

Gordon supercomputer in addition to its *Trestles* and *Triton* high-performance computer systems.

"/>

UC San Diego

UC San Diego News Center

June 05, 2012 | By Jan Zverina

SDSC Supercharges its ‘Data Oasis’ Storage System

Sustained Speeds of 100 GB/s Needed to Support SDSC’s ‘Big Data’ Initiatives

The San Diego Supercomputer Center (SDSC) at the University of California, San Diego has completed the deployment of its Lustre-based *Data Oasis* parallel file system, with four petabytes (PB) of capacity and 100 gigabytes per second (GB/s), to handle the data-intensive needs of the center’s new *Gordon* supercomputer in addition to its *Trestles* and *Triton* high-performance computer systems.

Using the I/O power of *Gordon*, *Trestles*, and *Triton*, sustained transfer rates of 100 GB/s have been measured, making *Data Oasis* one of the fastest parallel file systems in the academic community. The sustained speeds mean researchers could retrieve or store 64 terabytes (TB) of data – the equivalent of *Gordon*’s entire DRAM memory – in about 10 minutes, significantly reducing research times needed for retrieving, analyzing, storing, or sharing extremely large datasets.

“We view *Data Oasis* as a solution for coping with the data deluge going on in the scientific community, by providing a high-performance, scalable storage system that many of today’s researchers need,” said SDSC’s director Michael Norman. “We are entering the era of data-intensive computing, and that means positioning SDSC as a leading resource in the management of ‘big data’ challenges that are pervasive throughout genomics, biological and environmental research, astrophysics, Internet research, and business informatics, just to name a few.”

The term ‘big data’, as described by the National Science Foundation, refers to “large, diverse, complex, longitudinal, and/or distributed datasets generated from instruments, sensors, Internet transactions, email, video, click streams, and/or all other digital sources available today and in the future.” Many of those datasets are so voluminous that most conventional computers and software cannot effectively process them.

“Big data is not just about sheer size, it’s also about the speed of moving data where it needs to be, and the integrated infrastructure and software tools to effectively do research using those data,” added Norman. “The capability of *Data Oasis* allows researchers to analyze data at a much faster rate than other systems, which in turn helps extract knowledge and discovery from these datasets.”

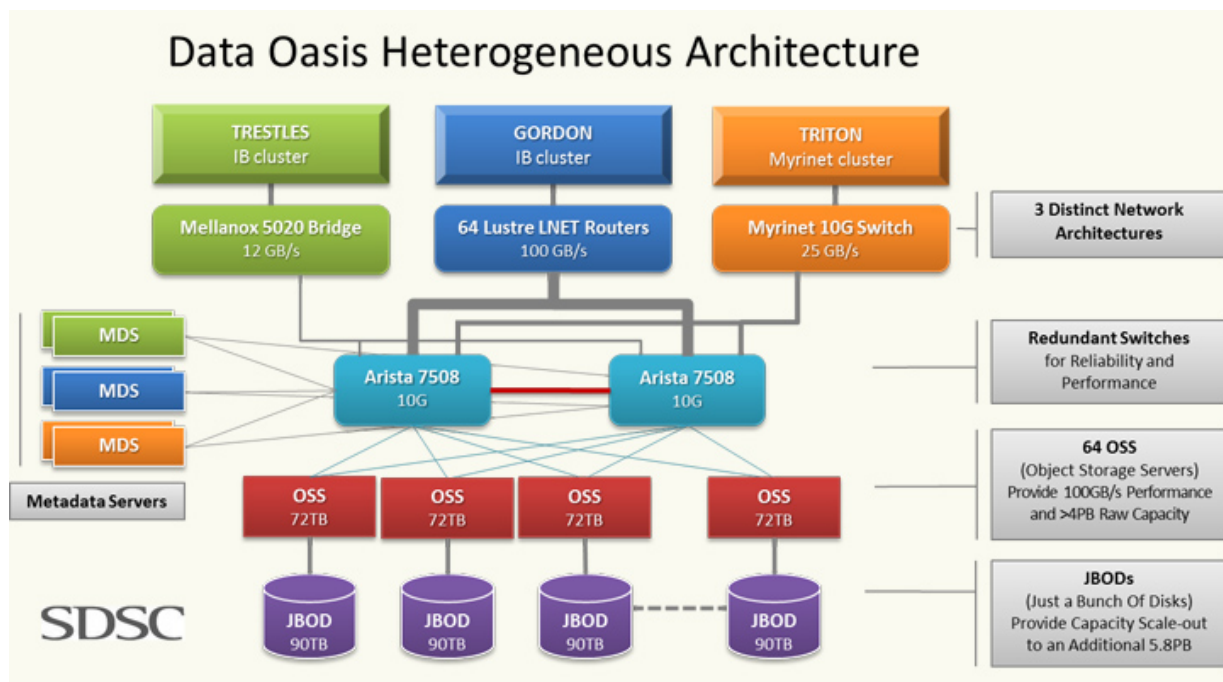


Image courtesy of SDSC/UC San Diego

Driven by *Gordon*

The sustained 100GB/s performance was driven by the requirement to have *Data Oasis* support the data-intensive computing power of *Gordon*, which went online earlier this year following a \$20 million National Science Foundation (NSF) grant. The sustained connectivity speeds were confirmed during *Gordon*'s acceptance testing earlier this year. With recent deployments, *Data Oasis* now serves all three SDSC HPC systems.

“We believe that this is the largest and fastest implementation of an all-Ethernet Lustre storage system,” said Phil Papadopoulos, SDSC’s chief technical officer who is responsible for the center’s data storage systems. “The three major client clusters use different bridging technologies to connect to *Data Oasis*: *Triton* uses a Myrinet-to-10 gigabit ethernet (GbE) bridge (320 gigabits per second), *Trestles* uses an Infiniband-to-Ethernet bridge (240Gb/s), and *Gordon* uses its I/O nodes as Lustre routers for more than one terabit-per-second of network ‘pipe’ to storage.”

Papadopoulos led the overall *Data Oasis* design, but turned to several specialists to provide selected hardware and to fine-tune the system’s software for optimal performance. [Aeon Computing](#) provided the 64 storage building blocks which constitute the system’s Object Storage Servers (OSS’s). Each of these is an I/O powerhouse in their own right. With dual-core Westmere processors, 36 high-speed SAS drives, and two dual-port 10GbE network cards, each OSS delivers sustained rates of over 2GB/s to remote clients. *Data Oasis*’ capacity and bandwidth are expandable with additional OSS’s, and at commodity pricing levels.

There also were some unique design elements of the storage servers to make management easier while improving overall I/O performance. “SDSC’s requirements for high performance, affordability, and maintainability really pushed the envelope,” according to Jeff Johnson, co-founder of Aeon Computing. “By working closely with SDSC’s engineers and systems staff, we were able to deliver a solution that meets the rigorous demands of data-intensive computing.”

Data Oasis’ backbone network architecture uses a pair of large [Arista 7508](#) 10Gb/s Ethernet switches for dual-path reliability and performance. These form the extreme performance network hub of SDSC’s parallel, network file system, and cloud based storage with more than 450 active 10GbE connections, and capacity for 768 in total. “The choice of 10GbE enables us to inter-operate with not only the three fundamentally different HPC interconnect architectures of *Triton*, *Trestles*, and *Gordon*, but also with high-speed campus production and research networks,” added SDSC’s Papadopoulos.

[Whamcloud](#) provided SDSC with consulting expertise in building SDSC’s *Data Oasis* Lustre environment, including hardware, networking, interfaces to various HPC systems, and software installation. In addition, Whamcloud is assisting in the preparation and installation of a ZFS-based Lustre file system, a “Sequoia-like” file system that combines restructured LNET code with the ZFS code base developed on Lawrence Livermore National Laboratory’s Sequoia system. The results will be provided to the Lustre community. Whamcloud will also share

expertise of Lustre SSD projects in order to grow SDSC's knowledge of how best to deploy SSDs/flash in a Lustre file system. *Gordon* will in part serve as a test bed for the Lustre community.

"We're excited to be working with SDSC to provide a Lustre knowledge-source as part of our technical collaboration," said Brent Gorda, CEO of Whamcloud. "SDSC is an acknowledged leader in big data, and this collaboration will directly benefit Lustre users."

"Completion of *Data Oasis* marks a significant accomplishment for the entire SDSC team and the expertise of our partners including Aeon and Whamcloud," said William Young, SDSC's lead for High-Performance Systems who oversaw the deployment of the production storage system. "Together we solved the challenges of interoperability with three different HPC systems while achieving the speeds, reliability, and flexibility needed."

SDSC's *Gordon* and *Trestles* systems and their storage systems are available for use to any researcher or educator at a U.S.-based institution and not-for-profit research through the NSF's Extreme Science and Engineering Discovery Environment, or XSEDE program. Industry-based research allocations and storage are also available. Industry researchers interested in using SDSC's resources or expertise, including *Triton*, should contact Ron Hawkins at rhawkins@sdsc.edu or 858 534-5045.

MEDIA CONTACT

Warren R. Froelich, 858 822-3622, froelich@sdsc.edu

UC San Diego's Studio Ten 300 offers radio and television connections for media interviews with our faculty, which can be coordinated via studio@ucsd.edu. To connect with a UC San Diego faculty expert on relevant issues and trending news stories, visit <https://ucsdnews.ucsd.edu/media-resources/faculty-experts>.