Matching Visualization with the Power of HPC

In scientific and engineering environments, high-performance computing (HPC) can produce vast amounts of data through simulations and modeling. To support innovation, design, research, and business efforts, this data must be analyzed. One of the most effective and easy to use analysis methods is visualization.

The need for visualization solutions spans key industries including aerospace, geosciences, automotive, manufacturing, research and defense, and medical and scientific research. Scientists and engineers often need to render complex 3D and 4D simulation data, using techniques such as direct volume rendering and isosurface extraction. Another common requirement is a platform that not only supports visualizing data statically, but also supports real-time interactivity and collaboration.

Visualization challenges

Finding visualization tools that match HPC capabilities can be difficult. Organizations face a number of challenges due to the amount and type of data being collected and rendered, for example:

- **Norsk Hydro:** Simulates production from reservoirs of various hydrocarbon components of oil and gas with pressure variations.
- **Western Geophysical:** Generates more than one petabyte of data per week for oil and gas exploration and production.
- **State University of New York, Stony Brook:** Renders interactively full-color Visible Female volume data of approximately 40 gigabytes.
- **SHARCNET, Canada:** Develops a volume rendering capability for astrophysical magneto hydrodynamic simulations that provides real-time interaction.

Today’s increased computational power can process and generate raw data in a timely fashion. However, it creates a new problem in that the large amount of generated data is difficult to analyze. Collaboration also becomes more difficult, and display size and resolution must increase.

Large data sets

Visualizing large HPC data sets are well beyond the capabilities of a typical desktop computer. High-end graphics systems are needed in each office that wants visualization capability. In addition, viewing the information in the office requires copying large amounts of data from the HPC system to your visualization system, slowing your workflow and stressing network and back-up systems.

Collaboration requirements

Technical problems require collaborative analysis. Visualization allows engineers, scientists, and other team members to analyze data and make decisions together. Multi-screen displays and higher resolutions are needed to provide the visualization tools for effective collaboration and analysis.

Multi-screen displays, multi-graphics systems

Large data sets create images that demand a lot of screen real estate. To work with large models and fine details, you may need displays with 4-8 times the resolution of a typical desktop system, or even more in the case of tiled displays, stereographic walls, and CAVEs. And rendering large data sets at interactive frame rates is difficult on a single graphics card, even if the image fills only a single display.

How can a graphics cluster help?

Much like using a cluster for computing needs, a graphics cluster system handles the visualization requirements of large data sets, matching display power with the compute power of HPC systems.

A graphics cluster consists of a group of workstations or servers that work together to perform visualization tasks. A graphics cluster typically uses commercial, industry-standard components, keeping costs down and providing support for future upgrades.

Graphics clusters can run applications that drive multi-screen displays and CAVE environments. They can also act as a shared pool of high-end graphics systems for individual remote users. Because they are flexible, you can use a graphics cluster independent of or integrated with HPC compute and storage resources.

What is the HP Scalable Visualization Array (SVA)?

The HP SVA is a cluster platform system. It can be a specialized, standalone system consisting entirely of visualization nodes, or it can be integrated into a larger cluster and share a single System Interconnect with the compute nodes and a storage system.

To illustrate how an HP SVA cluster works with an integrated compute and storage system, consider the following example:

- An automobile crash test simulation runs on the HPC compute nodes of the cluster.
- The large dataset generated on the compute nodes is stored in the storage servers for later retrieval or directed in realtime for rendering on the SVA portion of the system.
- Multiple users can log into the SVA, which allocates resources efficiently to meet the rendering and display requirements of each user application.
- Users' visualization applications use parallel programming techniques and visualization software, such as CEI's EnSight DR, to distribute graphical rendering across the SVA nodes, each of which in turn renders a portion of the output for the final image.
- Each portion of the final image is sent to a tile of a single or multi-tile display. The complete image is available for display locally. The complete image is also available for display remotely, but limited to single or two-tile output from a single graphics card.

Applications such as CEI's EnSight DR use distributed rendering to handle these very large data sets. They divide the data across multiple cluster nodes, and each node produces an image for the part of the data that was assigned to it. These partial images then get combined to produce the complete image. By distributing the load over multiple systems, an SVA cluster processes significantly larger data sets and can create larger images than any individual graphics card.

HP SVA Architecture

The HP SVA consists of a cluster of HP workstations or servers in visualization block configurations running open-source Linux. It uses industry-standard graphics cards and network adaptors, together with an integrated software system.

Display devices are not necessarily provided as part of the SVA. For example, a site can use projector display systems or immersive displays provided by third-party vendors. SVA hardware and software deliver images to digital or analog standard interfaces. The SVA depends on the graphics cards to drive the image output. This means that a wide range of display devices supported by the graphics cards are available for use.

SVA software provides mechanisms for describing multi-screen displays, for allocating cluster nodes to users, and for launching visualization applications on the cluster. Remote Graphics Software (RGS) delivers images to the desktop over standard Ethernet at interactive frame rates, supporting remote access and collaboration.

Effective HPC Visualization

SVA allows true interactive processing of large data sets. Coupled with distributed rendering software such as CEI’s EnSight DR, SVA harnesses the power of a graphics cluster to reduce rendering time, process large data sets, and create larger, more detailed images.