



Exploiting SSDs with ZFS

Brad Stone, VP Product Management
Nexenta Systems



Nexenta Systems

Mission: Enterprise-class storage for everyone

- Privately held
- Founded in 2005
- Founders developed iSCSI stack for Linux
- Core product: NexentaStor

“Fastest growth in commercial deployments in the last decade.”



What is NexentaStor?

File and Block Access



Software-based, unified storage appliance

Leading OpenStorage solution

- *Runs on standard hardware*

Key features:

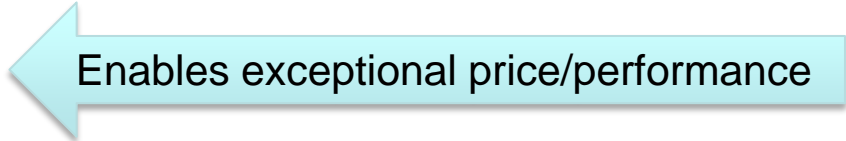
- End to end data integrity
- Detect and correct data corruption
- Unlimited file size & snaps
- Synchronous and asynchronous replication

 Superior storage for virtualized environments



Zetta File System

- 256-bit checksums
- Hybrid storage pools
- Instantaneous snapshots
- Thin provisioning
- In-line compression
- In-line and in-flight de-duplication
- In-line virus scan



Enables exceptional price/performance



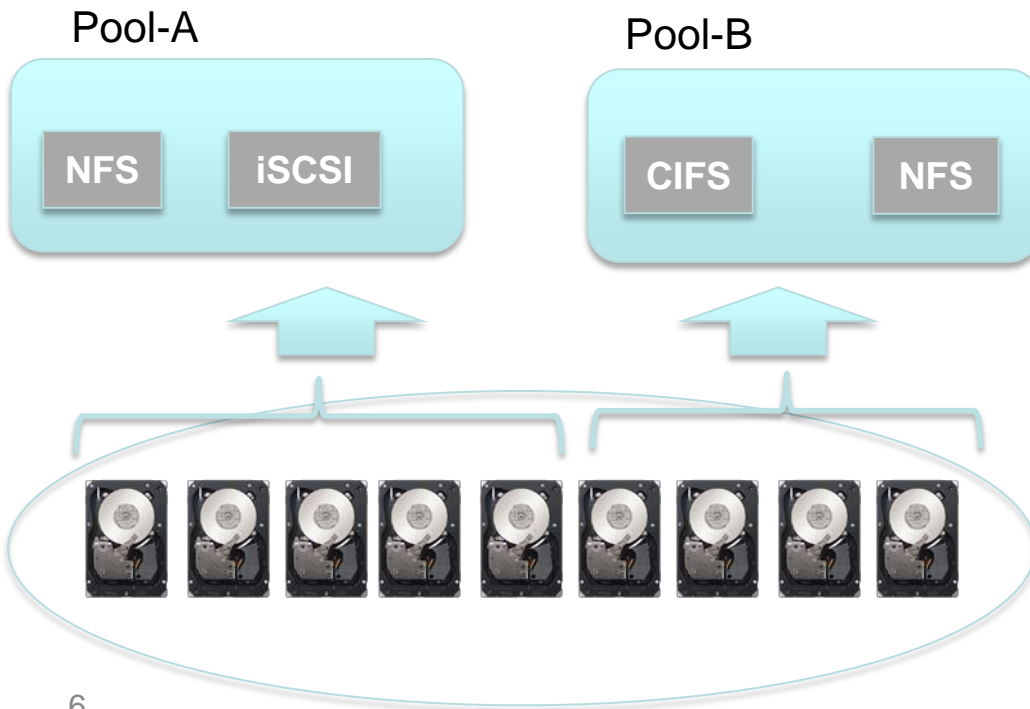
Efficient use of storage

ZFS – extraordinary scalability

Description	Limit
Number of data volumes on a system	2^{64}
Maximum size of a data volume	2^{78} bytes
Number of file systems in a data volume	2^{64}
Maximum size of a file system	2^{64} bytes
Number of devices in a data volume	2^{64}
Number of files in a directory	2^{56}
Maximum file size	2^{64} bytes
Number of attributes of a file	2^{48}
Maximum size of any attribute	2^{64} bytes
Number of snapshots of a file system	2^{64}

Unlimited snapshots with integrated search

ZFS Components



Discover logical disks
Determine RAID requirements

Create storage pools

Create datasets
Assign properties (e.g. de-dup)
Share (NFS, CIFS, iSCSI, FC)

Hybrid Storage Pools



Read Cache



Transaction log



Enables outstanding price/performance

ZFS Intent Log

- Choose write-optimized SSD for ZIL
- Deploy ZIL as transaction log
- Eliminates latency of writes to spinning disks
- Maintains data integrity

But it depends.
Asynchronous writes see no benefit.

- Can easily see 5-10x performance benefits

- Choose read-optimized SSDs for read cache
- Optimal size depends on workloads



NexentaStor Deployment Options

- High Availability Cluster
- Asynchronous Replication
- Synchronous Replication

Use Cases



Virtualization



Cloud Storage



VDI



Backup



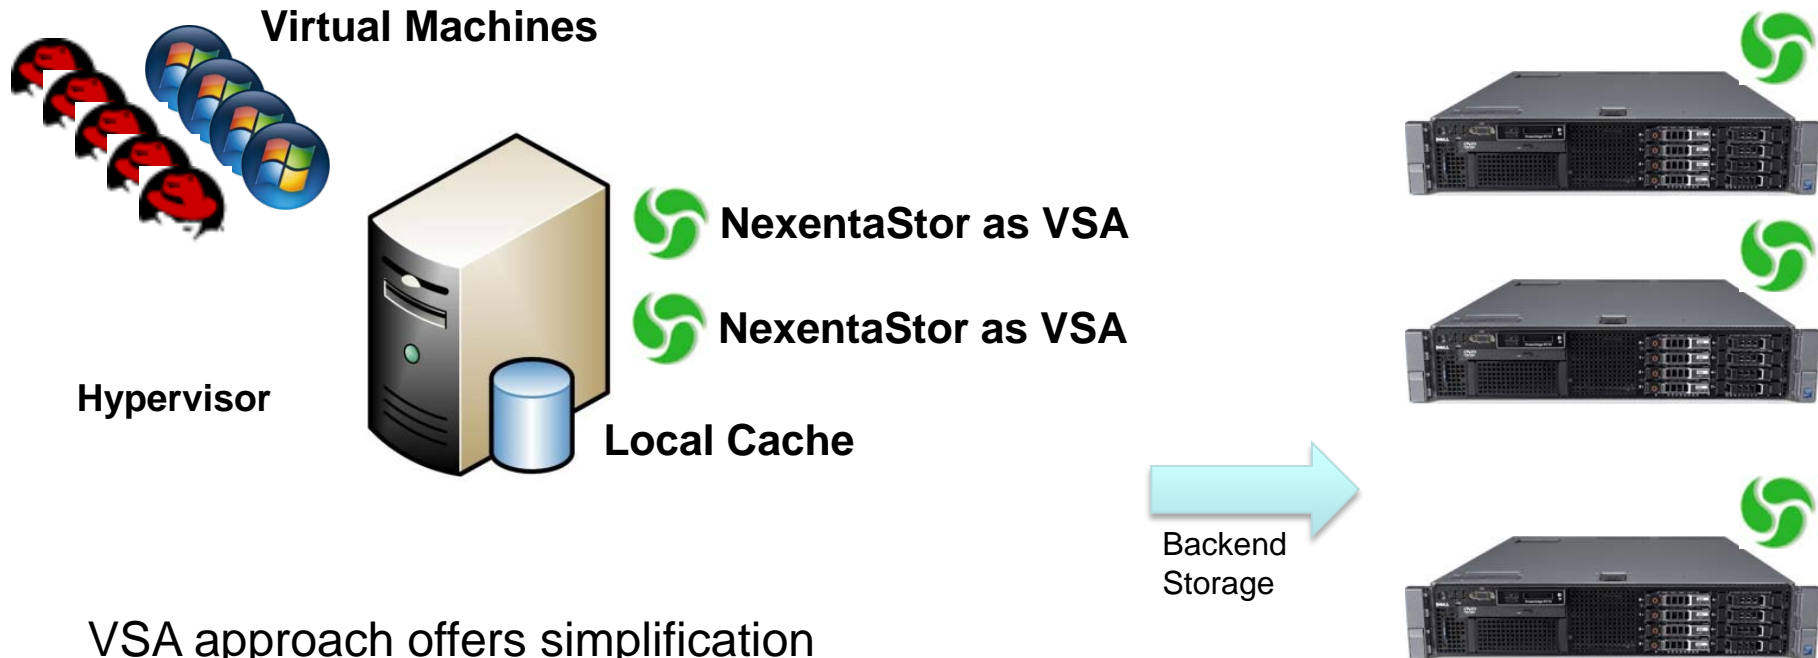
User Home Directories

Use Case: VDI

VDI is an intensive I/O workload (95% small random writes)

NexentaStor handles workload with ZIL devices

Together with Citrix achieved 600 microsecond latencies on 100-user test



VSA approach offers simplification
“VDI-in-a-Box”

Case Study: Virtualization

- Broadband Television Provider
- Business Problem:
 - 30 VMs on a server required high IOPS to Oracle database
 - Netapp FAS2020 couldn't meet requirements
- Solution:
 - NexentaStor with ZIL and L2ARC
 - Achieved 100,000 IOPS
 - Beat Netapp performance by 1000%

Case Study: Storage Backup

- Healthcare Provider
- Business Problem
 - Storage growth impacting ability to complete backups within maintenance window
- Solution
 - NexentaStor with ZIL
 - Backup window reduced more than 75%



THANK YOU!!

Some Notes about this template

- The first action you should take is to save this presentation
 - You have opened a design template (.pot)
 - Need to save as .ppt
- A master exists for:
 - Slides
 - Handouts - default is 3 to a page
 - You can print a different number, but no guarantees about appearance
 - Notes



Free-format slide title