



White Paper

Boosting Operational Efficiency by Streamlining File Storage Management

Managing file storage consumes resources, causes downtime, and places an added burden on administrators. Virtualizing the file storage environment can dramatically simplify day-to-day management, increasing efficiency and enabling administrators to focus on other business-critical initiatives.

by Janice Merk
Marketing Manager



Contents

Introduction	3
<hr/>	
Inefficiencies of File Storage Management	3
Scenario 1: A Data Migration Causes Downtime.	4
Scenario 2: A Data Migration Entails a Major Software Development Effort.	4
Scenario 3: Ad Hoc Data Migrations Become a Ritual.	4
Scenario 4: Endless Deletion of Files Is Necessary to Stay within Capacity Limits.	5
Scenario 5: Backups Are Unnecessarily Long and Cumbersome.	5
Scenario 6: PST Files Consume Valuable Storage in User Home Directories.	5
<hr/>	
An Intelligent Approach to File Storage Efficiency	6
File Management Efficiencies of F5 ARX	7
<hr/>	
Conclusion	10



Introduction

IT storage administrators often face an overwhelming challenge as they respond to demands placed on them from all directions—from users, the IT manager, and the business itself. Planned and unplanned outages foist a public relations role on administrators, who have to explain why data is unavailable and why users can't do their jobs. Administrators also play the role of enforcer, continuously monitoring storage and ensuring that users do not surpass their allocated limits. The storage administrator is expected to keep storage costs low and performance high, while striving to meet quarterly, annual, and strategic technology goals and business objectives.

Many of the demands faced by storage administrators are made more challenging by the manual and labor-intensive nature of storage management. Frequent tasks such as a vendor patch or a data backup can be disruptive to users and require administrators to work during non-standard hours. In addition, many tasks, such as a data migration, require manual and intricate reconfiguration of systems. Disruption and complexity can easily overburden administrators by increasing the need for intensive manual interventions, leading to operational inefficiencies, wasted time, and frustration.

The F5® ARX® file virtualization system can help simplify the management of file storage environments by eliminating the disruption typically associated with moving file data. ARX employs a technology called file virtualization to decouple the logical access to files from their storage device. This provides the ability to move files where and when desired, without impacting users. ARX file virtualization can be applied in several common scenarios to simplify tasks and make the jobs of storage administrators easier.

Inefficiencies of File Storage Management

Limited automation and a lack of data mobility require storage administrators to spend inordinate time performing manual, complex, and time-consuming tasks. As the amount of storage increases, the efficiency of the IT administrator decreases.



Scenario 1: A Data Migration Causes Downtime.

Data is constantly on the move. In many organizations, frequent data migrations result in business downtime, with manual intervention required to untangle and fix static file mappings between users and their file storage resources. Before administrators can rebuild these mappings (or mount points) and the data can go live, additional work is needed to move the target data offline, copy it to a new directory, and update user permissions in profiles or login scripts. Storage administrators are often forced to choose between working at night or on weekends or bringing systems down during the day—to the dismay of users. The process is not only tedious and complex but extremely inconvenient for all parties involved. In addition, the potential for errors is introduced as user documents often contain universal naming convention (UNC) paths, object links, or shortcuts that reference the file systems being migrated. To ensure continued access, any movement of data must preserve all of these paths, a task which is especially challenging.

Scenario 2: A Data Migration Entails a Major Software Development Effort.

A data migration can be even more complicated when the mappings between storage directories and business applications are hard-coded and the application has to be rewritten. This can add significant overhead, as each additional software release requires new design, development, and test-and-deployment work. Additionally, the new software release may introduce bugs into the customer-facing production environment.

Scenario 3: Ad Hoc Data Migrations Become a Ritual.

Many organizations are constantly running into the limits of their available storage capacity. When this happens, they're forced to temporarily alleviate the problem by running scripts to manually migrate data from one file system to another, with the temporary nature of the solution making this process a ceaseless ritual. Another common approach is the comingling of data between different corporate departments. When the available file storage of one department reaches capacity, storage is borrowed from another department. When capacity limits are hit again, storage is borrowed from yet another department, and so on. This "reactionary" approach to storage management leaves data from multiple groups unnecessarily entwined and difficult to untangle. Ad hoc data migrations are inefficient, requiring hands-on attention and constant monitoring by the storage administrator.



Scenario 4: Endless Deletion of Files Is Necessary to Stay within Capacity Limits.

When organizations run out of capacity, storage administrators often scramble to find free space elsewhere in the environment. This could involve identifying snapshots or files that can be safely deleted. The problem is compounded when the target files are embedded in core business or operational processes that may be needed again in the future or for compliance reasons. In some organizations, files are required to undergo a complex process of encoding, test, and quality control before they can be shared and made available to users. When those files are deleted to respond to a storage limitation and then needed again later, the file may undergo the same tedious process, lowering productivity.

Scenario 5: Backups Are Unnecessarily Long and Cumbersome.

Data backups are necessary for a range of reasons, from protecting data assets and disaster recovery to meeting regulatory requirements. For the storage administrator, however, the backup process may be long and cumbersome. Backups take a lot of time and require monitoring, whether they are incremental backups that happen overnight or full backups to tape taking place on weekends. The backup will be incomplete if the administrator encounters interruptions, and the margin for error is slim.

Full backups become unwieldy because all data is backed up to tape each time, even if the data has not changed. Even incremental backups are unwieldy when a full scan of all data is required to sift out the new data each time. For many organizations, up to 80 percent of the data on tier 1 storage has not changed in 90 days, yet those files continue to be backed up or scanned over and over again, increasing the backup window and costs.

Scenario 6: PST Files Consume Valuable Storage in User Home Directories.

Many organizations face challenges managing Microsoft Outlook .pst or personal folder (PST) files. They are great for users, allowing them to archive crucial information such as emails, contacts, and calendar events, but PST files take up more than a moderate amount of storage capacity. This often prompts a search for and deletion of large or unnecessary PST files by the IT team. (See Scenario 4.)



Further, archived PST files that remain unmodified are rolled up into regularly scheduled backup intervals, exacerbating backup times. (See Scenario 5.)

The scenarios described above produce constant anxiety and frustration for IT administrators. Some of these scenarios are workarounds for common file data management problems, requiring additional hours and continuous rework. Others describe manual, error-prone processes that require constant monitoring. Whether files are being deleted to avoid a catastrophe or hard-coded file paths are updated in an application as part of a data migration, storage administrators battle never-ending problems.

An Intelligent Approach to File Storage Efficiency

Even with the cost of storage capacity declining year after year, storage administrators often struggle to keep up with the growth in data without the right tools. F5 Networks offers a solution that entails a smart, intelligent approach to managing dynamic storage needs in an automated and efficient manner using file virtualization.

The F5 ARX system employs file virtualization to provide a layer of intelligence in the network that decouples logical access to files from the storage devices. ARX devices federate all of the file systems and file storage devices within an environment and logically represent them through virtual file systems. The ARX devices then stand proxy for logical file access through the virtual file systems to the proper physical locations on storage devices.

The benefits of this approach are twofold:

- Because ARX manages the physical mappings between clients and storage resources, regardless of whether the client is a user or an application, it eliminates much of the complexity inherent in managing file storage environments.
- Because file virtualization creates a proxy for logical file access, storage administrators can now move files between different storage devices or file systems without disrupting client access. This reduces, or even eliminates, the need for scheduled downtime as well as the additional overhead of client reconfiguration.



File virtualization does not introduce a new file system. Access to the existing file systems continues to take place through industry-standard Common Internet File System (CIFS) and/or Network File Systems (NFS) protocols, so it's not necessary to replace existing assets or learn new file storage interfaces. With file virtualization, data is free to move and storage resources are free to change on demand. Administrators no longer need to make updates to client mapping or monitor scripts to ensure data has been successfully moved from one device to another. Updates to applications and login scripts to reflect client mapping modifications are eliminated from the process completely. Storage administrators can use the data mobility afforded by file virtualization to readily manage rapidly growing and changing files.

In addition, ARX provides administrators with data management policies that can automate many of their most frequent and tedious tasks. From data migration to storage tiering to capacity balancing, these policies govern the movement and placement of files based on changing environmental conditions. With automated policies, ARX can liberate storage administrators from having to constantly respond to their network attached storage (NAS) or file server limits, and from applying homemade scripts to move files to free storage.

File Management Efficiencies of F5 ARX

Administrators can solve large and dynamic file management problems with F5 ARX. By eliminating much of the complexity found in traditional file storage environments, ARX can end disruption and drive significant efficiencies in the management of file data and file storage. In each of the six common file management scenarios, administrators can benefit from file virtualization.

Scenario 1: Data migrations cause downtime.

With ARX, managing client and resources mapping, updating embedded hard links, and reconfiguring clients are no longer necessary. Because clients now access files through virtual file systems, migrating files between different storage devices or file systems no longer disrupts client access. Storage administrators can move data at will and even schedule a file migration in the middle of a business day, without needing to negotiate downtime. In addition, ARX preserves UNC paths, object links, and shortcuts when files are migrated, greatly reducing the possibility of error.

Scenario 2: A data migration entails a major development effort.

Because ARX file virtualization decouples the logical access to files from their physical locations, moving data no longer requires rewriting application code.



This has allowed many F5 customers to scale back the duration of planned migrations from several months to merely days or even hours. Better yet, there is no chance of introducing new application bugs into a customer-facing production environment.

Scenario 3: Ad hoc data migrations become a ritual.

To manage the constantly varying and unpredictable nature of storage demands, administrators can prescribe policies to enable real-time responses to sudden storage needs. Storage administrators no longer need to constantly monitor storage resources and respond with labor-intensive manual actions. When one department runs low on storage capacity, administrators can enable additional storage from anywhere in the environment instead of merging data from different departments. The supply and demand of resources are dynamically balanced in an approach also known as capacity balancing, so ad hoc migrations are no longer needed daily, weekly or monthly to keep storage demands in control.

Scenario 4: Endless deletion of files is necessary to stay within capacity limits.

Storage administrators no longer need to exercise the search and delete method. By implementing a storage tier strategy, organizations can automatically migrate the majority of their file data to lower cost storage, eliminating the need to constantly delete large or old files from high-cost storage. In addition, policies can be established to automatically place files, as they are created, on various storage tiers, depending upon the business value of the data.

Scenario 5: Backups are unnecessarily long and cumbersome.

With a storage tier strategy, storage administrators gain the tools to differentiate between modified and unmodified data, as well as between different types of data. Instead of unnecessarily backing up both the modified and unmodified data in a target file system, ARX can automatically place unmodified data (or specific file types) on a secondary tier or even archive it in the cloud. Since unmodified data often consumes up to 80 percent of available capacity, moving it to another file system and backing it up less frequently significantly speeds the weekly full backup and reduces the amount of tape required. Scan times for the nightly incremental backup are reduced as well, because the amount of data being scanned is dramatically less. Separate backup policies can be established for each tier. As a result, backups for tier 1 storage no longer need to take an entire weekend.



Scenario 6: PST files consume valuable storage in user home directories.

By creating a file placement policy to automatically place PST files on less expensive secondary tiers, administrators can reduce costs and apply a separate policy to back up PST files less often. Administrators no longer need to search for large PST files and convince users to delete them. Unchanging PST files no longer have to be backed up every week. PST files can be conveniently placed in the appropriate storage tier and backed up as needed. This solution is not limited to PST files, either, but can apply to video files and any other data types that take up a lot of space.

File virtualization simplifies the management of a storage infrastructure. With intelligent file virtualization, there is no need to search for files to delete, run scripts to move your files from one directory to another, or manually re-configure client to storage mapping. Intelligent file virtualization is a simple, transparent, automated, and dynamic file storage solution that makes storage management efficient. With file virtualization technology, an organization can manage its storage infrastructure without downtime for users, weekend work for the IT team, or tedious and error-prone manual workarounds for problems.

Scenario	Before	After
Data migrations cause downtime	Administrators manually perform migrations	Administrators schedule data migrations using automated policies
	Administrators perform data migrations during nights/weekends	Migrations can be completed during the business day
	Manual reconfiguration of client to resource mapping	Automated reconfiguration of client to resource mapping
	Broken user shortcuts and document links	Preservation of UNC paths, object links, and shortcuts
Data migration entails a major software development effort	The need to rewrite code causes longer data migration time	Faster data migrations, since applications need no modification
	Increased risk of bugs from rewriting code for business applications	Less risk and no updates to customer-facing production systems
Ad hoc data migrations become a ritual	Unpredictable storage usage creates reactive environment	Proactive capacity balancing and seamless storage provisioning ARX balances file placements based on changing demands
	Homegrown scripts must be constantly developed to alleviate problems	Seamless movement of files between storage devices or storage tiers, with spare capacity from the storage environment available when existing file system is full

White Paper

Boosting Operational Efficiency by Streamlining File Storage Management

Endless deletion of files is necessary to stay within capacity limits	Files are manually identified and deleted to increase storage units, but those files are required again in the future	Storage tiering automatically moves non-critical data from expensive storage resources to lower-cost alternatives
Backups are unnecessarily long and cumbersome	Weekly full backups exceed the backup window	Full backup times reduced by up to 90%
	Unchanging data is redundantly backed up, driving up backup times	Different backup policies applied to changing and unchanging data or different types of data
PST files consume valuable storage in user home directories	Need to search for moderately large files to delete or for compliance reasons	PST files placed automatically on separate physical file systems
	Unnecessary backup of PST files	Different backup policies applied based on file type

Figure 1: Characteristics of common file management scenarios before and after ARX implementation

Conclusion

Files play a large role in both personal and professional personas. From the MP3 and JPG files on mobile phones to an organization’s PST and log files, files have been embraced as a part of everyday life. This explosion of data is difficult for organizations and their storage administrators to manage without inconvenience and inefficiency. Either users face downtime or the storage team has to work off-hours to handle complex and resource-intensive tasks such as data migrations or backups.

F5 ARX file virtualization technology empowers storage administrators with the appropriate tools to efficiently manage storage resources. ARX devices automate what are currently manual storage management tasks and eliminate the downtime associated with these tasks. With ARX, storage administrators are equipped with an automated, transparent, and policy-driven tool to meet the challenges of file data management today.

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

