

ReleaseNote44

July 2013, Apache Solr™ 4.4 available

The Lucene PMC is pleased to announce the release of Apache Solr 4.4

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 4.4 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 4.4 Release Highlights:

- * Solr indexes and transaction logs may be stored in HDFS with full read/write capability.
- * Schemaless mode: Added support for a mode that requires no up-front schema modifications, in which previously unknown fields' types are guessed based on the values in added/updated documents, and are then added to the schema prior to processing the update. Note that the below-described features are also useful independently from schemaless mode operation.
 - * New `Parse{Date,Integer,Long,Float,Double,Boolean}UpdateProcessorFactory` classes parse/guess the field value class for String-valued and unknown fields.
 - * New `AddSchemaFieldsUpdateProcessor`: Automatically add new fields to the schema when adding/updating documents with unknown fields. Custom rules map field value class(es) to schema fieldTypes.
 - * A new schemaless mode example configuration, using the above-described field-value-class-guessing and unknown-field-schema-addition features, is provided at `solr/example/example-schemaless/`.
- * Core Discovery mode: A new `solr.xml` format which does not store core information, but instead searches for files named 'core.properties' in the filesystem which tell Solr all the details about that core. The main example and the schemaless example both use this new format.
- * Schema REST API: Add support for creating copy fields.
- * A merged segment warmer may now be plugged into `solrconfig.xml`.
- * New `MaxScoreQParserPlugin`: Return `max()` instead of `sum()` of terms.
- * Binary files are now supported in ZooKeeper.
- * SolrJ's `SolrPing` object has new methods for `ping`, `enable`, and `disable`.
- * The Admin UI now supports adding documents to Solr.
- * Added a `PUT` command to the Solr `ZkCli` tool.
- * New "deleteshard" collections API that unloads all replicas of a given shard and then removes it from the cluster state. It will remove only those shards which are `INACTIVE` or have no range.
- * The Overseer can now optionally assign generic node names so that new addresses can host shards without naming confusion.
- * The CSV Update Handler now supports optionally adding the line number/row id to a document.
- * Added a new system wide info admin handler that exposes the system info

that could previously only be retrieved using a SolrCore.

Solr 4.4 also includes many other new features as well as numerous optimizations and bugfixes.

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

In the coming days, we will also be announcing the first official Solr Reference Guide available for download. In the meantime, users are encouraged to browse the online version and post comments and suggestions on the documentation:

<https://cwiki.apache.org/confluence/display/solr/Apache+Solr+Reference+Guide>

Note: The Apache Software Foundation uses an extensive mirroring network for distributing releases. It is possible that the mirror you are using may not have replicated the release yet. If that is the case, please try another mirror. This also goes for Maven access.