

CRAY XC40

Scaling Across the Supercomputer Performance Spectrum

Cray's XC40™ system leverages the combined advantages of next-generation Aries interconnect and Dragonfly network topology, Intel® Xeon® processors, integrated storage solutions and major enhancements to the Cray OS and programming environment. The Cray® XC® series supercomputer is a groundbreaking architecture upgradable to 100 petaflops per system and delivering:

- Sustained and scalable application performance
- Tight HPC optimization and integration
- Production supercomputing
- Investment protection – upgradability by design
- User productivity





The Cray XC40 supercomputer is a massively parallel processing (MPP) architecture focused on producing more capable HPC systems to address a broad range of user communities. The Cray XC series is targeted at scientists, researchers, engineers, analysts and students across the technology, science, industry and academic fields.



Adaptive Supercomputing Architecture

Building on Cray's adaptive supercomputing vision, the Cray XC series integrates extreme performance HPC interconnect capabilities with best-in-class processing technologies to produce a single, scalable architecture.

Understanding that no single processor engine is ideal for every type of user application, the Cray XC series highlights the flexibility of scalar processing, coprocessing and accelerators to build hybrid systems capable of leveraging the strengths of each technology into one "adaptive" HPC environment. Computing applications often need a balance of scalar and parallel operations to execute optimally. Merging x86 multicore benefits with coprocessor or accelerator many-core advantages addresses this need to target the best processing engine type for a specific function. With a robust hardware and software environment empowering customer choice, users can configure their Cray XC40 systems to their own unique requirements to meet their specific goals.

With the Cray XC40 system, the adaptive supercomputing concept applies to building flexibility into hardware and network upgradability, the comprehensive HPC software environment, optimized OS and ISV support, networking, storage and data management as well as power and cooling elements.

Holistic and Integrated Platform

Cray uses holistic and integrated methods to develop the world's most complete and robust HPC systems. This R&D process is a key benefit of Cray supercomputers and ensures that each product is developed, tested and validated against the most demanding real-world HPC applications. Rather than assembling cluster-like components and commodity networks that can degrade or act unpredictably at high node counts, the Cray XC series architecture improves upon the Cray history of "system-centric" development.

With a comprehensive scope of extreme performance interconnects, processing, packaging, cooling, power options, file systems, upgradability, supervisory systems, OS and software development environment, the Cray XC series delivers a quality, reliable HPC solution. Additionally, like all Cray systems, Cray XC40 supercomputers offer the ability to efficiently scale key software applications, easy and proactive future upgradability by design, and a tightly coupled interconnect and software environment.

Regardless of your application requirements, the Cray XC series scales across the performance spectrum — from smaller footprint, lower-density configurations up to the world's largest and highest-performing supercomputers.




Extreme Scalability and Sustained Performance

Cray has an established reputation for regularly running the biggest jobs on the largest numbers of nodes in the HPC industry. With the Cray XC series there's even more focus on solving extreme capability computational challenges. Designed to avoid the limitations of commodity assemblies, the Cray XC40 system scales hardware, networking and software across a broad throughput spectrum to deliver true sustained, real-world production performance. It all means users can now model even bigger datasets and simulate massive systems which previously needed to be partitioned down into numerous smaller-sized modules.

Aries Interconnect and Dragonfly Topology

To provide this breakthrough performance and scalability, Cray XC series supercomputers integrate the HPC-optimized Aries interconnect. This innovative intercommunications technology, implemented with a high-bandwidth, low-diameter network topology called Dragonfly, provides substantial improvements on all the network performance metrics for HPC: bandwidth, latency, message rate and more. Delivering unprecedented global bandwidth scalability at reasonable cost across a distributed memory system, this network provides programmers with global access to all of the memory of parallel applications and supports the most demanding global communication patterns.

The Dragonfly network topology is constructed from a configurable mix of backplane, copper and optical links, providing scalable global bandwidth and avoiding expensive external switches.




Adaptive supercomputing means a modular structure providing a customizable system for users to manage entry costs, and enables easy in-place upgrades for growing bandwidth requirements in the future.

The Aries ASIC provides the network interconnect for the compute nodes on the Cray XC40 system base blades and implements a standard PCI Express Gen3 host interface, empowering connectivity to a wide range of HPC processing compute engines. This universal nature of the Cray XC series open architecture allows the system to be configured with the best available devices today, and then augmented or upgraded in the future with the user's choice of processors/coprocessors utilizing processor daughter cards (PDCs), each with their own independent capabilities and development schedule.

The Cray XC series architecture implements two processor engines per compute node, and has four compute nodes per blade. Compute blades stack in eight pairs (16 to a chassis) and each cabinet can be populated with up to three chassis, culminating in 384 sockets per cabinet. Following the Intel Xeon processor roadmap and assuming 22 cores per processor, it sums to a possible 8,448 cores each, enabling up to 297 teraflops per cabinet and upgradable with the Intel schedules

to advance clock frequency and the number of embedded cores.

The open architecture of the Cray XC40 system empowers users to accelerate performance and reduce power consumption by running hybrid applications leveraging GPUs or coprocessors. Adaptive supercomputing means power-efficient performance and customer choice.



The Cray XC Series and Intel Xeon Processors

The Cray XC series marks the first time Cray is using the industry-leading 64-bit Intel Xeon processor E5 family members in a high-end supercomputer line. With these Intel processors, Cray XC40 systems can scale in excess of 1 million cores.

Custom or ISV Jobs on the Same System — Extreme Scale and Cluster Compatibility

Rather than being boxed in by a restricted system architecture, the Cray XC series provides complete workload flexibility. Based on generations of experience with both environments, Cray has leveraged a single machine to run both highly scalable custom workloads as well as industry-standard ISV jobs via the powerful Cray Linux® Environment (CLE). CLE enables a Cluster Compatibility Mode (CCM) to run out-of-the-box Linux/x86 versions of ISV software without any requirement for porting, recompiling or relinking.

Alternatively, Cray's Extreme Scalability Mode (ESM) can be set to run in a performance-optimized scenario for custom codes. These flexible and optimized operation modes are dynamic and available to the user on an individual job basis. CLE has been optimized to make the most of the advancements in the Aries interconnect and the Dragonfly topology without requiring user tuning. Adaptive supercomputing means supporting different techniques of code execution on the fly.

ROI, Upgradability and Investment Protection

Besides the customizable configuration of the exact machine a user requires, Cray XC40 supercomputer architecture is engineered for easy, flexible upgrades and expansion, a benefit that prolongs its productive lifetime and the user's investment. As

new technology advancements become available, users can take advantage of these next-generation progressions deep into the life cycle before ever considering replacing an HPC system. Adaptive supercomputing means longevity.

Cray XC40 System Resiliency Features

The Aries interconnect is designed to scale to massive HPC systems in which failures are to be expected, but where it is imperative that applications run to successful completion in the presence of errors. Aries uses error correcting code (ECC) to protect major memories and data paths within the device. The ECC combined with the Aries adaptive routing hardware (which spreads data packets over the available lanes that comprise each of the Dragonfly links) provide improved system and applications resiliency. In the event of a lane failure, the adaptive routing hardware will automatically mask it out. The HSS can even automatically

reconfigure to route around the bad links in the event of losing all connectivity between two interconnects.

Additionally, the Cray XC40 system features NodeKARE™ (Node Knowledge and Reconfiguration). If a user's program terminates abnormally, NodeKARE automatically runs diagnostics on all involved compute nodes and removes any unhealthy nodes from the compute pool. Subsequent jobs are allocated only to healthy nodes and run reliably to completion.


Innovative Cooling and Green Systems

Cray continues to advance its HPC cooling efficiency advantages, integrating a combination of vertical liquid coil units per compute cabinet and transverse air flow reused through the system. Fans in blower cabinets can be hot swapped and the system yields "room neutral" air exhaust. Innovative high-performance power and new cooling capabilities support powerful new high-wattage processors. Improve your TCO by reducing the number of chillers and air handlers. Eliminate the need for hot/cold aisles in your datacenter.



The Cray XC40 Supercomputer Delivers

- Sustained and scalable application performance
- Production supercomputing
- Tight HPC optimization and integration
- Investment protection
- Upgradability by design
- User productivity



Cray Provides

- A robust software environment
- Integrated storage solutions
- A broad partner ecosystem

Cray XC Series Targets



Environmental and Earth Sciences

- Climate
- Oceanography
- Weather



Manufacturing and CAE

- Aerospace CFD simulation
- Automotive safety and crash analysis
- Noise, vibration and harshness simulation
- Materials



Energy, Oil and Gas

- Seismic analysis
- Reservoir simulation



Government, Defense and National Security



Life Sciences

- Bioinformatics
- Proteomics
- Structural biology



Financial Services

The HPC Challenge

Higher performance or lower power? More flexibility or lower cost? Increased capabilities or improved greenness? Today's demanding HPC technology requirements are often in conflict with each other. The driving metrics around supercomputing are no longer purely about peak performance and faster clock rates. Now, users need to balance speed and power with other critical system capabilities


such as sustained real-world performance, economized power, cooling efficiency, tight packaging, and a robust partner ecosystem as well as full integration of OS, software and programming environments. Business requirements are equally important and need to address upgradeability and return on investment, operating and capital equipment costs, reliability and resiliency.

PDCs

Processor daughter card options providing the adaptive supercomputing flexibility of customer choice on processing engine technology are outlined in detail in separate product briefs covering CPUs, accelerators and coprocessors.

Software Environment

Cray XC series software environment details are described in a product brief covering the Cray Linux Environment and Cray Programming Environment, including discussions of programming paradigms, compilers, debuggers, science libraries and workload managers, as well as a wide variety of software, middleware and ISV partner offerings.



Data Storage

Cray offers a range of complementary high-performance storage optimized for the XC series line of supercomputers:

- The **Cray® DataWarp™** I/O acceleration option for the XC40 supercomputer utilizes flash storage to speed up storage performance to applications and compute nodes in a variety of scenarios. The Cray DataWarp accelerator reduces the cost of delivering storage performance by up to 5x over hard-drive disk (HDD) based file systems.
- The **Cray Sonexion** scale-out Lustre® system for big data and supercomputing provides everything needed to deploy fast scratch, temporary or primary storage for the XC system. Cray Sonexion storage scales efficiently, reducing the number of physical components by up to 25 percent over competing Lustre offerings — and scales performance in super-linear units, measured in GB/s to over one TB/s in a single file system.
- **Cray Tiered Adaptive Storage (TAS)**, powered by Varsity, is an all-open tiered data management system optimized for big data and supercomputing. Built on open systems and open formats, Cray TAS reduces the cost of managing data and long-term storage repositories, and it provides transparent data migration from high performance storage, such as Lustre, to deep tape archives.

Additional product briefs may be downloaded from Cray's website at www.cray.com/xc40



Cray XC40™ Series Specifications

Processor	64-bit Intel® Xeon® processor E5 family; up to 384 per cabinet
Memory	64-256 GB per node
	Memory bandwidth: up to 153 GB/s per node
Compute Cabinet	Up to 384 sockets per cabinet, upgradeable with processors of varying core counts
	Peak performance: initially up to 297 TF per system cabinet
Interconnect	1 Aries routing and communications ASIC per 4 compute nodes
	48 switch ports per Aries chip (500 GB/s switching capacity per chip)
	Dragonfly interconnect: low latency, high bandwidth topology
System Administration	Cray System Management Workstation (SMW)
	Single-system view for system administration
	System software rollback capability
Reliability Features (Hardware)	Integrated Cray Hardware Supervisory System (HSS)
	Independent, out-of-band management network
	Full ECC protection of all packet traffic in the Aries network
	Redundant power supplies; redundant voltage regulator modules
	Redundant paths to all system RAID
	Hot swap blowers, power supplies and compute blades
Reliability Features (Software)	Integrated pressure and temperature sensors
	HSS system monitors operation of all operating system kernels
	Lustre® file system object storage target failover; Lustre metadata server failover
	Software failover for critical system services including system database, system logger, and batch subsystems
	NodeKARE (Node Knowledge and Reconfiguration)
Operating System	Cray Linux Environment (includes SUSE Linux SLES11, HSS and SMW software)
	Extreme Scalability Mode (ESM) and Cluster Compatibility Mode (CCM)
Compilers, Libraries & Tools	Cray Compiler Environment, Intel Compiler, PGI Compiler, GNU compiler
	Support for the ISO Fortran standard (2008) including parallel programming using coarrays, C/C++ and UPC
	MPI 3.0, Cray SHMEM, other standard MPI libraries using CCM. Cray Apprentice and CrayPAT™ performance tools. Intel Parallel Studio Development Suite (option)
Job Management	PBS Professional job management system
	Moab Adaptive Computing Suite job management system
	SLURM – Simple Linux Unified Resource Manager
External I/O Interface	Infiniband, 40 and 10 Gigabit Ethernet, Fibre Channel (FC) and Ethernet
Disk Storage	Full line of FC, SAS and IB based disk arrays with support for FC and SATA disk drives, Sonexion® data storage system
Parallel File System	Lustre, Data Virtualization Service (DVS) allows support for NFS, external Lustre and other file systems
Power	90 kW per compute cabinet, maximum configuration
	Support for 480 VAC and 400 VAC computer rooms
	6 kW per blower cabinet, 20 AMP at 480 VAC or 16 AMP at 400 VAC (three-phase, ground)
Cooling	Water cooled with forced transverse air flow: 6,900 cfm intake
Dimensions (Cabinets)	H 80.25" x W 35.56" x D 62.00" (compute cabinet)
	H 80.25" x W 18.00" x D 42.00" (blower cabinet)
Weight (Operational)	3,450 lbs. per compute cabinet—liquid cooled, 243 lbs./square foot floor loading 750 lbs. per blower cabinet
Regulatory Compliance	EMC: FCC Part 15 Subpart B, CE Mark, CISPR 22 & 24, ICES-003, C-tick, VCCI
	Safety: IEC 60950-1, TUV SUD America CB Report
	Acoustic: ISO 7779, ISO 9296



Cray Inc. 901 Fifth Avenue, Suite 1000 Seattle, WA 98164 Tel: 206.701.2000 Fax: 206.701.2500 www.cray.com

© 2016 Cray Inc. All rights reserved. Specifications are subject to change without notice. Cray, the Cray logo and Sonexion are registered trademarks, and Cray XC and DataWarp are trademarks, of Cray Inc. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. All other trademarks mentioned herein are the properties of their respective owners. 20160305EMS