

Scaling Solr

Ere Maijala

first.last@helsinki.fi

VuFind Summit 2018



THE NATIONAL LIBRARY OF FINLAND – Library Network Services

docValues

- docValues are stored on disk
 - Cached by the file system
 - Less Java heap needed, thus less GC pressure
 - Look at fieldCache in Solr Admin
 - Less need for warmup (uninverting the inverted index)
- Use for all fields used for faceting AND sorting
 - Including id!
- "stored" can often be set to false
 - Does not preserve order in multivalued fields

Shards and Replica Types

- Sharding
 - Allows parallel processing of updates, faceting etc.
 - Faster recovery
- Replica Types
 - NRT = Near Real Time, the normal type
 - PULL = Pull replicas, not so real time
 - Uses the replication API for getting updates
 - TLOG = Somewhere between the two above
 - Indexing on NRT, others can be PULL
 - Less replicas having to do the indexing work

Deep Paging

- Deep paging = e.g. jumping to results page 100 000.
- Don't allow it
 - At least if sharded index
- Well, at least know what you're doing