

No-Compromise Storage for the Modern Datacenter

How Nimble Storage Delivers Performance and Capacity Efficiency,
Better Data Protection, and Stress-free Operations



Table of Contents

Modern Storage Challenges.....	3
Leveraging Innovation.....	4
A Fresh Approach to Storage.....	5
The Technology Behind No-Compromise Storage.....	6
Providing High Performance And Capacity a Low Cost.....	6
Eliminating the Tradeoff Between Performance and Capacity.....	7
Making Backup and Recovery Fast and Affordable.....	9
Simplifying Management Complexity	11
It's Time for No-Compromise Storage.....	13

Modern Storage Challenges

Managing today's exploding data growth is more than yesterday's storage infrastructures can handle. End users expect to have speedy access to their data—anywhere, any time, using any device. Technologies such as virtualization put heavy demands on storage in terms of performance, capacity, and data protection. Traditional storage technologies are unable to cope effectively with these demands, making primary storage and backup and recovery processes more costly and challenging to manage. The chances are good that you are dealing with one or more of the following storage challenges.

Challenge: Getting both performance and capacity

The limitations of traditional storage force the choice between performance and capacity. High-performance enterprise disks are expensive. High-capacity drives are less costly, but they are not fast enough by themselves to support most primary storage applications. Hybrid solutions that use a tiering model to blend tiers of storage are unable to move the right data to the right place at the right time, and cannot respond effectively to performance peaks.

Challenge: Improving business continuity and data availability

We live in an always-on world. Everyone expects data to be available any time, anywhere, on any device. Gone are the days when a company could shut down systems overnight or over the weekend for system upgrades or to perform a backup or restore.

Despite advances like disk backup and deduplication, traditional copy-based backup solutions consume a lot of computing resources and bandwidth. Meeting backup windows continues to be a struggle. Data recovery is slow and painful. Deduplication alone does not fully address these challenges. Traditional replication-based disaster recovery (DR) is bandwidth-intensive and expensive. As a result, it is used only for the most critical applications.

Challenge: Simplifying primary, backup, and disaster recovery storage management

Deploying, managing, upgrading, and supporting traditional storage can be an all-consuming effort that requires specialized training and expertise. In virtualized environments, storage management is even more complex.

At the heart of these challenges is the simple truth that traditional storage architectures are unable to keep up with today's storage demands. Modernizing these architectures to provide the performance, capacity, data protection, and manageability you need demands financial resources, time, and staff you may not have.

Leveraging Innovation

Every so often a technology comes along that revolutionizes an entire industry. Recent advances in flash memory promise to transform enterprise storage.

While there have been substantial improvements in disk density, CPU performance, and network bandwidth over the past decade, disk drive access time—a primary measure of performance—has improved the least. Poor gains in disk drive access time mean that high RPM drives are less able to keep up with demanding applications.

Table 1: Improvements in storage technology

Component	2000 – 2010 Increase
CPU speed	15x
DRAM speed	12x
Bus speed	30x
Network speed	100x
HDD capacity	16x
HDD random IOPS	1.2x
SSD random IOPS	100x

Until recently, the standard approach for delivering Input/output Operations per Second (IOPS) has been to deploy as many high-RPM drives as necessary to achieve the required performance. Increasingly, flash solid-state disk (SSD) drives are being used to deliver 50-100x better I/O performance than the fastest disk drive. Solid-state storage is becoming mainstream. However, adding SSD storage to existing arrays isn't simply a matter of replacing traditional hard disks with solid state drives. Strategies that optimize performance and reliability for HDDs are wasteful when applied to SSDs.

Table 2: Characteristics of SSD and HDD

Component	Flash (SSD)	Disk (HDD)
Random IOPS/\$	30x	1x
Sequential IO/\$	1x	2x
Capacity/\$	1x	20x
Reliability experience	Early	Mature

The ideal solution would combine the best of both worlds: use flash for random read performance needs and disk for random write, sequential throughput, and capacity requirements.

There are two ways technology vendors can incorporate disruptive technologies. One approach is to bolt them onto legacy architectures. This strategy typically fails to leverage the complete value of the technology and results in complex solutions.

The second way is to start from the ground up. Only solutions designed to incorporate disruptive technologies from scratch can cost-effectively deliver the maximum benefit of all technologies in a way that is simple to manage.

A Fresh Approach to Storage

History shows that true innovation rarely comes from established players. It is more likely to be introduced by newer players in the market who are not burdened by existing product architectures or shackled by the legacy needs of their largest customers. Newer companies can take a forward-thinking approach to the market's needs and emerging technologies.

Nimble Storage was founded by a team of industry veterans who developed groundbreaking primary storage and deduplication architectures at market-leading storage companies. Their experience gives Nimble a unique perspective to solving the industry's most pressing storage issues, starting from the ground up.

Before developing its new architecture, the Nimble team engaged with enterprise IT managers to build a deep understanding of their data storage and backup challenges. They asked IT managers what their "ideal" storage environment would look like. Invariably, the ideal solution was far less complex, much more powerful, and available for a fraction of the cost of today's solutions. In other words, these IT managers were looking for no-compromise storage.

All Nimble Storage solutions are built on a storage architecture designed to leverage a unique combination of innovative technologies to overcome storage challenges. Specifically, Nimble combines flash technology, high capacity disks, and multi-core, multi-threaded computing to deliver the industry's most efficient storage systems. In addition, Nimble is committed to the total customer experience over the lifecycle of its products. It starts with making systems easy to purchase, set up, and use. Proactive wellness and non-disruptive upgrades are features in all systems. Live heartbeats and secure tunnels connect each system to Nimble's world-class support organization that is standing by 24 x 7 to help keep data available and systems running efficiently.

The Technology Behind No-Compromise Storage

Nimble's breakthrough CASL™ architecture is designed to leverage flash memory in combination with high-capacity, low-RPM drives to deliver high performance at a cost-effective price. Just as server virtualization allows multiple workloads to be consolidated on the same server, Nimble's no-compromise approach allows organizations to consolidate multiple storage workloads on the same array. Nimble converges storage, backup, and disaster recovery into a single solution. WAN-efficient replication enables backup to a remote Nimble array, giving enterprises of almost any size a comprehensive, cost-effective disaster recovery solution.

To optimally leverage flash SSDs and low-cost multi-TB hard disk drives in the same system requires an evolved, hybrid architecture. Most storage architectures in use today were created over a decade ago. Even the most efficient of these systems has inherent limitations due to the file layout:

- Inability to adapt the data layout for new storage options (such as a mix of SSDs and low RPM disks), thereby limiting achievable performance and cost efficiencies
- Inability to move data rapidly to the right storage medium quickly, limiting responsiveness to rapid workload changes
- Inability to optimally vary storage block sizes based on application type

Providing High Performance and Capacity at a Low Cost

Cache Accelerated Sequential Layout (CASL) attacks these limitations head on to provide high performance and capacity at a low cost. A converged primary and backup storage solution must be optimized for performance as well as capacity. However, these optimizations are often at odds with each other. Combining them on an existing architecture can produce less than optimal results and may even be counter-productive.

To design its storage architecture, Nimble carefully developed a set of performance and capacity optimizations that are cost-effective *and* work well together. This architecture is called CASL, for *cache-accelerated sequential layout*. Sequential layout refers to writing data to disk sequentially, even when the writes (from the perspective of the application) are logically targeted to random addresses.

Cache-accelerated refers to leveraging a large cache on top of data stored on disk. Traditional DRAM storage caches are typically relatively small—providing less than 0.1% of the storage capacity of the disk subsystem. Such a small cache is not large enough to cache even the metadata (which typically requires more than 0.5% of the storage capacity), let alone cache the working set of application data. With CASL, cache is a combination of DRAM and a large amount of flash that is tightly integrated with storage.




Figure 1: CASL is optimized for performance, capacity, and data protection.

Eliminating the Tradeoff Between Performance and Capacity

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.

Lower Flash Usage to Reduce Costs.

Flash SSDs are expensive. By minimizing the amount of flash needed to accelerate performance, Nimble makes flash storage more affordable. To optimize the amount of data in flash, Nimble Storage starts with the block size of the file layout. When placing indexed or frequently read (“hot”) data in flash, Nimble tailors the logical block size to match the application. CASL allows very granular data placement—as small as 4KB blocks. For even more efficiency, all blocks are pre-compressed before being written to flash, and only one copy of data blocks is cached if they are shared between clones.

Prioritizing the hottest data into the flash cache also improves efficiency. Nimble systems accomplish this through real-time data placement. The system decides whether or not to place data on flash with every read and write.

Use Commodity MLC Flash SSDs.

Today’s SSDs degrade when burdened with continual patterns of random writes. When SSDs receive random writes, the write activity within the SSD is greater than the actual number of writes. This write amplification dramatically increases the number of write cycles that the SSD must support. Multi-level cell (MLC) flash is typically not suitable for traditional storage systems because it can only endure 5,000 to 10,000 write cycles. Instead, traditional systems must use single-level cell (SLC) SSDs and will soon begin using enterprise multi-level cell (eMLC) SSDs. SLC and eMLC technologies can endure up to 100,000 write cycles, but cost 4 to 6 times more than traditional MLC flash.

Nimble Storage approaches the problem of write amplification differently. The CASL file system is optimized to aggregate a large number of random writes into sequential I/O stripes. It only writes to flash in multiples of full-erase block width sizes. As a result, write amplification is minimized, allowing the use of lower-cost MLC SSDs.

Eliminate RAID Overhead Costs for SSDs.

Traditional architectures that use flash SSD as cache for Tier 0 performance typically write to flash before writing that data to disk. This approach increases the cost of using flash as a result of these penalties:

- Shorter SSD life
- 50% more capacity required for RAID protection
- Increased compute overhead for data movement

Flash is not a tier with Nimble. Because all data written to flash is first written to disk, there is no need for RAID on the SSDs. The only consequence of a lost SSD is a graceful reduction in flash cache acceleration capacity until the SSD can be hot-swapped—there is no data loss, no downtime, and no disruption to applications.

Increase Performance of High-Capacity HDDs

Nimble Storage dramatically increases the performance of high-capacity, low-RPM HDDs. This is a critical difference between the Nimble approach and other flash-as-cache approaches. High-capacity HDDs offer the lowest cost per GB but that cost comes at a price: random I/O performance is poor. On the other hand, they perform fairly well at sequential I/O. Nimble Storage achieves random IOPS on low-RPM HDDs that far exceed the performance of higher RPM disks by aggregating over a thousand random writes into a single sequential write operation. Coalescing random writes into a single RAID stripe enables CASL to improve write performance up to 100x compared to systems that use a fixed layout. By minimizing writes to disk, Nimble provides performance without relying on disk drive speed and number of spindles.

Reduce Capacity Requirements

Storage architectures that are based on fixed-sized blocks can only compress data that is offline or rarely accessed. The use of compression is also limited because data blocks compress at different rates, turning fixed-size blocks into variable-sized blocks after compression. The read-modify-write overhead associated with writing to compressed data is entirely eliminated with Nimble's CASL architecture.

“We achieved approximately 15 times the performance of our previous systems, and substantially greater performance than other systems we had considered. And we're getting data compression of 40 percent across the board.”

ERIC MYERS
SAN DIEGO CONVENTION
CENTER

Nimble is the only primary storage array that natively supports variable-size compressed blocks resulting in capacity savings of 30% to 75%. Using CASL, Nimble compresses all data in real time with no added latency.

Making Backup and Recovery Fast and Affordable

Traditional approaches to backup and recovery are stressed to the limit by unbridled data growth and the complexity created by virtualized desktop and server environments. Meeting backup windows is difficult. Data recovery is time consuming.

An additional challenge comes from infrequent recovery point objectives (RPOs). While deduplication allows organizations to retain the backups they need, the performance overhead of copying these large backup data sets means they often can't afford frequent backups. The result? Recovery points are few and far between. A typical RPO is one day. And recovery time objectives (RTOs) suffer as well. Because restores involve bulky copies and data format transformations, restores can take hours to reconstitute data.

When it comes to DR, business is always on and has little tolerance for programmed system downtime and unplanned outages. A traditional DR strategy relies on mirroring or replicating primary data to storage systems located in a different location. Recovery isn't instant and involves complex run-book processes. An additional challenge: the expense of duplicate systems puts disaster recovery out of reach for many business-critical applications.

Nimble's highly efficient snapshot and replication technologies eliminate many challenges associated with traditional backup and DR.

Efficient Snapshots and Replication for Extended Retention

Nimble's primary storage device can capture application consistent, nearly instant snapshots on a predefined schedule (every few minutes, hourly, or daily) without affecting application performance. Nimble snapshots achieve very high savings compared with traditional backup deduplication by only capturing new, compressed blocks and storing them on low cost disks. This approach is so efficient that most Nimble customers routinely retain many weeks of snapshots.

All snapshots or a subset can be replicated using very efficient replication to an offsite DR array. When needed, the data can be restored from snapshots within minutes. Applications can run directly off the backup/DR copies without any format conversion. With Nimble, there are no dedicated disk-based backup systems or backup windows to manage.

“Do you know how long it takes to back up a 1.5TB file server with over 30 million files using traditional software? It's real ugly and it takes more than 24 hours. During our DR test with Nimble we were able to restore our 1.5TB file server in 7 minutes.”

DARREN HOFFMAN
FIRST CHOICE
HEALTH NETWORK

Frequent Recovery Points (RPO)

Nimble's redirect-on-write snapshots consume 50-100x less space than full copies and have no performance penalty. Using variable logical blocks, Nimble matches the block size to specific applications. For example, just a 32k block size is required for Exchange 2010. Snapshots are also compressed, allowing thousands of snapshots to be retained. Compared to traditional systems Nimble offers significantly higher snapshot efficiency.

With Nimble, businesses can afford frequent recovery points and can store snapshot backups for as long as 120 days—far longer than with traditional snapshot architectures. Because 90% of restores occur within 30 days of a disaster, Nimble's backup/DR process simplifies backup and recovery for typical usage scenarios.

Shorter Recovery Time Objectives (RTO)

The goal of DR is quick application recovery. All Nimble systems are fully featured, whether they are used for primary or secondary storage. This makes the failover/failback process straightforward; simply mount and present the data from the secondary storage to the application. The time it takes to get the non-functional system back online, or RTO, is dramatically reduced.

Nimble also provides integration with Microsoft VSS (Volume Shadow-Copy Service), to provide granular backup and data recovery for all Microsoft applications. Integration with the ESX stack ensures VMware consistent snapshots.

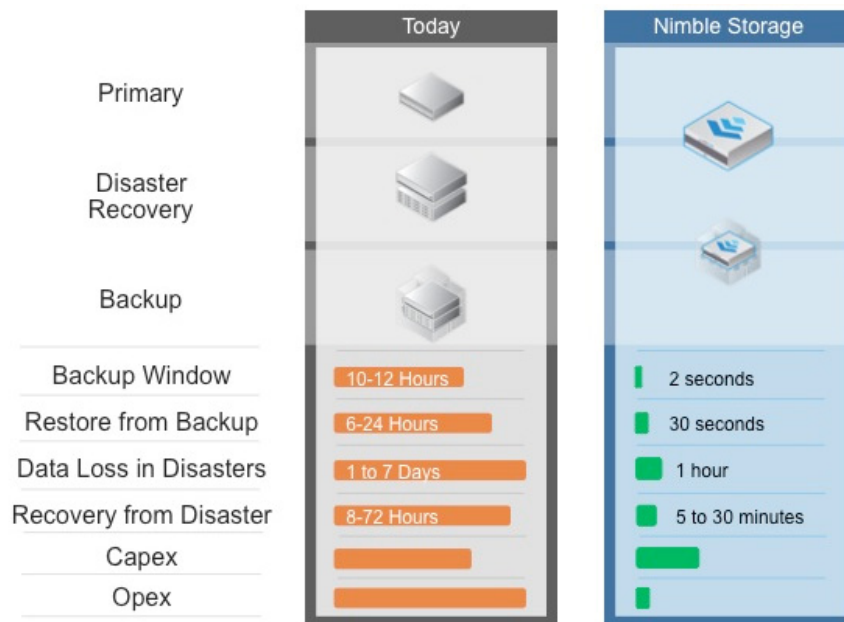


Figure 2: The Nimble advantage for backup and disaster recovery.

Simplifying Management Complexity

Traditional Storage Area Networks (SANs) are feature rich, but maintaining and managing them is not for the faint of heart.

Challenges include:

- System configuration is complex
- Provisioning processes are time-consuming
- Capacity planning is difficult
- It is impossible to optimize workloads continuously
- System upgrades often involve unplanned downtime
- Troubleshooting is painful, time-consuming, expensive, and requires onsite access

Reducing costs and reining in storage complexity are fundamental to Nimble Storage. Nimble's lifecycle approach to simplifying storage makes these systems simple to purchase, setup, use, and can easily be maintained by an IT generalist.

“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 infrastructure 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



Push-button deployment

- All inclusive features
- Integrated solutions
- Pre-tuned for applications

Predictive management

- Application and workload analysis
- Global peer insights
- Customized guidance

Proactive wellness

- Real-time health analysis
- Preemptive failure avoidance
- Rapid remote resolution

Figure 3: Nimble's lifecycle approach to simplified management.

Easy to Purchase and Set Up

Nimble systems contain all the software needed for storage, backup, and disaster recovery. There are no additional software licenses—all functionality is included with the system. The graphical user interface is user-friendly and simple to use. Storage, backup, and disaster recovery can each be configured in just three simple steps.

Built-in application profiles tune the system for specific workloads by setting logical block sizes and caching hints for critical workloads. Data protection policies are also predefined. This eliminates the manual tuning required with most other storage systems.

Deep Application Integration

Through deep application integration, Nimble systems simplify storage provisioning and backup processes. 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, reducing the reliance on administrator-driven restores.

Nimble Storage arrays are designed to integrate with:

- Microsoft applications: MS Exchange, MS SharePoint, and MS SQL Server
- Virtualization applications: VMware and Microsoft Hyper-V
- Oracle applications

The VMware vCenter plug-in allows the virtual datacenter administrator to perform storage operations from the convenience of the vCenter console. With Nimble’s plug-in, administrators can provision, clone, replicate, backup, monitor, and recover VMware datastores from the vCenter console. In addition, Nimble supports zero-copy clones, which save both time and capacity in VM environments.

“With Nimble it takes, literally, a few clicks of the mouse to establish a backup schedule and meet our data retention policy. Backups today take under a minute versus the hours or days it took in the past with the legacy hardware and software we had in place.”

MIKE BRESTER
MULVANNY G2

Proactive Wellness

Nimble’s built-in proactive wellness capabilities keep systems running in top condition. Frequent heartbeats track and monitor system health to identify failure conditions before they occur. The comprehensive command console simplifies troubleshooting and ongoing administration. When support is needed, the Nimble support infrastructure allows for easy data collection remotely, thereby greatly reducing triage and resolution time, without requiring onsite access.

Non-Disruptive Software Upgrades

Nimble’s rolling software upgrades mean that business doesn’t come to a standstill every time there’s a need to upgrade storage firmware or software. Software upgrades are non-disruptive and involve no downtime.

It's Time for No-Compromise Storage

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 Storage 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.

Nimble Storage solutions are available through a global network of world-class channel partners. For an in-person briefing, contact Nimble Storage at sales@nimblestorage.com or call 877-3NIMBLE (877-364-6253).

For more information, visit www.nimblestorage.com.



Nimble Storage, Inc.

2740 Zanker Rd., San Jose, CA 95134

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

© 2012 Nimble Storage 4.5.12