XT5m Supercomputer

Making Petaflops Performance Affordable

WHY THE TIME IS RIGHT FOR CRAY XT5M SUPERCOMPUTER

As supercomputing evolves, it is not enough to simply offer users more powerful processors. Users need new ways to harness that power. While it is physically possible to link thousands of processors together in rudimentary clusters, without the ability to make them work in parallel, manage them easily, or develop and run software applications efficiently, the effort is often a waste of time.

The need to harness supercomputing power is especially important as the capabilities and complexities inherent in high-end supercomputing become available in midrange systems. When midrange supercomputers reach high-end performance levels, unique scalability challenges arise. Cray is positioned to provide solutions which improve the scalability and usability of these computational resources, and maximize the value of such investments.

While Cray is the current high-end supercomputing market leader, with a 35% market share in 2008 for systems priced at $3 million and above (according to market analyst firm IDC), the release of the Cray XT5m supercomputer marks its entry into the midrange market. The "Mighty Mini," as it is affectionately known, is a high-performance system that is both powerful and affordable because it takes advantage of some of the hardware and software advancements of Cray's XT supercomputer family. And Cray knows a thing or two about achieving maximum performance. Cray's technologies form the backbone of the world's largest open science supercomputer, a petascale system currently in use at the U.S. Department of Energy's Oak Ridge National Laboratory. In this white paper, we'll look at those capabilities, the advantages they bring, and the markets and applications this groundbreaking system is already beginning to serve.

HARDWARE ADVANTAGES OF THE CRAY XT5M SUPERCOMPUTER

The specifications of the Cray XT5m supercomputer come close to those of its more powerful big brother, the Cray XT5. The Cray XT5m supercomputer can be configured with anywhere from one to six cabinets, each capable of seven to 12 teraflops of computing power. Each cabinet can accommodate up to 192 CPUs, such as the 64-bit Quad-Core and Six-Core AMD Opteron™ series 2000 processors. Its main memory, memory bandwidth, external I/O interfaces, reliability features (both in hardware and software), operating system, message passing libraries,

compilers, and power and cooling requirements all rival the Cray XT5. The XT5m is cost-optimized with a starting pricing of $500,000, so you get the most for your money.

With the current generation of Quad- and Six-Core AMD Opteron processors, the maximally configured Cray XT5m can boast almost 7,000 processing cores, a level of scalability that Cray has been regularly providing to its high-end customers for many years with industry-leading availability and reliability levels.

In creating a supercomputer that combines the best features of its high-end sibling and the value proposition technical engineers have come to expect from traditional clustered systems, Cray also endeavored to create a system that would encompass both initial affordability and a low total cost of ownership through its lifecycle. The Cray XT5m incorporates both leading-edge and traditional technologies to deliver the best of both worlds. These include:

ARCHITECTURE. The Cray XT5m uses a powerful, flexible torus architecture to support scalable multiprocessing, incorporating both compute and service nodes within the system. The compute nodes (comprised of two AMD Opteron processors) handle message-passing applications, and are linked with the system's high-performance interconnect. The service nodes are responsible for system and I/O connectivity; users compile and launch data processing work using these as login nodes. This design helps reduce (and may even eliminate) the scheduling complexities and asymmetric performance problems that large symmetric multiprocessing clusters can experience.

PROCESSORS. In terms of standardized components, the Cray XT5m takes advantage of AMD's HyperTransport technology by incorporating AMD Opteron processors. HyperTransport combines high-speed and low-latency protocols to increase the speed with which the AMD Opteron processors communicate with the network and within each node. This technology reduces system bottlenecks and helps processors take advantage of system memory and I/O more efficiently.

INTERCONNECT. The Cray XT5m supercomputer incorporates Cray's SeaStar interconnect technology, forming the backbone of the torus network. SeaStar boosts message-passing capabilities within the supercomputer and distributes I/O traffic to the file system. Cray designed the SeaStar interconnect technology to accommodate both performance and reliability in large-scale, distributed-memory supercomputing systems. It is an extremely scalable network and the basis for the first general-purpose petaflop system ever built.

The Cray XT5m system comes configured with four switch ports per SeaStar network chip configured in a two-dimensional torus. This network topology contrasts with the Cray XT5 product for larger systems, which is configured as a three-dimensional torus. This design helps reduce complexity and cost, thereby improving the total cost of ownership while still providing excellent network bandwidth to scale applications easily to 7,000 cores. The Cray XT5m network offers sustained bandwidth of up to 6 GB per second, with potential peak bandwidth of up to 9.6 GB per second. If the compute needs grow beyond the six-cabinet limit of the Cray XT5m, the system can be upgraded to the Cray XT5, complete with its three-dimensional torus, by adding a mezzanine card replacement, some additional cabling, and other elements.

I/O. Just as important as processing power is the ability to accommodate the data being analyzed. The Cray XT5m supports the same I/O software and subsystems that have been proven at Cray's largest customer sites by implementation of the XT5. The XT5m can configure file systems with the appropriate bandwidth for their specific project needs, and can utilize the Lustre file system technology as well.

UPGRADABILITY. Cray designed the XT5m supercomputer to be upgradeable over time so that customers could acquire it as an entry-level supercomputing system and add power as their needs change and their expertise increases. For instance, it currently uses Quad-Core AMD Opteron processors, but customers can easily upgrade to the more

powerful Six-Core AMD Opteron processors because the processors use the same power and thermal envelopes. In addition, the Cray XT5m can be expanded into a Cray XT5 configuration if customer requirements increase over time. Cray anticipates that customers who invest in the Cray XT5m based on AMD technology will have a viable upgrade path using a single cabinet for at least five years.

ENERGY EFFICIENCY. In these days of rising energy costs, a high-performance system requires the most flexible and efficient cooling available. The Cray XT5m is designed to be either air- or liquid-cooled. Liquid cooling incorporates Cray's EcoPhlex (Phase-change Liquid EXchange) technology, which was designed to work with highly dense systems. Each system cabinet is configured with coils designed to extract all heat before it can reach the datacenter, helping dramatically reduce the need for computer room air conditioning (CRAC) units. The heat exchange units work with chilled water and can handle a wide variety of water temperatures or datacenter conditions.

Overall, the hardware advancements in the Cray XT5m supercomputer represent a highly optimized and balanced system that will run scientific and engineering applications efficiently from hundreds to thousands of processors. It supports these applications with an optimized hardware infrastructure that keeps total cost of ownership and price/performance ratios lower than were previously available in the midrange supercomputing space.

SOFTWARE ADVANTAGES OF THE CRAY XT5M SUPERCOMPUTER

Cray has equipped the Cray XT5m supercomputer with software functionalities that take it well beyond those of simple clusters. Cray has optimized the base Linux operating system, system administration capabilities, and software development tools to provide specific support for scalable systems and applications. These software enhancements include:

SYSTEM ADMINISTRATION. A key facet of harnessing computing power is the ability to monitor uptime. Cray incorporates both its Hardware Supervisory System (HSS) and its System Management Workstation (SMW) software into the Cray XT5m system so that system administrators can easily manage and monitor it.

SOFTWARE DEVELOPMENT AND COMPILERS. Cray supports powerful parallel-programming models for developing and running high-performing applications, including MPI (message passing interface), C, C++, UPC, Fortran, Co-array Fortan, OpenMP, and its own SHMEM library for one-sided communication.

OTHER APPLICATION DEVELOPMENT TOOLS. Cray has partnered with leading HPC tool and middleware providers, including:

- Debuggers
- TotalView
- Allinea
- Workload Management
- Cluster Resources Inc., Moab
- Altair, PBS Pro
- Compilers
- Portland Group
- PathScale

VIRTUALIZATION. Cray is working to enhance portability of applications to the XT5m. A key innovation in this area is Cray's use of the Data Virtualization Service (DVS) technology. DVS allows for innovations like data projection to compute nodes, which makes the nodes appear as if they have their own local disks for applications that require local scratch or swap space.

TARGET MARKETS AND APPLICATIONS

The most exciting aspects of a midrange supercomputer are the opportunities it opens up for organizations across a myriad of industries. They can use the scalability of the XT5m to increase productivity and collaboration, and improve the quality of the science they do. (For more on how one academic institution uses its Cray XT5m as the basis for a strong public sector/private sector partnership, see the case study *Economical*

Supercomputing at the University of Stuttgart.)

The Cray XT5m will also make a strong contribution in the increasingly important area of energy research (already evidenced by the petascale Cray XT5 system in use by the Department of Energy at Oak Ridge). The ability to determine new ways of creating energy from biologically or organically based substances can require extensive capabilities in both parallel processing and scalability.

Increasingly, too, research in the field of climate change often requires the ability to analyze expansive amounts of data points and identify trends and patterns. With the debate continuing over the potential impact of shifting weather conditions, supercomputing systems such as the Cray XT5m can bring meteorological insight to a wider number of researchers.

On the business side, supercomputing at nearly high-end performance levels but mid-range cost can be of great importance to engineering and design companies that face increasingly prominent multidisciplinary challenges. In the automotive industry, for instance, engineers are designing vehicles that incorporate lighter weights for better mileage with stronger materials for safety—a task which can be much easier with nearly high-end supercomputing.

In the manufacturing industry, engineers are striving to create design efficiencies through a combination of intelligent electronics, electromechanical systems, motion control, and robotics. Designing optimized systems that combine multiple engineering disciplines is challenging, but it helps engineers achieve greater efficiencies because they are not forced to "over-design" equipment to accommodate reliability issues.

WHAT THE FUTURE HOLDS

Traditionally, processing power pushed down on the computing pyramid. Today's smartphones are as powerful as the earliest mainframes. But there are still sufficient high-end supercomputing challenges to require—and inspire— innovation. In the near future, these innovations are likely to come in two ways: increased access to petaflops-level power and increased affordability—both of which are possible with the Cray XT5m product.

With standard processors, operating systems, and optimized development tools, the Cray XT5m is a logical platform for the expanded development of scalable applications. And the less time researchers spend developing their own applications, the more time they can spend using them.

In terms of affordability, the Cray XT5m supercomputer based on the Quad-Core AMD Opteron™ processor provides multiple cost-of-ownership advantages. Its low entry-level cost brings supercomputing capabilities to more institutions and can build more industry/academia collaboration. At the same time, its built-in upgradeability for both processor and interconnect technology brings down total cost of ownership and increases return on investment, extending the value of the Cray system for many years.

With petaflops-level performance available to more researchers, the world's scientific and engineering challenges seem less daunting.