

# Candidate Test Harness

A harness to evaluate performance for Fedora Futures platform candidates.

<https://github.com/futures/ff-jmeter-madness>

## Test Process

1. Delete each test "object"
2. Create an object
3. Retrieve the object
4. Add "datastreams"
5. Modify the datastreams a specified number of times
6. Read the datastreams

## Test Software

[Apache Jmeter version 2.9](#) installed on futures1 (see [Test Platform](#) for details)

The jmeter test script [fedora.jmx](#) implements the above process for a fedora rest API.

## Test Data

The test data is generated at run time to produce a random binary data created from a stable set of file sizes, as explained in [Test corpora - The generated binary dataset](#)

The generation of files has been included in the jmeter test plan and does not need to be generated separately.

## Test Platform

### Test Results

The Jmeter test produces a csv file (one for each repeat) containing the following columns

(Jmeter reference: <http://jmeter.apache.org/usermanual/glossary.html>, <http://jmeter.apache.org/usermanual/listeners.html#csvlogformat>)

To view a log file generated by the Jmeter tests, have a look at the log files from the [fedora's tests](#).

Column name	Description
timeStamp	time of request (in milliseconds since 1/1/1970 )
elapsed	time measured from just before sending the request to just after the last response has been received (in milliseconds). It does not include time taken to render the response (as on a browser).
label	the name given to the CRUD operation (in milliseconds)
responseCode	The HTTP response code received for the request
responseMessage	The HTTP response message received for the request
threadName	This is derived from the Thread Group name and the thread within the group. The name has the format groupName + " " + groupIndex + "-" + threadIndex  groupName - name of the Thread Group element  groupIndex - number of the Thread Group in the Test Plan, starting from 1  threadIndex - number of the thread within the Thread Group, starting from 1
dataType	the type of data received
success	a binary value indicating the success or failure of the request

bytes	the response size
Latency	The time from just before sending the request to just after the first response has been received, including all the processing time needed to assemble the request as well as assembling the first part of the response (in milliseconds).

## Analyzing the test results

The stats visualizations were done using R (<http://www.r-project.org/>).

- For notes on installing R - follow the links for [Ubuntu](#), [Mac OS X](#), [Windows](#)
- The R code uses a few libraries - [ggplot2](#), [gridExtra](#), [tcltk](#), [RColorBrewer](#), [plyr](#)
  - Most of these libraries should be included in the base package, if not installing a library is very easy

```
install.packages("gridExtra", repos="http://R-Forge.R-project.org")
```

The code used to produce the graphs is in [fedora-jmx.r](#).

- To execute the code,

```
$ R  
> source('/path/to/the/file/fedora-jmx.r')
```

- The program will ask you to choose the directory which contains the test results (csv format)
- It will run through all the files, gather the data and produce 3 graphs and a summary of the data. These will be saved in your current working directory. (See [Fedora's test results](#) for an example of the plots generated)

The plots are used to measure the robustness of the software and the time it takes to respond to requests. This is tested for increasing loads to the system.