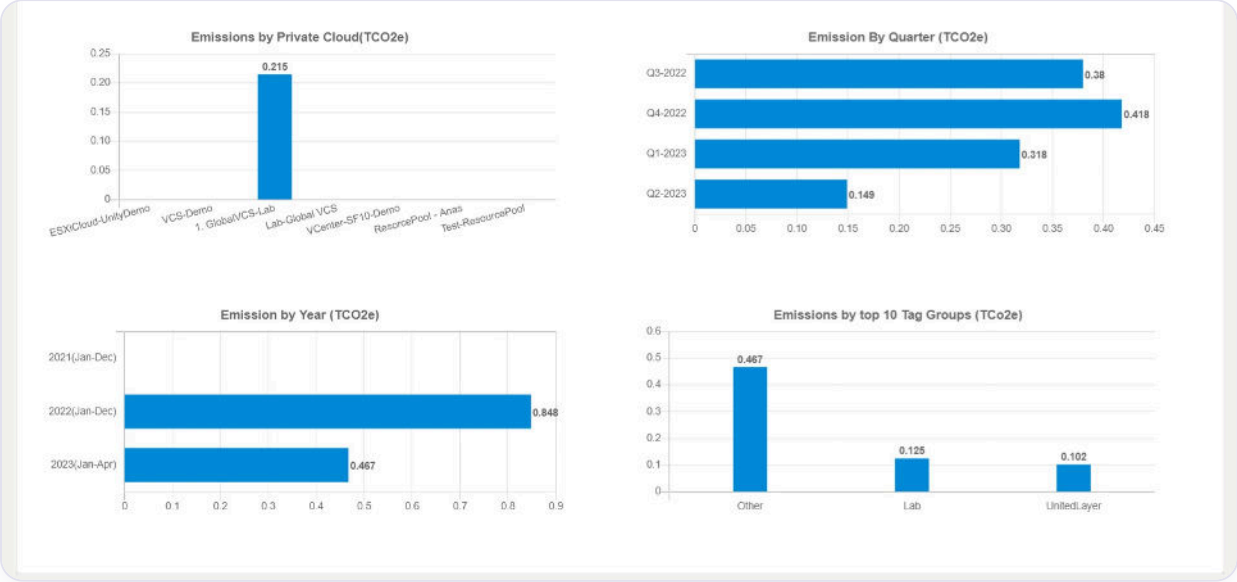
INDUSTRY CASE STUDY · AI & TECHNOLOGY

Run the AI factory with less energy and carbon

AI training and inference consume large amounts of GPU power. UnityOne GreenOps gives workload-level visibility and carbon-aware optimization, without slowing priority work.



Emissions by cloud, quarter, year and tag group.

The challenge

- GPU capacity distributed across on-premises, private and public cloud
- Energy-intensive workloads that were hard to optimize
- No workload-level view of utilization, power, cost, carbon intensity and criticality
- Risk of affecting priority AI workloads when reducing waste

How UnityOne GreenOps helps

- Unified GPU utilization, power, workload, cost and carbon telemetry
- Classified workloads by priority, flexibility, deadlines and resource demand
- Identified idle, underutilized and oversized GPU allocations
- Recommended carbon-aware scheduling, right-sizing, pooling and reclamation

© Copyright 2026 By UnityOne AI

CUSTOMER CASE STUDY

An AI research enterprise improves GPU utilization and cuts avoidable energy use

UnityOne AI partnered with an AI research enterprise running GPU-intensive training and inference across on-premises infrastructure, private cloud and public cloud.

CAPABILITIES USED

GPU efficiency
Find idle, underutilized and oversized allocations.

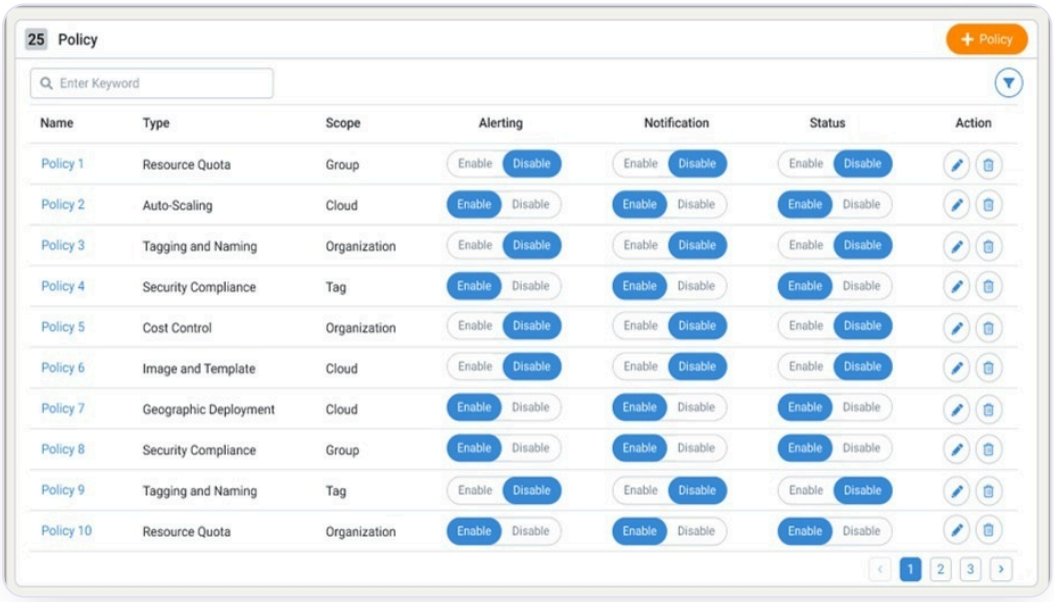
Carbon-aware scheduling
Run flexible jobs when and where carbon is lower.

Workload attribution
Carbon per model, container, VM and GPU.

Policy controls
Approvals, rollback and audit trails.

Hybrid AI estate
On-premises, private cloud, public cloud and Kubernetes.

Cost and carbon together
Balance carbon, cost, capacity and performance.



Policies with approvals and alerting for governed change.

© Copyright 2026 By UnityOne AI

RESULTS

What the enterprise achieved

- Improved GPU utilization and capacity for priority workloads
- Reduced avoidable energy consumption and carbon exposure
- Balanced carbon, cost, capacity and performance in workload decisions
- Governed optimization across data center, cloud and Kubernetes AI environments

Results as published on unityone.ai. The case study does not publish percentage results.

WHY UNITYONE

Built for AI
GPU, model and token efficiency.

Safe optimization
Priority workloads protected.

Proven
AI research enterprise case study.

GETTING STARTED

STEP 1
Explore
Try GreenOps in the UnityOne AI Playground.

STEP 2
Measure
Connect one site and one cloud to set a baseline.

STEP 3
Reduce
Act on recommendations and track results.

GET STARTED

See Cerne™ GreenOps in action

Explore the platform hands-on in the UnityOne AI Playground, or talk to our team about measuring and reducing the carbon impact of your hybrid estate.

▷ Enter Playground ✉ Contact Us

Email: contact@unityone.ai · Playground: cerne.unityone.ai · Sales: unityone.ai/contact-us

© Copyright 2026 By UnityOne AI