

# Back to the future: A new conceptual model for grid storage white paper



Introduction: Is there a grid in your storage future? .....	2
Bewildering storage challenges drive the storage grid.....	2
Storage grid today—utility storage redux.....	3
Storage utility .....	3
High-performance computing grid.....	3
Distributed ad-hoc compute grids .....	4
Virtualized open storage.....	4
Open storage dead end .....	4
If storage had smarts—the key to grid storage.....	5
New conceptual model for grid storage .....	5
Back to the future: Vendor oversees the grid backbone.....	6
Conclusion: Grow your storage grid naturally .....	7

## Introduction: Is there a grid in your storage future?

Grid storage is really about the management of storage, but you would not know that from listening to what many vendors say about grid storage. To them a storage grid resembles an idealized virtual storage area network (SAN) or some utility storage service.

At HP, grid storage is seen differently. It is about simplifying the management of storage so organizations can easily grow and change their storage at will, depending on the needs of the applications and the business. This is done by defining an intelligent unit of storage that is called a smart cell and connecting smart cells as needed in a peer-to-peer fabric. The organization builds as much or as little of whatever kind of storage it needs by assembling various smart cells. The smart cells communicate by way of the peer-to-peer fabric while policy-based automated management of the storage is built right into the grid through each smart cell. HP embeds management in a multi-level virtualized approach that encompasses not only storage and the fabric used to orchestrate the cells but advanced storage services, such as automated provisioning, and even server resources.

A distinctive aspect of grid storage is the way it uses the fabric to create a totally virtualized managed storage environment. This fabric, which provides the infrastructure essential for gridness, weaves in distributed control, resource pooling, sharing, resilience, and even virtualized services like automated provisioning and security. These storage-resident, fabric-based virtualized services bring ongoing value to storage operations, such as data movement among heterogeneous arrays or across long distances, zoning, partitioning, capacity, availability, and performance.

Grid storage, when viewed this way, solves the problems that have long troubled organizations with their storage: problems around scalability, flexibility, availability, and management. The grid moves management into the storage itself and enables the organization to scale storage predictably and reconfigure storage of various types as necessary. This is not practical with today's SANs no matter how well architected or virtualized.

Several vendors are promoting their concepts of grid storage. Some rely heavily on virtualization; others look much like a conventional SAN. These implementations raise more questions than they answer. The truth is that the storage grid at this point remains a work-in-progress.

This paper looks at the challenges driving interest in the storage grid. It reviews and challenges popular notions about open storage and virtualization as a path to the storage grid. Finally, it introduces a new conceptual model for a storage grid based on research done at HP, and in the process shows why organizations need a storage grid.

## Bewildering storage challenges drive the storage grid

Every organization that has tried to implement a SAN and manage an enterprise storage infrastructure encounters the challenges of management, scalability, performance, and complexity/flexibility.

- Management poses a constant challenge. Despite efforts to overlay highly abstracted unified management layers and virtualization, administrators must still invoke specific element managers for individual devices to perform many tasks.
- Scalability continues to present challenges. Simply plugging in a new device raises concerns. The devices themselves become a barrier as organizations try to expand, configure, and allocate logical pools of storage across multiple boxes.

- Performance predictability has become a critical issue even as storage and network performance in general have increased. Now organizations find themselves chasing moving bottlenecks as the performance increases of some components overwhelm downstream components. In the end, organizations cannot predict performance as they add or upgrade components.
- Increased complexity and the need for greater flexibility result as enterprise storage is called on to handle a host of new chores, which requires more systems and diverse systems. It is no longer enough to store, protect, and make accessible mission-critical production/transaction data. Organizations must store compliance data, email, records, intellectual property assets, and more under the umbrella of centralized, enterprise storage and make it readily accessible.

And these challenges must be met with an eye toward storage ROI and TCO. Although storage hardware costs continue to decrease, the cost of sophisticated storage software and the people needed to administer storage systems continues to increase, offsetting by far any gains from falling hardware prices. Thus, organizations must meet the preceding challenges in a way that enables them to reduce TCO while increasing storage ROI.

## Storage grid today—utility storage redux

Grid storage promises to overcome the challenges facing enterprise storage. Today, however, storage managers often confuse grid storage with other technology concepts, such as the storage utility, high-performance computing grids, widely distributed ad-hoc compute grids, and virtualized open storage.

### Storage utility

Many storage managers equate the storage grid with utility storage and the related storage-on-demand services. To them a storage grid is a replay of utility storage. As a utility, storage is regarded as a service, much like electricity or water. Plug into the outlet or turn on the faucet and out flows however much storage the organization needs. The organization pays only for the amount of storage it uses and holds the storage service provider accountable for performance through a service level agreement (SLA).

The utility service model is attractive to many organizations, but the storage utility is not a storage grid, although a storage grid certainly can deliver storage services as a utility. To start, the conventional storage utility delivers only one kind of storage service, usually primary block storage. Any other storage would require an altogether different storage utility service and a different outlet, faucet, or connection on the application's end. Also, the organization's management options are limited to mainly modest tweaking of storage allocation among applications. Finally and most importantly, the conventional storage utility lacks parallel processing both within and among groups of resources.

Such parallelism, it turns out, is a key attribute of grid storage, another aspect of its unique gridness. Through parallelism the storage cells act independently but in unison. This allows, for example, a storage grid to achieve perfect load balancing compared to RAID controllers, which offer no inherent load balancing.

### High-performance computing grid

Some managers confuse storage grids with high-performance computing grids. Although both connect elements in a grid matrix and have conceptually similar attributes, they are not the same. High-performance computing grids are designed to tap the processing power of many loosely coupled CPUs operating in parallel to solve complex compute-intensive problems. The applications typically must be specially designed to take advantage of the computing grid.

By contrast, the storage grid as a collection of diverse intelligent storage cells appears to the applications as a single pool of conventional storage. The storage is accessed in a consistent manner without requiring any change to the application. Plugging an application into the storage grid is the same as plugging into a conventional disk array or tape backup system.

## Distributed ad-hoc compute grids

Others confuse storage grids with distributed ad-hoc compute grids like the SETI project that captures radio noises from space and sends the signal to a voluntary network of PCs. The PCs use spare CPU cycles to process the signal for signs of intelligent life. Similarly, some managers vainly hope they can take advantage of the large amount of excess disk capacity residing on any system attached to their network for storage SETI-fashion, thereby increasing storage utilization. The idea sounds good on paper but from an enterprise storage standpoint, the drawbacks of such a strategy are so great in terms of backup, management, and performance that few organizations would even dare. It is hard enough to boost storage utilization with a conventional SAN.

Again, this is not what a storage grid does. In a storage grid, data is only stored in smart cells. There it is monitored, managed, and protected automatically.

## Virtualized open storage

For many the storage grid amounts to virtualized open storage. In this view, diverse standards-based storage systems sit behind a virtualization layer that masks the location and specifics of the various storage elements and provides a single system image for management of and access to the entire storage environment. This is not a storage grid. The inherent management limitations necessitating fallback to element-specific tools and the continued challenge of compatibility and interoperability even among so-called open systems preclude this as a viable grid storage option (see Open storage dead end).

## Open storage dead end

The storage grid is intended to solve the problems of storage scalability, flexibility, availability, and management. At one time open storage was thought of as a panacea for these problems. Despite years of pushing standards and working on open systems interoperability and compatibility, the problems remain far from resolved with open systems or even virtualized open systems.

In short, open heterogeneous storage is not a storage grid. To the contrary, as far as solving the problems of storage scalability, flexibility, availability, and management, open storage is a dead end. It simply is too difficult and costly to rein in the complexity of mixed open storage. The continuing challenges both large and small organizations face with heterogeneous SANs speak for themselves.

Standards, which barely enable adequate basic managed storage today, certainly cannot provide the ubiquitous and easily managed storage organizations need right now. And looking ahead, it is apparent that standards will not and cannot go far enough to create the kind of transparent storage environment required for the highly scalable and efficiently managed storage organizations really need.

Rather, with open storage the organization continues to wrestle with compatibility and interoperability deep within the workings of the storage device and throughout the SAN as a whole. Each vendor's device has its own unique attributes and characteristics that fall beyond the reach of standards. Ironically, the more heterogeneous an organization's open storage is, the harder it becomes to pull it together as a coherently managed storage service. There always ends up being functions that have to be treated differently on one device or another.

The goal of open storage—transparent multi-vendor storage that interoperates seamlessly to deliver a scalable and flexible storage service—is indeed highly desirable. However, open standards-based storage is unlikely to get an organization there anytime soon, if ever, unless the organization opts for the most rudimentary commodity storage. Achieving the goal of open storage requires grid storage.

## If storage had smarts—the key to grid storage

The difficulty organizations have overcoming the challenges previously described results in large part from the lack of intelligence in the storage. Storage components typically are not only lacking in but they have only limited abilities to communicate.

Since storage components lack deep management intelligence, organizations compensate by running a variety of storage applications outside the storage itself. With storage intelligence embedded in the storage elements themselves, as is the case of grid storage based on smart storage cells, much more is possible. Now the storage itself sees the storage situation, recognizes what needs to be done, and provides appropriate storage services, such as placing data, moving data, retrieving data, and ensuring data integrity as needed by applications. Using policies, for instance, the storage would inherently recognize when data had not been accessed for a period of time and move the data to other, more appropriate, smart storage cells within the grid. And such data movements could be accomplished to deliver quality of service goals.

## New conceptual model for grid storage

A storage grid is a peer-to-peer network of intelligent storage elements (smart cells) that performs all the functions of today's networked storage while scaling effortlessly and providing fully integrated management of the entire storage grid as though it were a single system. Different cells within the grid provide different storage services. The grid scales out simply by adding disks or other storage devices in self-contained building blocks called smart cells. It scales up by aggregating the performance and capabilities of its component modules. Whichever way you scale it, the storage grid always presents itself as just one entity to manage, and storage capabilities are always delivered from the resulting single system image.

Storage virtualization—large-scale storage pooling—is inherent and fully realized and continuously leveraged in the storage grid. The storage grid encompasses all types of storage and manages each appropriately. Administrators deal with business application requirements while the grid's built-in intelligence attends to intrinsic storage management tasks.

The new storage grid model, as envisioned by HP, has several features that distinguish it from other approaches to storage often mistaken for a storage grid.

- Communication within the grid—Storage cells communicate among themselves
- Management and control functions are distributed throughout the grid
- Storage specialization—Each cell has a storage specialty, a personality, such as block storage or tape backup
- Intelligence—Each cell has built-in intelligence enabling it to manage its storage, coordinate its actions, and recognize the needs of the applications automatically
- Policy-based management—Cells control their actions and manage their storage on the basis of policies
- Single system image—The grid appears to applications and administrators as a single storage system regardless of how large and diverse it may become

The storage grid itself consists of the following key components:

- Smart cell—The core element, contains a processor and, optionally, cache memory, storage devices, network connectivity, a set of common software, and additional software that further defines the cell's personality, such as access path management, data resilience, index/search/retrieval, and copy/replication services
- Fabric—The grid uses standards-based network fabrics, including Fibre Channel and Ethernet, for connections to storage clients (applications or users)
- Operating environment—The components of the storage grid share a common operating system framework enabling distributed intelligence
- Distributed peer-to-peer architecture—The cells communicate and coordinate their activities among themselves, which allows for such capabilities as failover in the event a cell is out of action
- Plug-in services—Beyond the traditional storage services already built into the smart cells, these services manipulate data or provide advanced forms of protection or even process and present information to applications or transfer information across systems

Pieces of the HP conceptual storage grid model are working and available today, implemented as part of the HP StorageWorks Reference Information Storage System.

The following table illustrates the operation of a storage grid when provisioning storage for an application by describing (left column) and how it differs from the conventional process (right column).

Storage grid process	Storage grid difference
Administrator inputs capacity, SLA, and other relevant information to the SAN Master Management application	SAN automatically examines its inventory for available resources (the Storage Administrator is not involved)
Embedded Management Application checks grid asset inventory for available resources (advising if resources not available)	If new storage assets are required, the grid management application recommends what is needed and how it might be configured (with an AGREE click option if desired)
Automatically provisions resources to the application in accordance to policies and SLA	Grid correctly provisions resources to the application automatically, imposing any additional requirements, such as ILM as necessary
Automatically tunes itself to maintain SLA and other conditions	Storage administrator need not be involved beyond the initial inputting of pertinent information
To add storage to an existing application, the server simply asks the grid for more storage, which is automatically delivered as above	

## Back to the future: Vendor oversees the grid backbone

As far as storage goes, the experience with open storage has proven disappointing. Despite years of plug-fests and compatibility and interoperability testing, organizations still cannot freely mix and match storage components from different vendors and sometimes even from the same vendors at will. The latest revision, enhancement, upgrade, or change in any component anywhere in the storage infrastructure can break painstakingly crafted interoperability.

Storage in the mainframe, which is open in a highly controlled way, works much better. This is due to the control the mainframe vendor imposes over the storage environment. Organizations can still shop among competing products but the choices are, admittedly, narrower. The tradeoff is more reliable and more consistent storage performance.

Grid storage presents a similar tradeoff. A vendor will define the grid backbone and key components of the grid. It will then screen and certify third-party products, hardware and software, that work within its grid framework. In this sense, grid storage represents a step back to the future.

Organizations have found this approach very effective with mainframe storage. Blade servers take a similar approach, as do the large disk array vendors. Organizations have no interest in digging into the internals of their enterprise disk arrays and mixing and matching various components. It is a high-risk exercise for which most lack the expertise. Organizations will feel the same about grid storage.

The HP concept of the storage grid, however, provides for openness based on APIs for third-party software. HP also understands the need to incorporate heterogeneous storage that conforms to key grid specifications. In effect, the storage grid, as HP envisions it, becomes a vendor-integrated storage environment.

Adoption of grid storage elevates the organization's approach to storage. Rather than wrestling with individual boxes, which has proven to be costly, organizations can focus on the long-term benefits of storage integration, availability, scalability, flexibility, reliability, and performance while the grid simplifies the management, acquisition, deployment, and maintenance of the storage environment. This ultimately increases storage ROI and reduces TCO.

## Conclusion: Grow your storage grid naturally

A key advantage of this model is how organizations can implement it gradually. They can start with one or a few smart cells to meet a specific set of storage needs and naturally grow their grid in capacity and capabilities by plugging in more or different smart cells as needed. They do not need to rip and replace their existing storage systems. Rather they can gradually replace them as they come off lease or otherwise prove inadequate simply by adding more smart cells to the grid.

Over time, organizations will be able to scale out and scale up their storage grid until it handles all their storage needs. They will be able to manage the entire storage infrastructure and service through a single management image. Administrators will be able to focus on business-level issues by defining policies. The storage grid itself will handle the operational work automatically.

© 2005 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

4AA0-1792ENW, August 2005

