



AHEAD

NVIDIA® DGX SUPERPOD™:

A Reliable AI Platform for Healthcare and Life Sciences

How AHEAD Helps Deploy NVIDIA AI Infrastructure

As the healthcare industry looks for new opportunities for innovation, AI is set to become a foundational technology. In fact, many medical providers and life sciences companies now recognize the need to leverage AI to adapt to a more competitive marketplace, meet the demand for digital services, and overcome other challenges the industry is currently facing.

NVIDIA® DGX SuperPOD™ offers an ideal AI infrastructure platform for healthcare companies to integrate AI and high-performance computing (HPC) into their organizations.






AI in Healthcare & Life Sciences

AI has enormous potential in the healthcare and life sciences industries, with many organizations already exploring use cases related to patient triaging, precision medicine, drug discovery, and more.

AI applications can transform the way healthcare providers deliver services, and in turn, lead to better patient outcomes and more efficient use of limited medical resources. For example, AI triaging can help hospitals prioritize patients according to the urgency of their illnesses, which can enhance care quality, reduce patient wait times, and lessen mortalities in emergency situations.

AI-powered precision medicine — which tailors treatments and drug management to an individual's unique circumstances — empowers healthcare providers to deliver more effective treatments. This data-driven approach to medicine can reduce healthcare costs by preventing complications, hospitalizations, and other adverse reactions that could occur from more generic treatments.

Machine learning algorithms can also be applied to a broad range of complex life sciences datasets and biomedical data, such as imaging and clinical information, to identify new medical trends and targets for novel drugs. This can accelerate clinical trials and drug discovery while also reducing research and development costs.




AI Infrastructure Challenges in Healthcare

While AI offers enormous potential in the healthcare industry, deploying it at scale comes with significant challenges. Some AI workloads in healthcare — particularly in areas like telemedicine and remote care — rely on low-latency computing, which requires optimized and scalable infrastructure.

The healthcare industry also tends to have substantial data privacy concerns and requirements for regulations like HIPAA that will keep a dominant portion of these AI workloads on-premise. This often requires optimizing existing data centers to handle more demanding computing workloads to roll out new AI use cases.

In addition, many healthcare providers lack the in-house AI and machine learning expertise necessary to build, deploy, and maintain models and infrastructure effectively. Partnering with AI experts, investing in training, and adopting turnkey technology solutions will be crucial for long-term success.



Exploring NVIDIA DGX SuperPOD

[NVIDIA DGX SuperPOD](#) provides a high-level blueprint for how to build scalable, performant, and reliable AI infrastructure, powered by NVIDIA DGX systems. The reference architecture outlines a full-stack data center platform that includes industry-leading compute, storage, networking, software, and infrastructure management optimized to work together and provide maximum performance at scale.

Ease of Use

NVIDIA DGX SuperPOD is a ready-to-run, turnkey supercomputer with high-performance compute, networking, storage, and software integration.

Scalability

NVIDIA DGX SuperPOD scales to tens of thousands of NVIDIA GPUs to handle training and inference for state-of-the-art AI models.

Developer Productivity

NVIDIA DGX SuperPOD includes enterprise-grade cluster and workload management software, libraries that accelerate infrastructure, and an operating system optimized for AI workloads.

Reliability

NVIDIA DGX SuperPOD is extensively tested and pushed to the limits with real-world enterprise AI/HPC workloads to ensure application performance.

There is also an adjacent NVIDIA DGX [BasePOD reference architecture](#), which provides infrastructure for enterprises to build and scale AI. Some organizations choose to start with NVIDIA DGX BasePOD before scaling into an NVIDIA DGX SuperPOD.

UNIVERSITY OF PENNSYLVANIA PARCC PARTNERS WITH AHEAD AND NVIDIA FOR AI & HPC

The University of Pennsylvania's Penn Advanced Research Computing Center (PARCC) is a major academic medical center, supporting large-scale healthcare, biomedical, and life sciences research. This initiative offers advanced support to interdisciplinary researchers across all 12 affiliated schools, next-generation technology deployments, and efficient computing resource allocation.


In partnership with AHEAD, the PARCC team embarked on plans to establish a robust shared AI and high-performance computing (HPC) research center that will empower computational and data-driven research across campus with research consulting services as well as support for hardware, software, and systems.

This project involved the deployment of an advanced NVIDIA DGX SuperPOD using AHEAD Foundry™ to help solve complex logistical challenges and reduce the implementation timeline to less than a year. AHEAD was also able to provide crucial insights into areas that the PARCC team did not have expertise.

By implementing the NVIDIA DGX SuperPOD, the University of Pennsylvania has doubled its compute capacity available for AI training and inferencing projects. This will greatly accelerate the groundbreaking healthcare and life sciences research the university is conducting.



Penn Advanced Research
Computing Center
UNIVERSITY of PENNSYLVANIA



Deploying AI Infrastructure from NVIDIA with AHEAD

Having an AI infrastructure partner you can trust allows you to concentrate on what applications and innovative business capabilities you want to build. AHEAD is a leader in AI with deep technical expertise in designing, building, and operationalizing AI and HPC infrastructure.

AHEAD has successfully designed and implemented several NVIDIA DGX SuperPOD across multiple industries.



Here's how AHEAD approaches an AI infrastructure engagement:




DEFINING REQUIREMENTS

Our AI experts work with NVIDIA and your team to determine the core elements of your NVIDIA DGX SuperPOD. Then we consider other customizations like security, data protection, and compute that may be needed, as well as the upstream network design.



RACK INTEGRATION

AHEAD has a team of experts in our [Foundry™ rack integration](#) facility that will work with your data center team or co-location provider to design how every piece of equipment will be racked and cabled in the real-world space based on the NVIDIA DGX SuperPOD reference architecture.



FLOOR PLAN DESIGN

It's important to make sure the space is a perfect fit for a custom NVIDIA DGX SuperPOD configuration. This means AHEAD will review your environment to ensure it meets all the requirements for deployment.



RACK ELEVATION & CABLING PLAN

All key components are designed at the same time to avoid project delays later on. This detailed design extends to all components adjacent to NVIDIA DGX SuperPOD, including storage, HPC servers, security devices, and more.




DEPLOYMENT

Once the final design is agreed on, AHEAD Foundry will build and deliver the NVIDIA DGX SuperPOD, with [AHEAD's Hatch® platform](#) providing full visibility into the supply chain for all procurement. Engineers and technicians from AHEAD and NVIDIA can be hands-on in the data center to take care of all the cross-rack cabling, performance testing, bring up, and other physical tasks to ready the technology for operations.



ONBOARDING & MANAGEMENT

After physical implementation and deployment, AHEAD can also provide oversight and accountability with ongoing program and lifecycle management for your NVIDIA DGX SuperPOD.



AHEAD for NVIDIA

In short, AHEAD has deep expertise in AI infrastructure and is experienced with NVIDIA technologies. Our team of data scientists, architects, engineers, and NVIDIA AI advisors collaborate with leading enterprises to implement world-class accelerated computing solutions. That's why we've been named [2025 NVIDIA Rising Star Partner of the Year](#).

We have the scale to stand up and manage NVIDIA DGX SuperPOD pre-integrated solutions through AHEAD Foundry and simplify asset management for clients with our AHEAD Hatch software and our professional and managed services to support full platform operations.

More specifically, AHEAD Foundry is our comprehensive integration facility for building, configuring, testing, and validating deployment capabilities using a plug and play approach. In addition, our full-stack integration facility with liquid-cooling capacity can support NVIDIA GB200 NVL72 systems and future chip generations requiring increased density.

The [AHEAD Hatch platform](#) for IT intelligence can track all components, design documentation, and the full asset lifecycle. Hatch also provides consolidated data at the site level for your infrastructure and can easily integrate with ServiceNow and other leading enterprise platforms to share data in real time. This accelerates every aspect of planning, executing, and supporting your AI infrastructure initiatives.

[Contact AHEAD](#) for your on-premise AI infrastructure
and to learn more about deploying NVIDIA DGX SuperPOD.



AHEAD

Combining cloud-native capabilities in software and data engineering with an unparalleled track record of modernizing infrastructure, we're uniquely positioned to help accelerate the promise of digital transformation.

Visit us at ahead.com.

National Hubs

CHICAGO

444 W. Lake Street
Suite 3000
Chicago, IL 60606

NEW YORK

500 5th Avenue
Floor 17
New York NY 10010

ATLANTA

1117 Perimeter Center
W406
Atlanta, GA 30338

SAN FRANCISCO

2000 Crow Canyon Place
Suite 250
San Ramon, CA 94583