

Case Study

Cloud
Workload Optimization



Nylas Realizes 35% Lower Google Cloud Infrastructure Costs with Granulate™ an Intel Company¹

Granulate uses autonomous, continuous optimization to improve cloud application performance and lower costs—with zero code changes and no additional R&D or maintenance efforts



“We couldn't believe how quickly we were able to realize real cost savings without having to change code, invest in research and development time, or go through service disruptions.”

—Christine Spang, Nylas cofounder and CTO

Nylas—the state-of-the-art communications platform giving developers universal access to email, calendar, and contacts providers through a single integration—empowers businesses to connect with, extract, and contextualize their own communications data to inspire innovative experiences for customers and employees. More than 150,000 developers worldwide use the Nylas platform to build productivity features—such as automated communications, smart scheduling, and contacts integration—into their applications.

Challenge: Rapid growth and a new customer base required a new cloud economics strategy

Since their founding in 2013, Nylas and their revolutionary approach to helping companies turn communications data into actionable insights through API-driven digital transformation have gained traction throughout the business world. Because of their success, Nylas not only experienced a rapid and dramatic increase in their traditional user base but also began seeing demand from new types of customers in the business-to-consumer (B2C) market.

“We were seeing steady growth in our more traditional user base, followed by an influx of requests from B2C-focused customers,” says Christine Spang, cofounder and CTO. “We were excited to have a new, growing customer base but soon realized their economic needs were drastically different than our existing B2B customers. Some of their requests were asking us to support a customer base that was 10x the size of our current user group! That’s when we knew we needed to make some significant changes to our platform to better serve our new customer base.”

The two legacy workloads at the core of the Nylas platform—the sync engine workload that constantly pulls and reprioritizes customer messages and the API workload that handles thousands of user requests per second—were costly and unable to scale to meet increased user demands. These challenges forced the Nylas team to begin several resource-intensive initiatives in an effort to optimize their workloads, including:

- Optimizing their existing infrastructure to realize cost savings
- Rewriting legacy workloads in Go and migrating them to a new provider
- Building an entirely new infrastructure for new and existing users

Solution: Granulate an Intel Company

Knowing the workload optimization initiatives would require their developers to focus on rewriting code—rather than working on core product development—the Nylas team began exploring alternative options for boosting performance. Having used Granulate in the past, the Nylas team was interested in seeing if the solution could once again deliver performance gains without requiring valuable development resources. Nylas reached out to their Intel support team to get the Granulate software implemented in their cloud environment.

“We already had a dedicated Intel team that we worked with, so they had a deep understanding of our application architecture and business priorities,” says Spang. “This close relationship with Intel made getting the Granulate software implemented an extremely easy process.”

Intel installed Granulate on a Nylas platform utilizing a Google Cloud N2 instance, and it began passively learning data flows, resource contentions, and processing patterns. Next, Intel activated the Granulate agents, and the program immediately started tailoring resource scheduling and prioritization decisions to adapt to Nylas’s service—resolving inefficiencies and increasing performance all while constantly adapting to changing traffic and new deployments.

Results: Nylas achieves 35 percent lower Google Cloud infrastructure costs with no additional development efforts

Within days of the Granulate agent being installed and activated in the Google Cloud public cloud environment, Nylas saw significant results. “From the initial Granulate installation through the training and activation phases, it took only around 20 days for us to see results,” says Spang. “We couldn’t believe how quickly we were able to realize real cost savings without having to change code, invest in research and development time, or go through service disruptions.”

With the impressive 35 percent cost reduction realized from using Granulate in Google Cloud, Nylas now plans to pass some of those savings on to their user base, including their new B2C-focused customers.¹

“From the get-go, we knew our new customers had specific economic requirements they needed to meet for their consumers. With the help of the Granulate software, we are able to deliver on their needs,” says Spang. The team is also exploring other ways to replicate the impressive results gained from the Granulate platform. They are identifying other workloads where performance can be optimized and costs can be reduced.



Granulate™ an Intel Company: Results for Nylas

↓35%
Google Cloud
infrastructure
cost reduction¹

Zero
code
changes

No
additional
R&D effort

How Granulate improves application performance and reduces costs for businesses

Granulate is an autonomous optimization software solution that loads into an existing data center or cloud environment and optimizes resource management without human intervention. Installed in minutes via a simple command line, the Granulate tool automatically learns a company’s resource usage patterns and the data flow of the application and operating system. The tool uses US-patented algorithmic models to identify instances of data bottlenecks and resource contention within workloads and then adjusts resource management decisions that accelerate data flow through an application.

This process involves adapting resource allocation at the operating system and runtime levels to continuously and autonomously optimize memory allocation, process swapping, thread scheduling, storage access, and network communications. Users perpetually experience performance increases as Granulate continuously recognizes usage patterns and optimizes each workload. Best of all, Granulate is fully autonomous and does not require adjustments or ongoing maintenance efforts, so there is very little overhead required for ongoing administration.

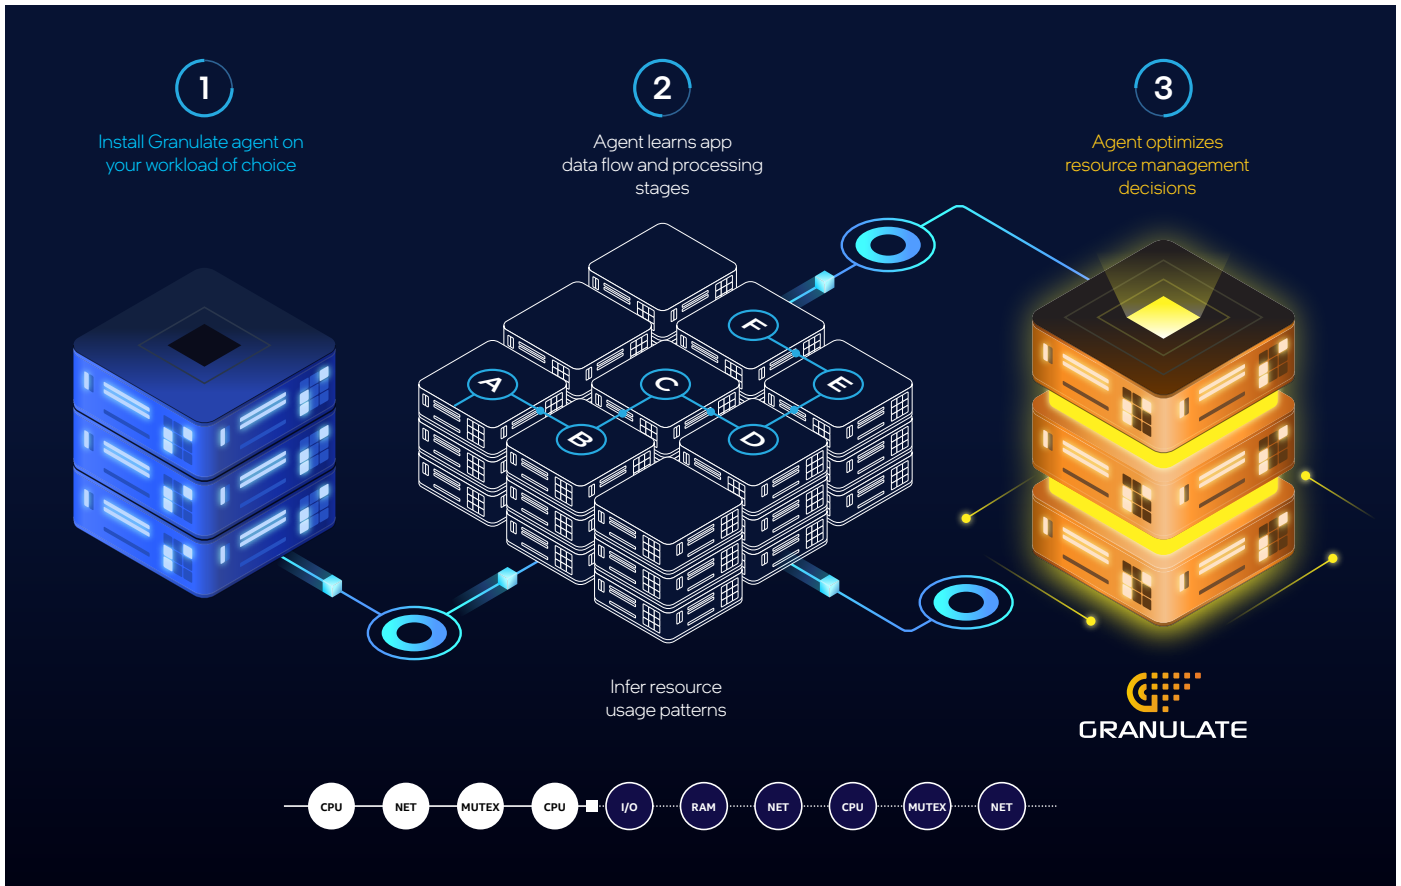


Figure 1: Granulate automatically identifies and resolves bottlenecks and resource contention within application data flow.

Potential benefits for businesses using Granulate



Time savings from seamless integration
 Granulate works with any cloud or on-premises environment and with containers or orchestrators like Kubernetes. Once deployed, the optimizer does not require any additional maintenance.



Continuous value from optimization
 After installation, Granulate will continuously identify new resource usage patterns as businesses make changes or updates to their workflows, resulting in ongoing performance improvements even as the applications change.



Greater agility for DevOps
 By freeing up resources that a DevOps team would otherwise spend on performance tuning and iterations, Granulate allows teams to focus on core product research and development.



Accelerated time to market
 Faster workload performance with fewer data bottlenecks and less resource contention allows businesses to develop new services and updates faster, shortening time to market and time to value.



Greater infrastructure stability
 Higher performance with fewer data bottlenecks allows businesses to achieve more-reliable uptime for their services, making it easier to meet service level agreements (SLAs).

Close collaboration and compute testing with Intel helps optimize Nylas's Google Cloud footprint

While Nylas explored the benefits of using the Granulate software for a second time, the team also worked alongside Intel to begin a large-scale, yearlong compute research and testing initiative that:

- Examined which cloud provider and instance configurations would provide the best possible experience to customers at the lowest cost
- Explored a way to further reduce latency for their application, which already handled millions of queries and messages in a small-time-frame window

Throughout the testing initiative, Nylas engaged with Intel experts to leverage their deep expertise in hardware, software, and cloud computing. Intel worked closely with the Nylas team to analyze their workloads, propose optimized instance configurations, and help them find the right path forward. However, the Nylas-Intel partnership was what made this testing initiative different than a normal collaboration.

“Our expectations of working with Intel were blown away with this project,” says Spang. “Intel was ‘in the trenches’ with us at every step of this journey. We learned just how deep their expertise in software and cloud goes and that they care about helping us grow and succeed more than anything.”

After more than a year of research and testing, Intel and Nylas were able to identify an ideal cloud provider and configuration that gave customers the best experience and cost possible while also improving the latency of Nylas's application.

[Learn more about Granulate >](#)

[Learn more about Nylas >](#)

[Register for a free trial of Nylas >](#)

[Learn more about cloud optimization tools from Intel >](#)



Notices and disclaimers

1. Source: Based on Nylas internal measurements Dec. 2022.

Intel is committed to respecting human rights and avoiding complicity in human rights abuses. See [Intel Global Human Rights Principles](#). Intel® products and software are intended only to be used in applications that do not cause or contribute to a violation of an internationally recognized human right.

Performance varies by use, configuration, and other factors. Learn more at [intel.com/PerformanceIndex](#).

Performance results are based on testing as of dates shown and may not reflect all publicly available updates. No product or component can be absolutely secure.

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

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

No product or component can be absolutely secure.

Your costs and results may vary.

© 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.

0323/ADS/CMD/PDF