




IT Knowledge • Business Results



# Clustered Storage: Valuable Today - Requisite Tomorrow

## An Essential Step to the Virtualized Data Center

By Tony Asaro  
Senior Analyst

March 2007

## Table of Contents

---

Table of Contents .....	1
Introduction.....	2
Users and Managers .....	3
The Root Cause, the Fluid Data Center and the Great Leap Forward .....	4
True Clustered Storage .....	5
True Clusters versus Management Clusters .....	6
True Clusters versus Global Namespace .....	6
Clustered Storage System Challenges.....	7
Valuable Today - Requisite Tomorrow.....	8
ESG's View .....	10

## Introduction

---

The biggest problem with managing storage environments is the complexity of the storage systems themselves. Most leading storage systems are complicated machines that require deep expertise and ongoing manual administration to maintain operational excellence. This complexity exists because these storage systems are not highly virtualized.

The more virtualized the storage system, the easier it is to manage and maintain operational excellence. Whenever we require a human, regardless of expertise and experience, to perform manual functions, the process decreases in efficiency and the chance of error significantly increases.

Virtualization provides both a logical view and control of physical assets. RAID is one of the most common forms of storage level virtualization. The storage system uses RAID to store data across a group of drives and then keeps copies of that data on one or more drives in the event of a drive failure. In some of its forms, RAID provides two main functions; increasing performance and providing data protection. However, while RAID is extremely valuable and in most cases a requisite storage system technology, it provides a somewhat low level of virtualization. RAID is just chapter one in the book of storage level virtualization.

What does virtualization have to do with clustering? Everything. Clustering is one of the most powerful and useful forms of virtualization. A cluster enables multiple physical systems to be seen as a single logical system (a lot like virtualization). As a result, clustered storage removes the limitations of individual physical devices, effectively removing the boundaries of the "box." Once these limitations are removed and the storage system becomes boundless, IT can transcend the daily grind of managing physical assets and become truly strategic in the business.

In most cases, IT is seen as overhead; a cost to the business. Unfortunately IT is often not considered strategic to the business. In some cases, IT can be its own greatest enemy in this regard. Because we have smart people pushing buttons, tuning devices and speaking "geek" - the business perceives that IT's value is in performing these activities. And this perception is justified, since in most cases, the IT people believe this as well. As an IT expert, it is comforting to know that you possess a wealth of great knowledge that no one else has and that the work you do provides tangible value to your company. This gives you job security and respect, people leave you alone and you have a clear focus on what to do today, tomorrow and forever. It doesn't matter that at times it's really hard and complex. So much the better. More proof of your value as an IT guru. But these are actually the wrong qualities to be valued for. Instead, IT should be a strategic part of the business that uses technology as a competitive advantage and a means to build new and better products and/or services.

The companies that understand that IT is strategic and that IT, in many cases, *is* the business are the ones that ultimately win. This is not achieved by making the same choices or implementing the status quo. Benjamin Franklin said the definition of insanity is doing the same thing over and over and expecting different results. I wouldn't go so far as saying that we are making insane choices in the data center; however, if we keep implementing large, complex, semi-virtualized storage systems that are expensive to acquire and maintain, we shouldn't expect to solve our problems any time soon.

## Users and Managers

---

There are two constituencies that are impacted by the limitations of our physical infrastructure: the users and managers of these bounded semi-virtualized systems.

Users just want things to work. They want to use their applications and access their data, without being encumbered by technology. Therein lays the paradox of IT. While technology is meant to improve our experience and efficiency, our reliance on it can paralyze us if that technology doesn't work as we would like. At this point, IT becomes the problem instead of the solution. The goal of IT is to be invisible to the end-user of its services and technologies, but that is typically not the case. In fact, the users are often all too aware of IT and their limitations.

IT managers are beset and besieged with a multitude of technologies, devices, tools, management applications and systems to contend with. The more complexity added into the environment, the more chance there is that IT will fail to meet the needs of the business and the users. As events change the limitation and bounded nature of physical storage, assets make themselves visible and apparent. That is when the trouble begins.

One of the most popular approaches to change management is to not change anything. That is why we often buy bigger systems than we need. We create large volumes that consume lots of capacity that is allocated, but frequently never used. We over-engineer our systems in order to minimize any change. It is often preferable to buy too much of something rather than deal with the consequences of not having enough. This leads to waste and ultimately creates more complication, not less. As new applications come online, we add more capacity to our systems, eventually buying new storage systems and increasing our physical infrastructure. The problem grows to the point where we need to consolidate these systems, which requires purchasing new, bigger systems. We then have to plan and implement a cut over from the old to the new. This takes time, money and resources - all things that are precious and never in great enough supply.

## The Root Cause, the Fluid Data Center and the Great Leap Forward

Ultimately, the problem should not and cannot be solved by storage management software. Masking the complexity of a storage system by developing GUIs and software is not the answer. That is like building a house on a faulty foundation. Eventually, a disaster will occur. That isn't to say that better GUIs and storage management software don't provide value. Clearly, they do. However, the place to solve storage system complexity is within the underlying architecture of the storage system itself. The complexity of individual storage systems is the root cause of the problem.

Using virtualization as the "glue" to aggregate physical components has historically proven to provide leaps in efficiency. I used RAID and clustering as examples of this earlier. In the case of both RAID and clustering, multiple discrete physical components are treated as a single, logical entity. By doing this, you increase performance and reliability. Another example of virtualization, and perhaps the most popular today, is the use of virtual servers, which essentially implements the inverse of clustering. End-users create multiple virtual servers on a single physical device. Even though clustering and virtual servers appear to have opposite value propositions, they fundamentally achieve the same goal: the ability to better utilize your physical IT assets in order to create a leap in efficiency.

For years, we have discussed the notion of a fluid data center. Clustering fully embodies the principle of the fluid data center. A true clustered storage system is a network unto itself, sharing communication and metadata between all of its physical components, thus creating a highly virtualized system. This is an essential point. Today, network storage is really an oxymoron. Storage systems are only partially networked. We have done a great job creating a partial network in which host systems share storage system resources. However, the storage systems themselves share no communication with one another. They are discrete and isolated devices that are networked to host systems but not to each other. In reality, network storage systems are not fully networked.

So, why do you need storage systems to be fully networked? The value of networking all IT assets should be abundantly clear. In the case of storage, it enables you to use a scalable system onto which you can keep adding resources as needed. You can add performance elements, capacity, and connectivity. You can move data between different tiers to maximize your investment. As you grow your storage environment, you can maintain ease of management. You can scale and load balance for performance. You can grow performance as needed. You can increase reliability to higher levels. A system with 1 TB of capacity is essentially as easy to manage as one with 1 PB of capacity. Even though the physical infrastructure grows and there are many hardware devices, it is one massively parallel, logical system that is easy to manage.

It is relatively easy to manage one system. It is hard to manage lots of systems. People don't scale. Therefore, we must scale our systems. There is little value in utilizing more people to support your burgeoning IT environment. The ability to manage multiple systems as a single logical element is how IT can scale. Clustering enables you to effectively manage multiple physical systems as a single logical system.

A true clustered storage system is logically one system. Once, it took one person to manage terabytes of data. That one person now can manage petabytes. This is an exponential leap in efficiency.

## True Clustered Storage

---

N-way clustered architectures support multiple storage controllers in a single cluster. Though these clusters may contain many storage controllers, they still appear to the applications and users as a single logical system. Clustering is a powerful form of virtualization that creates a logical system consisting of multiple physical components. Clustered storage systems allow users to add CPU, memory, and bandwidth transparently. It enables users to scale based on the needs of the business without having to buy a whole new storage system. In traditional dual node storage systems with fixed architectures, when a user's environment outgrows their storage system, they may be forced to buy another system that provides greater performance and scalability. With clustered architectures that allow for the aggregation of all hardware resources, performance can be scaled in a linear fashion, if and as needed.

Clustered storage systems can start out modestly and then grow to massively parallel systems. This is a cost-effective approach that doesn't require users to buy ahead of their needs, which is often the case with traditional storage systems that have fixed architectures. Additionally, by adding more storage controllers into the cluster, users can continually raise the performance ceiling. N-way clustered storage systems provide just-in-time scalability. But scalability should not create more complexity. Regardless of how big the cluster gets, it should still remain a single logical system to manage. Essentially, it should be just as easy to manage a clustered storage system with 100 nodes as it is to manage one with two nodes. The system contains more physical storage controllers but remains a single logical system.

Not all clustered storage systems are created equal. Some are what ESG refers to as "true clusters." True clustered storage provides a single logical storage system that allows any application to access any piece of data stored on it through any storage controller in the cluster. A true cluster provides for the aggregation of all its hardware resources into a single system.

A true cluster is a highly virtualized system. Each of the storage controller nodes communicate with one another, sharing a single "brain." The result of this is that host ports, CPU, memory, internal bandwidth and capacity are all logical pools of physical resources. Consider the power of multiple CPUs working concurrently to run various system and I/O operations: a pool of cache memory that can scale to hundreds of gigabytes, host ports that create a virtual pipe that reaches 10s of gigabits of bandwidth and a pool of disks with hundreds and even thousands of actuators all working in concert reading and writing data.

Operations such as handling read/write requests, performing local and remote data replication, and executing RAID rebuilds are serviced by a pool of resources that can be added to, if required. A true clustered storage system ultimately provides parallelism that can scale from small to massive. This is an important point. Enterprise-class storage systems also offer massive parallelism, but you have to buy the whole enchilada. You have to make a heavy investment at the onset, projecting what you will need for the next three to five years. With true clustered storage systems, you can add resources as needed.

It is important that the process of adding resources as needed is easy. The promise of a highly virtualized system that supports true clustered architectures is that the initial system is easy, adding resources is simple and the logical storage system remains easy to manage over its entire lifecycle, regardless of how massive the cluster gets.

Scale is all relative. Clustered storage is not just for large environments. When one IT person handles PCs, servers, the network, Internet connectivity, security, application support, servers and storage, you have a scalability problem. On the other side of the spectrum, in massive environments, you may have more resources, but you also have more infrastructure and user expectations to manage. Having easy to manage storage benefits every company and organization, regardless of size.

There is another important advantage of true clustered storage that is often overlooked. With a traditional dual node storage system, if one of the controllers is unavailable, then 50% of your resources are no longer

available, including system processors, cache and system memory and internal bandwidth. Consider a true clustered storage system in this scenario. With a three node storage cluster, you would be at 67% of your resources versus 50%. Four nodes would provide you with 75% of your resources, five nodes with 80% and so on. Clustered storage systems with true clustering can ensure that you keep your applications performing at acceptable levels even in the event of losing a controller node. It is important to consider the needs of your users and applications not only in best case situations, but also when something goes wrong. To balance this, you don't need to necessarily over-engineer a true clustered storage system, but it gives you the option to plan according to your needs. Additionally, true clustered storage systems often leverage commodity hardware and therefore are not cost prohibitive, enabling you to build larger clusters without an enormous price tag.

Today we buy more horsepower, memory, capacity and bandwidth than we need. We often design with peak requirements in mind and then add a little more on top of that, projecting out requirements for the next three to five years. As a result, we overspend. Every year CPU, memory, capacity and bandwidth decrease in price and buying these resources in advance is wasteful, especially since in some cases we never fully utilize them. It is not uncommon to have wasted disk capacity and nearly every end-user will tell you that they have some allocated but unused storage. In some cases, it is 50% or more of their total capacity. Additionally, we rarely talk about the wasted processing power, memory or bandwidth. These are harder to measure, but perhaps even more egregious than wasted disk capacity. True clustered storage eliminates this waste.

### **True Clusters versus Management Clusters**

Another type of storage clustering, referred to by ESG as "management clustering," provides ease of management but not scaleable performance or universal access. One popular approach of a management cluster is to combine two controller nodes with active-active clustering and connect this controller node pair to other pairs. Each controller node pair provides a single logical system, which will allow any host to access any piece of data via either node. Each dual node pair creates a logical cluster and will aggregate the hardware resources from the two nodes into a single system. All of the other nodes act in pairs as well and there is no aggregation of performance between the node pairs. The goal of management clusters is to simplify the management of the system, which is useful. However, management clusters don't provide the linear scalable performance of a true cluster.

### **True Clusters versus Global Namespace**

In NAS environments, a valuable virtualization technology, referred to as global namespace, is emerging. Global namespace is easily confused with true clusters since both are forms of virtualization that provide some overlapping value. Global namespace does present a logical layer that enables you to have a single view, access and control of multiple file systems. If you have two or more file systems, global namespace technology will present these multiple file systems to the users and applications as one file system. The value of global namespace is to simplify the management of multiple file systems. In some cases, companies have NAS systems with dozens, hundreds and even thousands of small file systems (ranging from 2 to 16 TB) that are each managed individually. A global namespace allows system administrators to manage numerous file systems and all of the mundane client side management more easily. Additionally, global namespace enables easier access to data since there is one logical path to it versus dozens, hundreds and thousands. Another benefit of global namespace is the ability to move data between different file systems transparently and online. Global namespace can be used to support existing heterogeneous NAS solutions or as part of a homogeneous NAS product. Regardless of how global name space is manifested, it is used to make NAS solutions easier to manage and more fully optimized.

True clustered storage provides similar value. *True clustered storage systems are next generation solutions that are designed to never become complex and to be perpetually optimal.* True clusters natively provide a single level of access. True clustered NAS systems present a single network drive or mount point that supports a massive amount of capacity and can store an enormous number of files therein. True clustered SAN systems offer the same value, but with virtual volumes. Like global namespace, both system administrators and users benefit from the ease of use and access with true clustered storage. The value

true clusters have that global namespace does not is the aggregation of the hardware resources. To re-emphasize, true clustered storage systems pool all of the hardware, including CPU, cache memory, bandwidth and capacity to create a massively parallel system, operating on all of the data regardless of where it is stored within the cluster. Some global namespace solutions do provide some level of scalable performance benefits through manual or automated load balancing, but it is not nearly the same thing.

It is important to note that true clustered storage systems and global namespace solutions can be complementary to one another. You can implement a true clustered storage system to implement new applications and leverage all of its value while using a global namespace solution to better manage your existing complex NAS environment. For example, you could use heterogeneous global namespace technology to move data off traditional NAS systems to a next generation true clustered NAS solution. This can be done for one time moves or as part of an intelligent tiered storage environment.

### **Clustered Storage System Challenges**

Clearly, I am a proponent of clustered storage. Confirmation bias can occur when we support a specific hypothesis and collect all evidence in favor of it while disregarding any evidence against it. In order to ensure that I am not guilty of confirmation bias, it is important to discuss the challenges of clustered storage systems.

While true clustered storage technology is compelling, in and of itself, it is not enough. It is important that we acquire and implement fully realized products and not simply any compelling technology. When evaluating these solutions, it is important to look at the products holistically. Evaluate its management software, features, additional capabilities and road map items. Don't look just at the product, but the company as well. Figure out if they would make a good partner and provide you with excellent support.

You need to understand how granular the true clustered storage systems can scale. You can scale some true clustered storage systems one storage controller node at a time, while others require multiple storage nodes to be added to the cluster in groups.

Most vendors that support true N-way clustered storage architectures typically focus on the scalable performance that it provides. While it's true that these storage systems may provide scalable performance, it may not be scalable for all types of data. Some systems may only provide excellent performance for large blocks of data. Others may provide good performance just for reads and not writes. Still others may be optimized for sequential data. Many clustered storage systems use commodity servers and components and avoid using any proprietary hardware. This approach lowers cost, but it often takes more commodity resources to achieve the same level of performance using proprietary hardware. While the price/performance of commodity hardware may be the same or more attractive, you may need more physical systems to achieve your performance requirements.

The ease of management over the lifecycle of the storage system is even more valuable than scalable performance. Even in environments that are not performance-intensive, clustered storage should be considered when evaluating storage systems. However, it is important that true clustered storage systems start out being easy to use and maintain that simplicity, no matter how large the cluster gets.

## Valuable Today - Requisite Tomorrow

---

Contrary to the common perception, we do not generally embrace innovation in IT. In fact, innovation is often resisted because of the disruption it causes in our current environments. We have invested heavily in what we have already implemented, spending time, money, and resources. We have become experts at managing these complex systems and developing best practices around them. The natural reflex is to continue to invest more time, money and resources doing more of the same, putting more software and hardware on top of the voluminous pile.

I refer to this as the Sisyphus Syndrome. Sisyphus was a brilliant man doomed to roll the same boulder up a hill for all eternity. Even in this scenario, we can at times make progress. Perhaps the boulder gets smaller or the hill less steep, but these are temporary conditions. Other events invariably occur over time that will make things more complex and difficult. And every day we push that boulder back up the hill.

A system administrator with a Fortune 100 company told me that what he feels is the ultimate goal is to get storage systems to reach consumer levels of ease of use. The leading storage systems are far from achieving this since they are complex, semi-virtualized devices that were not originally designed to be easy to use. Interestingly, this same system administrator has dozens of the most complex traditional storage systems and is not likely to replace them any time soon. I am sure he would ideally, like his existing vendors to make what he currently has easier and preserve what he already has in-house. But this is unlikely to happen. For him and many others out there, the notion of replacing all of these systems is not practical or realistic. Or is it?

Right now, the traditional transactional applications are the most associated with traditional storage systems. In large companies, this is still the reality and it is slow to change. However, large companies are implementing different tiers of storage. Disk-based storage archives that have clustered architectures are being implemented to a large extent. Some of these clustered storage systems are supporting hundreds of TBs and racing towards a PB of capacity in a single system.

Midsize companies are embracing clustered storage systems for their primary storage systems, having invested less in existing infrastructure. In many ways, they are able to innovate faster than their larger brethren. Even these companies support storage systems with upwards of 100 TBs. Since they have far fewer resources, they need to implement easier-to-manage storage in order to support the business.

The way we use storage is changing as well. Mission-critical applications encompass more than database, small transactional data types and e-mail. Online applications that are file-oriented and multimedia-based are now mission-critical and revenue generating. In some cases, millions of users are downloading and uploading these large files, images, video and audio. Consider iTunes, YouTube, MySpace and Kodak EZShare. Millions of consumers upload their data and download other content over the Internet. In fact, we are now getting to a point where more people are uploading more data – millions of files on a daily basis -- then they are downloading data.

The dynamics of business are changing. We have incorrectly taken a one size fits all approach, using the same type of storage system for all of our data. This is often the easy thing to do, but ultimately exacerbates the problem. This approach has a negative economic and operational impact.

For those companies that are invested heavily in existing solutions, a recommended strategy is to evaluate highly virtualized storage systems for new applications. That is a way to introduce innovation as non-disruptively as possible. Once you are comfortable with these solutions, you can begin to scale as needed. And eventually, you can begin to move away from the traditional complex and partially virtualized storage systems that populate your data center today and replace them with next generation highly virtualized solutions that were designed for ease of use initially and over its entire lifecycle.

If clustered storage is so valuable, why hasn't it become more prevalent? It just takes time. Any new technology or architecture has a multi-year journey towards dominance. It took 45 years for the automobile to achieve 25% market adoption in the US. We have many challenges, biases, politics, misinformation, confusion, invested time, money, resources and a general lack of education and awareness. This is all part of the maturation process of any technology.

## ESG's View

---

Today, highly virtualized clustered storage systems first have to address tactical needs that can move the business further down the field. However, the larger goal is to virtualize all layers of the data center to create fluidity of resources. Only by doing this can we break out of the "box" and become boundless.

ESG believes that all storage systems need to be highly virtualized. Instead of a one size fits all approach, you will have storage systems that can scale to any size you need. We need to evolve from fixed and finite to fluid and boundless.

The long term vision is that there should be one logical storage system that can meet all of your storage requirements from your most mission-critical, demanding applications to digital archiving and tertiary applications. This storage system will be highly virtualized, support file and block data, consist of lots of physical elements and different tiers of storage, and support intelligent data management and protection features. You can only achieve this vision with high levels of virtualization and clustered architectures. This next, next generation of storage systems needs to liberate IT from worrying about bits and bytes, bandwidth and processing power and nuts and bolts, and instead focus on being strategic to the business. This next, next generation of storage system must provide consumer levels of ease of use, while providing the highest levels of capability and reliability.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources the Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482-0188.