

ReleaseNote35

MMM???? 2011, Apache Lucene, 3.5.0 available
The Lucene PMC is pleased to announce the release of Apache Lucene 3.5.0.

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://www.apache.org/dyn/closer.cgi/lucene/java> (see note below).

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

Lucene 3.5.0 Release Highlights:

- * Added a very substantial (3-5X) RAM reduction required to hold the terms index on opening an IndexReader. (LUCENE-2205)
- * Added IndexSearcher.searchAfter which returns results after a specified ScoreDoc (e.g. last document on the previous page) to support deep paging use cases. (LUCENE-2215)
- * Added SearcherManager to manage sharing and reopening IndexSearchers across multiple search threads. Underlying IndexReader instances are safely closed if not referenced anymore. (LUCENE-3445, LUCENE-3558)
- * Added SearcherLifetimeManager which safely provides a consistent view of the index across multiple requests (e.g. paging/drilldown). (LUCENE-3558, LUCENE-3486)
- * Renamed IndexWriter.optimize to forceMerge to discourage use of this method since it is horribly costly and rarely justified anymore. (LUCENE-3454)
- * Added NGramPhraseQuery that speeds up phrase queries 30-50% when n-gram analysis is used. (LUCENE-3426)
- * Added a new reopen API (IndexReader.openIfChanged) that returns null instead of the old reader if there are no changes in the index. (LUCENE-3464)
- * Improvements to vector highlighting: support for more queries such as wildcards and boundary analysis for generated snippets (LUCENE-1824, LUCENE-1889)
- * IndexSearcher and IndexReader now perform additional checks to throw AlreadyClosedExceptions if searches are performed on a closed IndexReader. Performing searches on already closed reader can cause JVM crashes when invalid memory mapped files are referenced.
- * Several bugfixes, including a bug where closing an NRT reader after the writer was closed was incorrectly invoking the DeletionPolicy. See CHANGES.txt entries for full details.

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.