



Better Java for Big Data

Deploy big data solutions on Azul Zing® to deliver superior throughput and consistent performance, or choose Zulu® for fully supported open source Java.

High Scale Data Analysis

Big data includes many technologies, such as Elasticsearch®, Apache Hadoop™, Cassandra™, Spark™, and Solr™, that are used to create highly distributed, data-intensive solutions. Common use cases include ad targeting, web analytics, bioinformatics and real-time search. Results from these applications are critical to business success, driving real-time decisions that affect revenue, profitability and customer loyalty. Successful deployments often require very low latency, high scalability and fast processing of massive datasets.

Many big data applications and technology components are written in Java. Java provides lots of advantages for enterprises but can be the cause of performance issues. Sometimes extensive Java Virtual Machine (JVM) tuning can help the system reach acceptable performance, and sometimes even that isn't enough.

Delivering Consistent Performance

Big data technologies often use in-memory data storage and processing for fast results. However, response times can

spike and throughput stall if the JVM's memory starts to fill. Once memory usage reaches a certain threshold, the JVM stops processing to clean up old data and free up space, a process called garbage collection (GC). In practice the symptoms are as varied as read time degradation, increased time to data consistency, user timeouts and even system failures.

Zing is proven to solve JVM issues for big data technologies and improves performance of your entire JVM-centric solution stack.

Zing eliminates Java garbage collection as an issue and reduces peak latencies by up to three orders of magnitude, with minimal tuning. With GC problems out of the way, many performance issues are eliminated. Your system will be able to meet SLAs even under growing loads, and users will be delighted by its responsiveness. With Zing, you can finally realize the full value of your big data solution.

BENEFITS OF AZUL ZING FOR BIG DATA APPLICATIONS

- Maximize the number of requests your big data solution can handle and minimize the time it takes to serve them
- Meet the most stringent response time SLAs
- Improve quality of service by eliminating client disconnects due to timeouts
- Deploy with no coding changes to your applications





Azul Zing supports big data use cases not possible with other JVMs.

Perform real-time risk analysis across global portfolios, allow instantaneous drill-down in analytics systems, detect cyberattacks in complex networks, win more online ad auctions and provide a fully personalized experience for users.

Big data goes mainstream

With cheaper memory and storage, the rise of the Internet of Things and businesses taking a more analytics-based approach to everything from customer service to supply chain management, big data technologies are everywhere. This has created constant pressure on developers and IT to analyze more information faster and make the latest data available to users in real time. Only Azul Zing allows Java-based applications to meet these needs with consistently fast response times.

Zing is certified Java SE 8, 7 and 6 compliant. It requires no changes to your existing applications. Zing is also certified by our partners for use with Cloudera C5, Hortonworks Data Platform, DataStax and many other big data technologies. Zing is easy to deploy, and you'll see the value right away.

Azul Zulu: Your Fully Supported Open Source Java Option

Zulu is an enterprise-quality, certified build of OpenJDK™ that can be deployed across various operating systems (Windows, Mac OS X, Linux), Docker containers, hypervisors and Cloud platforms. Companies, Java developers, systems administrators and end users can now enjoy the full benefits of freely available 100% open source Java for big data solutions with the comfort of world-class support.

Selected Azul Big Data Partners



CUSTOMER SUCCESS

Feedzai Fraud Detection

Problem:

Processing pauses caused by a legacy JVM were causing fraudulent transactions to be missed, which created unplanned losses for card issuers.

Solution:

Azul partner Feedzai has a Cassandra-based real-time fraud detection system that uses Zing to ensure maximum streaming write throughput. With Zing, Feedzai can meet even the most demanding SLAs from some of the world's largest financial institutions.

“The real-time analysis of data to prevent fraud in the financial industry is key to predicting and preventing fraud. It's almost impossible to have ultra-low latencies – in the range of 5-10 milliseconds with a standard JVM – and our customers demand that. Azul powers the largest banks in the world and with peak load demands of up to 50,000 transactions per second, Zing will help ensure that we can deliver the best that artificially intelligent machines can offer.”

– Nuno Sebastiao, Chief Executive Officer of Feedzai

Contact us: info@azulsystems.com +1.650.230.6500 @AzulSystems azul.com/BigData