LCSE FACILITIES
JULY 23, 2005

NSF MRI Project

In September, 2004, NSF funded a project to build a prototype system in the LCSE for interactive analysis and visualization of multi-terabyte data sets. The team participating in this project, led by LCSE Director Paul Woodward, includes Profs. David Du, Ernest Retzel, David Yuen, Thomas Jones, Jon Weissman, Baoquan Chen, and Dr. David Porter. This prototype system is presently under construction. A system diagram is shown below. A 24-port Infiniband-4X switch from Topspin (recently bought out by Cisco) interconnects two 16 CPU Unisys ES7000 machines and 12 Dell PC workstations. Sustained data transfers of 350 MB/sec have been achieved with our software on Infiniband-4X. The Unisys SMPs were donated to the LCSE in March of 2003, with Intel donating the CPUs and Microsoft donating software.

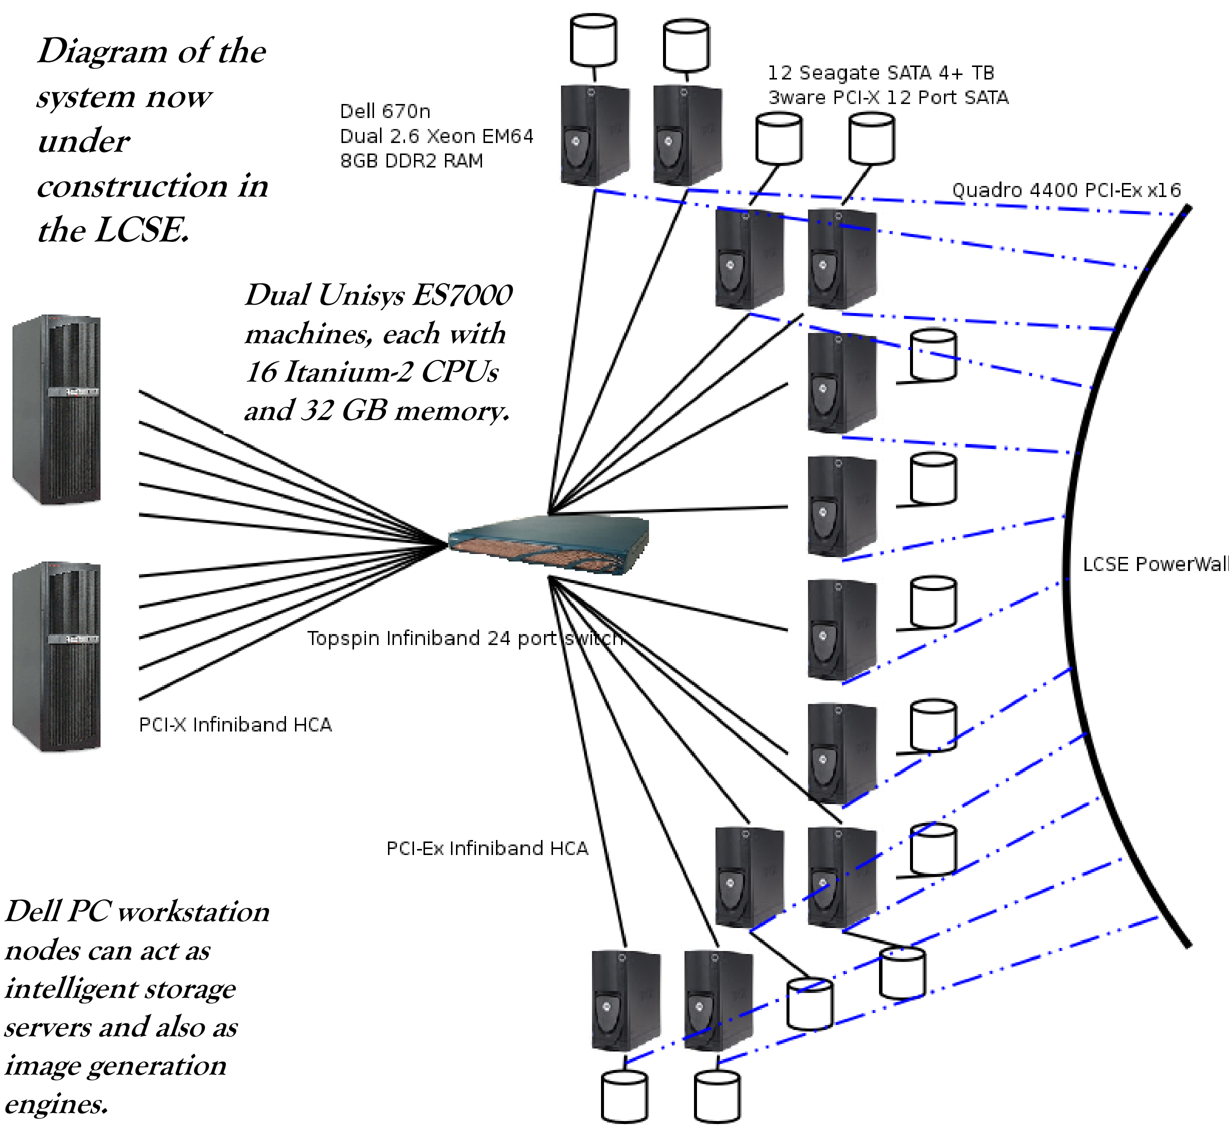

Each ES7000 has 16 Intel Itanium-2 CPUs running at 1.3 GHz and sharing 32 GB of memory. A detailed performance study of LCSE applications on these machines can be found at www.lcse.umn.edu. Each Dell PC workstation has dual Intel Xeon CPUs with 64-bit memory extensions running at 3.6 GHz and sharing 8 GB of memory. Graphics rendering power for delivery to the 10-panel LCSE PowerWall is provided through an nVidia Quadro 4400 graphics engine in each PC workstation. 12 Seagate 400 GB 7200 rpm SATA disks are attached to each Dell workstation using a 3Ware 12-channel controller. Sustained data rates of over 400 MB/sec have been achieved on a single machine with this disk subsystem. The disk speed and Infiniband interconnect speed are therefore well matched. We are grateful to Dell and to Seagate for arranging special pricing for these system components.

The goal of this system development is to take the LCSE applications to a much more interactive level. The LCSE team has been engaged in generating PowerWall movie visualizations for many years, but these movies have been generated in large batch-oriented processes. Over the years, the generation time has been reduced, even for our largest data sets, to an overnight batch rendering job (see ADCS lab project below). Our new system is intended to enable movie *generation* under interactive user control at full PowerWall resolution from multi-terabyte data too large to fit into a single renderer's memory. The prototype system will have only a single renderer for each PowerWall image panel, but for the largest data sets the 12 nodes will collaborate to generate images for only a single panel, so that an interactive frame rate can be maintained. A key concept is to replicate up to 2 TB of data on every PC workstation node, so that each node can stream into its graphics engine the subset of this data that it needs without any contention from data requests made by other nodes. The remaining 2 TB of disk storage on each node will be used for less volatile data. The LCSE's HVR volume rendering and movie generation software is being adapted to this new interactive mode of use. The potential also exists to perform interactive movie generation on the PowerWall from raw simulation data, without having first to process this data into voxel bricks for input to HVR. The Unisys SMPs, with their large shared memories and 16 Itanium-2 CPUs, are ideally suited to processing the more voluminous raw data, piping the voxel results over the Infiniband network to HVR running on the Dell nodes.

ADCS Student Lab Project

Through an earlier NSF-funded project, the LCSE team has collaborated with the University of Minnesota's Academic and Distributed Computing Services (ADCS) organization to leverage the ADCS investment in a large student lab in the Walter Library building that houses the Digital Technology Center, of which the LCSE is a part. This project team, led by LCSE Director Paul Woodward, includes Profs. Ernest Retzel, David Du, and Jon Weissman, as well as Prof. Ted Wetherbee of the Fond du Lac Tribal and Community College in Cloquet, Minnesota. The student lab in Walter Library contains 67 Dell PCs and 32 Macs. These are new machines, installed in June, 2005, that replace an earlier set now located on the St. Paul campus, where they are available to us over a Gigabit Ethernet link 24 hours per day. The new machines in the Walter Library building are available to us only when not in use by students. In the Figure above, taken in the ADCS lab very late at night, the machines (the set now on the St. Paul campus) that are not in use by students are seen rendering a PowerWall movie. Under this project we outfitted 52 of the original set of Dell PCs with 400 or 500 GB of disk storage each and with Gigabit Ethernet NICs and nVidia GeForce-FX graphics engines. We also purchased 20 TB of external FireWire-800 disk storage that is hot-pluggable and can be used in a flexible manner. Test clusters of Macs (for Retzel's effort in the CCGB) and PCs (for the LCSE effort) were also purchased, with the PCs attached to the LCSE PowerWall display. A Dell Gigabit Ethernet switching fabric was also installed interconnecting the LCSE machines with those in the ADCS student lab. Software was developed to automate management of disk files over these resources, accepting batch

requests for large sets of file moves that can execute late at night, when the ADCS PCs become available, without user interaction. Software for servicing large batch requests for automated movie generation on the ADCS PCs was also developed to exploit this special resource, which is generally unavailable during business hours.

Working together with the University's ADCS organization, and exploiting an ongoing relationship with Seagate to obtain special disk pricing, we used funds from this NSF project to double the memories of the 32 new Macs and 67 new Dell PCs in the ADCS lab to 2 GB and to purchase 400 GB 7200 rpm SATA disks from Seagate to install in the 67 PCs. Using the LCSE disk file management software, these disks provide an additional 26.8 TB of storage space. We have targeted the 32 dual 2.7 GHz G5 CPU Macs for use by Retzel's team for their bioinformatics applications and the 67 PCs, with single 3.8 GHz Pentium-4 CPUs, for fluid dynamical applications of the other project team members. The 250 GB SATA disks of the Macs provide enough space to create separate partitions for the student use and the bioinformatics use when the machines would otherwise be idle. The new ADCS cluster has roughly double the capability of the older one, and the 2 GB machine memories allow us to exploit this power for computation.

LCSE PowerWall Display:

The first LCSE PowerWall display and visualization system was built with support from the Army High Performance Computing Research Center (AHPCRC) at the University of Minnesota in collaboration with Silicon Graphics, Seagate, Ciprico, NPI, and Ampex as a prototype for demonstration in the SGI exhibit booth at Supercomputing 1994 in Washington, D. C. In 1995, an NSF MRI grant and an NSF MetaCenter Regional Alliance grant funded the construction in the LCSE of a second-generation system in which a single, more powerful SGI Onyx machine was able to drive all 4 image panels of the system, with 1600×1200 pixel resolution each. A few years later, support from the NSF PACI program through NCSA together with a DoE ASCI grant through the Lawrence Livermore Laboratory and a generous equipment loan from SGI/Cray allowed us to build a third generation PowerWall with 10 image panels of 1280×1024 pixels each driven by a cluster of SGI VisPCs or alternatively driven by 2 Onyx machines. This system was installed in the PowerWall room in the University of Minnesota's Digital Technology Center in 2001. It is shown in the image above, with David Porter and Paul Woodward in the foreground. Development for this system is on-going, but has been focused in recent years on the problem of convenient and rapid content generation (movies and images) and the associated management of many terabytes of data consisting of hundreds of thousands of individual files. The image generation has been done from disk and on parallel systems since the outset in 1994, but the numbers of disks, of network nodes, and of rendering engines has steadily increased. In 2005/2006 we hope to achieve the goal of interactive rendering from our largest data sets at full PowerWall resolution. This step will take us from the realm of visualization for data *presentation* to that of visualization for data *exploration*.