

IBM Turbonomic – AI-driven Optimization for Expensive Compute Resources

Maximize GPU & Compute Power While Minimizing Costs

AI workloads demand high-performance GPUs—but running them at peak efficiency is a challenge. **IBM Turbonomic ensures optimal GPU utilization, preventing over-provisioning and reducing costs.**

How It Works

- 🧠 **AI-Powered Optimization** – Allocates GPU resources dynamically based on actual demand.
- 🧠 **Avoids Idle Compute Costs** – Ensures that expensive GPUs aren't sitting idle.
- 🧠 **Intelligent Scaling** – Right-sizes workloads in public cloud environments like **Azure** and **AWS**.
- 🧠 **Proactive Cost Control** – Reduces the risk of excessive spending on high-performance computing.

The Impact

- ✅ **Eliminate GPU Waste** – Prevent unnecessary GPU over-provisioning.
- ✅ **Cut Cloud Compute Costs** – Avoid paying **€100+ per hour** for idle resources.
- ✅ **Ensure AI Performance** – Keep workloads running at peak efficiency.

Optimize Every AI & Compute Resource

Don't just buy more GPUs—use them smarter.

👉 **Take Control of AI Workload Costs – Request a Demo!**

More details: [GPU optimization - IBM Turbonomic](#)