A Practical Application for Grid Computing

A White Paper Prepared for Network Appliance
June 2007

EMA

Table of Contents

EMA

Introduction – Utility and Grid

Utility computing and utility storage are relatively new concepts. The meaning of "utility" derives from the way the word is used in most other circumstances: utility is a way of both spreading the cost of some necessary service over all of that service's users, and of providing more of those services than the individual users could ever provide for themselves. Because, for example, it is impractical for individual electricity users to build their own power plants, all contribute towards building and maintaining a common facility that provides power wherever and whenever a need arises. Costs and benefits are shared across the user population, with per-use charges perhaps being applied based on how much value each user extracts from the system. The value of such an on-demand system in a very real sense exceeds the sum of its parts, and in fact allows "oversubscribing" to services; a 100 megawatt power plant, for example, is quite capable of providing 100 megawatts of power to each of several sites – if those sites can schedule their use of that power in an efficient manner. As long as the demands on the system do not exceed the system's capabilities at any point in time, the utility functions well.

Power is distributed from plants to users over what is referred to as a power grid, which consists of switches, cabling and other necessary infrastructure. The power grid brings additional resources online when they are needed, perhaps tapping into supplementary power plants as necessary, and distributes the additional power to the users for as long as they need it. When demand decreases, the added resources are taken offline. In precisely the same way, computer grids[1] distribute computing assets and services to processes that require them and, as with power grids, enable oversubscription when efficiently managed. Resources such as computing power, storage, and communications capabilities are provisioned and de-provisioned as they are needed, withdrawing them from use and returning them to the common pool so that they may be shared by other authorized users.

Grids then are shared infrastructures, utilizing a combination of components that includes servers, storage, databases, networks, and management software which can be applied to and removed from specific computing problems as circumstances warrant. They form the underpinning of an on-demand computing environment, and provide the connectivity that ensures the total value delivered truly is greater than the sum of the individual parts. A properly constructed grid provides all necessary computing tools and infrastructure requirements for its applications, constantly adapting to the applications' changing requirements.

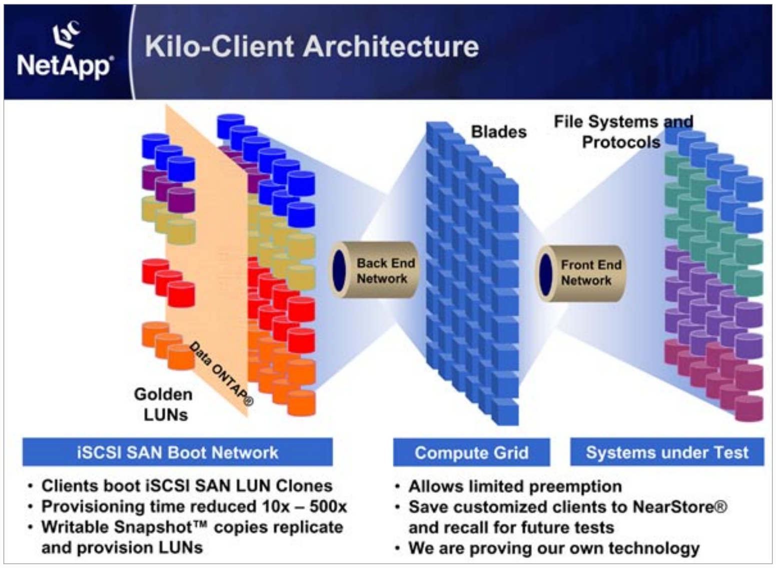

Figure 1: Kilo-Client Architecture

Commercial Grids?

Grids today appear mostly in high performance computing (HPC) environments, where hundreds of inexpensive compute nodes link together to solve parallel processing problems that previously could only have been addressed by using supercomputers. Storage grids were developed to supply the enormous amounts of data that such processing power was capable of producing and consuming. Properly managed, storage and throughput scale linearly with computing power.

Grids are a proven technology when it comes to HPC, but such sites are typically seen as a special case of computing, where computational performance is often the only criterion that really matters. Circumstances are

[1] Vendors and users use the term "grid" in many ways, some of which are inconsistent at best and misleading at worst. In this document, grid is defined generically as an infrastructure that optimizes the way distributed computing resources are shared across an organization, applying and removing assets to a problem so that computing, communications and storage all scale in a linear manner.

quite different at commercial sites, many of which are already implementing on-demand computing environments and would certainly receive value from what grids have to offer. Here however, while performance is still important, so are a number of other issues. Particularly significant is the need to strike the right balance between performance, operational efficiency and flexibility.

Many things contribute to operational efficiency, including ease-of-use and ease-of-management, system reliability, and the economical use of electrical power. Ideally, any complex system (a grid surely qualifies) shields both users and administrators from its own complexities. This means that the "plumbing" supporting applications should be transparent, and that administrators should have a simple management interface – one capable of being managed by locally defined policies – interposed between themselves and the complex assets they manage. System reliability should be commensurate with the required service levels, and power consumption should be minimized.

The NetApp Kilo-Client Lab

Network Appliance's KiloClient Lab at its Research Triangle Park, North Carolina site houses what is probably the world's largest commercial grid. Although it still serves its original purpose as a test-bed for large on-demand installations, KiloClient has gone well beyond the company's original objectives and today has evolved from a proof-of-concept experiment into a center for testing real-world applications under extreme on-demand conditions. Fundamental to KiloClient are the following:

Standard components mean less expense. Although originally designed as a test-bed for the shared development and testing of NetApp products – the company's new RAID 6 offering was developed here – the site also provides the company with an opportunity to test utility computing approaches in real-world environments, and to do so on a scale unavailable to other vendors. The site uses only commodity computing components (blade servers) rather than specialized hardware, and runs multiple OSs[2], which in turn connect to NetApp storage using iSCSI[3] over a 10-Gbit Ethernet backbone. Sixteen BladeCenter cabinets hold 4 chassis each, with 14 blades per chassis (blades within a chassis can fail over to one another) and 2 Ethernet ports per chassis. iSCSI was selected because of its ease of use, it runs over existing infrastructure, and it delivers the requisite performance and scalability at minimal cost.

Podules™. KiloClient at present is divided into six "Podules," groups of nodes sharing a common infrastructure of an Ethernet switch, a NetApp 6070HA (used as a boot appliance), and 16-20 chassis. Currently, each podule accommodates from 224-280 nodes, although a podule may be any size an application requires; the practical limits on the podule's size determined by available network connectivity and the bandwidth requirements of the applications that are running on it. Each podule is a self-sufficient subset of the overall grid, operating either independently or in concert with the other podules to service changing demand loads. Grouping similar components within each podule allows Network Appliance to streamline management of the podule infrastructure.

Automated provisioning. In on-demand environments, workload requirements change dynamically as various processes come online and go offline, a demand flux that requires efficient provisioning from the grid's virtualized compute and storage pools. Provisioning, a complicated and repetitive process that frequently falls victim to operator error, is only practical on such a large scale when the assets are accessed and de-accessed automatically. While policy-based automation is still very much a work in progress, several innovations make grid provisioning a much more streamlined operation. Automating 85% of the provisioning process appears to be an achievable goal. At this writing, Kilo-Client automation is estimated to take care of 60% of the grid's provisioning needs.

Centralized, optimized monitoring and management. Kilo-Client is far more flexible than typical IT systems. One centralized management console looks out over the whole of the grid, overseeing automated processes to ensure optimal asset use. As a result, management efficiency extends well beyond provisioning. Managers may apply a variety of Network Appliance management tools from this console: system images and

[2] Presently, KiloClient is running the following operating systems: RHEL 4, RHEL 5, SUSE Linux, Windows (2000, 2003, XP, Vista), VMware ESX, Solaris 10, DataOnTap simulator.

[3] NetApp has chosen iSCSI because of that technology's ease of use, because it runs over existing infrastructure, and because it delivers the requisite performance at minimal cost.

dataset clones may be moved from one part of the grid to another using SnapManager's snapshots[4] and snap restores, which offer point-in-time copies and restorations of whole databases with no performance penalties; when the need arises for different data transfer protocols, the default iSCSI environment may be withdrawn and replaced with NFS or FCP protocols; remote systems, both iSCSI and FibreChannel, may be brought online as demands shift; boot-from-SAN capability allows a single remote manager to recover servers quickly.

Golden images. When customers arrive to try out the on-demand environment, they have a choice of using one of the site's pre-configured "golden images" - predefined combinations of hardware, software, protocols, infrastructure, and applications typical of many commercial sites – or of building their own environments from scratch. Each golden image can be redeployed easily and quickly for future testing.

"Greening" of the grid. "Greening," the term given to reducing power demands in the IT room, has obvious ethical and business implications. KiloClient's novel approach to limiting cooling expenses moves beyond the current concept of hot and cold aisles to the fundamental design innovation of adjoining hot and cold *rooms*, designed to work together so that convection provides increasingly effective and cost-efficient temperature control. Thin provisioning, as it enables a relatively small number of storage devices to service a comparatively large number of client processes, "greens" the grid still further.

EMA Perspective

Enterprise Management Associates (EMA) looks at technology developments daily in the course of its research; KiloClient, however, is something quite apart from what we ordinarily see. Few companies actually do much real research these days, and fewer still are able to bridge the distance between their "long-view" research and their current development projects. Of those, only a handful makes their on-going Research and Development directly available to their partners and customers. KiloClient falls into this last group. It has clearly come quite a way – in its early days a test-bed, it now forms what may well be the heart of the ecosystem within which NetApp, its partners and customers all work.

[4] Snapshots capture the state of a system at various points in time. Network Appliance management software supports up to 255 snapshots.

KiloClient is a mix of technologies and services. It is built mostly from existing products, many of which come from NetApp, but it also has a broad mix of commodity products from other vendors that are common to data centers across the world today. The Ethernet backbone, for instance, is hardly exotic and has been a part of commercial data centers since the mid-1980s, while the iSCSI protocol running across the backbone is now the SAN protocol of choice for many data centers. Into this mix NetApp adds its own newer products (particularly in the area of management software), and some technologies that have not yet been released.

Certainly the way KiloClient brings the various grid components on- and off-line is a fine way to demonstrate the flexibility of the NetApp base product line – after all, many other vendors' product reconfigurations are a task for the brave and are not to be undertaken by the faint of heart. Here, that effort is taken care of by an automated central management console that oversees every node on the grid, and is greatly enhanced by NetApp's "thin" provisioning strategy. Because thin provisioning allocates minimal amounts of storage to each user or process (as opposed to the time-honored but wasteful over-provisioning IT has lived with for 30 years) less storage is wasted and more is available. The NetApp method saves on both initial investment and on power consumption down the line, which means that scaling is not just easier; it builds in added efficiencies over time.

All this flexibility enables configuration testing for numerous architectures, including both the ones NetApp itself needs to test and those needed by IT managers who need to test an environment that under other circumstances would be impossible to test. If NetApp's own "golden images" don't do the trick, customers can come in and build brand new configurations that accurately represent their own sites' anticipated future growth. The limits on how far this can scale out are unclear, but the grid continues to grow, and is already far larger than all but a few of the world's commercial data centers.

Most IT managers can only speculate on (or at best, model) this kind of growth; those having near-term strategic decisions to make regarding their storage, however, may need more. Now, rather than having to assemble the necessary hardware, software and staff resources, and then find some time when all are available at once, they have an alternative. For them, the KiloClient grid stands alone in the industry as an opportunity to do empirical testing before corporate managers make strategic investments. More than that, it also offers them the opportunity to quantify the time and resources needed to service a particular set of stakeholder processes, enabling them to quantify the value they can deliver to stakeholders, to understand what the impact of each change to the system will be, and to determine in advance what resources will be required to support each test process, and to calculate what the chargebacks are likely to be.

In sum, KiloClient is more than computing on steroids, although that is impressive enough considering it is built mostly from commodity parts. The real significance, however, is that this is a technology that delivers the day-to-day business efficiencies hoped for from utility computing, and makes those efficiencies not just a theoretical possibility, but makes them a practical consideration for commercial data centers. Enterprise Management Associates sees the following as key deliverables:

- Centralized management over the whole of the grid means that even the largest and most complex installations can be managed with minimal staff. In addition to lowering the required number of full-time employees (many of whom will be freed to perform a higher value functions), the highly automated nature of the management software removes much of the human element from the daily process, thus lowering the opportunity for operator-induced errors.
- The ability to over-oversubscribe to physical assets through thin provisioning squeezes maximum value from hardware and software investments; optimal use of existing equipment makes it easy for customers to defer capital expenditures.

Obviously this has lots to offer for sites heavily invested in infrastructure but which also find much of that infrastructure going unused for long periods of time. Sharing assets across their user community not only will optimize those assets' values, but will also enable a pay-as-you-go "billing" approach that makes it much more likely that stakeholders will get fair value for their IT investment. While KiloClient is clearly a case of "the greater your problem, the greater the value to be derived," it is easy to see this sort of technology benefiting businesses of any size, specifically including large, medium and small enterprises and, via more efficient service provider operations, to SMBs as well.

In fact, what companies would not benefit from the ability to redistribute performance and capacity on an as-needed basis? Dynamically re-provisioned assets will appeal to any site struggling to accommodate either unpredictable hot spots or cyclical shifts in demand. This would enable a business to share resources between its accounting, sales and manufacturing groups for example, or to make resources available to web-facing applications during times of high demand, and then return those assets to a general pool once demand subsides. ASPs in particular will like this; their business models can only operate efficiently by sharing a finite amount of services with a large number of clients.

KiloClient is the first viable example of a large-scale grid that is available for commercial applications. Any data center staggering under the load of servicing multiple shifting workloads (SAP, Oracle, and web applications, for example) on multiple disconnected platforms will have to pay attention.

About Network Appliance

Network Appliance is a world leader in unified storage solutions for today's data-intensive enterprise. Since its inception in 1992, Network Appliance has delivered technology, product, and partner firsts that simplify data management. Information about Network Appliance™ solutions and services is available at www.netapp.com

About Enterprise Management Associates, Inc.

Enterprise Management Associates is an advisory and research firm providing market insight to solution providers and technology guidance to Fortune 1000 companies. The EMA team is composed of industry respected analysts who deliver strategic awareness about computing and communications infrastructure. Coupling this team of experts with an ever-expanding knowledge repository gives EMA clients an unparalleled advantage against their competition. The firm has published hundreds of articles and books on technology management topics and is frequently requested to share their observations at management forums worldwide.

Corporate Headquarters:
5777 Central Avenue, Suite 105
Boulder, CO 80301
Phone: +1 303.543.9500
Fax: +1 303.543.7687
www.enterprisemanagement.com

1340.061507