



F5 White Paper

# Driving Data Migration with Intelligent Data Management

F5 and NetApp are reducing the cost and complexity of managing file storage. F5 ARX complements and enables transparent migration to the wide range of NetApp storage solutions.

**by Renny Shen**

Product Marketing Manager

**and Matt Quill**

Principal Solution Manager



# Contents

---

<b>Introduction</b>	<b>3</b>
---------------------	----------

---

<b>An Overview of the F5 and NetApp Solution</b>	<b>4</b>
Focus on Data Migration	4
About F5 ARX Solutions	4
About NetApp Storage Solutions	5

---

<b>A Closer Look at File Virtualization</b>	<b>5</b>
Seamless Migration to NetApp File Storage	7
Non-Disruptive Technology Refresh	8
Online Migration Among Multiple Types of Storage	8
Optimized NetApp Deduplication	9
Optimized Backup and Recovery	9
Storage Pooling and Aggregation	10

---

<b>Conclusion</b>	<b>11</b>
-------------------	-----------

---



## Introduction

Every IT team struggles to keep up with storage demands while minimizing downtime and disruptions. The rate of unstructured data growth has made this job increasingly challenging—pushing storage resources to the limit, increasing management complexity, and driving up costs. Additionally, as the volume of data has increased, so, too, has user demand for uninterrupted access to vital information.

The impact of this growth is felt at every level of the enterprise. Typically, data migrations and storage provisioning are manual processes that consume IT resources while disrupting users and applications. The constant growth of data means backups take increasing amounts of time to complete, and they are more costly to maintain. Utilization rates vary dramatically, as some file servers run out of capacity while unused capacity languishes on others.

Complicating matters further is the complexity and inflexibility of traditional storage infrastructures. Many storage infrastructures have evolved haphazardly over time, and now resemble discrete islands of storage rather than a unified storage pool. These islands are often composed of resources from multiple vendors that must be managed separately. In addition, users and applications are statically mapped to physical devices, creating configuration challenges whenever changes occur in the environment.

For years, IT has responded to data growth by continuously adding more storage. However, this strategy is no longer sustainable: While the cost of capacity has dropped, management costs continue to rise and now far outweigh acquisition costs. The reality is that data growth outpaces IT staffing and storage budgets by a wide margin. In the face of this data deluge, it is imperative for IT organizations to implement a smarter storage strategy that addresses the management complexity of data growth as well as the storage costs.

The right technology and data management approach can make data growth much more manageable. The key to simplifying data management is building an agile storage infrastructure that can easily adapt to evolving business and technology requirements. The joint solutions from F5® and NetApp can help eliminate restrictions on data mobility and make it possible for data to move freely, even automatically, without downtime or disruption. This gives IT the ability to address many of the issues associated with data growth today, including data migration, backup, and utilization.



# An Overview of the F5 and NetApp Solution

The definition of an agile IT environment is one that enables enterprises to respond quickly to change by adding, removing, and changing application and data services on demand. In the storage infrastructure, the key to agility lies with the ability to move data without incurring downtime. F5 and NetApp are collaborating to provide solutions that facilitate the transparent migration of data to and among NetApp file storage systems, enabling IT agility in the storage infrastructure.

## Focus on Data Migration

Many of the most common management tasks—such as data migrations, network-attached storage (NAS) consolidations, capacity balancing, platform upgrades, and technology refreshes—require data to be moved. However, moving data is an inherently disruptive operation that results in downtime and business interruptions. Affected client systems must be reconfigured to access data through new network shares and mount points, and IT often spends considerable time updating broken links and references in applications, documents, and user drive mappings. To minimize the impact, IT typically negotiates windows of downtime, often during nights and weekends, with business groups and users. As a result, migration projects—even those in small environments—can take weeks or months to implement.

F5® ARX® file virtualization devices and NetApp storage solutions work together in several ways to enable transparent data migration by removing the restrictions that hinder data movement and implementing automated intelligent data management policies. The joint solutions make it possible to move data between storage subsystems—even during business hours—without affecting applications or users, or requiring client reconfiguration. The result is reduced data migration times of up to 90 percent and the elimination of costly downtime.

## About F5 ARX Solutions

F5 ARX solutions virtualize the file storage environment, enabling intelligent network-based data management. ARX decouples the logical access to data from the physical location of that data, eliminating disruption and simplifying the way data is accessed, moved, and managed. In addition, with ARX policy-based data management capabilities, IT organizations can create “set and forget” policies to automate data movement, further simplifying management and reducing operational overhead.



## About NetApp Storage Solutions

The NetApp family of fabric-attached storage (FAS) systems helps organizations reduce costs and complexity, minimize risks, and control change by simplifying data management. The FAS product line provides storage solutions for a broad range of needs—from remote office applications with the FAS2000 series to the largest enterprise applications with the FAS6000 series.

## A Closer Look at File Virtualization

Central to the NetApp and F5 collaboration is the ARX file virtualization technology. File virtualization abstracts the physical storage environment, decoupling the logical access to file data from its physical location. Often referred to as a Global Namespace, this abstraction simplifies access to file data and hides physical changes from clients. Clients now access their file data logically through the Global Namespace instead of connecting to physical storage resources. The Global Namespace is not a new file system; it is simply an abstraction of the file systems already in place in the underlying storage layer. It uses industry-standard protocols—common Internet file system (CIFS) and network file system (NFS)—to communicate with the clients that access the data and the storage servers that handle the data.




Figure 1: ARX decouples logical access to file data from its physical location.

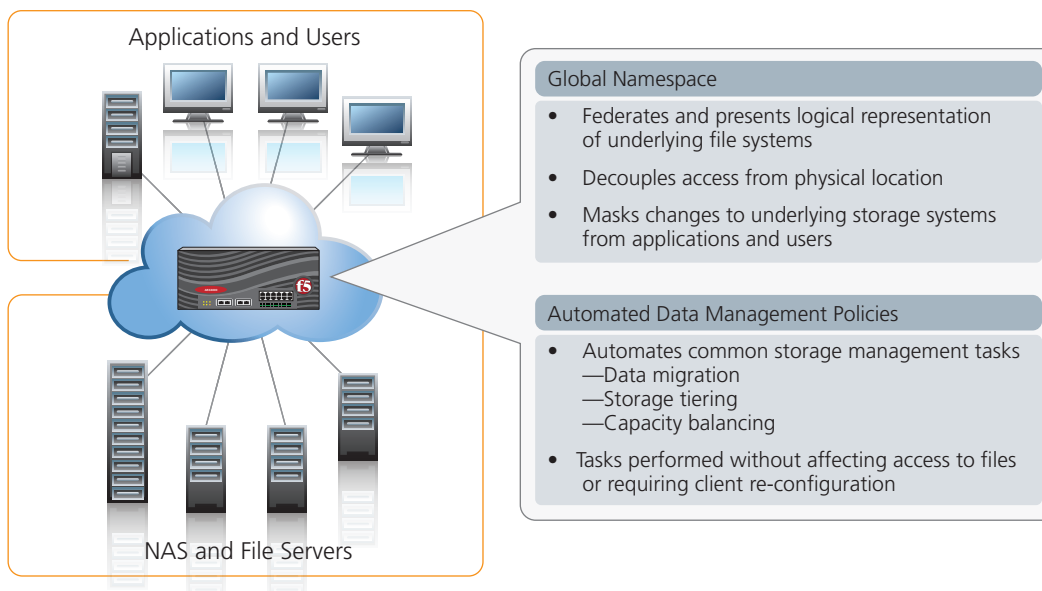


# ARX: Driving Data Migration in NetApp Environments

A Global Namespace is critical to the reducing complexity and inflexibility of traditional storage infrastructure. It hides physical differences between storage resources, and can consolidate hundreds or thousands of physical client-to-resource mappings to a much smaller and more manageable number of logical mappings. These logical mappings never change, which simplifies access to files and enables storage management tasks to take place without impacting users and applications.

In addition to a Global Namespace, ARX provides a powerful suite of real-time data management policies that simplify and automate the movement of data. These policies enable organizations to scale their storage environments more efficiently, introduce NetApp file storage with minimal interruption, and streamline overall NAS management. ARX solutions complement NetApp environments by enabling:

- Seamless migration to NetApp file storage.
- Non-disruptive technology refresh.
- Online migration between multiple types of storage.
- Optimized NetApp deduplication.
- Optimized backup and recovery.
- Storage pooling and aggregation.




**Figure 2: ARX provides a Global Namespace and a powerful suite of real-time management policies that simplify and automate the movement of data.**



## Seamless Migration to NetApp File Storage

Implementing new file storage in legacy environments presents many challenges. Migration and consolidation of legacy file storage devices is typically a labor-intensive endeavor, involving a considerable amount of downtime, manual scripts or utilities, post-migration validation of data and access, and cumbersome resolution of file conflicts. Ensuring that a full and reliable copy of file data has been migrated to the new storage platform using traditional methods can be a significant drain on already strained IT organizations.

By using ARX to hide physical changes from clients, even large-scale and complex data migrations can take place during business hours without disrupting user access. ARX employs powerful management policies to automate the movement of data—from individual files to entire file systems. These policies can dramatically simplify complex migration projects, including those that involve restructuring file systems or rebalancing data across those file systems. As a result, organizations can seamlessly migrate from alternate NAS platforms or consolidate multiple legacy file servers onto NetApp file storage in significantly less time, and with less disruption, than traditional methods.



**Figure 3: ARX enables organizations to migrate to NetApp file storage in significantly less time and with less disruption than traditional methods.**



## Non-Disruptive Technology Refresh

ARX can also eliminate many of the management challenges around refreshing existing NetApp environments. When upgrading existing systems, the affected systems must be shut down for maintenance. In order to preserve data access and minimize business downtime, IT typically migrates all of the affected data off of the target device to a temporary location. After the device has been upgraded to a newer version of software or firmware, IT then migrates the data back to the original device. This process can be highly disruptive to users and applications that need access to the data, and burdensome for IT staff responsible for managing any disruption. By eliminating the disruption from data migrations, ARX can dramatically simplify the process of upgrading or refreshing existing file storage devices.


Organizations looking to replace older NetApp systems with the latest offerings can also use ARX to seamlessly refresh their storage environments. Because ARX hides physical changes from users and applications, organizations can deploy the latest NetApp systems alongside their older ones, migrate their existing data to the new systems, and remove the older systems, all without impacting business.

## Online Migration Among Multiple Types of Storage

ARX provides a sophisticated policy engine that automatically classifies and moves data to the appropriate type of storage based upon its business value. IT can create migration policies based on the age and type of data, as well as many other characteristics. For example, policies can control the automatic movement of data that has not been modified in 90 days, or the migration of certain project directories or types of files to capacity optimized storage, ensuring that high-performance disk is reserved for mission-critical applications.

These capabilities provide network storage professionals with a transparent tool to automatically move static or non-mission-critical data from high-performance tiers to denser tiers of storage. In turn, this helps reduce management costs and streamline storage operations. (ARX also offers substantial relief to backup processes, which will be discussed in more detail in the sections below.) Automated policy-based data movement can also be used to implement file-level tiering capabilities in NetApp storage environments, as well as to selectively migrate certain applications, projects, or types of data that require the performance of NetApp storage systems to those systems in a non-disruptive manner.





**Figure 4: ARX can automatically migrate data among multiple types of storage based on policy.**

## Optimized NetApp Deduplication

The ARX device's automated policy engine can also be used to optimize the utilization of data deduplication in NetApp environments. Policies can be established to identify and automatically move data sets best suited for deduplication to volumes that have enabled NetApp deduplication capability. One consistent concern about implementing deduplication on primary storage volumes has been the perception that it might adversely impact performance when accessing mission-critical data. However, with ARX, movement of data to deduplicated volumes can be done while retaining mission-critical data on volumes non-deduplicated volumes. This maximizes deduplication efficiency and ensures that data sets that receive little or no benefit from deduplication are not subject to the overhead of the deduplication process. ARX enables organizations to realize the full storage efficiency benefits of deduplication, while minimizing any potential performance impact.

## Optimized Backup and Recovery

Data growth has fueled many challenges associated with backup, most notably unacceptably long backup timeframes and rising costs. In most environments, inactive data composes the majority of total data volume. The management of, and cost impact to, organizations backing up the same unchanged files weekly is substantial.

Using the ARX data management capabilities, organizations can establish policies that automatically move unchanged data to a separate location, where they can be backed up at less frequent intervals. The small percentage of regularly modified data



can be backed up weekly, while the majority of data that has not been modified can be backed up monthly or even quarterly. In either case, data movement remains transparent to applications and users and can be completed without administrator intervention. This reduces backup times, as well as backup infrastructure and media costs by up to 80 percent.

In addition, most NetApp file storage environments rely on snapshot technology as the first line of data protection. ARX integrates with NetApp snapshot technology to present users with a single virtual snapshot history across the environment. Virtual snapshots extend to other legacy systems to seamlessly integrate new NetApp devices into heterogeneous storage environments. ARX coordinates and federates snapshots simultaneously across NetApp and legacy platforms, maintaining consistency and facilitating simple and rapid recovery of individual files.

## Storage Pooling and Aggregation

In order to streamline their data management tasks, many organizations require a large single workspace for their users and applications. However, large workspaces introduce several challenges and risks in traditional storage environments, such as bottlenecked performance and long backup times.

ARX can help eliminate many of these concerns by presenting a single virtual namespace to clients that are composed of multiple physical storage resources. Using ARX, multiple file volumes can be federated into a single virtual volume or file system. This aggregates the capacity, throughput, and processing power of the storage resources without complicating user and application access.

Using ARX to virtualize the storage environment can dramatically simplify the provisioning or reprovisioning of additional storage capacity. Today, increasing existing storage volumes is a disruptive process that requires the migration of data to a new physical location. ARX enables new storage volumes to be provisioned into the existing virtual namespace and automatically added to capacity-balancing policies.

Capacity balancing with ARX can also bring significant benefits in data protection. By reducing the size of each physical volume, backup and recovery can be performed in less time. Multiple physical volumes can be backed up in parallel, further accelerating backup performance, reducing backup time, and speeding individual volume restores.

## Conclusion

Many of the challenges faced by today's IT organizations come from the disruption incurred when moving data. Data migrations, NAS consolidations, and platform upgrades take longer and require higher IT overhead. The difficulty of separating active and inactive data means storage efficiency features such as data deduplication are underused, while backups take longer and require more tape media. Additionally, uneven growth leads to constant storage reprovisioning and capacity balancing—tasks that require operational downtime.

F5 and NetApp are working together to help organizations create an agile storage infrastructure that will better respond to rapid data growth. F5 ARX solutions enable transparent data migration to and among NetApp storage systems, so data can be moved without affecting applications or users or requiring client reconfiguration. By enabling data movement to take place automatically and without disruption, the joint solutions offered by F5 and NetApp help organizations address many of the key issues around data migrations and storage provisioning.

**F5 Networks, Inc.** 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 [www.f5.com](http://www.f5.com)

F5 Networks, Inc.  
Corporate Headquarters  
[info@f5.com](mailto:info@f5.com)

F5 Networks  
Asia-Pacific  
[apacinfo@f5.com](mailto:apacinfo@f5.com)

F5 Networks Ltd.  
Europe/Middle-East/Africa  
[emeainfo@f5.com](mailto:emeainfo@f5.com)

F5 Networks  
Japan K.K.  
[f5j-info@f5.com](mailto:f5j-info@f5.com)

