

IBM Turbonomic

Application Resource Management
for Google Cloud

Cloud optimization you can continuously automate to prevent performance risk and 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 cloud spend
due to dynamic scaling and rightsizing¹

“We have lots of APM, management and monitoring tools but in terms of giving me the exact recommendations what do I need to do and what actions to take, there is nothing like that—except for Turbonomic. Humans are not sustainable for optimization, it requires Turbonomic.”

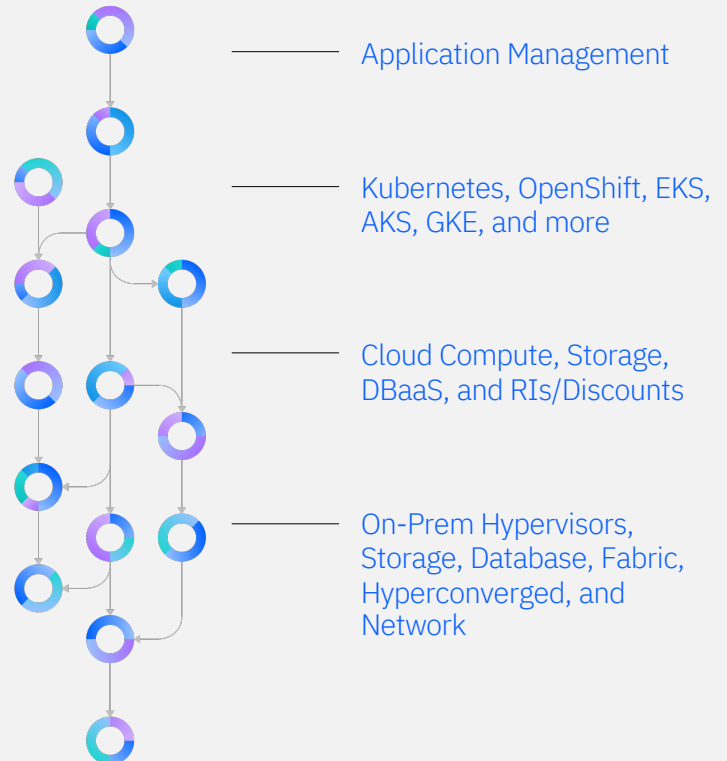
Director, Cloud Adoption and
Optimization at a Leading Multinational
Professional Services Company



Explore live sandbox environment at
turbonomic.com/try

Unlock application, cloud native, and
cloud elasticity anywhere

Our app-first, full-stack solution integrates
with a wide range of platforms to unlock
elasticity.



¹Forrester Total Economic Impact of IBM Turbonomic Application Resource Management

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

Trustworthy actions

App-first, demand-based analysis ensures actions can be safely automated across Kubernetes, Google Cloud, and more. Turbonomic delivers...

- Google Compute Engine optimization
- CUD-aware compute scaling
- GKE optimization

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 exactly how automating application resourcing impacts on the customer experience (response-time or other business SLOs).



Google Compute Engine

Automatically determines the correct VM types for Google Compute Engine. The following metrics are considered for every compute scaling decision:

- vCPU
- vMem
- IO Throughput
- IOPS

CUD-aware compute scaling

Generates specific scaling actions that assure performance, while accounting for your Committed Use Discounts (CUD).

Achieve maximum CUD utilization and coverage within regions or billing accounts.

Google Kubernetes Engine (GKE)

Optimizes the Kubernetes platform for performance and cost with the following actions:

- **Container rightsizing:** Scale container limits/requests up or down based on application demand
- **Continuous pod moves:** Move pods to avoid resource congestion and defragment the cluster
- **Intelligent cluster scaling:** Provision/suspend nodes based on real-time application demand
- **Container planning:** Simulate how to optimize the existing environment, onboard more applications faster.

Sustainable IT

Optimizing application resource consumption either in the datacenter, the public cloud, or both, improves an organization's long-term energy consumption profile.