

ReleaseNote42

January 2013, Apache Solr™ 4.2 available

The Lucene PMC is pleased to announce the release of Apache Solr 4.2

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 4.2 is available for immediate download at:

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

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

Solr 4.2 Release Highlights:

- * A read side REST API for the schema. Always wanted to introspect the schema over http? Now you can. Looks like the write side will be coming next.

- * DocValues have been integrated into Solr. DocValues can be loaded up a lot faster than the field cache and can also use different compression algorithms as well as in RAM or on Disk representations. Faceting, sorting, and function queries all get to benefit. How about the OS handling faceting and sorting caches off heap? No more tuning 60 gigabyte heaps? How about a snappy new per segment DocValues faceting method? Improved numeric faceting? Sweet.

- * Collection Aliasing. Got time based data? Want to re-index in a temporary collection and then swap it into production? Done. Stay tuned for Shard Aliasing.

- * Collection API responses. The collections API was still very new in 4.0, and while it improved a fair bit in 4.1, responses were certainly needed, but missed the cut off. Initially, we made the decision to make the Collection API super fault tolerant, which made responses tougher to do. No one wants to hunt through logs files to see how things turned out. Done in 4.2.

- * Interact with any collection on any node. Until 4.2, you could only interact with a node in your cluster if it hosted at least one replica of the collection you wanted to query/update. No longer - query any node, whether it has a piece of your intended collection or not and get a proxied response.

- * Allow custom shard names so that new host addresses can take over for retired shards. Working on Amazon without elastic ips? This is for you.

- * Lucene 4.2 optimizations such as compressed term vectors.

Solr 4.2 also includes many other new features as well as numerous optimizations and bugfixes.

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.