

Digital businesses need high-performing applications to retain customers and grow revenues. With IT budget, supply chain, and inflationary pressures rising, IT executives should focus on delivering great customer experiences that drive business value while controlling and optimizing teams and cloud costs.

Optimizing Java Application Performance for Improved Business Outcomes and Cloud Cost Efficiency

May 2023

Written by: Jevin Jensen, Research Vice President, Intelligent CloudOps

Introduction

CEOs want IT to be more agile and focused on tangible business outcomes. These same executives consistently tell IDC that improving the customer experience is the top priority for their business. However, in a January 2023 survey, IDC found that 73% of IT budgets are flat or negative for 2023 — a sign of tightened spending brought on by high inflation and a potential recession. So, where will the funds come from to improve the customer experience? How can IT address both challenges simultaneously?

Many companies quickly moved to a new digital business model during the pandemic. IT created new cloud-based customer-facing applications or migrated existing on-premises customer-facing applications to the cloud in record time. By moving to the public cloud and cloud-native application architectures, IT was able to meet the demands of the business during a challenging time. However, developers often wrote code without the scalability of the cloud in mind or migrated "as is" from on-premises legacy servers.

To meet the needs of the business and improve the customer experience required, IT must create and deliver more resilient and high-performing applications. Customers will quickly switch to competitors if a company's application is unavailable. Furthermore, a slow application can mean abandoned digital shopping carts or order transactions that fail to complete correctly. These issues impact the customer experience and lower a company's Net Promoter Score (NPS) and Customer Satisfaction Score (CSAT), which affects long-term revenue growth. In addition, as almost every product now relies on a set of digital technologies and processes delivered by applications, there is no margin for error for poorly performing applications. Unoptimized applications and error-heavy code can be company killers.

AT A GLANCE

KEY STAT

According to a January 2023 IDC survey, 73% of IT budgets are flat or negative for 2023.

WHAT'S IMPORTANT

C-suite executives tell IDC that customer experience (CX) is the number 1 priority of the business. At the same time, IT budgets have little or no funding to improve CX. Highly resilient and high-performing applications are important to delivering a positive customer experience.

KEY TAKEAWAY

Effective management and optimization of both costs and performance of Java applications are critical to the success of digital businesses.

With tightening IT budgets, enterprises must solve these application performance and resiliency issues *and* reduce the costs of operating applications once in production. To solve these challenges, enterprises should look to their unoptimized application code and how to improve runtime performance. Rewriting or replatforming entire applications can take years, but optimizing legacy Java code is one way to gain tangible business results today. Oracle estimates that over 60 billion Java virtual machines (JVMs) are currently used. With the recent Java SE subscription pricing change, existing Java applications can be much more expensive than they were just a few years ago. IDC estimates that 750 million new applications will be created by 2025, many of those in Java. Enterprises need an easy way to optimize the performance and licensing of their new and existing Java applications without the burden of rewriting or recompiling.

IT teams are facing the challenges of a tough budget environment and the need for high-performing applications. Optimized Java virtual machines offer a solution to both.

Many of these applications will exist in the public cloud. IDC estimates that enterprises waste 10–30% of their public cloud spending. Enterprises often overprovision servers as a quick way to handle slow application performance. Companies add resources to address Java's well-known slow start-up and warm-up performance issues. This slowness is typically seen before Java later adjusts to peak loads. The rush to the public cloud forced some IT shops to take a "lift and shift" approach, which did not optimize the application code or configurations for the cloud. This approach wastes more resources and increases costs. New "cloud native" applications frequently use popular services such as Kafka, Cassandra, and Solr, which may compound performance issues because these frameworks also run on the Java runtime.

Additionally, enterprises sometimes struggle with hybrid cloud support. Many enterprises tell IDC they run an on-premises cloud to support legacy applications and their public cloud provider. When you include the applications run at company-owned "edge" locations, over 60% of production transactions are run at a company-owned facility. This complexity can make it difficult for operation teams to support and identify performance issues. Enterprises often continue to run applications on premises for compliance or security reasons or simply because they have a backlog of legacy applications to modernize. Enterprises now develop and manage applications in a public and private cloud world, but the principles around cloud cost optimization remain the same in either environment.

Java has its inherent application performance design challenge too. The issue is referred to as the "stop the world" process. This process occurs during the Java garbage collection (GC) cycle, where every thread except the single GC thread is suspended. Program execution is stopped while the GC algorithm tracks Java objects, releases dead memory, and reclaims space for all objects in the heap. The time this process consumes, and its frequency, can directly impact the performance of an application and, by extension, the amount of cloud resources needed to meet service-level agreements (SLAs).

Enterprises must look for a comprehensive platform solution that improves Java performance, optimizes hardware resources, reduces cloud costs, and addresses resiliency issues such as start-up, warm-up, and "stop the world." Only then can enterprises deliver an exceptional customer application experience at a price the enterprise can afford.

Benefits

Enterprises looking to implement a Java optimization platform expect to solve multiple challenges in one integrated tool. The following is a list of expected benefits from a successfully implemented solution:

- » Ability to optimize the performance of custom Java applications and web services without any code changes
 - 2x throughput on the same infrastructure, reduce resources needed
 - Solve common Java issues of "slow start-up and warm-up" and "stop the world" architecture
- » Improved resiliency and availability with reduced performance tuning
- » Multiple cost improvements
 - 50% fewer compute/memory resources required for the same performance
 - Replace expensive Oracle Java SE subscriptions with lower-cost compatible option
 - Integrate Java apps as part of a FinOps strategy to manage future cloud costs and projects
- » Extended support for older versions (Java 6 and 7) found in many legacy apps
 - Enables legacy migration on the schedule of business while providing improved performance now
- » Back-end infrastructure improvement for Kafka and Cassandra, JBoss, Solr, Tomcat, and other cloud-native and Java-based technologies (Vendor reports up to 135% improvement in these Java applications.)
- » Ability for SaaS providers to improve the scalability of applications, passing savings along to end customers while improving their margins

Key Trends

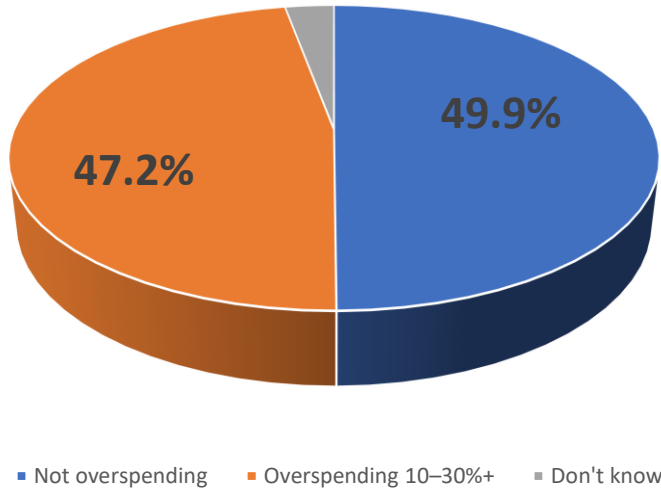
Enterprises are seeing explosive growth in applications to support digital business. Java is still very popular for modernizing applications. Popular use cases for Java include:

- » Data warehouse and analytics that need to support multi-terabyte heaps
- » Internet of Things (IoT) on the edge for data ingestion
- » Time-sensitive applications such as trading, crypto, risk/fraud, adtech, and gaming
- » SaaS providers supporting many customers with single-purpose applications
- » eCommerce retailers that need to scale out to handle seasonal business

Public cloud compute services are still growing and are a popular target for new Java applications. However, business executives are increasingly concerned about exploding cloud costs for resources needed to support these applications. See Figure 1 for evidence of respondents' concern about growing cloud costs.

FIGURE 1: **Cost Pressure Rising on Cloud Services**

Q Which of the following best matches your opinion when it comes to overspending on the following type of cloud resources: IaaS compute services?



n = 952

Source: IDC's Future Enterprise Resiliency and Spending Survey, Wave 2, March 2023

Considering Azul

Azul provides enterprise-grade solutions for running Java applications. The company has focused on delivering Java-based solutions since its founding in early 2002. Scott Sellers (now CEO), Gil Tene (current CTO), and Shyam Pillalamarri were the founders. These industry leaders identified enterprises needing to improve their Java performance and scalability in mission-critical applications. As a result, the company initially started work on a new type of processor for Java, which it supported until 2013. During this same time, the company pioneered a software-only approach to scaling and increasing the performance of Java applications. Launched in 2010, Azul Platform Prime (formerly Zing) was the company's first software-only solution.

Platform Prime is fully compliant and certified with Oracle's Java SE version and the OpenJDK Java Development Kit (JDK) and requires no code changes to use. Enterprises need not modify or recompile their Java code to utilize Platform Prime. Platform Prime JDK is available for multiple platforms, including Linux, Windows, and macOS, and supports a wide range of Java versions, from Java SE 6 to Java SE 17. It provides several features unavailable in other JVMs. Platform Prime's goal is to make applications start fast, go fast, and run fast — including eliminating the stop-the-world problem of Java. It achieves this by using Azul's patented Continuously Concurrent Compacting Collector (C4) garbage collector, which operates concurrently with the application threads and can manage large heap sizes with minimal impact on application performance. The company calls this "pauseless" garbage collection. This capability allows applications to scale to hundreds of gigabytes of memory or even multiple terabytes per JVM while maintaining predictable latency and high throughput. The heap sizes and garbage collection settings are dynamically adjusted, which can help improve application performance and improve resiliency.

Platform Prime also includes performance-improving features such as Falcon. Falcon is a just-in-time (JIT) compiler for Java that optimizes code for today's server hardware and is the default compiler in Platform Prime. This JIT compiler uses the low-level virtual machine (LLVM) open source code supported by many leading technology companies. This state-of-the-art compiler is continually improved to support every advance in processor architecture of the x64 and ARM64 (AWS Graviton) product lines. Falcon promises to reduce the amount of time developers spend on performance tuning and optimizes memory footprint when an application runs in the public cloud.

Another component of Platform Prime is its ReadyNow! feature. The goal of this feature is to address Java's "warm-up" issue. By default, Java starts fast but optimizes slowly over time to improve performance. Unfortunately, this can create problems because many companies have peak usage periods for their mission-critical applications. As a result, Java-based systems often experience slow performance for the first several minutes, which gives end users a slow start to each day. Java applications that are slow to respond can particularly impact financial and ecommerce companies, resulting in missed stock trades or abandoned shopping carts.

Some companies, such as high-performance trading platforms, have resorted to using synthetic orders or transactions to warm up systems each day, which can lead to other deoptimization issues. The goal of ReadyNow! is to deliver a practical long-term solution by providing IT teams with optimization profiles stored for future reuse. In addition, Azul offers a Cloud Native Compiler (CNC) with Platform Prime to improve speed and reduce cloud costs. CNC offloads compilation to the cloud instead of the local client — using the power of the cloud for compilation to enable faster code and greater efficiency. Additionally, ReadyNow! provides API for policy and configuration directives, allowing Java to participate fully in low-latency environments.

Azul launched Platform Core (formerly called Zulu) to enterprises starting in 2013. Based on OpenJDK, Platform Core is fully certified and compliant. It provides a direct replacement for Oracle Java SE. The company states that customers typically save 70% on license costs when replacing Oracle's Java. Platform Core is 100% open source and supports standard garbage collectors such as G1 GC or Parallel GC. Also, it uses industry-standard Hotspot JIT for its compiler. Platform Core, as well as Platform Prime, makes available two sets of quarterly patch updates. First, the Patch Set Update (PSU) includes critical security updates and a superset of updates adding more functionality and new features. Second, the Critical Patch Update (CPU) contains only fixes to security updates and critical bug fixes. Aside from Oracle, Azul is the only Java provider to deliver both sets of updates, typically within one hour of Oracle.

Last year, Azul announced a new security product called Azul Vulnerability Detection, which is designed to mitigate risks from vulnerabilities such as Log4j. This agentless cloud service gives enterprises visibility into their end-to-end software supply chain. In addition, Azul states that applications will see no performance penalty for this real-time usage monitoring against known vulnerabilities.

Azul also offers professional services to help customers optimize their Java applications for performance and scalability as well as migration services for moving off Oracle Java SE. For example, the company's experts can collaborate with customers to identify areas of their application that may be causing performance issues or scalability bottlenecks and provide recommendations for addressing these issues. In addition, Azul offers training and certification programs to help developers and IT professionals become Java performance optimization and management experts. Further, Azul offers out-of-band support and fixes for its JVM as well.

Challenges

Many enterprise IT teams across operations and development have existing performance management solutions. These are often called observability solutions and are used to address application performance issues. However, IT may be overconfident in solving Java application performance incidents while spending too much time performance tuning. Some IT groups don't connect the opportunity to increase performance with lowering application costs. Others might not have the required skill for proper Java application performance management capabilities, which might require training and partners to assist in developing a performance strategy for performance optimization and cost reduction. Some organizations don't focus enough on Java application performance and underestimate the importance of performance management across these often mission-critical, customer-facing applications.

Conclusion

Today, enterprise IT groups face multiple challenges in controlling costs while providing better-performing, more resilient applications to their customers. The long lead times necessary to optimize or replatform existing applications can mean businesses lose customers and revenue is negatively impacted. Enterprises looking to improve application performance and control cloud costs should consider a Java optimization platform. By speeding up the start-up of Java applications, improving scaling and runtime performance, and reducing the number of cloud resources needed for the environment, a Java optimization platform can be the silver bullet IT needs to meet these challenges.

About the Analyst



Jevin Jensen, Research Vice President, Infrastructure and Operations

Jevin Jensen is the Research Vice President covering IDC's Intelligent CloudOps Markets service. He covers infrastructure as code, GitOps, IT infrastructure automation, cloud cost transparency, FinOps, hybrid/public/multicloud management platforms, and edge management.

MESSAGE FROM THE SPONSOR

More About Azul Systems

Azul is the world's largest independent Java provider, supporting more versions of Java than any other vendor, including Oracle. Azul has nearly 1,000 customers and 300+ employees around the globe, with a 20-year history of Java innovation. Azul delivers faster Java application performance with smaller cloud footprints, helping companies improve cash flows.

Azul Platform Prime is ideal for optimizing cloud costs, and is used by FinOps, Performance Engineering, Infrastructure and Operations, Platform Engineering and many other teams to improve the performance of Java-based applications and infrastructure such as Apache Cassandra, Apache Kafka and Elasticsearch.

Azul Platform Prime is proven to reduce cloud costs by up to 50% via a simple equation: faster apps = less compute = lower cloud bill.

Azul customers include 35% of the Fortune 100 and leading brands like Avaya, Bazaarvoice, BMW, Credit Suisse, Deutsche Telekom, LG, Mastercard, Mizuho, Priceline, Salesforce, Software AG and Workday. Learn more at www.azul.com



The content in this paper was adapted from existing IDC research published on www.idc.com.

This publication was produced by IDC Custom Solutions. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Custom Solutions makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2023 IDC. Reproduction without written permission is completely forbidden.

IDC Research, Inc.
140 Kendrick Street
Building B
Needham, MA 02494, USA
T 508.872.8200
F 508.935.4015
Twitter @IDC
idc-insights-community.com
www.idc.com