Application Runtime Optimization

Powered By
GRANULATE
An Intel Company

Improving Performance, Reducing Cost.

The EVOTEK Application Runtime Optimization offering, powered by Granulate an Intel company, continuously and automatically learns an application's specific resource usage patterns and data flow to identify contended resources, bottlenecks, and prioritization opportunities. Then it optimizes on the runtime level to achieve performance improvements that lead to cost reduction.

Business Challenge

As more and more businesses shift their operations to the cloud, IT departments are facing a new set of challenges related to cloud optimization. One of the biggest challenges is managing costs with the saturation of available cost savings solutions. Additionally cost reduction initiatives require significant engineering and FinOps resources with long enablement processes. Finally, although workload and application performace optimization help drive cloud cost reduction, many of these solutions are not enablement or implemented at scale. With performance, security, and proper monitoring are now crucial factors regarding cloud optimization, it is important for businesses to stay on top of these challenges, and leverage the right tool to solve for the right problems.

A Proven Solution

EVOTEK and Granulate provide a unique solution giving businesses a tailored offering around real-time continuous optimization for workloads at scale. The EVOTEK Granulate Workload Optimization offerings is 100% environment & architecture agnostic, meaning it can run anywhere. With no configured, code or maintenance changes required, R&D efforts are a non-issue. In addition to being complimentary to other cost optimization methods, the installation is simple and deployable with a single command line.

5B OVER FIVE BILLION CORE HOURS SAVED

$200M OVER $200M CLOUD SPEND REDUCED

5X FIVE TIMES INCREASE IN THROUGHPUT

mobileye
An Intel Company

Mobileye deployed Granulate on an Elastic Compute Cloud (EC2) workload as container side-cars and began learning the workload, dataflow, and resource usage patterns within minutes. Following a few days of learning, the agents were ready to be activated to start optimizing the application's performance

200K CORES ON AWS

45% FASTER JOB COMPLETION

44% REDUCED COST FOR AWS CLOUD INSTANCES

1/2 THE TIME TO CREATE NEW MOBILEYE MAPS

0 ZERO CODE CHANGES WITH NO TUNING REQUIRED