varchar(503)</code>	NOT NULL, PRIMARY KEY	The internal ID of the resource creating the relationship
fedora_root_id	<code>varchar(503)</code>	NOT NULL	The internal ID of the archival group of the resource in fedora_id. These fields are the same for self contained resources.
ocfl_id	<code>varchar(503)</code>	NOT NULL	The OCFL id of the resource containing this Fedora resource.

ocfl_id_map_session_operations

The OCFL ID mapping operations being persisted in a transaction.

Field	Type	Constraints	Description
fedora_id	<code>varchar(503)</code>	NOT NULL, PRIMARY KEY	The internal ID of the resource creating the relationship
fedora_root_id	<code>varchar(503)</code>	NOT NULL	The internal ID of the archival group of the resource in fedora_id. These fields are the same for self contained resources.
ocfl_id	<code>varchar(503)</code>	NOT NULL	The OCFL id of the resource containing this Fedora resource.
session_id	<code>varchar(255)</code>	NOT NULL	The session ID this operation is from.
operation	<code>varchar(10)</code>	NOT NULL	The type of operation to be persisted. One of 'add' or 'delete'

Simple Search Indexes

The simple search indexes provide support for the various search features now available in Fedora.

simple_search

This index contains basic information about all the resources in the repository.

Field	Type	Constraints	Description
id	<code>bigint</code>	PRIMARY KEY, AUTO INCREMENT	A unique numeric id for this resource.
fedora_id	<code>varchar(503)</code>	NOT NULL, UNIQUE KEY	The internal ID of the resource.
created	<code>datetime</code>	NOT NULL	The datetime the resource was created.
modified	<code>datetime</code>	NOT NULL	The datetime the resource was last modified.
content_size	<code>bigint</code>		The size of the content.
mime_type	<code>varchar(255)</code>		The mime-type of the resource

search_rdf_type

This index contains rdf types that exist in the repository.

Field	Type	Constraints	Description
id	<code>bigint</code>	PRIMARY KEY, AUTO INCREMENT	A unique numeric id for this rdf type.
rdf_type_uri	<code>varchar(255)</code>	NOT NULL, UNIQUE KEY	The URI of the RDF type

search_resource_rdf_type

This index links resources from the simple_search table to rdf types in the search_rdf_types table.

Field	Type	Constraints	Description
resource_id	bigint	NOT NULL, FOREIGN KEY REFERENCES simple_search (id)	ID from the simple_search table
rdf_type_id	bigint	NOT NULL, FOREIGN KEY REFERENCES rdf_type_uri (id)	ID from the search_rdf_type table