

2017-01-16 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](#)
- [Colin Gross](#)
- [Kevin Ford](#)
- [Kieran Etienne](#)
- [Danny Bernstein](#)

Agenda

1. Topics on fedora performance/scaling
 - a. <https://docs.google.com/document/d/1AVYmsgbl-NuCtpmvfHYMqe2izR7leGEIcDx0Ya7q7YU/edit>
2. 2016 work summary
 - a. Review [Performance and Scalability Test Plans](#)
 - b. Graphing of JMeter results, [script](#)
 - c. I/O benchmarking of test environment, [script](#)
 - d. Collection of basic test environment characteristics, [script](#)
3. Hydra - Fedora performance (?)
4. Islandora - Fedora performance (?)
5. Next meeting time 02/13/2017 11am.

Minutes

Topics on Fedora Performance/Scaling

- Nick - none
- Colin - mocking large-scale tests, alt-methods of doing that
 - desire to test multi-TB loads
- Kieran - interested in testing processes and procedures
 - are both Tomcat and Jetty tested? - mostly Tomcat
- Kevin - interested in validating release candidates
 - now more generally interested in loading issues
 - multi-threading topics
- Danny - interested in scaling as it relates to HyBox
 - "many members" issue
 - multi-tenant support
 - Colin has done this in Hydra... but has not been able to backup single hierarchies
 - using different base-containers
- Yinlin - interested in large files

2016 work summary

- Are the defined tests adequate for the interests above?
 - Challenge in have a test environment that has enough disk space
 - Maybe use OpenSSH:rand to mock many files
 - Is running this on AWS infrastructure a possibility?
 - Who is going to pay for a 20-TB EBS volume?
 - Performant disk i/o can be expensive in AWS
 - AWS tests should include data around size/time/cost of previous tests
- We need to summarize the results of 2016 testing
 - Colin available to help in the visualization of the previous test results
 - Yinlin has created graphs of his tests
 - Colin has modified scripts based on team feedback
- Yinlin - need to write report for each test
- Interest in maintaining results in a GitHub
 - ACTION: create results repo in fcrepo4-labs
 - ACTION: write a report from last year's results
 - Current reports do not provide immediate clarity of the state of Fedora

- Volunteers:
 - Colin to work on scripts - fcrepo4-labs
 - Nick/Yinlin to work on drafting - will use GitHub
 - Colin to generate R-Script outputs that can be starting point for summary doc
 - Data result locations: post to S3

Actions

- Colin Gross to transfer fcrepo_perf_analysis to fcrepo4-labs
- Andrew Woods to post test outputs to S3
- Colin Gross to push outputs from R-scripts to GitHub
- Nick Ruest and Yinlin Chen to take above R-script output as starting point for performance test summary report