



F5 White Paper

# Creating Virtual Snapshots with F5 ARX

Storage snapshot technology is the foundation for modern data protection in the enterprise. The F5® ARX® file virtualization solution coordinates and federates snapshots across heterogeneous devices and volumes to simplify management and data protection in virtualized file storage environments.

**by Renny Shen**  
Product Marketing Manager



# Contents

<b>Introduction</b>	<b>3</b>
<hr/>	
<b>The Foundation: Snapshots</b>	<b>3</b>
Snapshot Policy	5
How It Works	6
Heterogeneous Environments	7
<hr/>	
<b>Virtualizing Snapshots with ARX</b>	<b>8</b>
Presentation	8
Configuring the Snapshot Rule	9
Creating a Virtual Snapshot	12
Reporting	13
<hr/>	
<b>Advantages of the F5 Approach</b>	<b>14</b>
Coordination	14
Crash Consistency	14
Other Advantages	16
<hr/>	
<b>Conclusion</b>	<b>17</b>



## Introduction

Snapshots, or point-in-time (PIT) copies, are becoming an essential technology for data protection in the enterprise. The risk associated with data loss has been steadily growing over the years. In response, companies have increasingly adopted snapshot technology as an integral part of their data protection strategy. Major storage vendors are taking advantage of this growing interest in snapshots, and they are increasingly incorporating this capability into their product offerings.

The proliferation of snapshot technologies has had an unintended consequence. In heterogeneous storage environments, companies are faced with the prospect of managing multiple proprietary snapshot technologies. This leads to a number of challenges:

- **Complexity.** Companies must maintain multiple mechanisms to perform the same task. Administrators must have knowledge of multiple proprietary systems with disparate interfaces and behavior.
- **Recoverability.** Data migrations typically result in the loss of historical snapshots from the client point of view. This is especially true when there is a change in storage vendor.
- **Crash consistency.** Because snapshots are created on a per-volume basis, it is difficult to coordinate snapshots across all file storage resources in order to create a crash-consistent image of the entire storage infrastructure.

Intelligent file virtualization is the obvious choice to manage snapshot technologies from multiple storage vendors in a heterogeneous storage environment.

F5® virtualization technology federates multiple disparate network-attached storage (NAS) devices and file servers into a single, unified Global Namespace so that users can access file data in a logical manner, regardless of where it physically resides. The virtual snapshot capability enables companies to manage their snapshots in the same manner, giving users the ability to recover single files from a logical presentation of the PIT copy or to back up a single global view of multiple file systems.

## The Foundation: Snapshots

A snapshot is a copy of a file system as it existed at a single point in time. This copy includes everything in the file system that was written to disk at the time that the snapshot was taken, including the directory tree structure, files, and file properties.



Transactions that were not written to disk, including any data cached on the client, are not present in the snapshot.

Because data is constantly changing, files preserved in a snapshot are quickly out of date. To provide data protection over time, companies integrate a snapshot policy in a two-tiered strategy. A typical snapshot policy will continuously create snapshot images at defined intervals and retain them for a period of time. For short-term data protection, snapshots are accessible and navigable by administrators and users. This is useful for recovering individual files that might have been lost, such as an older version of a file or an accidentally deleted file. For long-term data protection, companies create snapshots at a longer interval and use them as the source of a tape backup.

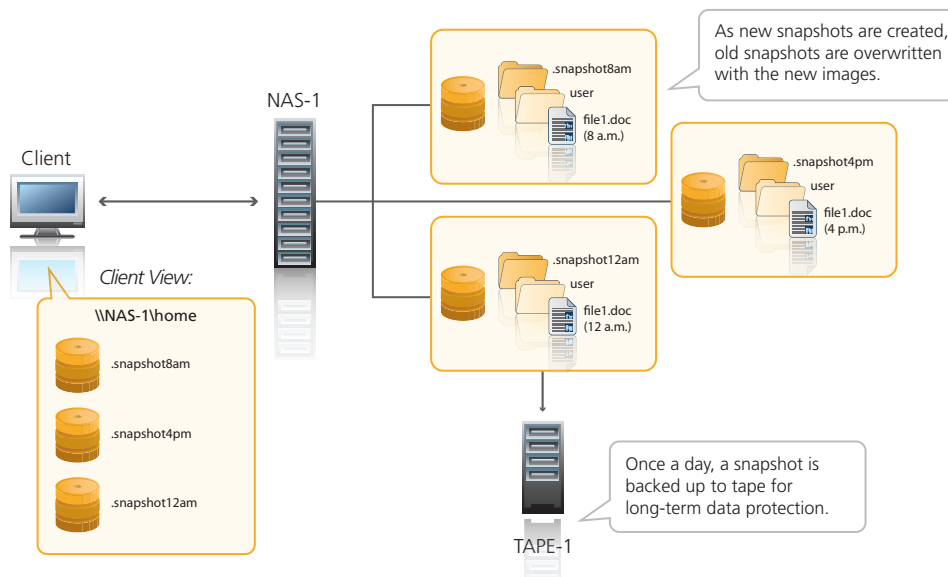


Figure 1: Example of snapshot policy

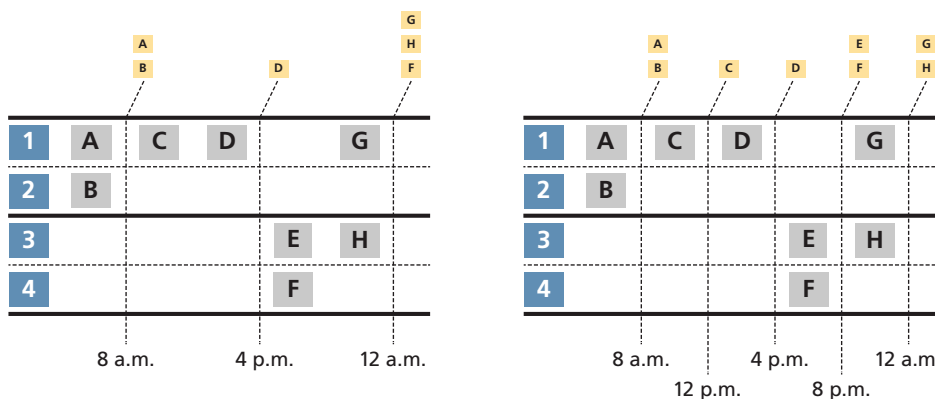
In Figure 1, the file storage device NAS-1 has been configured to create snapshots at eight-hour intervals beginning at 8 a.m., and to retain up to three snapshot images. An administrator or client can recover files from the snapshot images saved on disk by accessing .snapshot8am, .snapshot4pm, and .snapshot12am and navigating through the snapshot file system. In this example, file1.doc, which is modified and saved to disk several times throughout the day, is available for recovery in each of the snapshots. Once a day, the entire 12 a.m. snapshot is backed up for long-term data protection of the entire file system. The next day, the process begins again. At each interval, the image from the previous day is deleted and the space freed before a new snapshot is taken.



## Snapshot Policy

A snapshot policy can be used to implement one component of an organization's data protection strategy. A snapshot policy has two components that are related: period and retention. Period is the amount of time between successive snapshots. Retention dictates how long an individual snapshot is maintained. When determining snapshot policy, businesses must balance two opposing goals: the desire for an infinite number of recovery points versus the increase in storage capacity required to store those recovery points. In a changing file system, the longer a snapshot is retained the more capacity it will consume. Most implementations sacrifice recovery points in order to conserve storage capacity.

One of the key aspects of snapshot technology is that changes that occur between snapshots are not always captured. As a matter of design, only the state of the file system when the snapshot was created is preserved. If a file is modified multiple times between snapshots, the only versions that are preserved are the ones that existed when a snapshot was created. This behavior is illustrated in Figures 2a and 2b.



Figures 2a and 2b: Examples of snapshot policy

In Figure 2a, there are four files in the file system: 1, 2, 3, and 4. This example shows a snapshot policy with an interval of eight hours and retention of three. In the first snapshot at 8 a.m., the image contains changes to files 1 (A) and 2 (B). Between 8 a.m. and 4 p.m., two changes occur to file 1: (C) and (D). However, only the latest change (D) is preserved in the snapshot. Likewise, the final snapshot at 12 a.m. contains changes to file 1 (G), file 3 (H), and file 4 (F), but does not contain the earlier change to file 3 (E).

Figure 2b shows a snapshot policy with an interval of four hours and retention of five. The first snapshot at 8 a.m. is identical to the one in Figure 2a. However,



the addition of another snapshot at 12 p.m. captures the second change to file 1 (C) that was not captured in Figure 2a due to the longer interval between snapshots. Likewise, the snapshot at 8 p.m. captures the first change to file 3 (E) that is not captured with a larger interval.

The granularity of a snapshot policy is a function of its period and retention. A policy will schedule snapshots to be created at the intervals dictated by the period and retain sufficient snapshot images to provide the desired recovery points.

## How It Works

To create snapshots, vendors will use several techniques: copy-on-write, redirect-on-write, and split mirror.

### Copy-on-write

With this technique, storage is allocated to contain data changes. No data is actually copied at this time. Instead, the snapshot is populated with data changes as they occur. As shown in Figure 3, when a client issues a write request to disk, the original data is copied to the snapshot area before the new data is actually written. A read request to the active file system will proceed as normal, whereas a read request of the snapshot will go to the snapshot. If the data has changed from the original copy, the request will be served by the snapshot copy. If no changes have been recorded, it will be redirected to the original copy. Because copy-on-write does not copy data until it is modified, snapshots are created instantaneously and are very space-efficient.



**Figures 3a and 3b: Examples of copy-on-write**



## Redirect-on-write

With this technique, a snapshot copy is initially created by allocating storage space on the disk to contain new data changes. No data is actually copied at this time. Unlike with copy-on-write, write requests with redirect-on-write are directly issued to the new allocated space, and the original copy becomes the snapshot copy, as shown in Figure 4. Redirect-on-write has a performance advantage over copy-on-write, as only one write occurs with each write request, instead of two. Like copy-on-write, redirect-on-write does not copy data; therefore, snapshots are created instantaneously and are very space-efficient.




Figure 4: Example of redirect-on-write

## Split mirror

With this technique, a snapshot copy is created by cloning the original copy so it becomes the snapshot copy. Because of the initial data copy, creating split mirror copies is not instantaneous. Also, because each copy is a full replica of the original volume, split mirror copies are not space-efficient.

## Heterogeneous Environments

Because storage vendors create snapshots at the storage-device level, it is difficult to integrate snapshots in heterogeneous environments. Each vendor creates snapshots using proprietary techniques, and there is no mechanism to seamlessly tie together snapshots from multiple vendors. Instead, customers are left with individual pieces of the overall picture that they must manage separately and cannot put together.

Figure 5 illustrates a typical customer experience with snapshots in a heterogeneous environment. In this example, user home directories are split across two storage



devices from different vendors, NAS-1 and FS-1. Each device has been configured to create snapshots on a daily basis. Clients see all snapshots as distinct volumes in this environment. Scheduling consistent snapshots across the multiple home directory locations requires coordinating schedules on multiple devices.

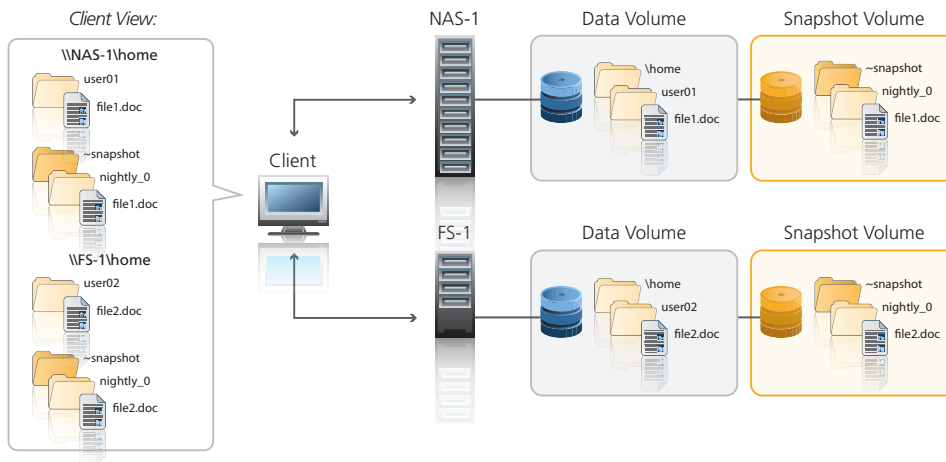


Figure 5: Snapshots without virtualization

## Virtualizing Snapshots with ARX

### Presentation

The F5 ARX file virtualization device can simplify the management of snapshots in a heterogeneous environment by federating multiple proprietary snapshots into a single virtual snapshot presentation. ARX devices virtualize snapshots on a per-managed-volume basis. Within a managed volume, the ARX device coordinates a physical snapshot creation operation on all backing file storage resources. When this is completed, the device aggregates them into a single virtual snapshot for presentation to clients.

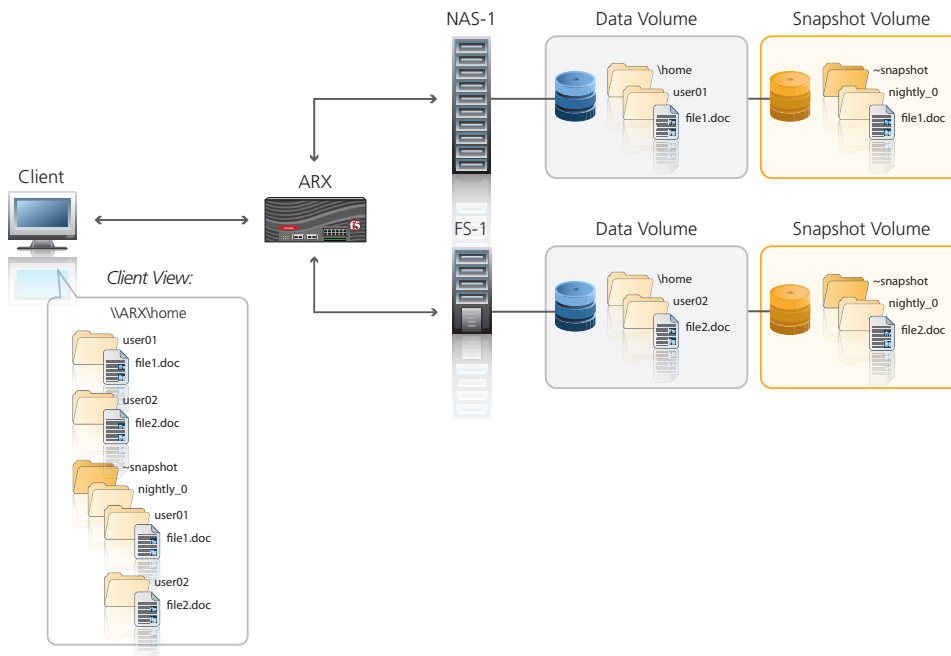


Figure 6: Snapshots with virtualization

Figure 6 illustrates this in an example in which a customer has two storage vendors or platforms and two different snapshot technologies. Because a virtual snapshot represents the state of a specific managed volume at a specific point in time, the ARX device creates the virtual snapshot logically within the managed volume. In this example, the managed volume \home contains physical volumes on two file storage devices: \home on NAS-1 and \home on FS-1. At a scheduled interval synchronized across both devices, the ARX device directs both NAS-1 and FS-1 to create a physical snapshot of their \home physical volume. The ARX device then aggregates the two separate physical snapshots into a single virtual snapshot that it presents to clients. Instead of two file systems, clients looking at the virtual snapshot see a single namespace with files from both storage devices. By default, the ARX device presents virtual snapshots in the ~snapshot directory in the root directory of the managed volume. Although administrators can choose any presentation schema, by default each snapshot has its own directory within ~snapshot named [rule name]\_[enumeration].

## Configuring the Snapshot Rule

Virtual snapshot creation is governed by snapshot rules. Before a snapshot can be created, an administrator must create a rule for each managed volume to be protected. Rules enforce the corporate data protection strategy and provide



companies with flexibility to balance the tradeoff between data protection and cost containment, according to their priorities. A rule applies only to the managed volume in which it was created, and a managed volume can have multiple rules simultaneously.

Snapshot rules have the following attributes:

- **Name.** Every rule has a name that is invoked when instantiating the snapshot rule. The rule name should be descriptive of the policy enforced and is also the name of the directory under which the snapshot is presented in the managed volume. For example, in Figure 6, the rule creating a nightly virtual snapshot is named “nightly.” Snapshots created by this rule are presented in the ~snapshot directory with the names nightly\_XXX.
- **Schedule.** Virtual snapshots must be scheduled before they will be created. A schedule specifies both the starting time and the interval between snapshots. As an example, companies wishing to provide three snapshots per day can schedule snapshots at eight-hour intervals starting at midnight. Scheduling is applied per rule across all shares and exports within the managed volume. This means that administrators can manage their snapshot scheduling centrally within the ARX device instead of across multiple storage devices.
- **Retention.** Every snapshot consumes storage space. The rule specifies the number of snapshots to retain before deleting the oldest image. The snapshot images in the ~snapshot directory are enumerated with 0 to [# retained – 1], where 0 always enumerates the most recent snapshot. Retention is applied per rule.
- **Reporting.** A rule can be enabled to generate reports recording the execution of a snapshot rule.

Figure 7 illustrates an example with a snapshot rule named “daily” that protects the contents of the \home managed volume. An administrator scheduled the “daily” rule to run every eight hours, starting at 8 a.m., and configured it to retain three snapshots to provide one day of rolling snapshot coverage. Snapshots are labeled daily\_XXX and can be found in the ~snapshot directory in the root directory of the managed volume. At 12 a.m., the rule has not yet been initiated and clients see an empty ~snapshot directory. The first snapshot is created at 8 a.m., labeled daily\_0. The second snapshot is created eight hours later at 4 p.m. The 8 a.m. snapshot, previously labeled daily\_0, has been renamed daily\_1, and the new 4 p.m. snapshot has been labeled daily\_0. The third snapshot is created at 12 a.m.



on the second day. With each new snapshot creation, the daily\_2 will be deleted, daily\_1 will be renamed daily\_2, and the new snapshot will be created under daily\_0.

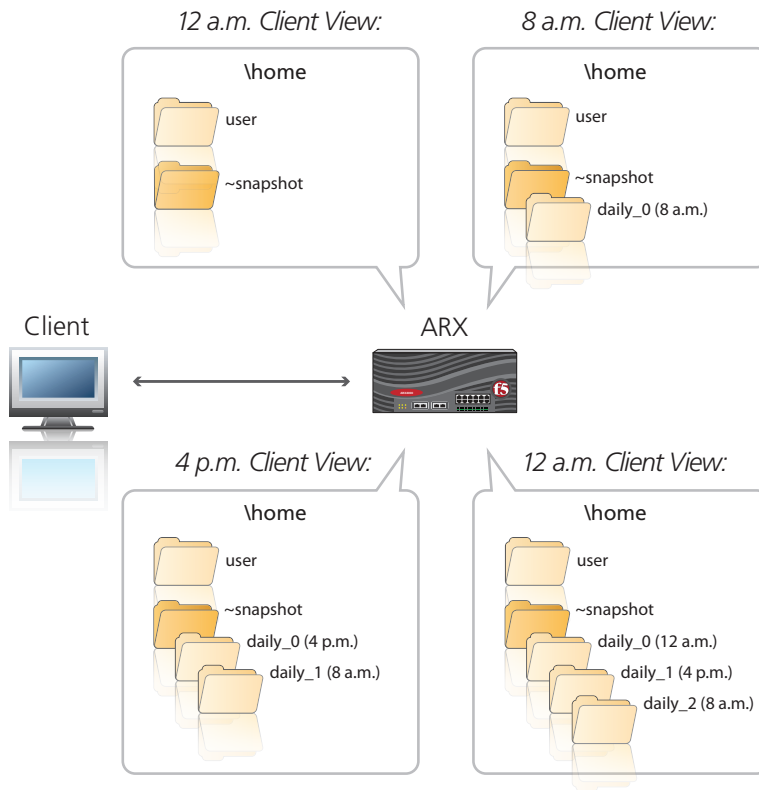



Figure 7: Instantiating a snapshot rule



## Creating a Virtual Snapshot

Figure 8 details the process in which the ARX device creates a virtual snapshot.



**Figure 8: Creating a virtual snapshot**

1. The virtual snapshot is the presentation of a managed volume at a particular point in time. The ARX device does not actually store any user data, but rather presents it in a logical fashion. To be accurate, the presentation must include all of the backing physical volumes contained in the managed volume.
2. Optionally, input/output (I/O) operations can be atomically fenced to the managed volume in order to ensure that the snapshot is crash consistent. The ability to fence file system I/O is most useful when the NAS devices are used to store application data. Because the ARX device is in line with



the managed volume, it is ideally situated to freeze file access to the target managed volume upon initiation of the virtual snapshot creation.

3. The ARX device sends a request to each storage device in the managed volume to create a physical snapshot. This is done in parallel to all storage devices to minimize the total time of operation. In addition, this operation is protected by a configurable timeout to minimize client impact. If any physical snapshot creation operation takes too long, the ARX device will timeout and roll back the snapshot creation process. This process typically takes one to a few seconds.
4. When every storage device has completed its physical snapshot creation operation, the ARX device aggregates all of the physical snapshots into a crash-consistent virtual snapshot.
5. If the I/O fence was enabled, the ARX device now removes the I/O fence and enables all delayed write requests to proceed to the managed volume.
6. In order to route read access to the contents of the virtual snapshot, the ARX device updates its internal snapshot metadata. At this point, it can present the virtual snapshot in the ~snapshot directory in the root directory of the managed volume.
7. The ARX device creates a snapshot report detailing the virtual snapshot creation operation, including the mapping from the virtual share to the backing volume.

## Reporting

ARX devices have extensive reporting capabilities, enabling administrators to monitor snapshot creation, verification, and removal. These reports can be created in either XML or CSV output formats. This data can be used to verify compliance of snapshots created against the corporate snapshot policy, including checking if the file server has deleted or reclaimed a physical snapshot within the virtual snapshot.



# Advantages of the F5 Approach


## Coordination

The ARX device's virtual snapshot capability minimizes the complexity of managing the snapshot process across multiple file servers. Today, administrators must manage snapshots on each file server independently, even if they are from the same vendor or platform. This task is even more daunting if the file servers are from different vendors or platforms in a heterogeneous environment. Now, administrators must have specialized knowledge of each snapshot technology.

With ARX, administrators no longer have to manage snapshots on a file-server-by-file-server basis. Instead they can manage snapshot policy holistically, without regard to the underlying physical storage. The ARX device's virtual snapshot capability coordinates snapshots from all network file shares on a managed volume basis. This makes the snapshot process more reliable and less costly to administer. In addition, ARX abstracts the interfaces to proprietary snapshot technologies, which enables companies to deploy or retain heterogeneous storage environments without complicating their data protection strategy.


## Crash Consistency

A major advantage of the ARX device's virtual snapshot capability is the ability to inherently provide crash consistency for snapshots in heterogeneous and multiple file server environments. F5's approach preserves write-ordering within a managed volume, including those that comprise multiple physical volumes, file servers, and even storage vendors. Crash consistency also handles "torn writes," meaning that a write is either complete in the snapshot image or not present at all.



**Figures 9a and 9b: Examples of crash consistency with snapshots**

Figures 9a and 9b demonstrate the need for crash consistency in snapshots across multiple physical volumes. In Figure 9a, a single application writes data to two physical volumes. Without a mechanism to provide crash consistency across both volumes, the resulting snapshot for files 1 and 2 can be out of sync with the snapshot for files 3 and 4. For example, the application simultaneously writes related data to files 1 (A) and 3 (B). However, because the snapshots are out of sync, the administrator can only recover the change to file 1 (A). The closest change recoverable for file 3 is (C). This could be unacceptable depending on the relationship between changes (A) and (B). If change (A) does not make sense without the context of (B), then the administrator cannot re-create the environment that existed at 8 a.m. Figure 9b shows a virtual snapshot with crash consistency across both physical volumes. Here, the 8 a.m. snapshot captures the related changes to files 1 (A) and 3 (B). Because of the inherent crash consistency, an administrator can recover the entire managed volume at a recovery point that is consistent with the application.



Figures 10a and 10b: Examples of crash consistency with ILM solutions

Figures 10a and 10b demonstrate the need for crash consistency in an information lifecycle management (ILM) solution in which files can be migrated from one tier to another, based on real-time policy. If the ILM solution is not snapshot aware and does not have a mechanism for handling files that are in the process of being migrated (or in flight), the resulting snapshot image can contain multiple copies of those files. Figure 10a demonstrates this situation. Just before 4 p.m., the ILM policy begins to migrate a very large file 2 from Tier 1 to Tier 2. Because of the file's size, the transaction is still in flight when the policy takes the 4 p.m. image. This results in two copies of file 2 in the 4 p.m. image: the original file 2a on Tier 1 and the incomplete file 2b on Tier 2. Furthermore, file 2b has a newer timestamp than file 2a, making it appear to be an updated version of the file. In the advent of a restore event, the administrator is faced with the decision to restore file 2 from the larger but older file 2a or from the smaller but newer (and incomplete) file 2b. The ARX device's virtual snapshot solution prevents this situation from occurring, as shown in Figure 10b. When creating a virtual snapshot, the ARX device terminates any file operation that is in flight and reinitiates the operation after the snapshot image has been created. Here, the 4 p.m. image contains only one copy of file 2: file 2a on Tier 1. The migration to file 2b is reflected in the next snapshot taken at 12 a.m.

## Other Advantages

The ARX device's virtual snapshot capability offers another advantage over physical snapshots in heterogeneous environments. Administrators have the freedom to apply a snapshot policy across an entire managed volume, and apply different policies to different managed volumes. Because managed volumes are typically

## White Paper

Creating Virtual Snapshots with F5 ARX

configured supporting applications, it enables the enforcement of data protection policy according to the business importance of applications.

# Conclusion

The market need for virtual snapshot technology has been driven by two trends: rapid data growth and increased requirements for data protection. As data protection has risen in strategic importance, IT users have increasingly identified snapshots as a core functionality for storage products, to the point where, today, no storage offering is considered complete without snapshot support. Meanwhile, rapid data growth has forced IT organizations to consider a range of alternatives in an attempt to contain their spiraling storage costs, including deployment of multiple storage vendors. Where these two trends intersect, they have created unique challenges with managing physical snapshots in the enterprise.

In a virtualized world, users need a solution that understands how to protect virtualized data. The ARX device's virtual snapshot solution provides a uniquely powerful approach to managing physical snapshots in virtualized storage environments. F5 applies its proven intelligent file virtualization technology to snapshots, integrating the presentation of physical snapshots across multiple backing volumes or file systems in a crash-consistent manner and with low impact to clients and applications.

**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)

