



Research Abstract

Java-Based Benchmarking Shines a Light on How Underlying Architecture Impacts Cloud Performance

Use Java Development Kit (JDK®) benchmarking results to make evidence-based management decisions that can help you achieve optimal performance from cloud-native applications.

The point of benchmarking is to translate lab results into actionable business decisions. For example, say you want improved performance from cloud-native, Java-based applications. Benchmarking can help you make evidence-based decisions about purchasing software and hardware and provisioning resources.

It is a given that improving workload performance in the cloud can help satisfy your customers' ever-growing appetites for fast, accurate data delivery. Customers who get stuck with a slow or faltering experience while performing search queries, watching media streams, or accessing medical or financial information might soon take their business elsewhere. Of course, improved performance is only one reason why so many organizations have migrated workloads such as Elasticsearch®, Apache Cassandra®, and Apache Spark™ to the cloud. Organizations of all sizes are using cloud-based platforms to stay competitive, meet service-level agreements (SLAs), and lower total cost of ownership (TCO).

Upgrading to **the latest long-term support (LTS) release of JDK®** can immediately boost the performance of cloud-native applications up to ...



Cloud-native applications run better on **3rd Gen Intel® Xeon® Scalable processors**, up to ...



... than Amazon Web Services® (AWS®) Graviton2 processors



... than AMD EPYC™ processors

Imagine What You Could Do with Faster Cloud Performance

Prowess benchmark results indicate 3rd Gen Intel® Xeon® Scalable processors can deliver better cloud performance than Amazon Web Services® (AWS®) Graviton2 or AMD EPYC™ processors for a variety of real-world examples.

Top social-media sites, news outlets, and other cloud-native companies understand that in the war for market share, service latency kills. These businesses rely on the Elasticsearch® search engine to sort through massive volumes of data and deliver answers in milliseconds.^{2,3}

To meet strict streaming-media requirements that include near-real-time delivery with extremely low failure percentages, social-media platforms turn to the Apache Cassandra® NoSQL database for high-speed data storage and access.⁴

An Apache Spark™ pipeline enables life-sciences researchers to speed up analyses of massive data volumes, such as high-throughput genome-sequencing.⁵

Applying benchmarking best practices, Prowess Consulting tested cloud-native applications to surface which underlying factors affect workload performance. The results demonstrate that Java Development Kit (JDK) version and processor type have the most impact. Benchmark testing of Java-based workloads in an Amazon® Elastic Compute Cloud™ (Amazon EC2®) environment revealed that using the latest LTS release of JDK improved the overall performance of Java-based applications across the board, and that cloud instances powered by Intel® Xeon® Scalable processors consistently outperformed those powered by Amazon Web Services® (AWS®) Graviton2 and AMD EPYC™ processors.

Get the full story by reading the technical research report,

“Java-Based Benchmarking Shines a Light on How Underlying Architecture Impacts Cloud Performance.”

¹ For system details, test results and analyses, see the full report. “Java-Based Benchmarking Shines a Light on How Underlying Architecture Impacts Cloud Performance.” www.prowesscorp.com/project/java-based-benchmarking-shines-light-on-how-underlying-architecture-impacts-cloud-performance

² These companies include Uber, Lyft, Tinder, Cisco, and Sprint. Source: “Elastic, search company for Uber and Tinder, nearly doubles in IPO.” October 2018. www.zdnet.com/article/elastic-search-company-for-uber-and-tinder-nearly-doubles-in-ipo/.

³ These companies include Wikipedia, The New York Times, Engadget, eBay, Sprint, Walmart, Meta (formerly Facebook), Walgreens, and Kreeger. Source: Sematext Group. “Elasticsearch Tutorial: A complete guide to getting started with the basic concepts: what it is, how it works, and what it’s used for.” <https://sematext.com/guides/elasticsearch/#what-is-elasticsearch-used-for-applications-examples>.

⁴ Instagram requires no more than nine-second latency for its streaming media with a failure of less than 0.001 percent. Source: Instagram engineering. “Open-sourcing a 10x reduction in Apache Cassandra tail latency.” May 2018. <https://instagram-engineering.com/open-sourcing-a-10x-reduction-in-apache-cassandra-tail-latency-d64f86b43589>.

⁵ The Broad Institute of MIT, in collaboration with Harvard University, uses an Apache Spark™ pipeline deployed on a Cray Urika®-GX supercomputer to reduce high-throughput genome-sequencing analysis time by up to 4.4x, compared to analysis without Apache Spark. Source: Data Centre Dynamics. “Cray announces Urika-GX, a supercomputing platform for big data.” May 2016. www.datacenterdynamics.com/en/news/cray-announces-urika-gx-a-supercomputing-platform-for-big-data/.



The analysis in this document was done by Prowess Consulting and commissioned by Intel.
Results have been simulated and are provided for informational purposes only.
Any difference in system hardware or software design or configuration may affect actual performance.
Prowess and the Prowess logo are trademarks of Prowess Consulting, LLC.

Copyright © 2022 Prowess Consulting, LLC. All rights reserved.
Other trademarks are the property of their respective owners.