



# Commodity-based Clustered Storage

Affordable storage for the Small and Midsized Business (SMB)

*Bill Mottram*

*Veridictus Associates, Inc.*

*September 2009*

---

**Abstract:** All storage administrators whether employed in large or small enterprises are continually seeking to reduce the costs associated with data storage while improving service levels.

However, if meaningful cost reductions and operational efficiencies are to be enjoyed, a break from traditional thinking such as investing in innovative storage architectures is required. Commodity-based clustered storage is one approach that is gaining traction in small and medium sized enterprises as well as data rich environments serviced by dedicated applications such as video surveillance and electronic medical record storage. This paper discusses how this innovative architecture can deliver many beneficial attributes which reduce the cost and complexity of storage. The technology discussed in this paper is particularly relevant to small and medium sized businesses that need efficient storage without compromising performance.

## 1: Introduction

The storage challenges faced by small and medium sized companies are no different than those experienced by larger enterprises. While IT consumers in Small and Midsized Businesses (SMBs) and larger enterprises have similar service level expectations, the SMB rarely has the specialized storage resources to manage the complexity of demanding Service Level Agreements (SLAs) that are constantly challenged by an aggressively growing storage infrastructure. SMB IT staffs, if available, tend to be generalists with a broad portfolio of IT issues to manage and are unlikely to have the in-depth knowledge of any one specialization, such as storage. Simplicity therefore is highly valued by the SMB.

SMBs have less capital to invest. With less to spend they need to get more out of their budget dollar. Faced with the realities of limited staff, expertise and finances, an SMB is challenged with finding a storage solution that is cost effective, easy to implement, manage, maintain and one that will do the job. They are looking for scalability and flexibility to meet future as well as current requirements.

So with that in mind what are SMBs looking for in a storage solution? What would characterize an ideal SMB solution? The concern that regularly percolates to the top is ease of use. If the complexity exceeds the capabilities of the resources available then failure is unavoidable. This suggests that the elimination of operational complexity is a must. However that is not the only challenge. SMB solutions must be:

1. readily scalable in capacity and performance to accommodate data growth;
2. scalable in time, meaning an existing investment can be easily upgraded with new technology as it becomes available;
3. self-maintaining and self-healing to minimize the necessity and complexity of scheduled and unscheduled service interruptions.

## *Key Attributes of Successful SMB Storage:*

- *Simplicity*
- *Affordable*
- *Scalable*
- *Flexible*
- *Reliable*
- *Data Integrity*
- *Appropriate Performance*

Irrespective of enterprise size, data protection and integrity are key to an organization's health. SMBs, just like their larger brethren, need to meet all the basic storage requirements of business and to the same rigorous availability and accessible standards.

This is the flexibility needed to protect the initial investment and is a fundamental characteristic of a "buy as needed" solution.

That being said, SMBs have not been without solutions. Direct Attach Storage (DAS) with its purchasing simplicity still dominates but as the volume of data and number of users grows, Network Attached Storage (NAS) and Storage Area Network (SAN) are making inroads into the SMB space. It is the improved operational economics of SAN and NAS such as consolidation, storage utilization and improved usability characteristics that is stimulating the increased market share of these technologies in the SMB space. Of the two, NAS with its simpler user experience is the more popular SMB option. However, just like original DAS solutions, the benefits of the relatively inexpensive NAS solutions are soon compromised as they meet the realities of unpredictable but dynamic data growth.

**DAS:** Direct Attached Storage. Storage that is attached to a server and only accessible to that server.

**NAS:** Network Attached Storage. Storage with a thin server operating system and IP network attachment. NAS storage handles file based data.

**SAN:** Storage Area Network. A special network comprising multiple hosts and storage peripherals. SAN storage handles block based I/O.


The purpose of this paper is to look at the conventional approaches that define today's solutions, to discuss their drawbacks and present the notion that storage consumers should be looking at the innovative clustered storage options as effective solutions to their general purpose storage needs.

## 2: The Traditional Approach

Conventional solutions whether DAS, SAN, or NAS have been monolithic architectures, or to use current parlance, scale-up. Monolithic is the traditional approach that traces its roots back to the early pioneering development efforts of IBM and is still the dominant storage architecture in commercial environments. This architecture is characterized by a limited number of controllers, each supporting dedicated storage. Systems tend to be complex with dedicated management resources required for all but the simplest of configurations. With today's challenging data growth exposing their lack of efficiency, lack of flexibility and disruptive scalability, storage consumers are looking for more effective and efficient options to meet their data storage and data management needs.

The traditional scale-up storage approach has four distinct disadvantages:

1. Tends to create inefficient “Islands of Storage” that increase cost and management complexity.
2. Creates bottlenecks that can challenge data access.
3. Storage is generally poorly utilized driving up the cost of usable storage.
4. Tends to be complex and have limited flexibility and disruptive scalability that drives defensive “poor” purchasing behavior.



*Monolithic or scale-up architectures are the traditional approach where storage sits behind a dedicated controller or server head.*

5. Have inherent single points of failure that challenge reliability and data access.

These architectural failings drive complexity and cost and as a consequence, they encourage poor purchasing behavior. An example would be the desire to avoid the service interruptions, unavoidable with the disruptive scalability of monolithic architectures, causing more storage to be purchased than the immediate business requirement demands. By purchasing now what is needed later, the costs are based on today's pricing. That's the wrong way to purchase storage considering disk drive cost is volume dependent. As the market volume increases, cost decreases at a historical rate of 8% to 10% per quarter. Storage density improvements also affect cost, as density increases cost decreases. Buying what's not needed now is essentially an overpayment for future storage. This poor purchasing behavior not only unnecessarily increases the acquisition cost but also the operational cost such as energy and cooling.

While these are all characteristics that work against the storage consumer by introducing cost, complexity, and angst into the routine tasks of managing storage and surviving data growth, monolithic structures will not disappear. They have a place and purpose particularly in large data centers, virtualized environments and in the support of dedicated applications.

However there is a growing interest in clustered storage which appears to offer an interesting set of alternative attributes that directly address the failings of monolithic options.

But before moving on to discuss clusters, a word of caution: beware of imposters. A clustered pair is not clustered storage; it is simply dual controllers with a failover application to provide a level of fault tolerance inherent in true cluster architectures. Beware of modular architectures that carry the cluster label but can't easily scale because of architecture limits or disruptive upgradability. These solutions are not clustered storage architectures.

The purpose of the remainder of this paper is to provide the reader with insight into the characteristics and user benefits of true cluster architectures and hopefully impart some wisdom to separate fact from marketecture.

### 3: Commercial Clustered Storage

#### The evolution of clustered storage

Clustered storage has its roots in the High Performance Computing (HPC) world where researchers needed to solve the cost/performance dilemma associated with their mainframe resources. With the dramatic evolution of commodity servers in price, performance and size, innovative thinkers realized that by aggregating the compute power of these individual resources their problem was solved.

Clustered storage followed a similar evolutionary path as clustered computing and while it is true the compute performance dilemma was mitigated, the unintended consequence was to move the performance bottleneck to storage. Large monolithic “big iron” storage solutions could no longer deliver the performance needed.



*Clustered or scale-out architectures - flexible and scalable meeting the challenge of aggressive capacity growth, unpredictable workloads and bandwidth requirements. Coherent peer relationships eliminate single points of failure, increasing data access and availability.*

**Clustered storage** solutions are a networked collection of nodes and irrespective of the number of nodes in the cluster there is only one logical management view.

The solution was to logically integrate (network) standard servers (storage servers/controllers) and basic storage units (JBODs), with software that manages the physical federation of these standard components (cluster nodes) and a unique file system (clustered file system) that manages the hosted data.

**Clustered File System** is a distributed file system that is not associated with a single server or particular group of clients but with a cluster of storage servers otherwise known as cluster nodes.

**Cluster Nodes:** The physical compute and storage components normally referred to as nodes that can be either compute (standard servers), storage (JBOD/Array) or hybrid nodes. Hybrids have both compute and storage resources in the same physical unit, simplifying linear performance and capacity growth.

**Clustered File System (CFS):** The intelligence that manages the data within the storage cluster. A clustered file system is a distributed file system that is not associated with a single server or particular group of clients but with a cluster of storage servers otherwise known as cluster nodes. The CFS services initiator requests irrespective of the node within the storage cluster receiving the request.

Cluster file systems have outstanding elasticity supporting a scalability that exceeds the architectural capabilities of comparable monolithic options. Clustered storage solutions can scale from the low terabyte (TB) range to very large pools of storage measured in petabytes (PBs). What tends to limit the upper limit is the physical capability of the storage architecture.

However, not all clustered file systems are created equal. The highest performing systems are parallel architectures that support parallel data access while allowing all nodes to concurrently access the same files. However, even parallel file systems are not created equal with differing approaches to concurrency control. The difference is two disparate approaches, namely symmetrical, that require clients to run metadata manager code, and asymmetrical that has dedicated metadata managers.

If any node or other component fails in a clustered architecture access to data is not compromised; there is no single point of failure.

Each node within a symmetrical cluster is a logical peer to all other nodes within the cluster. The file system is maintained across the entire cluster, so by definition this architecture delivers a robust, highly available data storage solution optimized for high availability. If any node or other component fails, access to data is not compromised, the node replacement is non-disruptive, and there is no single point of failure. This flexibility delivers improved reliability, accessibility, serviceability and upgradability. However, because of the distribution of the metadata, a symmetrical design has the potential to impact high-end performance, but unlikely to be an issue in general purpose commercial computing.

An asymmetrical cluster is similar but the need for dedicated metadata nodes compromises the symmetry and introduces a single point of failure, diluting the overall solution robustness.

## Clustered storage in the commercial world

Until recently, clustered storage was restricted to the HPC world but times change and innovation in clustered and data management technologies have enabled the growing adoption of clustered architectures in the general purpose computing domain. This acceptance is recognition that today's data challenges are requiring innovative solutions and the status quo is not an option. The new age data challenges are driven by the Internet, streaming media, digital photography, electronic storefronts, email, and compliance - once the exclusive enterprise challenges of the 1990s but are now being faced by the SMB.

So what are the attributes that differentiate a clustered storage solution and underlines its significance for commercial application, particularly for the SMB?

**Simplicity:** Independent of the number of nodes and the total capacity (TBs to PBs) that make up a cluster, there is only one management view which means significantly simplified management. A clustered solution will normally automate the traditionally manual task including data protection levels, load balancing and content rebalancing when the node count is changed. As the storage cluster grows from TBs to PBs the management view does not change and the burden on the administrator remains constant which means headcount requirements are independent of system scale.

**Data Accessibility:** The cluster intelligence is delivered by embedded cluster software that enables all nodes to work together seamlessly and act as a single system. Each cluster node is a coherent peer meaning each node is fully aware of all the other nodes. Should one node fail, requests are transparently pointed to another node for service. This assures data access without the traditional complexities of failover and multipathing.

**Reliability and Serviceability:** Clustered architectures are intrinsically robust and fault tolerant. Components can be replaced without the need for scheduled downtime. Self-healing techniques are becoming increasingly pervasive to protect against both device and data failure. These techniques contribute significantly to overall system robustness and serviceability.

**Common Storage Pool:** Contrary to most monolithic architectures, clustered solutions exploit a common storage pool that is accessible to all clients. This eliminates the phenomena of stranded storage that creates wasteful islands of storage and allows the consumer to use all the storage purchased.

**Flexibility and Scalability:** By their very nature, clustered architectures are unique in their ability to scale easily, transparently and in a near linear fashion. When nodes are added to the cluster, the file system transparently discovers the new node and aggregates the additional resources to achieve a near linear scalability in performance and capacity.

**Clustered storage** is unique in its ability to scale easily, transparently, and in a near linear fashion.

Accepting the notion that the aggregate power of a cluster configuration tends to be more than a comparable monolithic option and with single cluster instances that can scale beyond traditional monolithic limitations, holds significant promise for the data intensive SMB.

**Clustered storage**

A flexible solution that can simply support dynamic and unpredictable data growth.

The granularity that clustered architectures allow and the ease when adding, replacing or upgrading nodes without disruption to data access gives the administrator the flexibility to upgrade capacity and performance to match actual application needs and respond quickly to unpredicted demand.

This flexibility creates an opportunity to deliver a realistic “buy as needed” option where users can comfortably acquire storage as needed and not absorb the cost of buying forward. This practice of convenience is a waste of scarce CAPEX and OPEX dollars.

**Transparent Technology Refreshes:** Scalability presents a multi-dimensional challenge, requiring more thought than the traditional perspective of simply managing data growth.

**Investment protection**

Make sure the vendor of whatever system you choose guarantees at least one generation of backward compatibility of cluster nodes.

While data growth, increasing application demand, and increasing service level expectations are all elements considered, scalability through time is often missed. Simply put, during the life of an implementation, replacing compute or storage resources for the latest, fastest, densest, more efficient version is a likely scenario, and one that could add significant life to the initial investment. Consider the benefit of avoiding the traditional three to five year cycle of painful system upgrades and time consuming data migrations. Just make sure the vendor of whatever system you choose guarantees at least one generation of backward compatibility of cluster nodes (basic investment protection).

**Commodity Components:** Clustered solutions tend to use industry standard commodity components that are tightly integrated with cluster management software. The use of commodity components offers two distinct advantages; the first being an opportunity to ride a significant quarterly price decline in the order of 8% to 10%; the second is a dynamic innovation curve.

**Declining commodity price curve:**

Average price decline of hard disk drives is between 8% to 10% each quarter.

Figure 1, shows the historical pricing trend for magnetic storage from 2008 to 2010. This is a clear illustration of the 8% to 10% quarterly price decline lost to consumers who buy in advance of actual need without securing some price guarantees.




Figure 1: The historical pricing trend for magnetic storage from 2005 to 2010. (source: [www.archivebuilders.com](http://www.archivebuilders.com))

Another characteristic of commodity components is its dynamic innovation curve enabling innovative design teams to respond aggressively with new technology driven options, gaining a competitive advantage over vendors dealing with traditional designs that have invested in custom silicon.

#### 4: Unified Storage

While not unique to clustered storage architecture, unified storage has evolved as a novel way to reduce IT complexity, improve business flexibility and reduce costs. It is a technology implementation that holds a lot of promise for cost conscious storage consumers.

To understand the value of the concept, first consider the nature of enterprise data. Data can either be structured (block) or unstructured (file-based), with block IO needing SAN storage and file-based needing NAS storage. The challenge for the SMB is the complexity of managing disparate storage systems. However, the notion that unified storage does not care what type of data—block or file—you send to it, solves the dilemma.

Focusing on IP based storage, a unified solution supports both iSCSI (block) and NAS (file) protocols offering the consumer SAN and NAS options. This flexibility enables storage administrators to consolidate multiple workloads onto a single storage platform. Such systems simplify the complexity of managing disparate storage solutions for the SMB while the common storage pool concept of clustered storage significantly increases storage utilization: a synergistic blend of technologies. This increasingly popular concept is an interesting consolidation option and could be a boon for the resource and budget strapped SMB.

So what are the attributes of unified storage that are likely to resonate with storage administrators and their bosses?

1. **Flexibility:** enables storage administrators to consolidate multiple workloads onto a single storage platform.
2. **Hardware consolidation:** elimination of separate SAN and NAS hardware, replacing two physical systems with one.
3. **Simplicity:** with the elimination of siloed storage and having only one box to manage the day-to-day management is significantly simplified.
4. **Improved economics:** reduced hardware reduces capital and operational expenses.
5. **Improved storage utilization:** a shared storage pool is more efficient by eliminating stranded storage.
6. **Green:** obviously an approach with positive environmental benefits.

*Unified Storage  
does not care  
what type of  
data, block or file,  
you send to it.*

Unified storage is a concept that should appeal to all storage users, particularly the small and medium sized user. This community has all the same data management challenges as the bigger folks but has less budget dollars, less staff and less expertise which compounds the complexity of managing disparate storage systems.

## 5: A Pragmatic Approach to Clustered Storage for SMB

Commodity-based Clustered Storage (CCS) from Scale Computing is the latest introduction into the clustered, unified storage arena. This new participant has developed an innovative clustered storage solution which has all the attributes that resonate with SMBs who are looking for a reliable, affordable storage solution that meets the needs of their business and does not take a dedicated storage professional to manage.



The minimum Scale Computing configuration consists of three nodes networked through a dedicated switch. Nodes are based on a Supermicro server hosting four 500GB drives (SN1000), four 1TB drives (SN2000), or four 2TB drives (SN4000), delivering usable storage of 1TB, 2TB, or 4TB per node respectively. Additional nodes are easily and non-disruptively added to the cluster, linearly increasing both capacity and performance. Average time to add a node is five minutes, with an average time to first customer data of 20 minutes. Scale Computing clusters support mixed nodes ensuring a positive future upgrade path for initial investments as drive densities continue to grow and media technologies develop.

Intelligent Clustered Storage™ (ICM), the storage cluster management and file system, is a tight integration of licensed intellectual property, open source software and Scale Computing's proprietary technology. The ICM architecture automates storage services which reduces management complexity, supports what is called predictive hardware failure technology and enables a powerful clustered file system. All nodes are managed through a user interface (having recently watched a live demonstration, justifies a description of being intuitively simple).

ICM is a symmetrical clustered solution which as already explained, is optimized for high-availability and serviceability. Its foundation is IBM's GPFS which has already earned considerable credentials in the HPC world but in this instantiation has been adapted by the Scale design team to deliver not just its powerful functionality but to do so with the ease of use that enables it as a practical option for the commercial world.

The Scale Computing architecture falls into the unified storage category. Block and file protocols are natively supported on each node enabling the Scale Computing solution to consolidate multiple storage solutions into one. This feature significantly simplifies storage management by enabling storage administrators to consolidate block and file workloads into a single appliance; an interesting consolidation option for the SMB that was only previously available on much more expensive solutions.

## **Solution Highlights:**

- Nodes are based on 1U Supermicro storage servers with support for IBM platforms scheduled for 3Q2009
- Each node accommodates 4 drives with usable capacity configurations of 1TB (SN1000), 2TB (SN2000) and 4TB (SN4000).
- Minimum configuration is three nodes with a usable storage capacity of 3TB which Scale Computing calls their starter kit.
- No single point of failure in the cluster architecture with high availability (HA) features standard with all configurations.
- Automated data protection (RAID 10-mirror copies). No user intervention required.
- Self-healing: when a drive or node fails the data is rebalanced automatically. As mirroring is employed RAID rebuild times are not a factor.
- Current maximum configuration is 2.2PB, based on 1TB drives.
- ICM is a symmetrical, parallel clustered file system based on IBM GPFS.
- 210MB/s throughput with a three node cluster; scales linearly with additional nodes.
- Data migration-less scalability. Automatic load rebalancing as system capacity grows.
- Native multi-protocol support (iSCSI, NFS, CIFS/SAMBA).
- Two GigE ports per node.
- Pricing: At the time of writing the entry list price was \$13K (for minimum configuration of 3xSN1000); and a per “usable TB” cost that drops to \$2500 (when purchasing the SN2000); and a \$21,000 cost for a 12TB SN4000, usable TB cost of \$1,750. These prices will certainly create indigestion for the current vendor establishment.

## **6: Final Thoughts**

Clustered architectures are quickly emerging as the storage option that will deliver the flexibility needed in today's increasingly competitive business environment which despite current economic challenges will still experience data growth in the order of 40% to 50% CAGR. This is not just an ongoing management challenge but a significant budgetary one.

Clustered architectures can deliver a number of significant benefits including simplicity, accessibility, scalability reliability, serviceability and affordability to the storage consumer and to the SMB in particular. Unique are the scalability and non-disruptive upgrade attributes that enable sound purchasing decisions that optimize shrinking IT budgets and protect investments against early technology obsolescence.

True clustered storage is an architecture that should be considered by all storage professionals particularly those data rich small and medium sized enterprises, satellite offices and dynamic departmental applications.

**Bill Mottram** compliments his role as the managing partner of Veridictus Associates Inc, a high technology marketing consultancy, as a Principal Contributor for the Wikibon Project and as an analyst practitioner in the data storage industry with Data Mobility Group. He has over 25 years product development, marketing and sales experience in the data storage, information technology and medical device industries. Experienced with Fortune 500 companies such as StorageTek, Compaq and Hewlett Packard and smaller, entrepreneurial enterprises including Pillar Data Systems and COPAN Systems, he was the marketing leader responsible for the development, introduction and “go-to-market” activities for a number of innovative and highly successful data storage solutions. Although primarily a marketing professional, Mottram has considerable expertise in technology management, corporate acquisition and public company funding.

A graduate of Paisley College of Technology (electronic engineering) and holds a Masters of Business Administration from Pepperdine University.

Mottram can be contacted at [bill@veridictusassociates.com](mailto:bill@veridictusassociates.com)

Copyright © 2009 Veridictus Associates Inc, All Rights Reserved.

Reproduction of this publication without prior written permission is forbidden. Veridictus Associates Inc believes the statements contained herein are based on accurate and reliable information. However, because information is provided to Veridictus Associates Inc from various sources, we cannot warrant that this publication is complete and error-free. Veridictus Associates Inc disclaims all implied warranties, including warranties of merchantability or fitness for a particular purpose. Veridictus Associates Inc shall have no liability for any direct, incidental, special, or consequential damages or lost profits. The opinions expressed herein are subject to change without notice. All trademarks are the property of their respective companies.