



HPC Visualization Challenges & HP Solutions



Introduction	2
Visualization Challenges.....	2
HP Visualization Solutions.....	3
HP Scalable Visualization Array (SVA).....	3
The HP Parallel Compositing Library	4
For more information	4

Introduction

High-performance computing (HPC) often involves modeling and simulation, such as the flow of air around a vehicle, the effects of a head-on automobile collision, and weather modeling. These simulations produce vast amounts of data. Instruments and sensors also produce large amounts of data, such as seismic data for oil and gas exploration and data from medical instruments such as CT scanners. Consequently, large data sets need to be analyzed and visualized to gain insights that will lead to innovation, better design, advanced research and science outcomes and improved time-to-market.

Visualization Challenges

HPC visualization uses desktops varying in size and price, largely depending on the graphics card and amount of system memory. Usually, these desktops visualize data sets copied from HPC systems to the office systems. Tremendous advances in graphics technology have occurred in recent years, yet visualization requirements for very large HPC data sets are well beyond those of a typical desktop. For example, seismic data sets can range from 10's to 100's of Gigabytes. Putting an expensive desktop system in an office and copying such large data files from an HPC system to the desktop isn't necessarily the best approach.

The size of particular data sets and the collaboration requirements usually drive many of the requirements unique to HPC visualization. Cost is also an important factor, as customers want an approach that is based on industry standards and COTS technology, which results in economies of scale and performance.

Large data sets produce images that require of substantial screen real estate. Visualization users work with extremely large models with fine details that must be seen in context. Displays with 4 to 8 times the resolution of a typical desktop display or more may be needed, such as quad-HD projectors and LCD with resolutions of 3840x2160 (over 8 MPixel compared to typical desktop display of less than 2 MPixel). These displays are driven by up to 4 video inputs and display large data sets that have been rendered in great detail.

Collaboration is an equally important reason for large multi-screen displays. Because multi-projector or multi-LCD displays scale-up to even higher resolutions, they are used when people need to collaborate on a visualization project. Examples include researchers conducting joint analysis, engineers deciding where to drill for oil, or designers working together to create new aircraft models.

Other challenges also exist in visualizing increasingly large data-sets. Rendering large data sets can exceed what is feasible to do at interactive frame rates on a single graphics card, even when the image will only fill a single display. To achieve this goal, one must distribute the problem over multiple systems, typically a cluster of systems connected by a high-speed interconnect, such as InfiniBand. This approach is similar to distributing a computation across multiple nodes, only in the visualization scenario, the systems are equipped with graphics cards and the results are now images.

Another issue with large data sets is that they are inconvenient to copy from the HPC systems (where they were computed) to where they need to be visualized over a campus or miles apart. Sending these large data sets from their source slows the workflow and can stress backup systems. Users want to be able to visualize large data sets from their offices, without the time and cost of copying those data sets to their offices.

Companies would also prefer to avoid putting expensive high-end graphics systems in every office. Engineers and scientists may only need the high-end capabilities some of the time, but they want to be able to access them when needed.

A solution to both of these challenges is to pool together and share high-end graphics systems with HPC systems that have high-speed access to large data sets. The graphics systems can then be accessed by remote visualization. Remote visualization allows users to run a visualization application

on a high-end system in the lab but display the results locally on an office desktop. Users can interact with the application as if the keyboard, mouse, and display were plugged into the system in the lab.

To summarize the HPC challenges:

- HPC applications produce very large data sets that need to be visualized.
- Visualizing these data sets can require multi-screen displays and multiple graphics systems working together to drive these displays or distribute the work of producing the images.
- The HPC computing and storage systems are often located at some distance from the users' offices, but users want to visualize the results from their offices. Copying very large files and dedicating high-end desktops to every individual is a costly approach.

HP Visualization Solutions

To address these HPC challenges, HP provides a range of hardware and software solutions.

HP offers rack-mount servers with graphics cards that provide graphics systems in as little as 1U of rack-space. These systems include the management and monitoring features of a server and the reliability and price-performance that customers have come to expect from HP servers. For example, the HP DL160G5 Xeon-based server is particularly well suited for interfacing to graphics systems, since it supports PCI-Express 16x Gen2 cards.

HP also offers a 2U option, the Opteron-based DL385G5 server. Top-of-the-line graphics cards from NVIDIA, such as the FX 5600, are available. These graphics servers keep the data storage, high-end visualization, and computing systems consolidated as shared resources, while allowing engineers to visualize the data remotely from inexpensive desktops and even laptops.

HP also offers clusters of graphics systems that includes many systems managed together that share a high-speed interconnect. These systems can also share access to a large-scale parallel file system over this interconnect. The compute-systems can also be a part of the same interconnect, allowing the compute, visualization, and storage resources to all be integrated. These clusters can be used to run applications that drive the multi-screen displays (or CAVE environments) and scale-up to handle very large data sets by distributing image rendering across multiple systems. Or the cluster can be used to efficiently manage a shared pool of high-end graphics systems that are used remotely.

HP Scalable Visualization Array (SVA)

To tie all the pieces together, HP provides the Scalable Visualization Array (SVA) for managing and using graphics clusters. SVA is part of the Linux-based HP cluster system software and Unified Cluster Portfolio (UCP). SVA supports a range of hardware platforms and graphics cards and is available factory integrated and tested with HP worldwide support and service.

SVA enables scaling of interactive visualization to large data sets, using the Parallel Compositing Library (described later). SVA provides mechanisms that describe multi-screen displays, allocate cluster nodes to users, and launch visualization applications on the cluster.

Related to this, SVA provides features for running remote graphics software such as HP RGS sessions on a cluster. Consequently, the graphics cluster in the HPC environment becomes, in essence, a managed, shared pool of high-end graphics systems.

HP provides remote visualization via Remote Graphics Software (RGS). Customers use RGS to deliver excellent image quality to the desktop using standard Ethernet at interactive frame rates. SVA also provides similar support for running the open source remote visualization tools VirtualGL and TurboVNC. Besides giving remote access to high-end shared graphics systems, these tools enable remote collaboration. Other users can join a session, view, and interact with the application.

Applications use distributed rendering to scale-up to handle very large data sets. Basically, they divide the data across multiple cluster nodes, and each node produces an image for part of the data. Next, these partial images are combined to produce the complete image. The images are typically

combined, pixel-by-pixel, using either depth or transparency information associated with each pixel. A number of such applications are available commercially and as open source.

The HP Parallel Compositing Library

To help develop parallel cluster-aware applications, HP developed the Parallel Compositing Library. Using information about how the partial images overlap, the library routes pixels among nodes, combines the pixels, and delivers the resulting image to a display, which may be a multi-screen display. This library grew out of advanced development work that had its roots in an HP project to develop compositing hardware (SEPIA) and the work that HP did with others to define an API for parallel compositing. Within HP-CCN (HP Collaboration and Competency Network), HP continues to work with customers, research partners, and ISV's to enhance the library and incorporate it into applications and visualization libraries.

The Parallel Compositing Library simplifies the task of distributing an application by implementing optimal techniques for moving pixels within and between nodes. The Library has been optimized for reading and combining pixels and then transmitting those pixels over GigE and InfiniBand interconnects. This process frees the visualization application developers from learning and implementing interconnect-specific procedures, which may very well be outside their area of interest and expertise.

In a sense, the Parallel Compositing Library does for distributed visualization applications what MPI does for computational applications—it makes clusters a more approachable platform for scalable visualization.

Recognizing that application developers would be hesitant to have a dependency on a proprietary library, HP made this library open source for use and further enhancement by the open community. The Parallel Compositing Library is available on SourceForge.net. The project web page includes pointers to some of the applications that are using this library, including the source code.

In summary, HP is working with its customers and partners to understand the challenges unique to HPC visualization and to deliver cost-effective high-performance solutions:

- HP is providing graphics-enabled systems, based on industry-standard components. Visualization applications can scale-up to the data set and display sizes needed to handle the large data sets produced by HPC compute applications.
- HP also provides a remote visualization capability to help consolidate high-end systems and integrated them with other HPC storage and compute resources.
- With SVA and XC, HP is providing a tested and supported cluster software system that makes it easy to manage and use the shared visualization resources of a cluster.
- The HP-developed Parallel Compositing Library is available as open source. The Library makes it easier for developers to create distributed scale-up visualization applications to handle very large data sets.

For more information

www.hp.com/go/visualization

http://www.hp.com/techservers/hpccn/sci_vis/index.html

<http://sourceforge.net/projects/paracomp/>

© Copyright 2008 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Linux is a U.S. registered trademark of Linus Torvalds.

4AA1-9732ENW, May 2008

