

Meeting Notes 6/9

Theme: Shanghai library FOLIO project - ElasticSearch & checkout performance

Time: June 9, 2020 07:00pm (EST)

Attendees:

Vincent Bateau (Enterprise Architect, EBSCO)

Martin Tran (Principal Software Engineer, tech lead of FOLIO Performance Task Force, EBSCO)

Gang Zhou (Project manager, Shanghai library)

Sha Jiang (Technical Director, Jiayu)

Xiaoyang Chen, Kaiqin (Develop manager, China folio community)

Lucy Liu (Product Owner, Folio China)

Notes:

1. ElasticSearch

- a. Jiang Sha and Zhou Gang will represent the Shanghai team to join the ES discussion group.
- b. Next steps: The discussion group needs to reach some consensus on a solution proposal for a Search Engine (presumably using ElasticSearch). Then it will be run by the Technical Council for approval. After that, the team can start developing the app.
- c. The discussion will start with the current proposal. The Shanghai team will add some features.
- d. The tentative meeting time is 6pm (GMT+8), M/T/W/Th.
- e. Vince will create a channel on Slack for the Search Engine project and send invitations.

2. Performance Testing

- a. What led to the conclusion that Okapi was the problem? (Vince)
- b. Need to normally run 3 Okapis in a cluster because Hazelcast wants to achieve a quorum, which means when one of the servers goes down, it can still rebuild the configuration with the two servers that are left. If we run it with two servers only, it will try very hard to synchronize things more than normal.
- c. One comment from Texas A&M: you get a noticeable improvement in performance if you do not use the same database for applications and for Okapis. If you put Okapis database on a different server, you'll get an improvement.

- d. The testing results from Shanghai and from the community showed the performance was not good in the scenario of 20 and 400 users. Probably not a configuration issue.

3. Circulation App

a. Refactoring Analysis (Vince)

- Let the operations run in parallel. Introduce a checkout session and asynchronous processing. See the attached diagrams.
- Refactoring the inventory.
- Vince will propose the redesign of the circulation app to key stakeholders in the folio project and add it to TC blueprint items in the next couple of weeks. The next steps will be similar to the ES project.

b. Current improvement plans (Martin)

- Under the current architecture, Martin suggested
 - Test GET items API, take a look at the database, optimize queries and add index.
 - Look into GET loans to find the bottleneck.
 - Test POST check-in and POST check-out and profile them.
 - Switch logging level to bring down the time.
- GET items and GET loans will be common to both the current and the new designs of the circulation app. POST check-in and POST check-out might be different depending on how much the current code will be used in the new design.
- Okapi 3.0 with changes to RMB, permissions, tokens, etc. The Shanghai team will try Okapi 3.0.

c. Proposal for performance optimization (Shanghai)

- The Shanghai team suggested adding cache to reduce requests to mod-permission and mod-auth token.
- Vince's comments:
 - Caching can't improve performance significantly. It also has security risk.
 - It makes sense to reduce the number of calls. Okapi 3.0 has made some improvements on this.