Grid Computing Report

Robert Gravina

February 27, 2007

1 Evaluation of Kosha

Kosha [1] was created from the observation that many desktop and server machines in a typical corporate or academic environment have large unused disks which could be "harvested" for storage space and used in applications such as grid computing. To that end, it works well, with only a small overhead over NFS (around 6% overhead in performance for 16 nodes) and a reliability of 99.991% when using three levels of redundancy.

1.1 Cost-effectiveness

One important point to consider, which was raised during the presentation, is whether or not such a system makes sense from a cost perspective. In other words, would it be cheaper to simply equip NFS servers with larger disks when taking into consideration the power consumed by nodes running Kosha?

Assuming that machines running Kosha were kept on for that purpose alone, then the answer would most certainly be no. A typical university teaching lab or small business may have several hundred nodes which would be suitable for running Kosha, but these desktop machines would otherwise be turned off when not in use, or running in low-power (sleep) mode.

One thing to consider, however, is the vast amount of storage available on desktop machines. The authors of Kosha determined that of some 500 instructional machines at Purdue, 84% have 40GB drive of which 90% is free space. That's around 15 Terabytes of unused space! Comparing this to other ways of storing this data:

Drives Current drive costs are about US$400 for an Hitachi 1 Terabyte drive [2]. Assuming such drives are easy to add to a NFS server or computing cluster, drives can be added as necessary. Of course, if fault tolerance is required more expensive RAID systems would be necessary.

Services Services such as those provided by the San Diego Supercomputer Center (SDSC) [3] will host 1 Terabyte for around US$1500 per year. The additional costs over drives alone factor in the cost of maintenance, running servers, paying staff, purchasing backup media etc.

If machines running Kosha were used during periods where they would be on anyway (for a university computer lab, that may be from the morning to late at night) it could be a potentially cheap way of storing several terabytes of data. When these computers are switched off in the evenings, provided they are not all switched off at the same time, data could be migrated to higher capacity NFS servers. However, the system would not be feasible if all of these computers were powered off in the evenings,

since the NFS drives would have to be large enough to host data normally on these nodes, defeating the purpose of harvesting the storage space on these desktops.

Using Wake-on-LAN technology could enable Kosha desktops to consume little power while serving Kosha requests when required. If the Kosha mount point sees little activity overnight, then desktop nodes will not be woken.

Peer-to-peer distributed file systems may not be the most effective way to reliably serve large amounts of data to many nodes in all cases, but there clearly is a use case where this is viable, namely one with many nodes which would be on regardless of Kosha running or not, and more recently machines which can be woken up (or even booted) when requests are received over the network

References

[1] A. R. Butt, T. A. Johnson, Y. Zheng and Y. C. Hu, "Kosha: A Peer-to-Peer Enhancement for the Network File System" *Journal of Grid Computing* no.4, (2006): 323-341

[2] M. Kanellos, "Here comes the terabyte hard drive" *News.com* (Janurary 4, 2007): `http://news.com.com/2100-1041_3-6147409.html`

[3] J. Schiff, "Brother, can you spare a terabyte?" *StorageForum.com* (September 28, 2006): `http://www.enterprisestorageforum.com/hardware/features/article.php/3634881`