

ReleaseNote66

6 June 2017, Apache Solr 6.6.0 available

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 and analytics, rich document parsing, geospatial search, extensive REST APIs as well as parallel SQL. Solr is enterprise grade, secure and 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 6.6.0 is available for immediate download at:

. <http://lucene.apache.org/solr/mirrors-solr-latest-redir.html>

Please read CHANGES.txt for a full list of new features and changes:

. https://lucene.apache.org/solr/6_6_0/changes/Changes.html

Solr 6.6 Release Highlights:

Fields and field types:

- * Payload support with payload() value source and {!payload_score} and {!payload_check} query parsers
- * Solr support for SimpleTextCodec, via <codecFactory class="solr.SimpleTextCodecFactory"/> in solrconfig.xml (per-field specification in the schema is not possible)
- * Multi-field support to TermsComponent when requesting terms' statistics
- * Grouping support with PointFields
- * CollapseQParser support with PointFields
- * ExpandComponent support with PointFields

New API:

- * UPLOAD command (Config Set API) for uploading zipped configsets
- * MOVE_REPLICA command (Collections API) for moving a replica across nodes
- * LIST_ALIASES command (Collections API) to return a list of all collection aliases
- * STATUS command (Core Admin API) to emit collection details of each core

Script:

- * Basic authentication can be enabled/disabled using bin/solr|bin/solr.cmd
- * ls command to ZkCLI for listing only sub-directories
- * bin/solr script now prints warning when available system entropy is lower than 300
- * solr.in.sh installed by install script should be writable by solr user

Faceting:

- * Variance and Standard Deviation aggregators for the JSON Facet API
- * JSON Faceting now supports a query time 'join' domain change option

Streaming expressions and evaluators:

- * CartesianProductStream, which turns a single tuple with a multi-valued field into N tuples, one for each value in the multi-valued field
- * Basic math Streaming Evaluators
- * UUID Streaming Evaluator
- * Date/time Stream Evaluators
- * analyze Stream Evaluator to support streaming NLP
- * shuffle Streaming Expression
- * search Streaming Expression should work in non-SolrCloud mode
- * echo Streaming Expression
- * eval Streaming Expression
- * timeseries Streaming Expression
- * let, get and tuple Streaming Expressions
- * Correlation Stream Evaluator
- * stats Streaming Expression should work in non-SolrCloud mode
- * regress and predict Stream Evaluators
- * covariance Stream Evaluator
- * convolution Stream Evaluator
- * normalize Stream Evaluator

Examples:

- * Solr default/example uses WordDelimiterGraphFilterFactory and SynonymGraphFilterFactory
- * New DataImportHandler 'atom' example, replacing broken 'rss' example

* Redone DataImportHandler 'tika' example, removing all unused and irrelevant definitions

Metrics:

- * Expose cache statistics using metrics API
- * Improvements to metric reporters and API: support for "regex" parameter in /admin/metrics, "enabled" flag in reporter configurations, correct handling of "serviceUrl" in SolrJmxReporter, better handling of service clients for JMX, Ganglia and Graphite reporters

Upgrades/deprecations/removals:

- * Upgraded Metrics library to 3.2.2, Zookeeper to 3.4.10
- * Deprecated LatLonType, GeoHashField, SpatialPointVectorFieldType, and SpatialTermQueryPrefixTreeFieldType. Instead, switch to LatLonPointSpatialField or SpatialRecursivePrefixTreeFieldType or RptWithGeometrySpatialField
- * Deprecated PostingsSolrHighlighter. Use UnifiedSolrHighlighter instead.
- * Removed Solr contribs: map-reduce, morphlines-core and morphlines-cell

SolrJ:

- * CloudSolrClient can now be initialized using the base URL of a Solr instance instead of ZooKeeper hosts
- * SolrJ: Added SolrParams.toLocalParamsString() and ClientUtils.encodeLocalParamVal
- * LBHttpSolrClient.doRequest is now always wrapped in a Mapped Diagnostic Context (MDC)

Others:

- * New AtomicUpdateProcessor to convert normal update operations to atomic update operations
- * totalTermFreq support to TermsComponent
- * Hide keystore and truststore passwords from /admin/info/* outputs
- * Configurability for thread pool size to recoveryExecutor
- * Introducing sort=childfield(field) asc for searching by {!parent}
- * Transient core cache is now pluggable
- * Renamed getSortWithinGroup to getWithinGroupSort in search.grouping.Command

Optimizations:

- * facet.heatmap is now significantly faster when the docset (base query) matches everything and there are no deleted docs. It is also faster when the docset matches a small fraction of the index or none
- * Reduced heap consumption for filter({!join ... score=...})
- * JSON Facet API now uses hyper-log-log++ for determining the number of buckets when merging requests from a multi-shard distributed request
- * Better ZkStateWriter batching
- * Using cache for DistributedQueue in case of single-consumer

Further details of changes are available in the change log available at: http://lucene.apache.org/solr/6_6_0/changes/Changes.html

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

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 applies to Maven access.