2016-10-24 Performance - Scale meeting

Time/Place

- Time: 11:00AM Eastern Time US (UTC-5)
- Dial-in Number: (712) 775-7035
  - Participant Code: 479307#
  - International numbers: Conference Call Information
  - Web Access: https://www.freeconferencecallhd.com/wp-content/themes/responsive/flashphone/flash-phone.php

Attendees

- Nick Ruest
- Andrew Woods
- Yinlin Chen
- Daniel Lamb
- Colin Gross
- Danny Bernstein
- Bill Branan ⭐
- Youn Noh
- Osman Din
- David Wilcox
- Noah Botimer

Agenda

1. Update on graphs and summaries of completed tests

```
https://wiki.duraspace.org/display/FF/2016-08-15+Performance+-+Scale+meeting#id-2016-08-15Performance-
Scalemeeting-CurrentSummaries
```

2. Getting community sponsorship for testing infrastructure

   a. Developing a prospectus (goals, system requirements, description of tests, personnel)
   b. Institutions to approach
3. Establishing environment baselines: Colin Gross
4. Finalize initial test results: Danny Bernstein
5. ...

Minutes

1. Update on graphs and summaries of completed tests
   AW: Danny not present. Will postpone discussion. Results collected from previous tests documented on linked page. Final agenda item is to pull together clear summary of tests and simple graphs indicating how performance changes over course of tests. Danny wanted to take this on. R scripts were created years ago to consume output of JMeter tests: https://github.com/fcrepo4-labs/fcrepo4-jmeter. In the absence of any other tooling, may be useful. Is there anyone who has interest in pulling together results? Are there comments on results?
   CG: Experience with R. Will volunteer.
2. Getting community sponsorship for testing infrastructure
   AW: Folks at Art Institute of Chicago interested in participating in release testing. For each release, there is a list of testing activities. Once sanity testing is done, institutions that have the capacity would be in position to run scale tests. ARTIC ran scale tests against 4.7.0 that hit on a particular issue. Contacts: Stefano Cossu, Kevin Ford.
   NR: Access to servers with capacity to do scale testing on behalf of York.
   CG: Michigan's use cases are small.
   YN: Interested in finding out if Yale might be willing to contribute resources.
3. Establishing environment baselines
   CG: Model time it takes to run tests based on benchmarks of operating system. First pass: https://gist.github.com/grosscol/f997f2b3cef80edb640266a03f829a77. How long it takes to write to disk. How long it takes with Fedora. Model is a weight. Faster disks predict faster performance. Linear or logarithmic? Goal is to normalize testing across conditions. Currently working on I/O with JUnit tests. Possibly, other tests for memory, network, CPU.
   AW: Topic has come up a few times. Tests so far have been capacity tests (small files, large files, empty Fedora resources, Fedora resources with RDF). Another suite of tests could focus on demonstrating particular aspects of performance, to be run in a short amount of time but long enough to characterize performance. Slice off classes or methods in integration tests for the REST API.
   AW: Logic at ModeShape layer has been biggest determinant in previous testing. Lower level environment characteristics may not be visible in test results.
   CG: May only be, say, 10%, but will know how to consider in interpreting test results, or if bigger issue, say, at 20%.
   CG: Has tried HTTP API integration tests. Will try JMeter performance tests; may need help limiting runs.

AW: JMeter framework is comprehensive and should have straightforward way of limiting runs. Hardest part may be getting JMeter to run; raise questions on IRC.

CG: Will run tests on local machines; take a stab at regression modeling.

4. Finalize initial test results

AW: Goal is to bring together summaries of tests run across different environments; elevate to mailing lists. Summaries to be linked from: Performance and Scalability Test Plans

AW: Script for outputting characteristics of system: https://github.com/fcrepo4-archive/ff-jmeter-testResults/blob/master/2013-06-03-Fedora4Tests/gatherSystemInfo.sh

CG: Can produce HTML output from R using vector graphics. Will add system characteristics to output.

Actions

- [ ] Colin Gross will look into putting together a script for baselining hardware and network characteristics to be factored into each test run.

- [x] Nick Ruest: Add the log files to be added to the performance test results on the Test Plan page (https://wiki.duraspace.org/display/FF/Performance+and+Scalability+Test+Plans)

- [x] Unknown User (acoburn) will create a summary of what he would like to see out of a test (
  ↑ ~~FCREPO-2242~~ - Remove dependency on OpenRDF `CLOSED`).

- [x] Danny Bernstein: Summarize the existing test data.