

A Breakthrough Approach to Storage and Backup

Converged Storage and Backup Dramatically Cuts Costs,
Eliminates Complexity, and Delivers a Complete
Disaster Recovery Solution



Table of Contents

Executive Summary.....	3
Introduction.....	4
Enterprise Storage Challenges	5
Storage Technology Evolution	5
The Nimble Storage Approach	6
How It Works.....	7
Introducing the Nimble CS-Series.....	12
Summary	14

Executive Summary

Today's enterprises face significant IT challenges. Data storage requirements are increasing exponentially, and end-user expectations for fast and highly available applications continue to rise. Existing backup and disaster recovery schemes are at the limits of their capabilities, but budgets for new infrastructure are tight. As a result, many enterprises have run out of ways to cost-effectively address their IT challenges.

Nimble Storage has developed an entirely new approach to data storage that converges primary and backup storage into one array, dramatically cutting costs and complexity. And with a second Nimble array at a remote location, enterprises can also get a complete, cost-effective disaster recovery solution.




Figure 1: Nimble's breakthrough approach is based on converged primary storage, backup, and disaster recovery.

This white paper reviews the challenges faced by today's IT managers. The limitations of current storage architectures are examined, followed by a review of several key storage technologies that are in use today. The paper will conclude with an introduction to the Nimble Storage solution, describing the functionality, features, and benefits that enable enterprises of all sizes to solve their most significant storage challenges in a new, powerful, and cost-effective way.

"Nimble Storage has built what IT administrators might call the Holy Grail of high-performance data storage. For its fault tolerance, performance and innovative software, the *CRN* Test Center strongly recommends the CS240 from Nimble Storage."

EDWARD CORREIA
STAFF WRITER
CRN TEST LAB

Introduction

IT managers have difficult jobs. They must support the growing and continually changing infrastructure requirements of their businesses, and invariably do so with limited resources. They fight to keep pace with upcoming IT changes and technology trends, while managing major long-term initiatives and fighting fires on a weekly basis. They are underappreciated when things work well – but under pressure when IT challenges impact end users or the enterprise's key applications.

The typical IT team at a small or mid-size enterprise often consists of a handful of IT generalists who must wear many hats and cover a broad range of IT disciplines. Despite the lack of specialized expertise on every component of their storage infrastructure, they still must ensure that all key applications perform adequately and reliably, and that essential data is protected against a spectrum of potential failures.

The advent of virtualization has both been a boon and a curse for IT teams, as they now juggle an accelerating set of demands for new VMs, which in turn require more storage and protection against failures. The IT team not only has to manage these primary storage challenges, they must ensure all new primary data is backed up each night, with a complete backup of all systems over the weekend. Due to ever-increasing amounts of data, it is a challenge to complete these backups in time without impacting day-to-day business operations. Designated backup windows are often exceeded, requiring constant juggling of operations.

Many enterprises are acutely aware that their existing DR solutions are not adequately protecting the company against major site failures. How much important data might be lost forever if a disaster did occur? And how long would it take to get applications back online if they did go down? Unfortunately, the IT manager's ability to do better is constrained by budget, time, and logistical issues, including WAN bandwidth. As a result, most enterprises have already run out of options to keep up with the continued growth, cost, and complexity of their IT environments.

Fortunately, there is a new solution that can solve their most pressing storage problems quickly, efficiently, and cost-effectively.

Enterprise Storage Challenges

When IT managers are asked to describe the many barriers that keep them from solving the enterprise's storage challenges, their answers will usually focus on the following critical issues:

1. **Storage is too expensive.** Powerful new business and productivity applications require costly primary storage systems with high-RPM drives and secondary storage systems for backup. However, tight budgets limit IT's ability to deliver the performance and capacity needed to run these applications and cost-effectively protect growing amounts of data.
2. **It is increasingly complex to administer storage, backup, and DR.** Very few IT managers can afford to hire IT staff with the specialized training needed for every different system or server. Enterprises need easy-to-use storage solutions that can be efficiently managed by IT generalists.
3. **Backup and restore processes are slow and consume too many resources.** For the typical enterprise, there is simply too much data to back up during allotted maintenance windows. Current backup and restore processes take far too much time and consume excessive amounts of bandwidth.
4. **Disaster recovery solutions are either too costly or inadequate.** There are currently no practical disaster recovery solutions that will fit into the budgets of small- and mid-size enterprises. Inadequate DR puts businesses at serious risk for data loss and significant downtime in the case of a disaster.

Storage Technology Evolution

To better understand why existing storage systems cannot solve today's IT problems, it helps to step back a few years and examine several key storage technologies that are still in use today.

Ten years ago, enterprises used modular storage arrays consisting of a primary storage head unit with one or more shelf units attached to it. However, this storage model was not easily scalable. Enterprises had to predict their performance and capacity needs well ahead of time, or risk expensive forklift upgrades to accommodate growth. With their reliance on expensive, high-RPM drives and complex pricing structures, these storage arrays were expensive and complicated to operate without years of specialized training and experience.

Managing complex backup processes took far too much time, and failed backups presented a regular challenge. Restoring data quickly was an even bigger issue. Smaller enterprises typically relied on tape storage as the only affordable backup and DR solution. With the need

to retrieve tapes from offsite locations and then copy huge volumes of data from tape onto the primary systems, it would often take several days to get applications back online.

Storage technology made significant progress over the last decade with the introduction of the first scale-out storage systems. Enterprises were finally able to purchase extra storage incrementally without requiring forklift upgrades. This presented an appealing way to manage data growth, but storage costs were still too high. To get the capacity and performance required to run key applications, enterprises had to buy systems with many expensive, high-RPM drives.

The next step forward was the introduction of a new class of storage arrays designed specifically for disk-based backup instead of tape. Deduplication capabilities were built right into the backup systems, delivering up to 20x compression rates. Enterprises were now able to store 60 to 90 days' worth of backups in a reasonable amount of disk.




Figure 3: Traditional storage systems suffer from long backup windows and slow recovery times.

But deduplication did nothing to solve the problem of copying large amounts of data from primary storage to the backup systems. Processes still took far too much time, often exceeding designated backup windows. As such, companies couldn't replicate data to offsite locations more than once per day. If there was a disaster and the primary site went down, they could still lose a day's worth of data. Enterprise storage was better off than before, but despite these innovations, a number of significant problems remained.

The Nimble Storage Approach

Nimble Storage has developed an entirely new approach to storage, based on the concept of converged storage, backup, and disaster recovery. Nimble was founded by a team of industry veterans who helped develop groundbreaking primary storage and deduplication architectures at several market-leading storage companies. The Nimble Storage approach gives enterprises an innovative storage solution that solves their most pressing IT issues: reining in storage complexity, reducing costs, and implementing affordable backup and recovery.

There are two ways technology vendors can embrace innovation. One is to try to bolt new functionality onto existing legacy architectures. But this approach invariably includes costly tradeoffs and compromises, and results in very complex solutions. The other approach is to start from scratch. But for most large storage vendors, throwing away everything and starting fresh is not an option. Not only do these vendors need to protect their existing revenue systems, they are also constrained by their “legacy mindset” approach to new innovations.

Only new industry entrants have the luxury of starting with a clean slate – freeing them to use revolutionary ways to solve significant industry problems. Nimble Storage has taken this “blank slate” approach to innovation by designing a completely new storage solution from the ground up.

How It Works

Nimble Storage leverages the following recent advances in storage technology:

- The availability of affordable, high-density flash memory
- The use of cost-effective, high-density, low-RPM drives for primary storage
- Automated data placement technology
- Primary and backup storage capacity optimization

All Nimble Storage solutions are based on a fundamentally different architecture: Cache Accelerated Sequential Layout (CASL™). The following sections provide a high-level overview of Nimble CASL, describe how it works, and introduce the functionality and benefits it delivers to today’s enterprises.

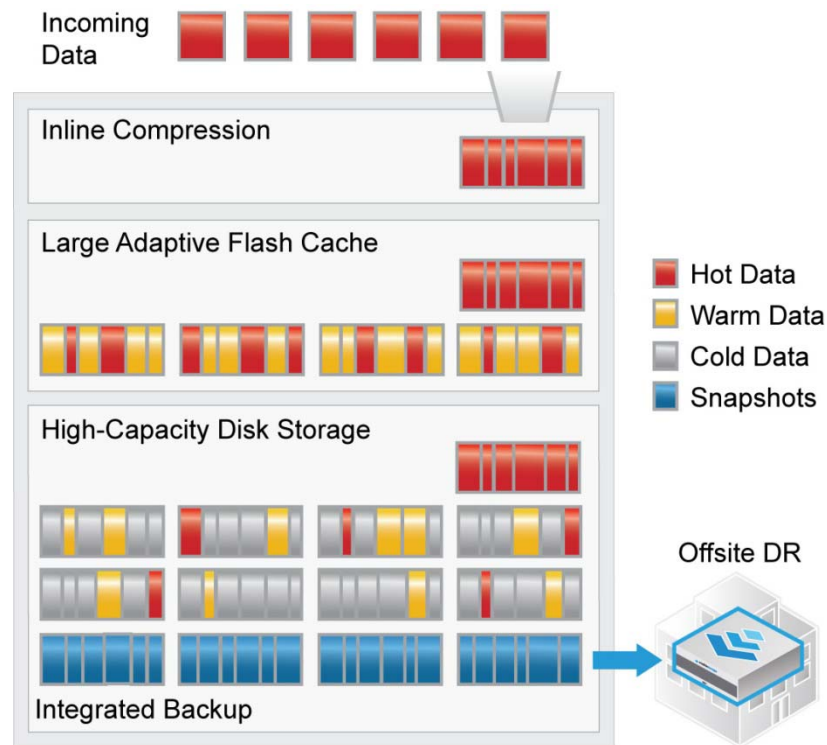


Figure 4: Nimble’s Cache Accelerated Sequential Layout (CASL).

Inline Data Compression for Primary Storage Optimization

Nimble Storage CASL provides in-line compression on all data. Existing primary storage architectures cannot effectively support inline compression because they store data as fixed-size blocks. In contrast, Nimble’s CASL architecture was designed to support variable-sized blocks inherent in compressed storage.

With CASL, compression is done in real time with no added latency as these random blocks of data are written to the storage array, typically reducing the amount of stored data between a factor of two and four. For email applications such as Microsoft Exchange, a 2x compression ratio is usually achieved, and with applications such as Microsoft SQL Server or other databases, 4x or even greater compression can be obtained. This immediately halves the cost of storage for most enterprises.

Combining Flash and High-Capacity Drives in One Device

Recent advances in flash memory hold the promise to transform enterprise storage. Flash-based devices, such as solid state drives (SSDs), do not have the mechanical limitations and latencies of hard drives. As a result, they can deliver very high I/O performance with a modest space and power footprint. However, new flash offerings from existing storage vendors have been too expensive, putting them well beyond the reach of most mid-size enterprises. These vendors have been constrained by trying to

“Nimble Storage’s breakthrough architecture will have a profound impact on the IT organization, on par with the server virtualization revolution. With Nimble, I can streamline my infra-structure so that I no longer have to silo storage from backup and disaster recovery. That’s world-changing for me.”

DAVE CONDE
IT DIRECTOR
eMETER

retrofit this technology into their traditional storage architectures that were never designed to utilize flash.

High-capacity, low-RPM drives can cost up to 85% less than high-RPM drives on a per GB basis. But systems built with low-RPM drives cannot deliver the random access performance needed by enterprise applications such as databases, limiting their use in existing architectures to archives, disk-based backup, and other sequential or low performance applications. Nimble's architecture is unique in being designed from the ground up to work with flash storage – in conjunction with low-RPM drives – enabling very fast performance at a very low cost.

Intelligent Data Optimization

Nimble's CASL architecture automatically optimizes data on a block-by-block basis. A copy of all active or "hot" data (and even "semi-active" data) can be held in the flash layer, enabling very fast random read performance. In addition, Nimble intelligently sequences random writes in a way that enables extremely high speed random write performance to both the flash layer as well as the high-capacity disk layer. Inactive or "cold" data resides only on compressed high-capacity hard drives, further reducing costs.

CASL architecture adapts to hot spot and application workload changes within milliseconds, compared with hours or days using tiered approaches. This maximizes the efficient utilization of the flash storage capacity. For example, there may be a terabyte of total database data, with only 3-5% being accessed in any given interval. But if the exact location of the active data changes periodically, the only solution with traditional architectures is to keep all of the database data on a terabyte of high-performance storage, wasting expensive storage capacity. In contrast, CASL dynamically moves hot data into the flash layer within milliseconds, serving that data with very high performance and very low latency.

Instant, Integrated Backups

Virtually all enterprise storage systems provide the ability to take snapshots, or point-in-time copies of data. Snapshots provide a very fast and efficient way to recover from user errors and application crashes. However, enterprises are usually limited to keeping at most a few days' worth of snapshots on disk, because storing multiple snapshots can consume large amounts of expensive and limited primary storage capacity. This mandates backing up primary data to a separate disk-based backup system frequently in order to store 60-90 days' worth of backups.

CASL enables instant, application-consistent backups on the same array with very efficient (up to 20x) backup capacity optimization. It captures compressed, incremental snapshots of the primary data on low-cost drives at pre-configured intervals. These backups are both extremely fast (on the order of milliseconds) and non-disruptive, because they do not impact

application or storage performance. It therefore becomes feasible for enterprises to capture frequent backups spaced just minutes apart (instead of the typical daily backup), enabling much finer recovery point objectives (RPOs) and improving data protection.

The Nimble approach significantly reduces costs by eliminating the need for a separate backup tier. It also eliminates the complexity and administrative overhead of managing massive data copies between tiers. Finally, it enables much faster restores, with recovery time objectives (RTOs) measured in seconds instead of hours in the event of application corruption or user error, further improving application availability.

Efficient, Cost-Effective Disaster Recovery

CASL architecture provides very efficient off-site replication capabilities for moving data to a second, remote Nimble Storage array. It only transfers highly compressed block-level changes, requiring low bandwidth to move data across the WAN. Furthermore, the remote replicas are primary data copies. They can be quickly accessed in the event of a disaster by simply failing over clients or applications to the remote array – without the intermediate step of having to restore the primary system from a backup solution.

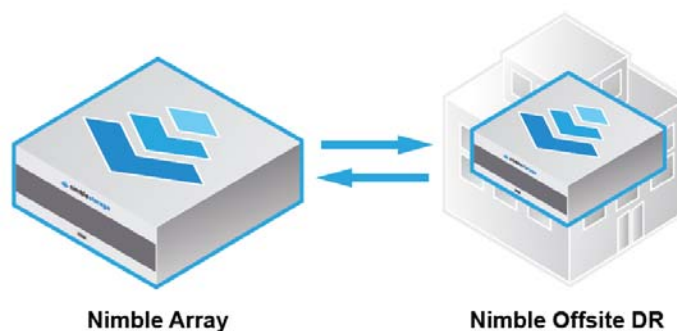


Figure 5: Nimble Storage arrays enable efficient, cost-effective DR.

Nimble Storage arrays support bi-directional and N-to-1 replication, allowing multiple sites to serve as a disaster recovery sites for each other. The arrays have the added flexibility of supporting asymmetric retention across sites, meaning the data retention period on a DR site can be different than the retention period at the primary site.

With Nimble arrays, IT managers can ensure application availability within an hour by simply failing over applications to use the remote replica of the primary data. When the primary site becomes available again, all intermediate data changes can be synchronized with the primary array, and applications failed back to the original site.

Local, integrated backups combined with offsite disaster recovery provide a comprehensive data protection solution for mid-size enterprises. The Nimble solution improves application

availability and provides business continuity at only a fraction of the cost of alternative solutions.

Dramatically Simplified Management

Designed for the IT generalist, Nimble's intuitive UI eliminates the complexity of provisioning and managing separate storage, backup, and disaster recovery devices. Using predefined application and data protection policies, new volumes can be created in just three simple steps. In addition, rich monitoring and reporting capabilities simplify troubleshooting and ongoing administration. The home page provides a system overview, including space usage and savings, system throughput and I/O activity, and recent alerts and events that may require attention. The volume summary page offers an at-a-glance view into the status and space usage of each volume.

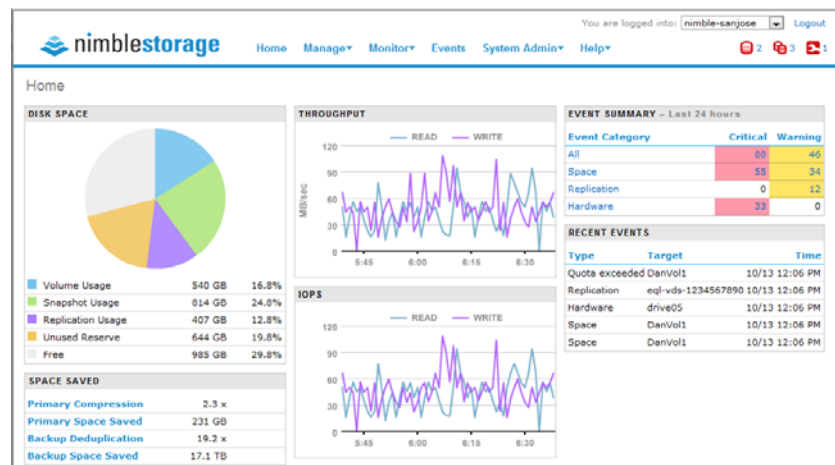


Figure 6: The Nimble Home Page

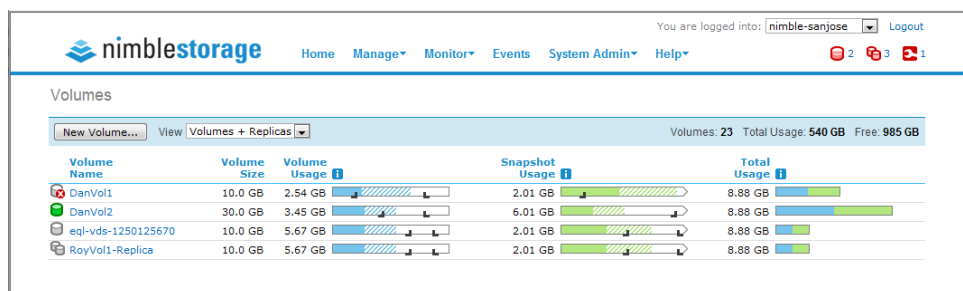


Figure 7: The Nimble Volume Summary Page

Storage Optimized for Microsoft and VMware Environments

Nimble Storage arrays are designed to tightly integrate with Microsoft applications including MS Exchange, MS SharePoint, and MS SQL Server, as well as VMware and Microsoft HyperV for virtualization. When creating a new volume for a specific application, users can

automatically optimize storage and data protection settings for the application type by simply picking the application from a drop-down list.

Nimble Storage arrays enable cost-effective, end-user driven restores of deleted files or emails through capabilities such as “Previous Versions” in MS Windows, or the “Deleted Item Restore” in MS Exchange, thus reducing the reliance on administrator-driven restores. Tasks such as efficiently cloning a VM can be accomplished right from the VMware vCenter user interface. By elegantly integrating a valuable capability like zero-copy clones, the Nimble Storage solution doesn’t just save time and money – it also greatly simplifies VM management.

Introducing the Nimble CS-Series



Nimble Storage CS-Series arrays are designed for enterprises whose applications demand both high performance (typically requiring expensive high-RPM drives for primary storage), and robust data protection (usually enabled by disk-based backup solutions and complemented by offsite replication). The arrays are highly available by design, with fully redundant and hot-swappable components including controllers, power supplies, fans, and drives. The arrays are protected by safeguards such as dual parity RAID, as well as end-to-end data checksums.

All advanced capabilities are included with the CS-Series at no additional charge, including thin provisioning, replication, efficient and fast cloning, and application-integrated data protection. And finally, Nimble Storage arrays are protected by robust diagnostics and support capabilities, with background data integrity checks, hardware diagnostics, as well as phone-home capabilities to enhance support levels.

All Nimble Storage arrays ship with built-in Nimble Protection Manager, which manages application-consistent backups, plus VM system-consistent backups and replicas for application data residing on the arrays. Pre-existing data protection policies (such as snapshot and replication schedules, retention goals, and cross-volume consistency) can

“Nimble Storage’s approach of utilizing flash and SATA in a way that combines primary storage and snapshot-based backup in a single solution, dramatically lowers equipment costs, reduces backup and restore times from hours to seconds, and streamlines processes.”

LAUREN WHITEHOUSE
SENIOR ANALYST
ENTERPRISE STRATEGY

be assigned to groups of data sets, and new policies can be easily created.

Backup integrity can be automatically verified using database verification tools at pre-configured schedules. Data can be instantly restored from any available backup by selecting it from the management user interface, speeding up and simplifying application recovery.

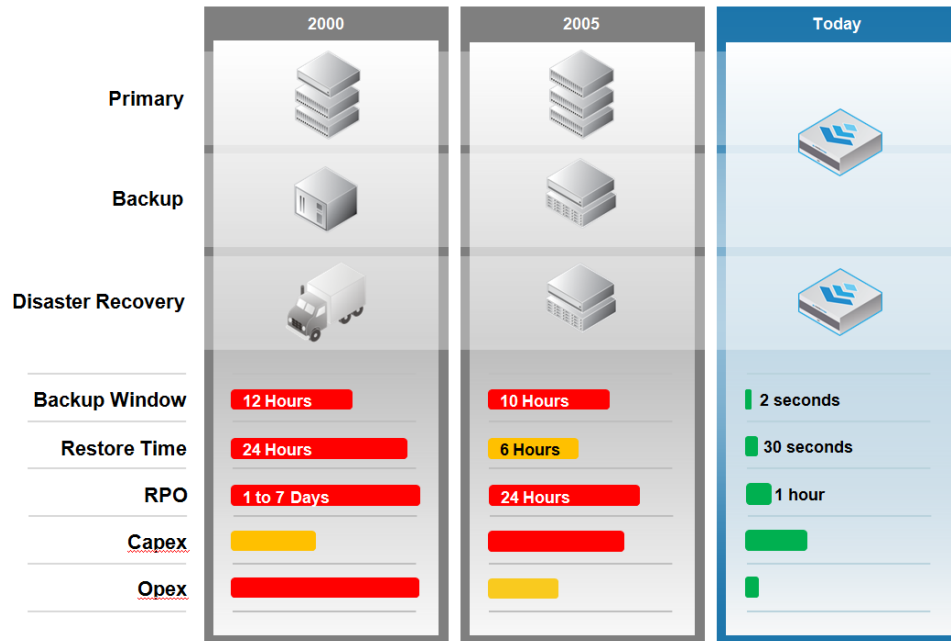


Figure 8. Nimble Storage significantly shortens backup windows and restore times, reduces capex, slashes opex, and delivers a complete disaster recovery solution.

Nimble Storage delivers the reliability and high availability that today's business-critical applications demand. All Nimble Storage arrays include fully redundant hardware components and high-availability software and support features, enabling enterprises of all sizes to increase service levels and improve the reliability of their critical applications and data.

Summary

With very small IT teams, limited budgets, complex applications, and rapidly growing data, today's IT managers are having serious difficulty just keeping pace with business demands. Unfortunately, current storage solutions based on existing architectures are not adequately addressing these challenges. It's time for a fresh approach.

Nimble has developed an entirely new approach to storage that converges storage, backup, and disaster recovery into a single solution. Nimble's breakthrough CASL architecture combines flash memory with low-cost, high-density drives, eliminating the need for expensive, high-RPM drives for primary storage and a separate disk-based backup solution. Nimble Storage reduces backups and restore times from days to seconds, and enables enterprises of any size to finally implement an affordable disaster recovery solution.

For more information on how Nimble Storage can help today's enterprises cost-effectively meet their storage, backup, and disaster recovery needs, visit www.nimblestorage.com, or call 877-3NIMBLE (877-364-6253).

"The large capacity of the Nimble CS240, along with its tight integration with Microsoft SQL Server and Exchange, allows us to completely rethink the backup process, significantly reducing complexity and costs associated with data protection."

CRAIG CONDIE
INFORMATION SYSTEMS DIRECTOR
MOHLER, NIXON & WILLIAMS



Nimble Storage, Inc.

2645 Zanker Rd., San Jose, CA 95134

Tel: 877-3NIMBLE (877-364-6253) | www.nimblestorage.com | info@nimblestorage.com

© 2010 Nimble Storage 11/17/10