I D C A N A L Y S T C O N N E C T I O N

Jie Wu
Research Manager, High-Performance Computing

HPC Directions In Digital Content Creation and Distribution

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 systems in the digital content creation and distribution (DCC) sector made up about US$678 million of the 2007 total and will grow to US$1.2 billion in 2012. Revenue growth on the computational side has been driven almost entirely 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 DCC sector, including growing software license fee costs, the need to manage the complexity of HPC resources, and the desire to handle some of today's batch-oriented jobs interactively.

The following questions were posed by representatives of Cisco, Hewlett-Packard, Intel and Microsoft to Jie Wu, Research Manager in IDC's High-Performance Computing practice, on behalf of the companies' HPC customers.

Q. What are the main application areas in DCC that need HPC power?

A. The major areas in the DCC market that require HPC power are the following:

- Digital content creation and distribution. Central to this market are applications in the areas of 2D and 3D animation – occasionally this includes the temporal dimension to become 4D animation – along with film and video editing and production, and multimedia authoring for CD-ROM and Web pages that utilize sophisticated graphics content. These DCC applications are heavily used to create animated films and movie special effects, but also for example to visualize hidden oil and gas deposits in the petroleum industry.

- Virtual reality visualization. Typical applications include product design and styling that require high-quality visualization. For instance, the automotive industry employs virtual reality visualization to help optimize the ergonomic placement of dashboard instruments and other elements inside the passenger cabins of vehicles.

- Large-scale gaming involving real-time rendering. Most of the world's major gaming sites that require HPC are located in the Asia/Pacific region.

- Transcoding involves rapidly encoding and decoding content across a wide range of formats and resolution levels. This is also related to remote visualization usage of the contents.

IDC Asia/Pacific: 80 Anson Road, #38-00 Fuji Xerox Towers Singapore 079907 P.65.6226.0330 F.65.6220.6116 www.idc.com.sg

Q. Which of these application areas will experience growth in the next two to three years?

A. All of these areas will see increased HPC adoption in the near-term future.

Q. Will any DCC applications running on workstations today require HPC servers in the future?

A. Historically, some problems that once required HPC servers have migrated over time to desktop workstations. In these cases, the workstations have grown sufficiently in performance to handle the former HPC-level problems, and the HPC systems have moved on to address a new generation of more challenging problems. But this historical movement also operates in reverse in some cases. As problem sizes and complexity grow, the demand for more computing power can exhaust the ability of contemporary workstations to run certain DCC applications at adequate performance levels. In these cases, users will need to move up to HPC servers in order to get their jobs done. IDC expects the movement of certain problems from desktop computers to HPC servers to occur in nearly all HPC market sectors. In the DCC sector, IDC sees important areas of desktop-to-server application migration as follows. IDC foresees strong growth for the Asia/Pacific region in both of these areas.

- Computer graphics (CG). The goal in this area is telling stories with state-of-the-art visuals. The top-grossing films of the past 20 years have almost all had state-of-the-art visuals and technical effects. This takes a lot of computational power and sophisticated software, and HPC is a key enabler for making this all come to life in films. It's also a key differentiator as audiences' expectations continually rise.

- Interactive rendering. Each frame of an animated film takes approximately 100 hours of compute time today. The largest computing task is rendering. Today, people sit at workstations and iterate on a piece of film, then at night submit it to a compute farm for overnight batch processing. The new goal is to have enough compute power to substitute interactive daytime processor for overnight batch processing. In the United States, DreamWorks Animation is aggressively pursuing this research goal by using one of the nation's largest HPC systems at the Department of Energy's Oak Ridge National Laboratory.

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

A. The movement toward industry-standard technologies/solutions and easy-to-use GUIs is extremely important. Most DCC users are neither computer scientists nor computational scientists. Ease-of-use and familiar, standard solutions are therefore crucial to help those users quickly move onto HPC servers and enable them to solve larger and more complex problems. Ease-of-use-oriented solutions will help graphic designers and content creators save tremendous amounts of time in the up-front learning curve and in the research and production phases. In addition, these solutions will provide investment protection benefits for end users – many DCC sites purchase large quantities of technical servers, and resources based on standard technologies are more affordable and easier to manage and upgrade.

Q. From both the hardware and software perspectives, what are the main challenges faced by HPC users in the DCC market?

A. From the hardware perspective, the main limitations are restrictions on I/O bandwidth, and total cost of ownership (TCO). As HPC hardware technologies continue to advance and systems built on multicore technology evolve and become mainstream, the requirement for higher I/O bandwidth capabilities is becoming increasingly important to achieve better overall system performance. TCO is also becoming a primary concern for many users in the DCC space. Key TCO considerations include costs for systems and

operations/maintenance, power and cooling, facilities, and storage and data management.

From the software perspective, one major hurdle is the increasing complexity of HPC clusters and related resources. Rapidly increasing average cluster sizes, multicore processors and accelerator options, and grids and other distributed computing environments all create important new challenges for software – from algorithms and applications to programming languages, operating systems and middleware. This growing complexity especially needs to be addressed for new and entry-level HPC users, who often lack in-house access to HPC expertise.

Q. What is being done to address these challenges?

A. HPC vendors are increasingly taming the challenges associated with purchasing, deploying, operating and upgrading HPC servers. The rapid global proliferation of standards-based HPC clusters since 2002 attests to the success of vendors' efforts to make these systems affordable and productive even for small businesses. More recently, leading HPC vendors have upped the ante on "ease-of-everything" by introducing reference architectures designed to ensure that the pre-integrated, pre-tested hardware and software elements of clusters will work well together "right out of the box." Some vendors are even bundling popular third-party (ISV) applications with their HPC systems based on reference architectures.

In addition, hardware and software vendors are partnering closely with accelerator vendors to augment the capabilities of standard x86 processors for certain problems with graphical processing units (GPUs), field programmable gate arrays (FPGAs) and other accelerators. Some early-stage visualization products are now available that integrate compute and storage capabilities. Operating environments are being enhanced to address the unique challenges of HPC clusters and computing environments, as are interconnect technologies and products.

ABOUT THIS ANALYST

Jie Wu, Research Manager within IDC's Technical Computing team, provides market analysis, research, and consulting in the Technical Computing space. Specific research areas include market sizing, segmentation, competitive analysis and forecasting of the Technical Computing market. With the expertise in building results from field studies, Ms. Wu provides insightful guidance and actionable recommendations to the Technical Computing community.

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

Simon Lim	Dennis Ang	Dave Chee	Deepak Setty
silim@cisco.com	dennis.ang@hp.com	dave.chee@intel.com	hpcinfo@microsoft.com
(65) 6317-5588	(65) 6727-6074	(65) 6213-1000	www.microsoft.com/hpc
www.cisco.com	www.hp.com	www.intel.com	

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