



F5 White Paper

Simplifying Disk-Based Backup and Restore with NetApp and F5 ARX

F5® ARX® solutions complement NetApp SnapMirror and SnapVault technologies to provide simple and effective data protection across virtualized file storage environments.

by Matt Quill

Principal Solutions Engineer

Renny Shen

Product Marketing Manager



Contents

Introduction	3
<hr/>	
Data Protection with NetApp	3
The Limitations of Traditional Backup	3
Deduplication Virtual Tape Library Solves Only Part of the Problem	4
Data Protection with NetApp	4
<hr/>	
Complementing NetApp Data Protection with ARX	5
File Virtualization	5
Virtual Snapshots	7
Snapshot-Only Shares	8
<hr/>	
Advantages of the F5 Approach	9
Simplified Federated Snapshots	9
Streamlined Offsite Data Protection	9
Deploy NetApp in Heterogeneous Storage Environments	10
Enable Streamlined Disk-Based Backup and Recovery	10
<hr/>	
Conclusion	10



Introduction

NetApp is at the forefront of streamlined disk-based data protection with its SnapMirror and SnapVault software technologies. SnapMirror enables synchronous or asynchronous replication of active volumes and Qtrees. SnapVault extends this capability by combining the point-in-time consistency of snapshots, versioning, and transport of snapshots to secondary or tertiary tiers of storage. With SnapVault, you can move older snapshots from mission-critical, high-cost storage controllers to lower-cost storage controllers.

The F5 ARX platform is the market leader in file virtualization technology. Similar to technologies that virtualize block-based storage, the ARX platform can be deployed in heterogeneous storage environments, and can seamlessly migrate and tier between storage platforms. The powerful ARX policy engine can move files to appropriate tiers of storage, inline and with no user disruption. It can also optimize primary storage, dramatically reduce backup windows, and employ the best use of deduplication technologies.

This paper will explain in detail how file storage professionals can combine the seamless data mobility of the ARX platform with the reliable protection of NetApp SnapMirror and SnapVault.

Data Protection with NetApp

The Limitations of Traditional Backup

Enterprises continue to struggle with large volumes of backup and recovery data. To negate risk, most organizations' strategies include combining nightly incremental backups with weekly full backups. As the volume of data continually increases, traditional network or SAN-based full backup schedules are becoming increasingly difficult to complete. Backup operations staff struggle with overabundant tape media and drive and cartridge failures, and by handling of offsite tape media, they incur significant management costs and risk. In addition to extended windows for backups, full recovery of large data volumes is also problematic in that recovery of data takes as long if not longer than backup. Although crucial for long-term and tertiary backup copies, tape infrastructures become increasingly unwieldy and difficult to manage as the amount of data grows.



Limited Recovery Point Objectives (RPOs) are another drawback of traditional backup methods. Given that even incremental backups can take hours to complete on large NAS platforms, tape backups typically occur once per day at most. With this arrangement, you can lose up to 24 hours of data in the event of a failure. For many organizations, losing a full day of data can be near catastrophic. Implementing a disk-based backup strategy is essential to ensuring the reliable backup and restore of mission-critical data. However, not all disk-based backup strategies are created equal.

Deduplication Virtual Tape Library Solves Only Part of the Problem

Deduplication technologies are found in two separate spaces in the storage environment. The first and most common location is in the area of backup and recovery. The second is in the primary and secondary storage space. The benefits of applying a disk-based deduplication appliance in the backup infrastructure are fairly clear: in sending largely identical data sets repeatedly to a disk target, elimination of redundancies will allow storage administrators to store backup data on disk at or near the cost of tape. When backup data resides on disk, backups and recoveries are typically faster and more reliable. Backup administrators can retain weeks or even months of data on disk and service most recovery operations from disk. This removes one of the largest challenges to maintaining a reliable backup and recovery environment—tape failure. Tape drives and media are frequently unreliable for restoration purposes because single backup copies can degrade or fail.

However, deduplication appliances still rely on network- or SAN-based backups and derive their greatest efficiency benefits from full backups. This means they do not substantially reduce the time required to complete a backup; rather the volume of backup data remains the same and organizations can realize only modest improvements in RTO/RPOs.

Data Protection with NetApp

NetApp snapshot technology greatly reduces backup time. NetApp snapshots present a full point-in-time image of your file systems, but without the overhead of maintaining full tape- or Virtual Tape Library-based backups. Only blocks that have changed between snapshot intervals are backed up, ensuring that the backup is completed in minutes rather than hours or days. Snapshots can be created in less than a second and scheduled for more frequent intervals during the day to improve RPOs.



Combining NetApp point-in-time snapshots with SnapVault provides the most efficient disk-based data protection on the market today. With SnapVault, organizations can retain more snapshots and for longer periods of time by offloading older snapshots to lower-cost secondary storage, either within the data center or offsite via the WAN. Because only the changed blocks are backed up with each snapshot, the SnapVault process is very efficient and can be substantially faster than even an incremental backup. On the designated secondary storage, you can enable deduplication to further optimize storage and retain more snapshots. More than 250 snapshots can be retained at a time, and you can increase the frequency of backup operations accordingly.

Combining the speed of snapshot-based data protection with the convenience of full backup images and the efficiency of NetApp SnapVault technology ensures you have full backup copies of data, but in a fraction of the time it takes to create tape-based full backups.

Complementing NetApp Data Protection with ARX

F5 ARX solutions complement the capabilities of NetApp Snapshot and SnapVault in virtualized file storage environments. ARX file virtualization can help organizations simplify data restoration by coordinating and federating snapshots across multiple NetApp storage systems and volumes, including those replicated to secondary storage systems through SnapVault.

File Virtualization

File virtualization abstracts the physical storage environment, decoupling the logical access to files from their physical location. Often referred to as a Global Namespace, ARX presents a federated view of the underlying file systems that simplifies access to file data and hides physical changes from clients. This simplification allows you to schedule migrations, tiering, and load balancing without downtime. Clients now access their file data logically through the Global Namespace instead of connecting to physical storage resources.

To understand how file virtualization works, consider Figure 1, which depicts the data flow without virtualization. In this example, when the client requests file1.doc,



NAS-1 sends the file to the client. Similarly, when the client requests file2.doc, FS-1 fulfills this request. The client has two separate drive mappings to the two network shares.

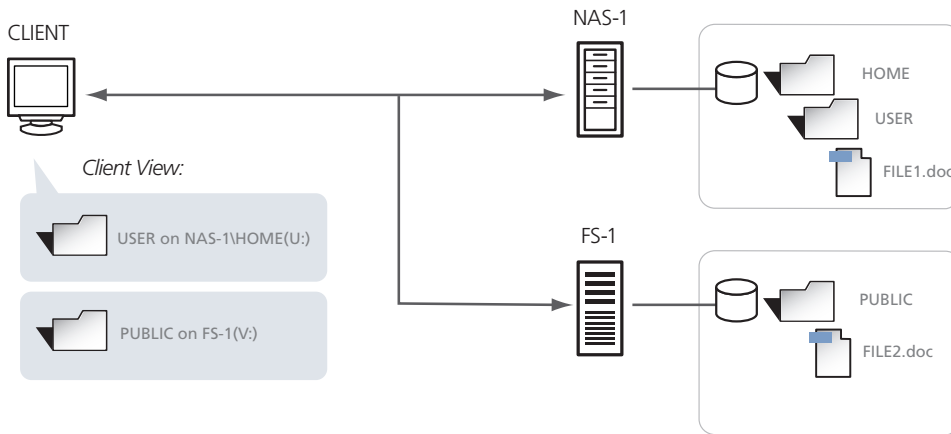


Figure 1: File access without virtualization

With file virtualization, the data flow is concealed from the user as shown in Figure 2. When the client requests file1.doc, the ARX device forwards that request to NAS-1. The response from NAS-1 is then forwarded to the client. When the client requests file2.doc, the ARX device forwards that request to FS-1. The response from FS-1 is then forwarded to the client. From the client perspective, both files appear to exist in the same location. The client maintains one drive mapping, which maps to the Global Namespace.

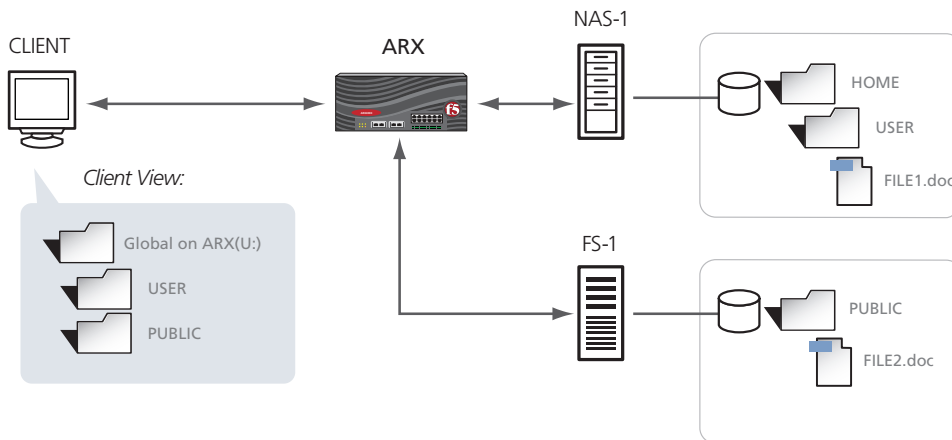


Figure 2: File access with virtualization



Virtual Snapshots

ARX can simplify the management of NetApp snapshots in virtualized environments through its virtual snapshot capability. With virtual snapshots, ARX coordinates the scheduling of physical snapshots for all of the underlying file systems in the Global Namespace. It then presents a federated view of the physical snapshots, as shown in Figure 3, to simplify the process of restoring either individual files or entire sets of files.

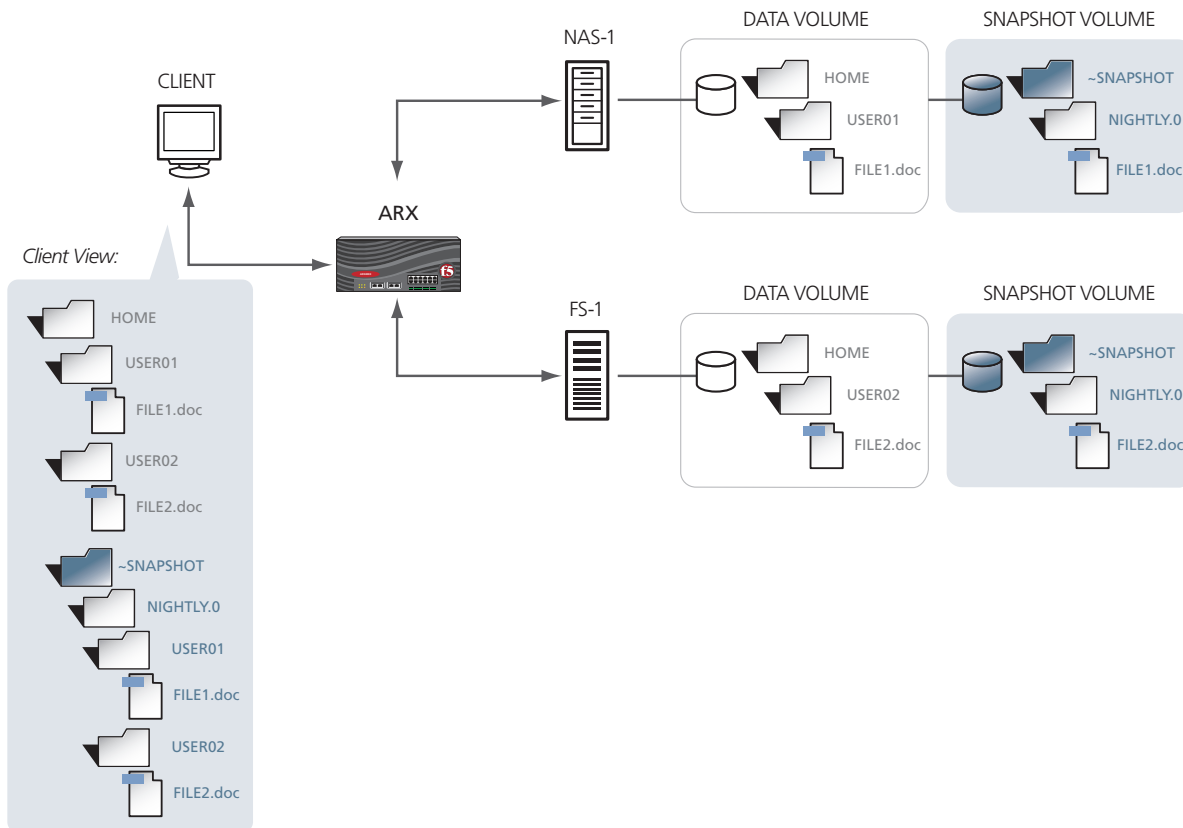


Figure 3: Virtual snapshots

Figure 3 demonstrates the power of the ARX virtual snapshot federation. In this illustration of consolidation from legacy Windows file servers to NetApp, snapshots from both platforms can be simultaneously coordinated and made consistent. As data is transitioned to NetApp from Windows, ARX maintains previous versions on Windows to protect against failure events or facilitate rollback if needed. Users retain a consistent view of their data via Previous Versions in Windows, which makes recovering older snapshots seamless and ensures that large-scale consolidations onto NetApp file storage are secure and protected.



Once consolidations are complete, you can coordinate snapshots across tiered NetApp file storage. For example, you can coordinate snapshots of a volume on a high-end NetApp controller with a snapshot of a volume on a lower-end NetApp controller. This ensures that ARX tiering does not adversely affect snapshots and that you maintain a consistent view across volumes.

Snapshot-Only Shares

SnapVault combines NetApp snapshot technology with replication. These technologies enable disk-based protection with efficient utilization. Transferring additional snapshots from higher-cost storage controllers to lower-cost storage controllers is seamless, because SnapVault only copies the changed blocks, which synthesizes full image capabilities with the efficiencies of incremental block level copies. ARX then enables you to seamlessly map these copies into the user shares. The snapshot-only shares functionality creates a non-metadata share view so you can extend previous versions of snapshots not only across storage platforms, but also as offsite SnapVault copies. People can now access share data from offsite snapshots as well as local snapshots, which simplifies recovery operations.

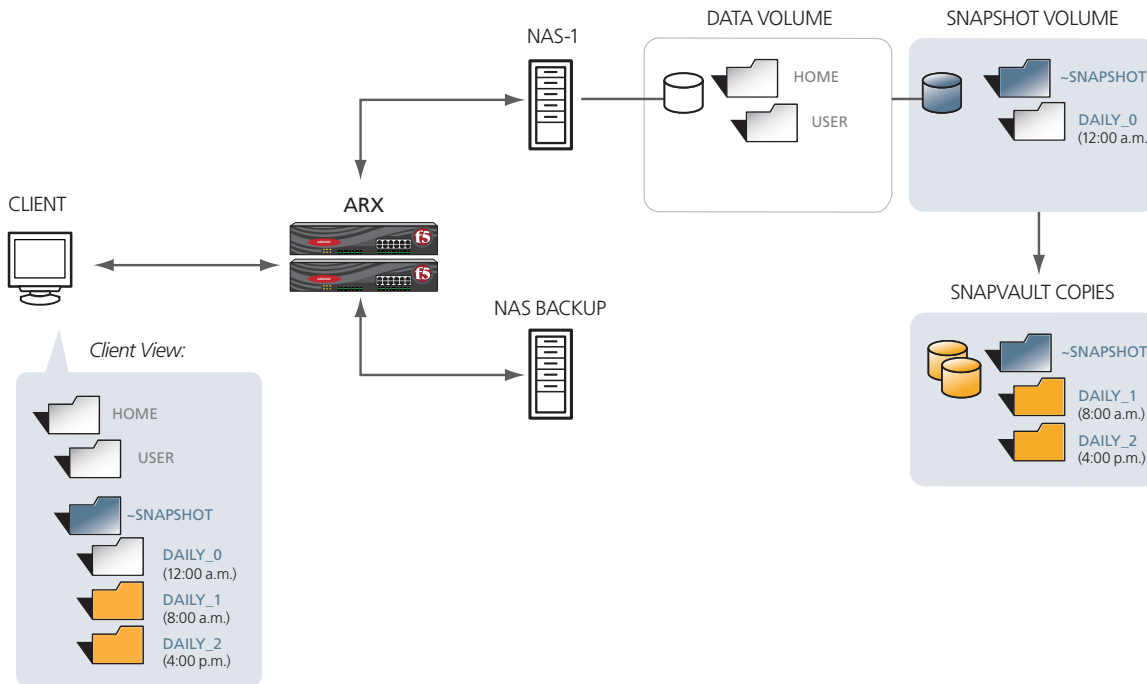


Figure 4: Snapshot-only shares



ARX virtual managed volumes now enable file storage administrators to manage snapshot-only shares in addition to actively managed shares. Once incorporated into the managed volume, users can browse and recover these offsite shares which simplifies recovery operations. Extending the promise and benefits of file virtualization to the offsite or secondary snap repository also extends the value of and interoperability with NetApp storage platforms. In Figure 4, both local daily snapshots and snapshots on the backup or offsite device can be presented via ARX. Users can now browse and restore their own backup information not only from the source NetApp controller, but also from the backup device. This functionality can be extended to Qtree SnapMirror as well for offsite or disaster recovery operations.

Advantages of the F5 Approach

Simplified Federated Snapshots

Administrators and clients can now move data across NetApp platforms, create tiers, and create a federated Global Namespace while maintaining snapshot consistency. This ensures that presentation of on-disk snapshots to clients is consistent across storage controllers and volumes, and it maintains a consistent presentation for user-initiated restores. By extending this functionality to backup or offsite storage via SnapVault or Qtree SnapMirror integration, you ensure that users can retrieve backup copies of snapshots across platforms. Organizations can create a consistent recovery point across the entire environment to minimize risk and reduce reliance on tape infrastructure for recovery.

Streamlined Offsite Data Protection

The ability to move snapshot copies to backup tiers of storage within the data center can also be extended by integrating ARX and SnapVault and Qtree SnapMirror snapshots. You can not only transport snapshots to a backup tier, but you can also replicate them to a second location. Snapshots across locations are presented for recovery through ARX via the Snap-Only Shares. Administrators and users can recover data from weeks- or even months-old snapshots. This ensures that NetApp snapshot, backup, and replication technologies are effectively coordinated across storage controllers as well as locations.

White Paper

Simplifying Disk-Based Backup and Restore with NetApp and F5 ARX

Deploy NetApp in Heterogeneous Storage Environments

Using snapshot federation, you can introduce NetApp file storage into legacy storage environments to effect large-scale consolidation or tiering. You can coordinate consistent snapshots between old and new file storage to ensure that migrations and tiering do not affect snapshots. For example, you can set up a large-scale Windows file server consolidation to move nightly snapshots to a NetApp backup target while a migration is occurring to new NetApp storage. In the event of a recovery request, users can browse snapshots from both the old Windows platform and the new NetApp platform for recovery. Clients can continue to retrieve file information from a snapshot regardless of where it resides.

Enable Streamlined Disk-Based Backup and Recovery

Using SnapVault gives you full backup capability on disk—and it's usually faster than performing a traditional incremental backup. Applying deduplication to the backup NetApp storage extends these efficiencies by reducing the amount of data stored on disk. While you can still schedule traditional tape backups for offsite or compliance purposes, eliminating the expense and overhead of deduplication Virtual Tape Library backups minimizes your reliance on tape for recovery. Integration with the ARX Snap-Only Shares switch ensures that you can undertake heterogeneous tiering and file storage consolidation projects without affecting the backup environment.

Conclusion

The ARX platform was designed to interoperate seamlessly with NetApp file storage. F5 and NetApp made significant joint development efforts to coordinate snapshots across controllers and platforms to ensure reliable and transparent migrations and tiering. F5 has extended this functionality to backups and offsite data protection with SnapVault snapshot-only shares. Customers now have the power to retrieve backup data from secondary platforms or offsite copies. This extends the full promise of ARX technology in NetApp environments to ensure full coordination of all backup procedures across controllers to substantially reduce backup and recovery management costs.

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com

