WHITE PAPER

EMC Centera

Optimizing Archive Efficiency

By Terri McClure
With Brian Garrett

January, 2009

Table of Contents

Growth of Digital Archiving Creates Storage Management Challenges

Digital archiving has become a hot topic for corporate IT—driven by unrelenting data growth, the need to meet corporate compliance guidelines, and increasing federal regulation. To cost effectively meet data growth, compliance, and governance challenges, many companies deploy tiered storage to offload static data from Tier 1 storage to a secondary archive tier. It is just too expensive to keep every bit of data on high-performance, transaction oriented arrays. But this increase in digital archiving capacity, while driving efficiencies for Tier 1 storage utilization, creates new challenges for the archive tier and drives increasing complexity in the burgeoning archive environment. ESG's data indicates that the next wave of archiving customers will seek more sophisticated solutions that can improve performance and reduce costs by using policy-based automation to better manage archives and hold storage capacity growth rates in check.

ESG estimates that digital archive capacity will grow at a 60% CAGR between 2008 and 2012, and reach 90 exabytes within the next three years (see Figure 1). Users will need tools to manage and store this data more effectively—or risk being buried in a digital avalanche.

FIGURE 1. DIGITAL ARCHIVE GROWTH THROUGH 2012

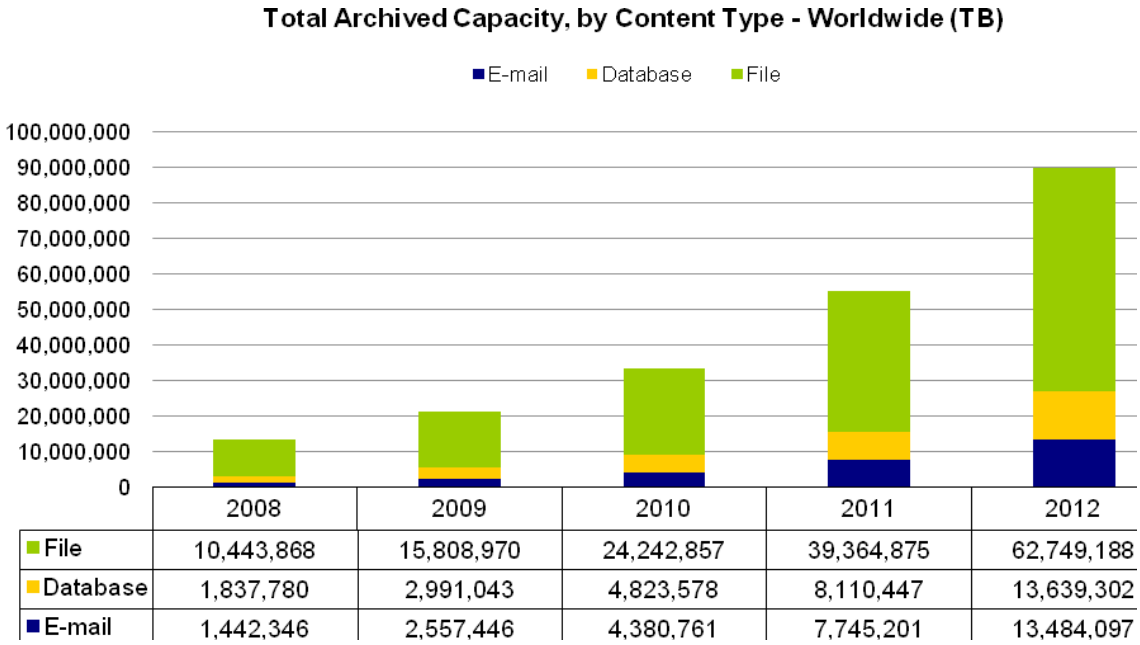

	2008	2009	2010	2011	2012
File	10,443,868	15,808,970	24,242,857	39,364,875	62,749,188
Database	1,837,780	2,991,043	4,823,578	8,110,447	13,639,302
E-mail	1,442,346	2,557,446	4,380,761	7,745,201	13,484,097

Source: ESG Report: 2007 File Archiving Survey: End-User Requirements & Priorities, December 2007

External disk systems are quickly becoming the digital archiving storage medium of choice for all types of archiving applications. For example, the growth of file archive capacity stored on *external disk systems* will exceed the overall market growth rate, increasing at an 85% CAGR between 2008 and 2012. Rapid growth in the use of external disk-based storage systems for file archiving means that the share of all file archive capacity stored on external disk will increase from 21% in 2008 to 37% in 2012.[1]

[1] Source: ESG Research Report: *2007 File Archiving Survey: End-User Requirements & Priorities*, December 2007.

Scale-Out Disk for File Archival

Scale-out disk—defined as a grid-based system with independent scale of processing power, bandwidth, and capacity—continues to grow in popularity as an archive platform. The most well known and widely deployed of these solutions is EMC Centera.

Scale-out disk carries a number of advantages over archiving to tape. First and foremost, it is online, which makes accessing archived data relatively fast. It is also more reliable. With scale-out disk, if a disk goes bad, there is typically some data protection mechanism in place—whether it is RAID, remote mirroring, or something in between. And the user will know the disk has failed and can correct the problem with a replacement disk. If there is a problem with a tape, nobody knows until the data is needed—and if the entire cartridge is bad, the data is lost.

There is a tradeoff when it comes to cost. Online storage uses power, space, and cooling, but it does not require the extensive capital and operational overhead associated with buying, handling, and storing tapes. While "high-profile" archiving considerations, such as regulatory compliance, garner many of the headlines in the IT and business press, factors related to IT operations were viewed by many respondents in ESG's 2007 digital archiving survey as key drivers in their decisions to implement digital file archiving solutions.[2] Interviews done for this paper clearly show that improving IT operational efficiencies are still the key consideration when choosing an archive storage platform. Cost efficiencies can be realized in scale-out systems thanks to the growing density of disk drives and much higher capacity arrays, while management efficiencies can be realized by keeping archive data online.

Offloading non-changing, persistent data from primary storage to an archive tier reduces the burden on primary storage systems in multiple ways while keeping data easily accessible. First, it reduces the amount of content to back up from the primary storage systems. The results are far shorter backup and recovery times as well as significantly lower backup-related capital expenditures. Second, additional backup optimizations accrue when archive data is deduplicated. Third, archiving reduces storage capital expenditures. Moving data off of Tier 1 storage to a scale-out disk solution frees up capacity on primary storage 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 the purchase of primary or Tier 1 capacity.

To net it out, implementing a disk-based archive tier reduces cost, increases primary storage system performance, and enables more data to be retained online for quick and reliable search and discovery. The end result is an increase in productivity, an increase in the ability to meet compliance requirements, and an increase in user satisfaction.

There are a number of scale-out disk archival options:

- Do it yourself with disk: Users create and manage network attached file systems.
- Do it yourself with NAS: Users manage one or more NAS appliances.
- Purpose-built NAS attached backup to disk appliance(s) with block-level dedupe.
- Purpose-built EMC Centera with object level deduplication.

There are pros and cons to each of these approaches. Do it yourself systems can quickly become complex as they scale, causing management costs to skyrocket. Scaling a purpose-built NAS environment is inherently more complex than scaling with Centera because it requires the same steps as configuring NAS for primary storage use: configuring RAID and LUNS, permissions, directory structure, etc. So the start-up time is extended for a NAS system compared to Centera, which has no concept of user-managed LUNS or RAID groups and no directory structure to deal with. We'll examine those differences later in the paper, but first consider the impact of deduplication in the archive environment.

[2] Ibid

Optimizing Information Archiving with Deduplication

Faced with ever-increasing amounts of archive data, customers are dealing with the next challenge: how to optimize utilization of the secondary or tertiary storage systems full of archived information. It is common for multiple exact copies of information and files to exist within a corporation. Users can save a lot of money, capacity, and time by saving just one copy of each. IT departments are quickly introducing systems that deduplicate redundant content, helping to improve utilization of these tiers. Implementing a data deduplication scheme reduces the overall disk capacity that needs to be purchased, managed, and maintained. It optimizes available storage capacity, allowing more data to be archived on disk, and helps users contain archive growth. Ultimately, it amplifies the benefits realized by deploying a scale-out disk archive solution.

To illustrate how deduplication works, let's look at general file archiving as an example. Consider Figure 2. Original capacity is pictured on the left, post deduplication data is represented on the right. The file data on the left includes multiple copies of files—for example, a PowerPoint file saved in home directories by two different users. As file data is moved into the archive, it goes through a deduplication engine to ensure only unique files are stored. Looking back at Figure 2, in this example, four of eight files were duplicates, so they were only stored once after deduplication, resulting in a 50% capacity savings. Block-based data deduplication works in much the same manner, but rather than looking at files, it looks at blocks of data within files. Block-based data deduplication is primarily useful for backup, where significant capacity savings over time can be realized.

FIGURE 2. FILE ARCHIVAL WITH DEDUPLICATION

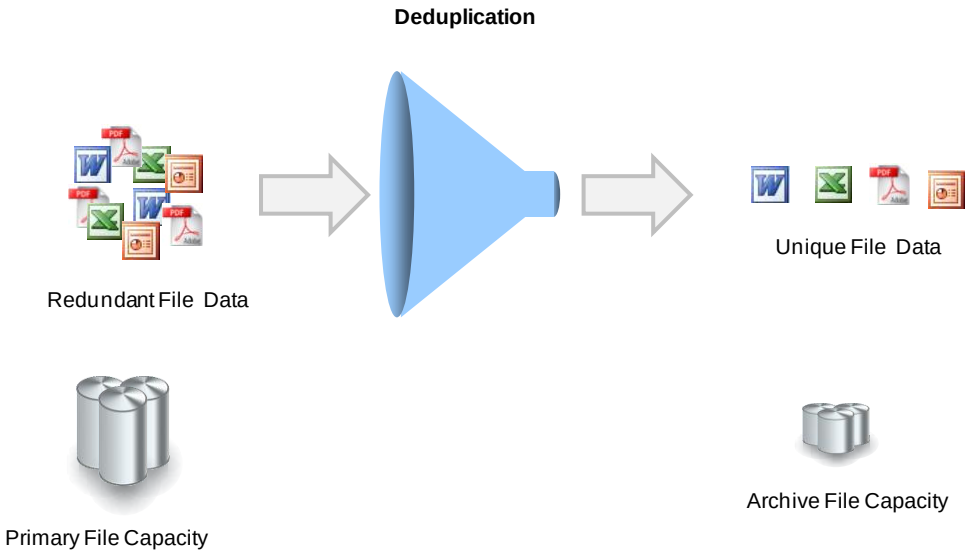

Source: Enterprise Strategy Group, 2008

It is important to understand how data deduplication for archival differs from deduplication for backup. The terms "backup" and "archive" are often used interchangeably, but there is a distinct difference in the processes. The process of archiving information involves physically *moving* information off of one platform to a new one, where it is retained for reference and to meet regulatory or business requirements. The backup process consists of *copying* information to a medium: disk, tape, or optical. In the backup process, the original information stays on the platform being backed up. Retained backups have a lot of information in common; 90% or more is typical because it is the repeated backup of primary data—this is where the big deduplication efficiency numbers the vendors tout come from. File-based archives have some information in common; up to 66% or more is possible. Because of these distinctions, deduplication of a disk-based archival is similar in characteristics to a *first* full backup. It is the optimization of that one data set. But it would be a mistake to compare the efficiencies realized in an information archive environment to what is realized in a retained backup environment.

Efficient Archiving with EMC Centera

Scale-out disk solutions supporting data deduplication continue to gain popularity as a preferred archive platform. One of the most widely known and deployed is EMC Centera. ESG interviewed a number of EMC Centera users to validate EMC's claims about Centera manageability, efficiency, and functionality versus managing a NAS-based archive. The users interviewed were from a number of industries, including financial services, public sector, higher education, and healthcare. The user feedback was 100% consistent. Every user cited the number one benefit of deploying Centera for archiving as hands off management. Of the organizations interviewed, managing archive capacities from tens to over four hundred terabytes, not a single user had more than .5 FTE (Full Time Equivalents) dedicated to management. The interview data, combined with ESG's own experience in the storage industry, confirmed that Centera holds an economic advantage when compared to purpose-built NAS appliances with block-level de-dupe.

The top benefits users cited related to using Centera as their archiving platform were:

- Management efficiency.
- Application integration.
- Reliability.
- Integrated remote replication.

> One user at a large international financial services firm manages 400 TB of archive information on Centera with .5 FTE. His NAS ratio for archived information was 3.5 FTE for 100 TB, a 28:1 difference in headcount required to manage NAS versus Centera.

Compared to NAS-based archive systems, Centera's flat address scheme significantly reduces management overhead. Consider the ease of managing a flat address scheme versus file systems. Essentially no management is required for a flat address scheme—there is no need to worry about physical and/or logical placement of information within the storage array like you'd have to with a NAS system. When you scale a NAS system, the differences are magnified. Imagine managing many file systems, worrying about directory structures, load balancing, performance optimization, network bottlenecks—the whole gamut associated with managing primary NAS systems—all for static, unchanging archive data. Bottom line: the management of archived information in a NAS environment is overkill for the business value derived from the information.

In our interviews, users cited scalability and density as factors in Centera's favor. The density and scalability of Centera means fewer discrete systems need to be deployed, monitored, and managed. Yes, some clustered NAS file systems can scale to the capacities required by deep online archives, but this magnifies the exception handling issues: The more moving parts you have (disk drives), the more failures you will experience. The self-healing nature of Centera means no impact for administrators and users when the inevitable failures occur.

In a Centera environment, the self-healing capability of the operating environment provides a very different user experience than replacing a failed disk drive in a NAS environment. Most NAS archives protect their data using RAID. A failed drive will expose a large number of data objects to risk of loss while the RAID group is being rebuilt using a hot spare. A self-healing, object-based environment like Centera exposes fewer objects to risk and will automatically bring the system back to health WITHOUT user intervention, in one pass. High availability is inherent in grid architecture—the grid architecture means higher availability and the ability to recover from a full node loss.

The grid architecture also provides significant flexibility. While Centera is an archive appliance and archives typically do not need to meet the I/O intensive performance demands that Tier 1 arrays do, some archive applications—such as digital imaging in the healthcare or financial services industries or call records in a call center—do need increased bandwidth and performance. Administrators can increase the ratio of access node profiles (processors and bandwidth) for those applications that require higher performance. For these types of applications, users reported excellent performance with Centera—citing single-digit millisecond response times. And this makes it easy for organizations to move static information from Tier 1 storage to the archive because the user experience remains the same. Centera also allows users to reduce the ratio of access node profiles to storage node profiles for more cost efficient bulk storage geared towards applications that can withstand longer response times. Since applications and data types vary significantly from company-to-company and application-to-application, actual mileage may vary, but the flexibility of the grid architecture to meet demands in either direction—rather than deploying an extra tier—adds significant value.

> "It is tough to discuss Centera management because we have to do so little."
>
> *Operations and Systems Analyst, Higher Education*

> "Centera is really hands off. It is just rock steady."
>
> *Data Administrator, Public Sector*

The implementation of a grid architecture also future-proofs—Centera can include multiple hardware generations in a single cluster. Many users cited a need to keep archive data available and easily accessible to meet lengthy regulatory retention requirements, while a number cited that they plan to keep data in the online archive "forever." Archiving to tape means lengthy processes to migrate to new tape technology as it evolves. And migrating NAS involves manually remapping directory structures and remounting file systems or buying a file virtualization appliance to seamlessly migrate at the expense of adding another set of products to manage. Migrating to newer technologies in a grid environment means plugging in new nodes and taking the old ones offline. The system automatically migrates the data behind the scenes. There would be no LUNS, RAID, mapping, or layouts to worry about, significantly reducing the labor associated with keeping up with technology advances over the life of the archived data.

The wide range of applications that are integrated with Centera was cited by a number of users as a key advantage, primarily because a user can pick and choose the products to use. But a potentially larger benefit is that, because of tight integration, both the application and storage know what each other is doing to the archived information. Applications integrated with Centera create an archive solution with a complete chain of information custody from the application through storage, which is highly beneficial for internal governance, e-discovery, and regulated industry use.

ESG Lab: Archive Deduplication Efficiency Analysis

ESG Lab evaluated the efficiency of Centera's object-based deduplication scheme (referred to by EMC as Single Instance Storage) against a NAS appliance with block-level deduplication. For the purposes of the test, the archive deduplication data set consisted of emulated corporate file share data made up of:

- 20,000 files,1,000 directories and a 10.99 GB data set that represented (see Figure 3):
  - 10 company divisions.
  - Each division has 10 departments.
  - Each department has 10 employees.
  - Each employee has 20 files: 10 programmatically unique, 10 from a common pool.
- Average shared file size = 700 KB (but ranged from tens of KB to several MBs).

FIGURE 3. CORPORATE FILESHARE DATA USED FOR TEST

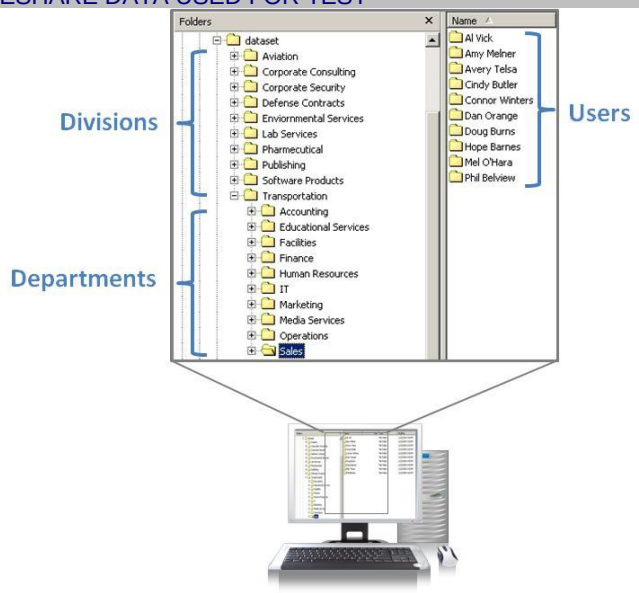

Source: Enterprise Strategy Group, 2008

The archive deduplication test process consisted of backing up the file data onto a NAS appliance using EMC Networker and copying to a Centera share through a CUA (Centera Universal Access, a NAS gateway to Centera) (see Figure 4). The NAS appliance was deduplicated using variable length blocks as it arrived at the appliance. The Centera deduplication was object-based (Single Instance Storage) as it arrived at the Centera through the CUA.

FIGURE 4. CONFIGURATION TESTED

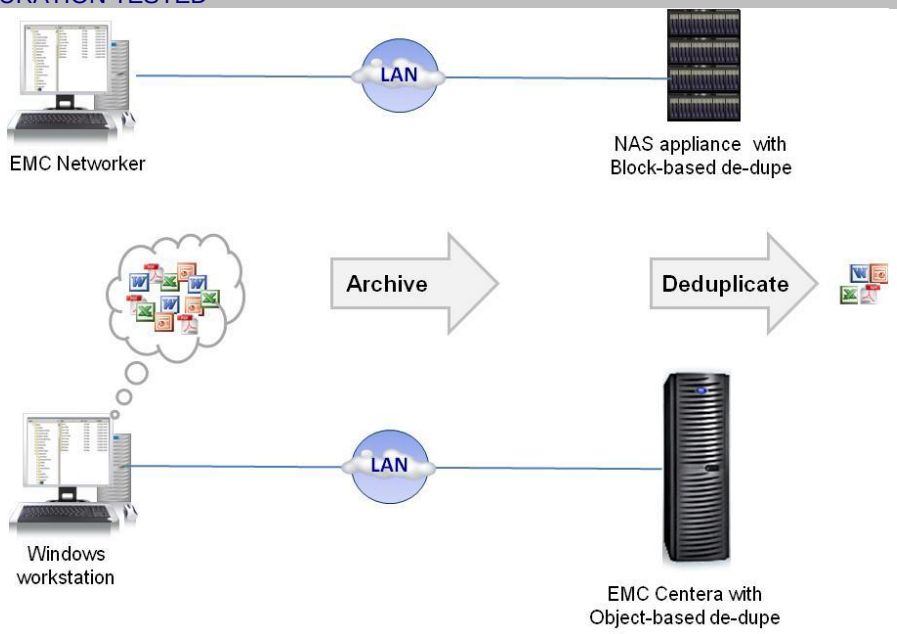

Source: Enterprise Strategy Group, 2008

The results of the test were not surprising. The NAS appliance realized a 69% capacity savings while Centera realized a 66% savings (see Figure 5).

FIGURE 5. CAPACITY SAVINGS COMPARISON: BLOCK-BASED VERSUS OBJECT-BASED DEDUPLICATION

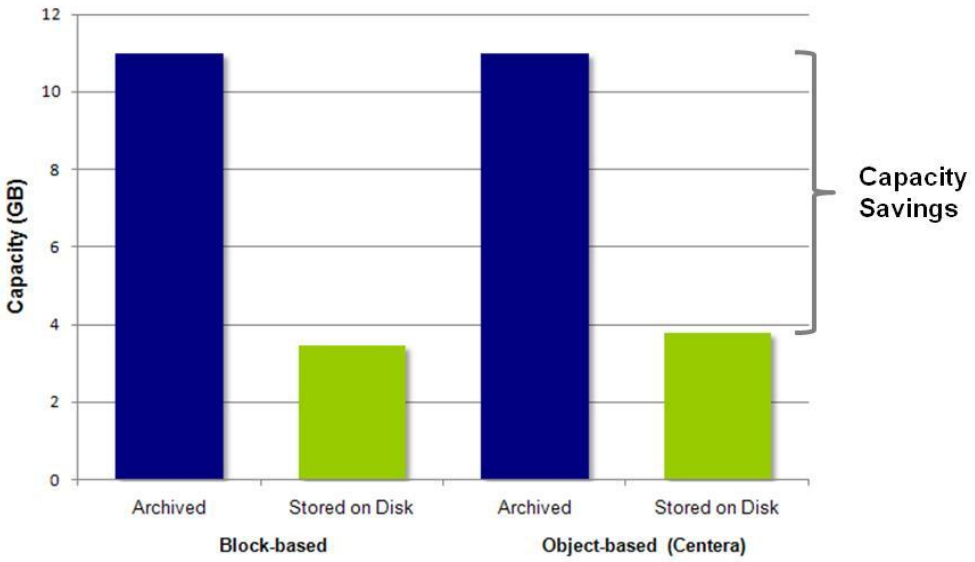

Source: Enterprise Strategy Group, 2008

The almost negligible difference in reduction rations—66% reduction for object-based data versus the 69% reduction for block-based data—while comparable, is not a perfect "apples to apples" comparison. ESG Lab has two possible reasons for the discrepancy. One possible reason is that, compared to Centera Single Instance Storage, block-level deduplication can identify some small chunks of data within files that are common and deduplicates those chunks. The odds are low for this happening in archive data; this is more likely to happen for deduplication of backup data. A more likely explanation is Centera's use of metadata. Unlike the NAS system

tested, Centera stores metadata in addition to the file, which is included in the measurement. Metadata is not stored in the block-level deduplication systems.

With storage-based metadata, the archive application and archive users can achieve a higher level of performance. When looking for a piece of information, the application looks through the metadata clips and does not have to actually open the files or content items. Also, if the file or content is exceptionally small, instead of single instancing it, Centera will store it in the metadata clip for each user or application that requested that it be stored (it embeds a copy in the metadata). That also helps performance when dealing with random requests for information from an archive. This active archive environment, for which Centera was purpose-built to serve, is one of the points where archive performance requirements diverge from backup performance requirements. Where block level deduplication was designed for reconstituting large quantities of contiguous, streaming information in a backup environment, only within the past year has the use case been expanded by vendors to include archiving. And most deduplication platforms hit performance bottlenecks when dealing with random requests rather than large streaming data processes. Centera, however, does not.

> For archiving, the capacity savings difference between Centera and NAS appliances with block-level de-duplication is negligible.
>
> *ESG Lab*

The above noted test results are similar to those achieved in first full backups done during previous ESG Lab Validations:

- 50% for an initial backup of home directory files. [3]
- 47% for another first full backup of production home directory data. [4]

Bottom Line

Based on this testing, as well as previous deduplication tests ESG has conducted, ESG can confidently state that the difference in capacity savings between Centera and NAS appliances with block-level deduplication is negligible. Both provide excellent levels of savings.

It should be noted that as the archive grows, the advantage swings to Centera because it can scale to be much larger than NAS appliances. Deduplication schemes do not span arrays; each time a new system is deployed, only the data within that system is deduplicated, reducing the efficiency of the overall deduplication effort. This limits the deduplication efficiency that can be achieved by deploying multiple NAS arrays versus one scaled-out Centera grid.

ESG's View

It is easy to be fooled by vendor claims of deduplication capacity savings of 20:1 or more, or 95% or more. These numbers apply primarily to retained backup data—where the same data sets are backed up over and over again—and not to archiving.

In an interview with ESG, one EMC Centera user reported a capacity reduction ratio of 5:1 when using Centera for e-mail archive (an 83% savings). Different applications will generate different savings and, for the example used in this paper, savings of 50% or more are typically achievable.

With a level playing field on the deduplication front, the evaluation criteria that users should consider for scale-out file archiving shifts from capacity reduction to overall management and ease of use. Users should evaluate disk-based archive solutions based on:

[3] Source: ESG Lab Validation Report, *Quantum Dxi-Series: Disk Backup and Remote Replication Appliances*, May 2007.
[4] Source: ESG Lab Validation Report, *ExaGrid Systems: Disk-based Backup with Data De-duplication*, March 2008.

- **Ease of deployment**: Evaluate how quickly the solution can be up and running and how much training is required.
- **Application and storage integration:** It is crucial for most businesses to find a solution that enforces an organization's information retention and disposition policies in storage. Applications and storage need to work together to create a complete chain of information custody from applications through storage and retrieval.
- **Ease of management:** Look for systems that are self-optimizing and self-managing—archive data is growing too fast to deal with tuning storage arrays and data paths.
- **Extreme scalability:** The digital archive is likely to be the fastest growing storage platform in the data center, thus extreme scalability is a core requirement to reduce the number of systems deployed and therefore reduce management costs. Also, deduplication schemes do not span arrays; each time a new system is deployed, only the data within that system is deduplicated, reducing the efficiency of the overall deduplication effort.
- **Resiliency:** Look for systems that are self-healing, able to withstand multiple component failures, and able to quickly rebuild data in the event of a failure. The longer a recovery takes, the longer data is at risk of loss, so object-based mirroring or distributed parity are requisite to speed rebuilds.
- **Protection from technology obsolescence:** The life of the data retained in the digital archive will exceed the useful life of the archive platform. Ask your vendor what the migration path is to new generations of hardware or higher capacity disk drives, and if there is downtime required or extra hardware needed to migrate the data.
- **Assured data integrity and authenticity:** Last, but not least, don't expose the corporate jewels to data corruption, silent errors, or lengthy file system checks. Locking down a system and setting retention policies is one thing, but ask your vendor how it assures the integrity of the data in the archive.

When you run down this list, it becomes clear why EMC's Centera platform has become a dominant player in the digital archive arena. Customer interviews confirmed that Centera has an advantage across the board. Users were unanimous in feedback that Centera is a hands-off system that "just runs." It is self-managing, self-healing, and integrated with over 250 applications. Because it is a grid-based, scale-out disk system that can run multiple generations of hardware within the cluster, it provides scalability that far outweighs scale-up NAS and a high level of future proofing.

But perhaps the most important point in Centera's favor when compared to NAS- or SAN-based archive solutions is assured data authenticity and integrity. In the real world, neither NAS- nor SAN-based storage systems innately test data for integrity; data will have to be backed up to ensure that copies of archived data are accessible in case of hardware failure or data corruption. This adds to overall management overhead and the backup pool. These integrity checks can add significant management overhead for NAS versus something that is basically hands-off in Centera. For a relatively small archive, for which data retention and integrity were not critical, locked NAS might be as hands-off as Centera, but not when scaled. When evaluated against criteria users are looking for in digital archive solutions, the score is clear: advantage, Centera.

20 Asylum Street
Milford, MA 01757
Tel: 508-482-0188
Fax: 508-482-0218

www.enterprisestrategygroup.com