

## Win by Trading Smarter with the Cray® Sonexion® Parallel File System

### Advantages of Parallel File Systems

Parallel file systems support multiple storage servers working together in parallel to increase storage throughput and capacity. Compute nodes, accessing a parallel file system architecture, can read and write data from multiple storage servers and devices in parallel, vastly improving performance over a traditional network attached storage (NAS) solution.

### Parallel File System vs. NAS

Parallel file system versus NAS performance ingesting 7 TB of historical trading data:

NAS	Parallel File System
1 model (7 TB) loaded in 5.56 hours	32 models (244 TB) loaded in 2:05 minutes



Hedge funds and quant trading firms compete to a large extent based on the quality of their analytics. Success is determined by two things: strategy innovation and strategy confidence.

Strategy innovation comes from understanding how events shape market movements. In analytic terms, this knowledge comes from massive unstructured data such as news feeds or wellhead data combined with massive market data. Strategy confidence comes from backtesting against years of fine-grained market data spanning different economic cycles and market conditions.

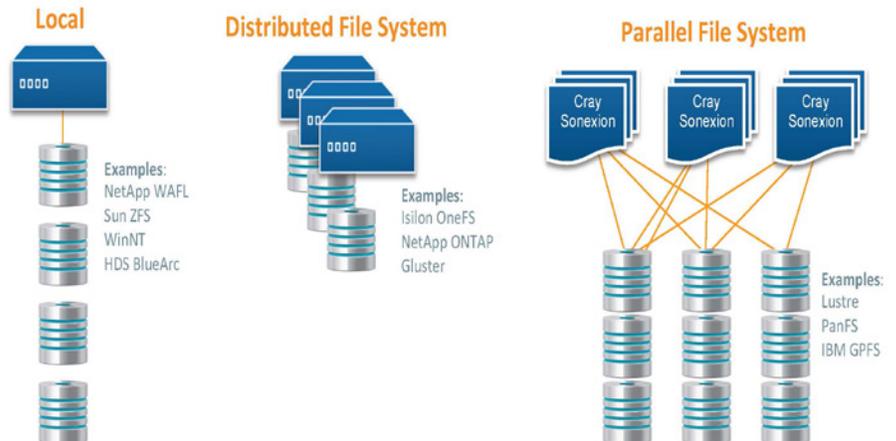
Quants model to the best of their analytic infrastructure's ability. If they can't get a good model turnaround for 10 years of data, then maybe one will do. You get a result, but can you be truly confident in the strategy? When you can't model fast enough to make a difference, it sounds like a compute problem. But it's not. It's almost certainly I/O — and it's solvable.

### Problem: Network Attached Storage Constrained by I/O

Traditional storage platforms are optimized for interactive and transactional processing, operating from cache whenever possible. When sizes of accessed data exceed cache sizes, performance drops sharply as I/O has to go directly to and from mechanical disk drives. Magnifying the issue, these accesses are serialized, which creates a bottleneck as various users have to wait their turn for limited resources.

### About Cray

Cray provides highly advanced systems and solutions that help you solve your most difficult computing, storage and data analytics challenges. The company's comprehensive portfolio includes expertly optimized cluster systems, extremely scalable, powerful supercomputers, advanced storage systems, and high performance data analytics and discovery platforms. Founded in 1972, Cray has focused exclusively on developing, building and supporting supercomputing technologies for over 40 years.



**Cray Inc.**  
 901 Fifth Avenue, Suite 1000  
 Seattle, WA 98164  
 Tel: 206.701.2000  
 Fax: 206.701.2500  
[www.cray.com](http://www.cray.com)

**NAS VS. PARALLEL:** Parallel file systems (right) allow all clients to access, read and write all files at the same time. In comparison, NAS (local and distributed file systems) operates in serial, limiting files to a single storage node and allowing access from only one network interface at a time.

## **Solution: Trade Smarter with the Cray® Sonexion® Lustre® Storage System**

The Cray Sonexion scale-out Lustre storage system enables you to trade smarter by eliminating I/O bottlenecks. It operates in parallel, allowing quants to test hundreds of terabytes of data in minutes. A Sonexion file system can easily scale I/O throughput from 7.5 GB/s to 1.7 TB/s.

Experts in parallel storage solutions, Cray has been developing high performance storage for over 40 years. And Cray systems have been deployed, tested and proven over and over again in the most demanding industries. In fact, Cray Sonexion scale-out Lustre storage holds the current record for the world's fastest production file system (the National Center for Supercomputing Applications' "Blue Waters" system).

You realize the benefit of this unmatched experience in the customized software, rigorous testing and expert support that means your storage solution is simple to manage, easily scales performance and capacity, and protects your valuable data. Combined with Cray's transparent service models and no hidden licensing fee policy, all you have to think about is maximizing your returns.

### **Simplify**

Attach Cray Sonexion scalable storage to any x86 Linux® cluster including Hadoop® clusters and get to work immediately. The preconfigured modular architecture means you deploy quickly. Additionally, its compact form factor reduces both the number of physical management points (drives, cables and other components) and your total hardware footprint by up to 50 percent. Less time spent managing means more resources for working.

The Sonexion system's simple design also translates into cost savings. Fewer drives and components reduce your capital costs as capacity grows. And configure only the combination of performance, capacity and feature functions you need today, with the knowledge that you can add or change easily as your needs shift. From a single rack to an entire file system, configurations can be performance- or capacity-optimized, or balanced for sheer scalability.

### **Scale**

Sonexion storage eliminates I/O bottleneck issues. It scales I/O performance incrementally from 7.5 GB/s to 45 GB/s in a single rack and to over 1.7 TB/s in a single file system. On the capacity side, Sonexion storage scales to over 2 PB in a single rack with up to 100 PB of usable capacity in a file system.

That same expert design that gives you scalable performance also ensures that performance is reliable, available and stable. The modular, redundant architecture of the Sonexion system and the exhaustive factory testing Cray conducts under demanding test conditions means your storage solution will perform as expected.

### **Protect**

Cray Sonexion offers higher levels of data protection than traditional storage. GridRAID provides software-based declustered parity to improve data protection and accelerate rebuild times by up to 3.5 times over traditional hardware RAID.