Software-Defined Storage at Microsoft

Jose Barreto
Principal Program Manager
Microsoft

Industry trends

What is Software Defined Storage?

Software intelligence delivering feature-rich cloud scale storage and economics built on industry standard hardware

Cloud inspired infrastructure and design

- Industry-standard hardware
- Integrating cloud design points in software
- Driving cloud cost efficiencies

Data explosion

- Device proliferation
- Modern apps
- Unstructured data analytics

Evolving technologies

- Improved SSDs, NVMe, NVDIMM
- Network performance improvements
- Maturity in software-based solutions

Scale out maturity

- Integrated solutions
- Rapid time to solution
- Policy-based management

Microsoft Azure: Enterprise-grade software-defined infrastructure that's proven at hyper-scale

Unparalleled productivity

Provision apps in minutes

57% of Fortune 500 companies run their applications on Azure

Maximum flexibility

Add1000 customers/day

Run 1000s of diverse apps

Extreme efficiency

10s of thousands of network changes every day

Enterprise-grade reliability

Financially backed SLA

Cutting-edge protection

Compliance with world class industry standards

Microsoft enterprise storage

Feature-rich software defined storage in your datacenter and in Azure

Basic storage capabilities have become mere “table stakes”

The value of storage solutions is shifting toward software

The cloud decentralizes your data and your storage

Identity and storage management solutions will evolve with the cloud

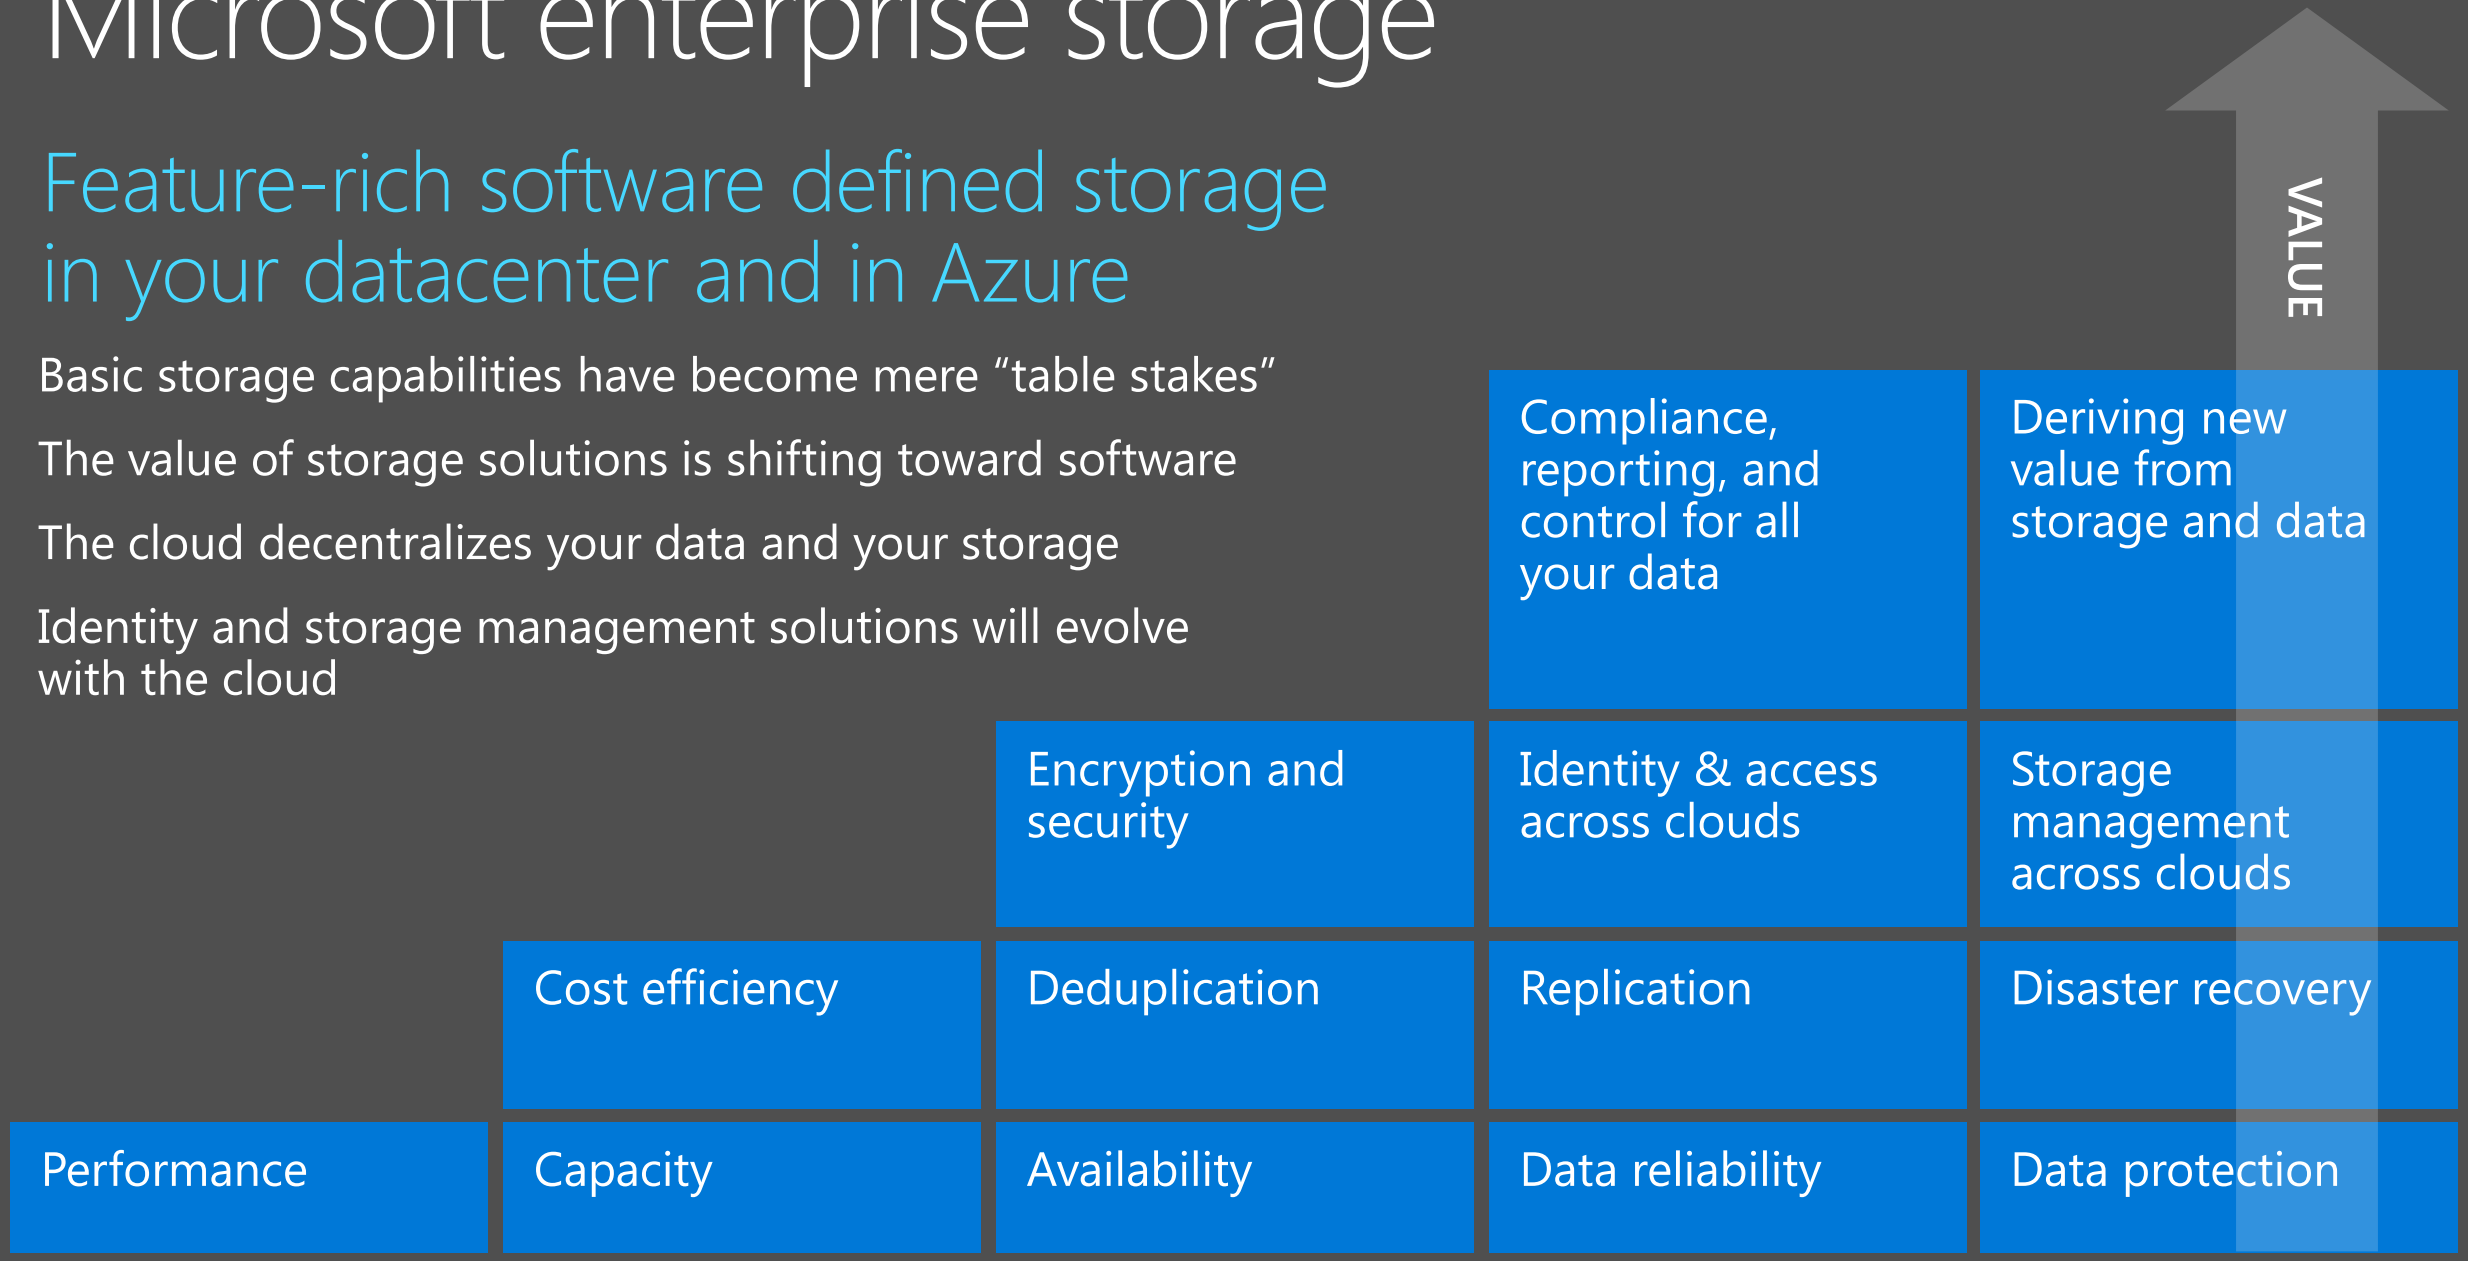

Customer choice

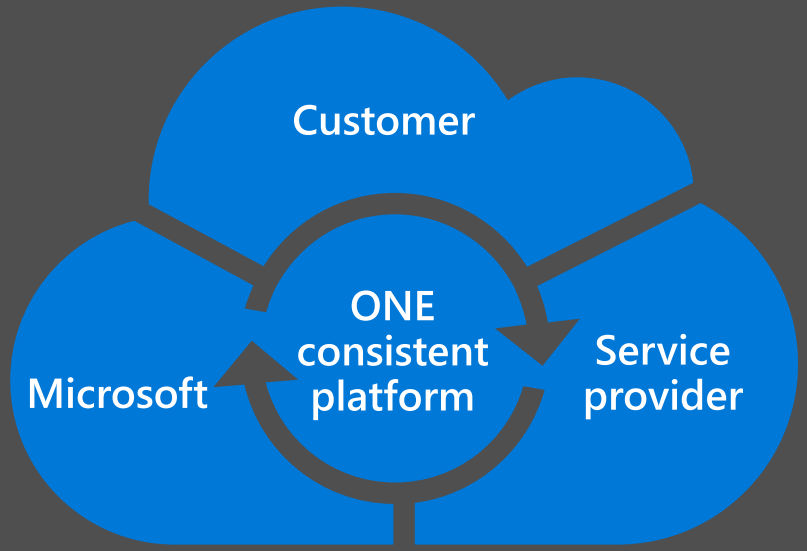

Private cloud with partner storage	Private cloud with Microsoft SDS	Hybrid cloud Storage	Public cloud Storage
SAN and NAS storage	Scale out file server with storage spaces	StorSimple with Azure storage	Azure storage

Microsoft Software Defined Storage (SDS)

Breadth offering, unified platform for Microsoft workloads and Linux
public cloud scale and cost economics for private cloud customers

Where is Microsoft today in the SDS journey?

Microsoft software defined storage deployment

Primary application data storage on cost effective, continuously available, high performance SMB3 file shares backed by tiered storage spaces

1. **Performance, scale:** SMB3 File Storage network
2. **Continuous availability and seamless scale-out** with File Server Nodes
3. **Elastic, reliable, optimized** tiered Storage paces
4. **Low cost** standard volume hardware
5. **System Center:** Unified storage management

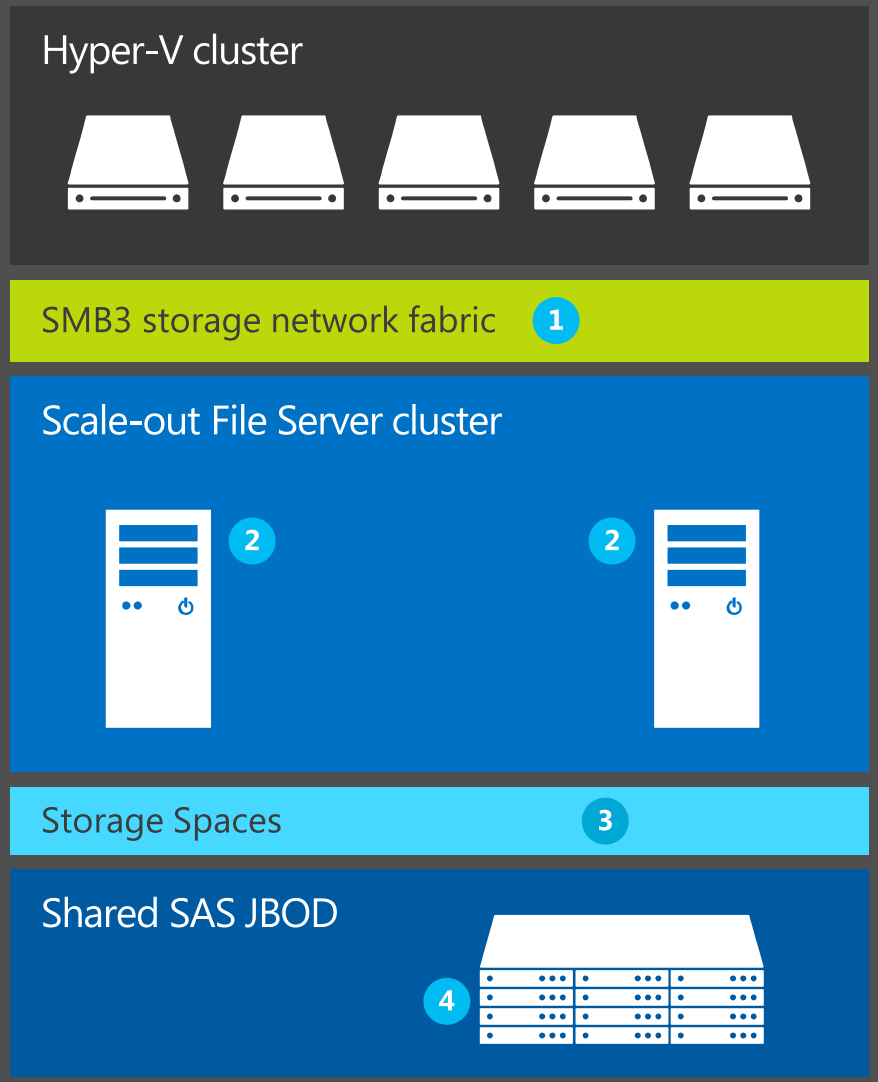

Storage Spaces and Scale-Out File Server

Virtualized storage—high performance storage for your private cloud

Spotlight capabilities

Virtualize storage: Transform low-cost, high volume hardware with Storage Spaces.

High performance: Combine HDDs and SSDs in tiered Storage Spaces to meet demand from intensive workloads.

Reduce costs: Eliminate complex and costly storage area network (SAN) infrastructure without sacrificing SAN-like capabilities.

Flexible: Independently scale capacity and compute to grow with your business demands.

Resilient: Multiple layers of redundancy across disk, enclosure, connectivity and file server nodes ensure highest availability.

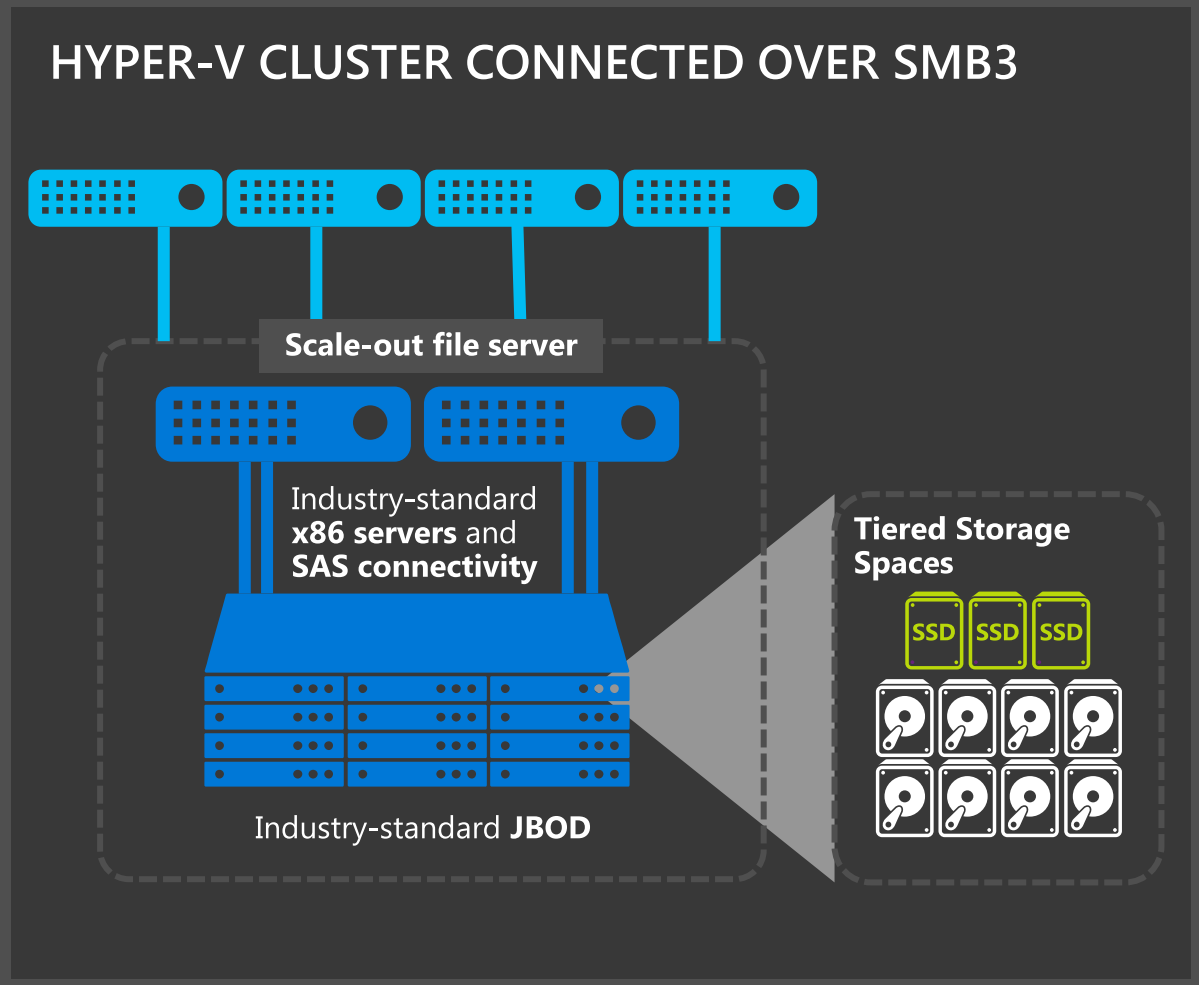

Unified storage management

Spotlight capabilities

Virtualized storage provisioning: Rapidly deploy Scale-Out File Servers, from bare-metal and provision file shares to Hyper-V hosts.

SAN integration: Centralized block and file-based storage area network (SAN) management, enabling discovery, classification, provisioning allocation and decommissioning.

Granular monitoring: Gain deep health and availability insight into storage infrastructure across Storage Spaces and SAN.

Operational insight: Get deep visibility into your data center capacity, pinpoint capacity shortages, investigate "what-if" scenarios and plan future storage needs.

Cloud-scale protection: Back up important data across physical or virtual file servers, to disk, tape or cloud.

Customer Stories

"By using Storage Spaces, we're getting storage performance of 450,000 IOPS and roughly a gigabyte per second in throughput—more than 20 times the IOPS and four times the throughput of our SAN."

— Daniel Weissenborn,
Solution Architect, ClearPointe

"In the past, a customer bought a SAN with expensive RAID [redundant array of independent disks] drives for an initial purchase price of more than $3,000 a terabyte. With Storage Spaces, our customers can use storage that costs less than $300 a terabyte."

— Rand Morimoto,
President, Convergent Computing

Not just bits. SDS solutions.

Cloud Platform System (CPS)
Integrated solution for HW and SW

Per rack
(1-4 racks)

512 cores
8TB RAM
262 TB usable storage

1360 Gb/s internal rack connectivity
560 Gb/s inter-rack connectivity

60 Gb/s external

2322 lbs.
42U
16.6 KW maximum

Networking
5 x Force 10 – S4810P (64 port @ 10GbE - Data)
1 x Force 10 – S55 (48 port @ 1GbE – Management)

Compute Scale Unit (32 x Hyper-V hosts)
Dell PowerEdge C6220ii – 4 Compute Nodes per 2U
- Dual socket Intel IvyBridge (E5-2650v2 @ 2.6GHz), 256 GB memory
- 2 x 10 GbE Mellanox NIC's (LBFO Team, NVGRE offload)
- 2 x 10 GbE Chelsio (iWARP/RDMA)
- 1 local SSD @ 200 GB (boot/paging)

Storage Scale Unit (4 File Servers, 4 JBODS)
Dell PowerEdge R620v2 (4 Server for Scale Out File Server)
- Dual socket Intel IvyBridge (E5-2650v2 @ 2.6GHz)
- 2 x LSI 9207-8E SAS Controllers (shared storage)
- 2 x 10 GbE Chelsio T520 (iWARP/RDMA)

PowerVault MD3060e JBODs (48 HDD, 12 SSD)
4 TB HDDs and 800 GB SSDs

Microsoft Azure StorSimple

StorSimple hybrid cloud storage delivers primary storage, data protection and disaster recovery in a single solution that integrates Azure storage services

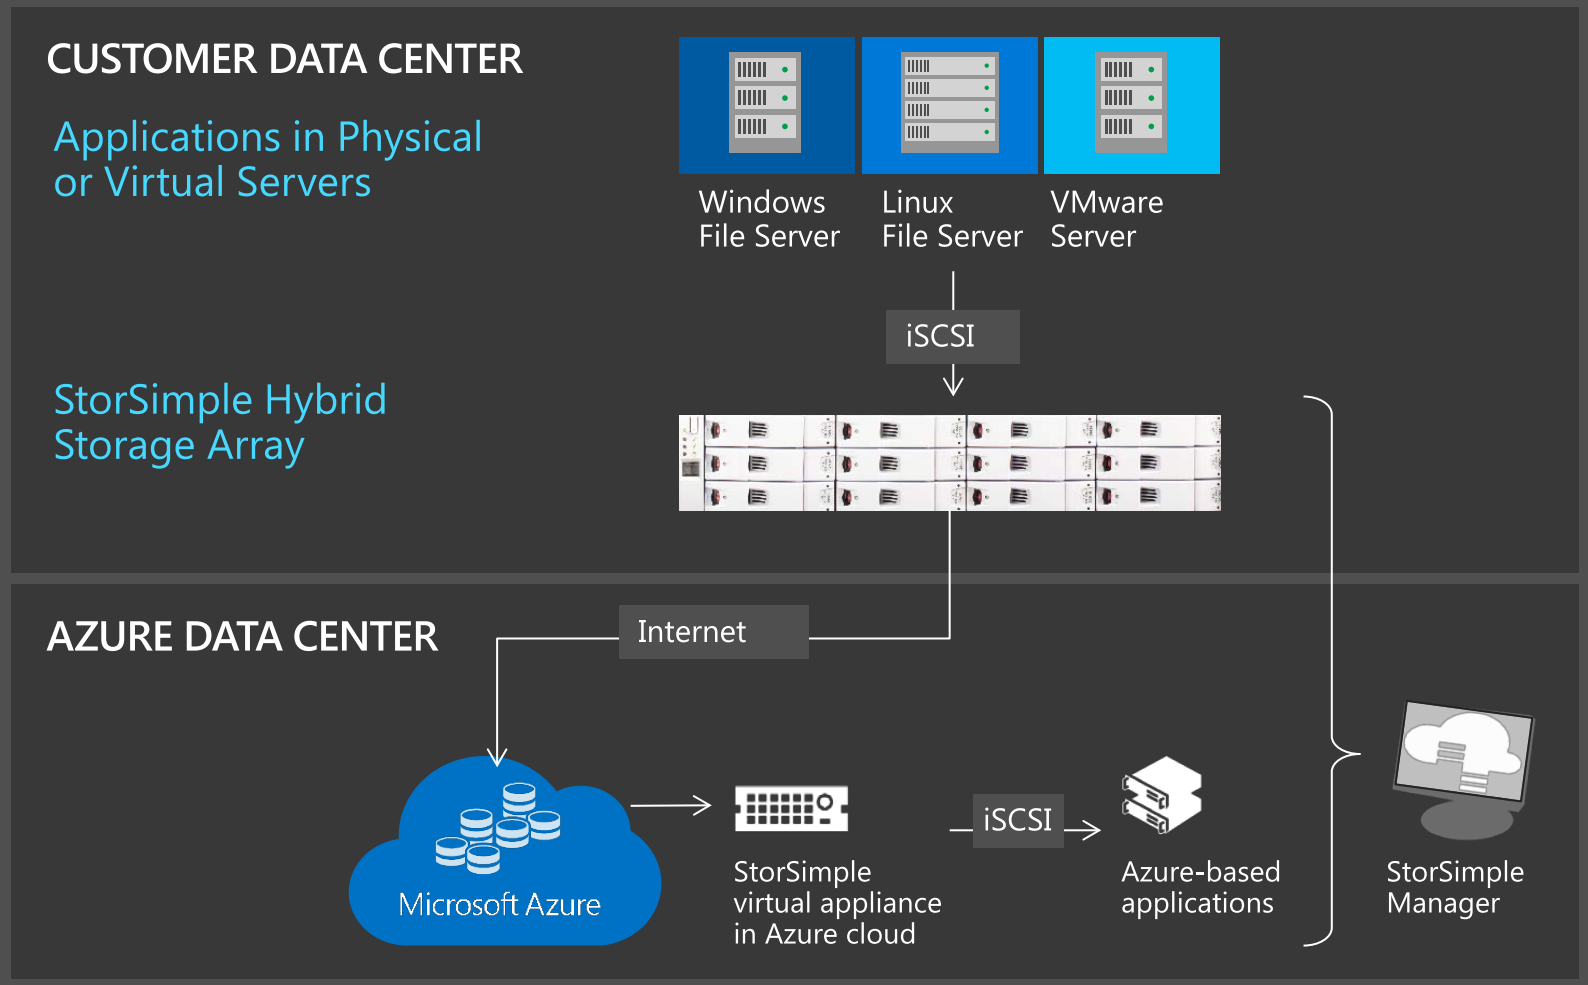

Enterprise SAN storage

Inline de-dupe, compression, and automatic tiering

Automated offsite data protection using cloud snapshots

Highly efficient, location independent disaster recovery

Data mobility for enterprise data

Consolidated storage and data management

What's in the Windows Server Technical Preview?

Continuing the SDS journey…

Reliability

Cross-site availability and disaster recovery

Improved tolerance to transient failures

Scalability

Manage noisy neighbors and demand surges

Deploy mixed workloads in a shared environment

Manageability

Easier migration to the new OS version

Improved monitoring and incident response

Reduced cost

More cost-effective by using volume hardware

Use SATA and NVMe, in addition to SAS

Storage Quality of Service (QoS): Greater efficiency
Control and monitor storage performance

Simple out of box behavior
- Enabled by default for Scale Out File Server
- Automatic metrics (normalized IOPs and latency) per VM and VHD

Flexible and customizable
- Policy per VHD, VM, service, or tenant
- Define minimum and maximum IOPs
- Fair distribution within policy

Management
- System Center VMM and Ops Manager
- PowerShell built-in for Hyper-V and SOFS

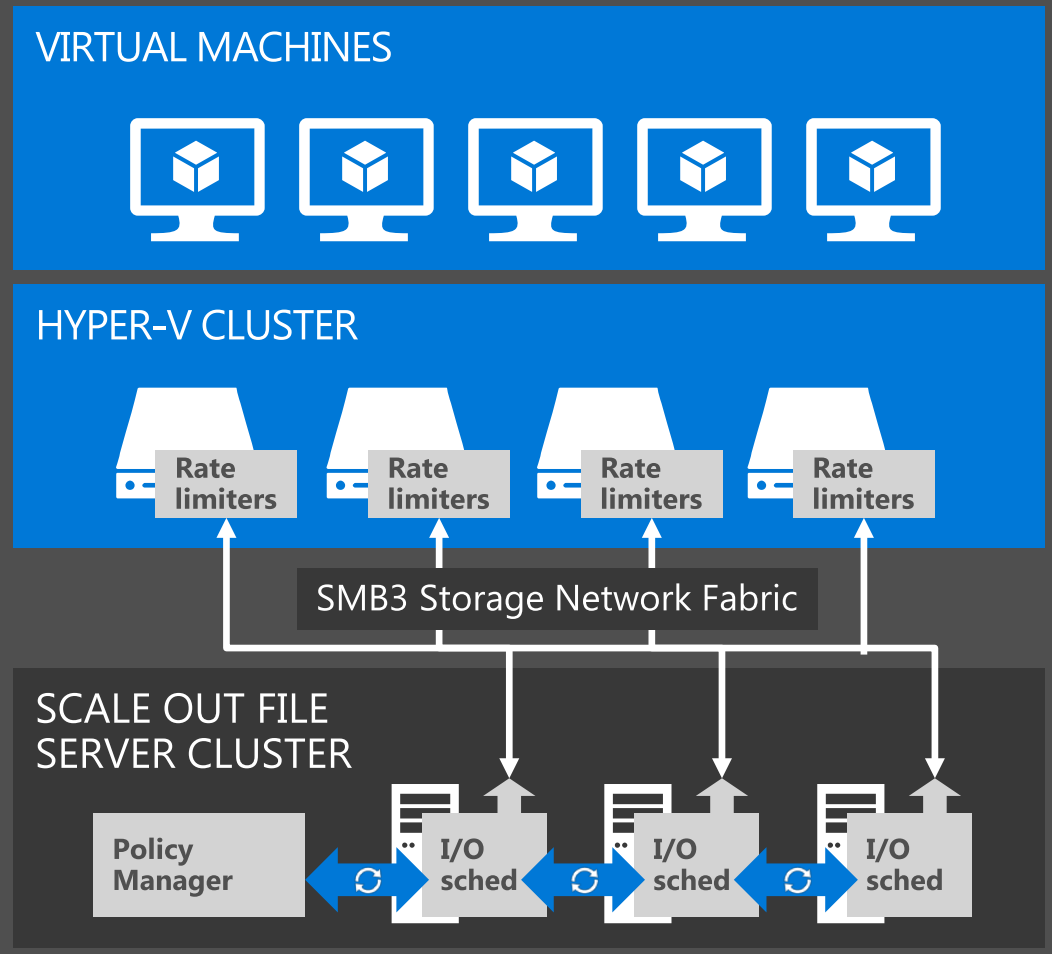

Rolling Upgrades – faster time to value

Simple

Rolling Upgrades with Win2012 R2 and vNext nodes within the same cluster

Easily roll in nodes with new OS version

Seamless

Zero downtime cloud upgrades for Hyper-V and Scale-out File Server

STORAGE (SCALE OUT FILE SERVER) OR HYPER-V CLUSTER

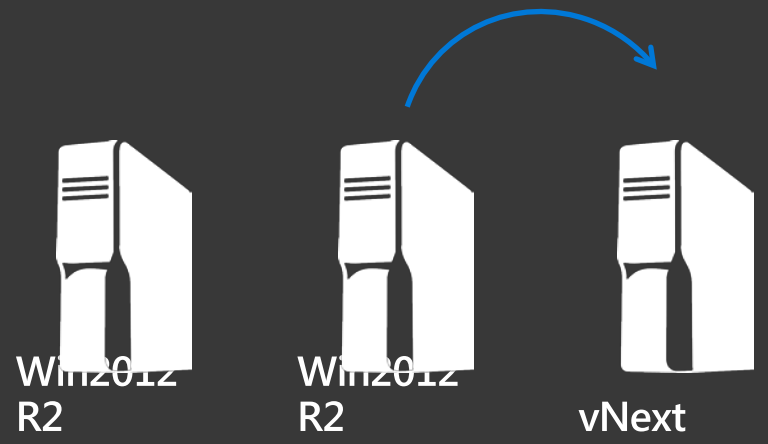

VM storage resiliency: Reliability

Resiliency

Designing for cloud scale with standard hardware

Preserve tenant VM session state in the event of transient storage disruption

Visibility

VM stack quickly notified on failure

Intelligent and quick VM response to block or file based storage infrastructure issues

Reliability

VM moved to PausedCritical state and will wait for storage to recover

Session state retained on recovery

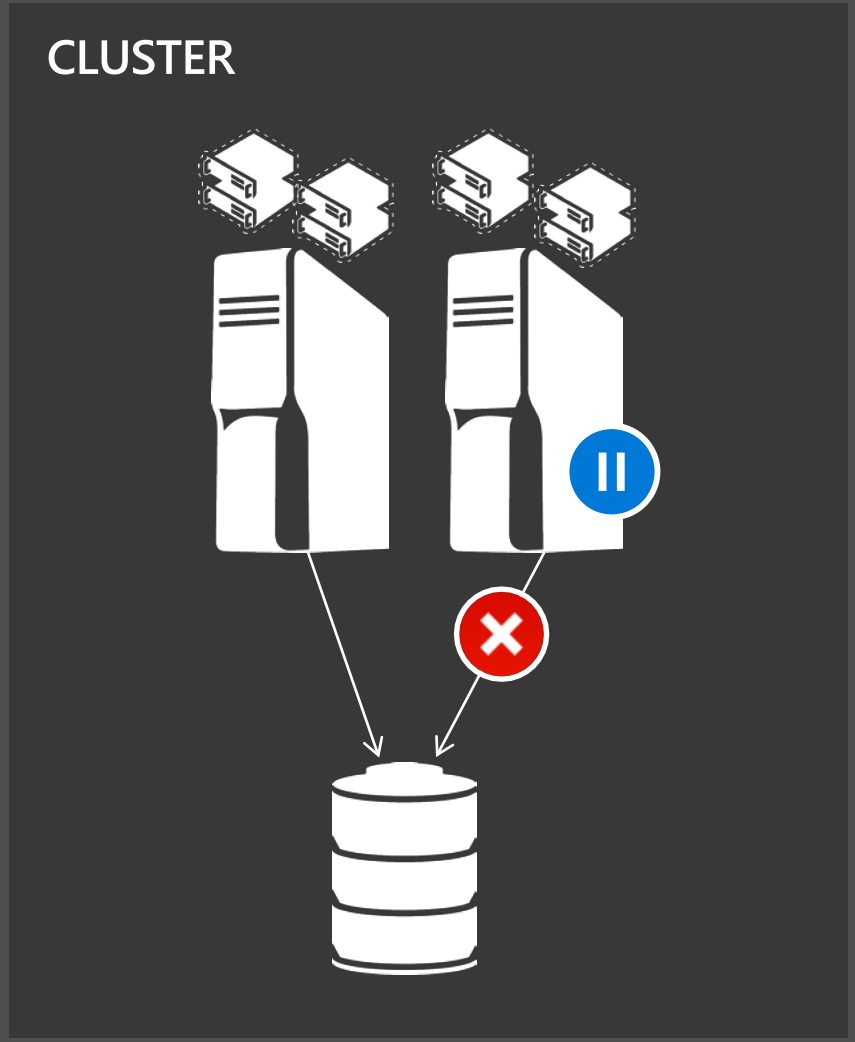

Storage replica

Protection of key data and workloads

Spotlight capabilities

Synchronous replication: Storage agnostic mirroring of data in physical sites with crash-consistent volumes ensuring zero data loss at the file system level.

Increase resilience: Unlocks new scenarios for metro-distance disaster recovery and multi-site failover clusters for automated high availability.

Complete solution: End to end for storage and clustering, including Hyper-V, Storage Replica, Storage Spaces, Cluster, Scale-Out File Server, SMB3, Deduplication and Resilient File System (ReFS), New Technology File System (NTFS).

Streamlined management: Graphical management for individual nodes and clusters through Failover Cluster Manager or Windows PowerShell.

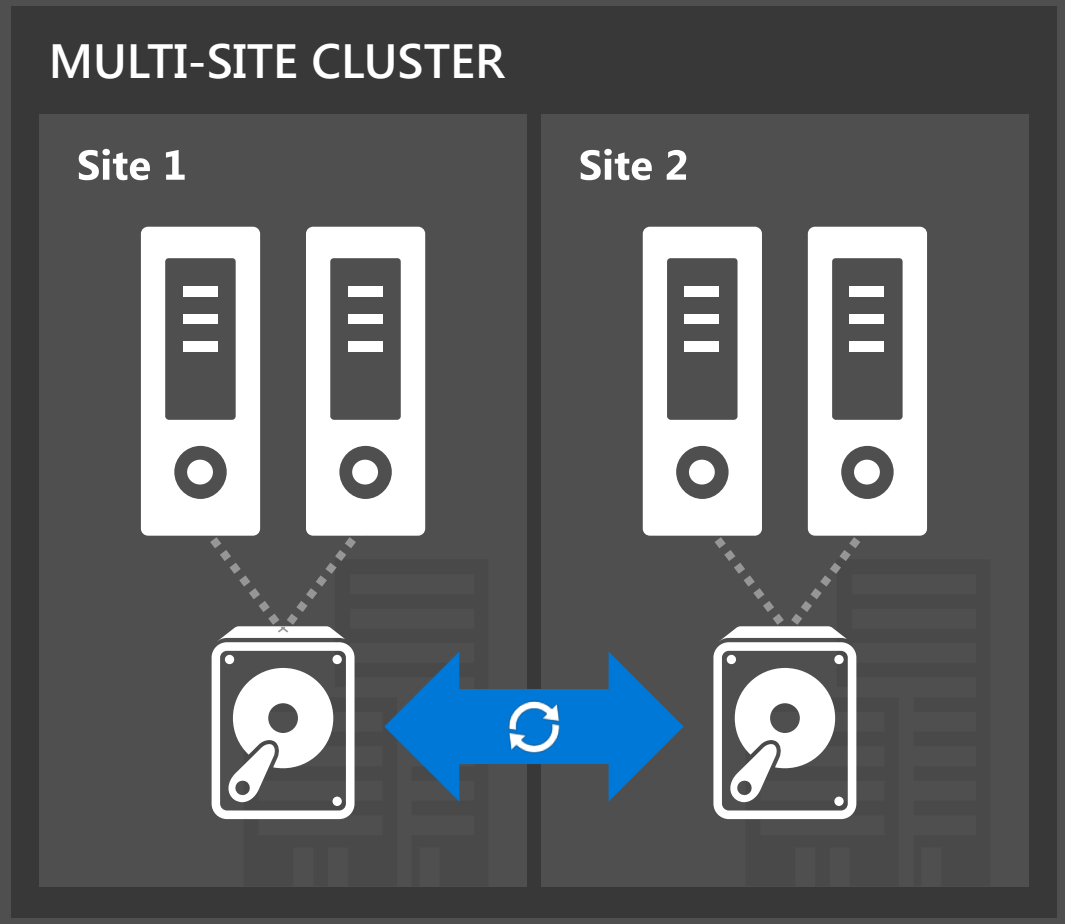

Storage Spaces Shared Nothing – Low cost

SoFS clusters with **no shared storage.** Doesn't need shared JBODs and SAS fabric behind Scale-Out File Server nodes

Cloud design points and management

- Prescriptive configuration. Reduced hardware costs with SATA drives.
- Deploy, manage, and monitor with SCVMM, SCOM.

Reliability, scalability, flexibility

- Fault tolerance to disk, enclosure, node failures.
- Scale pools to large number of drive.
- Fine-grained storage expansion.

Use cases

- Hyper-V IaaS storage.
- Storage for backup and replication targets.

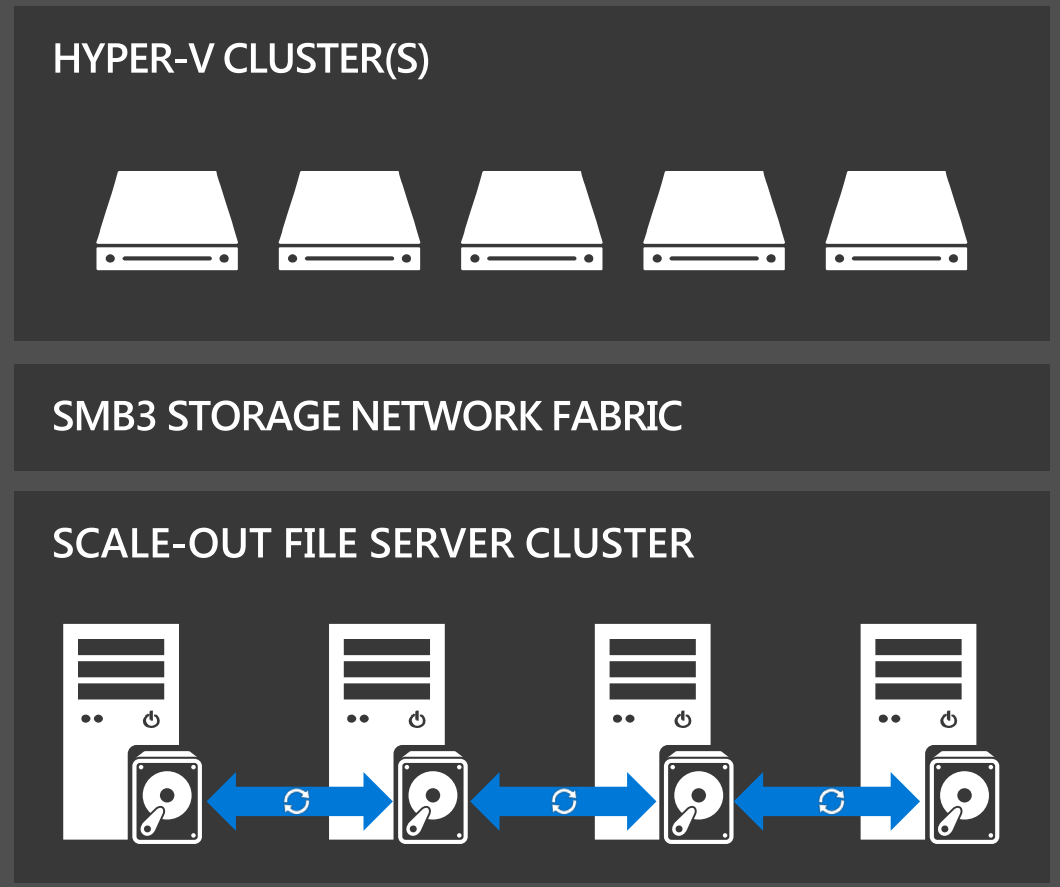

Customer Choice

Enterprise Storage choice: SAN or SDS

Traditional SAN

- Block protocol fabric
- Low latency network with FC
- Management of LUNs
- Data deduplication
- RAID resiliency groups
- Pooling of disks
- High availability
- Copy offload, Snapshots
- Storage tiering
- Persistent write-back cache
- Scale Up
- Storage QoS
- Replication
- Firmware Updates

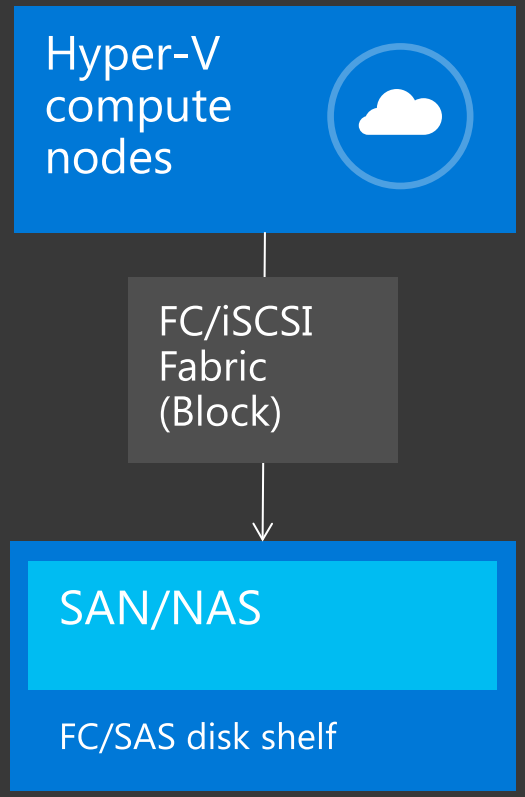

Microsoft SDS

- File protocol fabric
- Low latency with SMB3Direct
- Management of Shares
- Data deduplication
- Flexible resiliency options
- Pooling of disks
- Continuous availability
- SMB copy offload, Snapshots
- Performance with Tiering
- Persistent write-back cache
- Automatic Scale-Out Rebalancing

NEW in R2

- Storage QoS
- Storage Replica
- Rolling Cluster Upgrades
- Storage Spaces Shared Nothing

NEW in v.Next

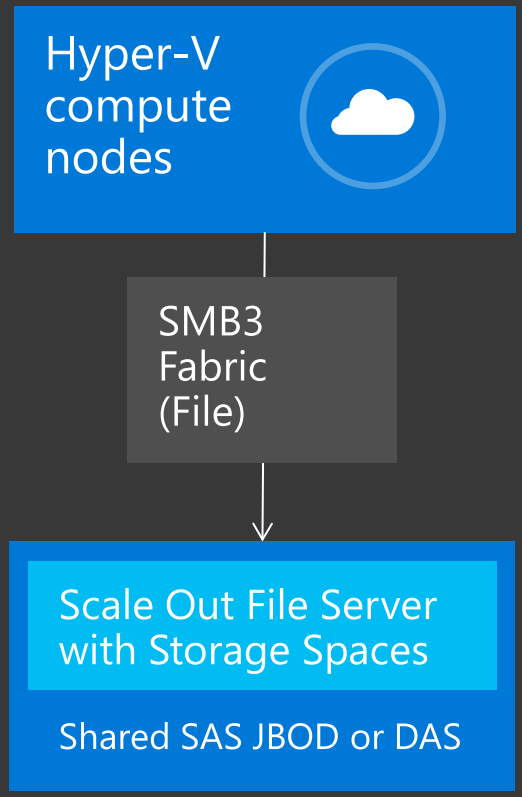

Customer choice

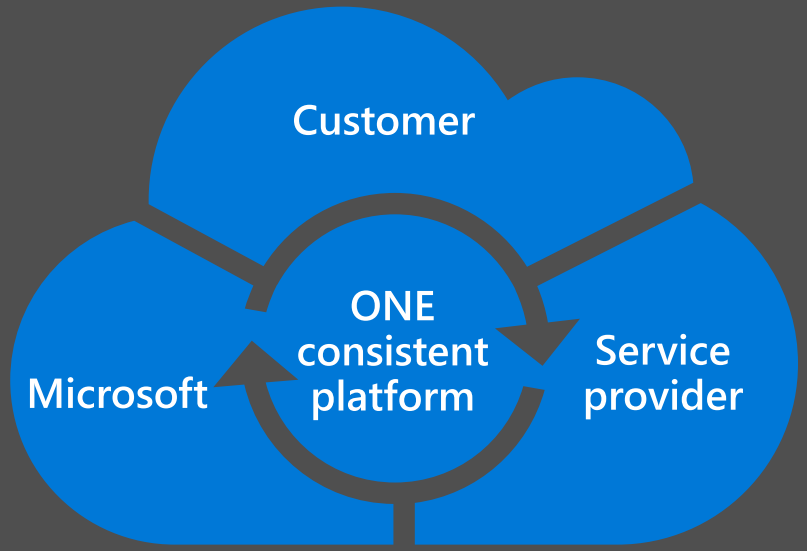

Private cloud with partner storage	Private cloud with Microsoft SDS	Hybrid cloud Storage	Public cloud Storage
SAN and NAS storage	Scale out file server with storage spaces	StorSimple with Azure storage	Azure storage

Microsoft Software Defined Storage (SDS)

Breadth offering, unified platform for Microsoft workloads and Linux
public cloud scale and cost economics for private cloud customers

Related content

Other DSI Sessions from Microsoft

- **Hybrid Cloud Storage with StorSimple with Mike Emard**
- **SMB Remote File Protocol (Including SMB 3.x) with Jose Barreto**

Sites and Blogs

- **http://www.microsoft.com/storage**
- **http://blogs.technet.com/b/josebda**
- **http://blogs.technet.com/b/filecab**
- **http://blogs.msdn.com/b/clustering**

Thank you!