



**NAS systems scale out to meet
growing storage demand.**

Contents

- 2 Introduction**
- 2 Understanding the traditional NAS role**
- 3 Gaining NAS benefits**
- 4 NAS shortcomings in enterprise settings**
- 5 Scale-out NAS**
- 8 Applications suited to scale-out NAS systems**
- 10 A logical progression—scale-out NAS as a service**
- 12 Meeting strategic storage demands**

Introduction

The efficient management of data is an ongoing struggle between access and scalability. Providing access to file-level data—the files associated with individual documents, multimedia files, database and other applications—becomes more difficult as more users are provided access and more data is stored. Achieving the scalability needed to respond to data volume growth also typically results in higher hardware and software costs, and greater management challenges.

Network attached storage (NAS) solutions provide simplicity, manageability and access. But until now they've lacked a single capability that has kept them from serving a role beyond the departmental level: scalability.

Emerging NAS systems, and managed storage services that are powered by these technologies, offer compelling business value to organizations because of their broad spectrum of applicability, competitive price/performance and scalability to meet growing and evolving data demands.

Understanding the traditional NAS role

The two approaches to shared data storage, storage area networks (SANs) and NAS have much in common, but they function at two different protocol layers. SANs communicate with applications using lower-level protocols, which offer fast performance, but little in the way of intelligent manageability. A SAN is a storage-class system that identifies data by block number and utilizes Fibre Channel protocol or InfiniBand technology to access the raw data blocks.

Highlights

NAS is a multilayer, data-class system that identifies information by file type, such as text, audio or video. As a full, multilayer solution (in fact, SAN can be one of its layers), NAS manages data security, user authentication and file locking. It utilizes standard protocols such as Network File System (NFS). Because NAS operates at these higher protocols for data transfer, it can be used to more intelligently store and access data. But there is a performance latency trade-off for gaining that intelligence. On a small scale, NAS also offers improved manageability. But again, performance and management suffer as the volume of file data grows.

NAS systems have traditionally been used for departmental access to file data.

Based on these profiles, SAN systems have typically been deployed for enterprise wide applications involving large volumes of data, whereas NAS systems have been deployed for departmental access to file data.

Gaining NAS benefits

NAS has several advantages that satisfy the needs of both end users of data and the IT administrators charged with delivering and managing data resources. The overall simplicity of saving, accessing, moving and archiving files between the device and its clients makes NAS ideal for file sharing tasks in support of workflows. Users basically can save to a centralized location and find files when and where they need them. Consolidation to an NAS device of all files associated with a particular application, department or group of users also helps facilitate management, because the file system is all in one place. Routine management tasks, such as performing backups, moving files within the file system to improve performance and sending files to archive media, are faster than in the distributed environment of an SAN.

NAS systems offer management advantages and fast access to data because of their centralized location.

Highlights

Intelligent routing of data at the operating system level can result in performance gains.

The problem with traditional NAS systems is their lack of scalability.

Applications that depend on NAS for storage can also perform better than their SAN counterparts because all of the file management tasks are handled on the NAS operating system, not on that of the application server. Not having to manage storage and file retrieval tasks allows the application server to dedicate more resources to processing application requests, which can result in performance gains.

NAS shortcomings in enterprise settings

IT departments needing to meet immediate demands for storage have typically responded by simply adding NAS devices to their networks: install a box, assign an IP address and start saving files. That approach is fine on a departmental level for similar types of files and for relatively small files. But the exponential growth of data that organizations of all sizes are experiencing means that the efficiency and manageability advantages of traditional NAS can be wiped out by an inability to scale. A file server that supports a group of 50 users in an accounting department, for example, becomes inadequate when a merger adds another 50 employees to the department. It is possible to add bigger file servers, but for organizations with tens of thousands of users, traditional NAS eventually becomes an ineffective solution.

Each NAS device has only so much performance capability, limited storage capacity and a finite number of file systems available to be divided in support of desired applications and users.

The typical NAS solution is limited to two nodes, so when capacities are reached, the logical solution is to add more devices. But this can result in disruptions to operations as well as place an added burden on IT staffs that must configure and deploy the device, reallocate file systems and create access permissions. Users must also reestablish shortcuts from their clients, potentially

Highlights

The scale-up approach increases complexity and cost and results in a single point of failure.

Scale-out NAS systems utilize scaling architectures that are common to Web serving—virtualized pools of inexpensive resources.

impacting productivity. Adding NAS devices can actually increase complexity and costs, and it can isolate data from broader access. This is not the best scenario for a growing organization.

Another alternative is to scale up, increasing the input/output size and storage capacity of the NAS device as data needs grow. But this has several disadvantages also. One is that you have significant cost for replacing a device with only an incremental increase in capacity (which can eventually be exceeded). Another is that you still have a single point of failure.

Scale-out NAS

The ideal solution is a system that combines the best in file system access and manageability yet can scale to enterprise needs. Fortunately, today's NAS solutions feature scalable architectures and enhanced management capabilities that allow them to offer the scalability demanded for traditional file sharing, enterprise file sharing and high-performance computing (HPC) applications.

These scale-out NAS systems utilize scaling architectures and technologies borrowed from server environments and apply them to the storage environment. The goal is to offer simplified access to data and the scalability to grow as dictated by the needs of individual users, applications, the IT department and the broader organization.

The scale-out NAS architecture is similar to that of Web environments. When visitors connect to a Web site, they aren't necessarily connecting to the same physical server to access the data each time. Inexpensive servers handle the traffic virtually. When the load increases, you don't deploy bigger servers; you deploy more servers and manage them virtually as one.

Highlights

Scale-out NAS represents a paradigm shift because it virtualizes storage devices, allowing greater scalability.

Clustering capabilities of scale-out NAS systems allow rapid addition of performance and total storage capacity.

Applying this scale-out design to NAS, which has historically been scaled up, represents a paradigm shift. The scale-up model is inherently limited by the size/capacity (and cost/reliability) of a single or failover pair of storage servers. A scale-out design is only limited by the software that combines the many inexpensive servers into one large virtual server.

These newer, more manageable NAS systems owe their increased scalability to technologies and practices borrowed from application servers, including:

- *Clustering*
- *Scalable file systems*
- *Storage pooling and provisioning.*

Clustering

Clustering is the use of multiple NAS device nodes to create a larger, logical storage system that can be maintained as a single entity. The advantages of clustering include enhanced storage system performance. Data requests and processing power are spread across a wider system, and the need to cache data saves and requests is reduced. Data is sent directly to the disks in a scale-out NAS system, improving the speed of requests and accommodating high volumes of requests.

Because the scale-out NAS system is managed as a single entity, when new nodes are added they are easily recognized by the management interface and can be put to work as needed. This ability to rapidly add storage capacity or improve application performance and resiliency contributes to a more efficient storage environment. The ability to add nodes easily and only when necessary is a significant improvement over traditional storage strategies that advocate overpurchasing, which leads to underutilization and higher total cost of ownership.

Highlights

A scale-out approach is inherently more available because it eliminates the single point of failure of traditional NAS.

IBM General Parallel File System is the software behind the new NAS capabilities.

Scalable file systems

The relatively limited size of the file systems has limited the use of NAS in enterprise settings.

In a scale-out design, a given NAS instance can be dynamically expanded (or contracted) at any time by simply adding inexpensive front-end servers to increase input/output performance or adding back-end disk (storage capacity) transparent to the end users, providing almost limitless scalability. Inherently, this dynamic architecture is also highly available. Should any server on the system fail or be removed, the remaining servers pick up the workload seamlessly.

This major shift is the result of two fundamental technologies:

- *Highly scalable shared file systems, which provide a single namespace for all files*
- *A method to virtualize the NFS service (or connecting protocol of choice), so that an end user sees one device rather than the individual servers that comprise the virtual NAS device*

IBM General Parallel File System™ (GPFS) is one of the most scalable commercial file systems available, installed in thousands of nodes and storing petabytes of data in discrete installations. Scale out file services enable the virtualization of the NFS service, transparently distributing service requests across multiple servers, which enables the creation of a scale-out NAS farm installation.

Highlights

IBM General Parallel File System allows different media types to be used in a single NAS system in support of information lifecycle management initiatives.

Scale-out NAS can be applied to almost any type of data, including large files, with the advantages of increased scale and improved performance.

Storage pooling and provisioning

Because of the limited number of nodes in a traditional NAS system, there has been no practical way to use different media on the same system to effectively employ information lifecycle management strategies. Scale-out NAS supported by GPFS the creation of pools of NAS that can span high-speed disk, low-speed disk and tape drives. It allows storage administrators to associate a particular file attribute—such as age of file, user ID, server ID, file type, directory or subdirectory—and then dynamically move the data around to the appropriate media based on association policies.

Creating pools of varied storage types allows increased control and manageability of data in support of information lifecycle management initiatives. Administrators can provision resources more dynamically in response to certain tasks, such as daily backups or batch processes done daily, quarterly or annually.

Applications suited to scale-out NAS systems

What kinds of data can you store on a scale-out NAS system? Few data types and applications are not suited for this approach. The initial NAS application of providing centralized storage for files, a place where users can easily find data, is still valid. Again the differences are scale and performance. Scale-out NAS can simply support more users, more and bigger files—and do so with improved performance. For personal productivity documents, spreadsheets, images and audio/video files, a scale-out NAS system can offer the performance and accessibility end users expect.

Highlights

Scale-out NAS supports advanced IT initiatives such as consolidation through virtualization and SOAs.

Industries with data volume and performance considerations are suited for scale-out NAS.

Scale-out NAS can also meet file requests from applications more efficiently than dispersed storage resources. It is particularly suited for high-volume, small-file-size content serving, such as delivering online services to consumers. Several high-profile online portals have started to utilize scale-out NAS to support their e-mail, multimedia storage and Web content services.

In an effort to more closely align IT services with business objectives, many organizations are undertaking advanced initiatives such as service-oriented architecture (SOA) and virtualization. Often the first step in these initiatives is the consolidation of dispersed resources to fewer, yet more scalable systems. Scale-out NAS supports consolidation by centralizing data to a single file system and streamlining management tasks associated with saving, moving and accessing files.

HPC applications require server processing power, but they also typically have high-volume, high-performance storage demands. Applications that typically analyze data or model results need a large amount of storage just to run and can quickly generate huge data volumes. Some industry applications suited to scale-out NAS include the following:

- *Oil and gas*
- *Geosciences*
- *Animation rendering*
- *Research*
- *Biotechnology research*

Highlights

Implementing a scale-out NAS system requires expertise.

A service approach can help organizations get up and running quickly.

IBM's scale out file services is an NAS solution for quickly implementing a scalable, global file space system.

A logical progression—scale-out NAS as a service

Fully realizing the strategic advantages and technical capabilities of scale-out NAS requires considerable expertise. Correctly designing and implementing a scale-out NAS system requires understanding and practical experience with high-end server environments and multiple operating system platforms, as well as storage design and management expertise that few organizations, even global enterprises, have access to today. And although a scale-out NAS system, once implemented, makes management of the storage environment more dynamic, it will require ongoing management expertise to maintain performance. If and when nodes need to be added to a scale-out NAS implementation, that will also require careful planning and execution to reduce the risk of impact on production environments.

For both IT management and strategic business reasons, some organizations may want to turn over the implementation of a scale-out NAS system to a capable service provider. IBM is leveraging several of its product and service capabilities to deliver a scalable file management solution that our clients can utilize to quickly implement a global file space system for storing and providing access to data.

A key component of the solution is the GPFS a high-performance shared-disk file system that can provide fast, reliable data access from all servers in a cluster. GPFS gives parallel applications simultaneous access to a set of files (or even a single file) from any node that has the file system mounted, while providing a high level of control over all file system operations.

Highlights

A graphical user interface included in the service offers a quick and clear view of NAS resources and performance.

The IBM solution offers scale-out capabilities of more than 500 nodes and capacity of up to one exabyte.

IBM's solution using scale out capabilities is made up of preconfigured IBM storage hardware devices and the GPFS, plus additional software for managing allocations in the file system. To simplify management for storage administrators, the solution also features a graphical user interface (GUI) that offers a quick and clear view of NAS resources and performance, enabling users to rapidly adjust to accommodate changing demands.

Although scale-out NAS is an evolving approach, most solutions in the market-place are still limited in scalability, typically only scaling up to eight nodes. Scale out file services from IBM have been tested to scale up to 500 nodes with the potential for more—making them feasible for larger organizations that have huge data volumes and want to consolidate on a single enterprise-class solution. IBM scale out file solutions are built on technology that offers significant performance capabilities and has been shown in real installations to deliver more than 130 gigabytes per second to a single file over IP connections. And in terms of total volume of data capable of being stored, scale out file services currently supports up to one exabyte—a billion gigabytes.

Available on client sites or hosted in IBM data centers, scale out file services also offers functional capabilities that enhance the value of the scale-out NAS system at its core. Examples of these include:

- *Support for information lifecycle management so that data is stored relative to its value*
- *Policy-based Quality of Service (QoS) services that offer storage and performance priorities to specific groups or applications*
- *Backup, restore and archiving for data security*
- *Snapshot and mirroring for high-availability systems requiring smaller backup windows*



- *Off-site disaster recovery that utilizes the global file system and enables recovered data to be sent to any number of IBM data centers*
- *Extensive monitoring and reporting down to user and project levels*
- *“Call home” hardware service agents to proactively maintain systems.*

Meeting strategic storage demands

Scale-out NAS systems can address scalability and performance challenges at a technical level, but ultimately they must support broader business goals.

Utilizing advanced architectures and consolidating to a single file system can help lower storage management requirements, improve utilization and potentially lower costs. Scale-out NAS systems can help improve the performance of applications and help deliver files to users more quickly and easily, supporting efforts to increase individual productivity and organizational responsiveness. These storage systems can also contribute to improved business resiliency and security, because their multinode design provides failover, especially when they are delivered as a managed service as part of a global file system.

Allocating the right amount of data storage to the right users at the right time will remain an ongoing challenge for organizations of all sizes. Scale-out NAS solutions can help alleviate some of those challenges, minimize the traditional performance trade-off and provide enhanced strategic advantage.

For more information

For more information about IBM storage and data services, visit:

ibm.com/services/storage

© Copyright IBM Corporation 2007

IBM Global Services
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
All Rights Reserved
09-07

IBM, the IBM logo and General Parallel File System are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Other company, product and service names may be trademarks or service marks of others.

References to IBM products and services do not imply that IBM intends to make them available in all countries in which IBM operates.