

ReleaseNote70

20 September 2017, Apache Solr™ 7.0.0 available

The Lucene PMC is pleased to announce the release of Apache Solr 7.0.0

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 7.0.0 is available for immediate download at: <http://lucene.apache.org/solr/mirrors-solr-latest-redir.html>

See http://lucene.apache.org/solr/7_0_0/changes/Changes.html for a full list of details.

Solr 7.0.0 Release Highlights:

- Replica Types - Solr 7 supports different replica types, which handle updates differently. In addition to pure NRT operation where all replicas build an index and keep a replication log, you can now also add so called PULL replicas, achieving the read-speed optimized benefits of a master/slave setup while at the same time keeping index redundancy.
- Auto-scaling. Solr can now allocate new replicas to nodes using a new auto scaling policy framework. This framework will in future releases enable Solr to move shards around based on load, disk etc.
- Indented JSON is now the default response format for all APIs, pass wt=xml and/or indent=off to use the previous unindented XML format.
- The JSON Facet API now supports two-phase facet refinement to ensure accurate counts and statistics for facet buckets returned in distributed mode.
- Streaming Expressions adds a new statistical programming syntax for the statistical analysis of sql queries, random samples, time series and graph result sets.
- Analytics Component version 2.0, which now supports distributed collections, expressions over multivalued fields, a new JSON request language, and more.
- The new v2 API, exposed at /api/ and also supported via SolrJ, is now the preferred API, but /solr/ continues to work.
- A new '_default' configset is used if no config is specified at collection creation. The data-driven functionality of this configset indexes strings as analyzed text while at the same time copying to a '*_str' field suitable for faceting.
- Solr 7 is tested with and verified to support Java 9.

Being a major release, Solr 7 removes many deprecated APIs, changes various parameter defaults and behavior. Some changes may require a re-index of your content. You are thus encouraged to thoroughly read the "Upgrade Notes" at http://lucene.apache.org/solr/7_0_0/changes/Changes.html or in the CHANGES.txt file accompanying the release.

Solr 7.0 also includes many other new features as well as numerous optimizations and bugfixes of the corresponding Apache Lucene release.

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.