

# BlueArc Titan

## Intelligent Multi-Tiered Storage System

By Tony Asaro  
Senior Analyst

May 2005

# Intelligent Multi-Tiered Storage

---

Intelligent multi-tiered storage addresses the performance requirements of mission-critical applications, the rapid growth of reference data, and the need to control the cost of storage infrastructure. On the revenue side of the customer's business intelligent multi-tiered storage first and foremost is about aligning the storage capabilities with the business\application needs. As the data flows through the business process and data life cycle, intelligent migration allows administrators to respond to the performance, capacity, and access characteristics necessary to support the user and application requirements at the time. Ultimately, intelligent multi-tiered storage also provides better storage utilization and great cost savings. Intelligent multi-tiered storage is about responding to demands of the business and delivering the right capability at the right time.

The reality of capital cost savings is also extremely important. ESG research recently conducted a study and found that on average storage budgets for hardware, software and services across all market segments were projected to increase by only three percent in 2005. With inflation in the United States at approximately 2.5%, this projected increase in IT storage budgets is essentially flat. However, storage capacity is still growing at a rapid rate but IT departments are not allowed to increase their budgets to store and manage it.

Intelligent multi-tiered storage can go a long way in addressing the "do more with less" edict imposed upon IT departments by combining high performance Fibre Channel (FC) drives and lower cost Serial ATA (SATA) disks. Intelligent multi-tiered storage should also provide tools to transparently migrate data between FC and SATA disks without impacting users. The goal is to reduce cost and complexity but not at the cost of application performance, data availability or data integrity.

## Hard Disk Technologies

ESG believes that the use of SATA disk drives in storage systems has made a major impact on the storage network market. SATA drives provide greater capacity at a lower cost than FC drives enabling customers to store more data for less cost. Today SATA drives have been implemented by tens of thousands of customers as a part of a tiered storage environment. SATA storage is often being used for data that is less frequently accessed or for applications that don't require the highest levels of performance.

FC disks play a major role as the high performance tier within a storage system. The hard drives are the slowest component in the I/O performance chain when compared to bandwidth, processors, cache memory and bus speeds. Having the fastest possible disk drives can have a significant impact on application performance especially for disk-intensive I/O workloads.

## Data Movement

Customers can combine the advantages of high performance storage (FC disks) with enormous data stores (SATA disks) while maintaining control over costs. However, having different disk types is not enough. It is important to have tools that will move data between these tiers without having to take applications offline or inconveniencing users. These tools turn a mixture of different disk drive types into an intelligent multi-tiered storage environment.

Moving data online between different drive types within a storage system can be a difficult process. Many storage systems do not provide the tools to perform this function online or transparently. The most popular method for data migration is performing a backup from one set of drives and then restoring to the new set. This is an offline process that can take a great deal of time. And while the data is being backed up and restored, new data has been created. Customers can also use host-based application to perform data migrations but this requires additional software and licenses for all of your servers that require data movement.

As we shall discuss in this report, BlueArc provides data migration tools that allow customers to seamlessly move data between different Tiers within a Titan storage system. This capability allows system administrators to perform data movements manually or automatically using definable policies.


The BlueArc Titan supports FC and SATA disk options, data migration software, and NFS, CIFS and iSCSI protocols creating an intelligent multi-tiered storage system. This report examines the benefits of intelligent multi-tiered storage using the BlueArc Titan storage system. The [ESG BlueArc Titan - Intelligent Multi-Tiered Storage System](#) report addresses the following:

- What is intelligent multi-tiered storage?
- What are the benefits of intelligent multi-tiered storage?
- What should customers consider when deploying a multi-tiered storage strategy?
- How does the BlueArc Titan support intelligent multi-tiered storage?
- How are BlueArc customers implementing intelligent multi-tiered storage?

# The BlueArc Titan and Multi-Tiered Storage

The BlueArc Titan is a next generation NAS storage system that provides a scalable file system, linear performance and a number of features and capabilities that make it a compelling solution. As seen in Figure One, the BlueArc Titan offers different tiers of disk capacity and storage protocols. The Titan is a flexible system that allows customers to combine these different options based on the needs of their environment.

Figure One: BlueArc Titan Multi-Tiered Storage System



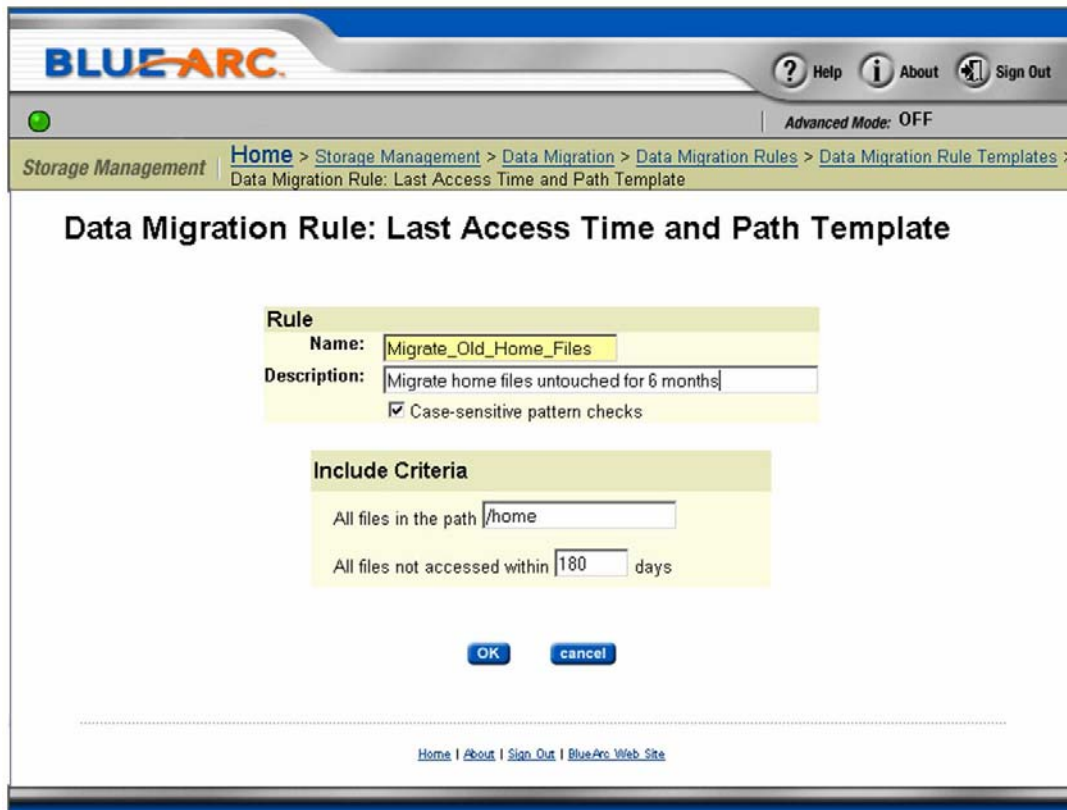
**Different applications can have data stored on different drive types appropriate to the application's requirements.** For example, customers can store performance intensive database applications on 146 GB FC hard drives running at 15,000 RPM. Additionally, customers can store non-transactional data, such as files (e.g. PDF, images, text) onto SATA drives. ESG is also aware of a large number of customers performing disk-to-disk (D2D) backups onto SATA drives.

The BlueArc Titan supports the following hard disk drive types:

<i>Hard Drive</i>	<i>Capacity</i>	<i>RPM</i>
FC - Tier One	146 GB	15,000
FC - Tier Two	146, 300 GB	10,000
SATA - Tier Three	400 GB	7,200

**Data can be migrated transparently and automatically to and from different tiers manually and/or using policy-driven metrics.** Having different hard disk types within a storage system is a critical first step towards productive storage tiers, but customers also need a way to transparently move data between them. BlueArc provides this ability through a feature they call Data Migrator. This is a feature that enables seamless movement between different tiers within a Titan storage system. System administrators can move data manually or setup policies for automated data movement (Figure Two).

Figure Two: Data Migrator Policies - Last Access and Path



Data Migrator offers various policies for movement data between tiers. In Figure Two, the customer is moving data within the /home directory that have not been accessed for 6 months. Other metrics include:

- Last Access
- Last Modified
- File Name
- File Type
- Path
- File Size
- And / or combinations of these attributes

In addition these rules can be created and managed via a policy that offers many options. Administrators can run rules to do a “what if” scenario, identifying the number files and amount of data that would be migrated, without actually moving any data. Timed policies with repetition at regular intervals can also be created. Additionally, policies based on trigger conditions, such as capacity utilization can be created. For example a rule can be executed when the storage capacity exceeds 80 percent.

**Administrators can make snapshot copies of data from FC disks to lower cost SATA disks.** Customers can cost justify putting mission-critical data on the fastest and most expensive FC drives, but often are forced to make copies of this data FC drives. BlueArc can make snapshot copies of data stored on FC drives onto lower cost SATA drives within a Titan storage system. This drives down the cost of copies and could actually motivate customers to make more snapshots, which results in greater data integrity and recoverability.

## Use Cases

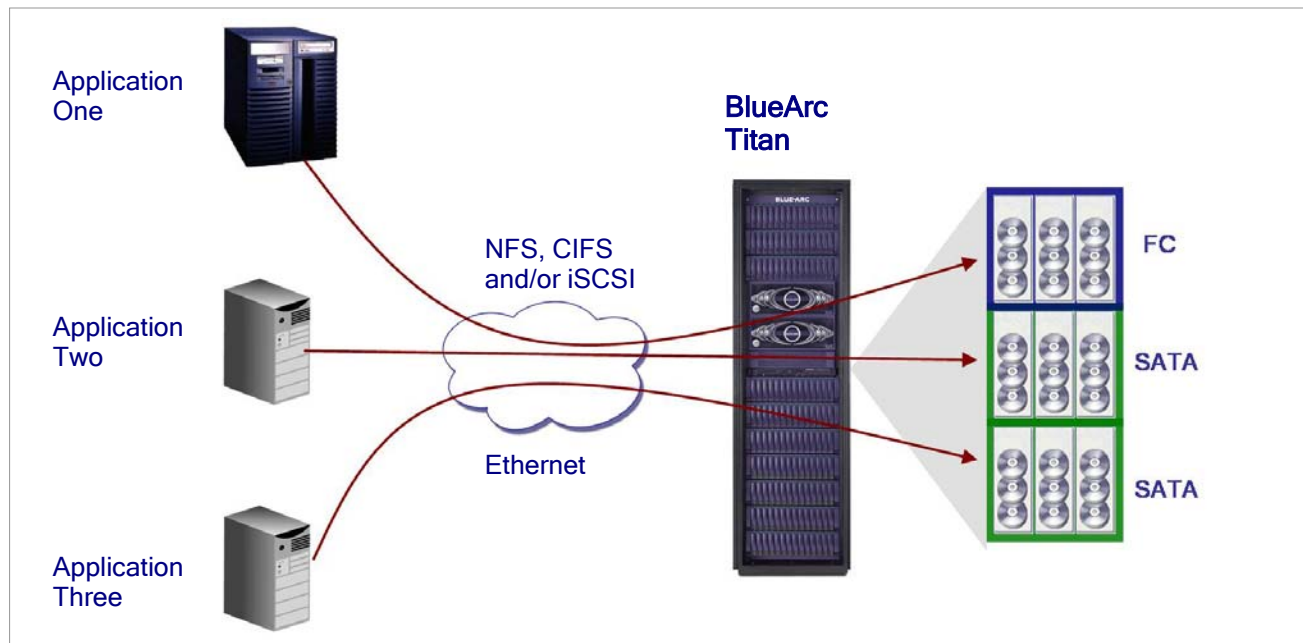
There are a number of ways to use an intelligent multi-tiered storage configuration. This section discusses different approaches to multi-tiered storage using the BlueArc Titan.

### Applications and Multi-Tiered Storage

One popular method for implementing a multi-tiered storage strategy is to store data from different applications onto various hard disk types. In Figure Three, the different applications could include a high performance application that is very transactional; the second level could be an online application that is not performance sensitive; and the third level can be used for archival and/or near-line applications. An example of how this could work:

- Application One can have a large number of files being created, read and edited with thousands of users accessing the data concurrently. This could be an online e-commerce application or a database application. This Tier could be accessed via NAS or SAN interconnects.
- Application Two could be an email application using iSCSI. For this example the customer's email application does not require maximum performance and is stored onto SATA drives.
- Application Three can be used as a disk-to-disk backup target and/or for storing unstructured archival data (e.g. images, scanned documents, PDF, spreadsheets, etc.). The BlueArc Titan also supports Write Once Read Many (WORM) functionality to meet corporate governance and regulatory compliance mandates.

Figure Three: Applications and Tiered Storage

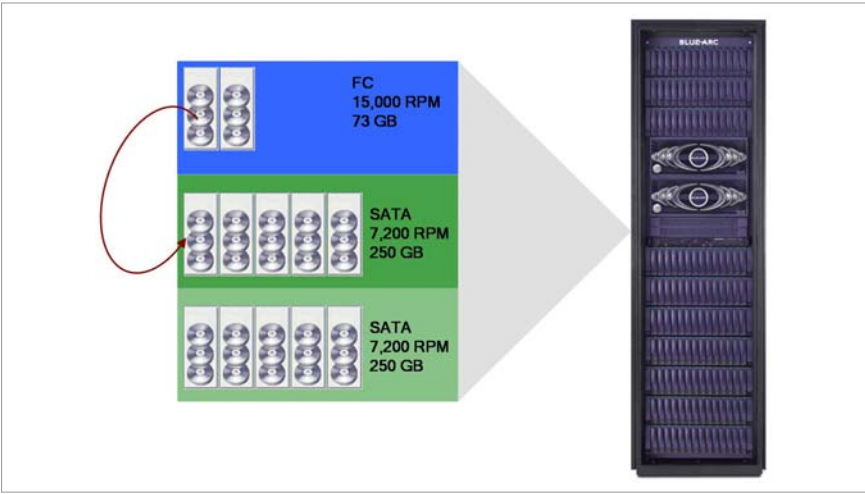


### Data Movement between Tiers

In addition to dedicating different applications to different Tiers, customers can move data between different tiers either manually or automatically using rules-based policies. For example, if data stored on Tier One has not been accessed for over a month system administrators can set a policy that moves it automatically to a lower Tier. The BlueArc Data Migrator moves data within the Titan online and transparently to the users. Data Migrator creates a pointer for the file on the FC tier, recovering this space, while directing users to the location of

the file on the SATA tier. The data is physically relocated, but remains in the same logical location in the file system.

Figure Four: Data Movement between Titan Tiers using Data Migrator

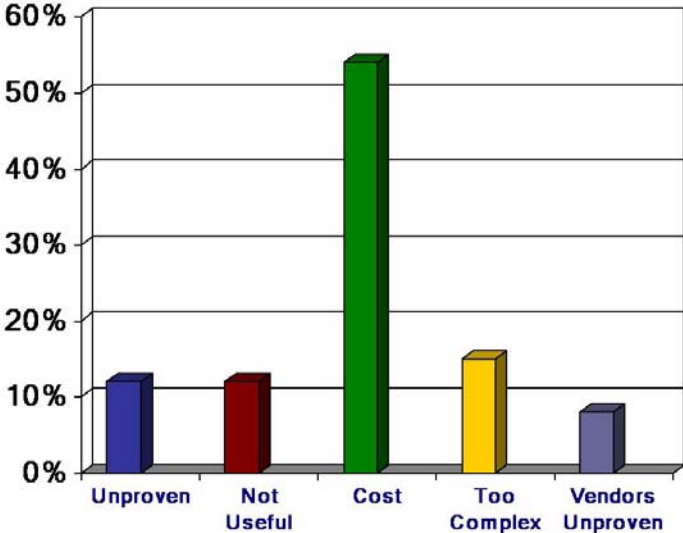


**Snapshot Copies**

Snapshots are one of the most useful and valuable features of storage systems. Snapshots are a static picture of data that is a frozen point in time. Snapshots are often used to recover files that have been accidentally deleted, corrupted or overwritten. Snapshots are attractive because they do not take as much disk space as full volume copies. Snapshots only make copies of changes and create a logical tree, if you will, of all of the data so that customers can recover full volumes of data if necessary. Snapshots only require as much space as any data written since the last snapshot, typically less than 20% of the capacity required for full volume copies.

ESG Research did a study<sup>1</sup> focusing on mid-size (\$50-499 million in annual revenue) and large Enterprise customers (\$500 million or greater in annual revenue) and asked them how many of them were using their storage systems snapshot capabilities. ESG was surprised to find out that only 20 percent of mid-size and 34 percent of large Enterprise companies were actually using snapshot technologies. The number one reason by far that customers cited for not using snapshots was cost (54%):

Chart One: Why Customers Do Not Implement Snapshot Technology



<sup>1</sup> ESG Research: The Evolution of Data Protection, December 2003

These customers felt that snapshot copies were too expensive because they had to be created on the same disks as the primary data. Using the BlueArc Titan customers can replicate snapshot copies of primary data stored on FC drives onto lower cost SATA disks. This reduces the cost of copies and as such system administrators can economically justify making more snapshots for greater data protection. The use of SATA can economically provide greater data protection by enabling customers to make more snapshot copies and reducing the chance of data loss.

## Customer Discussions

ESG interviewed two customers of BlueArc that are using different tiers of storage within their Titan systems. The following is an overview of these discussions, why they chose the BlueArc Titan, how they are using it and what the benefits they are experiencing.

### Customer One

The first customer that ESG spoke with didn't want to implement a FC SAN because of the cost and complexity associated with these solutions. However, when analyzing different NAS products they were disappointed at the performance until they evaluated the BlueArc Titan. They have three BlueArc Titan systems supporting approximately 40 TB of storage of mixed FC and SATA disks. The customer stores all of their production data onto the FC disks that can have up to 1,000 servers accessing the data concurrently. After a project is completed they move the data from FC disks to SATA.

This same customer has also begun to provide its end users with an option for FC or SATA for production data. If the customer chooses SATA then they will be allocated more capacity than if they picked FC. The trade off for them is capacity for performance.

### Customer Two

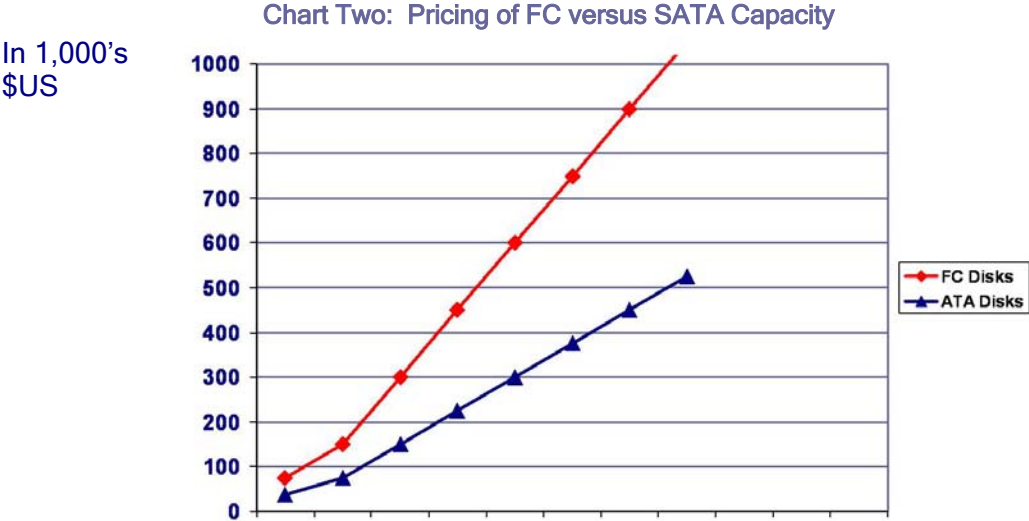
This BlueArc Titan customer has all of their production data stored on FC disks and uses the SATA drives as a disk-to-disk backup target. The FC disks are used for production data and the SATA disks are used to archive historical information. The tiered storage allows the customer to lower the overall costs of backup and yet improve access time in the event that data needs to be recovered from archive.

ESG research has found that customers are extremely interested in doing disk-to-disk backup for faster backup and recovery performance, minimize tape management, reduce human error associated with tape management, and perform backups on RAID protected storage.

BlueArc Multi-Tiered Storage allows the customer to align their storage capabilities with its overall business needs and provides an effective alternative to some portions of traditional tape archiving. "I have bought my last tape library," was his comment. They are keeping data retained on disk for fast recovery of data. They do use tape for off-site archive but he feels that at some point this may go away as well.

# The Economics of Tiered Storage

What is the financial impact of using SATA disks in conjunction with FC drives? To determine this let us assume that FC disks are 50% more in price than SATA. We can then draw two lines that show the relative cost difference. If a customer spends \$15,000 on FC drives they will pay \$7,500 for the equivalent capacity for SATA disks. As customers spend more money on disk capacity, the pricing becomes increasingly compelling. If a customer spends \$200,000 on FC capacity they will spend \$100,000 less for the equivalent SATA capacity. If you spend \$1,000,000 on FC then you will pay \$500,000 for SATA. The economics of SATA storage are straightforward and extremely attractive.



There are other reasons to buy multiple storage systems other than capacity. Customers may wish to have redundant systems for data protection or require more throughput for CPU-intensive or clustered applications. However, as customers implement different tiers of storage and want to use SATA for archival and backup data the additional capacity becomes extremely compelling. The above example is designed to show customers the impact of using higher density SATA disks and how they can strategize using the BlueArc Titan. Customers are using faster FC disks as Tier One and SATA as a lower Tier. This mixture of drive types leads to the best of both worlds combining performance, cost and capacity.

## Conclusion

---

Clearly the economics of network storage have to change. Storage budgets are still flat and data storage growth rates just continue to climb. Additionally, customers are retaining data for longer periods of time adding even more capacity to the data center. The use of SATA disks within network storage systems can significantly drive down costs. It is also essential to be able to manage data movement between different tiers without negatively impacting users and system administrators. The use of iSCSI SAN and NAS to provide shared network storage also lowers cost since specialized FC adapters and switches used in Fibre Channel SANs are not required to network storage.

The BlueArc Titan brings a lot to the table as a high-end IP-based storage system. The Titan has a scalable file system that supports up to 256 TB of multi-Tiered storage. Compare this to the leading NAS solutions from NetApp and EMC that support 16 TB in a single file system. Additionally, some other leading NAS vendors require customers to buy a separate system to add SATA storage to your storage network. The BlueArc Titan is known for its high performance and customers often choose it over the competition for just this reason. Additionally, the BlueArc Titan supports important software features including snapshots and remote replication.

Scalable file systems, performance and data protection are very important but the conversation shouldn't end there. What other capabilities, tools and features does the storage system bring to the table? Intelligent multi-tiered storage can go a long way to reducing costs in any environment. Whether it can save you \$10,000 or \$1 million, it is money that can be used elsewhere. Having more capacity at a lower cost can enable you to increase protection levels, retain data for longer periods of time, allow you to offer new services, perform disk-to-disk backup, etc. The BlueArc Titan multi-tiered storage offers unique functionality that customers should consider when determining how they are going to meet the challenges of rapidly growing storage with no additional budget.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482-0188.