

Microsoft®

Storage

Windows Server 2012

Table of contents

Copyright information	3
Introduction	4
Windows Server 2012: Vast Storage Capabilities for Everyone	4
Feature and Capability Summary.....	4
Storage.....	7
Storage Spaces.....	7
Rapid and Efficient Data Movement Using Intelligent Storage Arrays	9
File System	11
Resilient File System (ReFS)	11
Improved NTFS Availability with Online Scanning and Repair.....	12
Data Deduplication.....	14
Optimal Storage Use: Support for Thin Provisioning and Trim	16
Clustering	17
Cluster Shared Volume v2.....	17
Cluster-Aware Updating.....	18
File and Block Access.....	19
iSCSI Software Target.....	19
Diskless Boot.....	19
Block storage sharing	20
Deploying and Running VMware ESX Virtual Machines Using Server for NFS Data Stores ..	20
High-Performance, Highly-Available Storage Using SMB.....	21
Microsoft SQL Server over SMB	22
Hyper-V over SMB	23

SMB transparent failover.....	24
SMB Scale-Out	25
Information Worker file share improvements	25
SMB VSS for Remote File Shares	26
Online Backup	28
Networking	30
SMB Direct.....	30
SMB Multichannel.....	31
Hyper-V Replica	32
Virtualization.....	34
Virtual Fibre Channel in Hyper-V.....	34
Moving Virtual Machine Storage with Minimal Downtime.....	36
New Virtual Hard Disk Format	37
Hyper-V support for 4-KB Disk Sectors and 512e	38
Management	39
Comprehensive Windows Storage Management.....	39
Conclusion	41
List of charts, tables, and figures	42

Copyright information

© 2012 Microsoft Corporation. All rights reserved. This document is provided "as-is." Information and views expressed in this document, including URL and other Internet Web site references, may change without notice. You bear the risk of using it. This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes. You may modify this document for your internal, reference purposes.

Introduction

Microsoft is introducing several new storage features and capabilities with Microsoft Windows Server® 2012. These innovative features and capabilities extend functionality in profound ways, including the ability to leverage inexpensive storage to create highly available, robust, and high performing storage solutions. These new Microsoft storage capabilities add dynamic functionality on each server and can work together to further enhance functionality at scale in large enterprise environments. This paper outlines these new features and capabilities and how they integrate, co-exist, and complement one another to extend the capabilities of your entire storage infrastructure.

Windows Server 2012: Vast Storage Capabilities for Everyone

Windows Server 2012 storage capabilities are designed to provide enhanced reliability, availability, and performance at every layer of the Microsoft storage stack. They allow IT pros to combine features to build powerful new storage solutions, such as the ability to deliver feature-rich high performance storage functionality and performance without sacrificing budget. From improved compatibility with the latest storage technologies to forward-thinking innovations that make complex storage infrastructures as easy to manage as a single server, Windows Server 2012 has new storage capabilities and features for everyone.

Feature and Capability Summary

The following table provides a brief introductory look at the new capabilities available at each layer of the Windows Server 2012 storage solution stack.

Table 1: Capabilities available at each layer of the Windows Server 2012 storage solution stack

Layer	Feature	Description
Storage	Storage Spaces	Deploy cost-effective storage for business-critical data that leverages advanced Windows Server 2012 features to deliver capabilities like high availability, performance, reliability, and scalability features on inexpensive commodity storage.
	Offloaded Data Transfer (ODX)	Get the most out of investments in external storage arrays with the ability to move large chunks of data, or even entire virtual machines, between storage devices with speeds comparable to local storage, without taxing valuable server resources.

File system	Resilient File System (ReFS)	Maximize data availability, despite errors that would historically cause data loss or downtime. Data integrity helps make sure that business-critical data is protected from errors and available when needed. The architecture provides scalability and performance in an era of constantly growing dataset sizes and dynamic workloads.
	NTFS Availability	Deploy large NTFS volumes with the confidence that they are always available and resilient against downtime caused by potential file corruption.
	Data Deduplication	Save money by maximizing storage and data access efficiency with the use of sophisticated deduplication functionality that can be used on any server by itself or used together with Microsoft BranchCache to extend branch office storage capabilities.
	Thin Provisioning and Trim	“Get thin” and “stay thin” through new native support for thin provisioning and trim, which is the ability to provision storage as needed and reclaim storage that is no longer needed.
Clustering	Cluster Shared Volume version 2	Not just for Microsoft Hyper-V any longer, this popular shared file storage solution gives administrators the flexibility to house storage and applications in a consolidated cluster while taking advantage of many other new Windows Server 2012 features.
	iSCSI Software Target	Take advantage of low-cost servers providing block storage. One use case for iSCSI Target is the diskless Internet SCSI (iSCSI) boot services, which provide storage provisioning capabilities traditionally reserved for high-end storage devices.
File and Block Access	Support for VMware Virtual Machines and NFS 4.1	Reliably store and run VMware ESX virtual infrastructures and other third-party platforms with Network File System (NFS) support on Windows Server 2012, while using advanced Windows functionality for continuous availability.
	High Performance Highly Available Storage with SMB	Leverage high performance continually available file shares to deliver feature-rich storage for business-critical applications and virtual servers using advanced new functionality available with Windows Server 2012 SMB.
	SMB for Applications and Information Users	Deploy Microsoft SQL Server® data on flexible, high performance SMB file shares where administrators can analyze database storage performance and manage large sequential transfers, such as Microsoft SQL Server data warehouse, database backup, or restore operations.

	SMB Scale-Out	Deploy easily managed and scalable active/active file shares from multiple nodes of a cluster leveraging Cluster Shared Volumes (CSV) and the powerful new capabilities of SMB.
	VSS for Remote SMB File Shares	Leverage proven Volume Shadow Copy Service (VSS) to offload snapshot and backup operations to remote storage from VSS hardware providers within Hyper-V virtual machines.
Networking	SMB Direct	Implement high-performance storage capabilities that won't burden server CPU by using Remote Direct Memory Access (RDMA) hardware support for SMB. With RDMA and SMB, you can achieve Enterprise class performance without relying on expensive Fibre Channel storage area networks (SANs).
	SMB Multichannel	Enable aggregation of network bandwidth and network fault tolerance when multiple paths are available between SMB clients and file servers.
Virtualization	Virtual Fibre Channel Support	Improve Hyper-V virtual guest performance and availability by using direct connections to Fibre Channel host bus adapter (HBA) hardware from the virtual guest.
	Live Storage Migration	Maintain maximized service-level availability even while moving virtual machine storage between host servers independent of the actual virtual machine.
	New VHDX Standard	Deploy larger capacity virtual storage that can grow to 16 terabytes, is more resilient against corruption, and provides more efficient usage of new 4 KB sector hard disks with the new VHDX virtual storage standard.
Management	Comprehensive Storage Management	Bring all these features together through easy-to-use comprehensive storage management capabilities that let administrators manage remote storage on a single pane-of-glass management interface, automate storage management tasks with robust new Windows PowerShell command-line interface capabilities, and reach beyond Windows Server storage to manage storage arrays and SANs with SMI-S and new Storage Management Protocol from Microsoft.

Read on to learn more about these exciting new capabilities and see how to bring the vast storage service capabilities in Windows Server 2012 to all customer organizations, regardless of size.

Storage

Windows Server 2012 directly addresses the need to supply cost-effective storage solutions that provide the features required for business-critical storage while providing additional benefit for organizations that have already invested in external storage arrays.

Storage Spaces

The cost to acquire and manage highly available and reliable storage can represent a significant part of information technology budgets. Windows Server 2012 addresses this issue by delivering a sophisticated virtualized storage feature called Storage Spaces as part of the Windows Server 2012 storage platform. This provides an alternative option for companies that require advanced storage capabilities at a lower price point.

Storage Spaces introduces a new class of sophisticated storage virtualization enhancements to the storage stack that incorporates two new concepts:

- **Storage pools:** Virtualized units of administration that are aggregates of physical disk units. Pools enable storage aggregation, elastic capacity expansion, and delegated administration.
- **Storage spaces:** Virtual disks with associated attributes that include a desired level of resiliency, thin or fixed provisioning, automatic or controlled allocation on diverse storage media, and precise administrative control.

The Storage Spaces feature in Windows Server 2012 can leverage failover clustering for high availability and can be integrated with Cluster Shared Volumes (CSV) for scalable deployments.

Additionally, Storage Spaces includes these features:

- **Storage pooling:** Storage pools are the fundamental building blocks for Storage Spaces. IT administrators can flexibly create storage pools based on the needs of the deployment. For example, given a set of physical disks, an administrator can create one pool by using all available physical disks, or multiple pools by dividing the physical disks as required. Also, to maximize the value from storage hardware, the administrator can map a storage pool to combinations of hard disks in addition to solid-state drives (SSDs). Pools can be expanded dynamically by adding more drives, thereby seamlessly scaling to cope with increasing data growth as needed.
- **Multitenancy:** Administration of storage pools can be controlled through access control lists (ACLs) and delegated on a per-pool basis, to support hosting scenarios that require tenant isolation. Storage Spaces follows the familiar Windows security model and can be fully integrated with Active Directory Domain Services.
- **Resilient storage:** Storage Spaces supports two optional resiliency modes: mirroring and parity. Features such as per-pool hot spare support, background scrubbing, and intelligent error correction enable optimal service availability despite storage component failures.
- **Continuous availability through integration with failover clustering:** Storage Spaces is fully integrated with failover clustering to deliver continuously available service deployments. One or more pools can be clustered across multiple nodes in a single cluster. Storage Spaces can then be instantiated on individual nodes and will seamlessly migrate or fail over to a different node, either in

response to failure conditions or because of load balancing. Integration with CSV v2 enables scalable access to data on storage clusters.

- **Optimal storage use:** Server consolidation frequently results in multiple datasets sharing the same storage hardware. Storage Spaces supports thin provisioning to enable businesses to easily share storage capacity among multiple unrelated datasets, maximizing capacity use. Trim support enables capacity reclamation when possible.
- **Operational simplicity:** Fully scriptable remote management is permitted through the Windows Storage Management API, Windows Management Instrumentation (WMI), and Windows PowerShell. Storage Spaces can be managed easily through the File and Storage Services GUI in Server Manager or by using task automation with many new Windows PowerShell cmdlets.

For single-node environments, Windows Server 2012 requires:

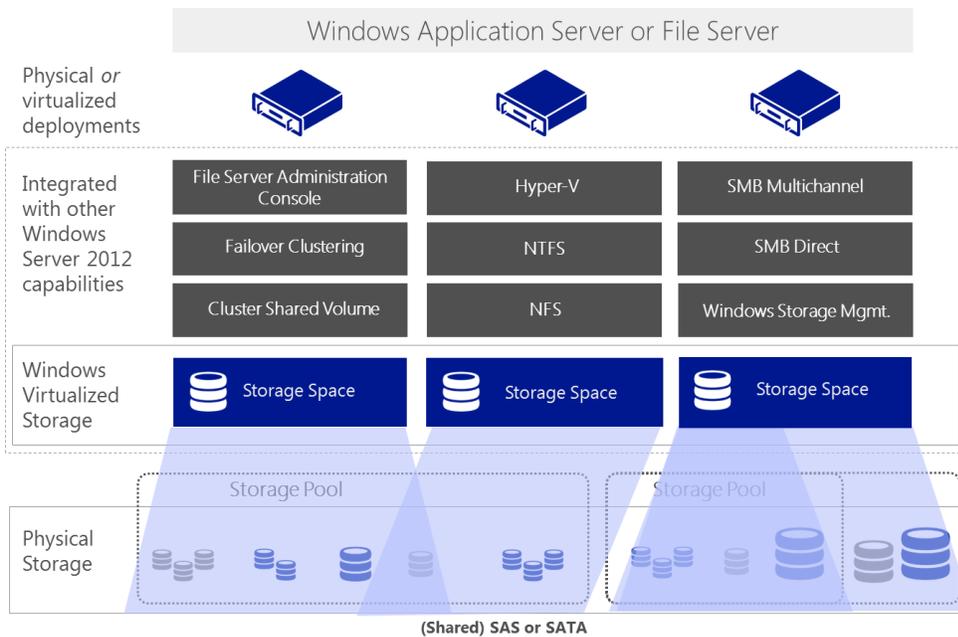
- Serial ATA (SATA) or Serial Attached SCSI (SAS) connected disks (in an optional just-a-bunch-of-disks [JBOD] enclosure).

For multiserver and multisite environments, Windows Server 2012 requires:

- Any requirements that are specified for Windows failover clustering and Windows CSV v2.
- SAS connected JBODs that help ensure compliance with Windows Certification requirements.

The following figure is one example of the Storage Spaces architecture in Windows Server 2012. It shows a subset of scenarios that are supported by Storage Spaces, to help illustrate both the capabilities of Storage Spaces and the flexible deployment options.

Figure 1: Storage Spaces conceptual deployment model



To understand the scenario depicted in the previous figure, consider:

- Physical storage is aggregated to create a storage pool. Depending on the needs of the deployment, you can create one storage pool or multiple storage pools from a set of physical disks.
- From a storage pool, you create virtual disks, which are referred to as “spaces”, when running on the Storage Spaces subsystem. You can configure these virtual disks with a variety of attributes, such as resiliency or just-in-time provisioning, depending on the needs of the deployment.
- Like a physical disk, virtual disks are partitioned and formatted with a file system, such as NTFS. Virtual disks fully integrate with other technologies that are part of Windows Server 2012, such as failover clustering and SMB Direct (RDMA).
- You can deploy Storage Spaces directly on an application server, such as Microsoft SQL Server, on a file server, or a virtual environment using Hyper-V.

The Windows Server 2012 Storage Spaces feature provides customers with a rich set of features that support a breadth of scenarios, as well as substantial reductions in costs associated with purchasing and operating sophisticated storage. By virtualizing commodity hardware, this feature offers a long-awaited solution to the problem of obtaining reliable and scalable storage at an affordable price.

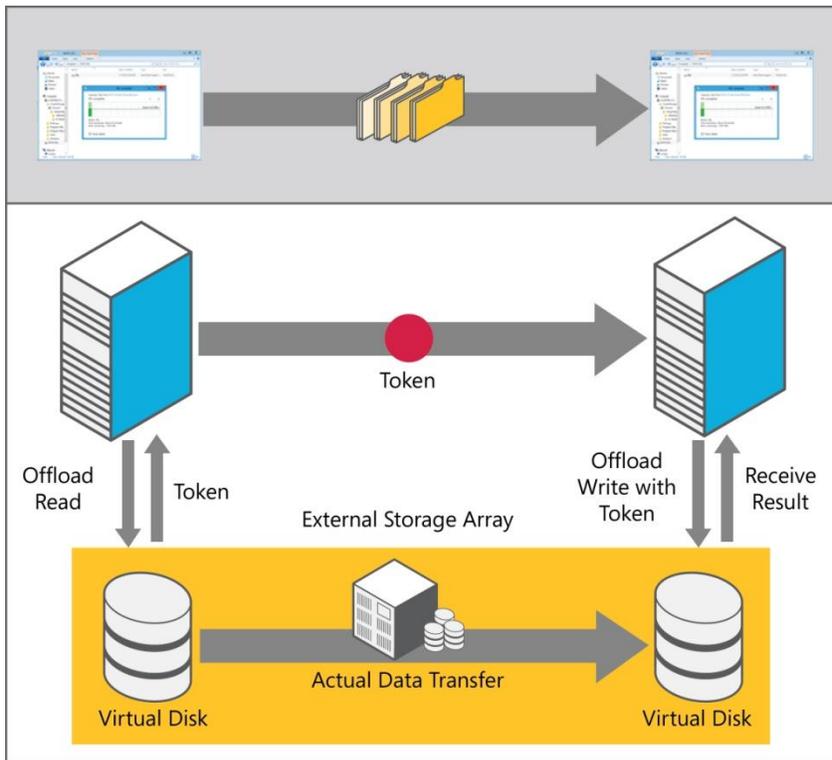
Rapid and Efficient Data Movement Using Intelligent Storage Arrays

Offloaded data transfer (ODX) in Windows Server 2012 enables you to accomplish more with existing external storage arrays by letting you quickly move large files and virtual machines directly between storage arrays, which reduces host CPU and network resource consumption.

ODX enables rapid provisioning and migration of virtual machines and provides significantly faster transfers of large files, such as database or video files. By offloading the file transfer to the storage array, ODX minimizes latencies, maximizes the use of array throughput, and reduces host resource usage, such as CPU and network consumption. File transfers are automatically and transparently offloaded when you move or copy files, regardless of whether you perform drag-and-drop operations in Windows Explorer or use command-line file copy commands. No administrator setup or intervention is needed.

To eliminate the inefficient and unnecessary steps required by traditional host-based file transfers, ODX uses a token-based mechanism for reading and writing data within or between intelligent storage arrays. Instead of routing the data through the host, a small token is copied between the source and destination. The token serves as a point-in-time representation of the data. For example, when you copy a file or migrate a virtual machine between storage locations, Windows Server 2012 copies the token representing the virtual machine file. This removes the need to copy the underlying data between servers.

Figure 2: Offloaded data transfer



In a copy operation that uses an ODX token, the steps to complete the copy operation are:

1. A user starts a file copy or move in Windows Explorer, a command-line interface, or a virtual machine migration.
2. Windows Server 2012 automatically translates this transfer request into an ODX token if it's supported by the storage array, and then receives a token representation of the data.
3. The token is copied between the source and destination systems.
4. The token is delivered to the storage array.
5. The storage array performs the copy internally and returns progress status.

ODX in Windows Server 2012 enables rapid provisioning and migration of virtual machines, lets you quickly move large files and virtual machines directly between storage arrays, and reduces host CPU and network resource consumption.

File System

Windows Server 2012 introduces four main advancements at the file system layer of the storage stack, including the introduction of a brand-new file system along with significant enhancements to the NTFS file system that improve data integrity and availability. Two other new features also help organizations save money by increasing the efficiency of how data is provisioned and stored, which means Windows Server 2012 helps improve your budget at the file system layer and improves your peace of mind as well.

Resilient File System (ReFS)

Windows Server 2012 introduces a new local file system called Resilient File System (ReFS). ReFS maximizes data availability and online operation, despite errors that would historically cause data loss or downtime. Data integrity helps protect business-critical data from errors and makes sure that the data is available when needed. ReFS architecture provides scalability and performance in an era of constantly growing dataset sizes and dynamic workloads.

ReFS was designed with three key goals in mind:

- Maintain the highest levels of system availability and reliability possible under the assumption that underlying storage may be unreliable.
- Provide a full end-to-end resilient architecture when used in conjunction with Storage Spaces so that these two features magnify the capabilities and reliability of one another when used together.
- Maintain compatibility with widely adopted and successful NTFS features while replacing features that provide limited value.

With these key goals in mind, Windows Server 2012 ReFS includes these capabilities:

- **Robust disk updating:** ReFS offers robust disk updating with an “allocate-on-write” transactional model (also known as “copy on write”). This approach maximizes reliability by avoiding problems associated with power failures during disk updates. This is accomplished by updating data with writes to different locations in an atomic manner, instead of updating data in-place.
- **Data Integrity:** All metadata is updated using allocate-on-write. Additionally, all metadata in ReFS is check-summed. This enables the detection of all forms of disk corruption, including degradation of data on media. ReFS also includes the optional ability to apply data integrity to user data. When this option, known as Integrity Streams, is enabled, ReFS always uses allocate-on-write for updates to the data and uses checksums to detect disk corruption. This option is useful in many scenarios, but is unsuitable in some cases. For example, some applications manage their file storage by relying on a specific file layout on the disk. Because Integrity Streams reallocates blocks every time file content is changed, those applications can no longer predict the file layout. This is why mechanisms and APIs are provided to enable or disable Integrity Streams settings at various levels of granularity including per-volume, per-directory, and per-file. When ReFS is used in conjunction with a mirrored Storage Space, detected corruption – both metadata and user data, when integrity streams are enabled – can be automatically repaired using the other copy provided by Storage Spaces.
- **Availability:** ReFS is designed to maximize online operation and data availability, even if corruption occurs. Although it is expected that many customers will use ReFS in conjunction with mirrored Storage Spaces to automatically repair corruption transparently. It is still possible, although very rarely, for such a volume to become corrupted, because of faulty memory passing corrupted data to

the disk for example. Additionally, some may choose not to use a mirrored storage space for their ReFS volumes. In these cases, when corruption does occur, ReFS implements “salvage”, this feature removes corrupted data from the namespace on a live volume so that good data is not adversely affected by non-repairable corrupted data. This means that volumes don’t have to be taken offline to correct errors and lets administrators restore files removed from the namespace by restoring them from backup.

- **Scalability:** As the amount and size of data that is stored on computers continues to increase, ReFS is designed to work well with extremely large data sets, petabytes and larger, without performance impact. The ReFS on-disk format is designed to support volume sizes up to 2^{78} bytes using 16 KB cluster sizes while Windows stack addressing allows 2^{64} bytes. This format also supports 2^{64} -1 byte file sizes, 2^{64} files in a directory and the same number of directories in a volume.
- **Application compatibility:** Maintains a high degree of compatibility with a widely adopted subset of NTFS features while deprecating others that provide limited value at the cost of system complexity and footprint. ReFS supports most NTFS features and Win32 APIs.
- **Proactive error identification:** ReFS integrates with a data scrubber that periodically scans the volume, trying to detect latent corruption and, when running on top of a Mirrored Storage Space, automatically repair the corrupted data.
- **Interoperability and Flexibility:** ReFS is designed to fit cleanly into the Windows storage stack with maximum flexibility and compatibility with other layers of the stack. Although ReFS is extensively tested for compatibility with supporting software, such as backup and antivirus applications, ReFS is designed to work well with features in other layers of the storage stack. One of the best examples of this is when ReFS is used in conjunction with Storage Spaces. ReFS seamlessly takes advantage of storage pools shared between multiple machines and virtual disks that can transition between them, providing additional resiliency to failures greater than what Storage Spaces or ReFS could do alone. ReFS can be installed on third-party storage subsystems as well.

Along with Storage Spaces, ReFS forms the storage foundation on Windows for the next decade and beyond with features that enable significant storage stability, flexibility, scalability, and availability.

Improved NTFS Availability with Online Scanning and Repair

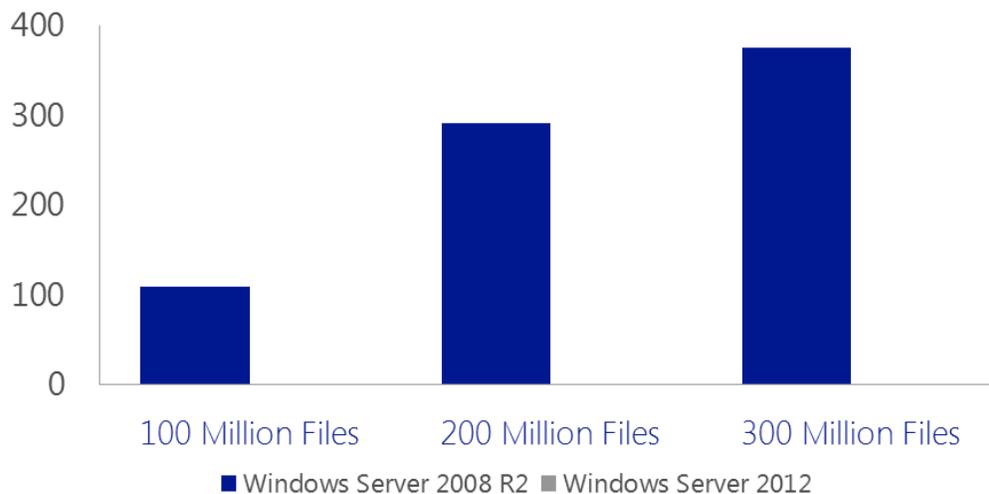
In Windows Server 2012 NTFS has been enhanced to maintain data integrity when using cost-effective industry standard SATA drives and provides online corruption scanning and repair capabilities that reduce the need to take volumes offline. When combined, all of these new capabilities let you deploy very large NTFS volumes with confidence.

There are two key enhancements that have been made to NTFS in Windows Server 2012. The first one targets the need to maintain data integrity on inexpensive commodity storage. This has been accomplished by enhancing NTFS to rely only on the “flush” command instead of “forced unit access” for all operations that require write ordering. This improves resiliency against metadata inconsistencies caused by unexpected power loss on commodity drives while enabling the disk to cache data as long as safely possible. This means that you can safely use cost-effective industry standard SATA drives.

NTFS availability is the focus of the second key enhancement and this is achieved through a combination of features that include:

- **Online Corruption Scanning:** Windows Server 2012 performs online corruption scanning operations as a background operation on NTFS volumes. This scanning operation identifies and corrects areas of data corruption if they occur, and includes logic that distinguishes between transient conditions and actual data corruption, which reduces the need for Chkdsk operations.
- **Improved Self Healing:** To further improve resiliency and availability, Windows Server 2012 has significantly increase online self-healing to resolve many issues on NTFS volumes without the need to take the volume offline to run Chkdsk.
- **Reduced Repair Times:** In the rare case of data corruption that cannot be fixed with online self-healing, administrators are notified that data corruption has occurred and can choose when to take the volume offline for a Chkdsk operation. Furthermore, because of the online corruption scanning capability, Chkdsk only scans and repairs tagged areas of data corruption. Since it doesn't have to scan the whole volume, the time needed to perform an offline repair is greatly reduced. In most cases, repairs that would have taken hours on volumes with a large number of files, now take seconds, as shown in the following chart.

Chart 1: Chkdsk - Runtime improvements (hours to seconds)



Source: "Microsoft Internal testing"

Chkdsk runs significantly faster in Windows Server 2012

Windows Server 2012 always provides the health state of the file system volume by using standard events that convey the current state of the volume. This gives administrators the flexibility to determine whether a volume has to be taken offline for a repair operation and when it is best to do this.

Windows Server 2012 improves data integrity while enhancing file system check and repair operations in ways that improve disk availability, and make the detection and repair of data corruption transparent and efficient.

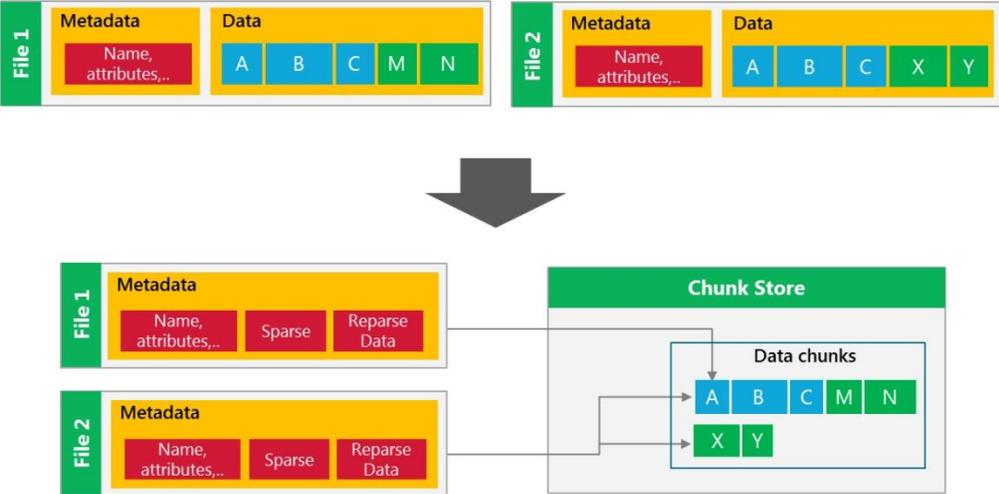
Data Deduplication

With the Data Deduplication feature of Windows Server 2012, organizations can significantly improve the efficiency of storage capacity usage. Storage efficiency is a true priority for most enterprise IT departments because storage costs have not dropped fast enough to offset the growth of data. To offset the need to increase storage, data storage efficiency has to increase wherever data is—in a data store or moving through a wide area network (WAN). In response to this growth, enterprise IT departments are consolidating file servers, and making capacity, scaling, and optimization their main goals for their storage consolidation platform.

Windows Server 2012 Data Deduplication addresses these issues and provides these features:

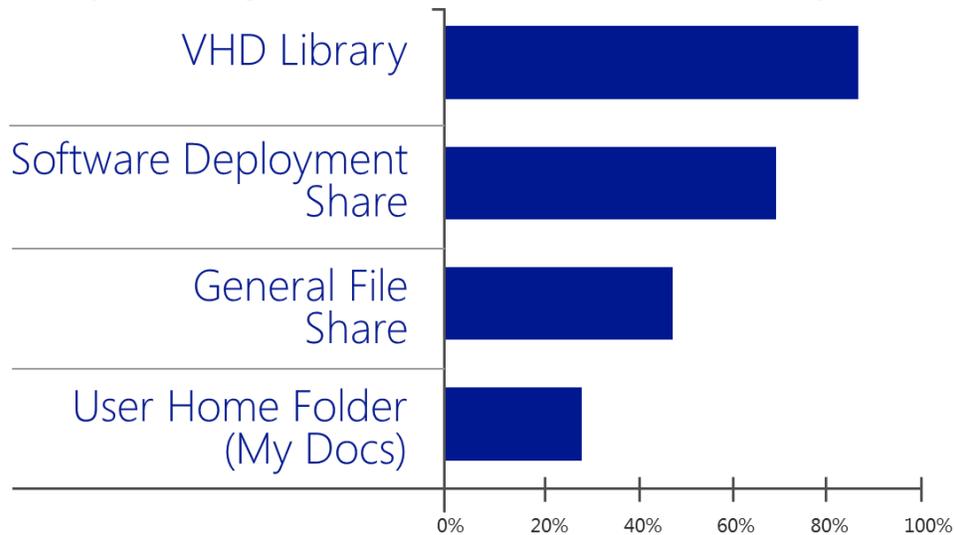
- **Capacity optimization:** Data Deduplication let you store more data in less physical space. You can achieve significantly better storage efficiency than previously possible with Single Instance Storage (SIS) or NTFS compression. Data Deduplication uses variable-size chunking and compression. Together they deliver optimization ratios of 2:1 for general file servers and up to 20:1 for virtualization VHD libraries.
- **Scalability and performance:** Data Deduplication is highly scalable, resource-efficient, and non-intrusive. It can run on dozens of large volumes of primary data simultaneously, without affecting other workloads on the server. Throttling the consumption of CPU and memory resources maintains a low impact on the server workloads. In addition, you have the flexibility to set times when Data Deduplication runs, specify the resources available to it, and establish policies on file selection for Data Deduplication.
- **Reliability and data integrity:** When you apply Data Deduplication, you must maintain data integrity. To help with data integrity, Windows Server 2012 leverages checksum, consistency, and identity validation. In addition, to recover data in the event of corruption, Windows Server 2012 maintains redundancy for all metadata and the most frequently referenced data.
- **Bandwidth efficiency alongside BranchCache;** Through integration with BranchCache, the same optimization techniques applied to improving data storage efficiency on the disk are applied to transferring data over the WAN to a branch office. The integration results in faster file download times and reduced bandwidth consumption.

Figure 3: On-disk transformation of files during Data Deduplication



Windows Server 2012 implements state of the art sub-file data chunking and compression to deliver optimal Data Deduplication. The result is an on-disk transformation of each file, as shown in the preceding diagram. Files are no longer stored as independent streams of data; they are replaced with stubs that point to data blocks stored within a common chunk store. Deduplication can be applied to a variety of workloads with typical savings of 50-90% as shown in the following chart. The implementation is transparent to the primary workload on the server – this is achieved by minimizing I/O impact through scheduled and selective optimization, and minimizing server impact through low resource (CPU/memory) usage. Further, the implementation is strongly optimized for high data integrity – this is achieved by performing data integrity validation on all data, as well as metadata, and by ensuring redundancy for all metadata as well as “popular” data chunks. Windows Server 2012 servers that enable deduplication as well as BranchCache apply the same algorithms to on-disk and on-the-wire savings, and benefit from faster download times and reduced bandwidth usage over wide area networks.

Chart 2: Typical savings when deduplication is applied to a variety of workloads



Source: “Microsoft internal testing”

With the constant growth in data and its impact on storage and infrastructure budgets, Data Deduplication in Windows Server 2012 is a technology with a clear, tangible, and measurable return on investment.

Optimal Storage Use: Support for Thin Provisioning and Trim

Like Data Deduplication, thin provisioning technology improves the efficiency of how we use and provision storage. Instead of removing redundant data on the volume, thin provisioning gains efficiencies by enabling you to allocate just enough storage at the moment of storage allocation, and then increase capacity as your business needs grow over time. Windows Server 2012 provides full support for thinly provisioned storage arrays, which lets you get the most out of your storage infrastructure. These sophisticated storage solutions offer just-in-time allocations known as "thin provisioning," and the ability to reclaim storage that is no longer needed, known as "trim."

The enhancements in Windows Server 2012 identify thinly provisioned virtual disks, provide standardized notifications when storage thresholds are reached, and let platform-enabling applications release storage when it is no longer needed for maximum use. Furthermore, as with many other new storage features in Windows Server 2012, thin provisioning enhances other capabilities when implemented together. One good example of this is that the Storage Spaces feature supports the use of thin provisioning and trim, allowing you to gain even more efficiency and capability from that already impressive storage solution.

The following enhancements and capabilities are included in Windows Server 2012:

- **Identification:** Windows Server 2012 uses a standard method to detect and identify thinly provisioned virtual disks. This enables the Windows storage stack and other storage management applications, both in the box and third-party, to treat this storage according to different needs.
- **Notification:** When certain physical storage thresholds are reached, Windows Server 2012 uses event notifications to allow you to take appropriate action as soon as possible. You can monitor these events by using sophisticated management applications, such as Microsoft System Center.
- **Optimization.** Windows Server 2012 provides a new API that enables applications to release storage when it is no longer needed. NTFS issues trim notifications when appropriate in real time. In addition, trim notifications are issued as a part of storage consolidation and optimization, which is performed regularly on a scheduled basis.

Thin provisioning and trim support in Windows Server 2012 lets you maximize the benefits of sophisticated storage infrastructure, and deliver full support and optimization of thinly provisioned storage arrays.

Clustering

Windows Server 2012 includes many new storage enhancements that improve how storage can be clustered. One of the great new storage clustering capabilities, that we will cover here, gives you the power to create active/active clustered file shares for increased performance, scalability, and availability.

Cluster Shared Volume v2

Cluster Shared Volume (CSV) was introduced in Windows Server 2008 R2 as a more efficient means for administrators to deploy storage for cluster-enabled virtual machines on Hyper-V clusters. Before CSVs, administrators had to provision a LUN on shared storage for each virtual machine so that machine had exclusive access to its virtual hard disks and to avoid mutual write conditions. With CSVs, all cluster hosts have simultaneous access to a single shared volume where storage for multiple virtual machines could be hosted, thus no more need to provision a new LUN every time you created a new virtual guest.

Windows Server 2012 introduces a number of new capabilities with CSV v2, some of these include:

- **Flexible Application and File Storage:** With CSV v2, cluster shared volumes extends its benefits beyond Hyper-V to support other application workloads and flexible file storage solutions. CSV v2 provides capabilities to clusters through shared namespaces to share configurations across all cluster nodes, including the ability to build continuously available cluster-wide file systems. Finally, application storage can be served from the same share as data, eliminating the need to deploy two clusters, an application and separate storage cluster, to support true high-availability application scenarios.
- **Integration with Other Windows Server 2012 RC Features:** Allows for inexpensive scalability, reliability, and management simplicity through tight integration with Storage Spaces. Gain high performance and resiliency capabilities with SMB Direct and SMB Multipathing, and create more efficient storage with thin provisioning.
- **Single Name Space:** CSV 2 provides a single consistent file name space where files have the same name and path when viewed from any node in the cluster. CSV volumes are exposed as directories and subdirectories under the "ClusterStorage" root directory.
- **Improved Backup and Restore:** CSV v2 supports several backup and restore capabilities including support for the full feature set of VSS and support for both Hardware and Software Backup of CSV volumes. CSVs also offer a distributed backup infrastructure for software snapshots and a new Software Snapshot Provider coordinates CSV v2snapshot creation, Point-In-Time Semantics at a cluster level, and the ability to do remote snapshots.

CSVs offer a rich new set of features that extend beyond the ease of management feature that has made it a popular choice for virtual storage today.

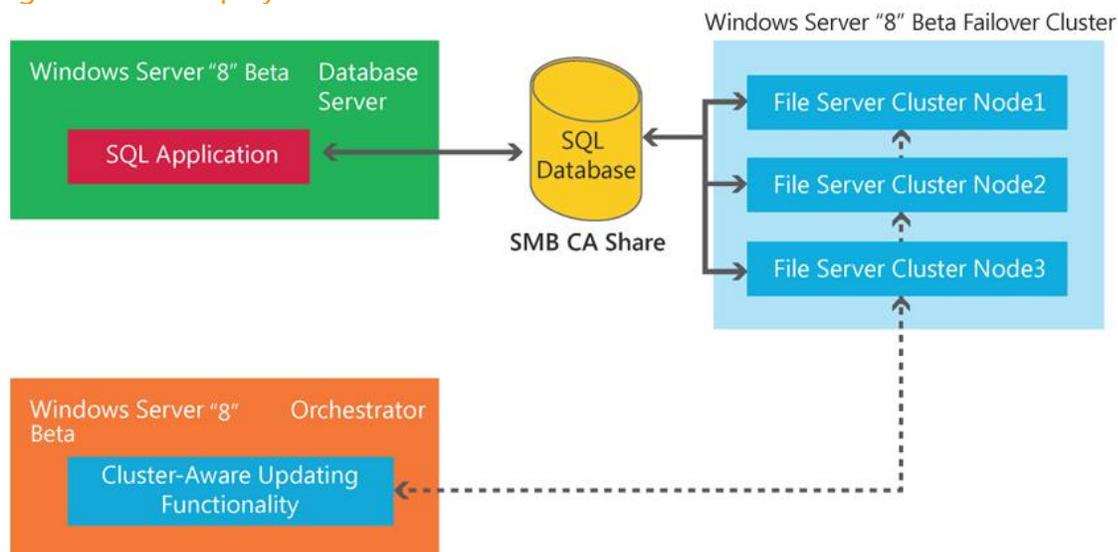
Cluster-Aware Updating

Cluster-Aware Updating (CAU) reduces server downtime and user disruption by allowing IT administrators to update clustered servers with little or no loss in availability during updates performed on cluster nodes. CAU transparently takes one node of the cluster offline, installs the updates, performs a restart if necessary, brings the node back online, and moves on to the next node. This feature is integrated into existing Windows update management infrastructure and can be further extended and automated with Windows PowerShell for integrating into larger IT automation initiatives.

CAU facilitates the cluster updating operation while running from a computer running Windows Server 2012 or Windows 8. The computer running the CAU process is called an orchestrator. CAU supports either of the two following modes of operation: remote-updating mode and self-updating mode. In remote-updating mode, a computer remote from the cluster being updated acts as an orchestrator. In self-updating mode, one of the cluster nodes being updated acts as an orchestrator, and is capable of self-updating the cluster on a user-defined schedule.

The following figure shows an example CAU deployment where a computer running SQL Server is accessing a database stored on a Continuously Available File Share on an SMB File Server failover cluster. In this example, because of the Continuous Availability features of SMB File Server and CAU orchestration, CAU can update the file server cluster in a completely transparent way to the SQL Server client.

Figure 4: CAU deployment



CAU deployment to a file server cluster with an uninterrupted SMB client

The end-to-end cluster update process by way of the CAU is cluster-aware and is completely automated. It integrates seamlessly with existing Windows Update Agent (WUA) and Microsoft Windows Server Update Services (WSUS) infrastructure. CAU also includes an extensible architecture that supports new plug-in development to orchestrate any node-updating tools, such as custom software installers, BIOS updating tools, and network adapter/HBA firmware updating tools. These tools, after integrated with CAU, can work across all cluster nodes in a cluster-aware manner.

Note: CAU is included with a full suite of Windows PowerShell cmdlets, and offers an intuitive GUI that is layered on top of the new and existing cmdlets.

File and Block Access

The largest number of feature enhancements introduced by Windows Server 2012 is at this layer of the storage stack. These include an in-the-box block-based storage solution that allows you to gain administrative flexibility, efficiency, and control by centralizing management of system boot volumes.

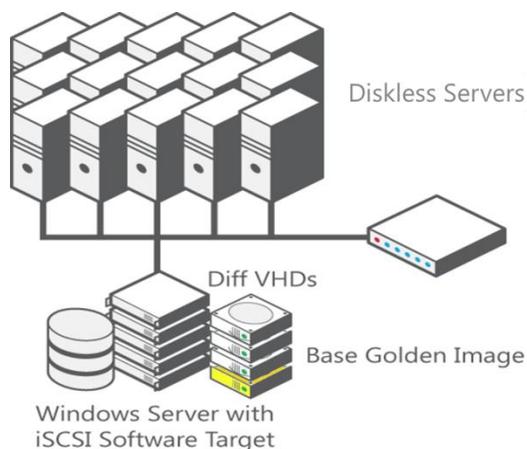
iSCSI Software Target

In Windows Server 2012, the iSCSI Software Target is available as a built-in option under the file and storage service role instead of a separate downloadable add-on to make it easier to deploy. Not only this, but iSCSI Software Target capabilities were enhanced to support diskless network boot capabilities and similar continuous availability configurations as those used by Continuous Availability SMB, again demonstrating how Windows Server 2012 storage enhancements are designed to enhance each other across all layers of the storage stack.

Diskless Boot

The following figure shows an example setup with iSCSI Target Server cluster supplying network boot services.

Figure 4: Windows Server 2012 efficiently provides highly available iSCSI block storage for diskless network boot



Network boot with iSCSI Software Target includes these benefits:

- **Cost savings on operating system storage:** By using differencing virtual disks, you can use a single operating system image to start many diskless computers. This capability offers several benefits including reduced administrative overhead and potential storage savings when boot images use differencing virtual disks.
- **Controlled operating system images that are more secure and straightforward to manage:** Some businesses require that data be secured by physically locking storage in a centralized location. In this scenario, servers running the workload access the data remotely, including operating system image data. With iSCSI Software Target boot, IT administrators can centrally manage the operating system boot images and control which applications to put in the golden image.

- **Rapid deployment:** Because the golden image is a system prepped operating system image, when the computers start off the golden image, they skip the file copying and installation phase, and go straight to the customization phase. In our testing, we deployed 256 machines in 34 minutes.
- **Fast recovery:** Because the operating system images are hosted on the iSCSI Software Target server, if the diskless client needs to be replaced, the new machine can point to the operating system image and boot up immediately.

In Windows Server 2012, the new iSCSI Software Target feature provides network boot capability for commodity hardware.

Network Boot with iSCSI Software Target does not require special hardware for functional verification. In environments with large-scale deployments, the design should be validated against specific hardware. For reference, Microsoft internal testing indicated that for a 256-iSCSI boot deployment, 24 KB by 15 KB-RPM (revolution per minute) disks in a RAID 10 configuration were required for storage. A network bandwidth of 10 gigabyte (GB) is optimal. A general estimate is 60 iSCSI boot servers per 1-GB network adapter. However, an iSCSI boot-capable network adapter is not required for this scenario. If the network adapter does not support it, a software boot loader can be used (such as iPXE open source boot firmware or netBoot/i from Double-Take Software).

Block storage sharing

iSCSI protocol enables the block storage sharing remotely. By being compliant with the standard, iSCSI Target could interoperate with any third-party implementations of iSCSI initiators.

In this release, the iSCSI Target has been continuously improving the performance and scalability. ultimately, the iSCSI Target feature built into Windows Server 2012 lets you create a SAN storage device on any hardware and gives you the power to store operating system images in a centralized location to improve efficiency, manageability, availability, and security.

Deploying and Running VMware ESX Virtual Machines Using Server for NFS Data Stores

File-based storage is a practical alternative to more SAN storage because it's straightforward to provision and has gained viability because it is simple to provision and manage. An example of this trend is the recent popularity of deploying and running VMware ESX/ESXi virtual machines from file-based storage accessed over the NFS protocol. To help you take advantage of this, Windows Server 2012 includes an updated Server for NFS that supports NFS version 4.1 and can leverage many other performance, reliability, and availability enhancements available throughout the Windows storage stack.

Some of the key features available with NFS for Windows Server 2012 include:

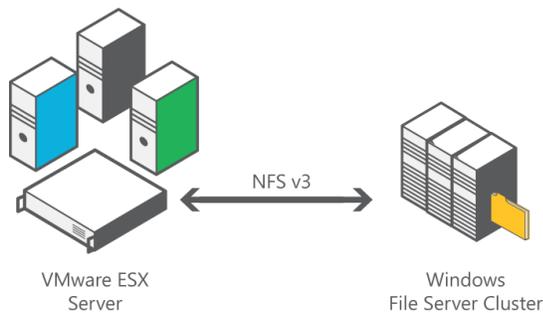
- **Storage for VMware Virtual Machines over NFS:** With Windows Server 2012, you can confidently deploy the Windows NFS server as a highly available storage back end for VMware virtual machines. We re-designed critical components of our NFS stack and focused on providing transparent failover semantics to NFS clients.
- **NFS 4.1 server:** The NFS 4.1 protocol is a significant evolution of the NFS protocol, and Microsoft delivers a standards-compliant server-side implementation in Windows Server 2012. Some of the Windows Server 2012 NFS 4.1 server features include: a flexible single-server namespace for easier share management, full Kerberos v5 support (including authentication, integrity, and privacy) for

enhanced security, VSS snapshot integration for backup, and Unmapped UNIX User Access to enable easier user account integration. Windows Server 2012 supports simultaneous Server Message Block (SMB) version 3.0 and NFS access to the same share, identity mapping using stores based on RFC-2307 for easier and more secure identity integration, and highly available cluster deployments.

- **Windows PowerShell everywhere:** In response to customer feedback, over 40 Windows PowerShell cmdlets provide task-based remote management of every aspect of the NFS server — from configuring NFS server settings to provisioning shares and share permissions.
- **Simplified identity mapping:** Windows Server 2012 includes a new flat file-based identity mapping store. Windows PowerShell cmdlets also replace cumbersome manual steps to provision Active Directory Lightweight Directory Services (AD LDS) as an identity mapping store and to manage mapped identities.

One of the best examples of how Windows Server 2012 for NFS works together with other storage capabilities involves deploying continuously available storage services for VMware ESX virtual machines.

Figure 5: Windows Server 2012 file server cluster with shared storage over NFS



For example, using VMware ESX 4.1 in conjunction with Windows Server 2012 Server for NFS support allows you to create NFS shares for VMware ESX virtual machines with the option of enabling continuous availability clustering configurations without the need for expensive storage arrays. Also, since the storage is deployed from Windows, you can quickly and easily provision an NFS share as a data store for VMware ESX servers and run virtual machines over the share through an easy-to-use and well-known storage management interface.

High-Performance, Highly-Available Storage Using SMB

The Server Message Block (SMB) Protocol is the file sharing protocol used by default on Windows computers. Windows Server 2012 provides a vast set of new SMB features, with an updated SMB protocol, that greatly enhance the reliability, availability, manageability, and performance of file servers.

The following are some of the key new features available with Windows Server 2012 SMB:

- **Active-Active file sharing with SMB Scale Out:** As already covered from the clustering layer, SMB lets you take advantage of the single namespace functionality provided by Cluster Shared Volume (CSV) version 2 to provide simultaneous access to shares from any node on a cluster. This allows for increased network bandwidth usage into a cluster and load balancing of client requests across all nodes.
- **SMB for Server Applications:** Many of the new SMB features are specifically designed for server applications that store the data on file shares—for example, database applications such as Microsoft SQL Server or virtualization software such as Hyper-V. This allows applications to take advantage of

advances in storage management, performance, reliability, and cost efficiency that come with SMB to deliver an application storage solution that rivals traditional Fibre Channel storage solutions in features and capabilities, but remains easier to provision and less expensive to implement.

- **Scalable, fast, and efficient storage access with SMB Direct:** SMB Direct (SMB over Remote Direct Memory Access (RDMA)) is a new storage protocol in Windows Server 2012. It enables direct memory-to-memory data transfers between server and storage, with minimal CPU usage, while using standard RDMA-capable network adapters. SMB Direct is supported on three types of RDMA technologies: iWARP, Infiniband, and RoCE. Minimizing the CPU overhead for storage I/O means that servers can handle larger compute workloads—for example, Hyper-V can host more virtual machines—with the saved CPU cycles.
- **Fast data transfers and network fault tolerance with SMB Multichannel:** With Windows Server 2012, customers can store application data (such as Hyper-V and Microsoft SQL Server) on remote SMB file shares. SMB Multichannel provides better throughput and multiple redundant paths from the server—for example, Hyper-V or Microsoft SQL Server—to the storage on a remote SMB share.
- **Transparent failover and node fault tolerance with SMB:** Supporting business critical server application workloads requires the connection to the storage back end to be continuously available. The new SMB server and client cooperate to provide transparent failover to an alternative cluster node for all SMB operations for planned moves and unplanned failures.

With so many new features, SMB offers a rich set of capabilities that, when combined, provide organizations with a robust high performance storage alternative to traditional Fibre Channel storage solutions at a much more affordable cost point from both an acquisition and operational perspective.

Microsoft SQL Server over SMB

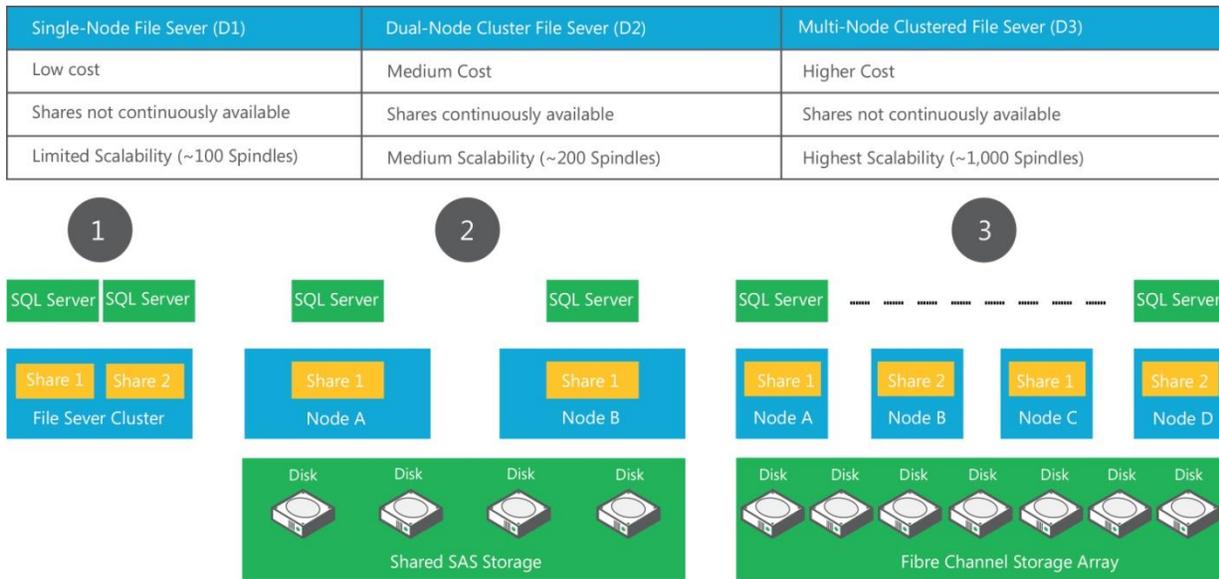
The SMB client and SMB server are both optimized for small random read/write I/O, to increase performance for common server application workloads, such as Microsoft SQL Server online transaction processing (OLTP). SMB also uses a large maximum transmission unit (MTU) feature, enabled by default, to significantly improve performance in large sequential transfers, such as those for Microsoft SQL Server data warehouse, in addition to database backup and restore operations.

Microsoft SQL Server over SMB has the following requirements:

- One or more computers running Windows Server 2012 with the File and Storage Services role installed.
- A common Active Directory infrastructure. The servers running Active Directory Domain Services do not have to run Windows Server 2012.
- Failover clustering on the Microsoft SQL Server side (with Microsoft SQL Server 2012 only), the File and Storage Services side, or both. Failover clustering is not required.

Microsoft SQL Server over SMB is flexible enough to support a variety of possible configurations, including single-node File Server, dual-node File Server, and multi-node File Server modes, as shown in the following figure.

Figure 6: Flexible configuration options for Microsoft SQL Server over SMB



Hyper-V over SMB

Prior to Windows Server 2012, remote storage options for Hyper-V were limited to expensive Fibre Channel SAN solutions that were difficult to provision for Hyper-V guests or other more inexpensive options that didn't offer many features. By enabling Hyper-V to use SMB file shares for virtual storage, administrators have a new option that is simple to provision with support for CSV v2 and inexpensive to deploy, but also offers performance capabilities and features that rival those available with Fibre Channel SANs. Windows Server 2012 also includes detailed SMB client performance counters that provide information about I/O size, I/O latency, IOPS, and so on. This lets a SQL Server database administrator or Hyper-V administrator analyze the performance of the SMB2 file shares where their data is stored.

Hyper-V over SMB requires:

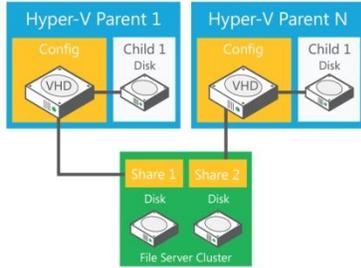
- One or more computers running Windows Server 2012 with the Hyper-V and File and Storage Services roles installed.
- A common Active Directory infrastructure. (The servers running Active Directory Domain Services do not have to run Windows Server 2012.)
- Failover clustering on the Hyper-V side, or the File and Storage Services side, or both. Failover clustering is not required.

As with SQL Server over SMB, Hyper-V over SMB supports a variety of flexible configurations that offer different levels of capabilities and availability, which include single-node File Server, dual-node File Server, and multi-node File Server modes as shown in the following figure.

Figure 7: Flexible configuration options for Hyper-V over SMB

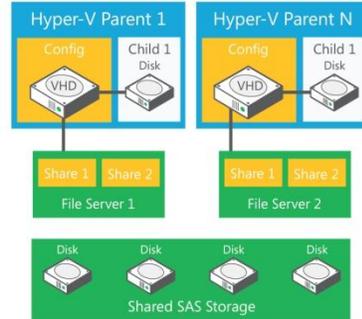
Single-Node File Server

- Lowest cost for shared storage
- Shares not continuously available



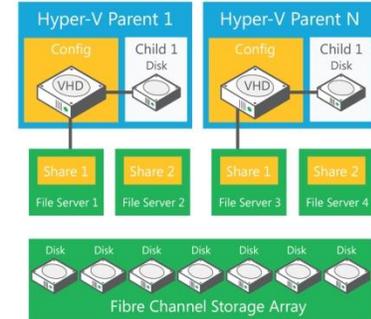
Dual-Node File Server

- Low cost for continuously available shared storage
- Limited scalability (up to a few hundred disks)



Multi-Node File Server

- Highest scalability (up to thousands of disks)
- Higher cost, but still lower than connecting all Hyper-V hosts with FC



With the ability to leverage innovative new SMB features like transparent failover, SMB Multichannel, and SMB Direct, you can now deploy storage for application workloads, such as Hyper-V and SQL Server, on cost-efficient, continuously available, high performance Windows Server 2012 File Servers.

SMB transparent failover

Windows Server 2012 reduces server downtime and application disruption by letting you store server application data on file shares and obtain a similar level of reliability, availability, manageability, and high performance that would typically be expected from a high-end Storage Area Network (SAN). New file server features include transparent failover, networking improvements for better bandwidth and resiliency, support for network adapters with RDMA capability, specific performance optimizations, and support for Windows PowerShell commands.

Server applications expect their storage to always be available and generally don't handle I/O errors or unexpected closures of file handles very well. Such events could lead to virtual machines stopping or databases going offline. Customers usually deploy hardware redundancy such as multiple network adapters, network switches, and computers in a Windows Server failover cluster. This helps mitigate hardware outages and helps keep the service running during hardware or software updates. With SMB Transparent Failover you can transparently move SMB file shares between the file server cluster nodes, without interruption of service for the SMB client. This is useful for planned events (for example, when you need to perform maintenance on a node) or surprise events (for example, when a hardware failure causes a node to fail). This is achieved regardless of the kind of operation that was underway when the failure occurred.

SMB transparent failover requires:

- **A failover cluster running Windows Server 2012 with at least two nodes:** The cluster must pass the cluster validation tests in the validation wizard.
- **File shares created with the Continuous Availability property:** This is the default setting that is set automatically when the components are installed with the cluster feature.
- **Computers running Windows 8 or Windows Server 2012:** Both computers must include the updated SMB client that supports high availability.

Note: Down-level clients can connect to file shares that have the Continuous Availability property, but transparent failover isn't supported for these clients.

SMB Scale-Out

One of the main advantages of file storage over block storage is the ease of configuration, paired with the ability to configure folders that can be shared by multiple clients. SMB takes this one step further by introducing the SMB Scale-Out feature, which provides the ability to share the same folders from multiple nodes of the same cluster. This is made possible by the use of Cluster Shared Volume (CSV), which in Windows Server 2012 supports file sharing. For example, if you have a four-node file server cluster using SMB Scale-Out, an SMB client will be able to access the share from any of the four nodes. This active-active configuration lets you balance the load across cluster nodes by allowing an administrator to move clients without any service interruption. This means that the maximum file serving capacity for a given share is no longer limited by the capacity of a single cluster node. SMB Scale-Out also helps to keep configurations simple, because a share needs to be configured only once to be consistently available from all nodes of the cluster. Also, SMB Scale-Out simplifies administration by not requiring cluster virtual IP addresses or by creating multiple cluster file server resources to take advantage of all cluster nodes.

SMB Scale-Out requires:

- **A failover cluster running Windows Server 2012 with at least two nodes:** The cluster must pass the cluster validation tests in the validation wizard.
- **File shares created on Cluster Shared Volume with the high availability property:** This is the default setting.
- **Computers running Windows 8 or Windows Server 2012:** Both computers must include the updated SMB client that supports high availability.



Note

Down-level SMB clients can connect to scale-out file shares for administrative purposes. Clients using an older version of SMB can't connect.

Information Worker file share improvements

While support for application workloads on SMB file shares is a brand new capability for the time-tested SMB protocol suite, support for housing user data is generally considered to be the traditional role of an SMB file share. However, this doesn't mean that there isn't anything new for that role for Windows Server 2012 SMB File Servers. There are a number of invaluable and forward-looking capabilities that enhance security and performance for user data on SMB file shares, including:

SMB Encryption

As more data is centralized in remote locations, mobile worker dependency on public networks will increase. Additionally, critical applications such as Hyper-V (and related scenarios such as storage migration, Virtual Desktop Infrastructure (VDI), cross-cluster migration, and so on) will start using SMB over potentially insecure networks. Moreover, regulatory compliance factors addressing files that contain high-business impact data are becoming increasingly important. Adding encryption as an inherent feature of the protocol eases the burden on administrators to implement an add-on security mechanism and make it easy to setup. SMB Encryption protects data in flight from unintentional snooping threats on untrusted network with no additional setup requirements.

Directory Leasing

The SMB Directory Leasing feature allows clients to cache directory metadata for extended periods of time. This helps a client in branch offices with high latency networks reduce roundtrips to the server. As a result, the traffic on WAN networks is reduced, and the end-user sees much improved latency when

accessing remote data. SMB Directory Leasing is optimized for Home Folder workloads, which are typically read/write scenarios with no sharing requirements and for Publication scenarios that have read only access patterns, but can be shared.

Primary Computers

Technologies such as Roaming User Profile and Folder Redirection have been successfully used to manage data for roaming users. They provide the benefits of centralized backup, consistent views across clients, and easy PC-reseating. The Primary Computers feature furthers the ability for administrators to ensure that user data gets roamed only on machines that are primarily used by those users. The advantage for the end-user is faster logon times on non-primary computers (such as shared conference machines). For administrators, it provides the benefit of reduced hassles associated with regular data cleanup, and increased security by ensuring that data is roamed only where it is needed.

SMB VSS for Remote File Shares

In Windows Server 2012, Hyper-V introduces support for SMB remote file storage. With this, Hyper-V can now store virtual machine files, including configuration files, and virtual hard disk files in shared folders that use the SMB protocol. Similarly, Microsoft SQL Server 2008 R2 introduced support for storing database and log files in shared folders that use the SMB 3.0 protocol.

To support these scenarios, and other applications that support storing data files in SMB shared folders, a new feature called VSS for SMB File Shares is being introduced. It enables backup applications that support VSS to perform application-consistent backup and restore of applications and store data on SMB File Shares.

VSS for SMB consists of four parts:

- Enhancements to the VSS infrastructure to support the new File Share Shadow Copy provider.
- New VSS provider "File Share Shadow Copy Provider." The File Share Shadow Copy Provider manages shadow copies on all SMB Universal Naming Convention (UNC) paths and relays the shadow copy request to the File Share Shadow Copy Agents. The File Share Shadow Copy Provider is being invoked on the computer running the application that supports VSS, which is storing data files on a remote UNC path.
- There is a new "File Share Shadow Copy Agent." The File Share Shadow Copy Agent manages a share to volume mappings and interacts with the VSS infrastructure on the file server to perform shadow copies of the volumes backing the SMB file shares where the applications that support VSS are storing their data files. The File Share Shadow Copy Agent is being invoked on the computer hosting the SMB file shares.
- The new File Share Shadow Copy Provider and the new File Share Shadow Copy Agent use a new remote procedure call (RPC)-based protocol to coordinate shadow copy requests of data stored on SMB file shares.

"VSS for SMB File Shares" requires:

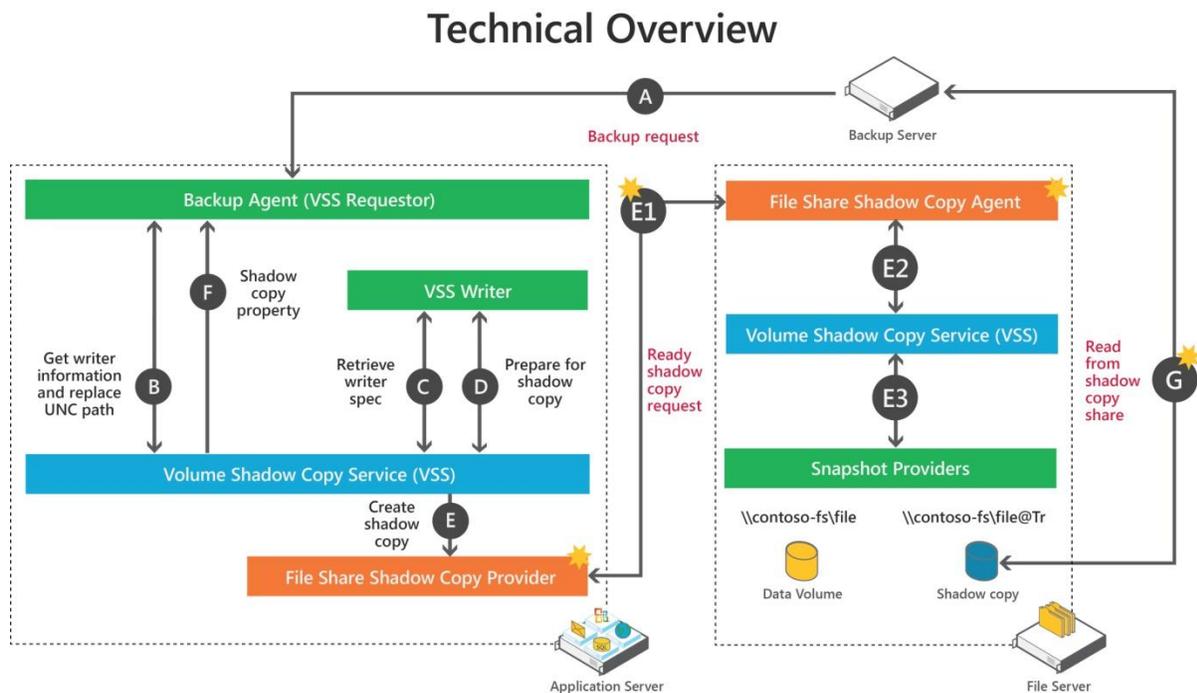
- Both application servers and file servers must be running Windows Server 2012.
- The "File Server VSS Agent Service" role service must be enabled on the file servers.
- Application servers and file servers must be joined to an Active Directory domain in the same Active Directory forest.

- The backup agent must run in a security context that has backup operators or administrators privileges on the file servers.

“VSS for SMB File Shares” supports:

- File servers in standalone, failover cluster, or Scale-Out cluster configurations.
- Application servers in standalone or failover cluster configuration.
- Normal, Continuously Available, and Scale-Out SMB file shares.
- SMB file shares as DFS-N link targets.
- Autorecovery.

Figure 8: VSS for SMB for Windows Server 2012



The following steps describe the high-level Shadow Copy sequence. For additional details on processing a backup under VSS, see [http://msdn.microsoft.com/en-us/library/aa384589\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/aa384589(VS.85).aspx).

Basic shadow copy sequence:

1. The Backup Server sends a backup request to its Backup Agent (VSS Requestor).
2. The VSS Requestor gathers writer information and resolves UNC path (SMB protocol only).
3. The VSS Service retrieves the writer metadata information and returns it to the VSS requestor.
4. The VSS Service sends Prepare Shadow Copy request to the VSS writers involved and the VSS writers flushes buffers and holds writes.

5. The VSS Service sends the Shadow Copy creation request to the File Share Shadow Copy Provider for any UNC paths involved in the Shadow Copy Set:
 - o The File Share Shadow Copy Provider relays the Shadow Copy creation request to the File Share Shadow Copy Agent on each remote file server involved in the Shadow Copy Set.
 - o The File Share Shadow Copy Agent initiates writer-less Shadow Copy creation request to the VSS Service on the file server.
 - o The VSS Service on the file server completes Shadow Copy request using the appropriate VSS hardware or system providers.
6. When the Shadow Copy creation sequence completes on the Application Server, the VSS requestor on the Application Server can retrieve the Shadow Copy properties from the VSS Service.

Based on the Shadow Copy device name from the Shadow Copy properties, the Backup Server can access the data on the Shadow Copy shares for backup.

If the shadow copy sequence fails at any point, it is canceled and the backup application will need to retry.

Online Backup

Microsoft Online Backup Service reduces server downtime and user disruption by offering IT administrators the option for backing up and protecting important data in an easily recoverable way from any location. Microsoft Online Backup Service delivers business continuity benefits to Windows Server customers by providing a backup solution that requires no up-front hardware cost other than a broadband Internet connection.

Shifting backup and restore operations into the cloud can be beneficial for smaller organizations, organizations with many branch offices, or departments within organizations that are responsible for their own IT operations. Online Backup options may also be enticing to organizations that don't have the budget for co-location or secondary sites but are interested in offsite data protection capabilities as part of their failure recovery plan.

Overall, Online Backup should be of interest to any organization that needs to reduce the costs associated with deploying and maintaining a separate backup/restore infrastructure, purchase and store backup media, and manage complicated backup maintenance tasks.

Secure and optimized

Online Backup performs block-level differential backups of files and folders to efficiently use available bandwidth and storage. This means that after a one-time full backup, only changed blocks are sent across to the online service.

Online Backup automatically compresses and encrypts data so backup data is optimized and secure as it is sent over the network and stored in the cloud. In the cloud, multiple points in times of the backup are optimally stored and automatically recycled based on a simple policy specified in Windows Server 2012. To maintain the security of the data, decryption of the backed up data is only done on the target server running Windows Server during the recovery process, and the encryption key (only known by the customer) is required to enable this.

Simple configuration and management

Microsoft Online Backup Service integrates with the familiar Windows Server backup utility to provide a seamless protection and recovery experience to a local disk, or to the cloud.

- Simple UI to configure and monitor the backups.
- Online service reduces backup management overheads when compared to disk backups. For example, additional storage seamlessly added for backup as data size grows.
- Integrated recovery experience to transparently recover data from local disk as well as Online Backup.
- Easily recover any data that was backed up to any server of your choice.
- Windows PowerShell command-line interface scripting capability.
- Microsoft System Center Operations Manager agent for centralized monitoring across multiple servers.
- Extensible to enable third-party management UI and corresponding service to be integrated into Windows Server Backup UI.

Networking

The vast array of improvements available with SMB don't stop at the file access layer, they extend well into the network layer with impressive new features that give you the power to create high-performance storage networks and let you leverage capabilities previously only available to expensive Fibre Channel storage devices.

SMB Direct

The SMB protocol in Windows Server 2012 includes support for Remote Direct Memory Access (RDMA) network adapters, which allows storage performance capabilities that rival Fibre Channel. RDMA network adapters enable this performance capability by operating at full speed with very low latency due to the ability to bypass the kernel and perform write and read operations directly to and from memory. This capability is possible since reliable transport protocols are implemented on the adapter hardware and allow for zero-copy networking with kernel bypass.

With this capability, applications, including SMB, can perform data transfers directly from memory, through the adapter, to the network, and then to the memory of the application requesting data from the file share. This means two kernel calls, one from the server and one from the client, are largely removed from the data transfer process, resulting in a greatly improved data transfer performance. This capability is especially useful for read/write intensive workloads such as Hyper-V or Microsoft SQL Server and results in remote file server performance comparable to local storage.

SMB Direct requires:

- **At least two computers running Windows Server 2012 RC:** No additional features have to be installed—the technology is available by default.
- **Network adapters with RDMA capability:** SMB Direct supports common RDMA capable adapter types including iWARP, Infiniband, and RDMA over Converged Ethernet (RoCE).

Note:

SMB Direct works in conjunction with SMB Multichannel to transparently provide a combination of exceptional performance and failover resiliency when multiple RDMA links between clients and SMB file servers are detected. Also, since RDMA bypasses the kernel stack, RDMA doesn't work with NIC Teaming, but does work with SMB Multichannel since SMB Multichannel is enabled at the application layer.

SMB Multichannel

When the SMB protocol was introduced, a single session connection path to file servers was more than sufficient for most file sharing needs since the types of usage workloads were limited. With the changes in how users and applications use SMB file shares, especially with the introduction of SMB application storage support, SMB file servers must contend with a wide variety of workloads simultaneously and with a degree of resilience that wasn't critical when SMB was initially developed.

To address this change in demands on file servers, Windows Server 2012 SMB includes a capability called SMB Multichannel that provides scalable and resilient connections to SMB shares that dynamically create multiple connections for single sessions or multiple sessions on single connections depending on connection capabilities and current demand. This capability to create flexible session to connection associations gives SMB a number of key features:

- **Connection Resiliency:** With the ability to dynamically associate multiple connections with a single session, SMB gains resiliency against connection failures usually caused by network interfaces or components. SMB Multichannel also allow clients to actively manage paths of similar network capability in a failover configuration that automatically switches sessions to the available paths if one path becomes unresponsive.
- **Network Usage:** SMB can leverage receive-side scaling (RSS)-capable network interfaces along with the multiple connection capability of SMB Multichannels to fully use high-bandwidth connections, such as those available on 10 GbE networks, during read and write operations with workloads that are evenly distributed across multiple CPUs.
- **Load Balancing:** Clients can adapt to changing network conditions to dynamically rebalance loads to a connection or across a set of connections that are more responsive when congestion or other performance issues occur.
- **Transport Flexibility:** Since SMB Multichannel also supports single session to multiple connection capabilities, SMB clients are flexible enough to dynamically adjust when new network interfaces become active. This is how SMB Multichannel is automatically enabled whenever multiple UNC paths are detected and can dynamically grow to use multiple paths as more are added, without administrator intervention.

SMB Multichannel gives SMB servers the capability to maximize throughput on a large variety of network connections and efficiently leverage multiple paths to supply the failover resiliency and load balanced performance capabilities required by demanding applications like Microsoft SQL Server and Hyper-V.

SMB Multichannel has these requirements, organized by how SMB Multichannel prioritizes connections when multiple connection types are available:

- **RDMA-capable network connections:** SMB Multichannel can be used with a single Infiniband connection on the client and server side or a dual Infiniband connection on each server connected to different subnets. While SMB Multichannel will offer scaling performance enhancements in single adapter scenarios through RDMA and RSS if available, it can't supply failover and load balancing capabilities without multiple paths. RDMA-capable adapters include iWARP, Infiniband, and RoCE.
- **RSS-capable network connections:** SMB Multichannel can leverage RSS-capable connections in 1-1 connection scenarios or multi-connection scenarios. As above, multichannel load balancing and failover capabilities aren't available unless multiple paths exist, but it does take advantage of RSS to provide scaling performance usage by spreading overhead between multiple processors by using RSS-capable hardware.

- **Load Balancing and Failover (LBFO) or aggregate interfaces:** When RDMA or RSS connections are not available, SMB prioritizes connections that use a collection of two or more physical interfaces. This requires more than one network interface on the client and server where both are configured as a network adapter team. In this scenario, load balancing and failover are the responsibility of the teaming protocol, not SMB Multichannel, when only one NIC Teaming connection is present and no other connection path is available.
- **Standard interfaces and Hyper-V virtual networks:** These connection types can use SMB Multichannel capabilities but only when multiple paths exist. For all practical intent, 1 GB Ethernet connections are the lowest priority connection type capable of using SMB Multichannel.
- **Wireless network interfaces:** Wireless interfaces aren't capable of multichannel operations.

When connections aren't similar between client and server, SMB Multichannel will take advantage of available connections when multiple connection paths exist. For example, if the SMB file server has a 10 GbE connection but the client server only has 4 1-GbE connections and all connections form paths to the file server, SMB Multichannel can create connections on each 1 GB interface to achieve better performance and resiliency even though the server's network capabilities exceed the client's.

Hyper-V Replica

You can use failover clustering to make virtual machines (virtual machines) highly available, but it doesn't protect businesses from outage of an entire datacenter without the use of hardware-based SAN replication across datacenters. Hyper-V Replica fills an important need by providing an affordable in-box failure recovery solution from an entire site, down to a single virtual machine. It provides asynchronous, unlimited replication of your virtual machines over a network link from one Hyper-V host at a primary site to another Hyper-V host at a Replica site, without reliance on storage arrays or other software replication technologies.

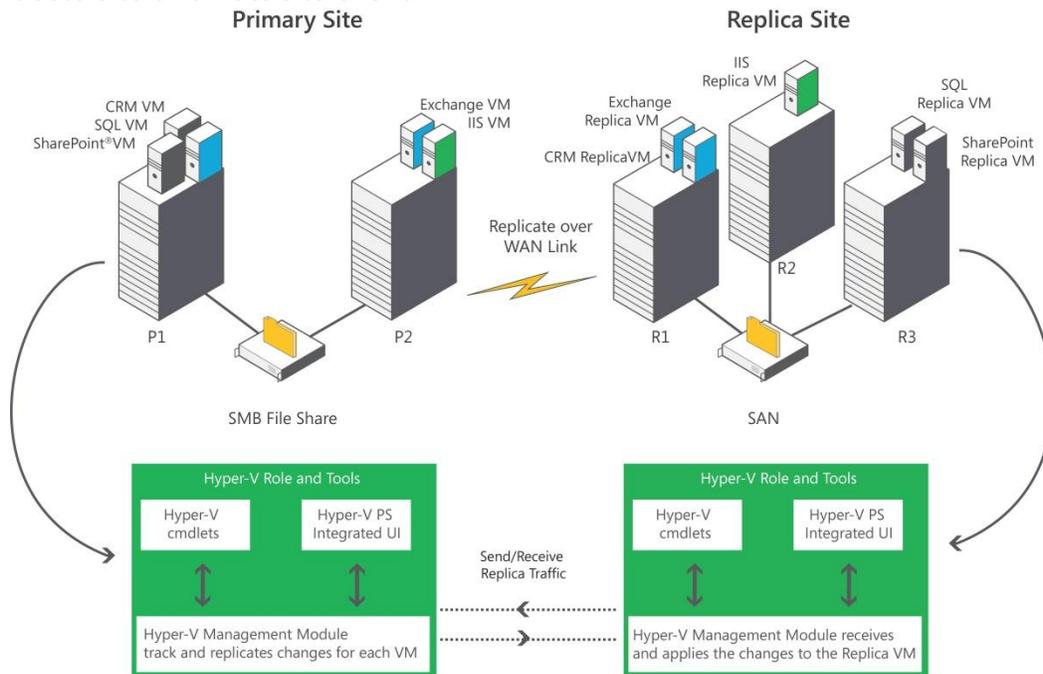
Virtualization

Virtualization opens new, effective possibilities to make sure that your network services are available at the most optimal performance levels. Hyper-V Replica provides a champion storage and workload solution that replicates efficiently, periodically, and asynchronously over IP-based networks, typically to a remote site. It also lets an IT administrator easily test the replica virtual machine without disrupting the ongoing replication. If a failure occurs at the primary site, IT administrators can quickly restore their business operations by displaying the replicated virtual machine at the Replica site. Hyper-V Replica provides a virtual machine-level, affordable, reliable, and manageable replication solution that is tightly integrated with Hyper-V Manager and the failover clustering feature in Windows Server 2012.

Replication of virtual machines

The following figure illustrates how IT administrators can use Hyper-V Replica to replicate virtual machines over a network link in a storage-independent manner.

Figure 9: More securely replicating virtual machines from a wide range of systems and clusters to a remote site over a WAN



For example, if a failure occurs at your company's primary site, IT administrators can manually failover the production virtual machines to the server running Hyper-V at the recovery site. During failover, the virtual machines would be brought online, and within minutes the rest of the network would access the virtual machines, with minimal impact on the business and the users. When the primary site or virtual machines are restored, IT administrators can easily revert the virtual machines back to the server running Hyper-V at the primary site.

Secure replication over the network

Hyper-V Replica tracks the write operations on the primary virtual machine and replicates these changes to the Replica server efficiently over a WAN. The network connection between the two servers uses the HTTP or HTTPS protocol and supports both integrated and certificate-based authentication. Connections configured to use integrated authentication are not encrypted. For an encrypted connection, use certificate-based authentication. Hyper-V Replica is closely integrated with Windows failover clustering and provides seamless replication across different migration scenarios in the primary and Replica servers.

Simple configuration and management

Hyper-V Replica includes these tools to simplify management:

- An integrated UI with Hyper-V Manager and the Failover Cluster Manager snap-in for the Microsoft Management Console (MMC).
- An extensible WMI interface.
- A Windows PowerShell command-line interface scripting capability.

Virtualization

Windows Server 2012 is being described as a “cloud-enabled operating system” so it should be obvious to find a number of new storage capabilities at this layer of the storage stack. From new support for unhindered access to Fibre Channel hardware from the virtual machine to new capabilities in virtual storage that better leverage current and next generation storage devices. It also gives you flexibility to freely move virtual storage transparently and without touching the virtual guest attached to that storage. There are many other new features for virtualization beyond the virtualization layer, including, for example, new support for deploying virtual storage on SMB file shares. To read more about virtualization improvements beyond storage, please go to the Windows Server 2012 Hyper-V Capability white paper at <http://go.microsoft.com/?linkid=9799217>.

Virtual Fibre Channel in Hyper-V

Windows Server 2012 provides Fibre Channel ports within the Hyper-V guest operating system, this lets you connect to Fibre Channel directly from virtual machines when virtualized workloads need to connect to existing storage arrays. This protects your investments in Fibre Channel, lets you virtualize workloads that use direct access to Fibre Channel storage, lets you cluster guest operating systems over Fibre Channel, and offers an important new storage option for servers hosted in your virtualization infrastructure.

Fibre Channel in Hyper-V requires:

- One or more installations of Windows Server 2012 with the Hyper-V role installed. Hyper-V requires a computer with processor support for hardware virtualization.
- A computer with one or more Fibre Channel host bus adapters (HBAs), each with an updated HBA driver that supports Virtual Fibre Channel. Updated HBA drivers are included with the in-box HBA drivers for some models, listed in the following table.

Table 2: Models that include updated HBA drivers with in-box HBA drivers

Vendor	Model
Brocade	BR415 / BR815
Brocade	BR425 / BR825
Brocade	BR804
Brocade	BR1860-1p / BR1860-2p
QLogic	Qxx25xx Fibre Channel HBAs

- Virtual machines configured to use a virtual Fibre Channel adapter, which must use Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 as the guest operating system.
- Connection only to data logical unit numbers (LUNs). Storage accessed through a virtual Fibre Channel connected to a LUN can't be used as boot media.

Virtual Fibre Channel for Hyper-V provides the guest operating system with unmediated access to a SAN by using a standard World Wide Name (WWN) associated with a virtual machine. Hyper-V users can now use Fibre Channel SANs to virtualize workloads that require direct access to SAN logical unit numbers (LUNs). Virtual Fibre Channel also allows you to operate in new scenarios, such as running the Windows failover clustering feature inside the guest operating system of a virtual machine connected to shared Fibre Channel storage.

Mid-range and high-end storage arrays are capable of advanced storage functionality that helps offload certain management tasks from the hosts to the SANs. Virtual Fibre Channel presents an alternate hardware-based I/O path to the Windows software virtual hard disk stack. This allows you to use the advanced functionality offered by your SANs directly from Hyper-V virtual machines. For example, you can use Hyper-V to offload storage functionality (like taking a snapshot of a LUN) on the SAN hardware by using a hardware Volume Shadow Copy Service (VSS) provider from within a Hyper-V virtual machine.

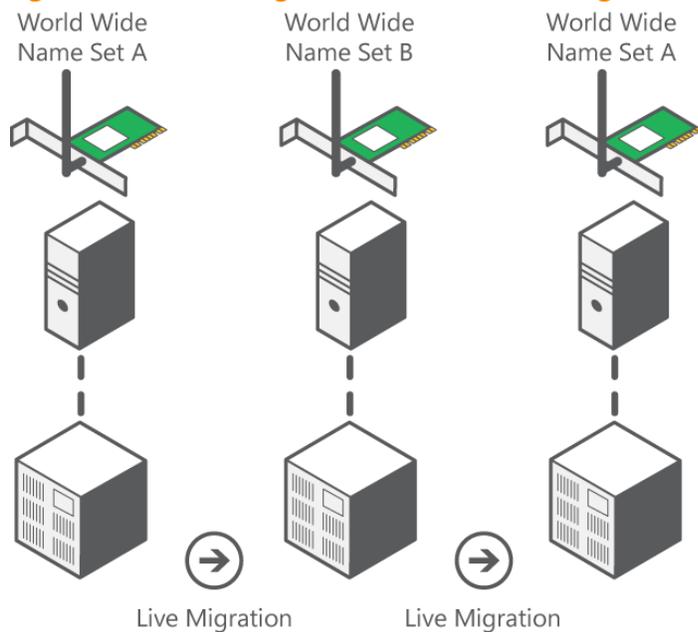
N_Port ID Virtualization (NPIV) Support: Virtual Fibre Channel for Hyper-V guests use the existing (NPIV) T11 standard to map multiple virtual N_Port IDs to a single physical Fibre Channel N_Port. A new NPIV port is created on the host each time you start a virtual machine that is configured with a virtual HBA. When the virtual machine stops running on the host, the NPIV port is removed.

Virtual SAN Support: Hyper-V allows you to define virtual SANs on the host to accommodate scenarios where a single Hyper-V host is connected to different SANs by using multiple Fibre Channel ports. A virtual SAN defines a named group of physical Fibre Channel ports that are connected to the same physical SAN. For example, assume a Hyper-V host is connected to two SANs—a production SAN and a test SAN. The host is connected to each SAN through two physical Fibre Channel ports. In this scenario, you might configure two virtual SANs—one named “Production SAN” that has the two physical Fibre Channel ports connected to the production SAN and one named “Test SAN” that has two physical Fibre Channel ports connected to the test SAN. You can use the same technique to name two separate paths to a single storage target.

You can configure as many as four virtual Fibre Channel adapters on a virtual machine and associate each one with a virtual SAN. Each virtual Fibre Channel adapter connects with one WWN address or two WWN addresses to support live migration. You can set each WWN address automatically or manually.

Live migration Support: To support live migration of virtual machines across Hyper-V hosts while maintaining Fibre Channel connectivity, two WWNs are configured for each virtual Fibre Channel adapter, as shown in the following figure: Set A and Set B. Hyper-V automatically alternates between the Set A and Set B WWN addresses during a live migration. This ensures that all logical unit numbers (LUNs) are available on the destination host before the migration and that no downtime occurs during the migration.

Figure 10: Alternating WWN addresses during a live migration



MPIO Support: Hyper-V in Windows Server 2012 can use the multipath I/O (MPIO) functionality to ensure continuous connectivity to Fibre Channel storage from within a virtual machine. You can use MPIO functionality with Fibre Channel to:

- Virtualize workloads that use MPIO. Install multiple Fibre Channel ports on the host, and use MPIO to provide highly available connectivity to the LUNs accessible by the host.
- Configure multiple virtual Fibre Channel adapters inside a virtual machine, and use a separate copy of MPIO within the guest operating system of the virtual machine to connect to the LUNs that the virtual machine can access. This configuration can coexist with a host MPIO setup.
- Use different device-specific modules (DSMs) for the host or each virtual machine. This approach enables live migration of the virtual machine configuration, including the configuration of DSM and connectivity between hosts and compatibility with existing server configurations and DSMs.

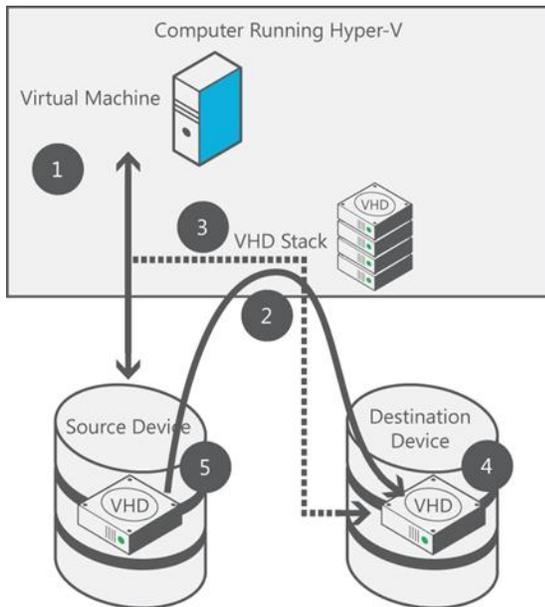
Moving Virtual Machine Storage with Minimal Downtime

Windows Server 2012 introduces live storage migration, which lets you move virtual hard disks that are attached to a running virtual machine. When you have the flexibility to manage storage without affecting the availability of your virtual machine workloads, you can perform maintenance on storage subsystems, upgrade storage appliance firmware and software, and balance loads while the virtual machine is in use.

The most common scenario in which live storage migration is used is upgrading the physical storage device that is hosting the source virtual hard disk. You may also want to move virtual machine storage between physical storage devices, at runtime, to take advantage of new, lower cost storage supported in this version of Hyper-V (such as server message block (SMB)-based) or to respond to reduced performance that results from bottlenecks in the storage throughput.

The steps that Hyper-V performs when you move virtual hard disks from one physical location to the other while the virtual machine is running are illustrated in the following figure.

Figure 11: Hyper-V moving virtual hard disks from one physical location to another



Windows Server 2012 provides the flexibility to move virtual hard disks both on shared storage subsystems and on non-shared storage as long as a network shared folder on Windows Server 2012 SMB is visible to both Hyper-V hosts.

Storage migration, combined with live migration, also lets you move a virtual machine between hosts on different cluster servers that are not using the same storage.

New Virtual Hard Disk Format

Hyper-V in Windows Server 2012 contains an update to the VHD format, called VHDX, that has much larger capacity and built-in resiliency. The VHDX format's principal new features are:

- Support for virtual hard disk storage capacity of up to 16 terabytes.
- Additional protection against data corruption during power failures by logging updates to the VHDX metadata structures.
- Improved alignment of the virtual hard disk format to work well on large sector physical disks.

The VHDX format also has the following features:

- Larger block sizes for dynamic and differential disks, which allows these disks to attune to the needs of the workload.
- 4-KB logical sector virtual disk that allows for increased performance when used by applications and workloads that are designed for 4-KB sectors.
- The ability to store custom metadata about the file that the user might want to record, such as operating system version, or updates that have been applied.
- Efficiency in representing data (called "trim"), resulting in smaller files size and allowing the underlying physical storage device to reclaim unused space. (Trim requires direct-attached storage or SCSI disks and trim-compatible hardware.)

Hyper-V support for 4-KB Disk Sectors and 512e

Most of the software industry has depended on disk sectors of 512 bytes, but the standard is moving to 4-KB disk sectors. To reduce compatibility problems that might arise from a change in sector size, hard drive vendors are introducing transitional 512-byte Emulation (512e) drives. These emulation drives offer some of the advantages offered by 4-KB native drives, such as improved format efficiency and an improved scheme for error correction codes (ECC), and they have fewer compatibility issues than would occur by exposing a 4-KB sector size at the disk interface.

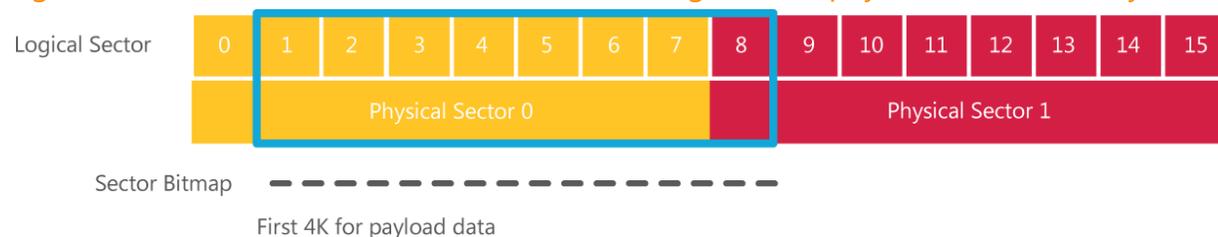
A 512e drive can perform a write only in terms of a physical sector—it can't directly write a 512-byte sector write issued to it. The internal process in the disk that makes this write possible follows these steps:

1. The disk reads the 4-KB physical sector into its internal cache, which contains the 512-byte logical sector referred to in the write.
2. Data in the 4-KB buffer is modified to include the updated 512-byte sector.
3. The disk performs a write of the updated 4-KB buffer back to its physical sector on the disk.

This process is called a "Read-Modify-Write," or RMW. This RMW process causes performance degradation in virtual hard disks for these reasons:

- Dynamic and differencing virtual hard disks have a 512-byte sector bitmap in front of their data payload. In addition, footer/header/parent locators all align to a 512-byte sector. It is common for the virtual hard disk drive to issue 512-byte writes to update these structures, resulting in the RMW behavior described above.
- Applications commonly issue read requests and write requests in multiples of 4-KB sizes (the default cluster size of NTFS). Because there is a 512-byte sector bitmap in front of the data payload block of dynamic and differencing virtual hard disks, the 4-KB blocks aren't aligned to the physical 4-KB boundary, as shown in the following figure.

Figure 12: Virtual hard disk 4-KB block (blue) not aligned with physical 4-KB boundary



Each 4-KB write issued by the current parser—to update the payload data—results in two reads for two blocks on the disk, which are then updated and subsequently written back to the two disk blocks.

The overall performance impact of the RMW process on the workloads usually ranged 30 to 80 percent and, at times, was even higher.

Hyper-V in Windows Server 2012 mitigates the performance effect of 512e disks on the virtual hard disk stack by preparing the previously mentioned structures for alignment to 4-KB boundaries in the VHD format. This avoids the RMW effect when accessing the data within the virtual hard disk file and when updating the virtual hard disk metadata structures.

Management

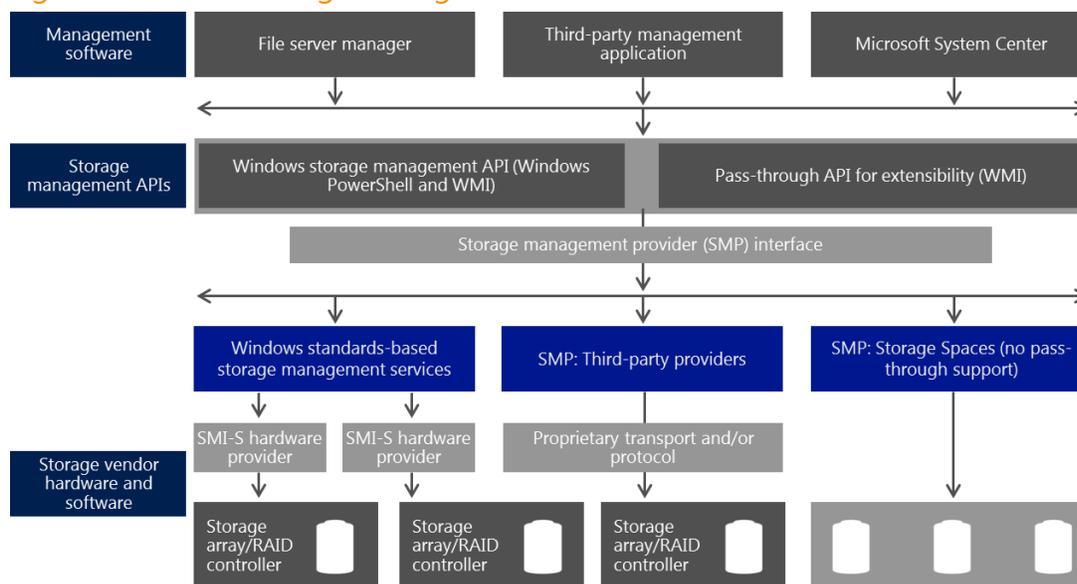
The management layer is where Windows Server 2012 brings it all together underneath a powerful and easy-to-use set of comprehensive management tools and capabilities that gives administrators the ability to manage storage for multiple servers, or even entire IT infrastructures, from a single pane of glass management interface or automate storage management tasks with powerful new Windows PowerShell cmdlets. Windows Server 2012 management capabilities extend even beyond Windows servers by leveraging industry standard Storage Management Initiative Specification (SMI-S) compatibility along with a new Storage Management Protocol (SMP) that can bring all storage management to your fingertips and bring you to entirely new levels of management ease and efficiency.

Comprehensive Windows Storage Management

To help tie all these amazing storage capabilities together, while making it simple to manage the growing complexity of your storage environment, Windows Server 2012 introduces a new, unified interface that uses Windows Management Instrumentation (WMI) for comprehensive management of physical and virtual storage, including third-party intelligent storage subsystems. The unified interface uses WMI to provide a rich experience to IT pros and developers using Windows PowerShell scripting, providing a diverse set of solutions. To illustrate the power of these solutions, and to help you take advantage of the WMI interface, Windows Server 2012 integrates powerful features such as high-performance computing (HPC), Windows PowerShell scripting, and Microsoft System Center. In addition, management applications can use a single Windows API to manage different storage types by using SMP or standards-based protocols such as SMI-S.

The unified interface for storage management provides a core set of defined WMI and Windows PowerShell interfaces, as well as features for more advanced management. The following figure shows the unified storage management architecture.

Figure 13: Unified storage management architecture



The unified interface is a powerful and consistent mechanism for managing storage, which can reduce complexity and operational costs. The storage interface provides capabilities for advanced management of storage in addition to the core set of defined WMI and Windows PowerShell interfaces. These capabilities offer specific advantages to distinct groups of users:

- **Enterprise system administrators:** The storage interface provides uniform, scriptable management that uses Windows PowerShell and a comprehensive set of cmdlets to help you automate and manage the following storage tasks:
 - Discovery
 - Thin provisioning support
 - Snapshot management
 - Replication
 - Masking and unmasking
 - Enumerating HBA ports
 - Creating pools, logical units, and volumes

For example, you can use a single script to configure host resources and then to configure and present the storage.

- **Independent Software Vendors (ISVs):** The storage interface includes the flexibility to manage any type of storage connected to a Windows-based system. With a WMI-based storage management provider, ISVs will appreciate a rich set of in-box storage management features included with the unified interface.
- **Storage manufacturers:** The seamless integration of devices with any storage management client provides a consistent experience for a manufacturer's customers.

The unified storage interface requires a storage management provider that is based on SMI-S or SMP. The interface integrates seamlessly with services offered by a SMI-S storage management provider, which allows you to manage third-party intelligent storage subsystems that have SMI-S providers.

Additionally, in Windows Server 2012, you can now use a single interface to remotely manage multiple file servers or clustered file server instances, as easily as managing one server.

Remote Storage Management

With the addition of the File and Storage Services role in Windows Server 2012, Server Manager includes a new functionality that allows you to remotely manage multiple file services, including managing file shares, the iSCSI Target along with storage pools and volumes, and storage.

The File and Storage Services management interface in Server Manager lets you to accomplish the following tasks:

- Create storage pools by grouping many physical disks on the network into easily managed volumes, and then manage those storage pools as a unit.
- Add and share volumes anywhere on the network by using SMB and NFS.
- Create and manage iSCSI virtual disks and targets.

Conclusion

With Windows Server 2012, Microsoft offers new storage solutions that provide continuously available services for organizations of all sizes. The new storage capabilities in Windows Server 2012 give you the power to leverage existing storage investments in ways that parallel the functionality and scalability of advanced storage area networks, while also reducing management costs and complexity. Features that are typically associated with high-end, expensive Fibre Channel SANs, such as thin provisioning, multipathing, deduplication, continuous availability, and live migration, are now within the grasp of organizations that otherwise might not have the budget or expertise to afford them.

Windows Server 2012 is a complete package that delivers proven, enterprise-class storage management and enables organizations to deploy business-critical applications with the same level of confidence previously reserved for traditional high-end software or hardware storage solutions.

List of charts, tables, and figures

Chart 1: Chkdsk - Runtime improvements (hours to seconds)	13
Chart 2: Typical savings when deduplication is applied to a variety of workloads.....	15
Table 1: Capabilities available at each layer of the Windows Server 2012 storage solution stack	4
Table 2: Models that include updated HBA drivers with in-box HBA drivers	34
Figure 1: Storage Spaces conceptual deployment model.....	8
Figure 2: Offloaded data transfer	10
Figure 3: On-disk transformation of files during Data Deduplication.....	14
Figure 4: Windows Server 2012 efficiently provides highly available iSCSI block storage for diskless network boot	19
Figure 5: Windows Server 2012 file server cluster with shared storage over NFS.....	21
Figure 6: Flexible configuration options for Microsoft SQL Server over SMB	23
Figure 7: Flexible configuration options for Hyper-V over SMB	24
Figure 8: VSS for SMB for Windows Server 2012.....	27
Figure 9: More securely replicating virtual machines from a wide range of systems and clusters to a remote site over a WAN	33
Figure 10: Alternating WWN addresses during a live migration	36
Figure 11: Hyper-V moving virtual hard disks from one physical location to another	37
Figure 12: Virtual hard disk 4-KB block (blue) not aligned with physical 4-KB boundary.....	38
Figure 13: Unified storage management architecture	39