

F5 and Data Domain Partner to Curb Storage Costs

Date: January, 2009

Author: Terri McClure, Analyst and Lauren Whitehouse, Analyst

Abstract: It isn't feasible to save all of your file data on high performance, expensive file servers. These Tier 1 storage systems are great for transactional content, but are too expensive for inactive or archived data. F5 and Data Domain have joined their respective solutions, forming a partnership designed to assist customers in deploying and realizing the benefits of tiered storage. By combining F5's tiered storage policy engine with Data Domain deduplication storage systems, mutual customers can realize the benefits of deploying tiered storage and, importantly, see dramatic reductions in the costs of storage.

Overview

No one can argue the benefits of tiered storage—deploying systems with various cost characteristics and aligning the value of data with the cost to store it. A piece of content may be stored on any number of storage tiers over the course of its lifecycle. One of the biggest challenges in deploying tiered storage, however, is the fear that moving data across different storage systems will disrupt access.

Tiered storage deployments also cause another concern for storage staff: dealing with the growing amount of data stored on secondary devices. Most data is not changing and rarely accessed, yet it is kept on high performance, expensive file servers. Companies should archive or migrate inactive data to less expensive devices in order to control costs. Thirty-seven percent of organizations recently surveyed by ESG said that their file archives are going to grow over 20% per year. This trend could equate to significant growth in storage costs and data protection inefficiencies unless companies find a better way to retain and manage their information.¹

It is easy to solve these problems independently. When moved with a solution such as F5's ARX series, data ends up on lower cost, dense storage systems without disruption to file access. With so much file data proliferating across today's enterprises, F5 ARX solves a big problem. Now, customers are dealing with the next challenge—how to optimize utilization of the many newly created secondary storage systems full of archived files. IT departments are quickly introducing systems that deduplicate redundant content, helping to improve utilization of secondary storage. Buoyed by its inline deduplication technology and remote replication solutions, Data Domain's portfolio has become a common secondary storage tier.

Data Domain and F5 realized that, together, they could help customers deploy tiered storage in file environments and dramatically reduce storage costs with deduplication. The combination of the two technologies can deliver considerably greater storage efficiencies than deploying each product separately.

About the Partnership

Before discussing the partnership, it's important to understand what each company does. F5's ARX product offers a robust set of data management policies that can move file data between storage tiers based on age, type, and size—amongst other criteria. ARX performs migrations transparently; there is no impact to users, thanks to a global namespace that abstracts the file's network name from the physical file location. Customers can use the ARX to migrate files that are subject to record retention requirements, non-transactional or unchanging in nature, or simply old but still containing useful data, to the most appropriate storage system. This allows customers to store data on the most cost-effective storage media while still meeting data access requirements for users.

Data Domain offers storage systems that deduplicate data inline, before it is written to the system. Customers can move or copy data to Data Domain systems using standard NAS file system interfaces (NFS or CIFS), via a

¹ Source: ESG Research Report, *2007 File Archiving Survey*, December, 2007.

Virtual Tape Library (VTL) interface using Fibre Channel (FC) or with the NetBackup OpenStorage interface. As a result, many organizations have used Data Domain as a disk-based data protection solution because backups typically involve a significant amount of duplicate information—all of which takes up expensive storage capacity. More recently, Data Domain deduplication systems have been deployed as archive solutions to reduce the storage requirements for information being retained in an archive for compliance, electronic discovery, and business reference purposes. In addition to performing deduplication to eliminate redundancies, Data Domain systems also compress the data as it is stored, further reducing data footprint and storage requirements even when there is not a lot of redundant data. The same Data Domain system can be used for both backup and archive use cases.

The F5 and Data Domain partnership enables F5 ARX solutions to move files from a primary storage system to a Data Domain environment, resulting in savings on storage and associated costs related to power, cooling, and floor space. The Data Domain file system interface can be identified and managed by the F5 ARX global namespace. As information is moved into a Data Domain system, it is deduplicated and compressed automatically. When an employee accesses the file in its original location, F5 ARX retrieves the file from the Data Domain system, transparent to the user. The two companies plan to continue joint interoperability development and publish best practices for integrated implementations. They already have reference customers successfully deploying F5 ARX in a tiered storage environment with Data Domain deduplication storage systems comprising one of the tiers.

The Benefits

Both F5 and Data Domain deliver extremely compelling benefits when deployed individually—primarily centered on more cost-effectively storing data, but through different approaches. However, when implemented together, customers can magnify the storage efficiencies, resulting in significant economic savings. Customers using a joint solution should look for:

- **More efficient data protection.** According to ESG research, 63% of organizations using file archiving solutions have seen an improvement in file server backups.² F5 ARX moves inactive data to Data Domain systems, reducing the amount of content to back up from the primary file server. The results are far shorter backup and recovery times, as well as significantly lower backup-related capital expenditure. Additional backup optimizations accrue from the fact that Tier 2 is deduplicated. Customers can even choose to back up the smaller primary file server data set to the same Data Domain system now housing the deduplicated data.
- **Disaster recovery support.** Once on the Data Domain system, the information can be automatically replicated to another Data Domain system over low-bandwidth connections for disaster recovery purposes. Since the system has already deduplicated and compressed the data, economies related to both reduced bandwidth and reduced remote storage capacity required for disaster recovery are realized.
- **Lower storage capital expenditures.** Using F5 ARX to move files to a Data Domain solution frees up capacity on primary file servers for newly created content. Between 70% and 80% of files created are persistent—they stop changing within the first 90 days, but consume expensive Tier 1 storage capacity. Moving persistent data to an archive tier frees up a significant chunk of Tier 1 storage, which allows customers to delay purchase of primary or Tier 1 file server capacity. With 64% of current data deduplication users experiencing a 10:1 or greater capacity reduction ratio, customers can also defer tertiary storage expenses because they can effectively store more data in a smaller system (that has deduplication).³ As companies begin storing more data on Data Domain systems, it is likely that deduplication rates will increase over time, driving additional storage savings.
- **Lower storage operating costs.** By transparently moving data between storage tiers with F5 ARX, employees can continue accessing data without disruption—there is no IT intervention required. Also, since data is deduplicated when it is moved, companies do not have to keep buying additional storage

² Source: ESG Research Report, *2007 File Archiving Survey*, December 2007.

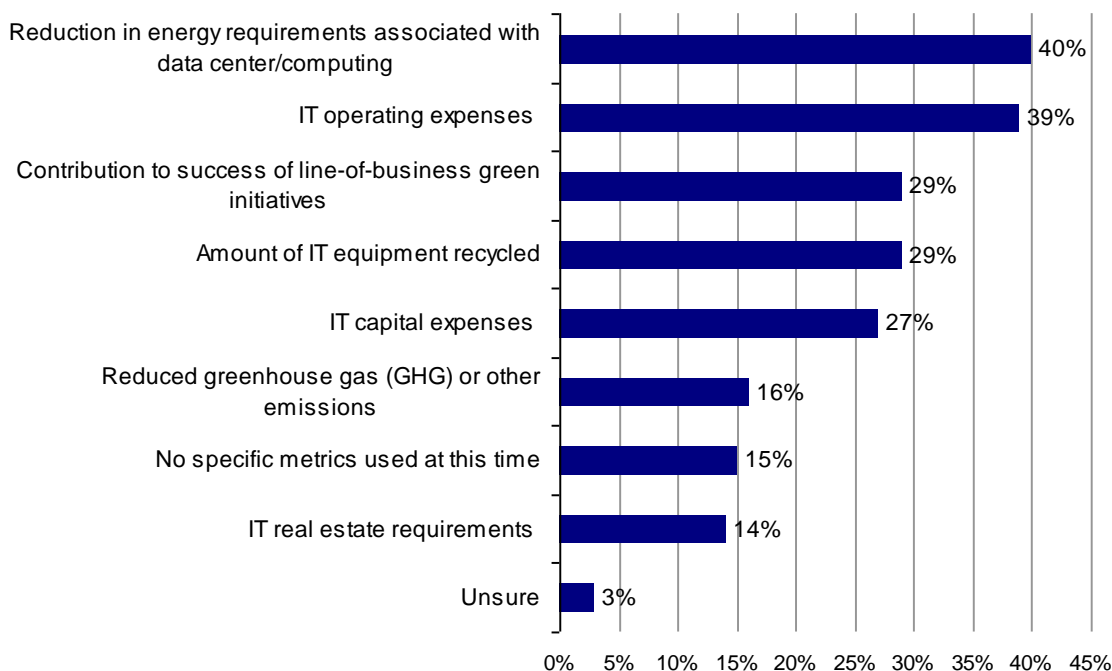
³ Source: ESG Research Report, *Data Protection Market Trends*, January 2008.

systems to save more information online for longer periods of time, which translates directly to greater power, cooling, and space efficiencies (PCSE). Given that data center space and power is hard to come by in many metropolitan areas, PCSE benefits can help IT meet corporate governance policies that include 'green initiatives' (see Figure 1).

- Seamless file archiving processes.** Many companies deploy tiered storage to facilitate archiving where data has to remain online and accessible for compliance, electronic discovery, and business reference reasons. ESG estimates that organizations will archive over 80 exabytes of file storage in the next three years and they will need tools to manage and store this data more effectively.⁴ To that end, customers can configure F5 products with policies to automatically identify archive candidates and move them to a Data Domain solution. Those organizations that need to comply with record retention regulations or evidence preservation requirements can leverage Data Domain Retention Lock software, which prevents data from being modified or deleted for the assigned retention period.

FIGURE 1. METRICS USED BY IT MANAGEMENT TO TRACK SUCCESS OF GREEN INITIATIVES

**What metrics does IT management most frequently use to evaluate the success of IT's contribution to your organization's green initiatives?
(Percent of respondents, N=223, multiple responses accepted)**



Source: ESG Report: *Global Green IT Priorities: Beyond Data Center Power and Cooling*, Enterprise Strategy Group, 2008

The Bottom Line

Emerging Web 2.0 applications, such as Wikis and blogs—combined with organic growth of traditional productivity applications—will drive the need for more file-based storage capacity. Companies can continue buying file servers to deal with incremental storage capacity growth and make do with existing backup and archive processes—or they can find a better way.

F5 and Data Domain are making it easy for customers to store more file data for longer periods of time at lower costs—which is a better way of doing things. Organizations have deployed these solutions independently and derived plenty of benefits. Now, they have an opportunity to increase storage capital and operating savings by

⁴ Source: ESG Research Report, *2007 File Archiving Survey*, December 2007

using them together. Like any other technology vendor partnership, commitment to interoperability and customer service is vital. So far, F5 and Data Domain have dedicated the appropriate resources. So long as this continues, customers trying to store file data more efficiently would benefit from investigating how the combined solution could help them.