

CRAY CLUSTERSTOR™

High-Performance Storage for
High-Performance Discovery



Storage for Every HPC Challenge

ClusterStor storage gives you **top performance, easy manageability and expert Cray support** to boost productivity and lower TCO. You get energy efficiency, a small datacenter footprint and 1 TB/s file system performance, plus a range of service and support options.

When growing performance demands meet budget constraints, you're under pressure to get the best ROI for your storage investment. Yesterday's technologies won't meet tomorrow's data storage requirements. Driven by ever-increasing performance and capacity requirements, storage is expected to account for 21 percent of HPC spend by 2021 — up from just 10 percent in 2008. That takes away from the compute, staff and applications you need to solve problems.

We balance the value equation with the right performance levels, speed, scalability, data protection and availability to fit your requirements and budget. ClusterStor systems reduce complexity with seamless design and easy installation, workload management and data flow. It all adds up to the industry's highest-performing and most efficient data storage platform with the lowest overall TCO.

More performance, better ROI

ClusterStor gives you enterprise-level performance with more capacity, fewer drives, less need for IT support and more data access. It supports the industry's fastest HPC-sustained throughput at 1.7 TB/s with data capacity over 80 PB.

Time and money saved

The ClusterStor system is designed for efficiency, with minimal components and efficient data flow. We build each system to your requirements so you don't waste time or money on technology you don't need.

Business-boosting insights

Faster discovery is a competitive advantage. The ClusterStor system speeds your time to insight by minimizing maintenance, performance gaps and mixed-workload I/O issues. It scales up to 100 PB with no loss of efficiency or performance.

The ClusterStor Difference

We run rigorous component, system and manufacturing test cycles to ensure excellent system-level integration and device-level reliability. Storage enclosures feature "cluster in a box" integration of dual high-availability object storage, metadata and management servers, monitored by the ClusterStor Manager software stack — a single management system for both hardware and software. This component-level approach creates modular building blocks for granular performance and capacity scalability. Its design reduces time from setup to production and optimizes system availability and efficiency.

- High-performance storage for productivity-critical applications
- Robust, enterprise-ready features and management tools
- Solution integration, system validation test and field support
- Cray disk drives, enclosures, object storage, metadata and management servers
- Cray expertise as the only vendor to build and support an end-to-end Lustre® solution

Who's Using ClusterStor Storage?

- 8 of the top 20 supercomputers — twice as many as the next vendor
- 6 out of 7 global oil and gas supermajors
- 8 out of the world's top 9 weather centers
- Over 80 universities, national labs and multinational companies worldwide



Cray ClusterStor L300

ClusterStor L300 all-HDD Lustre storage is optimized to deliver the performance you need with the smallest possible number of disks, enclosures and racks. This complete rack-scale solution delivers dense performance, accessible data and actionable insights.

It's a good choice for environments running several applications with large, sequential I/O workloads.

Cray ClusterStor L300N

ClusterStor L300N hybrid SSD/HDD storage with our flash-accelerated NXD software minimizes the unpredictable application performance caused by mixed I/O patterns. It selectively routes I/O and uses transparent flash acceleration to shield the application, file system and users from complexity.

It's your best value for mixed I/O workload performance.

Cray ClusterStor L300F

Our ClusterStor L300F model allows you to add a flash storage pool, creating a truly hybrid system. It's designed and optimized to overcome the latency experienced by rotating media — the remaining IOPS bottleneck for Lustre.

Unlike other flash solutions, the ClusterStor Lustre-based L300F system requires no specialized training, tuning or benchmarking.

Lustre users with high IOPS needs can dramatically reduce the run times of their applications, resulting in more iterations and faster time to insight.



Three Storage Solutions, Zero Challenges

ClusterStor L300 All-HDD Lustre Storage System

Achieves performance requirements with the lowest number of components by getting the maximum amount of performance from each storage device. Ideal for large, sequential I/O workloads.

ClusterStor L300N Hybrid (SSD/HDD) Storage System

Enabled by NXD flash acceleration software to cost-effectively solve the mixed I/O challenge while shielding the application, file system and users from complexity.

ClusterStor L300F Flash Hybrid Storage System

The ability to add a flash storage pool creates a truly hybrid system.



Cray® View for ClusterStor™ Makes Your Storage Experience Even Easier

Cray View for ClusterStor is the industry's first complete Lustre® storage performance analysis application. With always-on metrics you can maximize your resources, optimize system throughput, and provide users with the best results — without impacting performance. All the information you need is right there at your fingertips, in a single view:

- **Job Runtime Variability:** Real-time and historical views of data to help you understand what's impacting user jobs
- **Event Correlation:** A unified system view so you can correlate events that impact performance
- **Trend Analysis:** Data-driven visualization and analysis of historical data helps for identifying trends and shaping changes to the system
- **Alerting:** Threshold engine enables customized alerts on any metric

Cray® ClusterStor™ Specifications

File System Capacity (Raw)	Up to 4,920 TB per base rack using 10 TB SAS HDDs; up to 5,740 TB per expansion rack using 10 TB SAS HDDs
Scalable Storage Unit (SSU)	Up to 6 L300/L300N in the first rack in a system (base); up to 7 L300/L300N in the second/subsequent racks (storage base racks) Up to 16 L300F in any rack in a system (base & storage base racks) Intermix of SSU types within a rack is configuration dependent; contact Cray for configuration rules
Object Storage Servers	Up to 12 in base rack (first); up to 14 in storage expansion racks (second and greater)
Meta Data Management Unit (MMU)	Base configuration: high-availability server pair; 2U24 drive enclosure Expansion option: up to 8 2U24 MMUs configured with the Lustre® 2.5 Distributed Namespace functionality
System Management Unit (SMU)	High-availability server pair; 2U24 drive enclosure
Client Network Access	InfiniBand™ EDR, FDR, QDR; Ethernet 40/100, Intel® Omni-Path
Management Network	1 Gigabit Ethernet (dual management network with high-availability failover)
File System	Lustre 2.7 + supported enhancements
Maximum Nodes	Lustre 2.7 – up to 16 billion (requires optional distributed namespace servers)
SSU Drive Configuration	L300: 2 SSDs, RAID 1, 1+1, 82 HDDs, 3.5" 7.2K RPM with ClusterStor GridRAID L300N: 2 SSDs, 3.2 TB, 10 DWPD, RAID 1, 1+1; 82 HDDs, 3.5" 7.2K RPM or 82 HDDs 3.5" 10K RPM with ClusterStor GridRAID L300F: 24 SSDs, 1.6 TB / 3.2 TB, RAID 10
Expansion Storage Unit Configuration	82 HDDs, 3.5" 7.2K RPM with ClusterStor GridRAID
MMU Configuration	22 HDDs, 2.5" 10K RPM, 5+5
SMU Configuration	7 HDDs, 2.5" 10K RPM, 1+1 & 2+2, 1 spare, 5 HDDs, 2.5" 15K RPM, 1+1 x2, 1 spare