

WHITE PAPER

W I N T E R C O R P O R A T I O N

SCALABLE NETWORKED STORAGE


*Convergence of SAN and NAS
with HighRoad*

Specialists in the World's Largest Databases



www.wintercorp.com

SPONSORED
RESEARCH
PROGRAM



Explosive Data Growth: the “Content Big Bang”

A recent Study conducted by the School of Information Management and Systems (SIMS) at the University of California, Berkeley finds that explosive growth of information is only beginning. According to the study, humankind will generate more original information over the next three years than was created in the previous 300,000 years combined. The study estimated that in 2001, 6 exabytes (EB)¹ of new information will be generated. According to the study, that number is expected to double every year. This number only counts the initial copy and does not include replicas.

Businesses, Service Providers, and Governments will need to manage tremendous amounts of information. The difficulty will be in managing this information effectively. This is no easy task. Our ability to store and communicate information is quickly outpacing our ability to search, retrieve and present it. Information access, management and mobility may turn out to be three of the major challenges of the new century.

Business Implications

The business implications are profound, especially for organizations for which the focus is increasingly on large or rich content, such as Internet, telesales, video streaming, medical imaging, CAD/CAM and numerous other applications. In these environments, very large files need to be shared by a multitude of users. This is typically what Network Attached Storage (NAS) has handled over the IP network. For the vast majority of traditional file sharing applications, NAS has been and will continue to be the preferred solution.

But because of today's business requirements, the increasing richness of data and the coming explosion of information, some NAS-based applications are reaching their bandwidth limits. Even the expected improvements in Ethernet bandwidth will not solve the dilemma. Frequent updates and much larger user populations strain NAS capabilities even further. Storage Area Networks (SANs), on the other hand, provide the performance and scalability that the rich content explosion requires. SANs, however, lack file sharing capabilities.

This white paper explores a new approach that combines the file sharing of NAS and the performance and scalability of SAN and describes an EMC product, HighRoad, that implements this approach.

Actual Cases: Terabytes of Data, Millions of Files, Millions of Users

The Internet and other changes in technology have drastically increased the scale on which files are managed on behalf of multiple users and shared across networks of computers.

¹ An exabyte is a billion megabytes. This is 1,000,000,000,000,000 bytes, or 10^{15} bytes. If stored on a stack of personal computer diskettes, the stack would be more than 2000 miles high.

Electronic retailing. For example, an electronic retailer which presents its products to the public via 24-hour video programming has an online repository of thousands of files. Each file contains an hour of video and occupies 400 megabytes of storage, for a total of about 3 TB/year. Users within the company access the video files – in conjunction with other data – as they analyze which offers were successful and why: did we sell more units of product x when it was presented by a woman or by a man? When the setting was indoors or outdoors?

Web service with dynamic content. An Internet service provider has a large file frequently accessed by its millions of users. The file contains information that is updated throughout the day. The file is over 100GB. The company operates over 100 web servers which must have continual access to the content of the file so that all online users can benefit from it.

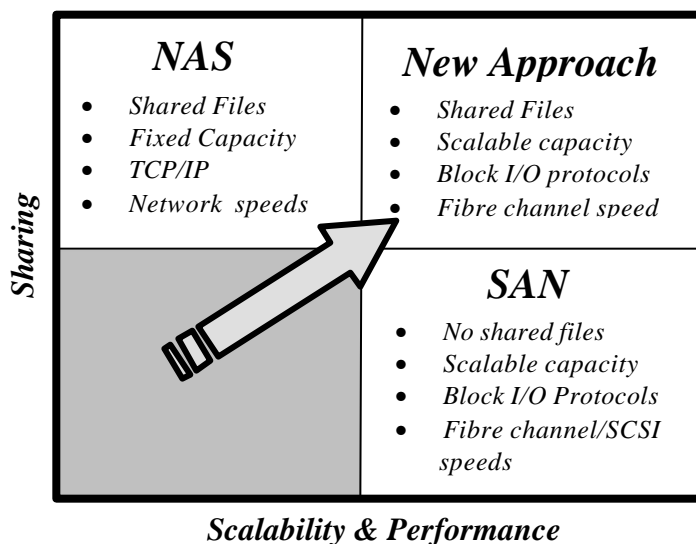


Figure 1: Convergence Requirements of Emerging Applications

The cumulative effect of all the requests from all of the web servers is an extraordinarily high workload on the shared file repository.

Impact. The combined impact of these increases in scale is remarkable. The total volume of data to be moved from storage to the user community is thousands to millions of times larger than comparable volumes of a few years ago.

SAN is best for DB and Other Transaction-Oriented Applications

A Storage Area Network (SAN) is the technology of choice for high performance applications that manage their own, typically large, data spaces. The SAN architecture provides flexibility in physical placement of data and fast direct access to any specific record or small set of records. SAN is often used for high-performance database access and transaction processing. Though SAN delivers high performance and scalability for shared storage resources, it does not provide for sharing of files among hosts.

NAS is best for Applications Requiring Shared Access to Files

NAS, Network Attached Storage, provides for the sharing of individual files among multiple, often heterogeneous, hosts on a network. NAS includes concurrency controls (typically locking) to protect the integrity of the files in the presence of reading and writing by multiple users. However, as requirements continue to grow, NAS

architectures are increasingly vulnerable to performance and scalability problems due to network overhead and bottlenecks. The requirements of today's networks – with their dynamic, content rich files, large user populations and large numbers of files – are pushing the bandwidth limits of NAS products.

The business applications described in the examples above are driven by a need to share access to the same large files. What is needed now is an approach that provides the file sharing function of NAS combined with the performance and scalability of SAN, as shown in Figure 1.

EMC has introduced a software product, HighRoad, which aims to provide exactly this convergence of SAN and NAS. HighRoad combines the file sharing, locking and simplicity of NAS with the high performance and speed of SANs, creating one unified storage network. It facilitates direct channel access from multiple clients to single instances of data.

HighRoad: NAS Network Request, SAN Channel Delivery

NAS is the preferred storage solution for shared access by large numbers of heterogeneous clients to a single file. But, many file-sharing applications demand greater performance than standard NAS file servers can deliver.

Figure 2: HighRoad Features

- Multi-Path File System: over channels and networks
- Multi-Protocol File System: standard file access protocols NFS & CIFS
- Multi-Ported File System: direct, concurrent, NAS and SAN access
- Data access approaching local disk speed for large file sequential access
- Heterogeneous clients share a single file system
- Highly available and scalable
- Transparent to applications
- Dynamic routing bridges the NAS and SAN domains
- Full file system functionality.

Source: EMC Corp.

HighRoad addresses this problem with new function currently available on EMC's Celerra file server. HighRoad increases data access performance by using separate mechanisms for control actions and data delivery. The user/client requests

information from the Celerra file server. HighRoad replies by delivering information directly to the application over the high-performance Fibre Channel SAN, as shown in Figure 3.

This topology optimizes both NAS and SAN, accelerating delivery of information to the user while placing the technical details behind the scenes. As a result, users benefit from shared data access at channel speeds, reduced network traffic and improved performance for applications that require many host servers to share access to very large data files in sequential access environments.

Traditional NAS environments provide for data sharing at the file level by channeling all data access and delivery through a file server, which can limit both performance and scalability in large file and/or rich content environments. Traditional SANs allow direct access to data in a common storage pool, but offer no provision for data sharing.

HighRoad: The New Approach to Networked Storage

The new approach implemented by HighRoad provides the opportunity to unify the storage infrastructure while satisfying growing requirements for performance and scalability in network file sharing.

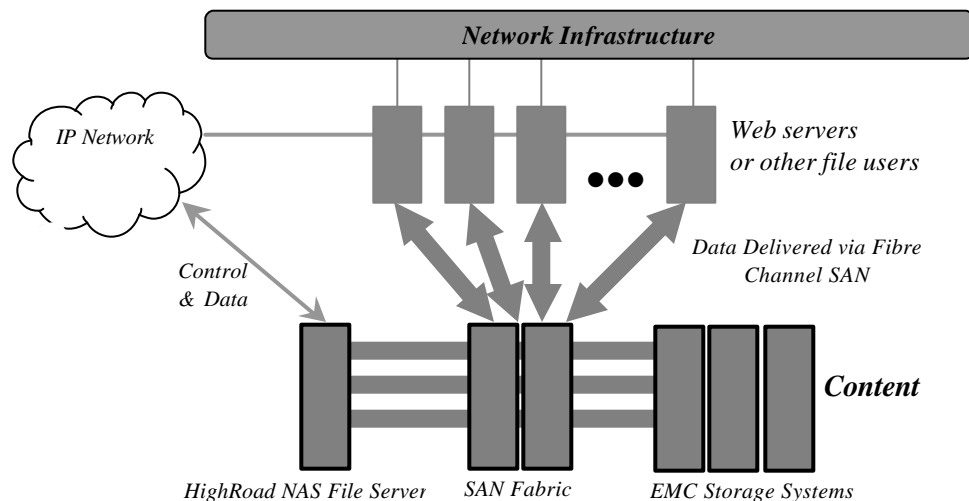


Figure 3: The New Approach to Networked Storage

Customers using HighRoad in production environments have measured large performance improvements, when compared to NAS. Others have found that to be the only way they could deliver certain new, mission-critical applications with the required availability and bandwidth. Some customers have stated that tasks that would previously have taken days with traditional techniques can now be completed on-line, a critical difference in the fast-paced knowledge economy.

Performance and Scalability Measurements. In performance tests at EMC, the HighRoad solution was compared to a NAS solution as shown in Figure 4. In all cases, client applications read a one gigabyte file. In the test represented by the leftmost bar in the bar chart, Celerra was used in a NAS configuration and achieved a throughput of about 20 megabytes/second. In the test represented by the middle bar, Celerra was used in the HighRoad configuration and data was delivered to the clients by means of a single Fibre Channel path, resulting in a threefold increase in throughput to about 60 megabytes/second. In the test represented by the right most bar, a second Fibre Channel path was added to the HighRoad configuration, resulting in a *further* doubling of throughput to 120 megabytes/second. Yet additional Fibre Channels paths can be added, and ought to result in yet higher throughput, although at some point other limits may come into play.

Thus, this particular test demonstrated a 3-6 times increase in throughput, when compared to using the same file server in a traditional NAS configuration. Similar tests

were conducted for both reading and writing files with record sizes ranging from 32 KB to 1 MB, with similar results.

Each of the companies mentioned in the cases above has adopted HighRoad to address its requirements. Here is what happened.

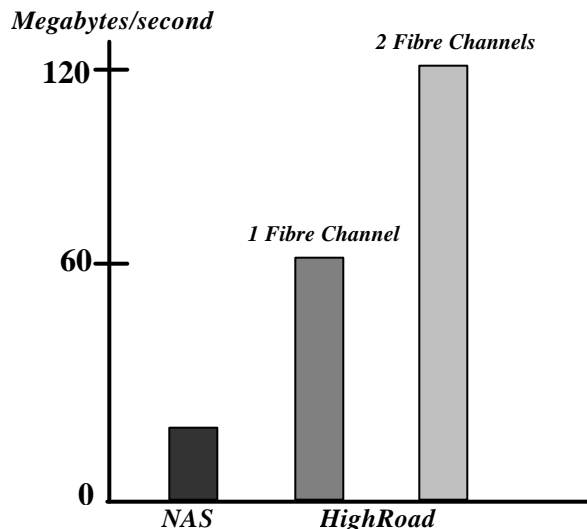


Figure 4: HighRoad throughput is 3-6 times NAS throughput in one test

Electronic Retailing: Large File Use Multiplied, New Revenue Stream.

This company, an early adopter of HighRoad, deployed it initially for internal business use. This enabled multiple business executives to view the files simultaneously whether they were in the same location or not. The company saw immediate benefit in multiplying the use of the same files by a great number of individuals, saving the company time and staffing costs. In the past, the programming was stored only on video tape.

Retrieval from the tape library was manual. Only one user at a time could examine a given tape. These files were not previously stored on

disk because NAS creates too much of a performance bottleneck and SAN does not allow file sharing.

As a result of success in its internal use, the company now plans to create a new revenue source based on the ability to rapidly serve large video streaming files to its Internet customers. Large video files will be placed on the web site for users to view and purchase products described in the videos.

Web Service with Dynamic Content: Serving Content to Millions of Users. The HighRoad solution enables the company to present the data to all the servers and deliver the 100GB file to each edge server so that, together, they can service the community of millions of users. The company saw major benefits in the solution. The system is scalable because servers all share the same information; there is no duplication of data across the network which allows better performance, fewer resources to manage the system. In addition, member services information is always served up-to-date and at channel speed. The company estimates that this solution will cut their staffing costs dramatically and improve performance by 3-6 times.

In summary, HighRoad is a solution based on new technology that provides high bandwidth and highly available, shared access to data among many clients. It provides direct, concurrent, and shared access to data allowing heterogeneous UNIX and Windows NT clients/hosts access via direct channels (SCSI or Fibre Channel). It increases data

throughput by using separate mechanisms for control data and file data. Control (metadata) operations utilize the network. Data is transferred over the Fibre Channel SAN.

Figure 5: HighRoad Typical Customer Environments

Customers with the following requirements:

- Applications requiring multiple host servers
- Shared access to large data files
- Primarily sequential storage access within files

Typical applications:

- Web hosting
- Media files
- Streaming video
- Advertising
- Image processing
- Satellite imaging
- Geologic/Seismic data
- Medical image processing
- Design applications
- CAD/CAM
- Software simulations/emulation
- Other applications
- Usenet news group
- Flat file databases
- Pre-press publishing
- Financial/economic modeling

Source: EMC Corp.

Conclusion

HighRoad delivers a powerful new capability: NAS functionality delivered with SAN performance for applicable environments. It also brings file sharing capabilities to SAN. It provides a capability for very large numbers of users to readily share large files with dynamic content. This new EMC solution will enable users to share large files or rich content on a large scale with large user communities, whether in e-commerce or enterprise applications.

More efficient file sharing means delivering large or rich content files faster, developing new applications sooner and getting them to market quicker. This can translate into seizing new market opportunities ahead of the competition, and increasing customer satisfaction, retention, and loyalty.

Recommendation. Winter believes that a large percentage of enterprises face extremely rapid growth in the demands for file sharing across networks of computers, both inside and outside of the organization. As a consequence, Winter recommends that companies find the most manageable, scalable, and unified SAN/NAS networked storage solution to satisfy enterprise requirements. Customers are therefore advised to pursue storage strategies that address all of these requirements in an integrated architecture.

*A leading center of expertise in very large databases,
Winter Corporation provides services in
research, consulting and education.*

*We help users and vendors understand their opportunities;
select their database and data warehouse platforms;
define and measure the value of their strategies, architectures and products;
plan, architect and design their implementations;
and manage their scalability, performance and availability issues.*

Our focus is databases near, at and beyond the VLDB frontier.



WINTER CORPORATION

411 WAVERLEY OAKS ROAD, SUITE 327
WALTHAM, MA 02452
617-695-1800

visit us at www.wintercorp.com