LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

UCRL-TR-226624

Purple Level - 1 Milestone Review Committee I/O and Archive Follow-up Demonstration

M. R. Gary

December 7, 2006

This work was performed under the auspices of the U.S. Department of Energy by University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

Purple Level-1 Milestone Review Committee
I/O and Archive Follow-up Demonstration
12/5/06

On July 7th 2006, the Purple Level-1 Review Committee convened and was presented with evidence of the completion of Level-2 Milestone 461 (Deploy First Phase of I/O Infrastructure for Purple) which was performed in direct support of the Purple Level-1 milestone. This evidence included a short presentation and the formal documentation of milestone #461 (see UCRL-TR-217288).

Following the meeting, the Committee asked for the following additional evidence:

1. Set a speed measurement/goal/target assuming a number of files that the user needs to get into the archives. Then redo the benchmark using whatever tool(s) the labs prefer (HTAR, for example). Document how long the process takes.

2. Develop a test to read files back to confirm that what the user gets out of the archive is what the user put into the archive.

This evidence has been collected and is presented here.

Speed Goal

We spoke with Bert Still, one of our largest archive users on Purple. While Bert was "very happy" with the rates to the archive from individual compute nodes (~105MB/s) we were interested in setting a target for login node transfers in order to simulate:

- the situation where users need to rapidly dump data from the GPFS file system to the archive in preparation for work on the file system,
- users having to quickly make room for newly computed data on GPFS,
- users archiving the results of a large computational effort.

Bert agreed that 1.5GB/s (which was also the goal of the original L2 milestone) was a good and appropriate target speed and was consistent with user requirements.

We also determined that it was the intention of the Committee that the demonstration closely mimic real-world conditions and that speed measurements be taken not as a sample of peak speed, but instead measure the end-to-end transfers of a large number of user files. In addition we chose to force the demonstration to compete with other archive and Purple activities rather than being performed during a downtime.

Files Used

We chose to use real user files, in this case those of Bert Still. We sampled 1,024 files totaling 668.6GB of data.

Demonstration

On Wednesday November 1st 2006, George Richmond performed the transfer demonstration using eight PFTP sessions across four Purple login nodes. The 1,024 user files transferred at an aggregate rate of 1.562GB/s. This test was a full end-to-end timing with the clock starting with the initial file launch and ending when the final file had been fully transferred. It was also concurrent with the production utilization of both Purple and the HPSS archive. Having reached the performance goal, George read back all of the files and did a bit-by-bit comparison. All files compared. In addition the files were force-migrated in HPSS and then pulled back from tape and again successfully compared on Purple.

Conclusion

Per the request of the Purple Level-1 Milestone Committee, we performed a test using 1,024 actual user files demonstrating an aggregate transfer rate from Purple to the HPSS archive of 1.562GB/s. This met our user and L2 transfer rate goal of 1.5GB/s to the archive. We performed comparisons of the files from HPSS disk cache and from tape and verified that the data had transferred correctly.

We reviewed the demonstration procedures and results with Bert Still who concurred that all demonstration goals were met.