GIGAIO | SDSC SAN DIEGO SUPERCOMPUTER CENTER

Low Latency Memory Fabric Powers Next Generation AI and HPC Prototype National Research Platform

CASE STUDY

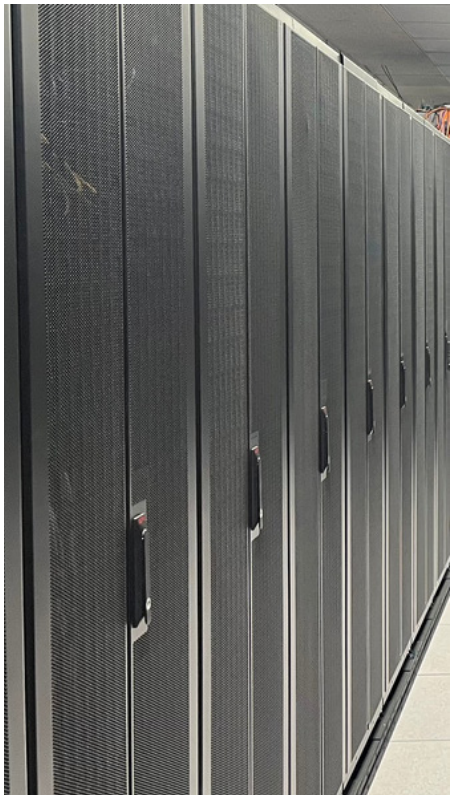

Bringing Composability to a Bold New National Research Platform

Responding to the National Science Foundation's call for a cyberinfrastructure ecosystem that meets the needs of today's data-intensive science, the San Diego Supercomputer Center (SDSC) located at UC San Diego has deployed the Prototype National Research Platform (PNRP), an innovative, all-in-one system anchored around the low-latency GigaIO Fabrex™ Memory Fabric platform.

PNRP brings together several cutting-edge technologies to meet the demands of next-generation science and engineering research. GigaIO's low-latency, high performance FabreX interconnect technology dynamically brings these elements to bear, significantly increasing their utilization over conventional networking — accelerating research, while significantly lowering costs. SDSC's system consists of 64 NVIDIA FP64 GPUs and 32 Xilinx Alveo FPGAs, connected by GigaIO FabreX with NVMe storage and x86 compute servers.

Driving Transformative Research

FabreX was selected for the PNRP because of its innovative ability to enable an all-in-one cyberinfrastructure ecosystem of computing resources, edge computing devices, and other instruments. This PNRP will expedite science and enable transformative discoveries in areas as diverse as public health, high-energy physics, and wildfire response.

"Our research requires that we aggregate disparate computational elements into highly usable and reconfigurable systems. FabreX makes it possible to dynamically bring these elements together in a very low-latency, high-performance interconnect while allowing for distinct, non-interfering workflows to co-exist on the same infrastructure."

— Dr. Frank Würthwein
Director, San Diego Supercomputer Center

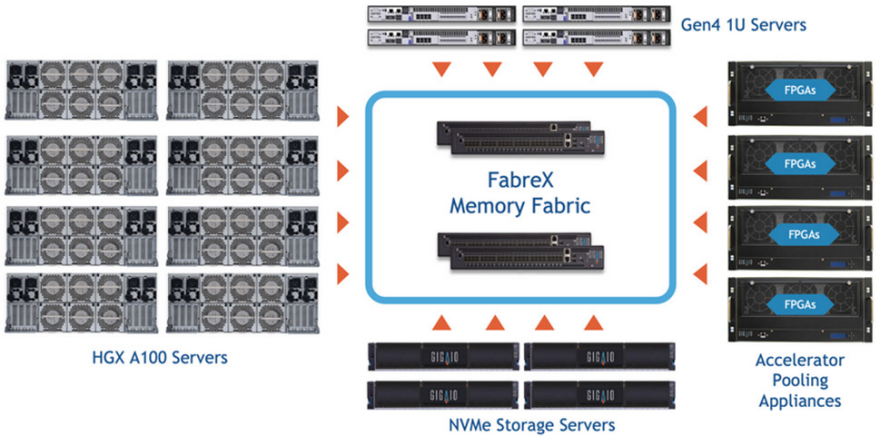

SDSC GigaIO Composable Disaggregated System Configuration

Accelerating Science with GigaIO's Composable Disaggregated Server Solution

SDSC's GigaIO system supports pathfinder research in astrophysics, molecular dynamics, health sciences microscopy, and imaging, including building multi-scale models of the brain and of the molecular components that make brain cells the most energy-efficient information mediators known. The GigaIO system will support time-critical research on the science of wildfires by providing urgent computing resources to support fire-perimeter predictions that can be used on the ground by fire managers to identify where to most effectively deploy firefighting resources in order to achieve fire containment.

GigaIO makes impossible configurations feasible for scientific experimentation by disaggregating and composing the various components of SDSC's mix of graphics processing units (GPUs) and field programmable gate array (FPGA) units with memory and storage, all connected with its fully integrated, extremely low-latency fabric. With the ability to connect resources across entire racks with PCIe, GigaIO's disaggregated architecture democratizes access to AI and HPC resources by enabling better access and utilization of the expensive computing resources previously trapped and stranded inside a static networked server infrastructure.

Accessing the power of GigaIO is as simple as using SDSC's existing job scheduling and cluster management tools — including NVIDIA Bright Cluster Manager, Slurm, and DevOps container automation tools like Chef and Puppet. GigaIO systems include NVIDIA BCM software to compose compute infrastructures that can be configured and scaled in minutes to coordinate complete dynamic infrastructure management, including schedulers and container management.

A Research Platform That Frees the GPUs to Achieve More Science for Less Money

FabreX unlocks the power of GPUs, increasing their utilization and efficiency, in order to run more workloads faster at a lower cost. What's more, there are no problematic NIC or OFED driver bugs or incompatibilities common to conventional HPC networking to stay on top of, and there are no more supply chain delays. And down the road, compute and accelerators can be added at the component level that plug-n-play with the existing environment, so every primary subsystem can operate on its own upgrade cycle. The total cost of the GigaIO platform is optimized over its entire lifecycle with modular upgrades and much higher utilization of all resources.

About SDSC

The San Diego Supercomputer Center is a leader and pioneer in high-performance and data-intensive computing, providing cyberinfrastructure resources, services, and expertise to the national research community, academia, and industry. Located on the UC San Diego campus and serving the University of California system and beyond, SDSC supports hundreds of multidisciplinary programs spanning a wide variety of domains, from astrophysics and earth sciences to disease research and drug discovery.