


[Next Article](#)
[Contents](#)
[Main Menu](#)
[NAS Home](#)

# NAS, LLNL Release PBS

by [Elisabeth Wechsler](#)

NAS and Lawrence Livermore National Laboratory (LLNL) released the first phase of Portable Batch System (PBS) on March 1. The result of a two-year collaborative software development effort, PBS provides job dependency and job synchronization features, as well as improved batch job scheduling, for massively parallel computers, clustered workstations, and "traditional" supercomputers.



The NAS-LLNL development team expects PBS to eventually replace Network Queuing System (NQS) because NQS "doesn't deal well with the demands of massively parallel machines," said Dave Tweten, project lead for the NAS side of the partnership. With PBS, it's also easier to run jobs remotely.

PBS has captured the attention of the Portable Operating System Interface (POSIX) standards committee, which will issue an IEEE standard for batch processing, based on the PBS design, later this spring. The POSIX standard defines a software interface so that system calls for job management and I/O, utilities, and job control language are the same for machines manufactured by any vendor.

## Expands Batch Processing

For NAS, the greatest benefits of PBS may be to significantly expand batch processing for its three massively parallel machines—Thinking Machine's CM-5, and Intel's iPSC/860 and Paragon—and also to act as a "lightweight client" by giving commands to run batch jobs from workstations. With PBS, a server has one job scheduler and a resource monitor for each workstation to control batch processing on all the workstations in the cluster. Unlike NQS, the job scheduler and resource monitor are "invisible" to the user, Tweten said. "The user can launch a job at the server and the server runs it on whatever machine is appropriate."

"With NQS, DQS [Distributed Queuing System], and DJM [Distributed Job Manager], you can't move jobs across these systems because they're totally incompatible," said Bob Henderson, also a member of the NAS team. PBS allows the migration of jobs to the least-loaded machine that satisfies the resource requirements, he explained.

## Jobs Held for 'Meaningful Interaction'

PBS emphasizes operational flexibility with simplified system software and increased usability through extensive man-page documentation. Job coordination is improved because jobs running on different machines are now held by PBS until meaningful interaction is likely.

The job dependency feature of PBS allows users to submit a series of jobs, and the execution of each job is dependent on the success or failure of the previous job in the series. PBS also breaks larger jobs into a sequence of smaller jobs, Henderson said.

The job synchronization feature allows PBS to run two or more jobs with parallel processes that need to share data and communicate over the network simultaneouslyÑand both jobs can be started at the same time.

Each side of the NAS-LLNL partnership contributed to the design and development of PBS. NAS developed the network protocols, the client-server application, the resource monitor, and the machine-specific, job-launching software. LLNL was responsible for the user utility interface, the applications programming interface library, and the job scheduler. The NAS-LLNL team meets twice a month, alternating between the two Bay Area sites, Moffett Field and Livermore.

## A More Risk-taking Culture

"LLNL has a more software risk taking culture than NAS," Tweten commented. "LLNL is very interested in running PBS on their [CRAY Y-MP] C90." This receptivity to use experimental software in the research environment was a strong inducement to forming the partnership, he said. When Tweten asked his counterparts at LLNL about their software testing policy, he was told: "When the programmer feels the software is ready, we put it into production." Tweten noted that NAS is "reluctant" to take this level of risk.

In terms of personnel, NAS has provided Tweten, Henderson, and Tom Proett (all of the NAS Systems Development Branch), to work on software design; and John Musch (NAS Computational Services Branch), to manage the testing program for developers and NAS users, scheduled to run through the summer. LLNL contributed two full-time employees (spread among several part-timers) from its Central Computing Facility and National Energy Research Supercomputing Center.

## Release Schedule

The first release of PBS supports the Cray, Paragon, and CM-5 machines, as well as Sun and SGI workstations. A second release, due in May, will provide the NQS gateway for passing jobs automatically to PBS -- a major challenge for the project team, according to Tweten. The third release, due in August, will support workstation clusters.

As a result of the NAS LLNL collaboration, each institutional partner will "own" a copy of PBS, which they will distribute to their respective user communities. For more information, send email to

[tweten@nas.nasa.gov](mailto:tweten@nas.nasa.gov).

[Next Article](#)[Contents](#)[Main Menu](#)[NAS Home](#)