

Future Grid: Accelerating Risk Analytics While Reducing TCO

Why It's Time to Reassess Compute Grids

New pressures in today's fast-moving capital markets are forcing banks to rethink their risk management systems architecture. The upshot? These new demands are **driving up the core count needs of banks' compute grids** enormously. Traditional compute clusters can't keep up. So what are the specific demands that are overtaxing banks' existing compute architectures?

For one, where each asset class used to have its own market data, pricing and risk models, the industry is now moving toward integration. Regulators as well as sound business practices expect the enterprise to work collaboratively on shareable models and trading systems from a bank-wide real-time risk view. Next, banks are starting to use dynamic programming languages to allow analysts to compose collaborative applications within a framework that requires a highly efficient compute core. Additionally, more complex scenario modeling — such as CCAR stress tests and Basel 4, as well as the drive to intra-day and near-real-time risking — multiply core count needs.

Combine these factors with today's business reality of low interest rates and exchange-traded derivatives at low margin and there's less money around to address compute requirements. The compute value chain must be viewed in terms of efficiency and capability together.

About Cray

Cray provides systems and solutions that help you solve your most difficult computing, storage and data analytics challenges. The company's comprehensive portfolio includes expertly optimized cluster systems, extremely scalable, powerful supercomputers, advanced storage systems and high-performance data analytics and discovery platforms. Founded in 1972, Cray has focused exclusively on developing, building and supporting supercomputing technologies for over 40 years.

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



How to Meet Today's Risk Analytics Demands: Cray Supercomputing

Cray is a world leader in creating some of the most powerful and flexible compute infrastructures with the lowest TCO. So what does the compute grid of the future look like, and why does it represent a better way?

- **Easy, grid-like architecture.** It starts with the Cray® XC™ supercomputer, a scale-out distributed shared-memory architecture. The XC system looks like a grid in that it is a collection of compute nodes, but it is managed as a single system, which saves on systems administration costs.
- **Standard components.** Cray compute nodes have standard Intel® Xeon® cores and memory but do not have SSD or hard disks. The result takes up less physical space, uses less power, and costs you less per node.
- **Money-saving interconnect.** The Aries™ interconnect delivers excellent bandwidth and latency so network-dependent workloads like Spark™ run faster. The interconnect's design makes data or compute placement within the grid immaterial and workflow design easier. Aries replaces Ethernet or InfiniBand node interconnects, eliminating the need to build an inter-node network fabric and saving you time, support and capital costs.
- **Reduced data duplication costs.** Cray's DataWarp™ I/O accelerator is fast flash memory that eliminates duplication costs. It sits directly on the Aries interconnect as a shared resource allowing you to share data at local SSD speeds. Most Monte Carlo simulations use the same initial conditions, market or reference data replicated to each node. This replication means SSD is needed for each of thousands of nodes, and jobs spend valuable time and network bandwidth in that data replication. The combination of Aries and DataWarp enables an efficient and cost-effective system.
- **Load-and-go applications.** Your applications run unmodified. The XC system runs existing Linux® Intel binaries without change. If desired, the source can be recompiled leveraging Cray's advanced programming environment to drive greater efficiencies. Additionally, each node is lightweight, with a stripped-down OS that runs applications typically 5 percent faster than an equivalent commodity grid node. Network-dependent workloads such as Spark can see 30 percent performance increases.
- **Accelerator options.** Nodes can include accelerators such as NVIDIA® Tesla™ or Intel® Xeon Phi™. Cray sells more accelerated flops than any other hardware vendor. We also have deep skills to assist in application re-architecture if needed.
- **Rapid time-to-value.** Cray sells systems rather than components. Your system arrives with your software stack tested and ready to go, eliminating costs associated with component assembly, configuration, optimization and test.

In the new world of more flexible and efficient risking, pricing and trading systems, a supercomputer represents a more flexible, higher capability and lower TCO response than traditional scale-out commodity grids. The Cray XC system lowers costs, but also opens up a range of new architecture opportunities.