hp success story

www.pnl.gov

The HP cluster in the Molecular Science Computing Facility, located in the William R. Wiley Environmental Molecular Sciences Laboratory, is operated by Pacific Northwest National Laboratory for the U.S. Department of Energy.

Pacific Northwest National Laboratory (PNNL) – Itanium® 2 and Linux-based systems

Environmental remediation, nuclear nonproliferation, genomics and proteomics; these are but a few of the many areas of science explored at the Pacific Northwest National Laboratory (PNNL), a U.S. Department of Energy (DOE) research facility located in Richland, Washington. Starting from the Hanford Works in 1940, and formally designated in 1965, PNNL originally provided research support for the DOE's Hanford Site (then the major producer of plutonium for U.S. nuclear weapons), and has since evolved into a multifaceted national laboratory.

In the 1980s, the focus of the laboratory turned from research and development related to nuclear energy and peaceful uses of nuclear materials to cleanup of the environmental hazards created by plutonium production of previous generations. Today, the primary topics of inquiry are reflected in PNNL's specific research areas: Energy Security and Information Technology;

Environmental Remediation; Fundamental Science with a focus on Chemistry, Nanoscience and Biology; and National Security. Each of the four divisions of the laboratory has its own mission and expertise.

finding a good balance

Naturally, such a facility faces certain challenges in matching technology to its needs. "Our biggest challenge is to quickly compute very complex, very long-running simulations. That means our software needs to run very efficiently on each processor and be scaleable across the entire system," notes Theresa Windus, technical group leader of molecular sciences software at the William R. Wiley Environmental Molecular Sciences Laboratory (EMSL), a user facility funded by the DOE. Furthermore, software for computational research, such as the massively parallel chemistry program NWChem, uses multiple computational methods requiring different actions from a computer. "Molecular dynamics, for example," says Windus, "requires low latency to global memory. Other methods rely heavily on having very good I/O to disk. So we need to have something that performs well for all those methods."

Early in 2002, following a thorough 15-month formal procurement process to determine the best equipment to meet user requirements, PNNL ordered an HP Linux-based Cluster 9.2-teraflop supercomputer to replace the old system. Windus says, "This system has a nice balance for the kind of computations we need to do." Dave Dixon, ESML associate director of theory, modeling and simulation, adds, "We purchased this machine because of its balance and the performance of each of its parts."

Particular strengths noted by both Dixon and Windus include very fast processors, fast access to large, local memory, a fast I/O subsystem, and cache and interprocessor communications systems compatible with types of simulations that will be performed on the new system.

getting more exact solutions

According to Dixon, "The new machine, enables people to do very different calculations than they could have done before.

"Something that was a challenging problem previously is no longer a challenging problem. This system will enable us to change the science we do and improve turnaround time."

Windus adds, "The system allows us to use fewer yet faster processors than previously required to solve the same problem. Even though we are using fewer processors, we actually can get the job done quicker. In addition to doing many more computations, we can also perform much larger calculations that solve our complex systems of equations much more accurately, leading to more reliable solutions."

The system is based on a cluster of HP rx2600 servers with a total of 1,438 Intel® Itanium® 2 processors. Each of the 700 nodes will have an independent connection to the 53TB Storage Area Network and will include a QSNet2/Elan4 interconnect from Quadrics. Due to a more efficient floating point processor and a better infrastructure than in previous systems, the supercomputer is capable of processing at close to peak efficiency.

In addition to the hardware, HP is providing engineering support to help install and integrate the system. This support consists of HP engineers working onsite with EMSL staff as they build the system and put it together, and assisting the lab with tuning its operating system and various application codes on the new machine. As always, HP is contributing to the whole solution, not just hardware.

"... This system will enable us to change the science we do and improve turnaround time."

Theresa Windus
Technical Group Leader
Molecular Sciences Software, EMSL

challenges

- **quickly compute complex, long-running simulations**
- **software that runs efficiently and is scalable across the entire system**

solution

- **world's fastest Linux-based supercomputer running on the Itanium 2 architecture**

results

- **use fewer, yet faster, processors than previously required to solve the same problem**
- **perform much larger calculations that solve complex systems of equations much more accurately, leading to more reliable solutions**

looking ahead

As of mid-2002, the supercomputer is still being tested and integrated at EMSL; when fully operational, it promises to enable its users to better address and solve the complex environmental problems facing the DOE and the U.S. The new machine promises to aid in creating solutions to such problems as the impact of environmental toxins on water, soil and the air we breathe as well as enabling us to better understand how cells work and how the environment impacts health. Looking five years out, Dixon predicts the system will have enabled EMSL's scientific user community to "make a big impact on environmental science and environmental remediation in the U.S. and in biology by solving large, complex problems that just couldn't be addressed with current technology."

customer at a glance:

industry sector: Government/Research

name: Pacific Northwest National Laboratory (PNNL)

headquarters: Richland, Wash.

founded: 1965

URL: www.pnl.gov and www.emsl.pnl.gov

technology highlights:

- HP Linux-based Cluster 9.2 teraflop Supercomputer
- 700+ rx2600 servers with Intel Itanium 2 processors
- 53 terabytes Storage Area Network
- HP Support Services

Published in USA