



Supercharge Your Enterprise Search Capabilities

Deploy Apache Lucene™ or Apache Solr™ on Advanced JVM Technology from Azul Systems for maximum performance and throughput

Accelerate your most demanding enterprise search workloads with Zing®

Apache Lucene is an Open Source, high-performance full-text search engine library written entirely in Java. An underutilized Lucene feature is its RAMDirectory, which provides the ability to search an index in memory instead of on the filesystem to improve performance.

Until Lucene 4 this functionality had generally been out of reach for many implementations because the Lucene RAMDirectory stresses the garbage collector of most commercially available Java Virtual Machines (JVMs). As the in-memory index grew, longer and more frequent garbage collection pauses killed performance. As a result Lucene developers used RAMDirectory only for small indices or testing. For production they let the operating system manage RAM by using MMapDirectory.

If you make the new generation size smaller throughput drops, but if you make it large you have long (sometimes very long) stop-the-world pauses.

Azul Zing is compatible with the Java SE standard and is a JVM that provides a far better alternative. Combining the Apache Lucene search engine with Zing allows very large indices (10s or even 100s of GB) to be held in-memory, increasing max (saturated) throughput by 2X and providing consistent response times even at high query loads. Whether your application is ad serving, driving product search for an eCommerce site or powering search for your enterprise site, Zing will allow you to keep larger indices in memory without garbage collection pauses. You can finally realize the full potential of in-memory search, providing consistent response times and faster search results for your users.

Zing is the only JVM that can elastically scale in memory and CPU cores and still guarantee response time consistency. Its unique C4 garbage collector has a fully concurrent new and old generation for pauseless operation, even with large in-memory datasets and high allocation rates.



BENEFITS OF SUPERCHARGING APACHE LUCENE WITH ZING

- **Delivers highly consistent response times, even under large user loads**
- **2X faster maximum (saturated) throughput**
- **Provides faster time-to-deployment with minimal JVM tuning**
- **Allows large indices to be held in-memory with pauseless operation**
- **Increased production-time visibility using the Zing Vision monitoring tool**



“Azul’s innovative Zing JVM and pauseless GC now enable Apache Lucene project developers to explore use cases requiring large heaps, such as holding an entire search index in memory for faster searching.”

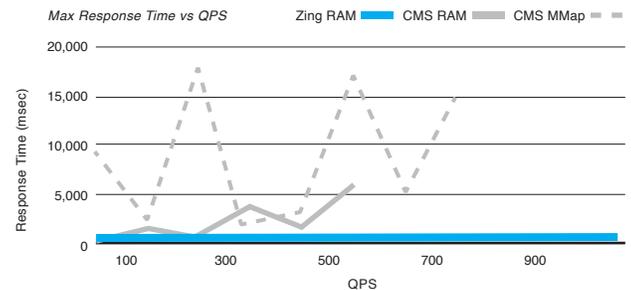
Michael McCandless
Apache Lucene committer and PMC member

Deploying Apache Lucene and Solr on Zing

- Complies with the Java SE standard is based on Oracle’s HotSpot
- Enables practical use of very large in-memory indexes (RAMDirectory) with pauseless operation
- Supports more simultaneous users and queries/sec with greater responsiveness
- Elastic memory – dynamic, automatic allocation of system memory based on application needs
- Requires no changes to the application
- Optimized for Linux and x86 deployments

Zing: the Best JVM for Your Apache Lucene or Solr Development and Deployment

Providing fast search results and handling heavy user and query loads are critical to your company’s or web site’s success. By developing and deploying Apache Lucene enterprise search on Zing, you will be able to hold very large indices in memory with pauseless operation. Zing supports more sustained throughput and higher query loads with greater responsiveness.



Worst case query latencies for the CMS garbage collector and Zing. Note that performance on Zing is remarkably consistent across a wide range of loads. Graph courtesy of Michael McCandless.

“...I remain impressed with Zing, and I wish its C4 collector were the default for Java! Then we all would stop having any GC worries and could freely use Java with very large heaps, fundamentally changing how we build software in this age of very cheap RAM.”

Michael McCandless

To get started, contact us:

Email info@azulsystems.com

Phone +1.650.230.6500

www.azulsystems.com/apache-lucene-solr

Monotype™

Copyright © 2015 Azul Systems, Inc. 1173 Borregas Avenue, Sunnyvale, CA 94089-1306 All rights reserved. Azul Systems, the Azul Systems logo, Zulu and Zing are registered trademarks, and ReadyNow! is a trademark of Azul Systems Inc. Java and OpenJDK are trademarks of Oracle Corporation and/or its affiliated companies in the United States and other countries. Apache, Apache Hadoop, Apache Lucene, Apache Solr, Apache Spark and their respective logos are either registered trademarks or trademarks of the Apache Software Foundation in the United States and/or other countries. Monotype is a trademark of Monotype Imaging Inc. registered in the United States Patent and Trademark Office and may be registered in certain other jurisdictions. The Monotype logo is a trademark of Monotype Imaging Inc. and may be registered in certain jurisdictions. Other marks are the property of their respective owners and are used here only for identification purposes. Products and specifications discussed in this document may reflect future versions and are subject to change by Azul Systems without notice.