

CRAY® SONEXION® 3000

The Cray® Sonexion® 3000 scale-out Lustre® storage system provides the cornerstone of Cray's performance-engineered, workflow-driven solutions for big data and supercomputing. The Sonexion 3000 system scales efficiently, delivers up to 38 percent more real-world throughput per rack unit, and reduces TCO by up to 25 percent compared to conventional Lustre solutions. Cray offers a single point of support for everything, including all hardware and software — from the compute systems to the storage.



CRAY SONEXION 3000


Cray® Sonexion® 3000 Storage Solution

Brought to you by Cray, the global leader in storage performance and I/O for supercomputing, the Sonexion system's modular, pre-integrated and compact design keeps costs low while delivering the right performance for analytics, compute clusters and supercomputers of all types. Performance and capacity **scale efficiently** in modular building blocks, reducing the number of hard disk drives needed to achieve sustained performance at scale. **Performance is optimized end to end**, from the compute clients to the network to the storage subsystem, based on the application workload.

Overall, **TCO is reduced by up to 25 percent** — using fewer components and racks than the competition to achieve your desired performance goals.

Data protection is provided through Grid RAID, a declustered parity form of data protection that speeds up rebuild times three and a half times over traditional RAID, while maintaining Lustre performance.

Management tools including the Cray Systems Snapshot Analyzer (SSA), Cray Sonexion Storage Manager (CSSM) and new management diagnostic infrastructure provide a comprehensive set of health monitoring and management tools, essential to maintaining and supporting Lustre at scale.



Scales Efficiently

Unlike monolithic controller-based SANs, the Sonexion system scales performance and capacity in balanced, modular building blocks. This ensures your desired performance levels are achieved using the fewest number of disk drives and external components possible — to reduce cost and complexity. Finally, the Sonexion system's compact form factor reduces the total storage footprint — including external cables, servers, HDDs and racks — by 30 percent on average over controller-based SANs.

The rack-based storage system utilizes high-performance scalable storage units (SSUs), which combine balanced levels of capacity and performance. The SSU consolidates and integrates everything (block storage, networking, OS and file system components) into a dense, easy-to-manage, rackable five-unit building block. SSUs deliver a finite amount of performance capacity. As SSUs are

added to the Sonexion system, performance scales near-linearly.

A capacity-optimized configuration of the Sonexion 3000 system delivers up to 3.36 PB of usable file system capacity and up to 67 GB/s of performance in a single 42-U rack, using 8 TB disk drives. By contrast, a performance-optimized configuration would deliver half the capacity per rack, but up to 100 GB/s of sustained throughput using the latest Seagate HPC drive technology.

The Cray Sonexion 3000 system efficiently scales single file system capacity and performance as needed, in modular units, using the fewest number of components and racks to reduce TCO. The overall result is that the Sonexion system delivers more bandwidth per rack unit, per gigabyte of capacity, than any other Lustre implementation.

CRAY SONEXION 3000

Planning and Deployment Challenges

- Planning and deployment time are measured in weeks to months on average. Often these costs are hidden in professional services which, based on \$3,000 per day, can add up to \$60,000 for one month of protracted services.
- Why professional services? To deal with the complexity of installing, configuring, mapping and managing component sprawl — from HBAs, cables and host mappings to logical unit names (LUNs) and RAID groups — and deploying Lustre across various compute systems.
- Cost and risk sourcing, integrating, testing and supporting components from multiple vendors.
- Developing and honing in-house Lustre expertise to offset the cost of professional services increases the liability of and dependency on individuals within the organization.
- Powering and cooling extraneous racks of servers and components of controller-based solutions.

Cray Reduces TCO By:

- Reducing deployment time from months to days at no charge to the customer
- Reducing the total datacenter footprint by 30 percent compared to the competition
- Providing a single point of support for everything — all hardware and software — which optionally includes compute capabilities
- Reducing power consumption by 15-20 percent per year, on average, over the competition



Performance Engineered End to End

Cray brings over 40 years' experience in storage performance engineering and workload profiling which started in supercomputing — and high-performance storage. Achieving optimal application performance requires deep holistic knowledge across a broad set of disciplines — from application scalability to compute to networking to storage. Cray's holistic systems expertise is unmatched. Cray end-to-end storage system architectures are compute integrated and performance optimized end to end, from the clients to the storage.

Cray's holistic systems expertise helps solve challenges such as:

- Characterizing and profiling workloads — and deep understanding of how to optimize performance for both file per process and shared file workloads
- Matching the required I/O to the desired workload and workflow on a per-customer basis
- Deploying Lustre at scale: We understand the complexity associated with configuring and tuning all aspects of Lustre storage — from clients to network to disk systems — which includes:
 - Client testing and deployment
 - Networking complexity using IB
 - Storage configuration complexity
- Lustre leadership: As a founder and promoter of OpenSFS, and liaison for large-scale customer requirements, Cray maintains a leadership position delivering Lustre at scale

How does this benefit the customer? Cray ensures:

- Applications perform as expected and jobs finish in predictable timeframes
- Support for a wider range of applications through performance tuning
- Increased user and admin productivity; less time spent configuring and troubleshooting storage

Protect Your High-Value Data

As file systems grow, so too does the need to protect your data due to drive failures, slow rebuild times and other unplanned outages. In most large-scale Lustre systems, there is no easy and efficient way to manage data to and from Lustre to nearline or archive media.

Cray Sonexion includes an advanced declustered parity data protection scheme called Grid RAID. Grid RAID improves data protection and accelerates rebuilds by up to three and a half times over conventional hardware RAID such as RAID6.

Manageability, Diagnostics and Support

The Cray Sonexion Storage Manager (CSSM) software application simplifies the end-to-end experience of deploying, managing and operating a large-scale Lustre solution. The CSSM offers system administrators an intuitive interface and alternative command-line set of tools to monitor and optimize the entire storage system.

CSSM provides status and control of all system components, including storage hardware, RAID, operating system and the Lustre file system in an integrated administrator interface. A web client hosted on one of the controller modules in the MMU interfaces with all distributed system manager component services.

CSSM also integrates a comprehensive set of community-developed tools that collect, index and analyze fast-moving data to help administrators keep the system stable and balanced. The CSSM is tightly integrated into the system stack — from storage and embedded server modules to the Lustre file system and the entire storage cluster — enabling rapid, accurate monitoring and diagnosis down to the component level. Systemwide software and firmware upgrades are executed through a simple single interface in the CSSM system, removing the complexity and risks of traditional large Lustre implementation.

For remote support, Cray includes the System Snapshot Analyzer (SSA), designed from the ground up to deliver improved service to customers. The SSA utilizes a call-home capability to proactively monitor for health, state and configuration changes, and provide faster response during a reported issue. SSA automates collection and reporting of support information. It can operate in a site-private stand-alone mode, or be enabled to call home to Cray support. This enables Cray to diagnose and respond to problems rapidly — and gives customers an easy, nondisruptive solution for remote support.


Reduces Total Cost of Ownership

When measuring all that goes into a large-scale high-performance storage solution (planning, designing, deploying and operating), Cray's end-to-end solution, built on the Sonexion system, reduces the TCO of petascale storage by up to 25 percent compared to external server- and SAN-based Lustre.



Cray® Sonexion® 3000 Specifications

Rack	Height	42U
	Width	600 mm
	Depth	1200 mm
	Data Switches	Dual 36-port InfiniBand EDR switches standard
	Management Switches	Dual 24-port Gigabit Ethernet switches standard Dual 48-port Gigabit Ethernet switches optional
	Standard Cooling	Passive
	Water-Cooled Option	Rear-door heat-exchange unit
	Full Rack Weight (Standard Air-Cooled)	1,138.9 kg (2,510.8 lbs)
	Full Rack Weight (Water-Cooled Door)	34.5 kg (76 lbs) additional
Metadata Management Unit	Metadata Controller Height	2U, high-availability server pair for metadata management
	2U24 Metadata Disk Enclosure Height	2U24 drive enclosure
System Management Unit	Metadata Controller Height	2U, high-availability server pair for system management
	2U84 Metadata Disk Enclosure Height	2U24 drive enclosure
Scalable Storage Unit (SSU)	Base/Expansion Unit Height	5U
	Base/Expansion Unit Data Drive Types	84 drive slots
	Base/Expansion Unit Data Drive Types	82 x 4 TB, 6 TB or 8 TB 7.2K RPM SAS 2 x SSDs
	IOR Read/Write Bandwidth	9-14 GB/s sustainable (InfiniBand) depending on drive type 11-16 GB/s peak (InfiniBand) depending on drive type
Power Consumption	Rack with Switches	<16 kilowatts
Heat Dissipation	Rack	<55,000 BTU
System Availability	Hot Swappable	Disk drives, power supply units, fans, power cooling modules, SBB controller modules
Software & Support Information	Software	CSSM, Linux® and Lustre® included, 1 year renewable
	Hardware	1 year renewable
General System Environmental Specifications	Altitude and Temperatures	-30 to 3,048m (-100 to 10,000ft)
		5-35° C
		De-rated by 1° C/300m above 900m below the specified maximum temperature
		20% to 80% non-condensing
	High Power (Preferred U.S.)	2 input PDUs per rack (2 total inputs per rack) Volts: 208V AC Type: 3 phase AC Amps: 60A rated Connector: 2 IEC60309 60A



Cray Inc. 901 Fifth Avenue, Suite 1000 Seattle, WA 98164 Tel: 206.701.2000 Fax: 206.701.2500 www.cray.com

© 2016 Cray Inc. All rights reserved. Specifications are subject to change without notice. Cray, the Cray logo and Sonexion are registered trademarks of Cray Inc. All other trademarks mentioned herein are the properties of their respective owners. 20160606EMS