



NVIDIA Mission Control

Run models, automate the essentials.



Building AI Factories

As AI rapidly transforms industries, driving unprecedented efficiency gains and innovation, enterprises that fail to establish the necessary mission-critical infrastructure risk falling behind competitors. However, building productive AI factories is complex work, filled with obstacles that many organizations find daunting. Even the most advanced enterprises face significant hurdles:

- > Recruiting and staffing specialized data center roles
- > Managing infrastructure disruptions
- > Optimizing AI workloads against shifting business demands
- > Maximizing resource utilization

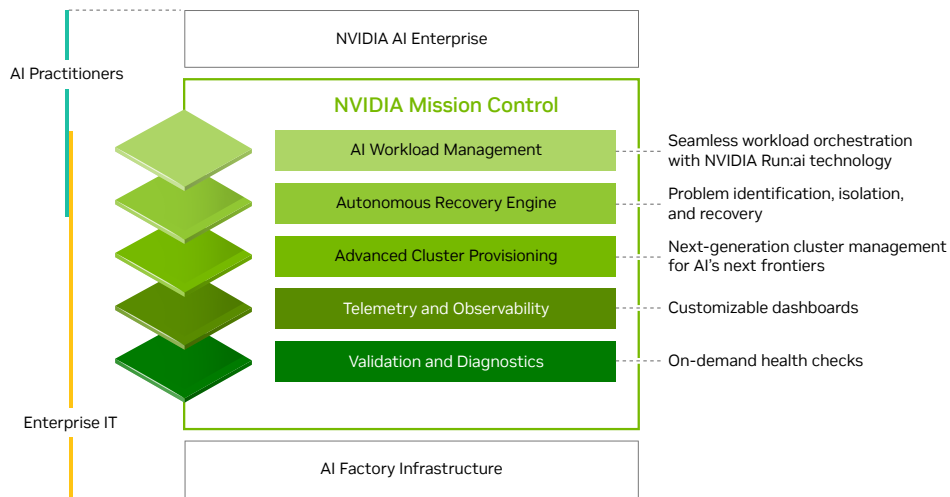
What enterprises truly need is a paradigm shift—a way to harness the power of AI infrastructure with the agility, efficiency, and scale of a hyperscaler.

Bringing the World's Most Advanced AI Factory Expertise to Every Business

NVIDIA Mission Control powers every aspect of AI factory operations—from developer workloads to infrastructure to facilities—with the skills of a world-class operations team, now delivered as software. It powers NVIDIA Blackwell data centers for the newest frontiers of AI, bringing instant agility to inference and training workloads and providing new full-stack intelligence that delivers world-class infrastructure resiliency. Mission Control lets every enterprise run AI with hyperscale-grade efficiency, accelerating AI experimentation.

Key Benefits

- > **Instant Agility:** Bring agility to mission-critical workloads with seamless orchestration, workload flexibility, and advanced cluster control.
- > **Hyperscale-Grade Efficiency, Without Complexity:** Get expert AI factory operations for intelligent 24/7 data center management, automating tasks and filling critical skill gaps.
- > **Gold Standard Infrastructure Resiliency:** Redefine infrastructure resiliency with proactive monitoring, rapid fault identification, and 10X faster time to recovery for training and inference runs.
- > **Accelerated AI Experimentation:** Maximize workload utilization and compute cycles, boosting developer productivity for a new standard of practical scale for enterprise AI.



State-of-the-art AI factory software stack

NVIDIA Mission Control in Action

- > **Simplified Cluster Bring-Up and Provisioning:** Accelerate deployment with new automation and standardized APIs, featuring integrated inventory management and visualizations.
- > **Seamless Workload Orchestration:** Empower model builders with effortless and simplified workload management with [NVIDIA Run:ai](#) functionality.
- > **Energy-Optimized Power Profiles:** Balance power requirements and tune GPU performance for various workload types with developer-selectable controls.
- > **Autonomous Job Recovery:** Identify, isolate, and recover from problems without manual intervention for maximized developer productivity and infrastructure resiliency.
- > **Customizable Dashboards:** Track key performance indicators with access to critical telemetry data about your cluster and easy-to-set dashboards.
- > **On-Demand Health Checks:** Validate hardware and cluster performance throughout the lifecycle of your infrastructure.
- > **Building Management Integration:** Enhance coordination with the building management system (BMS) for improved control for power and cooling events, including rapid leakage detection.

Manage and Run AI at Light Speed

Find out how NVIDIA Mission Control transforms AI data center operations and workload orchestration, enabling every enterprise to simplify how AI factories are deployed and operated throughout the entire cluster lifecycle.

Ready to Get Started?

To learn more about NVIDIA Mission Control, visit nvidia.com/mission-control

NVIDIA Mission Control includes NVIDIA Base Command™ Manager and NVIDIA Run:ai functionality as part of integrated software delivery across configuration, validation, and operations.

Cluster Management Features for Enterprise IT Teams

- > Advanced provisioning for secure, scalable, and standardized deployment of AI factories.
- > Cluster-wide control and serviceability.
- > Built-in playbooks to recover from hardware anomalies.
- > Unified observability across the entire AI factory—compute, networking, and management plane.
- > Built-in leakage policy for liquid-cooled systems.

Workload Management Features for AI Practitioners and MLOps

- > Centralized control plane for workload management.
- > Dynamic GPU allocation adapts compute resources to fluctuating AI demands.
- > Inference optimization reduces cold start times and scales inference workloads seamlessly.
- > Self-serve access to machine learning (ML) tools and compute provides on-demand access to ML environments, reducing setup time and increasing productivity.
- > Preconfigured, secure containers simplify deployment of popular open-source frameworks like TensorFlow, PyTorch, and Jupyter Notebooks.