



AHEAD

NVIDIA® DGX SUPERPOD™:

*A Scalable AI Platform* for  
Financial Services Industries

How AHEAD Helps Deploy AI and HPC  
with NVIDIA Infrastructure

There are massive opportunities for the financial services sector to leverage AI across their lines of business. In fact, [Accenture](#) suggests that banking will likely benefit more from generative AI than any other industry, with up to a 30% productivity boost across a wide range of use cases – from risk mitigation to customer service.

Financial institutions need a predictable and scalable platform to bring these AI use cases to market. The NVIDIA DGX SuperPOD offers the most powerful data center platform for financial services companies to integrate AI and high-performance computing (HPC) into their business.

**In this whitepaper, we'll discuss the benefits of deploying NVIDIA DGX SuperPOD to accelerate AI use cases in financial services, how AHEAD has planned and deployed this high-performance infrastructure with collaboration from NVIDIA, and why leading enterprises choose to partner with AHEAD.**





## AI in Financial Services

**From algorithmic trading and risk modeling to personalized financial recommendations and automated compliance, AI has enormous potential in the financial sector. Many financial services companies are already exploring use cases to optimize banking operations, detect fraud, and enhance customer experiences.**

While AI offers transformative potential, deploying it at scale comes with significant challenges. Some AI workloads in finance and capital markets – particularly in areas like real-time fraud detection and algorithmic trading – rely on low-latency computing, which requires optimized and scalable infrastructure.

The industry also tends to have substantial data gravity, IP sensitivity, and regulatory considerations 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.

In addition, many financial institutions 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 are crucial for long term success.

# Exploring NVIDIA DGX SuperPOD

[NVIDIA DGX SuperPOD](#) provides clients with a reference architecture to build scalable, performant, and reliable AI infrastructure powered by NVIDIA DGX systems. The reference architecture is 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 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 inferencing for state-of-the-art AI models.

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

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

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

## BANK OF NEW YORK PARTNERS WITH AHEAD AND NVIDIA FOR ENTERPRISE AI

[Bank of New York \(BNY\)](#) – one of the oldest banks in the U.S. that oversees nearly \$50 trillion in assets – was the first major bank to deploy an NVIDIA DGX SuperPOD. The long-standing financial institution wanted to accelerate its AI journey and needed an infrastructure platform to increase its processing capacity for demanding workloads.

BNY worked with AHEAD and NVIDIA to install and configure an NVIDIA DGX SuperPOD faster than typical installation timelines. The NVIDIA DGX SuperPOD now supports use cases like deposit forecasting, payment automation, predictive trade analytics, and end-of-day cash balances. In addition, BNY has since identified more than 600 opportunities for AI, and dozens are already in development.

In short, the new NVIDIA DGX SuperPOD enables BNY to expand its on-premise AI infrastructure, accelerating innovation within the firm and across the global financial system. This aligns with BNY's commitment to adopting new technologies that drive business growth.





# Deploying AI Infrastructure from NVIDIA with AHEAD

Although NVIDIA DGX SuperPOD is a powerful AI platform for financial organizations, it can still be difficult for enterprises to implement without the proper expertise. In fact, data center layout, power and cooling considerations, and network and storage all have to factor into an NVIDIA DGX SuperPOD deployment.

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 large-scale infrastructure. Read on to see the AHEAD approach to an AI infrastructure engagement.



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 HPC compute that may be needed, as well as the upstream network design.



### 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 the NVIDIA DGX SuperPOD, including storage, HPC servers, security devices, and more.



### 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.



### 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, and other physical tasks to operationalize the technology.



### FLOOR PLAN DESIGN

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



### 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 accelerated computing technology. Our team of data scientists, architects, engineers, and NVIDIA AI experts 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](#).

AHEAD has the scale to stand up and manage NVIDIA DGX SuperPOD deployments 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 Grace Blackwell GB200 NVL72 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](https://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