Solr4.0

Solr 4.0

Solr 4.0-ALPHA was released on 3 Jul 2012 Solr 4.0-BETA was released on 14 Aug 2012 Solr 4.0 GA was released on 12 Oct 2012

See http://lucene.apache.org/solr/solrnews.html to download the latest release.

- Solr 4.0
  - Overview
  - 4.0-ALPHA Release Announcement
  - Errata

Overview

Solr 4.0 refers to the first release planed using the shared 4x branch for Lucene & Solr.

This is slated to be the first release of Solr that will require the use of Java 1.6. (the primary motivation for adding this requirement being SolrCloud functionality)

The current roadmap discussed by the devs (as of May 2012) for 4.0 is:

- 4.0-ALPHA release 3 Jul 2012 **DONE**
- 4.0-BETA release no sooner then 30 days after 4.0-ALPHA. The BETA release may contain additional features & API changes compared to the SLPHA release, but should not change the index format unless absolutely necessary to fix a bug.
- 4.0 (final) release no sooner then 30 days after 4.0-BETA. The final release may contain additional features and API additions compared to the beta release, but should not change any APIs (or the index format) from the beta release unless absolutely necessary to fix a bug.

Please see Solr Development for links to nightly builds that contain the latest 4.x functionality.

If you found a link to this page in some documentation it was there to alert you to the fact that it described a feature that is expected to be included in the Solr 4.0 (final) when it happens. Jira can show you issues that people have (or hope to) incorporated into the 4.0 release.

4.0-ALPHA Release Announcement

Announcement Email

```
3 July 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?ver=4.0a

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)

Happy searching,

Lucene/Solr developers
```

Errata

Once Solr 4.0 has been released, this page may be updated if there are any errata for the release documentation that users should be aware of.