Microsoft®

Evaluation Guide

Windows Server 2012

Windows Server

Contents

Copyright Information

Abstract

This evaluation guide is designed to help you understand the design goals, feature set, and implementation of Microsoft Windows Server 2012. The guide provides an overview of key features and benefits provided by Windows Server 2012.

This guide is designed to help you perform a thorough and effective evaluation of Windows Server 2012, and is intended for anyone who is interested in learning more about Windows Server 2012. This guide is not a comprehensive explanation of the capabilities in Windows Server 2012; instead, it highlights the main feature areas that can help you understand and evaluate Windows Server 2012.

For the latest information about Windows Server 2012, including a link to the download location for the Windows Server 2012 release, go to http://www.microsoft.com/windowsserver2012/.

Introduction

The goal of this guide is to help you understand the value of Windows Server 2012 and to evaluate its features and benefits.

Cloud computing can present many opportunities for helping you solve your organization's fundamental business and IT challenges. Adopting a consistent, integrated approach is one of the most effective ways to build a private cloud, offer cloud services, or more securely connect to clouds. It is also beneficial to take advantage of new applications and services that can be deployed on-premises and in public cloud environments.

Windows Server 2012 brings the Microsoft cloud experience to you—from building and operating public clouds to delivering the most dynamic, available, and cost-effective server platform for the private cloud. It provides solutions to help you optimize your IT for the cloud so you can fully meet your organization's unique needs.

Windows Server 2012 delivers value in four key ways:

It takes you beyond virtualization. By offering a dynamic, multitenant infrastructure that can help you scale and secure workloads and cost-effectively build a Microsoft private cloud, Windows Server 2012 can help you provide:

- A complete virtualization platform – A fully isolated, multitenant environment features tools that can help guarantee service level agreements (SLAs), enable chargebacks through usage-based billing, and support self-service delivery.
- Increased scalability and performance – A high-density, highly scalable environment can be modified to perform at the optimum level based on your needs.
- Connectivity to cloud services – A common identity and management framework enables highly secure and reliable cross-premises connectivity.

It delivers the power of many servers, with the simplicity of one. By integrating a highly available and easy-to-manage multi-server platform, Windows Server 2012 can help you provide:

- Continuous availability – New and improved features offer cost-effective, high IT service uptime. They are designed to endure failures without disrupting services to users.
- Management efficiency – The automation of an even broader set of management tasks and the simplified deployment and virtualization of major workloads offer a path to full lights-out automation.
- Cost efficiency – The use of commodity storage, networking, and server infrastructure and increased power efficiency offer improved acquisition and operating economics.

It opens the door to every app on any cloud. By offering a broad, scalable, and elastic server platform that gives you the flexibility to build and deploy applications and websites on-premises, in the cloud, and in a hybrid environment, Windows Server 2012 can help you deliver:

- Flexibility to build on-premises and in the cloud – Developers can use a consistent set of tools and frameworks to build symmetrical or hybrid applications between the datacenter and the cloud.
- A scalable and elastic infrastructure – Windows Server 2012 offers service providers new features to increase website density and efficiency, and provides frameworks, services, and tools to increase scalability and elasticity for multitenant-enabled applications. Windows Server 2012 enables service providers to more effectively build, provision, and manage hosting environments.
- An open web and application development environment – Windows Server 2012 enables mission-critical applications and enhanced support for open standards, open source applications, and various development languages.

Modern workstyle, enabled. By empowering your IT professionals to provide users with flexible access to data and applications while simplifying management and maintaining security, control, and compliance, Windows Server 2012 can help you offer:

- Access to applications and data from virtually anywhere on nearly any device – Easier, on-demand access to virtualized work environments is available from virtually anywhere, including branch locations and public connectivity services.
- A more full Windows experience from virtually anywhere – A personalized and rich user experience is enabled from virtually any device, while adapting to different network conditions more quickly and responsively.
- Enhanced data security and compliance – Granular access to data and corporate resources is enabled by central audit and access policies based on strong identity, data classification, and simplified administration for remote access.

Converging business needs and technology innovations have opened the door for cloud computing—the on-demand delivery of standardized IT services running on shared resources. Cloud computing takes virtualization to the next level by treating compute, network, and storage resources as a flexible pool that can be allocated to any workload.

By providing the ability to increase and decrease usage based on demand, cloud computing can give you greater elasticity and scale at lower cost and with greater reliability. In addition, IT can empower your users to use self-service support, which frees up IT resources and gives users faster results.

IT professionals from organizations of all types and sizes—a small company, a government agency, a hosting provider, or a large enterprise—recognize that they need a cloud-optimized server platform that helps them implement these innovations so they can meet business requirements. That platform is Windows Server 2012.

Cloud Optimize Your IT with Windows Server 2012

When it comes to IT infrastructure, the message is clear: Organizations want a more consistent, integrated approach that helps them solve fundamental business and IT challenges. They also want to take advantage of new applications and services that can be deployed on-premises and in private and public cloud environments. The next release of Windows Server, Windows Server 2012, is a more dynamic, available, cost-effective server platform that allows organizations of all types and sizes to optimize their IT solutions for the cloud.

Converging needs and innovations

Business needs and technology innovations are converging, offering IT professionals a unique opportunity to take advantage of new innovations to meet business requirements. Consider these trending business needs:

- **Agility and flexibility.** Business owners need to be flexible and respond rapidly to market changes, while hosting providers must develop innovative new services to attract and maintain customers.
- **Efficiency.** Whether your organization is a small company, a government agency, a hosting provider, or a large enterprise, everyone must do their job more cost-effectively and efficiently to maintain margins.
- **Compliance.** Businesses and hosting providers need to protect customer and personal identity information, respond to changing regulations, and comply with industry and governmental policies.
- **Access.** Workers need access to data and information regardless of the infrastructure, network, device, or application used to deliver it. Businesses and hosting providers need to be able to offer anytime, anywhere access to IT services to satisfy worker and customer expectations.

To support these business needs, IT professionals themselves need an infrastructure that scales up and down quickly to meet changing business needs, minimizes downtime and failures, and maximizes management and cost efficiencies. Fortunately, several key technology innovations are making that more possible:

- **Virtualization.** IT departments can respond faster to requests from business units, reducing the time it takes to deploy infrastructure and services. Additionally, virtualization significantly reduces the number of physical servers required to support the business.
- **Security and identity management.** These critical technologies are evolving to provide highly secure and compliant environments that protect important assets and corporate and personal identities.
- **Cloud-based applications.** Anywhere access to critical applications helps to increase work productivity, improve communication, and increase customer contact, allowing organizations to improve their regular business rhythm and respond to market changes and opportunities.
- **Multitenancy and cross-premises integration.** These innovations help IT departments and hosting service providers maximize existing infrastructure investments while exploring new services, improved management, and higher availability.

Cloud computing opportunities

Converging business needs and technology innovations have opened the door for cloud computing—the on-demand delivery of standardized IT services running on shared resources. Cloud computing takes virtualization to the next level by treating compute, network, and storage resources as a flexible pool that can be allocated to any workload. When a cloud datacenter does this it becomes more dynamic and enables full decoupling of the physical infrastructure from the logical workloads.

By providing the ability to increase and decrease usage based on demand, cloud computing gives IT greater elasticity and scale at lower cost and with greater reliability. In addition, IT can empower users to use self-service support, which frees up IT resources and gives users faster results. Your organization can be more efficient because cloud computing is based on usage and driven by SLAs.

One important difference between private versus public clouds is distinguished by who can use the cloud services. A public cloud implements and offers services to many unrelated organizations, usually for a fee. A private cloud offers services within an organization. A hybrid cloud offers services from private and public clouds.

IT and business professionals are beginning to see the opportunities available through the cloud. Consider these two examples:

- Technology innovations such as virtualization, multitenancy, and Resource Metering make it easier for IT departments to provide organizations with flexibility, efficiency, and agility.
- Organizations must grant the right users access to information access while maintaining regulatory compliance. Technology improvements in security and identity management make it easier for IT to enforce compliance and give workers the right level of access to the right information from nearly anywhere.

Moving to cloud-optimized IT

To move toward cloud-optimized IT that can span on-premises and off-premises environments, you must have four key capabilities:

- A common virtualization platform that increases efficiency and performance across your infrastructure
- A common identity through directory services
- Common management through automation
- Common developer tools and platforms

Microsoft offers a common set of tools and services that provide these capabilities, all of which either are found in Windows Server 2012 or easily integrate with it. When combined with a set of management tools, such as Microsoft System Center, Windows Server 2012 offers a more complete private cloud solution. Windows Server 2012 provides the platform functionality that manages the physical servers, networking, and storage access, and enables the management layer built on top of it to expose these as a pool of compute, network, and storage resources.

For example, with Windows Server 2012, IT professionals can turn the cloud opportunities discussed earlier into reality:

- Network isolation and improved Hyper-V virtualization combined with new Resource Metering make it easier for hosting service providers and enterprises to implement private cloud services and usage-based billing for their customers and departments.

- Using new and improved features in Active Directory Doman Services (AD DS) along with claims-based authorization makes it easier for IT to grant access to information to the right people while maintaining a tighter audit trail for information governance and compliance with regulations such as the Health Insurance Portability and Accountability Act (HIPAA) and the Sarbanes-Oxley Act (SOX).

Capitalizing on existing skills and investments with Windows Server 2012

When you optimize your IT solutions for the cloud with Windows Server 2012, you capitalize on the skills and investment you have already made in constructing a familiar and consistent platform. Windows Server 2012 builds on that familiarity and optimizes your IT solutions for cloud infrastructure by helping you:

- Scale and protect workloads, cost-effectively build a private cloud, and more securely connect to cloud services.
- Efficiently manage infrastructure while maximizing uptime and minimizing failures and downtime.
- Build on an open and scalable web platform that supports cross-premises applications.
- Support a mobile and flexible work style that gives people access to data regardless of the infrastructure, network, device, or application they use to access it.

Windows Server 2012 is an ideal operating system for running multiple servers and the devices connecting them, whether they are physical or virtual, on-premises or off-premises.

With Windows Server 2012, you gain all the Microsoft experience behind building and operating private and public clouds, delivered as a more dynamic, available, cost-effective server platform. For example, Microsoft is the one of the only companies that has deep experience providing a comprehensive approach to cloud services at all levels:

- Cloud-based applications and experience through Microsoft Office 365 and Microsoft Dynamics
- Public cloud-based platforms through Windows Live Hotmail, Windows Live Messenger, Bing, Windows Azure, and Xbox LIVE
- Private cloud services for Microsoft itself

Consider these statistics:

- 9.9 billion messages a day delivered via Windows Live Messenger
- 600 million unique users every month on Windows Live and MSN
- 500 million active Windows Live IDs
- 5 petabytes of content served by Xbox LIVE during Christmas week
- Over 1 petabyte of updates served every month by Windows Update to millions of servers and hundreds of millions of PCs worldwide
- Tens of thousands of Windows Azure customers
- 40 million paid Microsoft Online Services in 36 countries

No other technology providers come close to this level of experience or completeness in a cloud offering. Although other providers tend to focus on providing specific cloud components, such as virtualization or development tools, Microsoft continues to incorporate lessons learned from implementing and managing cloud services to provide a cloud-optimized operating system with Windows Server 2012.

What if I'm not ready to move to the cloud?

Moving to Windows Server 2012 doesn't mean that you have to move to the cloud right away. It does mean that you're dedicated to staying current with a server platform that supports yesterday's technologies and today's business needs while preparing you for tomorrow's innovations.

Many of the new and improved features in Windows Server 2012 can help you meet demanding business requirements without ever giving thought to clouds—private or public. If you're working with virtualized environments, many of the improvements in Hyper-V and storage technologies will help you enhance your virtual environment and reduce storage costs. For example, Windows Server 2012 offers high-performance, continuously available file-share storage for server applications. Implementing this feature lets you store application data on inexpensive, easy-to-manage file shares and obtain nearly the same benefits you would expect from a storage area network (SAN).

Microsoft built Windows Server 2012 so you can capitalize on your existing investment. By choosing Windows Server 2012 now, you are preparing for the future. As you continue to grow and work within a heterogeneous environment, you can select the path that is recommended for your business. Whether implementing a private cloud is around the corner or over the horizon for you, Windows Server 2012 offers a great platform to prepare for and implement cloud-optimized IT now.

Preview for Windows Server 2012

This collection provides a high-level technical overview of the new and improved features in Windows Server 2012.

Active Directory

Active Directory

Dynamic Access Control

Failover Clustering

Cluster-Aware Updating

Failover Clustering

High-Performance, Continuously Available File Share Storage for Server Applications

File and Storage Services

Data Deduplication

Datacenter Diskless Boot

Encrypted Hard Drive

iSCSI High-Availability Block Storage

Multiterabyte Volumes

Rapid and Efficient Data Movement Using Intelligent Storage Arrays

Server for NFS Data Store

Storage Management

Storage Spaces

Thin Provisioning and Trim Storage

Unified Remote Management for File Services

User-Device Affinity

Hyper-V

Management

Networking

Remote Desktop Services

Remote Desktop Services

Server Core

Server Core and Full Server Integration

Active Directory

You can use Active Directory Domain Services (AD DS) in Windows Server 2012 to more rapidly and easily deploy domain controllers (on-premises and in the cloud), increase flexibility when auditing and authorizing access to files, and more easily perform administrative tasks at scale (locally or remotely) through consistent graphical and scripted management experiences. AD DS improvements in Windows Server 2012 include:

- Simplifying the on-premises AD DS deployment (formerly DCpromo) with a new streamlined domain controller promotion wizard that is integrated with Server Manager and built on Windows PowerShell.
- Providing greater support for the capabilities of public and private clouds through virtualization-safe technologies and the rapid deployment of virtual domain controllers through cloning.
- Providing a consistent graphical and scripted management experience that allows you to perform tasks in the Active Directory Administrative Center, automatically generating the required syntax for automating the task in Windows PowerShell.

Requirements

- Windows Server 2012
- AD DS server role

Technical overview

Active Directory and AD DS has been at the center of IT infrastructure for over 10 years, and its features, adoption, and business-value have grown release over release. Today, the majority of that Active Directory infrastructure remains on-premises, but there is an emerging trend toward cloud computing.

The adoption of cloud computing, however, will not occur overnight, and migrating suitable on-premises workloads or applications is an incremental and long-term exercise. New hybrid infrastructures will emerge, and it is essential that AD DS supports the needs of these new and unique deployment models that include services hosted entirely in the cloud, services that consist of cloud and on-premises components, and services that remain exclusively on-premises. These hybrid models will increase the importance, visibility, and emphasis around security and compliance, and they will compound the already complex and time-consuming exercise of ensuring that access to corporate data and services is appropriately audited and accurately expresses the business intent.

The following sections describe how Windows Server 2012 addresses these emerging needs.

Simplified deployment

AD DS deployment in Windows Server 2012 integrates all the required steps to deploy new domain controllers into a single graphical interface. It requires only one enterprise-level credential, and it can prepare the forest or domain by remotely targeting the appropriate operations master roles. The new deployment process conducts extensive prerequisite validation tests that minimize the opportunity for errors that might have otherwise blocked or slowed the installation. The AD DS Installation Wizard is built on Windows PowerShell, integrated with Server Manager, able to target multiple servers, and remotely deploy domain controllers, which results in a deployment experience that is simpler, more consistent, and less time consuming. The following figure shows the AD DS Configuration Wizard in Windows Server 2012.

Figure 1: AD DS Installation Wizard

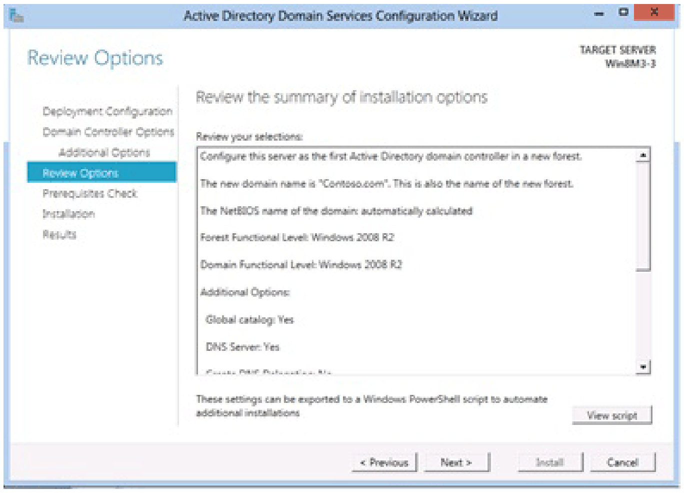

Review the summary of installation options

Review your selections:

Configure this server as the first Active Directory domain controller in a new forest.

The new domain name is "Contoso.com". This is also the name of the new forest.

The NetBIOS name of the domain: automatically calculated

Forest Functional Level: Windows 2008 R2

Domain Functional Level: Windows 2008 R2

Additional Options:

Global catalog: Yes

DNS Server: Yes

These settings can be exported to a Windows PowerShell script to automate additional installations

An AD DS installation includes the following features:

- **Adprep.exe integration into the AD DS installation process.** Reduces the time required to install AD DS and reduces the chances for errors that might block domain controller promotion.
- **The AD DS Installation Wizard, which is built on Windows PowerShell and can be run remotely on multiple servers.** Reduces the likelihood of administrative errors and the overall time that is required for installation, especially when you are deploying multiple domain controllers across global regions and domains.
- **Prerequisite validation in the AD DS Installation Wizard.** Identifies potential errors before the installation begins. You can correct error conditions before they occur without the concerns that result from a partially complete upgrade.
- **Configuration pages grouped in a sequence that mirror the requirements of the most common promotion options, with related options grouped in fewer wizard pages.** Provides better context for making installation choices and reduces the number of steps and amount of time required to complete the domain controller installation.
- **A wizard that exports a Windows PowerShell script that contains all the options that were specified during the graphical installation.** Simplifies the process by automating subsequent AD DS installations through automatically generated Windows PowerShell scripts.

Deployment with cloning

AD DS in Windows Server 2012 helps you to deploy replica virtual domain controllers by “cloning” existing virtual domain controllers. You can promote a single virtual domain controller by using the domain controller promotion interface in Server Manager, and then rapidly deploy additional virtual domain controllers within the same domain, through cloning.

The process of cloning involves creating a copy of an existing virtual domain controller, authorizing the source domain controller to be cloned in AD DS, and creating a configuration file that contains detailed promotion instructions (for example, name, IP address, Domain Name System [DNS] servers, and so on). Or you can leave the configuration file empty, which allows the system to automatically fill in the information. Cloning reduces the number of steps and time involved by eliminating repetitive deployment tasks, and it enables you to fully deploy additional domain controllers that are authorized and configured for cloning by the Active Directory domain administrator.

Safer virtualization of domain controllers

AD DS has been virtualized for several years, but features present in most hypervisors can invalidate strong assumptions made by the Active Directory replication algorithms. Primarily, the logical clocks that are used by domain controllers to determine relative levels of convergence only go forward in time. In Windows Server 2012, a virtual domain controller uses a unique identifier that is exposed by the hypervisor. This is called the *virtual machine GenerationID*. The virtual machine GenerationID changes whenever the virtual machine experiences an event that affects its position in time. The virtual machine GenerationID is exposed to the virtual machine’s address space within its BIOS, and it is made available to the operating system and applications through a driver in Windows Server 2012.

During boot and before completing any transaction, a virtual domain controller running Windows Server 2012 compares the current value of the virtual machine GenerationID against the value that it stored in the directory. A mismatch is interpreted as a “rollback” event, and the domain controller employs AD DS safeguards that are new in Windows Server 2012. These safeguards allow the virtual domain controller to converge with other domain controllers, and they prevent the virtual domain controller from creating duplicate security principals. For Windows Server 2012 virtual domain controllers to gain this extra level of protection, the virtual domain controller must be hosted on a virtual machine GenerationID–aware hypervisor such as Windows Server 2012 with the Hyper-V role.

Windows PowerShell script generation

The Windows PowerShell cmdlets for Active Directory are a set of tools that help you to manipulate and query AD DS to create scripts that automate common administrative tasks. The Active Directory Administrative Center uses these cmdlets to query and modify AD DS according to the actions performed within the Active Directory Administrative Center GUI.

In Windows Server 2012, the Windows PowerShell History viewer in the Active Directory Administrative Center, as shown in the following figure, allows an administrator to view the Windows PowerShell cmdlets as they run. For example, when you create a new fine-grained password policy, the Active Directory Administrative Center displays the equivalent Windows PowerShell cmdlets in the Windows PowerShell History viewer task pane. You can then use those cmdlets to enable automation for the process by creating a Windows PowerShell script.

Figure 2: Windows Server 2012 Windows PowerShell History viewer

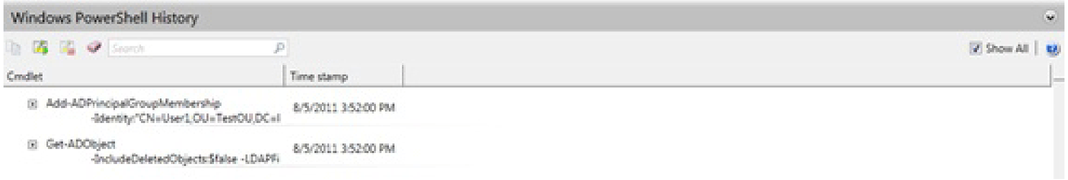

Cmdlet	Time stamp
⊞ Add-ADPrincipalGroupMembership -Identity:"CN=User1,OU=TestOU,DC=I	8/5/2011 3:52:00 PM
⊞ Get-ADObject -IncludeDeletedObjects:$false -LDAPFi	8/5/2011 3:52:00 PM

By combining scripts with scheduled tasks, you can enable automation for everyday administrative duties that were once completed manually. Because the cmdlets and required syntax are created for you, very little experience with Windows PowerShell is required. Because the Windows PowerShell cmdlets are the same as those that are run by the Active Directory Administrative Center, they are designed to function as expected.

Automation and Management

This section presents four scenarios to describe some of the improvements in automation and server management in Windows Server 2012.

Scenario #1: Increasing servers-per-administrator

Servers-per-admin is a long-standing IT benchmark that has recently become critically important due to the increased adoption of high-performance computing, server virtualization, and cloud computing, both private and public. A low servers-per-admin ratio can limit an IT organization's ability to support their company's business, or provide support only at an unsustainable expense.

Windows Server 2012 increases servers-per-admin by delivering a platform for robust, multicomputer automation for all the elements of a datacenter, including servers, Windows operating systems, storage, and networking. This platform can help you meet the following goals:

- Create distributed workflows that run small activities (in sequence or in parallel) that perform larger management tasks, such as multicomputer application provisioning.
- Design a system that can handle task failures and interruptions, such as retrying operations when a remote device is unavailable, or when a computer that is executing a long-running task needs to be restarted.
- Schedule jobs that run regularly or in response to an event to deliver standardized "lights-out" operations.
- Allow commands to be executed with a set of delegated credentials to minimize the privileges of each administrator.

These investments deliver an "operations ready" platform that enables automation of operating systems for high quality and high productivity.

Technical overview

Workflows

Windows PowerShell 3.0 goes beyond scripting and enables you to write workflows (long-running tasks can be repeated, made parallel, interrupted, and restarted). Workflows can be written in the Windows PowerShell language or in XAML, and they are executed by the Windows Workflow Foundation (WF) engine. Commands and scripts that restart a computer (for example, the **Restart-Computer** cmdlet) can wait for the computer to resume, or wait for a particular service on the restarted computer to be available.

Robust Sessions and handling failures

With the release of Windows PowerShell 3.0, sessions are not only persistent, they are resilient.

The Robust Sessions feature stores Windows PowerShell sessions on the remote (server-side) computer so users can search and reconnect to sessions at a later time.

Enhancements to the underlying protocol help ensure that connections are re-established in the event of intermittent or prolonged network failure. If sessions cannot reconnect, Windows PowerShell safely disconnects the session, saving the commands and the results until the user is able to reconnect.

Users can disconnect from and then reconnect to any session without losing state, whether they disconnect unintentionally through a network failure or intentionally by shutting down their computer and reconnecting from a different computer hours later.

Job scheduling

Windows PowerShell allows administrators to schedule jobs to run at a later time or according to a particular schedule. The Windows Task Scheduler is used to schedule and start the job, and a per-user job repository is used to store job output so that it is available later in any Windows PowerShell session on the computer.

RunAs capability

Remote administration is often delegated, and the session can ask users to perform tasks that require credentials that they do not possess. With Windows PowerShell 3.0, administrators can design session configurations that contain a defined set of commands that can be run by using the credentials of a different user. Credentials are stored securely in the WS-Management provider.

The following sections demonstrate the four features described above that enable improved remote multicomputer management: Robust Sessions, Workflows, RunAs, and Scheduled Jobs.

System requirements

This scenario requires the following programs and services:

- Windows PowerShell 3.0
- WS-Management 3.0
- Microsoft .NET Framework 4.0
- Windows PowerShell Integrated Scripting Environment (ISE)

To install the Windows PowerShell ISE feature (server only). Windows PowerShell ISE is installed by default on client computers running Windows operating systems. It is an optional feature in Windows Server operating systems.

To add the optional Windows PowerShell ISE feature, use one of the following methods.

- Start Windows PowerShell with the "Run as administrator" option.
- At the command prompt, type *Add-WindowsFeature PowerShell-ISE*

- OR -

- In Server Manager, start the Add Roles and Features Wizard.
- On the **Select features** page, click **Windows PowerShell ISE**.

Scenario summary

This section covered some of the new and enhanced features of Windows PowerShell remote management and background jobs.

The highlights include:

- **Robust Sessions** that can survive network failures.
- **Disconnect-Reconnect** enables users to disconnect from a remote session and reconnect later—even from a different computer by a different user.
- **Persistent commands and jobs** continue to run even if the session is disconnected.
- **Suspend-Resume** enables you to stop a running job and restart it later.
- **Workflows** bring the power of Windows Workflow Foundation to Windows PowerShell. Users can write workflows in XAML or in the Windows PowerShell language.
- **RunAs** is a new property of session configurations that enables commands in the session to run with an alternate credential by default.
- **Scheduled jobs** help you use Task Scheduler in Windows PowerShell to schedule define, trigger, and run scheduled jobs.

Scenario #2: Enables automation

Today's IT environments demand sophisticated production-quality automation to meet business objectives without needing an IT staff with advanced programming skills. Windows Server 2012 makes it easier for IT professionals to create custom automation solutions for unique problems by providing the right tools, making it easier to find the appropriate tool, and making it simpler to combine these tools to author a solution.

Simplicity is attained when a user can think about what they need, type it, and complete the task. Windows Server 2012 focuses on simplicity through the following investments:

- **Improved cmdlet discovery** makes it easier to find and run any of the 2,300 new high-level, task-oriented cmdlets that manage all elements of a datacenter.
- **Simplified language syntax** helps make scripts look less like code and more like natural language.
- **Show-Command** is a graphical search tool that helps you find the right cmdlet, view its options in a dialog box, and run it.
- **Productivity features** in the Windows PowerShell Integrated Scripting Environment (ISE) help you author clear, maintainable, production-ready scripts faster and more easily.
- **Update-Help** is a new cmdlet that updates Help files from the web, which enables cmdlet you to deliver accurate, real-world help by continuously incorporating feedback from the community.

Technical overview

Cmdlet discovery: Get-Command and Module Auto-Loading

The new cmdlet model hides the "module" packaging abstraction, which enables users to find and use cmdlets without having to search for them and import them. *Get-Command* has been updated to find all cmdlets that are installed on the system. Users can take advantage of the cmdlets immediately, because modules are imported automatically on first use.

Syntax simplification

The **ForEach-Object** and **Where-Object** cmdlets have been updated to support an intuitive command structure that more closely models natural language. Users can construct commands without script block, braces, the current object automatic variable ($_), or dot operators to get properties and methods. This means that the "punctuation" that plagued beginning users is no longer required.

New Windows PowerShell ISE features

The Windows PowerShell Integrated Scripting Environment (ISE) 3.0 includes many new features to ease beginning users into Windows PowerShell and provide advanced editing support for scripters.

- The **Show-Command** pane gives you the ability to find and run cmdlets in a dialog box.
- **IntelliSense** provides context-sensitive command completion for cmdlet and script names, parameter names and enumerated values, and property and method names.
- **Snippets** add reusable text to scripts and commands. The built-in snippets include templates for functions, parameters, and statements so users do not have to remember the syntax.
- **Collapsible regions** in scripts and XML files facilitate navigation in long scripts.

Updatable Help

Windows PowerShell 2.0 included extensive Help topics that were frequently updated online. But because the Help files were part of the Windows operating system, users could not update them, and the Help topics that were displayed at the command line were soon outdated. Non-Microsoft products had to convert online help to XML or display outdated help topics.

Help files for Windows PowerShell do not ship with the product. Instead, new *Update-Help* and *Save-Help* cmdlets download and install the newest Help files for each module. The cmdlets do all of the work, including finding the Help files on the Internet, determining whether they are newer than local files, unpacking them, and installing them in the correct location. The updated files are ready for immediate use in *Get-Help*, and users do not need to restart Windows PowerShell.

For large enterprises and users behind Internet firewalls, the *Save-Help* cmdlet downloads Help files to a file system location, such as a file share. The *Update-Help* cmdlet can download Help files from the file share, just as it does for Internet locations.

Updatable Help is available for all modules, including non-Microsoft modules, and it includes support for multiple languages.

System requirements

This scenario requires the following programs and services:

- Windows PowerShell 3.0
- WS-Management 3.0
- Microsoft .NET Framework 4.0
- Windows PowerShell Integrated Scripting Environment (ISE)

Windows PowerShell ISE is installed by default on Windows Server 2012 Full GUI and Minimum Server Interface. You can enable ISA on servers running Server Core, using the PowerShell cmdlet, *Add-WindowsFeature PowerShell-ISE*. This will install the Minimum Server Interface to that server.

Scenario summary

This section covered features that Windows PowerShell has made simpler, including:

- **Module auto-loading.** Modules are imported into the session on first use. The *Import-Module* cmdlet is still available but there is no need to use it.
- **Cmdlet discovery.** The *Get-Command* cmdlet now gets commands in all installed modules, even if they are not in the current session. When the command is used, it loads the module.
- **Syntax simplification.** The *ForEach-Object* and *Where-Object* cmdlets have new parameter sets that do not require script blocks or the current-object operator (*$_*), which makes the command format closer to natural language.
- ***Get-ChildItem*** improvements. The *Get-ChildItem* cmdlet includes new parameters for filtering files by file attributes.
- **Windows PowerShell ISE enhancements.** New features include IntelliSense, collapsible sections, snippets, and the *Show-Command* pane.

Scenario #3: Efficient deployment of workloads to remote servers and offline virtual hard disks

Industry shifts towards virtualized datacenters are driving an overall increase in deployed server images (virtual and physical). For administrators to realize the full value of virtualization, they require an efficient solution to deploy server workloads across the datacenter to keep up with accelerating scale.

In Windows Server 2008 R2, roles and features are deployed by using the Add Roles Wizard or Add Features Wizard in Server Manager running on a local server. This requires physical access to the server or Remote Desktop access by using Remote Desktop Protocol (RDP). Installing the Remote Server Administration Tool lets you run Server Manager on a Windows-based client computer, but functionality for adding roles and features is disabled, because remote deployment is not supported.

In Windows Server 2012, the deployment capabilities have been extended to support robust remote deployment of roles and features. By using Server Manager in Windows Server 2012, IT professionals can provision servers from their desktops without requiring physical access to the system or the need to enable RDP connections to each server. This improves the efficiency of server provisioning in several ways:

- The process of installing and configuring new servers is streamlined when servers are physically situated in remote datacenters.
- IT professionals can more easily deploy roles and features to minimal GUI servers such as Server Core.
- The process of provisioning new virtual server images is simplified with the ability to deploy roles and features directly to offline virtual hard disks (VHDs).
- Automation can be enabled for batch deployment of roles and features to multiple remote computers by using Windows PowerShell.

Technical overview

Enabling robust remote deployment

The Add Roles and Features Wizard in Server Manager are implemented by using new WMI providers that enable remote deployment and configuration. These providers run on Windows Server 2012, and they are managed by a workflow that runs on client computers based on Windows 8 remain robust after computer restarts.

When you enable the automation of remote deployments by using Windows PowerShell, the deployment cmdlets run locally on client computers based on Windows 8, and they communicate with native WMI providers on the deployment destination server, which eliminates the need for Windows PowerShell to be installed on minimal-GUI server images. Windows PowerShell cmdlets can also be used directly to enable the automation of batch deployment to Windows Server 2012 systems (including offline VHD images).

Streamlining server configuration and deployment

In Windows Server 2012, Server Manager includes configuration functionality that was previously provided in the *Initial Configuration Tasks* window. The result is a single surface for managing the configuration of Windows Server 2012 and its roles and features.

Deployment of both roles and features has been combined into a single Add Roles and Features Wizard. Although the process of installing roles is familiar and consistent with the Add Roles Wizard in earlier Windows Server releases, there are changes. To support remote deployment and installations on offline VHDs, some initial configuration tasks (formerly performed in the Add Roles Wizard during an installation) have moved to post-installation configuration wizards. For some offline VHD deployments, installation tasks have been deferred until the first time the resulting virtual machine is started.

Enabling batch deployment

In Windows Server 2012, the Add Roles and Features Wizard help you export configuration options to an XML file for use later with Windows PowerShell deployment cmdlets. By using the fan-out capabilities in Windows PowerShell, you can perform batch deployments of roles and features on multiple remote servers, and apply configuration settings that were saved during a previous wizard-based deployment.

Requirements

- Server Manager is installed on Windows Server 2012 by default, and it opens automatically when you first log on to your Windows Server 2012 system. Server Manager does not run on the Server Core installation option of Windows Server.

Scenario summary

This section covered the new features in Server Manager that simplify the processes of configuring new servers and deploying roles and features to Windows Server and offline VHDs. Remote deployment innovations have removed the need to be physically at a server to deploy roles, and batch deployment can be enabled through new Windows PowerShell cmdlets. These features help IT professionals to more efficiently provision new servers within their datacenters.

Scenario #4: Managing across multiple servers by using Server Manager

In Windows Server 2008 R2, Server Manager helps administrators understand the status and events on a single local or remote server, and the status and events for an installed role. As the size, number, and complexity of datacenters increase, managing individual servers becomes increasingly inefficient. IT professionals must manage workloads consisting of components that span multiple physical or virtual servers, instead of managing individual physical servers.

In Windows Server 2012, Server Manager provides a role-centric dashboard that helps IT professionals understand, at a glance, the state of their servers. Server Manager in Windows Server 2012 helps administrators manage groups of servers collectively from within a single, integrated console, which allows them to respond to business-critical issues with greater speed and agility.

As Server Manager moves toward a multiple-server management experience, existing single-computer tools (such as many existing MMC-based tools) are no longer as tightly integrated into Server Manager, but they can still be launched from Server Manager. As role-specific tools are enhanced in the future to support multiserver management, they will be fully integrated into Server Manager to provide a single launch point. In Windows Server 2012, the following multiple-server solutions are part of the Server Manager experience:

- **File storage management.** The management of file servers is improved by shifting from a single-server, single-service management model, to one in which multiple, individual file servers, or multiple failover clusters that are running the File Services role, can be managed remotely by using a single management application.
- **Remote Desktop Services.** Remote Desktop Services (RDS) provides session virtualization and virtual desktop infrastructure technologies that enable users to access session and virtual desktop collections. In Windows Server 2012, new management features in Server Manager simplify how RDS is deployed and managed in multiserver configurations. Scenario-based deployment reduces the complexity of installing RDS components across multiple servers, based on how RDS will be used. New multiple-server management tools simplify how administrators manage servers that are running RDS role services and virtual desktop infrastructures.
- **IP Address Management.** The IP Address Management (IPAM) feature is a central solution for managing IP addresses and associated infrastructure roles, such as DHCP server and DNS server, across a network. IPAM supports the discovery of addressing and naming servers. It provides a unified experience for tracking utilization trends and managing dynamic, static, and virtual IPv4 and IPv6 addresses. IPAM also supports monitoring the DNS Server service and the DHCP Server service, and multi-entity management of the DHCP Server service. IPAM tracks configuration changes and logs IP lease activity across the network.

Technical overview

The multiple-server experience of Server Manager

Server Manager is a standalone Windows Presentation Foundation (WPF) application that is layered over Windows PowerShell. It takes advantage of the remote management and workflow capabilities in Windows PowerShell to provide robust multicomputer management. Server Manager generates status views for multiple servers after polling servers for operational statistics, including which roles and features are installed, and events, service states, performance threshold alerts, and Best Practices Analyzer (BPA) scan results.

Integration with other management tools

Server Manager remains the key access point for server management tools. Where supported, Server Manager launches these tools in the context of the remote server that you are managing. New role-specific tools (such as File Storage Management, RDS, and IPAM) are easily integrated into the Server Manager console.

Performance impacts of Server Manager

The Server Manager dashboard does not provide live monitoring, and it has a default ten-minute polling cycle that users can modify in the console. By using a relatively infrequent default polling cycle, and returning only incremental data with each poll, the performance load impact on individual servers is minimized. Server Manager uses new WMI providers and Windows PowerShell cmdlets to pull updated status information from servers.

Note

Server Manager aims to minimize the performance impacts of polling across multiple servers and to provide the administrator with a level of control through customization options. Windows Server 2012 has been tested to manage up to ten remote servers. This number is slated to increase in future releases.

Managing earlier versions of Windows Server

The WMI providers that are used by Server Manager are installed by default on Windows Server 2012, and they will be available as an optional installation for Windows Server 2008 and Windows Server 2008 R2. Server Manager cannot fully manage Windows Server 2003 systems.

Note

These optional installation packages for Windows Server 2008 and Windows Server 2008 R2 are not available as part of Windows Server 2012.

Requirements

- Server Manager is installed on Windows Server 2012 by default, and it opens automatically when you first log on to your Windows Server 2012 system. Server Manager does not run on the Server Core installation option of Windows Server.
- To manage a server in Server Manager, administrator credentials for the managed server are required. For this evaluation, users can log on to all virtual machines in your test environment by using the built-in Administrator account.
- Users must enable the value of the *Remote management* property on the Local Server page in the Server Manager console on the server that you want to manage remotely.

Scenario summary

This section explained how the new features in Server Manager help administrators to better manage multiple server environments. The role-centric dashboard provides a clear understanding of the state of servers, and the integrated console enables administrators to respond to issues with more speed and agility.

BranchCache

Using servers to deliver content over wide area network (WAN) connections enables organizations to store content and applications in remote datacenters. Web servers, file servers, and application servers are being moved from main offices to remote locations, and employees depend on WAN connections to access their critical data.

In Windows Server 2012 and Windows 8, BranchCache can optimize the bandwidth over WAN connections between content servers and remote client computers. The following sections describe the components and new features in BranchCache.

BranchCache overview

BranchCache is a WAN bandwidth optimization technology that is included in the Windows Server 2012 and Windows 8 operating systems. To optimize WAN bandwidth, BranchCache downloads content from your content servers and caches the content at office locations, which enables client computers at office locations to access the content locally.

After a client computer has downloaded content one time, other clients that request the same content do not download it from the content servers over the WAN connection. Instead, they retrieve small identifiers, called *content information*, from the remote content servers. Clients use the content information to find the content in the local office. This content is cached on a computer running Windows Server or on other client computers, depending on the mode in which BranchCache has been deployed.

BranchCache modes

Hosted cache mode

When there is a server running Windows Server 2012 in the office location, BranchCache client computers are configured in hosted cache mode, and the server is called a *hosted cache server*.

Distributed cache mode

If your office does not have a server available to deploy as a hosted cache server, you can configure BranchCache in distributed cache mode on clients. In this mode, the client computers cache downloaded content and share it with other computers in the office.

BranchCache content servers

When you deploy BranchCache, you can deploy three types of content servers:

- **File server.** Supported file servers include computers that are running Windows Server 2012 or Windows Server 2008 R2 that have the File Services server role and the BranchCache for Network Files role service installed. These file servers use the server message block (SMB) protocol to exchange information. After you install your file server, you must also share folders and enable the generation of content hashes (also called *content information*) for shared folders. You do this by using Group Policy or Local Computer Policy to enable BranchCache.
- **Web server.** Supported web servers are computers that are running Windows Server 2012 or Windows Server 2008 R2 that have the Web Server (that is, Internet Information Services, or IIS) server role installed, and that use Hypertext Transfer Protocol (HTTP) or HTTP Secure (HTTPS). In addition, the web server must have the BranchCache feature installed.
- **Application server.** Supported application servers are computers that are running Windows Server 2012 or Windows Server 2008 R2 with Background Intelligent Transfer Service (BITS) installed and enabled. In addition, the application server must have the BranchCache feature installed.

BranchCache security

BranchCache security works more easily with your existing network security architectures. It does not require additional equipment or complex security configuration. BranchCache encrypts cached content and content information, and it does not allow unauthorized access to files in the cache. BranchCache can speed up encrypted communication through HTTPS or IPsec without compromising security.

BranchCache does not alter any Windows authentication or authorization processes. After you deploy BranchCache, authentication is still performed by using domain credentials, and authorization with Access Control Lists (ACLs) is unchanged. In addition, other configurations continue to function as they did before the BranchCache deployment.

BranchCache security is based on content information that is created on the content servers. This metadata, which is much smaller than the size of the actual content that it represents, takes the form of a series of hashes. After the content information is created, it is used in BranchCache message exchanges rather than the actual content, and it is exchanged by using the supported protocols (HTTP, HTTPS, and SMB).

BranchCache improvements

BranchCache in Windows Server 2012 and Windows 8 provide performance, manageability, scalability, and availability improvements over BranchCache in Windows Server 2008 R2 and Windows 7. This version of BranchCache provides the following improvements.

Larger branch office deployments

Multiple hosted cache servers

In the previous version of BranchCache, you could deploy only one hosted cache server per office location. Windows Server 2012 provides the ability to scale hosted cache mode deployments for offices of any size by allowing you to deploy as many hosted cache servers as are needed at a location.

Database performance

BranchCache now uses the Extensible Storage Engine (ESE) database technology that powers Microsoft Exchange Server. This enables a single hosted cache server that is running Windows Server 2012 to keep up with the demands of more people while using the same hardware that is utilized for a hosted cache server running Windows Server 2008 R2. It also allows a hosted cache server to store significantly more data (on the order of terabytes), which is necessary to provide high optimization for many people.

New tools and a simplified deployment model

- **Office configuration.** BranchCache no longer requires office-by-office configuration. There is no requirement for a Group Policy Object for each location, which streamlines deployment. A single Group Policy Object that contains a small group of settings is all that is required to deploy BranchCache in any size organization, from a large enterprise to a small business.
- **Computer configuration.** Client computer configuration is automatic. Computers can be configured through Group Policy as distributed cache mode computers by default. However, they will search for a hosted cache server. If one is discovered, clients automatically self-configure as hosted cache mode computers instead.
- **Data encryption.** Cache data is kept encrypted and hosted cache servers do not require server certificates. Previously, hosted cache servers were required to have a server certificate that was issued by a certification authority (CA) that client computers at the office location trusted. BranchCache security is improved with data encryption so that additional drive encryption technologies are no longer needed to protect cached data.
- **Manageability.** BranchCache is now manageable with Windows PowerShell and Windows Management Instrumentation (WMI). This enables scripting and remote management of BranchCache content servers, hosted cache servers, and client computers.
- **Preloaded content.** BranchCache provides tools to manipulate data and preload the content at remote locations ahead of time.
- **Integration with file server.** BranchCache is deeply integrated with the Windows-based file server, and it borrows the technology that is used to divide files into small pieces and eliminate duplicates. This increases the chance of finding duplicate pieces in independent files, resulting in greater bandwidth savings. BranchCache is also more tolerant of small changes in large files.
- **File calculations.** File division calculations are performed only one time, and they can be done offline. When a client computer that is running Windows 8 downloads content from a file server or web server that is running Windows Server 2012 and that is using disk deduplication technology, BranchCache does not spend CPU cycles calculating how to divide the content because the file server and web server have already made these calculations. Content information is calculated offline before a client computer requests a file. This provides faster performance and more bandwidth savings because content information is ready for the first client that requests the content, and the calculations have already been performed.

BranchCache requirements

You can deploy BranchCache in a domain-based or non-domain based environment.

You must provide network connectivity, such as a virtual private network (VPN) or DirectAccess connection, between the content servers and office locations so that client computers can access content on the content servers.

For distributed cache mode

You must deploy one or more content servers. For more information, see the "BranchCache content servers" section of this document.

In addition, you must deploy client computers that are running Windows 8 (or some editions of Windows 7). Client computers must have BranchCache distributed cache mode enabled.

For hosted cache mode

You must deploy one or more content servers as described previously in this section. For more information, see the "BranchCache content servers" section of this document.

In addition, you must deploy client computers that are running Windows 8 (or some editions of Windows 7). Client computers must have BranchCache hosted cache mode enabled; and you must deploy one or more hosted cache servers.

Note

BranchCache is supported on some editions of Windows Server 2008 R2 and Windows 7. For more information, see BranchCache on Microsoft TechNet.

Cluster-Aware Updating

Cluster-Aware Updating (CAU) is an automated feature that allows you to update clustered servers with little or no loss in availability during the update process. CAU transparently takes one node of the cluster offline, installs the updates, performs a restart if necessary, brings the node back online, and then moves on to the next node. This feature is integrated into the existing update management infrastructure in Windows Server 2012, and it can be further extended and automated with Windows PowerShell for integrating into larger IT automation initiatives.

Key benefit

Easily install updates to cluster nodes while maintaining availability.

Requirements

- Windows Server 2012
- Failover Clustering feature

Technical overview

CAU orchestrates the cluster updating operation while running from a computer running Windows Server 2012 or Windows 8. The computer running the CAU process is called an orchestrator.

The cluster updating orchestration includes the following:

1. Scanning for and downloading applicable updates on each cluster node.
2. Moving currently running clustered roles off each cluster node.
3. Ensuring that the cluster quorum is maintained throughout the cluster updating process.
4. Installing the updates on each cluster node.
5. Moving the clustered roles back to the original nodes.
6. Restarting cluster nodes if required by the installed updates.
7. Ensuring that all dependent updates are detected, downloaded, and installed on each cluster node, as determined by the previous set of updates installed on that cluster node.
8. Generating the necessary Updating Run reports and storing them on each cluster node.

The following figure shows the components involved in CAU.

Figure 3: CAU deployment to a file server cluster with an uninterrupted SMB client

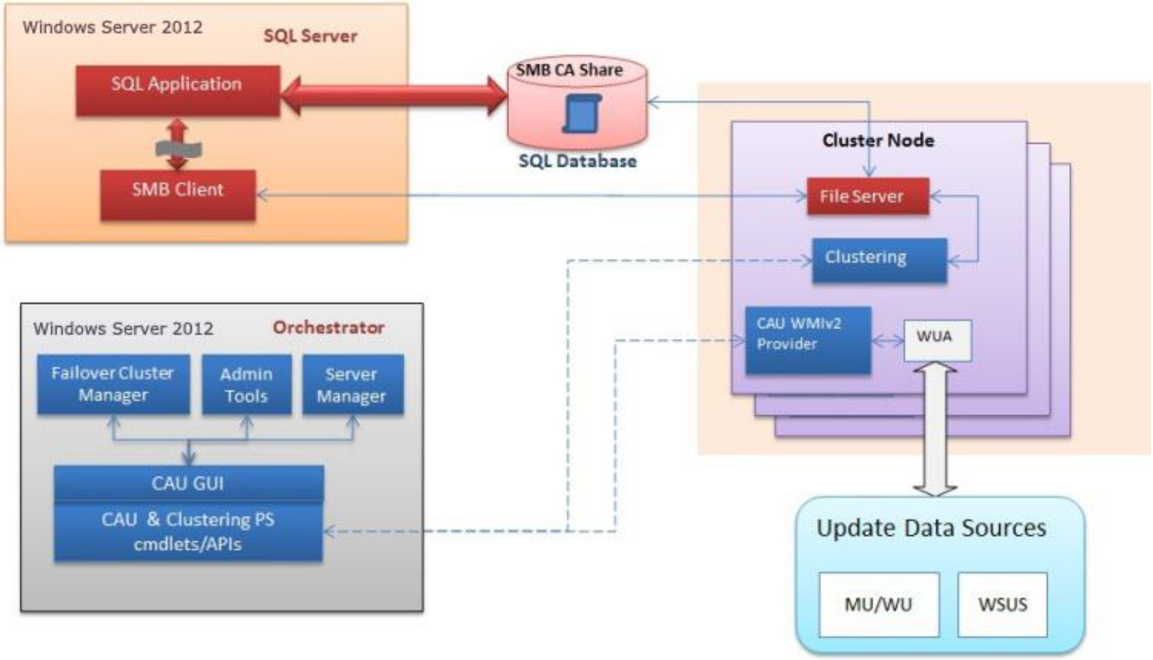

The end-to-end cluster update process through CAU is cluster-aware and completely automated, and it works more easily with existing Windows Update Agent (WUA) and Windows Server Update Services (WSUS) infrastructures. CAU also comes with an extensible architecture that supports new plug-in development to orchestrate node-updating tools—such as custom software installers, BIOS updating tools, and firmware updating tools for network adapters or host bus adapters (HBAs)—across all cluster nodes in a cluster-aware manner.

Note

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

Data Deduplication

The past decade has seen rapid growth in file-based data in enterprise environments. Although storage costs have been steadily dropping, they are not dropping fast enough to offset this growth, which makes storage efficiency a critical requirement for most enterprise IT departments. Further, efficiencies need to happen wherever the data is—whether it is sitting in a data store or moving through a wide area network. To cope with such growth, customers are rapidly consolidating file servers and making capacity scaling and optimization one of the primary requirements for a consolidation platform.

Data Deduplication is the act of finding and removing duplication within data without compromising its fidelity or integrity. Windows Server 2012 delivers innovative Data Deduplication, which provides the following benefits:

- **Capacity optimization.** Data Deduplication enables Windows Server 2012 to store more data in less physical space and gain storage efficiency that is significantly higher than was possible in previous releases of the Windows operating system that used Single Instance Storage (SIS) or NTFS file system compression. Data Deduplication uses variable size chunking and compression, which deliver optimization ratios of 2:1 for general file servers and up to 20:1 for virtualization data.
- **Scale and performance.** Data deduplication in Windows Server 2012 is highly scalable, resource efficient, and non-intrusive. It can run on dozens of large volumes of primary data simultaneously without impacting other workloads on the server. Low impact on the server workloads is maintained by throttling the CPU and memory resources that are consumed. In addition, users have the flexibility to set times when Data Deduplication should run, specify the resources available to deduplicate, and establish policies about file selection for Data Deduplication.
- **Reliability and data integrity.** When Data Deduplication is applied to data, the integrity of the data is maintained. Windows Server 2012 takes advantage of checksum values, consistency, and identity validation to ensure data integrity. In addition, for all the metadata and the most frequently-referenced data, Data Deduplication in Windows Server 2012 maintains redundancy to make sure that the data is recoverable if corruption occurs.
- **Bandwidth efficiency in conjunction with BranchCache.** Through integration with BranchCache, the same optimization techniques are applied to data that is transferred over the WAN to a branch office. This results in faster file download times and reduced bandwidth consumption.

Requirements

To take advantage of Data Deduplication in Windows Server 2012, the environment must meet the following requirements:

- **Server.** One computer running Windows Server 2012 or a virtual machine with at least one data volume
- **(Optional) Another computer.** One computer running Windows Server 2012 or Windows 8 that is connected to the server over a network

Technical overview

The goal of Data Deduplication is to store more data in less space by segmenting files into small (32-128 KB) and variable-sized chunks, identifying duplicate chunks, and then maintaining a single copy of each chunk. Redundant copies of the chunk are replaced by a reference to the single copy. In addition, chunks are also compressed for further space optimization.

The result is an on-disk transformation of each file as shown in Figure 4. Files are no longer stored as independent streams of data, but they are replaced with stubs that point to data blocks that are stored within a common store.

Figure 4: On-disk transformation of files during Data Deduplication

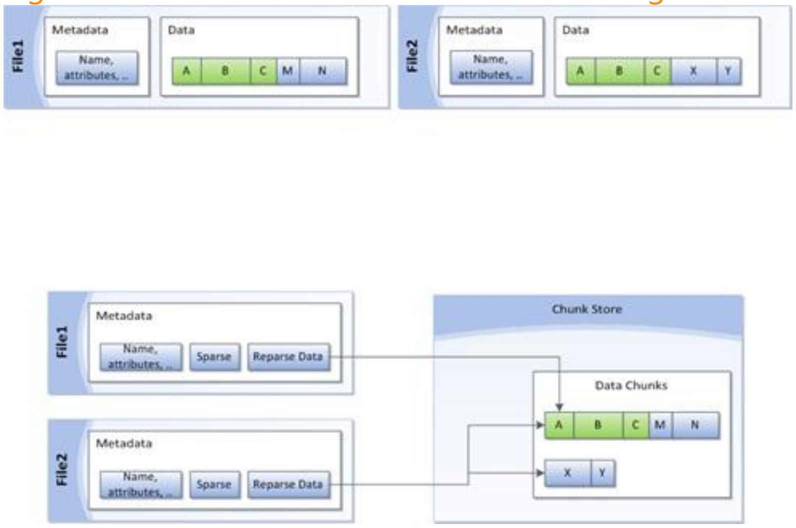

Datacenter Diskless Boot

The Microsoft iSCSI Software Target feature is built into Windows Server 2012, and it helps you create a Storage Area Network (SAN) device on hardware running the Windows operating system. iSCSI Software Target also enables you to network boot multiple computers from a single operating system image that is stored in a centralized location. This improves efficiency, manageability, availability, and security. iSCSI Software Target in Windows Server 2012 can boot hundreds of computers by using a single operating system image and it provides the following benefits:

- **Cost savings on operating system storage.** By using differencing virtual disks, you can use a single operating system image (the "golden image") to boot up to 256 computers. As an example, in a deployment of Windows Server 2008 R2 HPC edition, the operating system image is approximately 20 GB. A common deployment is to have two mirrored disk drives that act as the boot volume. Rounding the operating system storage to 40 GB per instance requires approximately 10 terabytes (TB) of storage—for only the operating system image—to boot 256 computers. With iSCSI Software Target boot, however, you will use 40 GB for the operating system base image and 2 GB for finding differences among virtual hard disks (VHDs) per server instance, totaling 552 GB for the operating system images. This provides a savings of over 90 percent on storage for the operating system images alone.
- **Increased security and simplified management with controlled operating system images.** Some enterprises require that data be secured by physically locking storage in a centralized location. In this scenario, servers access the data remotely, including the operating system image. With iSCSI Software Target boot, administrators can centrally manage the operating system boot images, and control which applications to put in the golden image.
- **Rapid deployment.** Because the golden image is an operating system image prepared using Sysprep, when the computers boot from the golden image, they skip the file copying and installation phase that occurs during Windows Setup, and they go straight to the customization phase. Microsoft testing resulted in 256 computers deployed in 34 minutes.
- **Fast recovery.** Because the operating system images are hosted on the iSCSI Software Target server, if the diskless client needs to be replaced, the new computer can point to the operating system image, and boot up immediately.

  A SAN boot is a solution that has been offered from various vendors. In Windows Server 2012, the new iSCSI Software Target feature provides this network boot capability on commodity hardware.

Requirements

iSCSI Software Target can be installed as part of the File Server role.

The easy management experience is provided through Server Manager. Corresponding Windows PowerShell cmdlets can be used for automation.

This feature does not require special hardware for functional verification. In datacenters with large-scale deployments, the design should be validated against specific hardware. For reference, Microsoft internal testing indicated that for a 256-iSCSI boot deployment, 24x15k-RPM (revolution per minute) disks in a RAID 10 configuration were required for storage. A network bandwidth of 10 GB is optimal. A general estimate is 60 iSCSI boot servers per 1 GB network adapter. However, an iSCSI boot-capable network adapter is not required for this scenario. If the network adapter does not support it, a software boot loader can be used (such as iPXE open source boot firmware).

Technical overview

The iSCSI Software Target feature in Windows Server 2012 supports diskless network boot without the need for special hardware or additional software. iSCSI Software Target fully complies with the iSCSI protocol in RFC 3720. A key component to this feature is that it supports differencing of virtual hard disks. This is critical in the boot scenario because multiple servers running Windows Server 2012 can boot by using only one base image.

iSCSI Software Target accomplishes diskless boot as follows:

- iSCSI Software Target supports creating differencing virtual disks based on a golden image.
- Each diskless client boots from its own differencing virtual disk.
- The diskless client reads from the golden image, and writes to its own differencing VHD.

Because the image is hosted in a centralized place, storage space savings are realized on operating system deployment. That can translate into freed-up storage space that was previously allocated to operating system boot images or budget savings because you will likely no longer need to purchase hard disks for this purpose and the power to cool them.

In Windows Server 2012, the scale of iSCSI boot has increased to 256. Therefore, the clustered target can support 256 boot clients, and it does not experience operating system errors with iSCSI Software Target service failover from one clustered node to another.

Dynamic Access Control

In Windows Server 2012, you can apply data governance across your file servers to control who can access information and to audit who has accessed information. Dynamic Access Control helps you:

- Identify data by using automatic and manual classification of files. For example, you can tag data in file servers across the organization.
- Control access to files by applying safety-net policies that use central access policies. For example, you could define who can access health information within the organization.
- Audit access to files by using central audit policies for compliance reporting and forensic analysis. For example, you could identify who accessed highly sensitive information.
- Apply Rights Management Services (RMS) protection by using automatic RMS encryption for sensitive Microsoft Office documents. For example, you could configure RMS to encrypt all documents that contain Health Insurance Portability and Accountability Act (HIPAA) information.

This feature set is based on infrastructure investments that can be further used by partners and line-of-business applications, and the features can provide great value for organizations that use Active Directory. This infrastructure includes:

- A new authorization and audit engine for Windows that can process conditional expressions and central policies
- Kerberos authentication support for user claims and device claims
- Improvements to the File Classification Infrastructure (FCI)
- RMS extensibility support so that partners can provide solutions that encrypt non-Microsoft files

Key benefit

You can comply with business and regulatory standards by using user claims and file classifications to centrally control access, audit access, and use RMS to protect information in files.

Requirements

This feature requires the following:

- Windows Server 2012
- Active Directory Domain Services
- File Server
- Rights Management Server (required only for the RMS extensibility support scenario)

Technical overview

Windows Server 2012 provides the following new ways to control access to your files while providing authorized users the resources they need:

- Central access control for information governance
- Policy changes and staging
- File access auditing for forensic analysis and compliance
- Access-denied remediation to troubleshoot file and shared folder access problems
- Classification-based encryption for sensitive Microsoft Office documents

The following sections describe these features in more detail.

Central access control

Central access policies for files help organizations to centrally deploy and manage authorization policies that include conditional expressions with user claims, device claims, and resource properties. Claims are assertions about the attributes of the object with which they are associated. For example, for accessing high business impact (HBI) data, a user must be a full-time employee, obtain access from a managed device, and log on with a smart card. These policies are defined and hosted in Active Directory.

The various organizational access policies are driven by compliance and business regulatory requirements. For example, if an organization has a business requirement to restrict access to personally identifiable information (PII) in files to only the file owner and members of the human resources (HR) department that are allowed to view PII information, this is an organization-wide policy that applies to PII files wherever they are located on file servers across the organization. In this example, you need to be able to:

- Identify and mark the files that contain PII.
- Identify the group of HR members who are allowed to view PII information.
- Have a central access policy that can easily be applied to all files that contain PII wherever they are located among the file servers across the organization.

The initiative to deploy and enforce an authorization policy may come for many reasons and from multiple levels of the organization. The following are some example policies:

- **Organization-wide authorization policy.** Most commonly initiated from the information security office, this authorization policy is driven from compliance or a high-level organization requirement and would be relevant across the organization. For example, HBI files should be accessible to only full-time employees.
- **Departmental authorization policy.** Each department in an organization has some special data-handling requirements that they want to enforce. For example, the finance department might want to limit access to finance servers to the finance employees.
- **Specific data-management policy.** This policy usually relates to compliance and business requirements, and it is targeted at protecting the correct access to the information that is being managed, such as preventing, modifying, or deleting files that are under retention or files that are under electronic discovery (e-discovery).
- **Need-to-know policy.** This is an authorization policy type and is typically used in conjunction with the policy types mentioned earlier. Examples include:
    - Vendors should be able to access and edit only files that pertain to a project they are working on.

- Financial institutions should implement information walls so that analysts do not access brokerage information and brokers do not access analysis information.

Real-life environments also teach us that every authorization policy needs to have exceptions so that organizations can quickly react when important business needs arise. For example, executives who cannot find their smart cards and need quick access to HBI information can call the Help Desk to get a temporary exception to access that information.

Central access policies act as security umbrellas that an organization applies across its servers. These policies enhance (but do not replace) the local access policies or discretionary access control lists (DACL) that are applied to files and folders. For example, if a DACL on a file allows access to a specific user, but a central policy that is applied to the file restricts access to the same user, the user cannot obtain access to the file.

A central policy rule has the following logical parts:

- **Applicability.** A condition that defines which data the policy applies to, such as Resource.BusinessImpact=High.
- **Access conditions.** A list of one or more access control entries (ACEs) that define who can access the data, such as Allow | Full Control | User.EmployeeType=FTE.
- **Exception.** An additional list of one or more ACEs that define an exception for the policy, such as MemberOf(HBIExceptionGroup).

The following two figures show the moving parts in a central access and audit policy.

Figure 5: Central access and audit policy concepts

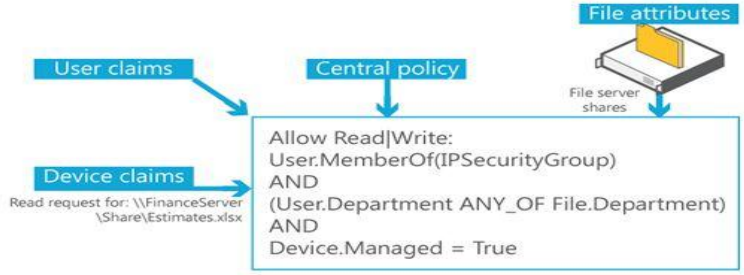
File attributes
User claims
Central policy
File server shares
Device claims
Read request for: \\FinanceServer \Share\Estimates.xlsx
Allow Read|Write:
User.MemberOf(IPSecurityGroup)
AND
(User.Department ANY_OF File.Department)
AND
Device.Managed = True

Figure 6: Central access policy workflow

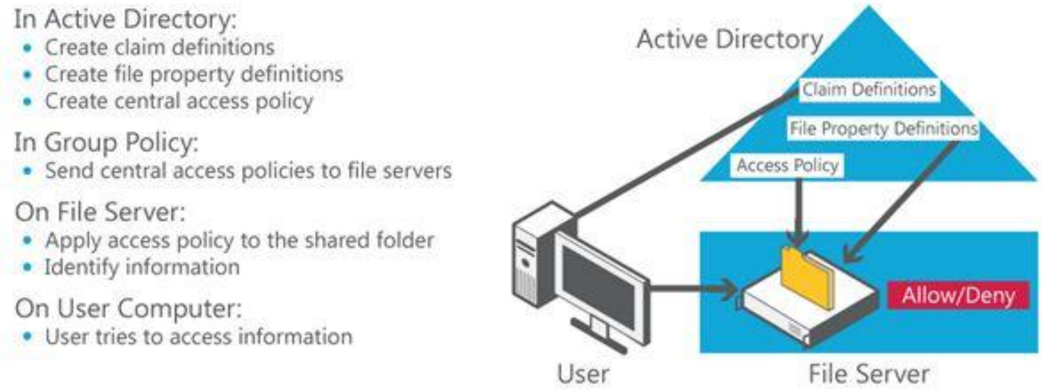

In Active Directory:
- Create claim definitions
- Create file property definitions
- Create central access policy

In Group Policy:
- Send central access policies to file servers

On File Server:
- Apply access policy to the shared folder
- Identify information

On User Computer:
- User tries to access information

The central authorization policy combines the following components:

- Centrally defined access rules that target specific types of information, such as HBI or PII.
- A centrally defined policy that contains a list of rules.
- A policy identifier that is assigned to each file on the file servers to point to a specific central list of policies that should be applied during the access authorization.

The following figure demonstrates how you can combine policies into policy lists to centrally control access to files.

Figure 7: Combining policies

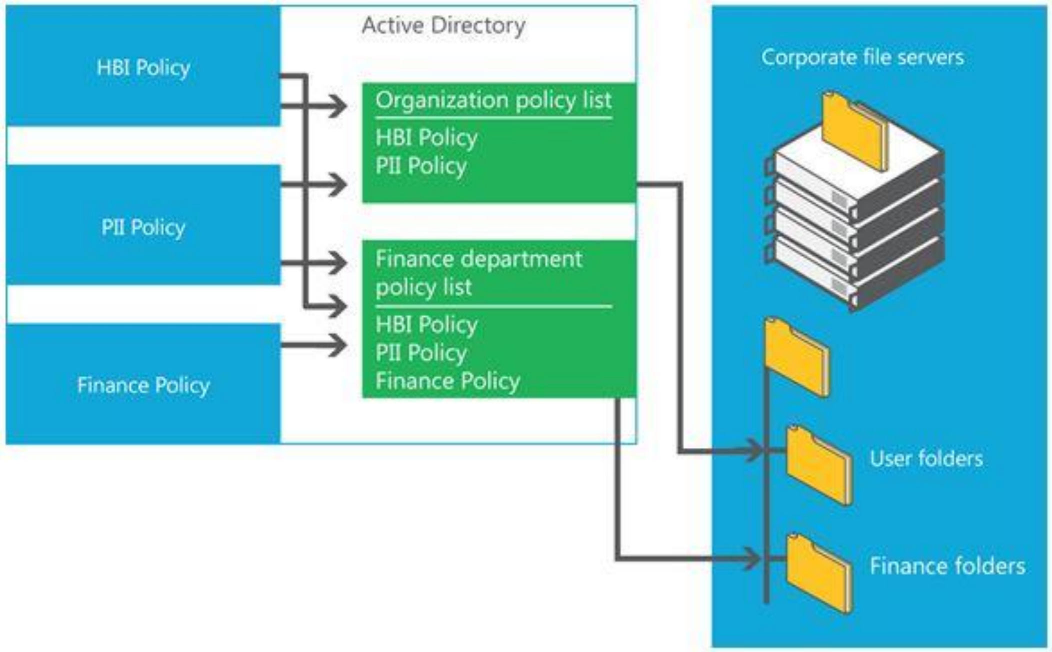

Policy changes and staging

When you want to change a policy, Windows Server 2012 allows you to use a proposed policy that runs parallel to the current access policy so that you can identify the consequences of the new policy without enforcing it. This feature, called policy staging, lets you measure the effects of a new access policy on your production environment.

When policy staging is enabled, Windows Server 2012 continues to use the current access policy to authorize a user’s access to files. However, if the result of the proposed policy is different than the access result of the current policy, the system logs an event with the details. You can use the logged events to determine whether the policy needs to be changed or is ready to be deployed.

File access auditing

Security auditing is one of the most powerful tools to help maintain the security of an enterprise. One of the key goals of security audits is regulatory compliance. For example, industry standards such as Sarbanes Oxley, HIPAA, and Payment Card Industry (PCI) require enterprises to follow a strict set of rules related to data security and privacy. Security audits help establish the presence or absence of such policies; thereby, proving compliance or noncompliance with these standards. Additionally, security audits help detect anomalous behavior, identify and mitigate gaps in security policy, and deter irresponsible behavior by creating a record of user activity that can be used for forensic analysis.

Audit policy requirements are typically driven at the following levels:

- **Information security.** File access audit trails are often used for forensic analysis and intrusion detection. Being able to get targeted events about access to high-value information lets organizations considerably improve their response time and investigation accuracy.
- **Organizational policy.** Organizations regulated by PCI standards could have a central policy to monitor access to all files that are marked as containing credit card information and PII.
- **Departmental policy.** Finance departments may want to restrict the ability to modify certain finance documents (such as a quarterly earnings report) to their own department, and thus would want to monitor all other attempts to change these documents.
- **Business policy.** Business owners may want to monitor all unauthorized attempts to view data belonging to their projects.

Additionally, the compliance department may want to monitor all changes to central authorization policies and policy constructs such as user, computer, and resource attributes.

One of the biggest considerations of security audits is the cost of collecting, storing, and analyzing audit events. If the audit policies are too broad, the volume of audit events collected rises, and this increases costs. If the audit policies are too narrow, you risk missing important events.

With Windows Server 2012, you can author audit policies by using claims and resource properties. This leads to richer, more targeted, and easier-to-manage audit policies. It enables scenarios that, until now, were either impossible or too difficult to perform. The following are examples of audit policies that administrators can author:

- Audit everyone who does not have a high security clearance and yet tries to access an HBI document. For example, Audit | Everyone | All-Access | Resource.BusinessImpact=HBI AND User.SecurityClearance!=High.
- Audit all vendors when they try to access documents that are related to projects they are not working on. For example, Audit | Everyone | All-Access | User.EmploymentStatus=Vendor AND User.Project Not_AnyOf Resource.Project.

These policies help regulate the volume of audit events and limit them to only the most relevant data or users.

After administrators have created and applied the audit policies, the next consideration for them is gleaning meaningful information from the audit events that they collected. Expression-based audit events help reduce the volume of audits. However, users need a way to query these events for meaningful information and ask questions such as, “Who is accessing my HBI data?” or “Was there an unauthorized attempt to access sensitive data?”

Windows Server 2012 enhances existing data access events with user, computer, and resource claims. These events are generated on a per-server basis. To provide a full view of events across the organization, Microsoft is working with partners to provide event collection and analysis tools, such as System Center Operation Manager Audit Collection Service (SCOM/ACS).

The figure below shows the moving parts of a central audit policy.

Figure 8: Central auditing experiences

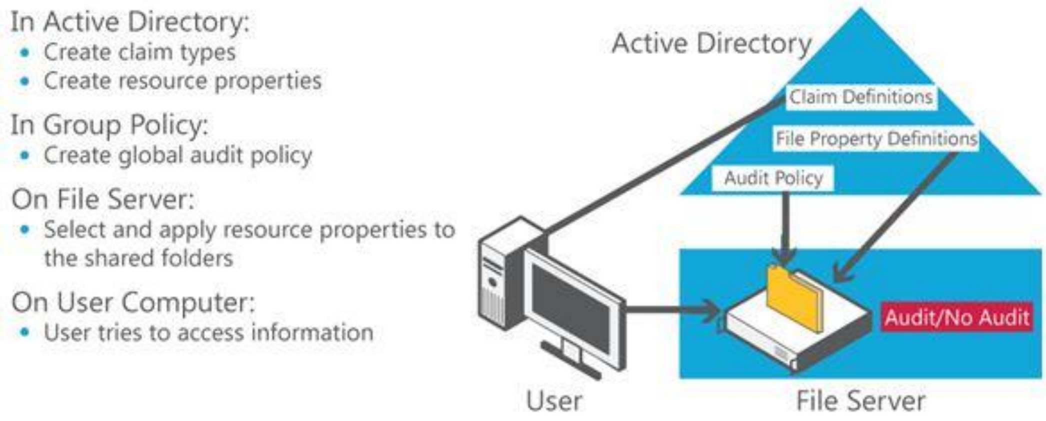
In Active Directory:
Create claim types
Create resource properties
In Group Policy:
Create global audit policy
On File Server:
Select and apply resource properties to the shared folders
On User Computer:
User tries to access information
Active Directory
Claim Definitions
File Property Definitions
Audit Policy
Audit/No Audit
User
File Server

Setting up and consuming security audits typically involves the following general steps:

1. Identifying the correct set of data and users to monitor.
2. Creating and applying appropriate audit policies.
3. Collecting and analyzing audit events.
4. Managing and monitoring the policies that were created.

Access-denied remediation

Access-denied remediation in Windows Server 2012 provides three processes to grant users access to the resources they need:

- **Self-remediation.** If users can determine what the issue is and remediate the problem so that they can get the requested access, the impact to the business is low and no special exceptions are needed in the organization access policy. Windows Server 2012 provides a general access-denied message that is authored by the server administrator for users so that they can try to self-remediate access-denied cases. This message can also include URLs to direct the users to self-remediation websites that are provided by the organization.
- **Remediation by the data owner.** Windows Server 2012 allows administrators to define owners of shared folders in the form of a distribution list so that users can directly connect with the data owners to request access. This is similar to the Microsoft SharePoint model where the data owner gets a request from the user to gain access to the files. In the File Server case, the remediation can range from adding the user rights for the appropriate file or directory to dealing with shared folder permissions. For example, if a DACL on a file allows access to a specific user, but a central policy restricts access to the same user, the user will not be able to gain access to the file and global policies. Users can request access, which results in sending an email with the request details to the data owner. Data owners can forward this information to the appropriate IT administrator if they cannot remediate the issues.
- **Remediation by Help Desk and file server administrators.** This type of remediation happens when the user cannot self-remediate the issue and the data owner cannot help. This is the most costly and time-consuming remediation. Windows Server 2012 provides a user interface where administrators can view the effective permission for users for a file or folder so that it is easier to troubleshoot access issues.

The following figure shows the steps that can be taken to remediate an issue if a user is denied access to a resource.

Figure 9: Remediation options

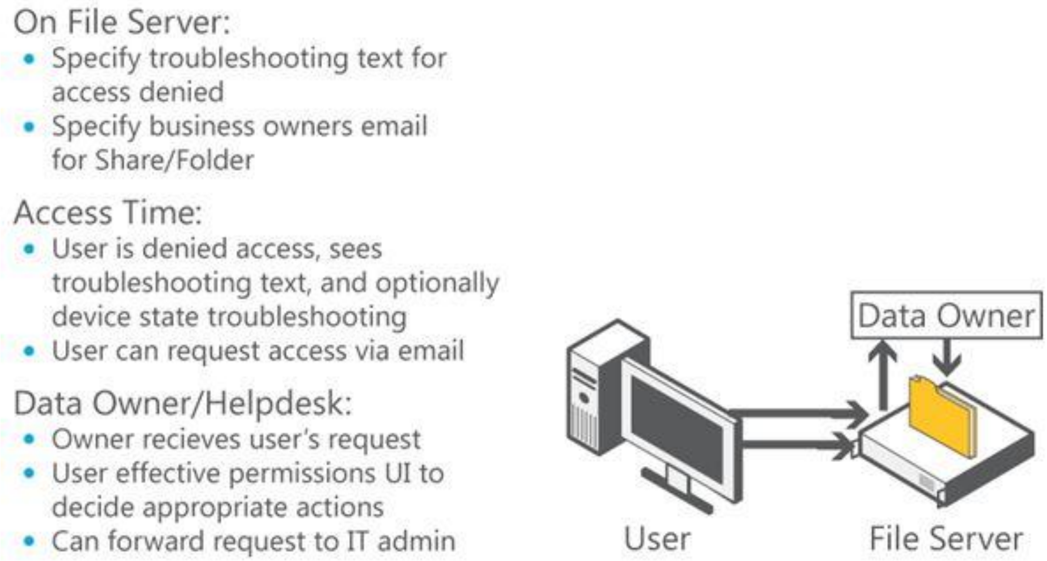

On File Server:

- Specify troubleshooting text for access denied
- Specify business owners email for Share/Folder

Access Time:

- User is denied access, sees troubleshooting text, and optionally device state troubleshooting
- User can request access via email

Data Owner/Helpdesk:

- Owner recieves user's request
- User effective permissions UI to decide appropriate actions
- Can forward request to IT admin

To make it easier on Help Desk and IT administrators, it is important to provide them with all the relevant access details and tools (such as effective access permissions); so that they can determine the issue and then make the configuration changes to satisfy the access request.

For example, users might follow this process to access a file that they currently do not have access to:

- The user attempts to read a file at \\financeshares, but the server returns an Access Denied error.
- Windows displays the **Access Denied** dialog box.
- Windows retrieves the access remediation information from Remediation Service on the file server.
- Windows displays access remediation options to the user with an option to request access.
- The user requests access to the resource.
- The server sends an email with access request information to the resource owner.

Applications such as those included in Microsoft Office can also take advantage of the new enhancements by directly using the new Windows UI or by calling the new access-denied remediation API.

Classification-based encryption for Windows Office documents

Protection of sensitive information is mainly about mitigating risk for the organization. Various compliance regulations, such as HIPAA or Payment Card Industry Data Security Standard (PCI-DSS), dictate encryption of information, and there are numerous business reasons to encrypt sensitive business information. However, encrypting information is expensive, and it might impair business productivity. Thus, organizations tend to have different approaches and priorities for encrypting their information.

To support this scenario, Windows Server 2012 provides the ability to automatically encrypt sensitive Windows Office files based on their classification. This is done through file management tasks that invoke RMS protection for sensitive documents a few seconds after the file is identified as being a sensitive file on the file server (continuous file management tasks on the file server).

RMS encryption provides another layer of protection for files. Even if a person with access to a sensitive file inadvertently sends that file through email, the file is protected by RMS encryption. Users who want to access the file must first authenticate themselves to an RMS server to receive the decryption key. The following figure shows this process.

Figure 10: Classification-based RMS protection

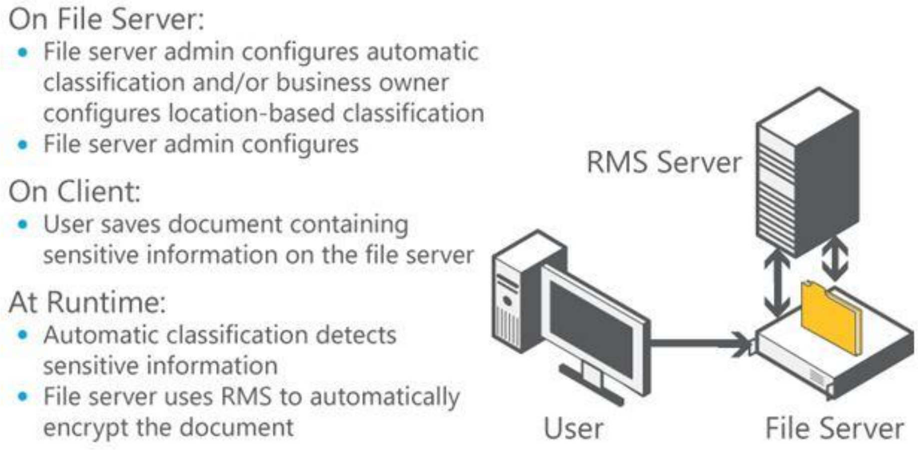

On File Server:

- File server admin configures automatic classification and/or business owner configures location-based classification
- File server admin configures

On Client:

- User saves document containing sensitive information on the file server

At Runtime:

- Automatic classification detects sensitive information
- File server uses RMS to automatically encrypt the document

Support for non-Microsoft file formats is available through non-Microsoft vendors. After a file has been protected by RMS encryption, data management features such as search- or content-based classification are no longer available for that file.

Encrypted Hard Drive

The Encrypted Hard Drive feature in Windows Server 2012 uses the rapid encryption that is provided by Microsoft BitLocker Drive Encryption to enhance data security and management. By offloading the cryptographic operations to hardware, Encrypted Hard Drive increases BitLocker performance and reduces CPU usage and power consumption. Because Encrypted Hard Drive encrypts data quickly, enterprise clients can expand BitLocker deployment with minimal impact on productivity.

Key benefit

Rapid encryption and decryption for BitLocker is the key benefit of the Encrypted Hard Drive feature.

Requirements

This feature requires the following:

- Windows 8 or Windows Server 2012
- For data disk – A disk that has not been provisioned and that supports Encrypted Hard Drive
- For startup disks – A disk that has not been provisioned and that supports Encrypted Hard Drive, and a computer that includes a Unified Extensible Firmware Interface (UEFI) (Class II or "Native" UEFI) with UEFI 2.3.1 support

Technical overview

Rapid encryption in BitLocker directly addresses the security needs of enterprises while offering significantly improved performance. In earlier versions of Windows, BitLocker required a two-step process to complete read/write requests, but in Windows Server 2012, Encrypted Hard Drive offloads the cryptographic operations to the drive controller for much greater efficiency. When Windows Server 2012 initializes Encrypted Hard Drive, it activates the security mode. This activation lets the drive controller generate a media key for every volume that the host computer creates. This media key, which is never exposed outside the disk, is used to rapidly encrypt or decrypt every byte of data that is sent or received from the disk.

Failover Clustering

Failover clustering provides high availability and scalability that can be applied to many different workloads that are running on Windows Server 2012. Clusters now support up to 4,000 virtual machines and up to 64 nodes. In addition, improvements in the Failover Cluster Manager snap-in make it easier to manage high availability services in a large network.

Key benefits

In Windows Server 2012, clusters provide increased scalability, easier management, faster failover, and more flexible architectures.

Requirements

- Windows Server 2012
- A network infrastructure that connects the nodes (servers) in the cluster and avoids having single points of failure
- Storage that is attached to nodes in the cluster
- Device controllers or appropriate adapters for the storage. These can be Serial Attached SCSI (SAS), Fibre Channel, or iSCSI

Technical overview

Windows Server 2012 includes the following enhancements to failover clustering:

- **Clustering improvements for Hyper-V and virtual machines.** Failover clustering now supports up to 4,000 virtual machines, and the improved Failover Cluster Manager snap-in makes it simpler to manage large numbers of virtual machines. For example, administrators can now perform large multiple-select actions to queue live migrations of multiple virtual machines. Administrators can also configure virtual machine priorities to control the order in which virtual machines are started and to ensure that lower-priority virtual machines automatically release resources if they are needed by higher priority virtual machines. Services running on clustered virtual machines can now be monitored. In addition, the Cluster Shared Volume (CSV) feature, which simplifies the configuration and operation of virtual machines, has been improved for greater security and performance. For more information about these and other related improvements, see the following sections of this document:
  - Clustering improvements for managing Hyper-V and virtual machines
  - Efficient automatic management of clustered virtual machines and other clustered roles
  - Monitoring of services that are running on clustered virtual machines
  - Innovations in Cluster Shared Volumes
  - Cluster validation improvements
  - Enhanced Windows PowerShell support for clusters

- **Clustering improvements for file servers.** For greater scalability of clustered file servers, a single failover cluster can now have as many as 64 nodes. Also, improvements to the validation wizard and the migration wizard in failover clustering make it easier to set up clustered file servers and to migrate existing clustered file servers to new clusters. Enhanced options such as BitLocker Drive Encryption for cluster disks and improved backup options can also be important for clustered file servers. For more information about these and other related improvements, see the following sections of this document:
    - Cluster validation improvements
    - Enhanced wizard for migrating settings from one cluster to another
    - Security and encryption support for traditional clustered disks and CSVs
- **Clustering improvements for continuously available file share storage for server applications (such as Hyper-V and SQL Server).** These improvements are described in the "High-Performance, Continuously Available File Share Storage for Server Applications" section of this document.
- **Improvements in cluster validation and deployment.** More tests are available in the cluster validation wizard. The wizard for migrating settings from one cluster to another has been enhanced, and Sysprep is now supported for deploying cluster nodes. For more information, see the "Improvements in cluster validation and deployment" section of this document.
- **Enhanced Windows PowerShell support for clusters.** Failover clusters in Windows Server 2012 include support for new Windows PowerShell cmdlets. For more information, see the "Enhanced Windows PowerShell support for clusters" section of this document.
- **Improvements in multisite failover clusters for failure recovery preparedness.** Multisite clusters can now use more flexible quorum options. Also, if a multisite failover cluster is configured to support virtual machines by using CSV, CSV traffic can now stream across multiple networks, which delivers improved I/O performance when in redirected mode. For more information, see the "Clustering improvements for multisite clusters" section of this document.

Clustering improvements for managing Hyper-V and virtual machines

The updated Failover Cluster Manager snap-in simplifies large-scale management of clustered virtual machines and other clustered roles. The new features include the following:

- **Features for managing large numbers of virtual machines or other clustered roles.** Search, filtering, and custom views in Failover Cluster Manager make it easier to manage large numbers of clustered virtual machines or other clustered roles.
- **Multiple-select feature.** Selecting multiple virtual machines and performing operations on them is now possible. Administrators can easily select a specific collection of virtual machines and then perform any needed operation on them (for example, live migration, save, shutdown, or start).
- **Simplified live migration and quick migration of virtual machines.** Live migration and quick migration are easier to perform from within Failover Cluster Manager.
- **New settings for queuing and priority setting for clustered applications, including virtual machines.** Failover Cluster Manager provides new settings for clustered virtual machines, such as queuing of live migrations, virtual machine priority, and preemption of lower priority virtual machines by higher priority virtual machines. For more information, see the "Efficient automatic management of clustered virtual machines and other clustered roles" section of this document.

- **Simpler configuration of Cluster Shared Volumes (CSV).** Configuring CSV is as simple as right-clicking in the Storage pane. CSV has additional enhancements, described in "Innovations in Cluster Shared Volumes".

Efficient automatic management of clustered virtual machines and other clustered roles

In a failover cluster that is running Windows Server 2012, administrators can configure settings, such as the relative priority of different virtual machines and other clustered roles, and then allow the clustered roles to be managed automatically by the cluster, so that manual intervention is rarely needed. The following table describes these settings:

Table 1: Failover cluster settings

Setting	Scope
Priority setting: High, Medium, Low, or No Auto Start	All clustered roles, including clustered virtual machines
Preemption of virtual machines based on priority	Clustered virtual machines
Memory-aware virtual machine placement	Clustered virtual machines
Queuing live migrations	Clustered virtual machines
Automated node draining	All clustered roles, including clustered virtual machines

The following sections provide more details about automatic management.

Priority settings for clustered virtual machines and other clustered roles

Administrators can now control the way in which the cluster handles virtual machines and other clustered roles by assigning a priority to each clustered role. The possible priorities are:

- High
- Medium
- Low
- No Auto Start

When a clustered role is created, the default priority is Medium.

By assigning priorities to clustered roles, administrators can influence:

- **Start order of roles.** Virtual machines or clustered roles with higher priority are started before those with lower priority.
- **Placement order of roles.** Virtual machines or clustered roles with higher priority are placed on appropriate nodes before virtual machines or clustered roles with lower priority.

When the whole cluster is restarted, multiple roles must be placed on multiple nodes in the cluster. If a node crashes or is evicted, multiple roles must be placed on the remaining nodes in the cluster. The placement order of these roles is determined by their priority setting.

If a No Auto Start priority is assigned to a clustered role, the role does not start automatically (that is, it does not come online) after it fails, which keeps resources available so other roles can start.

Preemption of virtual machines

In large-scale deployments and clusters that can include as many as 64 nodes, hundreds or even thousands of virtual machines can be configured in a cluster. If a network or node failure occurs or if a node must be taken offline, the workload must be supported by a reduced number of nodes. On an overcommitted cluster, high-priority virtual machines might not find enough available memory resources or other resources to successfully start.

In Windows Server 2012, administrators can configure their cluster so that higher priority virtual machines are the first ones started, so that they more easily obtain the memory and other resources they need. A lower priority virtual machine, which the cluster attempts to start later, cannot take resources from a higher priority virtual machine.

If high-priority virtual machines cannot find the necessary memory resources or other resources, the cluster service preempts (takes offline) lower priority virtual machines. This frees up resources. These freed-up resources can now be assigned to high-priority virtual machines to enable them to start successfully.

The preemption of virtual machines starts with the lowest priority virtual machines and continues, if necessary, with virtual machines in the next priority level. Also, the number of virtual machines that are preempted is kept to a minimum, which minimizes the impact of preemption.

Any virtual machines that are preempted are later restarted in priority preference order. This helps ensure that higher priority virtual machines can successfully obtain the memory and other resources that they need.

NUMA—memory-aware virtual machine placement

Failover clustering is an important infrastructure for the private cloud, and it can support large-scale deployments of virtual machines across large clusters. It becomes very important to ensure that virtual machines are placed on cluster nodes that can successfully host them. This includes placing virtual machines appropriately on physical servers that have multiple processors and use Non-Uniform Memory Access (NUMA).

In Windows Server 2008 R2, the operating system could obtain information on the NUMA configuration, and it could optimize memory allocation for various applications based on that information. In Windows Server 2012, Hyper-V extends that NUMA support to virtual machines.

With this feature, a failover cluster can place virtual machines appropriately after it evaluates the NUMA configuration of a given node, the workload already running on the node, and the available resources on the node. Before choosing the node on which to place the virtual machines, the cluster ensures that the node has enough available resources to successfully start the virtual machines. If no nodes exist that can host a higher priority virtual machine, the failover cluster tries to find nodes with enough lower priority virtual machines that could be preempted, which makes it possible to start the higher priority virtual machines.

The increased NUMA capabilities in Windows Server 2012 reduce the number of failover attempts before a virtual machine is successfully started, therefore increasing the uptime for virtual machines.

Queuing live migrations of virtual machines

In Windows Server 2012, Hyper-V supports multiple concurrent live migrations of virtual machines. Failover clustering extends this support by letting administrators queue live migrations of virtual machines. They can use the multiple-select feature (mentioned previously) and then initiate live migration of multiple virtual machines at the same time. The cluster initiates as many live migrations as possible, and then queues the remaining live migrations for later completion. If necessary, the cluster would also retry queued live migrations that did not at first succeed.

Queuing live migrations is integrated with NUMA-Memory-Aware virtual machine placement. As each live migration is initiated, the cluster selects the best possible node for that virtual machine.

In Failover Cluster Manager, you can view the detailed status of ongoing or queued live migrations.

Automated node draining

Before putting a node into maintenance mode and making necessary changes on the node, the clustered services and applications (also called clustered roles) that are running on that node must be moved to another node. In Windows Server 2008 R2, part of this process was manual. In Windows Server 2012, the process is automated.

Administrators can drain a cluster node with a single click in Failover Cluster Manager or drain the node by using the Windows PowerShell cmdlet, *Suspend-ClusterNode*. They can also specify the node to which they want to move all the workloads. In addition, administrators can use a related feature in Windows Server 2012, Cluster-Aware Updating (CAU), to apply software updates to cluster nodes. CAU uses an automated process to drain a node, apply software updates, move workloads back to the node, and repeat the process for each node until the entire cluster is updated.

As part of the process of draining, the failover cluster tries to live migrate or quickly migrate virtual machines, and it uses queuing and NUMA-related placement logic to place virtual machines appropriately on the cluster nodes.

Monitoring of services that are running on clustered virtual machines

For high availability of services running on clustered virtual machines, the state of those services is important, and in many cases needs to be monitored along with the state of the clustered virtual machines themselves. In clusters running Windows Server 2012, administrators can configure monitoring of these services on clustered virtual machines that are also running Windows Server 2012. If the service being monitored fails, the clustered virtual machine may be restarted or moved to another node, depending on service restart settings and cluster failover settings. This increases the uptime of these services.

For information about the Windows PowerShell cmdlets for monitoring services running on clustered virtual machines, see the "Enhanced Windows PowerShell support for clusters" section of this document.

Innovations in Cluster Shared Volumes

The Cluster Shared Volumes (CSV) feature was introduced in Windows Server 2008 R2 to simplify the configuration and management of clustered virtual machines. With CSV, multiple clustered virtual machines can use the same LUN (disk) and still live migrate or quick migrate from node to node independently of one another.

CSV overview

CSV is a distributed file access solution that provides multiple nodes in the cluster with simultaneous access to the same file system. Virtual machines or applications that run on CSV are no longer bound to storage, and they can share a common disk to reduce the number of LUNs, as shown in the following figure. Live migration of virtual machines becomes faster because volume ownership does not need to change.

Figure 11: Cluster Shared Volumes

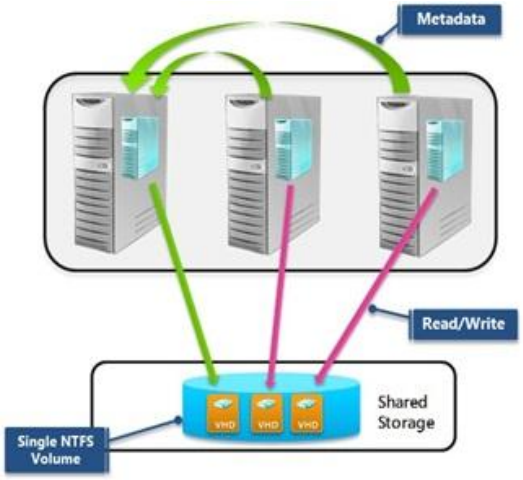

CSV does not impose any special hardware, file type, or directory structure restrictions. CSV uses the well-established NTFS file system.

What CSV delivers in Windows Server 2012

In Windows Server 2012, CSV supports new features in addition to the features that are supported in Windows Server 2008 R2.

CSV continues to support virtual machines, and it has been enhanced as follows:

- Expansion of CSV to support more workloads including File Server for scale-out application data, and possibly more
- Enhanced backup and restore of CSV volumes
- Integration with new file system features such as copy-offload, defragmenting of CSV volumes, and online corruption repair
- Multiple-subnet support
- Simplified set up for CSV by using Failover Cluster Manager
- Support for Server Message Block 3.0 (SMB) file-based storage for Hyper-V
- Better virtual machine creation and copying experience with copying that can be performed from nearly any node with the same high performance
- Support for BitLocker volume encryption for CSV volumes
- Improved I/O redirection to support block-level I/O and increased performance that results from support for multiple subnets
- Improved CSV performance by allowing direct I/O for more scenarios
- Support for memory mapped files

CSV proxy file system

CSV volumes now appear as CSVFS. The underlying technology is still the NTFS file system, and volumes are still formatted with NTFS. Because volumes appear as CSVFS, applications can discover that they are running on CSV, which helps improves compatibility.

CSV provides a single consistent file name space. Files have the same name and path when viewed from any node in the cluster. CSV volumes are exposed as directories and subdirectories under the ClusterStorage root directory.

In Windows Server 2012, CSV uses standard mount points for better interoperability with:

- Performance counters.
- System Center Operations Manager.
- Monitoring free space on CSV volumes.
- File system minifilter drivers (to allow improved interoperability with antivirus software and backup software).

Supports multiple subnet for Cluster Shared Volumes

In Windows Server 2012, CSV has been enhanced to integrate with SMB Multichannel. This helps you attain faster throughput for CSV volumes, especially when they are working in redirected mode. Applications can still have fast data access even with CSV in redirected mode.

CSV traffic can now stream across multiple networks, which delivers improved I/O performance when in redirected mode. CSV also takes advantage of SMB Direct (SMB over Remote Direct Memory Access) on network adapters that support this.

Figure 12: CSV streaming across multiple networks

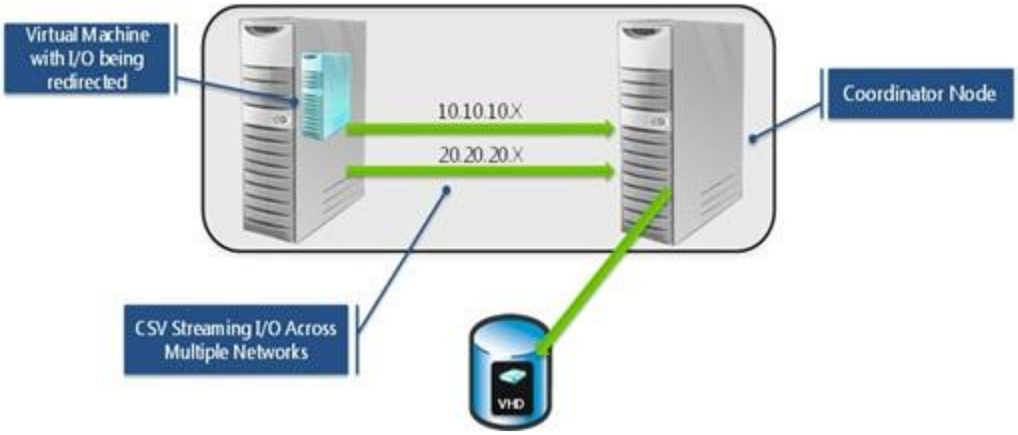
Virtual Machine with I/O being redirected
10.10.10.X
20.20.20.X
Coordinator Node
CSV Streaming I/O Across Multiple Networks
VHD

Security and encryption support for traditional clustered disks and CSVs

Failover clusters that are running Windows Server 2012 support BitLocker volume encryption for both traditional clustered disks and CSVs. Each node performs decryption by using the computer account for the cluster itself. This account is also called the Cluster Name Object (CNO). This action enables physical security for deployments outside secure datacenters and meets compliance requirements for volume-level encryption.

Figure 13: Encryption support

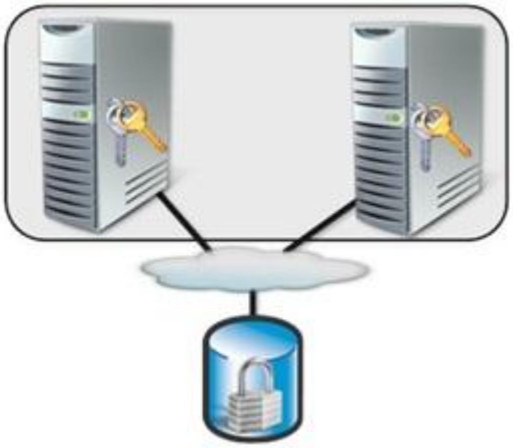

Ease of file backup with CSV

In Windows Server 2012, CSV helps provide a seamless backup experience for both backup applications and backup requestors. CSV includes:

- Distributed CSV backup infrastructure for Software Snapshot and coordination of CSV snapshot creation.
- Support for parallel backups across CSV volumes and across cluster nodes.
- Application-consistent and crash-consistent Volume Shadow Copy Service (VSS) snapshot support with a full feature set.
- Increased support for both hardware and software backup of CSV volumes.
- Full compatibility support for requestors that are running Windows Server 2008 R2.
- Direct I/O mode for snapshot and backup operations.
- Support for using Windows Server Backup (included in Windows Server 2012) for CSV backups.
- Support for incremental backups.

The following figure shows application-consistent and crash-consistent backup on CSV.

Figure 14: Seamless backup

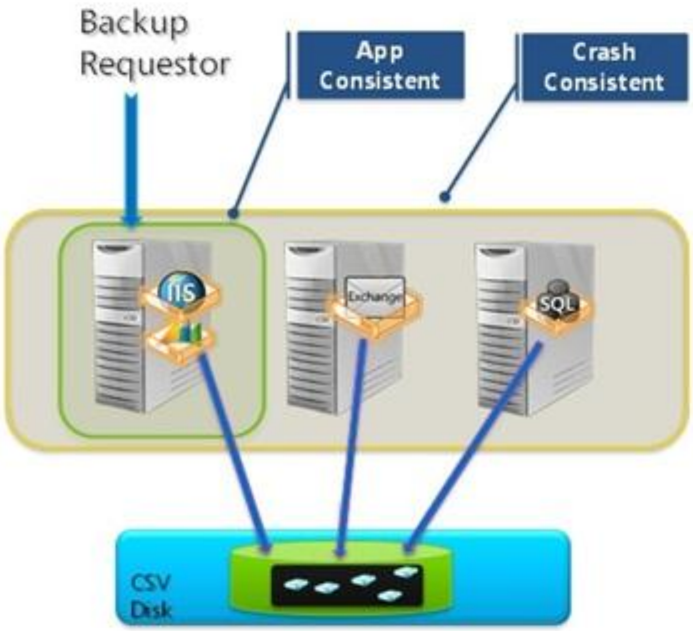

Improvements in cluster validation and deployment

The following improvements to failover clustering make clusters easier to deploy:

- **Improvements to failover cluster validation.** Validate a Configuration Wizard has been improved in multiple ways as described in the next section. One improvement to the wizard is the addition of tests to simplify the testing and validation of configurations that include Hyper-V. For more information about validation, see the "Cluster validation improvements" section of this document.
- **Improved Active Directory Domain Services integration.** Improvements to integration with Active Directory Domain Services (AD DS) simplify the deployment of failover clusters and the clustered roles within them. For more information, see the "Improved Active Directory Domain Services integration" section of this document.
- **Enhanced wizard for migrating settings from one cluster to another.** These enhancements make it easier to migrate the configuration settings for clustered services and applications from one cluster to another. For more information, see the "Enhanced wizard for migrating settings from one cluster to another" section of this document.
- **Support for Load Balancing and Failover (LBFO).** LBFO is a built-in network adapter teaming solution. Failover clusters work seamlessly with LBFO.
- **Support of Sysprep for installing failover cluster nodes.** Administrators can now use Sysprep when they install failover cluster nodes, for speed and consistency of installation.

Cluster validation improvements

The Validate a Configuration Wizard (inside Failover Cluster Manager) simplifies the process of validating hardware and software for use in a failover cluster. Validation tests examine hardware (servers, storage, and network) and other aspects of the configuration, such as the consistency of software updates across servers. Many aspects of validation have been improved, including the following:

- **Faster validation.** Validation tests, especially storage validation tests, now run significantly faster.
- **Targeted validation of new LUNs.** Administrators can target validation of a specific new logical unit number (LUN), rather than needing to test all LUNs every time they test storage.
- **Integration of validation with WMI.** Cluster validation status is now exposed through Windows Management Instrumentation (WMI), so that applications and scripts can programmatically consume it.
- **New validation tests for CSV.** Validation tests help administrators confirm that their configuration meets the requirements for CSV.
- **New validation tests for Hyper-V and virtual machines.** Validation tests help administrators determine whether the servers in their cluster are compatible for Hyper-V purposes, that is, whether the servers will support smooth failover of virtual machines from one host to another.

Improved Active Directory Domain Services integration

Significant enhancements have been made to improve working with Active Directory Domain Services (AD DS). These include:

- Administrators can now deploy clusters that have access only to a read-only domain controller (RODC). This enables cluster deployments in branch offices and in scenarios that use a perimeter network.
- Creation of cluster computer objects (computer accounts) in AD DS can now be targeted to a specific organizational unit for enterprises in which this is standard procedure.
- When a failover cluster is initially created and a corresponding computer object is created in AD DS, that object is configured to prevent accidental deletion.
- The cluster Network Name resource has additional health check logic, which periodically checks the health and properties of the computer object that represents the Network Name resource.

Enhanced wizard for migrating settings from one cluster to another

The wizard that lets administrators migrate the configuration settings for clustered roles (clustered services and applications) from one cluster to another has been enhanced to include the following:

- Support for copying the configuration information of multiple offline virtual machines from one cluster to another.
- Migration of configuration information from services and applications on clusters that are running Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012.

Enhanced Windows PowerShell support for clusters

Failover clusters in Windows Server 2012 include support for new Windows PowerShell cmdlets.

Table 2: Clustering PowerShell cmdlets

New Windows PowerShell cmdlets	Purpose
Add-ClusterCheckpoint Get-ClusterCheckpoint Remove-ClusterCheckpoint	Manage cluster registry checkpoints, including cryptographic checkpoints.
Add-ClusterScaleOutFileServerRole	Create a file server for scale-out application data. (This is a new type of file server in Windows Server 2012).
Add-ClusterVMMonitoredItem Get-ClusterVMMonitoredItem Remove-ClusterVMMonitoredItem	Monitor the health of services that are running inside a virtual machine.

New Windows PowerShell cmdlets	Purpose
Reset-ClusterVMMonitoredState	
Update-ClusterNetworkNameResource	Update the private properties of a Network Name resource. This also sends Domain Name System (DNS) updates.
Test-ClusterResourceFailure	Replace the Fail-ClusterResource cmdlet.

Clustering improvements for multisite clusters

In Windows Server 2008 R2, administrators could configure the quorum to include nodes. (A quorum is the collection of "voting" entities of which a majority must be online for the cluster to function.) However, if the quorum configuration included nodes, all nodes were treated equally in terms of their "votes." In Windows Server 2012, cluster quorum settings can be adjusted so that when the cluster determines whether it has a quorum, some nodes have a "vote" and some do not. This adjustment, which is shown in the following figure, can be useful when solutions are implemented across multiple sites.

Figure 15: Dynamic quorum management

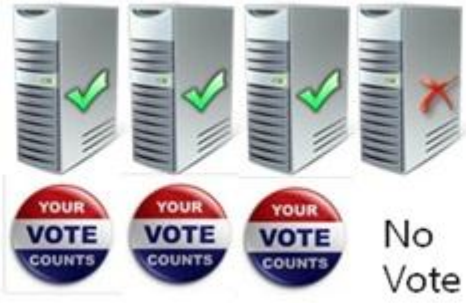

High Availability for DHCP Server Service

When you install the Dynamic Host Configuration Protocol (DHCP) server role, the DHCP server provides IP addresses and other network configuration parameters to host computers and network devices that are configured as DHCP clients. If DHCP becomes unavailable due to network, hardware, or other failures, it can result in a widespread loss of network connectivity for DHCP clients that are installed on subnets that are serviced by the DHCP server.

The DHCP server failover feature in Windows Server 2012 provides a method to help ensure continuous availability of the DHCP Server service to DHCP clients.

DHCP challenges

In Windows Server 2008 R2, there are two high availability options for DHCP server deployment. Each of these options has some challenges.

1. **DHCP in a Windows failover cluster.** This option places the DHCP server in a cluster with one or more additional servers that are configured with the DHCP server service. If the primary DHCP server fails, one of the other DHCP servers in the failover cluster assumes the load and provides the DHCP server service to clients. Despite this failover capability, the clustering deployment option uses a single, shared storage option. This makes the storage a single point of failure, and it requires additional investment in redundancy for storage. In addition, clustering involves relatively complex setup and maintenance.
2. **Split-scope DHCP.** Split-scope DHCP uses two independent DHCP servers that share the responsibility for a scope. Typically, 70 percent of the addresses in the scope are assigned to the primary server, and the remaining 30 percent are assigned to the backup server. If clients cannot reach the primary server, they can get an IP configuration from the secondary server. Split-scope deployment does not provide IP address continuity, and it is unusable in scenarios where the scope is already running at high utilization of address space, which is very common with IPv4 deployments.

DHCP solution

DHCP failover in Windows Server 2012 enables administrators to deploy a highly resilient DHCP Server service to support a large enterprise without the challenges of the options discussed earlier. The main goals of the feature include:

- Provide DHCP server service availability at all times on the enterprise network.
- If a DHCP server is no longer reachable, the DHCP client can extend the lease on its current IP address by contacting another DHCP server on the enterprise network.

The DHCP server failover feature allows two DHCP servers to assign IP addresses and DHCP options to DHCP clients that are located on the same subnet or that receive IP address leases from the same scope, which provides continuous availability of the DHCP servers to clients. The two DHCP servers replicate lease information between them, allowing either server to continue servicing clients for the entire subnet when the other server is unavailable. It is also possible to configure failover in a load-balancing configuration with the client requests distributed between the two servers in a failover relationship.

DHCP failover in Windows Server 2012 provides support for a maximum of two DHCP servers, and the failover relationship is limited to IPv4 scopes and subnets. Network nodes that use IPv6 typically determine their own IPv6 address by using stateless IP auto configuration. In this mode, the DHCP server delivers only the DHCP option configuration, and the server does not maintain any lease state information. A high availability deployment for stateless DHCP over IPv6 (DHCPv6) is possible by simply setting up two servers with identical option configurations. Even in a stateful DHCPv6 deployment, the scopes do not run under high address utilization, which makes split-scope DHCP a viable solution for high availability.

DHCP failover architecture

Administrators can deploy DHCP servers running Windows Server 2012 as failover partners in hot-standby mode or load-sharing mode.

Hot-standby mode

In hot-standby mode, two DHCP servers operate in a failover relationship where an active server is responsible for leasing IP addresses and configuration information to all clients in a scope or subnet. The secondary server continues servicing clients if the primary server becomes unavailable. A server is primary or secondary in the context of a subnet. For instance, a server that has the role of a primary server for a given subnet could be a secondary server for another subnet.

Hot-standby mode of operation is ideally suited to deployments where a central office or datacenter server acts as a standby backup server for a server at a remote site, which is local to the DHCP clients (for example, a hub and spoke deployment) as shown in Figure 16.

Figure 16: DHCP failover for multiple sites

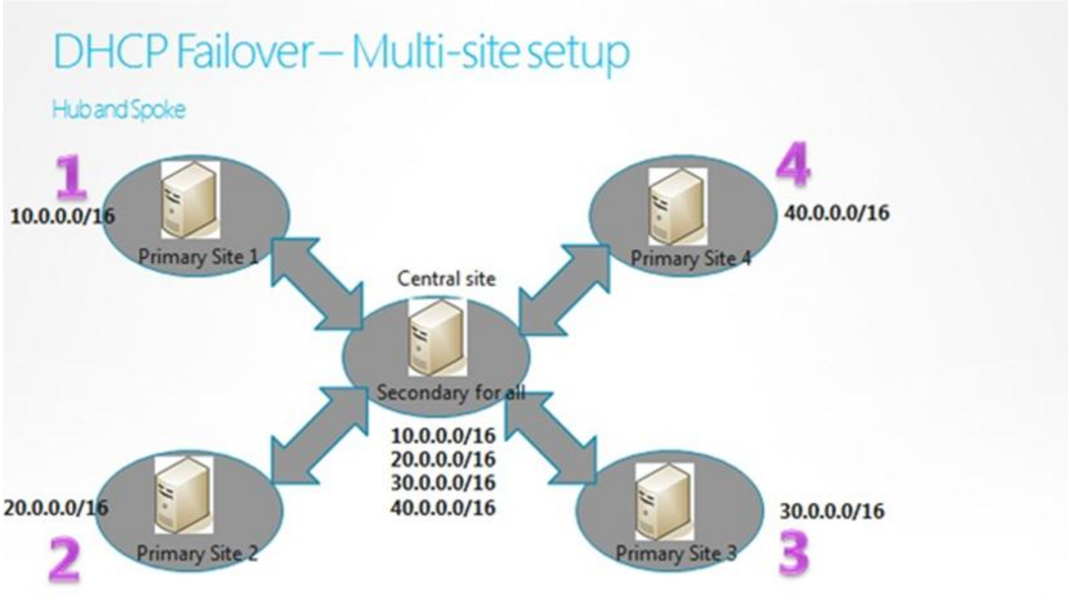

Load-sharing mode

In a load-sharing mode deployment, which is the default mode of operation, the two servers simultaneously serve IP addresses and options to clients on a specific subnet. The client requests are load balanced and shared between the two servers.

The load-sharing mode of operation is ideally suited to deployments where both servers in a failover relationship are located at the same physical site. Both servers respond to DHCP client requests based on the load distribution ratio that is configured by the administrator, as shown in Figures 17 and 18.

Figure 17: DHCP failover for a single site with a single subnet

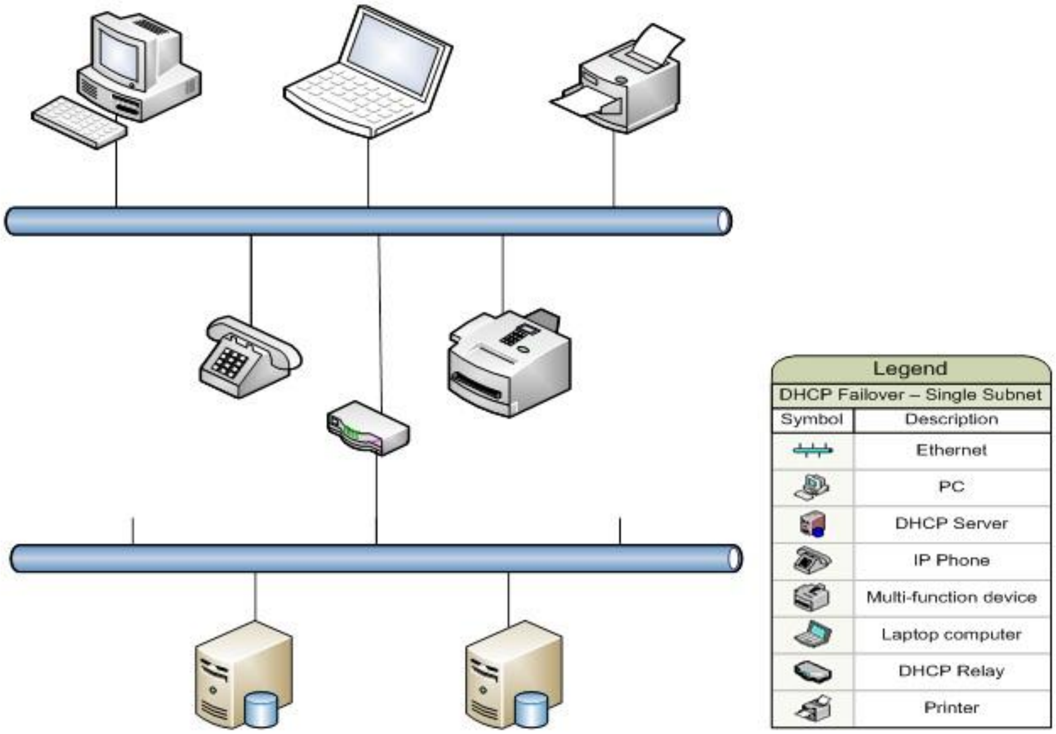

Figure 18: DHCP failover for a single site with multiple subnets

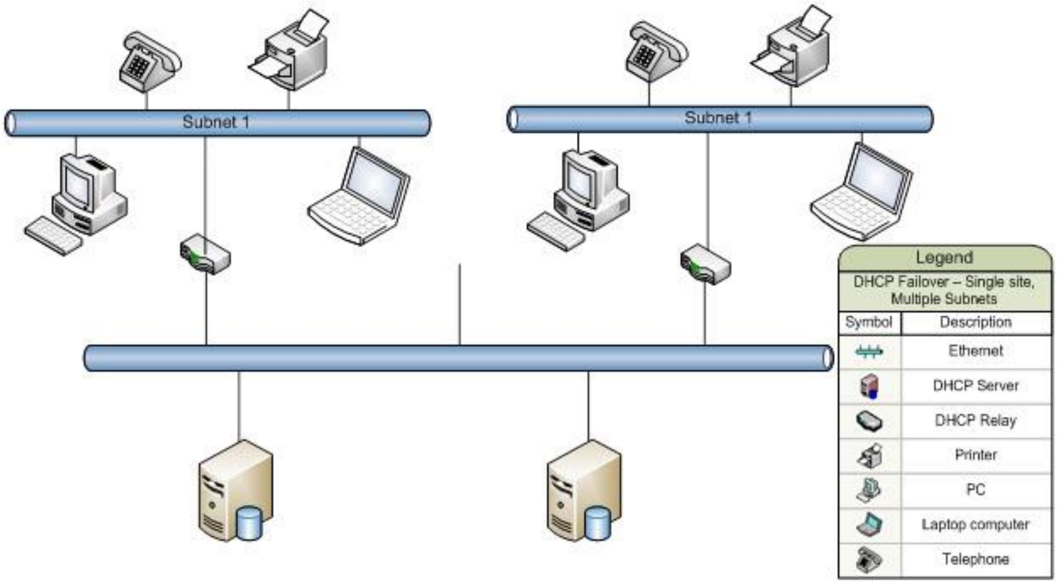

DHCP failover requirements

- **Software requirements.** DHCP failover requires two DHCP servers that are running Windows Server 2012.
- **Number of servers.** DHCP failover is not supported for more than two DHCP servers. The failover relationship always includes only two DHCP servers.
- **Domain membership.** In this topic, DHCP servers are considered domain member servers. However, you can also configure DHCP failover on workgroup computers.
- **Time synchronization.** For DHCP failover to function correctly, time must be continuously synchronized between the two servers in a failover relationship. Time synchronization can be maintained by deploying the Network Time Protocol (NTP) or an alternate mechanism. When the Failover Configuration Wizard is run, it compares the current time on the servers that are being configured for failover. If the time difference between the servers is greater than one minute, the failover setup process fails with a critical error that instructs the administrator to synchronize the time on the servers.
- **Hardware requirements.** For recommended server hardware specifications, see the "Windows Server 2008 R2 with SP1 System Requirements" section of this document.

High-Performance, Highly Available File Share Storage for Server Applications

Windows Server 2012 introduces new file server features that helps you store server application data on file shares and obtaining a similar level of reliability, availability, manageability, and high performance that you would expect from a storage area network (SAN). These new file server features include transparent failover, networking improvements for greater bandwidth and resiliency, support for network adapters with RDMA capability, specific performance optimizations, and support for Windows PowerShell commands.

Key benefits

Store application data on inexpensive, easy-to-manage file shares and obtain similar (or better) benefits of continuous availability, high performance, and manageability that you would expect from a SAN.

Technical overview

Windows Server 2012 introduces a set of new file server features that provide important improvements for server applications such as Microsoft SQL Server and Hyper-V, which can store data on file shares.

The following improvements have been added to Windows Server 2012:

- **SMB transparent failover.** You can now more easily perform hardware or software maintenance of nodes in a clustered file server by moving file shares between nodes without interrupting server applications that are storing data on these file shares. Also, if a hardware or software failure occurs on a cluster node, Server Message Block (SMB 3.0) 3.0 transparent failover lets file shares failover to another cluster node without interrupting server applications that are storing data on these file shares.
- **SMB multiple-channel.** This improvement allows aggregation of network bandwidth and network fault tolerance if multiple paths are available between the SMB client and the SMB server. Server applications can then take advantage of all available network bandwidth to be more resilient in the event of a network failure.
- **SMB direct.** This improvement uses a special type of network adapter that has remote direct memory access (RDMA) capability and can function at full speed with very low latency, while using very little CPU. For server roles or applications such as Hyper-V or SQL Server, this gives a remote file server performance comparable to local storage.
- **SMB performance counters for server applications.** Performance counters provide detailed information about I/O size, I/O latency, IOPS, and so on. This helps SQL Server database administrators or Hyper-V administrators analyze the performance of the SMB file shares where their data is stored.

- **SMB performance optimizations.** The SMB client and SMB server have been optimized for small, random read/write I/O, which is common in server applications such as SQL Server online transaction processing (OLTP). In addition, a large maximum transmission unit (MTU) is enabled by default, which significantly enhances performance in large sequential transfers, such as with a SQL Server data warehouse, a database backup or restore, or the copying or deployment of virtual hard disks.
- **SMB management with Windows PowerShell.** Organizations can use command line in Windows PowerShell, you can use the command line to manage SMB on a file server, end to end.
- **SMB remote file storage.** Hyper-V can now store virtual machine files (including configuration files, virtual hard disk files, and snapshots) in shared folders that use the SMB protocol. Support for storing database files in shared folders that use the SMB protocol was introduced in SQL Server 2008 R2.

The following figure shows an example of a two-node file server cluster that provides storage for Hyper-V and SQL Server.

Figure 19: Two-node File Server cluster providing storage for Hyper-V and SQL Server

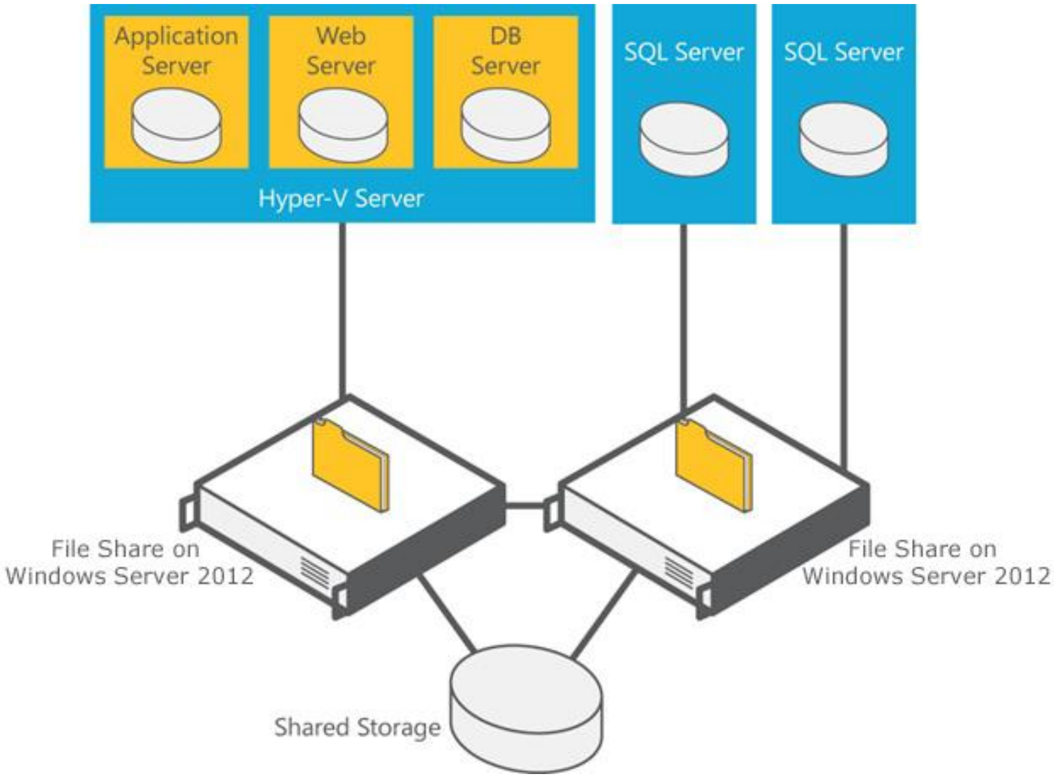

Following are the main advantages of storing server application data on shared folders in Windows Server 2012:

- **Ease of provisioning and management.** You can manage file shares instead of storage fabric and logical unit numbers (LUNs).
- **Increased flexibility.** You can dynamically relocate virtual machines or databases in the datacenter.
- **Ability to take advantage of existing investment in a converged network.** You can use your existing converged network with no specialized storage networking hardware.
- **Reduced capital expenditures.** Capital expenses (acquisition costs) are reduced.
- **Reduced operating expenditures.** You can reduce operating costs because there is virtually no need for specialized storage expertise.

Requirements

New SMB 3.0 features

SMB transparent failover has the following requirements:

- **A failover cluster running Windows Server 2012 with at least two nodes.** The cluster must pass the cluster validation tests in the validation wizard.
- **The "Services For Continuously Available Shares" role service installed on all cluster nodes.** This role service provides the persistent store that enables the file server to resume handles after a failover. It also provides a witness service that helps clients more quickly reconnect to a clustered file share after an unplanned failure.
- **File shares created with the Continuous Availability property.** This is the default setting.
- **Computers running Windows 8 (for client computers) or Windows Server 2012.** Both computers must include the updated SMB client that supports continuous availability.

Note

Down-level client computers can connect to file shares that have the Continuous Availability property, but transparent failover is not supported for these computers.

SMB multichannel has the following requirements:

- **At least two computers running Windows Server 2012.** No extra features need to be installed—the technology is available by default.
- The following network configurations are suggested:
  - **Single 10-GbE network adapters.** Each computer is configured with a single 10-Gigabit Ethernet (GbE) network interface.
  - **Dual 1-GbE network adapters.** Each computer must be configured with two 1-GbE network interfaces. Each SMB network adapter on a client computer communicates with an SMB network adapter on a server by using a different subnet.
  - **Dual 1-GbE network adapters in a team.** Each computer must be configured with two 1-GbE network interfaces configured as a Load Balancing and Failover (LBFO) team. Each SMB network adapter on a client and SMB network adapter on a server communicates by using teamed interfaces.
  - **Dual 10-GbE network adapters.** Each computer must be configured with two 10-GbE network interfaces. Each SMB client network adapter communicates with an SMB server network adapter by using a different subnet.
  - **Dual Infiniband network adapters.** Each computer must be configured with two Infiniband network interfaces. Each SMB client network adapter communicates with an SMB server network adapter by using a different subnet.
- SMB direct has the following requirements:
  - **At least two computers running Windows Server 2012.** No extra features need to be installed, and the technology is available by default.
  - **Network adapters with RDMA capability.** Currently, these network adapters come in three types: iWARP, Infiniband, and RDMA over Converged Ethernet (RoCE).

Hyper-V over SMB

Hyper-V over SMB has the following requirements:

- One or more computers running Windows Server 2012 with the Hyper-V role installed.
- One or more computers running Windows Server 2012 with the File Services role installed.
- A common Active Directory infrastructure. The servers running Active Directory Domain Services (AD DS) do not need to run Windows Server 2012.

Note

Although not required, Failover Clustering is supported on the Hyper-V side, the File Services side, or both.

The three most common file server configurations for Hyper-V over SMB are a single-node file server, a dual-node file server, and a multiple-node file server, as shown in the following figure.

Figure 20: Common configurations for Hyper-V over SMB

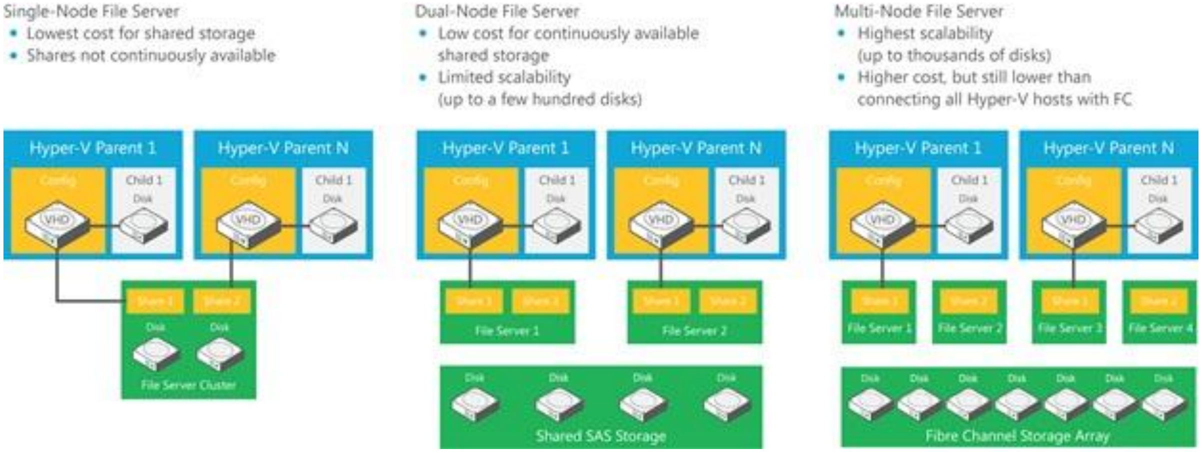
Single-Node File Server
• Lowest cost for shared storage
• Shares not continuously available
Dual-Node File Server
• Low cost for continuously available shared storage
• Limited scalability (up to a few hundred disks)
Multi-Node File Server
• Highest scalability (up to thousands of disks)
• Higher cost, but still lower than connecting all Hyper-V hosts with FC
Hyper-V Parent 1
Hyper-V Parent N
Config
Child 1
Disk
VHD
Share 1
Share 2
File Server Cluster
File Server 1
File Server 2
File Server 3
File Server 4
Shared SAS Storage
Fibre Channel Storage Array

SQL Server over SMB

SQL Server over SMB has the following requirements:

- One or more computers running Windows Server 2012 with SQL Server 2008 R2 or SQL Server 2012.
- One or more computers running Windows Server 2012 with the File Services role installed.
- A common Active Directory infrastructure. Servers running Active Directory Domain Services do not need to run Windows Server 2012.

Note

Although not required, Failover Clustering is supported on the SQL Server 2012 side, the File Services side, or both.

The three most common configurations for SQL Server over SMB are a single-node file server, a dual-node file server, and a multiple-node file server, as shown in the following figure.

Figure 21: Common configurations for SQL Server over SMB

Single-Node File Server (D1)	Dual-Node Cluster File Server (D2)	Multi-Node Clustered File Server (D3)
Low cost	Medium cost	Higher cost
Shares not continuously available	Shares continuously available	Shares continuously available
Limited scalability (~100 spindles)	Medium scalability (~200 spindles)	Highest scalability (~1,000 spindles)

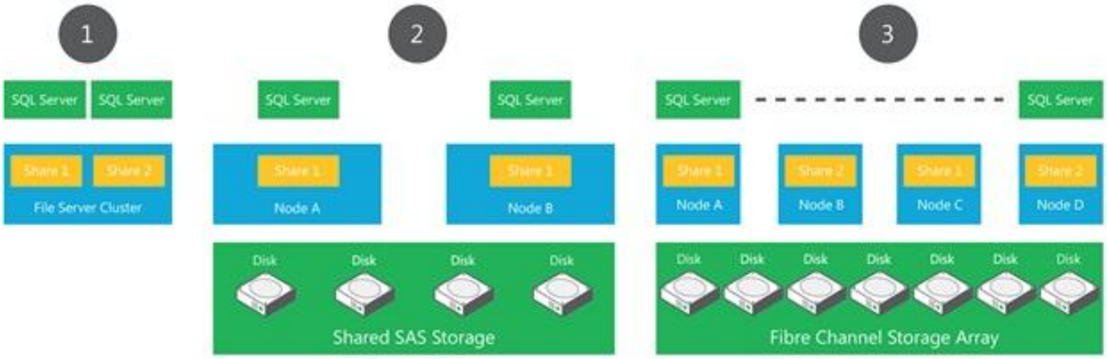

Hyper-V Automation Support

Windows PowerShell is the scripting solution for automating tasks in Windows Server. However, in current versions of Hyper-V, writing scripts for Hyper-V with in-box tools requires users to learn Windows Management Instrumentation (WMI). WMI provides a flexible set of interfaces, but they are designed for developers, not IT pros. Hyper-V in Windows Server 2012 addresses this issue by introducing more than 140 built-in Hyper-V cmdlets for Windows PowerShell. Administrators can use these new cmdlets to more easily enable automation of datacenter tasks that range from basic to complex.

Key benefits

The new Hyper-V cmdlets for Windows PowerShell provide IT pros with an easier way to enable automation of management tasks within Windows Server 2012. With the extensive number of Hyper-V cmdlets and the close integration with other parts of the operating system, administrators can now more easily enable automation of Hyper-V-related tasks in their environment.

Requirements

To try out the Hyper-V cmdlets, you will need the following:

- A computer running Windows Server 2012 with the Hyper-V role installed. Hyper-V requires a computer that has processor support for hardware virtualization.
- A user account that is a member of the Administrators group or Hyper-V Administrators group.

Technical overview

Hyper-V cmdlets enable administrators to perform the tasks that are available in the GUI of Hyper-V Manager—as well as a number of tasks that can be done exclusively through the cmdlets in Windows PowerShell.

Designed for IT pros

Windows PowerShell is designed intentionally for IT pros. This design decision is reflected in several ways:

- **Task-oriented interface.** Hyper-V cmdlets are designed so that it is easier for IT pros to go from thinking about the task to actually performing the task. The following table shows the task and the associated cmdlet syntax:

Table 3: Hyper-V PowerShell commands

Task	Windows PowerShell command to perform the task
Create a new virtual machine named “test”	New-VM –Name Test
Get a list of all virtual machines	Get-VM
Create a new virtual hard disk at d:\VHDs\test.vhd	New-VHD –Path D:\VHDs\test.vhd
Start all virtual machines whose name begins with “web”	Start-VM –Name web*
Connect the virtual network adapter on the “test” virtual machine to the “QA” switch.	Connect-VMNetworkAdapter –VMName test –SwitchName QA

- **Use of standard cmdlet verbs.** Hyper-V administrators often need to manage more than just Hyper-V. By using the same verbs as other Windows cmdlets, the Hyper-V cmdlets make it easier for administrators to extend their existing knowledge of Windows PowerShell. For example, administrators who are familiar with managing services through Windows PowerShell can reuse the same verbs to perform the corresponding tasks on a virtual machine, as shown in the following table:

Table 4: Hyper-V PowerShell commands

Task	Cmdlet for performing a task on a service	Cmdlet for performing a task on a virtual machine
Get	Get-Service	Get-VM
Configure	Set-Service	Set-VM
Create	New-Service	New-VM
Start	Start-Service	Start-VM
Stop	Stop-Service	Stop-VM
Restart	Restart-Service	Restart-VM
Suspend	Suspend-Service	Suspend-VM
Resume	Resume-Service	Resume-VM

There are similar examples with other core Windows PowerShell cmdlets as well:

Table 5: Hyper-V PowerShell commands

Core Windows PowerShell cmdlet	Hyper-V cmdlet
Import-CSV	Import-VM
Export-CSV	Export-VM
Enable-PSRemoting	Enable-VMMigration
Checkpoint-Computer	Checkpoint-VM
Measure-Command	Measure-VM

- **Consistent cmdlet nouns to simplify discoverability.** The nouns of the Hyper-V cmdlets are designed to make it easier for administrators to discover the cmdlets they need when they need them. All cmdlets in the Hyper-V module use one of three following noun prefixes:

Table 6: Hyper-V PowerShell commands

Prefix	Purpose
VM	Cmdlets for managing virtual machines
VHD	Cmdlets for managing virtual hard disk files
VFD	Cmdlets for managing virtual floppy disk files

Hyper-V Dynamic Memory

If you have idle or low-load virtual machines, as in pooled Virtual Desktop Infrastructure (VDI) environments, Dynamic Memory additions in Hyper-V in Windows Server 2012 enable you to increase consolidation and improve reliability for restart operations. You also gain agility in responding to requirement changes with these new capabilities.

Key benefits

With the Dynamic Memory improvements for Hyper-V in Windows Server 2012, you can attain higher consolidation numbers with improved reliability for restart operations. This can lead to lower costs, especially in environments that have many idle or low-load virtual machines, such as pooled VDI environments. Dynamic Memory run-time configuration changes can reduce downtime and provide increased agility to respond to requirement changes.

Technical overview

Dynamic Memory, introduced in Windows Server 2008 R2 Service Pack 1 (SP1), defined startup memory as the minimum amount of memory that a virtual machine can have. However, Windows requires more memory during startup than the steady state. As a result, administrators sometimes assign extra memory to a virtual machine because Hyper-V cannot reclaim memory from these virtual machines after startup. In Windows Server 2012, Dynamic Memory introduces a minimum memory setting, which allows Hyper-V to reclaim the unused memory from the virtual machines. This is reflected as increased virtual machine consolidation numbers, especially in Virtual Desktop Infrastructure (VDI) environments.

Windows Server 2012 also introduces smart paging for reliable virtual machine restart operations. Although minimum memory increases virtual machine consolidation numbers, it also brings a challenge. If a virtual machine has a smaller amount of memory than its startup memory and if it is restarted, Hyper-V needs additional memory to restart the virtual machine. Due to host memory pressure or virtual machine states, Hyper-V may not always have additional memory available. This can cause sporadic virtual machine restart failures. Smart Paging is used to bridge the memory gap between minimum memory and startup memory, and allow virtual machines to restart reliably.

Minimum memory configuration with reliable restart operation

As in the previous version of Dynamic Memory, you can configure a minimum memory amount for virtual machines in Windows Server 2012, and Hyper-V continues to ensure that this amount is always assigned to running virtual machines.

To provide a reliable restart experience for the virtual machines configured with less minimum memory than startup memory, Hyper-V in Windows Server 2012 uses smart aging. This memory management method uses disk resources as additional, temporary memory when more memory is required to restart a virtual machine. This approach has advantages and drawbacks. It provides a reliable way to keep the virtual machines running when there is no available physical memory. However, it can degrade virtual machine performance because disk access speeds are much slower than memory access speeds.

To minimize the performance impact of smart paging, Hyper-V uses it only when all of the following occurs:

- The virtual machine is being restarted.
- There is no available physical memory.
- No memory can be reclaimed from other virtual machines running on the host.

Smart paging is not used when:

- A virtual machine is being started from an "off state" (instead of a restart).
- Oversubscribing memory for a running virtual machine is required.
- A virtual machine is failing over in Hyper-V clusters.

When host memory is oversubscribed, Hyper-V continues to rely on the paging operation in the guest operating system because it is more effective than smart paging. The paging operation in the guest operating system is performed by Windows Memory Manager. Windows Memory Manager has more information than the Hyper-V host about memory usage within the virtual machine, which means it can provide Hyper-V with better information to use when choosing the memory to be paged. Because of this, less overhead to the system is incurred compared to smart paging.

To further reduce the impact of smart paging, Hyper-V removes memory from the virtual machine after it completes the start process. It accomplishes this by coordinating with Dynamic Memory components inside the guest operating system (a process sometimes referred to as "ballooning"), so the virtual machine stops using smart paging. With this technique, the use of smart paging is temporary and is not expected to be longer than 10 minutes.

Also note the following about how smart paging is used:

- Smart paging files are created only when needed for a virtual machine.
- After the additional amount of memory is removed, smart paging files are deleted.
- Smart paging is not used for this virtual machine again until another restart occurs and there is not enough physical memory.

Run-time Dynamic Memory configuration changes

Hyper-V in Windows Server 2012 enables users to make the following configuration changes to Dynamic Memory when the virtual machine is running:

- Increase the maximum memory.
- Decrease the minimum memory.

Hyper-V Network Virtualization

With the success of virtualized datacenters, IT organizations and hosting providers (providers who offer colocation or physical server rentals) are offering flexible virtualized infrastructures that make it easier to offer on-demand server instances to their customers. This new class of a service is referred to as infrastructure as a service (IaaS). Windows Server 2012 provides all the required platform capabilities to enable enterprise customers to build private clouds and transition to an IaaS operational model, and also to enable hosting providers to build public clouds and offer IaaS solutions to their customers.

Hyper-V Network Virtualization in Windows Server 2012 provides policy-based, software-controlled network virtualization, which helps reduce the management overhead that organizations face when they are expanding dedicated IaaS clouds. Network virtualization also helps provide better flexibility for cloud hosting providers, scalability for managing virtual machines, and higher resource utilization.

Network virtualization challenges

An IaaS scenario with multiple virtual machines from different divisions (dedicated cloud) or different customers (hosted cloud) requires secure isolation. Currently, virtual LANs (VLANs) are the mechanism most organizations use to provide address space reuse, and tenant isolation. A VLAN uses explicit tagging in the Ethernet frames, and it relies on Ethernet switches to enforce isolation and restrict traffic to network nodes of the same tag. The main drawbacks with VLANs are:

- Increased risk of an inadvertent outage due to cumbersome reconfiguration of production switches whenever virtual machines or isolation boundaries move in the dynamic datacenter.
- Limited scalability because typical switches support no more than 1,000 VLAN IDs (maximum of 4,094).
- VLANs cannot span multiple logical subnets, which limits the number of nodes within a single VLAN and restricts the placement of virtual machines based on physical location. Even though VLANs can be enhanced or stretched across physical intranet locations, the stretched VLANs must be all on the same subnet.

In addition to the drawbacks presented by VLANs, virtual machine IP address assignment presents other significant issues, including:

- Moving to a cloud platform typically requires reassigning IP addresses for the service workloads.
- Policies (security, management, and other) are tied to IP addresses.
- Physical locations determine the virtual machine IP address.
- The topological dependency of virtual machine deployment and traffic isolation.

The IP address is the fundamental address that is used for layer 3 network communication. In addition to being an address, there is semantic information associated with an IP address. For example, one subnet might contain specific services or be in a distinct physical location. Firewall rules, access control policies, and Internet Protocol security (IPsec) security associations are commonly linked to IP addresses. When moving to the cloud, organizations must change IP addresses to accommodate the physical and topological restrictions of the datacenter. This renumbering of IP addresses is burdensome because all of the associated policies based on IP addresses must also be updated.

When datacenter network administrators plan the physical layout of a datacenter, they must make decisions about where subnets will be physically placed and routed in the datacenter. These decisions are based on IP and Ethernet technologies that are 30 years old. These technologies influence the potential IP addresses that are allowed for virtual machines running on a specific server or server blade that is connected to a specific rack in the datacenter. When a virtual machine is provisioned and placed in the datacenter, it must adhere to the choices and restrictions regarding the IP address. Therefore, the typical result is that datacenter administrators assign IP addresses to the virtual machines, forcing the virtual machine owners to adjust all their policies that were based on the original IP address. This renumbering overhead is so high that many enterprises deploy only new services into their cloud platform, leaving legacy applications alone.

Networking virtualization solution

Network virtualization in Windows Server 2012 removes the constraints of VLAN and hierarchical IP address assignment for virtual machine provisioning. It makes IaaS cloud computing easier for customers to implement, and easier for hosting providers and datacenter administrators to manage. In addition, network virtualization maintains the necessary multitenant isolation and security requirements. The following list summarizes the key benefits and capabilities of network virtualization in Windows Server 2012:

- **Uncouples workloads from internal IP addresses.** Enables customers to keep their internal IP addresses while moving workloads to shared IaaS cloud platforms. Uncoupling workloads from internal IP addresses minimizes the configuration changes that are needed for IP addresses, DNS names, security policies, and virtual machine configurations.
- **Decouples server and network administration.** Server workload placement is simplified because migration and placement of workloads are independent of the underlying physical network configurations. Server administrators can focus on managing services and servers, while network administrators can focus on overall network infrastructure and traffic management.
- **Removes the tenant isolation dependency on VLANs.** In the software-defined, policy-based datacenter networks, network traffic isolation is no longer dependent on VLANs, but enforced within host computers running Hyper-V based on the multitenant isolation policy. Network administrators can still use VLANs to manage traffic for the physical infrastructure where the topology is primarily static.
- **Enables flexible workload placement.** Allows services and workloads to be placed or migrated to any server in the datacenter. They can keep their IP addresses, and they are not limited to a physical IP subnet hierarchy or VLAN configurations.
- **Simplifies the network and improves server and network resource utilization.** The rigidity of VLANs and dependency of virtual machine placement on physical network infrastructure results in overprovisioning and underutilization. By breaking the dependency, the increased flexibility of virtual machine workload placement can simplify network management and improve server and network resource utilization.

- **Works with existing infrastructure and emerging technology.** Network virtualization can be deployed in current datacenter environments, and it is compatible with the emerging datacenter “flat network” technologies, such as the Transparent Interconnection of Lots of Links (TRILL) architecture that is intended to expand Ethernet topologies.
- **Supports configuration by using Windows PowerShell and WMI.** You can use Windows PowerShell to enable scripting and automation of administrative tasks.

Requirements

Network virtualization in Windows Server 2012 requires Windows Server 2012 with the Hyper-V role.

Technical overview

Server virtualization is a concept that allows multiple server instances to run on a single physical host computer concurrently, yet isolated from each other, with each server instance operating as if it is the only server running on the physical computer. Network virtualization provides a similar capability, in which you can run multiple virtual network infrastructures, potentially with overlapping IP addresses, on the same physical network. With network virtualization, each virtual network infrastructure operates as if it is the only one running on the shared network infrastructure. The figure below shows this relationship.

Figure 22: Virtualization for servers and networks

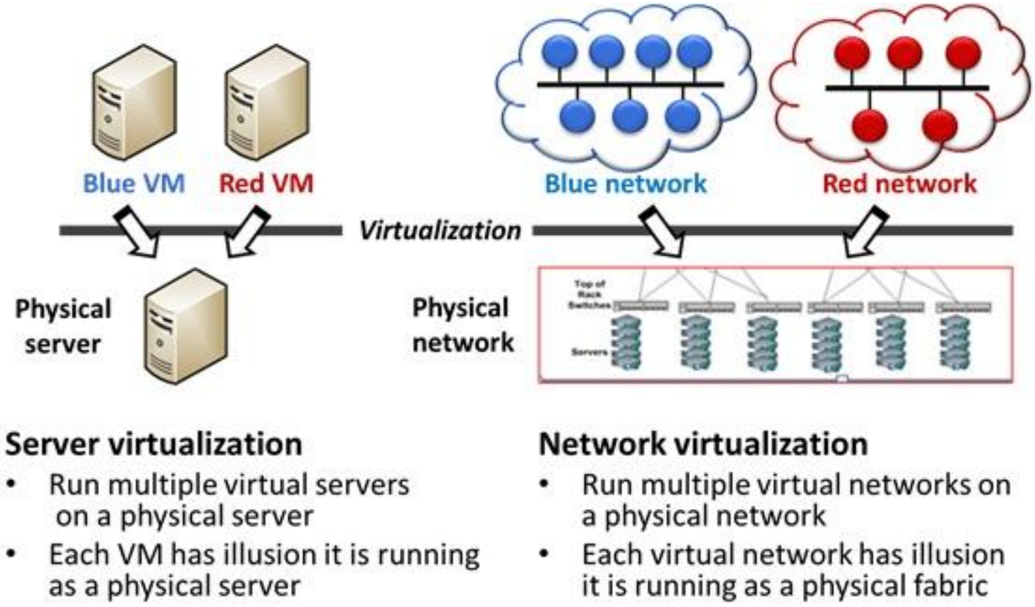

Server virtualization

- Run multiple virtual servers on a physical server
- Each VM has illusion it is running as a physical server

Network virtualization

- Run multiple virtual networks on a physical network
- Each virtual network has illusion it is running as a physical fabric

To virtualize the network with Windows Server 2012, each virtual machine is assigned two IP addresses:

- The customer address is the IP address that is assigned by the customer based on their intranet infrastructure. This address allows the customer to exchange network traffic with the virtual machine as if it had not been moved to a public or private cloud. The customer address is visible to the virtual machine and reachable by the customer.

- The provider address is the IP address that is assigned by the hosting provider based on their physical network infrastructure. The provider address appears in the layer-3 packets that are exchanged with the server that is running Hyper-V and is hosting the virtual machine. The provider address is visible on the physical network, but not to the virtual machine.

The layer of customer addresses maintains the topology of the customer network, which is virtualized and decoupled from the actual underlying physical network or physical network addresses, as implemented by the layer of provider addresses.

There are two techniques to virtualize the IP address of the virtual machine. One technique, called *IP rewrite*, modifies the customer IP address of the packets on the virtual machine before they are transferred on the physical network. IP rewrite is used to provide better performance because the networking offload technologies in Windows, such as Virtual Machine Queue (VMQ), continue to operate as they have in the past.

The other IP virtualization technique is *IP encapsulation*. In IP encapsulation, all of the virtual machines packets are encapsulated with a new header before being sent on the physical network. IP encapsulation offers better scalability because all of the virtual machines on a specific host can share the same provider IP address. The tenants are identified by examining the header of the encapsulated packet, which contains a tenant network ID. Because all of the virtual machines are sharing the provider IP address, the switching infrastructure only needs to know about the IP address and media access control (MAC) address for the provider address. This can result in a 20 times or greater reduction in the number of addresses that need to be learned by the network infrastructure, thus reducing the size of IP/MAC tables. Large IP/MAC tables are expensive because they require high-end switches.

With network virtualization in Windows Server 2012, any virtual machine workload can run unmodified on any server running Windows Server 2012 with Hyper-V, within any physical subnet. To do so, the server running Windows Server 2012 with Hyper-V must have the required policy settings that can map between the two addresses. This approach enables many benefits, which include cross-subnet live migration, customer virtual machines running IPv4 while the hosting provider is running an IPv6 datacenter (or vice-versa), and the use of IP address ranges that overlap between customers.

Consider the example in Figure 23.

Figure 23: Customer address virtualization

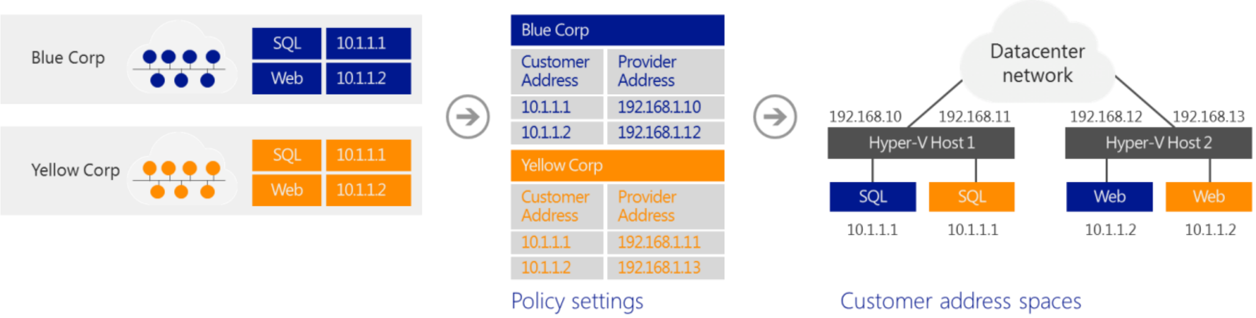

Blue Corp	
Customer Address	Provider Address
10.1.1.1	192.168.1.10
10.1.1.2	192.168.1.12

Yellow Corp	
Customer Address	Provider Address
10.1.1.1	192.168.1.11
10.1.1.2	192.168.1.13

Policy settings

Prior to moving to the shared IaaS service of the hosting provider, Blue Corp and Red Corp ran the following server configurations:

- A SQL Server (named SQL) at the IP address 10.1.1.1
- A web server (named WEB) at the IP address 10.1.1.2 that uses its SQL Server for database transactions

Blue Corp and Red Corp move their respective SQL and web servers to same shared IaaS service of the hosting provider where they run the SQL virtual machines in Hyper-V Host 1 and the web virtual machines in Hyper-V Host 2.

All virtual machines maintain their original intranet IP addresses (their customer addresses). Both companies are assigned the following provider addresses by their hosting provider when the virtual machines are provisioned:

- Provider addresses for Blue Corp virtual machines: SQL is 192.168.1.10, WEB is 192.168.1.12
- Provider addresses for Red Corp virtual machines: SQL is 192.168.1.11, WEB is 192.168.1.13

The hosting provider creates policy settings, which consist of an isolation group for Red Corp that maps the Customer Addresses of the Red Corp virtual machines to their assigned Provider Addresses and a separate isolation group for Blue Corp that maps the Customer Addresses of the Blue Corp virtual machines to their assigned Provider Addresses. The provider applies these policy settings to Hyper-V Host 1 and Hyper-V Host 2.

When the Blue Corp WEB virtual machine on Hyper-V Host 2 queries its SQL Server at 10.1.1.1, the following happens:

1. Hyper-V Host 2, based on its policy settings, starts with the addresses in the packet from:
    - Source: 10.1.1.2 (the Customer Address of Blue Corp WEB)
    - Destination: 10.1.1.1 (the Customer Address of Blue Corp SQL)
2. Hyper-V Host 2 translates those addresses to:
    - Source: 192.168.1.12 (the Provider Address for Blue Corp WEB)
    - Destination: 192.168.1.10 (the Provider Address for Blue Corp SQL)
3. When the packet is received at Hyper-V Host 1, based on its policy settings, it will accept the addresses in the packet from:
    - Source: 192.168.1.12 (the Provider Address for Blue Corp WEB)
    - Destination: 192.168.1.10 (the Provider Address for Blue Corp SQL)
4. Hyper-V Host 1 translates the addresses back to:
    - Source: 10.1.1.2 (the Customer Address of Blue Corp WEB)
    - Destination: 10.1.1.1 (the Customer Address of Blue Corp SQL)
5. Hyper-V Host 1 delivers the packet to the Blue Corp SQL virtual machine.
6. When the Blue Corp SQL virtual machine on Hyper-V Host 1 responds to the query, the following happens:
7. Hyper-V Host 1, based on its policy settings, starts with the addresses in the packet from:
    - Source: 10.1.1.1 (the Customer Address of Blue Corp SQL)
    - Destination 10.1.1.2 (the Customer Address of Blue Corp WEB)
8. Hyper-V Host 1 translates those addresses to:
    - Source: 192.168.1.10 (the Provider Address for Blue Corp SQL)
    - Destination: 192.168.1.12 (the Provider Address for Blue Corp WEB)
9. When it is received at Hyper-V Host 2, based on its policy settings, it will accept the addresses in the packet from:
    - Source: 192.168.1.10 (the Provider Address for Blue Corp SQL)

- Destination: 192.168.1.12 (the Provider Address for Blue Corp WEB)

10. Hyper-V Host 2 translates the addresses to:

  - Source: 10.1.1.1 (the Customer Address of Blue Corp SQL)
  - Destination: 10.1.1.2 (the Customer Address of Blue Corp WEB)

11. Hyper-V Host 2 delivers the packet to the Blue Corp WEB virtual machine.

A similar process for traffic between the Red Corp WEB and SQL virtual machines uses the settings in the Red Corp isolation group. With network virtualization, Red Corp and Blue Corp virtual machines interact as if they were on their original intranets, but never with each other—even though they are using the same IP addresses. The separate addresses (Customer Addresses and Provider Addresses), the policy settings of the hosts running Hyper-V, and the address translation between the Customer Address and Provider Address for inbound and outbound virtual machine traffic isolate these two sets of servers from each other.

The setting and maintenance of the network virtualization capabilities require the use of a policy management server, which can be integrated into the management tools that provide virtual machine management.

Hyper-V Offloaded Data Transfer

To take advantage of innovations in storage hardware that provide near-instantaneous copying of large amounts of data, Hyper-V in Windows Server 2012 introduces Offloaded Data Transfer (ODX). With this new feature, Hyper-V workloads use the offload semantics of the host hardware, as well as the virtual storage stack, to perform certain internal operations on virtual hard disks (VHDs) that require large amounts of data to be copied. Hyper-V performs these operations faster than was previously possible.

Key benefits

Hyper-V ODX is a new hardware feature that allows for copying large amounts of data from one location to another. By providing support for ODX in the Hyper-V storage stack, this feature makes it possible to complete these operations in a fraction of the time it would have taken without the support. Allowing the virtualized workload to use the semantics for ODX by passing from the workload to the host hardware helps the virtualized workload operate as efficiently as it would in a non-virtualized environment.

Requirements

Hyper-V ODX requires the following:

- ODX–capable hardware to host the VHD files. The hardware needs to be connected to the virtual machine as virtual SCSI devices or directly attached physical disks (sometimes referred to as pass-through disks).
- This optimization is also supported for natively attached, VHDX-based virtual disks.
- VHD-based or VHDX-based virtual disks attached to a virtual IDE controller do not support this optimization because IDE devices lack support for ODX.

Technical overview

Storage area network (SAN) vendors are working to provide near-instantaneous copy operations of large amounts of data. This storage is designed to allow the system above the disks to specify the move of a specific data set from one location to another (an "offloaded data transfer").

Crucial maintenance tasks for virtual hard disks—such as merge, move, and compact—depend on copying large amounts of data. The current method of copying data requires data to be read in and written to different locations, which can be a time-consuming process.

The storage stack of Hyper-V in Windows Server 2012 supports ODX operations so that these operations can be passed from the guest operating system to the host hardware. This ensures that the workload can use storage enabled for ODX as it would if it were running in a non-virtualized environment. The Hyper-V storage stack also issues ODX operations during maintenance operations for VHDs, such as merging disks and storage migration meta-operations where large amounts of data are moved.

Hyper-V Replica

With Hyper-V Replica, administrators can replicate their Hyper-V virtual machines from one Hyper-V host at a primary site to another Hyper-V host at the replica site. This feature helps reduce the total cost of ownership for an organization by providing a storage-agnostic and workload-agnostic solution that replicates efficiently, periodically, and asynchronously over IP-based networks across different storage subsystems and across sites. This scenario does not rely on shared storage, storage arrays, or other software replication technologies. The following figure demonstrates how Hyper-V Replica helps administrators easily replicate virtual machines to a remote site over a WAN link.

Figure 24: Replicated virtual machines

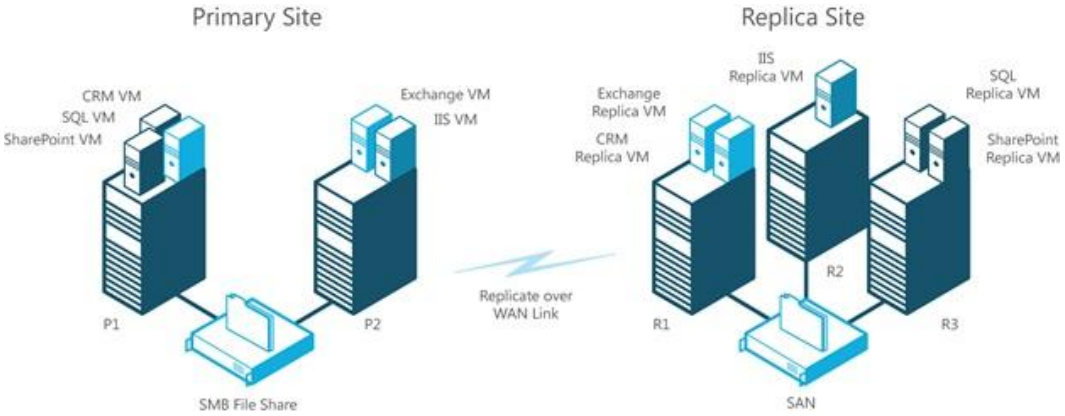

Additionally, administrators can use Hyper-V Replica to test the Replica virtual machine without disrupting the ongoing replication. If a failure occurs at the primary site, administrators can quickly restore their business operations by bringing up the replicated virtual machine at the Replica site.

Key benefits

Virtual machines can be easily replicated to different locations for greater protection and availability.

Key features

Hyper-V Replica tracks the write operations on the primary virtual machine and then replicates these changes to the replica server over a WAN. The network connection between the two servers uses the HTTP protocol and supports Kerberos authentication and certificate-based authentication, with optional support for encryption.

Hyper-V Replica is closely integrated with failover clustering in Windows Server 2012, and it provides replication across different migration scenarios in the primary and replica servers. This enables organizations to store VHDs in a different location to assist recovery in case the datacenter goes down due to natural disaster or other causes.

Hyper-V Resource Metering

IT organizations need tools to charge back business units that they support while providing the business units with the right amount of resources to match their needs. For hosting providers, it is equally important to issue chargebacks based on the amount of usage by each customer.

To implement advanced billing strategies that measure both the assigned capacity of a resource and its actual usage, earlier versions of Hyper-V required users to develop their own chargeback solutions that polled and aggregated performance counters. These solutions could be expensive to develop and sometimes led to loss of historical data.

To assist with more accurate, streamlined chargebacks while protecting historical information, Hyper-V in Windows Server 2012 introduces Resource Metering, a feature that allows customers to create cost-effective, usage-based billing solutions. With this feature, service providers can choose the best billing strategy for their business model, and independent software vendors can develop more reliable, end-to-end chargeback solutions on top of Hyper-V.

Key benefits

Hyper-V Resource Metering in Windows Server 2012 helps organizations to avoid the expense and complexity associated with building in-house metering solutions to track usage within specific business units. It enables hosting providers to quickly and cost-efficiently create a more advanced, reliable, usage-based billing solution that adjusts to the provider's business model and strategy.

Use of ACLs in network metering port

Enterprises pay for the Internet traffic in and out of their datacenters, but not for the network traffic within their datacenters. For this reason, providers generally consider Internet and intranet traffic separately for the purposes of billing. To differentiate between Internet and intranet traffic, providers can measure incoming and outgoing network traffic for any IP address range, by using network metering port ACLs.

Virtual machine metrics

Windows Server 2012 provides two options for administrators to obtain historical data on a client's use of virtual machine resources: Hyper-V cmdlets in Windows PowerShell and the new APIs in the Virtualization WMI provider. These tools expose the metrics for the following resources used by a virtual machine during a specific period of time:

- Average CPU usage, measured in megahertz over a period of time.
- Average physical memory usage, measured in megabytes.
- Minimum memory usage (lowest amount of physical memory).
- Maximum memory usage (highest amount of physical memory).
- Maximum amount of disk space allocated to a virtual machine.
- Total incoming network traffic, measured in megabytes, for a virtual network adapter.
- Total outgoing network traffic, measured in megabytes, for a virtual network adapter.

Movement of virtual machines between Hyper-V hosts—for example, through live, offline, or storage migrations—does not affect the collected data.

Hyper-V Support for Large Sector Disks

Hyper-V in Windows Server 2012 introduces support for large sector size disks to help ensure compatibility with emerging innovations in storage hardware.

The data storage industry is transitioning the physical format of hard disk drives from 512-byte sectors to 4096-byte sectors (also known as 4K or 4 KB sectors). This transition is driven by several factors, including increases in storage density and reliability.

However, most of the software industry has depended on disk sectors of 512 bytes in length. A change in sector size introduces compatibility issues in many applications. To minimize the impact on the industry, hard drive vendors are introducing transitional 512-byte emulation (512e) drives. These drives offer some of the advantages offered by 4 KB native drives, such as improved format efficiency and an improved scheme for error correction codes (ECC), but with fewer compatibility issues than would be experienced by exposing a 4 KB sector size at the disk interface.

Key benefits

The storage industry is introducing 4 KB physical format drives to provide increased capacity. This new format is expected to be widely adopted by customers. Updates to the virtualization stack in Hyper-V in Windows Server 2012 helps ensure compatibility for storage configurations where the underlying physical hard disk uses the new 4 KB format.

Requirements

This feature requires physical disk drives that use either the 512e or native 4 KB format.

Technical overview

Hyper-V in Windows Server 2012 provides the following:

Support for improved performance of VHDs on 512e disks

512e disks can perform write operations only in terms of a physical sector—that is, it cannot directly write a 512-byte sector write issued to it. The internal process in the disk that makes this write possible consists of the following steps:

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

This process is called a read modify write (RMW). This RMW process causes performance degradation in virtual hard disks (VHDs) because to the following reasons:

- Dynamic and differencing VHDs have a 512-byte sector bitmap in front of their data payload. In addition, footer/header/Parent Locators all align to a 512-byte sector. Therefore, it is common for the VHD driver to issue 512-byte writes to update these structures, which results in the RMW behavior described above.
- It is common for applications to issue reads and writes in multiples of 4 KB sizes (the default cluster size of NTFS). Because there is a 512-byte sector bitmap in front of the data payload block of dynamic and differencing VHDs, the 4 KB blocks are not aligned to the physical 4 KB boundary. Each 4 KB write issued by the current parser to update the payload data results in two reads for two blocks on the disk, which are then updated and subsequently written back to the two disk blocks.

The overall performance impact to the workloads was in the range of 30 percent to 80 percent and was even higher at times.

Hyper-V mitigates the performance impact of 512e disks on the VHD stack by padding the previously mentioned structures to be aligned to 4 KB boundaries in the VHD format. This mitigates the RMW impact when accessing the data within the VHD file, as well as when updating the VHD metadata structures.

Support for hosting VHDs on native 4 KB disks

The current VHD driver assumes a physical sector size of 512 bytes and issues 512-byte I/Os, which makes it incompatible with these disks. As a result, the current VHD driver cannot open VHD files on physical 4 KB sector disks. Hyper-V makes it possible to store VHDs on 4 KB disks by implementing a software RMW algorithm in the VHD layer to convert the 512-byte access and update request to the VHD file to corresponding 4 KB accesses and updates.

Hyper-V Virtual Fibre Channel

You need your virtualized workloads to connect easily and reliably to your existing storage arrays. Windows Server 2012 provides Fibre Channel ports within the guest operating system, which allows you to connect to Fibre Channel directly from within virtual machines. This feature protects your investments in Fibre Channel, helping you to virtualize workloads that use direct access to Fibre Channel storage, allows you to cluster guest operating systems over Fibre Channel, and provides an important new storage option for servers hosted in your virtualization infrastructure.

Key benefits

With this Hyper-V virtual Fibre Channel feature, you can connect to Fibre Channel storage from within a virtual machine. This allows you to use your existing Fibre Channel investments to support virtualized workloads. Support for Fibre Channel in Hyper-V guests also includes support for many related features, such as NPIV, virtual SANs, live migration, and MPIO.

Requirements

The virtual Fibre Channel feature in Hyper-V requires the following:

- One or more installations of Windows Server 2012 with the Hyper-V role installed. Hyper-V requires a computer with processor support for hardware virtualization.
- A computer with one or more Fibre Channel host bus adapters (HBAs) that have an updated HBA driver that supports Virtual Fibre Channel.
- Virtual machines configured to use a virtual Fibre Channel adapter, which must use Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 as the guest operating system.
- Connection only to data logical unit numbers (LUNs). Storage accessed through a virtual Fibre Channel connected to a LUN cannot be used as boot media.

Technical overview

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

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

NPIV support

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

Virtual SAN support

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

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

Live migration

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

Figure 25: Alternating WWN addresses during a live migration

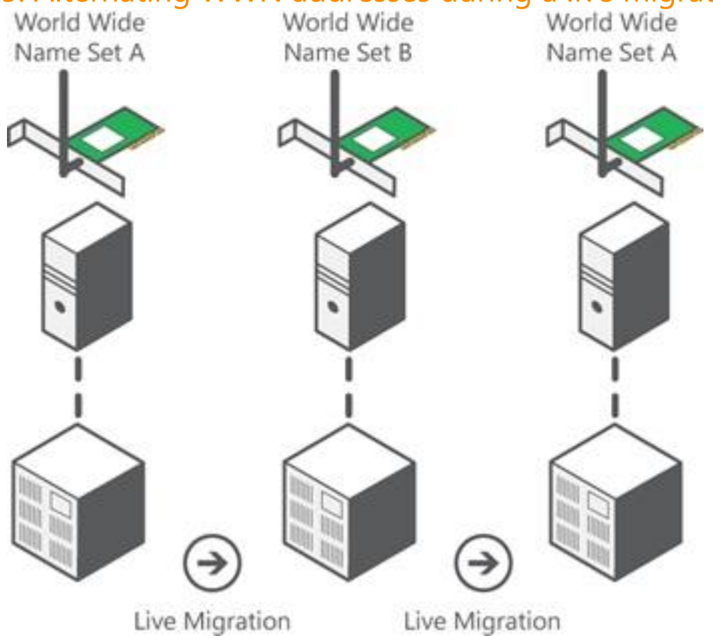

Multipath I/O functionality

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

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

Hyper-V Virtual Hard Disk Format

As enterprise workloads for virtual environments grow in size and in performance demands, virtual hard disk (VHD) formats need to accommodate them. Hyper-V in Windows Server 2012 introduces a new version of the VHD format called VHDX, which is designed to handle current and future workloads.

VHDX has a much larger storage capacity than the older VHD format. It also provides data corruption protection during power failures and optimizes structural alignments of dynamic and differencing disks to prevent performance degradation on new, large-sector physical disks.

Key benefits

The new VHDX format in Windows Server 2012 addresses the technological demands of an evolving enterprise by increasing storage capacity, protecting data, and ensuring quality performance on large-sector disks.

Technical overview

The main new features of the VHDX format are:

- Support for VHD storage capacity of up to 64 TB.
- Protection against data corruption during power failures by logging updates to the VHDX metadata structures.
- Improved alignment of the VHD format to work well on large sector disks.

The VHDX format also provides the following features:

- Larger block sizes for dynamic and differencing disks, which allow these disks to attune to the needs of the workload.
- A 4-KB logical sector virtual disk that allows for increased performance when used by applications and workloads that are designed for 4-KB sectors.
- The ability to store custom metadata about the file that the user might want to record, such as operating system version or patches applied.
- Efficiency in representing data (also known as “trim”), which results in smaller file size and allows the underlying physical storage device to reclaim unused space. (Trim requires physical disks directly attached to a virtual machine or SCSI disks, and trim-compatible hardware.

Hyper-V Virtual Switch

The Hyper-V virtual switch in Windows Server 2012 introduces a number of capabilities that are requested by customers for tenant isolation, traffic shaping, protection against malicious virtual machines, and easier troubleshooting of issues. This section focuses on the improvements in open extensibility and manageability for non-Microsoft extensions. Non-Microsoft extensions can be written to emulate the full capabilities of hardware-based switches and support more complex virtual environments and solutions.

The Hyper-V virtual switch is a layer 2 virtual network switch that provides programmatically managed and extensible capabilities to connect virtual machines to a physical network. The virtual switch provides policy enforcement for security, isolation, and service levels. With support for Network Device Interface Specification (NDIS) filter drivers and Windows Filtering Platform (WFP) callout drivers, the Hyper-V virtual switch enables you to use non-Microsoft extensible plug-ins that can provide enhanced networking and security capabilities.

The Hyper-V virtual switch enables you to implement and manage virtualized datacenters by providing the following:

- **Open platform.** The virtual switch is built on an open platform that allows independent software vendors to add or extend the capabilities provided natively in the virtual switch. The capabilities of the virtual switch function with the added capabilities of extensions.
- **Standard API.** The programming model for the extensible switch uses the same application programming interface (API) that is used for network filters and drivers in previous versions of Windows; that is, NDIS and WFP. There are new APIs and parameters added for virtual switch ports.
- **Windows reliability and quality.** The Windows platform and Windows Hardware Quality Logo (WHQL) program set high standards for extension quality.
- **Policy and configuration integration.** The management of extensions is integrated into the Windows management through Windows Management Instrumentation (WMI) calls and Windows PowerShell cmdlets, providing a standard management approach. Policies of extensions are automatically migrated with the virtual machine configuration during live migration.
- **Easy to troubleshoot.** Event logs and unified tracing are included with the virtual switch, which make it easier to diagnose and troubleshoot issues when they occur.

Requirements

Hyper-V virtual switch extensibility is built into the Hyper-V role, and it requires Windows Server 2012.

Technical overview

Hyper-V Virtual switch extensibility provides ways to extend the virtual switch for independent software vendors to add monitoring, filtering, and forwarding functionality without replacing all of the virtual switch functionality. Extensions are implemented by using NDIS filter drivers and WFP callout drivers, which are two public platforms for extending Windows networking functionality.

- **NDIS filter drivers.** An NDIS filter driver is a filtering service to monitor or modify network packets in Windows. NDIS filters were introduced with the NDIS 6.0 specification.
- **WFP callout drivers.** WFP, introduced in Windows Vista and Windows Server 2008, helps independent software vendors (ISVs) create WFP callout drivers that filter and modify TCP/IP packets, monitor or authorize connections, filter traffic protected by Internet Protocol security (IPsec), and filter remote procedure calls (RPCs). Filtering and modifying TCP/IP packets provides unprecedented access to the TCP/IP packet processing path. In this path, organizations can examine or modify outgoing and incoming packets before additional processing occurs. By accessing the TCP/IP processing path at different layers, organizations can more easily create firewalls, antivirus software, diagnostic software, and other types of applications and services. For more information, see "Windows Filtering Platform".

Extensions can extend or replace three aspects of the switching process: ingress filtering, destination look-up and forwarding, and egress filtering. In addition, extensions can gather statistical data by monitoring traffic at different layers of the virtual switch. Multiple monitoring and filtering extensions can be supported at the ingress and egress portions of the virtual switch. If organizations use a forwarding extension, only one instance of the extension can be used per switch instance, in which case it overrides the default forwarding of the virtual switch.

The following table shows the types of extensions, their purpose, examples, and how to implement them.

Table 7: Extension types for virtual switches

Extension	Purpose	Examples	Extensibility Component
Network packet inspection	View network packets for virtual machine to virtual machine traffic per virtual switch Cannot alter network packets	sFlow Network monitoring	NDIS filter driver
Network packet filter	Create, filter, and modify network packets that are entering or leaving the virtual switch and in virtual machine to virtual machine traffic	Security	NDIS filter driver
Network forwarding	Provide forwarding extension per virtual switch, which bypasses default forwarding (maximum of one per virtual switch)	OpenFlow, Virtual Ethernet Port Aggregator (VEPA) Proprietary network fabrics	NDIS filter driver
Intrusion detection or firewall	Filter and modify TCP/IP packets, monitor or authorize connections, filter traffic that is protected by IPsec, and filter RPCs	Virtual firewall Connection monitoring	WFP callout driver

The virtual switch includes the following attributes, which provide an open switch API that enables enhanced switch and management products to work with Hyper-V:

- **Extensible.** You can add new features and enhancements while retaining the native feature set.
- **Standard API.** You can build extensions on NDIS and the WFP, which are publicly available platforms that are familiar to Windows developers.
- **Live migration support.** You can use extensions in Hyper-V live migration.
- **Easy troubleshooting.** Support is provided for unified tracing for the virtual switch and its extensions.

Figure 26 shows the architecture of the virtual switch and the extensibility model.

Figure 26: Virtual switch architecture and extensibility

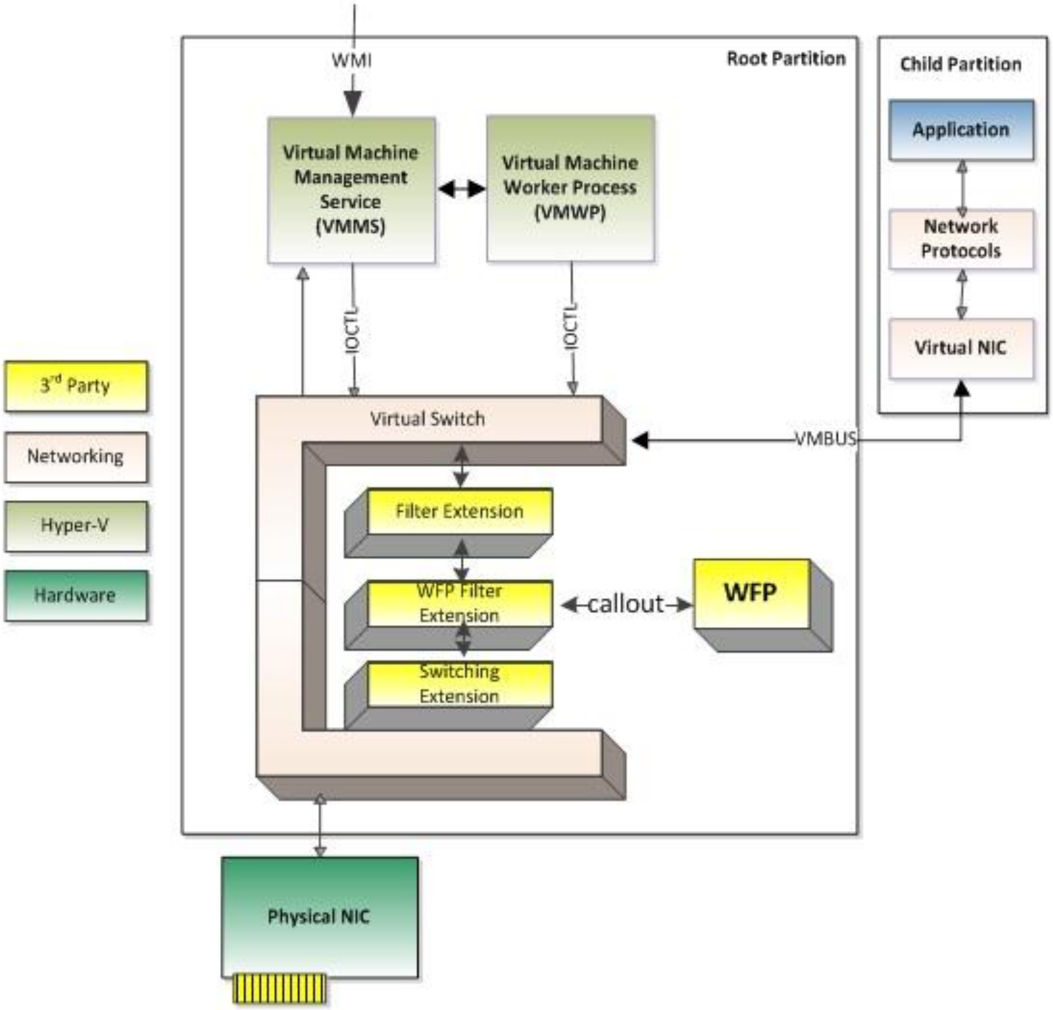

Manageability

With the Hyper-V virtual switch playing a similar role for virtual machines that physical network switches provide for physical devices, it is important that you can easily manage, troubleshoot, and resolve network issues. To accomplish this, Windows Server 2012 provides the following:

- **Windows PowerShell and scripting support.** Windows Server 2012 provides Windows PowerShell cmdlets for the virtual switch that you can use to build command-line tools or enable automation of scripts for setup, configuration, monitoring, and troubleshooting. Windows PowerShell also enables non-Microsoft developers to build tools to manage the virtual switch.
- **Unified Tracing and enhanced diagnostics.** Unified Tracing has been extended into the virtual switch to provide two levels of troubleshooting. At the first level, Event Tracing for Windows (ETW) enables organizations to trace packet events through the virtual switch and extensions, making it easier to isolate where an issue is occurring. The second level enables you to capture packets for a full trace of events and traffic packets.

Reliability and extension management

The Hyper-V virtual switch is an open platform that enables multiple vendors to provide extensions that are written to standard Windows API frameworks. The reliability of extensions is strengthened through the Windows standard framework and the reduction of non-Microsoft code required for functions, and the reliability is backed by the WHQL certification program. The virtual switch and its extensions are managed by using Windows PowerShell cmdlets, WMI calls, or the Hyper-V Manager user UI.

Hyper-V Support for Scaling Up and Scaling Out

The use of virtualization has gained widespread acceptance as a means for lowering costs through consolidation of multiple server roles, typically hosting utility and productivity applications. These workloads generally consume comparatively little in the way of CPU and memory resources, and generate only moderate amounts of I/O. To support these workloads, configuring a virtual machine to use 2 or 4 virtual processors and relatively modest amounts of memory is sufficient. Hyper-V in Windows Server 2008 R2 supported virtual machines with a maximum of virtual processors, and up to 64 GB of memory.

However, IT organizations are increasingly looking to take advantage of virtualization to deploy mission-critical, tier-1 business applications. These larger and more demanding workloads—including high-end, online transaction processing (OLTP) databases and online transaction analysis (OLTA) business intelligence solutions—are typically run on systems with 16 or more processors, and demand large amounts of memory. For example, for SQL workloads of this type, a general practice is to allocate 8 GB of memory per logical processor. For this class of workloads, fewer virtual machines would typically be run on each virtualization host, but each virtual machine would require more virtual processors and greater amounts of virtual machine memory than current less demanding workloads.

As workloads demand greater and greater system resources to run these business-critical applications, highly scalable servers continue to expand the limits of processor core counts, and offer increased system memory capacity. As multicore processors continue to evolve with increasing core counts, these high-end servers push the boundaries of scale support for operating systems and virtualization hosts.

Hyper-V in Windows Server 2012 supports running on large host systems through expanded support for host processors and memory, and it enables the virtualization of high-performance, scale-up workloads by supporting the configuration of large, high-performance virtual machines.

Key benefits

Hyper-V hosts in Windows Server 2012 provide support for 320 logical processors and 4 TB of system memory, which enables customers to run Hyper-V on the largest scale-up server systems currently available. Hyper-V can help customers to virtualize their most-demanding, mission-critical, tier-1 workloads by supporting large, high-performance virtual machines with up to 64 virtual processors and 1 TB of memory in a virtual machine. By projecting a virtual NUMA topology into large virtual machines, the guest operating system and applications such as SQL Server can leverage their existing thread scheduler and memory allocation optimizations, which ensures maximum performance and scalability of demanding workloads in a virtual machine. Support for SR-IOV–capable systems and network devices enables SR-IOV–capable network adapters to be assigned directly to a virtual machine, which maximizes network throughput while minimizing network latency as well as the CPU overhead required for processing network traffic.

Requirements

- One or more installations of Windows Server 2012 with the Hyper-V role installed. Hyper-V requires a server that is capable of running Hyper-V. Specifically, it must have processor support for hardware virtualization.
- The number of virtual processors that may be configured in a virtual machine depends on the number of processors on the physical computer. You must have at least as many logical processors in the virtualization host as the number of virtual processors required in the virtual machine. For example, to configure a virtual machine with the maximum of 64 virtual processors, you must be running Hyper-V on a virtualization host that has 64 or more logical processors.
- SR-IOV networking requires:
  - A host system which supports SR-IOV (for example, Intel VT-d), including chipset support for interrupt and DMA remapping, and proper firmware support to enable and describe the platform's SR-IOV capabilities to the operating system.
  - An SR-IOV–capable network adapter and driver in both the management operating system (which runs the Hyper-V role) and each virtual machine where a virtual function is assigned.

Technical overview

Hyper-V in Windows Server 2012 supports running on large server systems and enables the virtualization of high-performance, scale-up workloads by including the following changes and features:

Increased hardware support for the virtualization host. Hyper-V supports running on a host system with up to 320 logical processors and 4 TB of memory, which helps ensure compatibility with the largest scale-up server systems.

Support for large virtual machines. Hyper-V supports configuring a virtual machine with up to 64 virtual processors and up to 1 TB of memory, which is a significant increase from previous versions.

NUMA support in a virtual machine. Non-Uniform Memory Architecture, or (NUMA), is a computer architecture used in multiprocessor systems in which the time required for a processor to access memory depends on the memory's location relative to the processor. With NUMA, a processor can access local memory (memory attached directly to the processor) faster than it can access remote memory (local to another processor in the system). Modern operating systems and high-performance applications such as SQL Server have developed optimizations that recognize the system's NUMA topology and consider NUMA when scheduling threads or allocating memory to increase performance.

Projecting a virtual NUMA topology into a virtual machine enables optimal performance and workload scalability in large virtual machine configurations by allowing the guest operating system and applications such as SQL to leverage their inherent NUMA performance optimizations. The default virtual NUMA topology projected into a Hyper-V virtual machine is optimized to match the host's NUMA topology.

Note

If a virtual machine is configured to use Dynamic Memory, only one virtual NUMA node (that is, a flat NUMA topology) will be projected into the guest, which effectively disables virtual NUMA support.

Support for SR-IOV networking devices. Single Root I/O Virtualization (SR-IOV) is a standard introduced by the PCI-SIG. SR-IOV works in conjunction with system chipset support for virtualization technologies. This provides remapping of interrupts and DMA and enables SR-IOV capable devices to be assigned directly to a virtual machine. Hyper-V enables support for SR-IOV–capable network devices and enables an SR-IOV virtual function of a physical network adapter to be assigned directly to a virtual machine. This increases network throughput and reduces network latency, while also reducing the host CPU overhead required for processing network traffic.

Installation Options

This section summarizes the differences between the installation options available for Windows Server 2012, including the features that are installed with each option, the management options available after installation, and how to switch between the installation options during use. It also explains the differences between the Server Graphical Shell and the Minimal Server Interface and how to switch between them. In addition, it discusses how to use Features on Demand to further reduce the disk footprint by including the binary files for only the server roles organizations actually use.

Installation options description

When you install Windows Server 2012, you can choose between *Server Core Installation* and *Server with a GUI*. The Server with a GUI option is the Windows 8 equivalent of the Full installation option available in Windows Server 2008 R2. The Server Core Installation option reduces the space required on disk, the potential attack surface, and especially the servicing requirements, so we recommend that you choose the Server Core installation unless you have a particular need for the additional UI elements and graphical management tools that are included in the Server with a GUI option. For this reason, the Server Core installation is now the default. Because you can freely switch between these options at any time later, one approach might be to initially install the Server with a GUI option, use the graphical tools to configure the server, and then later switch to the Server Core Installation option.

An intermediate state is possible where you start with a Server with a GUI installation and then removes Server Graphical Shell, resulting in a server that comprises the Minimal Server Interface, Microsoft Management Console (MMC), Server Manager, and a subset of Control Panel. See the "Minimal Server Interface" section of this document for more information.

In addition, after installation of either option is complete, you can completely remove the binary files for server roles and features that you do not need, thereby conserving disk space and reducing the attack surface still further. For more information, see the "Features on Demand" section of this document.

For the smallest possible installation footprint, users should start with a Server Core installation and remove any server roles or features you do not need by using Features on Demand.

Server Core installation option

With this option, the standard UI (the Server Graphical Shell) is not installed; you manage the server using the command line, Windows PowerShell, or by remote methods.

- User interface consists of command prompt (Server Graphical Shell is not installed)
- Local installing, configuring, and uninstalling of server roles at a command prompt with Windows PowerShell
- Remote installing, configuring, and uninstalling of server roles with Server Manager, Remote Server Administration Tools (RSAT), or Windows PowerShell
- Microsoft Management Console not available locally
- Desktop Experience not available
- Server roles available:
    - Active Directory Certificate Services

- Active Directory Domain Services
- DHCP Server
- DNS Server
- File Services (including File Server Resource Manager)
- Active Directory Lightweight Directory Services (AD LDS)
- Hyper-V
- Print and Document Services
- Streaming Media Services
- Web Server (including a subset of ASP.NET)
- Windows Server Update Server
- Active Directory Rights Management Server
- Routing and remote access server

To convert to a Server with GUI installation with Windows PowerShell, follow the steps in the procedure below:

To use Windows PowerShell to convert from a Server Core installation to a Server with a GUI installation:

1. Create a folder to mount a Windows Imaging File (WIM) in with the command

   ```
   mkdir c:\mountdir
   ```

2. Determine the index number for Server Datacenter using this command at an elevated command prompt:

   ```
   Dism /get-wiminfo /wimfile:<drive>:sources\install.wim
   ```

3. Mount the WIM file using this command at an elevated command prompt

   ```
   Dism /mount-wim /WimFile:<drive>:\sources\install.wim /Index:<#_from_step_2>
   /MountDir:c:\mountdir /readonly
   ```

4. Start Windows PowerShell, and run this cmdlet:

   ```
   Install-WindowsFeature Server-Gui-Mgmt-Infra,Server-Gui-Shell –Restart –Source
   c:\mountdir\windows\winsxs
   ```

5. Alternatively, if you want to use Windows Update as the source instead of a WIM file, use this Windows PowerShell cmdlet:

   ```
   Install-WindowsFeature Server-Gui-Mgmt-Infra,Server-Gui-Shell –Restart
   ```

Server with a GUI option

With this option, the standard user interface and all tools are installed. Server roles and features are installed with Server Manager or by other methods.

- Standard graphical user interface (Server Graphical Shell). The Server Graphical Shell includes the new Windows Experience Start screen, but does not include support for apps with the new Windows UI. To enable support for apps with the new Windows UI, install the Desktop Experience feature.
- Local installing, configuring, and uninstalling of server roles with Server Manager or Windows PowerShell.
- Remote installing, configuring, uninstalling of server roles with Server Manager, Remote Server, RSAT, or Windows PowerShell.
- Microsoft Management Console installed.
- **Desktop Experience** can be installed with Server Manager or Windows PowerShell.
- **To convert to a Server Core installation with Windows PowerShell**, run the following cmdlet:
  ```
  Uninstall-WindowsFeature Server-Gui-Mgmt-Infra -restart
  ```

If you initially install with the Server with a GUI option and then use the above command to convert to a Server Core installation, you can later revert to a Server with a GUI installation without specifying a source. This is because the necessary files remain stored on the disk, even though they are no longer installed. For more information, and for instructions to completely remove the Server with a GUI files from disk, see the "Features on Demand" section of this document.

If you convert to a Server Core installation, Windows features, server roles, and GUI management tools that require a Server with a GUI installation will be uninstalled automatically. You can specify the `-WhatIf` option in Windows PowerShell to see exactly which features will be affected by the conversion.

Minimal Server Interface

In Windows Server 2012, you can remove the Server Graphical Shell, resulting in a "Minimal Server Interface." This is similar to a Server with a GUI installation, but Microsoft Internet Explorer 10, Windows Explorer, the desktop, and the Start screen are not installed. Microsoft Management Console (MMC), Server Manager, and a subset of Control Panel are still present.

If organizations start with a Server with a GUI installation, you can convert to the Minimal Server Interface at any time using Server Manager.

Note

When users change any of these options, they will have to restart the server for the change to take effect.

The table below for a summary of which selections to make in Server Manager (or cmdlets to use in Windows PowerShell) in order to get a given installation state:

Table 8: Installation options

To reach the installation state in each column	Server Core installation option	Minimal Server Interface	Server with a GUI installation option	Desktop Experience feature installed
Select these features in Server Manager	none	Graphical Management Tools and Infrastructure	Graphical Management Tools and Infrastructure Server Graphical Shell	Graphical Management Tools and Infrastructure Server Graphical Shell Desktop Experience
Run the Windows PowerShell install/uninstall commands with these values for the Name parameter:	none	Server-Gui-Mgmt-Infra	Server-Gui-Mgmt-Infra, Server-Gui-Shell	Server-Gui-Mgmt-Infra Server-Gui-Shell Desktop Experience

Features on Demand

In previous versions of Windows, even if a server role or feature was disabled, the binary files for it were still present on the disk, consuming space. In Windows Server 2012, not only can you disable a role or feature, but you can also completely remove its files, a state called "disabled with payload removed." To reinstall a role or feature that is disabled with payload removed, you must have access to an installation source.

To completely remove a role or feature, use `–Remove` with the `Uninstall-WindowsFeature` cmdlet of Windows PowerShell. For example, to completely remove Windows Explorer, Internet Explorer, and dependent components, run the following Windows PowerShell command:

```
Uninstall-WindowsFeature Server-Gui-Shell -remove
```

To install a role or feature for which the payload has been removed, use the Windows PowerShell `–Source` option of the `Install-WindowsFeature` Server Manager cmdlet. The `–Source` option specifies a path to a WIM mount point. If you do not specify a `–Source` option, Windows will use Windows Update by default.

Only component sources from the exact same version of Windows are supported. For example, a component source derived from the Windows Server Developer Preview is not a valid installation source for a server running Windows Server 2012.

Permissions might affect the system's ability to access Windows features for installation over a network. The Trusted Installer process runs within a system account that impersonates a user account. If you encounter network access issues, try issuing a net use command (for example, **net use * \\path\to\network**) to mount the network source and then point to the mounted network path.

Practical applications

These examples give you an idea of how you can choose the installation option that might be most appropriate for your deployment needs:

- Server Core installations require approximately 4 GB less space than a Server with a GUI installation. By using Server Core installations on virtual machines, you can save a significant amount of space by not having to store the GUI files on a virtual machine disk.
- Servers often have comparatively large amounts of memory and complex disk arrays, both of which can take a significant amount of time to initialize at startup. Because Server Core installations minimize the number of restarts required for updates, the frequency at which disk arrays and memory must be re-initialized is reduced.
- Certain server applications have dependencies on certain Windows services, libraries, applications, and files that are not available in Server Core installations, but the administrator wants to take advantage of the reduced need for updating typical Server Core installations. The Minimal Server Interface offers additional compatibility while still maintaining a reduced system footprint (though to a lesser extent than a Server Core installation).
- Features on Demand helps reduce the footprint of your virtual machine deployments by removing roles and features that will never be deployed in your virtual machines. Depending on the roles and features used in your virtual machines, it is possible to reduce the size by over 1 GB.

Reference table

This table summarizes which features are available locally depending on which installation option you choose.

Table 9: Management features available

	Server Core installation option	Minimal Server Interface	Server with a GUI installation option
Command prompt	Available	Available	Available
Windows PowerShell/ Windows .NET	Available	Available	Available
Server Manager	Not available	Available	Available
Microsoft Management Console	Not available	Available	Available
Control Panel	Not available	Not available	Available
Control Panel applets	Not available	Some available	Available
Windows Explorer	Not available	Not available	Available
Taskbar	Not available	Not available	Available
Notification area	Not available	Not available	Available
Internet Explorer	Not available	Not available	Available
Built-in help system	Not available	Not available	Available

IP Address Management

Monitoring and managing the IP infrastructure on a corporate network is a critical part of network administration, and it has become increasingly challenging as networks grow more dynamic and complex. Factors that add complexity to IP administration include:

1. Business expansions, such as new datacenters
2. Business alterations, such as the increased use of mobile devices and multiple devices per user
3. Distributed networks with multiple infrastructure servers
4. New technology and scenario adoptions, such as IPv6 and virtualization
5. Compliance requirements such as HIPAA and Sarbanes-Oxley
6. Limited personnel and financial resources
7. Utilization and reporting requirements for public IPv4 address space

To meet these needs, many IT administrators still track IP address allocation and utilization manually, using spreadsheets or custom database applications. This can be very time consuming and resource intensive, and is inherently prone to user error.

IP address management

Windows Server 2012 introduces a built-in framework for discovering, monitoring, auditing, and managing the IP address space that is used on a corporate network. IP Address Management (IPAM) in Windows Server 2012 provides components to administer and monitor your IP infrastructure, including:

- Automatic IP address infrastructure discovery
- Custom IP address space display, reporting, and management
- Audit of server configuration changes and tracking of IP address usage
- Monitoring and management of Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS) services

IPAM architecture

An IPAM server is typically a computer that is a member of a domain; however you cannot install the IPAM feature on an Active Directory domain controller.

There are two general methods to deploy IPAM servers:

1. Distributed – An IPAM server is deployed at every site in an enterprise.
2. Centralized – One IPAM server serves an enterprise.

An example of the distributed IPAM deployment method is shown in Figure 27, with one IPAM server located at the corporate headquarters and one at each branch office. There is no communication or database sharing between different IPAM servers in the enterprise. If multiple IPAM servers are deployed, you can customize the scope of discovery for each IPAM server, or filter the list of managed servers. A single IPAM server might manage a specific domain or location, perhaps with a second IPAM server configured as a backup.

Figure 27: IPAM network architecture

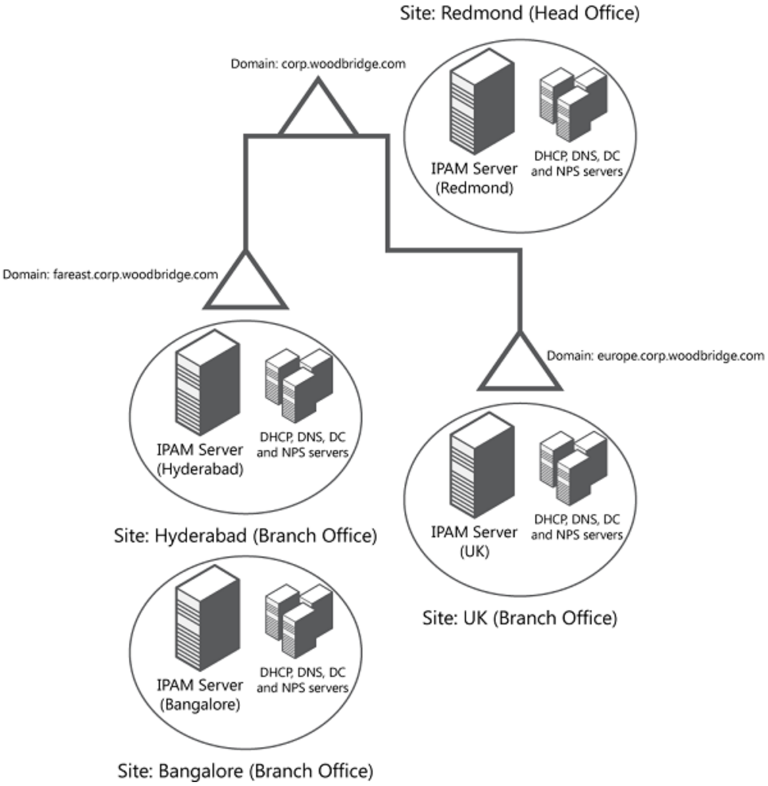
Site: Redmond (Head Office)
Domain: corp.woodbridge.com
IPAM Server (Redmond)
DHCP, DNS, DC and NPS servers
Domain: fareast.corp.woodbridge.com
Domain: europe.corp.woodbridge.com
IPAM Server (Hyderabad)
DHCP, DNS, DC and NPS servers
Site: Hyderabad (Branch Office)
IPAM Server (UK)
DHCP, DNS, DC and NPS servers
Site: UK (Branch Office)
IPAM Server (Bangalore)
DHCP, DNS, DC and NPS servers
Site: Bangalore (Branch Office)

IPAM will periodically attempt to locate Network Policy Server (NPS), DNS, DHCP servers, and domain controllers on the network that are within the scope of discovery that you specify. You must specify whether these servers are managed or not managed by IPAM. To be managed by IPAM, server security settings and firewall ports must be configured to allow the IPAM server access to perform the required monitoring and configuration functions. You can configure these settings manually or automatically by using Group Policy Objects (GPOs). If you select the automatic method, the settings are applied when a server is configured as *Managed*, and settings are removed when it is configured as ***Unmanaged***.

The IPAM server communicates with managed servers by using a remote procedure call (RPC) or Windows Management Instrumentation (WMI) interface, as shown in Figure 28. IPAM monitors domain controllers and NPS servers for IP address tracking purposes. In addition to monitoring functions, several DHCP server and scope properties can be configured by using IPAM. Zone status monitoring and a limited set of configuration functions are also available for DNS servers.

Figure 28: IPAM server communications

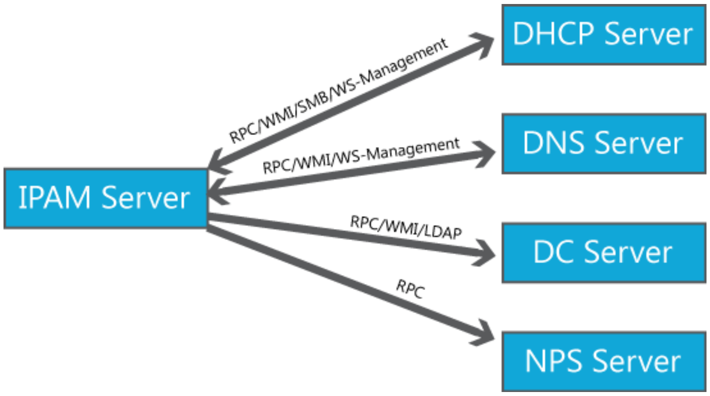

IPAM requirements

The scope of IPAM server discovery is limited to a single Active Directory forest. The forest itself might include of a mix of trusted and untrusted domains.

IPAM requires membership in an Active Directory domain, and it requires a functional network infrastructure environment to integrate with existing DHCP, DNS, domain controller, and NPS installations across the forest.

IPAM has the following specifications:

- IPAM supports only DHCP, DNS, domain controllers, and NPS servers that are running Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012.
- IPAM supports only domain member DHCP, DNS, and NPS servers in a single Active Directory forest.
- A single IPAM server is validated to support up to 150 DHCP servers and 500 DNS servers. More servers might be supported, depending on the environment.
- A single IPAM server can support up to 6,000 DHCP scopes, 40,000 static and dynamic IP address ranges, and 300 DNS zones (these are the limits tested).
- A single IPAM server is validated to support three million IPv4 and three million IPv6 addresses.
- IPAM stores three years of forensics data (such as IP address leases, host MAC addresses, and user logon and logoff information) for 100,000 users in Windows Internal Database. There is no database purge policy provided, and the administrator must purge the data manually as needed.
- IPAM does not support management and configuration of non-Microsoft network elements (for example, Windows Internet Name Service, DHCP relays, or proxies).
- IPAM supports only Windows Internal Database. No external databases are supported.
- IP address utilization trends are provided only for IPv4.
- IP address reclaiming support is provided only for IPv4.
- No special processing is performed for IPv6 stateless address auto configuration private extensions.
- No special processing is provided for virtualization technology or virtual machine migration.
- IPAM does not check for IP address consistency with routers and switches.
- IPAM does not support auditing of IPv6 address (stateless address auto configuration) on an unmanaged computer to track users.

For recommended IPAM server hardware specifications, see [Windows Server 2008 R2 with SP1 System Requirements]().

IPAM concepts

Server discovery. IPAM supports automatic discovery, based in Active Directory, of DNS and DHCP servers on the network. Discovery is based on the domains and server roles that you select when you configure the scope of discovery, as shown in Figure 29.

Figure 29: Discovery settings dialog box

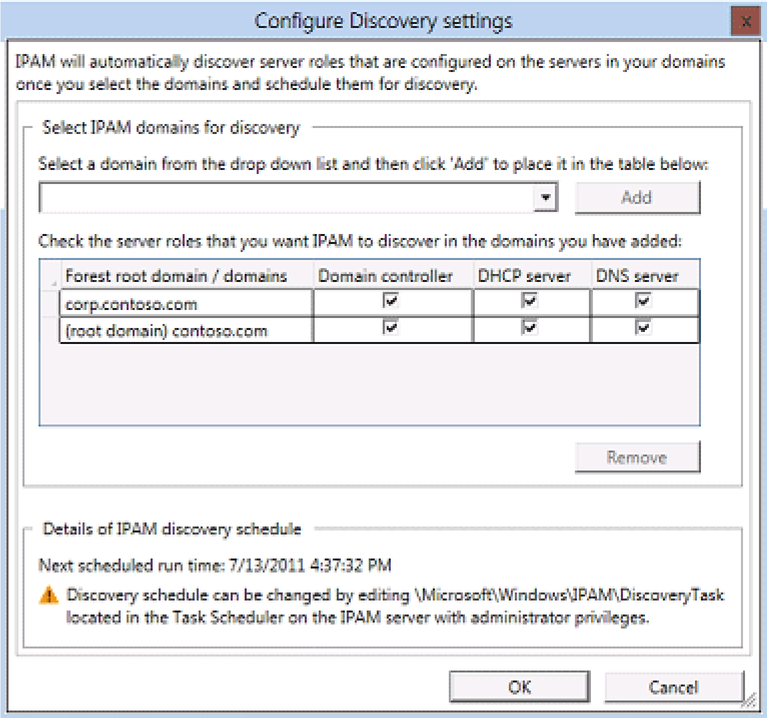

Configure Discovery settings

IPAM will automatically discover server roles that are configured on the servers in your domains once you select the domains and schedule them for discovery.

Select IPAM domains for discovery

Select a domain from the drop down list and then click 'Add' to place it in the table below:

Check the server roles that you want IPAM to discover in the domains you have added:

Forest root domain / domains	Domain controller	DHCP server	DNS server
corp.contoso.com	☑	☑	☑
(root domain) contoso.com	☑	☑	☑

Details of IPAM discovery schedule

Next scheduled run time: 7/13/2011 4:37:32 PM

Discovery schedule can be changed by editing \Microsoft\Windows\IPAM\DiscoveryTask located in the Task Scheduler on the IPAM server with administrator privileges.

IPAM discovers the domain controllers, DNS servers, and DHCP servers in the network and confirms their availability based on role-specific protocol transactions. In addition to automatic discovery, IPAM supports the manual adding of a server to the list of servers in the IPAM system as follows:

Managed servers. When you configure the manageability status of a server as *Managed*, this indicates that the server is part of the managed environment of the IPAM server. Data is retrieved from managed servers to display in various IPAM views. The type of data that is gathered depends on the server role.

Unmanaged servers. When you configure the manageability status of a server as *Unmanaged*, this indicates that the server is outside of the managed environment of the IPAM server. No data is collected by IPAM from these servers.

IPAM security groups. IPAM includes security groups with unique permissions settings that allow you to control the access level of specific users and groups. The available IPAM security groups are:

- IPAM Users – Members of this group have read-only permissions, and they can use all the views in the IPAM interface with the exception of the IP address tracking view.
- IPAM MSM Administrators – Members of the multi-server management (MSM) administrators group have read and write permissions to manage infrastructure servers (such as DHCP) and perform other IPAM common tasks. System administrators are typically members of this group.
- IPAM ASM Administrators – Members of the address space management (ASM) administrators group have read and write permissions to manage IP address space and perform other IPAM common tasks. Network administrators are typically members of this group.
- IPAM IP Tracking Administrators – Members of this group have permissions to view IP address tracking data on the network. You can use this group to protect the privacy information that might be contained in IP address tracking data.
- IPAM Administrators – Members of this group have permissions to view all IPAM data and perform all IPAM tasks.

IPAM data collection tasks. IPAM schedules the following tasks to retrieve data from managed servers to populate the IPAM views for monitoring and management. You can also modify these tasks by with Task Scheduler.

- Server Discovery – Automatically discovers domain controllers, DHCP servers, and DNS servers in the domains that you select.
- Server Configuration – Collects configuration information from DHCP and DNS servers to display in the IP address space and for server management functions.
- Address Utilization – Collects IP address space usage data from DHCP servers to display current and historical utilization.
- Event Collection – Collects DHCP and IPAM server operational events, as well as events from domain controllers, NPS, and DHCP servers for IP address tracking.
- Server Availability – Collects service status information from DHCP and DNS servers.
- Service Monitoring – Collects DNS zone status events from DNS servers.
- Address Expiry – Tracks IP address expiry state and logs notifications.

IP address mapping. In IPAM, an *IP address block* is a large chunk of IP addresses that is used to organize address spaces. *IP address ranges* are smaller chunks of IP addresses that typically correspond to a DHCP scope. IP address ranges are mapped to IP address blocks as shown in Figure 30.

Figure 30: IP block, range, and address

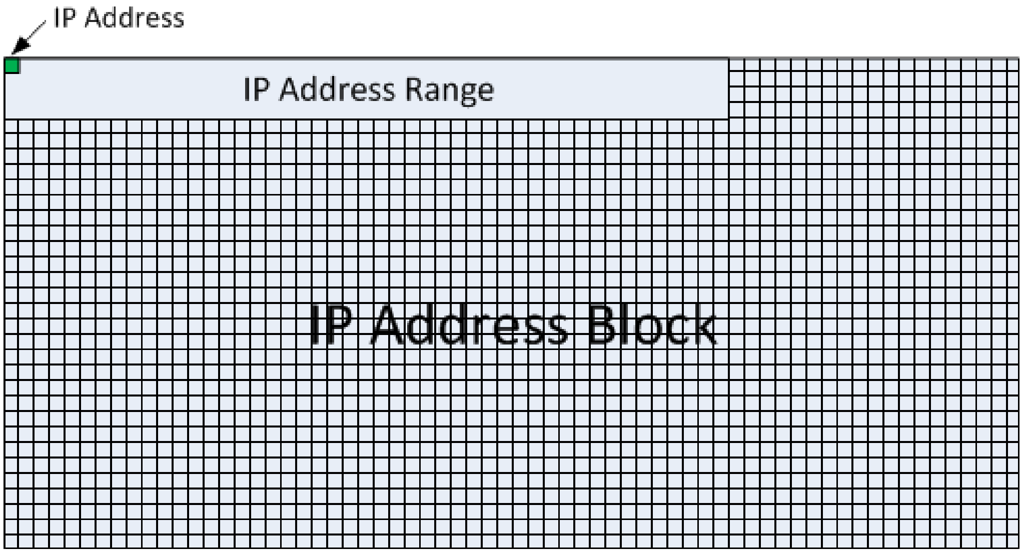

IP address blocks. IPAM automatically arranges IPv4 address blocks into public and private address spaces and IPv6 addresses into unicast global addresses. IP address blocks can be added, imported, edited, and deleted. If the start and end IP address of a block lies within the start and end IP address of another block, it is automatically arranged as a nested sub-block. IPAM automatically maps IP address ranges to the appropriate IP address block, based on the boundaries of the range. This enables a hierarchically organized view of the IP address ranges and a multilevel hierarchy of IP address blocks. IPAM rolls up utilization statistics and trends at the IP address block or IP address sub-block level, based on the ranges that are contained in the block.

IP address ranges. IP address ranges are the next hierarchical level of IP address space entities after IP address blocks. An IP range is an IP subnet that is marked by a start and end IP address. It is typically a DHCP scope or a static IPv4 or IPv6 address range or address pool that is used to assign addresses to hosts. IPAM enables you to centralize address ranges that might be spread across many heterogeneous systems, such as across multiple DHCP servers, virtual machine managers, or legacy spreadsheets.

IP addresses. IP addresses are the smallest entities that are provided under IP address ranges. IPAM enables end-to-end lifecycle management of IPv4 and IPv6 addresses, including record synchronization with DHCP and DNS servers. IPAM automatically maps an address to the appropriate range, based on the start and end address of the range. The value of the **Managed by Service** and **Service Instance** fields of an IP address should be identical to the IP address range to which it maps. If IPAM encounters duplicate addresses that are reported out of multiple systems, it can uniquely identify the addresses based on the **Managed by Service** and **Service Instance** fields that are associated with an IP address.

Unmapped address space. In the IP address space management view, an unmapped address space includes IP address ranges and individual IP addresses that are not mapped to the IP address blocks or that conflict with other IP address ranges or addresses in the IPAM system.

Managed by Service and **Service Instance.** Two IPAM settings that are associated with an IP address. *Managed by Service* refers to various types of systems that provide addressing services on the network, such as DHCP that is provided by Microsoft, non-Microsoft DHCP, or the Virtual Machine Management service. *Service Instance* refers to the specific server or network device.

Logical groups. By using the **Logical group** field in IPAM, you can assign IP addresses and ranges to unique, personalized, logical groups. This enables real-life visualization, management, and tracking of IP addresses.

IP address inventory. By default, IPAM defines a built-in logical group called the IP address inventory, which is organized by the device type field of the IP address.

Overlapped address space. IPAM supports overlapping address ranges managed by multiple systems. As long as the overlapping or conflicting IP address range maintains uniqueness based on the **Managed by Service** or **Service Instance** values, IPAM displays the ranges without any conflicts. IPAM prevents addition or import of conflicting address ranges that belong to the same **Managed by Service** or **Service Instance** value.

DNS and DHCP servers. IPAM displays DNS servers and DHCP servers by order of their network interface addresses (in /16 subnets for IPv4 and /48 subnets for IPv6). By using this view, IPAM enables periodic service availability monitoring for managed DNS servers and DHCP servers. You can view service status, such as running, stopped, or paused, displayed together with the duration of the service state. The same monitoring status can be filtered for only DHCP service or DNS service to visualize the overall service availability state. For DHCP server properties, the following actions are available for single or multiple DHCP servers: **Scope creation**, **Edit server properties**, **Configure server options**, **Define user classes**, **Define vendor classes**, **Set predefined options**, and **Export**.

DHCP scope properties. The DHCP scope properties view enables you to monitor scope utilization for DHCP scopes in the management list. DHCP scope utilization statistics are automatically and periodically polled by IPAM from managed DHCP servers. You can also track important scope properties such as scope name, subnet, and start and end IP addresses from the management list. The detailed list of all scope options that are configured is also available in the preview pane. The following actions are available for single or multiple DHCP scopes: Activate, Deactivate, Edit Scope properties and options, Delete, Export, and Duplicate scope.

IP Blocks view. This view for DHCP scopes enables you to visualize DHCP scopes organized in a multilevel hierarchy of IP address blocks. IP address blocks can be added, edited, or deleted from the IP range blocks node in IPAM. IPAM automatically maps a DHCP scope to the appropriate block, based on the start and end IP addresses of the scope. This enables a hierarchically organized view of the DHCP scopes by IP address blocks. IPAM compiles and presents utilization statistics at the IP address block or sub-block level.

DNS zone monitoring. IPAM enables DNS zone monitoring for DNS forward and reverse lookup zones. The zone status is derived by IPAM, based on zone events. From the Forward Lookup Zone node, IPAM displays a list of all forward lookup zones that are hosted by managed DNS servers. Their overall status is displayed, based on status from all the servers that are hosting that zone and the duration that the zone has been in that state. IPAM also displays a list of all authoritative servers for that zone in the preview pane, and it enables automatic hierarchical navigation of forward lookup zones. For each DNS server that is hosting the zone, IPAM displays the zone status on that server and the status duration. Other details such as the zone type, server availability, and IP address are also displayed. In addition, IPAM provides a catalog of all zone events from the server to assist with troubleshooting. With IPAM, you can visualize all IPv4 reverse lookup zones and IPv6 reverse lookup zones. A list of all authoritative servers that are hosting the selected reverse lookup zone is presented in the preview pane.

Custom fields. IPAM supports user defined extensible metadata that can be associated with IP address ranges, IP addresses, and servers. You can create metadata with multiple value types (such as Country/Region) or single value types (such as Building).

Custom logical groups. IPAM allows you to define the logical grouping of entities and visualize the utilization of address space based on these groups. Logical groups enable you to visualize IP address ranges from a business perspective rather than from a conventional hierarchy of IP address blocks. These logical groups can be customized and they can be hierarchical. Logical groups are defined by selecting the grouping criteria from built-in or user-defined custom fields. IPAM supports a multilevel hierarchy when you define a logical group. You can also create similar custom logical groups for IP addresses and servers. IPAM also presents utilization statistics and trends at the logical group level.

Trouble shooting. To detect IPAM access status of the servers managed by IPAM (access status of DHCP service, DNS service, and DNS event log readers group)

IPAM GPO-based provisioning. Support disjoint domain name space support, localization of DHCP users group, network share access handling, and removal of extra ACL'ing for DNS partition.

PowerShell Management. (get and add) for seamless import of meta-data as part of integration with other systems (System Center 2012, AD DS) and import of legacy address space tracking spreadsheet.

Integration with System Center 2012. PowerShell based integration module based on IPAM and System Center 2012 Virtual Machine Manager PowerShell 3.0 cmdlets. This integration helps enhance both cloud solution builder and IPAM solutions deployed by hosters, including:

- Centralization of address space required by Network Administrator for planning
- Utilization tracking and trend reporting
- Visualizing tenant-specific address allocation and utilization
- Overlapping address space management
- IP address lifetime tracking
- Synchronizing DHCP reservations and DNS records

IP address tracking. IPAM enables you to track IP addresses by using DHCP lease events and user logon events. IPAM periodically collects these events from managed DHCP servers, domain controllers, and NPS servers to provide a central repository of these relevant events that are the result of address activity in the enterprise network. The following filtered views are available:

- **By IP Address** – IPAM enables you to query DHCP lease events for a particular IP address over a specified time period. This displays the list of all relevant lease assignments, renew events, and release events for that IP address. You can also identify the host name and corresponding client ID (MAC address for IPv4 and DUID for IPv6) information that is associated with this IP address. IPAM also enables you to include user logon events for this IP address. This triggers advanced correlation logic, which displays all the user logon reports that are associated with the IP address.
- **By Client ID** – IPAM enables you to query DHCP lease events for a particular client ID (MAC address for IPv4 and DUID for IPv6) over a specified time period. This displays the list of all relevant lease assignments, renew events, and release events for that IP address. You can also identify the host name and the corresponding IP address that is associated with the client ID information. IPAM also enables you to include user logon events in the displayed results for this client ID. This triggers advanced correlation logic, which displays all the user logon reports that are associated with the client ID.

- **By Host Name** – IPAM enables you to query DHCP lease events for a particular host/computer name over a specified time period. This displays the list of all relevant lease assignments, renew events, and release events for that host name. You can also identify the IP address and corresponding client ID information that is associated with this host name through the lease event details. IPAM also enables you to include user logon events in the results for a host name. This triggers advanced correlation logic, which displays all the user logon reports that are associated with the host name.
- **By User Name** – IPAM enables you to query logon events for a particular user name from domain controllers and NPS servers over a specified time period. You can identify the IP address, client ID, and host name information associated with the user name by using lease event details. By using the correlation logic from the user name, IPAM displays the list of all relevant lease assignments, renew events, and release events.

iSCSI High-Availability Block Storage

In a virtualized environment, building a host cluster with Windows Server 2012 can limit downtime due to hardware failures. However, this scenario does not help if you need to take down a virtual machine for maintenance. If downtime of any kind is not an option, you should consider creating a virtual machine guest cluster on top of your host cluster—a feature that has become more relevant in Windows Server 2012. This enables you to service the workload by taking one virtual machine node down for maintenance while the others are still running.

In a host cluster only scenario, the cluster service runs inside the parent partition of the Hyper-V-enabled physical computer. The cluster manages the virtual machines, which reduces downtime when you need to make hardware changes or software updates to the parent partition. The cluster also provides the ability to migrate virtual machines to other servers to load balance.

Guest clustering runs inside a Hyper-V guest computer, and it supports high availability for workloads within the virtual machine. If you add guest clustering to the previous scenario, you can manage cluster-aware applications and services by moving applications between virtual machines within the cluster, without downtime. You can use this process for operating system, service, or application updates. In addition, shared block storage is directly presented to the virtual machine by using the Microsoft iSCSI Software Initiator.

Requirements

Enabling Microsoft iSCSI Software Target to provide block storage takes advantage of your existing Ethernet network. No additional hardware is needed. If high availability is an important criterion, organizations should consider setting up a high availability server. With a high availability server, you will need shared storage for the cluster—either hardware Fibre Channel storage or a serial attached SCSI (SAS) storage array.

If you enable guest clustering, you need to provide block storage. Any servers running Windows Server software with Microsoft iSCSI Software Target can provide block storage.

Technical overview

Microsoft iSCSI Software Target is a feature under the File Server role in Windows Server 2012, and it is in full compliance with the iSCSI protocol that is outlined in RFC 3720. The Microsoft iSCSI Initiator provides a way to enable an iSCSI initiator in virtual machines to connect to target servers.

Because iSCSI is an industrial standard, Microsoft iSCSI Software Target interoperates with non-Microsoft implementations of an iSCSI initiator. In a heterogeneous environment with other operating systems, Microsoft iSCSI Software Target can also provide storage access to systems such as Linux or Mac, which gives you more choices to access storage.

Windows Server 2012 includes a change in Microsoft iSCSI Software Target when it is used in a clustering configuration. This change improves scalability so that more initiators can connect to the target servers. This supports the goal to provide continuous availability and support workloads with less than 60 seconds of down time.

Moving Virtual Machine Storage

In Windows Server 2008 R2, you can move a running instance of a virtual machine using live migration, but you are not able to move a virtual machine's storage while that virtual machine is running.

Hyper-V in Windows Server 2012 introduces support for moving virtual machine storage without downtime by making it possible to move the storage while the virtual machine remains running. You can perform this task by using a new wizard in Hyper-V Manager or by using new Hyper-V cmdlets for Windows PowerShell.

You can add storage to either a stand-alone computer or to a Hyper-V cluster, and then move virtual machines to the new storage while the virtual machines continue to run.

The most common reason for moving a virtual machine's storage is to update the physical storage that is available to Hyper-V. You can also move virtual machine storage between physical storage devices, at run time, to respond to reduced performance that results from bottlenecks in the storage throughput.

Key benefits

Hyper-V in Windows Server 2012 makes it possible to move virtual machine storage while a virtual machine is running.

Requirements

You need to meet the following requirements to take advantage of the Hyper-V functionality that enables you to move virtual machine storage:

- One or more installations of Windows Server 2012 with the Hyper-V role installed
- A server that is capable of running Hyper-V (specifically, it must have processor support for hardware virtualization)
- Virtual machines that are configured to use only virtual hard disks for storage

Note

You cannot move the storage of a virtual machine when any of its storage is directly attached to a physical disk.

Technical overview

This new feature allows you to move the virtual hard disks of a virtual machine while those virtual hard disks remain available for use by the running virtual machine. When you move a running virtual machine's virtual hard disks, Hyper-V performs the following steps, as shown in the figure below:

Figure 31: Moving virtual hard disks

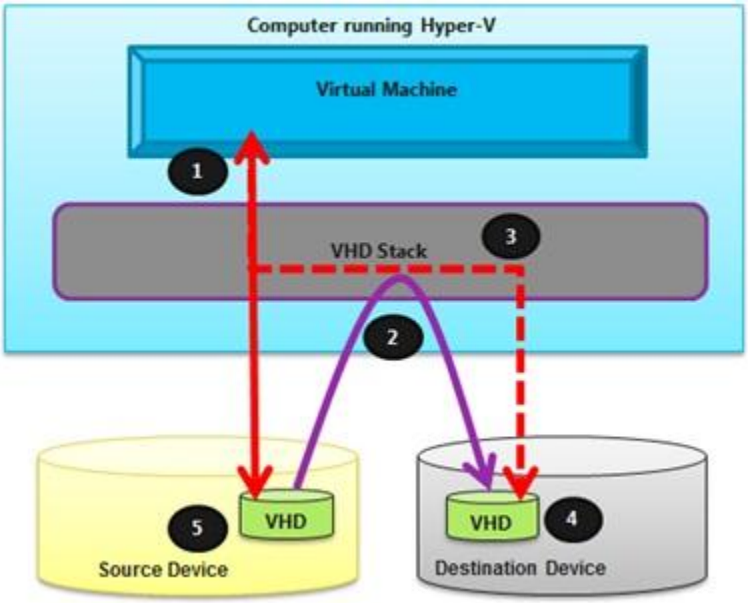

1. Throughout most of the move operation, disk reads and writes go to the source virtual hard disk.
2. While reads and writes occur on the source virtual hard disk, the disk contents are copied to the new destination virtual hard disk.
3. After the initial disk copy is complete, disk writes are mirrored to both the source and destination virtual hard disks while outstanding disk changes are replicated.
4. After the source and destination virtual hard disks are completely synchronized, the virtual machine switches over to using the destination virtual hard disk.
5. The source virtual hard disk is deleted.

Multiple-Terabyte Volumes

Windows Server 2012 and Windows 8 helps customers deploy multi-terabyte NTFS file system volumes, which supports consolidation scenarios and maximizes storage utilization. Historically, customers have not deployed large NTFS volumes. The average deployment sizes were approximately 500 gigabyte (GB) because in the event of file system corruption, the entire volume would be offline for an indeterminate period of time. Chkdsk in Windows Server 2012 and Windows 8 introduces a new approach that prioritizes volume availability and allows for the detection of corruption while the volume remains online with data available.

The new model of Chkdsk has the following benefits:

- Customers can more confidently deploy large volumes. Corruption-related downtime is now proportional to only the number of corruptions on the volume.
- Customers who are using clustered shared volumes see almost no downtime, even for correcting corruption events that would normally require a remount.
- Windows Server 2012 actively monitors the health state of the file system volume, and the health state to the administrator.
- Customers do not see any downtime for transient corruption events.
- Customers experience significantly fewer corruption events.

Requirements

- Chkdsk is enabled by default on Windows Server 2012
- No special hardware is needed

Technical overview

In the event of file system corruption in previous versions of Windows Server, the file system volume would be taken offline, and its data would be unavailable for an indeterminate period of time while Chkdsk ran and attempted to repair the corruption. The time for Chkdsk to complete was proportional to the number of files on the volume, and this time could not be accurately predicted by customers. For a single corruption issue, the entire volume and all of its files would still be analyzed offline. This model was problematic for customers because it potentially imposed significant downtimes, which made it challenging for customers to meet necessary business objectives. In addition, analysis indicated that a significant portion of file system corruptions were due to transient errors—Chkdsk would run, but it would not find any actual corruption.

Windows Server 2012 addresses this problem by prioritizing file system availability, even when corruption is present, and developing a new model for managing file system corruption. The model includes:

- Improved self-healing: NTFS instantaneously self-heals more issues online without requiring Chkdsk to run offline. This reduces the number of times that Chkdsk is required to run.
- Online analysis: In the previous model, almost all of the offline time required to run Chkdsk was spent scanning and analyzing the drive. In Windows Server 2012, the analysis phase of Chkdsk, which was responsible for the majority of the offline time, becomes an online background task. This enables the volume to remain online and available while the system determines whether there is corruption. There is also added logic in the model that verifies the type of corruption is not transient, preventing unnecessary analysis.
- Corruption correction: When the scan is completed, Windows Server 2012 informs the administrator (by using events and the management consoles), that the volumes need to be repaired and suggests a solution, such as performing a remount or reboot. Because the analysis phase has already completed, no additional scanning or detection is required. Chkdsk directly fixes the identified corruption, and the offline time is minimized to seconds. Therefore, the offline time for a volume is no longer proportional to the number of files on the volume, but rather to the number of corruptions on the volume.

Windows Server 2012 always provides the current health state of the file system volume by using standard events that indicate the current state of the volume. The following figure illustrates how significantly the new Chkdsk model changes downtime in the event of corruption.

Figure 32: Chkdsk downtimes

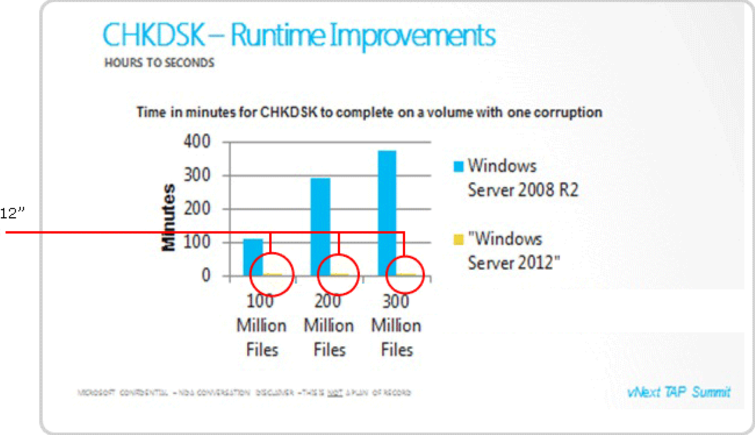

CHKDSK – Runtime Improvements

HOURS TO SECONDS

MICROSOFT CONFIDENTIAL – NDA CONVERSATION DISCLAIMER – THIS IS NOT A PLAN OF RECORD

vNext TAP Summit

In this benchmark, "Windows Server 2012" execution time <8 seconds

Network Adapter Teaming

Network adapter teaming, also known as load balancing and failover (LBFO), enables multiple network adapters on a computer to be placed into a team for the following purposes:

- Bandwidth aggregation
- Traffic failover to prevent connectivity loss in the event of a network component failure

This feature has been a requirement for independent hardware vendors (IHVs) to enter the server network adapter market, but until now network adapter teaming has not been included in Windows Server operating systems.

Requirements

Network adapter teaming requires the presence of a single Ethernet network adapter, which can be used for separating traffic that is using VLANs. All modes that provide fault protection through failover require at least two Ethernet network adapters. Windows Server 2012 supports up to 32 network adapters in a team.

Technical overview

Network adapter teaming technology includes many parts and options of network adapter teaming technology. This section includes the following sections:

- Network adapter teaming architecture
- Bandwidth aggregation (also known as load balancing) mechanisms
- Traffic distribution algorithms
- Network adapter teaming in virtual machines
- Incompatibilities
- Requirements

Network adapter teaming architecture

Today, nearly all network adapter teaming solutions on the market have a similar architecture, as shown in Figure 33.

Figure 33: Teaming solution architecture

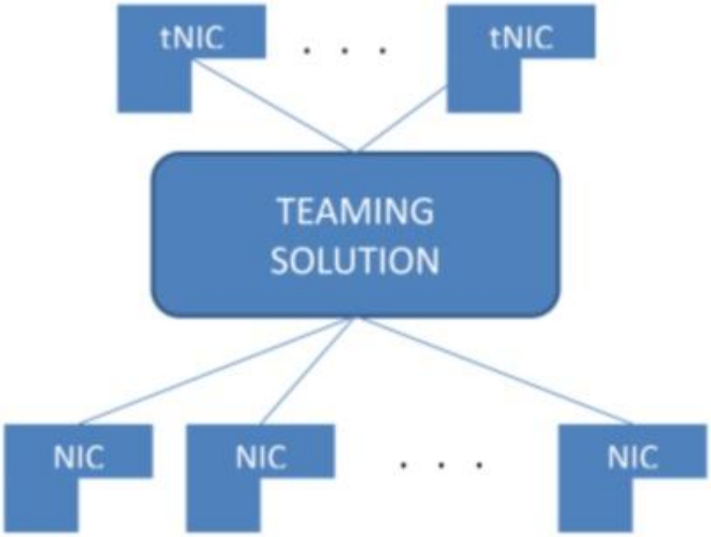

Two or more physical network adapters are connected to the network adapter teaming solution multiplexing unit, which then presents one or more virtual adapters (also known as *team network adapters*) to the operating system. There are several different algorithms that distribute inbound and outbound traffic between the physical network adapters. In current non-Microsoft network adapter teaming solutions, the team network adapters divide traffic by virtual LAN (VLAN) so that applications can connect to different VLANs at the same time. Technically, this separation of traffic is not part of network adapter teaming. However, because other commercial implementations of network adapter teaming have this capability, the Windows Server 2012 implementation also includes it.

Network adapter teaming configurations

There are two basic sets of algorithms that are used for network adapter teaming:

- Algorithms that require the switch to participate in the teaming, also known as *switch-dependent modes*. These algorithms usually require all the network adapters of the team to be connected to the same switch.
- Algorithms that do not require the switch to participate in the teaming, also referred to as *switch-independent modes*. Because the switch does not know that the network adapter is part of a team, the team network adapters can be connected to different switches. Switch-independent modes do not require that the team members connect to different switches; they merely make it possible.

There are two common choices for switch-dependent modes of network adapter teaming:

- **Generic or static teaming (IEEE 802.3ad draft v1).** This mode requires configuration on the switch and the computer to identify which links form the team. Because this is a statically configured solution, no additional protocol assists the switch and the computer to identify incorrectly plugged cables or other errors that could cause the team to fail. This mode is typically supported by server-class switches.
- **Dynamic teaming (IEEE 802.1ax, LACP).** IEEE 802.1ax uses the Link Aggregation Control Protocol (LACP) to dynamically identify links between the computer and a specific switch. This enables the

automatic creation of a team and, in theory, the expansion and reduction of a team simply by the transmission or receipt of LACP from the peer network adapter. Typical server-class switches support IEEE 802.1ax, but most switches require manual administration to enable LACP on the port.

Note

IEEE 802.1ax is also commonly known as IEEE 802.3ad because it was developed by the IEEE 802.3ad committee before being published as IEEE 802.1ax.

Both modes should result in inbound and outbound traffic approaching the practical limits of the aggregated bandwidth because the pool of links in the team functions as a single pipeline.

Traffic distribution algorithms

Outbound traffic can be distributed among the available links in many ways. Try to keep all packets that are associated with a single flow (TCP stream) on a single network adapter. This is needed to minimize out-of-order packet arrival scenarios.

Network adapter teaming in Windows Server 2012 supports the following traffic distribution methods:

- **Hyper-V switch port.** In cases where virtual machines have independent media access control (MAC) addresses, the MAC address of the virtual machine can provide the basis for dividing traffic. There is an advantage in using this scheme in virtualization. Because the adjacent switch can determine that specific source MAC addresses are on only one connected network adapter, the switch will balance the egress load (the traffic from the switch to the computer) on multiple links, based on the destination MAC address for the virtual machine. This is particularly helpful when used with virtual machine queue. However, this mode might not be specific enough to get a well-balanced distribution, and it limits a single virtual machine to the bandwidth that is available on a single network adapter.

  Note

  Windows Server 2012 uses the Hyper-V switch port as the identifier rather than the source MAC address, because in some instances, a virtual machine might be using more than one MAC address on a switch port.

- **Hashing.** This algorithm creates a hash based on components of the packet, and then it assigns packets that have that hash value to one of the available network adapters. This keeps all packets from the same TCP stream on the same network adapter. Hashing alone usually creates balance across the available network adapters. Some network adapter teaming solutions that are available on the market monitor the distribution of the traffic, and they reassign specific hash values to different network adapters in an attempt to better balance the traffic. The dynamic redistribution is known as *smart load balancing* or *adaptive load balancing*.

The components that can be used as inputs to the hashing function include the following:

- Source and destination MAC addresses
- Source and destination IP addresses, with or without considering the MAC addresses (double hash)
- Source and destination TCP ports, usually used along with the IP addresses (quadruple hash)

The quadruple hash creates a finer distribution of traffic streams, which results in smaller streams that can be independently moved between network adapters. However, it cannot be used for traffic that is not TCP or UDP traffic or where the TCP and UDP ports are hidden from the stack, such as traffic that is protected by Internet Protocol security (IPsec). In these cases, the hash falls back to a double hash.

Network adapter teaming in virtual machines

Network adapter teaming in Windows Server 2012 also works within a virtual machine. This allows a virtual machine to have virtual network adapters that are connected to more than one Hyper-V switch and still have connectivity even if the network adapter under that switch gets disconnected. This is particularly important when working with features such as Single Root I/O Virtualization (SR-IOV) because SR-IOV traffic does not go through the Hyper-V switch. Thus, it cannot be protected by a team that is under a Hyper-V switch. With the virtual machine teaming option, an administrator can set up two Hyper-V switches, each connected to its own SR-IOV-capable network adapter. At that point:

- Each virtual machine can then install a virtual function from one or both SR-IOV network adapters. Then, in the event of a network adapter disconnect, the virtual machine can fail over from the primary virtual function to the backup virtual function.
- Alternately, the virtual machine might have a virtual function from one network adapter and a non-virtual function network adapter to the other switch. If the network adapter that is associated with the virtual function gets disconnected, the traffic can fail over to the other switch without loss of connectivity.

Note

Because failover between network adapters in a virtual machine might result in traffic being sent with the MAC address of the other network adapter, each Hyper-V switch port that is associated with a virtual machine that is using network adapter teaming must be set to allow MAC spoofing.

Incompatibilities

Network adapter teaming is compatible with all networking capabilities in Windows Server 2012 with three exceptions: SR-IOV, remote direct memory access (RDMA), and TCP Chimney. For SR-IOV and remote direct memory access (RDMA), data is delivered directly to the network adapter without passing through the networking stack. Therefore, it is not possible for the network adapter team to look at or redirect the data to another path in the team. TCP Chimney is not supported with network adapter teaming in Windows Server 2012.

Performance Offloads for Accelerating Network IO

Enterprise administrators who are responsible for managing and maintaining their low latency applications have the following requirements from the server operating system:

Predictability. Admins require the operating system to respond in a predicable manner to changing IO requests over time

Scalability. Ability to handle increasing number of IOPS (reduce the latency to handle IOPS) over time.

Tier one applications are trying to pump a whole set of packets of data over the network, which needs to be processed by the server. In the past the network throughput was the limitation (that is, how much data can be sent across the network by the server). However with 10G networks becoming more prevalent in the enterprise datacenter, network throughput is no longer the issue and the limitation has moved to IOPS that can be processed by the server. Another issue is that packets are also decreasing in size, so higher packet rates and lower packet sizes cause larger overhead.

Power increase. Servers are increasing in power and so is the amount of processing that can be done on them. Data processing can be done using the CPU; however, in some cases the IO processing also requires CPU resources which can be offloaded to hardware to ease the burden on the CPU and allow the hardware ecosystem to perform some of the work.

The hardware ecosystem has evolved to create hardware network adapters that are capable of ODX without requiring overhead of the network stack or the operating system to be involved. This frees the operating system to perform the processing, while the data transfer is handled directly through hardware to enable the operating system to scale to process a higher number of transactions.

Absolute low latency. Decrease the total end-to-end transaction processing for mission-critical applications. Enterprises require these mission-critical applications to respond and process data with as much performance as possible.

Technical overview

The Windows Server 2012 changes at the system level include the following capabilities.

The administrator can configure the network for desired performance and scalability to process the optimum IOPS required for these mission-critical applications using the following:

Receive Side Scaling (RSS) improvements

- RSS is the optimum use of multiprocessor machines for handling IO interrupts.
- The RSS feature spreads the interrupts over multiple processors, thereby ensuring that a single processor is not burdened with handling all IO interrupts (as was the case in Windows Server 2003 and prior versions).
- In Windows Server 2012 a number of improvements have been made on top of the RSS feature that was first introduced in Windows Server 2008.

    - Active load balancing between the processors
    - The operating system will keep track of the load on the different CPUs and transfer the interrupts accordingly.
    - Admin can control which processors can be used for handling the requests
    - Admin can choose processors which are beyond 64 (K Group 0) for handling the IO requests, thereby taking advantage of some of the very high-end machines which have a large number of logical processors.
    - Interrupts are also NUMA aware, which means the operating system will transfer the interrupts within the same NUMA node to make it more efficient.
    - It works with inbox NIC Teaming or LBFO, making it possible to get failover and scalability, which removes the limitation of prior versions that admins had to choose between LBFO using hardware drivers or RSS.
    - In addition to TCP traffic the RSS in Windows Server 2012 will also work for UDP traffic.
    - In Windows Server 2012, the admin can also manage and debug the feature using built-in tools like WMI and Windows PowerShell.

Receive Segment Coalescing (RSC)

- This feature improves the scalability of the servers by reducing the overhead for processing a large number of network IO traffic by offloading some of the work to RSC-capable network adapters.
    - This feature is applicable in receive intensive workloads where the server is receiving a large number of small data packets that need to be combined before they can be processed. If all of these packets need to be handled by the server then the overhead will reduce performance and scale of the server.
    - RSC capable NICs can collect all these small/tiny packets and combine them into a single packet before passing them over to the CPU for processing. So the amount of overhead required by the CPU to process this data is significantly reduced.
- Overall expectation is that using RSC will reduce CPU utilization by 10 percent to 30 percent.

Note

This feature can also be managed using Windows PowerShell or WMI.

- From a competitive perspective this technology is generally referred to as Large Receive Offload (LRO) or Generic Receive Offload (GRO) on other operating systems and by external partners.

Requirements

Performance offloads using RSS are available on most server class 1-GbE and 10-GbE network adapters. Offload support for RSC is available only on 10-GbE network adapters. Currently, we support up to 32 network adapters.

Summary

By enabling low latency and mission-critical applications, Windows Server 2012 helps address the needs of some of the most demanding enterprise customers while maintaining the lowest infrastructure operating cost. Windows Server 2012 provides IT pros with the necessary higher scalability and improved performance to support low latency applications in both physical and virtual environments, while simplifying their management. Developers can now use Windows Server 2012 APIs to tune their applications to provide desired latency and scale for their most demanding applications.

For enterprises, it provides:

- Increased performance and scalability to meet the needs of the most demanding applications
- Added extensibility for developers and third-party ISVs to build mission-critical applications that can perform in low latency applications

Quality of Service

Windows Server 2012 includes new Quality of Service (QoS) bandwidth management features that enable cloud hosting providers and enterprises to provide services that deliver predictable network performance to virtual machines on a server running Hyper-V. In Windows Server 2012, QoS supports the management of upper-allowed and lower-allowed bandwidth limits, referred to in this document as maximum bandwidth and minimum bandwidth. Windows Server 2012 also takes advantage of hardware that is enabled for Data Center Bridging (DCB) to converge multiple types of network traffic on a single network adapter with a greater level of service to each type. With Windows PowerShell, you can configure all these new features manually or enable automation in a script to manage a group of servers, regardless of whether they stand alone or are joined to a domain.

For example, cloud hosting providers want to use servers running Hyper-V to host customers and still enable a specific level of performance based on service level agreements (SLAs). They want to ensure that no customer is impacted or compromised by other customers on their shared infrastructure, which includes computing, storage, and network resources. Likewise, enterprises have similar requirements. They want to run multiple application servers on a server running Hyper-V and be confident that each application server delivers predictable performance. Lack of performance predictability often drives administrators to put fewer virtual machines on a capable server or simply avoid virtualization, causing them to spend more money on physical equipment and infrastructure.

Furthermore, most cloud hosting providers and enterprises today use a dedicated network adapter and a dedicated subnet for a specific type of workload such as storage or live migration to enable network performance isolation on a server running Hyper-V. Although this deployment strategy works for those using 1-gigabit Ethernet network adapters, it becomes impractical for those who are using or plan to use 10-gigabit Ethernet network adapters. Not only does one 10-gigabit Ethernet network adapter (or two for high availability) already provide sufficient bandwidth for all the workloads on a server running Hyper-V in most deployments, but 10-gigabit Ethernet network adapters and switches are considerably more expensive than their 1-gigabit Ethernet counterparts. To better utilize 10-gigabit Ethernet hardware, a server running Hyper-V requires new capabilities to manage bandwidth.

Requirements

Every version of Windows Server 2012 includes the new QoS functionality. The minimum bandwidth that is enforced by the packet scheduler can nearly always be enabled, but it works better on 1-gigabit Ethernet network adapters or 10-gigabit Ethernet network adapters. We do not recommend that you enable QoS in Windows Server 2012 when it is running as a virtual machine within a virtualized environment. QoS is designed for traffic management on physical networks, rather than virtual networks.

For hardware-based minimum bandwidth, you must use a network adapter that supports DCB, and the miniport driver of the network adapter must implement the Network Driver Interface Specification (NDIS) of the QoS application programming interfaces (APIs). With these requirements met, you can use the new Windows PowerShell cmdlets to configure the network adapter to provide bandwidth guarantees to multiple types of network traffic.

DCB is a suite of technologies that help converge multiple subnets in a datacenter (such as your data and storage networks) onto a single subnet. DCB consists of the following:

- 802.1Qaz Enhanced Transmission Selection (ETS) to support the allocation of bandwidth among various types of traffic.
- 802.1Qbb Priority-based Flow Control (PFC) to enable flow control for a specific type of traffic.
- 802.1Qau Congestion Notification to support congestion management of long-lived data flows within a datacenter.

If a network adapter supports iSCSI offload or remote direct memory access (RDMA) over Converged Ethernet (RoCE), and it is used in a datacenter, the network adapter must also support ETS to provide bandwidth allocation to the offload traffic. In addition, RoCE requires a lossless transport. Because Ethernet does not guarantee packet delivery, the network adapter and the corresponding switch must support PFC. For these reasons, a network adapter must support ETS and PFC to pass the NDIS QoS logo test for Windows Server 2012 certification. 802.1Qau Congestion Notification is not required to obtain the logo. Furthermore, the ETS specifications from the IEEE include a software protocol called Data Center Bridging Exchange (DCBX) to allow a network adapter and a switch to exchange DCB configurations. DCBX is also not required to obtain the logo.

Technical overview

In Windows Server 2008 R2, QoS supports the enforcement of maximum bandwidth. Consider a typical server running Hyper-V in which there are four types of network traffic that share a single 10-gigabit Ethernet network adapter:

- Traffic between virtual machines and resources on other servers (virtual machine data)
- Traffic to and from storage (storage)
- Traffic for live migration of virtual machines between servers running Hyper-V (live migration)
- Traffic to and from a cluster shared volume (CSV) (cluster heartbeat)

If virtual machine data is rate-limited to 3 Gbps, this means the sum of the virtual machine data throughputs cannot exceed 3 Gbps at any time, even if the other network traffic types do not use the remaining 7 Gbps of bandwidth. However, this also means that the other types of traffic can reduce the actual amount of bandwidth that is available for virtual machine data to unacceptable levels, depending on whether or how their maximum bandwidths are defined.

In Windows Server 2012, QoS introduces a new bandwidth management feature: *Minimum bandwidth*. In contrast to maximum bandwidth, minimum bandwidth guarantees a specific amount of bandwidth to a specific type of traffic. Figure 34 provides an example of how minimum bandwidth works for each of the four types of network traffic flow in three different time periods: T1, T2, and T3.

Figure 34: How minimum bandwidth works

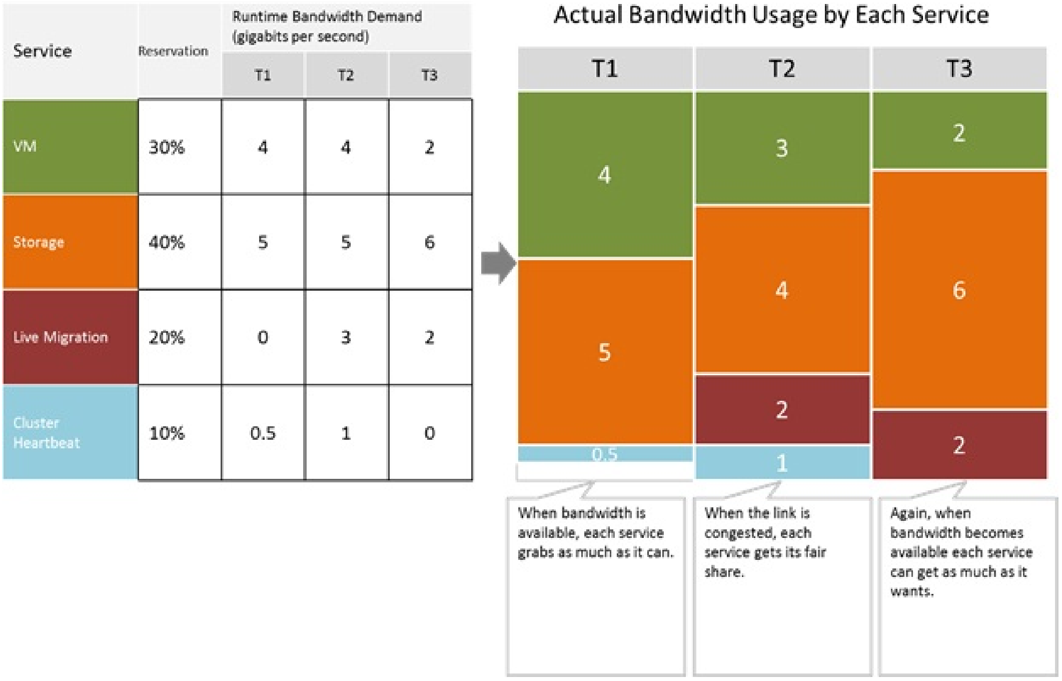

Service	Reservation	Runtime Bandwidth Demand (gigabits per second)		
		T1	T2	T3
VM	30%	4	4	2
Storage	40%	5	5	6
Live Migration	20%	0	3	2
Cluster Heartbeat	10%	0.5	1	0

When bandwidth is available, each service grabs as much as it can.

When the link is congested, each service gets its fair share.

Again, when bandwidth becomes available each service can get as much as it wants.

The left table shows the minimum amount of bandwidth that is reserved for a specific type of network traffic and the estimated amount of bandwidth that it needs in the three time periods. For example, storage is configured to have at least 40 percent of the bandwidth (4 Gbps of a 10-gigabit Ethernet network adapter) at any time. In T1 and T2, it has 5 Gbps worth of data to transmit; and in T3, it has 6 Gbps worth of data. The right table shows the actual amount of bandwidth each type of network traffic gets in T1, T2, and T3. In this example, storage is sent at 5 Gbps, 4 Gbps, and 6 Gbps, respectively, in the three periods.

The characteristics of minimum bandwidth can be summarized as follows:

- In the event of congestion, when the demand for network bandwidth exceeds the available bandwidth (such as in T2 period in the example), minimum bandwidth ensures each type of network traffic gets up to its assigned bandwidth. For this reason, minimum bandwidth is also known as *fair sharing*. This characteristic is essential to converge multiple types of network traffic on a single network adapter.
- If there is no congestion (that is, when there is sufficient bandwidth to accommodate all network traffic, such as in the T1 and T3 periods), each type of network traffic can exceed its quota and consume as much bandwidth as is available. This characteristic makes minimum bandwidth superior to maximum bandwidth in utilizing available bandwidth.

Windows Server 2012 offers two mechanisms to enforce minimum bandwidth:

- Through the newly enhanced packet scheduler
- Through network adapters that support Data Center Bridging (DCB)

In both cases, network traffic needs to be classified first. Windows Server 2012 classifies a packet or gives instructions to a network adapter to classify it. The results of classification are that a number of traffic flows are being managed and a specific packet can belong to only one of them.

For example, a traffic flow can be a live migration connection, a file transfer between a server and a client computer, or a remote desktop connection. Based on how the bandwidth policies are configured, the packet scheduler in Windows Server 2012 or the network adapter sends the packets that are included in a specific traffic flow at a rate equal to or higher than the minimum bandwidth that is configured for the traffic flow.

The two mechanisms have advantages and disadvantages:

- The packet scheduler in Windows Server 2012 provides a fine level of detail for classification. It is a better choice if you have many traffic flows that require minimum bandwidth enforcement. A typical example is a server running Hyper-V that is hosting many virtual machines, where each virtual machine is classified as a traffic flow.
- DCB support on the network adapter supports fewer traffic flows. However, it can classify network traffic that does not originate from the networking stack. A typical scenario involves a special network adapter called a converged network adapter (CNA) that supports iSCSI offload, in which iSCSI traffic bypasses the networking stack and is framed and transmitted directly by the CNA. Because the packet scheduler in the networking stack does not process this offloaded traffic, DCB is the only way to enforce minimum bandwidth.

Both mechanisms can be employed on the same server. However, do not enable both mechanisms at the same time for a specific type of network traffic. Enabling both mechanisms at the same time for the same types of network traffic reduces performance.

In Windows Server 2012, QoS policies and settings are managed by using Windows PowerShell. The new Windows PowerShell cmdlets for QoS support the QoS functionalities that are available in Windows Server 2008 R2 (such as maximum bandwidth and priority tagging) and the new features such as minimum bandwidth. Although you can only manually configure QoS policies by using Group Policy snap-ins (such as gpedit.msc or gpmc.msc) in Windows Server 2008 R2, you can program or enable the automation of QoS policies by using Windows PowerShell in Windows Server 2012. Windows Server 2012 facilitates static and dynamic configuration and enables you to manage virtualized servers that are connected to a converged network in your datacenter. Because Windows PowerShell has a remote computer management capability, you can manage QoS policies for a group of servers at one time, even if these servers are not joined to a domain.

Rapid and Efficient Data Movement Using Intelligent Storage Arrays

Windows Server 2012 maximizes an enterprise's investment in intelligent storage arrays by making it easier for data to move quickly within and between storage devices. Offloaded Data Transfer (ODX) enables administrators to rapidly provision and migrate virtual machines. ODX also provides significantly faster transfers on large files such as database files or video files.

By offloading the file transfer to the storage array, ODX minimizes latencies, maximally uses array throughput, and reduces host server resource usage, such as CPU and network consumption. Windows Server 2012 offloads file transfers transparently and automatically when you move or copy files, irrespective of whether you drag-and-drop files through Windows Explorer or use command-line file copy commands. No administrator setup or intervention is needed. ODX reduces the load on datacenter servers and the network, and it allows enterprises to get more value from their existing hardware infrastructure.

Key benefit

You can get more from your existing hardware infrastructure with the ability to rapidly move large files and virtual machines directly between storage arrays while eliminating CPU and network resource consumption on the host server.

Requirements

A storage array that supports ODX is the key requirement for moving data rapidly and efficiently.

Technical overview

In traditional host-based file transfers, the data to be transferred must be:

1. Read from the storage through the source server
2. Transferred across the network to the destination server
3. Written back to the storage through the destination server

To eliminate this inefficiency, ODX uses a token-based mechanism for reading and writing data within or between intelligent storage arrays. Instead of routing the data through the host, a small token is copied between the source server and destination server. The token serves as a point-in-time representation of the data. As an example, when you copy a file or migrate a virtual machine between storage locations (within or between storage arrays), a token representing the virtual machine file is copied, thereby removing the need to copy the underlying data through the servers.

The following figure explains the steps that are involved with a token-based copy operation.

Figure 35: Token-based copy operation

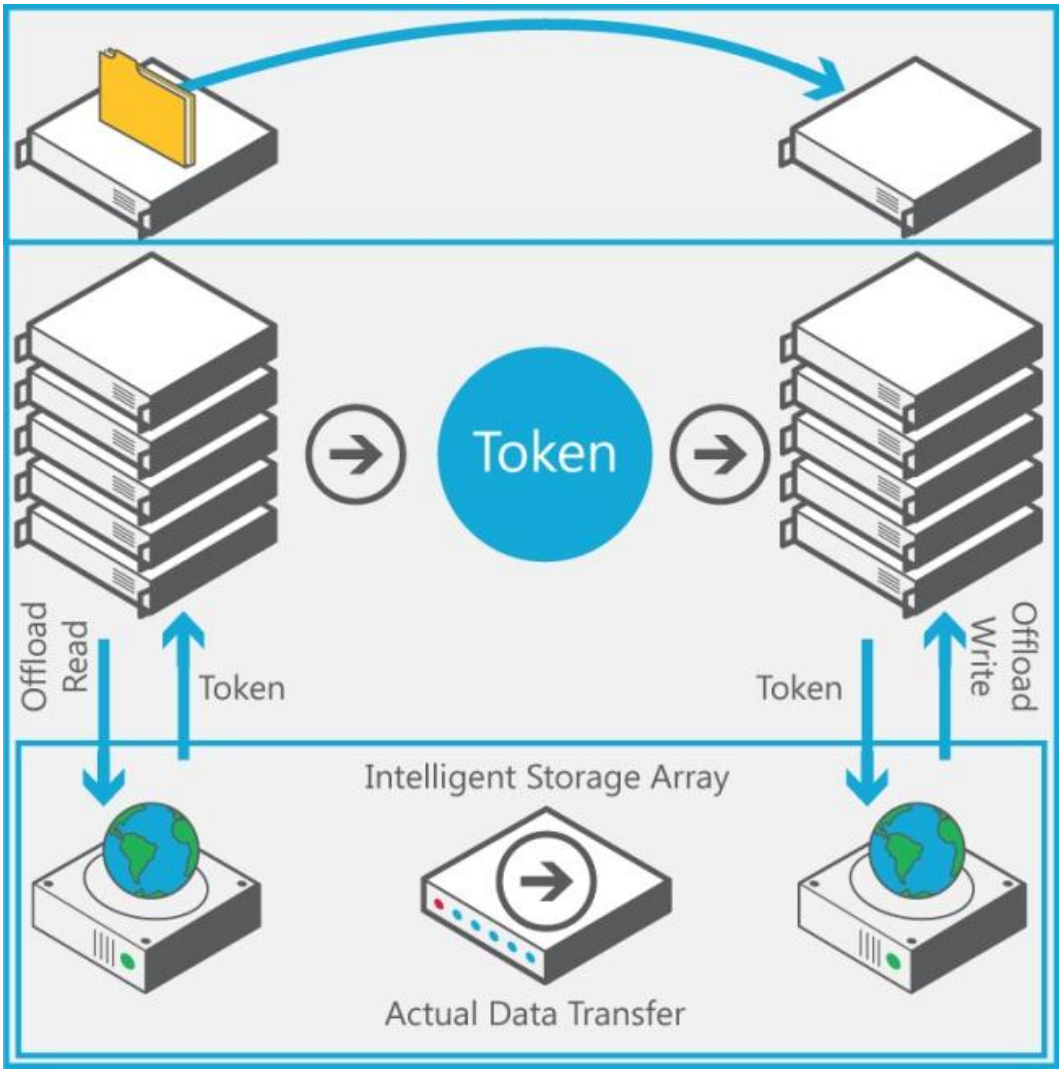

This procedure is described in the following steps:

1. A user copies or moves a file by using Windows Explorer, a command line interface, or as part of a virtual machine migration.
2. Windows Server 2012 automatically translates this transfer request into an ODX (if supported by the storage array), and it receives a token that represents the data.
3. The token is copied between the source server and destination server.
4. The token is delivered to the storage array.
5. The storage array internally performs the copy and provides status information to the user.

Remote Access

As increasing numbers of employees are required to remain productive while they are away from the office, the need for solutions that provide secure remote access to corporate networks has grown.

Windows Server 2012 provides an integrated remote access solution that is simple to deploy. Employees can access corporate network resources while working remotely, and IT administrators can manage corporate computers that are located outside the internal network.

To provide this functionality, remote access in Windows Server 2012 integrates DirectAccess and Routing and Remote Access Services (RRAS) VPN.

- DirectAccess was introduced in Windows Server 2008 R2. It allows managed computers located outside the corporate network to securely access internal resources without VPN connectivity. It establishes transparent connectivity to the corporate network every time a DirectAccess client computer connects to the Internet, even before the user logs on. In addition, DirectAccess allows administrators to easily monitor connections and remotely manage DirectAccess client computers located on the Internet. Computers running Windows Server 2012, Windows Server 2008 R2, Windows 8, and Windows 7 can be configured as DirectAccess client computers.
- RRAS provides remote access VPN connectivity among remote clients and servers, site-to-site connections between servers, and routing. A RRAS VPN provides a remote access solution for client computers that are unmanaged or running operating systems earlier than Windows 7.

In Windows Server 2012, DirectAccess and RRAS are integrated into a single Remote Access server role. The role is divided into two components: DirectAccess and VPN and Routing. DirectAccess and VPN can be configured together in the Remote Access Management console by using a single set of wizards. Other RRAS features can be configured by using the legacy Routing and Remote Management console. The new role allows for easier migration of RRAS and DirectAccess deployments from Windows 7, and it provides a number of new features and improvements.

Deployment requirements

DirectAccess deployment requirements include the following:

- **Server.** One or more servers running Windows Server 2012 with the Remote Access role installed. The server can be deployed at the edge or behind an edge firewall or other device.
- **Domain.** The server must be joined to an Active Directory Domain Services (AD DS) domain.
- **Network adapters.** The server must have at least one network adapter installed and enabled. If deployed with a single adapter, IP-HTTPS will be used for client connections.

  Note

  To use Teredo, two network adapters with two public consecutive IPv4 addresses on the external adapter are required.
- **Permissions.** The remote access administrator requires local administrator permissions on the server, domain user administrator permissions, and permissions to create a WMI filter (Domain Admins) on the domain controller. The WMI filter is required if the client Group Policy Object should be applied to only mobile computers in the domain.

- **Security groups.** An Active Directory security group that contains the computers you want to enable as DirectAccess clients.
- **DNS.** DNS server running Windows Server 2012, Windows Server 2008 with SP2, or Windows Server 2008 R2.
- **Client computer support.** DirectAccess client computer users must be members of an AD DS domain. DirectAccess client computers must be running Windows 8, Windows 7 Enterprise, Windows 7 Ultimate, Windows Server 2012, or Windows Server 2008 R2.
- **Certificate requirements.** A public key infrastructure (PKI) if the DirectAccess deployment requires NAP, two-factor authentication, or support for clients running Windows 7.

Technical overview

This section summarizes the benefits of remote access in Windows Server 2012, and new remote access features.

Improved management experience

- **Easy administration.** DirectAccess and VPN can be configured, managed, and monitored in a single location by using the new Remote Access Management console. Multiple remote access servers can be managed from the console.
- **Improved monitoring.** The Remote Access Management console in Windows Server 2012 provides detailed monitoring information as follows:
  - **Dashboard.** The dashboard provides top-level information about Remote Access servers and client computer activity. Reports can be generated quickly from the dashboard.
  - **Operations status.** Administrators can investigate the status of specific server components.
  - **User and client computer monitoring.** Administrators can view users and computers that are connected over VPN or DirectAccess at any time, and they can check the resources that clients are accessing.
  - **Accounting.** Data can be logged to a local Windows Internal Database or to a remote RADIUS server. The accounting log stores remote user information, operations statistics, server usage, and change history. Server usage logs provide server load statistics for the Remote Access server.
  - **Troubleshooting.** Detailed events and tracing are provided to help diagnose connectivity issues.
- **Network Connectivity Assistant (NCA).** NCA runs on DirectAccess client computers to provide a quick view of the DirectAccess connection status, links to corporate help resources, diagnostics tools, and troubleshooting information.
- **Windows PowerShell support.** Administrators can use Windows PowerShell command-line tools and automated scripts for Remote Access setup, configuration, management, monitoring, and troubleshooting.

Ease of deployment

- **Deployment modes.** In Windows Server 2008 R2, configuring DirectAccess for remote client management required manual modification of Windows Firewall rules. In Windows Server 2012, DirectAccess can be easily configured for remote client access and remote client management, or for only remote client management.
- **Simplified deployment.** DirectAccess in Windows Server 2012 provides a simpler configuration experience. Small and medium businesses can set up a working deployment with minimum requirements in only a few steps.
- **No certificate infrastructure.** For simple deployments, DirectAccess can be configured without requiring deployment of a certificate infrastructure.
- **Access to IPv4 servers.** DirectAccess in Windows Server 2012 supports access to internal servers that are running IPv4 only.
- **Simplified IPsec deployment.** Traditionally, DirectAccess requires the deployment of two IPsec tunnels. The first tunnel provides a connection to infrastructure servers that are required to authenticate and manage client computers. The second tunnel provides access to corporate resources after users log on. In a Windows Server 2012 deployment, DirectAccess can be deployed with a single IPsec tunnel.

New and improved deployment scenarios

- **Single network adapter support.** In Windows Server 2012, DirectAccess can be deployed on servers that are configured with a single network adapter running behind a firewall or network address translation (NAT) device.
- **Force tunneling.** By default, DirectAccess clients locate Internet access internal resources through DirectAccess and they locate Internet resources by using their local adapter settings. In Windows Server 2008 R2, forcing DirectAccess clients to connect to Internet resources through the DirectAccess server required manual manipulation of Group Policies. In Windows Server 2012, you can enable force tunneling directly in the Remote Access Management console.
- **NAP compliance.** In Windows Server 2008 R2, configuring Network Access Protection (NAP) to verify client compliance with corporate policies required manual editing of the Windows Firewall rules. In Windows Server 2012, you can enable NAP directly in the Remote Management console.
- **Multiple domain support.** In Windows Server 2008 R2, the DirectAccess server, clients, and internal servers had to belong to the same domain. This setting could only be modified by manually editing DirectAccess Group Policies. In Windows Server 2012, multiple domain support is integrated, and no manual editing is required.
- **Geographical location support.** In Windows Server 2012, Remote Access servers can be configured in a multiple site deployment that helps users in dispersed geographical locations to connect to the multiple site entry point that is closest to them. Traffic across a multiple site deployment can be distributed and balanced with an external global load balancer.
- **One-time password (OTP) client authentication.** In Windows Server 2008 R2, DirectAccess provided standard client IPsec authentication and two-factor authentication by using smart cards. Windows Server 2012 adds support for two-factor authentication by using a one-time password (OTP), which provides the ability to use OTP solutions that are provided by non-Microsoft vendors.

- **Virtual smart card support.** In addition to support for standard smart card authentication, DirectAccess can use the Trusted Platform Module (TPM)-based virtual smart card capabilities that are available in Windows Server 2012. The TPM of client computers can act as a virtual smart card for two-factor authentication, which removes the overhead and costs that are incurred in smart card deployment.
- **Behind edge device.** Remote Access servers can be placed behind an edge device such as a firewall or NAT router. This removes the requirement to have dedicated public IPv4 addresses for DirectAccess.
- **Off-premises client configuration.** In Windows Server 2008 R2, client computers must be connected to the corporate network to join a domain or receive domain settings. Windows Server 2012 provides the capability for computers to join a domain and receive domain settings remotely from the Internet.
- **Client computer support.** DirectAccess supports client computers running Windows Server 2012, Windows Server 2008 R2, Windows 8, and Windows 7.
- **Server core support.** Windows Server 2012 provides a minimal operating system installation known as a Server Core installation option. The Remote Access role can be installed and configured on a Server Core installation.

Scalability improvements

- **High availability and failover.** Remote Access in Windows Server 2012 provides support for more users, higher performance, and lower costs. Remote Access servers can be gathered into a load-balanced cluster for high availability and failover. Cluster traffic can be load balanced by using Windows Network Load Balancing (NLB) or a hardware load balancer.
- **Improved performance in virtualized environments.** With the shift toward virtualized datacenters and the reduced costs that are provided by virtualization, the Remote Access server role takes advantage of single root I/O virtualization (SR-IOV) for improved I/O performance when it is run on a virtual machine. In addition, Remote Access improves the overall scalability of the server host with support for IPsec hardware offload capabilities (available on many server interface cards that perform packet encryption and decryption in hardware).
- **IP-HTTPS NULL encryption.** IP-HTTPS provides DirectAccess client connectivity to internal IPv4 resources when other IPv4 transition technologies such as Teredo cannot be used. In Windows 2008 R2, IP-HTTPS performance is poor compared with other transition technologies because data that is already encrypted for DirectAccess by using IPsec is encrypted again using SSL. This incurs overhead. In Windows Server 2012, IP-HTTPS is implemented by using NULL encryption, which removes redundant SSL encryption during client communications and improves performance.
- **IP-HTTPS behind a proxy server.** IP-HTTPS runs in a system context rather than a user context. This context can cause connection issues. For example, if a DirectAccess client computer is located on the network of a partner company that uses a proxy for Internet access, and Web Proxy Autodiscovery Protocol (WPAD) detection is not used, the user must manually configure proxy settings to access the Internet. These settings are configured in Internet Explorer on a per user basis, and they cannot be retrieved in an intuitive way on behalf of IP-HTTPS. In addition, if the proxy requires authentication, the client provides credentials for Internet access, but IP-HTTPS will not provide the credentials that are required to authenticate to DirectAccess. In Windows Server 2012, a new feature solves these issues. Specifically, the user can configure IP-HTTPS to work behind a proxy that is not configured using WPAD, and IP-HTTPS will request and provide the proxy credentials that are needed for IP-HTTPS to request authentication.

Table 10: Remote access features

Benefit	Feature	Description
Improved management experience	Unified remote access Management console	DirectAccess and RRAS integrated into the remote access role Deployment of RRAS and DirectAccess on a single server Management of multiple servers in a single console Easy migration of RRAS and DirectAccess from Windows Server 2008 R2 to Windows Server 2012
	Detailed monitoring, logging, and reporting	Detailed monitoring of servers, clients, and user connections Accounting in multiple formats Detailed event logging Tracing and packet captures On-demand reporting
	Windows PowerShell scripting	Windows PowerShell scripting to configure, manage, and monitor Remote Access servers
	Network Connectivity Assistant (NCA) application	Integration with Windows Network Connection manager DirectAccess connectivity status Remediation for common failures Log collection for troubleshooting OTP connection options if OTP is enabled
Ease-of-deployment	DirectAccess deployment modes	Easier configuration of DirectAccess for client access and remote management or for only remote management
	Simplified DirectAccess	Getting Started Wizard with

Benefit	Feature	Description
	deployment	minimum requirements
	Deployment without a certificate infrastructure	DirectAccess client IPsec authentication with Active Directory credentials only (no computer certificate is required) Option to use a self-signed certificate that is created automatically by DirectAccess for authentication of the network location server and for IP-HTTPS Features that are not available without a certificate infrastructure include: Client compliance checking with NAP Support for client computers running Windows 7 Two-factor authentication
	Access to internal IPv4 support by using NAT64/DNS64	Support for client access to internal servers not running IPv6 DirectAccess deployment without upgrading IPv4 corporate servers
	Simplified IPsec deployment with single tunnel	DirectAccess clients access all resources through a single tunnel No requirement to manage a quarantine network of infrastructure servers that are only available over a single tunnel
New and improved deployment scenarios	Single network adapter support	Can deploy a server with a single adapter that is located behind an edge or NAT device Clients connect by using IP-HTTPS

Benefit	Feature	Description
	Force tunneling support	Easy configuration of force tunneling during DirectAccess configuration
	NAP support	Easy configuration of NAP during DirectAccess configuration
	Multiple domain support for DirectAccess	Ability to locate DirectAccess servers and clients in different domains
	Multiple geographical locations	Automatically connect clients to the DirectAccess server entry point closest to them Computers running Windows 8 can manually specify an entry point, overriding the automatic entry point that is assigned Support for failover from one DirectAccess entry point to another
	OTP client authentication	Support for two-factor authentication using OTP
	Virtual smart card support	Can leverage TPM on DirectAccess client computers to provide two-factor smart card authentication Can eliminate overhead that is associated with smart card deployment
	NAT support	Can deploy Remote Access servers behind an edge firewall or NAT device No requirement for the server to have an adapter connected directly to the Internet
	Off-premises client support	Client computers join a domain and retrieve domain settings

Benefit	Feature	Description
		through the Internet
	DirectAccess client support	Can install DirectAccess on computers running Windows Server 2012, Windows Server 2008 R2, Windows 8, and Windows 7 Ultimate or Windows 7 Enterprise Limitations for client computers running Windows 7: Cannot run the Network Connectivity Assistant (however, the DirectAccess Connectivity Assistant that was introduced in Windows Server 2008 R2 is supported) Must authenticate by using a computer certificate Cannot be automatically routed to a multiple site entry point in a multiple site deployment (each entry point must be statically configured to support client computers running Windows 7) Cannot select the multiple site entry point to which they connect
	Server Core installation support	Support for the Remote Access role on computers running a Server Core installation
Scalability	High availability	Can deploy multiple Remote Access servers in a cluster Can load balance the cluster by using Windows NLB or a hardware load balancer Windows NLB supports up to eight cluster members Hardware load balancing supports up to 64 cluster

Benefit	Feature	Description
		members Can add and remove servers from the cluster without interrupting connections that are in progress Support for all RRAS VPN protocols on server cluster deployments
	Virtualization improvements	SR-IOV virtualization for improved performance Support for IPsec Task Offload v2
	IP-HTTPS	Support for clients behind a proxy server that requires manual configuration of proxy settings Faster performance with NULL encryption

Remote Desktop Services

Remote Desktop Services enables the mobile work force to connect to desktops and applications from virtually anywhere they want to provide the full Windows experience.

Microsoft has made significant investments in the user and management experiences in Remote Desktop Services in Windows Server 2012. For the user experience, our goal was to ensure that that you have a rich experience, regardless of the type of device that is used to connect, the type of resources they are connecting to (for example, virtual desktops, RemoteApp programs, or session-based desktops), or whether they are connecting through a LAN or WAN. For the management experience, there is now a centralized console so you can manage the Remote Desktop Services role services and their associated resources from a single location.

In Windows Server 2012, Remote Desktop Services has enhanced support for the following scenarios:

- Unified, simplified, low-cost Virtual Desktop Infrastructure (VDI) deployments
- Unified, simplified, low-cost Session Virtualization deployments
- Centralized resource publishing
- Rich user experience with RemoteFX

The overall architecture for Remote Desktop Services in Windows Server 2012 is shown in Figure 36.

Figure 36: Remote Desktop Services architecture

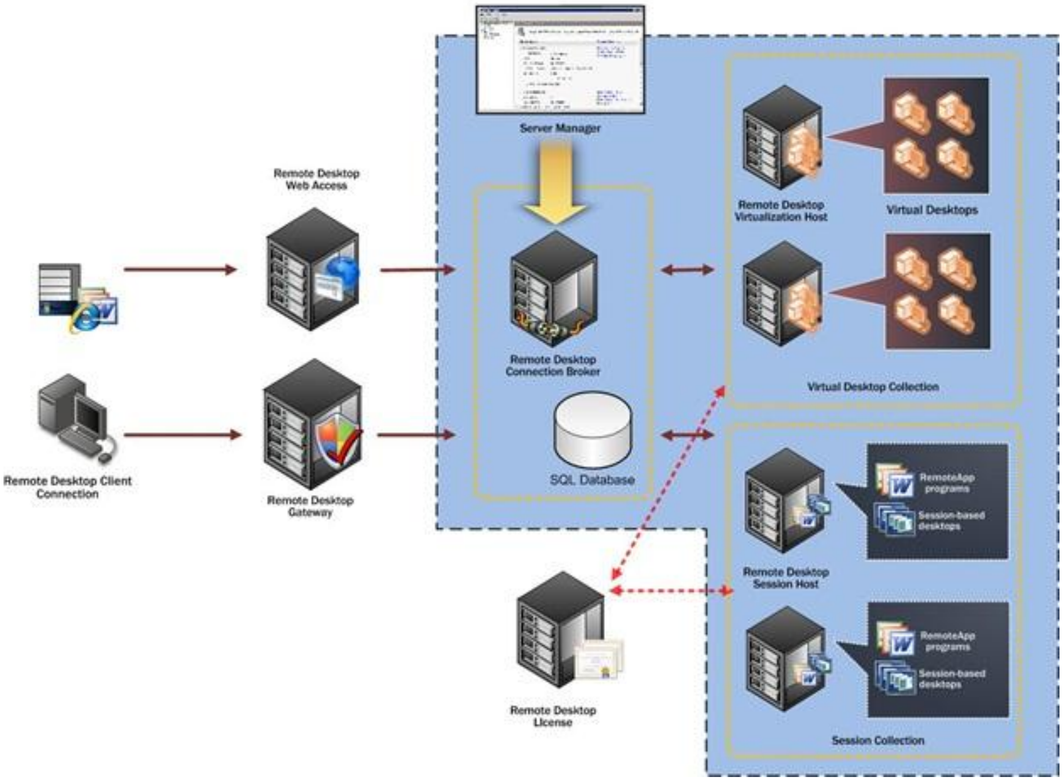

Unified, simplified, low-cost VDI deployment

Remote Desktop Services introduced a VDI deployment in Windows Server 2008 R2. In Windows Server 2012, VDI in Remote Desktop Services provides the following new ways to configure and manage your virtual desktops to reduce the total cost-of-ownership in traditional desktop deployments:

- **Unified central experience.** You can deploy VDI in less than few hours using commodity hardware and manage your pooled and personal virtual desktop deployments from a single pane of glass through a new unified central experience.
- **Simple single image management.** VDI enables automation for deploying and managing pooled virtual desktops with a virtual desktop template.
- **User personalization.** User profile disks help you to preserve user personalization settings for pooled virtual desktop deployments.
- **Less expensive storage.** You can use less expensive local storage by using the live migrate functionality among host computers for pooled virtual desktops, and personal virtual desktops can use less expensive central server message block (SMB) storage.

Virtual desktops

There are three types of virtual desktops in Windows Server 2012: pooled, personal, and session desktops. Pooled virtual desktops help you log on to any virtual desktop in the virtual desktop pool and get the same experience. With a pooled virtual desktop, user personalization is stored in a user profile disk. Personal virtual desktops are permanently assigned to a user account and the user logs on to the same virtual desktop each time to get a fully personalized experience. Session desktops take advantage of the session virtualization capabilities in Windows Server 2012.

The following table shows the available options for managing your pooled and virtual desktops.

Table 11: Options for managing pooled and virtual desktops

Virtual desktop creation	Virtual desktop updates	Pooled	Personal
Automatically create virtual desktops from a virtual desktop template	Virtual desktops are automatically updated	Yes (Recommended deployment scenario for pooled virtual desktops)	No
	Virtual desktops are individually updated	Yes (Supported but not recommended)	Yes (Recommended deployment scenario for personal virtual desktops)
Manually create or import virtual desktops	Virtual desktops are automatically updated	No	No

Virtual desktop creation	Virtual desktop updates	Pooled	Personal
	Virtual desktops are individually updated	Yes (Supported but not recommended)	Yes (Supported but not recommended)

Note
To increase performance, we recommend that you store your virtual desktop template for a pooled virtual desktop collection on a solid state drive (SSD).

VDI Quick Start

In Windows Server 2012, VDI Quick Start installs everything on one computer that you need test the Virtual Desktop Infrastructure deployment scenario.

VDI Quick Start does the following:

- Installs the RD Connection Broker, RD Virtualization Host, and RD Web Access role services on a single computer.
- Creates a pooled virtual desktop collection with two pooled virtual desktops that are based on a virtual hard disk that is the virtual desktop template.
- Creates a Hyper-V network switch named RDS Virtual.

Note
VDI Quick Start is not recommended for a production environment.

VDI Standard Deployment

Virtual Desktop Infrastructure Standard Deployment is new in Windows Server 2012, and it enables you to install the appropriate role services on separate computers. Unlike the VDI Quick Start, a standard deployment gives you more granular control of virtual desktops and virtual desktop collections by not creating them automatically.

High availability for your VDI deployment

Remote Desktop Services in Windows Server 2012 provides support for highly available virtual desktop collections by leveraging a clustered environment.

RD Connection Broker plays a central role with Remote Desktop Services scenarios by providing users access to the virtual desktops and RemoteApp programs in a virtual desktop collection.

With Windows Server 2012, you can easily deploy clustered instances of RD Connection Broker for high availability and improved scalability. With a centralized database, the RD Connection Broker service now only needs to be configured once. Any additional RD Connection Broker servers that are added to the deployment will be added as an active instance.

If one of the RD Connection Broker servers in the cluster fails, users are not impacted because they are redirected to another RD Connection Broker server automatically.

Unified, simplified, low-cost Session Virtualization deployment

Many customers are deploying RD Session Host servers to reduce the total cost of ownership involved with traditional desktop deployments. In previous versions of Windows Server, the deployment of an RD Session Host server farm could be a time consuming task because you had to preinstall the RD Session Host role service on each server. Additionally in Windows Server 2008 R2, ongoing management of the RD Session Host farm was done at a per-server level. In Windows Server 2012, a session virtualization deployment scenario was created to enable centralized installation and management from a "single pane of glass."

A session virtualization deployment consists of RD Session Host servers and infrastructure servers (such as RD Licensing, RD Connection Broker, RD Gateway, and RD Web Access).

Session collections (referred to as a farm in previous versions of Windows Server) are groupings of RD Session Host servers. A session collection is used to publish one of the following resources:

- Session desktops
- RemoteApp programs

A Session Virtualization deployment in Windows Server 2012 offers the following benefits:

- **Unified central experience.** You can deploy session virtualization in a few hours by using commodity hardware, and then you can centrally manage your session-based desktops and RemoteApp programs.
- **Simplified and centralized deployment.** Simple scenario-based installations enable you to create an entire session collection at one time.
- **User personalization.** User profile disks enable you to preserve user personalization settings for your session collections.
- **Centralized and unified management.** Manage all of the RD Session Host servers in your session collections from a single location.
- **Fairshare experience.** Provides a predictable user experience and ensures that one user does not negatively impact the performance of another user's session. To facilitate positive user experiences, the following features are enabled by default on RD Session Host servers:
  - **Network Fairshare.** Dynamically distributes available bandwidth across sessions to enable equal bandwidth utilization, based on the number of active sessions.
  - **Disk Fairshare.** Prevents sessions from over utilizing disk usage by equally distributing disk I/O among sessions.
  - **CPU Fairshare.** Dynamically distributes processor time across sessions, based on the number of active sessions and the load on these sessions. This was introduced in Windows Server 2008 R2, and has been improved for heavier loads in Windows Server 2012.

Session Virtualization Quick Start

Session Virtualization Quick Start installs everything on one computer that you need to evaluate the Session Virtualization deployment scenario.

Session Virtualization Quick Start does the following:

- Installs the RD Connection Broker, RD Web Access, and RD Session Host role services on a single computer
- Creates a session collection
- Publishes a session-based desktop for each RD Session Host server in the collection
- Publishes RemoteApp programs

Note
Session Virtualization Quick Start is not recommended for a production environment.

Session Virtualization Standard Deployment

Session Virtualization Standard Deployment enables you to install the appropriate role services on separate computers. Unlike session virtualization Quick Start, this deployment type gives you more control over the session collection and published RemoteApp programs by not creating them automatically.

Enabling high availability for your Session Virtualization deployment

Remote Desktop Services in Windows Server 2012 provides support for highly available session collections by leveraging a clustered environment.

The RD Connection Broker role plays a central role in Remote Desktop Services scenarios by providing users access to their session-based desktops or RemoteApp programs that are hosted on RD Session Host servers.

With Windows Server 2012, you can deploy clustered instances of RD Connection Broker for high availability and improved scalability. With a centralized database, the RD Connection Broker service now needs to be configured only once. Additional RD Connection Broker servers that are added to the deployment will be added as an active instance.

If one of the RD Connection Broker servers in the cluster fails, users are not impacted because they are redirected to another RD Connection Broker server automatically.

Centralized resource publishing

In Windows Server 2008 R2, publishing and managing applications on pooled and personal virtual desktops is a time consuming and costly process. RemoteApp programs only partially integrate with the native Windows experience, so they add to the management cost because there is virtually no way to organize published RemoteApp programs for end users.

Remote Desktop Services in Windows Server 2012 enables you to publish and manage resources (such as RemoteApp programs, session-based desktops, and virtual desktops) from a centralized console. With the new publishing features in the centralized console, you can see an historic view of resources that are assigned to end users, you can change published resources for any virtual desktop collection or session collection, and you can edit the properties of published resources. Centralized resource publishing provides end users with an experience that can replace locally installed applications.

In addition to publishing and managing resources from the centralized console, you can now configure a RemoteApp and Desktop connection URL by using Group Policy, and then make this URL automatically available to end users by having them enter their email address.

Rich user experience with RemoteFX

Windows Server 2008 R2 and Windows 7 with Service Pack 1 introduced Microsoft RemoteFX, which enables the delivery of a full Windows user experience to a range of client devices including rich clients, thin clients, and ultrathin clients. Windows Server 2012 builds on this platform to enable a richer and more seamless experience on all types of networks and all types of devices. Specifically, Remote Desktop Protocol (RDP) in Windows 8 Release Candidate enables a consistent and seamless user experience when connecting to centralized workspaces, even on networks where bandwidth is limited and latency is high.

RDP in Windows 8 Release Candidate introduces new capabilities and enhances several existing capabilities that deliver a great end-user experience. Some of the key capabilities include:

- Superior wide area network (WAN) performance with RemoteFX for WAN
- Full three-dimensional (3-D) and Windows Aero remoting experience with a software GPU
- Rich desktop remoting experience for all content types with RemoteFX Adaptive Graphics
- Smooth media playback experience with RemoteFX Media Streaming
- Rich client-side device support with RemoteFX multiple-touch and RemoteFX USB Redirection

RemoteFX for WAN

RDP in Windows 8 has been optimized to work better over low bandwidth, high-latency connections through the addition of the following features:

- **User Datagram Protocol (UDP) transport.** RDP in Windows 8 intelligently chooses between the TCP and UDP transports, depending on the content type and the quality of the connection. When Remote Desktop is enabled on a computer, UDP for port 3389 is automatically enabled in the Windows firewall. For enhanced performance, make sure that this port is enabled on your network.
- **RemoteFX Network Auto Detect.** RemoteFX Network Auto Detect determines the amount of available bandwidth between the client and server, and then uses this information to optimize the user's experience.
- RemoteFX Network Auto Detect is automatically enabled by using Remote Desktop Connection.

RemoteFX Adaptive Graphics

RemoteFX Adaptive Graphics is a new feature in Windows 8, and it enables a seamless delivery of virtual desktop and RemoteApp programs by using the Windows Aero and 3-D experience across a variety of networks, including networks where bandwidth is limited and latency is high.

Two of the key features that are enabled by RemoteFX Adaptive Graphics are:

- RemoteFX progressive download
- Windows Aero and 3-D experience through Microsoft Basic Render Driver (a software GPU)

By default, the RemoteFX graphics processing pipeline will adaptively determine the optimal RDP experience level, based on available bandwidth and server resource availability. You can change the RDP experience level by using Group Policy.

RemoteFX Media Streaming

RemoteFX Media Streaming enables a smooth multimedia experience on networks where bandwidth is limited and latency is high. The key features in RemoteFX Media Streaming are smooth video playback that uses H.264-encoded video streaming and audio video synchronization.

RemoteFX Media Streaming requires the Desktop Experience feature.

RemoteFX multiple-touch

RDP in Windows 8 introduces true multiple-touch and gesture remoting with support for up to ten simultaneous touch inputs. This enables users to use the new touch and gesture-enabled applications in Remote Desktop environments.

RemoteFX USB redirection

RDP in Windows 8 supports RemoteFX USB Redirection for RD Session Host. In Windows Server 2008 R2 with SP1, RemoteFX USB Redirection was only supported within virtual desktops that were using the RD Virtualization Host role service. Configuring RemoteFX USB Redirection for RD Session Host uses the same configuration steps as configuring RemoteFX USB Redirection for RD Virtualization Host. For more information about RemoteFX USB Redirection for RD Virtualization Host, see Configuring USB Device Redirection with Microsoft RemoteFX Step-by-Step Guide.

The following is a list of devices that can be used with RemoteFX USB Redirection:

- All-in-one printer
- Scanner
- Biometric
- Webcam
- VoIP telephone and headset

RemoteFX vGPU

RemoteFX vGPU enables an administrator to share a physical GPU on a Hyper-V server across multiple virtual machines. This feature enables a desktop graphics experience that is similar to that on a physical computer because applications running in a virtual machine can access GPU resources.

Hardware requirements

The Hyper-V server and the RemoteFX-enabled virtual desktop must meet the RemoteFX hardware requirements. The requirements have not changed since the release of Windows Server 2008 R2 with SP1. They include:

- Ensure that the hyperthreading technology is enabled on the Hyper-V server.
- Configure the proper memory as required. Per the Windows 8 requirements, if you are using an x86-based virtual machine, you must configure at least 1024 megabytes (MB) of RAM. If you are using an x64-based virtual machine, you must configure at least 2048 MB of RAM.
- Ensure that the builds match on the RemoteFX server, the virtual desktop, and the client computer.
- For more information about the hardware requirements for deploying RemoteFX vGPU, see Hardware Considerations for RemoteFX.

Resilient File System

Resilient File System (ReFS), a new local file system introduced in Windows Server 2012, maximizes data availability and online operation, despite errors that would historically cause data loss or downtime. The data integrity built into ReFS ensures that business critical data is protected from errors and available when needed, while the ReFS architecture is designed to provide scalability and performance in an era of constantly growing data set sizes and dynamic workloads.

Bottom line

ReFS provides a cost-effective platform that maximizes data availability, scales efficiently to very large data sets across diverse workloads, and guarantees data integrity via resiliency to corruption (regardless of software or hardware failures).

Requirements

- Windows Server 2012
- No special hardware is needed

Technical overview

ReFS is designed with three key goals in mind:

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

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

Robust disk updating. ReFS offers robust disk updating with an allocate on write transactional model (also known as "copy on write"). This approach maximizes reliability by avoiding problems associated with power failures during disk updates. This is accomplished by updating data with writes to different locations in an atomic fashion, rather than updating data in-place.

Data integrity. To enable detection of all forms of disk corruption, all file system metadata is updated using allocate-on-write and is additionally protected with checksums. ReFS includes the optional ability to apply data integrity to user data as well. When this option, known as Integrity Streams, is enabled, ReFS always uses allocate-on-write for updates to user data and uses checksums to deterct disk corruption. While this option is useful in many scenarios, it is not appropriate in some cases. This is why mechanisms and APIs are provided to enable or disable Integrity Streams settings at various levels of granularity including per-volume, per-directory, and per-file. When ReFS is used in conjunction with a mirrored Storage Space, detected corruption—both metadata and user data, when integrity streams are enabled—can be automatically repaired using the alternate copy provided by Storage Spaces.

Availability. ReFS is designed to maximize online operation even if corruption occurs. While it is expected that many customers will use ReFS in conjunction with mirrored Storage Spaces to automatically repair corruption transparently, it is still possible, though rare, for such a volume to become corrupted. Additionally, some customers may choose not to use a mirrored storage space for their ReFS volumes. In these cases, when corruption does occur, ReFS implements "salvage" which is a feature that will remove corrupt data from the namespace on a live volume so that good data is not adversely affected by non-repairable corrupt data. This ensures that volumes do not have to be taken offline to correct errors and enables administrators to simply restore files removed from the namespace by restoring them from backup. There is no Chkdsk with ReFS.

Scalability. As the amount and size of data that is stored on computers continues to increase rapidly, ReFS is designed to work well with extremely large data sets, petabytes and larger, without performance impact. While practical concerns surrounding system configurations (such as the amount of memory), limits by various system components and the time taken to populate data sets or backup times may define practical limitations. The ReFS on-disk format is designed to support volume sizes up to 2^78 bytes using 16-KB cluster sizes while Windows stack addressing allows 2^64 bytes. This format also supports 2^64-1 byte file sizes, 2^64 files in a directory and the same number of directories in a volume.

Application compatibility. ReFS maintains a high degree of compatibility with a subset of NTFS features that are widely adopted while deprecating others that provide limited value at the cost of system complexity and footprint. ReFS supports most NTFS features and Win32 APIs.

Proactive error identification. ReFS integrates with a data scrubber that periodically scans the volume, trying to detect latent corruption and, when running on top of a Mirrored Storage Space, automatically repair the corrupted data.

Interoperability and Flexibility. ReFS is designed to fit cleanly into the Windows storage stack with maximum flexibility and compatibility with other layers of the stack. While extensively tested to ensure compatibility with supporting software, such as backup and antivirus applications, ReFS is designed to work well together with features in other layers of the storage stack. One of the best examples of this is when ReFS is used in conjunction with Storage Spaces. ReFS can easily take advantage of storage pools shared between multiple machines and virtual disks that can easily transition between them, providing additional resiliency to failures above what Storage Spaces or ReFS could do alone. ReFS can be installed on third-party storage subsystems as well.

Conclusion

By utilizing an integrated storage stack comprising ReFS and Storage Spaces, you can now deploy the most cost-effective platform for available and scalable data access using commodity storage.

Secure Naming Services

Domain Name System Security Extensions (DNSSEC) is a suite of extensions that add security to the DNS protocol. The core DNSSEC extensions are specified in Internet Engineering Task Force (IETF) Request for Comments (RFCs) 4033, 4034, and 4035, with additional RFCs providing supporting information. Specifically, DNSSEC provides origin authority, data integrity, and authenticated denial-of-existence. In addition to several new concepts and operations for DNS servers and DNS clients, DNSSEC introduces four new resource records to DNS: DNSKEY, RRSIG, NSEC/NSEC3, and DS.

Problem

Domain Name System (DNS) is an important building block of the Internet. DNS translates readable text names, such as www.example.com, to IP addresses, such as 198.51.100.100, which a computer can understand. Since its inception, DNS has played a key role in the operation of the Internet, and it continues to provide scalable and extensible naming infrastructure for computer networks. In addition, DNS plays a critical role in Active Directory Domain Services (AD DS) enterprise networks by providing vital name resolution services that enable client computers and other devices to find domain controllers, other clients, and other resources in enterprise networks.

However, DNS as originally defined in RFCs 1034 and 1035 has no security component. As the Internet grew in scale and importance, it became clear that leaving DNS unsecured presented a major security risk. By spoofing DNS responses, an attacker can redirect clients to malicious servers, which in turn compromises the security of the DNS client. Because almost all network communication begins with a name resolution, compromising DNS results in compromising all forms of network communication.

Solution

DNSSEC is a suite of additions to DNS that helps protect DNS traffic from man-in-the-middle and spoofing attacks. By validating a digital signature that is attached to each DNS response, the resolver can verify the authenticity of the DNS data (even if it is coming from an untrusted DNS server).

DNSSEC support in Windows Server 2008 R2

Windows Server 2008 R2 introduced support for DNSSEC, and it provides the ability to generate keys and host a signed zone. However, there are several limitations to the support that is provided in Windows Server 2008 R2:

- Zones can only be signed offline, and only by using a file-based copy of the zone. It is not possible to generate signatures or update signatures on a zone while the zone is online.
- The processes of key generation and zone signing are manual, and require the command-line tool dnscmd.exe.
- Dynamic updates to DNS records are not supported.
- Lack of support for automatic key rollover.
- Lack of support for some security standards, such as NSEC3, RSA/SHA-2 signing algorithms, and automatic trust anchor rollovers.

DNSSEC support in Windows Server 2012

Windows Server 2012 introduces support for online signing and enables automation of key management.

New supported features on the authoritative DNS server include:

- Support for DNS scenarios that are integrated in Active Directory, including DNS dynamic updates in DNSSEC-signed zones.
- Support for updated DNSSEC standards, including NSEC3 and RSA/SHA-2.
- Enables automation for trust anchor distribution through Active Directory Domain Services.
- Enables automation for trust anchor rollover support per RFC 5011.
- Updated user interface with deployment and management wizards.
- Windows PowerShell command-line interface for easier management and scripting.

New supported features on the non-authoritative DNS resolver include:

- Validation of records that are signed with updated DNSSEC standards (NSEC3, RSA/SHA-2).
- Enables automation for trust anchor rollover support per RFC 5011.
- Extraction of the root trust anchor.

DNSSEC in Windows Server 2012 primarily supports enterprise customers who want to protect their internal DNS infrastructure, ensure that names resolved from the Internet are secure, and ensure that DNS communication between their enterprise networks and those of other enterprises is secured.

DNSSEC deployment is a phased process that begins with signing DNS zones. After the zones are signed and the stability of the DNS infrastructure is ensured, validation of the DNS responses is enabled on caching resolvers. A number of settings are available in Windows Server 2012 to secure the link between the DNS client and the local caching resolver.

The following sections discuss some key aspects of the DNSSEC support in Windows Server 2012 that enable the signing of DNS zones in Active Directory environments and the validation of DNS responses at the caching resolver.

The Key Master

In a multiple master DNS deployment, which is typical in enterprise Active Directory environments, each master server is identical to all of the other master servers. However, for DNSSEC, a single master server must perform key generation and key management. Windows Server 2012 introduces the concept of a *Key Master* (KM) for DNSSEC. Any authoritative DNS server that hosts a primary copy of the zone can be designated as the KM, and virtually any server can be the KM for multiple zones if it hosts a primary copy of the zone. An administrator selects the KM to be the DNS server in charge of key management for a specific zone. Different DNS servers can also function as key masters for different zones. The KM is in the context of a specific zone and is not global across multiple servers, zones, domains, or forests. When the administrator initially performs DNSSEC operations on a zone, the current server automatically becomes the KM for that zone. DNSSEC is configured on the KM, and the KM is responsible for key generation and key management for the zone.

The KM generates all keys for the zone, and it is responsible for distributing private keys and zone signing information. Only the KM can perform any operations with or on the key signing key (KSK). The KM is fully responsible for performing KSK and zone signing key (ZSK) rollovers, and for polling child zones to keep signed delegations up-to-date. The server designated as the KM must be online, and it must be a high availability server to ensure uninterrupted service for key signing operations.

Automation of post-signing tasks

A primary limitation of DNSSEC in Windows Server 2008 R2 is that any changes made to a DNSSEC-signed zone require manual offline resigning of the zone. Dynamic updates are therefore disabled on any signed zone, which severely limits the deployment of signed zones that are integrated in Active Directory. In addition, there was no way to re-sign zones automatically when signatures neared expiration.

In Windows Server 2012, the DNS server can perform all of these operations to keep signatures up-to-date automatically without the need for administrative intervention. This significantly reduces the cost-of-ownership of a DNSSEC-signed zone.

Examples of how automation works in Windows Server 2012:

- **Dynamic updates.** Dynamic updates can be enabled on a signed zone, and any DNS server that is authoritative for the zone can accept dynamic updates. When a DNS server running Windows Server 2012 receives an update, it automatically generates signatures for the update and adds the update and its signatures to the zone. Through Active Directory replication, the unsigned update replicates to all other authoritative servers, and each authoritative server generates signatures for the update and adds it to its own copy of the zone.
- **Updates to the zone through management tools.** Updates to zone data through management tools such as DNS Manager, dnscmd.exe, and Windows PowerShell no longer require a manual re-signing of the zone.
- **Keeping signatures up to date.** The DNS server running Windows Server 2012 keeps track of signature expiration, and it automatically refreshes signatures that are about to expire to ensure continuous availability of zone data.
- **Scavenging.** Scavenging stale records in a signed zone works exactly like it does in unsigned zones.

Key rollovers

DNSSEC keys do not have a permanent lifetime, and they require periodic replacement. The longer a key is in use, the greater the risk of compromise. The key replacement process is known as a *key rollover*. Performing key rollovers are a vital part of operating a DNSSEC-signed zone. In Windows Server 2008 R2, administrators must perform the rollover process manually.

In Windows Server 2012, DNS enables automation for key rollover management, including provisioning and configuration of the supported methods for ZSK and KSK rollover, and the actual process of performing a key rollover.

Hashed authenticated denial-of-existence (NSEC3)

When the original DNSSEC standards were defined (in RFCs 4033, 4034, and 4035), the specifications included a mechanism to provide authenticated denial-of-existence through a resource record called *Next Secure* or NSEC. To provide this functionality, the DNS server sorts the contents of the zone in the canonical fashion before signing.

A canonical name order is required to construct the NSEC name chain and to construct and verify RRSIG resource records. The DNS server orders the owner names by sorting the names according to their most significant labels. For example, the following names are sorted in canonical DNS name order:

- example
- a.example
- yljkjljk.a.example
- Z.a.example
- zABC.a.EXAMPLE
- z.example

When the zone is signed, an NSEC resource record is generated at each name, which contains a pointer to the next name, thus providing a linked list of zone data. This NSEC record proves that nothing exists between the two names. However, NSEC enumerates the contents of the zone and allows a technique known as *zone walking*. A malicious resolver can issue several non-existent queries and enumerate the entire zone by analyzing the NSEC records returned.

To address this security concern, RFC 5155 creates the NSEC3 resource record, which offers authenticated denial-of-existence; but in addition, it protects against undesired zone enumeration. To provide protection against zone enumeration, the owner names that are used in the NSEC3 resource record are cryptographic hashes of the original owner name that is prepended as a single label to the name of the zone. NSEC3 also allows a signed parent zone to host unsigned delegations and not require a re-signing whenever an unsigned delegation changes. This is a highly attractive feature for large zones like .com that are expected to have hundreds of thousands of unsigned delegations.

Windows Server 2012 introduces support for NSEC3 to the DNS server for its authoritative and non-authoritative operations, including full support for the NSEC3 and NSEC3PARAM resource records. Likewise, the DNS client can recognize the NSEC3 and NSEC3PARAM records. During the initial signing process, administrators can select NSEC or NSEC3 for the preferred denial-of-existence mechanism.

Trust anchors

In cryptographic terms, a *trust anchor* (TA) is an authoritative entity that is represented by a public key and associated data. In DNSSEC terms, a trust anchor is a configured DNSKEY resource record or it is a DS resource record hash of a DNSKEY resource record. A resolver uses this public key or hash as a starting point for building the authentication chain to a signed DNS response. The TA for a signed zone must be configured on every DNS resolver that attempts to validate DNS data from that signed zone.

During the configuration process, if the DNS server that is hosting the zone is a domain controller, you can distribute trust anchors automatically to other DNS servers in the domain. This includes file-backed zones that are hosted on domain controllers. If the KM is a stand-alone DNS server that is not integrated in Active Directory (such as a member server), this option is not available.

Enables automation to update DNSSEC trust anchors

When an administrator uses dnscmd.exe to sign a DNS zone in Windows Server 2008 R2, the server creates a text file that contains the trust anchor (TA) DNS records. If the current zone does not have a parent zone that is signed, and it represents the root of an island of security, the administrator can optionally publish or distribute the TAs to administrators of other organizations.

To validate the DNSSEC protected data, DNSSEC-aware resolvers must have knowledge of the TA for that zone. In Windows Server 2008 R2, creating the TA at the trust point and adding the TA to the DNSSEC-aware resolvers is a manual process. If the public key expires or is compromised, the TA has to be updated manually by deleting the current TA and adding a new TA at the trust point. DNSSEC-aware resolvers have to manually update their local copies of the TA. In Windows Server 2012, DNSSEC-aware resolvers poll for TA changes and automatically keep their copy of the TA updated.

Hosting signed zones on read-only domain controllers

In Windows Server 2008 and Windows Server 2008 R2, DNS servers that run on read-only domain controllers (RODCs) host copies of all zones that are integrated in Active Directory. However, the DNS server cannot make any updates to the zones that it hosts. It is considered unsafe to replicate private keys to an RODC, because RODCs are designed to operate in physically insecure environments.

In Windows Server 2012, an RODC loads unsigned zones in Active Directory without a change in functionality from Windows Server 2008 R2. However, if the RODC finds a signed zone in the directory, it does not load the zone as if it is integrated in Active Directory. Instead, it creates a secondary copy of the zone, and then configures the closest writeable domain controller for the domain as the primary server.

The RODC then attempts to perform a zone transfer. If zone transfers are enabled on the primary DNS server that the RODC selected, then the transfer succeeds and the RODC begins to respond to queries for the zone. However, if zone transfers are not enabled, the transfer operation fails. The RODC logs an error event and takes no further action. In this scenario, the administrator must manually enable zone transfers on the primary DNS server selected by the RODC. Alternately, the administrator can reconfigure the RODC to point to a different primary server that has zone transfers enabled.

DNSSEC requirements

Software requirements

Windows Server 2008 R2 supports DNSSEC-signed zones. However, if a zone is signed with NSEC3, authoritative DNS servers must be running Windows Server 2012. You cannot sign or host a zone that is signed with NSEC3 by using a DNS server running Windows Server 2008 R2.

Client computers must be running Windows 8 to perform DNSSEC validation on data from a zone that has been signed with NSEC3.

Server Core and Full Server Integration

In Windows Server 2012, the Server Core installation option is no longer an irrevocable selection that is made during setup. In Windows Server 2008 R2 and Windows Server 2008, if your requirements changed, there was no way to convert to a full installation or a Server Core installation without completely reinstalling the operating system. An administrator now has the ability to convert between a Server Core installation and a full installation as needed.

There are several scenarios in which this capability is especially useful:

- An administrator installed and is running a full installation option of Windows Server 2012, but exclusively using the roles that run on a Server Core installation. The administrator can convert the servers to Server Core installations to reduce the image size and increase servicing advantages without having to reprovision all of their servers.
- An administrator installs a Server Core installation and now needs to make a change or troubleshoot something that is not possible with the remote GUI. The administrator may not know how to make the change from the command line or cannot find a command-line equivalent. The administrator can convert the server to a full installation, perform the changes, and then convert it back to a Server Core installation to reduce the image size and maintain servicing advantages.
- An administrator wants to use the GUI for all of the initial configuration steps to make the initial configuration experience as easy as possible, yet wants to reduce the image size and maintain the servicing advantages that a Server Core installation provides. The administrator can install a full installation, configure the server as needed, and then convert it to a Server Core installation.
- An enterprise mandates a single server operating system image, so it cannot use a Server Core installation because it requires two images. Windows Server 2012 integrates the Server Core installation and the full server installation options. Now the enterprise can use a single server operating system image to deploy full installations of Windows Server 2012 and then convert them to Server Core installations to reduce the image size and provide the servicing advantages that it offers.

Requirements

You need a Server Core installation or a full installation of Windows Server 2012. No special hardware is required.

Technical overview

In Windows Server 2008 R2 and Windows Server 2008, the Server Core installation and the full installation options were selections that an administrator made at the time of installation.

In Windows Server 2012, the installation options are integrated, and three large optional features are provided. An administrator can install or uninstall these options to move between Server Core and full server installations.

Server for NFS Data Store

In the past several years, organizations have been increasingly pressured to reduce their IT spending. Reduced IT budgets drive a business need for IT to find new ways to decrease costs. Because the cost of storage constitutes a sizeable portion of IT budgets, this trend has resulted in the shift towards converged storage fabric. Virtualization is a large driver of the increase in demand for storage capacity. File-based storage has gained credence as a viable alternative to more expensive storage area network (SAN) block-based storage because it is simple to provision and manage. An example of this trend is deploying and running VMware ESX and VMware ESXi virtual machines from file-based storage accessed over the Network File System (NFS) protocol.

The need to store and run critical virtualized enterprise workloads from file-based storage demands a guarantee for increased availability and reliability. Windows Server 2012 includes several key improvements to the file system, storage, and networking stack. Server for NFS is the NFS server implementation that is included with Windows Server operating systems. For Windows Server 2012, Server for NFS has been updated to support continuous availability. These improvements make it possible to reliably store and run VMware ESX on a virtual machine that is running Windows Server 2012 by using the NFS protocol to share files. Additional improvements in failover clustering make it possible to deploy Server for NFS in a clustered configuration and take advantage of better resilience to hardware and software outages that may afflict individual cluster nodes. At the same time, failover clustering isolates VMware ESX hypervisor hosts from intermittent glitches during failovers.

Figure 37: NFS sharing files

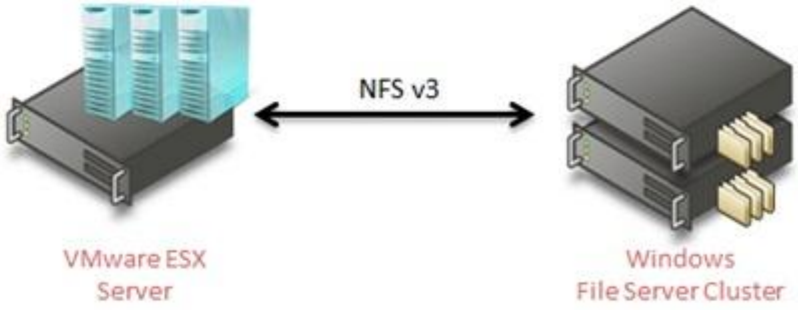

Requirements

The following items are required to deploy and run VMware ESX on virtual machines that are using Server for NFS as a data store:

- A computer running Windows Server 2012 with the failover clustering feature installed and shared storage
- A server running VMware ESX 4.1
- A computer that is used for management (such as a computer running Windows 7) with VMware vSphere Client version 4.1 installed on it
- Installable media, such as a CD or ISO image of the guest operating system that you want to run on the VMware ESX hypervisor host

On the computer running Windows Server 2012, you need to use Server Manager to install the following features:

- The File Services role
- The Server for NFS role service
- The failover clustering feature

This feature set includes Windows PowerShell cmdlets that you can use to set NFS to share files. You can also use these tools to quickly provision the shared files as a data store for servers running VMware ESX and to run virtual machines on the shared files by using the NFS protocol.

Technical overview

Windows Server 2012 includes a new Open Network Computing Remote Procedure Call (ONCRPC), which consists of an External Data Representation (XDR) runtime built on the Winsock Kernel (WSK) interface. This configuration replaces the previous runtime that was built on the legacy Transport Driver Interface (TDI). The new XDR runtime is equipped with a dynamic autotuned thread pool that does not require performance tuning using registry keys, and it features a more efficient XDR encoding and decoding implementation. In addition, it implements dynamic endpoints and consumes Plug and Play notifications to support the addition and removal of network interfaces at runtime. With this new design, the XDR runtime infrastructure is capable of performing faster and more targeted failovers to provide continuous availability. In Windows Server 2012, Server for NFS and Client for NFS have been ported to run on this new scalable, high-performance XDR runtime infrastructure.

The failover cluster resource DLL is completely rewritten in Windows Server 2012, which enables Server for NFS to fail over much faster than on previous versions of Windows Server. The previous per-share cluster resource for NFS has been replaced by a consolidated per-virtual disk cluster resource, which results in fewer cluster resources on the failover cluster and faster end-to-end failovers. The failover logic within the NFS cluster resource has also been reworked for performance. The failover clustering infrastructure has undergone several key performance enhancements, which has resulted in faster failovers of the disk and network resources. This helps reduce the total amount of time for a failover to occur, and it makes it easy to deploy shared files as a reliable file-based storage backend for VMware ESX hypervisor hosts.

In addition, Windows Server 2012 includes a comprehensive set of task-oriented Windows PowerShell cmdlets that make it easier to provision and manage the NFS shares.

Storage Management

Windows Server 2012 enables storage management that is comprehensive and fully scriptable, and administrators can manage it remotely. A WMI-based interface provides a single mechanism through which to manage all storage, including non-Microsoft intelligent storage subsystems and virtualized local storage (known as Storage Spaces). Additionally, management applications can use a single Windows API to manage different storage types by using standards-based protocols such as Storage Management Initiative Specification (SMI-S).

Key benefit

Windows Server 2012 reduces complexity and cost through a comprehensive storage management interface.

Requirements

Your storage infrastructure should support the Storage Management Provider interface or support an appropriate version of SMI-S.

Technical overview

The unified interface for storage management in Windows Server 2012 provides not only a core set of defined WMI and Windows PowerShell interfaces, but also features for more advanced management, as shown in the following figure.

Figure 38: Unified storage management architecture

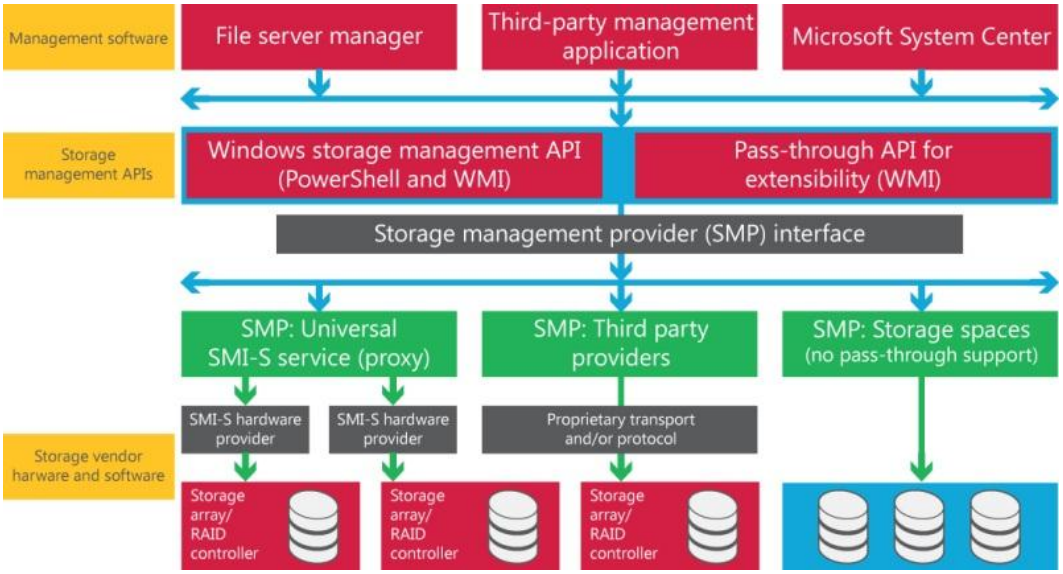
Management software
File server manager
Third-party management application
Microsoft System Center
Storage management APIs
Windows storage management API (PowerShell and WMI)
Pass-through API for extensibility (WMI)
Storage management provider (SMP) interface
SMP: Universal SMI-S service (proxy)
SMP: Third party providers
SMP: Storage spaces (no pass-through support)
SMI-S hardware provider
SMI-S hardware provider
Proprietary transport and/or protocol
Storage vendor harware and software
Storage array/ RAID controller
Storage array/ RAID controller
Storage array/ RAID controller

This design offers specific advantages for the following users:

- **Enterprise system administrators.** Uniform, scriptable management by using Windows PowerShell with a comprehensive set of cmdlets that can be used for discovery; thin provisioning support; snapshot management; replication; masking and unmasking; enumerating HBA ports; and creating pools, logical units, and volumes. As an example, an administrator can use a single script to configure host resources and to configure and present the storage to Windows.
- **ISVs.** Flexibility to administer virtually any type of storage that is connected to a Windows system.
- **Storage manufacturers.** Easier integration of devices with any storage management client, which ensures consistent experience for users.

The unified storage management interface requires the use of a storage management provider that is based on SMI-S or WMI. The interface integrates more easily with services offered by SMI-S providers, which enables administrators to manage non-Microsoft intelligent storage subsystems that have SMI-S providers. For WMI-based services, the interface includes a rich set of built-in storage management features that ISVs will find particularly helpful.

Storage Spaces

Storage Spaces in Windows Server 2012 enables cost-effective, optimally used, highly available, scalable, and flexible storage solutions for business-critical (virtual or physical) deployments. Windows Server 2012 delivers sophisticated storage virtualization capabilities, which empower customers to use industry-standard storage for single-node and scalable multiple-node deployments.

Key benefit

Confidently deploy highly available and scalable storage infrastructure at a significantly lower cost.

Requirements

For single-node deployments:

- Windows Server 2012
- Serial ATA (SATA) or Serial Attached SCSI (SAS) connected disks (in an optional just-a-bunch-of-disks [JBOD] enclosure)
- For multiple-node clustered shared-storage deployments:
    - Two or more servers running Windows Server 2012
    - Requirements as specified for failover clustering and Windows CSV
    - SAS connected JBODs that comply with Windows Certification requirements

Technical overview

Core technologies

Storage Spaces delivers storage virtualization capabilities within Windows Server 2012. The storage stack has been fundamentally enhanced to incorporate two new abstractions:

- Storage pools are administrative units of physical disks. Pools permit storage aggregation, elastic capacity expansion, and delegated administration.
- Storage spaces are virtual disks with associated attributes such as a desired level of resiliency, thin or fixed provisioning, automatic or controlled allocation on heterogeneous storage media, and granular administrative control.

Storage Spaces are manageable through the Windows Storage Management API in Windows Management Instrumentation (WMI) and Windows PowerShell, and through the File Services GUI in Server Manager. Storage Spaces is completely integrated with failover clustering for high availability, and it is integrated with CSV for scale-out deployments.

Features

Storage Spaces includes the following features:

- **Storage pools.** Storage pools are the fundamental building blocks for Storage Spaces. Administrators are already familiar with this concept, so they will not have to learn a new model. They can flexibly create storage pools based on the needs of the deployment. For example, given a set of physical disks, an administrator can create one pool (by using all the available physical disks) or multiple pools (by dividing the physical disks as required). Furthermore, to maximize the value from storage hardware, the administrator can map a storage pool to combinations of hard disks as well as solid-state drives (SSDs). Pools can be expanded dynamically by simply adding additional drives, thereby more easily scaling to cope with unceasing data growth.
- **Multitenancy.** Administration of storage pools can be controlled through access control lists (ACLs) and delegated on a per-pool basis, thereby supporting hosting scenarios that require tenant isolation. Storage Spaces follows the familiar Windows security model; therefore, it can be fully integrated with Active Directory Domain Services.
- **Resilient storage.** Storage Spaces support two optional resiliency modes: mirroring and parity. Per-pool support for disks that are reserved for replacing failed disks (hot spares), background scrubbing, and intelligent error correction enable continuous service availability despite storage component failures.
- **Continuous availability.** Storage Spaces is fully integrated with failover clustering, which enables it to deliver continuously available service deployments. One or more pools can be clustered across multiple nodes within a single cluster. Storage spaces can then be instantiated on individual nodes, and the storage will more easily fail over to a different node when necessary (in response to failure conditions or due to load balancing). Integration with CSVs permits scale-out access to data.
- **Optimal storage use.** Server consolidation often results in multiple data sets sharing the same storage hardware. Storage Spaces supports thin provisioning to allow businesses to easily share storage capacity among multiple unrelated data sets and thereby maximize capacity use. Trim support permits capacity reclamation when possible.
- **Operational simplicity.** Management that is fully remote and scriptable is permitted through the Windows Storage Management API, WMI, and Windows PowerShell. Storage Spaces can be easily managed through the File Services role in Server Manager.

Thin Provisioning and Trim Storage

Sophisticated storage solutions offer just-in-time allocations (also known as thin provisioning) and the ability to reclaim storage that is no longer needed (also known as trim). Windows Server 2012 integrates with these sophisticated storage solutions to enable organizations to get the most out of their storage infrastructures.

Key benefit

Administrators can maximize the benefits of sophisticated storage infrastructure that is accessed through Windows Server 2012.

Requirements

- Enabled by default in Windows Server 2012
- Storage infrastructure that complies with the certification that is required for Windows Server 2012
- Standards-compliant hardware for identification

Technical overview

To confidently deploy sophisticated storage solutions that support just-in-time allocation, you need to know that you can provision additional capacity as needed. Windows Server 2012 identifies thinly provisioned virtual disks, provides standardized notifications when use thresholds are crossed, and provides a platform that enables applications to give up storage when it is no longer needed, thereby ensuring maximal use.

The following capabilities are included:

- **Identification.** Windows Server 2012 uses a standardized method to detect and identify thinly provisioned virtual disks, thereby enabling additional capabilities delivered by the storage stack that is provided in the operating system and through storage management applications.
- **Notification.** When configured physical storage use thresholds are reached, Windows Server 2012 notifies the administrator through events, which enables the administrator to take appropriate action as soon as possible. These events can also be used for automated actions by sophisticated management applications, such as Microsoft System Center.
- **Optimization.** Windows Server 2012 provides a new API that lets applications return storage when it is no longer needed. NTFS issues trim notifications in real time when appropriate. Additionally, trim notifications are issued as part of storage consolidation (optimization), which is performed regularly on a scheduled basis.

Unified Remote Management for File Services

Managing file servers on your network previously required that you connect to each file server individually. In Windows Server 2012, the new File Services management functionality in Server Manager enables you to use a single interface to remotely manage multiple file servers or clustered file server instances.

Key benefits

Administrators can use Server Manager to manage multiple file servers, including clusters, as easily as managing one server. You can use storage pools to group many physical disks into easily managed volumes.

Requirements

Unified Remote Management requires one or more servers running Windows Server 2012.

Technical overview

Server Manager includes a new functionality for the File Services role that enables you to remotely manage multiple file servers, including managing file shares, volumes, and storage. File Services helps you to create pools of volumes on the network, and then manage those storage pools as a unit. You can also use File Services to add and share volumes anywhere on the network.

User-Device Affinity

With User-Device Affinity in Windows Server 2012, you can map a user to a limited set of computers where Folder Redirection, roaming user profiles, or both are used. This provides a simple and powerful method for associating particular user profiles with particular computers or devices, simplifying administrator oversight, improving data security, and helping to protect user profiles from corruption. Windows Server 2012 provides User-Device Affinity by expanding the Active Directory user schema.

Key benefit

Administrators can control which computers roaming user profiles and offline files are stored on.

Requirements

This feature requires the following:

- Windows Server 2012
- Windows 8
- An Active Directory domain that is running the Windows Server 2012 schema

Technical overview

In Windows Server 2012, the Active Directory user schema is extended to support primary computers that are associated with a user. An administrator can populate the ms-DS-Primary-Computer attribute with a list of computer names for a user or group object.

There are three major benefits to this approach:

- The administrator can specify which computers people can use to access their redirected data and settings. For example, the administrator can choose to roam user data and settings among a user's desktop and laptop, but to not roam the information when that user logs on to a conference room computer.
- User-Device Affinity reduces the security and privacy risk of leaving residual personal or corporate data on computers where the user has logged on. For example, a general manager who logs on to an employee's computer for temporary access does not leave behind any personal or corporate data.
- User-Device Affinity enables the administrator to mitigate the risk of an improperly configured or otherwise corrupt profile, which could result from roaming between differently configured systems, such as between x86-based and x64-based computers.

To make User-Device Affinity possible, Folder Redirection, and roaming user profile components incorporate the following additional logic checks when a user logs on to a computer:

1. The Windows operating system checks the new Group Policy setting to determine if the ms-DS-Primary-Computer attribute in Active Directory Domain Services (AD DS) should influence the decision to roam the user's profile or apply Folder Redirection.
2. If the policy setting enables primary computer support, Windows verifies that the AD DS schema supports the ms-DS-Primary-Computer attribute. If it does, Windows determines if the computer that the user is logging on to is designated as a primary computer for the user as follows:
   - If the computer is one of the user's primary computers, Windows applies the roaming user profile and redirected folders.
   - If the computer is not one of the user's primary computers, Windows loads a cached local profile for the user, if present, or it creates a new local profile. Windows also removes any existing redirected folders according to the removal action that was specified by the previously applied policy setting, which is retained in the local Folder Redirection configuration.

Figure 39: User-Device Affinity Group Policy setting

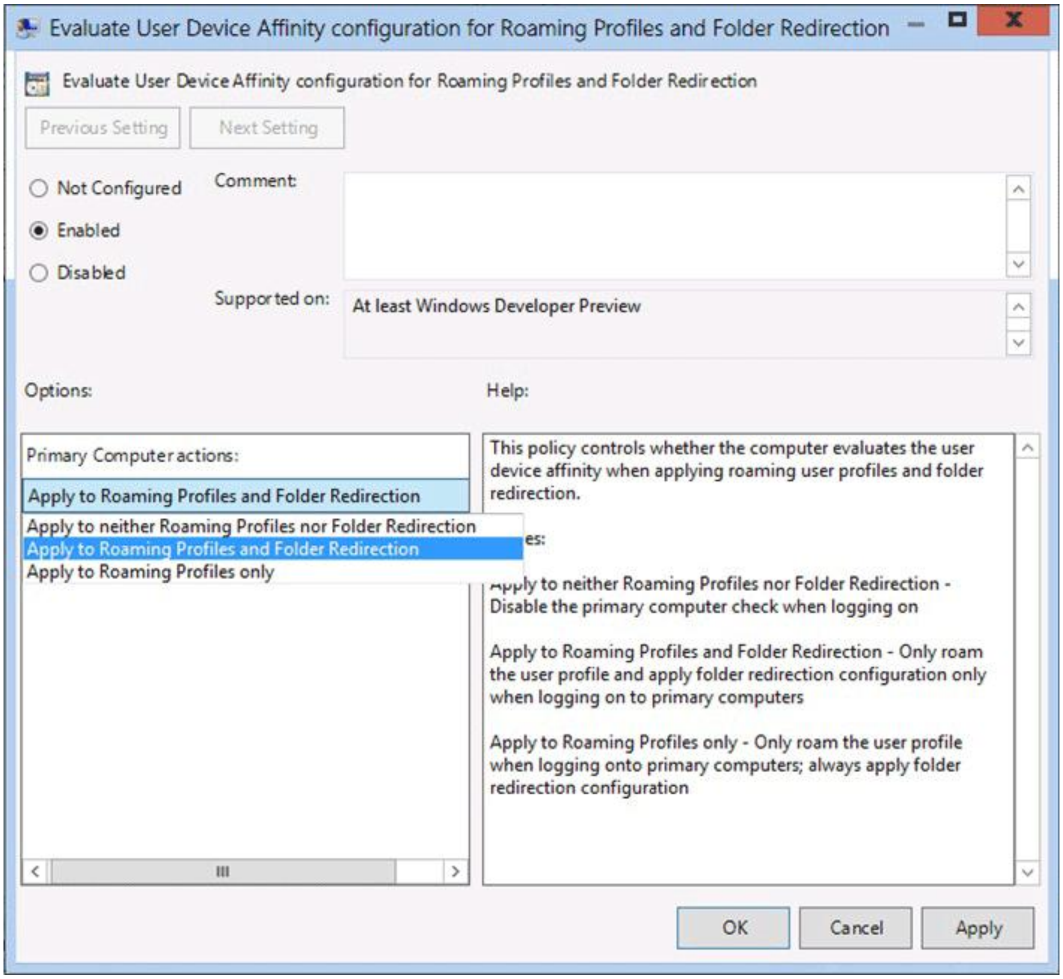
Evaluate User Device Affinity configuration for Roaming Profiles and Folder Redirection
Previous Setting
Next Setting
Comment:
Options:
Help:
OK
Cancel
Apply

Evaluate User Device Affinity configuration for Roaming Profiles and Folder Redirection

○ Not Configured
● Enabled
○ Disabled

Supported on: At least Windows Developer Preview

Primary Computer actions:

Apply to Roaming Profiles and Folder Redirection

- Apply to neither Roaming Profiles nor Folder Redirection
- Apply to Roaming Profiles and Folder Redirection
- Apply to Roaming Profiles only

This policy controls whether the computer evaluates the user device affinity when applying roaming user profiles and folder redirection.

es:

Apply to neither Roaming Profiles nor Folder Redirection - Disable the primary computer check when logging on

Apply to Roaming Profiles and Folder Redirection - Only roam the user profile and apply folder redirection configuration only when logging on to primary computers

Apply to Roaming Profiles only - Only roam the user profile when logging onto primary computers; always apply folder redirection configuration

Virtual Machine Live Migration

Live migration of virtual machines is a key Hyper-V feature in Windows Server 2008 R2. Hyper-V in Windows Server 2012 introduces the following live migration improvements:

Faster and simultaneous migration

Live migrations are now able to utilize higher network bandwidths (up to 10 Gigabit) to complete migrations faster. You can also perform multiple simultaneous live migrations to enable you to move many virtual machines in a cluster quickly. These changes allow you to implement high levels of mobility and flexibility in private cloud solutions.

Live migration outside of a clustered environment

In Windows Server 2012, you can configure a virtual machine so that it is stored on an SMB file share. You can then perform a live migration on this running virtual machine between non-clustered servers running Hyper-V, while the virtual machine's storage remains on the central SMB share. This allows users to gain the benefits of virtual machine mobility without having to invest in the clustering infrastructure if they do not need guarantees of availability in their environment.

You can also perform a live migration of a virtual machine between two stand-alone servers running Hyper-V when you are only using local storage for the virtual machine. In this case, the virtual machine's storage is mirrored to the destination server over the network, and then the virtual machine is migrated, while it continues to run and provide network services.

This functionality enables live migration in the most basic deployments and in more advanced scenarios, such as performing a live migration for a virtual machine among multiple, separate clusters to load balance across an entire datacenter.

Key benefits

Live migration of virtual machines in Windows Server 2012 delivers improved performance and flexibility. It is also now available inside and outside of clustered environments—both with and without shared storage.

Requirements

Common requirements for any form of live migration

- Two (or more) servers running Hyper-V that:
  - Support hardware virtualization.
  - Are using processors from the same manufacturer (for example, all AMD or all Intel).
  - Belong to the same Active Directory domain.
- Virtual machines must be configured to use virtual hard disks or virtual Fibre Channel disks (no physical disks).
- Use of a private network is recommended for live migration network traffic.

Requirements for live migration in a cluster

- Windows Failover Clustering is enabled and configured.
- Cluster Shared Volume (CSV) storage in the cluster is enabled.

Requirements for live migration using shared storage

- All files that comprise a virtual machine (for example, virtual hard disks, snapshots, and configuration) are stored on an SMB share.
- Permissions on the SMB share have been configured to grant access to the computer accounts of all servers running Hyper-V.

Requirements for live migration with no shared infrastructure

- No extra requirements exist.

Technical overview

Hyper-V live migration moves running virtual machines from one physical server to another with no impact on virtual machine availability to users. By pre-copying the memory of the migrating virtual machine to the destination server, live migration minimizes the transfer time of the virtual machine. A live migration is deterministic, which means that the administrator, or script, that initiates the live migration determines which computer is used as the destination for the live migration. The guest operating system of the migrating virtual machine is not aware that the migration is happening, so no special configuration for the guest operating system is needed.

After initiating a live migration, the following process occurs:

1. **Live migration setup occurs.** During the live migration setup stage, the source server creates a TCP connection with the destination server. This connection transfers the virtual machine configuration data to the destination server. A skeleton virtual machine is set up on the destination server and memory is allocated to the destination virtual machine.
2. **Memory pages are transferred from the source node to the destination node.** In the second stage of a live migration, the memory assigned to the migrating virtual machine is copied over the network to the destination server. This memory is referred to as the "working set" of the migrating virtual machine. A page of memory is 4 KB.

   For example, suppose that a virtual machine named "test virtual machine" configured with 1024 MB of RAM is migrating to another server running Hyper-V. The entire 1024 MB of RAM assigned to this virtual machine is the working set of virtual machines called "test virtual machine." The utilized pages within the "test virtual machine" working set are copied to the destination server.

   In addition to copying the working set, "test virtual machine," to the destination server, Hyper-V monitors the pages in the working set for "test virtual machine" on the source server. As memory pages are modified by "test virtual machine," they are tracked and marked as being modified. The list of modified pages is simply the list of memory pages "test virtual machine" has modified after the copy of its working set has begun.

During this phase of the migration, the migrating virtual machine continues to run. Hyper-V iterates the memory copy process several times, with each iteration requiring a smaller number of modified pages to be copied. After the working set is copied to the destination server, the next stage of the live migration begins.

3. **Modified pages are transferred.** The third stage of a live migration is a memory copy process that duplicates the remaining modified memory pages for "test virtual machine" to the destination server. The source server transfers the CPU and device state of the virtual machine to the destination server.

   During this stage, the network bandwidth available between the source and destination servers is critical to the speed of the live migration. Using a 1 Gigabit Ethernet or faster is important. The faster the source server transfers the modified pages from the migrating virtual machines working set, the more quickly the live migration is completed.

   The number of pages transferred in this stage is determined by how actively the virtual machine accesses and modifies the memory pages. The more modified pages there are, the longer it takes to transfer all pages to the destination server.

   After the modified memory pages are copied completely to the destination server, the destination server has an up-to-date working set for "test virtual machine." The working set for "test virtual machine" is present on the destination server in the exact state it was in when "test virtual machine" began the migration process.

   Note
   You can cancel the live migration process at any point before this stage of the migration.

4. **The storage handle is moved from the source server to the destination server.** During the fourth stage of a live migration, control of the storage associated with "test virtual machine," such as any virtual hard disk files or physical storage attached through a virtual Fibre Channel adapter, is transferred to the destination server. (Virtual Fibre Channel is also a new Hyper-V feature in Windows Server 2012.)

5. **The virtual machine is brought online on the destination server.** In the fifth stage of a live migration, the destination server now has the up-to-date working set for "test virtual machine," as well as access to any storage used by "test virtual machine." At this point "test virtual machine" is resumed.

6. **Network cleanup occurs.** In the final stage of a live migration, the migrated virtual machine is running on the destination server. At this point, a message is sent to the network switch. This message causes the network switch to obtain the new MAC addresses of the migrated virtual machine so that network traffic to and from "test virtual machine" can use the correct switch port.

   The live migration process completes in less time than the TCP time-out interval for the virtual machine being migrated. TCP time-out intervals vary based on network topology and other factors. The following variables may affect live migration speed:

   - The number of modified pages on the virtual machine to be migrated—the larger the number of modified pages, the longer the virtual machine will remain in a migrating state.
   - Available network bandwidth between source and destination servers.
   - Hardware configuration of source and destination servers.
   - Load on source and destination servers.
   - Available bandwidth (network or Fibre Channel) between servers running Hyper-V and shared storage.

The live migration process for a virtual machine inside a cluster (when the virtual machine is stored on a CSV volume) and for a virtual machine outside of a cluster (when the virtual machine is stored on an SMB share) is practically identical.

When performing a live migration of a virtual machine between two computers with no shared infrastructure, the first thing that Hyper-V does is perform a partial migration of the virtual machines storage, as follows:

1. Throughout most of the move operation, disk reads and writes go to the source virtual hard disk.
2. While reads and writes occur on the source virtual hard disk, the disk contents are copied over the network to the new destination virtual hard disk.
3. After the initial disk copy is complete, disk writes are mirrored to both the source and destination virtual hard disks while outstanding disk changes are replicated.
4. After the source and destination virtual hard disks are completely synchronized, the virtual machine live migration is initiated, following the same process that is used for live migration with shared storage.
5. Once the live migration is complete and the virtual machine is successfully running on the destination server, the files on the source server are deleted.

Virtual Machine Import Wizard

Administrators often think of a virtual machine as a single, stand-alone entity that they can move around to address their operational needs. However, a virtual machine consists of several parts, which administrators do not normally need to think about:

- Virtual hard disks, stored as files on the physical storage.
- Virtual machine snapshots, stored as a special type of virtual hard disk file.
- The saved state of the different, host-specific devices.
- The memory file for the virtual machine or its snapshot.
- The virtual machine configuration file, which organizes all of those parts and arranges them into a working virtual machine.

Each virtual machine and every snapshot associated with it must be unique, so globally unique identifiers are used. Additionally, virtual machines store and use some host-specific information, such as the path information for virtual hard disk files. When Hyper-V tries to start a virtual machine, it goes through a series of validation checks before being started. Problems such as hardware differences that might exist when a virtual machine is moved to another host can cause these validation checks to fail. That, in turn, prevents the virtual machine from starting. The administrator is left with files on the disk that take up space and are not useful.

Hyper-V in Windows Server 2012 introduces a new Import wizard that detects and fixes more than 40 different types of incompatibilities. The Import wizard walks you through the steps of addressing incompatibilities when you import the virtual machine to the new host—so this wizard can help with configuration that is associated with physical hardware, such as memory, virtual switches, and virtual processors.

Also, you no longer need to export a virtual machine to be able to import it. You can simply copy a virtual machine and its associated files to the new host, and then use the Import wizard to specify the location of the files. This "registers" the virtual machine with Hyper-V and makes it available for use. You can copy a virtual machine to an NTFS-formatted USB drive, and you can recover virtual machines in cases where the system drive fails but the data drive that stores the virtual machines is intact.

In addition to the new wizard, automation support is available. The new Hyper-V module for Windows PowerShell includes cmdlets for importing virtual machines.

Key benefits

The new Import wizard in Windows Server 2012 provides a simpler, improved way to import or copy virtual machines.

Requirements

To try out the Import wizard, you will need the following:

- Two installations of Windows Server 2012 with the Hyper-V role installed (Hyper-V requires a computer that has processor support for hardware virtualization)
- A virtual machine
- A user account that belongs to the local Hyper-V Administrators group

Technical overview

To import a virtual machine, the wizard does the following:

1. **Creates a copy of the virtual machine configuration file.** This is done as a precaution in case an unexpected restart occurs on the host, such as from a power outage.
2. **Validates hardware.** Information in the virtual machine configuration file is compared to hardware on the new host.
3. **Compiles a list of errors.** This list identifies what needs to be reconfigured and determines which pages appear next in the wizard.
4. **Displays the relevant pages, one category at a time.** The wizard explains each incompatibility to help you reconfigure the virtual machine so it is compatible with the new host.
5. **Removes the copy of the configuration file.** After the wizard does this, the virtual machine is ready to be started.

Windows Server 2012 Edition Overview

The Windows Server 2012 product line-up has been streamlined and simplified, making it easier for customers to choose the edition that is right for their needs.

- Datacenter edition for highly-virtualized private cloud environments.
- Standard edition for non-virtualized or lightly virtualized environments.
- Essentials edition for small businesses with up to 25 users running on servers with up to two processors.
- Foundation edition for small businesses with up to 15 users running on single processor servers.

Edition	Feature Comparison	Licensing Model
Datacenter	Unlimited virtual instances All features	Processor + CAL*
Standard	Two virtual instances All features	Processor + CAL*
Essentials	2 processor Limited features	Server 25 User Limit
Foundation	1 processor Limited features	OEM Only

*Client Access Licenses (CALs) are required for every user or device accessing a server.

Licensing overview

The packaging and licensing structure for Windows Server 2012 Datacenter edition and Windows Server 2012 Standard edition has been updated to simplify purchasing and reduce management requirements.

- Two editions differentiated only by virtualization rights – two virtual instances for Standard edition and unlimited virtual instances for Datacenter edition.
- A consistent processor-based licensing model that covers up to two physical processors on a server.

Windows Server 2012 Essentials edition and Windows Server 2012 Foundation edition remain unchanged.

Client Access License (CAL)

Windows Server Standard and Datacenter editions will continue to require Windows Server CALs for every user or device accessing a server. Some additional or advanced functionality will continue to require the purchase of an additive CAL. These are CALs that you need in addition to the Windows Server CAL to access functionality, such as Remote Desktop Services or Active Directory Rights Management Services. Note: You must have a Windows Server 2012 CAL to access an instance of Windows Server 2012.

Determining the number of licenses for Datacenter and Standard editions

Each license covers up to two physical processors on a single server. The minimum number of licenses required for each server is determined by the number of physical processors. For Standard edition you can add more virtual instances by assigning additional licenses to the server (two incremental virtual instances are added per license).

Licensing examples	Datacenter licenses required	Standard licenses required
One 1-processor, non- virtualized server	1	1
One 4-processor, non- virtualized server	2	2
One 2-processor server with three virtual OSEs	1	2
One 2-processor server with 12 virtual OSEs	1	6

Planning for Windows Server 2012

If you are planning to deploy Windows Server 2012, remember to select the edition of Windows Server 2012 based on your virtualization needs and edition features:

- Datacenter edition for highly-virtualized private clouds.
- Standard edition for lightly or non-virtualized environments.
- Essentials edition for small businesses with up to 25 users, running on servers with up to two processors.
- Foundation edition for small businesses with up to 15 users buying single processor servers from OEMs.
- Renewing Software Assurance is the best way to protect investments while gaining access to new versions, technical assistance and Deployment Planning Services.
- The Microsoft Enrollment for Core Infrastructure (ECI) will continue to offer the best value for private cloud and datacenter management pricing.
- Core CAL and Enterprise CAL Suites will continue to be the most cost effective way to purchase Windows Server CALs to access Windows Server 2012 Standard and Datacenter editions.

List of charts, tables, and figures