Welcome to Solr

Table of contents

1. Introduction

Solr is an open source enterprise search server based on the Lucene Java search library, with XML/HTTP and JSON APIs, hit highlighting, faceted search, caching, replication, and a web administration interface. It runs in a Java servlet container such as Tomcat.

See the complete feature list for more details, then check out the tutorial.

2. News

2.1. 02 October 2007 - Solr at OSSummit Asia

Lucene and Solr tutorials!

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

- November 26: Lucene Boot Camp by Erik Hatcher (originally by Grant Ingersoll). An all-day training focusing on getting started with Lucene - the core library under Solr.
- November 27: Solr in a Day by Erik Hatcher. All you need to know to use Solr effectively.
- November 30, 11:30 am: Lucene Case Studies by Erik Hatcher. A rapid series of examples of many Lucene and Solr using applications.

2.2. 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.

2.3. 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.

2.4. 17 January 2007: Solr graduates from Incubator

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

2.5. 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.

2.6. 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.

2.7. 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.

2.8. 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.

2.9. 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.