



Dan Lee
Research Vice President, Technical Computing



Steve Conway
Research Vice President, High-Performance Computing

High-Performance Computing in the Public Sector

July 2008

The global market for high-performance computing (HPC) servers has experienced strong growth each year since 2003. In 2007, sales of HPC servers reached US\$11.6 billion. IDC estimates that worldwide sales of HPC sever systems in the public sector (including government labs, defense, life sciences, and university /academic research) represented over US\$7 billion in 2007. IDC projects that this market will grow by around 10% per year over the next five years. Public sector server growth has been heavily driven by clustered servers that provide new levels of price/performance. Clusters recently crossed the 65% market share line to become the dominant species of HPC servers. At the same time, there are a number of customer pain points in the public sector, including growing software license fee costs; power, cooling, and facility space; and performance on application codes.

The following questions were posed by representatives of Cisco, Hewlett-Packard, Intel, and Microsoft to Dan Lee, Research Vice President in IDC's Technical Computing group, and to Steve Conway, Research Vice President in IDC's High-Performance Computing practice, on behalf of the companies' HPC customers.

Q. What are the main uses for HPC in public sector market segments today (physical sciences, life sciences, government and defense research, weather and climate research)?

A. For decades, researchers have been relying on computer simulations and modeling to advance the boundaries of science and engineering. Today, this dependence is stronger than ever. HPC has become an indispensable tool that enables scientists and engineers to solve complex problems, expand the frontiers of science and engineering, and improve the quality of life for humankind.

As a strong benefactor/funder for HPC, the public sector has been playing a commanding leadership role since the inception of this technology. What is perhaps less well known is the fact that a number of market segments in the public sector are also among the primary beneficiaries of the advantages that HPC has helped to deliver.

Historically, various branches of physical sciences such as plasma physics and astrophysics, high-energy and nuclear physics, and chemistry, along with government initiatives, programs and services in basic sciences, energy, materials science, defense, intelligence, aeronautics and space, weather forecasting and climate modeling have constituted the majority of public sector uses for HPC. Further, since the 1990s, the fields

of life sciences (in particular, computational biology and bioinformatics) and nanotechnology, each holding breakthrough promises, have rapidly moved into the limelight in the HPC arena.

Q. What major trends are driving HPC growth in the public sector?

- A. Collectively, the adoption of HPC in the public sector is motivated by a diversity of factors such as the drive for scientific leadership, national defense, industrial competitiveness, national health, and the mitigation of severe weather and climate change.

Moreover, the landscape of HPC has been evolving toward one that consists of both global competition and collaboration, a trend that's especially notable in Asia/Pacific and Europe, as opposed to previously where the market saw heavy competition between the United States and Japan. For example, some of the world's largest supercomputers can now be found in India and China, while in Europe there are ambitious initiatives that call for the establishment of a pan-European HPC infrastructure that will consist of multiple petascale systems at the top tier of the infrastructure.

The architectures of HPC have gone through several waves of evolution. In this aspect, one of the major trends of HPC has been the emergence of standards-based cluster computing. Driven primarily by its cost-effectiveness and the availability of open-source software, cluster computing using standards-based CPUs, interconnects, memories, I/O and disk components as the leading HPC architecture today. It continues to gain market momentum and is expected to maintain its leading status for the foreseeable future, both in the overall HPC market and in the public sector.

Q. How important are industry-standard technologies/solutions and easy-to-use GUIs?

- A. The adoption of industry-standard technologies and solutions in HPC helps to drive down costs, improve interoperability and shorten times-to-solution. HPC hardware and software solutions can benefit from tracking the fast growth of these technologies to shorten the time-to-market of their development cycles.

Still, in today's challenging research environment that demands operational efficiency and effectiveness, the majority of scientists and engineers want to focus on their core competency, which is to solve scientific and engineering problems.

One of the most severe pain points for users is that HPC systems, from entry-level to high-end versions, are too hard to use and parallel programming is still an arcane task that often turns the delivery of performance and efficiency into a monumental undertaking. Therefore, ease-of-use has always been a highly desirable HPC user requirement. Easy-to-use, cost-effective standard graphical user interfaces (GUIs), programming tools, and technologies such as x86 processors, standard operating systems and interconnects, can become key competitive differentiators that help to address this common pain point by greatly simplifying and accelerating parallel programming and development efforts.

Today, ease-of-use is an even more important requirement because of the growing complexity associated with different programming models, software implementations and hardware architectures, and the rapid increase in systems sizes. HPC reference architectures that provide pre-integrated, pre-tested hardware and software are becoming increasingly attractive for this reason.

Q. What hardware, software and networking challenges do public sector HPC users face?

- A. Today, public sector HPC users face a number of challenges. Major advances in parallel computing, memory bandwidth, and I/O, as well as networking, data storage and

management will be needed to enable scientists and engineers to simulate larger problems with finer resolution and/or more iterations and engage in global collaboration. A quick review of today's major HPC trends can help to elaborate.

First of all, software development has been lagging behind hardware development, and the gap is widening as a result of the industry's adoption of multicore technology.

Until recently, CPU chip designers have relied on increasing a chip's transistor density to improve processor performance. But the process of shrinking the physical dimensions of circuits is proving to be more difficult with technologies that are currently available or foreseeable in the near future. In the past, users typically were able to achieve application performance gains simply by upgrading to systems powered by new processors with higher clock frequencies. But as dictated by physics, processors with higher transistor densities and clock frequencies also consume more power and generate more heat. The escalating power and heat issues on single-core processors eventually led to the emergence of mainstream multicore processor architectures with per-processor core counts expected to double following the trend defined by Moore's Law.

From the software viewpoint, as the number of cores on a processor and the number of nodes in a parallel HPC system increase concurrently, pressure mounts for the introduction of new parallel programming technologies, software tools, algorithms, and highly parallel applications to take advantage of the added raw hardware power. This problem is further complicated by growing resource contention, because without breakthrough innovation, per-core memory bandwidth in today's multicore servers will suffer a 50% reduction every time the number of cores on a processor doubles.

The availability, scalability, reliability and affordability of efficient compilers, tools, systems and applications software are collectively the foremost challenge. For public sector HPC users, there is an additional challenge – this sector will be the predominant early adopters of the extremely large petascale systems that are on the horizon. There is little doubt that these leading-edge users will face new challenges related to scalability, programmability, system reliability and fault tolerance in this new HPC regime.

As the industry advances toward the peta-computing era, another important HPC trend will be the explosive growth of data driven by the need to run more refined simulations with more iterations for better results and time-to-solution. Together, these hardware and software issues present great challenges to the general and public sector HPC communities.

Q. What is being done to address these challenges?

- A. Many truly remarkable discoveries and achievements would have been impossible without HPC, yet the technology has the potential to deliver many more breakthroughs and promises. To press ahead, the HPC community must adequately address the challenges described earlier. Doing this would require HPC vendors and users to work in unison in several critical areas.
- First, the HPC community must work toward an overall long-term strategy that focuses more on software instead of only on "big iron" machines.
 - Second, vendors should recognize "ease-of-everything" as one of the most desirable user requirements and strive to offer programs, solutions and services that help to shield users from the increasing complexity associated with HPC hardware and software.
 - Third, the public sector should augment its leadership role by significantly boosting HPC research and development efforts within national labs, universities, and key agencies.
 - Finally, a healthy and vibrant HPC ecosystem cannot scale and be sustained without a brand-new agenda for education in parallel programming.

ABOUT THESE ANALYSS

Daniel Lee is a research vice president within IDC's Technical Computing group. He is actively involved in all aspects of IDC's high-performance computing research, consulting, and business development. A veteran of both the supercomputer and the storage industries, Mr. Lee conducts in-depth research in key HPC and storage areas including global market trends, HPC architectures, applications, storage and data management.

Steve Conway is a research vice president within IDC's High-Performance Computing group. He plays an important role in directing and implementing HPC research related to the worldwide market for technical servers and supercomputers. A 20-year veteran of the HPC and IT industries, Mr. Conway co-authors key IDC studies, reports and white papers, helps organize and advance the HPC User Forum, and provides thought leadership and practical guidance for users, vendors and other members of the HPC community.

Those interested in more information can contact the following company representatives:

Simon Lim
silim@cisco.com
(65) 6317-5588
www.cisco.com

Dennis Ang
dennis.ang@hp.com
(65) 6727-6074
www.hp.com

Dave Chee
dave.chee@intel.com
(65) 6213-1000
www.intel.com

Deepak Setty
hpcinfo@microsoft.com
www.microsoft.com/hpc

ABOUT THIS PUBLICATION

This publication was produced by IDC Go-to-Market Services. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Go-to-Market Services makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

COPYRIGHT AND RESTRICTIONS

Any IDC information or reference to IDC that is to be used in advertising, press releases, or promotional materials requires prior written approval from IDC. For permission requests, contact the GMS information line at 65-6829-7757 or gmsap@idc.com. Translation and/or localization of this document requires an additional license from IDC. For more information on IDC, visit www.idc.com. For more information on IDC GMS, visit www.idc.com/gms.

IDC Asia/Pacific, 80 Anson Road, #38-00, Singapore 079970 P.65.6226.0330 F.65.6220.6116 www.idc.com.

Copyright 2008 IDC. Reproduction is forbidden unless authorized. All rights reserved.