IBM **Turbonomic**

Application Resource Management
for Public Cloud

Cloud optimization you can continuously automate to prevent performance risk and cost overruns.

Software 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 rightsizing [1]

Safely accelerate cloud migrations
Optimize on-prem workloads first, then assess appropriate cloud configurations.

Unlock cloud elasticity with continuous optimization
Automate application resourcing across compute, storage, DBaaS, and Kubernetes.

Maximize ROI of next-gen Kubernetes platforms
Continuous optimization from apps to platform to infrastructure unlocks elasticity at every layer.

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

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.

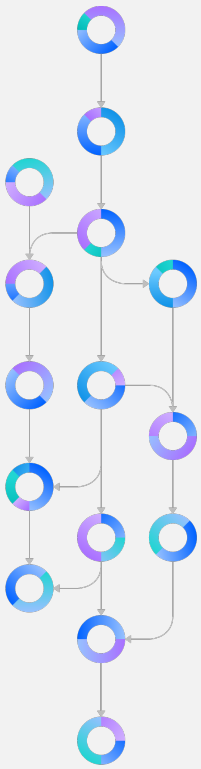

- Application Management
- Kubernetes, OpenShift, EKS, AKS, GKE, and more
- Cloud Compute, Storage, DBaaS, and RIs/Discounts
- On-Prem Hypervisors, Storage, Database, Fabric, Hyperconverged, and Network

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

Unified platform delivers optimization for all app resources

Compute optimization

Automatically determines the correct instance type for cloud application workloads, accounting for the following with every compute scaling decision:

- VCPU
- VMem
- Network & Storage IO
- Throughput
- Reserved Instance Inventory
- Pricing/Discounts
- Disk count, quota, available region capacity, and more

The only solution that simultaneously considers IOPs, and discounts.

Storage optimization

Considers IOPS and throughput, to determine when you need to...

- **Scale between cloud tiers** for performance (IOPS, throughput) and cost
- **Size up volumes for performance** (IOPS, throughput)
- **Modify capacity of IOPS & throughput limit** for IOPS limits

Increase volume sizes to improve performance. Identify & delete unattached volumes. Always, use exactly what you need.

DBaaS Optimization

Azure SQL Database Scaling

Scale Between Azure Database Tiers: Move between Azure SQL DB Tiers based on utilization (DTU*) with near-zero downtime.

Size Up/Down Database Volumes: Non-disruptively increase or decrease disk size (for used space)

AWS RDS Scaling

Considers storage & compute when generating RDS scaling actions.

Continuously analyzes vCPU, vMem, DB Cache Hit Rate, Storage Amount, & IOPS, generating specific scale up / down actions, which include:

- Change in compute tier
- Change in storage tier
- Change in Storage Amount
- Change in Provisioned IOPS (for the io1 storage type)
- Combination of these actions

RIs & Discounts

RI-aware compute scaling actions increase existing RI inventory utilization.

Demand-based RI purchasing actions maximize reservation-to-VM coverage.

RI optimization helps you:

- Purchase the first order of reservations
- Optimize the use of your current reservations
- Re-purchase soon-to-be-expired reservations
- Continue reservation purchases for growing environments

Kubernetes optimization

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
- **Cluster scaling:** Provision/suspend nodes based on real-time application demand
- **Container planning:** Simulate how to optimize the existing environment, onboard more applications faster.

Supports all upstream versions of Kubernetes including EKS, GKE, AKS, and Red Hat OpenShift.

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, major cloud providers, and on-prem VMware.

- Intelligent sizing
- Continuous placement
- Dynamic scaling
- Start/stop

Operationalized →

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

- **Ansible**
- **Azure DevOps**
- **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).

IBM

Minimize risk with automation

Cloud elasticity and operational excellence at scale requires a preemptive approach that mitigates risk to application performance *and* cloud budgets. Only Turbonomic provides specific actions that prevent performance bottlenecks and minimize cloud waste. Automation is the key to mitigating these risks—and only Turbonomic provides cloud optimization you can automate.

- ☑ PREVENT PERFORMANCE BOTTLENECKS
- ☑ PREVENT CLOUD WASTE
- ☑ AUTOMATE DYNAMIC RESOURCING
- ☑ ACHIEVE CLOUD ELASTICITY

Traditional Approaches Stall Cloud Success

Reactive and siloed tools simply cannot continuously assure performance while safely reducing cloud waste.