

CRAY DATAWARP APPLICATIONS I/O ACCELERATOR

The Cray® DataWarp® applications I/O accelerator delivers application-ready storage I/O to high-velocity, data-driven workflows. Through a balanced and cohesive I/O acceleration tier engineered for Cray® XC™ supercomputers, DataWarp software optimizes performance, maximizes efficiency, and reduces TCO.

THE APPLICATIONS I/O CHALLENGE

As data-intensive applications demand more from compute infrastructures, the gap widens between computational resources and disk-based storage performance. Disk-based storage can't keep up.

Here's the upshot. If you rely on conventional disk and file system technologies, delivering sustained performance on a range of applications is becoming increasingly costly and impractical.

BALANCING COMPUTE, MEMORY AND STORAGE

Meeting compute and storage needs requires a balanced compute-memory-storage system architecture with a cache tier featuring SSD and in-memory flash. To be effective, the tier should accommodate "bursty" application I/O patterns through a combination of high-bandwidth, low-latency I/O forwarding to deliver thousands to millions of aggregate IOPS to applications.

HOW DATAWARP SOFTWARE DOES IT

Cray DataWarp applications I/O accelerator delivers a balanced, cohesive system architecture from compute to storage. DataWarp software allocates storage dynamically in either private (dedicated) or shared modes. It can provide storage performance quality-of-service to individual applications based on the site's configuration, the user's request, and

the currently available resources. The benefit is simple: streamline workflows that "just work." They process faster and access storage on demand.

The DataWarp accelerator technology can be used as a storage cache for parallel file systems such as Lustre®, GPFS™, and PanFS™. In these scenarios, the applications I/O accelerator capabilities drive up the overall utilization of the parallel file system by buffering performance across the new tier. As a result, it boosts pure performance by delivering the right storage performance to applications. DataWarp improves overall application performance by decoupling application I/O from the corresponding parallel file system I/O.

Taking advantage of the quicker price evolution of nonvolatile flash, we designed I/O blades with fast SSD close to the powerful XC compute blade nodes. By doing so, we introduce a buffer platform tier that complements the balance of existing on-board memory and external disk storage. This adaptive architecture addresses the throughput/capacity gap, enabling a caching platform that can support burst buffer capabilities and ensure peak performance of a parallel file system. With PCIe interfaces, the XC series I/O blades can be configured with a variety of flash SSD options, giving each job an optimal balance of speed, performance, and cost.



MORE PERFORMANCE, SAME COST

The DataWarp accelerator provides 5 times more adaptive, application-ready I/O performance than a disk-based parallel file system at the same cost. It reduces cost per IOPS up to 4 times by scaling from 70,000 to 40 million aggregate IOPS in a single supercomputing system.

You get high bandwidth with virtually no impact on other I/O executing in the system. This ability ensures quality of service and sustained bandwidth to specific applications — and a performance improvement of up to 10 times over PFS-based scratch systems. Effectively, rebalancing compute-storage systems with optimized levels of memory and flash SSD reduces some of the disk arm seek thrashing that can cause performance losses.

ADAPTIVE FLEXIBILITY, EASE OF USE

While increasing compute and storage utilization for both peak and sustained workloads, the DataWarp applications I/O accelerator reduces:

- **Cost per IOPS.** New SSD-based tiers deliver two to five times the IOPS of disk-based systems.
- **Cost of bandwidth.** The new SSD-based tiers deliver 10 times the bandwidth of disk-based systems.
- **Overall system TCO.**

You get a new level of flexibility to allocate the right type and amount of data storage and I/O movement per job or node. Storage is dynamically allocated to maximize compute and storage utilization across the entire system — adaptive flexibility means you can put the best resources where you need them the most. For example, you can put a scratch file system on every node, or a burst capability close to compute nodes for faster checkpoint restart.

Leveraging a worldwide installed base of leadership-caliber systems and Cray's deep expertise in

HPC-optimized, high-density design and integrated hardware/software packaging, the I/O accelerator delivers top performance with fewer cabinets, cables, and disks. Efficiency for a large, capability-class machine can be increased by 70 to 90 percent, saving millions in compute costs. Storage is automatically and dynamically provisioned, driving up its utilization and driving down its operating cost.

The DataWarp software stack includes HPC tools that streamline application processing in a range of environments. It provides low-level (kernel mode) POSIX-compliant I/O forwarding techniques. By cutting down on overhead and improving workload throughput by 10 to 15 percent, users reclaim compute time and improve overall performance. Wall clock execution time is optimized, meaning more runs per day and iterations per compute session, faster time to results and higher productivity.

DATAWARP DELIVERS:

- Flash SSD I/O blades and Aries interconnect
- Data virtualization software, management software, and I/O forwarding mechanisms
- Ability to scale up a data pool to customer requirements
- Performance gains in I/O, total system compute, bandwidth, and IOPS
- Reduced storage system costs
- Improved TCO
- System balance across memory-flash-disk
- Ease of use advances