

Exablox

Company and Business Update

Douglas Brockett, CEO

Launched April 2013

(Founded 2010)

OneBlox - scale-out storage for the enterprise - converged storage for primary and backup / archival data

OneSystem - manage on-premises storage exclusively from anywhere - visibility, control, and security without cost / complexity of traditional management

100+ customers, 30+ billion objects

typical customer had 50TB a few years ago, now they need 200TB and 3x performance. And in a few years, they're going to need 800TB and 10x performance

if we could reimagine storage - what would it look like?

scale out

dedupe

snapshot

replicate

be simple yet powerful

be managed from everywhere

Unique Architecture [photo]

Nearly POSIX-compliant file system on top of an object store

"Scale out NAS on the outside, Object storage on the inside"

Standardized access (SMB / NFS)

horizontal scaling

multiple copy protection

advantages of traditional NAS

- suited for dynamic, transactional data

- simple, traditional data access

- compatible with current enterprise apps with broad applicability

advantages of object storage

- excellent scalability

- simple storage mgmt with unified storage access

- data management features such as inline dedupe, encryption - automatic

OneBlox

only exablox provides

- scale-out dedupe

- scale-out continuous snapshots

- scale-out RAID-less capacity
- scale-out site-to-site DR
- bring any drive - one at a time at retail pricing

Auto-clustering - each node adds

- capacity
- performance
- resiliency

Appliance

- 8 * 3.5" drive bays - up to 48 raw TB
- SATA & SAS, mix and match

OneSystem

<https://onesystem.exablox.com>

- subscription SaaS
- Multi-tenant - VARs, MSPs, Departments
- workflow - point-n-click, drag-n-drop, directly actionable
- storage mgmt - provision, manage, monitor, trending, subscribe

Exablox

- combines the "web scale" scale out approach with plug and play app compatibility
- combines off the shelf appliance with bring your own drive economics
- smashes the \$/TB, initial cost, \$/Feature and TCO barriers for enterprise storage

Powerful, simple, protected

OneSystem

Tad Hunt, CTO & Co-founder

Technology Vision

SDS + Storage mgmt

- manage policy, not technology
- SDS "wrapped in tin" for the mid market
- Eliminate complexity
- PnP

- Next Gen features

NAS features atop object storage

- without metadata servers
- without bolt on NAS gateways
- without separate data and metadata servers
- scale capacity, performance, or resilience: just add a node

Intellectual Property

- cloud based management
- distributed filesystems
- distributed object storage systems

Technology Benefits

Create scale out NAS + object clusters atop mixed media - HDD, SSD, Shingled drives

Deliver benefits of object storage technology to traditional apps

- standard file protocols

- eliminate forklift upgrade - single namespace across the scale of the cluster

RAID free data protection

- self-healing from multiple drive and node failures

- rebalancing time proportional to quantity of objects on the failed drive

- mix and match drive types, capacities, technologies

- introduce next gen drives without long validation cycles

Ability to scale capacity from TB to PB easily

- zero config expansion

- manage from anywhere

Able to support all NAS workloads well

- other object stores designed primarily for large files

- OneBlox 3308: 1B objects

1/4 the memory needed per TB of capacity

Ease of Use

- all nodes perform all functions: storage, control, NAS interface

- single failure domain - node

OneSystem

[photo]

Features and Benefits (FCAPS)

Fault Monitoring

- zero config email alerts

- fine grained, actionable intelligence: "Replace the drive in slot 5"

- predictive failure intelligence gathered across global customer base

Configuration Management

- policy based rather than tech based

- Exablox manages config backups, tracks changes, etc

- simple: creating a DR replica is a drag-and-drop operation

Accounting

- usage tracking

- report generation

- facilitates billing for subscription storage services

Performance Management

- global install base telemetry

- capacity and utilisation trending

- hotspot identification

- predictive analytics

Security Management

- AD configuration

- cross organisational views, drilling down to individual drives

- fine-grained, hierarchical admin rights

- audit trails

need OneSystem connectivity to perform upgrades

Management in the cloud, data on-premises

User-initiated upgrades

OneBlox FileSystem

SmashFS Features and Benefits

Protocol

SMB - NAS + Backup

NFS - NAS + Virtualisation + Backup

Distributed Filesystem

POSIX + Snapshots - NAS

Global Namespace - Reduced operational complexity

RAFT / Consensus - zero-config cluster management + scale-out

Distributed Object Store

Small Objects - NAS + Deduplication + IoT

Content Addressability - Deduplication, snapshot, object migration

Consistently Hashed - scale-out + fast repair

High-level OneSystem Overview [photo -

SFD7_Day3_Exablox_OneSystem_Overview_High-level_Architecture]

Global Namespace

Ring-wide POSIX semantics

- managed via Conspiracy

Per Share

- policies

- features

- workload specific layouts

No filesystem limits

- unlimited nodes

- unlimited user data objects

- billions of objects per node!

Shared nothing

- no dedicated metadata server

- peer to peer object and metadata distribution

Ring to Ring replication

- snapshot based

- WAN optimised

Raft Consensus Manager - Conspiracy

Enables Large Clusters

- robust failover

- automatic discovery and growth

- manages ring membership (nodes, drives)

Manages Persistent Configuration

- users, groups

- shares

- Encryption keys

Manages Volatile Operational States

- Lock state

- namespace delegations

- drive telemetry

- environmental state (RPM, temps)

Manages Persistent Operational State

- snapshot state
- Drive membership

CDP Snapshots

Inline CDP Snapshots

Reduplicated against

- live filesystem
- all other snapshots

End user access via virtual namespace hierarchy

Per share snapshot policies

Snapshot replication

- Async snapshot-by-snapshot WAN optimised

On demand Promotion

- R/O >> R/W is O(1)

OneBlox - Object Store

Object Distribution

Object Router

- Multiplexes objects across constituent ODBs
- Manages object durability
- Manages object lifetimes

Object Database

- Persistently stores objects and Attributes
- represents a replica of a shard

Benefits

dedupe

data protection

shared-nothing, scale-out architecture

Consistent Hashing and Content Addressability

It starts with objects

- objects have a key, and a value
- the value is the content

Consistent Hashing

- Maps the key to a database shard via a function

Content Addressability

- Generates the key via a function applied to the object data
- Same data -> Same key
- Different Data -> different key

Why do we care?

- no need for metadata servers
- deduplication
- end-to-end data integrity

Consistent Hashing - Exablox Style

Hash function

- cryptographic hash

- good distribution

Shards

- dynamically assigned to ODBs based on ODB capacity
- Mapping stored in Conspiracy

Store

- Hash(Key) -> Shard
- Write data to the ODB
- Write replicas to adjacent shards

Retrieve

- Hash(Key) -> Shard
- Retrieve data from the device connected to the shard

Recover

- Assign failed shard to a replica device
- Rebalance

Object DB Journal

Fast, persistent attribute DB

- key for enabling dedupe
- key to scalable garbage collection and object iteration across huge clusters

Append only journal

- key to supporting converged primary and backup workloads
- enables high performance sequential and random writes and overwrites
- enables continuous snapshots
- co-residing data and metadata enables small files performance avoiding seeks

OneBlox and OneSystem Demo

Sean Herrington, Senior Director Product Management

Typical power usage is 100W for a fully loaded system

Embargo Section