

ReleaseNote40alpha

MMM??? 2012, Apache Solr™ 4.0-alpha available
The Lucene PMC is pleased to announce the release of Apache Solr 4.0-alpha.

Solr is the popular, blazing fast, open source NoSQL search platform from the Apache Lucene project. Its major features include powerful full-text search, hit highlighting, faceted search, dynamic clustering, database integration, rich document (e.g., Word, PDF) handling, and geospatial search. Solr is highly scalable, providing fault tolerant distributed search and indexing, and powers the search and navigation features of many of the world's largest internet sites.

Solr 4.0-alpha is available for immediate download at:
<http://lucene.apache.org/solr/mirrors-solr-latest-redir.html>

See the CHANGES.txt file included with the release for a full list of details.

Solr 4.0-alpha Release Highlights:

The largest set of features goes by the development code-name "Solr Cloud" and involves bringing easy scalability to Solr. See <http://wiki.apache.org/solr/SolrCloud> for more details.

- * Distributed indexing designed from the ground up for near real-time (NRT) and NoSQL features such as realtime-get, optimistic locking, and durable updates.
- * High availability with no single points of failure.
- * Apache Zookeeper integration for distributed coordination and cluster metadata and configuration storage.
- * Immunity to split-brain issues due to Zookeeper's Paxos distributed consensus protocols.
- * Updates sent to any node in the cluster and are automatically forwarded to the correct shard and replicated to multiple nodes for redundancy.
- * Queries sent to any node automatically perform a full distributed search across the cluster with load balancing and fail-over.

Solr 4.0-alpha includes more NoSQL features for those using Solr as a primary data store:

- * Update durability - A transaction log ensures that even uncommitted documents are never lost.
- * Real-time Get - The ability to quickly retrieve the latest version of a document, without the need to commit or open a new searcher
- * Versioning and Optimistic Locking - combined with real-time get, this allows read-update-write functionality that ensures no conflicting changes were made concurrently by other clients.
- * Atomic updates - the ability to add, remove, change, and increment fields of an existing document without having to send in the complete document again.

There are many other features coming in Solr 4, such as

- * Pivot Faceting - Multi-level or hierarchical faceting where the top constraints for one field are found for each top constraint of a different field.
- * Pseudo-fields - The ability to alias fields, or to add metadata along with returned documents, such as function query values and results of spatial distance calculations.
- * A spell checker implementation that can work directly from the main index instead of creating a sidecar index.
- * Pseudo-Join functionality - The ability to select a set of documents based on their relationship to a second set of documents.
- * Function query enhancements including conditional function queries and relevancy functions.
- * New update processors to facilitate modifying documents prior to indexing.
- * A brand new web admin interface, including support for SolrCloud.

This is an alpha release for early adopters. The guarantee for this alpha release is that the index format will be the 4.0 index format, supported through the 5.x series of Lucene/Solr, unless there is a critical bug (e.g. that would cause index corruption) that would prevent this.

Please report any feedback to the mailing lists (<http://lucene.apache.org/solr/discussion.html>)