

OneBlox and OneSystem Technical Overview

The Business Challenge

Some estimates indicate that, in 2012, 2.5 quintillion bytes of data was generated globally – that’s two billion terabytes. Every day, that rate of generation increases as more and more people store more and more data in new and different ways. Enterprises have been forced to continue to “throw hardware” at this storage problem in order to maintain the rates of capacity growth. As a result, many enterprises are supported by storage platforms that might better be called “Frankenstorage.” Moreover, the foundational storage arrays used by these enterprises is likely not as scalable and feature-rich as necessary for them to be fully reliant on these systems.

In short, traditional storage systems do not provide the breadth and depth of functionality that modern enterprises demand. But, that’s only a part of the equation.

Storage Technology has Fallen Behind

It’s easy to see that traditional storage technology has fallen woefully behind what is required in the modern enterprise. In so many ways, storage actively works against an organization, requiring constant attention, expensive consultants, and specialized third party monitoring tools. Some storage systems look much like a car with no tires, no dashboard, and a square steering wheel. Sure, it would be possible to drive on the rims, but it would be a poor driving experience and, without the dashboard, it would be impossible to figure out how fast the car is driving.

Data Management has Become Increasingly Difficult

Besides the very real problem of ongoing capacity upgrade needs, enterprises also face a multitude of challenges around ongoing management of the data and of the very systems that house the data. While traditional storage vendors have attempted to bring ease of use and ease of administration to their products, they are saddled with the burdens of legacy architectures that practically force complexity into the administrative process. As such, storage management has become a specialized niche for many IT staffers that are tasked with maintaining an organization’s storage assets. The problem here is that every minute that a storage administrator has to spend on unnecessary tasks is a minute delay in getting a product to market or a minute lost in focusing on business-facing organizational goals.

What Would Perfect Storage and Storage Management Look Like?

In a modern enterprise, storage is the lynchpin upon which all services depend. Employees and customers alike are affected by the storage choices that are made. Because of this and after studying the major storage challenges that affect enterprises, Exablox decided to ask a very simple question: “What would the perfect storage system and

storage management look like?” It’s a short question, but with a number of different answers. At the same time, we scanned the storage market and realized that there can be many correct answers but no one has brought together what we feel is the perfect solution... until now. Exablox has cherry-picked the best features from best of breed solutions across the storage landscape and brought them into a single product.

Exablox will:

- Scale-out seamlessly to accommodate capacity, performance, and/or availability requirements
- Deduplicate information to maximize storage utilization
- Snapshot information to provide granular file level recovery
- Fail over and replicate resources enabling continuous information availability

On the management front, Exablox also identified the cream of the crop with regard to how certain products are managed and envisioned and built a storage management system that has the following characteristics:

- Be intuitive and easy to use
- Configure and integrate with cloud-based services
- Proactively diagnose and report problems
- Strive for zero configuration for storage management tasks
- The end result of this effort is a powerful, yet simple, storage management appliance, OneBlox®, coupled with a cloud-based management service, OneSystem. In fact, the OneBlox and OneSystem® combination is so intuitive that a typical IT staff person can have the initial node up and running in just 5 minutes.
- But that simplicity doesn’t come by way of a scaled-back product. The Exablox solution is incredibly powerful and feature-filled, as you will discover in the rest of this paper.

The Exablox Ecosystem and Components

In order to achieve the lofty goals set forth for this new vision of storage, Exablox has built, from the ground up, a brand new and unique storage ecosystem. Exablox provides organizations with massively scalable storage with enterprise-class features and single-instance cloud management that brings the entire environment under a single administrative umbrella, regardless of how large it ultimately grows.

Exablox takes a building block approach to architecture, starting with a node, expanding to a ring, and eventually becoming a mesh. Each object type is described in the following sections.



OneBlox

A OneBlox appliance is a single hardware device intended to store a company’s data. OneBlox is a 2U appliance provides eight hot swappable 3.5” hard drive slots and supports up to 48 TB of raw capacity when each slot is outfitted with 6 TB drives. As larger capacity drives become available,

OneBlox will immediately support those configurations. In order to meet the unique performance goals across a variety of customers, the OneBlox appliance supports the use of SATA and SAS disk drives and drive capacities can be mixed within a given OneBlox.


A OneBlox appliance exposes administrator-managed shares through the Server Message Block (SMB) protocol and appears on the network as a Network Attached Storage (NAS) device.

Each OneBlox appliance connects to the network with four gigabit Ethernet (GbE) ports, all found on the back of the device. In addition, administrators enjoy at-a-glance operating information thanks to the 3 inch LCD screen found on the front of each appliance and the LED lights that provide immediate appliance diagnostic details.

With the OneBlox appliance, customers can start small and grow as the need arises. Start with just a few hard drives and add more as capacity (either drives or OneBlox) as needs change, all with no downtime and zero-configuration. This is the basic way that the OneBlox system easily scales to meet growing customer needs, but without the complexity that accompanies traditional storage systems.

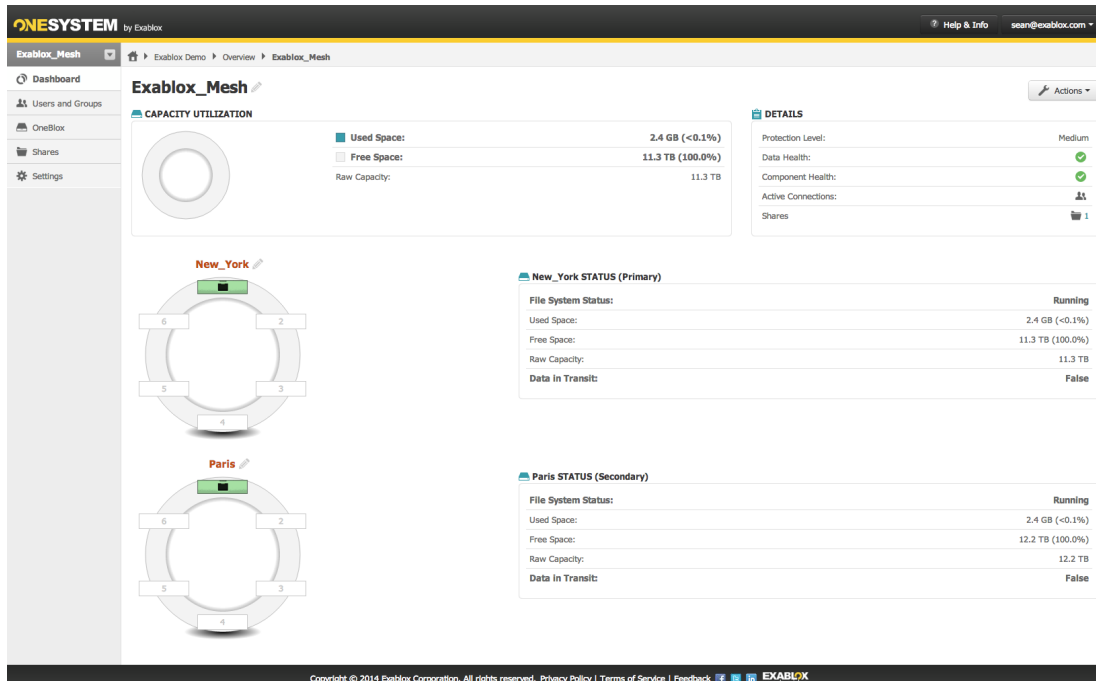
Ring

Scalability is in the Exablox DNA. With a single appliance, customers can grow to support up to 48 TB of storage. But it doesn't stop there. Customers can connect multiple OneBlox appliances together to form a ring, which presents a single file system to applications and to users. In other words, this multi-node cluster appears as a single nearly 300 TB global file system and the whole system is still managed through a single OneSystem account, as discussed later. As new appliances are added to a ring, behind-the-scenes data protection technology kicks in to automatically spread data among the various nodes in order best to protect the environment in the event of a hardware failure. This data protection technology works through replication, with three copies of all data stored in various locations in the ring.



Mesh

Capacity and scalability are critical to many organizations. But without comprehensive data protection, a storage system can't be considered "simple" and "easy to use" since administrators will need to make significant investments of time and expense in separate technologies to provide data protection. Through the use of a mesh configuration, customers can create collections of two or more rings that are automatically mirrored or replicated to one another, all behind the scenes.



Lest anyone think that replication has to be a difficult service to implement, it's important to note that the OneBlox mesh is created through a simple, intuitive, GUI-based drag and drop process akin to moving a file from one folder to another in an operating system. It's simply done in 4 steps and 4 minutes.

OneSystem

If the hardware Exablox is bringing to bear seems impressive, the software suite behind it – called OneSystem – is nothing short of magical. Instead of relying on the traditional storage management paradigm, the entire OneBlox ecosystem – the appliance, rings, and meshes – are managed from a central, cloud-based management service. No matter where customers are, they are able to access the web-based OneSystem, shown in the Figure below.

Besides being eminently accessible, OneSystem is also simple and intuitive, from initial storage configuration to ongoing management. By putting pertinent information right in front of the administrator, there's no digging through screens and screens of information in an attempt to figure out how much space is left in the environment and whether or not there are any alarms that require resolution.

OneSystem can also seamlessly integrate with an on-premises Active Directory (AD) environment in order to provide users and groups to the management environment. This allows an administrator to leverage existing assets, which further simplifies the administrative experience.

Inspired by the simplicity of a Bluetooth-style pairing, OneBlox with the OneSystem management is a snap. When a new OneBlox appliance is deployed, it connects to OneSystem to announce its availability. From there, an administrator logs in to OneSystem and, using details that helpfully appear on the OneBlox LCD display, pairs the appliance with OneSystem. The entire process takes mere minutes and once accomplished, the full administrative experience becomes available to the administrator. The use of both a username and password to log in to OneSystem and a unique pairing code provides multifactor security to keep the right hands on the data.

To maintain as much ease of use as possible, all of OneBlox features are enabled and configured by default and do not present individual configuration parameters.

Distributed Object-based File System

To achieve the stated goals of the OneBlox system, Exablox has created a groundbreaking new file system. OneBlox's object-based file system handles all aspects of data storage and protection, including protecting all storage in a number of ways. OneBlox protects against:

- Data deletion or corruption
- Drive removal
- Drive failure
- Appliance failure
- Ring failure

Simple doesn't mean ordinary. Exablox's vision required the design and implementation of many enterprise-class storage features architected to be as powerful as those features found in traditional storage systems, but with exposure to the administrator that retains the desire to provide an easy-to-use solution. As such, significant engineering has gone into building world-class features for organizations more interested in solving business problems than in fighting with storage.

OneBlox's file system is built on top of a distributed content addressable content store. As new data goes into the file system, they are chunked up into 32KB blocks and distributed around the ring. Once distributed, additional replicas are created for data protection. Clients accessing content stored in a ring are load balanced across multiple nodes, which increases performance and reliability of the storage environment.

To maintain performance levels that are necessary for large amounts of data, each node maintains its object metadata in an index stored on a solid state disk (SSD) which is built into each OneBlox.

The Global Namespace

The Global Namespace is the single presentation of the OneBlox environment to the user base and is the file system that powers OneBlox. With the OneBlox file system, administrators do not create traditional storage objects such as LUNs, volumes and aggregates. Nor does the OneBlox system use RAID for data protection. As storage capacities have grown, RAID has proven to be the Achilles' Heel for many, as rebuild times in the event of a disk failure grow longer, increasing the potential for data loss.

Instead, OneBlox uses a comprehensive replication engine and other features, including Continuous Data Protection (CDP), to protect data and to present to users a single, unified view of the storage environment.

Local Replication



The cornerstone for many enterprise data protection efforts, replication is a key asset in the Exablock arsenal and is used throughout the ecosystem, from single appliances all the way up interconnected meshes. In fact, data protection is so critical in the Exablock vision that, by design, multiple copies of data are always kept, but the exact process depends on how many OneBlox nodes and Rings exist in the environment:

- **Single node.** OneBlox creates three replicated copies of the information on the single node across 3 drives, which enables the loss of up to two physical hard drives in the appliance, without loss of data.
- **Three or more nodes in a ring.** OneBlox leverages all nodes in the ring as replication targets, enabling a scenario in which two nodes can fail and users will still have access to their information.

Inline Deduplication

Data protection always comes with some kind of cost. When data is replicated, as is the case with OneBlox, additional capacity is required to store that replicated data. However, to maximize the overall usable capacity of a OneBlox appliance or ring, Exablock has imbued the system with powerful inline data deduplication. Each time a file or portion of a file is targeted to the OneBlox ring, it checks to see if that data already exists in the environment. If it does, that second copy is simply discarded and a very small pointer to the original copy of the file is put into that space instead, which uses just a fraction of the storage that would otherwise be used.

To put it simply, each time data is written to the environment, OneBlox asks a simple question: Do I need to store the data or just create a pointer?

In a traditional storage environment, deduplication is often disabled and not used for primary storage due to the processing overhead that is introduced inside these legacy environments. With OneBlox, not only is deduplication enabled by default, it can't be disabled. With the OneBlox object-based file system, global deduplication essentially 'comes for free'.

Continuous Data Protection Through Snapshots

Although fast disk-based technologies have been available for many years, many companies will rely on backup windows to keep long-term archives of their critical business data. Most often, these backup windows take place over night. This presents a significant opportunity for data loss in the organization if, for example, a data loss event takes place at the end of a workday. In this case, all of the work performed since the most recent backup – which took place in the wee hours of the morning – is lost.

OneBlox provides a powerful data recovery feature known as Continuous Data Protection (CDP). Throughout the day, OneBlox continuously takes snapshots of new data written in the system and keeps those snapshots as recovery points in the event of data loss, which could be something as simple as an accidentally deleted file. This snapshot-based approach to Continuous Data Protection eliminates the data loss window, enabling much faster recovery in the event of a serious data loss issue.

It might seem that the snapshots feature could use massive amounts of disk space. Thanks to OneBlox's deduplication engine, snapshots use only a fraction of the space that they otherwise would, enabling very granular data recovery at a minuscule cost. All snapshots are deduplicated against the primary data.

In a "keeping IT focused on the business" feature, these snapshots are actually available to users who can access them through Windows Explorer or the Mac Finder and recover older versions of their own files without IT having to be involved. OneBlox provides the ability to go back to the last file written (which could be the last few seconds), as well as further back in time: the last minute, hour, day, week, month, or year(s).

CIFS/SMB

The Microsoft Windows-standard CIFS/SMB file server protocol has been around for a long time and enjoys wide operating system support. As such, OneBlox exposes file shares via the Global Namespace using SMB. This enables immediate access to OneBlox shares from Windows-, Linux-, and Mac-based systems in the organization.

Additionally, OneBlox provides OneSystem with the ability for administrators to authenticate with Active Directory environments so that existing user accounts can be used for OneBlox access. This negates the need to maintain a separate user directory in OneSystem if users elect to use Active Directory.

Remote Replication

Sometimes, disaster strikes. Mother Nature has a knack for attacking without warning and entire parts of cities can be wiped off the map. To protect against site failures such as this, many companies deploy disaster recovery sites that are physical separated from the main data center.

For those concerned about potential site failure, a ring, which can consist of multiple interconnected OneBlox, may be replicated to another ring in what is known as a mesh. This replication feature replicates the information across the two rings and only replicates the changed deduplicated objects on an ongoing basis. In the event of one site going offline, the other site has additional copies of the information and users can access the information from the alternate ring. OneBlox enables customers to build reliability and disaster recovery for a low price and without the complexity that often dooms these implementations. Moreover, all information replicated to each ring is already deduplicated and encrypted, so the network utilization is maximized.

Increasing Capacity

No company is a static entity. Organizations grow and change every day and this growth becomes a part of the organization culture. Why is it then that adding additional storage capacity with a traditional storage system is always an "event" that requires planned downtime and can carry with it the risk of an outage?

There's a much better way to handle capacity increases. With zero configuration additional disk drives or OneBlox can be added in support of an organization's growth. With OneBlox, there are three primary ways to add capacity:

- Add new hard drives to an existing OneBlox
- Upgrade existing hard drives in an existing appliance with drives of a higher capacity
- Add new OneBlox to an existing ring

Adding or Upgrading Hard Drives

With traditional storage arrays, adding capacity often means buying a huge, expensive shelf fully populated with hard drives even if just a few TB of extra capacity is needed. With OneBlox, if an appliance has available hard drive slots, an administrator can simply insert a new hard disk until it clicks in place. That's it. It's a complete tool-less process and it doesn't matter what the capacity is of the new hard drive. OneSystem simply detects that a new drive has been added and works behind the scenes to integrate its capacity into the environment. Organizations can always use any hard drive purchased from a computer store. OneBlox inherently supports the latest drive technology.

If a customer wishes to upgrade the capacity of an existing appliance, OneSystem and OneBlox will handle the integration of the new capacity and the redistribution of existing data on the appliance and in the ring. Whether adding additional drives or additional OneBlox, it requires zero configuration. The additional capacity is dynamically pooled with the existing global file system and applications and users immediately have more capacity. With a traditional storage array, this drive-by-drive replacement process would be extremely tedious, and most likely not supported by the storage array vendor.

Adding OneBlox

As more capacity is needed in an environment, customers can simply add more OneBlox to an existing ring and scale capacity, performance, and availability. As new devices are deployed, they are automatically discovered by OneSystem and are seamlessly added to the storage environment. Similar to drive expansion, adding OneBlox to a ring requires zero configuration. From a user perspective, the effective storage capacity of the environment simply grows as new OneBlox are added. No changes to users or applications is required to take advantage of the additional capacity.

Summary

Exablox uses three words – **Powerful**, **Simple**, and **Protected** – to describe what is a top to bottom reinvention of how storage works at the most fundamental levels. By abandoning legacy RAID concepts and traditional scale-up storage architectures, Exablox has incorporated best of breed ideas into a single integrated solution and management experience. Exablox is enabling organizations to solve complex business challenges with powerful, yet simple and elegant, storage.