IDC TECHNOLOGY SPOTLIGHT

Efficiently Delivering Enterprise-Class File-Based Storage

December 2009

Adapted from *File Systems Define the Future for Managing Storage* by Noemi Greyzdorf and Richard L. Villars, IDC #215463

Sponsored by GreenBytes

Unstructured or file-based data will continue to create new challenges for organizations. In light of the need to share data among multiple entities, organizations responsible for larger amounts of file-based data are facing the need to make this data available and accessible and do so in an efficient manner. This means managing infrastructure, acquisition, and maintenance costs without sacrificing security and performance.

Specifically, organizations with large amounts of file-based data have the following storage requirements in order to achieve efficiencies:

- *Capacity optimization use to lower the data footprint necessary for storage (This impacts both disk consumption and datacenter resource consumption such as facilities and energy.)*
- *Security for stored files and security from unlawful access to data as well as securing the availability of data over time*
- *Scalability without creating additional complexity to manage the environment*
- *Ability to manage information assets and not just data based on business requirements and established retention parameters*

This Technology Spotlight explores the challenges facing organizations with growing amounts of data and discusses the new technologies that address the associated storage needs. In addition, this paper highlights the role that GreenBytes plays in this emerging market.

Solving the Storage Efficiency Challenge

The advent of the Internet has changed not only how we do business but also with whom. The globalization of markets has increasingly placed emphasis on information sharing among customers, vendors, and partners. The emerging information age demands application availability, access to data, and the protection of information assets. More information is being created and stored in digital format to comply with market demands. Additionally, improvements in technology are creating larger data sets, while regulatory and governing bodies are demanding that data be retained for longer periods of time. These trends, coupled with the down economy, are creating new challenges for organizations that are struggling to accommodate more data with existing or smaller budgets. The main challenges facing datacenters include:

- Managing more data and associated storage with existing administrative resources
- Improving resource utilization to contain exploding demands on storage resources

- Reducing operational complexity and achieving greater data availability, integrity, and reliability
- Improving the storage provisioning process and reducing management complexity of storage when deployed with server and desktop virtualization

To achieve these goals, organizations are looking for processes and technologies that deliver greater overall efficiency. Some approaches being implemented to achieve greater efficiency in the datacenter are as follows:

- Deploying server and desktop virtualization to address the operational complexity of server and desktop management, to improve physical resource utilization and reduce the datacenter footprint, and to contain costs
- Deploying management tools to gain visibility into information assets and storage environments that provide intelligence about information creation, access to information patterns, and storage resource allocations to various information assets
- Leveraging innovative approaches to storage infrastructure such as capacity optimization, which includes compression, single instancing, deduplication, thin provisioning, dynamic storage tiers, and snapshots
- Implementing persistent storage architectures that deliver higher levels of reliability, accessibility, and data protection while offering a more flexible architecture to accommodate changing performance and capacity requirements
- Augmenting existing resources with offerings from service providers to reduce idle resources or accommodate spikes in demand

When selecting an approach or a technology, organizations need to consider any solution as part of the greater ecosystem; often, the decision to implement one technology can have negative implications in other areas of the environment. Therefore, organizations must look at problems and solutions holistically; the solution selected must support the existing and emerging ecosystem of infrastructure technologies and applications. Some examples are:

- Though virtualization of servers delivers a faster way to provision resources for new applications and improves server resource utilization, true operational benefits can't be realized without addressing storage requirements. Virtualization does bring some complexity to the storage side of the equation, such as performance load balancing, timely provisioning of resources, and greater demands on capacity.
- Thin provisioning may prevent organizations from allocating and provisioning storage resources before there is demand for them and therefore delay or eliminate unnecessary purchases, but it doesn't work in all environments. If capacity growth is unpredictable or if there is a need for a greater spindle count, thin provisioning won't be applicable.
- Selecting the right approach to capacity optimization is also important. Not all approaches deliver the same results in all environments. Using capacity optimization against the wrong data set can result in performance degradation without the desired reduction in data footprint results.

Benefits

The thoughtful implementation of new approaches and technologies that can reduce complexity and costs and improve resource utilization and operational efficiency has enormous benefits to an organization's ability to address the emerging challenges of managing information assets and storage resources under a tsunami of data being created and stored. The benefits are both operational and financial, and they are often interdependent:

- Reducing infrastructure complexity not only lowers the cost of resource acquisition and support but also decreases the management overhead required to maintain these resources.
- Improved resource utilization through implementation of persistent architectures and virtualization not only reduces the capital investment in storage but also decreases operational complexity and extends the life of the infrastructure.
- Deploying management tools to gain visibility into information assets and storage resources enables administrators to make intelligent decisions regarding data retention, storage, and resource planning, resulting in less waste and a better service level to the organization.

GreenBytes Product Profile

Ashaway, Rhode Island–based GreenBytes Inc. was founded in 2007 to address the challenges associated with the growth of data coupled with the need for greater energy efficiency. In an effort to address the emerging storage challenges facing enterprises, GreenBytes has designed a storage system that offers an architecture and features that deliver storage efficiency and performance. GreenBytes' system is designed to deliver access to storage resources via file protocols, NFS and CIFS, and iSCSI. The internal architecture is laid out to support inline capacity optimization without performance degradation while providing storage capacity in a compact, energy-efficient form factor. Additionally, GreenBytes has integrated with Citrix StorageLink to simplify the provisioning of virtual machines in the Xen and Hyper-V environments together with necessary storage, making the beginning-to-end process more seamless, efficient, and timely. This tool allows storage provisioning to align with virtual machine provisioning, bringing the efficiency proposition of virtualization full circle.

System Design

The system has the following layers, each contributing to the overall function of the system:

- The **file system** layer provides the directory structure for storing data in file format. Data is accessed via NFS or CIFS through this layer.
- The **object layer** is responsible for mapping objects to blocks. Objects are constituted of ACL, files, directories, metadata, file attributes, and block devices (volumes). Block-level access via iSCSI is delivered directly from this layer without going through the file system.
- **Adaptive replacement cache (ARC)** is the next layer and caches both reads and writes. All transactions, read and write, are performed through ARC. If an object is not in cache at the time of a read request, the data is first written into cache and then served to the requester. This allows for data with the highest probability of being accessed again to be placed into cache. This is a critical component of the architecture that delivers performance and efficiency with writes and reads. This is achieved as follows:
    - All writes are first cached in ARC, which allows the system to aggregate all the writes and write data to disk based on IOPS requirements. Aggregating all writes also eliminates the random nature of some transactions and the dependence on spindle count. In this approach, all writes become sequential and only the throughput of the drives remain as a performance determinant. The benefit is the ability to deploy the system with slower-spinning, higher-density drives that consume less datacenter space, require less energy to power and cool, and deliver greater capacity per dollar.
    - The current systems available on the market come with either 12GB or 24GB of cache, but the architecture is not limited to this cache size. It can be scaled to the size of the 64-bit addressable space, which offers the architectural flexibility to alter hardware components to suit the specific workload being placed on the system.

- When caching writes, data is written to open cache, then it is prepared to be emptied, and then the writes are committed to disk. In a read cache scenario, the system has two levels of cache. Level 1 is RAM. This is where the hottest data is placed for read access. To compensate for the limited capacity of RAM and yet accommodate all reads that can benefit from being cached, there is a Level 2 cache that can be any low-latency media such as SSD. There is a "heartbeat" between the two levels of cache so that as data cools off a bit on RAM but is still being accessed, it can be moved down to the SSD. The Level 2 cache can be as big as a terabyte. This allows the system to handle a large volume of requests for the same data set simultaneously without performance degradation. Once a request for data comes in, it is placed into RAM and stored in the SSD until the demand for it cools off. One application of such an architecture is to support virtual desktops at the time of system boot, commonly referred to as the boot storm. Another use case for such an architecture is when many requesters are accessing the same piece of data in parallel. The data is always served from cache.

- **Pools** are next in the stack. These are RAID sets; there is typically one pool per system consisting of multiple RAID sets. RAID sets consist of 24 drives where there are two allocated for parity and one as a global space. The data is striped across all the drives. Pools are dynamically expandable, allowing the system to scale with the demands of the application. The expansion of the pool doesn't create downtime or require a reboot of the system. A typical pool can consist of multiple RAID sets, each RAID set can vary in size depending on the hard drives used. As drive size increases, additional RAID sets can utilize higher-density drives for more capacity. All drives within a RAID set have to be the same size. (For details on GreenBytes technology, see Figure 1.)

Figure 1

GreenBytes Flow Diagram

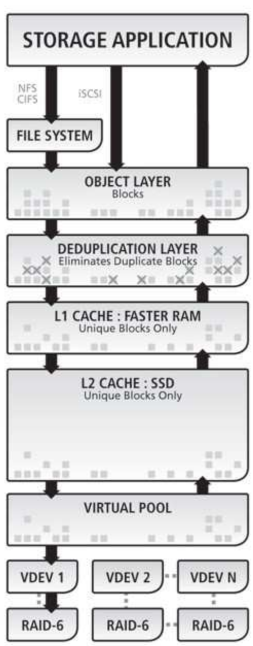

Source: GreenBytes Inc., 2009

Capacity Optimization

The architecture designed by GreenBytes also supports inline deduplication of data at the file system layer as well as at the iSCSI block layer. This is where the design of the system allows for deduplication to be applied to all data, inline, without affecting the overall performance of the system. The result is that the system optimizes capacity for production environments running live applications as well as for more secondary tiers of online storage. GreenBytes achieves this by designing deduplication into the overall architecture of the system; deduplication is as much a component of the system as ARC.

- Deduplication occurs between the ARC and pools. A cryptographic-quality hash is applied to each incoming block of data to create a unique identifying fingerprint of the block. GreenBytes' hash algorithm is optimized for commodity processors and is computationally very fast, sustaining network line speeds.
- The data structures containing the hash codes, dcache, are stored in both conventional cache, as well as on disk. The dcache data is stored on disk and mirrored for fast access in the ARC/L2ARC. On startup, the appliance quickly warms up the the ARC with dedupe entries so that the system is immediately ready for high performance data ingest.
- All incoming blocks' hash codes are compared against all others contained in the dcache for deduplication operations. GreenBytes' patent applied for technology is unique in that the search operation is again computationally very efficient, and is agnostic to the size of the dcache. In other words, GreenBytes' deduplication engine scales to very high capacities, hundreds or more terabytes in a single namespace.
- Deduplication is global across the pool of storage, and simultaneously supports both NAS filesystems and iSCSI SAN volumes. For iSCSI volumes, optimal deduplication characteristics are achieved when the block size of the volume is harmonized with the file system writing to it.

Product Benefits

GreenBytes has designed a system that addresses the emerging challenges of storage administrators. It eliminates lock-in to a file versus block storage device by offering access via NFS, CIFS, or iSCSI. In many midmarket organizations, this can offer capital as well as operational cost savings. The intelligence of the system is hardware independent, allowing the hardware configuration to change with the demands of the market. It is currently deployed on an x86 platform with standard components, using 2.5in. drives for higher density and power efficiency. Though the current hardware configurations can support up to 260TB, this is a limitation of the hardware connections, not of the architecture. The result is a system that can be flexible to support the evolving demands of the business. For example, an increase in the number of cores directly increases the number of threads the system can support, increasing the overall performance of the system.

Another example of how this architecture benefits storage administrators is in its use to support desktop virtualization. At the time employees boot their systems, often referred to as boot storm, there is enormous demand on a storage system for IOPS. The use of ARC supports boot storms without having to deploy excess capacity just to achieve the necessary number of spindles for the required IOPS. The efficiency of this approach directly translates into savings in capacity costs as well as the overall datacenter footprint.

In addition to the design of the solution, GreenBytes has identified a use case where its technology helps address administrative challenges associated with server virtualization. The tight integration with StorageLink Manager, coupled with the architecture of the system, delivers a simple, consistent way of provisioning virtual machines and associated storage. This is achieved by leveraging templates that can be created by administrators who have all the storage parameters already defined.

Challenges

The market has focused on addressing the need for file-based storage in the enterprise with features that lower storage costs by increasing utilization of existing resources, simplifying management of the systems, and delivering greater efficiency through capacity optimization, dynamic storage tiering, and thin technologies. Unfortunately, there isn't a standard way of evaluating and comparing these options for the best desired result.

GreenBytes has the technology to address the existing demand for capacity optimization in the data protection and disaster recovery space. In addition to data protection and disaster recovery, GreenBytes' technology is well positioned to address other types of application workloads with more performance-oriented requirements. As end users gain experience with this technology and opt to leverage it with other applications in their environment, GreenBytes will need quickly adapt to these new transitions.

Conclusion

Market dynamics are driving enterprises to demand more from their IT. Enterprises want a more efficient datacenter so that they can deliver increasingly higher levels of service without incurring significant costs. This can be achieved only through the adoption of processes and technologies that help create operational efficiencies as well as improve resource utilization and lower the total cost of ownership. Operational efficiency can be achieved through simplification of management, increased visibility into the environment, and a more automated way to manage resources. Costs can be cut by deploying innovative technologies that more intelligently utilize the physical resources available. This includes virtualization, capacity optimization, information governance, and dynamic storage tiers.

GreenBytes is responding with a clean-slate storage architecture that has the promise to address a variety of emerging challenges.

A B O U T   T H I S   P U B L I C A T I O N

This publication was produced by IDC Go-to-Market Services. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Go-to-Market Services makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.