

Computer Technology Review®

TECHNOLOGY SOLUTIONS FOR SYSTEMS INTEGRATORS, VARS & OEMS

Grid Storage for Grid Computing

By Eric R. Schott

Grid computing has evolved into a well-understood architecture that provides users and applications the unimpeded use of a large pool of IT resources. The computer grid connects computers (servers or desktops) and storage across an enterprise and virtualizes heterogeneous and geographically dispersed components into one useable system. The goal of grid computing is to allow the sharing of computing and data resources for a number of workloads and to enable collaboration both within and across organizations. Because grids unleash latent power of installed resources, they can give companies a huge gain in productivity, significantly accelerating compute-intensive processes.

Beyond the acceleration of large-scale computing tasks, grid computing allows for the dynamic assignment of workloads to the most efficient resources for the job, so a large number of tasks run more effectively. Grid computing hides complexity, affording all of its users a single, unified experience. Users are allowed to share not only file data but computing resources, enabling

full collaboration toward common goals with its more flexible, open, and resilient operational infrastructure. With grid computing, users have immediate access to compute and data resources in a fluid environment that responds fluidly to their demands.

Storage requirements for the compute grid

The storage requirements for such a dynamic computing environment are demanding. An enterprise compute grid requires transparent access to a similarly dynamic and scalable pool of storage. Storage for grid computing requires a global filing system in order to present a single storage space to the workloads - providing both transparency of location as well as access to data. In fact, many storage grid products in use today are really network-attached storage (NAS). While NAS provides transparency, this solution limits what applications can utilize the grid; it limits scale and storage management capabilities. While file systems seem to be a common abstraction, there is still a lack of standardization and compatibility, since different grids implement different file system protocols - proprietary, NFS, or CIFS. This limits choice in the

types of operating systems and applications that participate in the grid. The best solution is for the grid to standardize at a lower level of storage, allowing choice of file system in the same way the grid allows choice in applications.

While the discussion about grid file systems is important, the nature of the storage underneath them is equally critical, if not more so. A global file system provides access to the storage layer, but it does not control its behavior. It is clear that direct-attached storage (DAS) is not enough and that some kind of consolidated storage is called for. So what are the necessary attributes of the storage layer, and what is the best solution for grid computing?

To understand the ideal attributes of storage for grids, it is important to look at the way grid computing operates. A compute grid is, in essence, a job shop. On any given day, users log on and give it something to do. Sometimes these jobs are very large and require the collaboration of many components of the grid, drawing on all available resources, compute cycles, bandwidth, and storage as needed to get the job done. At the end of the day - or at the end of the week - when the task

is completed, the application withdraws its demands and leaves grid resources open to the next job. This goes on constantly, with large and small jobs running throughout the grid for variable periods of time at fluctuating levels of intensity - so the impact on the storage infrastructure is random and chaotic.

The global filing systems provide transparent access to storage, but are not equipped to provide optimized and efficient access. As workloads change dynamically throughout the grid, so must storage. Data location (provisioning, volume creation) and load balancing need to take place on the fly - responding to the demands of the applications initiating, running and completing their tasks. The storage resource not only must be virtualized, but it also must be intelligent and self-managing in order to continually provide optimal data services to the computing tasks that rely on it.

The evolution of the storage grid

Over the years, there has been much industry discussion about the future direction of corporate computing and the evolving architecture of a new, more responsive data center.

Going by many titles - on-demand computing, utility computing, grid computing - they all point to the same vision: a virtual pool of data center resources dynamically responsive to user demands. How this vision evolved is no mystery. Any IT manager can tell you that what has become standard operating procedure in the conventional data center is tedious hands-on work. So the trend towards consolidation and virtualization is already underway, driven by a real business requirement for efficiency and responsiveness. And nowhere is this hunger for simplicity and automation more apparent than in the area of storage.

Businesses that could afford it have been consolidating servers and migrating from DAS to storage area networks (SANs). Fibre Channel SANs were deployed in large, well-staffed data centers to centralize enterprise storage to ease data management and data protection for a few servers. However, the initial phases of server and storage consolidation did not do much to virtualize, simplify and automate basic storage management functions. If anything, Fibre Channel SANs introduced as much new complexity as they eliminated. Provisioning storage for new applications, load balancing, creating volumes, configuring RAID sets, and so on, remained a time-consuming, manual task. Not to mention the expense of Fibre Channel connectivity, including switches and HBAs, and the cost of training staff to administer the SAN.

With the advent of iSCSI-based SANs, IT managers have

an opportunity to introduce the simplicity, scale, and ease-of-use that IT managers have wanted for some time - and, beyond that, iSCSI will accelerate the development of the storage grid.

In taking a hard look at the frustrations of IT managers with traditional, monolithic storage solutions, innovative storage system engineers have adopted key concepts from grid computing and applied them to iSCSI SANs - delivering modularity, automation, scalability, and data protection right in the storage. In the process of addressing the most pressing frustrations inherent in traditional SANs, the IP storage grid was born.

Grid storage for grid computing

Given the unique demands of the compute grid on its storage infrastructure, storage for the grid must be unusually flexible. DAS is simply not an option. Virtualization is a start, providing the single unit behavior that the global filing system requires to present data stores to the compute grid, so SAN architectures are called for; however, the scale of these SANs is beyond the capabilities of Fibre Channel. With scale comes many management problems that constrain the fluid operation of the grid. New management paradigms are needed to remove manual administration and the over-provisioning required today to stay ahead of the constant demand for more storage and the changing workloads of the grid.

The ideal enterprise-class iSCSI-based SAN built for modular scalability, on the other hand, provides many

advantages for grid computing. First, its IP-based communications protocol provides scalability, ease of management, and integration of the storage grid, which is connected throughout by a ubiquitous Ethernet network. Second, it provides a single, centralized pool of storage that can expand as needed with the addition of instantly integrated modules with no application disruption or performance degradation. Third, much like the compute grid, the storage grid acts like a single resource, intelligently sharing the workload to ensure that no single part of the storage grid is overloaded.

Beyond these basic requirements - all of which are borrowed from the grid concept itself - the storage grid must be designed for high availability, with redundant and hot-swappable components as well as automated snapshot and replication capabilities for continuous data protection. The compute grid cannot go down, so the storage grid must provide the highest levels of data protection.

Another highly desirable feature of the storage grid is its support of SAN boot, which can completely unfetter servers from a specific task and allow them to be deployed and redeployed as needed for whatever application they are best suited. When a server can boot from the SAN, it is liberated from local disk and can boot from any volume on the storage grid. This introduces further flexibility and protection within the grid by allowing malfunctioning servers to be replaced quickly with little or no downtime.

Grid storage for everyone

Just as iSCSI has made enterprise-class SANs and storage management accessible to everyone, it has also made the enterprise storage grid a reality for small and medium-size businesses as well as for large enterprises. All of the attributes of the storage grid that are a necessity for grid computing make everyday data center operations a lot easier.

A paradigm shift is occurring, but not just because of iSCSI, though it was a catalyst. The old monolithic storage arrays have been challenged by a new generation of robust, enterprise-ready modular arrays with intelligent management tools already built in - these new arrays happen to have iSCSI for a communications protocol because it serves their ease of use and scalability. But what gives this new breed of storage platform its real value is its grid architecture that offers enterprise-quality performance and reliability, intelligent automation, and seamless virtualization of a single, scalable pool of storage. These new IP SAN solutions are modular and affordable for small and medium-size businesses to deploy in increments while being robust enough for large enterprises to scale to hundreds of terabytes.

Eric R. Schott is director of product management, EqualLogic, Inc., Nashua, NH

www.equallogic.com



EQUALLOGIC[®]

SIMPLIFYING NETWORKED STORAGE™

EqualLogic, Inc.

9 Townsend West

Nashua, NH 03063

phone: 603-579-9762 / 888-579-9762

fax: 603-579-6910

info@equallogic.com