

2016-08-15 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](#) ★
- [Daniel Lamb](#)
- [Andrew Woods](#)
- [David Wilcox](#)
- [Yinlin Chen](#)

Agenda

[Performance and Scalability Test Plans](#)

1. Testing updates
 - a. Retrieving objects that link to a large number of repository objects
2. Finalizing initial tests
 - a. Completing Tests 1, 5 and 6
 - b. Reporting out
3. Next set of tests
 - a. Versioning
 - b. Batch Atomic Operations

Current Summaries

Test 1 - large files

1. VTech - LevelDB, 250MB, out of disk space

Test 2 - small files

1. Princeton - LevelDB, 523k resources, end:too many open-files error
2. Princeton - PostgreSQL, 3.9M resources, end:OOM GC
3. Princeton - PostgreSQL (UseG1GC), 4.1M resources, end:Arjuna timeout
4. Princeton - PostgreSQL (Mode5), 7.7M resources, end:out of disk space
5. VTech - LevelDB, 250k resources, end: 500 response
6. VTech - MySQL, 757k resources, end: 500 response
7. VTech - PostgreSQL, 5.6M resources, end>manual stop

Test 3 - many files

1. Princeton - LevelDB, 584k resources, end:Arjuna timeout
2. Princeton - PostgreSQL, 1M resources, end:out of disk space
3. Princeton - PostgreSQL (Mode5), 979k resources, end:out of disk space
4. Princeton - PostgreSQL (Mode5.1), 975k resources, end:out of disk space
5. VTech - LevelDB, 3.5M resources, end: out of disk space
6. VTech - MySQL, 428k resources, end: 500 response
7. VTech - PostgreSQL, 3.5M resources, end: out of disk space
8. York - LevelDB, 343k resources, end:500 response (no log)
9. York - MySQL (Mode5), 1.3M resources, end:MySQL lost connection

Test 4 - many containers

1. Princeton - LevelDB, 737k resources, end:?
2. Princeton - PostgreSQL, 3.7M resources, end:network interruption
3. Princeton - PostgreSQL (Mode5), 17M resources, end:SQL persistence error
4. Princeton - PostgreSQL (Mode5.1), 4.1M resources, end>manual stop
5. UWMadison - ?
6. VTech - LevelDB, 13M resources, end>manual stop
7. York - LevelDB, 9.1M resources, end:too many open files

Test 5 - many RDF containers

1. Princeton - PostgreSQL (Mode5.1), 7.2M resources, end:out of disk space
2. Princeton - LevelDB, 689k resources, end:Arjuna timeout
3. VTech - LevelDB, 350k resources, end: NoHttpResponseException
4. VTech - MySQL, 533k resources, end: 404,Not Found
5. VTech - PostgreSQL, 2.2M resources, end>manual stop
6. York - MySQL, ?

Test 6 - files and containers

1. None

Minutes

Testing updates

- [Nick Ruest](#) is still running Test #5 at York University
- [Benjamin Armintor](#) is working on two branches to address this issue (lower level; Modeshape level)
- [Esmé Cowles](#) is working on the issue at the Hydra level; adding inverse relationships (two pull requests are in, CurationConcerns and Sufia)
- There will be an update coming out very soon

Finalizing initial tests

Reviewed current summary:

- Test #1 - Need a test setup with enough disk space
- Test #1 - How much disk space do we need?
- Test #2 - Review what's going on with messaging; Arjuna timeout error
- Test #2 - How much disk space do we need to run these tests?
- Test #2 - Another MySQL test would be ideal ([Nick Ruest](#) will run it)

- Virginia Tech tests can only run for 7 days given the resource available. Test 1 will use another machine which doesn't have 7 days limit, but can only be used to conduct Test 1.
- Test #3 - Review what's going on with messaging; Arjuna timeout error
- Test #3 - How much disk space do we need to run these tests?
- Test #3 - [Nick Ruest](#) will look for his logs
- Test #4 - [Nick Ruest](#) will look at the reason York test failed
- Test #5 - Review what's going on with messaging; Arjuna timeout error

What concrete things can we do for further investigation?

- Determine what log files should be shared; jmeter, catalina, mysql, postgres
 - jmeter; test log, perf.log, and csv file for test
- Investigate the Arjuna timeout error; [Daniel Lamb](#) to do some preliminary investigation
- Document server configuration for too-many open files ([Nick Ruest](#) first crack at it)
- Document MySQL and PostgreSQL tuning
- Look closely at the Modeshape 5 tests
- Establish common MySQL and PostgreSQL configuration, or just use default?
 - PostgreSQL might have very conservative default, and may not be indicative of possible performance. [Daniel Lamb](#) to investigate.
- [Andrew Woods](#) to look at generating graphs, and what is needed.

What tests should we re-run?

- All of them, with the goal of getting highest number of resources. This would involve restarting Tomcat, or the server update test condition failure, and starting the test again.
- Identify performance trends by way of graphing