

ReleaseNote64

23 January 2017, Apache Solr 6.4.0 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 6.4.0 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_4_0/changes/Changes.html

Solr 6.4 Release Highlights:

Streaming:

- * Addition of a HavingStream to Streaming API and Streaming Expressions
- * Addition of a priority Streaming Expression
- * Streaming expressions now support collection aliases

Machine Learning:

- * Configurable Learning-To-Rank (LTR) support: upload feature definitions, extract feature values, upload your own machine learnt models and use them to rerank search results.

Faceting:

- * Added "param" query type to facet domain filter specification to obtain filters via query parameters
- * Any facet command can be filtered using a new parameter filter. Example: { type:terms, field:category, filter:"user:yonik" }

Scripts / Command line:

- * A new command-line tool to manage the snapshots functionality
- * bin/solr and bin/solr.cmd now use mkroot command

SolrCloud / SolrJ

- * LukeResponse now supports dynamic fields
- * Solrj client now supports hierarchical clusters and other topics marker
- * Collection backup/restore are extensible.

Security:

- * Support Secure Impersonation / Proxy User for Solr authentication
- * Key Store type can be specified in solr.in.sh file for SSL
- * New generic authentication plugins: 'HadoopAuthPlugin' and 'ConfigurableInternodeAuthHadoopPlugin' that delegate all functionality to Hadoop authentication framework

Query / QueryParser / Highlighting:

- * A new highlighter: The Unified Highlighter. Try it via `hl.method=unified`; many popular highlighting parameters / features are supported. It's the highest performing highlighter, especially for large documents. Highlighting phrase queries and exotic queries are supported equally as well as the Original Highlighter (aka the default/standard one). Please use this new highlighter and report issues since it will likely become the default one day.
- * Leading wildcard in complexphrase query parser are now accepted and optimized with the `ReversedWildcardFilterFactory` when it's provided

Metrics:

- * Use metrics-jvm library to instrument jvm internals such as GC, memory usage and others.
- * A lot of metrics have been added to the collection: index merges, index store I/Os, query, update, core admin, core load thread pools, shard replication, tlog replay and replicas
- * A new /admin/metrics API to return all metrics collected by Solr via API.

Misc changes:

- * The new config parameter 'maxRamMB' can now limit the memory consumed by the FastLRUCache
- * A new document processor 'SkipExistingDocumentsProcessor' that skips duplicate inserts and ignores updates to missing docs
- * FieldCache information fetched via the mbeans handler or seen via the UI now displays the total size used.

* A new config flag 'enable' allows to enable/disable any cache

Further details of changes are available in the change log available at: http://lucene.apache.org/solr/6_4_0/changes/Changes.html

Please note, this release cannot be built from source with Java 8 update 121, use an earlier version instead! This is caused by a bug introduced into the Javadocs tool shipped with that update. The workaround was too late for this Solr release. Of course, you can use the binary artifacts.

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.