


Cisco UCS C845A M8 Rack Server

Contents

Product overview	3
Features and benefits	5
Product specifications	6
Ordering information	7
Warranty information	7
Product sustainability	7
Cisco and partner services	7
Cisco Capital	7
Document history	8

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 (2) to eight (8) NVIDIA or AMD PCIe GPUs and NVIDIA AI Enterprise software, it delivers high performance for a wide range of AI workloads – including Generative AI fine-tuning, RAG and inference.

Product overview



The Cisco UCS C845A M8 Rack Server is designed to address the most demanding AI workloads. Now an integral component of Cisco® AI PODs (Cisco Validated Designs for AI), the Cisco UCS C845A M8 provides a robust foundation for modern AI infrastructure, allowing organizations to easily scale their AI capabilities with confidence.

With support for 2, 4, 6, or 8 NVIDIA GPUs—including the NVIDIA RTX Pro 6000 Blackwell, H100 NVL, H200 NVL, L40S, and also supports AMD Instinct MI210 GPUs this system offers unparalleled flexibility to meet the diverse needs of enterprises. Leveraging the sophistication of the MGX modular reference design, this platform is also future-ready, with next-generation NVIDIA GPUs expected to seamlessly integrate as they become available.

As part of Cisco AI PODs, the UCS C845A M8 is designed to accelerate workloads such as large-scale deep learning, Large Language Model (LLM) training, model optimization, and AI inference. It supports hybrid AI workflows across cloud and edge environments, making it the perfect solution for enterprises working on cutting-edge applications such as AI agents, conversational AI, intelligent search, and virtual desktops. Enhanced by 5th Gen AMD EPYC "Turin" CPUs and NVIDIA BlueField-3 DPUs, this server eliminates performance bottlenecks, ensuring optimal CPU and GPU utilization for data-intensive workloads.

With seamless integration into Cisco Intersight®, customers can manage their entire infrastructure—including compute, storage, and networking—through a unified interface. The Cisco UCS C845A M8, now part of Cisco AI PODs, is the ideal platform for organizations seeking to scale their AI initiatives with confidence and efficiency.

What use cases does the Cisco UCS C845A M8 address?

The Cisco UCS C845A M8, a highly scalable and customizable server integrated into Cisco AI PODs, is engineered to drive a multitude of AI workloads. Its flexible GPU configurations enable it to address the most demanding AI challenges, including large deep learning, Large Language Model (LLM) training, model fine-tuning, large model inferencing, and Retrieval-Augmented Generation (RAG).

The platform's versatility is further enhanced by its support for various GPUs, each optimized for specific market needs:

NVIDIA H100 NVL:

- Large-scale High-Performance Computing (HPC) and AI workloads
- Advanced AI research and high-performance computing
- Generative AI inference for large language models
- High-performance AI training

NVIDIA H200 NVL:

- High-performance LLM inference
- Generative AI training and fine-tuning

NVIDIA RTX Pro 6000 Blackwell:

- Agentic and physical AI: powering autonomous systems, smart factories, and robotics for real-time decision-making
- Advanced scientific computing and rendering: accelerating complex simulations, medical imaging, and engineering analysis across hybrid environments
- High-fidelity 3D graphics and video: driving content creation, post-production, and immersive VR/AR experiences
- Hybrid AI/ML workflows: enabling seamless training, inferencing, and AI-assisted graphics across on-premises and cloud
- Edge-to-core AI applications: supporting real-time AI at the edge with centralized management and cloud integration

AMD Instinct MI210:

- High-Performance Computing (HPC): accelerating scientific research, simulations, and complex modeling
- Energy-efficient AI/ML workloads: supporting deep learning training and inferencing with a focus on power efficiency
- Large-scale data analytics: speeding up data processing and analytics for enterprise applications
- Hybrid AI inference: providing balanced compute and efficiency for inferencing tasks in diverse environments
- Specialized simulation workloads: enhancing performance for engineering, manufacturing, and bioinformatics simulations



NVIDIA L40S:

- Generative AI foundation model fine-tuning
- Deployment of intelligent chatbots and search tools
- Language processing and conversational AI
- Graphics, rendering, and NVIDIA Omniverse applications
- Virtual desktop infrastructure

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• 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 9355 32 cores, 3.55GHz (max boost 4.4 GHz) or 2x 5th Gen AMD EPYC 9375F 32 cores, 3.85GHz (max boost 4.8 GHz) or 2x 5th Gen AMD EPYC 9455 48 cores, 3.15GHz (max boost 4.4 GHz) or 2x 5th Gen AMD EPYC 9475F 48 cores, 3.65GHz (max boost 4.8 GHz) or 2x 5th Gen AMD EPYC 9555 64 cores, 3.2GHz (max boost 4.4 GHz) or 2x 5th Gen AMD EPYC 9575F 64 cores, 3.3GHz (max boost 5 GHz) or 2x 5th Gen AMD EPYC 9655 96 cores, 2.6GHz (max boost 4.5 GHz)
Memory	32 DIMM slots (16 DIMMS per CPU) 32, 64, 96, 128GB DDR5 at up to 5200 MT/s for up to 4TB memory system capacity Refer to the Cisco UCS C845A M8 AI Memory Guide .
GPUs	Up to 8x NVIDIA RTX Pro 6000 Blackwell Server GPU or Up to 8x NVIDIA H200 NVL or Up to 8x NVIDIA H100 NVL or Up to 8x NVIDIA L40S or Up to 8x AMD Instinct MI210
Boot drives	2x 960GB M.2 SATA SSDs (includes hardware RAID) 2x 240GB M.2 SATA SSDs (includes hardware RAID)
Internal storage	Up to 20x 1.9TB E1.S 15 mm NVMe PCIe drives Up to 20x 3.8TB E1.S 15 mm NVMe PCIe drives Up to 20x 7.6TB E1.S 15 mm NVMe PCIe drives
PCIe slots	One PCIe Gen5x16 FHHL for north-south traffic – NVIDIA ConnectX-7 (2x200G), NVIDIA BlueField-3 B3220 DPU 4x PCIe Gen5x16 FHHL for east-west traffic – NVIDIA ConnectX-7 (1x400G) 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

Part #	Product description
UCS-MGPUM8-MLB	Cisco UCS-845A M8 Rack Server chassis Major Line Bundle (MLB) This MLB consists of the server node (UCSC-C845A-M8) with software and Cisco Intersight options. Use this PID to begin a new configuration.

For more information, check the [Cisco UCS C845A M8 Rack Server Ordering Guide](#).

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](#).

Our experts recommend

- [BIOS Performance and workload: Tuning guide for Cisco UCS M8 Platforms White Paper](#)
- [Cisco AI PODs At-a-Glance](#)
- [Cisco UCS Servers with AMD EPYC](#)

Document history

New or Revised Topic	Described In	Date
New Product overview, specifications and Ordering information	UCS C845A M8 Data Sheet	July 2025

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)