

Case Study

Intel® Software Center of Excellence
Intel® Xeon® Processors



Intel® Software Center of Excellence Significantly Scales GoPay Payment Processor Performance on Google Cloud

Optimizations for the Barito logging platform deliver up to 15% higher throughput on Google Cloud with 2nd Gen Intel® Xeon® processors and up to 87% performance improvement with 3rd Gen Intel® Xeon® processor-based instances.¹

Solution Summary

- Google Cloud N2 instances with 3rd Gen Intel Xeon processors
- Intel optimizations for Barito logging platform
- Optimizations by Intel® Software Center of Excellence



Executive Summary

GoPay, a part of GoTo Financial, is one of the subsidiaries of GoTo Group, the largest digital ecosystem in Indonesia. GoTo uses the Barito logging platform running on Google Cloud to monitor and log all transactions. The Barito platform presented a bottleneck for scaling up transactions for GoPay. The company turned to Google Cloud and the Intel® Software Center of Excellence (CoE) to explore performance improvement opportunities for the platform.

Developers in the Software CoE recommended optimizations for Barito that resulted in an immediate improvement in Barito CPU utilization and an increase in transaction capacity. The recommendations also included benchmarking Barito on N2 instances with 3rd Gen Intel® Xeon® processors to understand how capacity could be further improved. These optimizations allowed GoTo to increase throughput without adding cluster costs on existing N1 instances and to significantly scale capacity using newer Google Cloud N2 instances.

Challenge

[GoPay](#) is an e-wallet based in Indonesia. GoPay, a part of GoTo Financial, is one of the subsidiaries of [GoTo](#) Group, the largest digital ecosystem in Indonesia. Using the GoPay or [Gojek](#) app available on Apple and Android devices, users can make transactions with any merchants that accept GoPay as a payment method.

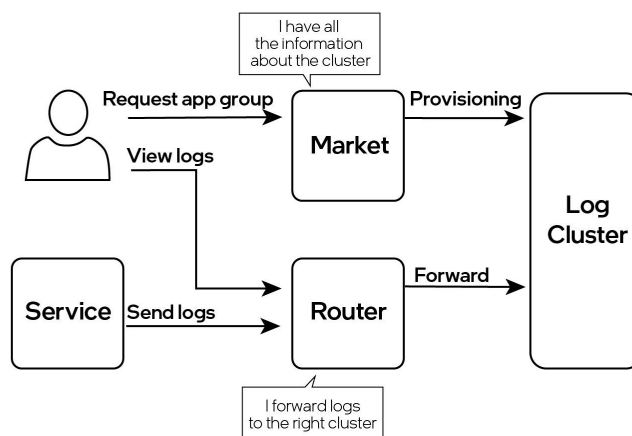


Figure 1. The Barito logging system logs activities across all Goto properties, including GoPay.

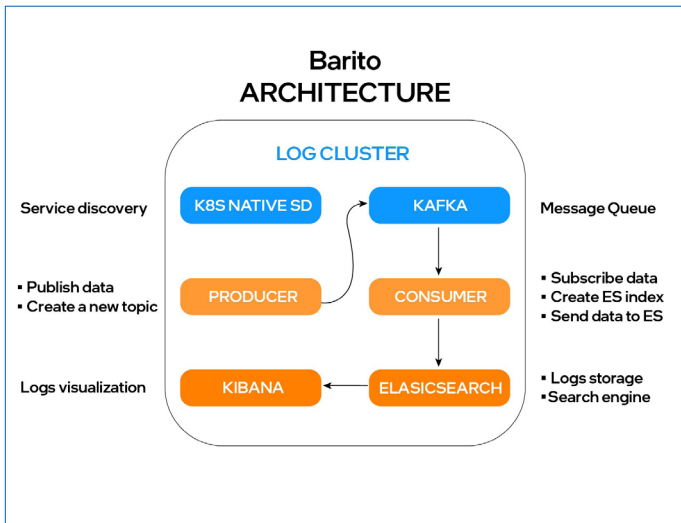


Figure 2. The Barito logging platform integrates several applications and tools.

An intelligent logging platform called **Barito** (Figure 1) logs transactions across all GoTo’s properties. Barito is critical to maintaining strong and efficient operations across GoTo’s services.

The Barito cluster integrates several tools and applications, including Apache Kafka, Kibana, and other tools to produce, track, search, and visualize all transactions as they happen (Figure 2). GoTo’s IT and customer support teams can access the logs to monitor services and troubleshoot issues for the company and customers. Barito runs on clusters of N1 instances on Google Cloud using Intel® Xeon® processors.

With over thousands and more of transactions coming every minute, the need to get better visibility and monitoring real-time was crucial. GoPay needed more processing or CPU headroom to process more transactions. Rather than scale instances, which would lead to higher costs, GoTo engineers turned to Google Cloud and the Intel® Software Center of Excellence (CoE) to help isolate issues and offer solutions to improve performance.

“We’re excited about this collaboration with the Intel CoE program. In the short term we’re able to identify and improve CPU utilization by up to 10 percent for one of our components. We are looking forward to working closely so we can achieve more optimal utilization and improvement in the future.”

*—Eka Risky, Principal Engineer, Payments Infrastructure
GoTo Financial*

Solution

The **Software Center of Excellence (CoE)** is a service of Intel in collaboration with Google Cloud. The CoE helps customers running Google Cloud instances based on Intel architecture and technologies to optimize performance of their workloads.

Intel Software Center of Excellence (CoE) Overview

Intel Software CoE’s engineers are experts in Intel optimization tools and the technologies and capabilities of Intel processors, chip sets, and other Intel products. Working with Google, these developers have helped Google Cloud customers improve performance by modifying code to optimize for parallelism, reduce time spent in tasks, or take advantage of Intel technologies, such as the built-in Intel® Accelerator Engines integrated in all 4th Gen Intel® Xeon® processors. The CoE services enhance the value customers get from their Intel Xeon processor-based instances. As a Software CoE customer, users get:

- Direct access to Intel engineers
- Guidance for improving the price performance and the operational performance of Intel-based Google Cloud instances
- Code-level recommendations for software running on Intel processors so customers can attain the most benefit possible from their Google Cloud investments

An engagement with the Software CoE includes a Performance Report with optimization recommendations, an action plan, and an implementation plan for customers to execute. Recommendations for potential support from Google Cloud Professional Services Organization (PSO) can also be included.

Barito Analysis

Developers in the CoE analyzed Barito flame graphs provided by GoPay (Figure 3) to understand where latencies occurred during the execution of services. Developers discovered that the Prometheus data scraper feeding Kafka consumed approximately 22 percent of available CPU cycles. This seemed high for the workload.

The Software CoE Performance Report recommended two actions:

1. Decrease the Prometheus interval by half, which would reduce CPU utilization for Prometheus.
2. Test Barito on Google Cloud N2 instances with 3rd Gen Intel Xeon processors.

For testing Barito performance between N2 and N1 instances, the Barito software stack was deployed on Google Kubernetes Engine (GKE) clusters with the client running on N1 instances. Locust was used to drive traffic to the GKE




Figure 3. The Intel Software CoE analyzed Barito flame graphs to uncover performance enhancement recommendations.

clusters with a varying load to study the characteristics and response of the workload. The usercount was set to a rate of 5000 at 40 users per second. The length of the logs was set at 1100, and Wikipedia was the text source for benchmarking.

Result

For the optimizations on N1 instances, by cutting the Prometheus scraping interval by half, Kafka CPU utilization dropped significantly by 10 percent.² Throughput of Barito logging increased by up to 15 percent,² allowing more transactions to occur.

For the N1 vs N2 benchmarks, the higher performance 3rd Gen Intel Xeon processor-based instances enabled the following improvements:

- 20 percent overall reduction in CPU usage
- Fewer nodes needed—from 186 N1 to 150 n2-standard-32 nodes

With N2 instances, Intel showed 87 percent more responses per second compared to N1—nearly twice the response rate of N1 (Figure 4). This allowed 50 percent more users to be added with requests being processed nearly 10 percent faster, while using fewer Google Cloud resources.³

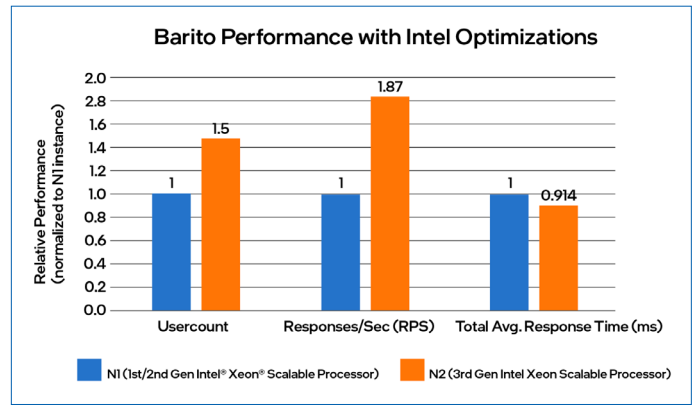
These optimizations and migrating to N2 instances with a smaller overall configuration due to needing fewer nodes allowed GoTo to improve Barito logging throughput **cut costs about 9.8 percent.**

The benchmarks resulted in low CPU utilization on both worker nodes of the GKE cluster, suggesting GoPay could potentially reduce spending by moving to lower core count Google Cloud instances on the worker nodes, while maintaining transaction quality to customers. Intel also recommended GoPay to upgrade Java, Golang, and Ruby to the latest versions to get the best performance from Intel architecture for their Barito software stack.

Solution Summary

GoPay transaction throughput suffered due to the Barito logging system. Barito runs on Google Cloud N1 instance clusters. Working with the Intel Software Center of Excellence, GoPay implemented the optimization recommendations of the Software CoE developers, resulting in an immediate significant improvement in CPU utilization. The optimizations allowed GoPay to improve transaction throughput without upgrading their cluster instances or increasing vCPU count.

Software CoE developers also recommended benchmarking Barito on N2 instances with 3rd Gen Intel Xeon processors. The benchmark revealed the N2 instances could support up to 50 percent more users with 87 percent more responses per second, while cutting average response time by nearly 10 percent.² These improvements could all be achieved on 36 fewer instances, potentially reducing costs for running Barito.



The higher performance 3rd Gen Intel Xeon processor-based instance had nearly twice the response rate compared to the 2nd Gen Intel Xeon processor-based instance.

Where to Get More Information

Learn more about [Intel® Xeon® processors](#).

Discover more about the [Intel® Software Center of Excellence](#).

Find out more about [GoPay](#).

Solution Ingredients

- Google Cloud N2 instances with 3rd Gen Intel Xeon processors
- Intel optimizations for Barito logging platform
- Optimizations by Intel Software Center of Excellence



¹ Performance improvement achieved by a combination of utilizing the Google N2 instances and implementing the optimizations that the Intel Software COE recommended. Testing by GoPay as of 6/20/2023.

² Performance claim provided by GoPay, who achieved this by cutting the Prometheus scraping interval by half. Testing by GoPay as of 6/20/2023.

³ According to GoPay, the higher performance 3rd Gen Intel Xeon processor-based instances enabled the 20 percent overall reduction in CPU usage. Fewer nodes were needed—from 186 N1 to 150 n2-standard-32 nodes. With N2 instances, Intel showed 87 percent more responses per second compared to N1. This allowed 50 percent more users to be added with requests being processed nearly 10 percent faster, while using fewer Google Cloud resources. Testing by GoPay as of 6/20/2023.

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

For workloads and configurations visit www.Intel.com/PerformanceIndex. Results may vary.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© 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.