NASnews

▶ Main Menu July – August 1997 • Volume 2, Number 25

Within This Article...

In This Issue...

Langley-Ames SP2 Metacenter Enjoys Success, Plans for New Challenges

by James P. Jones and Mary Hultquist

Two members of the Langley-Ames metacenter team report on progress and plans for the future. The project, which connects two IBM SP2 testbed systems at NASA Langley Research Center (LaRC) and the NAS Facility at Ames Research Center, is funded by NASA's High Performance Computing and Communications Program.

The Langley-Ames metacenter team has achieved numerous goals since the project's inception in late 1995: getting faster job turnaround; balancing the workload across both systems; providing a wider range of resources; and migrating jobs automatically, giving users the ability to direct or limit migration.

These successes have attracted the attention of other sites, including Wright-Patterson Air Force Base and the U.S. Army Corps of Engineers, Vicksburg, MS. Team members at LaRC and Ames have begun transferring critical software used in the metacenter to these sites, including the Portable Batch System (PBS, the metacenter's batch queuing system), NAS Site Wide Accounting (ACCT++), and the PBS external job schedulers written at NAS and LaRC.

Focus on Access, Software

Since October, when the metacenter officially became the primary means for users to access the IBM SP2s, software development efforts have focused on two major areas: improving access to user files and data (which may reside on either system) and identifying critical software to be transferred to the next version of the metacenter testbed.

Fast, Reliable Access Not So Easy

The largest technical challenge has been to provide users with fast, reliable access to files located across the metacenter. Currently, users must track their file locations, as those files can't be seen from other

systems. The best solution would be to have all users' SP2 files visible and accessible from anywhere within the metacenter. However, given the immaturity of software on the testbed systems, providing this capability is easier said than done.

The two most promising distributed filesystems, NFSv3 and DCE/DFS, are not yet available and not mature enough, respectively, to reliably fill this need. Metacenter team members are working with IBM and other vendors to provide this global accessibility. As an interim solution, the team is using PBS, which provides a file-staging capability (a mechanism for users to identify files that are needed by their batch jobs). PBS handles the copying of files to the appropriate system within the metacenter, and returns solutions to the location specified by the user.

In addition, the team has modified PBS, the job scheduler, and the network configuration within and between the two SP2s to improve file transfers. While this approach has worked, it is not the most graceful solution, because it requires that users specify all files to be used during a job. The team is pursuing the elusive goal of a global shared filesystem as it works with vendors and develops local software.

Sharing Experiences, Planning for Future

While continuing work on these enhancements, the metacenter team is sharing its experiences with other interested sites through invited talks and visits. The team is also planning the transition to the next HPCC-funded testbed, which will replace the current SP2s with yet-to-be-determined systems. This "new" metacenter will expand to a include a third parallel system located at NASA Lewis Research Center. All three systems are scheduled to arrive by the end of the year.

Most options for the next system include architectures with a single system image -- that is, each supercomputer appears to users as one system, with contiguous processors and memory, rather than as a cluster of machines, each with their own CPUs and memory.

A shift in the underlying architecture alone could alleviate much of the difficulty experienced with the current metacenter. In addition, solutions for locating and transferring data files for processing will be designed into the new configuration. (The SP2s were originally configured independently before becoming part of the metacenter.) The next testbed will be specifically configured with the metacenter in mind.

- More information on the metacenter project.