

PNY®



Enabling AI  
with Software.

PNY®

NVIDIA®

# Visual Computing

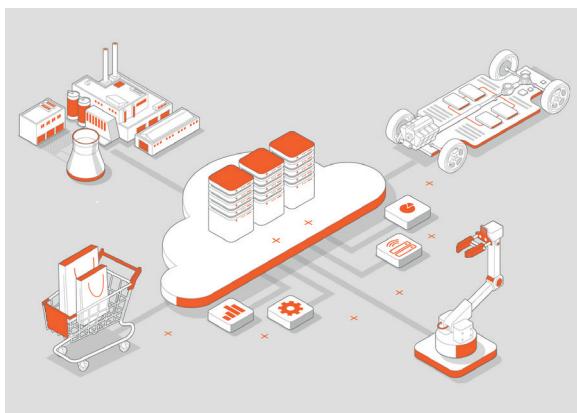
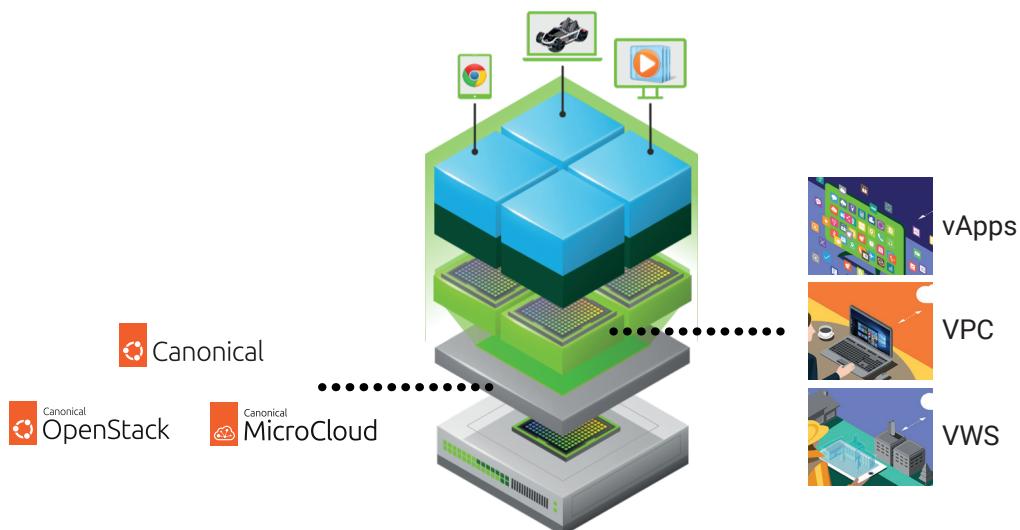
## POWERING PERFORMANCE WITH GPU VIRTUALIZATION

Visual Computing is the engine that drives today's most demanding AI and visualization workloads. By leveraging GPU acceleration, organizations can maximize efficiency, run complex simulations and deliver advanced 3D and data visualization.

### Virtualization

Unlock Next Level Performance with Virtual GPUs

The GPU virtualization, leveraged by NVIDIA technologies and Canonical's open-source solutions, provides optimal flexibility and scalability for AI infrastructures, enabling efficient resource management and maximum performance.



### Canonical MicroCloud & Canonical OpenStack

#### OPEN SOURCE SCALABLE CLOUD INFRASTRUCTURE

Delivers lightweight, automated and secure cloud infrastructure. Deploy clusters in minutes, scale from edge to data center and keep costs low with open-source efficiency. Built on Ubuntu with integrated security and updates, it ensures flexibility and reliability for AI workloads.



### NVIDIA Virtual GPU (vGPU)

#### FLEXIBLE AND SCALABLE GPU POWER FOR EVERY WORKLOAD

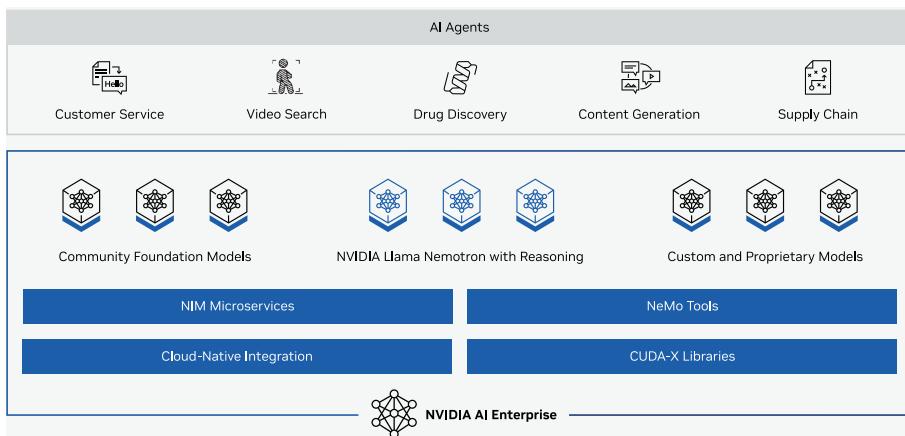
NVIDIA vGPU boosts user density and reduces costs by offloading workloads from CPU to GPU. It supports everything from daily apps to AI and data science, with flexible GPU sharing or pooling for demanding tasks. Enterprise-grade software ensures security, performance and long-term ROI.



# PNY AI Factory

## ORCHESTRATING THE AI LIFECYCLE

The PNY AI Factory is the backbone of AI production, a software-defined environment where data, compute and orchestration converge. It enables organizations to accelerate development, streamline deployment and scale AI workloads with confidence.



## NVIDIA AI Enterprise

### Cloud-Native Software Platform

NVIDIA Enterprise enhances AI scalability and performance with tools like cluster management, vGPU and optimized networking. CUDA enables fast data processing and model training, accelerating data science workflows.

## NVIDIA Run:AI & NVIDIA Mission Control

Delivering Accelerated AI Operations with Dynamic Orchestration Across NVIDIA

NVIDIA Run:ai brings advanced orchestration and scheduling to NVIDIA's AI platforms, enabling enterprises to scale AI operations with minimal complexity and maximum performance.

The diagram shows the NVIDIA Run:AI architecture. It consists of four main layers: 'Data Engineering', 'Experiment Tracking', 'Modeling Training', and 'Model Serving' at the top; 'NVIDIA AI Enterprise / AI Software' in the middle; 'NVIDIA Run:ai' (with sub-modules: AI Lifecycle Integration, Resource Management, Workload Orchestration, Policy Engine, and GPU Orchestration); and 'Infrastructure' (with partners: NVIDIA, Cloud Service Providers, OEMs, and NVIDIA Cloud Partners) at the bottom. The bottom section also features the 'NVIDIA Mission Control' interface.

**NVIDIA Run:AI**  
THE ENTERPRISE PLATFORM FOR AI WORKLOADS AND GPU ORCHESTRATION

NVIDIA Run:AI accelerates AI operations with dynamic orchestration across the AI life cycle, maximizing GPU efficiency, scaling workloads and integrating seamlessly into hybrid AI infrastructure with zero manual effort.



## NVIDIA Mission Control

### NEXT-LEVEL AI OPERATIONS FOR INFRASTRUCTURE AND WORKLOADS

Mission Control is NVIDIA's central management platform for AI infrastructure. It unifies tools like Run:AI to deliver dynamic orchestration, maximize GPU efficiency and provide seamless visibility and control across clusters empowering teams to scale AI workflows faster with less effort.

# NVIDIA DGX Cloud

Fully managed AI training platform

NVIDIA DGX™ Cloud is a unified AI platform on major clouds, delivering high performance, software optimizations and AI expertise. Fully managed by NVIDIA, it ensures top-tier performance and enables organizations to seamlessly move AI workloads from development to deployment in the age of agentic and physical AI.



NVIDIA DGX Cloud  
Create



NVIDIA DGX Cloud  
Serverless Inference



Video Curation & Post-Training on DGX Cloud



NVIDIA DGX Cloud  
Benchmarking Catilica

## Canonical

AI-Optimized Solutions with NVIDIA Technologies and Seamless MLOps Integration

Canonical is the publisher of Ubuntu and a provider of security, support and services for open-source solutions. PNY and Canonical together offer solutions based on NVIDIA virtualization and orchestration technologies, with seamless integration of data science and MLOps tools for building AI models, optimized for AI clustering.



Ubuntu - the world's  
favorite enterprise Linux



Data center cloudification



Simplify your data  
management and MLOps



Managed service

Canonical has been selected by NVIDIA as a preferred infrastructure partner.

Ubuntu OS comes installed by default on NVIDIA servers.

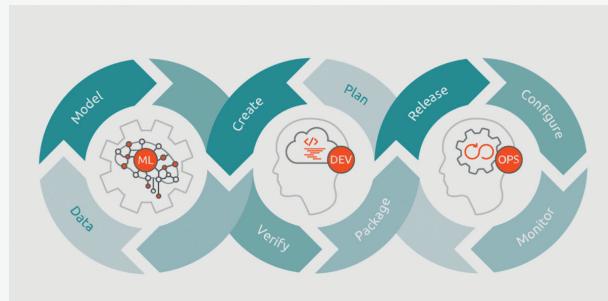


### Canonical Data solutions

#### KEY BENEFITS

- ✓ Open-source portfolio
- ✓ Integrated security
- ✓ Multi-cloud support
- ✓ Automated scalability

Canonical Data solutions deliver a secure, automated and scalable open-source portfolio from PostgreSQL and MySQL to MongoDB, Kafka, Spark and OpenSearch, designed for seamless operation across hybrid and multi-cloud environments.



### Open source MLOps

#### KEY BENEFITS

- ✓ Fully open-source
- ✓ Integrated automation
- ✓ Scalable workflows
- ✓ Reliable lifecycle management

Canonical delivers a fully open-source MLOps stack that simplifies the entire machine learning lifecycle, enabling scalable, reliable and reproducible AI workflows across any cloud. With integrated tools like Charmed Kubeflow, Charmed MLFlow and Charmed Feast, organizations can deploy, track and manage AI models across any cloud.

FOR MORE INFORMATION:

Visit: [WWW.PNY.EU](http://WWW.PNY.EU)

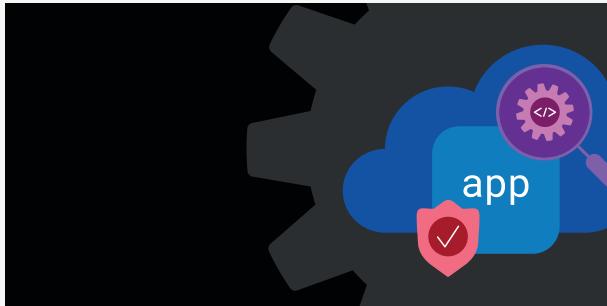
The PNY logo is a registered trademark of PNY Technologies, Inc. All other trademarks are the property of their respective owners. All rights reserved. - © 2025 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, NVIDIA DGX, NVIDIA DGX™ Cloud, NVIDIA Omniverse and Run:AI are trademarks and/or registered trademarks of NVIDIA Corporation. All company and product names are trademarks or registered trademarks of the respective owners with which they are associated. Features, pricing, availability and specifications are all subject to change without notice. - © 2025 Canonical Ltd. Ubuntu and Canonical are registered trademarks of Canonical Ltd. - © 2025 Scality. All rights reserved. - © 2025 WekalQ, Inc. All rights reserved.

ACMK0931

# ☒ Networking – AI Performance and Security

Secure, Scalable Networking for AI and Cloud

AI workloads demand networks that are automated, high-performance and secure. Modern solutions combine cloud-style automation to manage diverse switch fabrics with advanced load balancing and security to maximize efficiency and protect critical applications.



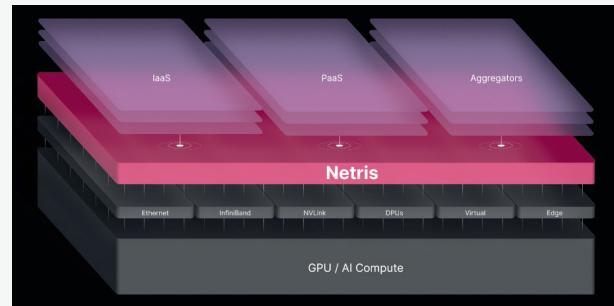
## F5 BIG-IP

INDUSTRY-LEADING APPLICATION DELIVERY AND SECURITY SERVICES

### KEY BENEFITS

- ✓ Enterprises running AI and microservices
- ✓ Kubernetes-based environments
- ✓ Secure, performance-driven multi-tenant workloads

Offloading traffic improves GPU efficiency, cuts latency and lowers costs. Zero Trust security—firewalls, DDoS protection, workload isolation—enables secure scaling of AI, microservices and 5G.



## Netris

CLOUD-STYLE AUTOMATION

### KEY BENEFITS

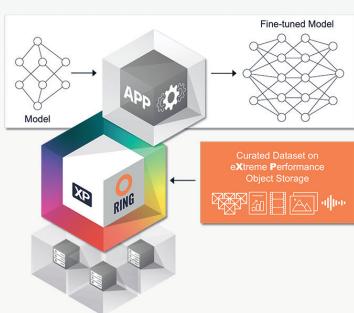
- ✓ Cloud Service Providers (CSPs)
- ✓ Multi-tenant infrastructures
- ✓ Operators of large-scale GPU fabrics

Built for CSP multi-tenant networks, Netris streamlines data center and GPU fabric operations with cloud-like automation. It manages NVIDIA, Arista, Cisco and Broadcom SAI switches and provides core services including VPCs, load balancers and firewalls.

# ☒ Data – The Foundation of Scalable AI

High-Performance Data Architecture for AI

AI relies on data – storing it securely, accessing it quickly and delivering it to GPUs at scale. Modern AI infrastructures require both massive storage capacity for raw data and high-performance pipelines to feed training and inference workloads in real time.



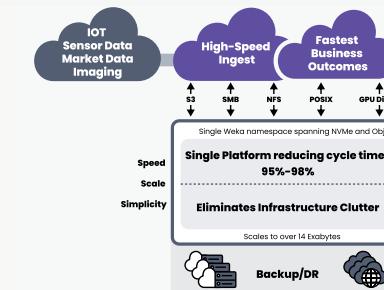
## Scality

SCALABLE STORAGE FOR AI DATA LAKES

### KEY BENEFITS

- ✓ Petabyte-scale capacity
- ✓ Unified multi-protocol storage
- ✓ Reliable and cost-efficient

Scality delivers reliable, petabyte-scale storage for images, video and logs. With unified object, file and block support, it is the ideal foundation for AI data lakes and long-term archives.



## Weka

HIGH-PERFORMANCE DATA FOR AI WORKLOADS

### KEY BENEFITS

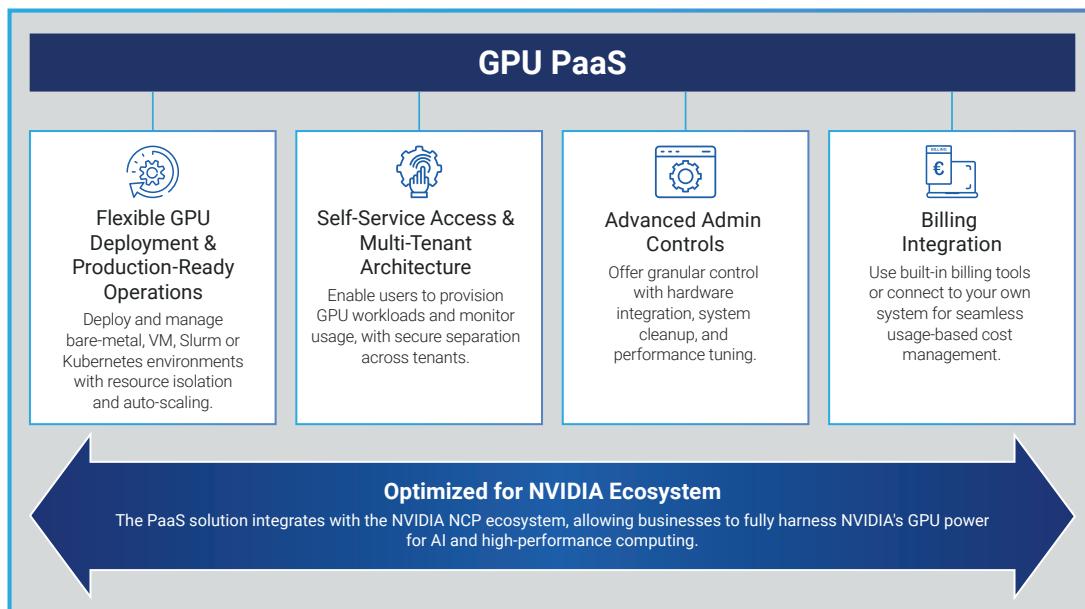
- ✓ Extreme throughput for AI/ML
- ✓ Low latency for real-time pipelines
- ✓ Scalable, simple architecture

Weka provides ultra-fast throughput and low latency, ensuring GPUs are always fed with data for training, inference and real-time analytics.

## ☒ PaaS – Platform As A Service

Accelerated, Scalable and Secure GPU Infrastructure for Modern AI Workloads

A GPU Platform as a Service (PaaS) simplifies cloud deployment and management of GPU-heavy applications, using orchestration tools for efficient resource use. It supports scalable and secure AI workloads with features like multi-tenancy, self-service access and integrated monitoring.



# AI Applications

## FROM VISUALIZATION TO INDUSTRY USE CASES

AI applications connect data, models and visualization into real-world solutions. From Omniverse for collaboration and digital twins, to NVIDIA models and microservices for generative AI and RAG, these tools accelerate deployment and innovation. Across industries like retail, smart cities and manufacturing, AI applications transform infrastructure into measurable impact.

## ☒ AI Use Cases

NVIDIA's models and microservices make it possible to rapidly deploy advanced AI solutions across industries.



### Agentic AI

Develop agents that can reason, plan and act autonomously



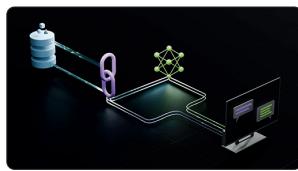
### AI Data

Leverage next-generation enterprise infrastructure optimized for AI



### Conversational AI

Deliver natural, personalized interactions with real-time speech capabilities



### Cybersecurity

Strengthen your systems and infrastructure with AI-driven security solutions



### Conversational AI

Deliver natural, personalized interactions with real-time speech capabilities

# NVIDIA Blueprints

## REFERENCE ARCHITECTURES FOR FASTER AI DEPLOYMENT

### NVIDIA Provides a Full Stack for AI Factories

Built on NVIDIA customer-validated data center reference architectures

Intelligence, Tokens and Business Outcomes

SOFTWARE

NVIDIA Blueprints

NVIDIA AI Enterprise

NVIDIA Omniverse

AI Workload and GPU Orchestration

Infrastructure Management

COMPUTE

CPU GPU NVLink

NETWORK

Infiniband Ethernet

STORAGE

Data

NVIDIA-Certified Systems

Electricity

\*NVIS = NVIDIA Infrastructure Specialists

NVIS\*  
and/or  
NVIDIA  
Partner  
Services

NVIDIA AI Blueprints are customer-validated guides for building real-world AI solutions with the full NVIDIA stack. They connect infrastructure, orchestration and applications to accelerate results. By showing how NIM microservices and AI models form scalable, production-ready systems, they reduce complexity and speed deployment from concept to production.

## ↗ Visualization Use Cases

AI-powered visualization transforms industries by improving decision-making and collaboration.



### Industrial

Real-time design validation, predictive maintenance



### Retail

Immersive customer experiences and store optimization



### Smart City

Traffic, energy and urban planning simulations



### Omniverse

Collaborative digital twin and simulation workflows



### ISV Solutions

Partner applications optimized for NVIDIA GPUs

## ↗ Visualization

NVIDIA Omniverse™ - Revolutionize design collaboration and Simulation

NVIDIA Omniverse™ is a platform of APIs, SDKs and services that enable developers to easily integrate Universal Scene Description (OpenUSD) and RTX rendering technologies into existing software tools and simulation workflows for building AI systems.

Transforming creativity and innovation with NVIDIA Omniverse

### 3D Avatar

#### KEY BENEFITS

- ✓ Unified data pipelines
- ✓ Simulation
- ✓ Python-based modular development platform



### Digital Twin

#### KEY BENEFITS

- ✓ Previsualization
- ✓ Real-Time Dailies
- ✓ Content Collaboration





PNY orchestrates AI software environments, accelerating AI development and deployment.

## PNY Services



Consulting



Development



Deployment & Support

## Benefits



### Faster AI

Accelerate deployment and reduce time-to-value.



### Enterprise-Grade

Reliable, secure and supported at every step



### AI Infrastructure

From development to production, ready for scale.



### Cloud & Hybrid

Simple, flexible deployments anywhere you need.



### AI Tools & Use Cases

Appliance-ready solutions designed for real-world AI.

#### PNY Technologies Europe

Zac du Phare  
9 rue Joseph Cugnot - BP 40181  
33708 Mérignac Cedex, France  
Tel: +33 (0)5 40 240 240

#### PNY Technologies GmbH

Schumanstraße 18a  
52146 Würselen  
Germany  
Tel: +49 (0)240 540 848-0

#### PNY Technologies Middle East Fze

Jafza View 19 308, Jebel Ali Free Zone  
PO Box 263897  
Dubai  
Tel: +97148814966

CONTACT US

**SALES@PNY.EU**