The Storage Grid: Based on a Series of Articles by Mike Karp, Enterprise Management Associates

Today, many businesses use network storage to meet their diverse and dynamic storage needs. SAN and NAS technologies provide scalability, efficient resource sharing, system availability, and more. These technologies and capabilities continue to be critical, but SAN and NAS are just steps in the evolution of storage.

In a world awash with data, we are challenged every day to turn data into information, information into knowledge, and knowledge into a competitive edge. Storage today is not about where information lives, it is about putting information to work.

That's why, going forward, storage will need to be simpler to acquire and manage, easier and faster to provision, and provide incredible scalability. HP believes that the storage grid is the next step in the evolution of storage.

The following was first published in *Network World Fusion* by Mike Karp, Senior Analyst at Enterprise Management Associates. In it, he analyzes both the promise and the potential pitfalls of a storage grid. We hope you find it useful.

Ta Da! The Enterprise IT System of the Future.

Two technologies are likely to exert a profound influence on storage and storage management during the next few years: information lifecycle management, and storage grids. They are two important and interrelated parts of what may well be the enterprise IT system of the future.

The Point is…….

In a very real sense, the planets (or stars, if you have a more galactic point of view) are now in alignment when it comes to supporting ILM. I mean by this, that the following indicators show that the time is ripe for the values that ILM delivers:

Point one: Pent-up demand. Many large IT shops have been deferring major storage investment for several quarters now. The result is that many, rightly or wrongly, already see themselves as being pushed to their limits when it comes to providing for their upcoming capacity requirements.

Point two: Budgets are getting bigger, but guarded spending will still be the order of the day. IT investment, while apparently a bit looser this year, is still going to be closely watched and stringently monitored. It is therefore reasonable to expect that storage investments will, whenever possible, try to extract maximum value, and that the concept of "good enough" will often replace "best in class" for hardware purchases.

Point three: "Good enough" hardware, when used appropriately, really is good enough for enterprise IT. SATA drives have already achieved widespread acceptance on the floor of the enterprise IT shop in virtual tape libraries and for storing non-mission critical data. (Some vendors are also trying to position a low-cost Fibre Channel drive - FATA - as a value-based alternative to the high priced Fibre Channel drives currently in use. FATA has yet to make an appearance on the marketplace, however.)

Point four: The opportunity for flexible storage is about to emerge. Serial-Attached SCSI (SAS) drives are interchangeable with SATA drives (in SAS cabinetry), and one can be swapped out for the other without restriction. Soon, a single array will be able to provide two or more tiers of storage.

Point five: Moving data is not quite as painful as it once was. Data migration tools, while still apparently not widely used, are available from several vendors and allow some level of policy-driven automation to be applied to the process. Automated data movement between the various tiers of storage is the most crucial element for ILM.

Point six: For larger enterprises, storage grid technology now seems to be a lot closer than we thought. These have the capability of providing real-time scalability in terms of both capacity and throughput. If this ability to scale – all the automated provisioning capability that necessarily must accompany it – lives up to its promise, storage grids are likely to play a key role in IT, and in the largest enterprise shops may well provide the most efficient implementation of a site's top-tier storage.

Optimized Capacity? What's That?

The technology behind storage grids is likely to play a significant role both in the enterprise and in enterprise ILM implementations. Here's why:

For companies that make widespread use of very large IT systems, the storage grid holds the promise of making real-time scalability and near-infinite capacity always available. Theoretically, at least, it can provide this capacity at any tier of storage a manager might care to define.

The job of grids – and this applies to both storage grids and compute grids – is to make capacity available when needed. Note, however, that capacity in this case needs to be redefined just a bit. With grids, capacity does not equal raw capacity, by which I mean the sheer amount of storage or computing power a company can have at its fingertips. Rather, it refers to what might be termed "optimized" capacity.

Optimized capacity? In terms of storage, this means two things:

First, it means that as additional capacity is brought online, taken off-line or accessed, nothing is added to the management load. Thus when provisioning or using storage devices, file systems will not have to be mounted by users, and admins will not have to "twiddle the dials" on their management consoles to bring additional systems online or to perform other management tasks. All decisions regarding rights of access and so forth have already been determined by existing policies.

The first goal of a storage grid: If a qualified user needs additional capacity, it is made available with no appreciable management overhead.

Optimized capacity also means that whenever storage capacity is added, throughput keeps pace. The importance of this point cannot be overemphasized; it does us little good to add new capacity if, by adding the new storage, we also slow down the system.

To place this second point in a context that will be familiar to many of you, consider what happens on a 16-bit wide SCSI bus. While you are certainly free to populate every address on the bus with a storage device, throughput tops off when the fifth device is added. Thus, even though you may add devices to your heart's content (up to 15), you can only access five of the devices at a time at peak efficiency.

With grids, such a situation will be unacceptable: capacity and throughput must scale proportionally so that every time additional storage is added, throughput also increases.

This objective of grids – matching easily managed and scalable capacity with equally scalable throughput – addresses a storage goal that we have been looking to achieve for over a decade. Beginning in the mid-1990s, system designers began taking compute cycles devoted to I/O processing away from the central CPU and deploying them out to intelligent controllers (if you have a sense of history, this is what the I2O initiative was all about, and to a large degree what TCP/IP offload engines are concerned with now).

Grids, of course, must do all this on a much larger scale, and in an easy-to-manage fashion. Also, their aim is to make all this pervasively available, with no single point of failure within the system.

So the major expectations of a storage grid are that storage scales in a near-limitless fashion, that throughput scales at the same pace as does capacity, that all this capacity and throughput must always be available, and that the whole system must be easily manageable. How is this accomplished?

How a Storage Grid Works – a Diagram in Words

A picture that shows the topography of a grid is often drawn as lattice, with each intersection representing a node, or cell. Within each of these cells is the storage, of course, but there is also a great deal more.

The node, or cell, is the basic unit of the grid. Because each node on a storage grid must support the added throughput that goes along with increased storage, each node also has a local processor devoted to I/O operations, cache (when necessary) to support the processor's operations, and a communications link to the other nodes in the system.

Each cell within the grid is connected to the other cells by multiple pathways (picture the lattice), and is part of a single virtual image that the user sees. The virtualization software that looks across the grid will be capable of balancing loads across cells in addition to other management functions.

The flexibility of such nodes will most likely vary substantially from one vendor to another. Some early implementations undoubtedly will be quite simple, while the more sophisticated ones will probably turn out to be extremely flexible in the kind of value-add they give to stored data.

For example, HP's implementation of the storage grid is based on a network of what it calls "smart cells," each of which can use either proprietary or third-party software to be tuned to perform specific tasks (hence, presumably, the "smart" in smart cell). Some specialized smart cells might, for example, be configured as NAS file servers, while others might be dedicated to providing block storage. Yet other cells might be content sensitive, providing WORM storage or another type of service that addresses issues of regulatory compliance. When the cells have to be repurposed, the grid management system will be flexible enough to change the "personalities" of these cells on the fly.

The management software within each smart cell can be dynamically loaded, and changed easily with a software download.

Repurposing cells is a way to maximize their value. For example, a high performance cell might contain management software with data copy, access path management and other management software appropriate to its need. As the cell's hardware ages, the cell may be repurposed to serve an archiving function. In such a case, for instance, some of the existing software might be exchanged for data migration software. As long as there is enough intelligence within the cell, cells can be allocated and reallocated to whatever tier of storage best suits the lifecycle management needs of the site.

How much of all this is likely to be proprietary and how much commodity product? Good question.

Grid Lock

If grid components are going to include storage, capability for both communications and processing, and enough memory to handle whatever demands are put on each particular cell within the grid, how many of these components will be generic? Or to get right to the point, how much of this will lock you in to dealing with a single vendor whose underlying technology is key to the environment, and perhaps allows it to function as a "grid provider"?

Keep in mind two things. First, what comes next is only likely to apply to storage grids, and may have no relationship at all to compute grids. And second, because we are talking about the future, all of this comes from my spending time gazing into the crystal fish bowl – far more accurate than a standard crystal ball at the county fair, it's true, but still noticeably short of being infallible.

A natural assumption is for us to expect that grid cells (and indeed, the concept of grid as well) will, in the course of time, evolve through several phases, but I think it is pretty clear that they will begin as a series of essentially identical nodes from a single vendor. This of course means the earliest appearances of storage grids will likely be completely proprietary systems. As we have discussed in an earlier column, vendor lock-in of this sort may not be such a bad thing if the proprietary environment provides a level of service (via a high performance file system or some other advantage, for example) that would be otherwise unavailable.

Eventually, grid cells are likely to become quite flexible, resembling an interconnected network of standardized, modular building blocks. The contents of these modules will be able to be swapped out as the cells are repurposed. In such cases, for example, mass amounts of memory might be added to support specific application requirements or a serial-attached SCSI (SAS) storage array might be demoted to serial ATA (SATA).

Ultimately, if the vendors do it properly, very little will be proprietary within each cell; storage will be able to come from any vendor, and IT managers will have the opportunity to continue playing their vendors against one another, just as many of the more clever ones do today. I fully expect, however, that the part of the puzzle that provides the fundamental grid management – mostly software, but perhaps including some proprietary hardware as well – will always come from a single vendor.

Grids then will represent a danger of sorts to some managers, and because of their proprietary aspects (or simply because you don't need the value-add they have to offer), storage grids may well prove to be inappropriate for many of you. For the rest, grids are likely to provide a significant part of the infrastructure of larger IT environments and will likely find good use both for large sites getting into information lifecycle management, and most particularly for those of you who will need to provide on-demand ("utility") storage services in a cost-efficient manner.

Living Up to the Hype

Storage grids potentially represent a massively available hardware infrastructure that can be placed beneath extremely large virtualized storage environments, but if they are to provide the greatest effectiveness, their assets must also be capable of being reassigned (provisioned) dynamically as needs change. But what will dynamic provisioning really involve?

Even the earliest grids, if they are to be effective, will require that provisioning be driven by automated policies rather than "hand tuning." Such early grids will consist of what are essentially identical (or at least very similar) cells, and although management systems will handle these individual cells within the grid as discrete entities, management also will mostly treat them monolithically, adding and subtracting cells as a complete unit rather than bit by bit. Additional storage will thus be aggregated to processes on a cell-by-cell basis in an action that is purely quantitative.

With more sophisticated grids, however, provisioning cells will go beyond just adding more devices and will begin to take on a qualitative dimension as well.

As storage grids grow in sophistication and as large enterprises come to rely on them to increasing degrees, many of the cells are likely to become optimized in favor of certain functions and will begin taking on "personalities" of their own. An implication of this is that as cells are provisioned to one specialized process over another, the personalities that define the cells may also come to require personality adjustments. In such cases for example, a cell that had previously been performing a file serving function for video-on-demand (VOD) might be reallocated to perform near line storage duties. In addition to swapping out high performance disks for cheaper ones (or perhaps just reusing the original – but now obsolescing – disks), the large read ahead buffers used for VOD would no longer be necessary and would either be physically removed or perhaps dynamically reassigned.

When this occurs, management functionality would be dynamically loaded (again, controlled by policies) and needed adjustments would be implemented in a theoretically straightforward manner through software downloads.

What is pretty clear is that storage grids, if they have any hope of living up to their potential, will have to support all the storage tiers needed to manage the information lifecycle, and will need the ability to change dynamically (but fortunately, not yet organically) to changing needs and changing hardware assets. Grids will evolve to become the infrastructure of the largest on-demand systems.

Today, HP StorageWorks is delivering elements of the Storage Grid and has articulated a roadmap that clearly and realistically defines a future that will enable you to put information to work.

For more information about the HP StorageWorks Grid, visit us on the web at: http://www.hp.com/go/storageworksgrid.

You can also e-mail your comments and questions to: storagegrid@hp.com.

To learn more about HP and the company's many other offerings, visit: www.hp.com.

ENTERPRISE MANAGEMENT ASSOCIATES™
Phone: 303.543.9500
www.emausa.com

867.11052004