

FUJITSU


shaping tomorrow with you



FUJITSU Server PRIMERGY CX400 M4

Workload-specific power in a modular form factor

FUJITSU Server PRIMERGY CX400 M4




Workload-specific power in a compact and modular form factor

The FUJITSU Server PRIMERGY CX400 M4 is a modular enclosure for the Fujitsu multi-node ecosystem that combines the density and efficiency of blade-like servers with the simplicity and cost benefits of rack-based systems. The PRIMERGY CX server nodes provide the power for ***VDI, infrastructure virtualization, High Performance and Technical Computing, Web Services, Cloud Computing, deep learning but also as a foundation for hyper-converged infrastructures.***

FUJITSU Server PRIMERGY CX400 M4



Density optimized server infrastructures for HPC, cloud and hyper-converged computing.



PRIMERGY CX400

PRIMERGY CX2550

PRIMERGY CX2560

PRIMERGY CX2570

Workload-specific power in a compact and modular form factor



Cloud/HPC optimized half-width 1U server node




All-round half-width 1U server node



HPC/VDI optimized half-width 2U server node




Comprehensive support services available

FUJITSU Server PRIMERGY CX400 M4

Server Node Overview



PRIMERGY CX400 M4 Enclosure

Front view



Rear view



Optimized server nodes for a variety of workloads

PRIMERGY CX2550 M4

- Use case: HPC node
- Up to 4 nodes/chassis
- 2x CPU's (3 UPI links / up to 205W)
- 16x DIMMs
- Omni-Path support
- 2x 2.5" storage drives per node
- 2x PCIe Gen 3 slots
- Liquid cooling (option)



PRIMERGY CX2560 M4

- Use case: All-round/Cloud node
- Up to 4 nodes/chassis
- 2 CPU's (3 UPI links / up to 150W)
- 16x DIMMs
- 6x 2.5" storage drives per node
- DynamicLoM OCP
- 2x PCIe Gen3 slots
- Air-cooling




PRIMERGY CX2570 M4

- Use case: HPC/VDI node
- Up to 2 nodes/chassis
- 2 CPU's (3 UPI links/ up to 205W)
- 16x DIMMs
- 6x 2.5" storage drives per node
- 2x GPU (PCIe) or 4x GPU (SXM2)
- DynamicLoM OCP
- 1x PCIe slots
- Liquid cooling (option)



FUJITSU Server PRIMERGY CX2550 M4



Air cooling /
optional liquid cooling



Usage Scenarios


- Scale-out cloud infrastructure
- High Performance and Technical Computing
- Mainstream enterprise workloads



Overview

- Cost optimized half-width 1U server node
- Up to 4x CX2550 M4 per CX400 M4 chassis
- 2x Intel® Xeon® Processor Scalable Family
 - Omni Path support for CPU1 (*planned*)
 - 3x Ultra Path Interconnect (UPI) links
- 16 DIMMs per node (up to 2,048 GB DDR4 memory)
- 2x 1GbE (Mgmt. LAN / Onboard NIC)
- 2x PCIe slot (Gen3 x16)
- 2x 2.5" SATA drives
- Internal Boot Devices: 1x mirrored USB / 2x M.2 slots (PCIe/SAS)
- 1x TPM / 1x MicroSD for BMC

FUJITSU Server PRIMERGY CX2560 M4



Air cooling



Usage Scenarios

- Standard / All-round server node (e.g. mainstream enterprise workloads, web serving, hosting, virtualization)
- High Performance and Technical Computing



Overview

- Condensed half-width 1U server node
- Up to 4x CX2560 M4 per CX400 M4 chassis
- 2x Intel® Xeon® Processor Scalable Family (3x Ultra Path Interconnect (UPI) links)
- 16 DIMMs per node (up to 2,048 GB DDR4 memory)
- 1x DynamicLoM OCP + 2x PCIe slot (Gen3 x16)
- 2x 1GbE (Mgmt. LAN / Onboard NIC)
- Up to 6x 2.5" SAS/SATA (thereof 2x PCIe)
- Internal Boot Devices: 1x mirrored USB / 2x M.2 slots (PCIe/SAS)
- 1x TPM / 1x MicroSD for BMC

FUJITSU Server PRIMERGY CX2570 M4



Air cooling /
Liquid cooling for SXM2



Usage Scenarios

- Virtualized Desktop Infrastructures
- High Performance and Technical Computing




Overview

- Half-width 2U server node
- Up to 2x CX2570 M4 per CX400 M4 chassis
- 2x Intel® Xeon® Processor Scalable Family (3x Ultra Path Interconnect (UPI) links)
- 16 DIMMs per node (up to 2,048 GB DDR4 memory)
- 2x GPU (PCIe) or 4x GPU (SXM2 with LC)
- 1x DynamicLoM OCP + 1x PCIe slot (Gen3 x16)
- 2x 1GbE (Mgmt. LAN / Onboard NIC)
- Up to 6x 2.5" SAS/SATA (thereof 2x PCIe)
- Internal Boot Devices: 1x mirrored USB / 2x M.2 slots (PCIe/SAS)
- 1x TPM / 1x MicroSD for BMC

Boost performance for HPC and hyperscale workloads

- Accelerate most demanding HPC, hyperscale data center workloads as well as improve virtual desktops and applications
- Horsepower needed to run bigger simulations faster than ever before
- PRIMERGY CX2570 M4 supports up to 4 GPU accelerators with SXM2, or 2 GPUs with PCIe Full Height/Length form factor
- Available options
 - NVIDIA Tesla P100/V100 (for SXM2)
 - NVIDIA Tesla P100/V100 (for PCIe)
 - NVIDIA Tesla M60
 - NVIDIA Tesla M10



Graphic Processor Options

NVIDIA Tesla P/V100 (for SXM2) NVIDIA Tesla P/V100 (for PCIe)

World's most advanced data center GPU

- max. 4x per system
- Superior strong-scaling performance for HPC and hyperscale applications
- 3584 NVIDIA CUDA Cores
- GPU Memory: 16 GB CoWoS HBM2
- Double-Precision Performance: 5.3 TeraFLOPS
- Single-Precision Performance: 10.6 TeraFLOPS



- max. 2x per system
- Boost throughput and save money for HPC and hyperscale data centers
- 3584 NVIDIA CUDA Cores
- GPU Memory: 16GB CoWoS HBM2
- Double-Precision Performance: 4.7 TeraFLOPS
- Single-Precision Performance: 9.3 TeraFLOPS



NVIDIA Tesla M60

Designed for data centers that are virtualizing desktop graphics

- max. 2x per system
- Performance-Optimized Graphics Virtualization
- 4096 NVIDIA CUDA Cores (2048 per GPU)
- 16 GB of GDDR5 Memory (8 per GPU)
- PCIe 3.0

NVIDIA Tesla M10

- max. 2x per system
- Density-Optimized Graphics Virtualization
- 2560 NVIDIA CUDA Cores (640 per GPU)
- 32 GB of GDDR5 Memory (8 per GPU)
- PCIe 3.0

