Taiwan Chemical Engineering Research Team Uses NeoSapphire All-Flash Arrays to Accelerate Parallel Computing Performance and Increase Research Efficiency

Customer Profile

The Taiwan Chemical Engineering Research Team focuses on chemistry research and testing various materials, with fundamental research into innovative surfaces, contact surfaces and nanoscience at the core. It is also involved in the research fields of renewable energy, molecular science and engineering, and environmentally friendly, clean production technologies. The scope of its research is fairly extensive, and it is one of the leading teams in Taiwan's academic chemical engineering community.

Challenge

The research team carries out painstaking research into various chemical materials and tests them repeatedly, and will often use various computational chemistry software to prove research assumptions and certainties. These calculation results can explain a lot of chemical experiments, and can also predict the intermediate products produced during certain tests, which helps in understanding the reaction mechanisms of substances. As a result, the research team needs to make extensive use of quantum chemistry and materials science simulation software, such as Gaussian, $DMoL^3$, CASTEP and VASP. However, as the complexity of the experiments increases, and due to different researchers needing to share the same IT equipment, it often takes a long time to produce experiment results. If the research results need to be adjusted or verified after they are produced, more time is needed for further computation. AccelStor Senior Product Manager Alex Ho, who is assisting this research team, points out: "Competition within the academic community is actually the same as ordinary commercial competition. The research team is able to present experiment results and publish these in academic papers at an earlier stage. However, the IT resources in a lot of Taiwanese laboratories are quite limited. Sometimes one set of IT equipment has to be shared by professors or students with different research topics. This is already relatively time consuming. If IT equipment is inefficient, this will be even more of a burden on laboratory members."

The research team initially increased system computing performance by increasing compute nodes. However, after increasing to 15 compute nodes, there is still an obvious performance bottleneck. AccelStor Senior Product Manager Alex Ho says: "One side effect of increasing compute nodes is that the quantity of data in random read/write operations originally generated by 1 compute node increases significantly after increasing to 15 compute nodes. Random read/writes are the most difficult operation for a storage device to process, and can also cause serious computation bottlenecks." In particular, when using computer software that requires a lot of data access, such as Gaussian or $DMoL^3$, system performance will be unable to be upgraded due to waiting for data access. This causes the research team to waste valuable time waiting for computation results.

Solution

Following discussions between AccelStor and the research team, it was realized that the laboratory IT equipment did not actually lack computing power, but that the performance slowdown was in the back-end storage devices after all data had been computed and processed. The traditional hard disk storage devices originally used by the laboratory were built using an SAS HDD RAID 0 configuration. The laboratory used 15 compute nodes and integrated these with the hard disk storage devices using an NFS protocol. AccelStor suggested replacing the back-end storage devices with a NeoSapphire 3401 (NS3401), and adding a 10GbE network switch to increase the data exchange bandwidth (see Figure 1). The NS3401 is an all-flash array that uses a solid-state drive (SSD) as a storage unit.

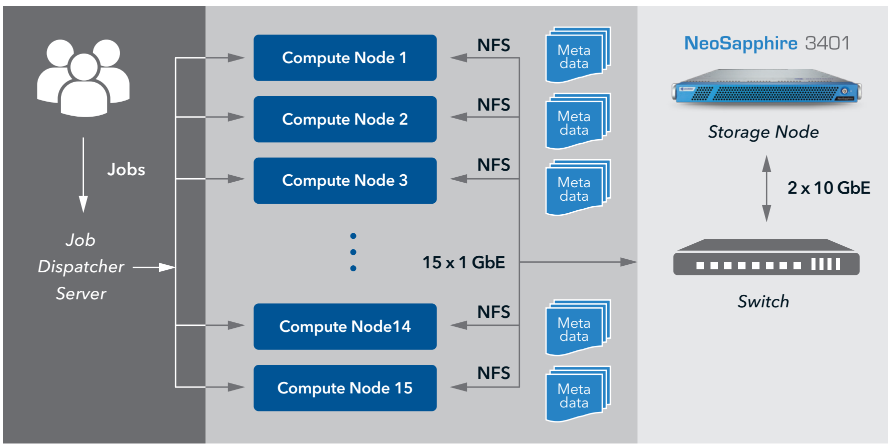

Figure 1: Laboratory IT equipment architecture

The main feature of an SSD is that its access speed is faster than that of a traditional hard drive. The NS3401 is also equipped with AccelStor's exclusive FlexiRemap software technology, which targets the random write data that is likely to cause system storage performance bottlenecks, and improves performance acceleration operations.

Benefits Introduced

Using the Gaussian 09 software frequently adopted by the research team, AccelStor applied the coupled cluster method to compare the overall computing time of a traditional hard disk array with the NS3401 all-flash array under the same structural model. The traditional hard disk array required 32 minutes 25 seconds, while the NS3401 only required 5 minutes 47 seconds, making it 5.4 times faster. Even when using different calculation methods, the NS3401 still has a significant lead. This is because using the NS3401 to resolve storage bottlenecks in the original architecture enables the compute node CPU to fully implement stoichiometry. The AccelStor team also trialed and compared different quantities of compute nodes. When there was only one compute node, the traditional hard disk array's performance was about the same as that of the NS3401. However, when compute nodes were increased and their quantity of data also increased, it was obvious that the NS3401 could still execute all operations within a certain period of time, while the time required by the traditional hard disk array increased significantly (See Figure 2). Obviously, after system compute node clustering increases, system performance bottlenecks will often be transferred from computing node to storage node. The level of scattered data also increases due to the growing number of the nodes. Traditional hard disk arrays will usually become performance bottlenecks.

In addition, there was a high level of scattered data write when using the original layout of multiple compute nodes, making it more likely for write times to slow down and system life to be shortened. AccelStor Senior Product Manager Alex Ho explains: "Laboratory IT equipment budgets are actually extremely limited. Although the cost of using a single all-flash array is relatively high, it really can reduce the enormous cost of building compute nodes. It can also significantly reduce future operating costs, including electricity bills. Combined with its reduced computing times, it can assist laboratories in performing more analysis more accurately."

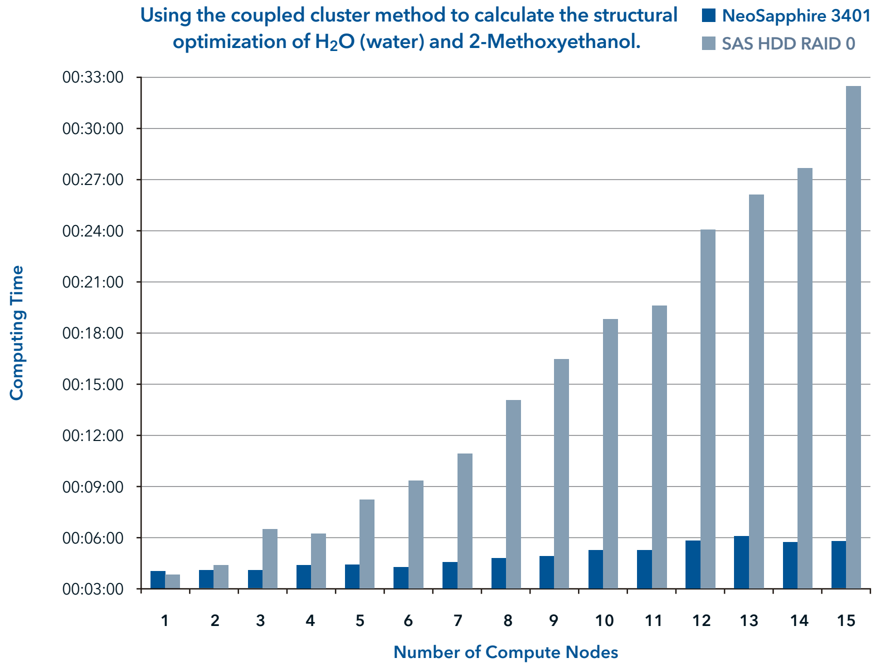

Figure 2: Times required by NS3401 and traditional hard disk array for different numbers of compute nodes when running Gaussian 09

Conclusion

With the quantity of data being analyzed increasing on a daily basis, the academic community's requirements for research accuracy are becoming increasingly stringent. This has intensified the current burden on research organizations' IT equipment. Academic institutions particularly hope to take advantage of the gains from today's high-performance computing (HPC) to help increase the quality and quantity of their research. AccelStor Senior Product Manager Alex Ho reveals: "Of the academic institutions we've been in touch with, there isn't a single one that doesn't want to improve the speed and quality of their current data analysis. The NeoSapphire 3401 has demonstrated significantly improved system performance in this area, and it can also help to achieve excellent budget control over the construction of an entire machine room."

Accomplishments

- The overall computing time of a traditional hard disk array was compared to that of the NS3401 all-flash array under the same structural model. The NS3401 was 5.4 times faster
- Significantly reduces overall system rollout costs and future operating costs
- Resolves storage device performance bottlenecks, releasing the existing performance of compute nodes

Products Introduced

- NeoSapphire 3401 All-Flash Array