

Avoiding storage bottlenecks and increasing performance

*Is a scale-out NAS system the same as a scale-out file system?
Knowing the difference is critical.*



Highlights

- Choose the most impactful solution for your storage environment
- Easily scale your storage in both capacity and performance
- Enable applications to read from all storage nodes in parallel
- Avoid the typical bottlenecks associated with network attached storage
- Provide maximum data availability, integrity and security

Organizations are generating and collecting volumes of data and analyzing it to produce business insight. Different approaches to storing and analyzing the data have emerged, each one suggesting a different concept of where to place the data and where to run the analytics. Distributed systems such as Hadoop and Apache Spark now complement traditional data analytics applications, with both of these systems requiring rapid access to a large quantity of data from all sources and types.

Facing rapid growth in data and an insatiable need for faster time to results, organizations are making greater use of scale-out architectures to dynamically scale both performance and capacity to support future demand. Scale-out storage architectures are built by adding independent nodes that provide processing power along with capacity.

However, evaluating different scale-out architectures can be confusing. For example, there is a significant difference between accessing data over scale-out network attached storage (NAS) and using scale-out file systems. This paper explains that difference to help you choose an impactful solution for your storage environment and avoid data-related bottlenecks that can slow performance.

The difference between a file system and NAS

Everyone who uses a computer is familiar with file systems, even if they are not aware of it. The file system is the technology that stores and retrieves data, and provides the structure for locating data. In the case of files, that structure consists of directories and folders presented to the user. For objects, it is a unique name. A file system is responsible for data integrity and consistency, managing user access and writing data to disk or other media.

NAS allows remote computers to access the files on a remote server. Typically, the computers do this with an industry-standard protocol such as network file system (NFS) or Microsoft Windows server message block (SMB). Underneath the NAS protocol, the server can use any file system. NAS provides central sharing of data across heterogeneous systems, and NAS servers can deliver good performance if the network capacity and latency are sufficient. A NAS file system can be scaled, but only to certain limits.

Scale-out file systems such as **IBM® Spectrum Scale™** provide the unique option of using a local client that appears as a local file system, but is actually part of a more global, highly scalable file system that can be shared across heterogeneous systems. As a result, it combines the application simplicity of a local file system and the sharing benefit of NAS. Unlike NAS, however, the IBM Spectrum Scale client can directly access disks in parallel.

The IBM General Parallel File System (GPFS™) protocol at the heart of IBM Spectrum Scale can scale in both capacity and performance, avoiding the typical bottlenecks associated with NAS. IBM Spectrum Scale provides a single namespace across clients within different systems over the network.

A proven, high-performance data management solution, IBM Spectrum Scale is used extensively across multiple industries worldwide. IBM has invested in IBM Spectrum Scale enterprise features that make it reliable, easy to use and suitable for mission-critical storage of all types, ensuring maximum data availability, integrity and security. IBM Spectrum Scale is part of the **IBM Spectrum Storage™** family of software defined storage offerings designed to address today's operational storage challenges and business needs for maximum storage optimization.

Scale-out NAS and remote file systems

A scale-out NAS system is comprised of multiple NAS storage nodes with internal or external storage (Figure 1). Each NAS storage node represents the same file systems to user applications. Data is stored across all NAS storage nodes, and adding capacity can be achieved by adding more storage nodes. Data coming in through a NAS protocol such as NFS typically has to be converted to another file system protocol before the data blocks are stored on disk.

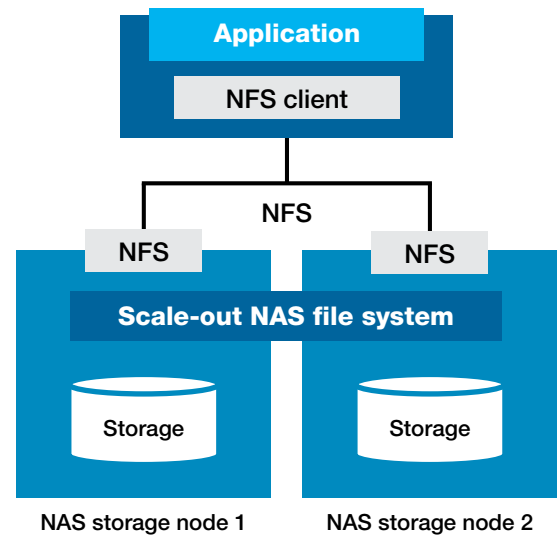


Figure 1. Scale-out NAS system architecture with NFS access.

A scale-out file system (Figure 2) has a similar architecture as a scale-out NAS system and is also comprised of multiple storage nodes with internal or external storage. But unlike a scale-out NAS system, access to the file system is granted through the IBM GPFS protocol and not through NAS protocols such as NFS. The IBM GPFS protocol is truly parallel and block-oriented. As a result, the scale-out file system avoids the possibility of an NFS bottleneck, which can occur with a scale-out NAS system. In addition, data written by the application does not have to be converted to another protocol, because it is directly stored on disk.




Figure 2. Scale-out file system architecture.

The NFS performance bottleneck

To explore the bottleneck issue, assume that an application accesses a scale-out NAS system through standard NFS version 3 or 4. File system access through NFS is available on all NAS storage nodes. However, for a single file streaming operation, the application is bound to a single NFS node of the scale-out NAS system (Figure 3). In other words, for a single file operation such as reading a large file for backup, there is a point-to-point connection between the NFS client and a single NAS storage node. That single NAS storage node is the source of the bottleneck, because it can easily become overloaded.

In contrast, when an application accesses a scale-out file system such as IBM Spectrum Scale, the file system is accessed directly through the file system client. With IBM Spectrum Scale, the client does not write a file in the file system; instead, it stripes file blocks directly across all storage nodes. So when the application performs a single file streaming operation—such as reading a large file for backup—the application reads from all storage nodes in parallel and is not bound to a single node.

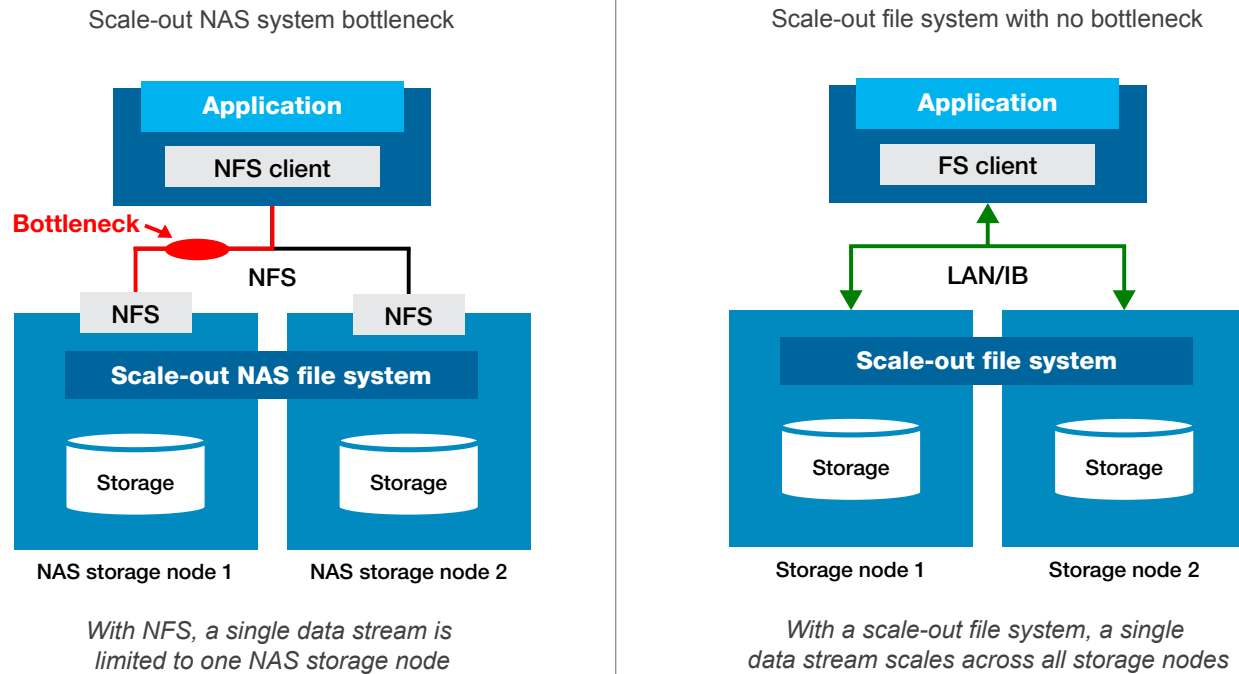


Figure 3. The NFS bottleneck does not exist with a scale-out file system.

Improved performance with IBM Spectrum Scale

With a scale-out file system such as IBM Spectrum Scale, performance can be scaled to the limits of the network. IBM has tested this performance using IBM Spectrum Protect™, which enables advanced data backup and data recovery. The test showed that a single IBM Spectrum Protect application installed on an IBM Spectrum Scale file system node can write at a speed of 5+ GB/sec to the IBM Spectrum Scale file system.¹

Free trial offer

Avoid NFS bottlenecks. Try IBM Spectrum Scale free for 90 days. See details and download at ibm.com/systems/storage/spectrum/scale/trial.html

This result is possible only because IBM Spectrum Scale uses all available storage nodes for single file streaming operations. Traditional scale-out NAS systems cannot benefit from this model because of the underlying point-to-point connection problem with NFS.

The IBM Spectrum Scale performance advantage is not limited to the NFS bottleneck. The parallel file system architecture not only spreads the load across multiple disks; it enables multiple nodes to be accessed at the same time, thereby accelerating performance within the overall storage environment. In addition, with a file system protocol such as NFS, performance is limited due to the overhead involved. This overhead does not exist with a scale-out file system such as IBM Spectrum Scale.

IBM Spectrum Scale running on [IBM Elastic Storage™ Server](#) is the first product with published results for the new SPEC SFS 2014 benchmark, the industry standard for measuring storage throughput, which has shown to be 100 times faster than the SPEC SFS reference system for the VDA workload.²

Use cases for NAS and IBM Spectrum Scale

NAS systems and scale-out file storage solutions such as IBM Spectrum Scale each have merits, depending on the use case and the implementation and optimization decisions made by the IT team. While both NAS and IBM Spectrum Scale exist to provide centralized storage for users and applications servers, the emphasis in each solution is different:

- With NAS, the emphasis is on easy connectivity to as many desktops and application servers as possible using the appropriate authentication and protocols for the use case. Because users rarely work on the same file or directory and have infrequent updates, performance is good enough.
 - With IBM Spectrum Scale, the emphasis is on achieving the highest performance and availability to the application servers and workstations by eliminating the NAS protocols and using servers, drives and network connections in parallel as much as possible.
-

Summary

It is important to differentiate between “scale-out NAS systems” and “scale-out file systems.” A scale-out NAS system does not scale performance like a true scale-out file system.

Techniques are available to mitigate the NFS performance bottleneck—for example, using Domain Name Service (DNS) to load balance. With load balancing techniques, application I/O requests to the NAS system are spread across NAS storage nodes, but do not balance single file streaming operations. Single file stream operations still land on one NAS storage node that can become a bottleneck. Parallel NFS (pNFS) standards have been created to address this known issue. However, the standards have not yet matured with extremely limited market adoption.

IBM Spectrum Scale allows organizations to scale storage performance while avoiding the NFS bottleneck and the overhead inherent in other approaches. IBM Spectrum Scale can manage file and object data to help accelerate applications with:

- Policies to automatically move data to different storage pools and tiers based on your criteria, such as the “heat” of the data
- Advanced data caching to speed access to frequently used data
- Unified storage with multi-protocol support, including NFS, SMB, OpenStack SWIFT and S3 object storage
- Drop-in replacement for the Hadoop Distributed File System (HDFS) to provide an enterprise-class data management solution for Hadoop and Spark applications

Active File Management in IBM Spectrum Scale, coupled with advanced routing and caching, accelerates applications across the data center or around geographical locations by providing local read/write performance at remote sites.

Get it your way—as software, cloud service or appliance

A mix of deployment models enables IBM Spectrum Storage software to be deployed on commodity servers for ultimate low cost with the control of on-premises deployment, consumed as a cloud service for tremendous flexibility, or as an appliance to complement traditional infrastructure deployments.

For more information

To learn more about IBM Spectrum Scale, contact your IBM representative or IBM Business Partner, or visit:

ibm.com/systems/storage/spectrum/scale

Read IBM Spectrum Scale success stories at:

ibm.com/systems/storage/spectrum/scale/resources.html

Download the ESG smart paper “IBM: The Optimal Storage Platform for Big Data” at: ibm.com/systems/storage/spectrum/scale/big-data-storage/index.html



© Copyright IBM Corporation 2016

IBM Systems
Route 100
Somers, NY 10589

Produced in the United States of America
June 2016

IBM, the IBM logo, ibm.com, GPFS, IBM Elastic Storage, IBM Spectrum Protect, IBM Spectrum Scale, and IBM Spectrum Storage are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary. THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.

¹ Nils Haustein and Frank Kraemer, “Is a scale out NAS system the same as a scale out file system?”, September 2015, ibm.com/developerworks/community/blogs/storageneers/entry/Is_a_scale_out_NAS_system_the_same_as_a_scale_out_file_system?lang=en

² The Storage Community, “What is the source of competitive advantage for businesses now-a-days?”, April 2016, <http://storagecommunity.org/easyblog/entry/what-is-the-source-of-competitive-advantage-for-businesses-now-a-days>



Please Recycle