REPORT REPRINT

Exablox scores in primary storage, aims for wider uses

TIM STAMMERS

24 DEC, 2015

The company's fast-growing sales are already proving the appeal of its low-cost scale-out NAS storage. Usage is currently dominated by backups, but Exablox says it has an expansion into that virtualization market firmly in mind.

451 Research®

 |

Exablox's storage is an unusual combination of an object-based scale-out foundation, with a front end or top layer providing native NAS-style file access. The company's mostly SMB customers appreciate the simplicity and low cost of object storage, combined with SMB and NFS support for existing applications. Near-continuous data protection and easy deployment and management are also part of the picture. The next stop in the company's roadmap is more usage of its systems to support virtualized servers. To get there, it will need to boost its random I/O performance. In the summer of 2015 it released second-generation hardware that it says is up to 10 times faster than previous Exablox gear, and it is working on a software update that will provide another performance boost. The company has also just improved the data de-dupe function in its systems, reinforcing their strengths as backup targets.

THE 451 TAKE

Putting a NAS front end on object storage appears an to be obvious way to overcome one of the biggest objections to object storage, which is that existing enterprise applications cannot use object protocols. However, judging by the scarcity of object storage with native NFS and CIFS support, it is not easy to create. Exablox appears to have done the job well, resulting in appealing, low-cost scale-out storage for backup, archive, and other large-file or non-performance-sensitive data. If it can deliver on its promise to extend that appeal into mainstream virtual server and desktop support, it will not only be addressing a bigger market, but also offering a single storage platform that suits multiple uses.

CONTEXT

Based in Sunnyvale, California, Exablox was founded in 2010, and has raised $45.5m in VC. During 2015, Exablox expanded its list of investors to include Dell Ventures and Toshiba America Electronic Components, which joined previous investors DCM, Norwest Venture Partners and US Venture Partners.

Total headcount is just below 70. Exablox CEO Doug Brockett is an industry veteran, and was previously VP and GM at firewall vendor SonicWALL, and VP of marketing at networking specialist Nexsi. Exablox's VP of engineering, Ramesh Iyer Balan, was previously senior director at Data Domain, where he led the de-dupe engineering team. We note Exablox's wide usage as a backup target that competes with Data Domain. Before that, Balan led file system development for Veritas. The VP of marketing at Exablox, Shridar Subramanian, previously held the same post at PCIe flash specialist Virident. Exablox's VP of sales, Ezra Hookano, was previously VP of channel sales at PCIe pioneer Fusion-io, which was bought for $1.3bn by SanDisk in 2014.

SALES GROWTH

Exablox's sales are growing quickly. The company began shipping its storage system in 2013, and now claims more than 200 customers and 15PB of total installed raw capacity. Compared with the fourth quarter of 2014, the customer count has more than doubled, and the installed capacity has tripled.

The company says that roughly half of its customers use its systems to store primary or working copies of data for second- or third-tier applications, such as x-ray images, home directories and file shares, legal documents, and other 'digital assets.' Many of those customers also use the same systems to store backup data. Indeed, backup storage is the single most common application. Roughly 75% of customers use Exablox systems for backups – whether or not they also use the same systems to store primary working data. Some customers also use Exablox systems to store live VM images.

Although Exablox says it is beginning to win large customers, so far buyers have typically been SMBs, with headcounts ranging up to around 1,000, and many with 250-500 employees. But the company has chalked up its first Fortune 500 customer, which has deployed Exablox storage at multiple locations. Also, some deployments by midsize organizations are large – for example, a customer with 500 employees has installed 1PB of Exablox capacity for linguistics analysis.

Other customers include universities, law firms and healthcare providers. The 30 or so customer logos on Exablox's website include the Johns Hopkins University, New York University, Kawasaki Motors and Princeton Radiology. Some small service providers are using Exablox storage, but the company says they are not a significant part of its business.

Making a conservative assumption of 6:1 data reduction from de-dupe and compression, Exablox says its prices have been as low as $0.10 per effective GB over three years, for the largest possible Exablox cluster. With the addition of variable-length de-dupe, even lower prices are possible because of the improved de-dupe ratio, according to Exablox.

The largest deployment currently comprises 14 nodes – or around $160,000 in hardware costs – and is being used for PACS (picture archiving communications systems) medical data. Exablox says it is restricting its channel to large distributors. It has signed its first reseller in Japan, and is currently in trials with a distributor in the UK.

PRODUCT

Exablox nodes are called OneBlox, and can be used in clusters of one to seven nodes. The company is preparing to qualify larger clusters (or 'rings' in Exablox terminology), but says it is not aiming to compete with object-storage systems that run to hundreds of nodes.

The pricing and purchasing model is unusual. Each node carries a list price of $11,995, including three-year warranty and maintenance contract, but no disk drives. Customers source their own disk drives at retail prices, avoiding the markups applied by other full storage system vendors. Customers also have to pay a software subscription of $1,795 per year for Exablox's cloud-based management service, which is called OneSystem.

The larger of the two types of node offered by Exablox includes 12 bays for 3.5-inch SAS or SATA drives. Drives can be a mix of SAS and SATA, and around 90% of customers opt for SATA drives. The drives can also be a mix of capacities, currently up to 8TB, allowing customers to step up to the latest drive capacities when they appear, and have been qualified by Exablox.

Because of the object-oriented underpinning, adding nodes to an existing cluster does not require any reconfiguration of RAID groups or other parameters. This distinguishes Exablox's systems from conventional scale-out NAS, according to Exablox. Unlike the latter, Exablox's systems automatically discover new nodes and redistribute data across them, and have no RAID groups or LUNs to manage.

Alongside what Exablox claims is a very rapid 10-minute deployment and subsequent scaling out of clusters, the company highlights its near-continuous data protection. By default, an Exablox system takes snapshots every 10 seconds. Conventional RAID-based systems cannot do this because of the processing overhead of maintaining the very large number of live, pointer-based snapshots that would result, according to Exablox. In contrast, the company says its object underpinnings allow a single OneBlox node to create and handle up to two billion data objects, and more than 10 million files in a single directory.

Rollbacks take less than five minutes, and end users can make self-service file retrievals using Windows Previous Versions or Mac Finder, according to Exablox. Cluster-to-cluster and site-to-site replication are available for disaster recovery, and like the de-dupe, compression and snapshots, the replication can be configured for individual file shares.

VARIABLE-LENGTH DE-DUPE

To protect data against drive or node failures, Exablox does not use RAID striping as used in conventional file or block storage, or the erasure coding that is used in some other object storage systems. Instead, the Exablox system creates three copies of data objects and distributes them across a cluster. This carries a much higher capacity overhead than, for example, RAID 6. However, overall costs per GB of effective capacity are still low, not least because of Exablox's use of in-line data de-duplication and compression.

The de-dupe is across an entire cluster rather than across individual nodes or domains, as in other primary and backup-only disk systems, and currently uses fixed block sizes. In combination with the in-line compression that Exablox introduced in September, the typical data size-reduction ratio has been around 6:1. The company recently introduced variable block size de-dupe, which will give greater reduction ratios, especially for backup data. Exablox says this will reduce prices to as low as $0.03 per effective GB for the largest seven-node clusters – assuming a 2:1 compression ratio, and a 10:1 variable block de-dupe ratio.

MOVING UP IN PRIMARY STORAGE

Use of Exablox's systems with virtual servers and desktops is mostly restricted to storing backup copies of VMs. The company's products are certified by third-party backup vendors such as Veeam, Veritas (formerly Symantec), CommVault and Unitrends.

However, in 2014 Exablox declared its intention to see its systems used more frequently to store virtual server and desktop images. One of the steps the company says it will soon take to make this happen is to gain certifications for VMware vSphere and Microsoft Hyper-V. Exablox has already begun to address the performance of its storage when handling the small-block random I/O generated by virtualization.

The company launched the second of the two types of OneBlox node it now offers over the summer. Called the OneBlox 4312, the new device offers up to 10 times the performance of its predecessor, according to Exablox. A major reason for the speed boost was the addition of 10Gb Ethernet ports; another was a change of the controller processor – from Cavium ARM-based processors to Intel Haswell x86 processors. This also allowed the maximum number of disk drives in a node to be increased from eight to 12.

Exablox cites third-party performance tests that show that a single instance of the newer OneBlox 4312 node, using 10Gb Ethernet ports, delivers up 1GB/sec sequential, 64KB write throughput from eight clients. Read speed was twice that, at up to 2GB/sec for eight clients. One GB/sec equals a backup ingestion rate of 3.6TB per hour of raw data. This is single-node performance. As a comparison, EMC's entry-level Data Domain 2200 backup appliance handles up to 3.8TB per hour.

However, VDI and many other virtualized applications need strong performance for random I/O rather than sequential I/O of backups, and for smaller I/O block sizes. Currently, Exablox is not citing any performance numbers for such workloads, and says that it is not yet pitching its systems for general virtualization storage. An exception is virtualized servers running applications such as CAD/CAM, which involve large files for which Exablox's system can already provide good performance.

Exablox says that in 2016 it will launch a software update that will provide strong random I/O for 8KB and 16KB block sizes. Although each Exablox 4312 node uses a 1TB flash drive to store metadata, it does not use flash to store data, and there is a tiering or flash caching mechanism.

COMPETITION

As a vendor of disk targets for backups, Exablox's rivals include EMC, ExaGrid, Quantum and Dell. Against these suppliers, Exablox highlights its ability to scale out its systems without any reconfiguration of backup applications, its global de-dupe across an entire cluster and its low cost.

In the low-end to midrange market for NAS or file-level storage, the dominant suppliers are NetApp, with the FAS 2500 series, and EMC, with the VNX 5200 and VNXe devices. Against this competition, Exablox has a clear price advantage, although as mentioned previously, random I/O performance is currently likely to be a weakness.

Exablox has been very unusual among object storage vendors in its decision to provide native NFS and CIFS access, and no direct object-level data access. As far as 451 Research is aware, Scality is the only other major object vendor that has taken this route rather than use a separate gateway. However, Scality is not a direct rival to Exablox – it is operating much further upmarket. This situation could change as other object vendors follow Exablox's example and address the fact that customers frequently cite lack of NFS and CIFS support as an objection to object storage. The open source Ceph storage software project is currently looking to add file services to its existing block- and object-level data access.

SWOT ANALYSIS

STRENGTHS
Exablox is reporting rapid sales growth, resulting from the virtues of its products in the areas of cost, and ease of scaling and management.

WEAKNESSES
The company has not yet proved that it will find significant use for its storage systems outside of their current field of backups, archive and non-performance-sensitive applications.

OPPORTUNITIES
As data volumes continue to grow, Exablox's core virtues of low cost, scalability and simplicity of management will become increasingly important to enterprises of all sizes.

THREATS
The types of data that Exablox's systems are currently being used to store are also strong candidates for off-premises storage in public clouds. Exablox will have to maintain its low prices to avoid this threat.