LITE-ON Storage is now Solid State Storage Technology Corporation

High Performance Computing and Choosing the Right SSD

What is High Performance Computing?

High Performance Computing, or HPC, is the aggregation of computing power from multiple sources with the goal of delivering higher performance in order to solve complex problems. These tasks are compute-

intensive and traditionally handled in the realm of the expensive, elaborate supercomputer; HPC instead offers a flexible, cost-effective alternative. Such a configuration can also be leveraged for the cloud, making it more relevant than ever. The building blocks are found within existing hardware – a network of GPUs, CPUs, and storage.

Although demand for this type of problem-solving was once relatively niche, it has become much larger with applications in fields such as chemistry, biochemistry, physics, climate modeling, geology, prototyping, logistics, linguistics, and much more. The ability to utilize a wider range of hardware in a modular way, balancing for performance, efficiency, and cost, means that businesses and academic establishments now have other options for serious computational work. The inherent scalability means even smaller organizations can produce multiplicative benefits for their designs.

How HPC Works

HPC works by combining the power of multiple machines into an array, allowing them to work together on complicated problems. Each resource is referred to as a node and generally, the array is built in parallel where data is available to any node at any given time. Data is regularly updated and depending on the problem there may be large, small, or both types of data requests, and the array must be capable of handling these with reasonable and consistent latency. This array must also scale for increasing data demands and amount of nodes.

Many of the complex problems HPC can manage involve huge amounts of data, in the petabyte or exabyte range, often with the need for high bandwidth and especially sustained reads. This means managing a lot of power but also requires the presence of power loss protection (PLP) and error correction code (ECC). The overall structure is distributed – much as is seen with the cloud, traditional distributed computing, etc. – with shared data and metadata. The resources in use are typically for processing, that is CPUs and GPUs, all applied to work towards a shared goal.

Choosing the Right SSD for HPC

Data, and therefore storage, is the backbone of the HPC array. Beyond the need for redundancy and error correction, power efficiency remains a strong concern especially as the array scales upwards. These facets tie into the total cost of ownership (TCO) which makes SSDs an easy choice over hard drives. This is especially true as only non-volatile memory has the type of performance needed for this type of configuration – reliably low latency coupled with high throughput. As the amount of nodes accessing the data at any given time can be quite high, having a device capable of many IOPS is also critical.

HPC arrays may have both small and large I/O requests which also require the flexibility of a solid state drive. Typically, the most important, hottest data will be closest to the compute nodes – the nodes working on computation – while colder data can remain in a separate storage array. This means the highest-performing SSDs should be closest to the compute nodes in that configuration while less-performant options are useful for long-term storage. Drives of any sort must be scalable and ready to distribute large amounts of data as needed.

Conclusion

High Performance Computing is expanding at a very rapid pace, bringing immense computational power to even relatively small organizations. This creates more efficient workflows and better designs through solutions to traditionally complex problems. The HPC array offers a modular approach with hardware,

leveraging disparate nodes together without the need for a high-maintenance supercomputer. Storage remains the spine of the array as data is all-important and must be shared broadly, quickly, and reliably between nodes with regular updates. Solid state drives and their support of high-performance protocols offer a solution that HDDs cannot match for HPC.

Our SSD Solutions

PCIe™ - Our ED1 Series is a powerful, high performance SSD made for edge storage applications. It comes in M.2 and U.2 form factors.

SATA - Our ER2 SSD Series delivers affordability and performance with superior random read/write speeds of up to 90,000/45,000 IOPS. It comes in M.2 and 2.5" form factors.

Specifications and features are subject to change without prior notice. Images are samples only, not actual products. Please check with a Solid State Storage Technology Corp. representative for details.

ABOUT US

A subsidiary of KIOXIA Corporation, **Solid State Storage Technology Corporation** is a global leader in the design, development, and manufacturing of digital storage solutions. We offer a comprehensive lineup of high-performance customizable SSDs for the Enterprise, Industrial, and Business Client markets. With various form factors and interfaces, our SSD solutions help businesses simplify their storage infrastructures accelerating variable workloads, improving efficiency, and reducing total cost of ownership.

Learn more at www.ssstc.com