

ReleaseNote50

20 February 2015, Apache Solr™ 5.0.0 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, 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 5.0 is available for immediate download at:

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

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

Solr 5.0 Release Highlights:

- * Usability improvements that include improved bin scripts and new and restructured examples.
 - * Scripts to support installing and running Solr as a service on Linux.
 - * Distributed IDF is now supported and can be enabled via the config. Currently, there are four supported implementations for the same:
 - * LocalStatsCache: Local document stats.
 - * ExactStatsCache: One time use aggregation
 - * ExactSharedStatsCache: Stats shared across requests
 - * LRUSStatsCache: Stats shared in an LRU cache across requests
 - * Solr will no longer ship a war file and instead be a downloadable application.
 - * SolrJ now has first class support for Collections API.
 - * Implicit registration of replication, get and admin handlers.
 - * Config API that supports paramsets for easily configuring solr parameters and configuring fields. This API also supports managing of pre-existing request handlers and editing common solrconfig.xml via overlay.
 - * API for managing blobs allows uploading request handler jars and registering them via config API.
 - * BALANCESHARDUNIQUE Collection API that allows for even distribution of custom replica properties.
 - * There's now an option to not shuffle the nodeSet provided during collection creation.
 - * Option to configure bandwidth usage by Replication handler to prevent it from using up all the bandwidth.
 - * Splitting of clusterstate to per-collection enables scalability improvement in SolrCloud. This is also the default format for new Collections that would be created going forward.
 - * timeAllowed is now used to prematurely terminate requests during query expansion and SolrClient request retry.
 - * pivot.facet results can now include nested stats.field results constrained by those pivots.
 - * stats.field can be used to generate stats over the results of arbitrary numeric functions.
- It also allows for requesting for statistics for pivot facets using tags.
- * A new DateRangeField has been added for indexing date ranges, especially multi-valued ones.
 - * Spatial fields that used to require units=degrees now take distanceUnits=degrees/kilometers miles instead.
 - * MoreLikeThis query parser allows requesting for documents similar to an existing document and also works in SolrCloud mode.
- * Logging improvements:
- * Transaction log replay status is now logged
 - * Optional logging of slow requests.

Solr 5.0 also includes many other new features as well as numerous optimizations and bugfixes of the corresponding Apache Lucene release.

Detailed change log:

http://lucene.apache.org/solr/5_0_0/changes/Changes.html

Also available is the Solr Reference Guide for Solr 5.0. This 535 page PDF serves as the definitive user's manual for Solr 5.0. It can be downloaded from the Apache mirror network: <https://s.apache.org/Solr-Ref-Guide-PDF>

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 goes for Maven access.

