

STORAGE SWITZERLAND BRIEFING REPORT

HIGH PERFORMANCE AND DEDUPLICATION FOR PRIMARY STORAGE



George Crump, Senior Analyst

Primary storage is the next great frontier for deduplication, but unlike backup this tier of storage has a unique set of demands. A primary storage system with deduplication has to perform well, has to maintain high availability and it has to scale. The [GreenBytes HA-3000](#) is one of the first integrated storage systems that meets these demands, delivering the capacity saving economics of deduplication in a high performance and highly available system.

Primary storage must deliver an economical solution based on three criteria, like the sides of a triangle: High Availability, Capacity Efficiency and Performance. Each of these sides is critical when selecting a solution that can meet the demands and budget realities of the modern data center.

HA Economics

The first side of the triangle is the requirement for a highly available (HA) storage system. The HA-3000 is an IP-based block storage system that can present iSCSI volumes. The unit features dual controllers powered by two Xeon LC5528 Quad Core Embedded CPUs for 8

threads of processing. It also has 24GB of cache per controller, each of which has a Quad 1GbE and Dual 10GbE connections.

HA is of course critical for most primary storage installations. The secret is in how that availability is delivered to the customer. Many NAS or universal storage systems use an expensive dual headed design, which means two controllers share access to the same set of disks. The secondary controller is simply there as a stand-by in case the primary fails. Not only does this method cost more it also doesn't fully leverage the performance capabilities of the two heads.

The HA-3000 has the capability for an active-active controller architecture. Each controller is active to its own pool of storage, while providing failover capability for each other. What this means is that for a small additional cost, since a whole processing/controller head does not need to be purchased, the customer gets HA, while increasing IOPS and supported capacity. In other words the system is faster, more scalable, more reliable and only a little more expensive than a non-HA system.

Capacity Economics

The second side of the triangle for primary storage systems is capacity efficiency, or the ability to pack more data into each TB of physical storage. A key challenge facing storage managers today is how to provide primary storage without breaking the IT budget. GreenBytes has features that reduce the effective cost per GB and help meet this challenge. First, the HA-3000 provides inline deduplication, meaning that redundant data is identified and eliminated prior to each write cycle. The holy grail for primary storage systems is to provide this deduplication without performance loss. Being inline means that all unnecessary writes are eliminated upfront and don't need to be cleaned up later. This makes the storage system easier to manage since you don't have to deal with a deduplicated and non-deduplicated state. It also saves wear and tear on the disk media. Most importantly it reduces the data amplification of an unneeded write.

Data amplification is a term for additional copies of data segments that are generated from processes subsequent to the original write. First, depending on RAID level, one or two parity bits must be created. Second, if a snapshot or clone of a volume is taken all that redundant data has to be tracked and copied out as it changes. Third, in the event of replication all that data is sent to a remote site. Inline deduplication stops this amplification at its source. The original write, if it's redundant, never occurs. The net impact is that deduplication, depending on the environment, should produce anywhere from 5X to 15X in improved capacity efficiency.

Deduplication is not enough for a fully optimized infrastructure because it only achieves its efficiency if there is redundant data. While there is plenty of redundancy, especially in virtualized environments, there is also some data that's not redundant. And of course there are the original data segments that the deduplication tree is built from. All of these data are candidates for compression which is also included in the GreenBytes HA-3000. By

compressing the data, the HA-3000 gains another 2X in capacity efficiency.

The HA-3000, like other unified storage systems, provides thinly provisioned volumes. A challenge is that when those systems are used as block storage (iSCSI) they often can't reclaim unused space. In other words the storage may start thin but it will always "gain weight" as files are added and deleted. Most storage systems can't perform thin reclamation which is the ability to recognize deleted data, reclaim it and return it to the global storage pool.

The HA-3000 also has the unique ability to reclaim space on its thin provisioned iSCSI volumes. It can provide this support for Linux, Windows and VMware attached hosts. In other words as data moves through its life cycle on a connected host and is eventually deleted, the volume can be put on a diet and "re-thin" itself.

Primary Inline Deduplication, Compression and Thin Reclamation are key capabilities in the effort to control storage growth. Finding all of those features in a single system and not having them impact performance is very unique.

Performance Economics

The third side of the triangle for primary storage systems is the requirement to deliver high performance. The HA-3000 meets this need by leveraging three SSDs in the system, two of which are SLC-based for long life and high performance and the third is a high performance ram drive for write-back, non-volatile cache. The GreenBytes Hybrid Storage Architecture (HSA) uses SSD to provide high performance when needed, yet maintains SAS drive economics when it's not.

HSA first uses the SSD tier for meta-data staging and for scaling up the deduplication engine and write-back, synchronous I/O transactions. The use cases provide a significant improvement on storage I/O performance that many other storage systems overlook. The rest of the SSD storage is used as a cache to accelerate the most recently accessed data and serve it directly from this memory based storage. GreenBytes has shown to scale linearly as SSD and capacity is added to the system.

Ideal For Virtual Environments

The HA-3000 is an ideal solution for virtual environments. It has the required availability options, the capacity saving features and the performance that these environments need. An excellent example is the Virtual Desktop Infrastructure (VDI) use case. A typical virtual desktop needs about 15GB of storage and 1,000 desktops would need 15TB. While hypervisors have thin provisioning and cloning capabilities that should help reduce capacity needs, their use often comes at a significant performance loss forcing many IT organizations to stop using them. GreenBytes' low impact capacity optimization takes this responsibility away from the hypervisor allowing greater VDI density.

VDI performance is also an issue. Each mechanical disk can deliver about 120 IOPS, enough to sustain 3 to 4 virtual desktops each. At this rate, an environment with 1000 virtual desktops would require more than 250 disks. This means significant capacity would be wasted in order to generate the required performance, since 250TB (assuming one TB drives) is around 10x more capacity than is needed for the VDI environment. GreenBytes, with their prudent use of SSD, can sustain 100s of virtual desktops with a fraction of this spindle count.

Finally, all those virtual desktops are going to need a NAS to use as a file share. The GreenBytes HA-3000 has the performance, scalability and reliability to support both needs from a single system, and support a virtual server infrastructure as well.

Summary

Unified storage is not really unified if you need separate systems for server virtualization, desktop virtualization, OLTP applications and file sharing. Even if some of these functions could be combined, most environments will require two or three systems to satisfy the performance and capacity requirements they generate. The HA-3000 truly brings mid-range data centers one storage system that can do it all and do so economically.

About Storage Switzerland

Storage Switzerland is an analyst firm focused on the virtualization and storage marketplaces. For more information please visit our web site: <http://www.storage-switzerland.com>

Copyright © 2011 Storage Switzerland, Inc. - All rights reserved