

Maintaining the Index

There should be very little effort needed to maintain the index. Solr maintains the index and the HUBzero CMS will instruct Solr to add, remove, or update records inside of its index.

Solr saves its index on the filesystem of the server which allows the retention of data if the server needs to reboot or the Solr process crashes. This prevents having to rebuild the index from scratch in such events.

IMPORTANT NOTE: Due to the large amount of processing power needed to convert database content into a searchable document and the need to communicate with a system outside of the CMS, changes to the index WILL NOT be reflected immediately. The queue will be worked on a first-in-first-out basis. This means that the oldest item in the work queue will be processed first. The amount of time it takes to perform indexing operations depends on the amount of data contained on the hub. If there is a large amount of content, the time to perform a full index will be greater. Once the full index is built, indexing operations should be noticeably quicker