

Welcome to Solr

Table of contents

1 What Is Solr?.....	2
2 Get Started.....	2
3 News.....	2
3.1 10 November 2009 - Solr 1.4 Released.....	2
3.2 20 August 2009 - Solr's first book is published!.....	3
3.3 18 August 2009 - Lucene at US ApacheCon.....	3
3.4 09 February 2009 - Lucene at ApacheCon Europe 2009 in Amsterdam.....	4
3.5 19 December 2008 - Solr Logo Contest Results.....	5
3.6 03 October 2008 - Solr Logo Contest.....	5
3.7 15 September 2008 - Solr 1.3.0 Available.....	5
3.8 28 August 2008 - Lucene/Solr at ApacheCon New Orleans.....	5
3.9 03 September 2007 - Lucene at ApacheCon Atlanta.....	5
3.10 06 June 2007: Release 1.2 available.....	6
3.11 17 January 2007: Solr graduates from Incubator.....	6
3.12 22 December 2006: Release 1.1.0 available.....	6
3.13 15 August 2006: Solr at ApacheCon US.....	6
3.14 21 April 2006: Solr at ApacheCon.....	6
3.15 21 February 2006: nightly builds.....	7
3.16 17 January 2006: Solr Joins Apache Incubator.....	7

1. What Is Solr?

Solr is the popular, blazing fast open source enterprise search platform from the Apache Lucene project. Its major features include powerful full-text search, hit highlighting, faceted search, dynamic clustering, database integration, and rich document (e.g., Word, PDF) handling. Solr is highly scalable, providing distributed search and index replication, and it powers the search and navigation features of many of the world's largest internet sites.

Solr is written in Java and runs as a standalone full-text search server within a servlet container such as [Tomcat](#). Solr uses the [Lucene Java](#) search library at its core for full-text indexing and search, and has REST-like HTTP/XML and JSON APIs that make it easy to use from virtually any programming language. Solr's powerful external configuration allows it to be tailored to almost any type of application without Java coding, and it has an extensive plugin architecture when more advanced customization is required.

See the complete [feature list](#) for more details.

For more information about Solr, please see the [Solr wiki](#).

2. Get Started

- Download Solr [here](#)
- Check out the [tutorial](#)
- Read the [Solr wiki](#) to learn more
- Join the [community](#)
- [Give Back!](#) (Optional, but encouraged!)

3. News

3.1. 10 November 2009 - Solr 1.4 Released

Solr 1.4 has been released and is now available for public [download](#)! New Solr 1.4 features include

- Major performance enhancements in indexing, searching, and faceting
- Revamped all-Java index replication that's simple to configure and can replicate config files
- Greatly improved database integration via the DataImportHandler
- Rich document processing (Word, PDF, HTML) via Apache Tika
- Dynamic search results clustering via Carrot2
- Multi-select faceting (support for multiple items in a single category to be selected)

- Many powerful query enhancements, including ranges over arbitrary functions, nested queries of different syntaxes
- Many other plugins including Terms for auto-suggest, Statistics, TermVectors, Deduplication

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

3.2. 20 August 2009 - Solr's first book is published!

David Smiley and Eric Pugh are proud to introduce the first book on Solr, "[Solr 1.4 Enterprise Search Server](#)" from Packt Publishing.

This book is a comprehensive reference guide for nearly every feature Solr has to offer. It serves the reader right from initiation to development to deployment. It also comes with complete running examples to demonstrate its use and show how to integrate it with other languages and frameworks.

To keep this interesting and realistic, it uses a large open source set of metadata about artists, releases, and tracks courtesy of the MusicBrainz.org project. Using this data as a testing ground for Solr, you will learn how to import this data in various ways from CSV to XML to database access. You will then learn how to search this data in a myriad of ways, including Solr's rich query syntax, "boosting" match scores based on record data and other means, about searching across multiple fields with different boosts, getting facets on the results, auto-complete user queries, spell-correcting searches, highlighting queried text in search results, and so on.

After this thorough tour, you'll see working examples of integrating a variety of technologies with Solr such as Java, JavaScript, Drupal, Ruby, PHP, and Python.

Finally, this book covers various deployment considerations to include indexing strategies and performance-oriented configuration that will enable you to scale Solr to meet the needs of a high-volume site.

3.3. 18 August 2009 - Lucene at US ApacheCon

[ApacheCon US](#) is once again in the Bay Area and Lucene is coming along for the ride! The Lucene community has planned two full days of talks, plus a meetup and the usual bevy of training. With a well-balanced mix of first time and veteran ApacheCon speakers, the [Lucene track](#) at ApacheCon US promises to have something for everyone. Be sure not to miss:

Training:

- [Lucene Boot Camp](#) - A two day training session, Nov. 2nd & 3rd

- [Solr Day](#) - A one day training session, Nov. 2nd

Thursday, Nov. 5th

- [Introduction to the Lucene Ecosystem](#) - Grant Ingersoll @ 9:00
- [Lucene Basics and New Features](#) - Michael Busch @ 10:00
- [Apache Solr: Out of the Box](#) - Chris Hostetter @ 14:00
- [Introduction to Nutch](#) - Andrzej Bialecki @ 15:00
- [Lucene and Solr Performance Tuning](#) - Mark Miller @ 16:30

Friday, Nov. 6th

- [Implementing an Information Retrieval Framework for an Organizational Repository](#) - Sithu D Sudarsan @ 9:00
- [Apache Mahout - Going from raw data to Information](#) - Isabel Drost @ 10:00
- [MIME Magic with Apache Tika](#) - Jukka Zitting @ 11:30
- [Building Intelligent Search Applications with the Lucene Ecosystem](#) - Ted Dunning @ 14:00
- [Realtime Search](#) - Jason Rutherglen @ 15:00

3.4. 09 February 2009 - Lucene at ApacheCon Europe 2009 in Amsterdam



Lucene will be extremely well represented at [ApacheCon US 2009](#) in Amsterdam, Netherlands this March 23-27, 2009:

- [Lucene Boot Camp](#) - A two day training session, March 23 & 24th
- [Solr Boot Camp](#) - A one day training session, March 24th
- [Introducing Apache Mahout](#) - Grant Ingersoll. March 25th @ 10:30
- [Lucene/Solr Case Studies](#) - Erik Hatcher. March 25th @ 11:30
- [Advanced Indexing Techniques with Apache Lucene](#) - Michael Busch. March 25th @ 14:00
- [Apache Solr - A Case Study](#) - Uri Boness. March 26th @ 17:30
- [Best of breed - httpd, forrest, solr and droids](#) - Thorsten Scherler. March 27th @ 17:30

- [Apache Droids - an intelligent standalone robot framework](#) - Thorsten Scherler. March 26th @ 15:00

3.5. 19 December 2008 - Solr Logo Contest Results

Many great logos were submitted, but only one could be chosen. Congratulations Michiel, the creator of the winning logo that is proudly displayed at the top of this page.

3.6. 03 October 2008 - Solr Logo Contest

By popular demand, Solr is holding a contest to pick a new Solr logo. Details about how to submit an entry can be found [on the wiki](#). The Deadline for submissions is November 20th, 2008 @ 11:59PM GMT.

3.7. 15 September 2008 - Solr 1.3.0 Available

Solr 1.3.0 is available for public download. This version contains many enhancements and bug fixes, including distributed search capabilities, Lucene 2.3.x performance improvements and many others.

See the [release notes](#) for more details. Download is available from a [Apache Mirror](#).

3.8. 28 August 2008 - Lucene/Solr at ApacheCon New Orleans



Lucene will be extremely well represented at [ApacheCon US 2008](#) in New Orleans this November 3-7, 2008:

- [Lucene Boot Camp](#) - A two day training session, November 3rd & 4th
- [Solr Boot Camp](#) - A one day training session, November 4th
- An [entire day of Lucene sessions](#) on November 5th

3.9. 03 September 2007 - Lucene at ApacheCon Atlanta

Lucene will once again be well represented at ApacheCon USA in Atlanta this November 12-16, 2007.

The following talks and trainings are scheduled for this year's conference:

- November 12: [Lucene Boot Camp](#) by Grant Ingersoll. An all-day training focusing on getting started with Lucene.
- November 16, 9:00 am: [Apache Solr out of the Box](#) by Chris Hostetter. Introduction to Solr.
- November 16, 10:00 am: [Building a Vertical Search Site using Apache Software](#) by Ken Krugler. Will cover many Lucene-based projects.
- November 16, 3:00 pm: [Apache Lucene Performance](#) by Grant Ingersoll. Tips and techniques for improving Lucene performance.
- November 16, 4:00 pm: [Advanced Indexing Techniques with Apache Lucene](#) by Michael Busch. Information on payloads and advanced indexing techniques.

3.10. 06 June 2007: Release 1.2 available

This is the first release since Solr graduated from the Incubator, bringing many new features, including CSV/delimited-text data loading, time based autocommit, faster faceting, negative filters, a spell-check handler, sounds-like word filters, regex text filters, and more flexible plugins.

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

3.11. 17 January 2007: Solr graduates from Incubator

Solr has graduated from the Apache Incubator, and is now a sub-project of Lucene.

3.12. 22 December 2006: Release 1.1.0 available

This is the first release since Solr joined the Incubator, and brings many new features and performance optimizations including highlighting, faceted search, and JSON/Python/Ruby response formats.

3.13. 15 August 2006: Solr at ApacheCon US

Chris Hostetter will be presenting "[Faceted Searching With Apache Solr](#)" at ApacheCon US 2006, on October 13th at 4:30pm. See the [ApacheCon](#) website for more details.

3.14. 21 April 2006: Solr at ApacheCon

Yonik Seeley will be presenting "**Apache Solr, a Full-Text Search Server based on Lucene**" at ApacheCon Europe 2006, on June 29th at 5:30pm. See the [ApacheCon](#) website for more details.

3.15. 21 February 2006: nightly builds

Solr now has nightly builds. This automatically creates a [downloadable version of Solr every night](#). All unit tests must pass, or a message is sent to the developers mailing list and no new version is created. This also updates the [javadoc](#).

3.16. 17 January 2006: Solr Joins Apache Incubator

Solr, a search server based on Lucene, has been accepted into the Apache Incubator. Solr was originally developed by CNET Networks, and is widely used within CNET to provide high relevancy search and faceted browsing capabilities.