



Highlights

- Is speed costing you money? A financial brokerage firm found it lost USD 4 million per millisecond even if it was only 5 milliseconds behind its competition.
 - What if you could significantly improve the service delivery of your enterprise Java applications? Could it improve your revenue?
 - Java applications run faster on IBM Z® and LinuxONE™ servers based on concurrent garbage collection through the Guarded Storage Facility.
 - The IBM z14™, Rockhopper II™, and Emperor II™ reduce garbage collection time by 90-92 percent compared to a Lenovo SR650
 - Simultaneously, they delivered three times or more throughput and 80-92 percent faster response time
-

Significantly improve the service delivery of enterprise Java applications

Speed matters. A lot.

Speed can cost you customers, your reputation and revenue.

Amazon found that every 100 milliseconds of latency cost the company 1 percent of sales. That's about USD 1.8 billion today. A financial brokerage firm discovered it could lose USD 4 million in revenue per millisecond if it was only 5 milliseconds behind its competition. And Google internal studies found that a half-second delay in generating search results essentially killed user satisfaction and caused traffic to drop by 20 percent.¹

How does speed affect your business? A half second seems like a pretty small timeframe—but can you afford it? Do you dare take the gamble?

Consider enterprise Java business applications. Java is one of the most popular programming languages used today. It runs everything, from laptops to games to supercomputers, and is particularly used extensively in large-scale client-server deployments. One of the most unique and important features of Java—and arguably a feature that leads to its popularity—is that it performs automatic “garbage collection.” In other words, the Java runtime, that is, the Java Virtual Machine (JVM), automatically allocates and deallocates memory so programmers don't have to, unlike other languages like C and C++ where memory management must be explicitly implemented in the code. But one of the unfortunate drawbacks of this feature is that all program threads stop at once when garbage collection takes place. And for some Java programs, the act of garbage collection can become quite noticeable, with detrimental “stop-the-world” pause times that affect processing and service delivery.



IBM Systems

Executive Summary

It's not necessarily easy to predict whether a Java workload might be subject to long garbage collection pause times, although, generally, these programs are ones with high transaction or user volume and large memory use. Typically, businesses will know if Java garbage collection is getting in their way. Or they can easily analyze their system to find out. Some JVMs have more than one scheme to use to address garbage collection, depending on the parameters of the program. Over the years, administrators have had to become adept at managing heap sizes and setting configuration parameters to attempt to minimize the effects of garbage collection. Sometimes they have been successful, and sometimes not.

IBM delivers advancements in Java garbage collection

What if you could significantly improve the service delivery of your enterprise Java applications? Well, the new IBM z14 mainframe and the IBM LinuxONE enterprise Linux server give you that option.

With all the IBM z14, IBM LinuxONE Emperor II and IBM LinuxONE Rockhopper II models, IBM has introduced a unique hardware feature that's designed to greatly reduce garbage collection pause times, and significantly improve response times for workloads that do a lot of garbage collection. This feature isn't available on x86 servers. The feature is called the Guarded Storage Facility (GSF) and is enabled by specifying an additional option (-Xgc:concurrentScavenge) to the JVM. The GSF essentially enables garbage collection to run almost completely concurrently with the application threads, thus dramatically reducing the pause times and allowing the workloads to run much more smoothly.

Comparing the IBM z13[®] mainframe without GSF and the IBM z14 Models M01-M05 with GSF, an IBM team showed that pause times were reduced by a factor of up to 10 times. In addition, throughput was improved by a factor of up to three times when the service-level agreement (SLA) was held at a common response time. The test used a large Java e-commerce workload running on the IBM z/OS[®] operating system.

Java applications run faster on IBM mainframe servers

The IBM Competitive Project Office team was asked to take this test one step further and compare the pause times, response times and throughput for the same test workload running on Linux on z14 and IBM LinuxONE servers versus running on Linux on an x86 server. The team tested using both an IBM z14 Model M03 (equivalent to an IBM LinuxONE Emperor II), and a new IBM z14 Model ZR1 (equivalent to the new IBM LinuxONE Rockhopper II), both with GSF enabled. The x86 server used was a Lenovo SR650. The results were impressive.²

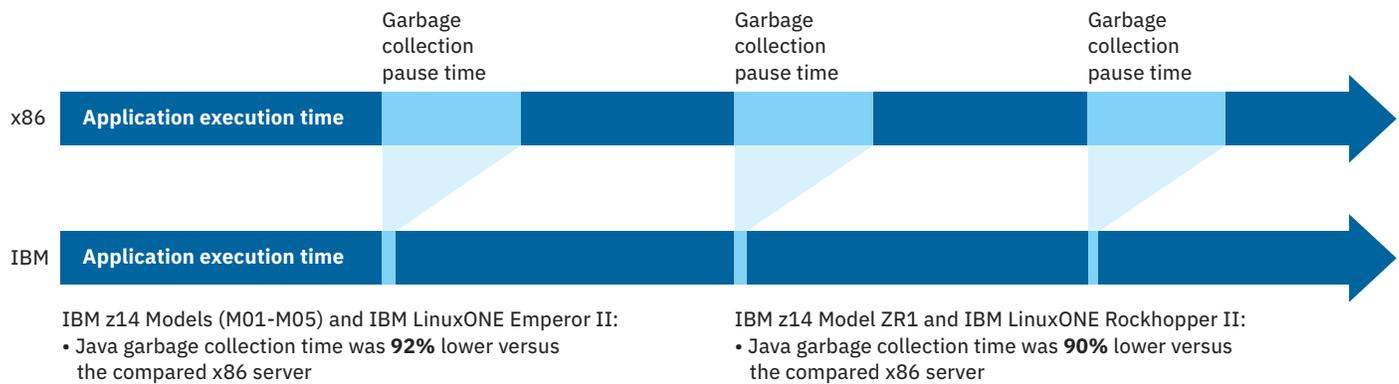


Figure 1: Under internal tests, Java garbage collections times were significantly reduced.

IBM z14 Models M01-M05 (IBM LinuxONE Emperor II):

- At a common response time, IBM z14 Models M01-M05 delivered 3.8 times more throughput versus the compared x86 server.
- At a similar throughput rate, the IBM Z[®] mainframe delivered a response time 92 percent lower versus the compared x86.

IBM z14 Models ZR1 (IBM LinuxONE Rockhopper II):

- At a common response time, IBM z14 Models M01-M05 delivered 3.0 times more throughput versus the compared x86 server.
- At a similar throughput rate, the IBM Z mainframe delivered a response time 80 percent lower versus the compared x86.

While the Guarded Storage Facility may not impact all Java workloads, IBM strongly recommends you examine your applications to understand if Java garbage collection is an issue. You could see significant improvements in response times and throughout with this new feature.

For more information

If you have questions or wish to learn more about the IBM z14 mainframe and the IBM LinuxONE enterprise Linux server, contact your IBM representative or IBM Business Partner. To learn more about the Guarded Storage Facility feature, visit ibm.com/developerworks.



© Copyright IBM Corporation 2018

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America
September 2018

IBM, the IBM logo, ibm.com, IBM LinuxONE, IBM LinuxONE Emperor II, IBM LinuxONE Rockhopper, IBM Z, z13, z14, and z/OS are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

Intel and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary. It is the user's responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs. THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

¹ blog.gigaspace.com/amazon-found-every-100ms-of-latency-cost-them-1-in-sales

² The IBM z14 system included 8 Integrated Facilities for Linux (IFL) in one logical partition (LPAR), Red Hat Enterprise Linux (RHEL) 7.4 with Linux Guarded Storage Facility kernel patch, and Java 1.8 SR5. The x86 system (Lenovo ThinkSystem SR650, 2 Socket, 16 cores total, Intel Xeon Gold 6134 CPU @ 3.20 GHz) ran with 8 cores, no hypervisor, RHEL 7.4 and Java 1.8 SR5. Tests used a Java e-commerce application that demonstrated average garbage collection pause times of approximately 300 milliseconds when running without GSF enabled. Data was then collected again when GSF was enabled on the IBM z14 server. The response time SLA requirement was that 99 percent of responses had to be received within 25 milliseconds. For Model M01-M05 test, the throughput rates were 23,551 transactions per second (TPS) on the Intel server and 22,889 TPS on the IBM z14 server. For the Model ZR1 test, the throughput rates were 19,902 TPS on the Intel server and 19,652 TPS on the IBM z14 ZR1 mainframe server.



Please Recycle