

Overview

Cray ClusterStor E1000 Storage Systems

Is your storage slowing down your HPC compute cluster?

The Cray ClusterStor E1000 storage solution is purpose-engineered to meet the demanding input/output requirements of supercomputers and HPC clusters in a very efficient way. The E1000 parallel storage solution typically achieves the given HPC storage performance requirements with significantly fewer storage drives than alternative storage offerings. That means for HPC users with a fixed budget for the HPC system can spend more of their budget on CPU/GPU compute nodes accelerating time-to-insight. The Cray ClusterStor E1000 storage solution embeds the open source parallel file system Lustre® to deliver this efficient performance. Hewlett Packard Enterprise provides enterprise-grade customer support in-house - and scales out (nearly) linearly, without software licensing for the file system per terabyte capacity or per storage drive. This allows customers to reap the benefits of the open source community while getting enterprise-grade support from Hewlett Packard Enterprise.

The next-generation ClusterStor storage system, ClusterStor E1000, is designed on NVMe generation 4 hardware building blocks based on the latest storage technologies to provide enhanced flexibility and resiliency. It is an I/O and storage subsystem consisting of a global POSIX single name space file system that is configured to provide access between compute clients and ClusterStor E1000 storage nodes using the Lustre parallel file system.

ClusterStor E1000 offer three basic configurations:

- All Flash
- All HDD
- Tiered with both HDD and flash (shown below)

The all flash configuration is suitable for lots of small, random read I/O, high sequential performance required in many AI applications including Machine Learning and Deep Learning. The base rack provides up to 1,440 GB/s read and up to 900 GB/s write with Multi-Rail LNet. The expansion rack provides up to 1,600 GB/s read and up to 1,000 GB/s write Multi-Rail LNet.

The all HDD configuration is used mainly for large sequential I/O and modeling and simulation. The base rack provides up to 90 GB/s read and write up to approximately 7.5 PB of usable capacity, depending on the HDD capacity used. The expansion rack provides up to 120 GB/s read and write up to 10PB of usable capacity, depending on the HDD capacity used.

Customers can mix and match the disk and flash building blocks in a tiered configuration to tailor to their workloads.

What's New

- Up to 80/50 GBs Read/Write aggregate file system throughput from 24 NVMe Gen 4 SSD in a 2 rack unit form factor.
 - Up to 30/30 GBs Read/Write aggregate file system throughput from 212 7.2K RPM SAS HDD in a 10 rack unit form factor.
 - Up to 40/40 GBs Read/Write aggregate file system throughput from 424 7.2K RPM SAS HDD in a 18 rack unit form factor.
 - More than one terabyte per second aggregate file system performance in a single rack.
 - More than 10 petabyte usable storage capacity in a single rack (using 16TB HDDs)
 - Attaches to any supercomputer or HPC cluster that support either Cray Slingshot, HDR/EDR InfiniBand or 200/100 Gigabit Ethernet.
 - Inhouse, enterprise-grade support for performance-accelerating Lustre features like Data on Metadata Targets (DoM), Progressive File Layout (PFL), Multi-Rail LNet, and many more advanced Lustre features .
-

Overview



Cray ClusterStor E1000 – Closed



Overview



Example Cray ClusterStor E1000 – Open (doors removed)



Standard Features

Cray ClusterStor E1000 Building Blocks

ClusterStor E1000 is configured from storage building blocks along with surrounding network, rack and power infrastructure. These building blocks are:

Cray ClusterStor E1000 System Management Unit (SMU)

The ClusterStor E1000 system management unit contains a pair of embedded storage management nodes with SSDs to hold configuration and logging data. The SMU delivers Lustre and ClusterStor system management services. ClusterStor 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, and reporting on system status. The SMU runs a web and CLI server that allows the administrator to monitor, configure, and administer the file system, the ClusterStor management framework, and ClusterStor hardware.

Cray ClusterStor E1000 MetaData Unit (MDU)

ClusterStor E1000 MDUs stores and manages global Lustre metadata with a pair of Lustre metadata servers (MDS) and flash-based metadata targets (MDTs). An MDU is populated with 24 SSDs, two partitions with approximately half the capacity are each configured and are assigned to the two embedded controllers as MDTs using LDISKFS. MDUs can be configured with multiple MDTs, with each controller configured with one or two high speed networks adapters, depending on required network resiliency per MDS.

ClusterStor E1000 configures MDUs to support Lustre DoM, which allows small files or initial portions of files to be stored with their metadata, improving small file performance, stat performance to small files, and reducing interference between small file I/O and streaming I/O. The metadata tier stores metadata and optional DoM in RAID volumes.

ClusterStor E1000 can scale additional MDUs using Lustre Distributed Name Space (DNE) to match specified performance targets.

Cray ClusterStor E1000 Scalable Storage Unit – All Flash Array (SSU-F)

ClusterStor E1000 SSU-F provides flash-based file I/O data services and network request handling for the file system with a pair of Lustre object storage servers (OSS) each configured with a one or more Lustre object storage target(s) (OSTs) to store and retrieve the portions of the file system data that are committed to it. Two OSTs are distributed evenly between the two OSSs in each SSU-F so that both OSSs are active concurrently (that is, OSSs are active-active), each operating on its own exclusive subset of the available OSTs (that is, each OST is active-passive).

A ClusterStor E1000 SSU-F is populated with 24 SSDs. For a throughput optimized configuration, approximately half the capacity are each configured with ClusterStor's GridRAID declustered parity and sparing RAID solution using LDISKFS. For an IOPs optimized SSU-F configuration, a different RAID scheme is used to improve small random I/O workloads. Each controller can be configured with two or three high-speed network adapters configured with Mult-Rail LNet to exploit maximum throughput performance per SSU-F.

ClusterStor E1000 can be scaled to many SSU-Fs and/or combined with SSU-Ds to achieve specified performance requirements. ClusterStor SSU-D features are mentioned below.




Standard Features

Cray ClusterStor E1000 Scalable Storage Unit – Disk (SSU-D)


ClusterStor E1000 SSU-D provides HDD-based file I/O data services and network request handling for the file system with similar OSS and OST features mentioned above for ClusterStor E1000's SSU-F.

ClusterStor E1000 has multiple SSU-D configurations with either one, two or four ClusterStor ultra-dense HDD enclosures. Each ultra-dense disk enclosure is configured with 106 SAS HDDs and contains two Lustre OSTs, each configured with ClusterStor's GridRAID declustered parity and sparing RAID solution using LDISKFS. Each SSU-D controller is configured one high-speed network adapter and redundant SAS adapters to connect to one or more ClusterStor ultra-dense HDD enclosures.


ClusterStor E1000 can be scaled to many SSU-Ds and/or combined with SSU-Fs to achieve specified performance requirements.



SSU-D1



SSU-D2



SSU-D4

ClusterStor E1000 Scalable Storage Unit – Disk (SSU-D)



Service and Support

Warranty

1-0-0 warranty

This product is covered by a global limited warranty and supported by Hewlett Packard Enterprise Services. Hardware diagnostic support and repair is available for one year from date of purchase. Support for software is available for 90 days from date of purchase. Enhancements to warranty services are available through HPE or customized service agreements.

Solid State Drives are subject to maximum usage and or maximum supported lifetime limitations, whichever occurs first. Maximum Supported Lifetime is the period in years set to equal the warranty for the device. Maximum usage limit is the maximum amount of data that can be written to the device before reaching the device's write endurance limit.

Recommended Services

HPE Proactive Care with 24x7 coverage

HPE Proactive Care gives customers an enhanced call experience. When your products are connected to HPE, Proactive Care helps prevent problems and maintains IT stability by utilizing personalized proactive reports with recommendations and advice. This Service combines three years proactive reporting and advice with our 24x7 coverage, four hour hardware response time when there is a problem. <https://www.hpe.com/h20195/v2/GetPDF.aspx/4AA3-8855ENW.pdf>

HPE Foundation Care 24x7

HPE Foundation Care 24x7 gives you access to HPE 24 hours a day, seven days a week for assistance on resolving issues. This service includes need based Hardware onsite response within four hours. In addition, collaborative software support is included in this service that provides troubleshooting assistance on industry leading software running on your HPE server. Simplify your support experience and make HPE your first call to help resolve hardware or software problems.

<https://www.hpe.com/h20195/V2/GetDocument.aspx?docname=4AA4-8876ENW&cc=us&lc=en>

HPE Foundation Care NBD

HPE Foundation Care Next Business Day connects you to HPE during business hours for assistance on resolving issues – This service features need based next business day hardware onsite response and software call back within two hours. In addition, Collaborative software support and provides troubleshooting assistance on industry leading software running on your HPE server. Simplify your support experience and make HPE your first call to help resolve hardware or software problems.

<https://www.hpe.com/h20195/V2/GetDocument.aspx?docname=4AA4-8876ENW&cc=us&lc=en>

Other Related Services

HPE Hardware Installation and Startup Service for Cray ClusterStor

This service provides onsite Installation and Startup for a Cray ClusterStor E1000 according to product specifications including options. An assigned Installation Project Manager coordinates end to end tasks from factory integration to onsite deployment, installation and test of the Cray ClusterStor E1000.

HPE Technology Services Support Credits

Offer flexible services and technical skills to meet your changing IT demands. With a menu of service that is tailored to suit your needs, you get additional resources and specialist skills to help you maintain peak performance of your IT. Offered as annual credits, you can plan your budgets while proactively responding to your dynamic business.

HPE Datacenter Care

Helps improve IT stability and security, increase the value of IT, and enable agility and innovation. It is a structured framework of repeatable, tested, and globally available services “building blocks.” You can deploy, operate, and evolve your datacenter wherever you are on your IT journey. With HPE Datacenter Care, you benefit from a personalized relationship with Hewlett Packard Enterprise via a single point of accountability for HPE and others' products. For more information, visit

<http://www.hpe.com/services/datacentercare>



Service and Support

HPE Education Services

Keep your IT staff trained making sure they have the right skills to deliver on your business outcomes. Book on a class today and learn how to get the most from your technology investment. <http://www.hpe.com/ww/learn>

HPE Support Center

The HPE Support Center is a personalized online support portal with access to information, tools and experts to support HPE business products. Submit support cases online, chat with Hewlett Packard Enterprise experts, access support resources or collaborate with peers.

Learn more at <http://www.hpe.com/support/hpesc>

Options

Notes: Included options are covered under the HPE Service Contract applied to the HPE Server and/or Storage. No separate HPE support services need to be purchased.

Warranty and Support Services will extend to include HPE options configured with your server or storage device. The price of support service is not impacted by configuration details. HPE sourced options that are compatible with your product will be covered under your server support at the same level of coverage allowing you to upgrade freely. Installation for HPE options is available as needed. To keep support costs low for everyone, some high value options will require additional support. Additional support is only required on select high value workload accelerators, high-speed switches, InfiniBand and UPS batteries over 12KVA. See the specific high value options that require additional support [here](#)

For more information visit: <http://www.hpe.com/services>

Parts and Materials

Hewlett Packard Enterprise 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.

The defective media retention service feature option applies only to Disk or eligible SSD/Flash Drives replaced by Hewlett Packard Enterprise due to malfunction.



Technical Specifications

Cray ClusterStor E1000

- **Filesystem**
Lustre 2.12 with supported enhancements
- **Data Path**
ClusterStor GridRAID with LDISKFS

Scalable Storage Unit – Flash (SSU-F)

- **Storage Controller**
High availability, dual node pair with 24 NVMe Gen 4 SSDs
- **SSD Options**
1 DDPD: 1.92 / 3.84 / 7.68 / 15.36 TB
3 DPWD: 1.6 / 3.2 / 6.4 / 12.8 TB
- **Maximum SSUs per Rack** (dependent on type of network supported)
 - 18 per Base Rack
 - 20 per Expansion Rack
- **Maximum usable capacity per rack**
Up to 5PB (using 15.36 TB SSDs)

Scalable Storage Unit – Disks (SSU-D)

- **Storage Controller**
High availability, dual controller pair with 2 NVMe Gen 4 SSDs
- **Disk Enclosure**
Ultra-dense SAS JBOD with 106 HDDs in 4U
- **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/ 6 / 10 / 12 / 14 / 16 TB SAS HDDs
- **Maximum SSUs per Rack** (dependent on type of SSU-D configuration and of network supported)
 - 4 per Base Rack
 - 4 per Expansion Rack
- **Maximum usable capacity per rack**
Up to 10PB (using 16TB HDDs)

MetaData Unit (MDU)

- **Base Configuration**
High-availability, dual controller pair with 24 NVMe Gen 4 SSDs RAID protected
- **Number of files**
 - Up to eight billion per MDU in standard configuration with LDISKFS
- **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 4 SSDs RAID protected for system management, logging, and boot services

Client Network Options

Cray Slingshot, HDR / EDR InfiniBand, 200/100 Gb/s Ethernet



Summary of Changes

Date	Version History	Action	Description of Change
05-Oct-2020	Version 5	Changed	Overview, Standard Features, sections were updated
14-Sep-2020	Version 4	Changed	Standard Features section was updated
08-Sep-2020	Version 3	Changed	Overview and Service and Support sections were updated. Minor improvements/changes
03-Aug-2020	Version 2	Changed	Overview, Standard Features, Service and Support and Technical Specifications sections were updated Added new HDD and SSD options Minor improvements throughout
06-Jul-2020	Version 1	New	New QuickSpecs



Copyright

Make the right purchase decision.
Contact our presales specialists.



Chat



Email



Call



[Get updates](#)



© Copyright 2020 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

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

a00062172enw - 16535 - WorldWide - V5 - 05-October-2020