

SEPTEMBER 2024

## Cisco Nexus Hyperfabric – The Rest of the Story

Jim Frey, Principal Analyst, Networking

**Abstract:** Cisco has announced early availability of Cisco Nexus Hyperfabric and the 6000 Series switches, including a scope that goes well beyond what was announced in June for AI clusters. The new Nexus Hyperfabric is much broader and much more revolutionary for Cisco, greatly expanding the types of data center network deployments that can be addressed.

### Overview: Networking for AI

When Cisco Systems announced its Nexus Hyperfabric solution at the Cisco Live event in June 2024, the focus was very much on how the new approach would address building and deploying generative AI clusters, in collaboration with GPU maker NVIDIA. Since AI was a major theme of the Cisco Live event, and AI hype was at a fevered peak, that announcement was an essential statement of strategy and direction, keeping pace with the rest of the market and the continuous stream of similar announcements by other networking and infrastructure vendors. A key aspect of that initial announcement was that Hyperfabric would vastly simplify and manage the full lifecycle of data center networks from the cloud.

The networking switch element that Cisco brought to the mix with this new offering is called the Cisco 6000 Series. Not to be confused with the end-of-life Cisco Nexus 6000 series of switches, the new Cisco 6000 Series would be a new design that would meet the specific and unique new needs of extreme high performance (400G and 800G), RDMA over Converged Ethernet. Also included were the automated configuration and management capabilities necessary to support the anticipated needs of deploying and operating GPU cluster architectures.

At the Cisco Live event, details on the Cisco 6000 were limited, and while well-rounded and comprehensive in scope, the Nexus Hyperfabric announcement was very much forward looking, promising delivery in 2025. Since the solution was aimed at enabling enterprise AI buildouts, as opposed to the hyperscalers that are the earliest generative AI adopters and deployers, the timing for solution delivery recognized that most organizations were still in the research and planning phases.

### Analysis: The Story Unfolds: Nexus Hyperfabric Broadens

Fast forward three months to September 2024, and Cisco has reached beta/early availability for Nexus Hyperfabric, although not yet for the AI-centric flavor of Nexus Hyperfabric. Details now reveal the broader strategy for Nexus Hyperfabric, and it goes far beyond what was described during the opening act. While the AI-centric version continues to develop, Cisco intends to deliver Nexus Hyperfabric for various use cases, including re-platforming existing modest-sized data centers for easier cloud-managed operational lifecycles, edge/remote data centers (such as light-touch smaller fabrics in hospitals, manufacturing and power sites, retail offices, development sites, and transportation hubs), and fabrics deployed in colocation facilities. These scenarios have not been ideally suited for the existing Cisco Nexus switch product lines, which are designed for on-premises operational models and rich customization options. In order to deliver Hyperfabric, the Cisco team has developed a new switch and fabric architecture that has very new and different characteristics:

- **Cloud-managed.** Cisco's new Cloud SaaS Controller facilitates design, including automatically prescribing a full bill of materials (equipment, cabling, power, etc.) and a deployment plan, automates configuration (switches "phone home" when connected and are automatically configured by the controller), automates validation testing to ensure consistent application of configurations and policies, and then continuously monitors the switches for health, activity, and performance. Cisco has built the UI using the same framework as the Meraki Dashboard, a proven and highly rated cloud network management solution.
- **Open and remotely programmable.** The new 6000 series switches are Cisco's first to support SONiC standard networking, supporting open network architectures. Although built on SONiC—Cisco is a consortium member—customers will not have direct OS access on these switches. The Cisco 6000 switches are only configured and managed utilizing the Cloud SaaS Controller, using either the web-based console or a full JSON API, the latter enabling integration with DevOps tools such as Ansible and Terraform. This means that there will be no command-line interface, and no reason—or means—for opening direct sessions on each individual switch.
- **Easier to use.** The new Nexus Hyperfabric is a switching solution intended to meet the common needs of the many rather than the unique needs of the few. The fabric underlay is an EVPN and VXLAN fabric which fully supports dual-stack IPv4 and IPv6, but these details are largely abstracted away from the user. Configuration is limited to a simple Layer 2 or Layer 3 setup, and all configuration details are determined and deployed by the Controller. The design objective is to enable a broader range of IT pros to engage in building, deploying, and operating the solution, and not only those with deep networking expertise.
- **Scales gracefully.** Target deployments can utilize leaf-spine architectures as well as simple mesh configurations, providing flexibility that is especially vital for distributed or edge data center deployments, such as transport hubs, hospitals, and power and manufacturing sites. These deployments can scale down to a single standalone switch but importantly, if more capacity is needed, additional switches can be added and will be automatically recognized and incorporated. This means that it will no longer be necessary to pre-plan and deploy a complete network fabric to anticipate future computing needs—rather, networking can be added as needed along the way, just in time.
- **Integrated services.** Cisco owns and manages the cloud controller and ensures that the underlay fabrics are properly configured through automation. Customers are responsible for the devices that connect to the fabric, such as hosts and upstream routing, and the networks they use. There are no on-premises or air gap options at this time, although those are under consideration for the future.

The Nexus Hyperfabric is targeted at the growing number of distributed smaller data and compute clusters that are commonplace in geographically diverse organizations, including increasing interest in edge computing stacks. It is not meant to be an alternative to networking solutions for large core data centers, where Nexus switch fabrics are still the primary offering and continue to support all of the customizations necessary for purpose-built, highly tuned and optimized connectivity. That said, it is possible to connect Nexus Hyperfabric to other Nexus fabrics, by installing a Hyperfabric switch in the non-Hyperfabric data center as a gateway.

## Conclusion

The Cisco 6000 Series switches and the Nexus Hyperfabric they will comprise are in many respects a new and fresh approach to networking for Cisco, bringing Cisco switching directly into relevance for deployment in smaller, distributed data center environments and edge computing clusters. This will enable organizations that have chosen Cisco as their strategic networking supplier to extend their Cisco solutions into parts of the network that were not a great fit for Nexus switching in the past. The new "as a service" model will accelerate ROI and reduce the cost and complexity of operations, smoothing over the burgeoning networking skills gaps faced by so many organizations.

The Nexus Hyperfabric initiative demonstrates that Cisco has been listening to market demands and is responding with open, cloud-oriented solutions that align with the evolving role of the network in the AI-native app economy. General availability will be here soon—late 2024—and Cisco is poised to deliver this new solution for data center fabrics that can be used to achieve, perhaps for the first time, a true data center networking as a service.

©TechTarget, Inc. or its subsidiaries. All rights reserved. TechTarget, and the TechTarget logo, are trademarks or registered trademarks of TechTarget, Inc. and are registered in jurisdictions worldwide. Other product and service names and logos, including for BrightTALK, Xtelligent, and the Enterprise Strategy Group might be trademarks of TechTarget or its subsidiaries. All other trademarks, logos and brand names are the property of their respective owners.

Information contained in this publication has been obtained by sources TechTarget considers to be reliable but is not warranted by TechTarget. This publication may contain opinions of TechTarget, which are subject to change. This publication may include forecasts, projections, and other predictive statements that represent TechTarget's assumptions and expectations in light of currently available information. These forecasts are based on industry trends and involve variables and uncertainties. Consequently, TechTarget makes no warranty as to the accuracy of specific forecasts, projections or predictive statements contained herein.


Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of TechTarget, is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact Client Relations at [cr@esg-global.com](mailto:cr@esg-global.com).

---

#### About Enterprise Strategy Group

TechTarget's Enterprise Strategy Group provides focused and actionable market intelligence, demand-side research, analyst advisory services, GTM strategy guidance, solution validations, and custom content supporting enterprise technology buying and selling.

 [contact@esg-global.com](mailto:contact@esg-global.com)

 [www.esg-global.com](http://www.esg-global.com)