Next Generation Object Storage with the Accessibility of Network Attached Storage

Introduction

Legacy storage systems have issues when it comes to scaling. Integrated RAID disk configurations do not lend themselves to expansion, and conventional file systems aren't much better. But the issue of scaling has been solved in cloud storage systems by the use of object stores. These store data as an object in a distributed database, eliminating scalability and management issues.

But object stores aren't compatible with the majority of conventional applications and operating systems. That's why an integrated approach is right: Bring the benefits of object storage with the compatibility of conventional networkattached storage protocols.

Challenges

Managing storage systems is a key issue for IT administrators, and this is especially true as the environment grows. As multiple storage arrays are added, it is difficult to rebalance client workloads across multiple systems. This is why most storage solutions are monolithic: They may scale-up (by adding additional disk shelves) but they do not truly scale-out, bringing additional controller resources like memory, ports, or CPU power.

Flexibility has long been an afterthought in enterprise storage. Arrays were designed primarily for reliability, with advanced features, manageability, and flexibility being lower priorities. But decades of work have made reliability a given for most systems, and attention is turning to these other areas.

Conventional storage systems simply cannot keep up as demand grows. They may scale for a while by adding disk shelves (the so-called "modular" or "scale-up" approach) but the controller will eventually be over-taxed by all this capacity. The next alternative has been clustering multiple controllers together, but this has limits as well: It is still difficult to re-balance workloads across nodes of a cluster, and increasing inter-node traffic puts a limit to the size a cluster can grow to.

Cloud storage systems eliminate these scaling issues by building a distributed object store with many small loosely connected nodes. This has been proven to reach massive scale for web applications and are secure, reliable, and feature-rich. But conventional operating systems do not have cloud storage protocol clients, making these impossible to integrate with existing applications.

Solution

An ideal storage solution would combine the future-proof scalable storage architecture of a cloud storage system with conventional client-access protocols. Unlike a scale-up or clustered storage array, a distributed object store allows all aspects of storage system architecture (capacity, compute, memory, I/O) to grow as needs change. Storage is managed as a pool rather than in "silos" and resources can be assigned to clients based on requirements. It is not easy to utilize conventional access protocols like NFS and SMB along with an object store, but this combination brings many benefits.

US: sales@storagecraft.com
Europe: sales@storagecraft.eu
Australia: sales@storagecraft.com.au
New Zealand: sales@storagecraft.co.nz
Asia: asiasales@storagecraft.com.au

Object Storage

OneBlox is reimagining storage by integrating a flexible object storage back-end with industry-standard client protocols. The connected clients see a file server, but data is protected and managed using state-of-the-art object storage. OneBlox can be accessed with standard NFS and SMB protocols bringing the best of both worlds together.

With OneBlox, industry-standard network-attached storage protocols like NFS and SMB are completely integrated with the object storage layer and with each other. Multiple OneBlox are organized into a Ring, with data and access distributed across all participating OneBlox. Regardless of the protocol used, data is stored and accessed through the unified object store. This means that NFS (UNIX) and SMB (Windows) have access the same pool of storage capacity and IT organizations can seamlessly create NFS or SMB shares satisfying a variety of applications accessibility requirements.

A cryptographic hash is generated for each 32 KB data chunk stored, uniquely identifying it in the object store and enabling advanced features. Hashes can be compared on write and read to ensure data integrity, and data can be de-duplicated in-line to increase storage capacity efficiency. OneBlox does not require a separate metadata server like some object stores, and the hash lookup is accelerated with flash storage rather than expensive system RAM.

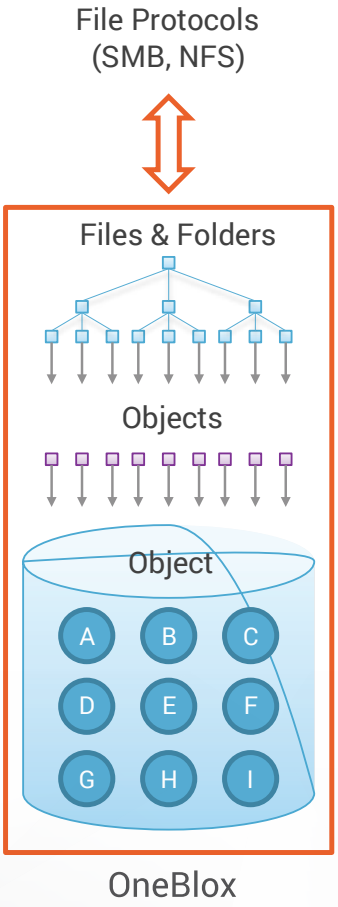

The OneBlox object store gives some other unique advantages as well. Storage management is simplified, since the entire OneBlox Ring(s) is managed as an entity through OneSystem, our cloud-based management service. OneBlox is flexible enough to support a Ring with varying storage capacity in each OneBlox and even varying hard disk drive capacities optimally. Customers can use whichever drives are available and the system will correctly protect and balance data access across them.

Conclusion

The OneBlox solution marries the compatibility of industry-standard protocols with the flexibility of object storage, bringing unique advantages in terms of storage management and scalability over time. No other storage system is as flexible or as scalable, and it all starts with an affordable “bring your own drives” solution.

US: sales@storagecraft.com
Europe: sales@storagecraft.eu
Australia: sales@storagecraft.com.au
New Zealand: sales@storagecraft.co.nz
Asia: asiasales@storagecraft.com.au