


Intel® Granulate™ for GCP

Improve GCP performance to cut costs by up to 45%

Intel Granulate empowers Google Cloud Platform users with autonomous, continuous app-level performance optimization and capacity management, significantly reducing cloud costs while requiring no code changes.





GCP ready from day one

Intel Granulate is listed in GCP's commercial marketplace, seamlessly integrated from day one. Intel Granulate provides broad support for all GCP resources and services, regardless of compute type, Linux distribution, or development language.


[Find Intel Granulate in the Google Cloud Marketplace](#) >

Intel Granulate supports GCP by running on:

Easy installation with no maintenance required

-  **Simple setup**
 By entering just one line of code in the command line, organizations can manually install Intel Granulate in minutes. Standard provisioning tools such as Google Cloud Deployment Manager, Chef, Ansible, and Puppet are fully supported as well.
-  **No code changes or manual efforts**
 Intel Granulate monitors and then automatically and continuously updates resource allocation to reflect the application's needs - without human intervention, code changes, or R&D efforts.

Use Cases




Kubernetes orchestration and optimization on GKE

Gain full visibility into GKE clusters, seamlessly complement HPA scaling policies, and achieve your cost performance goals by applying custom rightsizing recommendations based on actual usage in production with Intel Granulate Capacity Optimization. This solution is deployed as a Daemonset by default and supports additional installation methods including Helm Chart, Docker, CLI, and more



Optimizing Big Data workloads on Dataproc

Process large data sets on Dataproc faster with autonomous and continuous optimization across various key aspects of Big Data workloads, including YARN resource allocation, Spark executor dynamic allocation, improved dynamic scaling, crypto and compression acceleration, memory arenas, and JVM runtime execution



Runtime Optimization on GCP services

Boost application performance with Intel Granulate to automatically optimize key runtime features and capabilities including thread scheduling, lockless networking, inter-process communication, connection pooling, congestion control, and memory arenas. Intel Granulate autonomously and continuously learns your application's specific resource usage patterns and data flow. The solution identifies contended resources, bottlenecks and prioritization opportunities by analyzing CPU scheduling order, oversubscribed locks, memory, network and disk access patterns.

Why Intel Granulate?

Intel Granulate supports GCP customers by offering a suite of cloud optimization solutions, running on GCE instances and GKE Kubernetes services, and Dataproc and Databricks big data workloads. Intel Granulate provides DevOps teams with optimization solutions for all major runtimes, such as Python, Java, Scala, Go, and more. Google Cloud customers are seeing improvements in their job completion time, throughput, response time, and carbon footprint while realizing **up to 45% cost savings**



45% Cost Reduction

CUSTOMER SPOTLIGHT

Nylas Achieves Over 35% Cost Reduction on GCP Leveraging Intel Granulate



"The product is awesome, we don't need to do anything and it just makes everything faster, it's super cool."

Caleb Geene,
Staff Site Reliability Engineer at Nylas

[Watch the Session](#) >