INFOSTOR®
news analysis & trends

readerexpert I/O

Sharing eases angst about slow performance

Q: Our research group is running analyses on files ranging from 50MB to several hundred megabytes in size. Performance was fine when a single computer was able to do all the work, but as we've moved the files to a file server and added more computers, performance has become unbearably slow. Can we use storage area network (SAN) technology to improve file-server performance?

Jacob Farmer is the CTO of Cambridge Computer Services. He can be reached at jacobf@cambridgecomputer.com.

This is a common problem for research—as well as video editing, medical imaging, and pre-press—applications. In short, your file server is buckling under the demands of your client computers.

The best way for you to improve file-sharing performance is to share a disk device over a SAN rather than using a network file server. SANs make it easy for multiple computers to connect to the same disk system over a high-speed connection.

Although conventional file systems and application software were not designed to share disk resources, with the aid of special software it is possible for multiple computers to share files over a shared disk connection.

One way to do this is to use *volume-sharing software*, which lets multiple client computers read a common disk volume, but typically only allows one to write the data (though write privileges can often be passed among computers). If you can live with this limitation, then this may be a good, inexpensive solution for you since your performance should be pretty close to that of direct-attached storage (DAS).

Volume-sharing software tends to cost less than $1,000 per seat and some products even allow sharing across platforms (e.g., between Windows PCs and Macs). Be sure to check for compatibility with your operating systems and service packs before you buy.

Also, if you only have a handful of computers, you might not need to implement a Fibre Channel SAN; a multi-port SCSI disk array may be sufficient.

The next step up in the food chain is a class of products I would classify as *SAN file systems* (see Reader-Expert I/O, *InfoStor*, August 2002, p. 6). These products replace the conventional file systems on your computers and allow full read-and-write file-sharing to a common storage device. Some SAN file systems are cross-platform, but you'll want to make sure that your exact OS version is supported.

SAN file systems come in out-of-band (metadata and regular data travel over different paths) and in-band configurations. Out-of-band products (sometimes called multi-path file systems) are actually pretty mature, and they often come with storage device management features common to disk virtualization products. So, you get management benefits on top of performance improvements.

In-band file systems do more or less the same thing. Which type is better for you depends on a variety of factors—most notably, the task at hand. So, if you are shopping for a SAN file system, you will want to select one that meets your current application criteria.

And, finally, you could implement a *NAS server that is based on a SAN file system*. Such a beast would actually be made up of multiple computers that all serve the same file system. Think of it as a multi-node, clustered file system. The device would have the combined throughput of multiple file servers, but it would serve the same set of data. In addition, it would be fault-tolerant (i.e., if a node failed, clients would roll over to another node without interruption).

The benefit of buying SAN file-system technology that is packaged as NAS is that it's all pre-qualified and turnkey. All you would need to do is connect the NAS head to your SAN and LAN and off you go. Cross-platform interoperability is achieved through network file access protocols such as NFS and CIFS.

The one downside of this type of approach is that NFS and CIFS introduce some latency—perhaps too much for your application. A good test would be to see if your application runs okay with just one or two computers reading files over NFS or CIFS. If the performance is acceptable, then that means that network file services are not the bottleneck and that a super-fast file server would actually solve your problem.

NAS heads built on SAN file systems are an imminent reality. I've actually been kicking around live products in my lab. Stay tuned on this subject. You will see some great new products as early as this fall. ❑

If you have a question you would like to ask one of our experts, please e-mail Heidi Biggar at heidib@pennwell.com.