



07 June 2015, Apache Solr™ 5.2 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 5.2 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/5\\_2\\_0/changes/Changes.html](https://lucene.apache.org/solr/5_2_0/changes/Changes.html)

Solr 5.2 Release Highlights:

- \* Restore API allows restoring a core from an index backup.
- \* JSON Facet API
  - \* `unique()` is now implemented for numeric and date fields
  - \* Optional flatter form via a "type" parameter
  - \* Added support for "mincount" parameter in range facets to suppress buckets less than that count
- \* Multi-select faceting support for the Facet Module via the "excludeTags" parameter which disregards any matching tagged filters for that facet.
- \* `hll()` facet function for distributed cardinality via HyperLogLog algorithm.
  - See examples at <http://yonik.com/solr-count-distinct/>
- \* A new "facet.range.method" parameter to let users choose how to do range faceting between an implementation based on filters (previous algorithm, using "facet.range.method=filter") or DocValues ("facet.range.method=dv")
- \* Rule-based Replica assignment during collection, shard, and replica creation.
- \* Stats component:
  - \* New 'cardinality' option for stats.field, uses HyperLogLog to efficiently estimate the cardinality of a field w/bounded RAM. Blog post: <https://lucidworks.com/blog/hyperloglog-field-value-cardinality-stats/>
  - \* stats.field now supports individual local params for 'countDistinct' and 'distinctValues'. 'calcdistinct' is still supported as an alias for both options.
- \* Solr security
  - \* Authentication and Authorization frameworks that define interfaces, and mechanisms to create, load, and use authorization/authentication plugins have been added.
  - \* A Kerberos authentication plugin which would allow running a Kerberized Solr setup.
- \* Solr Streaming Expressions
  - See <https://cwiki.apache.org/confluence/display/solr/Streaming+Expressions>
- \* `bin/post` (and `SimplePostTool` in `-Dauto=yes` mode) now sends rather than skips files without a known content type, as "application/octet-stream", provided it still is in the allowed filetypes setting.
- \* HDFS transaction log replication factor is now configurable
- \* A cluster-wide property can now be added/edited/deleted using the `zkcli` script and doesn't require a running Solr instance.
- \* New spatial `RptWithGeometrySpatialField`, based on `CompositeSpatialStrategy`, which blends RPT indexes for speed with serialized geometry for accuracy. Includes a Lucene segment based in-memory shape cache.
- \* Refactored Admin UI using AngularJS. It isn't the default, but a parallel UI interface in this release.
- \* Solr has internally been upgraded to use Jetty 9.

Solr 5.2 also includes many other new features as well as numerous optimizations and bugfixes of the corresponding Apache Lucene release.

For upgrading from 5.1, please look at the "Upgrading from Solr 5.1" section in the change log.

Detailed change log:

[http://lucene.apache.org/solr/5\\_2\\_0/changes/Changes.html](http://lucene.apache.org/solr/5_2_0/changes/Changes.html)

Also available is the Solr Reference Guide for Solr 5.2. This PDF serves as the definitive user's manual for Solr 5.2. 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.

