

Welcome to Apache Lucene!

Table of contents

| | |
|---|----|
| 1 What Is Apache Lucene?..... | 2 |
| 2 The Apache Software Foundation..... | 2 |
| 3 News..... | 2 |
| 3.1 27 November 2011 - Lucene Core 3.5.0 and Solr 3.5.0 Available..... | 2 |
| 3.2 26 October 2011 - Java 7u1 fixes index corruption and crash bugs in Apache Lucene Core and Apache Solr..... | 3 |
| 3.3 14 September 2011 - Lucene Core 3.4.0 and Solr 3.4.0 Available..... | 4 |
| 3.4 28 July 2011 - WARNING: Index corruption and crashes in Apache Lucene Core / Apache Solr with Java 7..... | 5 |
| 3.5 1 July 2011 - Lucene Core 3.3 and Solr 3.3 Available..... | 6 |
| 3.6 4 June 2011 - Lucene Core 3.2 and Solr 3.2 Available..... | 7 |
| 3.7 31 March 2011 - Lucene Core 3.1 and Solr 3.1 Available..... | 8 |
| 3.8 3 December 2010 - Lucene Java 3.0.3 and 2.9.4 available..... | 9 |
| 3.9 25 June 2010 - Solr 1.4.1 Released..... | 10 |
| 3.10 18 June 2010 - Lucene Java 3.0.2 and 2.9.3 available..... | 10 |

1 What Is Apache Lucene?

The Apache Lucene™ project develops open-source search software, including:

- [Apache Lucene Core](#)™ (formerly named Lucene Java), our flagship sub-project, provides a Java-based indexing and search implementation, as well as spellchecking, hit highlighting and advanced analysis/tokenization capabilities.
- [Apache Solr](#)™ is our high performance enterprise search server, with XML/HTTP and JSON/Python/Ruby APIs, hit highlighting, faceted search, caching, replication, distributed search, database integration, web admin and search interfaces.
- [Apache PyLucene](#)™ is a Python port of the the [Lucene Core](#) project.
- [Apache Open Relevance Project](#)™ is a subproject with the aim of collecting and distributing free materials for relevance testing and performance.

2 The Apache Software Foundation

The [Apache Software Foundation](#) provides support for the Apache community of open-source software projects. The Apache projects are defined by collaborative consensus based processes, an open, pragmatic software license and a desire to create high quality software that leads the way in its field. Apache Lucene, Apache Solr, Apache PyLucene, Apache Open Relevance Project and their respective logos are trademarks of The Apache Software Foundation. All other marks mentioned may be trademarks or registered trademarks of their respective owners.

3 News

3.1 27 November 2011 - Lucene Core 3.5.0 and Solr 3.5.0 Available

The Lucene PMC is pleased to announce the availability of Apache Lucene 3.5.0 and Apache Solr 3.5.0.

Lucene can be downloaded from <http://www.apache.org/dyn/closer.cgi/lucene/java/> and Solr can be downloaded from <http://www.apache.org/dyn/closer.cgi/lucene/solr/>

Highlights of the Lucene release include:

- 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.

Highlights of the Solr release include:

- Bug fixes and improvements from Apache Lucene 3.5.0, including a very substantial (3-5X) RAM reduction required to hold the terms index on opening an IndexReader. ([LUCENE-2205](#))
- Added support for distributed result grouping. ([SOLR-2066](#), [SOLR-2776](#))
- Added support for Hunspell stemmer TokenFilter supporting stemming for 99 languages. ([SOLR-2769](#))
- A new contrib module "langid" adds language identification capabilities as an Update Processor, using Tika's LanguageIdentifier or Cybozu language-detection library ([SOLR-1979](#))
- Numeric types including Trie and date types now support sortMissingFirst/Last. ([SOLR-2881](#))
- Added hl.q parameter. It is optional and if it is specified, it overrides q parameter in Highlighter. ([SOLR-1926](#))
- Several minor bugfixes like date parsing for years from 0001-1000, ignored configurations when using QueryAnalyzer with SpellCheckComponent and many more. See CHANGES.txt entries for full details.

3.2 26 October 2011 - Java 7u1 fixes index corruption and crash bugs in Apache Lucene Core and Apache Solr

Oracle released [Java 7u1](#) on October 19. According to the release notes and tests done by the Lucene committers, all bugs reported on July 28 are fixed in this release, so code using

Porter stemmer no longer crashes with SIGSEGV. We were not able to experience any index corruption anymore, so it is safe to use Java 7u1 with Lucene Core and Solr.

On the same day, Oracle released [Java 6u29](#) fixing the same problems occurring with Java 6, if the JVM switches `-XX:+AggressiveOpts` or `-XX:+OptimizeStringConcat` were used. Of course, you should **not** use experimental JVM options like `-XX:+AggressiveOpts` in production environments! We recommend everybody to upgrade to this latest version 6u29.

In case you upgrade to Java 7, remember that you may have to reindex, as the unicode version shipped with Java 7 changed and tokenization behaves differently (e.g. lowercasing). For more information, read `JRE_VERSION_MIGRATION.txt` in your distribution package!

3.3 14 September 2011 - Lucene Core 3.4.0 and Solr 3.4.0 Available

The Lucene PMC is pleased to announce the availability of Apache Lucene 3.4.0 and Apache Solr 3.4.0.

Lucene can be downloaded from <http://www.apache.org/dyn/closer.cgi/lucene/java/> and Solr can be downloaded from <http://www.apache.org/dyn/closer.cgi/lucene/solr/>

If you are already using Apache Lucene 3.1, 3.2 or 3.3, we strongly recommend you upgrade to 3.4.0 because of the index corruption bug on OS or computer crash or power loss ([LUCENE-3418](#)), now fixed in 3.4.0.

Highlights of the Lucene release include:

- Fixed a major bug ([LUCENE-3418](#)) whereby a Lucene index could easily become corrupted if the OS or computer crashed or lost power.
- Added a new faceting module (`contrib/facet`) for computing facet counts (both hierarchical and non-hierarchical) at search time ([LUCENE-3079](#)).
- Added a new join module (`contrib/join`), enabling indexing and searching of nested (parent/child) documents using `BlockJoinQuery/Collector` ([LUCENE-3171](#)).
- It is now possible to index documents with term frequencies included but without positions ([LUCENE-2048](#)); previously `omitTermFreqAndPositions` always omitted both.
- The modular `QueryParser` (`contrib/queryparser`) can now create `NumericRangeQuery`.
- Added `SynonymFilter`, in `contrib/analyzers`, to apply multi-word synonyms during indexing or querying, including parsers to read the wordnet and solr synonym formats ([LUCENE-3233](#)).
- You can now control how documents that don't have a value on the sort field should sort ([LUCENE-3390](#)), using `SortField.setMissingValue`.

- Fixed a case where term vectors could be silently deleted from the index after `addIndexes` ([LUCENE-3402](#)).

Highlights of the Solr release include:

- SolrJ client can now parse grouped and range facets results ([SOLR-2523](#)).
- A new `XsltUpdateRequestHandler` allows posting XML that's transformed by a provided XSLT into a valid Solr document ([SOLR-2630](#)).
- Post-group faceting option (`group.truncate`) can now compute facet counts for only the highest ranking documents per-group. ([SOLR-2665](#)).
- Add `commitWithin` update request parameter to all update handlers that were previously missing it. This tells Solr to commit the change within the specified amount of time ([SOLR-2540](#)).
- You can now specify `NIOFSDirectory` ([SOLR-2670](#)).
- New parameter `hl.phraseLimit` speeds up `FastVectorHighlighter` ([LUCENE-3234](#)).
- The query cache and filter cache can now be disabled per request. See [this wiki page](#) ([SOLR-2429](#)).
- Improved memory usage, build time, and performance of `SynonymFilterFactory` ([LUCENE-3233](#)).
- Added `omitPositions` to the schema, so you can omit position information while still indexing term frequencies ([LUCENE-2048](#)).
- Various fixes for multi-threaded `DataImportHandler`.

3.4 28 July 2011 - WARNING: Index corruption and crashes in Apache Lucene Core / Apache Solr with Java 7

Oracle released [Java 7](#) today. Unfortunately it contains hotspot compiler optimizations, which miscompile some loops. This can affect code of several Apache projects. Sometimes JVMs only crash, but in several cases, results calculated can be incorrect, leading to bugs in applications (see Hotspot bugs [7070134](#), [7044738](#), [7068051](#)).

Apache Lucene Core and **Apache Solr** are two Apache projects, which are affected by these bugs, namely all versions released until today. Solr users with the default configuration will have Java crashing with `SIGSEGV` as soon as they start to index documents, as one affected part is the well-known Porter stemmer (see [LUCENE-3335](#)). Other loops in Lucene may be miscompiled, too, leading to index corruption (especially on Lucene trunk with pulsing codec; other loops may be affected, too - [LUCENE-3346](#)).

These problems were detected only 5 days before the official Java 7 release, so Oracle had no time to fix those bugs, affecting also many more applications. In response to our questions, they proposed to include the fixes into service release u2 (eventually into service release u1, see [this mail](#)). **This means you cannot use Apache Lucene/Solr with Java 7 releases**

before Update 2! If you do, please don't open bug reports, it is not the committers' fault! At least disable loop optimizations using the `-XX:-UseLoopPredicate` JVM option to not risk index corruptions.

Please note: Also Java 6 users are affected, if they use one of those JVM options, which are **not** enabled by default: `-XX:+OptimizeStringConcat` or `-XX:+AggressiveOpts`.

It is strongly recommended not to use any hotspot optimization switches in any Java version without extensive testing!

In case you upgrade to Java 7, remember that you may have to reindex, as the unicode version shipped with Java 7 changed and tokenization behaves differently (e.g. lowercasing). For more information, read `JRE_VERSION_MIGRATION.txt` in your distribution package!

3.5 1 July 2011 - Lucene Core 3.3 and Solr 3.3 Available

The Lucene PMC is pleased to announce the availability of Apache Lucene 3.3 and Apache Solr 3.3.

Lucene can be downloaded from <http://www.apache.org/dyn/closer.cgi/lucene/java/> and Solr can be downloaded from <http://www.apache.org/dyn/closer.cgi/lucene/solr/>

Highlights of the Lucene release include:

- The spellchecker module now includes suggest/auto-complete functionality, with three implementations: Jaspell, Ternary Trie, and Finite State.
- Support for merging results from multiple shards, for both "normal" search results (`TopDocs.merge`) as well as grouped results using the grouping module (`SearchGroup.merge`, `TopGroups.merge`).
- An optimized implementation of `KStem`, a less aggressive stemmer for English
- Single-pass grouping implementation based on block document indexing.
- Improvements to `MMapDirectory` (now also the default implementation returned by `FSDirectory.open` on 64-bit Linux).
- `NRTManager` simplifies handling near-real-time search with multiple search threads, allowing the application to control which indexing changes must be visible to which search requests.
- `TwoPhaseCommitTool` facilitates performing a multi-resource two-phased commit, including `IndexWriter`.
- The default merge policy, `TieredMergePolicy`, has a new method (`set/getReclaimDeletesWeight`) to control how aggressively it targets segments with deletions, and is now more aggressive than before by default.
- `PKIndexSplitter` tool splits an index by a mid-point term.

Highlights of the Solr release include:

- Grouping / Field Collapsing
- A new, automaton-based suggest/autocomplete implementation offering an order of magnitude smaller RAM consumption.
- KStemFilterFactory, an optimized implementation of a less aggressive stemmer for English.
- Solr defaults to a new, more efficient merge policy (TieredMergePolicy). See <http://s.apache.org/merging> for more information.
- Important bugfixes, including extremely high RAM usage in spellchecking.
- Bugfixes and improvements from Apache Lucene 3.3

3.6 4 June 2011 - Lucene Core 3.2 and Solr 3.2 Available

The Lucene PMC is pleased to announce the availability of Apache Lucene 3.2 and Apache Solr 3.2.

Lucene can be downloaded from <http://www.apache.org/dyn/closer.cgi/lucene/java/> and Solr can be downloaded from <http://www.apache.org/dyn/closer.cgi/lucene/solr/>

Highlights of the Lucene release include:

- A new grouping module, under lucene/contrib/grouping, enables search results to be grouped by a single-valued indexed field
- A new IndexUpgrader tool fully converts an old index to the current format.
- A new Directory implementation, NRTCachingDirectory, caches small segments in RAM, to reduce the I/O load for applications with fast NRT reopen rates.
- A new Collector implementation, CachingCollector, is able to gather search hits (document IDs and optionally also scores) and then replay them. This is useful for Collectors that require two or more passes to produce results.
- Index a document block using IndexWriter's new addDocuments or updateDocuments methods. These experimental APIs ensure that the block of documents will forever remain contiguous in the index, enabling interesting future features like grouping and joins.
- A new default merge policy, TieredMergePolicy, which is more efficient due to being able to merge non-contiguous segments. See <http://s.apache.org/merging> for details.
- NumericField is now returned correctly when you load a stored document (previously you received a normal Field back, with the numeric value converted string).
- Deleted terms are now applied during flushing to the newly flushed segment, which is more efficient than having to later initialize a reader for that segment.

Highlights of the Solr release include:

- Ability to specify `overwrite` and `commitWithin` as request parameters when using the JSON update format.
- `TermQParserPlugin`, useful when generating filter queries from terms returned from field faceting or the terms component.
- `DebugComponent` now supports using a `NamedList` to model `Explanation` objects in its responses instead of `Explanation.toString`.
- Improvements to the UIMA and Carrot2 integrations.
- Highlighting performance improvements.
- A test-framework jar for easy testing of Solr extensions.
- Bugfixes and improvements from Apache Lucene 3.2.

3.7 31 March 2011 - Lucene Core 3.1 and Solr 3.1 Available

The Lucene PMC is pleased to announce the availability of Apache Lucene 3.1 and Apache Solr 3.1. The version number for Solr 3.1 was chosen to reflect the merge of development with Lucene, which is currently also on 3.1. Going forward, we expect the Solr version to be the same as the Lucene version. Solr 3.1 contains Lucene 3.1 and is the release after Solr 1.4.1.

Lucene can be downloaded from <http://www.apache.org/dyn/closer.cgi/lucene/java/> and Solr can be downloaded from <http://www.apache.org/dyn/closer.cgi/lucene/solr/>

Highlights of the Lucene release include:

- Numerous performance improvements: faster exact `PhraseQuery`; merging favors segments with deletions; primary key lookup is faster; `IndexWriter.addIndexes(Directory[])` uses file copy instead of merging; various `Directory` performance improvements; compound file is dynamically turned off for large segments; fully deleted segments are dropped on commit; faster snowball analyzers (in contrib); `ConcurrentMergeScheduler` is more careful about setting priority of merge threads.
- `ReusableAnalyzerBase` makes it easier to reuse `TokenStreams` correctly.
- Improved Analysis capabilities: Improved Unicode support, including Unicode 4, more friendly term handling (`CharTermAttribute`), easier object reuse and better support for protected words in lossy token filters (e.g. stemmers).
- `ConstantScoreQuery` now allows directly wrapping a `Query`.
- `IndexWriter` is now configured with a new separate builder API, `IndexWriterConfig`. You can now control `IndexWriter`'s previously fixed internal thread limit by calling `setMaxThreadStates`.
- `IndexWriter.getReader` is replaced by `IndexReader.open(IndexWriter)`. In addition you can now specify whether deletes should be resolved when you open an NRT reader.
- `MultiSearcher` is deprecated; `ParallelMultiSearcher` has been absorbed directly into `IndexSearcher`.

- On 64bit Windows and Solaris JVMs, MMapDirectory is now the default implementation (returned by FSDirectory.open). MMapDirectory also enables unmapping if the JVM supports it.
- New TotalHitCountCollector just counts total number of hits.
- ReaderFinishedListener API enables external caches to evict entries once a segment is finished.

Highlights of the Solr release include:

- Numeric range facets (similar to date faceting).
- New spatial search, including spatial filtering, boosting and sorting capabilities.
- Example Velocity driven search UI at <http://localhost:8983/solr/browse>
- A new termvector-based highlighter
- Extend dismax (edismax) query parser which addresses some missing features in the dismax query parser along with some extensions.
- Several more components now support distributed mode: TermsComponent, SpellCheckComponent.
- A new Auto Suggest component.
- Ability to sort by functions.
- JSON document indexing.
- CSV response format.
- Apache UIMA integration for metadata extraction.
- Leverages Lucene 3.1 and it's inherent optimizations and bug fixes as well as new analysis capabilities.
- Numerous improvements, bug fixes, and optimizations.

3.8 3 December 2010 - Lucene Java 3.0.3 and 2.9.4 available

Both releases fix bugs in the previous versions:

- [2.9.4](#) is a bugfix release for the Lucene Java 2.x series, based on Java 1.4.
- [3.0.3](#) has the same bug fix level but is for the Lucene Java 3.x series, based on Java 5.

New users of Lucene are advised to use version 3.0.3 for new developments, because it has a clean, type-safe API.

This release contains numerous bug fixes and improvements since 2.9.3 / 3.0.2, including:

- a memory leak in `IndexWriter` exacerbated by frequent commits
- a file handle leak in `IndexWriter` when near-real-time readers are opened with compound file format enabled
- a rare index corruption case on disk full

- `NumericRangeQuery` / `NumericRangeFilter` sometimes returning incorrect results with bounds near `Long.MIN_VALUE` and `Long.MAX_VALUE`
- various thread safety issues
- Lucene 2.9.4 can now also read indexes created by 3.0.x

Both releases are fully compatible with the corresponding previous versions. We strongly recommend upgrading to 2.9.4 if you are using 2.9.x; and to 3.0.3 if you are using 3.0.x.

See [3.0.3 CHANGES](#) and [2.9.4 CHANGES](#) for details. **Binary and source distributions are available [here](#).** Maven artifacts are available [here](#).

3.9 25 June 2010 - Solr 1.4.1 Released

Solr 1.4.1 has been released and is now available for public [download!](#) Solr 1.4.1 is a bug fix release for Solr 1.4 that includes many Solr bug fixes as well as Lucene bug fixes from Lucene 2.9.3.

See the [release notes](#) for more details.

3.10 18 June 2010 - Lucene Java 3.0.2 and 2.9.3 available

Both releases fix bugs in the previous versions:

- [2.9.3](#) is a bugfix release for the Lucene Java 2.x series, based on Java 1.4.
- [3.0.2](#) has the same bug fix level but is for the Lucene Java 3.x series, based on Java 5.

New users of Lucene are advised to use version 3.0.2 for new developments, because it has a clean, type-safe API.

Important improvements in these releases include:

- Fixed memory leaks in `IndexWriter` when large documents are indexed. It also uses now shared memory pools for term vectors and stored fields. `IndexWriter` now releases `Fieldables` and `Readers` on close.
- `NativeFSLockFactory` fixes and improvements. Release write lock if exception occurs in `IndexWriter` ctors.
- Improve concurrency of `IndexReader`, especially in the context of near real-time readers.
- Near real-time readers, opened while `addIndexes*` is running, no longer miss some segments.
- Performance improvements in `ParallelMultiSearcher` (3.0.2 only).
- `IndexSearcher` no longer throws `NegativeArraySizeException` if you pass `Integer.MAX_VALUE` as `nDocs` to search methods.

Both releases are fully compatible with the corresponding previous versions. We strongly recommend upgrading to 2.9.3 if you are using 2.9.x; and to 3.0.2 if you are using 3.0.x.

See [3.0.2 CHANGES](#) and [2.9.3 CHANGES](#) for details. **Binary and source distributions are available [here](#).** Maven artifacts are available [here](#).