



Supercharge Your Search Capabilities

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

Accelerate your most demanding search workloads with Zing®

Elasticsearch is a distributed, RESTful search analytics engine based upon Lucene, and (as is Apache Solr) written in Java. It allows for the performance and combination of many types of searches – structured, unstructured, geo, metric, while its aggregations allow for the exploration of trends and patterns in the data. A great many sites, from Global 1000 to startups use Elasticsearch both for internal use and for Cloud-based search service offerings.

Unfortunately, many organizations aren't able to realize the full potential of their Elasticsearch deployment, and they struggle with Java-based stalls and timeouts caused by CMS or G1-based garbage collection pauses. These pauses and stalls are not related to the search technology itself, but are caused by the underlying Java infrastructure. Often many hours of tuning are required to reach acceptable performance, and periodic spikes in performance issues affect customer satisfaction, use up valuable engineering time and potentially cause operational alarms.

Azul's Zing can help to dramatically improve the performance of your Search deployment. We have years of experience running Search (Lucene, Elastic, Solr) technology on Zing, and thousands of Java-based Search nodes running on Zing in production.

Combining Search with Zing allows very large indices (10s or even 100s of GBs, up to 2 TB total) 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.

With Zing you can achieve the full potential of in-memory search, provide consistent response times and realize faster search results for your users. It is the only JVM that can elastically scale in memory and CPU cores and still guarantee response time consistency. Zing's powerful 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 SEARCH WITH ZING

- Delivers highly consistent response times, even under large user loads
- Improved median search latency PLUS 20X+ lower latency at 90th percentile and beyond
- Provides faster time-to-deployment with minimal JVM tuning
- Allows large indices to be held in-memory – delivering far better QoS



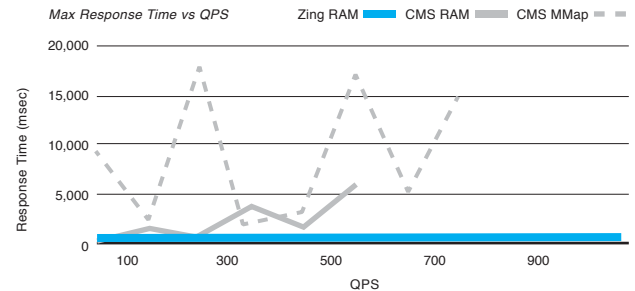
Why Zing for search workloads?

- GC pauses hurt search cluster stability.
- Long (stop-the-world) pauses can make search nodes appear unreachable and lost to the cluster, causing other nodes to stop following the master pause while they elect a new one.
- Long GC pauses can get data nodes ejected from a cluster and forcing shard reallocation can lead to the master removing the node from the cluster and reallocating the paused node's assigned shards, greatly increasing network traffic and disk I/O across the cluster,

Zing *eliminates* Java's stop-the-world GC pauses.

Deploying Elastic, Lucene or Solr on Zing

- Improve QoS eliminating Java GC-caused pauses, stalls and failures
- Enable practical use of very large in-memory indexes with pauseless operation
- Stop struggling with JVM tuning
- Deploy Zing with zero coding changes to your applications
- Meet your SLA targets with few AWS instances or fewer servers, reducing Opex and Capex



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.

Zing: the Best JVM for Elastic, 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 Java-based 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.

“Without Zing we would not have been able to deploy Apache Solr for our production system. Our customers could have experienced long pauses when searching for critical documents.”

Mou Nandi
Search Engineer and Architect
NetDocuments

To get started, contact us:

Email info@azul.com

Phone +1.650.230.6500

www.azul.com/elastic-lucene-solr

Monotype

Copyright © 2017 Azul Systems, Inc. 385 Moffett Park Drive, Suite 115, Sunnyvale, CA 94089 All rights reserved. Azul Systems, the Azul Systems logo, Zulu ReadyNow! and Zing are registered trademarks of Azul Systems. 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.