Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

Virtualized Grid-Based Storage Architecture for Evolving User Needs

Sponsored by: NEC

Robert C. Gray Laura DuBois

March 2007

EXECUTIVE SUMMARY

IDC finds that virtualized grid-based storage will be the underlying framework enabling future reduction of infrastructure complexity, administration cost, and management overhead. Through the adoption of these new storage approaches, firms will satisfy expectations for managing their information and content to satisfy regulatory compliance, corporate governance, and stakeholder requirements. This white paper identifies the practical challenges facing today's storage architecture in meeting current and future business and technical requirements. Grid-based storage architectures offer a solution to these challenges. This white paper distinguishes attributes required for best-of-breed, virtualized grid-based products. The overarching benefits to firms moving through the stages of storage infrastructure development can be substantial.

SITUATION OVERVIEW

Convergence of Infrastructure and Information

IT organizations find themselves at a crossroads. Complexities in today's datacenter storage environments have driven datacenter managers to operate at the level of managing infrastructure. Business units and functional stakeholders are now driving IT executives to manage content and information in addition to the infrastructure on which it resides.

The complexity facing IT today in the datacenter is the result of storage heterogeneity with disparate tools to manage each environment as well as a high rate of change of IT infrastructure. Datacenters are increasingly dynamic where servers and storage must be configured, provisioned, or upgraded. Technology refreshes every three to five years require time-intensive and disruptive migrations, which businesses can no longer tolerate. IT professionals must manage exponential storage growth as well as new SAN and virtualization technology while dealing with interoperability challenges.

Compounding the infrastructure management challenges are complex business requirements for tighter information controls, access, and integrity. Simultaneous management of infrastructure and information is difficult for even the most sophisticated and advanced firms. Emerging approaches such as virtualized grid-based storage strategies frame a solution that enables a reduction of infrastructure complexity, administration cost, and management overhead. The IT focus, then, moves to meeting increasingly stringent SLAs and business policies on information control, protection, storage, and retrieval.

IT Organizations Evolve Through Stages of Storage Infrastructure Development

Based on extensive primary research, IDC finds that firms move through stages of storage infrastructure development (see Table 1). Firms move from lower stages to higher stages of storage infrastructure development one stage at a time, much like they do with stages of organizational development. IDC finds that these stages of storage infrastructure development are an effective means to segmenting users based on the current and future needs that result from their current and future levels of infrastructure. IDC's segmentation distinguishes five stages of storage infrastructure development. Although there is a correlation between the scale of storage infrastructure and the stage of development, there are multiple dimensions that tip an organization from one stage to another. Size is only one aspect.

TABLE 1

Stages of Storage Infrastructure Development

Stage	Focus	Characteristics
1	Internal or DAS storage	• Applications tend not to be 24 x 7. • Organizations have limited staff and financial resources and limited dependency on technology. • Storage is not a focus. • Internal storage or DAS is good enough. • Firms expect next stage to be a SAN. • Capacity tipping point is around 2TB. • Backup and restoration is mostly an individual server operation, but LAN-based backup to tape may be in operation.
2	Build SANs	• Applications are complex; at least one application requires high availability and at least one mission-critical application has redundancy and requires 24 x 7 operation. • Organizations have more than five IT employees, are committed to networked storage, and have an approved budget for a SAN or at least one SAN installed. • Internal storage and/or DAS is used for older applications but is not good enough for future needs. • Firms are gradually pulling the storage out of servers during upgrades, using snapshots for backup, and can make basic configuration changes using the SAN with minimal stress on IT staff resources.
3	Reduce TCO of SANs	• Organizations have installed multiple SANs and have been through multiple generations of the products. • Applications are critical business processes. • Organizations typically have tens to hundreds of terabytes and concerns about administrative inefficiencies in the provisioning and management of SAN-based storage capacity. • Firms are looking for help in local and remote SAN interconnect and are showing heavy and growing use of disk-based data protection and recovery. • Organizations recognize the need for disaster recovery sites, have made recent investments in server solutions such as VMware, and are looking to see how the solutions will affect their storage environments.
4	Manage data	• Companies are leading-edge buyers that gain advantage from technology. • Organizations have test/development labs in place for final qualifications. • Pain points extend beyond storage mechanics and infrastructure to federating data, organizational control of data, and matching data value to access/storage tier. • Information has senior value with a focus on management and access to management for new applications.

TABLE 1

Stages of Storage Infrastructure Development

Stage	Focus	Characteristics
5	Create data and information services	• Bleeding-edge firms partner with suppliers, may originate specs, and do some of their own integration and design. • Prime focus and value is centered on the data/information service. • Applications are a component of services. • Users want to pay only for services used. • Users are willing to pay for instant gratification in regard to new applications, etc.

Source: IDC, 2007

IDC estimates that a larger percentage of firms are in the earlier stages of storage infrastructure development, while fewer firms have reached the more advanced stages of 4 or 5. However, today's solution, with compounded growth, becomes tomorrow's problem requiring a new approach to simplification. In mature stages 4 and 5, organizations are able to manage information better than infrastructure; however, IDC finds that more firms are being asked to manage both information and infrastructure. Today's storage infrastructure inhibits this.

Networked Storage Technologies SAN and NAS Present Practical Challenges

The adoption of SAN and NAS has been significant. SANs and NAS have provided the ability to manage larger capacities with few to no staff increases. The terabytes managed per storage administrator have increased dramatically. However, as SAN and NAS networked storage have sought to address shared access, reduced administration, improved backup, and other infrastructure-related issues, the architectures have themselves introduced a set of new challenges.

SANs Are Fixed and Static

SANs hardwire servers and storage into a storage network. Static bindings remain between a host and the LUNs to which that host has access. SANs as a storage architecture have no concept of dynamic adjustments to account for performance, reliability, availability, and resource utilization, and they make dynamic, automated provisioning a challenge at best.

SANs Have Practical Scale Limits and Required Planned Downtime

Today's SANs expose the RAID arrays behind the network and their differences. They tend not to be designed for continuous operation and sometimes require outages for preventive maintenance and/or data migration to accommodate hardware rollover to new generations. These limits and others have led to a proliferation of SANs in most enterprises and a failure to obtain the promised benefits of a single storage pool.

SAN Evolution Increases Interoperability Challenges

Each SAN evolution is based on new inventions and technology. It is, therefore, difficult for older technology to interoperate with new architectures. First-generation SANs were basic point-to-point connectivity networks using early SCSI protocols over Fibre Channel. The growth in SAN adoption and advancements in technology resulted in disparate SAN islands based on heterogeneous and sometimes noninteroperable devices. Today these disparate SAN islands must be consolidated into single, homogeneous physical SANs while supporting the logical segregation needed for different business functions and units.

SAN Standards Are Evolving with Invention

Evolving requirements have led to an ever-changing set of Fibre Channel protocols and standards for functionality and interoperability. As a result, the users of each generation of SANs have experienced the pain and instability associated with any newly invented and complex product.

Complexity, Cost, and Management Overhead Remain

SANs only minimally address the need to reduce complexity, cost, and management overhead. SANs have been continuously in catch-up mode to the needs of the most demanding storage users. Their proliferation is increasingly the source of a new complexity and increased management cost. Many of today's storage systems installations have time-consuming manual operations. They also require independent management of the storage arrays. Despite the level of storage consolidation provided by SAN and NAS, users report siloed operational barriers, especially when it comes to replication, backup, and archiving.

IT Infrastructure Paradigm Shift Is in Progress

An overall IT infrastructure paradigm shift is under way. The fixed and static dedicated IT infrastructure that originally offered clarity and simplicity now involves unit proliferation and layers of complexity. The new emerging virtualized grid-based storage architecture offers a renewed simplicity of continuous operation, shared use, and automated management through virtualization and a grid-based core, especially for both application servers and storage solutions. This shift is the result of the following practical competitive business and technology realities.

Information Outlives the Media or Technology

Today firms must meet duties to retain and preserve information based on regulatory and legal requirements. Regulatory compliance requirements in certain industries can require firms to retain information for 20 years, which is far longer than the three- to five-year technology refresh cycles that are common today. Band-Aids and add-on incumbent data migration technologies, which require application outages and potential data integrity concerns, are a pain.

Driven by Business Rules, Firms Have No Tolerance for Downtime

In an increasingly 24 x 7 environment, and in a society where information is always available and systems are always on, there is little to no room for application and user downtime. IT maintenance tasks, such as storage hardware or software upgrades or data migrations, can no longer tolerate downtime for these normal data management processes.

Application and Geography Scale Needs Are Underserved by Single Servers and Disk Drives

Current and emerging applications routinely have hundreds of thousands of users. Some have millions. The Internet and VPNs have permitted these users to be distributed around the globe. These applications are typically designed around a tiered and scale-out architecture that provides a high degree of availability and, more important, provides the capability of near-uniform application response under very dynamic load conditions. Uniform data access and storage for these mega-applications is required on a scale that traditional disk drives and storage architectures are strained to provide cost-effectively.

Storage Is Changing from Scale-Up to Scale-Out Infrastructure

Much like the path to cost-effective server infrastructure, the path to cost-effective storage infrastructure is through the exploitation of quantities of commodity disk drives and commodity standard servers that are stitched together in a scale-out storage architecture. All indications are that content with its monetizing and data mining will only increase as IT focuses on personal data and entertainment.

Computing Power Is Changing from Single-Server Focus to Community of Intelligent Nodes

In the future, computing power will come from communities of intelligent nodes. Some communities will serve applications in a flexible virtualized manner. Some will provide connectivity, as the Internet and intranet do. Many will be linked to provide a seamless and infinitely scalable storage capability.

These grid-based infrastructures have an analogy in regional healthcare networks. Service is provided seamlessly at each satellite center. If a particular healthcare center is inaccessible, full service is still available from the other centers. Thus, an integrated, scale-out healthcare system provides the best of both worlds — local services (regardless of location), the ability to serve many patients simultaneously, and very high availability of services.

IT's Focus on Infrastructure and Costs Is Shifting Toward Providing a Service

Today's applications require always-on performance. It is an age of convenience in which consumers, customers, and even employees are not conditioned to wait. Policies and SLAs on performance, protection, availability, retrieval, and migration must be at the highest levels.

The Two Models of Distributed Computing

Distributed computing architectures focus primarily on clusters and grids with variations in the way that each is deployed. Each model consists of multiple intelligent nodes functioning in combination. In both concepts, applications continue to function as nodes are dropped or added to the architecture. However, scale and degree of functional integration as well as expectations about "when" the application will be processed distinguish the two architectures and the benefits possible with the architecture.

Cluster Model

The distributed computing architecture known as clusters dates back to the 1980s, when minicomputers began to be deployed as clusters, sharing a common storage source and interconnects organized through the use of failover availability software. A "heartbeat" between the server nodes allows the linked servers to know when any one of the server nodes goes offline — triggering a failover to another server in the cluster and accessing shared storage to restart the application and retrieve the shared data.

This style of clustering, known as failover availability clustering, is widely used. However, it often requires customers to integrate the solution themselves or to write "scripts" that make custom applications cluster-aware. Alternatively, customers look to vendors to "cluster-enable" packaged applications through the use of "agents," which are provided with the clustering software product, or to ship preconfigured clusters, which contain server, storage, interconnect, and software elements, directly from the factory.

Grid Model

The grid architecture is based on a multinode distributed operating system, often referred to as grid software, which distributes fragments of the application or workloads among the nodes with the aim of directing workloads to the nodes where they can be processed best.

The grid-enabling software must be aware of the compute or data resources that are available on the grid. A task scheduler is typically provided that assigns specific workloads to specific "target" nodes. The assignment of workloads happens across a network, much as a server operating system would direct specific compute resources within a PC or server to process specific tasks.

It is important to note that one or more of these target resources may, in fact, be a cluster. Clusters and grids can coexist in this way, with clusters included as resources within a grid. Setup and operation of a grid is typically more automatic and open-ended than it would be for a failover-style cluster. In that regard, the grid is somewhat more self-organizing than the failover cluster.

The target design for a grid typically scales beyond the size of clusters and may reach across countries or continents to form the grid topology. (Clusters with WAN support between server nodes also exist and are called geoclusters.) The grid typically may lose multiple nodes with minimal impact on the application, which is another indication of its open-ended design.

Workload balancing clusters and scalability clusters can also redirect workload to available nodes, even in the event that a single node goes offline. The processing is moved to alternate server nodes and is continued. The results are gathered (and returned to the end users) by the software that is controlling the cluster. Thus, a grid-based architecture is one that reaches beyond the scope of most clusters to include all of the computing resources on a single, networked entity.

Considerations in a Grid-Based Storage Architecture

As firms shift away from traditional storage solutions to virtualized grid-based storage architectures, users should use specific metrics in judging grid-based storage suitability. IDC finds that a grid-based storage architecture should include the attributes noted in the following sections.

Protection: Data Management Services Integrated Within Architecture

Incumbent SAN technology typically segregates data management services as expensive, separate add-ons. Copy and data mobility functions, snapshots, CDP, replication, migration, and single instancing are operation-intensive, discrete, and separately bought and managed functions.

Data security is currently an afterthought. As an add-on, encryption has a separate management interface. Activation is specific and subject to high operational costs and to errors. In some cases, the function requires separate capacities and storage areas or tiers. Functions such as secure-erase risk data assets by failure to execute policy consistently.

We believe grid-based solutions will supply many of these functions as an integral part of the core architecture with substantially lower acquisition and operation costs.

Low TCO and Simplified Operation: Executed Through Automation

Integrated functions in future grid-based solutions will have a high degree of automation, providing new expanded PB (petabyte) levels of administrator capability. Organizations should expect automated, transparent provisioning and capacity expansion. Further, they should expect automated, nondisruptive self-tuning and load balancing to optimize performance and utilization as the system expands or contracts. Additional enhancements will include nondisruptive, online data migration to generations of new technology, self-detection of errors, and automated healing-correction handling.

In the future, differentiated SLAs will be possible on an application level. Automation and embedded data management functions will be utilized to optimize the cost and performance on an application basis.

Tiers of Resiliency: Extending into Higher Levels

Unlike SANs, grid-based storage must have no planned downtime. Preventive maintenance, firmware upgrades, and even next-generation migration should be expected to occur live without having to bring down the application.

Grid-based solutions should provide for an extended range of data resiliency, including true continuous operation, even under disaster conditions. The grid-based infrastructure not only should be able to dial in the number of tolerated element failures but also should provide for embedded and default site-level failure and recovery. Today's focus on disaster recovery will shift to disaster resiliency or disaster prevention. Data will be able to survive site-level failure, hence eliminating today's complex and costly failover/failback process.

Massive Scaling: Without Limits

A well-architected grid-based storage solution should provide balanced scaling of both performance and capacity to keep pace with data growth and business requirements. Capacity should be shared without silos.

Efficient Use of Raw Capacity

Next-generation, grid-based storage architectures should provide for efficient use of raw capacity, utilizing overserving compute cycles to eliminate duplicate data instances. Allocation of capacity should be virtualized from physical and actual use.

Architecture with Foundation Capability to Evolve

Leading grid-based storage architectures will be identified partly by the provision for growth of a rich ecosystem of advanced functions. The design will enable, not restrict, extended possibilities, partnerships, and collaborations.

Technologies Required to Obtain Attributes from Grid-Based Storage

IDC outlines a set of stringent requirements for true next-generation, grid-based storage architectures and for the ability to meet standards in the largest and most mature enterprise environments. This description gives a clear demarcation between last-generation SANs and next-generation, grid-based technology. The requirements imply specific technology mechanisms to provide the environment to supply needed metrics. The following sections describe technologies that distinguish grid-based storage that can deliver on the needed attributes.

Automatic Data Mobility for Load Balancing and Data Migration

Organizations should look for separate and distinct data movement engines and policy engines. Control should be separate, allowing independent adjustments as the system expands or contracts or as different resources are consumed and needed.

The technology must also allow nondisruptive expansion and contraction of the system. New nodes in the grid must be automatically discovered and added into the community without manual intervention.

These grid node additions or deletions need to be handled in the background and must not impact the performance or availability of the system from the applications perspective.

Rules-Based Control

The control engine should offer intelligence and a flexibility of rules optimization. The policy should support auto-discovery, auto-acquisition, and auto-termination. The rules-based policy should cover resiliency as well as the integrated data management services offered by the system. Policy levers should provide for differentiated SLAs. Finally, the control should perform smoothly as if there were a degree of self-awareness.

Data Distribution over Nodes

For performance and resiliency scaling, there should be a sophisticated and distributed index that distributes data over varying numbers of nodes for ranges of resiliency. The distribution needs to be self-adjusting as the system expands or contracts so that performance and resiliency are always maximized.

Distributed Architecture and Computing

A fully distributed architecture can contain no centralized resources, which would constitute single points of failure and hence be a fatal flaw in attempts to provide continuous availability and resilience. A highly scalable distributed architecture cannot have fixed numbers of structures or resources or fixed table sizes that would limit scale. Finally, an ideal grid-based architecture would be capable of sustaining various element failures and capable of growing to a massive scale, both occurring while general performance levels are maintained.

No Reliance upon RAID

The grid-based architecture will not be limited by the past century's now restrictive RAID-based data protection conventions. While innovative by the standards of a quarter century ago in the environment of simple ranks of disk drives, RAID now finds itself at odds with scaling objectives in an expansive grid-based storage solution. RAID was designed simply to protect against an HDD failure (or two) within one node, not across multiple nodes. Computing parities across multiple failure domains (nodes and even sites) would quickly become the storage grid's choke point well before desired and distinguishing scale and distance breakthroughs are met.

The breakthrough for grid-based storage solutions can be had from enhanced data resiliency technology. Such enhanced technology would create sufficient parity based on user-defined resiliency levels and intelligently distribute data and parity across disks and nodes to maximize resiliency at any given time based on the grid's configuration. This enhanced resiliency needs to support levels beyond two failures (RAID 6) and should be capable of supporting entire site failures in a geodistributed environment. Beyond specifying the resiliency level, all of this provisioning methodology should be automatic and nearly hidden.

Deduplication

A virtualized grid-based storage architecture providing high resiliency and continuous availability would naturally become a repository for snapshots and replicated data retained for application recovery and other uses. Because backups, archives, and inactive files tend to have replicated data, the embedded capability of the grid-based architecture to include deduping would be highly valued.

A Platform for Evolution and Growth

With a mandate of continuous operation, the architecture must provide for rolling node upgrades. The grid system needs to support open and heterogeneous nodes utilizing different hardware and different disk sizes. Nodes of unlike revision levels must still associate and be proper grid members. This allows customers to easily take advantage of new technology within the grid. Buyers can ride a technology curve without having a forklift upgrade, even as evolution and growth change the topology, resilience, and capacity.

Industry-Standard Interfaces

The grid-based solution is nothing unless it seamlessly interfaces and supports the expected storage communications protocols and industry standards. Without common language, critical application support would not exist. Organizations should look for data protocols such as NFS, CIFS, WebDAV, and HTTP and for management control protocols such as XAM and SMI-S.

OPPORTUNITIES

Robust, virtualized grid-based storage solutions will provide a set of benefits well matched to the pains and needs of storage users moving into more advanced stages of infrastructure development. For the most demanding storage requirements, users can expect unprecedented reduced cost of operations, expanded availability and resiliency levels, and unlimited scalability.

Users and datacenters that are making complementary advances in server infrastructure to blade servers and hypervisors (VMware, XenSource, Microsoft) will find grid-based storage solutions to be natural complements.

With secondary storage for protection, retention, and archiving and for content data that is increasing as a proportion of total storage capacity, the ability to integrate multiple service levels into a uniform architecture will gain increasing value. A virtualized grid-based foundation architecture is capable of providing the unique requirements of extended availability and capacity over performance. Users should first consider virtualized grid-based storage for expanding secondary storage requirements.

CHALLENGES

Companies will find solutions for their emerging enterprise storage needs in the form of virtualized grid-based architectures, but they will have to overcome the following implementation challenges:

- **Uncoupling from the image of compute grids and grid computing.** Distributed computing carries a legacy of complexity and a position outside the mainstream. While grid computing may be moving toward the mainstream, it is not typically held to be ready for production environments and is seen as the darling of the research high-performance computing set.

- **Magnitude of value proposition versus risk for deployment.** Early adopters will be those with a premium need for the benefits promised by grid-based technology with either a high-risk tolerance or a payoff that merits moving forward. Over time we expect risks to be reduced and more users to find resonance with the grid-based storage value propositions.

- **Integration into rest of infrastructure.** Storage is the data in the datacenter and hence is of core importance. Supporting existing protocols and management frameworks and being compatible with existing staff skills are important aspects of a nondisruptive, seamless integration. The ability to hide the grid complexity with operations simplicity will be crucial.

- **Meeting performance needs.** Performance has been one of the challenges not only to grid-based architectures but to all distributed architectures. Traditional approaches have frequently bottlenecked with additive locking times as the grid expands. Successful virtualized grid-based storage architectures will have to demonstrate solutions and also offer data access and transfer times that are comparable to those of traditional storage architectures.

- **Standards and open grid-based storage.** The "one storage pool for all" concept of a storage grid eventually requires open standards. Buyers will demand multiple sources for grid nodes and storage arrays. The successful technology and architecture owner will need to build an ecosystem that includes second-source suppliers that permit the eventual formation of multibrand, heterogeneous storage grids.

CONCLUSION

IDC believes virtualized grid-based storage is the next big thing in storage architectures. Leading storage users need a strategy to be early users of the technology, and they need to understand the technology so that they are positioned to exploit it for the substantial benefits. 2007 will be a good year to make initial investments.