The bridge to possible

Cisco Compute June 2024 Announcement

© 2024 Cisco and/or its affiliates. All rights reserved.

The bridge to possible

Overview

Q: What did Cisco® announce on June 4, 2024?

A: Cisco announced the:

- Cisco Compute Hyperconverged X-Series System, Cisco Compute Hyperconverged X-Series Direct, and Cisco Compute Hyperconverged X210c M7 All NVMe Nodes.
- Expanded Compute-only support leveraging existing Cisco UCS blade and rack infrastructure for the Cisco Compute Hyperconverged with Nutanix solution.
- Cisco Compute Hyperconverged with Nutanix GPT-in-a-Box solution.
- Cisco UCS X215c M8 Compute Node.
- Cisco UCS C245 M8, C225 M8 Rack Server.
- Cisco Intersight updates.

Cisco Compute Hyperconverged with Nutanix

Q: What did Cisco Announce in June 2024?

A: The Cisco Compute Hyperconverged X-Series System, Cisco Compute Hyperconverged X-Series Direct, and Cisco Compute Hyperconverged X210c M7 All NVMe Nodes.

Q: When will these new models be orderable?

A: Early orderability will begin in June 2024.

Q: How is the Cisco Compute Hyperconverged X-Series System managed?

A: The Cisco Compute Hyperconverged X210c M7 All NVMe Nodes is managed via the Cisco Intersight Infrastructure Service – Intersight Managed Mode (IMM). IMM support is planned for Q4CY'24*.

* All dates shared are roadmap targets dates and are subject to change based on program execution without prior notice.

Q: What hypervisors will be supported?

A: Nutanix Acropolis Hypervisor (AHV) and VMware vSphere (ESXi).

Q: What are the use cases for Cisco Compute Hyperconverged X-Series System and Cisco Compute Hyperconverged X-Direct?

A: Any standard workload can be supported with x86 CPUs. Some targeted use cases are:

- Virtual Desktop Infrastructure (VDI) that you can deploy close to users and ensure a smooth user experience with optional GPU rendering.
- Mission-critical enterprise applications and databases.
- AI/ML inferencing, with up to six compact GPUs per server, or up to two full-length, full-height GPUs.
- Cloud-native, containerized environments.
- Virtualized environments such as those supported by VMware vSphere (ESXi) or Nutanix Acropolis Hypervisor (AHV).

Please refer to the Cisco Compute Hyperconverged X210c M7 All NVMe Nodes specification sheet for configuration and component details.

Q: What are the key solution components for Cisco Compute Hyperconverged X-Series with Nutanix?

A: Cisco Compute Hyperconverged X-Series with Nutanix is made up of the following items:

- Nutanix Cloud Platform (NCP), which includes Nutanix Cloud Infrastructure (NCI) and Nutanix Cloud Manager (NCM).
- Cisco UCS X9508 Chassis.
- Two 9108 Intelligent Fabric Modules for unified Ethernet fabric.
- One or more of the following nodes: up to eight two-socket Cisco Compute Hyperconverged X210c M7 All NVMe Nodes.

© 2024 Cisco and/or its affiliates. All rights reserved.

  - 5th Gen mLOM -25Gps and 100Gbps are available.
- 1-node, 2-node, and 3+ node+ clusters (restrictions apply).
- Hypervisor support: Nutanix Acropolis Hypervisor (AHV) and VMware vSphere (ESXi).
- Systems management:
  - Cisco Intersight Infrastructure Service (Intersight Managed Mode).
- Optional components:
  - Cisco UCS X-Fabric modules.
  - PCIe node with GPU support used in conjunction with the Cisco UCS X-Fabric modules.
  - GPU support: NVIDIA L4, L40S, and A16.

Q: What specific drive options are available for the Cisco Compute Hyperconverged X210c M7 nodes?

A: Please refer to specification document for drive types and capacities.

- Cisco Compute Hyperconverged X210c M7 All NVMe Node data sheet.

Q: How are the nodes and clusters managed?

A: Together, Cisco and Nutanix enable you to see and control your entire hyperconverged infrastructure fleet through a single pane of glass—spanning clusters globally—with a simple and accessible cloud infrastructure operations platform: Cisco Intersight. The Cisco Compute Hyperconverged with Nutanix solution helps you deploy HCI on a global scale with automated workflows. The solution uses Cisco Intersight® to deploy and manage physical infrastructure and Nutanix Prism Central to manage your hyperconverged environment.

Q: Is there GPU support?

A: Yes, reference the specification document for specific GPU model support.

- Cisco Compute Hyperconverged X-Series System Data Sheet.
- Cisco Compute Hyperconverged X-Direct Data Sheet.
- Cisco Compute Hyperconverged X210c M7 All NVMe Node Data Sheet.

Q: What is the Cisco Compute Hyperconverged with Nutanix GPT-in-a-Box solution?

A: The Cisco Compute Hyperconverged with Nutanix GPT-in-a-Box solution takes the complexity out of adopting generative AI by providing prescriptive steps for deploying the underlying infrastructure for Nutanix GPT-in-a-Box. This reference architecture combines Cisco servers and SaaS operations with Nutanix software, utilizing the most popular Large Language Models (LLMs) to produce a fully validated AI-ready platform that can simplify and jumpstart your AI initiatives from the data center to the edge.

Q: When will the reference architecture be available?

A: A Cisco Validated Design is planned for June 2024*.

* All dates shared are roadmap targets dates and are subject to change based on program execution without prior notice.

Q: What is expanded Compute-only (CO) mode support?

A: A compute-only (CO) node allows you to expand the computing capacity (CPU and memory) of your Nutanix cluster seamlessly and efficiently. A CO node is part of a Nutanix hyperconverged cluster, and there is no Controller VM (CVM) running on the CO node or local storage. The Nutanix hyperconverged cluster uses the resources (CPUs and memory) of a CO node exclusively for computing purposes. With the Cisco Compute hyperconverged with Nutanix solution we can now accelerate infrastructure modernization, and service and manage diverse workloads leveraging existing Cisco UCS blade and rack servers with the Nutanix AHV hypervisor.

© 2024 Cisco and/or its affiliates. All rights reserved.

Q: How does Compute-only (CO) mode support benefit customers?

A: By reusing your existing UCS blade and rack infrastructure as compute-only nodes with Nutanix, it maximizes customers' return of investment of compute resources independently from storage resources. It also provides more flexibility and choice to modernize your compute infrastructure with a next-generation UCS X-Series modular compute system while retaining the flexibility to leverage X-Series as part of hyperconverged or converged infrastructure build outs, all managed through Cisco Intersight.

Q: When will Compute-only node support be available?

A: Support is planned for Q4CY24*

*** All dates shared are roadmap targets dates and are subject to change based on program execution without prior notice.**

Q: Where can I go to learn more?

A: For more information, go to: https://www.cisco.com/site/us/en/products/computing/hyperconverged-infrastructure/index.html.

Cisco UCS X215c M8 Compute Node

Q: What is the X215c M8 Compute Node?

A: The X215c M8 is the fourth compute node for the UCS X-Series Modular System. It is a single slot, two-socket server supporting 4th Gen AMD EPYC.

Q: How many cores do the 4th Gen AMD EPYC processors support?

A: The 4th Gen AMD EPYC processors support from 16 to 128 cores depending upon the model.

Q: What are some of the benefits of UCS X215c M8 Compute Node?

A: With its high core count, compute-intensive applications can see significant performance improvements as well as other benefits like power and cost efficiencies.

Q: How much memory is supported?

A: Up to 6 TB of DDR5 memory with speeds up to 4800 MT/s is supported.

Q: How many drives are supported?

A: Up to six Small-Form-Factor (SFF) front-loading hot-pluggable drives – NVMe/SAS/SATA.

Q: What network adapters are supported?

A: Cisco UCS 15000 Series Virtual Interface Cards (VICs).

Q: How is this server managed?

A: The server can be managed two ways: Cisco Intersight and Cisco UCS Manager.

Q: Can I mix and match the X215c M8 Compute Node with M6 and M7 compute nodes?

A: Yes, you may have M6, M7, and M8 compute nodes in the same chassis.

Q: Can the X215c M8 Compute Node be used with the UCS X440p PCIe Node?

A: Yes, the UCS X440p PCIe Node is supported for adding Intel or NVIDIA GPUs.

Q: Does the X215c M8 Compute Node support the GPU mezzanine?

A: Yes, the GPU mezzanine is supported.

© 2024 Cisco and/or its affiliates. All rights reserved.

The bridge to possible

Cisco UCS C225 M8 Rack Server

Q: What is the C225 M8 Rack Server?

A: The C225 M8 is a 1RU, single socket server supporting 4th Gen AMD EPYC.

Q: How many cores do the 4th Gen AMD EPYC processors support?

A: The 4th Gen AMD EPYC processors support from 16 to 128 cores depending upon the model.

Q: What are some of the benefits of UCS C-Series C225 M8 server with AMD EPYC 4th Gen CPUs?

A: Supporting double the number of cores per socket than AMD 3rd Gen EPYC-based M6 servers and delivering up to 2.8X more performance, compute-intensive applications will see significant performance improvements and will reap other benefits like power and cost efficiencies.

Q: How much memory is supported?

A: Up to 3 TB of DDR5 memory with speeds up to 4800 MT/s is supported.

Q: How many drives are supported?

A: Up to 10 x 2.5-inch SAS and SATA HDDs, SSD, NVMe drives (up to four) or up to 10 x 2.5-inch NVMe PCIe SSDs.

Q: What network adapters are supported?

A: Cisco UCS 15000 Series Virtual Interface Cards (VICs) and a variety of third-party adapters are supported.

Q: How is this server managed?

A: The server can be managed three ways: Cisco Intersight, Cisco UCS Manager, or with the Cisco Integrated Management Controller (IMC).

Cisco Intersight

Q: What new servers are supported?

A: Cisco UCS C245 M8 and UCS X-Series Direct support has been added in the SaaS and Virtual Appliance versions of Intersight.

Q: What monitoring enhancements have been added to the SaaS version?

A: A device Metrics Tab along with SFP adapter and NIF metrics.

Q: What monitoring enhancements have been added to the Virtual Appliance version?

A: Topology Views, data collection, Metrics Tab, Metrics Explorer, and a guided walkthrough.

Q: Are there new integrations with Nutanix?

A: Yes support for Nutanix with X-Series has also been added (SaaS).

Q: What other updates have been made to the Intersight Virtual Appliance?

A: Additional VMware integration and Incident Management integration with ServiceNow.

© 2024 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: 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)

C67-4384530-00 06/24