AI, Unboxed in Minutes with Cisco Nexus Hyperfabric Full-Stack AI infrastructure Option

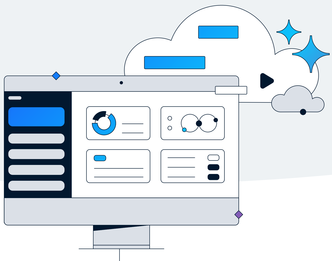

Accelerate your AI journey with simplicity and speed.

A Gartner® survey showed, “Approximately 60% of respondents reported that a typical AI initiative takes anywhere between six months to a year in progressing from prototype to production.”

Source: Gartner®, Infographic: 2025 Technology Adoption Radar for Data and Analytics: 18 July 2025
GARTNER is a trademark of Gartner, Inc. and/or its affiliates.

1

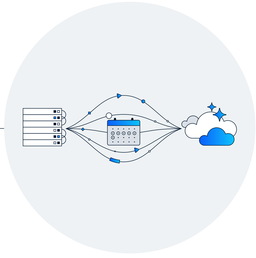

2

Cisco Nexus® Hyperfabric™ full-stack AI infrastructure

“Once our materials arrived in Texas, we fully built and configured the network in just 6 days. Processes that used to take months took just days.”

Mohammed A Jameel,
Lead AI Network Architect, Cisco Systems IT

You need AI infrastructure that works as fast as your business moves.

As per a Gartner® report, “Respondents cited technical incompatibility and/or architectural complexity as the biggest risk factor associated with deploying technologies, linking it to approximately 63% of the technologies.”

Source: Gartner®, Infographic: 2025 Technology Adoption Radar for Data and Analytics: 18 July 2025

3

4

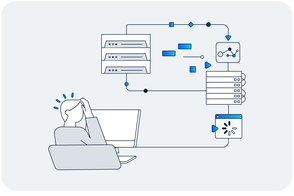

Old ways can’t keep up:

- Complexity of AI infrastructure deployment and management
- Difficulty when scaling network, compute, and storage
- Lack of integrated visibility and monitoring for AI stack
- Challenges ensuring low latency, high throughput, and reliable AI workloads
- Need for Enterprise Reference Architecture (ERA) compliant, future-ready AI infrastructure

With Cisco Nexus Hyperfabric full-stack AI infrastructure, get your first AI token in minutes, not months.

Cloud-managed, pre-integrated, and built for rapid scale–so you can focus on outcomes, not infrastructure.

- Cloud-managed fabric as a service with on premises hardware
- Launch AI workloads in days not months
- Manage everything from a single cloud dashboard
- Scale on your terms–across data center
- Reduce complexity and cost
- Rely on Cisco’s global expertise and support

5

6

Simplify AI infrastructure deployment with the cloud-managed Cisco Nexus Hyperfabric.

Accelerate AI lifecycle from design to scale with automated provisioning and monitoring.

Achieve high throughput, low latency, and scalable AI infrastructure optimized for AI workloads.

Gain full visibility and control over network, compute, GPUs, and storage components.

Reduce operational overhead and speed time-to-value for AI clusters.

Compliant with NVIDIA ERA for seamless AI model training and deployment.

Ready to accelerate your AI?

Watch analyst discussion

Visit the webpage

© 2026 Cisco and/or its affiliates. All rights reserved. Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

C26-4765569-00 02/25