Multi-Tiered Storage - Consolidating the differing storage requirements of the enterprise into a single storage system

By Louis Gray

BlueArc Corporation
Corporate Headquarters
225 Baypointe Parkway
San Jose, CA 95134
USA
info@bluearc.com
Voice 1-408-576-6600
FAX 1-408-576-6601
http://www.bluearc.com

Executive Summary

Performance has always come at a steep cost in the world of enterprise storage, with flexibility often taking a back seat when it comes to IT purchasing decisions. In fact, storage has typically been bought and deployed as a "one size fits all" solution, regardless of the kinds of applications the network is running. Such an approach has required multiple servers to address multiple needs, resulting in significantly higher costs and a much greater degree of management complexity.

BlueArc's Multi-Tiered Storage (MTS) solution changes all this. MTS offers storage performance and consolidation while supporting different types of storage within the same network-attached storage (NAS)-based system, according to the specific requirements of the applications. For the first time, a combination of high-performance online, moderate performance nearline and infrequently accessed archival data can be configured in a single, seamless NAS platform—a BlueArc SiliconServer.

This paper looks at the issues surrounding today's growing storage needs in enterprise and project applications, examines the need for a multi-tier storage solution, explains BlueArc's Multi-Tiered storage platform and demonstrates how the system is applied to storage applications.

Today's Issues

Online data is doubling each year, but disk prices are falling rapidly. Now storage access, infrastructure and software costs are the dominant factors in the rising cost of storage.

Today's business need for scalable storage broadly fits into two categories - enterprise and project storage. In both of these applications, storage needs are growing exponentially.

Enterprise Storage

Enterprise user home directories are increasing more than two-fold annually, as presentations and documents become more graphics rich, email traffic increases and as employees need to keep more data at their fingertips to help their company prosper in a highly competitive and demanding marketplace. For example, it was not too long ago that a 2 Megabyte PowerPoint presentation was considered large, while now we see large presentations approaching 10 Megabytes apiece, and the average size nearing 2 Megabytes. Factors such as this, combined with employees' needs to retain additional storage in their home directories and mail archives is driving the average enterprise IT manager to prepare for a four-fold increase in storage space simply to keep pace with the organization's needs.

Project Storage

Work on digitized feature films, particle physics projects, bioinformatics/genetic research, seismic research, digital design, broadcast, medical imaging, CAD and other projects already require very significant levels of storage. Business demands and new enabling technology (e.g. low-cost Linux cluster super-computing) are driving for faster project completion and the analysis of even more detail, driving the need for performance network access to multiple terabytes of storage. The recent adoption of low-cost ATA disk arrays in many data centers offers affordable online storage or caching of archive data. In many applications this brings strong productivity benefits. Immediate online availability of archived data enables it to be quickly searched and re-used, rather than abandoned on low-performance archive tapes. In a typical enterprise, this combination promotes growth beyond tens of terabytes of online storage.

Infrastructure costs are becoming dominant in scaled storage applications

Attempts to deal with this level of growth are stretching current storage architectures to the limit. While Fibre Channel storage arrays are still a significant cost factor in expanding storage needs, the simple cost of storage arrays themselves is now not the prime contributor to growth costs. The infrastructure costs surrounding storage expansion (backup strategies, server proliferation and administration costs) are poised to become highly dominant factors for growth.

24 by 7 access to data

24 by 7 access to storage is being threatened by extending backup windows. Elimination of backup windows is the much-needed solution.

A difficult issue that enterprise and project IT managers face is the business mandate for 24 by 7 access to stored information. In the face of expanding storage needs, IT managers would like to be in a position to extend backup windows to accommodate for increased growth rather than contract them, enabling administrators to scale each server further, rather than proliferating servers. But backup activity cannot spill into peak business hours

if acceptable service is to be maintained. Traditional server architectures cannot run backup copy activity and continue to serve files without a severe drop in performance, which only becomes more difficult as service levels remain high around the clock.

Additionally, fast restore times and disaster recovery are key – the cost of downtime to businesses is high and IT managers need to strike the right balance in this area of the IT strategy, devising what amount of history to keep online and where it should be stored.

The introduction of large low-cost ATA storage arrays has changed the landscape for data restore and disaster recovery. With large, low-cost storage as the target for backup data and prime source for immediate backup history, both backup and restore time are reduced when compared to traditional tape backup. While online disk storage does not replace the high integrity holding data copies off site, disk storage offered as "virtual tape" or replicated intermediate storage is highly beneficial for reducing backup times and restore times. Savings and benefits from low-cost storage must not be competing with a higher cost, more complex infrastructure to access it.

Changing the rules for Storage

Today's issues demand a different architectural approach to storage.

- The infrastructure costs associated with scaling storage must be reduced.

  Enterprise and project storage needs will continue to expand. Raw storage costs are reducing significantly, but infrastructure costs associated with increasing storage capacity remain high. Proliferation of servers, backup and restore times, backup costs, storage networking complexity, and the administration overhead associated with storage infrastructures need to be tackled and contained.

- Match storage to applications without imposing large infrastructure costs.

  The arrival of resilient arrays of low-cost and high capacity ATA disks heralds the introduction of staged backup, which reduces both backup and restore times and lets servers scale much further. Low cost of storage and very high capacity are key elements for this application, since even the lowest performing disks are significantly faster than tape access. ATA storage fits this need. However, the performance characteristics of this type of storage are not high enough for database and transaction intensive applications. In these cases, Fibre Channel disks deliver appropriate performance at a higher cost.

- 24 Hour availability.

  Set-aside backup windows are no longer acceptable in today's around the clock businesses. They must be removed or significantly reduced, requiring a storage platform that can be backed up without affecting user access performance, even in

peak activity periods. Replication options must also operate without affecting user activity. Restore capabilities need to be very fast and disaster recovery needs to be quick and simple. Access to storage needs to be highly available, but redundant path architectures deployed to achieve this must be uniform across all types of storage without increasing infrastructure costs.

- Scaling by server proliferation is too expensive.

  Performance limitations in traditional servers restrict the amount of storage supported by a single server. Physical scaling limits for direct or SAN-attached storage cannot ultimately be avoided, even by offloading backup from the server. NAS appliances and NAS gateways extend these limits and are becoming more important as installed levels of storage continue to expand. Traditional NAS is more cost effective but lacks full integration

The role of NAS is changing. Switched Gigabit Ethernet as a network backbone expands the role of NAS in today's networks, while resilient network switch architectures support vast bandwidths with high availability. A Gigabit-switched backbone can support user traffic, application server disk access and archive traffic with ease. Furthermore, network traffic can be shaped and prioritized with standard tools in the Gigabit backbone switch, should the need arise.

With all this functionality included, the cost of commissioning a second and independent storage network needs careful thought. By removing file and archive storage from the SAN, server and SAN networking costs can be saved. For this reason, NAS has emerged as the optimal architecture for primary file storage and for online archive applications, which the availability of low-cost ATA storage has enabled. NAS also challenges the role of SAN as prime storage for database and other application servers.

Traditional NAS solutions have dedicated functionality. Performance NAS filers support primary file storage, and cost-competitive systems support general storage needs with the option of dual server heads offering high availability and fast fail-over in the event of a network or NAS component failure. Slower, cost-effective single head NAS systems are used for archive applications, built to support low-cost ATA storage generally in blocks of 2- 12 terabytes.

For a customer to require a combination of archive storage and primary storage is quite common, but to further reduce costs, an integrated approach, with all storage types supported in a single NAS system, has been sought but has not been found, as regrettably such an approach would demand unsustainable scalability and performance from the software-based architecture on which traditional NAS filers are based. The approach would also introduce performance issues related to backup, since is it hard to conceive that tape archive copying would be completed during today's shrinking backup windows. Decreased user access performance when running archive traffic during business hours adversely affects business productivity.

BlueArc - Scalable, high performance Multi-Tiered Storage with low Infrastructure costs

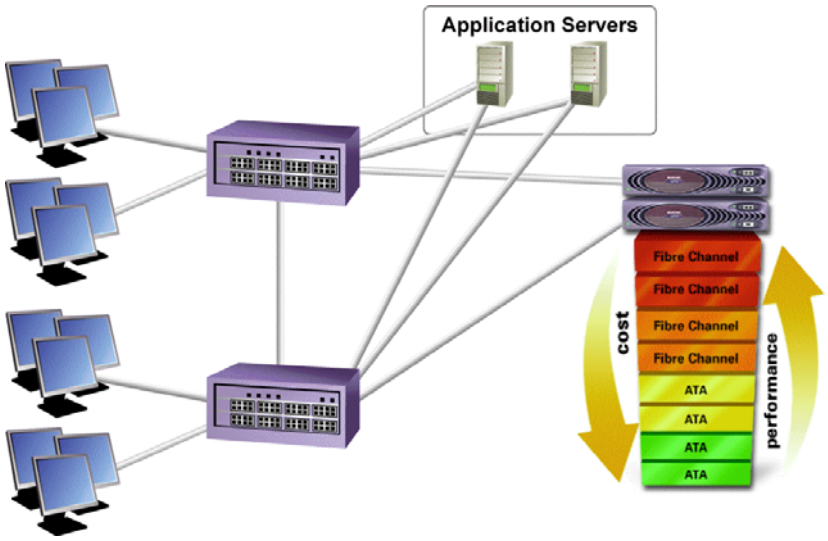

BlueArc Network Attached Storage is based on the SiliconServer, a NAS head with hardware data movement, similar to network backbone switches. BlueArc's architecture has the performance and scalability to support an integrated storage solution, with hardware data movement to maximize throughput.

BlueArc also solves the backup traffic problem, because of the high internal bandwidth in the server and because movement of data between network and storage is prioritized over disk to disk and disk to tape activity, should contention for resources arise.

BlueArc's Multi-Tiered Storage, with SiliconServer scalability and performance, consolidates storage to a single system. The system changes network storage rules by:

- Reducing the infrastructure costs of scaling
- Offering multiple tiers of varied cost and performance storage, without imposing large infrastructure costs
- Reducing the number of servers through consolidation and replacement
- Enterprise class continuous availability storage

The enhanced productivity derived from MTS is seen by customers in fields such as Broadcasting, Post Production, Manufacturing Genetic Research, Government Research, Mapping, University projects and many more.

The use of hardware to move data within the SiliconServer is extended to backup and disk-to-disk copies over Fibre Channel. The SiliconServer uses hardware data movement in the implementation of NMDP, the standards based backup protocol for NAS. Coupled with the ability to schedule a snapshot of the file system prior to making a backup, this capability has three distinct advantages:

1. The traditional problem associated with slowed server performance when running backup activity is removed and there is no need for a backup window.

    - Network data movement is prioritized over NDMP traffic
    - NDMP is supported by all backup software applications
    - Backups run faster because of speed of hardware data movement

2. Tape drive efficiency
    - Backups of a frozen on-disk snapshot can be spooled to tape as a background activity, if necessary running 24 hours per day to get maximum scalability out of tape drives. This delivers two to three times the scalability from each tape drive compared with running the drive only during a backup window. This mode of operation is fully supported by standard data backup software applications.

3. Backup copies can be stored online for immediate access

    - Fast restore and disaster recovery by referencing an online snapshot of the backup, rather than referencing the archive tape. SiliconServer file system snapshots store a complete image of the file system at the time the snapshot was taken. The snapshot retains a block level image of the file system. As files subsequently change, the snapshot stores only the blocks that have been modified since the snapshot was taken. Snapshots are highly efficient in terms of the amount of disk space they occupy, therefore it costs little to retain a week or more of daily backup copies online for fast restore should it be necessary. Backup copies on disk are not only highly desirable for business continuance in the event of a disaster or accidental deletion, but they also reduce the number of required tape library tape slots - if backup history is held online, there is no need to leave recent backup tapes online in the tape library.

Shared File Storage

BlueArc's Multi-Tiered storage is mounted as a network share/export. This means that multiple application servers can share the same data using standard Windows and UNIX locks. The SiliconServer also supports secure sharing between UNIX and Windows environments, with permissions mapping and lock integration between the two environments. Data sharing is vital for imaging, Web applications, design, engineering, and research applications where centralized data needs to be shared between a number of application servers and clients. BlueArc's Multi-Tiered Network Attached Storage is the optimal, strategic choice for these applications.