

Elasticsearch Runs Up to 1.58x Faster on AWS M6i Instances than AWS M6g Instances

Third-party benchmarking reveals that Amazon Web Services (AWS) cloud instances with 3rd Gen Intel® Xeon® Scalable processors outperform instances with AWS Graviton2 processors.

Considering migration to accelerate Elasticsearch performance?

Update your application with Intel and get better performance than the competition.¹

"To better handle heavy workloads, add more nodes to your cluster or upgrade your existing nodes to increase capacity."

— Elastic²

Up to
1.58x
faster

Elasticsearch workloads in Amazon Web Services (AWS) instances with Intel significantly outperform instances with Graviton.¹

Are you getting all the speed that Elasticsearch offers?

Every day, social media sites, news outlets, and other cloud-native companies enter fierce competitions for market share. The rules are simple—deliver the most information and be the fastest. Oftentimes, the winners at the end of the day are using Elasticsearch to deliver answers in milliseconds.^{3,4}

When Elasticsearch handles too much data without adequate compute power, latency can become an issue.² This workload latency could be caused by provisioning too few vCPUs or by opening cloud instances with under-powered CPUs.

You can win at the data-delivery game with smart planning and cloud resourcing best practices. As a starting point, select hardware that helps Elasticsearch run at top speed. In a [study](#) sponsored by Intel, Prowess Consulting conducted workload testing to evaluate hardware performance. The results of the study reveal that Elasticsearch workloads run up to 1.58x faster using AWS instances with Intel compared to instances with Graviton2.¹

Using a best-in-class cloud approach

Prowess Consulting's methodology helped ensure that the workload testing emulated the conditions of production deployments and that the results could be applied to real-world conditions in the cloud. To help understand the impact of processor family selection, testing was performed on AWS's M6i instances with 3rd Gen Intel Xeon Scalable processors, compared against M6g instances with AWS's ARM-based Graviton processor. The engineers ran tests using application workloads instead of synthetic benchmarks and incorporated network-security practices, such as user authentication and Transport Layer Security (TLS)/Secure Sockets Layer (SSL) encryption.

The results demonstrate how your choice of hardware can significantly improve Elasticsearch performance. The study tested two popular cloud-computing architectures built on 3rd Gen Intel Xeon Scalable processors and AWS Graviton2 processors. To emulate production runtime conditions, the application was configured to use the latest long-term support (LTS) release of Java Development Kit (JDK), JDK 17. The cloud instance used for testing was the Amazon Elastic Compute Cloud (Amazon EC2) M6 instance family, an inherently stable instance that supports a variety of general-purpose workloads. The Elasticsearch workload was benchmarked using three M6 instance sizes (see Figure 1), and the results verify that performance gains occur independent of instance size.¹




Figure 1. Elasticsearch workloads running in M6i instances with Intel outperformed M6g instances with Graviton across all instance sizes¹

The right hardware for faster cloud workloads

After testing Elasticsearch across three instance sizes, the results showed that AWS instances with Intel consistently outperformed instances with Graviton. The m6i.16xlarge instance size showed the greatest improvement—outperforming the m6g.16xlarge instance by up to 1.58x. The m6i.4xlarge instance size was 1.25x faster than m6g.4xlarge, and the m6i.xlarge instance was 1.03x faster than m6g.xlarge.¹

Among the main reasons you use Elasticsearch is for its fast access and smooth handling of massive volumes of data. Applying lab results to the real world, you could use an accelerated Elasticsearch search engine to quickly troubleshoot and resolve messaging failures across a global network of hundreds of millions of users.⁵

Conclusion

The cloud is a noisy place, with lots of variables that can impact workload performance. Use these benchmarking results to find clarity and manage Elasticsearch for exceptional performance. Win over your customers with data-delivery leadership in AWS instances with Intel.

Learn more

Read the technical research report, “[Java-Based Benchmarking Shines a Light on How Underlying Architecture Impacts Cloud Performance](#).”

Get great [performance, cost, and flexibility](#) from cloud computing with Intel processors in AWS.

Learn how Intel Xeon Scalable processors can unlock the potential of your [cloud computing](#).



¹ Normalized performance results based on latency testing. Prowess Consulting. “[Java-Based Benchmarking Shines a Light on How Underlying Architecture Impacts Cloud Performance](#).” Commissioned by Intel. 2022. See endnote 4 and Table 2 for details.

² Elastic. “[Troubleshooting: Fix common cluster issues—high CPU usage](#).”

³ ZDNet. “[Elastic, search company for Uber and Tinder, nearly doubles in IPO](#).” October 2018.

⁴ 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](#).”

⁵ Towards Data Science. “[An Overview on Elasticsearch and its usage](#).” March 2019.

See our complete [Legal Notices and Disclaimers](#).

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

0423/TB/PRW/PDF