

CASE STUDY: ORACLE AND TWITTER

How Oracle GraalVM Supercharged Twitter's Microservices Platform

**Basic Research on Computing Fundamentals Delivers
Once-in-a-Lifetime Performance Improvements**



Holger Mueller
Vice President and Principal Analyst

Copy Editor: Jim Donahue

Layout Editor: Aubrey Coggins

TABLE OF CONTENTS

- At a Glance 3
- The Company 4
- The Challenges 4
- The Solution 5
- The Impact..... 7
- The Takeaways..... 8
- The Recommendations 9
- Analyst Bio11
- About Constellation Research.....12

AT A GLANCE

Twitter runs one of the most visible social networks in the world. Its move to a microservices infrastructure has addressed scalability challenges it faced in the past. Because Twitter is a free service (apart from the premium features it offers to enterprise customers), cost-effectiveness is a necessity that must be balanced with uptime.

As such, the company is constantly looking into ways to increase availability of the platform while keeping an eye on costs. Twitter saw Oracle GraalVM, a language-independent compiler engine and virtual machine, and decided to try it. Average CPU savings for compiler innovation are in the 1-2 percent range, but using Oracle GraalVM, Twitter realized between 8 and 11 percent CPU savings, depending on the microservice ported. This enables rare, once-in-a-lifetime savings for compiler innovation.

Twitter

- **Headquarters:** San Francisco
- **2018 Revenue:** \$3.04 billion
- **No. of Employees:** 3,920
- **Industry:** Social media
- **Other:** Founded in March 2006 by Jack Dorsey (now CEO), Noah Glass, Biz Stone and Evan Williams

Business Theme



Technology Optimization

AT A GLANCE

Problems	<ul style="list-style-type: none">• A leading, free social service that, because it's free, needs a cost-efficient platform.• Existing Scala C2 compiler is known for its inefficiencies.
Solutions	<ul style="list-style-type: none">• Move to Oracle GraalVM compiler.• Gradually move microservices to the new compiler.
Benefits	<ul style="list-style-type: none">• Substantial, CPU utilization savings of 8-11 percent.• Creation of a new, language-agnostic platform.

THE COMPANY

Twitter is one of the most prominent social networks worldwide. Twitter shares usage numbers as monetizable Daily Active Usage (mDAU); the company reported 145 million mDAUs during its recent third-quarter 2019 earnings call. Year over year, those users have increased their mDAU use of Twitter by almost 17 percent. After many years of steep losses, Twitter has managed to have eight back-to-back quarters in the black, with a profit of \$37 million in the third quarter of 2019.

The combination of user growth, active use and drive to profitability has applied pressure on the Twitter infrastructure team to keep an eye on costs. Back in 2010, the engineering team realized that because of performance issues caused by massive use of Twitter during the FIFA World Cup, a rearchitecture of the platform and code was necessary. As a result, Twitter went from being one of the largest Ruby on Rails users to a Scala shop, from MySQL to Hadoop and Manhattan, and created its in-house remote procedure call service with Finagle. Effectively, Twitter has moved to a microservices infrastructure, and it has seen substantial gains in scalability and performance.

Twitter has shown substantial reliability across different percentile distributions. The number of physical machines involved in serving the Twitter site has decreased between 5 and 12 percent.

Having moved to a reliable and scalable new platform, Twitter is looking for other strategies to reduce costs—without changing the codebase. Touching code that is working and scaling well is seldom a good idea in software engineering because of the costs and risks.

THE CHALLENGES

Because its service is free to consumers, Twitter is always under pressure to reduce costs. The more efficient Twitter's infrastructure is, the better for Twitter's profitability. And because financial analysts value profitability, Twitter must focus on cost efficiency in its operations.

At the same time, Twitter—in its role as an advertisement service—cannot skimp on uptime and reliability. When Twitter is down, no ads can be served and no advertisement revenue flows. With advertising representing more than 80 percent of total revenue for Twitter, any downtime has immediate implications on the bottom line.

Beyond immediate cost pressures, Twitter also needs to account for its future viability as a platform and its attraction as an employer to developers. The future of social media is not cast in stone. That means different services and media may be needed a few years from now. At the same time, information science graduates often change their preferences for what development languages they want to use in their first professional jobs. Therefore, a more language-agnostic platform would benefit Twitter in both investment security and talent acquisition.

THE SOLUTION

Given these three challenges—cost pressures, future-proofing and talent acquisition—Twitter staff engineer Chris Thalinger decided to try the Oracle GraalVM compiler. He was familiar with the GraalVM project from previous work experience and believed it would address Twitter’s use case. Oracle had created Oracle GraalVM as a new compiler written in Java, provided it first in JDK 8 and allowed for Java- and any Java bytecode-compatible languages to be executed. These include Scala, in which almost all of Twitter’s microservices are written.

But Oracle GraalVM was new software and, as such, faced scrutiny regarding quality. Thalinger was surprised to find only a few issues, of which only two resulted in more work to be fixed. All were resolved swiftly together with the Oracle team.

When Thalinger put Oracle GraalVM to a test in its own cluster, based on `jvmci-0.30` and `GraalVM-vm-0.22`, he saw about 11 percent lower CPU consumption to handle service requests compared to running the same microservice in the C2 compiler. The result seemed almost too good to be true because compiler improvements and changes usually yield 1 percent or sometimes 2 percent.

For the services tested, it was clear that Oracle GraalVM beat the popular but venerable C2 compiler mainly due to better inlining (the execution of functions in one line of source code) and better escape analysis (determining the dynamic scope of pointers).

After additional validation, Thalinger faced the ultimate challenge implementing GraalVM: He had to convince the owners of the microservices to implement them in production.

Not surprisingly, the owners were concerned about changing compilers, but after some prodding, a few gave it a try. As they showed positive results, more and more microservice owners have switched over to Oracle GraalVM, to the point that most of Twitter runs on Oracle GraalVM today.

Of course, Twitter started with smaller microservices to test the performance and stability of Oracle GraalVM but quickly gained the confidence to move the quintessential Tweet microservices to Oracle GraalVM. As of summer 2019, between 30 and 40 major services of Twitter—for instance, Tweet, Social and News—have been running in production powered by the Oracle GraalVM compiler.

Moreover, this is likely just the beginning of benefits that Twitter will gain from Oracle GraalVM because further optimization is possible. Early experiments have shown a potential total savings of 13 percent CPU time compared with the baseline.

Any good technologist reading this case study should wonder how the Oracle GraalVM compiler has done so much better than the C2 compiler. Apart from the benefits of being a recent R&D investment and being a modern software product, the main performance drivers are in the area of better inlining and escape analysis. Inlining refers to the ability of a compiler to process single-line functions. Escape analysis involves optimizing the dynamic scope of pointers. With Java- and Java bytecode-compatible code, using inline functions and pointers, a more efficient implementation of these basic compiler functions provided by Oracle GraalVM had significant performance gains. That they would be around 10 percent is a very positive and welcome surprise.

The Technology

- Oracle GraalVM

“Being able to see 8–11 percent performance improvements on the same hardware without having to change the underlying code is a once-in-a-lifetime event, leading to substantial cost savings and future flexibility for Twitter.”

—Chris Thalinger,
Staff Engineer, Twitter

Finally, Oracle GraalVM is also being developed further by Oracle Labs. More optimization and performance improvements are not guaranteed but are likely to come because the product is in its very early life cycle.

THE IMPACT

Twitter is seeing substantial benefits from using Oracle GraalVM, as expected on the hardware side, in tangible cost savings, but also longer term in terms of platform flexibility and talent acquisition.

- **Server savings.** Twitter's hardware and IT costs are proprietary and cannot be shared. But with GraalVM, Twitter can run about 10 percent more load on the same infrastructure. This is a remarkable result that normally can be achieved only with changes to the underlying code. The impressive results delivered by Oracle's GraalVM compiler, applied to the average CPU cost of a server run by a public cloud provider, is about \$127 per CPU/year. For enterprises operating hundreds or thousands of servers, this quickly results in substantial savings. The remarkable part is that these savings come without touching a line of code, ensuring that functional capabilities are the same and no retesting (for instance, by the business experts) is required.
- **Platform benefits.** By moving to Oracle GraalVM, Twitter gains the flexibility to create capabilities beyond its current assets that are written in Scala. Future requirements may create the need to support other languages that are more suitable to the use case to be addressed. In the fast-paced world of social media, acquisitions happen often, and integrating new capabilities with existing ones is a key success factor. With the multilanguage support in Oracle GraalVM, Twitter gains substantial platform freedom and flexibility. Time to market of in-house-built and potentially acquired new capabilities will be shorter while keeping the economies of scale in place that a common platform provides.
- **Talent acquisition gains.** If there is one thing that is certain about the future, it's that recruiting developers will get increasingly difficult, especially in Silicon Valley. At the same time, the programming languages of choice are changing fast. What used to

be C became Java and is now Python. Running on a platform that supports multiple languages adds substantial benefits because Twitter will be able to tap into a larger talent pool for acquiring new employees. At the same time, the company likely won't have to retrain new employees because they are able to program in the language they already know best.

THE TAKEAWAYS

Lessons Learned

There are three key lessons to be learned from Twitter's use of Oracle GraalVM.

- 1. Challenge conventional thinking.** Platform limitations are traditionally responded to by rewrites of the application logic to take advantage of a platform. Twitter moved off Ruby on Rails and MySQL with a rewrite in Scala, creating key platform elements such as Finagle and Manhattan while gaining substantial reliability and performance advantages. The adoption of Oracle GraalVM, on the other hand, lowered Twitter's CPU utilization around 10 percent, usually not feasible without changing code and platform. CxOs must realize that substantial performance gains can also be reached by simply changing compilers.
- 2. Realize it is new technology.** As with all new products, caution is critical and Twitter did well by first testing Oracle GraalVM. That there were only few issues that could be overcome quickly is a good sign and should give CxOs looking for benefits from Oracle GraalVM some confidence. But prudence is key for all engineering success, and Oracle GraalVM needs to be tested thoroughly before you unleash the compiler into production.
- 3. Invest into future-proofing your architecture.** Platforms and technology move faster than ever before, but market and competitive requirements are moving even more quickly. Investing in future-proofing a platform against potential and likely eventualities is good technology stewardship for any enterprise. It's up to the most senior CxO in the room to provide the funding, time and discipline for an enterprise to future-proof its platform with moves such as evaluating and potentially adopting offerings like Oracle GraalVM.

ROI

As with any technology move, the adoption of Oracle GraalVM has quantitative and qualitative benefits. On the quantitative side, average CPU savings of 10 percent result in substantial server savings—or additional capacity. Of course, savings will vary by application, platform, language and utilization. But given that no code needs to be touched, savings in this range justify an isolated functional and then performance test.

On the qualitative side, using a polyglot compiler like Oracle GraalVM is a key step toward future-proofing a platform, both in terms of load flexibility in the short run and, more strategically, regarding programming language artifacts running on the same infrastructure and server fabric.

THE RECOMMENDATIONS

Sometimes, technology offerings that didn't take off in the past suddenly become relevant. Historically, polyglot compilers have seldom worked too well. They usually show a performance penalty for the portability and typically have been limited in the number of supported programming languages. But once a technology giant like Oracle tries to redefine the polyglot compiler, enterprises should pay attention. And while Oracle Labs didn't know if the GraalVM project would work out—with a goal of establishing no performance penalty compared with native compilers—it is now clear that it succeeded.

With this background, CxOs and architects should start paying attention to Oracle GraalVM, especially when they run Java and Java bytecode-compatible languages. When a company with the demands of Twitter can find great success with Oracle GraalVM, the initial scalability and quality concerns that typically surround new software are much less relevant. Therefore, it comes back to justifying a trial and, if there are positive outcomes, a full scalability test. With the average CPU savings of 10 percent seen at Twitter, it's easy to justify that trial. Then it comes down to the savings seen, and a decision must be made whether they warrant a full scalability test. If they do, the next step involves starting production-near environments or even trying some production loads—always, of course, with the appropriate caution for any substantial and vital platform changes.

This case study puts any CxO who runs large Java- and JVM bytecode-compatible loads on notice: They are likely running their next-generation applications at a higher total cost of ownership than they could and should.

Moving to Oracle GraalVM also has long-term benefits because it opens software creation to additional programming languages, insulating enterprises in regard to running heterogeneous code ecosystems and enabling them to hire programmers proficient in other languages.

ANALYST BIO

Holger Mueller

Vice President and Principal Analyst

Holger Mueller is vice president and principal analyst at Constellation Research, providing guidance for the fundamental enablers of the cloud, IaaS, PaaS, with forays up the tech stack into big data, analytics and SaaS. Holger provides strategy and counsel to key clients, including chief information officers (CIO), chief technology officers (CTO), chief product officers (CPO), investment analysts, venture capitalists, sell-side firms and technology buyers.

Prior to joining Constellation Research, Holger was VP of products for *NorthgateArinso*, a KKR company. He led the transformation of products to the cloud and laid the foundation for new business-process-as-a-service (BPaaS) capabilities. Previously, he was the chief application architect with *SAP* and was also VP of products for *FICO*. Before that, he worked for *Oracle* in various management functions—both on the application development (CRM, Fusion) and business development sides. Holger started his career with *Kiefer & Veitinger*, which he helped grow from a startup to Europe's largest CRM vendor from 1995 onwards. Holger has a Diplom Kaufmann from University of Mannheim, with a focus on Information Science, Marketing, International Management and Chemical Technology. As a native European, Mueller speaks six languages.

 [@holgermu](https://twitter.com/holgermu)  constellationr.com/users/holger-mueller  linkedin.com/in/holgermueller

ABOUT CONSTELLATION RESEARCH

Constellation Research is an award-winning, Silicon Valley-based research and advisory firm that helps organizations navigate the challenges of digital disruption through business models transformation and the judicious application of disruptive technologies. Unlike the legacy analyst firms, Constellation Research is disrupting how research is accessed, what topics are covered and how clients can partner with a research firm to achieve success. Over 350 clients have joined from an ecosystem of buyers, partners, solution providers, C-suite, boards of directors and vendor clients. Our mission is to identify, validate and share insights with our clients.

Organizational Highlights

- Named Institute of Industry Analyst Relations (IIAR) New Analyst Firm of the Year in 2011 and #1 Independent Analyst Firm for 2014 and 2015.
- Experienced research team with an average of 25 years of practitioner, management and industry experience.
- Organizers of the Constellation Connected Enterprise—an innovation summit and best practices knowledge-sharing retreat for business leaders.
- Founders of Constellation Executive Network, a membership organization for digital leaders seeking to learn from market leaders and fast followers.

 www.ConstellationR.com

 [@ConstellationR](https://twitter.com/ConstellationR)

 info@ConstellationR.com

 sales@ConstellationR.com

Unauthorized reproduction or distribution in whole or in part in any form, including photocopying, faxing, image scanning, e-mailing, digitization, or making available for electronic downloading is prohibited without written permission from Constellation Research, Inc. Prior to photocopying, scanning, and digitizing items for internal or personal use, please contact Constellation Research, Inc. All trade names, trademarks, or registered trademarks are trade names, trademarks, or registered trademarks of their respective owners.

Information contained in this publication has been compiled from sources believed to be reliable, but the accuracy of this information is not guaranteed. Constellation Research, Inc. disclaims all warranties and conditions with regard to the content, express or implied, including warranties of merchantability and fitness for a particular purpose, nor assumes any legal liability for the accuracy, completeness, or usefulness of any information contained herein. Any reference to a commercial product, process, or service does not imply or constitute an endorsement of the same by Constellation Research, Inc.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold or distributed with the understanding that Constellation Research, Inc. is not engaged in rendering legal, accounting, or other professional service. If legal advice or other expert assistance is required, the services of a competent professional person should be sought. Constellation Research, Inc. assumes no liability for how this information is used or applied nor makes any express warranties on outcomes. (Modified from the Declaration of Principles jointly adopted by the American Bar Association and a Committee of Publishers and Associations.)

Your trust is important to us, and as such, we believe in being open and transparent about our financial relationships. With our clients' permission, we publish their names on our website.

San Francisco | Belfast | Boston | Colorado Springs | Cupertino | Denver | London | New York | Northern Virginia
Palo Alto | Pune | Sacramento | Santa Monica | Sydney | Toronto | Washington, D.C.