How to Migrate Data to Oracle ZFS Storage Appliance Using Shadow Migration

February, 2024, Version [1.0]

Public

Disclaimer

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described in this document remains at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

Table of contents

List of figures

Purpose

This document provides an overview of the Shadow Migration feature of Oracle ZFS Storage Appliance. It is meant for storage administrators and users to understand the concept, implementation, and benefits of using Shadow Migration for data migration tasks between legacy systems and the Oracle ZFS Storage Appliance.

Introduction

Oracle ZFS Storage Appliance is a purpose-built storage system meticulously engineered to seamlessly integrate with Oracle software, delivering unparalleled storage performance and efficiency.

It offers an essential tool to Oracle ZFS Storage Appliance administrators known as Shadow Migration. Shadow Migration allows the storage administrator to efficiently move data residing in legacy systems using NFS protocol. This feature enables the seamless migration of file systems, previously stored on legacy systems to the Oracle ZFS Storage appliance, along with the benefits of its unique features and capabilities.

In the context of this document, file systems are categorized as follows:

- **Legacy File System:** This represents the original data source at the start of the data migration process, residing on the legacy storage system.
- **Replacement File System:** This serves as the migration destination and is usually located on the replacement storage system.

After the migration process successfully completes, the replacement file system becomes the definitive data source and the legacy file system storage can be repurposed or taken out of service.

Data Migration Methods

Data migration projects traditionally involve two primary methods: synchronization and interposition. In the following section, we will explore these different methodologies along with their main advantages and disadvantages.

Synchronization Method

The synchronization method involves copying data directly from the legacy source to the replacement source while still allowing data access by clients through the primary route.

Figure 1. Architecture of data migration using synchronization.

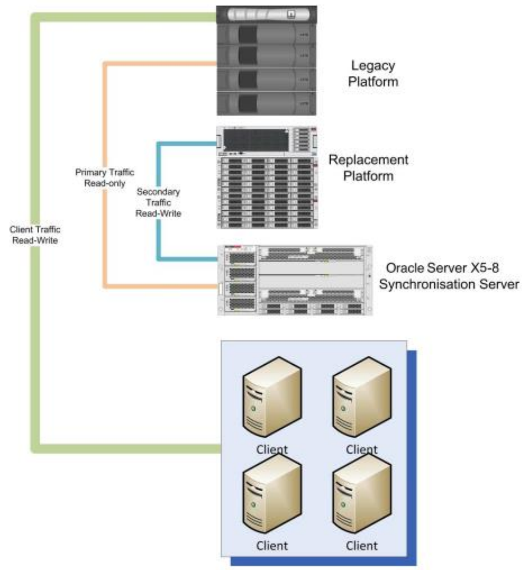

Synchronization utilities analyze the data state in both legacy and replacement sources to identify the minimum amount of data that needs to be copied. Only the changes that have occurred since the last synchronization run are copied, including newly created or updated files, deleted files, newly created directories, and deleted directories in the replacement source.

Some synchronization utilities divide the files to be copied into chunks which are then checksummed on both the legacy and replacement sources. These checksums are compared to determine the minimum amount of data to be copied. If the checksums match, the chunk doesn't need to be copied from the legacy source to the replacement source. If the checksums do not match, the chunk is copied across the network and applied to the replacement file.

Upon initiation of the synchronization run, a 'plan of action' is generated outlining the list of files that require copying from legacy to the replacement systems. In Figure 2, the plan is represented by dashed lines, indicating the files scheduled for transfer to the previously empty replacement file system.

Figure 2. File copy 'plan of action'

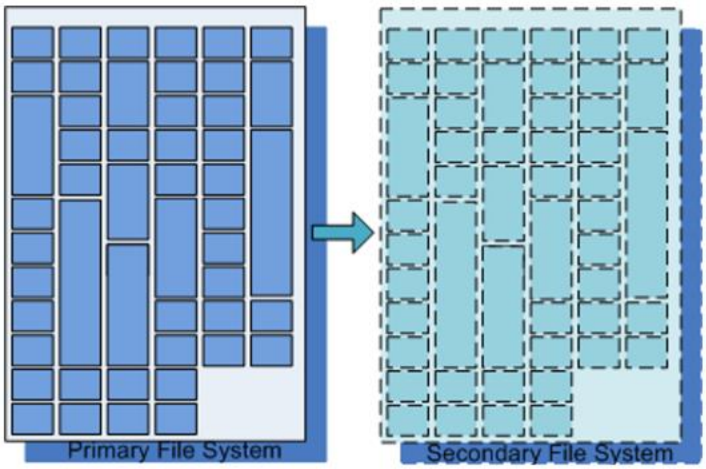

Once the plan of action has been formulated, the actual copying of files takes place between the legacy and replacement file systems. Any new additions made to the legacy file system after the plan's creation will not be included in the 'plan of action' and thus will not be transferred between the file systems.

Figure 3. Synchronization 'copy' phase

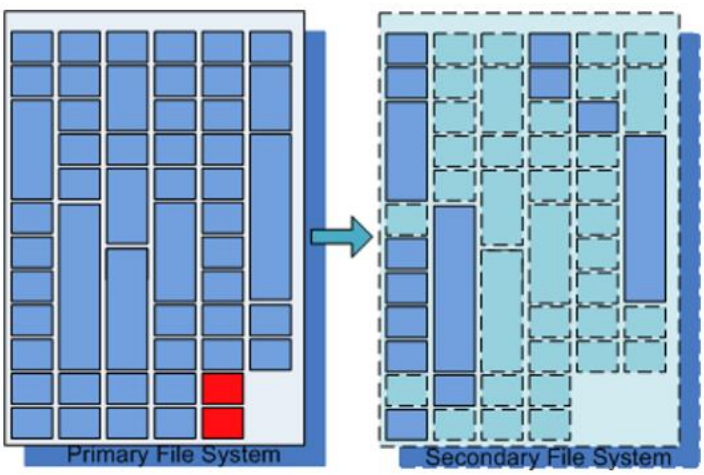

In Figure 3, two new files are shown as red boxes, indicating that they were created in legacy file system. However, these files do not have corresponding entries in the 'plan of action'.

During synchronization process, any changes to the legacy data source, whether during or after the file selection, are effectively lost for the duration of the synchronization run. If a file included in the 'plan of action' undergoes changes in the legacy system before the copy phase is executed, those changes will be copied across. However, if any changes occur after the synchronization process handles the file, those changes will be lost.

To address this situation, it might be necessary to run the synchronization process multiple times to ensure that previously missed changes are captured and dealt with. The number of iterations in the synchronization process is determined by the volatility of the file source, as it is an iterative process aiming to achieve complete data consistency between the legacy and replacement systems.

At some point, it becomes essential to halt access to the legacy source to guarantee that all changes have been successfully applied to the replacement source. To ensure a seamless transition, no update actions should be permitted on the legacy source. If update access times to files in the legacy source is necessary for auditing or compliance requirements, it is advisable to disallow all access, including read-only access. This ensures that the data remains consistent and up to date during the migration process.

During the downtime period, the final synchronization should be performed, ensuring that the synchronization process is applied until no further changes are needed between the legacy and replacement sources.

Once the synchronization run is successful and all data is consistent, the legacy source should be disabled, and all clients should start referring to their replacement source as the authoritative version. To minimize administrative burden, the administration can facilitate the switch by swapping IP addresses between the sources.

It is important to note that for all files to be migrated, they must be readable by the interposer device, which should have effective root access to the legacy file system. Additionally, to ensure seamless migration, the permission and access control lists for the files and folders must be duplicated on the replacement file system. Both the legacy and replacement platforms, along with the interposer, must have the same name service configuration to ensure proper data access and management during the migration process.

Advantages of Synchronization Migration Method

Advantages of Synchronization Migration Method:

- Data Redundancy: The synchronization method ensures that two copies of data exist in both the legacy and replacement sources after the initial synchronization process. This redundancy provides an added layer of data protection. While the replacement source may not be as up to date as the legacy source, it is still current up to the latest synchronization run. This can be sufficient for disaster recovery policies, ensuring data availability in case of unexpected events.
- Flexibility in Timing: The synchronization process can be executed multiple times, allowing administrators to fine tune the data transfer. This final synchronization downtime and switchover decision can be made just before the last synchronization runs, offering greater flexibility and control over the migration process.

Disadvantages of Synchronization Migration Method

Disadvantages of Synchronization Migration Method:

- Volatility Impact: The duration of all synchronization runs, except the initial one, is solely determined by the volatility of the file system. This makes it challenging to accurately quantify the downtime required for the final synchronization, as the entire file system contents might have changed with additional files created.
- Downtime Challenges: The final downtime needed for synchronization may be difficult or even impossible to estimate effectively. In some cases, the required downtime window might not be tolerable for continuous client access.
- Delayed ROI: Until the final synchronization process is completed, clients will continue using the older storage or servers, while the new replacement storage with its enhanced features remain idle. As a result, the return on investment won't begin until after the final synchronization process.

- Network Overhead: If the legacy platform lacks the capability to perform synchronization, an additional network hop is introduced. This involves the synchronization server reading source data over the network and then transmitting the same data to the replacement platform, leading to potential network overhead.

Interposition Method

When employing interposition method, an interposing device is positioned between the legacy source and the client users of the source. The device seamlessly handles the migration from legacy to replacement sources in the background, without affecting the clients. During the migration process, the legacy file system must remain unchanged, as any modifications could lead to issues or data corruption. It is essential to avoid sharing the legacy file system with other devices, and the interposer should be the sole point of interaction during this phase.

Figure 4. Synchronization 'copy' phase

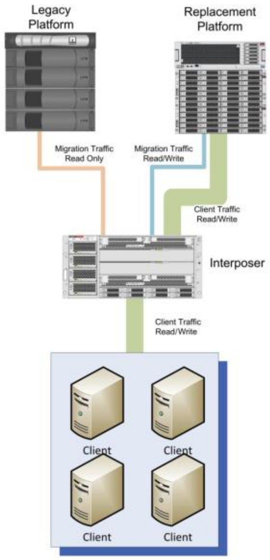

During the migration process using interposition methodology, some downtime is necessary to direct clients to the interposition equipment. The interposing device intelligently selects the data source and the type of access requested (read-only or read/write).

If a file has already been copied to the new platform, any I/O request is fulfilled from here. However, if the file hasn't been migrated yet, the interposition process migrates the file to the new platform before allowing the I/O request to proceed, served by the replacement platform.

This approach may introduce some latency to I/O requests, but only if the file hasn't been migrated or recently accessed since migration. Larger files might encounter more noticeable latency, while smaller files usually copy over quickly. If a specific application relies heavily on numerous small files, there might be increased latency during the first access.

After the migration is complete, the interposing equipment can be removed. Some downtime may be required to redirect clients to fully migrated replacement source, which becomes the authoritative data copy.

During the interposition migration, clients access the legacy file system through the replacement file system, although no actual data transfer has occurred yet.

Figure 5. Interposition Initiation

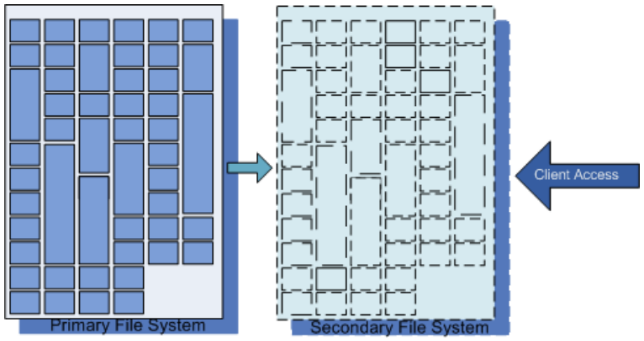

As the migration proceeds, files are transferred from legacy to replacement file systems. When a file is requested for reading, one of the two scenarios unfolds, depending on whether the file has already been migrated to the replacement file system or not:

- If the file is already migrated to the replacement file system, the I/O request is promptly processed and fulfilled from the replacement file system (as shown in figure 6).

Figure 6. Interposer Read-Case 1: The file has already migrated to replacement file system

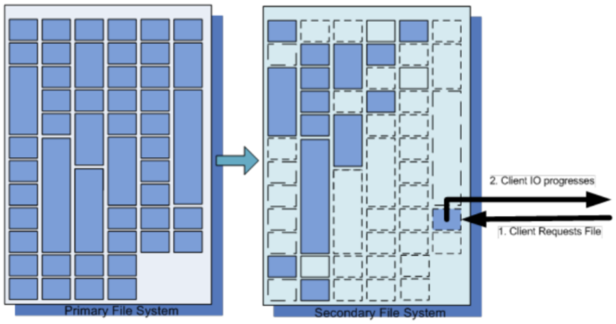

- If the file has not yet been migrated, the I/O request is temporarily paused while the interposer begins the file migration. Once the migration is complete, the I/O request is allowed to continue and is satisfied from the replacement file system (as shown in figure 7).

Figure 7. Interposer Read-Case 2: The file has not yet migrated from primary/legacy file system

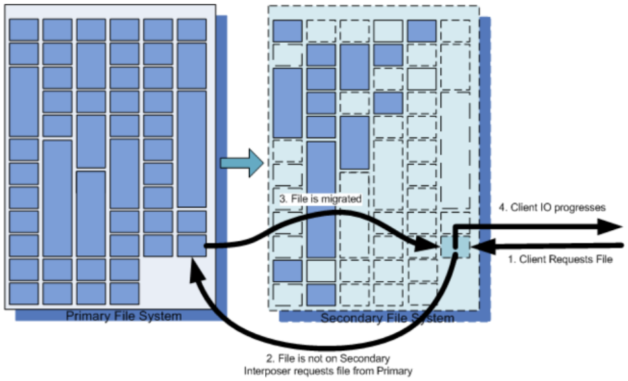

Any changes made to a file initiate the migration process, copying the file system from the legacy to the replacement file system (if not already copied). These modifications are performed on the replacement file system without affecting the legacy file system (as depicted in figure 8). Throughout this process, client access is served from the replacement file system.

Figure 8. Interposition in progress, showing new files created directly on secondary/replacement file system

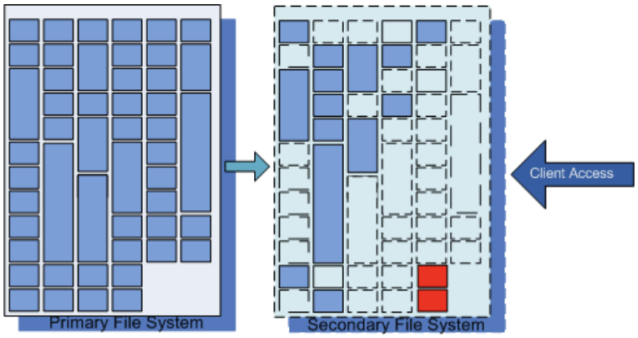

Once all files are successfully copied, the interposition migration automatically stops, and the legacy and replacement file systems no longer require further maintenance.

Advantages of Interposition Method

The interposition method offers several advantages:

- Immediate Feature Access: Clients can immediately access the features and services of the destination platform once the data migration begins.
- Minimal Downtime: The only downtime needed is to modify the client configuration to use the interposition device as the client access path. In some cases, downtime might be required to remove the interposition device and further configure clients to refer directly to the replacement source.
- Seamless Migration: The migration process occurs transparently in the background, without disrupting client access to the data.
- Flexibility: The interposition method allows for more flexibility in handling the migration, and the final switchover can be performed at the administrator's convenience.
- Efficient Resource Utilization: Unlike traditional methods, the replacement platform with its enhanced features is actively utilized by clients during the migration, making the migration process more efficient.
- Simplified Rollback: If needed, rolling back to the legacy system is easier since the legacy system remains untouched during the migration, reducing the risk of data loss.
- Reduced Network Load: Unlike traditional methods, where data is copied between systems, the interposition method minimizes network load as the data is directly accessed from the source.

Disadvantages of Interposition Method

While the interposition method provides a smoother, more streamlined migration process with minimal impact on client access and system performance, it has some disadvantages:

- Prolonged Device Usage: Interposition devices may remain in place long after the migration project is completed due to the downtime required to switch clients to the replacement storage platform. This can lead to underutilization of the interposer device.
- Latency and Indirection: The interposer introduces a small amount of latency for each request, as data is routed through it. While the latency is minimal, it can accumulate over time and impact overall system performance.
- Commitment Point: The decision to proceed with the migration project is made when the interposer device is deployed in the data path between clients and legacy/replacement platforms. Once in place, it becomes challenging to undo the migration and synchronize any new or modified files back to the legacy platform.
- Wasted Latency Post-Migration: After the migration is completed, the additional latency introduced by the interposer becomes unnecessary unless the device provides unique data services not available on the replacement platform.

Which Method is Better?

Selecting the best method for data migration is subjective and relies on various factors. It depends on the access patterns and file size distribution within the file system. Both synchronization and interposition methods can be deployed in a data migration project, but they cannot be used simultaneously on the same file system.

For situations where large files dominate and read-only access is prevalent, the synchronization approach might be more suitable. On the other hand, if the migration source contains numerous volatile files, the interposition method might be preferred due to its manageable downtime requirements.

There is no one-size-fits-all solution, and the choice between methods should be based on local policies and specific access patterns for each file system being migrated. Careful consideration of these factors will help determine the approach with the least impact on the overall migration process.

Accomplishing Data Migration Using the Oracle ZFS Storage Appliance Shadow Migration Feature

The Oracle ZFS Storage Appliance serves as an interposer device with a significant advantage: it does not require removal at the end of the migration process. The capability is made possible by the service called Shadow Migration, which allows for smooth data migration from legacy platforms using NFS. It also facilitates data migration between storage pools when the source is defined as local to the Oracle ZFS Storage appliance.

Figure 9. Oracle ZFS Storage Appliance Shadow Migration Feature

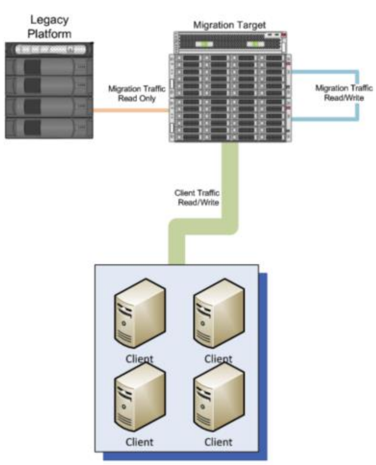

Using Shadow Migration to Move Data from an NFS Source

Shadow Migration feature can be enabled from Oracle ZFS Storage Appliance's browser user interface (BUI) under Configuration > Services > Shadow Migration.

Figure 10. Shadow Migration service accessible from Oracle ZFS Appliance's BUI

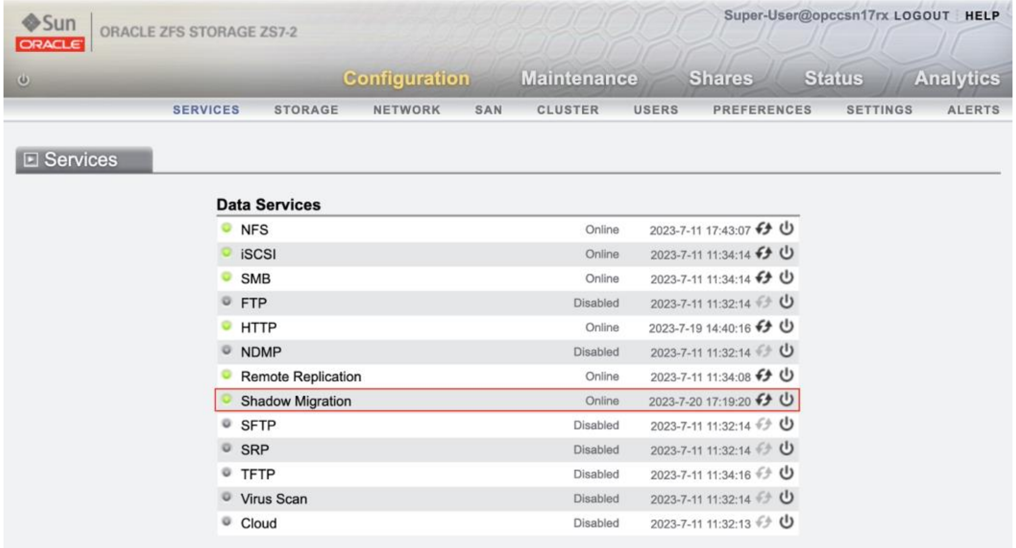

When employing Oracle ZFS Storage Appliance for data migration from legacy NFS source, it is important to optimize the performance of the legacy platform to ensure maximum bandwidth to Oracle ZFS Storage Appliance. Redirecting all client access to the appliance and maintaining a contention free, high bandwidth network link between the legacy platform and Oracle ZFS Storage Appliance can be achieved by setting up a dedicated Ethernet switch exclusively for these platforms. During the migration process, it is advisable to allocate multiple interfaces on each platform.

To facilitate the migration, the Oracle ZFS Storage Appliance replaces the legacy platform as the client-side service. Ensure that any file system being migrated are configured as read only, as previously mentioned. To disable client access to the legacy platform, it should be removed from any switch or router accessible to the clients. This may require reconfiguring the IP addresses on both platforms to accommodate the temporary network created for the migration process. After this step, the clients can be reconfigured to utilize the services of the replacement Oracle ZFS Storage Appliance, or the appliance can be configured to adopt the networking addresses of the legacy equipment.

Performance Considerations

The performance of Shadow Migration depends on several factors:

- File Size Profiles
- Number of files
- Network contention on the migration network
- Network configuration of the migration network
- Destination storage pool data protection configuration
- Log and cache configuration
- Number of devices in the destination pool
- Performance of the legacy device
- Volatility of file system during migration
- Access pattern of existing data during migration process

To achieve optimal migration performance, the first step involves ensuring a congestion-free network path between the source and destination devices with high bandwidth network interfaces and devices. Using a dedicated network switch solely for migration traffic, as shown in figure 10, can be a temporary solution.

Figure 11. Possible physical network configuration

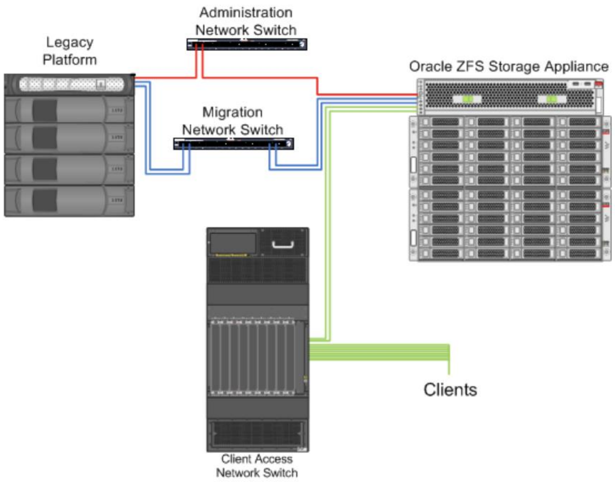

Throughout the migration process, maintain administrative access on both legacy and replacement platforms to monitor performance and detect hardware failures.

Consider configuring multiple interfaces as a single virtual interface, such as IP multipathing (IPMP) or LCAP aggregation link interface in Oracle ZFS Storage appliance family, to take advantage of these setups. However, keep in mind that the legacy platform must also support a similar configuration. The specific methodologies and terms used may vary depending on the manufacturer.

Before allocating all network devices to the migration process, assess the client access requirements during this phase.

Keep in mind that the file size profile significantly impacts migration performance. Migrating a large number of small files can take longer compared to the same overall size of very large files due to the overhead of creating files with correct attributes and handling individual file open and close operations.

When conducting data migration, consider testing various file profiles, thread configurations, and NFS versions to understand the potential variations in migration times.

Disable deduplication and compression and maintain consistency by migrating file systems from the same set of disks and to the same pool configuration.

Pay attention to the Maximum Transmission Unit (MTU) size, but be cautious when changing it, as it can impact performance in high-contention environments and lead to data stream corruption. The recommended transmission unit (MTU) size for Shadow Migration is 9,000 bytes.

For NFSv4 configuration, prioritize configuring the legacy platform, especially if complex Access Control Lists (ACLs) require maintenance.

Ensure that log and cache devices within Oracle ZFS Storage Appliance are well-configured to optimize migration performance and data access by clients during migration.

Figure 12. Time taken to migrate 64MB of data with and without log devices

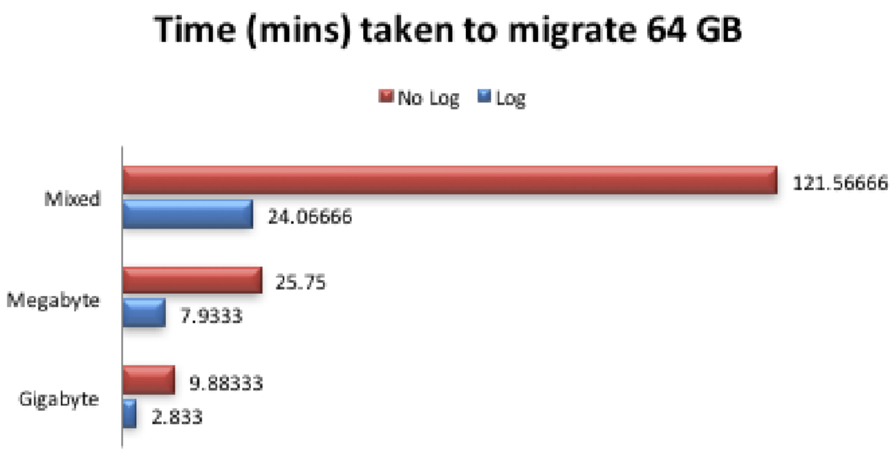

Using Shadow Migration for Local Data Transfer

Shadow Migration is not limited to migrating data from legacy sources; it can also be utilized to transfer data within an Oracle ZFS Storage Appliance from one storage pool to another. This feature is particularly beneficial when moving data from a high-performance pool to a high-capacity pool, where the data's access patterns have changed, and high-performance support is no longer required.

Additionally, Shadow Migration enables the modification of settings that are otherwise unchangeable, such as adjusting the file system record size or creating a compressed version of the file system.

Regardless of the specific use case, the data migration process remains the same. The source file system must be accessed in read-only mode or not accessed at all, as Shadow Migration is not designed to handle dynamic source file systems.

Performance tests for local migrations were conducted using the same file size distributions as in NFS-based Shadow Migration. The test results should be interpreted as relative times for varying file size distributions rather than absolute performance indicators.

Figure 13. Time taken to migrate 128GB of data from local Source

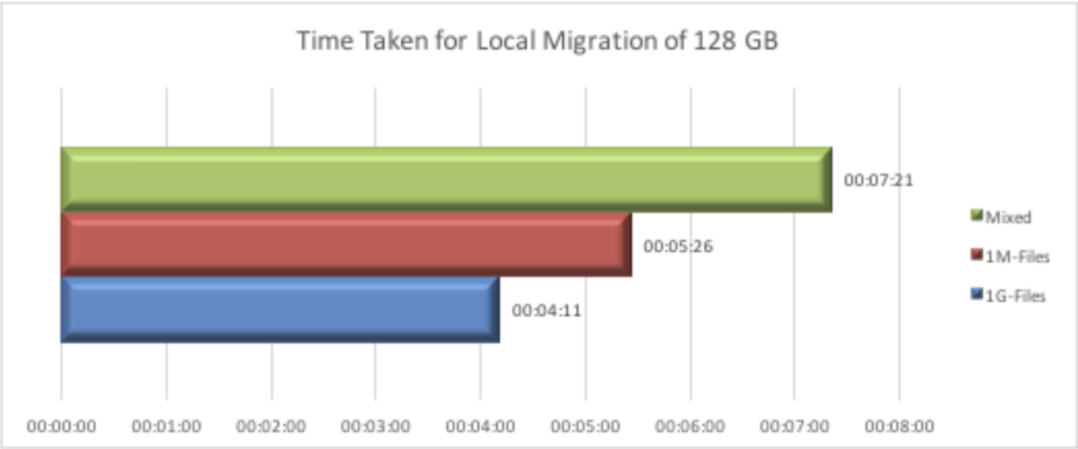

Figure 14. Throughput to migrate128GB of data from local Source

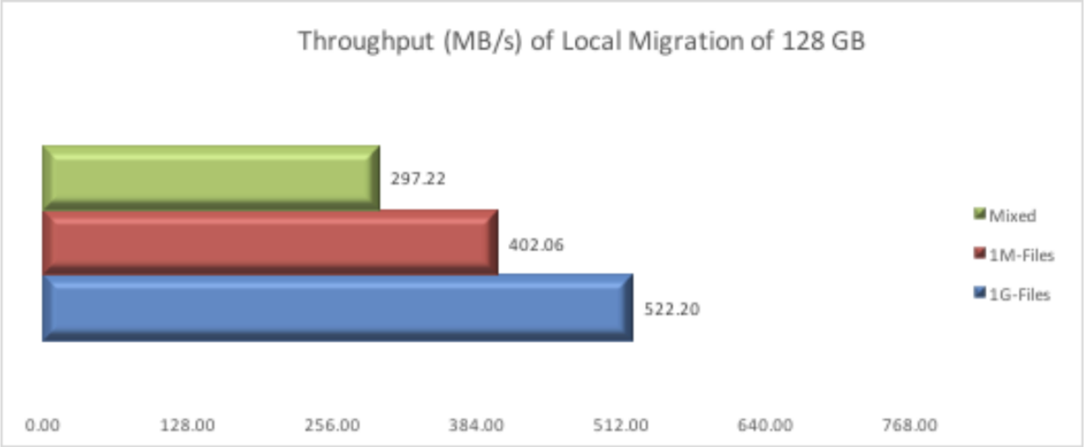

Monitoring Shadow Migration Using Analytics

Oracle ZFS Storage Appliance is equipped with robust analytics that enable storage administrators to closely monitor the system's performance. These analytics provide valuable insights into various metrics related to Shadow Migration:

Shadow Migration bytes:

- Broken down by file name
- Broken down by project
- Broken down by share
- Represented as raw statistics

Shadow Migration ops:

- Broken down by file name
- Broken down by project
- Broken down by share
- Broken down by latency
- Represented as raw statistics

Shadow Migration requests:

- Broken down by file name
- Broken down by project
- Broken down by share
- Broken down by latency
- Represented as raw statistics

These performance metrics provide administrators with a comprehensive view of the data migration process and facilitate detailed analysis to optimize the efficiency and effectiveness of Shadow Migration. Additionally, monitoring the network load during migration is vital for a successful migration process.

Figure 15. Shadow Migration analytics display in Oracle ZFS storage Appliance BUI

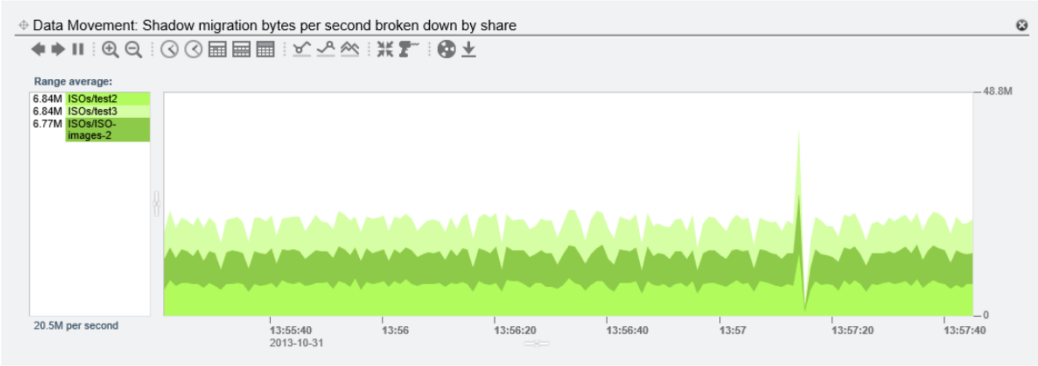

Conclusion

Shadow Migration offered by Oracle ZFS Storage Appliance proves to be a valuable addition to the administrator's toolkit. It automates and streamlines data migration from legacy platforms to Oracle ZFS Storage Appliance, as well as within the appliance family when needed. However, it's important to note that while Shadow Migration is highly beneficial for monitoring migration performance, it may not be a one-size-fits-all solution for every migration scenario, given the diversity of file sizes and types.

What sets Shadow Migration apart is its advanced analytics capability, which allows administrators to monitor progress at both the individual file level and by share, while also gaining insights into client I/O access patterns. This detailed information enables accurate estimation of the total migration time, empowering administrators to plan and execute migrations efficiently.

Connect with us

Call +**1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

blogs.oracle.com facebook.com/oracle twitter.com/oracle

Author: Scott Ledbetter, Sheetal Sabharwal