

Scale Bigger
Store Smarter
Accelerate Everything

How You Can Manage Data As A Vital Business Asset By Taking An Intelligent Approach To Tiered Storage

There are few assets more valuable to a business than its data. This has always been true, but it is particularly relevant today as we continue to move into an age where all business assets — from seemingly mundane emails to annual reports to videos and complex technical renderings — are being created and saved in a digital format.

For businesses of all types this raises significant challenges: The cost of storing data can be significant and the idea of just adding additional systems to handle increased demand is both inefficient and impractical. For organizations with large data storage requirements of 30 terabytes to multiple petabytes the need for cost-efficient solutions is particularly glaring.

But the challenges go way beyond dealing with the rising costs of storing more data. The reality is that the data being stored must be managed as a strategic asset. Think of it this way: If your organization is going to the expense of saving the data, then the data must have value to the business. But if the people in your organization can't get to the data when they need it, or if the performance of the data is compromised in any way, then you are not taking advantage of this asset. Worse, you are probably wasting a lot of money on inefficient storage management solutions.

That's where intelligent tiered storage comes in. Intelligent tiered storage, as delivered by innovators such as BlueArc, enables organizations to manage their data across the entire enterprise as a strategic asset by moving the data transparently between storage tiers, depending upon its value to the organization at any particular time.

Here's what this can mean: Data of any value — even if it hasn't been viewed in years — can be accessed simply, efficiently and transparently as soon as the organization deems it necessary. What's more, by managing the data in this manner, performance is not compromised one iota, even on data that has been sitting in an archive or secondary tier. And the organization will always know where every piece of data is sitting at any particular point in time.

Why Intelligent Tiering is Essential

Most businesses with large data storage requirements have recognized the need to manage their data strategically through the implementation of tiered storage solutions. According to a recent survey by *Storage Magazine*, approximately 60 percent of IT professionals said they are using tiered storage.

But simply having a tiered storage solution does not mean you are managing your data assets strategically or intelligently. Moving data to different tiers may save money, but it doesn't guarantee you will be able to access data when you need it and it could, in fact, add significant management overhead if you are doing these processes manually. In that same survey by *Storage Magazine* more than 35 percent of respondents said they are moving data manually from tier to tier.

The only way to truly manage data as a strategic asset is to manage it intelligently and transparently, so that the data moves automatically to the appropriate tier and can be moved up to a higher tier easily when it is needed. Transparency means the IT department doesn't have to reconnect applications after data has been moved and it means users don't have to be notified that their data may reside in a different location.

For certain businesses, it is simple to see why managing data in this manner is absolutely essential. For example:

- Media and entertainment companies that need access to old episodes or movies and are also reusing video or music
- Energy companies that need to reuse complex, data-intensive, high-performance seismic images or analyses
- Healthcare, biomedical and pharmaceutical companies that need access to computational analysis workflows and complex analytical research

But for all businesses with large data requirements, the need for high-performance intelligent tiered storage is equally essential. Not only are organizations producing more data than ever before, they are producing data in a wider variety of formats and they are keeping it for much longer periods of time. What's more, they are often producing and reproducing some of the same data and storing it over and over again in many different places. Think about a corporate video or PDF that is sent to hundreds of employees and is saved by each of them: How much valuable storage space is that wasting?

What's more, changes in regulatory compliance and Electronic Discovery laws mean that data not only must be stored longer, it must be quickly and easily accessible wherever it is stored. Companies with more than \$500 million in revenue generally face five or more litigation matters each year, according to *InformationWeek*: No one wants to have to tell a judge in a criminal or civil case that your company is unable to provide information because it is archived somewhere and you can't find it.

It's also important that performance is not compromised on data that is either in secondary storage or archives. A complex 3-D CAD drawing is a complex 3-D CAD drawing no matter where it is stored. The same with other high-performance data such as high-definition video: If your CEO is showing an older advertisement or corporate video to a key customer, you don't want pixilation problems or audio out of synch with video.

A Breakthrough In Data Management

The approach to intelligent tiered storage offered by BlueArc is the state of the art in storage management because it enables organizations to truly treat their data as a business asset. That's because the data is managed automatically and transparently –based on its inherent value to the organization at any time, and not based on the storage tier where it happens to reside at any particular moment.

BlueArc's SiliconFS provides a hardware-based architecture that enables the IT organization to automate the process of migrating data to and from secondary storage without impacting the application or the experience of end users. The BlueArc approach enables critical functions such as archiving, deduplication and compression to be handled automatically and transparently. It also means all of these functions can be extended to NAS systems other than BlueArc and to multiple storage technologies.

Central to this groundbreaking approach by BlueArc are two software features that are embedded in SiliconFS called Data Migrator and Cross-Volume Links.

Data Migrator is a policy-based engine that enables administrators to implement data movement policies that work automatically and transparently. To the user, the data appears as if it were still on its original storage tier. The integrity of the data is maintained during the migration or recall and, as far as end users and applications are concerned, the data has not moved at all. Among the key functions this enables are content-aware deduplication and optimization as well as granular user-definable compression.

Cross-Volume Links support this same type of seamless data migration and extends it to external devices, including those offered by deduplication vendors and third-party NAS systems. Data can be moved from tier to tier and to external tiers and still be managed and presented as a single, cohesive whole. Functions such as archiving, deduplication and compression are also extended to third-party devices as well.

The value to a business delivered by this level of advanced functionality can be staggering:

- **ROI:** Organizations can optimize their existing infrastructure by moving secondary or tertiary data to less costly disks while freeing up capacity on their high-performance disks. This enables organizations to maximize performance without having to purchase additional high-performance drives as frequently. This not only reduces expenditures, it also reduces energy consumption and consolidates space in the data center. The cost savings are both immediate and long-term.
- **Efficiencies:** By automating functions such as archiving, deduplication and compression, the management overhead on the IT organization is reduced significantly and the entire storage infrastructure operates much more efficiently.
- **Scalability:** When it's time to scale up to handle increased storage requirements Data Migrator and Cross-Volume Links enable this process to be handled smoothly and efficiently, enabling the organization to choose best-of-breed technology solutions at any time.
- **Performance:** The software available in SiliconFS enables the architecture to automatically manage spikes in demand on secondary or archived storage so performance is assured whether the data is residing on a high-performance Fibre channel drive or a lower-performance SAS or SATA drive.

Further enhancing the high-performance capabilities of the BlueArc Intelligent Tiering is a feature called **Dynamic Read Caching**, which extends highest performance features to data that is in secondary tiers or is archived on disk. With Dynamic Read Caching, space on higher-performance storage tiers is made available for the caching of “hot” files. As soon as a file is accessed off the secondary tier it is “hot” and loaded into the cache of the primary tier. Subsequent access of the file while it is in the read cache is from the primary tier and access would be nearly instantaneous. This is a feature that is particularly useful in applications such as eDiscovery, entertainment, Internet services and life sciences.

Real-World Success Stories

All of this functionality is critical for businesses that aim to manage data as a strategic asset. And, in today's world, if you are not managing your data in this way, how are you managing it? As noted, all businesses must deal with issues such as ROI, performance, efficiency, compliance and eDiscovery. Beyond that, here are some of the innovative ways in which BlueArc's Tiered Storage solution with Data Migrator and Cross-Volume Links is being applied in a couple of specific applications:

Rainmaker Entertainment Inc., based in Vancouver, is one of the world's leading producers of CG animation. It has produced more than 30 feature-length films and generated more than a half-billion dollars worth of video sales. For Rainmaker, BlueArc's Tiered Storage solution is critical in managing deduplication and archiving, among other important functions.

As a studio, Rainmaker is challenged with having a finite amount of data storage space and the need to keep data online for as long as possible both during production and after it finishes. Once a show is finished, the company archives it and removes it from disk — yet producers and artists are always seeking access to previous material. Without the BlueArc solution, retrieving data from storage and returning it to disk could take several days or more. With BlueArc, artists can access any previous show in near real time. This means lower production costs and faster speed to market.

Another BlueArc customer is a major service firm that handles document and data management for law firms and corporate legal departments. The firm is using Blue Arc intelligent tiering with Data Migrator and Cross-Volume Links in a disaster recovery application with an emphasis on compression and deduplication to dramatically reduce bandwidth requirements. The firm was able to enhance its eDiscovery functionality by moving data transparently across tiers and enabling the company — and its customers — to get the information it needs in an efficient and timely manner.

Managing Data As a Vital Business Asset

Managing data as a vital business asset involves a lot more than saving money by moving data around between storage tiers. Saving money is important: Data storage requirements are mushrooming — with no end in sight — and even now most businesses are spending up to 20 percent of their IT budgets on data storage, according to the analyst firm IDG.

Which is all the more reason why businesses need to treat their data as a business asset and manage it cost-effectively, efficiently and with an eye toward future growth. After all, if data is a business asset, growing it as an asset should be a good thing for the business. But if the growth is not managed strategically, then the asset can quickly become a liability.

With BlueArc's Tiered Storage solutions featuring Data Migrator and Cross-Volume Links, businesses can be assured that they will be able to tackle any data storage challenge now while future-proofing their storage infrastructure for any new challenges that may be looming on the horizon.

About BlueArc

BlueArc is a leading provider of high performance unified network storage systems to enterprise markets, as well as data intensive markets, such as electronic discovery, entertainment, federal government, higher education, Internet services, oil and gas and life sciences. Our products support both network attached storage, or NAS, and storage area network, or SAN, services on a converged network storage platform.

We enable companies to expand the ways they explore, discover, research, create, process and innovate in data-intensive environments. Our products replace complex and performance-limited products with high performance, scalable and easy to use systems capable of handling the most data intensive applications and environments. Further, we believe that our energy efficient design and our products' ability to consolidate legacy storage infrastructures, dramatically increases storage utilization rates and reduces our customers' total cost of ownership.



BlueArc Corporation
Corporate Headquarters
50 Rio Robles Drive
San Jose, CA 95134
t 408 576 6600
f 408 576 6601
www.bluearc.com

BlueArc UK Ltd.
European Headquarters
Queensgate House
Cookham Road
Bracknell RG12 1RB, United Kingdom
t +44 (0) 1344 408 200
f +44 (0) 1344 408 202