

# ReleaseNote70

## 20 September 2017, Apache Lucene™ 7.0.0 available

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

Apache Lucene is a high-performance, full-featured text search engine library written entirely in Java. It is a technology suitable for nearly any application that requires full-text search, especially cross-platform.

This release contains numerous bug fixes, optimizations, and improvements, some of which are highlighted below. The release is available for immediate download at:

<http://lucene.apache.org/core/mirrors-core-latest-redir.html> Please read CHANGES.txt for a full list of new features and changes:

[https://lucene.apache.org/core/7\\_0\\_0/changes/Changes.html](https://lucene.apache.org/core/7_0_0/changes/Changes.html)

Lucene 7.0.0 Release Highlights:

- Doc values switched from random access to iterators.
- The 7.0 codec now sparsely encodes sparse doc values and length normalization factors ("norms"), which also translates to optimization in both indexing, and search on sparse values. With these changes, you finally only pay for what you actually use with doc values, in index size, indexing performance, etc.
- Index time boost for documents is now removed.
- Substantial performance gains for delete and update heavy Lucene usage; see <http://blog.mikemccandless.com/2017/07/lucene-gets-concurrent-deletes-and.html> for details
- Query scoring is now simpler with removal of coord factor, and query normalization.
- Classic query parser no longer splits on whitespaces. This enables better multi-word synonym support.
- The version of Lucene that created the index segment would be recorded, along with the version that last modified the index.
- IndexWriter, used to add, update and delete documents in your index, will no longer accept broken token offsets sometimes produced by misbehaving token filters.
- IndexReader exposes methods that are typically used to manage resources whose lifetime needs to mimic the lifetime of segments/indexes, typically caches. They have been made much less trappy.
- The dimensional points API now takes a field name up front to offer per-field points access, matching how the doc values APIs work.
- The PostingsHighlighter was removed. Migrating to the [UnifiedHighlighter](#) should be straight-forward.

Apache Lucene was tested to be fully compatible with the release of Java 9 and its module system Jigsaw, coming out tomorrow on September 21th!

Further details of changes are available in the change log available at: [http://lucene.apache.org/core/7\\_0\\_0/changes/Changes.html](http://lucene.apache.org/core/7_0_0/changes/Changes.html)

To read more about the changes, also see: <http://blog.mikemccandless.com/2017/03/apache-lucene-70-is-coming-soon.html>

Please report any feedback to the mailing lists (<http://lucene.apache.org/core/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.