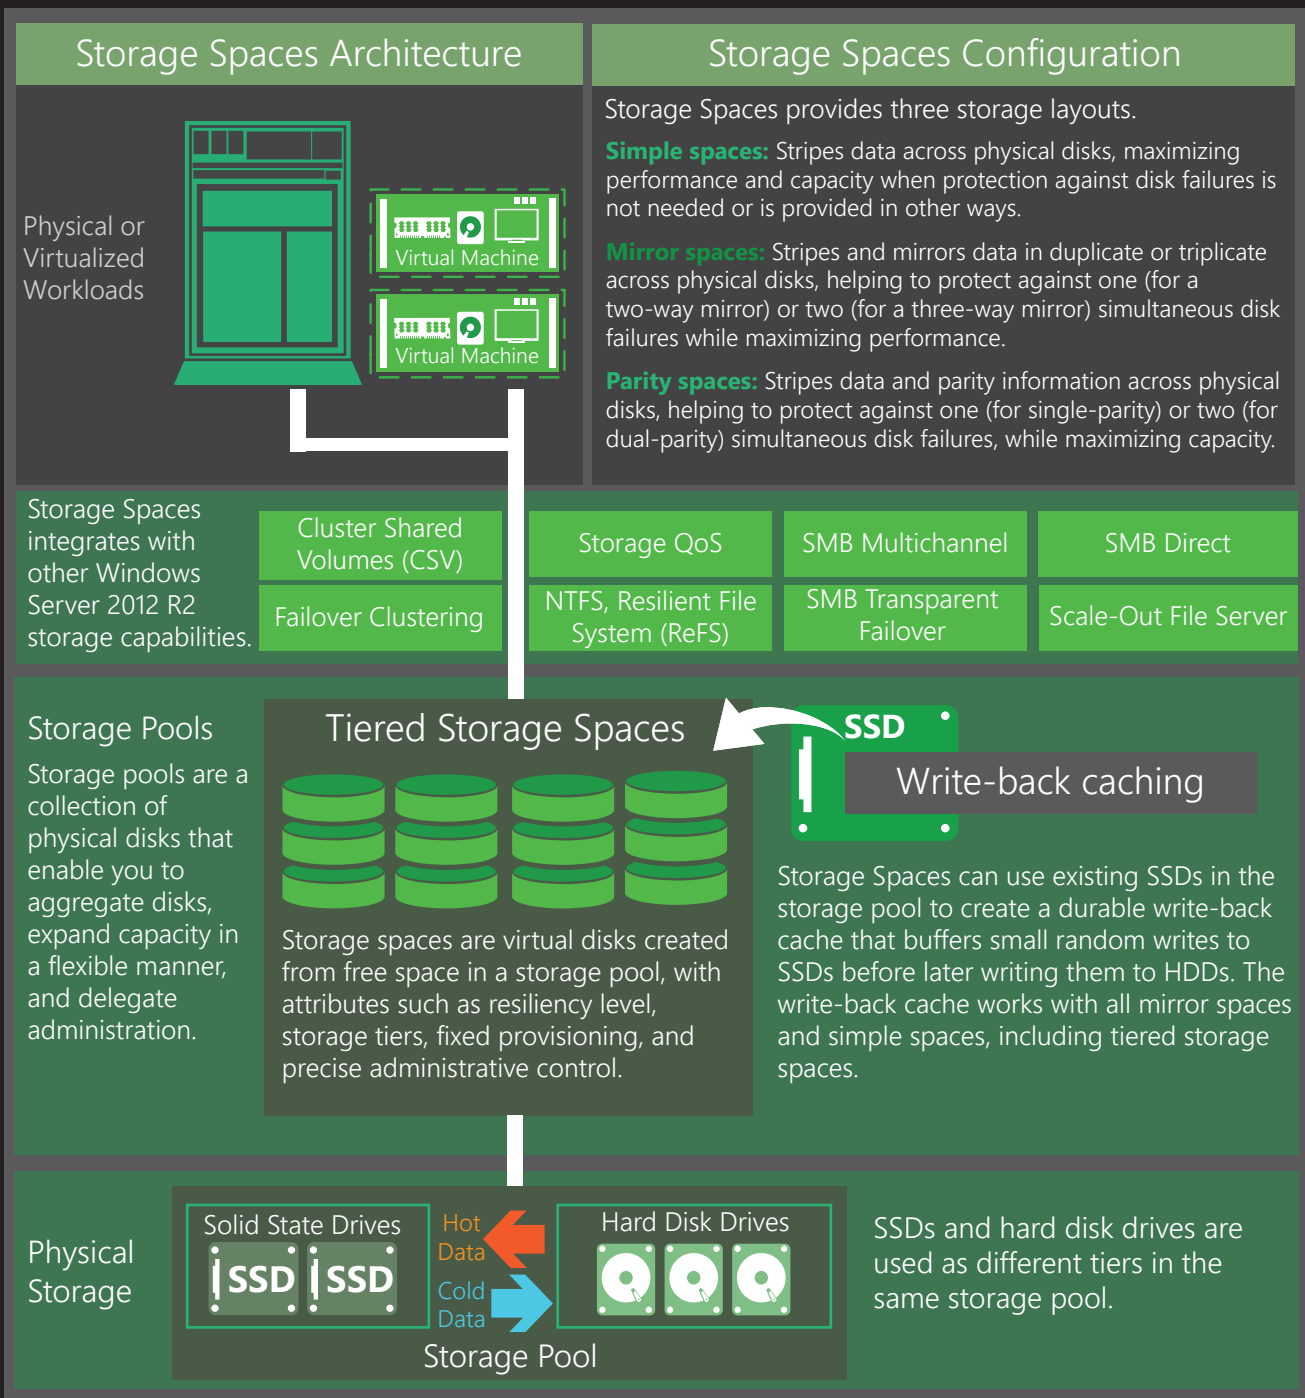# Storage Spaces

Tiered Storage Spaces

Storage Spaces is a storage virtualization technology in Windows Server 2012 R2 and Windows Server 2012 that enables you to group industry-standard disks into storage pools and then create virtual disks called storage spaces from free space in the pools.

In Windows Server 2012 R2, you can create tiered storage spaces with a mixture of solid state drives (SSD) (fast tier) and hard disk drives (standard tier). Frequently accessed (hot) data in a tiered storage space gets moved to the SSDs in the fast tier, while infrequently accessed (cold) data is moved to the high-capacity hard disks in the standard tier. This enables you to balance capacity and performance.
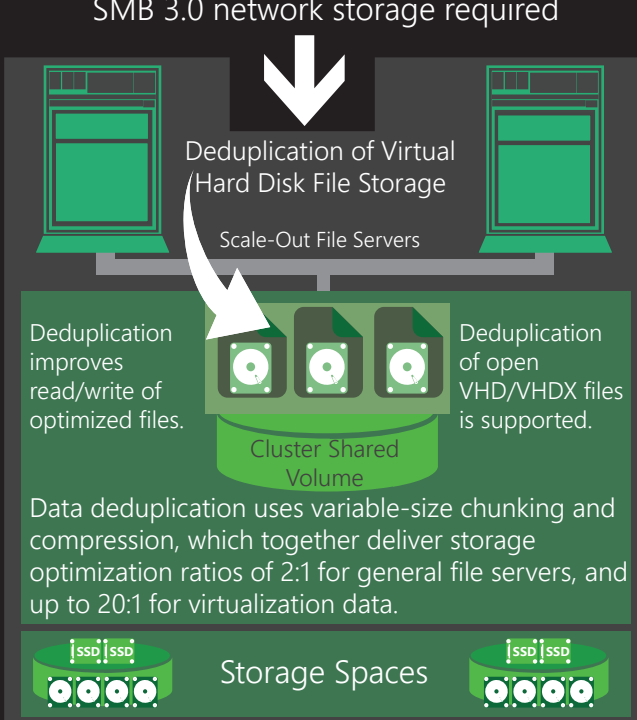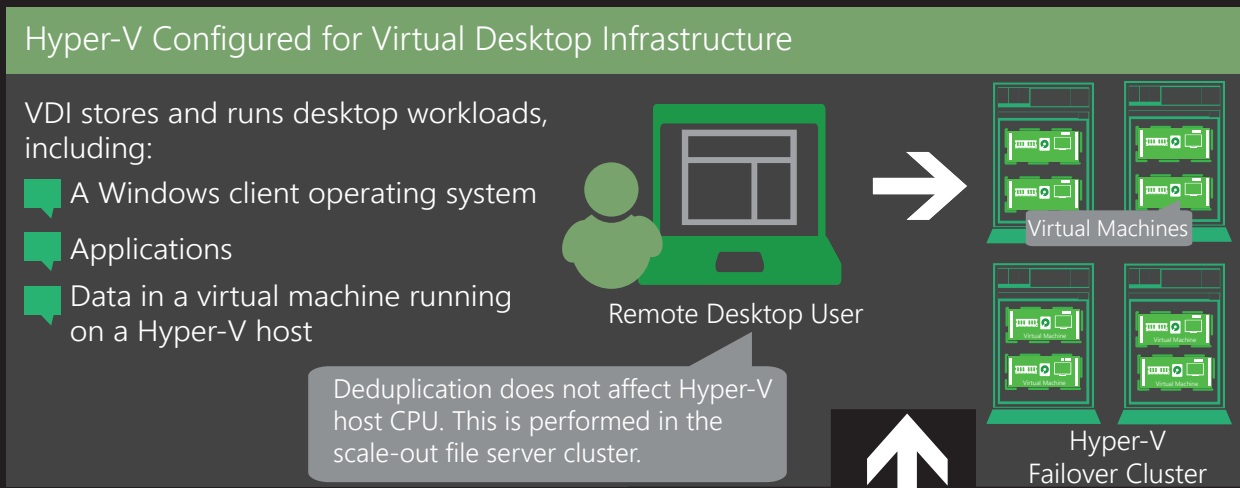
Storage Spaces Architecture

Physical or Virtualized Workloads

Virtual Machine

Virtual Machine

Storage Spaces integrates with other Windows Server 2012 R2 storage capabilities.

Cluster Shared Volumes (CSV)	Storage QoS	SMB Multichannel	SMB Direct
Failover Clustering	NTFS, Resilient File System (ReFS)	SMB Transparent Failover	Scale-Out File Server

Storage Spaces Configuration

Storage Spaces provides three storage layouts.

Simple spaces: Stripes data across physical disks, maximizing performance and capacity when protection against disk failures is not needed or is provided in other ways.

Mirror spaces: Stripes and mirrors data in duplicate or triplicate across physical disks, helping to protect against one (for a two-way mirror) or two (for a three-way mirror) simultaneous disk failures while maximizing performance.

Parity spaces: Stripes data and parity information across physical disks, helping to protect against one (for single-parity) or two (for dual-parity) simultaneous disk failures, while maximizing capacity.

Storage Pools

Storage pools are a collection of physical disks that enable you to aggregate disks, expand capacity in a flexible manner, and delegate administration.

Tiered Storage Spaces

SSD

Write-back caching

Storage spaces are virtual disks created from free space in a storage pool, with attributes such as resiliency level, storage tiers, fixed provisioning, and precise administrative control.

Storage Spaces can use existing SSDs in the storage pool to create a durable write-back cache that buffers small random writes to SSDs before later writing them to HDDs. The write-back cache works with all mirror spaces and simple spaces, including tiered storage spaces.

Physical Storage

Solid State Drives

SSD SSD

Hot Data

Cold Data

Hard Disk Drives

Storage Pool

SSDs and hard disk drives are used as different tiers in the same storage pool.

Deduplication

Deduplication for Virtual Desktop Infrastructure

You can significantly reduce your storage capacity requirements by using the data deduplication feature on your Clustered Shared Volume. In Windows Server 2012 R2, data deduplication is supported only in a Scale-Out File Server deployment for Virtual Desktop Infrastructure (VDI) workloads with separate storage and compute nodes. The desktop virtual hard disks (VHDs) are stored on a scale-out file server and accessed using Server Message Block (SMB 3.0).

Data deduplication is a storage efficiency feature in Windows Server 2012 R2 that helps address the growing demand for file storage. Data deduplication can drastically reduce how much disk space your files consume by automatically finding identical parts (chunks) of files and replacing duplicate parts with a reference to the single copy. The chunks are compressed and then organized into special container files in the System Volume Information folder.

Hyper-V Configured for Virtual Desktop Infrastructure

VDI stores and runs desktop workloads, including:

- A Windows client operating system
- Applications
- Data in a virtual machine running on a Hyper-V host

Remote Desktop User

Virtual Machines

Deduplication does not affect Hyper-V host CPU. This is performed in the scale-out file server cluster.

Hyper-V Failover Cluster

SMB 3.0 network storage required

Scale-out File Server with CSV

In Windows Server 2012 R2, data deduplication now supports open virtual hard disk files (both VHD and VHDX) and deduplication of data stored on Cluster Shared Volumes used in a scale-out file server.

You must install deduplication on a remotely connected file server. This ensures performance is maintained on your Hyper-V VDI infrastructure.

Data deduplication in Windows Server 2012 R2 has been enhanced to include performance improvements, which results in faster read/write of optimized files and improved optimization speed.

Deduplication is supported only for data volumes, and not for boot or system volumes. ReFS volumes do not support using deduplication.

Deduplication of Virtual Hard Disk File Storage

Scale-Out File Servers

Deduplication improves read/write of optimized files.

Deduplication of open VHD/VHDX files is supported.

Cluster Shared Volume

Data deduplication uses variable-size chunking and compression, which together deliver storage optimization ratios of 2:1 for general file servers, and up to 20:1 for virtualization data.

SSD SSD   Storage Spaces   SSD SSD

More information ....