

Introducing

Windows Server 2012 R2

Technical Overview

PUBLISHED BY
Microsoft Press
A Division of Microsoft Corporation
One Microsoft Way
Redmond, Washington 98052-6399

Copyright © 2013 Microsoft Corporation

All rights reserved. No part of the contents of this book may be reproduced or transmitted in any form or by any means without the written permission of the publisher.

Library of Congress Control Number: 2013945007
ISBN: 978-0-7356-8278-8

Microsoft Press books are available through booksellers and distributors worldwide. If you need support related to this book, email Microsoft Press Book Support at mspinput@microsoft.com. Please tell us what you think of this book at <http://www.microsoft.com/learning/booksurvey>.

Microsoft and the trademarks listed at <http://www.microsoft.com/about/legal/en/us/IntellectualProperty/Trademarks/EN-US.aspx> are trademarks of the Microsoft group of companies. All other marks are property of their respective owners.

The example companies, organizations, products, domain names, email addresses, logos, people, places, and events depicted herein are fictitious. No association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred.

This book expresses the author's views and opinions. The information contained in this book is provided without any express, statutory, or implied warranties. Neither the authors, Microsoft Corporation, nor its resellers, or distributors will be held liable for any damages caused or alleged to be caused either directly or indirectly by this book.

Acquisitions Editor: Anne Hamilton

Developmental Editor: Karen Szall

Project Editors: Valerie Woolley and Carol Dillingham

Editorial Production: Christian Holdener, S4Carlisle Publishing Services

Copyeditor: Andrew Jones

Indexer: Jean Skipp

Contents

<i>Introduction</i>	<i>ix</i>
Chapter 1 Cloud OS	1
The big picture	1
Journey to the Cloud OS	2
Let's begin!	3
Chapter 2 Hyper-V	5
Previous enhancements to Hyper-V	5
Generation 2 virtual machines	7
Automatic Virtual Machine Activation	11
Remote access over VMBus	12
Cross-version live migration	14
Faster live migration	16
Online VHDX resize	17
Live export	18
More robust Linux support	20
Managing Hyper-V hosts running previous versions of Windows Server	21
Hyper-V Replica enhancements	21
Additional Hyper-V improvements	25
Learn more	25

What do you think of this book? We want to hear from you!
Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

microsoft.com/learning/booksurvey

Chapter 3	Storage	27
	Previous enhancements to storage	27
	Microsoft's vision for storage	29
	Building the solution using Windows Server 2012 R2	29
	Enabling the solution using System Center 2012 R2	31
	Storage Management API	32
	Storage QoS	33
	iSCSI Target Server enhancements	34
	SMB 3.0 enhancements	36
	Data deduplication enhancements	38
	Storage Spaces enhancements	39
	Storage Spaces in Windows Server 2012	41
	Storage Spaces in Windows Server 2012 R2	42
	Work Folders	50
	Learn more	60
Chapter 4	Failover Clustering	63
	Previous enhancements to Failover Clustering	63
	Guest clustering using shared virtual disks	64
	Hosting highly available workloads	64
	Separating virtual resources from physical infrastructure	65
	Understanding shared virtual disks	67
	Using shared virtual disks	69
	CSV and SoFS enhancements	70
	Failover Clustering and SoFS	70
	Optimized ownership of CSV disks	71
	Increased CSV resiliency	72
	Improved CSV cache allocation	73
	CSV and other storage features	73

Changes to heartbeat threshold	74
Detecting the health of virtual machines	75
Virtual machine drain on shutdown	78
Dynamic witness	81
Active Directory-detached clustering	82
Learn more	82

Chapter 5 Networking 85

Previous enhancements to networking	85
Virtual RSS	88
Windows NIC Teaming enhancements	90
NIC Teaming in Windows Server 2012	91
NIC Teaming in Windows Server 2012 R2	92
Choosing the right teaming mode	94
Choosing the right load-balancing mode	95
Improved network diagnostics	96
Network diagnostic cmdlets in Windows Server 2012	97
Network diagnostic cmdlets in Windows Server 2012 R2	102
IPAM enhancements	105
IPAM in Windows Server 2012	107
IPAM in Windows Server 2012 R2	107
Hyper-V Network Virtualization enhancements	112
How Hyper-V Network Virtualization works	113
Hyper-V Network Virtualization enhancements in Windows Server 2012 R2	115
Hyper-V Virtual Switch enhancements	119
Hyper-V Virtual Switch in Windows Server 2012	119
Hyper-V Virtual Switch in Windows Server 2012 R2	121
Learn more	121

Chapter 6	Active Directory	123
	Previous enhancements to Active Directory	124
	Workplace Join	125
	Implementing Workplace Join	126
	Multi-factor access control	127
	Web Application Proxy	127
	Implementing Web Application Proxy	128
	Improved LDAP search performance.....	131
	Learn more.....	136
Chapter 7	Group Policy	137
	Previous enhancements to Group Policy	137
	Group Policy caching	138
	Group Policy Preferences and IPv6	139
	New policy settings	142
	Start screen	142
	Multi-monitor display	143
	Apps	144
	Sync your settings	145
	Work Folders	145
	Kerberos authentication	146
	Logon scripts	148
	Windows Update	149
	Windows Runtime apps	149
	Microsoft accounts	150
	Automatic sign-in	150
	Windows SkyDrive	151
	Learn more.....	152

Chapter 8	IIS	153
	Previous enhancements to IIS	153
	Dynamic Site Activation	155
	Idle Worker Process Page-out	156
	Custom IIS logging fields	158
	IIS ETW logging	159
	Learn more	162
Chapter 9	Remote Desktop Services	165
	Previous RDS enhancements	165
	Improved RemoteApp experience	167
	DX11.1 support	168
	Improved video experience	169
	Seamless display handling	169
	Quick Reconnect	169
	Session Shadowing	170
	VDI and data deduplication	172
	Restricted Admin mode	173
	Learn more	173
Chapter 10	Windows PowerShell	175
	Previous enhancements in Windows PowerShell 3.0	175
	Windows PowerShell Desired State Configuration	176
	Save-Help	184
	Yet more new cmdlets	186
	DISM cmdlets	186
	DHCP server cmdlets	189

DNS server cmdlets	190
Hardware certification cmdlets	191
Hyper-V cmdlets	191
iSCSI target cmdlets	195
Net TCP/IP cmdlets	196
Network Address Translation cmdlets	196
Network event packet capture cmdlets	197
Network Virtualization cmdlets	197
Physical Computer System View (PCSV) cmdlets	198
Software Inventory Logging cmdlets	198
Start screen cmdlets	198
Windows Deployment Services cmdlets	200
Windows PowerShell Web Access cmdlets	201
Windows Search cmdlets	201
Learn more.....	202

Chapter 11 Windows Server Essentials **203**

Learn more.....	214
-----------------	-----

<i>Appendix</i>	215
-----------------	-----

<i>Index</i>	217
--------------	-----

What do you think of this book? We want to hear from you!
 Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

microsoft.com/learning/booksurvey

Introduction

This book is intended to provide you with an overview of the new features and enhancements introduced in Windows Server 2012 R2. The intended audience for this book is IT pros who deploy, manage, and maintain Windows Server workloads in data center, private cloud, and hosting provider environments.

We assume that you are at least somewhat familiar with the features and capabilities of the previous platform Windows Server 2012. If you are not familiar with all the new features and enhancements Microsoft introduced previously in Windows Server 2012, we recommend that you first read *Introducing Windows Server 2012 RTM Edition* (Microsoft Press, 2012). This e-book is available as a free download from Microsoft in three formats:

- PDF from <http://aka.ms/682788pdf>
- EPUB from <http://aka.ms/682788epub>
- MOBI from <http://aka.ms/682788mobi>

A key feature of this book is the technical sidebars that have been contributed by Microsoft insiders. These sidebars were written by experts who have been closely involved in the Windows Server 2012 R2 development process and include Program Managers, Support Escalation Engineers, Technical Consultants, Data Center Specialists, and others who work at Microsoft in various capacities.

Acknowledgments

Three groups of people have helped make this book possible, and as the author I'd like to thank them all here. First, the following experts at Microsoft have contributed sidebars that explain and demonstrate different aspects of Windows Server 2012 R2:

- Deepak Srivastava
- Erez Benari
- Gene Chellis
- Jason M. Anderson
- Jeff Butte
- John Marlin
- Justin Turner
- Mark Gehazi
- Nir Ben Zvi

Second, the following Microsoft insiders have peer reviewed various portions of this book to help us ensure our content is as accurate as possible:

- Aanand Ramachandran
- Adam Carter
- Ben Armstrong
- Bryan Matthew
- CJ Williams
- Clinton Ho
- Deepak Srivastava
- Elden Christensen
- Erez Benari
- Gabriel Silva
- Guang Hu
- Jason Gerund
- Jeff Woolsey
- John Savill
- Jose Barreto
- Matthew John
- Raghavendran Gururajan
- Roiy Zysman
- Shivam Garg
- Symon Perriman
- Vijay Tandra Sistla
- Vijay Tewari
- Yang Cao Sun
- Yuheng Cai sun

Finally, I'd also like to thank Valerie Woolley, Content Project Manager at Microsoft Press; Christian Holdener at S4Carlisle Publishing Services; and copyeditor Andrew Jones.

Errata & book support

We've made every effort to ensure the accuracy of this content and its companion content. Any errors that have been reported since this content was published are listed on our Microsoft Press site:

<http://aka.ms/introwinsrv2012R2/errata>

If you find an error that is not already listed, you can report it to us through the same page.

If you need additional support, email Microsoft Press Book Support at *mspinput@microsoft.com*.

Please note that product support for Microsoft software is not offered through the addresses above.

We want to hear from you

At Microsoft Press, your satisfaction is our top priority, and your feedback our most valuable asset. Please tell us what you think of this book at:

<http://aka.ms/tellpress>

The survey is short, and we read every one of your comments and ideas. Thanks in advance for your input!

Stay in touch

Let's keep the conversation going! We're on Twitter:

<http://twitter.com/MicrosoftPress>.

Storage

Storage is a key part of any IT infrastructure. For many organizations, storage is also a major cost center that consumes a large portion of the budget of the IT department. Maximizing the performance and efficiency of storage while helping to reduce costs was a major goal of Windows Server 2012, and the numerous new features and capabilities introduced in that platform now provide enterprises with new ways of squeezing the most out of shrinking IT budgets when it comes to storage.

Windows Server 2012 R2 takes these previous enhancements a step further and enables organizations to re-envision how storage infrastructure can be built and managed. This chapter examines the new storage features in Windows Server 2012 R2, focusing mainly on improvements to file- and block-based storage in the platform. The next chapter will build on this one by examining how failover clustering has been enhanced in Windows Server 2012 R2 and how Hyper-V virtualization can benefit from the improvements to storage and failover clustering. But first we'll briefly review some of the storage enhancements that were previously introduced in Windows Server 2012.

Previous enhancements to storage

A number of new storage capabilities were introduced in the previous version of Windows Server, and space doesn't allow us to describe each of them in detail. But as a quick summary, the following might be characterized as some of the more significant storage improvements in the platform:

Storage Spaces Storage Spaces provided storage virtualization capabilities that allow you to group industry-standard disks (such as Serial ATA or Serial Attached SCSI disks) into storage pools. You could then create virtual disks called "storage spaces" or "spaces" from the available capacity in the pools and provision resilient storage volumes as they were needed. This allowed you to make more efficient use of disk capacity, provision storage quickly and easily without impacting users, delegate the administration of storage, and provide cost-effective storage for business-critical applications that make use of low-cost, commodity-based just-a-bunch-of-disks (JBODs).

SMB 3.0 A new version of the Server Message Block (SMB) protocol, a network file sharing protocol that enables applications on a computer to read and write to files and to request services from server programs on a network, was introduced in

Windows Server 2012. These improvements included SMB Direct, SMB Multichannel, SMB Transparent Failover, and other enhancements that enabled new scenarios such as storing Hyper-V virtual machine files and Microsoft SQL Server database files on file shares on an SMB 3.0 file server (called a Scale-out File Server) instead of having to store these files on local storage, Direct Attached Storage (DAS), or a Storage Area Network (SAN) array.

ReFS The new Resilient File System (ReFS) introduced in Windows Server 2012 provided enhanced integrity, availability, scalability, and error protection for file-based data storage. ReFS supported volume sizes up to 18 exabytes and could be especially useful on file servers storing large amounts of data or running disk-intensive applications that require high levels of performance. ReFS in Windows Server 2012 did not support certain NTFS features, however, such as disk quotas.

Data deduplication Data deduplication allowed more data to be stored in less space without compromising the integrity or fidelity of the data stored on the volume. It accomplished this by segmenting files into small, variable-sized chunks, identifying any duplicate chunks present, and maintaining only a single copy of each unique chunk of data. One scenario where this was useful was when data was transferred over the WAN to a branch office using the BranchCache feature of Windows Server 2012.

iSCSI Target Server The iSCSI Target Server provides block storage to servers and applications on the network using the Internet SCSI (iSCSI) standard. When combined with other availability technologies in Windows Server 2012, iSCSI Target Server provided continuously available storage that previously required organizations to purchase expensive, high-end SAN arrays.

ODX Offloaded Data Transfer (ODX) functionality in Windows Server 2012 enabled ODX-capable storage arrays to bypass the host computer and directly transfer data within or between compatible storage devices. The result was to minimize latency, maximize array throughput, and reduce resource usage, such as CPU and network consumption on the host computer. For example, by using ODX-capable storage arrays accessed via iSCSI, Fibre Channel, or SMB 3.0 file shares, virtual machines stored on the array could be imported and exported much more rapidly than they could without ODX capability being present.

Chkdsk Windows Server 2012 introduced a new Chkdsk model that allowed organizations to confidently deploy large, multiterabyte NTFS file system volumes without worrying about their availability being compromised should file system corruption be detected on them. The new version of Chkdsk ran automatically in the background and actively monitored the health state of the file system volume. Should file system corruption be detected, NTFS now instantaneously self-healed most issues online without requiring Chkdsk to run offline. This means that the amount of time needed for running Chkdsk on multiterabyte data volumes can be reduced from hours to only a few seconds, plus in many scenarios you won't even need to take the disk offline and run Chkdsk on it at all.

Storage management improvements Beginning with Windows Server 2012, you could now use the File and Storage Services role in Server Manager to remotely manage multiple file servers running Windows Server 2012, including their storage pools, volumes, shares, and iSCSI

virtual disks, all from a single user interface. You could also use the new Windows PowerShell cmdlets in Windows Server 2012 to automate the same storage management tasks.

Now that we've reviewed the storage improvements introduced previously in Windows Server 2012, let's move on and look at some of the new storage capabilities and enhancements added in Windows Server 2012 R2.

Microsoft's vision for storage

As you can see from the previous section, Windows Server 2012 introduced a lot of new storage features and capabilities to the Windows Server platform. Together with System Center 2012 SP1, Windows Server 2012 provided organizations with a cost-effective solution for building and deploying private clouds using file-based storage access composed of low-cost commodity storage accessed over a standard Ethernet network.

Although the R2 release of Windows Server 2012 adds a number of incremental improvements to both file- and block-based storage and to how storage is managed on the Windows Server platform, it also represents something more. Microsoft's vision and goals with respect to storage for this new release are threefold:

- To greatly reduce the capital and operational storage and available costs for organizations deploying Infrastructure-as-a-Service (IaaS) services for private clouds, hosted clouds, and cloud service providers.
- To disaggregate compute and storage resources so they can be independently managed and scaled at each layer of cloud infrastructure.
- To allow enterprises to build software-defined storage solutions using inexpensive, industry-standard servers, networks, and shared JBOD storage.

With this focus in mind, the incremental improvements to storage capabilities in Windows Server 2012 R2 are designed to specifically target the above three goals.

Building the solution using Windows Server 2012 R2

To understand how Windows Server 2012 R2 can be used to implement the above vision for cloud computing, let's look at an example. Figure 3-1 shows the compute, networking, and storage components of a simple private cloud solution built using the Windows Server platform. You can think of this solution as having four layers as follows:

- **Compute layer** At the top are several Hyper-V hosts joined together in a failover cluster. These hosts use commodity server hardware to provide cost-efficient scale-out capabilities. For example, if the solution needs more processing power to run more workloads running in virtual machines, you can add another commodity server to the Hyper-V cluster. Utilizing the scale-out approach like this is often a more cost-effective solution for organizations than using a scale-up solution that involves only two expensive high-end host machines, where you need to add another processor to each host if you want to run more workloads.

- **Network layer** A low-cost industry-standard Ethernet network is used to connect the Hyper-V cluster that provides compute resources for the solution with the Scale-out File Servers (SoFS) that provide virtualized storage resources for the cloud. This kind of approach can be a lot more cost-effective for many organizations than utilizing a proprietary SAN for their storage layer. That's because you don't need to install expensive host bus adapters (HBAs) in the Hyper-V hosts to enable them to connect to storage volumes (logical unit numbers, or LUNs) provisioned on the SAN. This approach is only possible, however, because of new capabilities in version 3.0 of the SMB protocol that was first introduced in Windows Server 2012.
- **Virtualized storage layer** The virtual machines running on the clustered Hyper-V hosts have their virtual machine files (virtual hard disks, configuration files, snapshots, and so on) stored on SoFS. SoFS was first introduced in Windows Server 2012 and represented clustered file servers that allow you to store server application data, such as Hyper-V virtual machine files or SQL Server database files, on file shares while maintaining a similar level of reliability, availability, manageability, and high performance as using a SAN for storing such files. With a SoFS, all file shares are online on all nodes simultaneously in an active-active cluster configuration. Again, this kind of approach can be much more cost-effective for organizations than using a SAN for storing the virtual machine files for a Hyper-V cluster. And again, using a scale-out approach instead of scale-up can be a more cost-effective solution, and as we'll see in the next chapter, Windows Server 2012 R2 increases the scale-out capabilities of the SoFS role service.

To enable the virtualization of storage resources, Storage Spaces can be used on SoFS. This allows the physical storage resources for the solution to be pooled together to provide resiliency in case of failure of a storage device. Storage Spaces was also introduced in Windows Server 2012 and provides two types of resiliency—mirroring and parity. Storage devices can be selectively reserved as hot spares so they can automatically replace devices that fail, thus ensuring that the integrity of data is preserved in the event of a power interruption or hardware failure. The storage pools you create using Storage Spaces can be used to provision virtual disks (virtualized storage) on which you can create new volumes and shares for your solution.

- **Physical storage layer** Depending on the performance needs of the workloads running in the solution, the virtual machine files for the Hyper-V clusters can be stored on different types of storage devices. Supported drives include Serial ATA (SATA) and Serial Attached SCSI (SAS) disks, which can be either hard disk drives (HDDs) or solid-state drives (SSDs). These disks can either be internal to the SoFS, directly connected as DAS, or within JBOD enclosures.

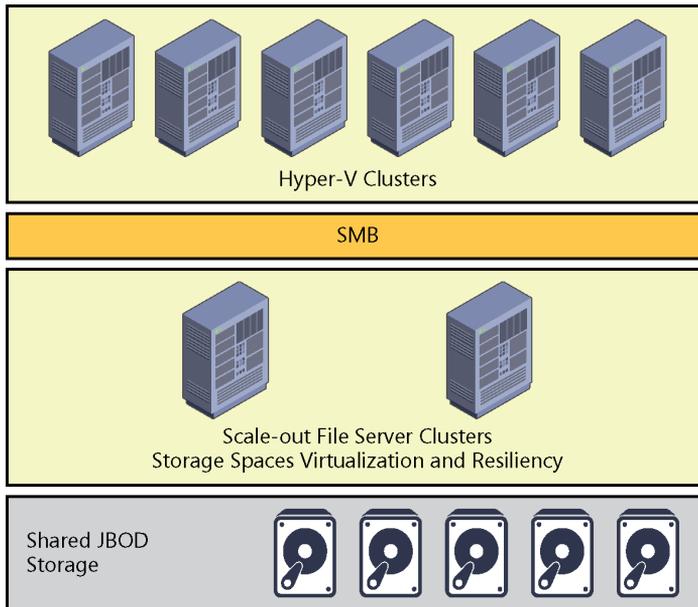


FIGURE 3-1 Microsoft’s vision for storage can be implemented by using Windows Server 2012 Hyper-V, SoFS clusters, and shared JBOD storage.

Although this kind of cloud solution is already possible using Windows Server 2012, the enhancements to virtualization, storage, and networking in Windows Server 2012 R2 now make it possible to optimize such solutions to achieve enterprise-quality levels of performance, reliability, and availability. That’s why Windows Server 2012 R2 is being positioned by Microsoft as a cloud-optimized server operating system.

Enabling the solution using System Center 2012 R2

Windows Server 2012 R2 only represents the foundation for a cloud solution. To manage the cloud solution and the workloads running on it, you also need System Center, and in particular you need System Center Virtual Machine Manager (VMM). What truly makes this new release of Windows Server a cloud-optimized operating system is that it represents the first time that Microsoft has synchronized the product release cycles of Windows Server and System Center. The goal of doing this is to ensure that Microsoft can deliver to both its enterprise and service provider customers a completely integrated solution for building and deploying both private and hosted clouds.

To achieve this goal, the R2 release of System Center 2012 (particularly in VMM 2012 R2) also includes numerous enhancements, particularly in the areas of storage performance, provisioning, and management. For example, System Center 2012 R2 now supports:

- Faster enumerations through its Storage Management Initiative - Specification (SMI-S) storage service
- Real-time updates for out-of-band changes using Common Information Model (CIM) indications
- Fibre Channel fabric discovery and zone provisioning
- Support for Hyper-V Virtual Fibre Channel
- ODX optimized virtual machine deployments
- Rapid provisioning using differencing disks

Although the focus of this book is on Windows Server 2012 R2 and its new features and enhancements, System Center 2012 R2 (and particularly VMM) should really be considered the default platform going forward for managing a cloud solution built using Windows Server 2012 R2 as its foundation.

Storage Management API

One of the key storage management improvements introduced in Windows Server 2012 and System Center 2012 was the Storage Management application programming interface (SM-API). SM-API is a Windows Management Infrastructure (WMI)-based programming interface that provides a standards-based way of managing storage on the Windows Server platform, and it supersedes the Virtual Disk Service (VDS) API used in previous versions of Windows Server. And in VMM 2012, the new Windows Standards-Based Storage Management service—which utilizes SM-API—replaces the Microsoft Storage Management Service used in previous versions of VMM.

Figure 3-2 shows how SM-API can be used to manage different kinds of storage providers, arrays, and devices by using Server Manager, VMM, or a third-party storage management tool. (You can also use Windows PowerShell to manage storage on Windows Server 2012.) Some examples of different types of storage providers and arrays you can manage using SM-API include:

- Older symmetric multiprocessing (SMP) based internal storage subsystems
- Newer SMI-S-based internal storage subsystems
- SMI-S-based network attached storage (NAS) devices
- SMI-S-compliant Fibre Channel switches on SAN arrays using CIM pass through
- JBODs that are compatible with Storage Spaces using CIM pass through

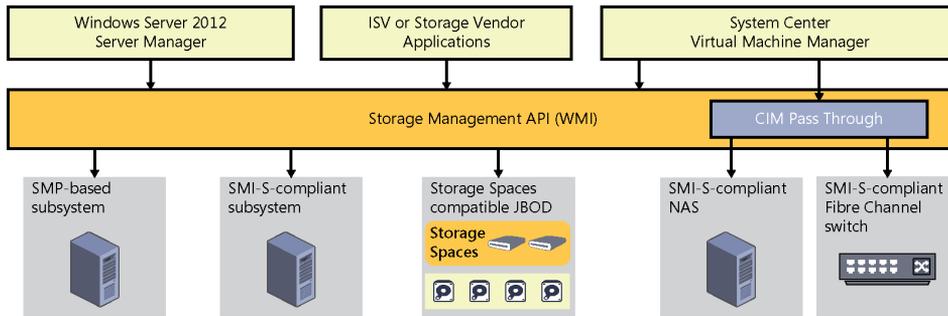


FIGURE 3-2 SM-API can be used to manage a wide range of storage providers and arrays.

Although SM-API isn't new in this R2 release of Windows Server 2012, it has been enhanced in several ways, and especially in the area of performance. Specifically, the new version of SM-API includes:

- A new architecture that performs enumerations of storage resources 10 times faster than previously
- The addition of remoting and cluster-awareness when managing Storage Spaces
- Support for new Storage Spaces features like write-back caching and storage tiering which are described later in this chapter
- The ability to use VMM to manage Storage Spaces and SoFS using SM-API

Storage QoS

Storage Quality of Service (QoS) is another new feature of file-based storage introduced in Windows Server 2012 R2. Storage QoS is enabled at the VHDX layer and allows you to limit the maximum IOPS allowed to a virtual disk on a Hyper-V host. It can also allow you to set triggers to send notifications when a specified minimum IOPS is not met for a virtual disk.

Possible usage scenarios for this feature include:

- Configuring different service-level agreements (SLAs) for different types of storage operations within your infrastructure. For example, a hoster might use this feature to configure Bronze, Silver, and Gold SLAs for storage performance available for different classes of tenants. You can even set alerts that trigger when virtual machines are not getting enough IOPS for storage access.
- Restricting the disk throughput for overactive or disruptive virtual machines within your environment that are saturating the storage array. Hosting providers will especially love this capability since it means they won't have to worry about one tenant consuming excessive storage fabric resources at the expense of other tenants.

As Figure 3-3 shows, Storage QoS can even be configured while the virtual machine is running. This allows organizations to have a lot of flexibility in how they manage access to the storage fabric from workloads running in their cloud environments.

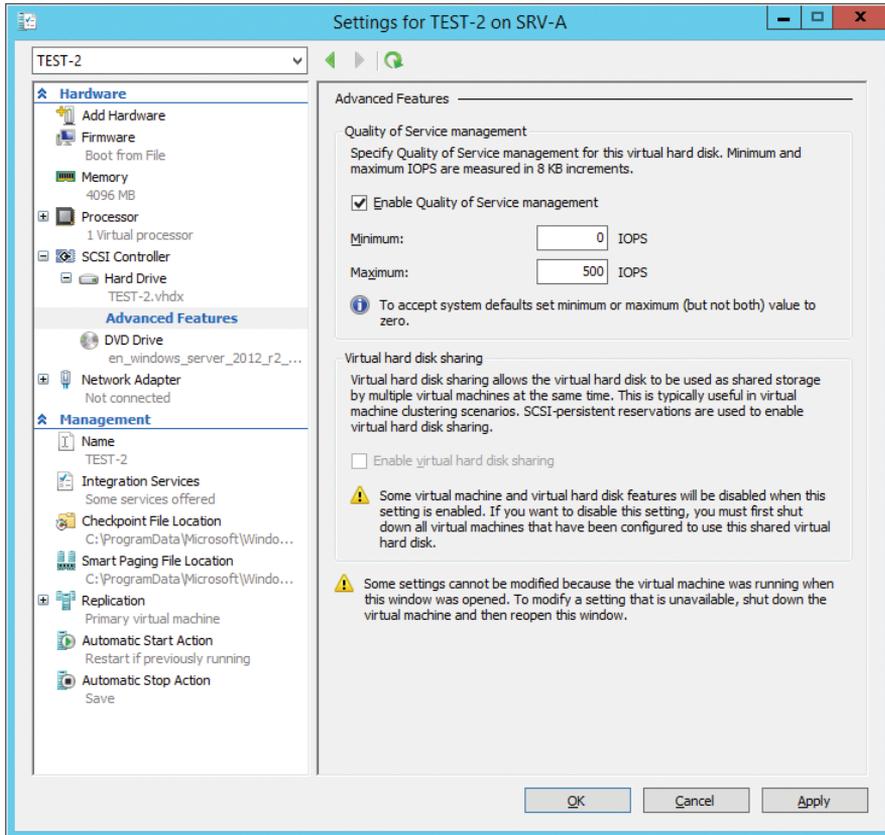


FIGURE 3-3 You can configure Storage QoS for a virtual machine using Hyper-V Manager.

iSCSI Target Server enhancements

Internet Small Computer System Interface (iSCSI) is an industry-standard protocol that allows sharing of block-level storage over a TCP/IP network. Block-level storage is typically used in SANs and is supported by the iSCSI, Fibre Channel, and SAS connection mechanisms. File-level storage involves using network shares on volumes that have been formatted using a file system like NTFS or ReFS.

iSCSI is designed to transmit and receive Small Computer System Interface (SCSI) commands and data encapsulated as TCP packets. This enables servers to utilize storage on an iSCSI-based storage device, such as an iSCSI SAN, even when the servers and SAN are in different locations.

Fibre Channel SANs can be prohibitively expensive for a small or midsized business because they require specialized connection hardware such as HBAs and cabling. By contrast, iSCSI needs no specialized connection hardware or special-purpose cabling because it can use a standard Ethernet network for connecting servers with the storage array. This means that iSCSI storage can be deployed using an organization's existing network infrastructure, which helps keep the cost of the iSCSI approach low.

Beginning with Windows Server 2012, a built-in role service (iSCSI Target Server) and client component (iSCSI Initiator) are included and can be used to implement an iSCSI-based storage solution without the need of purchasing a third-party iSCSI SAN. Using these new features, organizations can deploy iSCSI storage without the need of purchasing any additional storage hardware or software.

Some of the usage scenarios for iSCSI storage include:

- Deploying diskless servers that boot from iSCSI virtual disks over the network
- Providing block storage to applications that require or can benefit from it
- Creating iSCSI storage test environments where you can validate applications before deploying them onto a third-party iSCSI SAN

Because Microsoft iSCSI technologies are based on industry standards, you can also deploy Windows Server–based iSCSI storage solutions together with third-party solutions.

One of the key components of the iSCSI implementation in Windows Server 2012 is iSCSI Target Server, a role service under the File and Storage Services role. In a typical iSCSI storage scenario, an iSCSI initiator (a service running on the server consuming the storage) establishes a session (consisting of one or more TCP connections) with an iSCSI target (an object on the target server that allows an iSCSI initiator to establish a connection with the target server) in order to access an iSCSI virtual disk (storage backed by a virtual hard disk file) on the iSCSI Target Server (the server or device, such as a SAN, that shares storage so that users or applications running on a different server can consume the storage). Figure 3-4 shows how these different iSCSI components work together on the Windows Server platform.

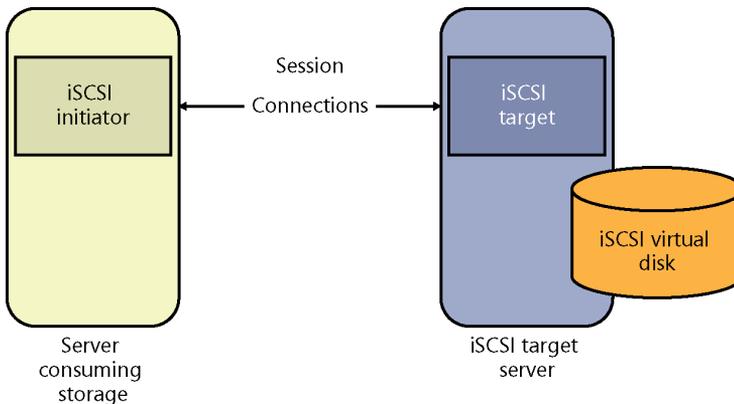


FIGURE 3-4 The basic iSCSI components work together on Windows Server.

In the R2 release of Windows Server 2012, the iSCSI Target Server role service has been enhanced in a couple of important ways:

- VHDX support is now included, which enables the provisioning of larger LUNs up to 64 TB in size. VHDX support also means you can now expand or shrink iSCSI LUNs while they are online (see the section “Online VHDX resize” in Chapter 2, “Hyper-V”) and dynamically grow them for greater scalability and flexibility. And VHDX is now the default virtual disk format when creating new iSCSI LUNs.
- iSCSI Target Server can now be fully managed using SMS-S. This means that you can now perform end-to-end management of your iSCSI storage system using VMM.

SMB 3.0 enhancements

SMB 3.0 is at the core of the new SoFS functionality introduced previously in Windows Server 2012 and enables file-based storage solutions (file servers) to provide network storage for servers that have similar performance to expensive, proprietary SAN solutions. Using SMB 3.0 and the Windows Server 2012 platform, organizations can build low-cost scale-out storage fabrics that can meet the needs of a broad range of private and hosted cloud solutions.

Some of the key features of SMB 3.0 include:

- **SMB Scale Out** Allows you to create file shares using Cluster Shared Volumes (CSV) to provide simultaneous access to data files, with direct I/O, through all nodes in a file server cluster. This results in improved utilization of network bandwidth, load balancing of SMB 3.0 clients, and optimized performance for server applications.
- **SMB Transparent Failover** Allows you to perform hardware or software maintenance of nodes in a clustered SMB 3.0 file server (an SoFS) without interruption to server applications storing data on the file shares. If a hardware or software failure occurs on a cluster node, SMB 3.0 clients can transparently reconnect to another cluster node without interruption as well.
- **SMB Direct** Supports using network adapters that have Remote Direct Memory Access (RDMA) capability, which allows them to function at full speed with very low latency and very low CPU utilization. When used with workloads, such as Hyper-V or SQL Server, the result is that a remote SMB 3.0 file server can resemble local storage in its performance.
- **SMB Multichannel** Provides aggregation of network bandwidth and network fault tolerance when multiple paths are available between the SMB 3.0 client and the SMB 3.0 file server. This results in server applications taking full advantage of all available network bandwidth and being more resilient to network failure.
- **SMB Encryption** Provides end-to-end encryption of SMB data to protect data from eavesdropping without the need for configuring Internet Protocol security (IPsec), deploying specialized hardware, or utilizing WAN accelerators. Encryption can be configured on a per share basis or for the entire SMB 3.0 file server.

With the release of version 3.0 of SMB, the SMB protocol has become more than just a network file sharing protocol used for copying files over a network. Some of the additional uses for SMB now include:

- A protocol transport for CSV that enables I/O forwarding between the nodes of a failover cluster.
- A storage protocol that enables Hyper-V hosts to access and run virtual machine files stored on file servers on the network.
- A protocol transport for performing live migrations of virtual machines between clustered and nonclustered Hyper-V hosts.

A number of improvements have been made to SMB 3.0 in Windows Server 2012 R2. For example, the performance of SMB Direct has been enhanced to provide a 50 percent improvement for small IO workloads when used with RDMA-capable network adapters. For example, 8KB data transfers have now increased from about 300K I/O operations per second (IOPS) to about 450K IOPS per interface.

A second improvement is the increased efficiency and density of hosting workloads with small I/Os, for example, when running an online transaction processing (OLTP) database workload inside a virtual machine. SMB Direct in Windows Server 2012 R2 also includes optimizations for using 40 Gbps Ethernet and 56 Gbps InfiniBand for network transport.

SMB connections can also now be managed per share on SoFS instead of per file server as in Windows Server 2012. But we'll defer further discussion of this until we reexamine SoFS in Chapter 4, "Failover Clustering."

Another new feature of SMB 3.0 in Windows Server 2012 R2 is SMB Bandwidth Management. Because SMB now has so many different functions in a network and storage infrastructure built using Windows Server 2012, it now represents a common infrastructure component in many environments. That means it's important to be able to control how much bandwidth SMB uses when it's performing many different tasks within an infrastructure. As Figure 3-5 shows, you can now configure bandwidth limits for different categories of SMB usage. For example, the figure shows that three categories of bandwidth limits have been configured to ensure optimal performance of the various infrastructure components present in this infrastructure:

- **Default** A limit of 100 MB/s has been configured for Hyper-V host 1 to use SMB when performing file copies from the file server used for library storage by VMM 2012 R2.
- **VirtualMachine** No limit has been set for the amount of bandwidth that Hyper-V host 1 can utilize when using SMB to access virtual machine files stored on the SoFS.
- **LiveMigration** A limit of 500 MB/s has been set for SMB to use when performing live migrations of virtual machines from Hyper-V host 1 to Hyper-V host 2.

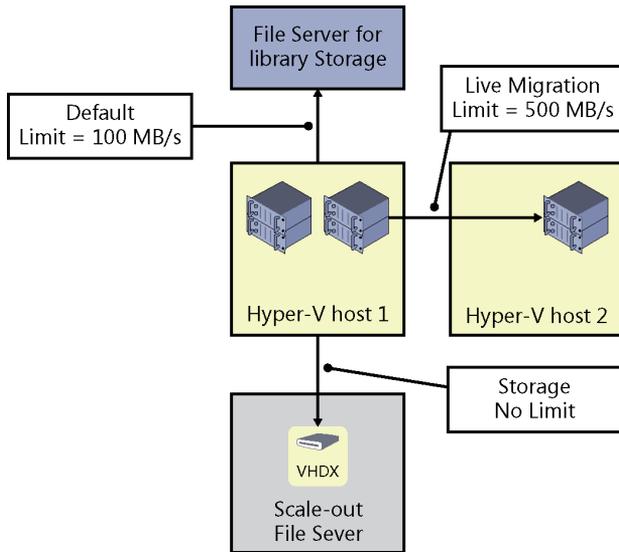


FIGURE 3-5 SMB 3.0 now supports bandwidth management.

Data deduplication enhancements

Data deduplication was introduced in Windows Server 2012 to help enterprises cope with exponentially increasing growth of data storage in their environments. Data deduplication allows Windows Server 2012 to store more data in less physical space to optimize the capacity of their storage fabric. Data deduplication is highly scalable, resource efficient, and nonintrusive in Windows Server 2012 and can run on multiple volumes simultaneously without affecting other workloads running on the server. Checksums, consistency, and identity validation are used to ensure data integrity, and redundant copies of file system metadata are maintained to ensure data is recoverable in the event of corruption.

Windows Server 2012 R2 includes several important improvements to the way data deduplication works. For example, in the previous version deduplication could only be used with files that are closed, such as virtual machine files stored in the VMM library. With this new release, however, deduplication can now be used even with open virtual hard disk files (both VHD and VHDX) for virtual desktop infrastructure (VDI) scenarios using SMB 3.0 remote storage.

Deduplication in Windows Server 2012 was also incompatible with CSVs. This meant that deduplication couldn't be used to optimize storage of virtual machine files stored on SoFS. This limitation has now been removed in the R2 release of Windows Server 2012 with support for deduplication of data stored on CSVs used by SoFS. For example, Figure 3-6 shows a failover cluster of Hyper-V hosts with the virtual machine files being stored on CSVs used by a two-node SoFS. In order for this scenario to work, SMB 3.0 must be used as the network

storage protocol. Although SoFS running Windows Server 2012 can provide this functionality, using SoFS running Windows Server 2012 R2 enables deduplication to be turned on for the CSVs, which enables space savings as high as 90 percent on the CSVs. This type of space savings can be especially beneficial for Windows Server 2012 or Windows Server 2012 R2 Hyper-V or VDI deployments using remote storage on the Windows Server 2012 R2 SoFS.

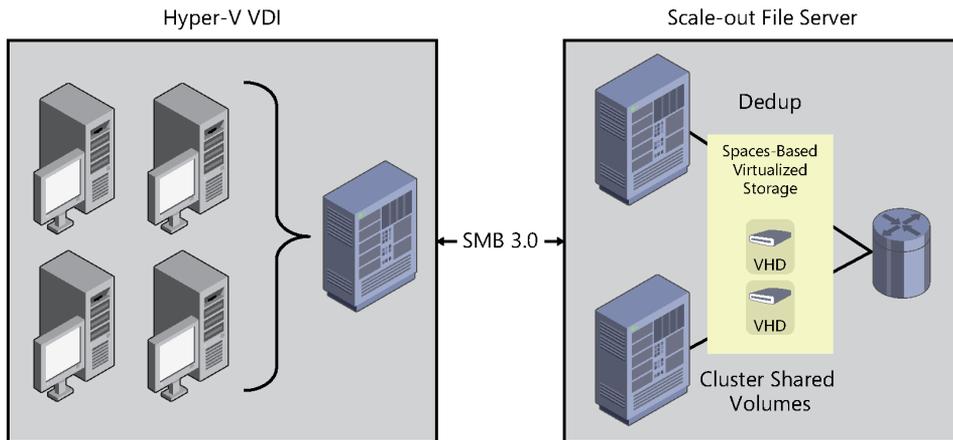


FIGURE 3-6 Windows Server 2012 R2 now supports data deduplication on CSV volumes and live VDI image files on SMB 3.0 storage.

Other improvements to data deduplication in Windows Server 2012 R2 include performance enhancements resulting from faster read/write of optimized files and improved optimization speed. Deduplication is supported only for data volumes, however, and not for boot or system volumes. In addition, ReFS volumes do not support using deduplication.

Storage Spaces enhancements

Until Windows Server 2012 was released, implementing storage virtualization required purchasing proprietary third-party SAN solutions that were expensive and required using their own set of management tools. Such solutions also required special training to implement, manage, and maintain them effectively. Storage Spaces, first introduced in Windows Server 2012, was designed to make storage virtualization affordable even for small businesses. Storage Spaces is simple to deploy and manage, and it can provide businesses with shared storage that can grow on demand to meet an organization's changing needs.

Some of the benefits of using Storage Spaces include:

- **Increased scalability** Additional physical storage can easily be added and used to meet increasing business demands.
- **Increased flexibility** New storage pools can be created and existing ones expanded as the need arises.

- **Increased efficiency** Unused storage capacity can be reclaimed to enable more efficient use of existing physical storage resources.
- **Increased elasticity** Storage capacity can be preallocated by using thin provisioning to meet growing demand even when the underlying physical storage is insufficient.
- **Lower cost** Low-cost, commodity-based storage devices can be used to save IT departments money that can be better allocated elsewhere.

To understand how Storage Spaces might be used for private cloud solutions, Figure 3-7 compares a traditional SAN-based storage solution with one built using Storage Spaces in Windows Server 2012. On the left side is a cluster of Hyper-V hosts whose virtual machine files are stored in LUNs on the SAN. These LUNs are backed by enterprise-class SAS disks (which can be HDDs or SSDs) mounted in disk shelves in the SAN chassis. Establishing connectivity between the Hyper-V host cluster and the SAN requires installing Fibre Channel or iSCSI HBAs in these hosts (depending on type of SAN involved) and data is transferred between the SAN and the Hyper-V host cluster using either Fibre Channel or iSCSI as a block-level storage protocol. Proprietary technology is required for this solution in the form of HBAs, cabling, and the SAN chassis.

By comparison, using Storage Spaces on the right requires no use of proprietary technology. Instead, all of the components of this solution can use off-the-shelf commodity-based server hardware. The SAN chassis is replaced with a cluster of file servers running Windows Server 2012 on enterprise-level server system hardware and rack-based JBOD enclosures containing the same kind of enterprise-class SAS disks (HDDs or SSDs) that might be used in the traditional SAN approach. Connectivity between the Hyper-V host cluster requires only a standard Ethernet network (typically 10 GbE) and high-performance network interface cards installed in the Hyper-V hosts and uses SMB 3.0 as a file-based storage protocol. Everything is commodity hardware here, and there's no vendor lock-in as there is with the traditional SAN approach.

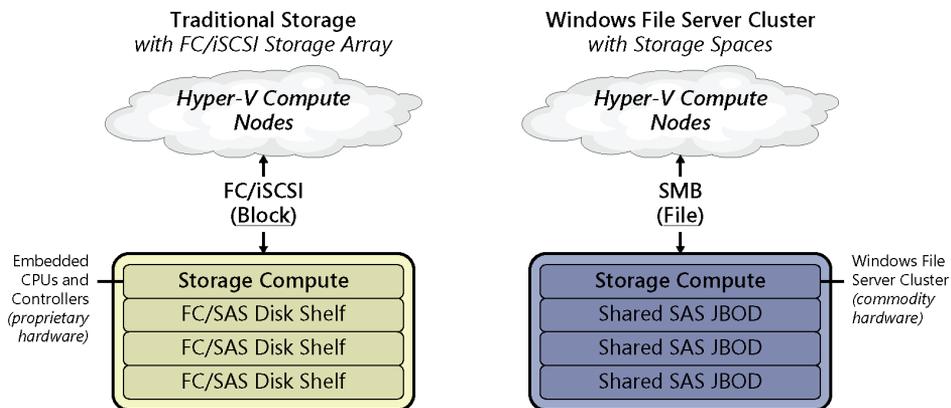


FIGURE 3-7 A comparison of traditional SAN storage with one based on Storage Spaces.

Storage Spaces in Windows Server 2012

The reaction when Storage Spaces was announced was somewhat qualified, especially by large enterprises that run workloads requiring the highest levels of performance involving millions of IOPS and massive throughput. The reason for this was because it was initially assumed that the performance of a virtualized storage solution based on Storage Spaces would fall short of what a typical SAN array can deliver. However, Microsoft soon proved its critics wrong with a demonstration performed at TechEd 2012, where a three-node high-performance server cluster was connected to a 24-bay JBOD filled with enterprise-grade SSDs. When Storage Spaces was used to present storage from the JBOD to the cluster nodes, the performance resulted in an aggregate sequential throughput of 12 GB/s and 1.45 million IOPS. When a second JBOD was added to the environment, the IOPS were increased to 2.7 million! Clearly, Storage Spaces is an enterprise-ready storage virtualization technology and its usage scenarios are not limited only to smaller deployments.

The challenge, however, with the Windows Server 2012 version of Storage Spaces is deciding whether you want to optimize performance or storage capacity when building your storage virtualization solution. For example, if you use Storage Spaces to create storage pools backed by low-cost, large-capacity commodity HDDs, you get a capacity-optimized storage solution, but the performance might not be at the level that some of your applications require. This is typically because large-capacity HDDs are optimized for sequential data access, whereas many server applications perform best with random data access. On the other hand, if you create pools using more expensive SSDs, you can easily achieve the kind of random I/O performance your applications require, but you probably won't have enough room in your budget to meet your capacity requirements for storage.

The logical solution is to use a mix of low-cost, large-capacity commodity HDDs together with expensive, high-performance enterprise-class SSDs. Building a Storage Spaces solution along these lines can provide you with the best of both worlds and deliver high levels of IOPS at a relatively low cost compared to using a SAN. This means that there are three ways you can build a virtualized storage solution using Storage Spaces in Windows Server 2012:

- **Capacity-optimized approach** Uses only low-cost, large-capacity commodity HDDs to provide high capacity while minimizing cost per terabyte
- **Performance-optimized approach** Uses only expensive, high-performance enterprise-class SSDs to provide extreme performance, high throughput, and the largest number of IOPS per dollar
- **Balanced approach** Uses a mix of HDDs and SSDs to achieve good performance and reasonable capacity at an acceptable cost

Unfortunately, there's a problem with the balanced approach. This is because although most enterprise workloads have a relatively large data set, the majority of data in this working set is often cold (seldom-accessed) data. Only a minority of data is typically in active use at a given time, and this hot data can be considered the working set for such workloads. Naturally, this working set also changes over time for the typical server workload. Since the working set is small, it would seem natural to place the hot data (the working set) on high-performance

SSDs while keeping the majority of the data (which is cold data) on high-capacity HDDs. But the working set changes over time, so how do you seamlessly ensure that hot data is placed on SSDs and cold data on HDDs when you use Storage Spaces to create pools containing both SSDs and HDDs?

The answer is you couldn't do that with Storage Spaces—until now.

Storage Spaces in Windows Server 2012 R2

As Figure 3-8 shows, the Windows Server 2012 R2 version of Storage Spaces now allows you to create a tiered storage solution that transparently delivers an appropriate balance between capacity and performance that can meet the needs of enterprise workloads. The result is that the workload's most frequently accessed data (the working set) will automatically be stored on the SSD tier while the rest of the workload's data is stored on the HDD tier.

How does Storage Spaces accomplish this? By having the file system actively measure the activity of the workload in the background and then automatically and transparently move data to the appropriate tier (SSD or HDD) depending on how hot or cold the data is determined to be. Storage Spaces in Windows Server 2012 R2 can thus ensure that the workload's hot data is always stored on the SSD tier to take advantage of the high performance of this tier, and its cold data on the HDD tier to make use of the high capacity of this tier. If a portion of the data for a particular file becomes hotter (is accessed more frequently), then it gets moved from the HDD tier to the SSD tier. And if the portion of data becomes cooler (is accessed less frequently), then it gets moved from the SSD tier to the HDD tier.

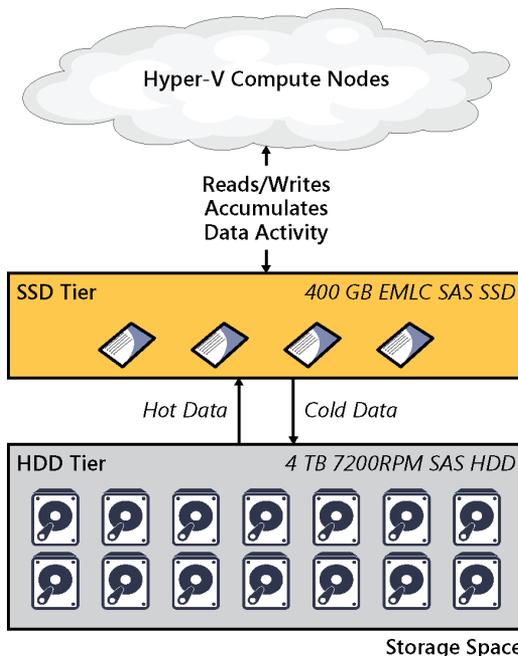


FIGURE 3-8 Storage Spaces in Windows Server 2012 R2 now supports data tiering.

This seamless movement of data between tiers is configured by default to happen daily in 1 MB chunks, but you also have the option of configuring the scheduled task for this operation to run as frequently as you want. Data moves between tiers in the background and has minimal impact on the performance of the storage space. If needed, you can use Windows PowerShell to assign certain files to a specific tier, thereby overriding the automatic placement of data based on heat. For example, the parent virtual hard disk file for a collection of pooled virtual machines in a VDI environment might be assigned to the SSD tier to ensure the file always remains pinned to this tier. The result of doing this can be to achieve significant improvements in the boot times of the hundreds or thousands of virtual desktops derived from this parent virtual hard disk.

Tiered Storage Spaces is an exciting new capability in Windows Server 2012 R2, and when this book is revised and enlarged in the RTM timeframe, we'll take a closer look at how it can be implemented. But there's one additional new feature of Storage Spaces that we'll briefly describe here, and it's called write-back caching. Whereas the goal of tiering is to balance capacity against performance, the purpose of write-back caching is to smooth out short-term bursts of random writes. Write-back caching integrates seamlessly into tiered volumes and is enabled by default. The write-back cache is located on the SSD tier of a storage space and services smaller, random writes; larger, sequential writes are serviced by the HDD tier. You can also enable write-back caching on non-tiered volumes.

Now let's hear from one of our insiders at Microsoft to learn more about the Storage Spaces enhancements in Windows Server 2012 R2 and watch these new features in action.

Storage Spaces in action

Let's examine the new features of Storage Spaces in Windows Server 2012 R2 and see some of them in action.

Storage tiering

It has been observed that SSDs with 15k speed are expensive and low capacity, but with some good performance and poor \$/GB rate. On the other hand, 7.5K drives can give much better capacity at a great \$/GB rate, but they're slow. We wanted to give our customers a new way to balance scalability and performance, and we have added tiered storage functionality in Windows Server 2012 R2 Storage Spaces.

Using storage tiering, the most active data (hot data) can be placed on SSD and less active data (cold data) will be placed on normal HDD, and hence it provides great throughput for I/O operations.

Tiered storage helps maximize IOPS/\$ because it:

- Utilizes best characteristics of SSDs and HDDs in single storage
- Provides balance between capacity and performance
- Allows admins to assign files to specific storage

When creating a virtual disk, we can assign some capacity from the fast tier and slow tier as well. Windows will automatically analyze the data on your virtual disk to identify hot spots (Most Active Data) and cold spots (Less Active Data) and move them into fast and slow tiers, respectively. This capability is very useful in multiple scenarios and provides maximum throughput. For example:

- Assume we are using Storage Spaces for storing VDI virtual machines. The parent VHDX will be pinned to the fast tier, while the child differential disks live in the slow tier.
- During the boot process, SSD provides the good IOPS bandwidth to reduce logon time.

Write-back cache

The write-back caching feature creates a separate write-back cache for each accelerated volume from available SSDs in the Storage Spaces pool. There are many scenarios when we observe spikes in disk I/O and these can reduce write performance. To overcome this, a virtual disk leverages tiered storage to absorb the spike in write activity. Whenever a spike is detected, the SSDs in the fast tier are used, which offers greater IOPS and can take the pressure that the normal HDD tier cannot.

Dual parity and parity space support for failover clusters

This provides resiliency to **two-disk** failures for archival workloads, while minimizing \$/TB and space utilization for resiliency. It also provides:

- Efficient rebuild times
- Windows Clustering support
- Integrated journaling for power fault resiliency

Faster storage space rebuild

A storage space and a RAID array can be configured with the traditional hot spare (one or more). When a disk fails, the other available disk is used to replace it. This will offer more IOPS to normal production storage activity and leave your business less vulnerable during a repair process.

Let's watch this in action. The following is a demo showcasing the dual parity capability and the disk repair process for maximizing uptime. We'll start with a Windows Server 2012 R2 that has seven SSDs of different capacities totaling 230 GB that is direct-attached via SAS to the server:

PHYSICAL DISKS
Primordial on SRV01

Filter

Name	Capaci...	Bus	Usage	Media T
PhysicalDisk1 (SRV01)	20.0 GB	SAS	Automatic	SSD
PhysicalDisk2 (SRV01)	20.0 GB	SAS	Automatic	SSD
PhysicalDisk3 (SRV01)	30.0 GB	SAS	Automatic	SSD
PhysicalDisk4 (SRV01)	30.0 GB	SAS	Automatic	SSD
PhysicalDisk5 (SRV01)	30.0 GB	SAS	Automatic	SSD
PhysicalDisk6 (SRV01)	50.0 GB	SAS	Automatic	SSD
PhysicalDisk7 (SRV01)	50.0 GB	SAS	Automatic	SSD

As you can see, there is only the primordial storage pool:

STORAGE POOLS
All storage pools | 1 total

Filter

Name	Type	Version	Managed by	Available to	Read-Write Server	Capacity	Free Space
Storage Spaces (1)							
Primordial	Available Di...	Windows Server 2012 R2	SRV01	SRV01	SRV01		

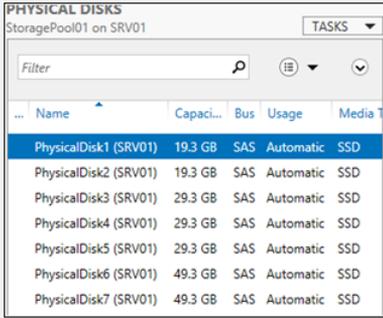
Now let's use all seven physical disks in the primordial pool to create a new storage pool named StoragePool01:

STORAGE POOLS
All storage pools | 1 total

Filter

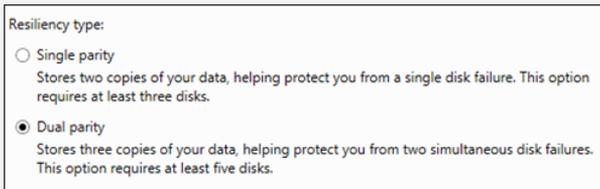
Name	Type	Version	Managed by	Available to	Read-Write Server	Capacity	Free Space
Storage Spaces (1)							
StoragePool01	Storage Pool	Windows Server 2012 R2	SRV01	SRV01	SRV01	225 GB	223 GB

Note that the total capacity of the StoragePool01 is 225 GB but the free space is only 223 GB because some disk space is reserved by Windows for internal purposes. The capacity of the seven physical disks changes to reflect their addition to the new pool:



Name	Capacity	Bus	Usage	Media
PhysicalDisk1 (SRV01)	19.3 GB	SAS	Automatic	SSD
PhysicalDisk2 (SRV01)	19.3 GB	SAS	Automatic	SSD
PhysicalDisk3 (SRV01)	29.3 GB	SAS	Automatic	SSD
PhysicalDisk4 (SRV01)	29.3 GB	SAS	Automatic	SSD
PhysicalDisk5 (SRV01)	29.3 GB	SAS	Automatic	SSD
PhysicalDisk6 (SRV01)	49.3 GB	SAS	Automatic	SSD
PhysicalDisk7 (SRV01)	49.3 GB	SAS	Automatic	SSD

Now let's create a new virtual disk with dual parity mode:

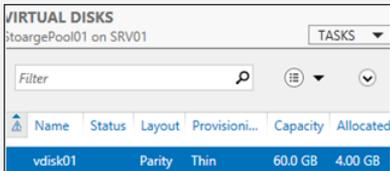


Resiliency type:

- Single parity
Stores two copies of your data, helping protect you from a single disk failure. This option requires at least three disks.
- Dual parity
Stores three copies of your data, helping protect you from two simultaneous disk failures. This option requires at least five disks.

As you can see, dual parity mode stores three copies of your data and can protect against simultaneous failure of up to two physical disks in the underlying pool.

Here's a screenshot showing the newly created virtual disk named vdisk01:



Name	Status	Layout	Provisioning	Capacity	Allocated
vdisk01	Parity	Thin	60.0 GB	4.00 GB	

The Health Status of the new virtual disk is Healthy while its Operational Status is OK. All seven disks are in use:

Health

Health status: Healthy
Operational status: OK
Storage layout: Parity

Physical disks in use:

Slot	Name	Status	Capacity	Free Space	Bus	Chassis	Media Type
	PhysicalDisk1	OK	19.3 GB	17.0 GB	SAS		SSD
	PhysicalDisk7	OK	49.3 GB	47.0 GB	SAS		SSD
	PhysicalDisk6	OK	49.3 GB	47.0 GB	SAS		SSD
	PhysicalDisk5	OK	29.3 GB	28.0 GB	SAS		SSD
	PhysicalDisk2	OK	19.3 GB	17.0 GB	SAS		SSD
	PhysicalDisk4	OK	29.3 GB	27.0 GB	SAS		SSD
	PhysicalDisk3	OK	29.3 GB	27.0 GB	SAS		SSD

After we created the new virtual disk, we created a new volume E: on it as shown here:

Devices and drives (4)

Floppy Disk Drive (A:)	Floppy Disk Drive		
Local Disk (C:)	Local Disk	29.6 GB	18.3 GB
DVD Drive (D:)	CD Drive		
New Volume (E:)	Local Disk	59.8 GB	59.7 GB

Now let's check out the dual parity capability. First, let's detach two physical disks (disks 1 and 2) from the server to simulate the simultaneous failure of two physical disks. When we do this, Storage Spaces shows a warning message for the pool:

STORAGE POOLS

All storage pools | 1 total

Filter

Name	Type	Version	Managed by	Available to	Read-Write Server	Capacity	Free Space
Storage Spaces (1)							
StoragePool01	Storage Pool	Windows Server 2012 R2	SRV01	SRV01	SRV01	225 GB	210 GB

In addition, the two detached physical disks display Unknown for their Bus type:

PHYSICAL DISKS

StoragePool01 on SRV01

Filter

Name	Capacity	Bus	Usage
PhysicalDisk-1 (SRV01)	19.3 GB	Unknown	Automatic
PhysicalDisk-1 (SRV01)	19.3 GB	Unknown	Automatic
PhysicalDisk3 (SRV01)	29.3 GB	SAS	Automatic
PhysicalDisk4 (SRV01)	29.3 GB	SAS	Automatic
PhysicalDisk5 (SRV01)	29.3 GB	SAS	Automatic
PhysicalDisk6 (SRV01)	49.3 GB	SAS	Automatic
PhysicalDisk7 (SRV01)	49.3 GB	SAS	Automatic

The virtual disk also shows a warning message:

Name	Status	Layout	Provisioni...	Capacity	Allocated
vdisk01	Warning	Parity	Thin	60.0 GB	4.00 GB

The Health Status of the virtual disk has changed from Healthy to Warning, and the Operational Status has changed from OK to Degraded:

Health status:	Warning
Operational status:	Degraded
Storage layout:	Parity

However, the volume E: that we created from the virtual disk is still accessible:

Device	Type	Capacity	Allocated
Floppy Disk Drive (A:)	Floppy Disk Drive		
Local Disk (C:)	Local Disk	29.6 GB	18.3 GB
DVD Drive (D:)	CD Drive		
New Volume (E:)	Local Disk	59.8 GB	59.7 GB

As you can see, the dual parity capability of Storage Spaces in Windows Server 2012 R2 can compensate for two-disk failures—the volume created previously is in good status for I/O operations.

Now let's see what happens when we remove a third physical disk from the server to simulate the simultaneous failure of three physical disks. When this is done, the virtual disk is now in a failed state since dual parity only tolerates up to two simultaneous disk failures:

Name	Status	Layout	Provisioni...	Capacity	Allocated
vdisk01	Failed	Parity	Thin	60.0 GB	4.00 GB

The Health Status for the virtual disk now changes to Unhealthy, and its Operational Status becomes Detached:

Health	
Health status:	Unhealthy
Operational status:	Detached
Storage layout:	Parity

If we had additional disks in the pool and wanted Storage Spaces to automatically compensate for disk failures, these additional disks could be configured in Hot Spare mode as per recommended practices. When this is done, a spare disk will switch to Automatic mode whenever a disk failure is detected to ensure maximum uptime. You can change the existing usage mode of a physical disk to Hot Spare mode either using the Storage Spaces canvas of Server Manager or by using Windows PowerShell as shown here:

```
PS C:\Users\Administrator> Set-PhysicalDisk Physicaldisk1 -Usage hotspare
PS C:\Users\Administrator> Get-PhysicalDisk Physicaldisk1
```

FriendlyName	CanPool	OperationalStatus	HealthStatus	Usage	Size
PhysicalDisk1	False	OK	Healthy	Hot Spare	29.25 GB

Conclusion

Storage Spaces in Windows Server 2012 supports mirrored storage and single parity, which allows single disk failure, but with Windows Server 2012 R2 two disk failures are supported if you configure the inbuilt dual parity mechanism. Additionally, Windows Server 2012 R2 introduces a parallelized repair process, where the remaining healthy disks move around and take responsibility of data that was stored on the failed disk. It is extremely quick because all disks in the storage space are involved. Hence, Storage Spaces delivers sophisticated storage virtualization capabilities, which empower customers to use industry-standard storage for single computer and scalable multinode deployment.

Deepak Srivastava
Partner Consultant, Microsoft

Work Folders

Work Folders is another new File and Storage Services feature introduced in Windows Server 2012 R2 that can help organizations address the bring your own device (BYOD) scenario that so many of them are currently facing as the modern workplace evolves. Work Folders provides a consistent way for users to access their work files from their personal computers and devices anywhere from centrally managed file servers on the corporate network. Work Folders can also be deployed alongside existing deployments of other Microsoft data synchronization solutions such as Folder Redirection (FR), Offline Files (also known as client-side caching, or CSC), and even home folders.

To learn more about how Work Folders works, how you set it up, and how you use it, let's now hear from some insiders on the team who developed this exciting new feature.

Understanding and implementing Work Folders

During Windows Server 2012 R2 planning, we noticed two converging trends around managing and protecting corporate data:

- Users “I need to work from anywhere on my different devices”
- IT “I'd like to empower my information workers (users) while reducing information leakage and keeping control of the corporate data that is sprawled across devices”

Work Folders is a new file server-based sync solution in Windows Server 2012 R2 and Windows 8.1 that is designed to address these two trends. Work Folders enables IT administrators to provide information workers the ability to sync their work data on all their devices wherever they are while remaining in compliance with company policies. This is done by syncing user data from devices to on-premises file servers, which are now extended to include a new sync protocol.

Work Folders as experienced by an information worker

To show how this works, here's an example of how an information worker, Joe, might use Work Folders to separate his work data from his personal data while having the ability to work from any device. When Joe saves a document on his work computer in the Work Folders directory, the document is synced to an IT-controlled file server. When Joe returns home, he can pick up his Surface RT (where the document is already synced) and head to the beach. He can work on the document offline, and when he returns home the document is synced back with the file server and all of the changes are available to him the next day when he returns to the office.

Look familiar? Indeed, this is how consumer storage services such as SkyDrive and business collaboration services such as SkyDrive Pro work. We kept the user

interaction simple and familiar so that there is little user education required. The biggest difference from SkyDrive or SkyDrive Pro is that the centralized storage for Work Folders is an on-premises file server running Windows Server 2012 R2, but we'll get to that a little later in this post.

Work Folders as experienced by an administrator

IT administrators can use Work Folders to gain more control over corporate data and user devices and centralize user work data so that they can apply the appropriate processes and tools to keep their company in compliance. This can range from simply having a copy of the data if the user leaves the company to a wide range of capabilities such as backup, retention, classification, and automated encryption.

For example, when a user authors a sensitive document in Work Folders on his or her work PC, it gets synced to the file server. The file server then can automatically classify the document based on content, if configured using File Server Resource Manager, and encrypt the document using Windows Rights Management Services before syncing the document back to all the user's devices. This allows a seamless experience for the user while keeping the organization in compliance and preventing leakage of sensitive information.

Capabilities of Work Folders

Our main design focus around Work Folders was to keep it simple for the information workers while allowing IT administrators to use the familiar **low-cost, high-scale** Windows file server with all the rich functionality available on the back end from high availability to comprehensive data management. Some of the functionality provided by Work Folders includes the following:

- It provides a single point of access to work files on a user's work and personal PCs and devices (Windows 8.1 and Windows RT 8.1, with immediate plans to follow up with Windows 7 and iPad support and other devices likely in the future).
- It enables users to access work files while offline and sync with the central file server when the PC or device next has Internet or network connectivity.
- It helps maintain data encryption in transit, as well as at rest, on devices and allow corporate data wipe through device management services such as Windows Intune.
- It uses existing file server management technologies, such as file classification and folder quotas, to manage user data.
- You can configure security policies to instruct user PCs and devices to encrypt Work Folders and use a lock screen password, for example.
- You can implement Failover Clustering together with Work Folders to provide a high-availability solution.

Considerations for Work Folders

When planning for the implementation of Work Folders, you need to be aware of the following considerations:

- Backend storage is provided by on-premises file servers and Work Folders must be stored in local storage on the file server (for example, data can be on local shares on a Windows Server 2012 R2).
- Users sync to their own folder on the file server—there is no support for syncing arbitrary file shares (for example, sync the sales demos share to my device).
- Work Folders doesn't provide collaboration functionality such as sharing sync files or folders with other users (we recommend using SkyDrive Pro if you need document collaboration features).

Comparison between Work Folders and other data sync solutions

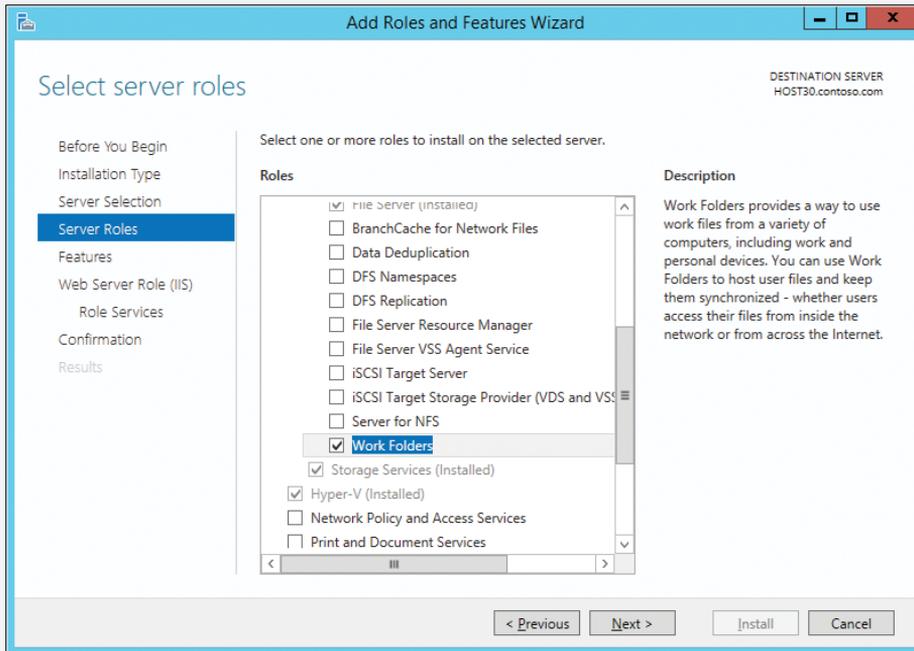
The table below compares the different underlying technologies used and intended usage scenarios for Work Folders, FR/CSC, Microsoft SkyDrive, and Microsoft SkyDrive Pro.

Sync solution	Data location	Access protocol	Intended usage scenarios	Support for personal devices
Work Folders	File server	HTTPS	The user's work data	Yes
FR/CSC	File server	SMB (on-premises or VPN)	The user's work data	No
SkyDrive	Public cloud	HTTPS	The user's personal data	Yes
SkyDrive Pro	SharePoint Online in Office 365 or on-premises SharePoint 2013	HTTPS	The user's work data and team work data	Yes

Setting up Work Folders

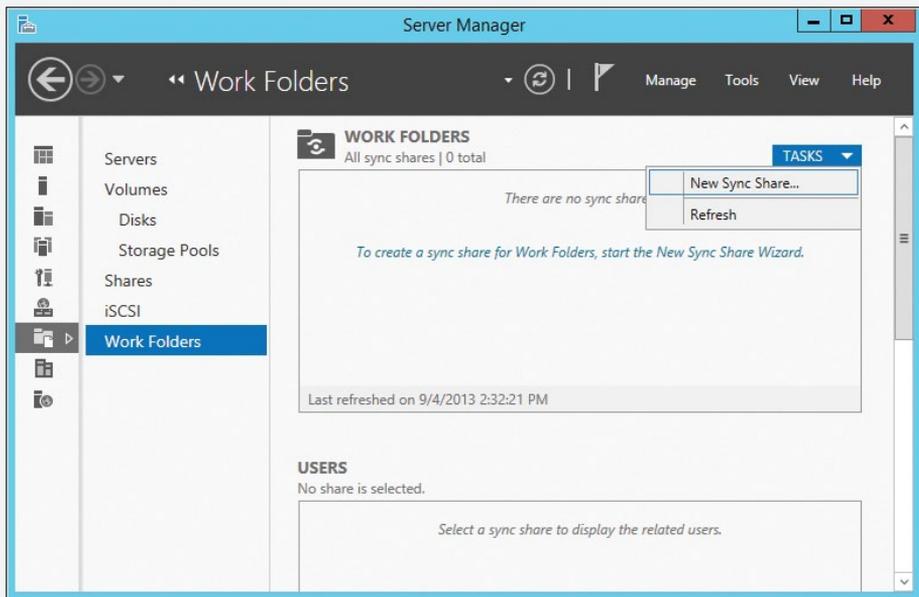
The following is a brief overview of how to set up and use Work Folders. Note that additional steps are required to enable the client computer to be able to retrieve Work Folders settings and connect to the server across the Internet. For more detailed information see the Work Folders Test Lab Deployment Guide referenced in the Learn More section at the end of this chapter.

Work Folders requires an Active Directory environment. Start by installing the Work Folders role service on a file server running Windows Server 2012 R2. The Work Folders role service is part of the File and iSCSI Services role service for the file and Storage Services role:

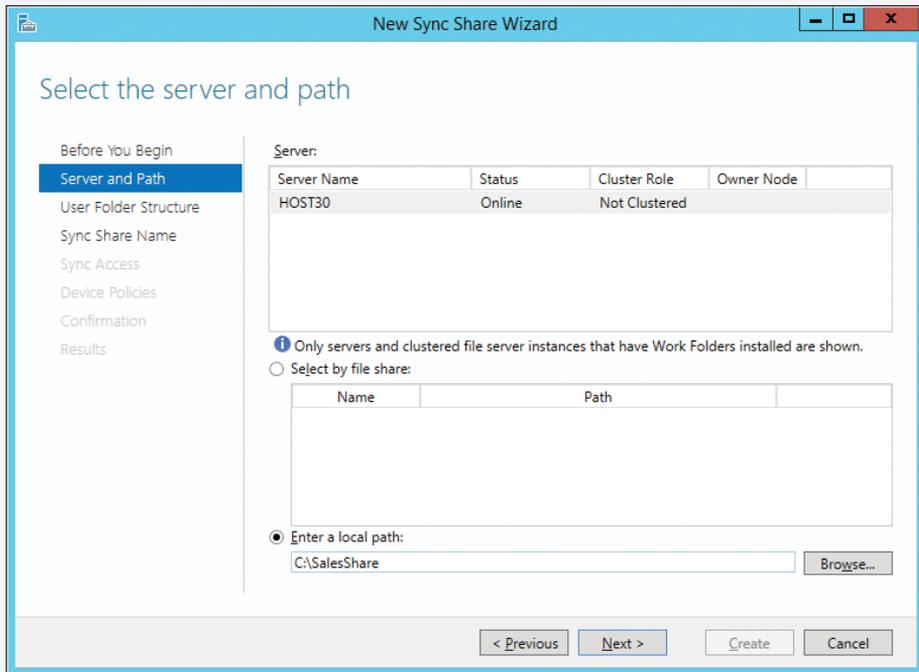


Installing the Work Folders role service will also install some components of the IIS (Web Server) role.

A new Work Folders page will be added to the File and Storage Services canvas in Server Manager. To configure Work Folders, you first need to create a sync share, which maps to a local path where all the user folders will be hosted under and to the group of users who can access the sync share:



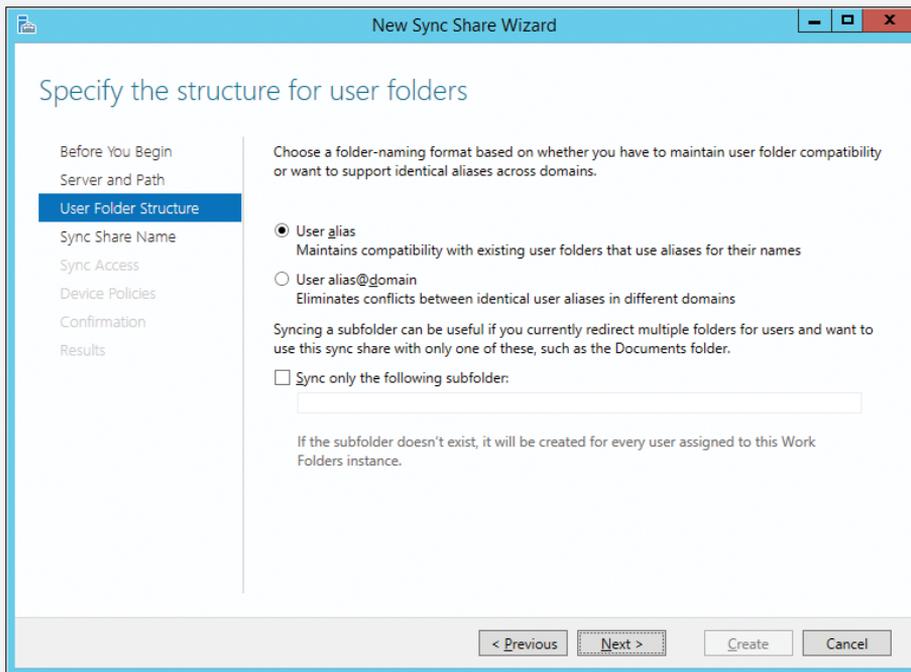
Creating a sync share simply allows the user to access the data hosted on the file server through the Sync protocol. Selecting the New Sync Share task, as shown previously, launches the New Sync Share Wizard:



If you already have a local path that has been configured as an SMB share, such as a Folder Redirection share, you can simply select the first option, Select By File Share. In this walkthrough, we have selected the Enter A Local Path option to create a new SMB share for the folder C:\SalesShare on the file server.

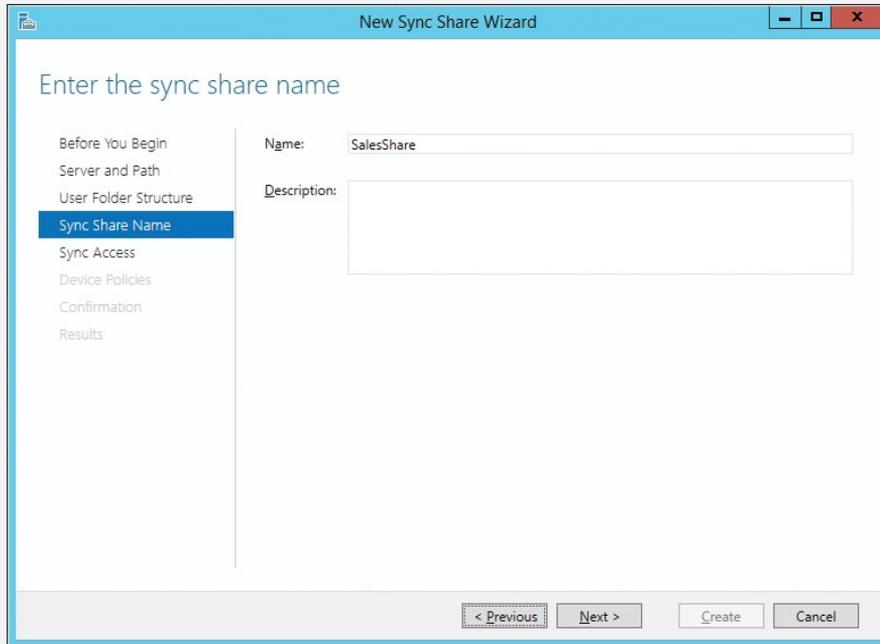
The next page of the wizard allows us to specify the structure for user folders. There are two options here:

- **Using user alias** This option is selected by default and is compatible with other technologies such as Folder Redirection or home folders. For example, the folder for Karen Berg will be named after her alias kberg.
- **Using a combination of user alias and domain information** This option ensures the uniqueness of the folder name for users across domains. For example, let's say your organization has two users named Karen Berg, one in the hq.contoso.com domain and the other in the research.contoso.com domain. The folder for the first Karen will be named kberg@hq.contoso.com and the folder for the second Karen will be named kberg@research.contoso.com, thus ensuring each Karen has her own unique folder.

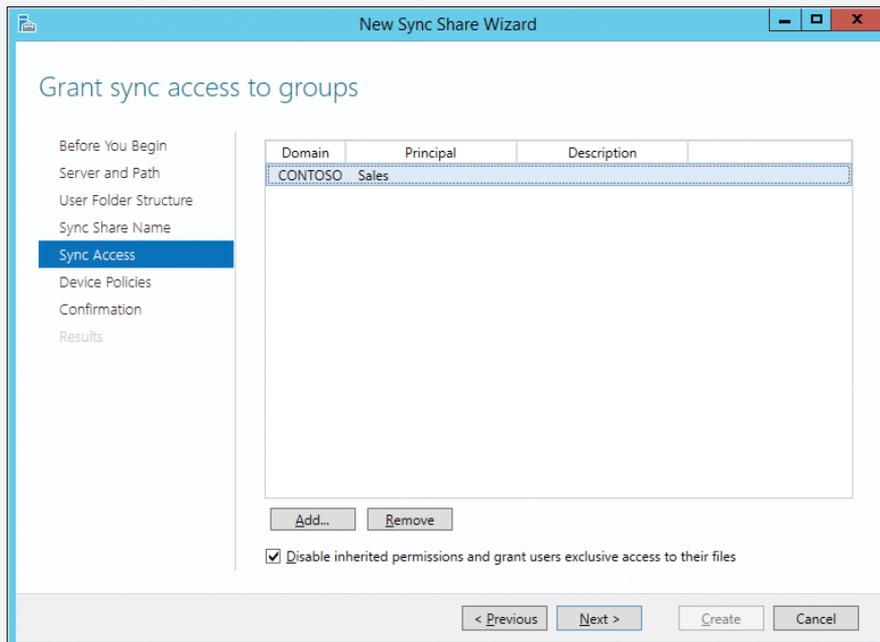


By default, all the folders/files under the user folder will be synced to the devices. The **check box** allows the admin to specify a single subfolder to be synced to the devices.

The next page allows you to specify the sync share name and an optional description:

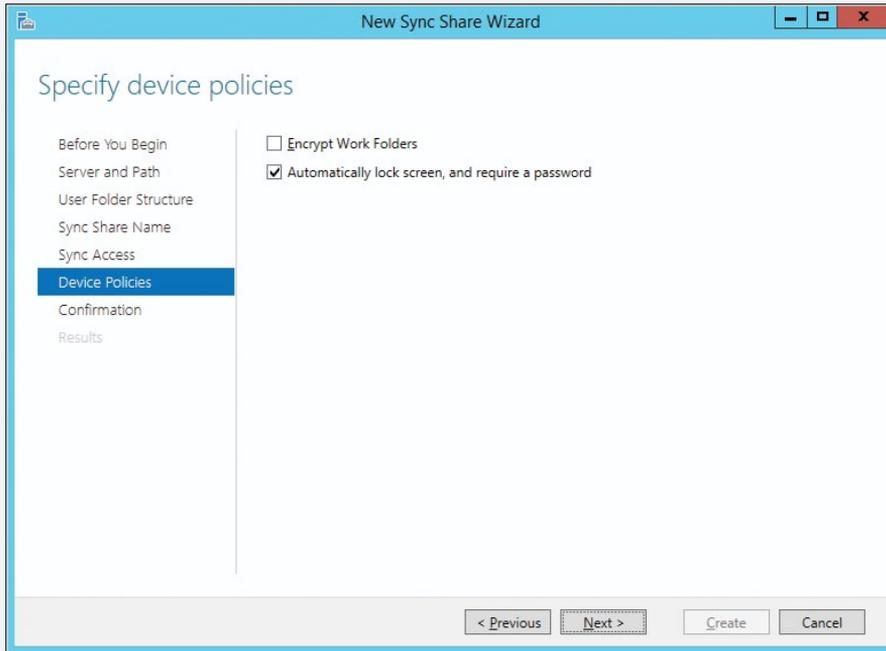


The next page is where you can assign specific security groups for sync share access:



Note that, by default, administrators will not be able to access user data on the file server. If you want to enable administrator access to user data, clear **Disable Inherited Permissions And Grant Users Exclusive Access To Their Files**.

The next page is where you can configure device policies:



Selecting the **Encrypt Work Folders Policy** will ensure the document under the work folders on the client devices is encrypted with the Enterprise ID. By default, the Enterprise ID is the user primary SMTP email address (which is also the proxyAddresses attribute of the user object in Active Directory). Encryption using a different key ensures that deleting the user's Work Folder documents will not delete the user's personal documents on the same device.

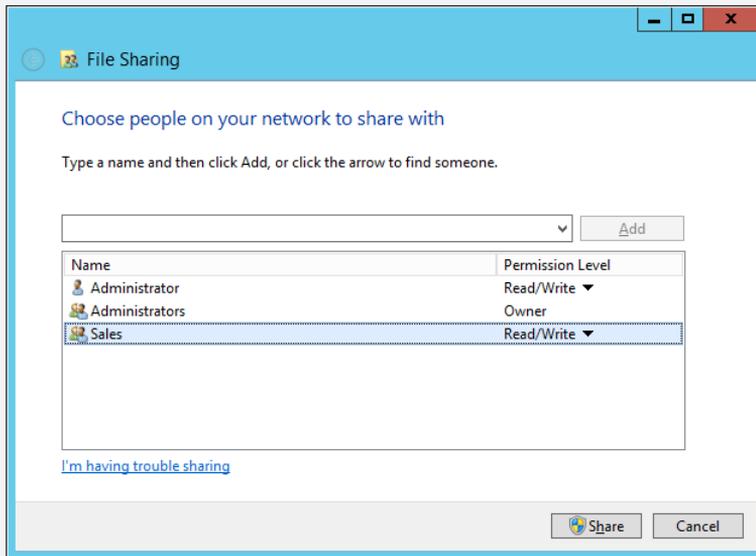
Selecting the **Automatically Lock Screen And Require A Password Policy** (this is selected by default) enforces the following policy on the user's device:

- Minimum password length of 6
- Autolock screen set to be 15 minutes or less
- Maximum password retry of 10 or less

If the user's device doesn't meet the configured policies, the user will not be able to set up Work Folders.

Once you've completed the wizard, we can enable the sync share for SMB access by right-clicking the SalesShare folder in Windows Explorer and selecting **Share With**,

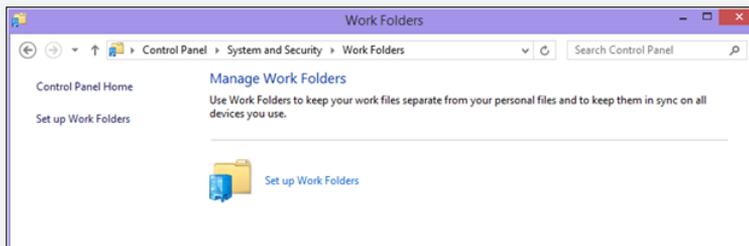
and then Specific People. Then in the File Sharing dialog, add the Sales group and give it Read/Write access, as shown here, and then click Share:



All of the above can also be done of course with Windows PowerShell commands, if desired.

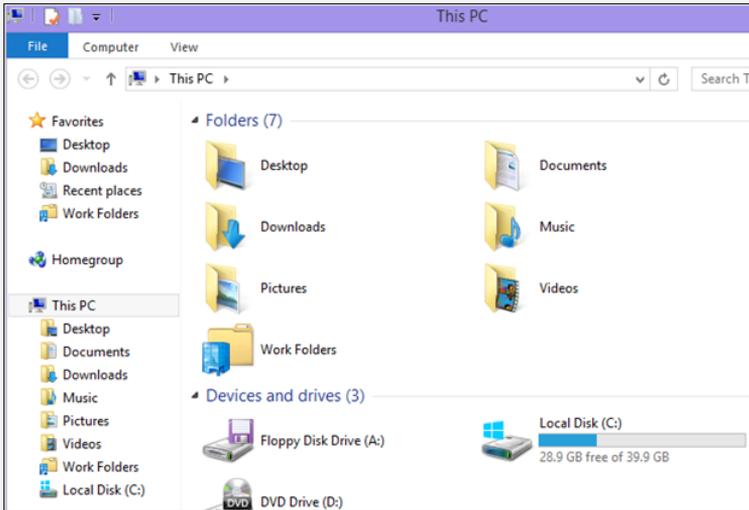
By default, the Work Folders client always **connects** to the file server using SSL, which requires the server to have an SSL certificate installed and configured.

Once Work Folders has been set up on the file server, the user can now set up Work Folders on client computers and devices. For example, on a computer running Windows 8.1 the user would open Control Panel, select System and Security, and then select Work Folders. Doing this displays the Manage Work Folders item as shown here:

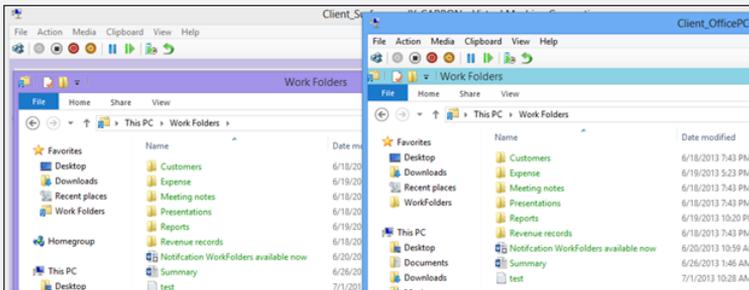


The user then provides his or her email address and domain credentials, accepts or changes where the Work Folders will be stored (by default in the user profile), and consents to the configured Work Folders policies being applied.

At this point a new icon labeled Work Folders will be displayed under This PC in Windows Explorer on the user's computer:



At this point Work Folders has been configured on both the server and client sides and is ready for use. To test Work Folders, the user could use two client PCs running Windows 8.1. The user begins by using Notepad to create a new text file named "test" on the first PC and saves the file to the Work Folders location (which, as the previous screenshot shows, has been added under Favorites in Windows Explorer). When the file is saved, Work Folders automatically syncs the file to the Work Folders location on the second PC as shown here:



For a more detailed demonstration of how to configure Work Folders on both the server and client sides, and of how the user can store and access data using Work Folders, see the Work Folders Test Lab Deployment guide, which is referenced in the "Learn more" section at the end of this chapter.

Nir Ben Zvi
Principal Program Manager, WSSC WS PM USA

Learn more

You can learn more about the new storage features and enhancements in Windows Server 2012 R2 by checking out the following topics on Microsoft TechNet:

- "New and changed functionality in File and Storage Services" at http://technet.microsoft.com/en-us/library/4cb00829-8d05-4499-8adc-7506e159f857#BKMK_NEW.
- "What's New for iSCSI Target Server in Windows Server 2012 R2" at <http://technet.microsoft.com/en-us/library/dn305893.aspx>.
- "What's New for SMB in Windows Server 2012 R2" at <http://technet.microsoft.com/en-us/library/hh831474.aspx>.
- "Storage Quality of Service for Hyper-V" at <http://technet.microsoft.com/en-us/library/dn282281.aspx>.
- "Storage Spaces Overview" at <http://technet.microsoft.com/en-us/library/hh831739.aspx>.
- "Work Folders Overview" at <http://technet.microsoft.com/en-us/library/dn265974.aspx>.

The following Microsoft TechNet and MSDN blog posts also have lots of information about the new storage features and enhancements in Windows Server 2012 R2:

- "Storage Transformation for your Datacenter" at <http://blogs.technet.com/b/windowsserver/archive/2013/06/26/storage-transformation-for-your-datacenter.aspx>.
- "What's New for SMI-S in Windows Server 2012 R2" at <http://blogs.technet.com/b/filecab/archive/2013/07/31/what-s-new-for-smi-s-in-windows-server-2012-r2.aspx>.
- "Introducing Work Folders on Windows Server 2012 R2" at <http://blogs.technet.com/b/filecab/archive/2013/07/09/introducing-work-folders-on-windows-server-2012-r2.aspx>.
- "Work Folders Test Lab Deployment" at <http://blogs.technet.com/b/filecab/archive/2013/07/10/work-folders-test-lab-deployment.aspx>.
- "Work Folders Certificate Management" at <http://blogs.technet.com/b/filecab/archive/2013/08/09/work-folders-certificate-management.aspx>.
- "Extending Data Deduplication to New Workloads in Windows Server 2012 R2" at <http://blogs.technet.com/b/filecab/archive/2013/07/31/extending-data-deduplication-to-new-workloads-in-windows-server-2012-r2.aspx>.
- "Deploying Data Deduplication for VDI storage in Windows Server 2012 R2" at <http://blogs.technet.com/b/filecab/archive/2013/07/31/deploying-data-deduplication-for-vdi-storage-in-windows-server-2012-r2.aspx>.
- "Storage and File Services PowerShell Cmdlets Quick Reference Card for Windows Server 2012 R2 [Preview Edition]" at <http://blogs.technet.com/b/filecab/archive/2013/07/30/storage-and-file-services-powershell-cmdlets-quick-reference-card-for-windows-server-2012-r2-preview-edition.aspx>.

- "Windows Server 2012 R2 Storage: Step-by-step with Storage Spaces, SMB Scale-Out and Shared VHDX (Physical)" at <http://blogs.technet.com/b/josebda/archive/2013/07/31/windows-server-2012-r2-storage-step-by-step-with-storage-spaces-smb-scale-out-and-shared-vhdx-physical.aspx>.
- "Step-by-step for Storage Spaces Tiering in Windows Server 2012 R2" at <http://blogs.technet.com/b/josebda/archive/2013/08/28/step-by-step-for-storage-spaces-tiering-in-windows-server-2012-r2.aspx>.
- "iSCSI Target Server in Windows Server 2012 R2" at <http://blogs.technet.com/b/filecab/archive/2013/07/31/iscsi-target-server-in-windows-server-2012-r2.aspx>.

Be sure also to check out the following videos from TechEd 2013 on Channel 9:

- "Storage Changes in Windows Server 2012 R2" at <http://channel9.msdn.com/Events/TechEd/NorthAmerica/2013/C9-14>.
- "Reduce Storage Costs with Data Deduplication" at <http://channel9.msdn.com/Events/TechEd/NorthAmerica/2013/MDC-B342>.
- "Storage Spaces: What's New in Windows Server 2012 R2" at <http://channel9.msdn.com/Events/TechEd/NorthAmerica/2013/MDC-B218>.
- "Storage and Availability Improvements in Windows Server 2012 R2" at <http://channel9.msdn.com/Events/TechEd/NorthAmerica/2013/MDC-B333>.

Index

Symbols and Numbers

.wim files, 186–187

A

access control

- multi-factor, 127
- role-based, 108–110

activation

- dynamic site, 155–156

activation, virtual machine, 11–12

Active Directory

- detached clustering, 82
- enhancements in Windows Server 2012, 124
- Kerberos authentication, 146–148
- LDAP search performance, 131–135
- multi-factor access control, 127
- Web Application Proxy, 127–131
- Windows Azure, 206
- Work Folders in, 53
- Workplace Join, 125–126

Active Directory Administrative Center (ADAC), 124

Active Directory Certificate Services (AD CS), 123

Active Directory Domain Services (AD DS),
112, 123, 131–135

Active Directory Federation Services (AD FS),
123, 127, 130

Active Directory Lightweight Directory Services
(AD LDS), 123

Active Directory Rights Management Services
(AD RMS), 123

AD CS (Active Directory Certificate Services), 123

AD DS (Active Directory Domain Services),
112, 123, 131–135

AD FS (Active Directory Federation Services),
123, 127, 130

AD LDS (Active Directory Lightweight Directory
Services), 123

AD RMS (Active Directory Rights Management
Services), 123

ADAC (Active Directory Administrative Center), 124

Add-DhcpServerSecurityGroup cmdlet, 190

address assignment, 114

address filtering, 154

Address Hash, 93, 95–96

address management. *See* IP Address
Management (IPAM)

address translation, 196

Add-WindowsImage cmdlet, 186–187

administrators

- access to user files, 57
- use of Work Folders, 51

aliases, 55

Anderson, Jason, 214

Application Initialization, 154, 156

application pools, 154, 157

apps

- group policy settings, 144–145, 150
- My Server, 209–212
- publishing, 130–131
- Start Screen view, 143
- Window Store, 150

architecture

- extensible switch, 119–120
- Hyper-V extensible switch, 115–117
- IP Address Management (IPAM), 109
- private cloud solution, 30–31

authentication. *See also* preauthentication
compound, 146–147

device, 125, 127

multi-factor, 127

automatic setup, of Work Folders, 145–146
Automatic Virtual Machine Activation (AVMA), 11–12
availability, of virtual machines, 78–80

B

backgrounds, start screen, 143
backups, 20, 24, 207
bandwidth management, 37–38, 86, 96, 120
Benari, Erez, 94
block-based storage, 28–29, 34, 40, 68
booting
 fast, 137
 Generation 2 virtual machines, 7
 secure, 7–8, 10
BranchCache, 87, 204
business needs and cloud computing, 1–3
Butte, Jeff, 186

C

CA (Customer Addresses), 114, 116
cache
 allocation, 73
 write-back, 33, 43
capacity versus performance, 41–42
capturing extensions, 116, 119–120
CAU (Cluster-Aware Updating), 64
Chkdsk, 28
Client Full System Restore, 213
cloud computing
 Microsoft’s vision for, 1, 29–31
 platform synchronization, 1–3
 Windows Server Essentials 2012 R2, 204–205
cloud service integration, 205
cluster nodes
 draining, 78–80
 heartbeat thresholds, 75
 transparent failover, 36
 updates, 64
Cluster Shared Volumes (CSV)
 and deduplication, 38
 failover clustering and, 71–73
 for block storage, 68
 overview, 63
 ownership of disks, 71–72
 resiliency enhancements, 72–73

Cluster-Aware Updating (CAU), 64
cold data, 41–44
cold requests, 156–157
color, start screen accent, 143
communication, IP Address Management (IPAM), 110
compression, 16–17
compute layer, in cloud computing solutions, 29
configuration drift, 177–178
Configuration Instance Document script, 178–179
Configuration Manager, 178
configurations
 network, 100–101
 Work Folders, 146
configuring
 ETW logging, 159
 virtual machine health protection, 76–78
 Web Application Proxy, 128–129
connectivity
 testing, 102–105, 191
 virtual machine health protection, 76–78
Connector software, 213
Content protection, 149–150
Copy-VMFile cmdlet, 192
CPU throttling, 154
credentials
 DNS, 189
 restriction of delegation, 148
CSV (Cluster Shared Volumes)
 and deduplication, 38
 failover clustering and, 71–73
 for block storage, 68
 overview, 63
 ownership of disks, 71–72
 resiliency enhancements, 72–73
Customer Addresses (CA), 114, 116

D

Dashboard, 205–207
Data Center Bridging (DCB), 86
data center management, 177–183
data collection, 110–111
data protection, 213
data tiering, 42–44
Datacenter edition, 11–12
Debug-VM cmdlet, 192–195
deduplication, 28, 38–39, 73, 172

- delegation, credential, 148
- Deployment Image Servicing and Management (DISM)
 - cmdlets, 186–189
- Desired State Configuration (DSC), 176–183
- Desired State Configuration (DSC) Pull Server, 178–179, 181, 183
- Device Manager, 9–10
- device policies, 57
- devices. *See* mobile device management; personal devices
- DHCP (Dynamic Host Configuration Protocol), 86, 105–106, 120, 189–190
- disaster recovery, 21–24
- disconnection, 75, 78, 175
- DISM (Deployment Image Servicing and Management)
 - cmdlets, 186–189
- displays
 - handling, 169
 - multi-monitor, 143–144
- distribution groups, 206
- distros, 20
- DNS credentials management, 189
- domain controllers, 124
 - search performance, 132–135
- Domain Name System (DNS), 87, 105–106, 190
- drain on shutdown capability, 78–80
- DSC. *See* Desired State Configuration (DSC)
- dual parity, 44–49
- DX11.1, 168
- Dynamic Host Configuration Protocol (DHCP), 86, 105–106, 120, 189–190
- Dynamic Memory, 6, 20
- Dynamic Site Activation, 155–156
- Dynamic Teaming. *See* Link Aggregation Control Protocol (LACP)
- Dynamic Witness, 81

E

- email distribution groups, 206
- Emulex, 115
- enabling
 - ETW logging, 159–160
 - Idle Worker Process Page-Out, 157
 - Protected Network on Hyper V host clusters, 75–76
 - replication frequencies, 22

- Enterprise ID, 149–150
- event tracing, 120, 161–162, 197
- Event Tracing for Windows, 159–160
- Expand-WindowsImage cmdlet, 186–187
- Export-StartLayout cmdlet, 198–199
- Export-WindowsImage cmdlet, 187
- extensible virtual switches, 6, 115–117, 119–120
- extensions
 - DNS Security, 87
 - third-party forwarding extensions, 121
 - types of, 115–116, 119–120

F

- Failover Clustering
 - Active Directory-detached clustering, 82
 - cache allocation, 73
 - CSV ownership distribution, 71–72
 - draining resources off nodes, 78–80
 - dynamic witness feature, 81
 - enhancements in Windows Server 2012, 63–64
 - health monitoring, 73
 - heartbeat thresholds, 74–75
 - Scale-Out File Servers (SoFS), 70–71
 - virtual machine health protection, 76–78
- Fast Startup, 137
- Federation Proxy Service, 128
- Fibre Channel, 7, 32, 35, 40
- File and Storage Services, 28, 35, 53
- file storage
 - group policy settings, 151
 - Resilient File System (ReFS), 28, 73
 - SoFS, 68
- files
 - copying to virtual machines, 192
 - remote access to, 209–212
- filtering
 - address, 154
 - extensions, 116, 119–120
 - network packet, 197
- firmware configuration, 191
- folders
 - remote access to, 209–212
- forwarding extensions, 116, 119, 121
- FR/CSC, 52

G

- gateway solutions, 116–118
- Gehazi, Mark, 15
- Generation 2 virtual machines, 7–11
 - benefits of, 10–11
 - booting, 7
 - characteristics of, 7–8
 - creating, 9–10
 - deployment of, 8
 - Guest operating system on, 7–8
 - hardware device changes, 8
 - Security in, 10
- GenerationID, virtual machine, 124
- Generic Teaming, 95
- Get-StartApps cmdlet, 199
- Get-WindowsImageContent cmdlet, 187–188
- Get-WindowsSearchSetting cmdlet, 201
- graphics display
 - using DX11.1, 168
- Group Policy
 - advanced graphics for RDS, 168
 - enhancements in Windows Server 2012, 137–138
 - infrastructure status, 137
 - IPv6, 139–141
 - new policy settings for, 142–145
 - preferences, 139–141
 - starter GPOs, 138
 - Windows PowerShell Remoting, 179
 - Windows Runtime, 149–150
- Group Policy settings
 - apps, 144–145, 150
 - authentication, 146–148
 - automatic sign-in, 150–151
 - caching, 138–139
 - logon scripts, 148
 - start screen, 142–144
 - Window Store apps, 150
 - Windows SkyDrive, 151–152
 - Windows Update, 149
 - Work Folders, 145–146
- guest clustering, 7, 25, 65–68, 115
- guest operating system, 7–8, 12–14

H

- hardware certification, 191
- HCK (Windows Hardware Certification Kit), 191
- heartbeat thresholds, 74–75
- help files, 184–186
- high availability, 63–65
- History Viewer, 124
- host clustering, 65
- host management, 21, 64–67
- hot data, 42–44
- hot spares, 30, 44, 49
- Hyper-V
 - cluster node heartbeat thresholds, 75
 - enhancements in Windows Server 2012, 5–7
 - extensible switch architecture, 115–116
 - improvements to, 25
 - live export feature, 18–19
 - PowerShell cmdlets, 191–195
 - support for Linux, 20–21
 - switches, 25, 88, 119–121
 - WinDbg in, 193
- Hyper-V Manager, 9, 13, 17–18
- Hyper-V Network Virtualization
 - Hyper-V Virtual Switch, 115–116, 119–121
 - in Windows Server 2012 R2, 115–118
 - IP addresses on virtual network machines, 115
 - overview, 112–114
 - performance enhancements, 115
 - Windows Server Gateway, 116–118
- Hyper-V Port, 93, 96
- Hyper-V Recovery Manager, 23–24
- Hyper-V Replica, 7, 21–24

I

- Idle Worker Process Page-Out, 156–157
- IEEE 802.3ad teaming, 95
- IEEE 802.lax teaming, 95
- IIS 8.0, 153–155
- IIS 8.5

- cold requests in, 156–157
- configuring ETW logging in, 159
- Dynamic Site Activation, 155–156
- ETW logging, 159–162
- logging custom fields, 158–159
- IIS logging, 158–160
- implementing
 - guest clustering using shared virtual disks, 69–70
 - Web Application Proxy, 128–131
 - Work Folders, 52
- Import-StartLayout cmdlet, 199
- infrastructure
 - hoster, 67
 - key management, 11–12
 - network, 113–114
 - separating virtual resources from physical, 65–67
 - status, 137
- Input/Output Operations Per Second. *See* IOPS
(Input/Output Operations Per Second)
- Intelligent Platform Management Interface (IPMI), 198
- Internet Information Services. *See* IIS 8.0; IIS 8.5
- Internet Protocol version 6 (IPv6), 139–141
- Internet Small Computer System Interface (iSCSI)
 - Initiator, 35
 - server storage, 34–36
 - target cmdlets, 195
 - Target Server, 28, 35–36
- IOPS (Input/Output Operations Per Second),
25, 33, 41, 43–44
- IP Address Management (IPAM), 87, 105–112
 - Active Directory integration, 112
 - architecture, 109
 - communication, 110
 - data collection tasks, 110
 - in Windows Server 2012, 107
 - monitoring capabilities, 110–111
 - role-based access control, 108–110
 - security groups, 110
 - storage, 110–111
 - virtual address management, 107–108
 - Windows PowerShell and, 112
- IPv6 (Internet Protocol version 6), 139–141
- iSCSI (Internet Small Computer System Interface)
 - Initiator, 35
 - server storage, 34–36

- target cmdlets, 195
- Target Server, 28, 35–36

J

- JBOD storage, 27, 29–33, 40
- jobs, scheduled, 176

K

- Kerberos authentication, 146–148
- key management infrastructure, 11–12

L

- LACP (Link Aggregation Control Protocol), 92, 94–95
- layouts, start screen, 142
- LBFO (load balancing and failover), 86, 91
- licensing, virtual machine, 11–12
- Lightweight Directory Access Protocol (LDAP),
123, 131–135
- Link Aggregation Control Protocol (LACP), 92, 94–95
- Linux Integration Services (LIS), 20
- Linux support, 20–21, 195
- live export feature, 18–19
- live migration
 - cross-version, 14–16
 - in previous version of Hyper-V, 5
 - network disconnection during, 78
 - SMB 3.0 protocol, 37–38
 - speed of, 16–17
 - vm availability during, 78–80
- load balancing and failover (LBFO), 86, 91
- load balancing modes, 92–93, 96, 121
- Local Configuration Manager, 178, 180
- logging
 - custom fields, 158–159
 - ETW, 159–162
 - software inventory, 198
- logical unit numbers. *See* LUNs
- logon scripts, 148
- LUNs, 36, 40, 66–67, 70–72

M

MAC address spoofing, 120
 Managed Object Format (MOF) files, 178, 180–182
 Marlin, John, 78
 memory allocation, 20, 73
 memory management, 6, 153, 168, 192–193
 Message Analyzer, 160–162
 Microsoft, 1, 29–31
 Microsoft Office 365, 206
 Microsoft SQL Server databases, 110–111
 Microsoft Windows IIS Logging, 160
 migration. *See also* live migration
 in previous version of Hyper-V, 5–6
 mobile device management, 127–131,
 206–207, 209–212
 MOF files. *See* Managed Object Format (MOF) files
 monitoring
 health, 73
 packet, 120
 using IPAM, 110–111
 virtual machine health, 75–78
 multicast scopes, 189
 My Server apps, 209–212

N

NAT (Network Address Translation) cmdlets, 196
 network adapters
 for guest operating systems, 8
 RDMA-capable, 16–17
 teaming modes, 91
 testing, 97–101, 191–192
 Network Address Translation (NAT) cmdlets, 196
 network event capture cmdlets, 197
 network interruptions, 74–75
 network layer, in cloud computing solutions, 30
 network packet filtering, 197
 network switches, 119–121
 network traffic, 116–118, 121
 network virtualization, 87, 112–118
 networking
 address management, 105–112
 diagnostics, 96–105
 enhancements in Windows Server 2012, 85–88
 Hyper-V Virtual Switch, 121
 load balancing modes, 95–96

 software-defined, 112–120
 virtual RSS (vRSS), 88–90
 Windows NIC Teaming, 90–95
 New-WindowImage cmdlet, 188
 NIC Teaming, 90–96
 benefits of, 90–91
 function of, 86
 in Windows Server 2012, 91
 in Windows Server 2012 R2, 92–96
 load-balancing modes, 92–93, 95–96
 network diagnostics, 96
 teaming modes, 91–92, 94–95
 NMI. *See* Non-Maskable Interrupt (NMI) method
 nodes
 in Configuration Instance Document script, 179
 Non-Maskable Interrupt (NMI), 192–193, 195
 Non-Uniform Memory Access (NUMA), 86, 153, 168

O

Offloaded Data Transfer (ODX), 28

P

packet filtering, 197
 packet monitoring, 120
 paged queries, 134–135
 page-outs, 156–157
 parity space support, 44
 PCSV (Physical Computer System View) cmdlets, 198
 Performance Monitor, 162
 performance versus storage capacity, 41–42
 personal devices
 Workplace Join, 125
 Physical Computer System View (PCSV) cmdlets, 198
 physical storage, in cloud computing solution, 30
 policy caching, 138–139
 port ACLs, 120–121
 port mirroring, 120
 Powershell
 cmdlets
 Get-WindowSearchSetting, 201
 Set-WindowSearchSetting, 201
 PowerShell
 allocating block cache, 73
 changing usage mode of physical disk, 49

cmdlets

- Add-DhcpServerSecurityGroup, 190
- Add-WindowsImage, 186–187
- Copy-VMFile, 192
- Debug-VM, 192–195
- Deployment Image Servicing and Management (DISM), 186–189
- DNS credentials management, 189
- DNS server, 190
- Dynamic Host Configuration Protocol (DHCP), 189–190
- Enable-PSRemoting, 179
- Expand-WindowsImage, 186–187
- Export-CliXml, 186
- Export-StartLayout, 142, 198–199
- Export-WindowsImage, 187
- Get-Cluster, 73, 81
- Get-ClusterGroup, 78
- Get-Command, 89
- Get-DscLocalConfigurationManager, 180
- Get-NetAdapter, 97
- Get-NetAdapterRss, 89
- Get-NetIPAddress, 98–100
- Get-NetIPConfiguration, 100–101
- Get-StartApps, 199
- Get-VMNetworkAdapter, 77–78, 97
- Get-WindowsImageContent, 187–188
- hardware certification, 191
- Hyper-V, 191–195
- Import-StartLayout, 199
- Invoke-Command, 185–186
- Invoke-GPUUpdate, 137
- iSCSI target, 195
- isolation settings for virtual network adapters, 192
- Move-VM, 9
- multicast support, 189
- Net TCP/IP, 196
- Network Address Translation (NAT), 196
- network packet filtering, 197
- Network Virtualization, 197
- New-PSSession, 175
- New-VM, 9
- New-WindowsImage, 188
- Physical Computer System View (PCSV), 198
- record repair, 190

- Remove-WindowsImage, 188
- Repair-DhcpServerv4IPRecord, 190
- routing domains and virtual subnets, 192
- security groups, 190
- Set-DscLocalConfigurationManager, 180
- Show-Command, 176
- Software Inventory Logging, 198
- Split-WindowsImage, 188–189
- StartDscConfiguration, 182
- Test-NetConnection, 102–105
- Test-VMNetworkAdapter, 105, 191
- virtual machine firmware, 191
- Windows Deployment Services, 200–201
- configuring vm health protection, 76–78
- Desired State Configuration (DSC), 176–183
- domain controllers deployment, 124
- enhancements in Windows PowerShell 3.0, 175–176
- Group Policy updates, 137
- History Viewer, 124
- in previous version of Hyper-V, 6
- Invoke-Command, 185
- IP Address Management (IPAM), 112
- IPAM integration with, 112
- managing virtual Receive Side Scaling (vRSS), 89–90
- network diagnostics, 97–105
- Save Help, 184–186
- Start Screen layout, 142, 198–199
- tiered storage, 43
- Web Access, 176, 201
- preauthentication, 130
- preferences
 - Group Policy, 139–141
- Protected Network setting, 75–78
- Provider Addresses, 114, 116
- PSRemoting, 179
- Pull Server. *See* Desired State Configuration (DSC) Pull Server

Q

- Quality of Service (QoS) management, 25, 33–34, 86
- queries, paged, 134–135
- query optimization, 131–135
- Quick Reconnect, 169
- quorum functionality, 64, 81

R

- RDMA (Remote Direct Memory Access), 16
- RDS (Remote Desktop Services)
 - data deduplication, 172
 - enhancements in Windows Server 2012, 165–167
 - Quick Reconnect, 169
 - RemoteFX functionality in, 168
 - Restricted Admin mode for, 173
 - Session Shadowing, 170–172
 - thumbnail preview, 167
 - video performance, 169
 - virtual desktop interface, 172
 - VM Connect, 13
- Receive Side Scaling (RSS), 86, 88–90
- reconnection, session, 169
- ReFS (Resilient File System), 28, 73
- remote access
 - for workers and BYOD, 209–212
 - over VMBus, 12–14
- Remote Access server role, 127–131
- Remote Desktop Services (RDS)
 - data deduplication, 172
 - enhancements in Windows Server 2012, 165–167
 - Quick Reconnect, 169
 - RemoteFX functionality in, 168
 - Restricted Admin mode for, 173
 - Session Shadowing, 170–172
 - thumbnail preview, 167
 - video performance, 169
 - virtual desktop interface, 172
 - VM Connect, 13
- Remote Direct Memory Access (RDMA), 16
- Remote Group Policy updates, 137
- Remote Web Access (RWA), 213
- RemoteFX, 168
- Remove-WindowsImage cmdlet, 188
- Repair-DhcpServerv4IPRecord cmdlet, 190
- replication
 - extended, 23–24
 - frequencies, 21–22
 - in Windows Server 2012, 7
- resiliency
 - Cluster Shared Volumes (CSV), 72–73
 - in previous version of Hyper-V, 5
- Resilient File System (ReFS), 28, 73
- resize, virtual hard disk, 17–18, 20
- restore, client full system, 213
- Restricted Admin mode, 173
- Restricted mode policy setting, 148
- role services
 - Desired State Configuration (DSC), 176–183
 - Federation Proxy Service, 128
 - File and Storage Services, 28, 35, 53
 - Remote Desktop Services (RDS), 165–166
 - Web Application Proxy, 127–131
- router guards, 120
- RSS (Receive Side Scaling), 86, 88–90
- RWA (Remote Web Access), 213

S

- Save function, group policy setting, 152
- Save Help, 184–186
- scale out capability, 29–30
 - Generation 2 virtual machines, 10
 - in Failover Clustering, 63
 - NUMA-aware, 153
 - SMB, 36
 - under Windows Server 2012, 5
- Scale-Out File Servers (SoFS), 6, 27, 30, 68, 70–72
- SCSI controllers, 7–8
- SCSI disks, 7
- search cmdlets, 201
- Secure Boot, 7–8, 10–11
- security
 - groups, 190
 - in Generation 2 virtual machines, 10
- server management, 72–73, 112, 201, 207–209
- Server Message Block (SMB), 36–38
 - file servers. *See* Scale-Out File Servers (SoFS)
 - in Windows Server 2012, 6, 27, 88
 - SMB Bandwidth Management, 37–38
 - SMB Direct, 16, 36
 - SMB Encryption, 36
 - SMB Multichannel, 16, 36
 - SMB Transparent Failover, 36
- Server Name Indication (SNI), 153
- Session Shadowing, 170–172
- sessions
 - disconnected, 175
- Set-WindowsSearchSetting cmdlet, 201
- shadowing. session, 170–172
- Sharepoint online library management, 206
- shutdown, virtual machine, 78–80

- single-root I/O virtualization (SR-IOV), 6, 88
 - site
 - activation, 155–156
 - SM-API (Storage Management API), 32–33
 - Smart Paging, 6
 - SMB (Server Message Block), 36–38
 - file servers. *See* Scale-Out File Servers (SoFS)
 - in Windows Server 2012, 6, 27, 88
 - SMB Bandwidth Management, 37–38
 - SMB Direct, 16, 36
 - SMB Encryption, 36
 - SMB Multichannel, 16, 36
 - SMB Transparent Failover, 36
 - SMS-S, 36
 - SNI (Server Name Indication), 153
 - SoFS. *See* Scale-Out File Servers (SoFS)
 - Software Inventory Logging cmdlets, 198
 - solid-state drives (SSDs), 30
 - Split-WindowImage cmdlet, 188–189
 - SQL Server databases, 110–111
 - SR-IOV (single-root I/O virtualization), 6, 88
 - Srivastava, Deepak, 49
 - SSDs (solid-state drives), 30
 - SSL certificates, 154
 - Start Screen
 - cmdlets, 198–199
 - group policy settings, 142–144
 - startup. *See* booting
 - Static Teaming mode, 91, 94–96
 - status, Group Policy, 137
 - storage
 - block-based, 28–29, 34, 40, 68
 - capacity versus performance in, 41–42
 - enhancements in Windows Server 2012, 27–29
 - file-based, 28, 68, 73, 151
 - iSCSI-based, 34–36
 - JBOD, 27, 29–33, 40
 - layers, 30
 - management, 28, 32–33
 - Microsoft's vision for, 29
 - migration, 6
 - of SSL certificates, 154
 - Quality of Service (QoS), 33–34
 - snapshot, 6
 - tiered, 33, 43–44
 - traditional versus WFS cluster, 40
 - Work Folders, 50–59
 - Storage Management API (SM-API), 32–33
 - storage pools, 27, 30, 39, 41
 - creating, 45–46
 - storage spaces
 - in storage pools, 27
 - rebuilding, 44
 - Storage Spaces, 39–49
 - benefits of, 39–40
 - capacity versus performance in, 41–42
 - cluster-shared volumes and, 73
 - compared to SAN-based storage, 40
 - data tiering, 42–44
 - dual parity, 44–49
 - in Windows Server 2012, 27, 40–42
 - in Windows Server 2012 R2, 42–49
 - Scale-Out File Servers (SoFS), 30–31, 33
 - storage resiliency and parity using, 30
 - write-back cache feature, 33, 43–44
 - storage tiers, 43–44
 - storage volumes. *See* LUNs
 - Switch Dependent teaming mode, 94–95
 - Switch Independent teaming mode, 91–94
 - switches, network, 119–121
 - sync share
 - device policies for, 57
 - names, 56
 - purpose of, 53–55
 - security groups for access, 56–57
 - sharing files with, 57–59
 - user folder structure, 55
 - Sync Your Settings, 145
 - synchronization, data, 50–51
 - System Center, 2–3, 31–32
 - System Center Virtual Machine Manager (VMM), 31–33, 107–108
- ## T
- TCP/IP cmdlets, 196
 - testing
 - connectivity, 102–105, 191
 - network adapters, 97–101
 - thin provisioning, 40
 - third-party extensions, 6, 121
 - third-party solutions
 - iSCSI-based storage and, 35
 - touch functionality, 209, 213
 - traffic, network, 116–118, 121

- troubleshooting
 - network connections, 96–105
 - network tracing, 121
- trunk mode, 120

U

- UEFI (Unified Extensible Firmware Interface), 7
- Unified Extensible Firmware Interface (UEFI), 7
- Unix, 20
- updates
 - failover cluster nodes, 64
 - Group Policy, 137
- usage mode, 49
- user aliases, 55
- user group management, 206
- users, use of Work Folders, 50–51

V

- VHDX format, 6, 36
- video drivers, 20
- video experience, 168
- virtual desktop interface, 172
- virtual disks
 - dual parity and parity space support for, 44–49
 - guest clustering using shared, 65–67, 69–70
 - shared, 67–68
 - tiered storage, 44
- virtual hard disks
 - resizing of, 17–18, 20
- Virtual Machine Connection (VM Connect), 12–14
- Virtual Machine Manager (VMM), 31–33, 107–108
- Virtual Machine Queue (VMQ), 85
- virtual machines
 - automatic activation of, 11–12
 - bandwidth management, 37–38
 - connectivity, 102–105, 191
 - debugging, 192–195
 - draining nodes prior to shutdown, 78–80
 - file shares, 6
 - firmware configuration, 191
 - generation options for, 9–10
 - health monitoring, 75–78
 - importation, 6
 - licensing, 11–12

- live export, 18–19
- maintaining availability, 78–80
- networks, 113–115
 - remote access to, 12–14
 - replication, 7, 21–24
- virtual Receive Side Scaling (vRSS), 88–90
- virtual resources
 - separating from physical infrastructure, 65–67
- virtualization. *See also* Hyper-V
 - in Windows Server Essentials 2012 R2, 204
 - network, 87, 113–114
 - of storage resources, 30
 - Windows PowerShell 4.0 cmdlets, 197
- VLANs, 120
- VM Connect, 12–14
- vm2dmp.exe tool, 193
- VMBus, 12–14
- VMM (System Center Virtual Machine Manager), 31–33, 107–108
- VMQ (Virtual Machine Queue), 85
- vRSS (virtual Receive Side Scaling), 88–90

W

- Wait and Verbose switches, 182
- Web Application Proxy, 127–131
- Web Services Management (WSMAN) devices, 198
- WFP (Windows Filtering Platform), 116
- WinDbg, 193
- Windows 8, 7, 209–212
- Windows Activation Service, 155
- Windows Azure, 2–3, 20
- Windows Azure Active Directory, 206
- Windows Azure Backup, 20, 207
- Windows Azure Hyper-V Recovery Manager, 23–24
- Windows Filtering Platform (WFP), 116
- Windows Hardware Certification Kit (HCK), 191
- Windows Intune, 207
- Windows Management Framework, 175
- Windows NIC Teaming. *See* NIC Teaming
- Windows Phone, 209–212
- Windows PowerShell. *See* PowerShell
- Windows PowerShell Workflow, 176
- Windows Runtime, 149–150
- Windows Search, 201
 - cmdlets, 201

Windows Server 2012

- Active Directory enhancements in, 124
 - Failover clustering enhancements in, 63–64
 - Group Policy enhancements in, 137–138
 - Hyper-V enhancements in, 5–7
 - Hyper-V Virtual Switch in, 119–121
 - IIS 8.0 enhancements in, 153–155
 - IP Address Management (IPAM), 107
 - network adapters in, 97–101
 - network diagnostics, 97–101
 - networking enhancements in, 85–88
 - NIC Teaming in, 91
 - PowerShell enhancements in, 175–176
 - Remote Desktop Services (RDS) enhancements in, 165–167
 - storage enhancements in, 27–29
 - Storage Spaces in, 27, 40–42
- Windows Server 2012 R2
- Hyper-V, 7–25
 - Hyper-V Network Virtualization enhancements in, 115–118
 - Hyper-V Virtual Switch, 121
 - IP Address Management (IPAM), 107–112
 - network diagnostics, 102–105
 - policy settings, 142
- Windows Server and cloud computing, 2–3
- Windows Server Database Edition, 207–209
- Windows Server Essentials 2012, 203
- Windows Server Essentials 2012 R2, 203–214

- Dashboard, 205–207
 - Experience server role, 207–209
 - features of, 203–214
 - in cloud computing, 204–205
 - versions, 203
 - virtualization in, 204
- Windows Server Gateway (WSG), 116–118
- Windows Server Standard Edition, 207–209
- Windows SkyDrive, 52, 151–152
- Windows Store apps, 150
- Windows Update, 149
- Work Folders, 50–59
- comparison to other data sync solutions, 52
 - data synchronization, 53–59
 - functionality, 51
 - Group Policy settings for, 145–146
 - implementing, 52
 - installing, 52–59
 - Kerberos authentication, 146–148
 - naming, 55
 - testing, 59
 - versus Windows SkyDrive, 52
- worker processes, idle, 156–157
- workflows, 176
- workloads, 64–65, 72
- Workplace Join, 125–126
- write-back cache, 43–44
- WSG (Windows Server Gateway), 116–118
- WSMAN (Web Services Management) devices, 198

