

# Introduction to the OneBlox Architecture

## Introduction

OneBlox is a storage solution for unstructured data without any legacy thinking, or simply “storage reimaged”. But what does this mean from an architecture perspective? How is OneSystem and OneBlox different from legacy solutions? This paper is the second in a series and continues the discussion from, “*OneBlox Solution Overview.*”

## OneSystem

The OneSystem multi-tenant cloud-based management service is simple and direct, focusing on enabling administrators to quickly configure the system rather than requiring a large amount of storage expertise.

The fundamental design principles were to produce an affordable scale-out Network Attached Storage (NAS) solution with enterprise feature that is easy to manage and grow over time. OneBlox is advanced, yet simple, due in large part to OneSystem. By leveraging next generation cloud-based architectures, organizations no longer have to be concerned with managing, yet another, management server in the data center. Administrators can manage their OneBlox appliances through a web browser from anywhere, and all OneBlox appliances in a Ring, or multiple rings, are managed together from a single interface.

The OneSystem application is used to combine OneBlox appliances into Rings as well as configuring shares, user access, and remote replication. The OneBlox Ring can be joined to an Active Directory domain to leverage existing enterprise user accounts. This makes configuring OneSystem just as easy as other components of the OneBlox solution.

OneSystem is also used to configure remote replication. If local availability isn't sufficient to meet recovery point objectives or compliance with corporate data protection policies, OneBlox can replicate data from one site to another to provide off site backup copies of your data.

## OneBlox

### Storage Reimagined

With a peer-to-peer scale-out architecture, OneBlox appliances can grow from small storage capacities to larger capacities with zero configuration by the IT administrator. OneBlox is designed to optimize performance while lowering acquisition and ownership costs. OneBlox's architecture does not rely on a central metadata service for data access. This allows for scaling in small or large units with the cost of scaling being linked to the amount of scaling. Inspiration was sought in emerging technologies like peer-to-peer file sharing and next generation cloudbased storage rather than in legacy storage design.

OneBlox addresses two major storage problems that have arisen in recent years. The first is the plagues backup/recovery challenge. Organizations continue to face recovery challenges as storage capacities increase, application recovery time objectives decrease, and legacy tape architectures fail to satisfy business requirements. OneBlox, as a disk-based backup/recovery target solves these challenges, while reducing acquisition and ongoing management costs.

The second storage problem that OneBlox addresses is the large number of small, unstructured files that have become pervasive across many businesses and have proliferated through organizations and richer

media types that have become a part of business: documents full of pictures, photos recording everything from travel to insurance claims, and increasing volumes of videos. These types of media are consuming vast amounts of valuable resources on file servers and expensive SANs in large organizations.

The requirement to store and protect terabytes of files is a headache for many organizations, particularly as they are trying to do more for the business with ever-shrinking budgets. In smaller organizations without proper content protection plans, this data is at risk of being destroyed or lost forever. OneBlox provides secure and reliable storage for disk-based data protection and file storage for organizations of all sizes.

### **Easy to Grow**

OneBlox has eight hot swap disk bays on the front. These can be populated with your choice of SAS or SATA disks. Different spindle speeds or disk sizes are no problem. OneBlox will detect new disks and add them to the existing OneBlox Ring automatically. Data will be dynamically redistributed amongst the disks and more capacity will be available to the applications.

As OneBlox are added to an existing clustered configuration, or a *Ring*, they will automatically rebalance the existing data across all the nodes to ensure information is protected to the greatest possible extent. OneBlox creates three copies of the data to ensure that if any two drives or OneBlox fail, the information is still protected. Existing systems can also have additional disk drives added and the Ring rebalances the data within the Ring without administrator actions.

If a disk fails, and it is not replaced, the reverse process will occur. The global storage capacity will dynamically shrink and the data that is no longer stored in redundant locations will be copied to new locations on the Ring until there are three copies of every object.

Inside the OneBlox is standard hardware that calculates each object as information is written to the Ring. Up to four gigabit Ethernet ports can be used to connect each OneBlox node to your network and a standard cable provides the power.

### **Object Storage with Inline deduplication**

The files stored on OneBlox are stored as objects, which is very different than legacy RAID storing fixed blocks in a given volume and LUN. Even the internal data structures used to manage the OneBlox filesystem are treated as objects. Objects are content addressable: the object can be located by the Object ID, which is a SHA-1 hash of its data block, and is calculated at the time information is written to the Ring.

No central database of object to disk location or shared file table is necessary in OneBlox. All that is needed is the data and a math calculation that is done in dedicated hardware. That math calculation also identifies if the Ring is already storing a copy of that object, resulting in inline deduplication for every object across the global filesystem. With that information OneBlox is able to eliminate storage waste, resulting in built-in deduplication for every object.

Furthermore, when the data is read back, the object is compared to the Object ID ensuring that the data that is read back from disk is the same as the data we wrote to the disk. This eliminates the risk of bit-rot and other drive errors that jeopardize the validity of information delivered to applications and users in legacy RAID storage systems. Information read from OneBlox is mathematically guaranteed. If the hash is not the same, OneBlox will automatically retrieve the object from another location in the Ring and deliver the information to the application.

### **Continuous Data Protection**

Continuous snapshots are an integral part of OneBlox, are enabled by default, and have nearly zero performance overhead on the filesystem. This is because OneBlox tracks all changes made to the files

stored on the Ring as they are made. As this process occurs automatically there is no need to pause file system operations to take a snapshot. The integrated continuous data protection functionality also ensures that all file versions are available rather than those changes made since the last manual snapshot.

Users can easily find old copies of data in these folders without administrator intervention. Restoring a deleted file is as simple as opening the folder for the date and time before the deletion and copying the file back into the share. This process does not generate a helpdesk call, freeing administrators to concentrate on other important tasks. In addition, since the deleted files is already being stored on the OneBlox appliance there is no reduction in available space when the file is restored.

### Replication

OneBlox has local replication enabled by default. It uses local replication to ensure three copies of the data resides on three different disks or three different OneBlox within the Ring to provide protection against two disk or node failures. The inherent integrity check ensures that only one copy of the data is needed to ensure it is the expected data. Protecting against site failure can be enabled by replicating the information within a Ring to another Ring in another location, whether it is a neighboring building or another city.

In the past, offsite data protection was the job of tape backup technology. Full copies of every piece of data needed to be written to tapes each night and taken to a remote location. The huge growth in data storage capacity has not been matched by an equally large growth in data transfer rates. The average amount of data stored is now usually too much to restore from tape in an acceptable time to meet modern business recovery point objectives.

To counteract this trend, file restores and disaster recovery have had to move to replication and storage snapshots. With the combination of OneBlox's continuous data protection and remote replication, information is available for rapid restores, either locally or at a remote location. Additionally, because OneBlox is built on our object-based file system, all information that is replicated to another Ring is deduplicated—minimizing bandwidth requirements.

### Bring Your Own Drives

The BYOD (Bring Your Own Disk) strategy for OneBlox enables organizations to mix and match different types, speeds and capacities of retail-priced disk drives, driving costs down by a factor of 10. According to Gartner, disks can easily account for 50%, or more, of a disk storage system's hardware costs.

### Security

This same hardware is used by the OneBlox to quickly identify the stored data objects using calculated SHA-1 hashes as object IDs. This hardware-based calculation is used to identify which OneBlox appliance has a piece of data, allowing data to be retrieved after calculation rather than a metadata query. The use of calculation to locate data objects is a key to the OneBlox scalability, avoiding a central metadata directory that can become a bottleneck at large scale.

The same hardware is used to verify the data when it is read back from the disks giving you certainty that your data isn't being compromised by bit-rot. Since OneBlox keeps three copies of your data on separate disks, the system ensures that one copy will still be correct in the event of disk or appliance failure.

OneBlox provides an enterprise scale-out NAS that is easy to grow, manage, and affordable to all organizations. The content addressable object based storage is a reimagining of how a file storage system should work. Eliminating legacy RAID and utilizing patented object-based file system gives OneBlox the ability to deliver enterprise features, without the cost and complexity of legacy storage solutions.