

PANASAS® PARALLEL STORAGE: HPC BENEFITS FOR COMPUTER AIDED ENGINEERING APPLICATIONS

The combination of scalable application software and HPC clusters with Panasas storage technology has demonstrated new and significant productivity advantages for workflows in CAE applications. The benefits range from dramatic HPC cost-performance improvements to “non-practical” high-fidelity CAE simulation solutions on commodity clusters.


INTRODUCTION: I/O BOTTLENECKS IN THE CAE WORKFLOW

As the CAE industry continues an aggressive platform migration from proprietary Unix servers to commodity HPC clusters, CAE models also become more realistic, requiring clusters to handle ever-increasing volumes of I/O and the movement of large files. Today, as organizations rapidly expand their cluster deployments as a cost-effective way to grow HPC resources for CAE demand, many encounter I/O bottlenecks when using legacy network attached storage (NAS) architectures.

Initially these NAS systems offered advantages such as shared storage and simplified IT administration which reduced costs, but today few provide the scalability required for effective I/O performance in parallel CAE simulations. Scalable clusters and parallel CAE software have advanced over the years, while NAS scalability has not progressed. Recently a new class of shared parallel storage technology was developed to remove serial bottlenecks and scale I/O, therefore extending the overall scalability of CAE simulations on clusters.

Panasas parallel storage is the leading solution of parallel NAS, that enables the most advanced and I/O demanding CAE challenges to become practical applications. Examples included high-fidelity transient CFD, large eddy simulation (LES), aeroacoustics, large DOF structural dynamic response, parameterized non-deterministic CAE simulations for design optimization, and the coupling of CAE disciplines such as fluid-structure interaction (FSI).

CAE workflows are overburdened with lost productivity when engineers and scientists must wait for serial I/O operations and large file transfers to complete. Further, as simulation and workflow performance degrades, so does CAE analyst efficiency and effective workgroup collaboration. Panasas eliminates the I/O bottlenecks with a cost-saving solution that restores productivity and drives analyst creativity.



FEATURES AND BENEFITS

PanFS Shared Parallel File System:


- **Enhanced CAE Workflow**
Eases collaboration tasks such as pre- and post-processing. Drives productivity into CAE throughput requirements of multi-job instances and coupling of CAE disciplines.

Panasas DirectFLOW® parallel access protocol and support for NFS and CIFS:

- **Maximizes Performance**
Shared storage for all data and platforms, with maximized performance for each protocol that is used for I/O transfer.

Single Global Namespace:

- **Reduces IT Overhead**
Simplifies storage and data management to provide seamless scalability as CAE model size and number of jobs grow.



PANASAS SHARED STORAGE AND PARALLEL FILE SYSTEM

Panasas developed a parallel storage architecture that combines key advantages of legacy NAS, yet eliminates the drawbacks that have made them unsuitable for large distributed clusters. Panasas parallel storage offers the high-performance of direct access to disk, in addition to the benefits of shared files and metadata. Just as clusters are designed to distribute computational work evenly across compute nodes, Panasas parallel storage is designed to distribute data evenly across storage devices for shared parallel data access directly between cluster nodes and the parallel file system.

As CAE analysts continue to increase the number of compute cores per simulation to keep pace with growing CAE model fidelity, I/O operations must be performed in parallel to realize the essential benefits of overall simulation scalability. With the Panasas ActiveStor 5000 parallel storage, each node on a cluster has direct access to the shared storage and file system for concurrent data reads and writes during a computation. Once the simulation is complete, end-users have direct access to results on the same file system for post-processing and visualization.

EXAMPLE: FLUENT TRANSIENT CFD AND PARALLEL I/O

The benefits of Panasas parallel storage are demonstrated in an example of a transient CFD model using FLUENT, the market leading commercial CFD software from ANSYS Inc. (www.ansys.com). An engineering collaboration between ANSYS and Panasas that began in 2006 has developed a parallel I/O scheme for FLUENT 12 (mid-2008 release), based on MPI-IO under the MPI-2 standard. Such capability will enable overall FLUENT scalability for even the heaviest I/O applications, including large steady state models with frequent checkpoints, and even more challenging transient CFD that can often require 100 solution writes per single simulation.

A production case of a 90 million cell external aerodynamics model from FLUENT 6.2 (with serial I/O) was used as the base-

line for the parallel I/O comparison. The case was run for 500 iterations on a Linux cluster using 64, 128, and 192 cores with a Myrinet interconnect between cluster nodes, and using a conventional NAS with NFS. The steady state case (single input read and single solution write) exhibited increased I/O time as processing cores increased from 64 to 192. The overall effect is that total solution time increased with increasing cores, causing performance degradation at 192 cores and beyond. For the corresponding 90 million cell transient case requiring 100 solutions writes, the total time at 192 cores was estimated to increase by 36-fold over the steady case, which was considered impractical and never attempted.

The benefits of FLUENT 12 and parallel I/O (vs. 6.2 serial) were demonstrated with concurrent writes of the local domain solutions into the global solution file on the Panasas parallel file system. FLUENT 12 for the steady (1 write) case at 192 cores improved the total time by nearly 3-fold over FLUENT 6.2, and more importantly the simulation continued to scale beyond 192 cores. This gain means 3x additional utilization of existing FLUENT software licenses already purchased, along with the ability to grow model size and complexity without limitations on I/O scalability. An even greater benefit was demonstrated in the once-impractical transient (100 write) case. FLUENT 12 with parallel I/O only required a 50% increase in total time at 192 cores going from 1 write to 100 writes, which was estimated to improve the FLUENT 6.2 transient case total time by 64-fold!

INVESTMENTS IN ISVS FOR THE CAE WORKFLOW

Often CAE workgroups have separate storage infra-structures for various stages in a CAE workflow of computation and pre/post-processing. Panasas storage offers a unified approach to data access and storage management for all simulation and workflow tasks. Panasas understands that CAE workflow solutions require close ISV collaboration and have developed alliances with industry leaders Altair, ANSYS, CD-adapco, ESI, LSTC, Metacomp, MSC.Software, and SIMULIA. Contact Panasas today to learn more about the benefits for your CAE workflow and HPC objectives.