



PNY®

AI & DATA CENTER SOLUTIONS





BENEFITS OF PNY DATA CENTER SOLUTIONS


Artificial intelligence has become the go-to approach for solving difficult business challenges. Whether improving customer service, optimizing supply chains, extracting business intelligence, or designing cutting-edge products and services across nearly every industry, AI gives organizations the mechanism to realize innovation. More recently, Generative AI has significantly accelerated the integration of AI across various industries by utilizing AI and Machine Learning algorithms that use existing content to learn and generate new content, such as text, images, music, audio, and videos. It is applied in art, medicine, video games, finance, music, and many other. As a pioneer in AI infrastructure, PNY Data Center Solutions provide the most powerful and complete AI platform for bringing these essential ideas to life. By incorporating generative AI, PNY enhances its capabilities, enabling businesses to harness the full potential of AI-driven.

One System for all AI Infrastructure

AI INFRASTRUCTURE RE-IMAGINED, OPTIMIZED, AND READY FOR ENTERPRISE AI-AT-SCALE

Flexible AI Infrastructure that adapts to the pace of enterprise:

- One universal building block for the AI data center
- Any workload on any node : any time
- Limitless capacity planning with predictability great performance with scale



European Sovereign AI Infrastructure

A new generation of sovereign AI infrastructure is emerging across Europe, enabling organisations to access world-class AI infrastructure while maintaining full control over their data, operations, and strategic independence. In Sweden, a large-scale AI infrastructure is now being established. Built on NVIDIA's latest DGX SuperPOD technology, these platforms provide the performance, scalability, and reliability required to support advanced AI development and deployment.

PNY Technologies has played a pivotal role in this project, leading the way in AI innovation. As a key player, PNY Technologies supplied, installed, and leveraged its expertise in the NVIDIA SuperPOD infrastructure, building on enduring partnerships, especially with AI NVIDIA leader.

AI infrastructure is set up with premium NVIDIA AI technologies:

Cluster of 16 NVIDIA DGX GB300 Tensor Core, 1152 GPUs NVIDIA Blackwell Ultra, up to 23040 PFLOPS (FP4), 12,28 PB of VAST storage, InfiniBand 1.6 Tb/s switch.

PNY KEY SELLING POINTS

PNY provides the most comprehensive computing, storage and networking portfolio for Datacenters.

Its vast expertise and experience enable to support the challenges of the Datacenter coming from a broad range of industries such as Healthcare, Retail, Higher Education and more.

With its constant focus on innovation, quality, reliability, and support of excellence, the company helps its customers to empower and manage their infrastructure in the most technical and financial efficient ways.



LOGISTICS
STRENGTH AND
ADAPTABILITY



LOCAL
TECHNICAL
SUPPORT





EXPERIENCE &
EXPERTISE

PNY SOFTWARE SOLUTIONS

OPTIMIZE AI PRODUCTION

At the forefront of AI production, orchestrate massive data processing, optimize hardware resource utilization, and secure every exchange. Unlock AI’s potential with cutting-edge software solutions designed for efficiency, scalability, and security.




NVIDIA AI Enterprise

Deploy AI from pilot to production with NVIDIA AI Enterprise platform. Streamline workflows, optimize model performance with microservices, and ensure secure, seamless transitions from prototype to production.



Container and Management Integration

Start with Ubuntu supported by Canonical and get trusted open source and run your entire Artificial Intelligence/Machine Learning lifecycle on a single integrated stack.



Workload Orchestration

Transform your AI infrastructure to accelerate development, optimize resources, and lead the race in AI innovation with our workload orchestration solutions.



Consulting & Development

Embrace and accomplish AI projects with consulting and customized development services, delivering innovative and effective solutions.

Global AI Expertise



AI Solutions Aggregator

PNY delivers optimized applications for complex and evolving AI technologies through an aggregation of partners within the same technology landscape, leveraging NVIDIA GPU solutions to ensure high performance and reliability.




AI System & Service Advisor

Empower every stage of the customer lifecycle to deliver significant value, accelerate decision-making, and facilitate fast and easy deployment with our integrated solutions and services, spanning from audit to installation.

NVIDIA DGX™ SYSTEMS

SETTING THE BAR FOR ENTREPRISE AI INFRASTRUCTURE

NVIDIA DGX™ systems are purpose-built to meet the demands of enterprise AI and data science, delivering the fastest start in AI development, effortless productivity, and revolutionary performance for insights in hours instead of months.



NVIDIA DGX Rubin NVL8

Supercharged performance for agentic AI

DGX Rubin NVL8 is a turnkey infrastructure solution purpose-built to accelerate any AI workload and deliver intelligence at scale.



NVIDIA DGX Vera Rubin NVL72

Building the next frontier of AI

DGX Vera Rubin NVL72 delivers industry-leading performance per watt and tokens per watt for any AI workload, allowing enterprises to dramatically scale intelligence while improving token economics.

System Specifications

	NVIDIA DGX Rubin NVL8	NVIDIA Vera Rubin NVL72
GPUS	8x NVIDIA Rubin GPUs	72 NVIDIA Rubin GPUs
Total GPU Memory	2.3 TB 176 TB/s	20.7 TB 1,580 TB/s
Performance	NVFP4 Inference: 400 PFLOPS NVFP4 Training: 280 PFLOPS* FP8/FP6 Training: 140 PFLOPS*	NVFP4 Inference: 3,600 PFLOPS NVFP4 Training*: 2,520 PFLOPS FP8/FP6 Training*: 1,260 PFLOPS
System Power Usage	~24 kW	Power per rack (rated)~150kW
CPU	2x Intel® Xeon® 6776P processors	3,168 custom NVIDIA Olympus cores (Arm® compatible) 36 NVIDIA Vera CPUs
System Memory	2 TB	54TB of LPDDR5X memory
Networking	8x OSFP ports serving 8x single-port NVIDIA® ConnectX®-9 VPI up to 800 Gb/s NVIDIA InfiniBand and Ethernet 2x 400G QSFP112 NVIDIA BlueField®-4 DPUs Up to 800 Gb/s NVIDIA InfiniBand and Ethernet	>144x OSFP single-port NVIDIA® ConnectX®-9 VPI with 800 Gb/s NVIDIA InfiniBand and Ethernet > 18x dual-port NVIDIA BlueField®-4 VPI with 400 Gb/s NVIDIA InfiniBand and Ethernet
Management network	1GbE RJ45 host baseboard management controller (BMC) 1GbE onboard NIC with RJ45	Host baseboard management controller (BMC)
Storage	OS: 2x 2 TB NVMe M.2 Data storage: 8x 4 TB NVMe E1.S	Each compute tray provides a total of 4x 7.68TB E1.S NVMe as RAID0 for local storage and 1x 2TB E1.S NVMe for the OS image
Software	NVIDIA AI Enterprise (optimized AI software) NVIDIA Mission Control (AI data center operations and orchestration with NVIDIA, Run:ai technology) NVIDIA DGX OS (operating system), Support Ubuntu	NVIDIA AI Enterprise (optimized AI software) NVIDIA Mission Control (AI data center operations and orchestration with NVIDIA, Run:ai technology) NVIDIA DGX OS (operating system), Support Ubuntu

* Dense specification.

NVIDIA DGX BASEPOD™ & SUPERPOD™

BUILDING ENTERPRISE AI REQUIRES
A SOLUTION THAT'S ENTERPRISE-READY

A solution that eliminates the complexity of designing AI infrastructure, integrated into your existing environment without upheaval, and that enables a simplified, fast deployment experience. NVIDIA DGX BasePODs and SuperPODs solution delivers the world's best platform for enterprise AI innovation, from AI infrastructure building block to Data Center as a product. Start Small, Scale Predictably in Response to Business Demand.



NVIDIA DGX BasePOD™

A reference architecture that incorporates best practices for compute, networking, storage, power, cooling, and more, in an integrated AI infrastructure design built on NVIDIA DGX, from 2 to 16 nodes.

NVIDIA DGX SuperPOD™

AI data center infrastructure platform that enables IT to deliver performance without compromise.

DGX SuperPOD offers leadership-class accelerated infrastructure and agile, scalable performance for the most challenging AI and high-performance computing (HPC) workloads.



Complete Lifecycle
of Expertise



Ever-Improving Advanced
Infrastructure Software



Powered by NVIDIA Base
Command



NVIDIA DGX Vera Rubin NVL72

NVIDIA DGX™ Vera Rubin NVL72 provides enterprises with a turnkey, ready-to-deploy AI infrastructure solution built upon the NVIDIA Rubin platform, purpose-built to be deployed at scale to accelerate the most complex AI models.



Maximum Efficiency With NVIDIA
Vera Rubin Superchips



The Universal Engine for
AI at Scale



Accelerating Operations With
Intelligent Software



FOR DISCOVERING ALL OUR RECOMMENDATIONS:
Please visit our PNY Data Center Solutions on:
WWW.PNY.COM/EN-EU/PROFESSIONAL/

PNY AI Factory

UNIFIED INFRASTRUCTURE FOR AI EXCELLENCE

Cloud Service Providers and Enterprises need AI infrastructure that's fast, scalable, and built for multi-tenancy.

This solution maximizes NVIDIA GPU efficiency, reduces costs, and accelerates AI workloads, empowering CSPs and Enterprises to deliver high-performance AI services while boosting profitability and growth.

NETWORKING
NVIDIA Networking

COMPUTE
NVIDIA DGX / HGX B200 / B300
NVIDIA GB200 / GB300 NVL72

STORAGE
Dedicated Flash FS AI storage
+ Object S3 storage

ADVANCED AI SOFTWARE
A full-stack AI software suite including workload orchestration, containerization, and AI management tools.

Large-Scale AI Model Training

Real-Time AI Inference

Massive Data Processing

AI Cloud & Multi-Tenant AI Workloads

PNY, YOUR KEY SOLUTION PARTNER IN THE AI ECOSYSTEM

PNY Technologies is leading the way in AI innovation, setting new standards in the industry. As a key partner, PNY offers end-to-end solutions in the AI ecosystem, including compute, storage, networking, and complete software stack. With our deep expertise, we can define what are the best applications and storage strategies that will help you on your journey to AI.

BENEFITS OF AI STORAGE



Data Protection and Security

AI storage solutions come with robust security features. This includes encryption, access controls, and auditing capabilities.



Flexibility and Adaptability

AI storage solutions are designed to be flexible and adaptable to different types of AI workloads.



Advanced Analytics and Data Insights

Optimized AI storage can help support those analytics, and more accurate, insightful, and responsive to real-time data.



Automation of Routine Tasks

AI storage can automate many routine storage management tasks, such as provisioning, data migration, and backup.

FOR MORE INFORMATION:

Visit: 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. - © 2026 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, NVIDIA DGX, NVIDIA DGX BasePOD, NVIDIA DGX SuperPOD, NVIDIA LinkX Cables, ConnectX, NVLink, NVIDIA BlueField DPU, NVIDIA Base Command, NVIDIA NVSwitch and Spectrum-X 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. - © 2026 Scaleway SAS. All rights reserved. - © 2026 illad Group. All rights reserved. - © 2026 DataDirect Networks, All Rights Reserved. - © 2026 Canonical Ltd. Ubuntu and Canonical are registered trademarks of Canonical Ltd.

ACMK0969

UNLOCK THE FULL POTENTIAL OF YOUR DATACENTER

Networking for AI Data Centers

TWO TYPES OF AI DATACENTERS




AI Factories



AI Cloud

THE NETWORK DEFINES THE DATA CENTER



Cloud

- Multi-tenant
- Variety of small-scale workloads
- Traditional Ethernet network can suffice


Generative AI Cloud

- Multi-tenant
- Variety of workloads including larger scale Generative AI
- Traditional Ethernet network for North-South traffic
- **NVIDIA Spectrum-X Ethernet for AI fabric (East West)**

AI Factories

- Single or few users
- Extremely large AI models
- **NVIDIA NVLink and InfiniBand gold standard for AI fabric**

Complete Infiniband and Ethernet portfolio



SmartNICs and DPU

Up to 800Gb/s the Ethernet Network Adapters, InfiniBand Host Channel Adapters (HCA) and Data Processing Unit (DPU) provide accelerated networking and security for the most advanced cloud and AI workloads.




Switches

The NVIDIA Spectrum™ Ethernet switch family spans 1GbE to 800GbE, ideal for AI fabrics with NVIDIA GPUs and end-to-end cloud data center networks. NVIDIA Quantum InfiniBand switches offer up to 800Gb/s per port, smart acceleration, and a robust architecture for high-performance computing, AI, and hyperscale cloud infrastructures, with reduced cost and complexity.



Interconnect

Explore NVIDIA's comprehensive line of 100G/200G/400G/800G InfiniBand and Ethernet interconnects, offering low latency, low power, and high reliability for AI and accelerated computing. NVIDIA LinkX® Cables connect within Quantum and Spectrum™ architectures for switch-to-switch, top-of-rack to ConnectX® smart network adapters, and BlueField® DPUs in servers and storage.



Spectrum™ -X

First ethernet networking platform for AI, leveraging the close integration of NVIDIA Spectrum™-6 Ethernet switch with the NVIDIA BlueField®-4 data processing unit DPU, Spectrum-X achieves peak performance in AI, machine learning, natural language processing, various industry applications.



PNY provides unsurpassed service and commitment to its Data Center customers.

SOLUTIONS FOR EACH MARKET



HPC



Data Center



AI/Machine Learning



Clouds



Storage



Security



Telecom



Media & Entertainment

PNY ADVANTAGE

- Over 20 years of expertise delivering NVIDIA GPU solutions
- Strong partnerships with leading technology suppliers
- Dedicated head count for Sales, Marketing and Support
- Local pre-sales and post-sales assistance
- Direct access to technical support hotlines
- Pre-sales tools, support and configuration assistance
- Dedicated Field Application Engineers providing value-added on-site installation
- Published product support and training materials
- Advanced replacement options for mission-critical deployments
- Long product life cycles and availability
- Loyalty partner channel programs
- Dedicated support programs
- Robust logistics and operational capabilities
- Equipment loan for strategic opportunities
- PNY LAB: a technology center supporting the development of AI, HPC, and VDI solutions

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 QUADRO® GmbH

Schumanstraße 18a
52146 Würselen
Germany

Tel: +49 (0)2405/40848-0

PNY Technologies Middle East Fze

Jafza View 19 308, Jebel Ali Free Zone
PO Box 263897
Dubai

Tel: +97148814966

CONTACT US

SALES@PNY.EU