

Fedora 6 Testing

This document details both the categories of tests that help in understanding Fedora's performance and scalability characteristics, as well as the tools used to perform those tests. This documentation is an aggregation of various testing initiatives from within the Fedora community over the last several years. Although the tooling has been updated to work with the most recent version of Fedora (6.0), it is important that the results produced through the current tests be compared to historical results in order to better understand the changes (improvements or degradations) in Fedora's performance characteristics over time.

Note: This testing approach was initially structured by the [Product Technology Subgroup's: Fedora 6 Test Plan](#)

- [Test Categories](#)
 - [API specification compliance](#)
 - [OCFL compliance](#)
 - [Ingest performance](#)
 - [Read performance](#)
 - [Versioning performance](#)
 - [Search performance](#)
 - [Migration performance and validation](#)
 - [Rebuild performance](#)
- [Test Tooling](#)
- [Test Results](#)
- [Historical Results](#)

Test Categories

API specification compliance

What is Fedora's degree of compliance with the [Fedora API specification](#)?

OCFL compliance

What is Fedora's degree of compliance with the [OCFL specification](#)? Tests related to this category should be run against both the OCFL produced by Fedora, as well as the OCFL produced by the Fedora 3 to 6 [migration tooling](#).

Ingest performance

What are Fedora's ingest performance characteristics? Specifically, how much content can reasonably be stored in Fedora (number of objects, number of petabytes)? How long does it take to ingest a large corpus of data into Fedora?

Read performance

What are Fedora's read performance characteristics? Specifically, how does the response time of various read requests change as the size of the repository grows?

Versioning performance

How do Fedora's ingest and read performance characteristics change with auto-versioning on or disabled?

Search performance

What are Fedora's query response time characteristics for common and complex queries as the size of the repository grows?

Migration performance and validation

What are the migration times from Fedora 3 to an OCFL representation for a variety of corpus sizes (number of objects, number of petabytes)? What are the rebuild times for Fedora 6 layered over the OCFL representation?

Rebuild performance

What is the speed of the rebuild with different data profiles, database engines, and instance types?

How is rebuild performance affected by instance size (cpu count), database engine, database instance size?

Test Tooling

1. Bash scripts used for measuring timing of fixed-scope operations.
 - a. <https://github.com/fcrepo4-labs/fcrepo-performance-test-scripts>
 - b. Relevant for: *Ingest performance, Read performance, Versioning performance, Search performance*
 - c. [Detailed usage documentation](#)
2. **JMeter** scripts used for measuring limits of scale and stress. These tests run until the application dies.
 - a. <https://github.com/fcrepo4-labs/fcrepo4-jmeter>
 - b. Relevant for: *Ingest performance, Read performance, Versioning performance*
 - c. [Detailed usage documentation](#)
3. Fedora API Test Suite is used to measure compliance of the Fedora application with the [Fedora API specification](#).
 - a. <https://github.com/fcrepo/Fedora-API-Test-Suite>
 - b. Relevant for: *API specification compliance*
 - c. [Detailed usage documentation](#)
4. Python OCFL validator used to measure compliance of an OCFL structure with the [OCFL specification](#).
 - a. <https://github.com/zimeon/ocfl-py>
 - b. Relevant for: *OCFL compliance, Migration performance and validation*
 - c. [Detailed usage documentation](#)
5. Fedora 3 to 6 migration tooling used for migrating F3 objects into OCFL over which F6 can be layered.
 - a. <https://github.com/fcrepo-exts/migration-utils>
 - b. Relevant for: *Migration performance and validation*
 - c. [Detailed usage documentation](#)

Test Results

Current test results can be found on the [Fedora 6 Test Results](#) page.

Historical Results

One of the benefits of running these tests is to measure the change in performance characteristics in relation to previous versions of Fedora. Ideally, one of the products of this testing initiative will be to have a clear depiction, across the test categories, of performance result summaries from the past and present.

Although a consolidated summary of previous tests is not available in a single location, the following references point to the locations where previous test results can be found:

1. Results from "bash scripts"
 - a. [Many Members Performance Testing](#)
2. Results from "JMeter scripts"
 - a. [Test 1 Results Summary](#)
 - b. [Test 2 Results Summary](#)
 - c. [Test 3 Results Summary](#)
 - d. [Test 4 Results Summary](#)
 - e. [Test 5 Results Summary](#)
3. Results from "Fedora API Test Suite"
 - a. [spec-tests](#)
4. Results from very early, ad hoc testing
 - a. [2014-08-20 - Performance Summary](#)