



Hewlett Packard
Enterprise

The next generation storage solution

Migrating EMC Isilon NFS users to the HPE Scality RING

Contents

Introduction.....	2
Addressing typical Isilon pain points.....	2
Use cases.....	2
Qualifying the opportunity.....	3
Preparing for data migration.....	3
Data migration and user cutover.....	4

Introduction

This document describes how to migrate NFS-based users from a traditional NAS infrastructure to an HPE density-optimized storage solution. All of the examples are based on migrating from an EMC Isilon device to an HPE-Scality density-optimized storage solution. With minor modifications, the same process can be used to migrate NAS data from any traditional NAS vendor to an HPE-Scality density-optimized storage solution.

This document does not try to duplicate existing documentation. It provides pointers to all of the pre-existing documents on sizing, installing, and configuring an HPE-Scality NFS-based solution and focuses on the aspects of NAS data migration not covered in any other document.

Addressing typical Isilon pain points

EMC Isilon came to market to address an issue of scale, but the core technology was not designed for the challenges of today's hyperscale environments. Legacy file system concepts cripple the architecture, introducing inherent limitations for many standard deployments.

Today's data center is software-defined, and unstructured data is managed in a single namespace built on an infinitely scalable layer of build-as-you-need industry standard servers such as the HPE Apollo 4000 family. EMC restricts hardware expansion options based on the initial NAS configuration selected the customer, while the HPE-Scality solution allows a customer to grow capacity using any HPE ProLiant model that makes sense, with no hardware certification required for the customer's platform of choice.

Hard disk drives are changing rapidly, with new technologies driving higher capacities at lower price points. EMC restricts drive choices available for storage expansion, while the HPE-Scality solution allows a customer use different drive types and sizes as new technologies are introduced into the market.

Isilon systems use a protocol to propagate metadata updates, respond to requests, and distribute data that consumes a large amount of network bandwidth on a sustained basis. This means an expensive InfiniBand back end is required for an Isilon cluster creating financial viability issues (InfiniBand port count) and technical limitations (network saturation). Customers commonly experience these technical issues when a single cluster approaches 4 PB of usable capacity.

The HPE-Scality solution uses efficient data placement and protection methods, allowing a high performance platform to be deployed using standard 10 Gbps Ethernet technology. The core enabling technology is based on object-storage key and value concepts, removing any built-in theoretical capacity limits such as Isilon's thresholds of 144 nodes or 30 PB raw maximum. Scality deployments have expanded into the hundreds of nodes and hundreds of PBs, in active/active configurations across data centers, and with an ability to continue growing without any practical bound.

Isilon systems are not designed to handle mixed workloads of both small and large files efficiently. Minimum block sizes on an Isilon system can potentially waste space when storing many small files, and EMC recommends deploying separate systems to handle separate workloads. HPE servers running Scality RING abstract the storage space away from the hard drives, eliminating the concept of block size. The RING also offers multiple data protection methods with configurable policies to address data on a per-file basis, ensuring the best performance for each data type.

Use cases

As data centers are upgraded and traditional storage models are being challenged, the storage market has been converging into two tiers. The two-tier topology specifically addresses the challenges of scale, silos, availability, and flexibility.

One tier supports a modest quantity of very hot (latency sensitive) data with high-power, low-latency direct-attached storage (DAS) or SAN. The second tier stores the remaining 80 percent of data across multiple workloads in a giant and expandable storage pool, increasingly collapsing traditional Object Storage and NAS into a flexible pool of software-based storage running on industry standard servers.

The HPE-Scality solution supports access protocols using standard file system interfaces like NFS and CIFS in addition to local Linux® file system access from an application server directly to the Scality RING nodes. The HPE-Scality solution can also be configured to serve data via standard object protocols like HTTP/REST, CDMI, Amazon S3, and OpenStack® Swift. Finally, the system may also be used for tier-2, block-based virtualization workloads like Cinder running on KVM.

Customers have found this type of solution to be very attractive for the following use cases:

- **Active Archives:** Scality enables organizations to build exabyte-scale Active Archives that have the performance benefits of disk storage at a cost point that is comparable to tape, and may even present long-term cost advantages over tape. Leveraging industry-standard servers and storage optimization features such as erasure coding, companies can build exabyte-scale Active Archives without the management and migration costs of tape.
- **Content Distribution:** Content Distribution (or Content Delivery) refers to providing digital data to geographically distributed end users. Content Distribution can be in support of scale-out Web applications or for digital media streaming services, including new non-linear content distribution models such as network personal video recorder (NPVR), network digital video recorder (NDVR), or video on demand (VOD). In addition to—potentially—using a content distribution network (CDN), content providers need a scale-out storage tier, which can be provided by CDNs or managed by the content providers themselves.
- **Web and Cloud:** The Scality RING was developed using some of the key design principles used by giants such as Google™ and Amazon. As a result, Scality enables Web application and cloud service providers to build petabyte-scale storage infrastructures that meet all the performance, availability, and durability requirements for scale-out Web applications at a much lower cost than public cloud offerings. The Scality RING features a wide range of interfaces and protocols to enable easy application integration.
- **Distributed Computing:** While Scality is not ideal for all scientific computations, it is ideal for archival of data for highly distributed or parallel computing. Customers use the RING as a secondary tier to a high-performance distributed file system like Lustre or GPFS.

Qualifying the opportunity

Use the following “gateway questions” to determine if an HPE-Scality active archive solution can provide equivalent functionality to the customer’s EMC Isilon:

- Is the EMC Isilon system used as an active archive device? If the application is a structured database or the I/O latency requirements are less than 10 ms at peak utilization, then an HPE SAN should be used instead of SL4540 or Apollo 4500 servers to provide the storage capacity.
- Does the amount of data to migrate exceed 200 TB? If so, it is a candidate to migrate to the HPE-Scality solution. Data can come from a single EMC Isilon device or by collapsing multiple NAS silos into a single HPE-Scality density-optimized solution using NAS file access.

If the answers to these two questions are positive, you will also need to ask the following questions in order to determine the number of Scality NFS connector servers required for the actual deployment:

- What is the peak I/O throughput requirement?
- What is the peak IOPS requirement?
- Is there a specific latency requirement?

Preparing for data migration

Once all of the following actions are complete, the NAS data migration can begin.

1. Contact your HPE customer account representative to work with you to size and order a HPE-Scality solution based on your storage and throughput requirements.
2. The HPE account representative will work with you to plan and schedule the actual hardware and software configuration once the equipment has arrived at the customer site. The Scality RING software will be installed by a field engineer to ensure correct deployment and configuration.
3. Using the Scality Supervisor UI, configure an initial volume inside the Volumes tab as described within docs.scality.com/display/R50/Supervisor+Volumes+Tab.

4. Get the group and user ID information for the Isilon storage device using the following command:

```
# isi auth users list --provider="<provider-type>:<provider-name>"
```

5. Define the group(s) and UID(s) used by the Isilon device to access NFS files on the Scality NFS connector nodes using standard Linux tools.

6. The network infrastructure between the Isilon cluster and the HPE-Scality cluster must be 10 Gbps to achieve a reasonable transfer rate. If possible, the MTU size should be set to 9,000 bytes. It is also recommended that data migration be performed over an isolated network. If the data migration must use the same network as normal NFS clients, then use the rsync transfer rate limitation option to ensure the data migration doesn't affect everyday operation.

Data migration and user cutover

Data migration is a simple three-step process.

First, a test should be run on a small set of files to confirm that average throughput is as predicted, time stamps are maintained, and the Linux user and group permissions are preserved.

Once the test phase works perfectly, then all of the files on the Isilon device can be migrated to the HPE-Scality cluster.

The transfer process must be run repeatedly to keep the two storage configurations in sync until a scheduled maintenance window when all the users can be cutover from the Isilon system to the HEP-Scality cluster.

The data migration workflow description makes the following assumptions:

- The standard Isilon NFS export mount point named /ifs is mounted on the Scality system doing the transfer as /mnt/ifs.
- The NFS mount point for the Scality NFS connector is /mnt/scality_ifs.

To perform the migration successfully, execute the steps listed in the order list that is as follows:

1. Identify a small number of files and directories that will be used to verify that the group and user information has been successfully copied to the server running the Scality NFS connector.
2. Run rsync to transfer the files. For example, to synchronize the Scality-HPE storage cluster with the Isilon device using the NFS mount points described in the assumptions previously at a bandwidth not to exceed 200 MB/s (network I/O threshold known not to affect normal operation at point during the day) execute:

```
# rsync -avz --delete --progress --bwlimit=200000 /mnt/ifs /mnt/scality_ifs
```

If you wish to see what rsync will do prior to taking action, execute:

```
# rsync -avz --dry-run --delete --progress --bwlimit=200000 /mnt/ifs /mnt/scality_ifs
```

3. Periodically run rsync to minimize the number of updates on the cutover day.
4. Prior to the cutting over to the HEP-Scality storage, do the following:
 - a. Confirm that each user ID can access the files it is permitted to access via the NFS mount point(s) for the HPE-Scality storage.
 - b. Verify user and group owners plus file permissions match by executing:

```
# ls -lR /ifs; egrep -v '\/ifs|total ' t1 > t2
```

```
# ls -lR /scality_ifs > t3; egrep -v '\scality_ifs|total ' t3 > t4
```

```
# diff t2 t4 > diff.log
```

5. During the maintenance window:
 - a. Execute “rsync -avz -delete --progress /mnt/ifs /mnt/scality_ifs”
 - b. Disconnect the Isilon device from the network
 - c. Stop NFS on HPE-Scality connectors
 - d. Create /mnt/ifs on each Scality connector
 - e. Change the export mount point from /mnt/scality_ifs to /mnt/ifs for each Scality connector
 - f. Restart all of the Scality connector servers

Execute pre-planned tests to verify the cutover has been successful. If the answer is “yes,” nothing more needs to be done. If the answer is “no,” the actions in steps B through F mentioned previously must be reversed to bring the Isilon device back online.

Learn more at

hpe.com/servers/bigdata



Sign up for updates

★ Rate this document

© Copyright 2015–2016 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Google is a registered trademark of Google Inc. OpenStack Word Mark is either a registered trademark/service mark or trademark/service mark of the OpenStack Foundation in the United States and other countries and is used with the OpenStack Foundation’s permission. We are not affiliated with endorsed or sponsored by the OpenStack Foundation or the OpenStack community. Pivotal and Cloud Foundry are trademarks and/or registered trademarks of Pivotal Software, Inc. in the United States and/or other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries.