

Design - Simple, Synchronous Query in Fedora

The lack of a simple synchronous search within Fedora eliminates the possibility of certain kinds of interaction patterns deemed desirable by some members of the community. On the other hand, having extremely sophisticated and comprehensive search support built into Fedora pushes the application towards the "monolithic" design pattern that negatively impacts the ability to optimize for performance and scale considerations. By making simple synchronous query support an optional specified extension to Fedora, with a lightweight implementation in the [community Fedora implementation](#) we can support these alternate interaction patterns without imposing undue burdens on future Fedora implementations or significant performance impacts on the community implementation.

Desired Features

Query Example	Purpose Served	Notes
Find the item(s) with a given dc:identifier.	When performing a migration or large ingest operation that may be halted or interrupted, being able to identify whether a particular item was ingested without storing information outside of the repository would allow resumption of the process. For example, your ingest script could be roughly: <div>1. does item X exist in the repository (where X is identified by having a particular dc:identifier)</div> <div>2. if not, ingest it</div> <div>3. if so, go to the next item</div>	While the example uses dc:identifier, it wouldn't be a problem to expand this to any directly applied user-managed triple.

Known Implementation Issues

The Modeshape implementation dynamically generates many of triples in its output. For example, rdf:type triples based on whether a resource is a container or a binary, membership triples, and triples generated by Direct or Indirect Containers. These generated triples are not persisted where it would be easy to search them.