

2016-07-18 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
- Esme Cowles ★

Agenda

1. Recent test results
 - a. [Virginia Tech Results - Test 2](#)
 - b. [Virginia Tech Results - Test 3](#)
 - c. [Virginia Tech Results - Test 5](#)
 - d. [York University - Test 3 Results - LevelDB](#)
 - e. [York University - Test 3 Results - MySQL](#)
 - f. [Princeton Results - Test 5](#)
 - g. York University - Test 5 (In progress)
2. Reporting out - enough tests have been run to create a summary picture for the community
 - a. Action: Someone to draft a summary message to be added to the wiki and posted to the list
 - b. Test 4; Complete
 - i. [Test 4 Results Summary](#)
 - ii. <https://groups.google.com/d/msg/fedora-tech/NoBhA0E14PA/bqHia6I6AwAJ>
 - c. Test 2
 - d. Test 3
 - e. Test 5
 - f. Complete report
3. Retrieving objects that link to a large number of repository objects (e.g., collections with thousands of members)
 - a. Slow with LevelDB: 10 sec vs. 200 msec for non-repository URI objects
 - b. Much slower with Postgres/MySQL: 30+ sec
 - c. See [Real World Performance \(Retrieving Objects With Many Links to Repository Objects\)](#)

Minutes

1. Additional test results have come in from [Yinlin Chen](#). [Nick Ruest](#) is also running tests on an ongoing basis.
 - a. In general, LevelDB performs well, but does not scale to large numbers of resources.
 - b. PostgreSQL and MySQL scale much better.
 - c. Database configuration and setup have a big impact on performance (unsurprisingly).
2. We should write up summaries of each of the test, like the Test 4 summary
 - a. [Nick Ruest](#) will write up summaries of Tests 2, 3, and 5
 - b. Then we can move on to write up the overall summary
 - c. Would be nice to get Test 1 results with large files, but need more storage to test effectively.
 - d. Also good to run Test 6, which combines various operations.
3. [Esme Cowles](#): I found that it's very slow to retrieve objects with many links to repository objects.
 - a. Will investigate this week and create JIRA tickets when I get to specific issues.