

Intel Granulate for AWS

Improve AWS performance to cut costs by 45%

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

45%
Costs Reduction

40%
Reduced Response Time


20%
Increase Throughput

0
Code Changes

AWS ready, out-of-the-box

Intel Granulate is a proud AWS Partner Network (APN) member. Intel Granulate is listed on the AWS commercial marketplace to deliver greater customer value while providing broad support for all AWS resources and services, regardless of compute type, Linux distribution, or development language.

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



Simple 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 AWS CloudFormation, EMR bootstrap, 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 EKS

Gain full visibility into EKS 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, and CLI.



Optimizing Big Data workloads on EMR

Process large data sets on EMR 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 AWS 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 AWS customers by optimizing the most popular compute services including EC2, EKS, ECS, EMR, and top AWS Partner Solutions such as Databricks, Cloudera, and RedHat. With Intel Granulate, AWS customers are seeing improvements in their job completion time, throughput, response time, and carbon footprint, while realizing **up to 45% cost savings**.



CUSTOMER SPOTLIGHT

Snap Inc achieves 13% average cost reduction on EKS workloads with Intel Granulate



"The ease of use and short time to value have been significant motivating factors for Snap to continue the expansion of Intel Granulate's Kubernetes optimization solution to more internal services at a rapid pace."

Tom Brown, Software Engineer at Snap Inc

[Watch the Session >](#)