CRAY XC SERIES INTERCONNECT AND NETWORK TECHNOLOGY

The HPC-optimized Aries™ interconnect and Dragonfly network topology integrated into Cray® XC™ supercomputers drive breakthrough advantage in performance and scalability.

COST-EFFICIENT, SCALABLE BANDWIDTH

Cray XC supercomputers run the world's most challenging workloads. And fundamental to giving them the performance and scalability to do it is the Aries interconnect.

This innovative intercommunications technology, implemented with a high-bandwidth, low-diameter network topology called Dragonfly, provides substantial enhancements on all of the network performance metrics for HPC: bandwidth, latency, message rate, and more. Delivering global bandwidth scalability at reasonable cost across a distributed memory system, this network gives programmers global access to all the memory of parallel applications and supports the most demanding global communication patterns.

HOW THE ARIES INTERCONNECT ACHIEVES A PERFORMANCE ADVANTAGE

The Aries interconnect is a Cray-developed ASIC designed with a holistic system approach. It integrates four NICs with two router tiles for the interconnect. The PCIe-3 compute node connections are flexible, supporting different processor and accelerator types and configurations.

The Dragonfly topology is bandwidth-rich, scalable, and upgradable. It addresses the toughest network patterns — from nearest neighbor-type communication to irregular to all-to-all communications.

Dragonfly groups of 384 compute nodes are constructed from 96 blades over six chassis in a pair of cabinets. An electrical all-to-all network — constructed from backplane links — connects the blades in each chassis. A second all-to-all network — constructed using electrical cables — connects the six chassis in each group. The group-level interconnect is common to all large XC systems. The full system is constructed to connecting Dragonfly groups with an optical all-to-all network.

With this network technology, global bandwidth can be tuned by varying the number of optical cables in cabinet group connections. Dragonfly helps run workloads at very high system utilization efficiency as jub placement is insensitive. Jobs execute from perfectly grouped placement to worst-case placement with very little runtime variance.

Additionally, we have developed local shared-memory optimizations in the XC software stack libraries for MPI collectives. These optimizations significantly improve performance on high core count compute nodes. Through implementing dynamic allocations (connectionless transfers) rather than static allocations, you get reduced MPI

memory footprint and overhead and improved performance.

With fewer global traffic hops, no external director switches, and half the optical cables required for a comparbale fat-tree implementation, the Aries/ Dragonfly technology offers performance and cost advantages over a commodity-based equivalent.

READY FOR THE MOST DEMANDING APPLICATIONS

The Aries interconnect and Dragonfly network integrated into Cray XC supercomputers provide cost-effective, scalable global bandwidth. This architecture — together with the Cray-optimized software stack, programming environment, and advanced reliability features — gives you a system that can scale to millions of parallel compute cores.

Overall, the XC system integrated with the Aries interconnect delivers consistent, scalable, and sustained computing performance so you can deliver results.

Network rank-1 backplane, rank-2 electrical cables (black), rank-3 optical cables (green)