

Also check out Cormac's post - <http://cormachogan.com/2015/02/18/a-closer-look-at-springpath/>

Corporate Overview

Out of stealth about 3 weeks ago

Krishna (Architect for VMFS) and Mallik (invented VxLAN, SR—IOV) founders
Experienced team with 200+ patents in their careers

Product Overview

Ravi Parthasarathy - VP, Product Management

Essentially storage sw deployed on commodity hw providing an enterprise class solution

- enterprise grade, scale out
- maximum simplicity
- completely software

Enterprise Grade

- robust, resiliency and data integrity
- data mirroring and automatic rebalancing
- flash / memory perf
- native, space efficient snapshots
- VM / VVOL / File granularity
- inline dedupe and compression
- lower \$/GB using high capacity 7.2K RPM drives

Maximum Simplicity

- leverage existing mgmt tools
- zero learning curve
- no legacy storage complexity
- rapid provisioning of applications
- cloud based auto-support monitoring
- proactive alerts and rapid resolution

SW economics

- choose your (prescribed) servers
- choose your platform (VMware 5.5 and above, OpenStack / KVM will be offered in beta shortly)
- annual subscriptions, per server, including support
- any server, any capacity
- upgrade servers without sw tax
- scale out compute or perf or cap
- just-in-time scaling in small increments

Architecture Deep Dive

Mallik Mahalingam (Founder, CEO and CTO) and Krishna Yadappanavar (Founder)

[photo]

Data Platform

100% sw
elastic scaling
enterprise grade
integrates into existing mgmt tools

data mgmt storage sw on commodity hw, without compromising features, scale or performance

Design Goals

- scale out perf and cap linearly
- scale out caching tier independently from cap tier, with losing data mgmt features
- leverage flash for perf and low speed hard disks for capacity
- maximise utilisation of free space in flash or hard disks, when nodes appear / disappear in cluster
- maximise space usage using inline compression and inline deduplication in all tiers
- pointer based file level snapshots and clones
- support a variety of platforms (VMware, KVM, Hyper-V, Containers ...)
- leverage existing mgmt applications and frameworks

Scale out and Distributed File System

- Start with as few as 3 servers
- Springpath sw cluster installs in minutes
- add servers, one or more at a time
- distribute and rebalance data across servers automatically
- retire older servers

Independent scaling of compute, cache or capacity

HALO Architecture

[photo]

Hardware agnostic log-structured object

*Data Access Layer

VMware

ESXi

NFS/VAAI/VVOL

Springpath data Platform

KVM

NFS/Cinder/Nova/Glance

Springpath data Platform

Hyper-V

SMB

Springpath data Platform

*Data Distribution

Avoid controller hotspots

leverage cache across all SSDs in the cluster

*Data Virtualisation - Caching
Striping across and within the VMs

Take a stripe and route it to one of the cache vNodes

Wanted to “decouple the ability to server the data from the location that you’re serving it from”

Rebalance cache on node addition or removal

Data Virtualisation

Write back caching to SSDs with mirroring

- all writes to cache vNodes go to a write log on SSD
- synchronously mirror one or two copies for HA
- acknowledge after mirror writes are complete

Max write size is 64K

De-staging of write log (write log is currently 2GB)

- writes are de-staged from write log to data and metadata vNodes
- data and metadata are mirrored to one or two nodes for HA
- data can be de-staged to a local or different server based on available space

Uniform Space Utilisation

- utilise free capacity when new nodes are added
- faster rebuilds

Read caching

- data is cached in both memory and SSD for reads
- misses are fetched from HDDs from any node in the cluster

*Data Optimisation

Inline dedupe and compression

- inline, dedupe of memory, SSD and HDD
- striping enables dedupe across files
- inline compression on SSD and HDD

Data Management

Native Snapshots

- Pointer Based Snapshots - fast creations and deletions, no consolidation overhead
- Fine-grained or coarse-grained - VM-level or VM folder level
- VAAI / Cinder integrated - quiesced and crash-consistent
- Use vCenter Snapshot Manager
- Policy Based - schedules, retention period

Native Clones

- pointer based writeable snapshots
- VM-level
- VAAI integrated
- Batch version GUI - clone names, use customisation spec

Integrated cluster mgmt approach

- extend compute mgmt consoles - no separate we mgmt GUI
- vSphere Web client plug-in - for specific features like performance metrics and cluster status

*Summary

Technology

- Log structured layout
- Data virtualisation
- data distribution
- data services
- integrated mgmt

Benefits

- Flash endurance, compression friendly, faster rebuilds
- scale performance and capacity independently, eliminate hotspots
- granular scaling and rebalancing
- fast efficient snapshots and clones
- reduced management ...

Product Demo

Will and Sean

Closed Session