IBM **Turbonomic**
for Google Kubernetes Engine

Optimization you can continuously automate to *prevent* performance risk & cost overruns.

Software (not people) continuously makes complex resourcing decisions to ensure all applications get exactly what they need to perform.

Improve application performance

Increase IT productivity

33% reduction in public cloud spend due to dynamic scaling and workload resizing[1]

75% improved infrastructure utilization and avoided annual refresh costs by 75%[1]

70% with understanding of application demand, avoided required infrastructure growth spend by 70%[1]

Automation unlocks the business value of Google Kubernetes Engine

Application-aware and full-stack. See how dynamic resourcing impacts app response time.

Multidimensional analysis ensures actions are actionable and trustworthy—safely automate!

Operationalize automation by integrating actions into pipelines, processes, and workflows.

[1]Forrester Total Economic Impact of IBM Turbonomic Application Resource Management

Confidently assure application performance on GKE

Increase deployment frequency

Accelerate the onboarding of more applications to the platform, IBM Turbonomic ensures apps continuously perform at the lowest cost.

Scale apps & infra based on SLOs

Horizontally scale microservice applications and the underlying cluster based on app response time or business SLOs. It's the smarter way to drive elasticity in the cloud.

Minimize manual labor

DevOps and SRE teams avoid the guesswork of rightsizing containers and setting pod (HPA) or infrastructure autoscaling policies. Instead, software correctly determines Kubernetes hundreds of cluster and configuration parameters.

Full-stack GKE optimization minimizes cloud cost

Optimization of GKE ensures apps and clusters are appropriately sized for performance, while minimizing cloud cost. IBM Turbonomic also optimizes persistent volumes and other volumes for IOps, ensuring continuous performance at the lowest cost across the stack.

Accelerating sustainable outcomes

Applications consume exactly what they need to perform, minimizing cloud cost and materially reducing your carbon footprint.

Turbonomic supports all upstream versions of Kubernetes.

Live sandbox at
ibm.com/products/turbonomic

Achieving real business outcomes requires continuous optimization to be automated at scale

ACTIONS

- **Container Rightsizing**
- **Pod Moves**
- **Cluster Scaling**
- **Container Planning**
- **Pod Scaling** (SLO Assurance)

OPERATIONALIZED

Integrate with any pipeline, IaC, ITSM, or communication tool in your organization!

- **Ansible**
- **GitHub**
- **GitLab**
- **Jenkins**
- **Puppet**
- **Slack**
- **Terraform**

...and more!

BUSINESS IMPACT

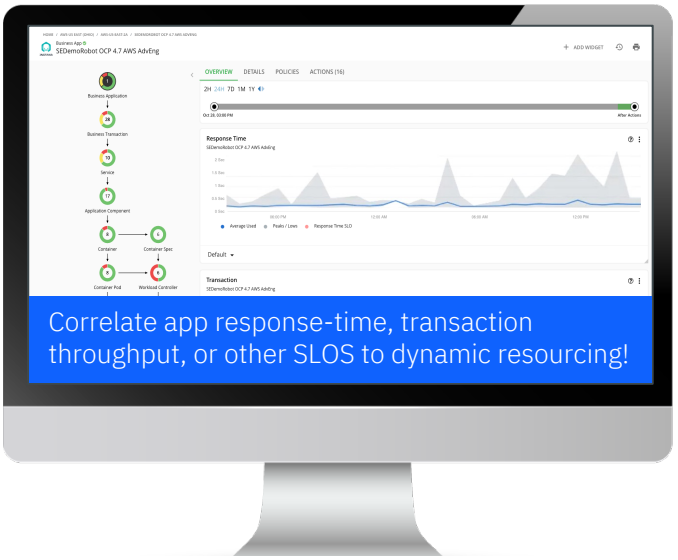

Build trust with AppDev by showing automation's impact on the client experience.

App-first, full-stack automation unlocks cloud native elasticity

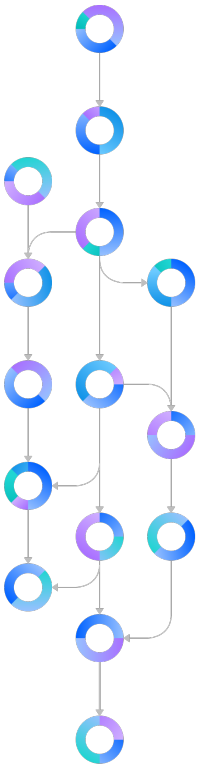

Application SLOs

- Application (horizontal pod) scaling
- ...based on response-time or other service level objectives

Google Kubernetes Engine

- Container rightsizing
- Pod moves
- Cluster scaling
- ...all based on application demand

Google Cloud IaaS

- Optimize persistent volumes and other volumes for IOps performance and cost

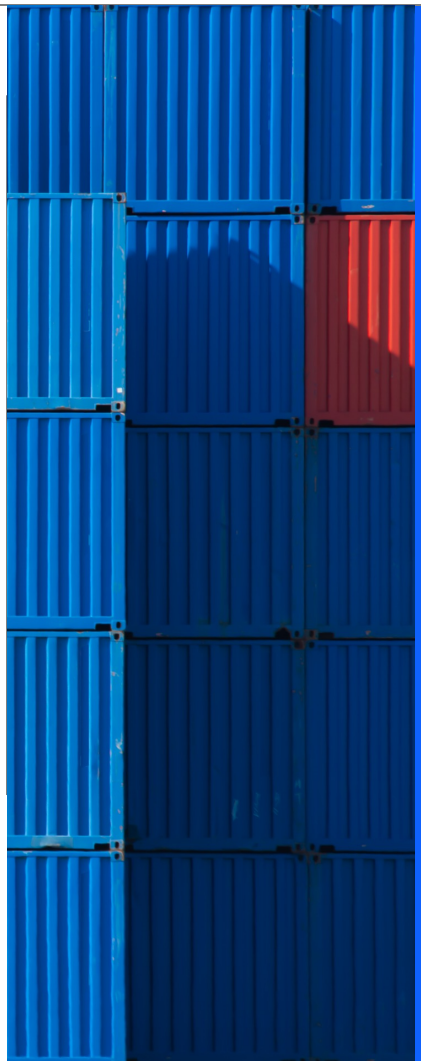

KEY DIFFERENTIATORS

- App-first, demand-based analysis ensures actions can be safely automated
- AVOID THROTTLING—The silent killer of response time! CPU throttling metrics inform CPU limit rightsizing
- Historical data for the life of the service
- Percentile analysis to protect peaks

IBM