Lucene3.6

Lucene 3.6

Lucene 3.6 was released on April 12, 2012

Release Announcement

```
12 April 2012, Apache Lucene™ 3.6.0 available
The Lucene PMC is pleased to announce the release of Apache Lucene 3.6.0.

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 (see
note below).

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

Lucene 3.6.0 Release Highlights:

  * In addition to Java 5 and Java 6, this release has now
    full Java 7 support (minimum JDK 7u1 required).

  * TypeTokenFilter filters tokens based on their TypeAttribute.

  * Fixed offset bugs in a number of CharFilters, Tokenizers and TokenFilters
    that could lead to exceptions during highlighting.

  * Added phonetic encoders: Metaphone, Soundex, Caverphone,
    Beider-Morse, etc.

  * CJKBigramFilter and CJKWidthFilter replace CJKTokenizer.

  * Kuromoji morphological analyzer tokenizes Japanese text, producing
    both compound words and their segmentation.

  * Static index pruning (Carmel pruning) removes postings with low
    within-document term frequency.

  * QueryParser now interprets '*' as an open end for range
    queries.

  * FieldValueFilter excludes documents missing the specified field.

  * CheckIndex and IndexUpgrader allow you to specify the
    specific FSDirectory implementation to use with the new -dir-impl
    command-line option.

  * FSTs can now do reverse lookup (by output) in certain cases and
    can be packed to reduce their size.  There is now a method to
    retrieve top N shortest paths from a start node in an FST.

  * New WFSTCompletionLookup suggester supports finer-grained
    ranking for suggestions.

  * FST based suggesters now use an offline (disk-based) sort, instead
    of in-memory sort, when pre-sorting the suggestions.
```

```
  * ToChildBlockJoinQuery joins in the opposite direction (parent down
    to child documents).

  * New query-time joining is more flexible (but less performant) than
    index-time joins.

  * Added HTMLStripCharFilter to strip HTML markup.

  * Security fix: Better prevention of virtual machine SIGSEGVs when
    using MMapDirectory: Code using cloned IndexInputs of already
    closed indexes could possibly crash VM, allowing DoS attacks to
    your application.

  * Many bug fixes...

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.

Happy searching,

Apache Lucene/Solr Developers
```

Errata

This page may be updated if there are any errata for the release documentation that users should be aware of.