

# QuickSpecs

## Seagate ClusterStor L300

### Overview

#### Seagate ClusterStor L300

The Seagate ClusterStor L300 is an integrated hardware and software implementation of the Lustre® file system instantiated as a rack scale appliance supporting High Performance Computing (HPC) and High Performance Data Analytics (HPDA) applications in technical computing markets such as Oil and Gas, Life Science / Public Health, BioTech, Finance, Weather Forecast / Climate Prediction, Electronic Design Automation, Engineering Design / Manufacturing, National Science / Research, Government and Defense.

#### Benefits

The ClusterStor L300 is an ideal scale-out high performance data storage solution for productivity critical public sector and commercial HPC and Big Data application environments, which require industry leading performance efficiency and reliable storage capacity that scale linearly with the addition of modular storage processing building blocks. Specifically, the ClusterStor L300 excels in user environments requiring high availability, sustained performance supporting hundreds to tens of thousands of client compute nodes, as well as low latency parallel file system access to massive data sets with data storage capacity ranging from peta-bytes (PB) to multiple tens of PBs. With the Seagate ClusterStor L300 solution, technical computing organizations gain...

- Production proven linear scalability from tens of GB/sec to over 1TB/sec sustained throughput performance
- Industry highest performance efficiency per rack yielding low Total Cost of Ownership (TCO)
- Support for over 25,000 client compute nodes per single file system with support for up to 16 billion files
- Seagate Enterprise Edition Lustre v2.7
- Distributed Namespace (DNE) support providing scalable and reliable metadata performance
- End to end solution management and ease of use, critical to attain low OPEX

All ClusterStor L300 modular hardware and embedded software elements are pre-installed, pre-configured, and pre-tested prior to its delivery to customer end users. Seagate ClusterStor is the only Lustre solution that is completely manufactured, tested, delivered and supported by a single vendor, responsible for all aspects of hardware and software design, architectural integration, deployment, service and support. Seagate's holistic engineered solution includes enterprise disk drives, high availability storage enclosures, embedded file system servers, high speed networking, as well as the industry's best practice in Lustre file system software tuning tailored to optimize solution level productivity, failover robustness as well as sustained throughput performance efficiency and near linear modular system scale.

#### The ClusterStor L300 solution delivers:

- Purpose-built scale-out high performance data storage solution - provides the performance and capacity you need with minimal administration requirements; future growth in performance and/or capacity is easily satisfied with the addition of modular building blocks
- Factory installation, pre-configuration and burn-in ensures a working solution on delivery, and maximum uptime for high performance production storage environments
- Rack & roll deployment in hours, versus months, ensuring faster client acceptance
- One-call for all hardware and software support from a single point of contact



## Platform Information

### ClusterStor L300 Architectural Elements

ClusterStor L300 provides a pre-installed and pre-configured deployment of the Lustre file system, based on Seagate Enterprise Edition Lustre version 2.7 or higher designed to deliver sustained high performance throughput, with integrated high availability and resiliency features, end-to-end system management, and the ability to expand existing systems. Due to industry leading performance efficiency, ClusterStor uniquely offers these features with minimal use of energy consumption, cooling, and data center footprint. Truly scaling file system performance efficiently means aggressively reducing the space, power, cooling and time required to achieve HPC performance and results.

Lustre is open-source client/server based software that uses a distributed architecture and is designed for large-scale compute and I/O-intensive, performance-sensitive applications. The Lustre architecture is used for many different types and sizes of Linux clusters, but it is best known for powering many of the Top 500 HPC clusters in the world.

Each ClusterStor L300 system includes a base rack, containing a pair of redundant, top of rack (TOR) management switches providing high availability, out-of-band management of the ClusterStor solution. Each base rack also includes the Cluster Management Unit (CMU), which is provided for Lustre cluster Management Server functions, (MGT), and Metadata Server (MDS), functions of the ClusterStor solution. Redundant high availability server node pairs with RAID 1 and 10 HDD data protection storage in a 2U24 chassis is used to support management and metadata functions. Metadata storage is provided to store metadata information on each file and directory in ClusterStor, for the primary Metadata Target (MDT).

### Lustre Metadata Server (MDS)

Within the CMU, the Lustre Metadata Server (MDS) manages all metadata associated with a Lustre file system and responds to Lustre client requests to access file system data as needed. Metadata managed by the MDS includes the file system namespace (file and directory names), permissions, file layout information, and network request handling. The MDS stores file system metadata in one or more Metadata Target(s) (MDT) associated with the file system.

Within ClusterStor L300 solution, a pair of MDSs runs in a primary/secondary active-passive configuration on a pair of physical servers that are connected to the MDT shared storage which is hosted as a dedicated RAID 10 volume. The MDS is only active on one server node at a time as the default configuration. In the event the primary MDS node fails, the failover process ensures that the secondary MDS node is activated, takes control of the MDT, and takes over all MDS functions

### Lustre Distributed Namespace (DNE)

Within the ClusterStor L300 architecture, Lustre Distributed Namespace (DNE) capability provides additional MDT targets within Seagate's Enterprise hardened Lustre 2.7 or higher, enabling enhanced MDS performance and increased inode support, with improvements in load balancing metadata operations for increased reliability. This feature allows multiple MDTs operating through multiple MDS nodes, to be configured and operated as part of a single file system. ClusterStor support for the DNE component, referred to as the Additional DNE Unit (ADU) feature, adds additional MDS/MDT capability to a file system cluster. ClusterStor L300 systems with Seagate Enterprise Edition Lustre 2.7 or above may be configured with up to 8 ADUs. ClusterStor ADUs may be added, providing up to 16 additional MDTs to the primary MDTs, significantly scaling metadata performance and inode capacity as needed.


### Lustre Object Storage Servers (OSS)

The basic building block of ClusterStor is the Scalable Storage Unit, or SSU. The SSU is five rack units (5U) high and delivers up to 12GB/sec of sustained file system performance as verified by the IOR benchmark suite over Mellanox EDR high speed InfiniBand networks. The SSU contains two Lustre Object Storage Servers (OSSs) in a high availability (HA) configuration, utilizing high performance Intel processors. Each OSS controls one Object Storage Target (OST), for a total of two OSTs per SSU using Seagate GridRAID with T10-PI type 2 formatted drives for data protection and integrity. The two OSTs are configured with Seagate GridRAID as 4+1 Drives per OST, using 8+2 Data plus Parity blocks, and 2 integrated sparing drives distributed as spare blocks within the GridRAID architecture, providing the fastest recovery possible for a failed hard drive. The ClusterStor L300 base rack may contain up to six SSUs.

Platform Information



**L300 Base Rack**



**L300 Expansion Rack**

The ClusterStor L300 further expands with the addition of expansion racks. Each expansion rack contains a pair of redundant, top of rack (TOR) management switches and up to seven SSUs. With the addition of each SSU (or addition of expansion rack(s) containing SSUs), the ClusterStor L300 increases performance and capacity.

## Platform Information

### STEP A : Start with Base Rack



The Seagate provides Base Rack includes Top of Rack Switches, ClusterStor Management Unit (CMU) and from one to six Scalable Storage Units (SSUs)

### STEP B: Add Expansion Rack



The Seagate provided Expansion Rack includes Top of Rack Switches and from one to seven Scalable Storage Unit (SSUs)

### Step C: Add additional Expansion Racks



The Seagate Industry unique pre-configured and factory tested high availability rack-scale Lustre solution provides a fully integrated high performance Storage solution with comprehensive end to end management and monitoring simplifies deployment, expansion and Support

The ClusterStor architecture removes complexity and is naturally more efficient by design. The ClusterStor SSU integrates operating system, data protection, the Lustre® file system and management into a single high availability building block that consolidates storage, network and server processing. The result is an easy to deploy, easy to use and easy to manage solution. There is no need to guess at how to scale; each SSU is a balanced performance building block delivering a predictable level of performance and storage capacity.

Overall system performance is directly proportional to the number of SSUs due to ClusterStor's efficient internal optimization that yields industry leading linear performance scalability. End users simply configure the ratio of SSUs and/or Expansion Storage Units (ESUs) to satisfy their performance and/or data capacity needs. See the Installation Guide for details on supported configurations including SSUs and attached ESUs.

### ClusterStor L300 Scalable Storage Unit (SSU)

The Scalable Storage Unit (SSU) in the ClusterStor L300 system is based on the Seagate designed, developed and manufactured industry leading enterprise 5U84 3.5 inch drive platform, featuring:

- Up to 82 SAS hard disks and 2 Solid State Drives (SSDs) per enclosure
- Enterprise level Reliability, Availability and Serviceability
- Platinum Certified PSUs
- Adaptive cooling technology

## Standard Features

### ClusterStor L300 Scalable Storage Unit (SSU)



1.- Two Drawers each containing 42 drive bays

1.- Embedded Server Modules (ESM)

The Seagate ClusterStor L300 SSU enclosure is pre-integrated into the Seagate provided rack and offers a high density design, with no single point of failure. At the front of the enclosure there are two drawers, each containing 42 drive bays, supporting 3.5 inch dual-ported SAS drives housed within carriers, with redundant data paths from each server/controller module to each drive. Each drawer can be easily accessed from the front of the enclosure, eliminating the need to remove the SSU from the rack. Drives are hot-swappable, enabling the system to remain fully operational, within normal thermal specifications, for several minutes when a drive is replaced. The platform offers fully redundant power and cooling. The enclosure is designed to maximize the reliability of the High Density Drive (HDD), with vibration dampening and temperature monitoring technology throughout the system. The front panel includes an operator control panel located on the front bezel, which is used to indicate the shelf number where the enclosure is located in a rack, and the general health of the enclosure. The rear of the enclosure houses two Embedded Server Modules (ESMs), two power supplies and five cooling modules.

The SSU minimizes energy use through advances in Seagate technology including: individual drive power control; advanced adaptive cooling technology and Platinum Certified PSUs that are 92% efficient at 50% load. Additionally, The SSU is designed to meet and exceed stringent worldwide requirements for recycling and environmental friendliness.

#### SSU Embedded Server Module (ESM)

The Embedded Server Module (ESM), also referred to as a storage application controller or Lustre OSS controller, is a dedicated x86 hardware server module that utilizes the Storage Bridge Bay (SBB) industry standard form factor, enabling it to plug directly into the Seagate (OneStor) family of enclosures. The ESM combines the attributes of a standard server, a RAID controller and a JBOD module into one modular design capable of running standard operating systems. Up to two ESMs can be installed in a single SSU enclosure to provide embedded server and storage application functions.

#### ClusterStor Manager (CSM)

ClusterStor Manager (CSM) is a browser-based tool that provides a unified system management view, presenting all necessary information from the different sources of information on the platform including: hardware, Lustre and storage. CSM supports the Chrome, Firefox, and Internet Explorer browsers and provides functions for use in system configuration, deployment, and day to day end-user operations. CSM also provides the ClusterStor CLI (CSCLI), which includes a robust set of CLI commands enabling many system actions and functions.

The ClusterStor Installer guides the administrator through the automated deployment of all ClusterStor software, configuration and setup of the cluster, including verification tests prior to customer shipment. Once delivered and post-configured at the customer site, CSM provides a browser-based or command line interface feature set that monitors and manages the entire storage cluster and provides assistance with diagnosing operational issues. This browser-based tool is accessible through a direct connection between the customer's IP network and the ClusterStor Management Server nodes. CSM is fully routable and can be accessed over the Internet if the customer network is configured to support this.

#### ClusterStor Serviceability

All systems logs are consolidated across the entire cluster into the central management server, this includes logs from every layer of the system from disk drive data through reporting information to Lustre error messages. This enables administrators to easily identify anomalies in the ClusterStor

## Standard Features

operation and proactively service the system to avoid downtime/performance impacts. The hardware incorporates a modular design that enables commonly failing components to be field serviceable. The most commonly anticipated failure items are disk drives, controllers/server modules, power supplies and cooling modules, all of which can be easily serviced without downtime.

SSU disk drives are located in a pull out drawer and are released by the press of a button. LED indicators help service personnel identify the correct disk to be repaired. The maintenance operation to remove a failed disk drive and install a replacement drive in an SSU typically take no more than 2 minutes to complete. After a disk drive is replaced, the system uses a discovery process to identify the new drive and include this new hardware in the available storage inventory. HDD replacement on the CMU metadata storage is identical to the SSU disk drive replacement procedure. The controllers themselves, identified as critical to the system are individually serviceable, with a failed module being easily replaced even while data is still being served and all OSTs being accessed thru its partner ESM

## Additional Features

### Key Features

#### **Seagate Enterprise Edition Lustre 2.7 Distributed Namespace - Expands Metadata Performance and Scalability**

With Lustre 2.7, customers may optionally add Distributed Namespace (DNE) metadata servers. Lustre client metadata operations can now be allocated across multiple metadata servers. This enables greater scalability in the size of the namespace with support for up to 77 billion files and up to 700% faster metadata performance.

Expand metadata performance and scalability with up to sixteen Distributed Namespace metadata servers, each configured in an active / active high availability pair within Seagate 2U enclosures. Each 2U Distributed Namespace Server provides two metadata servers and metadata storage capacity, along with seamless integration including manageability.

#### **ClusterStor GridRAID – Faster Drive Rebuilds**

As a standard feature in the ClusterStor L300 GridRAID provides up to 400% faster data reconstruction rates, compared to legacy RAID 6 methods

GridRAID consolidates the number of Object Storage Targets (OSTs) by 4 times compared to legacy RAID 6 (8+2) methods, thus reducing management tasks and costs

#### **ClusterStor Manager - Comprehensive System Management**

ClusterStor Manager, a comprehensive system management application, is included with every system. ClusterStor Manager is web browser based, part of ClusterStor's distributed management framework and is responsible for pulling everything together as a singly managed system. ClusterStor Manager consolidates management of the storage infrastructure, RAID data protection layer, operating system, and the file system into a single, easy-to-use, administrator interface bringing unprecedented system visibility to the system administrator.

### Engineered Solution

Factory installation, pre-configuration and burn-in ensures a working solution on delivery, and maximum uptime for high performance production storage environments  
Rack & roll deployment in hours, versus months, ensuring faster client acceptance  
One-call for all hardware and software support from a single point of contact

### Removes Complexity

ClusterStor's architecture removes complexity and is naturally more efficient by design. The ClusterStor Scalable Storage Unit integrates operating system, data protection, the Lustre® file system and management into a single high availability building block that consolidates storage, network and server processing. The result is an easy to deploy, easy to use and easy to manage solution. There is no need to guess at how to scale; each Scalable Storage Unit is a balanced performance building block delivering a predictable level of performance and storage capacity.

### Linear Scale

Industry highest performance efficiency per rack yielding low TCO  
Overall system performance is directly proportional to the number of Scalable Storage Units due to ClusterStor's efficient internal optimization that yields industry leading linear performance scalability. End users simply add Scalable Storage Units and/or Expansion Storage Units to satisfy their performance and/or data capacity needs

## Warranty

### Warranty Information

- **Hardware Warranty**  
12 month limited hardware warranty included with original hardware purchase  
Provides return to base hardware repair / replacement / credit via standard process on confirmed failures
- **Software Warranty**  
Ninety (90) days limited software warranty included with original hardware purchase  
Delivers remote support on the standard base software version and filesystem which shipped with your system for the first 90-Days

### Extended Software and Hardware Support

Tiered offering with all options available from 1 through 5 year extended support.

<b>Software Extended Support</b>		
Options available for 8x5 standard business support in geography or 24x7 business		
Bronze Hardware Support NBD Part Replacement Services	Silver Hardware Support NBD Engineer Repair Services + Part	Gold Hardware Support 4hr Engineer Repair Services + Part



## Summary of Changes

**NOTE:** Sizing for storage servers should be tailored to the specific workload requirement. Data protection strategy and average file size should be taken into account for optimal configuration. As such it is recommended that you engage with your Hewlett Packard Enterprise representative to document your individual solution needs and design an environment with required services to deploy.

<b>General Information</b>	<b>Parallel File System Performance</b>	Up to 96 GB/s bandwidth performance per 42RU height base rack Up to 112GB/s bandwidth performance per 42RU height expansion rack	
	<b>File System Capacity (raw)</b>	Up to 4,920 TB per rack using 10TB SAS HDDs Up to 5,740 TB per expansion rack using 10TB SAS HDDs	
	<b>Scalable Storage Unit (SSU)</b>	Up to 6 in base rack (first) Up to 7 in storage expansion racks (second rack or greater)	
	<b>Object Storage Servers</b>	Up to 12 in base rack (first) Up to 14 in storage expansion racks (second rack or greater)	
	<b>Metadata Management Server / Distributed Namespace Server Management Server</b>	High availability, embedded application controller pair, in a 2U 24 drive enclosure High availability, embedded application controller pair, in a 2U24 drive enclosure	
	<b>Client Access</b>	InfiniBand QDR, EDR or FDR, Ethernet 40GE or 100GE Omni Path 100Gbit	
	<b>Management Network</b>	1 Gigabit Ethernet (dual management network)	
	<b>File System</b>	Seagate Enterprise Edition Lustre® 2.7 + Seagate supported enhancements	
	<b>Maximum inodes</b>	Lustre® 2.7 - Up to 77 Billion inodes (requires optional Distributed Namespace Servers)	
	<b>Hierarchical Storage Management (HSM)</b>	Lustre Robinhood Change Log Management Ready	
	<b>Disk Drives</b>	<b>SSU Hard Disk Drive</b>	Dual ported 12Gb/s SAS, 7.2K RPM drives (4, 6, 8 or 10 TB capacity per drive)
		<b>SSU SSD</b>	Dual ported 800GB SAS drives
<b>Metadata Management Unit Hard Disk Drive</b>		Dual ported 12Gb/s SAS, 10K RPM drives ( 900GB or 1.8TBcapacity per drive)	
<b>System Management Unit Hard Disk Drive</b>		Dual ported 12Gb/s SAS drives (300 GB 15K RPM and 900GB 10K RPM capacity per drive)	
<b>Dimensions</b>	<b>Height</b>	1,991 mm (78.4 in)	
	<b>Width</b>	600 mm (23.62 in)	
	<b>Depth</b>	1,200 mm (47.24 in)	

## Summary of Changes

<b>Systems Availability</b>	<b>Weight</b>	1,141 Kg (2,510 lbs)
	<b>Hot Swappable</b>	Disk Drives, Power Supplies, Fans, Power Cooling Modules and Server Modules
<b>Power Consumption</b>	<b>Power</b>	Redundant Power Supplies and Power Cooling Modules
	<b>SSU Power Cooling Modules</b>	5 redundant fan modules per SSU, each with dual fans
	<b>Base Rack Configuration</b>	11.93 Kilowatts Nominal
<b>Heat Dissipation</b>	<b>Storage Rack Configuration</b>	12.27 Kilowatts Nominal
	<b>Base Rack Configuration</b>	40,721 BTU Nominal
<b>Altitude and Temperatures</b>	<b>Storage Rack Configuration</b>	41,882 BTU Nominal
	<b>Operational Altitude</b>	-30 to 3048m (-100 to 10,000ft)
	<b>Operational Temperature Range</b>	5°C to 32°C
	<b>Temperature Variance</b>	De-rated by 1°C/300m above 900m below the specified maximum temp.
	<b>Humidity</b>	20% to 80% non-condensing

## Summary of Changes

Date	Version History	Action	Description of Change
11-Jul-2017	From Version 1 to 2	Changed	Update and correct the content throughout the QuickSpecs
15-Aug-2016	Version 1	Created	Create QuickSpecs for Seagate ClusterStor L300



Sign up for updates



**Hewlett Packard  
Enterprise**

© Copyright 2017 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.

Windows and Microsoft are registered trademarks of Microsoft Corp., in the U.S.

c05069175- 15596 - Worldwide – V2 - 11-July-2017