

REPORT REPRINT

Eying prime time, Exablox steps into cheap, flashy but virtuous storage

TIM STAMMERS

27 SEP 2016

The startup has already built a growing customer base for its unusual object-based NAS storage. With a view to encouraging more tier-one usage of its systems, the company is poised to ship an all-flash but nevertheless low-cost version of the same device.

THIS REPORT, LICENSED EXCLUSIVELY TO EXABLOX, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR REPOSTED, IN WHOLE OR IN PART, BY THE RECIPIENT, WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.



©2016 451 Research, LLC | WWW.451RESEARCH.COM

Since it began shipping its product in 2013, Exablox has been building a business selling object-based, scale-out storage that allows data access via the NFS and CIFS/SMB file protocols. That NAS front end allows the Exablox system to store data for a wide range of applications not usually supported by object-based storage systems. Compared to conventional NAS boxes that do not have object underpinnings, Exablox claims a number of virtues for its system, including low management overheads, near-continuous data protection and simple scaling.

Currently, the major applications for Exablox's OneBlox system are as secondary storage – for backups, and for data that is not performance-sensitive. But since 2015, the company has been working to move upmarket into primary storage, by increasing the performance of its systems, with storage for virtual servers and desktop as the principal targets.

In the next quarter, Exablox will continue this process by shipping the first all-flash version of OneBlox, and releasing a software update that will boost performance for virtual server and desktop storage across all Exablox systems. The company says its all-flash OneBlox will be about as fast as rival all-flash arrays (AFAs), but claims it will be several times cheaper per gigabyte of capacity. That is because the device will be populated with flash drives bought independently by customers, at retail prices, rather than at what are often highly inflated prices when vendors only support drives resold by themselves.

THE 451 TAKE

Stepping into all-flash storage is a natural development that fits nicely with Exablox's ambitions to expand into primary storage. The plan to allow customers to buy their own flash drives to fit to the AFA is critical, because it will ensure very competitive purchase prices for the overall system, even compared to hybrid, disk-and-flash storage systems. This practice isn't new to Exablox, which already allows customers to source their own drives for its existing disk systems. Also, the AFA's competitive pricing will compensate for the fact that the Exablox AFA may not quite match the performance of other AFAs, because of its object origins. Nevertheless, the device will certainly be fast as well as inexpensive, and will share the scalability and management advantages that have already created a growing business for Exablox. While the company is citing server and storage virtualization as the principal applications, 451 Research presumes that it will also be able to target a number of other file-level, performance-sensitive applications.

BACKGROUND

Based in Sunnyvale, California, Exablox was founded in 2010, and has raised \$45.5m in VC funding. Investors include Dell Ventures, DCM, Norwest Venture Partners, Toshiba America Electronic Components and US Venture Partners.

When we last published a report about Exablox in December 2015, the company was claiming rapid sales growth and more than 200 customers. Three months later, it issued a press release claiming more than 400% year-on-year revenue growth in 2015. When we spoke to Exablox this month, the company declined to state a customer count, but said that sales growth continues to be strong, and to include much repeat business.

STRATEGY

Exablox's target customers include small to mid-sized organizations with up to about 1,000 employees, and many with 250-500 employees. Departmental deployments by very large organizations are also in the mix, and early this year, Exablox named Lockheed Martin, Virgin Media, Kawasaki Motors Manufacturing and Southwest Research Institute as customers.

Backup storage is the most common application, and late last year, Exablox estimated that about three-quarters of its customers use its devices for that purpose. About half also use the devices to store working data for non-performance-sensitive applications, such as medical imaging. But Exablox aims to expand into primary storage for virtualized servers and desktops. At present, the company estimates that less than 20% of its customers are using its systems for that purpose. The forthcoming AFA and software update are expected to drive up that percentage.

The OneBlox is a NAS-only system that currently provides access to data only via the CIFS/SMB and NFS protocols. Currently, much AFA usage involves block-level data access, and this might appear to be a problem for Exablox. However, a sizeable share of AFA usage also involves virtualized servers and desktops, for which data is often handled via NFS (VMware vSphere) and CIFS/SMB (Microsoft Hyper-V).

SOFTWARE UPDATE AND AFA

Fitting its current applications, the OneBlox system has been tuned to handle large IO block sizes. But to handle data for virtualized servers and desktops, it also needs to perform well with small 8KB and 4KB blocks. A software update designed for that purpose is currently being used by early access customers of the forthcoming AFA. When the software update becomes fully GA, it will be available for all versions of OneBlox.

The OneBlox storage system is detailed in our previous report. Currently, all variants of the OneBlox are disk devices, which use flash drives only to store metadata. The forthcoming all-flash version is called the OneBlox 5210, and Exablox describes it as its third-generation device.

Like Exablox's existing OneBlox devices, multiple 5210 units can be scaled out in a cluster of up to seven nodes. Each 1U OneBlox 5210 is based on a Dell server, like the existing OneBlox models. The 5210 can be fitted with up to 10 mixed SAS and SATA 2.5-inch flash drives, and Exablox has qualified a range of Intel, Samsung and Toshiba devices, including 2TB and 4TB drives. A seven-node cluster fitted with 4TB drives would present 280TB of raw capacity. Factoring in a 5:1 data-reduction ratio for Exablox's inline compression and variable and fixed length dedupe, the company says the price for that configuration will be around \$0.40 per effective gigabyte.

OneBlox's object underpinnings rely on replicas for protection against drive failures, which impose much greater capacity overheads than the RAID protection used in other AFAs. This issue is lessened by the fact that although the all-disk OneBlox variants maintain three copies of data, the 5210 maintains only two copies, reducing the capacity overhead. This is because the 5210's flash drives allow it to recover more quickly from a drive failure by redistributing data across the remaining drives, reducing its period of vulnerability to a second drive failure during that process.

Nevertheless, the 5210's capacity overhead is still much greater than for RAID-based AFAs. Exablox says that despite this, the 5210 is cheaper because its customers buy flash drives at retail prices, which are far lower than the OEM prices charged by other AFA makers for drives. In some cases, incumbent storage makers charge more than 20 times the retail price for rebranded flash drives to be fitted to their arrays.

Exablox claims that a single 5210 can handle 50,000 IOPS (8KB blocks, 70% write.) The company isn't yet declaring a latency to go with that IOPS number, but says it will be a few milliseconds or less, as compared to the 1ms or less latency usually claimed for block-level AFAs. For many applications, latency is more critical to performance than IOPS, and Exablox acknowledges that the latency of the 5210 may prevent it from suiting the top 10% of performance-sensitive applications. 451 Research believes it will suit many mainstream applications, and we note that even if the OneBlox 5210 imposes a latency of low single-digit milliseconds, it will still be around 10 times faster than the capacity disk in hybrid storage systems.

The OneBlox operating system includes a number of features designed to boost performance. The NAS front end includes a B-plus file structure, which uses less space than a traditional B-Tree structure, and so can more easily fit into DRAM and flash, significantly increasing look-up speeds. Having translated a file address to a content-based object handle, Exablox says that only two more IOs are required – one to access the data offset, and one to access the data itself.

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's existing disk systems enjoy a clear price advantage, although their random I/O performance is likely to be a weakness.

The AFA sector is led by EMC and its XtremIO device, while HDS, HP, NetApp and Pure Storage also enjoy strong AFA sales. Although the initial target applications for the all-flash OneBlox 5210 are storage for virtual servers and desktops, 451 Research believes the 5210 might also present a lower-cost alternative to emerging all-flash NAS systems such as Pure's forthcoming FlashBlade device, which are aimed at performance-sensitive file-based applications, such as analytics and scientific computing.

Among other object storage vendors, 451 Research is aware of only one that has taken the same decision as Exablox, and is providing native file-level access and no direct object-level data access. That vendor is Scality, which is not, however, a direct rival to Exablox, as 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/SMB 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-level and object-level data access.

SWOT ANALYSIS

STRENGTHS

Exablox's AFA is set to carry low purchase costs. The company has already proven the appeal of its object-based platform, and the practicality of customers' independent sourcing of disk drives.

WEAKNESSES

Exablox has not yet fully quantified the latency of its AFA, but says it might not be low enough to suit the top 10% of most performance-sensitive applications.

OPPORTUNITIES

The OneBlox 5210 AFA could be exactly the stepping stone that will extend the range of applications for Exablox's system to include data storage for performance-sensitive applications.

THREATS

The price gap between an existing Exablox disk array and the forthcoming AFA may dissuade existing customers, who will continue to perceive Exablox only as a backup target, or repository for very cool data.