Overview

HPE Cray Supercomputing Storage Systems E2000

The HPE Cray supercomputing storage systems E2000 is engineered to efficiently meet the demanding input/output requirements of supercomputers and HPC clusters. The E2000 parallel storage system achieves HPC-level storage performance with significantly fewer storage drives than alternative storage offerings. To the HPC user, this means more of a fixed HPC system budget can be spent on CPU/GPU computes nodes, accelerating time-to-insight. The HPE Cray E2000 storage system embeds the Lustre® open-source parallel file system backed by enterprise-grade customer support. The combination of engineered hardware and system design along with a test hardened open source file system and a tuned storage software stack allows the HPE Cray E2000 to scale out to accommodate the largest compute systems on the planet, without software licensing by storage drives, by capacity levels, or by storage system. HPE Cray storage customers reap the benefits of open-source community collective software development with world class customer support from Hewlett Packard Enterprise.
The E2000 storage system is an I/O and storage subsystem consisting of a single, global POSIX name space file system configured to provide access between compute clients and E2000 storage nodes using the Lustre® parallel file system. E2000 hardware is designed on NVMe generation 5 building blocks utilizing the latest storage technologies to provide top performance with enhanced flexibility and resiliency.

E2000 is offered in four configurations:

- All Flash
- All Disk
- Hybrid
- Tiered with both Flash and Disk

The E2000 all-flash configuration is suitable for the demands of small block, random read I/O, and the high sequential I/O performance required in many AI applications including Machine Learning and Deep Learning. An E2000 all-flash base rack can provide up to 2,000 GB/s read and up to 1,600 GB/s write using the Multi-Rail feature of LNet . An E2000 expansion rack can provide up to an additional 2,200 GB/s read and up to 1,800 GB/s write with the Multi-Rail feature of LNet.

The E2000 all-HDD configuration is optimized for large block sequential I/O and the demands of modeling and simulation workloads. In this configuration, the base rack provides up to 104 GB/s read and write up to approximately 11.3 PB of usable capacity, depending on the HDD capacity used. The expansion rack provides up to 140 GB/s read and write up to 12.6PB of usable capacity, depending on the HDD capacity used.

The E2000 hybrid configuration combines the best of flash and HDD performance into one configuration. The hybrid base rack provides up to 420 GB/s read and up to 340 GB/s write from flash pools with approximately 11.3 PB of usable HDD capacity, depending on the HDD capacity used. The hybrid expansion rack provides up to 570 GB/s read and up to 490 GB/s writes from flash pools with approximately 15 PB of usable HDD capacity, depending on the HDD capacity used.

Customers can utilize hybrid building blocks to combine flash and disk into one filesystem or mix and match the flash and disk building blocks in a tiered configuration to tailor to their workloads.

What's New

- Data-encryption-at-rest (DEAR) and internal key management.
- Tiering for E2000
- E2000 Monitoring
- Managed Gateway nodes
- Updated Drive capacities for Flash and Disk
- Attaches to any supercomputer or HPC cluster that supports either HPE Slingshot, NDR InfiniBand or 400 Gigabit Ethernet.
- In-house, enterprise-grade support for performance-accelerating Lustre® features like Data on Metadata Targets (DoM), Progressive File Layout (PFL), Self-Extending Layouts (SEL), Pool Quotas, Pool Spillover, Multi-Rail LNet, Distributed Name Space (DNE), and many more advanced Lustre® features.

Overview

HPE Cray Supercomputing Storage Systems E2000 – Closed

Example E2000 – Open (doors removed)

Standard Features

E2000 Building Blocks

E2000 is configured from storage building blocks and a surrounding network complete with rack and power infrastructure. These building blocks are:

E2000 System Management Unit (SMU)

The E2000 system management unit contains a pair of embedded storage management nodes with SSDs to hold configuration and logging data. The SMU delivers Lustre® and E2000 system management services while hosting the Tiering and Monitoring service, and can be configured as a data mover. E2000 system management is required for managing hardware configuration, software images, boot up of the underlying hardware system (servers, devices, software stack), monitoring system metrics and system health, reporting on system status, and advanced Lustre® performance monitoring. The SMU runs a web and CLI server that allows the administrator to monitor, configure, and administer the file system, the E2000 management framework, and E2000 hardware. The SMU has a full Redfish/Swordfish APIs to allow administrators to export monitoring and health data to be consumed by existing data center monitoring tools.

E2000 MetaData Unit (MDU)

The E2000 metadata units store and manage global Lustre® metadata. Each MDU contains a pair of Lustre® metadata servers (MDS) and flash-based metadata targets (MDTs). An MDU is populated with 32 SSDs split into two partitions. Each partition contains approximately half the total MDU capacity. Each partition is configured and assigned to one of the two embedded MDS controllers as an MDT using LDISKFS or OpenZFS. MDUs can be configured with multiple MDTs, with each controller configured with two high speed networks adapters.

E2000 configures MDUs to support Lustre® DoM, data on metadata, which allows small files or the initial portions of larger files to be stored with their metadata. The Lustre® DoM feature can improve small file I/O performance, improve stat performance to small files, and reduce interference between small file I/O and streaming I/O. The metadata tier stores metadata and optional DoM in RAID volumes.

E2000 can scale additional MDUs using Lustre® Distributed Name Space (DNE) to match specified performance targets or metadata inode requirements.

E2000 Scalable Storage Unit – All Flash Array (SSU-F)

The E2000 scalable storage units for an all flash array (SSU-F) provide flash-based Lustre® I/O data services and network request handling for the file system. The SSU-F contains a pair of Lustre® object storage servers (OSS), each configured with two Lustre® object storage target(s) (OSTs) to store and retrieve the portions of the file system data that are committed to it. Within the SSU-F, the four OSTs are distributed evenly between the two OSSs and both OSSs are active concurrently, and are said to be in an active-active configuration. With OSSs operating in active-active mode, each operates on its own exclusive subset of the available object storage targets, but either is able to take over the partner node's OSTs should the partner node fail.

A E2000 SSU-F is populated with 32 SSDs. For a throughput optimized configuration, the drives are partitioned and configured with HPE's parity declustered RAID (GridRAID) using LDISKFS or OpenZFS and its parity declustering RAID support. For an IOPs optimized SSU-F configuration, a different RAID scheme with LDISKFS is used to improve small random I/O workloads. Each controller is equipped with two network adapters configured with Multi-Rail LNet to exploit maximum throughput performance per SSU-F.

E2000 can be scaled to many SSU-Fs and/or combined with SSU-Ds or SSU-Ms to achieve specified performance and capacity requirements. E2000 SSU-D and SSU-M units are described below.

Standard Features

E2000 Scalable Storage Unit – Disk (SSU-D)

The E2000 scalable storage units for disk (SSU-D) provide HDD-based Lustre® I/O data services and network request handling for the file system with similar OSS and OST features mentioned above in the description of E2000 's SSU-F building block.

E2000 has multiple SSU-D configurations utilizing either one, two or four E2000 ultra-dense HDD enclosures. Each ultra-dense disk enclosure is configured with 106 SAS HDDs and contains two Lustre® OSTs, each configured with HPE's parity declustered RAID (GridRAID) using LDISKFS or OpenZFS and its parity declustering RAID support. Each SSU-D controller is configured with one high-speed network adapter and redundant SAS adapters to connect to one or more E2000 ultra-dense HDD enclosures.

E2000 can be scaled to many SSU-Ds and/or combined with SSU-Fs to achieve specified performance and/or capacity requirements.

The E2000's configurations with SSU-D are factory integrated with a purpose built rack to accommodate weight and serviceability. For more information about this rack contact your local HPE reprecentative.

E2000 Scalable Storage Unit – Disk (SSU-M)

The E2000 has a second high density HDD enclosure option which is less dense than that managed by the SSU-D, but this enclosure fits within a more standard rack. We refer to this second scalable storage unit for disk as SSU-M. The SSU-M provides HDD-based Lustre® I/O data services and network request handling for the file system with similar OSS and OST features mentioned above for E2000 's SSU-D and SSU-F building blocks.

E2000 has multiple SSU-M configurations with either one, two or four E2000 high-density HDD enclosures. Each high-density disk enclosure is configured with 84 SAS HDDs and contains two Lustre® OSTs, each configured with HPE's parity declustered RAID (GridRAID) using LDISKFS or OpenZFS and its parity declustering RAID support. Each SSU-M controller is configured one high-speed network adapter and redundant SAS adapters to connect to one or more E2000 HDD enclosures.

E2000 can be scaled to many SSU-Ms and/or combined with SSU-Fs to achieve specified performance requirements.

E2000's SSU-M configurations allow for factory integration with HPE G2 Enterprise and Advanced Series 42U shock racks and select HPE G2/3 Metered and Switched Power Distribution Units.

2000 Scalable Storage Unit – Hybrid (SSU-H)

The E2000 scalable storage unit – hybrid (SSU-H) provides both flash-based and HDD-based Lustre® I/O data services and network request handling for the file system with similar OSS and OST features mentioned above for E2000 's SSU-F, SSU-D and SSU-M building blocks.

E2000 has two SSU-H configurations supporting either two E2000 ultra-dense or high-density HDD enclosures. Each SSU-H building block is configured with HPE's parity declustered RAID (GridRAID) using LDISKFS or OpenZFS and its parity declustering RAID support. Each SSU-H controller is configured two high-speed network adapters and redundant SAS adapters to connect to two E2000 HDD enclosures. Each controller can be configured with Mult-Rail LNet to exploit maximum flash throughput performance per SSU-H.

E2000 can be scaled to many SSU-Hs to achieve specified performance requirements.

Standard Features

SSU-D1

SSU-D2

SSU-D4

E2000 Scalable Storage Unit – Disk (SSU-D)

SSU-M1

SSU-M2

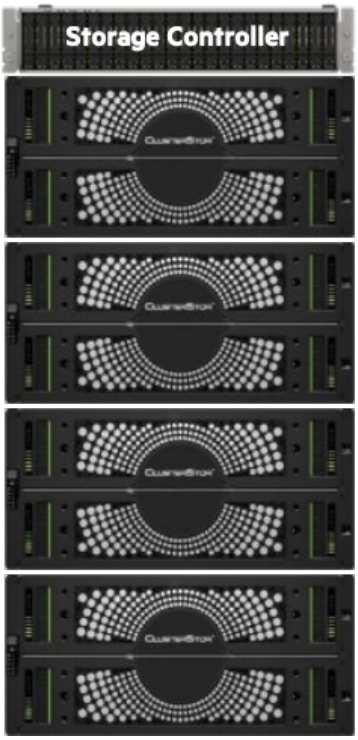

SSU-M4

E2000 Scalable Storage Unit – Disk (SSU-M)

Standard Features

SSU-HD1 **SSU-HD2** **SSU-HM1** **SSU-HM2**

E2000 Scalable Storage Unit – Disk (SSU-H)

Management Switches

E2000 utilizes HPE Aruba Networking switches for a private management network between all building blocks in the system. The private management network is deployed as a highly-available configuration. Depending on the configuration of the system, HPE Aruba Networking CX 8360 switches, HPE Aruba Networking CX 6300M switches or a combination of HPE Aruba Networking CX 8325 and CX 6300M switches will be factory integrated with each E2000 . Only back-to-front (also called power-to-port or reverse air flow) HPE Aruba Networking switches are supported for E2000's private management network.

Racks and PDUs

E2000 storage systems are compatible with HPE G2 Advanced and Enterprise 42U and 48U Shock Racks and Cray custom racks. HPE requires racks with a minimum depth of 1200mm to best accommodate the length of the E2000 building blocks mentioned above. For E2000 configurations using SSU-M and SSU-F, HPE G2 600x1200mm shock racks provide the necessary room for cabling and serviceability. For E2000 configurations using SSU-D, HPE G2 Enterprise 800x1200mm shock racks or Cray custom racks are required for cabling and serviceability.

E2000 storage systems are compatible with select rack-mounted HPE G2/3 PDUs and Cray PDUs. A minimum of two PDUs are required with E2000 for power redundancy. The desired HPE G2 or Cray custom rack and E2000 configuration could limit what PDU options are available.

E2000 Software, Services, and Tools

E2000 is configured with powerful software tools and services to create a complete, production-ready storage system that is designed to meet the unique requirements of HPC and AI applications. This software, services, and tools are:

Lustre® for E2000

E2000 comes pre-configured with the industry's most powerful and scalable parallel file system to assure that your compute cluster avoids waiting for I/O and is computing with maximum efficiency. Lustre® for E2000 is a carefully tested and curated version of the open source community Lustre® based on version 2.15 release. Each Lustre® feature is exhaustively tested at scale by HPE and any defects corrected before it is released for production use. Lustre® for E2000 is offered with either OpenZFS or with LDISKFS as the local object store using HPE's parity declustered RAID (GridRAID) to deliver maximum application performance and up to 4x faster recovery time from failed drives. Lustre® for E2000 is fully supported by HPE Services and E2000 Engineering resources guaranteeing fast and thorough responses to any issues found by our customers. For deployments in GPU based compute systems, HPE Cray Lustre® clients has full support of Nvidia's Magnum I/O (a.k.a., GPU Direct Storage) features.

Standard Features

Tiering for E2000

Tiering in E2000 is a set of software features that make a hybrid storage system comprised of flash and disk easier to use and manage. It is designed to monitor and maintain file layouts and free space on flash tiers and disk tiers. Tiering has multiple capabilities that can be customized by the administrator via data management policies, including fine-grained indexing controls, file migration between OSTs or pools, restriping, purging, or generating filesystem reports. These actions are invoked automatically based on timers or system conditions (e.g. capacity threshold).

Tiering uses a scale-out architecture to efficiently move large volumes of data. Tiering components utilize the SMU to perform all functions, and additional data mover nodes can be configured to increase data throughput.

Scalable Search

The Scalable Search service maintains a efficiently-searchable indexed database files of filesystem metadata, used by policies to quickly search the filesystem namespace with multiple selection criteria to find matching candidate files amongst billions of files. The databases are maintained automatically as a background process.

Transparent Tiering

Transparent migration optimizes file layouts automatically based on policies created and set by the administrator. The policy engine automates file migration between the flash and disk tiers based on triggers such as the available capacity in the flash tier or a fixed schedule, and can migrate files based on path, name, file size, modification time, or any other file metadata. This is particularly useful to help maintain a flash tier at a maximum percentage full by moving older files to disk, for example. All files remain completely available for read or write during this process.

Parallel Data Movers

Additional data mover nodes can be configured to scale out the Tiering services from the SMU, providing additional bandwidth to migrate data and index the filesystem. The additional nodes are pre-configured HPE servers that can be added and monitored by E2000 management nodes.

Purging Data

Administrators can create policies to delete files from the filesystem based on various conditions, such as age of the file, location of the file, last access of the file, size of the file, to name a few examples.

Reporting

Tiering services can automatically generate reports of any filesystem metadata query in JSON or CSV form, timestamped and stored in the designated directory. Reports can be used to track usage information, find users or groups consuming significant capacity, or provide information to fine-tune other policies.

Data Encryption at Rest (DEAR) with Key Manager

E2000 supports encrypting data at rest with key management using Secure Encrypted Drives (SED) for NVMe and SED or FIPS compliant HDDs. E2000 supports a built-in key manager to handle the encryption keys. When enabled, E2000 enrolls all file system storage devices to participate in encrypting data. The key manager is used to authenticate all devices to allow access to encrypted data. After drive power loss, access to data is prohibited until the system retrieves unique keys from the key manager and uses them to unlock the devices.

Managed Exporter Nodes

E2000 supports the following gateway nodes to extend the access of the Lustre® parallel filesystem to additional clients using NFS and or S3. E2000 will manage up to 4 nodes for NFS to allow clients to mount the Lustre® filesystem over NFS or 4 nodes for S3 to ingest S3 objects as POSIX accessible objects and in turn export Lustre® data to S3.

Standard Features

E2000 Monitoring

E2000 monitoring provides always-on metrics to give customers the insight needed to maximize resources, optimize system throughput, and provide users with the best results and experience without impacting performance. All the information users need is available in a single view. Using streaming telemetry with customizable visualization of real-time and historical data, E2000 monitoring software tracks job performance on Lustre® file systems and enables the correlation of system events that cause job performance issues. E2000 monitoring provides administrators with visibility to understand application performance variability, to isolate problems faster, to enable trend analysis, and to provide alerts for hands-off monitoring.

The landing page for E2000 monitoring offers a quick overview of the utilization and activity on the E2000 storage system giving administrators the ability to quickly identify potential issues and application performance problems.

Sample of the E2000 Monitoring Landing Page

E2000 monitoring provides visualization of performance data for overall storage systems and individual jobs, including throughput and metadata activity. Performance data is stored persistently and can be monitored in near real-time or historically, to triage performance problems, study trend analysis, and plan for capacity increases. Monitoring data sources include E2000 system logs and InfiniBand fabrics, enabling administrators to correlate system events with storage performance. E2000 monitoring supports customized alerts and graphs to tailor the application to each customer's needs. Other notable features:

- Job Runtime Variability – Real-time and historical views of data to help you understand what is impacting user jobs
- Event Correlation – A unified view of the system, providing administrators with the ability to correlate systemic events that impact performance
- Trend Analysis – Data-driven analysis and visualization from historical data can help identify trends that can then be used to shape changes to the system
- Alerting – Threshold engine enables customized alerts based on any metric

Job Monitoring Capabilities

Job performance metrics are recorded and monitored in real-time to detect inefficient utilization of resources such as small I/O requests or metadata request inundation that can impact the responsiveness of the entire system. Jobs can be sorted and filtered as needed to pinpoint specific jobs by attribute, or to identify jobs running within certain timeframes.

Standard Features

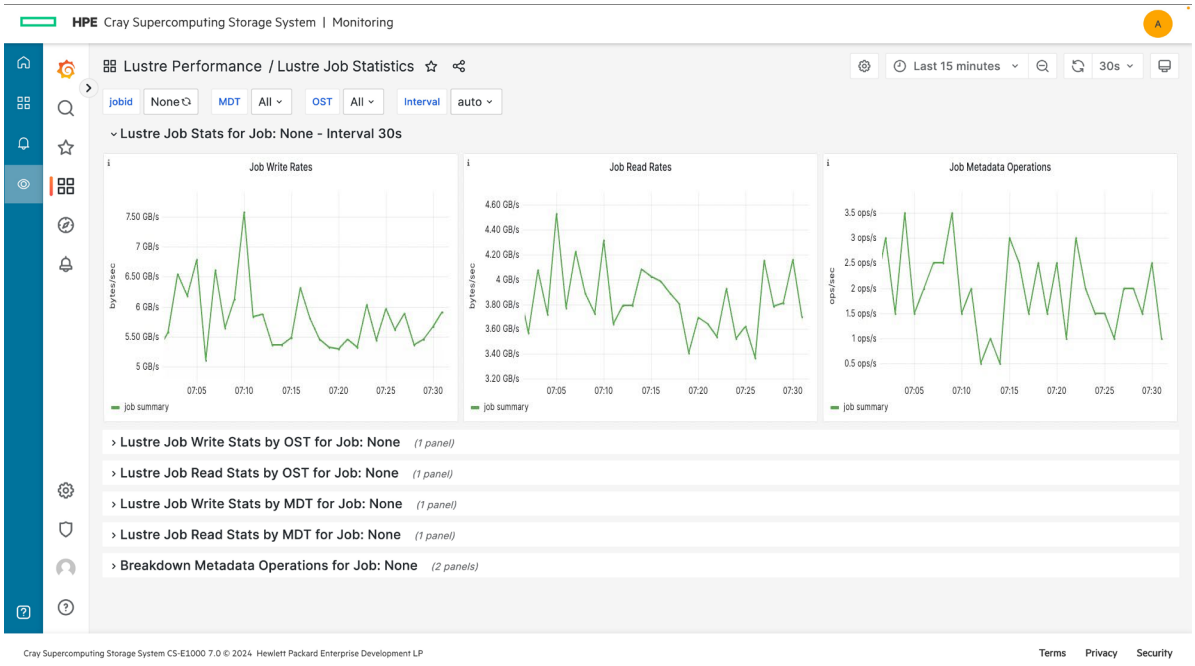

Sample of the Jobs Statitics

From the Jobs page, administrators can drill down to job details and compare its activity to that occurring on the rest of the E2000 storage system to provide system level context to the target job.

Job details also include per-device reporting for Lustre® object storage targets (OSTs) and metadata targets (MDTs) including metadata operations.

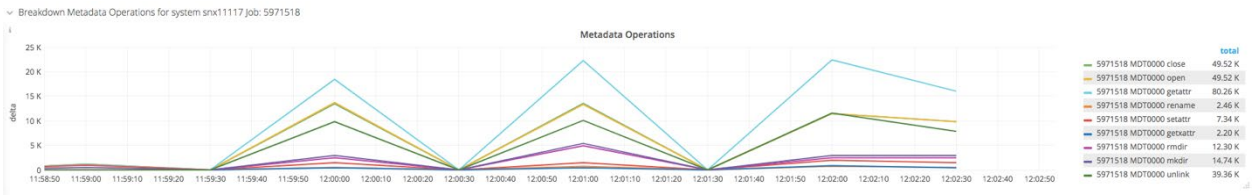

Sample Breakdown of the Metadata Operations Per Device

Dashboards

E2000 monitoring includes visualization dashboards based on Grafana covering the E2000 system overview and Lustre® device metrics for an in-depth view of the storage environment. Available dashboards include:

- E2000 Storage Overview – Displays an overview of performance, jobs, and capacity for one or more E2000 storage systems.
- E2000 Server Metrics – Displays CPU and memory utilization for all E2000 servers.
- Lustre® File System Capacity, Metadata, and Performance – Separate displays covering per device capacity, performance, and metadata operations.
- Customized Dashboards - supports the creation of custom dashboards from a wealth of documented metrics.

Standard Features

For more information about Cray E2000 , HPE and HPE Aruba Networking products mentioned in this quickspec, please visit:

- Cray E2000 (Technical White Paper):
  https://www.hpe.com/psnow/doc/a50001954enw
- Cray E2000 (Business White Paper):
  https://www.hpe.com/psnow/doc/a50002042enw
- HPE Slingshot Interconnect
  https://www.hpe.com/psnow/doc/a50002546enw
- HPE NDR InfiniBand Switches:
  https://www.hpe.com/psnow/doc/PSN1014668195SEEN
- HPE Aruba Networking CX 6300M Switch:
  https://www.hpe.com/psnow/doc/a00073540enw
- HPE Aruba Networking CX 8325 Switch:
  https://www.hpe.com/psnow/doc/a00056519enw
- HPE Aruba Networking CX 8360 Switch:
  https://www.hpe.com/psnow/doc/a00021857enw.pdf
- HPE G2 Enterprise Series Racks:
  https://www.hpe.com/psnow/doc/a00002907enw
- HPE G2 Advanced Series Racks:
  https://www.hpe.com/psnow/doc/c05324689
- HPE G2 Metered and Switched Power Distribution Units:
  https://www.hpe.com/psnow/doc/a00002910enw

Service and Support

HPE Services

No matter where you are in your digital transformation journey, you can count on HPE Services to deliver the expertise you need when, where and how you need it. From planning to deployment, ongoing operations and beyond, our experts can help you realize your digital ambitions.

https://www.hpe.com/services

Consulting Services

No matter where you are in your journey to hybrid cloud, experts can help you map out your next steps. From determining what workloads should live where, to handling governance and compliance, to managing costs, our experts can help you optimize your operations.

https://www.hpe.com/services/consulting

HPE Managed Services

HPE runs your IT operations, providing services that monitor, operate, and optimize your infrastructure and applications, delivered consistently and globally to give you unified control and let you focus on innovation.

HPE Managed Services | HPE

Operational services

Optimize your entire IT environment and drive innovation. Manage day-to-day IT operational tasks while freeing up valuable time and resources. Meet service-level targets and business objectives with features designed to drive better business outcomes.

https://www.hpe.com/services/operational

HPE Complete Care Service

HPE Complete Care Service is a modular, edge-to-cloud IT environment service designed to help optimize your entire IT environment and achieve agreed upon IT outcomes and business goals through a personalized experience. All delivered by an assigned team of HPE Services experts. HPE Complete Care Service provides:

- A complete coverage approach -- edge to cloud
- An assigned HPE team
- Modular and fully personalized engagement
- Enhanced Incident Management experience with priority access
- Digitally enabled and AI driven customer experience

https://www.hpe.com/services/completecare

HPE Tech Care Service

HPE Tech Care Service is the operational support service experience for HPE products. The service goes beyond traditional support by providing access to product specific experts, an AI driven digital experience, and general technical guidance to not only reduce risk but constantly search for ways to do things better. HPE Tech Care Service delivers a customer-centric, AI driven, and digitally enabled customer experience to move your business forward. HPE Tech Care Service is available in three response levels. Basic, which provides 9x5 business hour availability and a 2-hour response time. Essential which provides a 15-minute response time 24x7 for most enterprise level customers, and Critical which includes a 6-hour repair commitment where available and outage management response for severity 1 incidents.

https://www.hpe.com/services/techcare

Service and Support

HPE Lifecycle Services

HPE Lifecycle Services provide a variety of options to help maintain your HPE systems and solutions at all stages of the product lifecycle. A few popular examples include:

- Lifecycle Install and Startup Services: Various levels for physical installation and power on, remote access setup, installation and startup, and enhanced installation services with the operating system.
- HPE Firmware Update Analysis Service: Recommendations for firmware revision levels for selected HPE products, taking into account the relevant revision dependencies within your IT environment.
- HPE Firmware Update Implementation Service: Implementation of firmware updates for selected HPE server, storage, and solution products, taking into account the relevant revision dependencies within your IT environment.
- Implementation assistance services: Highly trained technical service specialists to assist you with a variety of activities, ranging from design, implementation, and platform deployment to consolidation, migration, project management, and onsite technical forums.
- HPE Service Credits: Access to prepaid services for flexibility to choose from a variety of specialized service activities, including assessments, performance maintenance reviews, firmware management, professional services, and operational best practices.

Notes: To review the list of Lifecycle Services available for your product go to:
https://www.hpe.com/services/lifecycle

For a list of the most frequently purchased services using service credits, see the **HPE Service Credits Menu**

Other Related Services from HPE Services:

HPE Education Services

Training and certification designed for IT and business professionals across all industries. Broad catalogue of course offerings to expand skills and proficiencies in topics ranging from cloud and cybersecurity to AI and DevOps. Create learning paths to expand proficiency in a specific subject. Schedule training in a way that works best for your business with flexible continuous learning options.
https://www.hpe.com/services/training

Defective Media Retention

An option available with HPE Complete Care Service and HPE Tech Care Service and applies only to Disk or eligible SSD/Flash Drives replaced by HPE due to malfunction.

Consult your HPE Sales Representative or Authorized Channel Partner of choice for any additional questions and services options.

Parts and Materials

HPE will provide HPE-supported replacement parts and materials necessary to maintain the covered hardware product in operating condition, including parts and materials for available and recommended engineering improvements.

Parts and components that have reached their maximum supported lifetime and/or the maximum usage limitations as set forth in the manufacturer's operating manual, product quick-specs, or the technical product data sheet will not be provided, repaired, or replaced as part of these services.

How to Purchase Services

Services are sold by Hewlett Packard Enterprise and Hewlett Packard Enterprise Authorized Service Partners:

- Services for customers purchasing from HPE or an enterprise reseller are quoted using HPE order configuration tools.
- Customers purchasing from a commercial reseller can find services at **https://ssc.hpe.com/portal/site/ssc/**

Service and Support

AI Powered and Digitally Enabled Support Experience

Achieve faster time to resolution with access to product-specific resources and expertise through a digital and data driven customer experience

Sign into the HPE Support Center experience, featuring streamlined self-serve case creation and management capabilities with inline knowledge recommendations. You will also find personalized task alerts and powerful troubleshooting support through an intelligent virtual agent with seamless transition when needed to a live support agent.
https://support.hpe.com/hpesc/public/home/signin

Consume IT On Your Terms

HPE GreenLake edge-to-cloud platform brings the cloud experience directly to your apps and data wherever they are—the edge, colocations, or your data center. It delivers cloud services for on-premises IT infrastructure specifically tailored to your most demanding workloads. With a pay-per-use, scalable, point-and-click self-service experience that is managed for you, HPE GreenLake edge-to-cloud platform accelerates digital transformation in a distributed, edge-to-cloud world.

- Get faster time to market
- Save on TCO, align costs to business
- Scale quickly, meet unpredictable demand
- Simplify IT operations across your data centers and clouds

To learn more about HPE Services, please contact your Hewlett Packard Enterprise sales representative or Hewlett Packard Enterprise Authorized Channel Partner. Contact information for a representative in your area can be found at "Contact HPE"
https://www.hpe.com/us/en/contact-hpe.html

For more information
http://www.hpe.com/services

Technical Specifications

Cray E2000

- Filesystem
  Lustre® 2.15 with supported enhancements
- Data Path
  E2000 GridRAID with LDISKFS or OpenZFS and its RAID support

Scalable Storage Unit – Flash (SSU-F)

- Storage Controller
  High availability, dual node pair with 32 NVMe Gen 5 SSDs
- SSD Options
  - 1 DWPD: 1.92 / 3.84 / 7.68 / 15.36 TB
  - 3 DPWD: 1.6 / 3.2 / 6.4 TB
  - All SSD options above are self-encryption (SED) capable
- Maximum SSUs per Rack (dependent on type of network supported)
  - 14 per Base Rack
  - 16 per Expansion Rack
- Maximum usable capacity per rack
  Up to 15.5 PB (using 15.36 TB SSDs)

Scalable Storage Unit – Disks (SSU-D)

- Storage Controller
  High availability, dual controller pair with two NVMe Gen 5 SSDs
- Disk Enclosure
  Ultra-dense 4U SAS JBOD with 106 HDDs
- Configurations
  - SSU-D1: Storage controller plus one 4U106 disk enclosure in 6U
  - SSU-D2: Storage controller plus two 4U106 disk enclosures in 10U
  - SSU-D4: Storage controller plus four 4U106 disk enclosures in 18U
- HDD Options
  - 4 / 12 / 16 / 20 / 24 TB SAS SED HDDs
- Maximum SSUs per Rack (dependent on type of SSU-D configuration and network supported)
  - Three per Base Rack
  - Four per Expansion Rack
- Maximum usable capacity per rack
  Up to 15PB (using 24TB HDDs)

Scalable Storage Unit – Disks (SSU-M)

- Storage Controller
  High availability, dual controller pair with two NVMe Gen 5 SSDs
- Disk Enclosure
  High-density 5U SAS JBOD with 84 HDDs
- Configurations
  - SSU-M1: Storage controller plus one 5U84 disk enclosure in 7U
  - SSU-M2: Storage controller plus two 5U84 disk enclosures in 12U
  - SSU-M4: Storage controller plus four 5U84 disk enclosures in 22U
- HDD Options
  - 4 / 12 / 16 / 20 / 24 TB SAS SED HDDs
- Maximum SSUs per Rack (dependent on type of SSU-M configuration and network supported)
  - Three per Base Rack
  - Four per Expansion Rack
- Maximum usable capacity per rack
  Up to 11.9 PB (using 24TB HDDs)

Technical Specifications

Scalable Storage Unit – Hybrid (SSU-M)

- Storage Controller
  High availability, dual controller pair with 32 NVMe Gen 5 SSDs
- SSD Options
  Same SSD options for SSU-F mentioned above
- Disk Enclosure
  Ultra-dense 4U SAS JBOD with 106 HDDs
  High-density 5U SAS JBOD with 84 HDDs
- Configurations
  - SSU-HD1: Storage controller plus one 4U106 disk enclosure in 6U
  - SSU-HD2: Storage controller plus two 4U106 disk enclosures in 10U
  - SSU-HM1: Storage controller plus one 5U84 disk enclosure in 7/U
  - SSU-HM2: Storage controller plus two 5U84 disk enclosures in 12U
- HDD Options
  - For SSU-HD, same HDD options for SSU-D mentioned above
  - SSU-HM, same HDD options for SSU-M mentioned above
- Maximum SSUs per Rack (dependent on type of SSU-H configuration and network supported)
  - Three per Base Rack
  - Four per Expansion Rack

MetaData Unit (MDU)

- Base Configuration
  High-availability, dual controller pair with 32 NVMe Gen 5 SSDs RAID protected
- Number of files
  Up to 17.2 billion per MDU in standard configuration with LDISKFS[1]
- SSD Options
  - 1 DWPD: 1.92 / 3.84 / 7.68 TB
  - 3 DPWD: 1.6 / 3.2 / 6.4 TB
  - All SSD options are self-encryption capable
- Expansion Option
  Up to 10 MDUs per file system configured with Lustre® Distributed Namespace functionality

System Management Unit (SMU)

High-availability, dual controller pair with NVMe Gen 5 SSDs RAID protected for system management, logging, and boot services

Client High-Speed Network Options [2]

HPE Slingshot200/400, NDR InfiniBand, HPE M-series 400 GbE Switches
It is highly recommended that E2000 utilize a high-speed network that is deployed in a highly-available configuration.

Notes:

– [1] 1.92 TB or higher capacity SSD required to achieve eight billion files per MDU

Technical Specifications

Building Blocks	SMU	MDU	SSU-F	SSU-D/M Controller
HPE SKU(s) (includes common options)	S1H80A	S1H81A	S1H83A	S1H82A
Dimensions	88.9mm x 486mm x 803mm (H x W x D)			
Weight	27.7kg | 61 lbs	31.3kg | 69 lbs	31.3kg | 69 lbs	27.7kg | 61 lbs
AC Power Input	200-240VAC			
AC Power Input Frequency	50-60Hz			
Typical Power (kW | kVA)	1.295 | 1.439	1.900 | 2.111	1.900 | 2.111	1.283 | 1.426
Max Power (kW | kVA)	1,824 | 2.027	2.463 | 2.737	2.463 | 2.737	1.807 | 2.008
Power Cables[1]	Two 3m C19/C20			
Operating Temperature	5°C-35°C | 41°F-95°F			
Shipping Temperature	-40°C-60°C | –40°F-+140°F			
Operating Altitude	-200-10,000ft | -61m-3,000m			
Shipping Altitude	-200-40,000ft | -61m-12,192m			
Operating Humidity	20% to 80% non-condensing			
Non-operating Humidity	10% to 90% non-condensing			
Operating Vibration	0.10G at 5 Hz to 500 Hz			
Non-operating Vibration	0.5G at 5 Hz to 500 Hz			
Operating Shock	3G at 11ms, 1/2 sine wave pulse			
Non-operating Shock	10G at 11ms, 1/2 sine wave pulse			
Maximum Exhaust Air Flow (CFM)	162	218	218	160
Thermal (BTU)	6220	8399	8399	6162
Management network cables	Default of four 10ft (3m) CAT6 Ethernet cables. Lengths may vary when built and shipped in a datacenter rack by HPE.			

Notes:

– [1] HPE will determine which power cable is utilized when built and shipped in a datacenter rack by HPE. Both power cable types are provided otherwise.

Technical Specifications

Building Blocks	SSU-H Controller
HPE SKU(s) [1] **(includes common options)**	S1H84H
Dimensions	88.9mm x 486mm x 803mm (H x W x D)
Weight	31.3kg | 69 lbs.
AC Power Input	200-240VAC
AC Power Input Frequency	50-60Hz
Typical Power (kW | kVA)	1900 | 2.111kVA
Max Power (kW | kVA)	2463 | 2.737kVA
Power Cables[2]	Two 1m C19/C20
Operating Temperature	5°C-35°C | 41°F-95°F
Shipping Temperature	-40°C-60°C | –40°F-+140°F
Operating Altitude	-200-10,000ft | -61m-3,000m
Shipping Altitude	-200-40,000ft | -61m-12,192m
Operating Humidity	20% to 80% non-condensing
Non-operating Humidity	10% to 90% non-condensing
Operating Vibration	0.10G at 5 Hz to 500 Hz
Non-operating Vibration	0.5G at 5 Hz to 500 Hz
Operating Shock	3G at 11ms, 1/2 sine wave pulse
Non-operating Shock	10G at 11ms, 1/2 sine wave pulse
Maximum Exhaust Air Flow (CFM)	219
Thermal (BTU)	8399
Management network cables	Default of four 10ft (3m) CAT6 Ethernet cables. Lengths may vary when built and shipped in a datacenter rack by HPE.

Notes:

- [1] Must utilize HPE optical SAS cables to connect to E2000 SSU-H enclosures.
- [2] HPE will determine which power cable is utilized when built and shipped in a datacenter rack by HPE. Both power cable types are provided otherwise.

Technical Specifications

SSU-D Ultra-Dense 4U106 Enclosure[1,2]	4TB SED	12TB SED	16TB SED
HPE SKU(s)	S3J04A	TBD	R6Q54A
Dimensions	176.4mm x 441mm x 1139mm (H x W x D)		
Weight	135.1 kg | 298.5 lbs.	139.6 kg | 308.5 lbs.	137.2 kg | 303.2 lbs.
AC Power Input	200-240VAC		
AC Power Input Frequency	50-60Hz		
Typical Power (kW | kVA)	1317W | 1.463kVA	1398W | 1.553kVA	1227W | 1.363kVA
Max Power (kW | kVA)	1834W | 2.038 kVA	1918W | 2.131kVA	1706W | 1.896kVA
Power Cables	Two 1m C19/C20		
Operating Temperature	5°C to 35°C | 41°F to 95°F - derate 1°C for every 300m above 900m		
Shipping Temperature	–40°C to +70°C | –40°F to +158°F		
Operating Altitude	-200-10,000ft | -61m-3,000m		
Shipping Altitude	-200-40,000ft | -61m-12,192m		
Operating Humidity	10% to 80% noncondensing		
Non-operating Humidity	5% to 100% noncondensing		
Operating Vibration	0.18 Grms, 5-500Hz, 30 min per axis		
Non-operating Vibration	0.54 Grms (in Z) 0.25 Grms (in X & Y), 6-200Hz		
Operating Shock	3 G, 11ms (per axis)		
Non-operating Shock	15 G, 7ms, 10 shock pulse		
Maximum Exhaust Air Flow (CFM)	163	174	152
Thermal (BTU)	6255	6653	5821
Cable Management Arm (CMA)	Included and required for serviceability. CMA will consume up to 1.5m of cable length for all cables connected.		

Notes:

- [1] Hewlett Packard Enterprise has not tested or validated the E2000SSU-D Ultra-Dense 4U106 Enclosure with any third-party racks. Before installing the E2000 SSU-D Ultra-Dense 4U106 Enclosure in a third-party rack, be sure to properly scope the limitations of the rack. Before proceeding with the installation, consider the following:
  - You must fully understand the static and dynamic load carrying capacity of the rack and be sure that it can accommodate the weight of the E2000 SSU-D Ultra-Dense 4U106 enclosure.
  - Be sure sufficient clearance exists for cabling, installation and removal of PDUs, installation and removal the enclosure and enclosure FRUs, actuation of the rack doors and full articulation of the drive drawers.
- [2] Must utilize HPE optical SAS cables to connected to R4M27A, R7K93A or R7K89A.

Technical Specifications

SSU-D Ultra-Dense 4U106 Enclosure[1,2]	20TB SED	24TB SED
HPE SKU(s)	S1H96A	S1H97A
Dimensions	176.4mm x 441mm x 1139mm (H x W x D)	
Weight	137.2 kg | 303.2 lbs.	
AC Power Input	200-240VAC	
AC Power Input Frequency	50-60Hz	
Typical Power (kW | kVA)	1213W | 1348kVA	
Max Power (kW | kVA)	1685W | 1873kVA	
Power Cables	Two 1m C19/C20	
Operating Temperature	5°C to 35°C | 41°F to 95°F - derate 1°C for every 300m above 900m	
Shipping Temperature	–40°C to +70°C | –40°F to +158°F	
Operating Altitude	-200-10,000ft | -61m-3,000m	
Shipping Altitude	-200-40,000ft | -61m-12,192m	
Operating Humidity	10% to 80% noncondensing	
Non-operating Humidity	5% to 100% noncondensing	
Operating Vibration	0.18 Grms, 5-500Hz, 30 min per axis	
Non-operating Vibration	0.54 Gs rms (in Z) 0.25 Grms (in X & Y), 6-200Hz	
Operating Shock	3 G, 11ms (per axis)	
Non-operating Shock	15 G, 7ms, 10 shock pulse	
Maximum Exhaust Air Flow (CFM)	150	
Thermal (BTU)	5749	
Cable Management Arm (CMA)	Included and required for serviceability. CMA will consume up to 1.5m of cable length for all cables connected.	

Notes:

- [1] Hewlett Packard Enterprise has not tested or validated the E2000 SSU-D Ultra-Dense 4U106 Enclosure with any third-party racks. Before installing the E2000 SSU-D Ultra-Dense 4U106 Enclosure in a third-party rack, be sure to properly scope the limitations of the rack. Before proceeding with the installation, consider the following:
  - You must fully understand the static and dynamic load carrying capacity of the rack and be sure that it can accommodate the weight of the E2000 SSU-D Ultra-Dense 4U106 enclosure.
  - Be sure sufficient clearance exists for cabling, installation and removal of PDUs, installation and removal the enclosure and enclosure FRUs, actuation of the rack doors and full articulation of the drive drawers.
- [2] Must utilize HPE optical SAS cables to connected to R4M27A, R7K93A or R7K89A.

Technical Specifications

SSU-M High-Density 5U84 Enclosure[2,3]	4TB SED	12TB SED	16TB SED
HPE Cray Supercomputing Storage Systems SAS 12G 5U 84-disk LFF Drive Enclosure	R6Q55A		
HPE HDD Bundle SKU[1]	R6Q49A	HPE Cray Supercomputing Storage Systems 16TB SAS 12G 5U 84-disk LFF LP SED 42-pack HDD Bundle	R6Q51A
Dimensions	220mm x 483mm x 933mm (H x W x D)		
Weight	129.3 kg | 285.8 lbs.		132.9 kg | 293.8 lbs
AC Power Input	200-240VAC		
AC Power Input Frequency	50-60Hz		
Typical Power (kW | kVA)	1222W | 1.358kVA		1287W / 1.430kVA
Max Power (kW | kVA)	1691W / 1.879kVA		1728W | 1.920kVA
Power Cables	Two 1m C19/C20		
Operating Temperature	5°C to 35°C (41°F to 95°F) derate 1°C for every 300m above 900m		
Shipping Temperature	–40°C to +70°C | –40°F to +158°F		
Operating Altitude	-200-10,000ft | -61m-3,000m		
Shipping Altitude	-200-40,000ft | -61m-12,192m		
Operating Humidity	10% to 80% noncondensing		
Non-operating Humidity	5% to 100% noncondensing		
Operating Vibration	0.21 Grms, 5-500Hz		
Non-operating Vibration	1.04 Grms, 2-200Hz		
Operating Shock	5 G, 10ms, half sine		
Non-operating Shock	20 G, 10ms, half sine		
Maximum Exhaust Air Flow (CFM)	150		154
Thermal (BTU)	5769		6085

Notes:

- [1] Must order quantity two with each R6Q55A
- [2] Hewlett Packard Enterprise has not tested or validated the E2000 SSU-M High-Density 5U84 Enclosure with any third-party racks. Before installing the E2000 SSU-M High-Density 5U84 Enclosure in a third-party rack, be sure to properly scope the limitations of the rack. Before proceeding with the installation, consider the following:
  - You must fully understand the static and dynamic load carrying capacity of the rack and be sure that it can accommodate the weight of the E2000 SSU-M High-Density 5U84 enclosure.
  - Be sure sufficient clearance exists for cabling, installation and removal of PDUs, installation and removal the enclosure and enclosure FRUs, actuation of the rack doors and full articulation of the drive drawers.
- [3] Must utilize HPE optical SAS cables to connected to R4M27A, R7K93A or R7K89A.

Technical Specifications

SSU-M High-Density 5U84 Enclosure[2,3]	20TB SED	24TB SED
HPE Cray Supercomputing Storage Systems SAS 12G 5U 84-disk LFF Drive Enclosure	R6Q55A	
HPE HDD Bundle SKU[1]	HPE Cray Supercomputing Storage Systems 24TB SAS 12G 5U 84-disk LFF LP SED 42-pack HDD Bundle	S1H94A
Dimensions	176.4mm x 441mm x 1139mm (H x W x D)	
Weight	135 kg | 298 lbs.	131 kg | 289.5 lbs.
AC Power Input	200-240VAC	
AC Power Input Frequency	50-60Hz	
Typical Power (kW | kVA)	1338 W | 1487kVA	1151W | 1279kVA
Max Power (kW | kVA)	1857W | 2063kVA	1590W | 1767kVA
Power Cables	Two 1m C19/C20	
Operating Temperature	5°C to 35°C | 41°F to 95°F - derate 1°C for every 300m above 900m, 20°C/hr max rate of change	
Shipping Temperature	–40°C to +70°C | –40°F to +158°F	
Operating Altitude	-200-10,000ft | -61m-3,000m	
Shipping Altitude	-200-40,000ft | -61m-12,192m	
Operating Humidity	10% to 80% noncondensing	
Non-operating Humidity	5% to 100% noncondensing	
Operating Vibration	0.21 Grms, 5-500Hz	
Non-operating Vibration	1.04 Grms, 2-200Hz	
Operating Shock	5 G, 10ms, half sine	
Non-operating Shock	20 G, 10ms, half sine	
Maximum Exhaust Air Flow (CFM)	165	142
Thermal (BTU)	6337	5425

Notes:

- [1] Must order quantity two with each R6Q55A
- [2] Hewlett Packard Enterprise has not tested or validated the E2000 SSU-M High-Density 5U84 Enclosure with any third-party racks. Before installing the E2000 SSU-M High-Density 5U84 Enclosure in a third-party rack, be sure to properly scope the limitations of the rack. Before proceeding with the installation, consider the following:
  - You must fully understand the static and dynamic load carrying capacity of the rack and be sure that it can accommodate the weight of the E2000 SSU-M High-Density 5U84 enclosure.
  - Be sure sufficient clearance exists for cabling, installation and removal of PDUs, installation and removal the enclosure and enclosure FRUs, actuation of the rack doors and full articulation of the drive drawers.
- [3] Must utilize HPE optical SAS cables to connected to R4M27A, R7K93A or R7K89A.
- [4] Contact your sales representative about availability

Technical Specifications

Data Mover		
SKU Number	S3J02A	
Model Name	HPE Cray Supercomputing Storage Systems Data Services Mover Node	
Processor	EPYC 9124 (16-Core, 3.0 GHz, 200W)	
Number of Processors	One processor	
Memory	32 GB RDIMM SR 1Rx4 4800 MT/s (1x 32 GB)	
Network Controller	HPE 1GbE 4-port Base-T OCP3 Adapter plus the choice of 1 or 2 PCIe cards for Lustre® and NFS Networks	
Internal Storage	2x 480GB SATA M.2	
PCI-Express Slots	1 PCIe x16 Primary Riser 1 PCIe x16 Secondary Reiser	
Power Supply	2x 1000W HPE FlexSlot Power Supply	
Fans	7- standard fans	
Management	Default: HPE iLO Standard with Intelligent Provisioning, HPE OneView Standard (requires download)	
Energy Star	3.0 certified	
Form Factor	1U Rack	
Warranty	3-year parts, 3-year labor, 3-year onsite support with next business day response.	
Dimensions		
8SFF chassis	4.29 X 43.46 X 64.94 cm 1.69 X 17.11 X 25.57 In	
Package	24.2 X 60 X 91.6 cm 9.53 X 23.6 X 36.06 In	
Weight Maximum	8 SFF chassis with 1 processor, 2 power supply, 1 standard heatsink, 12 DIMM, 2 High Speed Network Cards, and 7 standard fans. o	15.54 kg o 34.27 lb
	Package	4.21 kg or 9.281 lb
Rated Line Voltage	100 to 120 VAC 200 to 240 VAC	

Technical Specifications

NFS Protocol Node		
SKU Number	S3J03A	
Model Name	HPE Cray Supercomputing Storage Systems Data Services NFS Protocol Node	
Processor	EPYC 9124 (16-Core, 3.0 GHz, 200W)	
Number of Processors	One processor	
Memory	32 GB RDIMM SR 1Rx4 4800 MT/s (1x 32 GB)	
Network Controller	HPE 1GbE 4-port Base-T OCP3 Adapter plus the choice of 1 or 2 PCIe cards for Lustre® and NFS Networks	
Internal Storage	2x 480GB SATA M.2	
PCI-Express Slots	1 PCIe x16 Primary Riser 1 PCIe x16 Secondary Reiser	
Power Supply	2x 1000W HPE FlexSlot Power Supply	
Fans	7- standard fans	
Management	Default: HPE iLO Standard with Intelligent Provisioning, HPE OneView Standard (requires download)	
Energy Star	3.0 certified	
Form Factor	1U Rack	
Warranty	3-year parts, 3-year labor, 3-year onsite support with next business day response.	
Dimensions		
8SFF chassis	4.29 X 43.46 X 64.94 cm 1.69 X 17.11 X 25.57 In	
Package	24.2 X 60 X 91.6 cm 9.53 X 23.6 X 36.06 In	
Weight Maximum	8 SFF chassis with 1 processor, 2 power supply, 1 standard heatsink, 12 DIMM, 2 High Speed Network Cards, and 7 standard fans.	15.54 kg o 34.27 lb
	Package	4.21 kg or 9.281 lb
Rated Line Voltage	100 to 120 VAC 200 to 240 VAC	

Summary of Changes

Date	Version History	Action	Description of Change
03-Mar-2025	Version 1	New	New QuickSpecs

Copyright

Make the right purchase decision.
Contact our presales specialists.

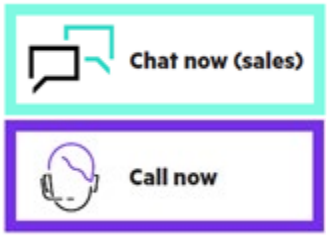

Get updates

Hewlett Packard Enterprise

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

a00038968enw - 16134 - WorldWide - V1 - 03-March-2025