



NetApp®



Solution Brief

Red Hat and NetApp: Parallelization of File Services with pNFS



KEY BENEFITS

pNFS

- Eliminates single-controller bottlenecks
- Achieves better scalability and manageability
- Overcomes NFS sprawl
- Coexists with other NFS protocol versions

Red Hat and NetApp Partnership

- Offers accelerated innovation through collaboration
- Offers thought leadership in file-based storage

The Challenge

Relentless data growth continues to be a major challenge for organizations of all sizes. To cope with this data growth, IT administrators focus on filling capacity holes and spend even less time with strategic storage decisions. Managing traditional storage infrastructures in this environment of massive data growth can be very difficult. Scaling beyond a few storage systems dramatically increases the burden of manageability, while load and performance balancing, servicing hardware, and backup and recovery are all increasingly difficult when scaling traditional storage infrastructures.

NFS server environments have grown from standalone servers hosting NFS to enterprise-class storage controllers, such as NetApp® FAS storage systems. Because of the simplicity of NFS as a storage solution, the number of deployments has grown significantly. NFS is no longer relegated to use with home directories, but is widely deployed for virtualization and mission-critical applications. We refer to this explosion of growth coupled with the increased burden of manageability as NFS sprawl,

which causes controller bottlenecks, inefficient data mobility, and challenges in load balancing across controllers.

The Opportunity

What if you could massively scale your storage environment with your existing IT resources and skill sets? What if you could deploy a dynamic application and storage environment that allows you to distribute data across platforms non-disruptively to eliminate bottlenecks in performance and capacity imbalance? And what if your data paths were automatically optimized to minimize latency and deliver the most efficient path from application to data as you relocate data on the fly?

NetApp and Red Hat® have been working together on industry-standard Parallel Network File System (pNFS) to develop solutions that solve the problems associated with NFS sprawl. The pNFS solution needs both client and server components to address the requirements of scalability and manageability.

NetApp addresses the challenges and opportunities of scale with the added capabilities of storage clustering and

NetApp and Red Hat have a long-standing relationship delivering solutions that help customers achieve greater performance, scalability, and business agility.

pNFS available in the NetApp Data ONTAP® 8.1 operating in Cluster-Mode storage platform. NetApp FAS and V-Series storage running Data ONTAP 8.1 operating in Cluster-Mode can scale from just a few terabytes of data to over 50 petabytes of data, all of which can be managed as a single storage entity when deployed as a cluster. Storage controllers and disk shelves are clustered together in a way that simplifies management, optimizes performance, and can eliminate planned downtime.

Red Hat, in turn, has delivered the first pNFS client to take advantage of the scalability, flexibility, simplified management, and optimized data paths in pNFS. By providing pNFS in Red Hat Enterprise Linux®, Red Hat has enabled application workloads to take full advantage of the benefits of pNFS without modification, allowing a seamless transition for existing applications.

The combination of NetApp storage and Red Hat Enterprise Linux delivers a first-to-market pNFS solution.

What is pNFS?

pNFS is the latest iteration of the widely deployed NFS standards for accessing file data over a network. Previously, NFSv4 provided many features and capabilities over NFSv3, with a focus on security and heterogeneity. The pNFS extensions included in NFSv4.1

(RFC 5661 3) offer full backward compatibility with NFSv4 features and functionalities and have been approved and ratified by industry-standards bodies. pNFS requires both client and storage components built into both the server and storage operating systems.

With traditional NFS versions 3, 4, and 4.1, the metadata and data input/output (I/O) shared the same I/O path. With pNFS, the metadata and data I/O can be split across multiple I/O paths. A metadata server handles all the metadata activities from the client while the data servers provide a direct path for data access. The flexibility to separate the metadata and data I/O paths allows pNFS solutions to deliver compelling benefits. Some of the key benefits of pNFS include:

- Elimination of single-controller bottlenecks for file data
- Improved scalability and manageability
- Direct data access path for all data
- Coexistence with other NFS protocol versions

The pNFS client shipped with Red Hat Enterprise Linux requires little configuration to provide intelligent coordination of the data path between client and server. This coordination between client and server enables improved management, I/O path optimization,

and increased I/O performance. pNFS also offers additional reliability, availability, and scalability, without the requirement for multipathing software and the associated incremental costs. Ultimately, pNFS makes sure that there is no effect on performance when a data path is moved from one controller to another; pNFS reoptimizes data paths with reduced latency for NFS data access.

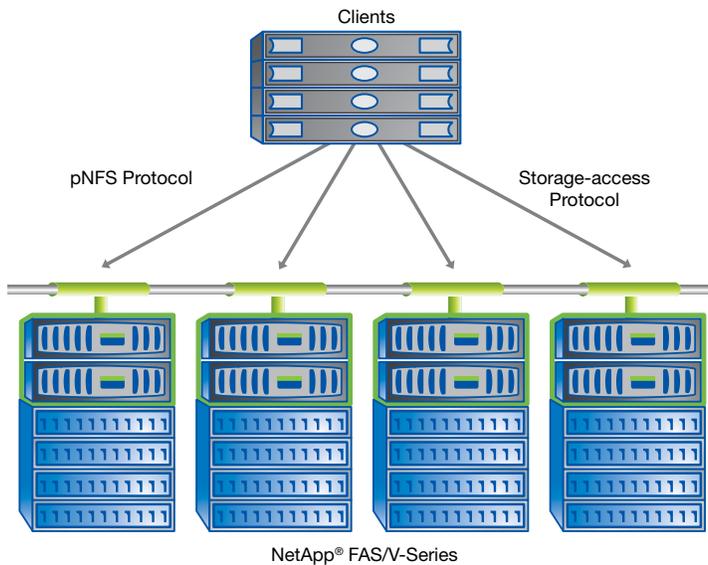
NetApp supports pNFS for file protocol access as part of Data ONTAP 8.1 operating in Cluster-Mode. The NetApp pNFS implementation leverages storage clustering technology to provide the Red Hat Enterprise Linux pNFS client with an optimized path to data that is served by the network addresses hosted on the physical storage controller.

pNFS Use Cases

All of these different pNFS capabilities make it a good fit for multiple use cases.

Business-critical applications

Business-critical applications by definition require the highest service levels. These applications demand that storage bandwidth and capacity grow seamlessly with server requirements. As NetApp storage volumes are transparently migrated to more powerful controllers in the NetApp cluster, the Red Hat Enterprise Linux



pNFS client automatically follows the data movement, self-adjusts, and reoptimizes the data path. The net result is near-zero downtime with no server or application reconfiguration required.

Multi-tenant storage solutions

Having parallel data access means that multi-tenant, heterogeneous workloads benefit directly from pNFS. The data resides on the NetApp cluster and is not tied to a specific NetApp controller. With pNFS, the Red Hat Enterprise Linux servers find the optimal data path and automatically adjust for optimum throughput.

High SLA and QoS workloads

NFSv4.1 and pNFS can provide flexibility for mounting the file system from anywhere in the cluster namespace. Clustered applications can be mounted over pNFS while legacy applications can still be mounted over NFSv3. File systems that are exported from storage can have clients mounted over different flavors of NFS so that they can coexist without making any significant changes to the applications that access the data. This level of flexibility reduces the overhead of frequent change management.

Virtualization environments

Hypervisors and virtual machines utilizing the Red Hat Enterprise Linux

pNFS client are able to maintain multiple connections per session, which spreads the load across multiple network interfaces. Think of it as multipathing for NFS, without requiring a separate multipath driver or configuration.

The Red Hat and NetApp Partnership

NetApp and Red Hat have a long-standing relationship delivering solutions that help customers achieve greater performance, scalability, and business agility.

Red Hat and NetApp have collaborated on the Linux NFS client stack for many years. Partnerships between the engineering organizations of NetApp and Red Hat have produced significant innovations and success stories. Although several companies have contributed to pNFS, joint development between NetApp and Red Hat delivered the first-to-market pNFS client and server tandem.

Delivering Value

In addition to stringent testing and validation, the joint pNFS solution benefits from the backing of world-class support organizations at Red Hat and NetApp. Whether the two companies deploy business-critical applications with strict SLAs or high-performance virtualization solutions, customers know that Red Hat and NetApp seamlessly support customers' needs.

The technical leadership and innovation delivered with Data ONTAP operating in Cluster-Mode and pNFS are the storage foundation of choice for customers who seek to move beyond reactive IT. NetApp and Red Hat are working together to deliver intelligent solutions that incorporate clustered server and storage. By overcoming traditional scaling and performance limitations, pNFS makes storage no longer a limiting factor, but an enabling one.

Additional Resources

TR-4063: "Parallel Network File System Configuration and Best Practices for Data ONTAP 8.1 Cluster-Mode"

TR-3982: "Data ONTAP 8.1 Operating in Cluster-Mode: An Introduction"

About NetApp

NetApp creates innovative storage and data management solutions that deliver outstanding cost efficiency and accelerate business breakthroughs. Discover our passion for helping companies around the world go further, faster at www.netapp.com.

Go further, faster®



www.netapp.com

© 2012 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, and Data ONTAP are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Red Hat is a registered trademark of Red Hat, Inc. in the U.S. and other countries. Linux is a registered trademark of Linus Torvalds. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. DS-3340-0612