Overview

HPE Cray EX Supercomputer

The HPE Cray EX supercomputer is a liquid cooled blade-based, high-density clustered computer system designed from the ground up to deliver the utmost in performance, scale, and density. The basic building blocks of the HPE Cray EX Supercomputer is the EX4000 liquid cooled cabinet or EX2500 liquid cooled rack. The EX4000 cabinet is a sealed unit using closed-loop cooling technology and does not exhaust heated air into the data center. The EX2500 offers the same sealed closed-loop cooling capabilities as the EX4000. However, it also offers the ability to use part of the rack for air cooled hardware in certain configurations. In both configurations, direct attached liquid cooled cold plates provide for efficient heat removal from high power devices including processors, GPUs, and switches via a cooling distribution unit (CDU).

HPE Cray EX Supercomputer - Closed and With Trim

Standard Features

HPE Cray EX4000 Details

A single cabinet can accommodate up to 64 compute blade slots within 8 compute chassis. The cabinet is not configured with any cooling fans. All cooling needs for the cabinet are provided by direct liquid cooling and the CDU. This approach to cooling provides greater efficiency for the rack-level cooling, decreases power costs associated with cooling (no blowers) and utilizes a single water source per CDU

One cabinet supports the following:

- 8 compute chassis
- 4 power shelves with a maximum of 8 rectifiers per shelf- 32 total 15kW rectifiers per cabinet. **Notes:** 1 rectifier per shelf is used for redundancy
- 4 PDUs (1 per power shelf)
- 4 power input whips (3-phase)
- Maximum of 64 compute blades
- Maximum of 64 HPE Slingshot switch blades

Single Cabinet – Doors Closed Without CDU

Standard Features

HPE Cray EX2500 Details

A single rack can accommodate up to 24 compute blade slots within 3 compute chassis. In this configuration, the rack is not configured with any cooling fans. All cooling needs for the cabinet are provided by direct liquid cooling and the CDUs. This approach to cooling provides greater efficiency for the rack-level cooling, decreases power costs associated with cooling (no blowers) and utilizes a single water source per CDU

One rack supports the following:

- Up to 3 compute chassis
- Up to 3 power shelves with 4 rectifiers per shelf - 12 total 15kw rectifiers per rack. **Notes:** 1 rectifier per shelf is used for redundancy
- 3 PDUs (1 per power shelf)
- 3 power input whips (3-phase)
- Maximum of 24 compute blades
- Maximum of 24 HPE Slingshot switch blades
- Up to 2 in-rack CDUs for a max of 141KVA cooling capacity **Notes:** 1 CDU is included by default and will not show up as a line item in the configuration. The second CDU will show up as a line item in the configuration.

Single EX2500 rack – Doors Closed

Standard Features

Compute Chassis

The compute chassis is a mechanical assembly that provides power, cooling, system control, and network fabric for up to 8 compute blade slots. 8 chassis are installed in the EX4000 and up to 3 chassis are installed in the EX2500.

The features of the compute chassis are as follows:

- 8 compute blade slots
- 8 HPE Slingshot switch blade slots
- One power/signal midplane

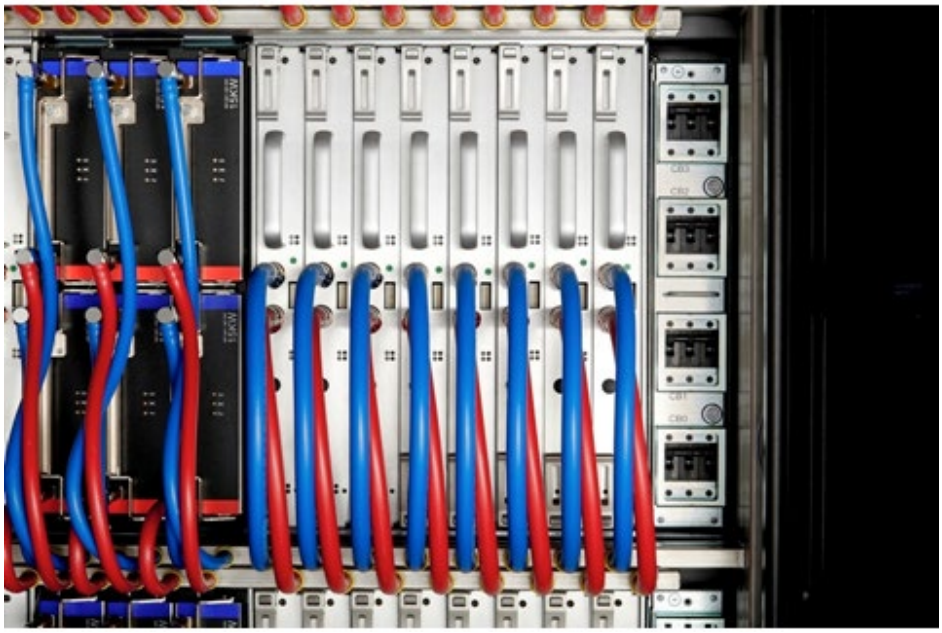

Compute Blades

Blades have three basic sections: computation, memory, and I/O and consume one blade slot in the compute chassis. The following blades are designed for the HPE Cray EX Supercomputer

HPE Cray EX425

The features of this compute blade are as follows:

- Four 2-socket CPU nodes
- Support for the full 2nd Gen AMD EPYC™ 7002 or 3rd Gen AMD EPYC™ 7003 series processor stack
- 8 DIMMs per CPU socket (1DPC)
- Up to 64 GB DIMMs at up to 3200 MT/s
- Up to 8 HPE Slingshot 200Gbit/sec ports per blade
- 2 Board Management Controllers (BMC) per blade
- Cooled with cold plate

Standard Features

HPE Cray EX4252

The features of this compute blade are as follows:

- Four 2-socket CPU nodes
- Support for the full 4th Gen AMD EPYC™ 9004 series processor stack
- 12 DIMMs per CPU socket (1DPC)
- Up to 64 GB DIMMs at up to 4800 MT/s
- Up to 8 HPE Slingshot 200Gbit/sec ports per blade
- 0 or 1 local NVMe M.2 SSD per node (up to 4 per blade)
- 2 Board Management Controllers (BMC) per blade
- Cooled with cold plate

HPE Cray EX420

The features of this compute blade are as follows:

- Four 2-socket CPU nodes
- Support for the full 4th Gen Intel Xeon scalable processor stack
- 8 DIMMs per CPU socket (1DPC)
- Up to 64 GB DIMMs at up to 4800 MT/s
- Up to 8 HPE Slingshot 200Gbit/sec ports per blade
- 0 or 1 local NVMe M.2 SSD per node (up to 4 per blade)
- 2 Board Management Controllers (BMC) per blade
- Cooled with cold plate

HPE Cray EX235n

The features of this compute blade are as follows:

- Two 4-socket A100 GPUs x 1-socket CPU nodes
- Support for the full AMD 3rd Gen AMD EPYC™ 7003 series processor stack
- 8 DIMMs per CPU socket (1DPC)
- Up to 64 GB DIMMs at up to 3200 MT/s
- Up to 8 HPE Slingshot 200Gbit/sec ports per blade
- 2 Board Management Controllers (BMC) per blade
- Cooled with cold plate

HPE Cray EX235a

The features of this compute blade are as follows:

- Two 4x AMD Instinct™ M250X Accelerators x 1-socket CPU nodes
- Support for the 3rd Gen AMD EPYC™ processor
- 8 DIMMs per socket (1DPC)
- Up to 64 GB DIMMs at up to 3200 MT/s
- Up to 8 HPE Slingshot 200Gbit/sec ports per blade
- 2 Board Management Controllers (BMC) per blade
- Cooled with cold plate

Standard Features

Switch Chassis

The Slingshot switch blades are in the switch chassis and mounted to the rear of the compute chassis. The purpose of the switch chassis is to provide a structure for orthogonally mounting the switch blades to the compute chassis. There is no active backplane connecting the switches to the compute blades. Each compute blade directly connects to one or more switch blades in the switch chassis using a cableless connection. The switch chassis supports a maximum of eight Slingshotswitch blades.

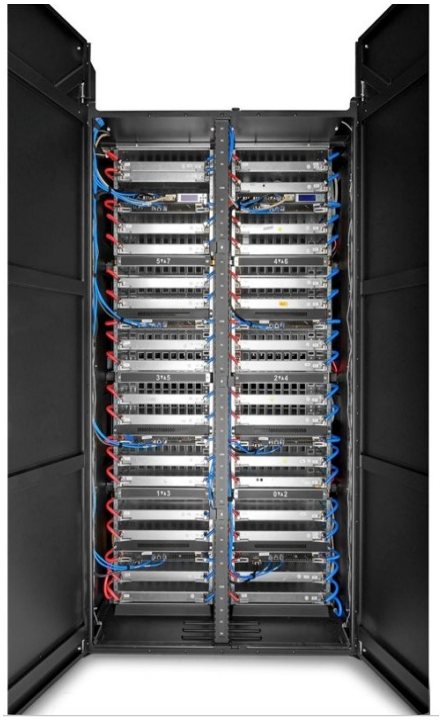

Compute Cabinet Rear View with Doors Open

HPE Slingshot switch blade

The following switch blade types are supported

The HPE Slingshot Switch blade is a 64 x, 200Gb/sec port switch designed for the switch chassis. Switch blades are inserted into the rear of the cabinet and provide the high-speed network interface for the compute blades. Each switch blade connects to all eight compute blades through orthogonal connectors and provides fabric connections using copper and optical cables through its faceplate to expand the network.

Each switch blade has 8 local connectors, with two ports per connector for a total of 16 downlink ports to the eight compute blades. Each switch blade also has 24 QSFP-DD connectors with two 200 Gb/sec ports per connector for a total of 48 copper or optical connections to the other switch blades in the cabinet and other cabinets to create the high-speed fabric. The number of switch blades in a chassis depends on the number of NIC injection points per compute node. For two-socket CPU blades such as the HPE Cray EX425 or EX4252 or EX420, a typical configuration would be two switch blades in the switch chassis to support the 4 nodes on the compute blade via the two port NIC Mezzanine card. These two switches would service this configuration across all 8 compute blades in the chassis which consists of 32 nodes . For GPU based blades, a typical configuration is four switch blades in order to server one NIC per GPU.

Standard Features

HPE Slingshot NIC Cards

Compute blades are outfitted with a mezzanine card that provides compute blade NIC connectivity to the Slinghot switch blades. Each mezzanine card support two NICs through two PCIe buses. Each mezzanine card is 100% direct liquid cooled through a cold plate. An "L0" copper cable connects each NIC mezzanine card to the connector on the back of the compute blade. The compute blade will be the determinant for which mezzanine cards are supported and the number of mezzanine cards that are required per blade and whether there are options to configure the injection bandwidth per node (meaning the number of NICs).

The following mezzanine card type are offered:

Dual port 100Gbit/sec Mellanox ConnectX-5 Ethernet Mezzanine Card

This mezzanine card is a two-port NIC design. Each of the two 100Gb/s Ethernet NIC port connects to a host system using a PCIe Gen3 a x16 channel. Solutions with this NIC can utilize programming environment support for Cray MPI, OpenSHMEM, Chapel, UPC, as well as third-party software that utilizes the RDMA-Over-Converged-Ethernet (RoCE) protocol.

Dual port 200 Gbit/sec HPE Slingshot Mezzanine Adapter

The HPE Slingshot 200 Gbit/sec mezzanine card is based on HPE silicon designed to deliver supercomputer-class performance and scalability with the HPE Slingshot Switches. The NICs incorporate hardware features that extensively offload message processing across the range of HPC workloads and message sizes to ensure compute resources operate with high efficiency due to less interruption and latency from message and communications processing. The NICs also seamless complement operation with high performance networking technologies built into the Slingshot switches to optimize the end-to-end solution for congestion management, fine grain flow control, and traffic classes. The host interface is PCIe Gen4 x16 with support for extended speed mode at 25GT/s (where supported by the CPU or GPU).

A full suite of HPC and Ethernet features is supported including strong progression of MPI message matching and data transfer, along with programming environment support for Cray MPI, OpenSHMEM, Chapel, UPC, and industry software that can support the HPE Slingshot NIC through the Libfabric interface.

Topology

The HPE Cray EX Supercomputer operates using the dragonfly topology for the high speed network fabric created using the Slingshot switches. . This provides a lower cost, efficient, and highly scalable approach over alternative topology designs. Any two endpoints in even the largest supercomputer can be reached with 3 or fewer switch to switch hops which reduces latency. The design maximizes use of high-speed copper cables by utilizing all-to-all connected switches in an EX cabinet as a dragonfly "group", each of which in turn uses optical cables to connect all the groups in different cabinets. This creates substantial efficiency by reducing the use of expensive optical cables and enabling multiple paths to deliver very high global bandwidth.

Various dragonfly switch group sizes are available and selected depending on the requirements to optimize across lowest system cost (by maximizing the number of copper cables), maximum system scale, and the desired scalable unit of deployment. For EX4000 configurations, it is typical to configure the entire cabinet as one switch group Here the two typical configurations are the 16 switch group (2 switches in each of the 8 chassis in a cabinet) or a 32 switch group 4 switches in each of the 8 chassis in a cabinet). It is possible to implement more than one group in a cabinet, such as configuring a cabinet that had 4 switches in each of the 8 chassis in the cabinet as two 16-switch groups instead of a single 32-switch group.

Standard Features

In the EX2500 cabinet, because the maximum system size is usually smaller than for EX4000. Therefore it is typical to configure each of the chassis as a group, although for larger systems build with EX2500, groups sizes that aggregate switches in several chassis together - within the same or adjacent racks – are supported.

The chart below provide a reference for the group sizes supported:

Switches per group	Maximum Number of Groups (with no global bandwidth tapering)	Maximum scale (here captured as number of NIC endpoints)	Example Use In Cray EX Supercomputers
2	33	1056	One EX 2500 chassis populated with CPU blades
4	49	3,136	One EX 2500 chassis populated with CPU blades
8	81	10,368	Multi-chassis group to use EX2500 for large system
16	145	37,120	Fully populated (8 chassis) EX4000 cabinet with CPU blades
32	257	263,168	Fully populated (8 chassis) EX4000 cabinet with GPU blades

CDU (Cooling Distribution Unit) for EX4000

The cooling distribution unit (CDU) is a liquid-to-liquid heat exchanger that is used to remove heat from HPE Cray EX Supercomputer. The CDU uses a secondary loop to circulate a heat transfer liquid to the cold sinks. The heat captured in the secondary loop is transferred to the facilities primary loop via a liquid-to-liquid heat exchanger.

The CDU is designed to circulate and control the heat transfer fluid to the manifolds that are in each chassis in the cabinet. The CDU is rated for 1.6MW of cooling. One CDU supports a maximum of four cabinets

The CDU consists of a cabinet that includes a heat exchanger, circulating pumps, control valve, sensors, controller, valves, and piping. The CDU monitors room conditions and prevents condensation by maintaining the secondary loop at a temperature above the room's dew point.

All functions, such as switching pumps (if applicable), controlling water temperature, etc., are managed by the controller using user defined settings.

Standard Features

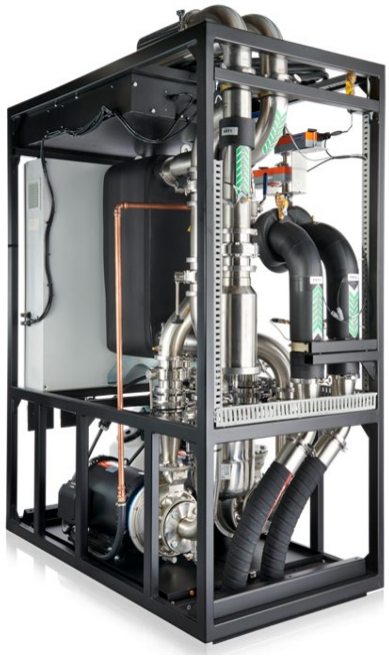

CDU Side View with no Doors or Skins

Software Stack

HPE Cray EX supercomputers are complete solutions with software and hardware that are tightly integrated and performance-tuned to offer the best system performance while bringing new standards in flexibility, manageability, and resiliency to supercomputing.

Cray supercomputer software stacks address the needs of both system administrators, developers, and end-users.

Administrative Software

HPE Cray supercomputer users now have the option to choose either HPE Cray System Management or HPE Performance Cluster Manager.

HPE Cray System Management - a built-for-scale system management solution offering administrators all functionalities they need to keep the HPE Cray EX system healthy, utilized to the maximum and accommodating wide range of workload requirements via –aaS experience. The software is built to manage systems which can scale to Exascale deployments featuring:

- Comprehensive monitoring and management of all aspects of the system: CPU/GPU, network (integrated HPE Slingshot Fabric Manager), storage as well as power management and monitoring combined with provisioning for operational efficiency.
- Partitioning and batch or container orchestration enable customers to run a variety of HPC/AI/HPDA workloads the way that makes the best use of their system without logistical constraints.
- REST APIs & standard protocols enable full interoperability with existing monitoring, management, and automation toolsets.

HPE Performance Cluster Manager - a comprehensive, flexible HPC system management solution that enables fast setup, provisioning and monitoring including the following features

- Hardware discovery and Linux operating system installation for compute and service nodes
- Inventory management
- Telemetry data collection and analysis
- Alert monitoring and component diagnosis
- Power resource monitoring and management
- Software image management **Developer Software**

Standard Features

HPE Cray Programming Environment – is a fully integrated software development suite offering programmers comprehensive set of tools for developing, porting, debugging, and tuning of their applications so they can shorten application development time and accelerate their performance.

The programming environment is designed to make porting of existing applications easier with minimal recording and changes to the existing programming models to simplify transition to the new hardware architectures and configurations, such as HPE Cray EX systems.

Operating System

HPE Cray OS is a compute operating system based on SLES with enhancements. The enhancements provide customers with capabilities specific to supercomputing and high-performance computing fully supported by HPE Pointnext. These modifications don't alter the ability to run standard Linux applications, but rather enhance it for performance, scale, and reliability. We integrate and test these materials together and package releases.

Overall

While HPE Cray System Management and HPE Cray Operating System are designed to support HPE Cray EX systems with HPE Slingshot, HPE Cray Programming Environment product also supports other HPE and HPE Cray HPC systems (using InfiniBand interconnect).

The software stack is supported by HPE Pointnext Services.

Features

HPE Cray EX Supercomputer	
Operating system	• HPE Cray Operating System or RedHat RHEL (HPE Performance Cluster Manager only)
System Management and Fabric software	• HPE Cray System Management or HPE Performance Cluster Manager • HPE Slingshot Network Manager
Workload Management and Orchestration	• Altair® PBS Professional • Slurm Workload Manager • Containers: Singularity & Docker
Software and Application Development Tools:	**HPE Cray Programming Environment** • **Development** – Compiling environment – Communication Libraries: HPE Cray MPI, SHMEM – Scientific Libraries: LAPACK, ScaLAPACK, BLAS, libsci, IRT, FFTW 3.0 – I/O Libraries: NETCDF, HDF5 – 3rd party programming environments: o AMD ROCm and AOCC o NVIDIA HPC SDK o GNU Compilers • **Performance analysis tools** – Tools for performance analysis and optimization – versions for both experienced and novice users – Code parallelization assistant for application optimization via code restructuring – Visualization tool for quick assessment of severity of issues – Debuggers: GDB for HPC, Valgrind for HPC, tools for stack trace analysis & abnormal termination processing – 3rd party debugger support: Arm® Forge, TotalView™ by Perforce
DL/AI Tools:	• **Deep learning plugin**

Service and Support

Product warranty

HPE offers a 13-month warranty on all HPE Cray-branded hardware components that begins at the time of shipment and provides replacement or repair of failed hardware at HPE's discretion. This warranty provides only the most basic customer hardware support and is designed for highly skilled customers that intend to maintain their own systems. This HPE limited warranty does not provide any support or warranty obligation for software, even if sold, delivered, or installed by HPE.

Installation

The HPE Cray EX system requires the following installation services:

- Pre-installation activities and solution implementation:
  HPE and the Customer determine all installation activities that must be completed prior to System installation. The Customer agrees to complete all of the pre-installation activities required. This includes HPE site engineering work onsite as may be described in the HPE Cray Site Preparation Guide.
  Solution Implementation: upon completion of the pre-installation activities, HPE will provide the software components as set forth in the applicable system purchase agreement or bill of materials. Any additional software installation or configuration will need to be documented separately and will incur additional charges. This configuration service does not include any customer specific configuration, customization or testing unless otherwise specified.
- System testing and performance validation:

The HPE installation personnel will conduct tests to verify the health and performance of the System. The tests are not intended to demonstrate application performance; the tests verify that the system infrastructure is working properly and delivering the intended performance level. HPE manufacturing tests and diagnostics will be used by installation personnel while onsite to validate that all hardware is functional, meeting the same performance and functional specifications as tested at the factory. Any additional testing that is required should be specified in a separate mutually agreed writing.

Hardware Maintenance Service Features

The HPE Cray EX system benefits of HPE's highest level of support for high-performance compute 'HPE Complete Care - Cray' that may include HPE presence onsite.

- This service level offers customers access to the HPE Cray customer portal. Case logging is available 24x7 by telephone or via this customer portal.
- There is a choice of two maintenance coverage windows: 9x5 or 24x7. Onsite response time options are Next Business Day, 4 hours, 2 hours, or 1 hour. When an issue is reported, an HPE technical representative will arrive onsite within the response time window to identify and begin resolving the issue.
- HPE provides critical spare parts to reduce any downtime associated with failures or maintenance. Critical spare parts may be located either onsite or at professionally managed regional spare part depots that provide rapid transportation of spare parts to customer sites. Customers may elect to supplement the HPE-owned spare parts inventory by purchasing additional spare parts.
- HPE Cray EX customers have access to the support snapshot analyzer that collects, analyzes, and reports support information for HPE air cooled and liquid cooled HPC, and Cray ClusterStor systems.
- HPE Remote Support providing remote access and support. Capabilities range from a customer having full control of a remote screen-sharing session to the HPE Pointnext Services team having the ability to log in securely as needed to resolve issues and perform administrative functions.

Technical Specifications

Software Support Features

Support for HPE developed software includes the following features:

- Access to self-help resources on customer portal
  - Ability to open and submit a support case
  - Access to HPE knowledge articles
  - Ability to download:
    - Software releases and updates, including BIOS and FW
    - Software Patches
- Notification of key operational items through the field notice (FN) process
- Assistance from HPE Pointnext Services to resolve issues within the service level coverage window for the hardware contract; assistance includes:
  - Triage to investigate/analyze issues
  - Confirmation whether the issue is hardware or software

Confirmation if the issue is related to an HPE-supported product or a third-party-supported product. If the issue is with a HPE-supported product, HPE Pointnext Services may provide configuration recommendations, possible work arounds, and directions to install a later version or patch, and/or submit a bug to get the issue fixed. For HPE products, HPE reserves the right to determine whether and how an issue will be resolved.

Customized Software

Support is provided for products sold by HPE and with a valid HPE Pointnext Services support agreement. Support for third-party products without a related HPE Pointnext Services support agreement requires the user to contact the third-party vendor for assistance. If customers modify HPE-delivered software without authorization from HPE, any issues resulting from the unapproved modifications fall outside of the standard support service agreement and HPE is not responsible for any resulting defects, damage, failure, performance degradation, or issues of any kind, or correction or remedy of same. HPE may require the user to remove custom modifications to confirm that a modification is not the source of the issue. Customers may request that HPE Pointnext Services assist in making modifications to a product. HPE Pointnext Services will do its best to implement the request via a billable statement of work (SOW).

API and CLI Support

Support is available for HPE published APIs. Unpublished APIs are not eligible for support. Documentation outlining published API best practices and limitations is available at support.hpe.com , accessible either directly or through the HPE Cray customer portal. HPE will assist in determining if the API is working correctly, if the documentation is incorrect, or if the issue is an enhancement request.

HPE Application Programming Interface (API) and Command Line Interface (CLI) features allow the flexibility to configure and customize your system to optimize operations in your environment. These tools have the ability to significantly alter your system operations. If not properly tested and implemented in a controlled manner, they can introduce significant problems in your environment. When using these features or otherwise modifying or altering APIs, customers take on the responsibility to resolve or mitigate any issues they have introduced into the system.
HPE Pointnext Services is not available to provide support to resolve issues that arise from the use of CLIs or APIs in a form not identical to those published by HPE.

Customer Training

Training courses are taught by HPE system experts and combine lectures with hands-on labs to enhance understanding and retention. The courses cover all aspects of using and maintaining an HPE system, from system administration to application development, porting, and optimization. A full listing of the standard HPE Cray training courses, along with their descriptions, can be found at **https://education.hpe.com/ww/en/training/portfolio/servers.html.**

Subject to separate ordering arrangement, classes are scheduled on regular cycles at the HPE training facilities and can be scheduled for onsite delivery. HPE also offers customized training courses and can provide quotes for these courses based on the customer's needs.

Technical Specifications

HPE Cray EX4000 Supercomputer (cabinet)	
Dimensions with overhead cable trays	98 x 46.5 x 68.5 in (H x W x D) 2489 x 1181 x 1740 mm (H x W x D)
Weight: (Maximum)	Up to 8000 lbs. (3629kg)
Floor Loading (Flat Base)	362 lbs./sq ft (1767 kg/qs m) (Operational)
Compute blade chassis	8 compute blade chassis with integrated compute trays, switches, and power Up to 7 + 1 redundant 15kW power supplies per 2 compute blade chassis
Cooling	Closed-loop airflow with direct liquid cooling for high wattage components and room neutral up to 32°C data center supply water.
Power Requirements (Max)	Up to 400KVA with 480V Up to 350KVA with 400V

HPE Cray EX2500 Supercomputer (rack)	
Dimensions with overhead cable trays	90.65 x 35.43 x 67.86 in (H x W x D) 2302 x 900 x 1719 mm (H x W x D)
Dimensions without cable trays	78.78 x 35.43 x 67.86 (H x W x D) 2000 x 900 x 1719 (H x W x D)
Weight: (Maximum)	Up to 3225 lbs. (1462 kg)
Floor Loading (Flat Base)	806 lbs per caster (365 kg per caster)
Compute blade chassis	8 compute blade chassis with integrated compute trays, switches, and power Up to 3 + 1 redundant 15kW power supplies per compute blade chassis
Cooling	Closed-loop airflow with direct liquid cooling for high wattage components and room neutral up to 32°C data center supply water.
Power Requirements (Max)	Up to 141 kVA with 480V Up to 128kVA with 400V

Compute Blade Options – (HPE Cray EX425)	
Form factor	single-slot blade for the HPE Cray EX425 compute chassis assembly
Processors	AMD 2nd Gen AMD EPYC™ 7002 series and 3rd Gen AMD EPYC™ 7003 series
Compute blade	Four 2-socket CPU nodes
Memory/node	Up to 1024 GB per node, 16 DIMM slots (8 per CPU socket) per node
Memory technology	16, 32, and 64 GB DDR4 3200 MT/s ECC Registered DIMMs
Local storage	none
Fabric options	HPE Slingshot (1 or 2 injection ports per node)

Compute Blade Options – (HPE Cray EX4252)	
Form factor	single-slot blade for the HPE Cray EX4252 compute chassis assembly
Processors	AMD 4th Gen EPYC™ 9004 series processor stack
Compute blade	Four 2-socket CPU nodes
Memory/blade	Up to 1536 GB per node, 24 DIMM slots (12 per CPU socket) per node
Memory technology	16, 32, and 64 GB DDR5 4800 MT/s ECC Registered DIMMs
Local storage	0 or 1 local NVMe M.2 SSD per node (up to 4 per blade)
Fabric options	HPE Slingshot (1 or 2 injection ports per node)

Compute Blade Options – (HPE Cray EX420)	
Form factor	single-slot blade for the HPE Cray EX420 compute chassis assembly
Processors	Intel 4th Gen Xeon scalable processor stack
Compute blade	Four 2-socket CPU nodes
Memory/blade	Up to 1024 GB per node, 16 DIMM slots (8 per CPU socket) per node
Memory technology	16, 32, and 64 GB DDR5 4800 MT/s ECC Registered DIMMs
Local storage	0 or 1 local NVMe M.2 SSD per node (up to 4 per blade)
Fabric options	HPE Slingshot (1 or 2 injection ports per node)

Technical Specifications

Compute Blade Options – (HPE Cray EX235n)	
Form factor	single-slot blade for the HPE Cray EX235 compute chassis assembly
Processors	AMD 3rd Gen AMD EPYC™ 7003 series
Compute blade	Two 4-socket A100 GPUs x 1-socket CPU nodes
Memory/blade	Up to 512 GB per node, 8 DIMM slots per node
Memory technology	16, 32, and 64 GB DDR4 3200 MT/s ECC Registered DIMMs
Local storage	none
Fabric options	HPE Slingshot (2 or 4 injection ports per node)

Compute Blade Options – (HPE Cray EX235a)	
Form factor	single-slot blade for the HPE Cray EX235 compute chassis assembly
Processors	3rd Gen EPYC™ processor
Compute blade	Two 4x AMD Instinct™ M250X Accelerators x 1-socket CPU nodes
Memory/blade	Up to 512 GB per node, 8 DIMM slots per node
Memory technology	16, 32, and 64 GB DDR4 3200 MT/s ECC Registered DIMMs
Local storage	Two M.2 NVME SSD per node
Fabric options	HPE Slingshot (4 injection ports per node)

Integrated HPE Slingshot Switch Blade	
Ethernet Ports	64 Ethernet ports (16 Host; 48 QSFP-DD ports)
Port capability	100/200 Gb/s per port
Switch fabric capability	12.8Tb/s
Messages/s capability	1.2B/s

HPE Slingshot 200Gb 2-port Mezzanine Adapter (R4K44A)	
Form factor	EX Blade Mezzanine form factor
Dimensions	3.311 x 7.362 (H x W) (84.09mm x 186.99)
Weight:	0.396lbs (180g)
Physical connectivity	Two L0 connectors to blade switch ports
Host connectivity	Two PCIe Gen 4 x16 (Supports ESM speeds depending on CPU or GPU availability)
Port capability – 200 Gbps	IEEE 802.3cd/bs (200 Gbps) Ethernet over 4 x 50 Gbps (PAM-4) lanes Proprietary cabling
Throughput per port	Line rate
Port status	Depending on blade type
Power and cooling	48W max 30W typical Liquid-cooled

Summary of Changes

Date	Version History	Action	Description of Change
06-Mar-2023	Version 12	Changed	Standard Features and Technical Specifications sections were updated
19-Sep-2022	Version 11	Changed	Service and Support and Technical Specifications sections were updated.
06-Sep-2022	Version 10	Changed	Added Cray EX2500 information
16-May-2022	Version 9	Changed	Service and Support and Technical Specifications sections were updated.
04-Apr-2022	Version 8	Changed	Standard Features and Technical Specifications sections were updated
04-Oct-2021	Version 7	Changed	Standard Features section was updated
07-Sep-2021	Version 6	Changed	Standard Features section was updated
06-Jul-2021	Version 5	Changed	Updated Software Development Tools. Standard Features and Technical Specifications sections were updated.
17-May-2021	Version 4	Changed	Standard Features section was updated
06-Apr-2021	Version 3	Changed	Overview, Standard Features and Technical Specifications sections were updated.
05-Oct-2020	Version 2	Changed	Service and Support section was updated.
03-Aug-2020	Version 1	New	New QuickSpecs

Copyright

Make the right purchase decision.
Contact our presales specialists.

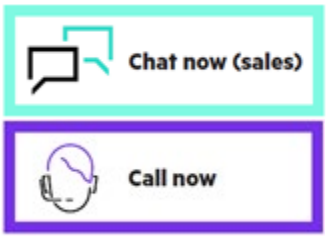

Get updates

For hard drives, 1GB = 1 billion bytes. Actual formatted capacity is less.

a00094635enw - 16619 - WorldWide - V12 - 06-March-2023