

Primary Data

Company Background

Shannon McPhee, VP, Communications

Founded August 2013, exited stealth nov 2014

\$60 million in VC funding

Chief Scientist - Woz

Employ more than 80 ppl

Expecting to ship later this year

Who is Primary Data?

Lance Smith, CEO

[photo]

Automated data mobility through data virtualisation

Intelligent Policy-Driven Environments

Dynamic Data Mobility - see and place data across all storage resources within a single global data space

policy-driven agility - non-dis policy-based data movement

intelligent automation - automated real-time infrastructure dynamically aligns supply demand

linear scalability - perf and cap linearly and incrementally

global compatibility - single hardware-agnostic solution enhances coexisting legacy it and modern scale-out and hybrid cloud architectures

What is Data Virtualisation?

David Flynn, CTO

Whiteboard session

3 needs - Performance, Price, Protection (Fast, Safe, Cheap). You can't have all three at the same time

This isn't going to change

Managing data via the storage container that holds it. This is the tail wagging the dog.

Separate the metadata from the data

Take the routing function from out of the storage array (data hypervisor)

Uses pNFS client

File on File, on Block, on Object, on DAS

Block as File

Object

Splits the metadata and control path off to the side

80% of IOPS to primary storage is to storage that doesn't need to exist after a crash (temp, scratch, swap, etc)

When VMware did virtualisation

1. Utilisation - doorknocker

2. Manageability - this is what made them stick

--

Doing Data Virtualisation also includes

3. Performance

Once you've virtualised the data, the problem becomes setting the objectives for the storage and the needs of the data

Synchronous local replication

Primary Data UX Preview

Sidharth Saxena, User Experience Architect

Unless you can visualise what you're trying to control through automation it's not going to be as useful

Directories / Pattern, Service Level Objectives (SLO), Performance profile, availability, security, Off-site
Granularity is at a file level

Back to David Flynn

Question around backup (from David). It's a lot of work. What if we re-think it?

Something a bit similar to SnapVault

Same namespace

But if you pull out the metadata, if you lose that the data is no good to you.

Periodically places copies of the metadata into the storage container so that it's self-describing

Why Do I need Data virtualisation?

Kaycee Lai, SVP Products & Sales

policy-based automation and orchestration

Customer use cases

Large global financial institution

(100000 VMs on 20 - 30000 hosts; 140000 virtual desktops)

Challenges

- claiming only 20% utilised storage across all tiers
- storage over-provisioning, labor intensive data mgmt ops - at scale and across heterogeneous env

Primary Data

- reduce over provisioning and increase storage util
- improved responsiveness
- simplify mgmt of storage for VMs

Large Semiconductor manufacturing company

Challenges

- cost to manage data is high (600TB is unmanaged by IT)

- management complexity in using multiple point solutions

Primary Data

- reduce cost of storage and heterogeneous solutions
- protect UX with data-driven methodologies over data management

Top 100 Health Science Company

Challenges

- Manage cost associated with exponential data growth, looking to reclaim expensive t1 storage
- onsite archiving platform insufficient to deal with growth

Primary Data

- Automate lifecycle of data growth from production to archive
- Simplify Management across storage resources from numerous vendors

Fortune 150 High Tech company

Challenges

- Managing data across 8 DCs is a “nightmare
- Exponentially more complex in 50PB environment

Primary Data

- Simplify Management by scaling out storage with non-disruptive data mobility
- reduce costs by reducing over provisioning and increasing storage utilisation
- accelerate innovation as development data follows the sun