

Cisco UCS C845A M8 Rack Server

Contents

| | |
|----------------------------|---|
| Product overview | 3 |
| Features and benefits | 4 |
| Product specifications | 4 |
| Ordering information | 5 |
| Warranty information | 5 |
| Product sustainability | 5 |
| Cisco and partner services | 6 |
| Cisco Capital | 6 |

The Cisco UCS C845A M8 Rack Server is a highly scalable, flexible, and customizable AI system based on the NVIDIA MGX reference design for accelerated computing. With support for two to eight NVIDIA PCIe GPUs and NVIDIA AI Enterprise software, it delivers high performance for a wide range of AI workloads – including Generative AI fine-tuning and inference.

Product overview

Built on the NVIDIA MGX modular reference design, the Cisco UCS C845A M8 Rack Server is a flexible, scalable, and customizable AI system capable of growing as your AI needs grow. Configure with 2, 4, 6, or 8 PCIe GPUs to address a multitude of workloads ranging from generative AI, graphics and rendering, to virtual desktop infrastructure.



UCS C845A M8 servers can be configured with two to eight NVIDIA GPUs. Depending on the configuration, customers can choose between the PCIe-based NVIDIA H100 NVL, H200 NVL, or L40S GPU. Thanks to the sophistication of the MGX design, other next-generation NVIDIA GPUs are planned to be introduced on this same platform as these GPUs become available.

With a compute node powered by AMD’s new high-end Turin (5th Gen) CPUs, designed specifically for AI workloads, the UCS C845A M8 provides a no-compromise solution for CPU or GPU performance required to avoid bottlenecks within an AI server. Another benefit is the ability to configure the server with the NVIDIA ConnectX-7 NIC and/or NVIDIA BlueField-3 DPUs to handle data traffic in and out of the server.

Systems equipped with NVIDIA H100 NVL or H200 NVL GPU come with a five-year license for NVIDIA AI Enterprise, a cloud-native software platform that streamlines development and deployment of production-grade AI solutions, including AI agents, generative AI, computer vision, speech AI, and more. Easy-to-use microservices optimize model performance with enterprise-grade security, support, and stability, ensuring a smooth transition from prototype to production for enterprises that run their businesses on AI.

Using Cisco Intersight®, customers can see, control, and automate their compute, storage, and networking infrastructure, wherever it is, from one place just as they can do today with their other Cisco UCS® blade or rack servers.

Features and benefits

Table 1. Summary of features and benefits of the Cisco UCS C845A M8 Rack Server

| Feature | Benefit |
|--|---|
| Scalable and flexible | <ul style="list-style-type: none"> • Rapid deployment of AI and accelerated computing technologies • Perfectly suited for any data -center environment |
| Built on NVIDIA MGX design | <ul style="list-style-type: none"> • Modular and future-ready design that can adapt to evolving technology needs • Next-generation GPUs can be integrated without purchasing a new platform |
| Dense design | <ul style="list-style-type: none"> • Drive 8x NVIDIA GPUs, 2x AMD Turin CPUs, and 5x PCIe slots in a 4RU chassis |
| Cisco Intersight support | <ul style="list-style-type: none"> • See, control, and automate your compute, storage, and networking infrastructure throughout its lifecycle |
| E1.S drives for local storage | <ul style="list-style-type: none"> • Increased storage density, improved thermal management, and hot-swappable drives • High performance, lower power consumption, and compact form factor |
| Addition to NVIDIA MGX reference design | <ul style="list-style-type: none"> • Improved power delivery, fewer PCBs, and improved cable routing |

Product specifications

Table 2. Specifications of the Cisco UCS C845A M8 Rack Server

| Item | Specifications |
|-------------------------|--|
| Form factor | 4RU Air-cooled rack server |
| Processors | 2x 5th Gen AMD EPYC 9475F 48 cores, 3.65GHz (max boost 4.8 GHz) or 2x 5th Gen AMD EPYC 9655 96 cores, 2.6GHz (max boost 4.5 GHz) or 2x 5th Gen AMD EPYC 9575F 64 cores, 3.3GHz (max boost 5 GHz) |
| Memory | 8x 64GB DDR5 up to 5,200 MT/s RDIMMs @1DPC or 16x 64GB DDR5 up to 5,200 MT/s RDIMMs @1DPC or 16x 96GB DDR5 up to 5,200 MT/s RDIMMs @1DPC or 32x 64GB DDR5 up to 4,400 MT/s RDIMMs or 32x 96GB DDR5 up to 4,400 MT/s RDIMMs |
| GPUs | Up to 8x NVIDIA H200 NVL or Up to 8x NVIDIA H100 NVL or Up to 8x NVIDIA L40S |
| Boot drives | 2x 960GB M.2 SATA SSDs (includes hardware RAID) |
| Internal storage | Up to 20x E1.S NVMe PCIe drives |

| Item | Specifications |
|---|---|
| PCIe slots | 1x PCIe Gen5x16 FHHL for north-south traffic – NVIDIA ConnectX-7 (1x400G, 2x200G), NVIDIA BlueField-3 B3220 DPU, NVIDIA BlueField-3 3240 (double slot) 4x PCIe Gen5x16 FHHL for east-west traffic – NVIDIA ConnectX-7 or NVIDIA BlueField-3 3140H SuperNIC |
| OCP 3.0 slot | Intel® X710-DA2 (dual-port 10GBaseT) for x86 management |
| Power supply | 4x 3.2kW 80Plus 12V M-CRPS hot-swappable redundant PSUs with N+1 Redundancy |
| Management | Cisco Intersight |
| Hardware and software interoperability | See the Cisco Hardware and Software Interoperability lists for a complete listing of supported operating systems and peripheral options |

Ordering information

This information will be provided in the second quarter of 2025 when the UCS C845A M8 will be orderable.

Warranty information

Cisco UCS C845A M8 Rack Servers have a three-year Next-Business-Day (NBD) hardware warranty and a 90-day software warranty.

Product sustainability

Information about Cisco’s Environmental, Social, and Governance (ESG) initiatives and performance is provided in Cisco’s CSR and sustainability [reporting](#).

Table 3. Cisco environmental sustainability information

| Sustainability topic | | Reference |
|----------------------|--|---|
| General | Information on product-material-content laws and regulations | Materials |
| | Information on electronic waste laws and regulations, including our products, batteries, and packaging | WEEE Compliance |
| | Information on product takeback and reuse program | Cisco Takeback and Reuse Program |
| | Sustainability inquiries | Contact: csr_inquiries@cisco.com |
| Material | Product packaging weight and materials | Contact: environment@cisco.com |

Cisco and partner services

Cisco and our industry-leading partners deliver services that accelerate your transition to Cisco UCS solutions for AI and High-Performance Computing (HPC). Cisco Unified Computing Services™ can help you create an agile infrastructure, accelerate time to value, reduce costs and risks, and maintain availability during deployment and migration. After deployment, our services can help you improve performance, availability, and resiliency as your business needs evolve, and help you further mitigate risk.

For more information, visit <https://www.cisco.com/go/unifiedcomputingservices>.

Cisco Capital

Flexible payment solutions to help you achieve your objectives

Cisco Capital® makes it easier to get the right technology to achieve your objectives, enable business transformation and help you stay competitive. We can help you reduce the total cost of ownership, conserve capital, and accelerate growth. In more than 100 countries, our flexible payment solutions can help you acquire hardware, software, services and complementary third-party equipment in easy, predictable payments. [Learn more](#).

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

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)