

Introduction

When it comes to the increasingly complex task of managing data storage, many small and midsize organizations face even greater challenges than large, global enterprises due to relentless growth in business information that needs to be managed and secured. In a recent survey by Storage Magazine, about one-fourth of the respondents pegged their unstructured file growth at 50% or higher year-over-year.

The consumerization of IT with fading boundaries between personal and workplace computing along with the unstoppable march to a virtualized IT infrastructure is pushing IT personnel to find new ways to manage their storage infrastructure. Some of their biggest hurdles are:

- Scaling storage as the business grows rapidly
- Containing the rising expense of data storage capacity
- Dealing with the complexity of storage management and architecture
- Devoting precious staff time efficiently to manage storage and data backup

With data volumes growing faster than anticipated, legacy scale-up storage architecture with two controllers sharing the workload and providing failover just doesn't stack up today. In a scale-up storage architecture, disks and shelves are added to expand capacity, only memory may be added to scale performance. Eventually, and often rather quickly, a scalability limit is reached, and the only available options are to either add another storage system or do a forklift upgrade and replace the existing system. The consequences are sprawling storage silos that are complicated to manage with no investment protection that ultimately leads to spiraling costs.

In addition, small and midsize organizations that are looking to implement storage solutions with enterprise-class features to meet their business needs face complicated installation and management challenges. And in most cases, they are not equipped to handle these difficulties due to IT resource constraints and a lack of appropriate in-house expertise.

Fortunately, there are new approaches to storage on the market for growing organizations that can help address their storage needs without requiring dedicated management resources while at the same time driving down their costs.

A New and Effective Approach to Storage

Exablox has designed a unique storage solution from the ground-up for small and midsize enterprises to help address their escalating data storage needs and deal with runaway storage costs and information management nightmares. Exablox offers OneBlox, a scale-out storage solution that is an on-premises hardware appliance combined with OneSystem, an integrated, enterprise-grade cloud-based storage management service.



OneBlox is a purpose-built appliance with all enterprise-grade features built-in such as continuous data protection, compression, inline deduplication, and disaster recovery and is completely automated for the IT personnel to setup and operate. The unique scale-out, converged storage appliance consolidates primary data as well as the backup and archival data.

Eliminate Complexity


OneBlox is plug and play scale-out storage with the ability to “bring your own disks” to the appliance – paying retail pricing rather than a 2-5X vendor premium. OneBlox provides additional flexibility by allowing users to mix-and-match drive types (SAS, SATA) and capacity within the same OneBlox and within a ring with zero-configuration. Drives can be removed, recycled, reused, and replaced dynamically while the file system restores the desired levels of data protection automatically.



The engine that drives OneBlox storage appliance is a distributed object-based file system, which provides ubiquitous data access by presenting a CIFS/SMB or NFS file share to users and applications. Object level replication protects data from disk or node failures while eliminating the need for legacy RAID protection schemes. Provisioning is instantaneous since there is no RAID, volume or LUNs to configure. Data replication is in

real time and by default, in triplicate. If there is more than one node in the local cluster called a “ring”, it ensures that at least one of the replicates is on a different node. Failed disks (or nodes) can simply be removed and replaced with no disruption to data services to the application and users — OneBlox will dynamically re-replicate all the objects on the node as needed to maintain the appropriate protection level. There are no vulnerable or offline rebuild windows as with legacy RAID approaches and no worrying about LUNs or volume sizes.


Scale Effortlessly



Multiple OneBlox appliances can be aggregated into a single clustered ring by simply powering them up on the same LAN. They automatically discover the other OneBlox appliances and join the ring with zero-configuration. Once nodes are clustered, not only does OneBlox automatically redistribute any data objects across the ring, it presents all the aggregated storage capacity in the same global namespace. When either a disk or a new node is added, the current network share expands automatically and non-disruptively to include the new capacity. There is no need to manually manage storage pools, RAID sets, manage volumes or migrate data to consume the increased capacity.

Optimized Storage Capacity


OneBlox includes inline data deduplication and compression to dramatically improve storage utilization that helps lower costs. It applies deduplication for all written data with no questions to answer, no radio buttons to click and no management. The deduplication is performed across the entire ring or global namespace, and compression to user-selected shares which significantly reduces the total storage space required.



No-Compromise Data Availability



In addition to the ring level replication, OneBlox provides advanced continuous data protection in the form of automatic snapshots. At any point, an immediate online recovery of previous versions of a file can be achieved without time-consuming restoration from backups. Users can easily find and recover files themselves by navigating through a dedicated “Snapshots” folder through Mac Finder or Windows Explorer. The snapshots are deduplicated against the primary data as well as compressed to keep the storage requirements low, are directly accessible and cannot be modified.



Exablox also offers cost-efficient, real-time, cross-site replication for disaster recovery to protect against failure of an entire ring or site. Since the data is already compressed and deduplicated with only incremental changed objects needing remote replication, WAN usage is optimized.

Simple Storage Management

Exablox provides OneSystem – an innovative, multitenant, cloud-based management service – that eliminates the need for dedicated servers and software on the customer premises to manage storage. OneSystem allows for ubiquitous accessibility and simplified storage management workflow. With OneSystem, there is no need to install local storage tools or configure VPNs for remote management access.

OneBlox administrators simply log into OneSystem from any browser and manage their on-premises OneBlox systems. Management workflow is visual, point-and-click, drag-and-drop and directly actionable. An intuitive browser interface is used to configure new shares, users, groups, integration with Active Directory, and setting up disaster recovery. OneSystem proactively tracks and monitors the usage, capacity and health, and issues alerts if necessary.



Increase Flexibility

According to Forrester, one of the greatest culprits behind data growth is actually disk-based backup data. Between 2010 and 2012, the average enterprise server backup data store grew by 42%, archival data grew by 45%. With more and more mobile workers, PC backup storage is also growing at an explosive rate, almost 100% over the past two years. To contain this spiraling backup and archival costs, Forrester advises the IT organizations to look for following key criteria in their storage systems:

- Data reduction capability
- Continuous data protection
- Ease of setup and implementation
- Ease of on-going management

The Exablox storage solution unquestionably delivers on all of these metrics:

- Compression and Inline deduplication keeps the effective capacity growth in check
- Scale-out capability helps to grow capacity dynamically as the data grows
- Continuous data protection safeguards the “last-resort” data, and
- The solution has been designed from ground-up for simplified management workflow

In fact, it can consolidate both the primary/ production file data as well as the backup/ archive data within a single cluster depending on the organizational needs. The continuous data protection capability eliminates the need for frequent full-backups of the data stored on the OneBlox ring. The compression and inline deduplication helps keep the storage footprint in check when using the OneBlox ring as a target for backup/ archive of the production data from the existing storage systems in the environment.

