Microsoft

Networking

Windows Server 2012

Windows Server 2012

Table of Contents

Copyright information

Addressing the Challenges Facing IT Pros

Technology advances that are transforming the current business landscape are also affecting how IT pros do their jobs. These technologies include private and public cloud computing, mobile workforces, and assets dispersed across LANs, wide area networks (WANs), and datacenters.

Managing networking and network assets is a complex task that is similar to managing a power grid: the goal is to make sure smooth transmissions (data and electricity, respectively), that all components are connected properly, and that customers get the services they need when they need them.

In its ongoing communications with enterprise customers, Microsoft has identified three key networking issues where IT pros see the biggest challenges in the near future. These include: managing private clouds in an efficient, flexible way; easily and securely extending private clouds to leverage public cloud services; and easily connecting end users to IT resources, regardless of location. Let's take a brief look at each issue.

Managing private clouds efficiently

An important aspect of ensuring that private clouds deliver on their business potential is to make sure that they don't add to IT management burdens. For this to happen, a networking platform needs the following attributes:

- **Efficient automation** with plenty of options for automating essential network tasks to streamline private cloud deployment and management.
- **Flexibility** to scale private clouds up or down at will without being hampered by networking and storage resources and associated administrative personnel.
- **Straightforward consolidation capabilities** that allow, for example, more workloads to run on existing hardware resources, the convergence of data and storage network resources on the same physical infrastructure, and more options for using lower-cost storage arrays and switches.

Linking private clouds with public cloud services

To generate even greater benefits from private clouds, organizations need to find ways of augmenting private cloud functionality with public cloud services that are provided by the organization itself or by third parties. However, connecting private and public clouds needs to be accomplished easily and in a secure fashion. These features can help make that happen:

- **Efficient multitenancy:** Virtualization technology helps to consolidate resources through multitenant environments; it also provides IT departments with the means to implement greater security. This includes isolating tenants, line-of-business applications, and data, which makes it easier and friendlier to link private cloud tenants and resources with public cloud services.
- **Using public clouds to extend the datacenter:** With the correct technology in place, public clouds can be used to extend private cloud resources and processes such as data stores.
- **Seamless communication between private and public clouds:** Integrating the functionality of private and public clouds requires straightforward communications and connection of IT resources between the two environments.

Connecting users easily to IT resources

More and more organizations today are geographically dispersed and often have highly mobile workforces. In these scenarios, it is essential for users to connect quickly and easily to IT resources. These features are need for this to work:

- **All IT assets need to be available to end users, when they need them, regardless of location:** The physical location of the users and assets such as databases and private clouds should be irrelevant to a fast, straightforward connection.
- **All users need to be able to communicate with all clouds:** This can be restricted, of course, based on an organization's identity and security polices and individual or departmental requirements.

Windows Server 2012: Enhanced networking capabilities

Windows Server 2008 R2 introduced several new networking-related features, including Virtual Desktop Infrastructure (VDI) deployments, remote application (RemoteApp) publishing, peer-to-peer caching optimizations, improvements for optimizing network traffic performance, measures to avoid data duplication, and WAN optimizations.

Windows Server® 2012 builds on these advances with an array of new and enhanced features that help reduce networking complexity while lowering costs and simplifying management tasks. With Windows Server 2012, IT administrators have tools to automate and consolidate networking processes and resources. It provides the functionality to more easily connect private clouds with public cloud services. And Windows Server 2012 provides more, and better, tools for connecting users to IT resources and services across physical boundaries and private and public cloud environments.

This paper provides a closer look at the features of Windows Server 2012 that facilitate more efficient and cost-effective networking management.

Managing Continuously Available, Secure, and Efficient Private Clouds

Windows Server 2012 includes many new and enhanced features to simplify management, consolidate resources, and automate tasks. This helps to lower costs and increase flexibility for administrators dealing with changing and expanding network environments across geographically dispersed areas.

These features are discussed in this section:

- Reliability
  - NIC Teaming technology
  - Hyper-V Replica
  - SMB 3.0 Multichannel and SMB 3.0 Direct
  - Dynamic Host Configuration Protocol (DHCP) Failover
- Performance
  - Single Root I/O Virtualization (SR-IOV)
  - Dynamic Virtual Machine Queue
  - Accelerating network I/O, including receive segment coalescing and receive side scaling
  - IPsec task offload
- Manageability
  - IP Address Manager (IPAM)
  - Resource Metering
- Security
  - Domain Name System Security Extensions (DNSSEC)

Reliability

NIC Teaming technology

Windows Server 2012 lets you determine fault tolerance on your network adapters without buying additional hardware and software. One additional feature is NIC Teaming, which allows multiple network interfaces to work together as a team, preventing connectivity loss if one network adapter fails. It allows a server to tolerate network adapter and port failure up to the first switch segment. It also allows you to aggregate bandwidth from multiple network adapters. For example, four 1-gigabyte (Gb) network adapters can provide an aggregate of 4 Gbps of throughput.

The advantages of a built-in Windows teaming solution are that it works with all network adapter vendors, eliminates potential problems caused by proprietary solutions, provides a common set of management tools for all adapter types, and is fully supported by Microsoft.

Technical description

Today, all available NIC Teaming solutions have a similar architecture, as shown in the following figure.

Figure 1: Standard NIC Teaming solution architecture

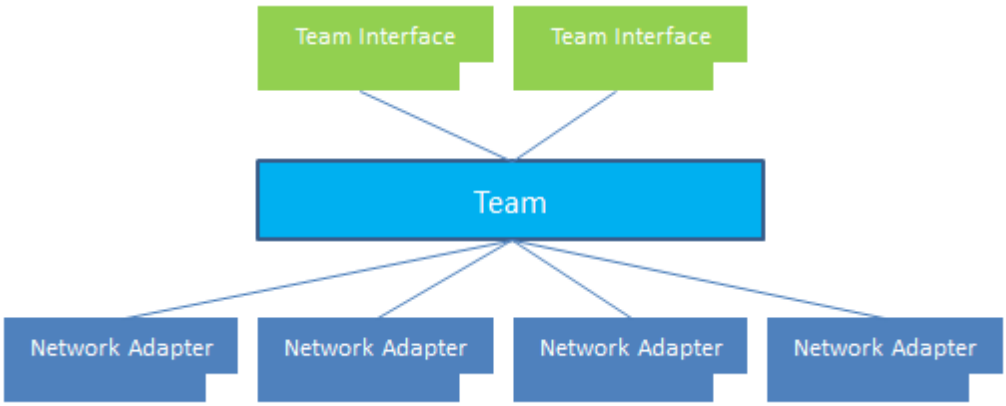

Teaming network adapters involves:

- **NIC Teaming configurations:** On or more physical network adapters connect to the NIC Teaming solution's multiplexing unit and present one or more "virtual adapters" (team network adapters) to the operating system.
- **Algorithms for traffic distribution:** Several different algorithms distribute outbound traffic between the network adapters. Traffic can be spread according to a hash of the source and destination addresses and ports, or it can be spread according to the originating port for virtual machines, as shown in the following figure. The NIC Teaming Solution can also divide traffic by virtual LAN (VLAN), so that applications can connect to different VLANs simultaneously.

Figure 2: NIC Teaming in a virtual machine configuration

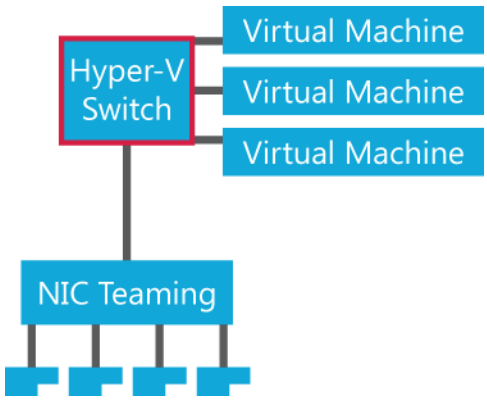

NIC Teaming configurations

NIC Teaming uses two sets of configuration algorithms:

- **Switch-independent modes:** These algorithms make it possible for team members to connect to different switches because the switch doesn't know that the interface is part of a team at the host. These modes don't require the switch to participate in the teaming.
- **Switch-dependent modes:** These algorithms require the switch to participate in the teaming. Here, all interfaces of the team are connected to the same switch.

There are two common choices for switch-dependent modes of NIC Teaming:

- **Generic or static teaming (IEEE 802.3ad draft v1):** This mode requires configuration on both the switch and host to identify which links form the team. Because this is a statically configured solution, there is no additional protocol to assist the switch and host to identify incorrectly plugged cables or other errors that could cause the team to fail to perform. Typically, this mode is supported by server-class switches.
- **Dynamic teaming (IEEE 802.1ax, Link Aggregation Control Protocol (LACP)):** This mode is also commonly referred to as IEEE 802.3ad because it was developed in the IEEE 802.3ad committee before being published as IEEE 802.1ax. It works by using the LACP to dynamically identify links that are connected between the host and a specific switch. Typical server-class switches support IEEE 802.1ax, but most require administration to enable LACP on the port. There are security challenges to allowing an almost completely dynamic IEEE 802.1ax to operate on a switch. As a result, switches today still require the switch administrator to configure the switch ports that are allowed to be members of such a team.

Either of these switch-dependent modes results in inbound and outbound traffic that approach the practical limits of the aggregated bandwidth. This is because the team's pool of links is seen as a single pipe.

Incompatibilities

NIC Teaming is compatible with all networking capabilities in Windows Server 2012, with these exceptions:

- SR-IOV
- Remote Direct Memory Access (RDMA)
- TCP Chimney Offload

For SR-IOV and RDMA, data is delivered directly to the network adapter without passing it through the networking stack. Therefore, the network adapter team can't see or redirect the data to another path in the team. In this release, TCP Chimney Offload isn't supported with NIC Teaming.

Requirements

The NIC Teaming feature requires:

- Windows Server 2012.
- At least one network adapter
  - If there are two or more network adapters they should be of the same speed.
  - Two or more network adapters are required if you are seeking bandwidth aggregation or failover protection.
  - One or more network adapters if you are only seeking VLAN segregation for the network stack.

Summary

You can use NIC Teaming to perform three important tasks:

- Aggregate bandwidth.
- Prevent connectivity loss in case of a network component failure.
- VLAN segregation of traffic from a host.

It works with both physical servers and virtual machines. You now have the option to use this feature and receive full support for your configuration from Microsoft, regardless of your network adapter vendor.

Windows 8 NIC Teaming is managed with Windows PowerShell and through the NIC Teaming configuration UI. Both the Windows PowerShell and the UI work for both physical and virtual servers. The UI can manage physical and virtual servers at the same time.

Hyper-V Replica

One of the primary reasons to set up a cloud, and for customers to move to the cloud, is to permit quicker and easier failure recovery. Despite increasing outages and the high cost per incident, IT administrators often postpone implementing solutions because of the complexity and cost. Organizations need an affordable and reliable business-continuity solution that helps them recover from a failure.

Virtualization opens new, effective possibilities to help make sure that services are optimally available. Windows Server 2012 introduces Hyper-V Replica, a feature that provides asynchronous replication of virtual machines for the purposes of business continuity and disaster recovery.

Technical description

Hyper-V Replica is a new feature in Windows Server 2012. It lets you replicate your Hyper-V virtual machines over a network link from one Hyper-V host at a primary site to another Hyper-V host at a replica site with little reliance on storage arrays or other software replication technologies. The following figure shows a more secure replication of virtual machines from different systems and clusters to a remote site over a WAN.

Figure 3: More securely replicating virtual machines from a wide range of systems and clusters to a remote site over a WAN

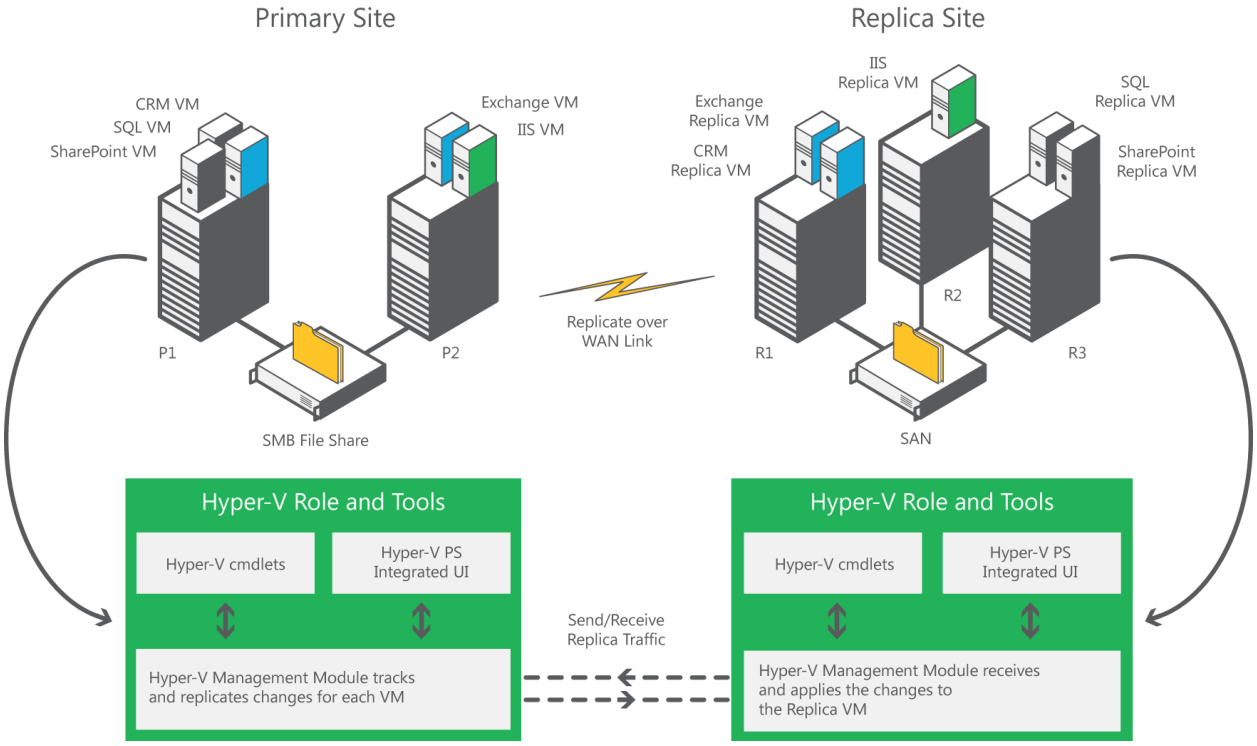

Hyper-V Replica tracks the write operations on the primary virtual machine and replicates these changes to the Replica server efficiently over a WAN. The network connection between the two servers uses the HTTP or HTTPS protocol and supports both Windows-integrated and certificate-based authentication. For an encrypted connection, you should choose certificate-based authentication. Hyper-V Replica is closely integrated with Windows Failover Clustering and provides easier replication across different migration scenarios in the primary and Replica servers.

To simplify management, Hyper-V Replica includes these tools:

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

Requirements

To use Hyper-V Replica, you need two physical computers configured with:

- Windows Server 2012.
- Hyper-V server role.
- Hardware that supports the Hyper-V role.
- Sufficient storage to host the files that virtualized workloads use. Additional storage on the Replica server, based on the replication configuration settings, may be needed.
- Sufficient network bandwidth among the locations that host the primary and Replica servers and sites.
- Firewall rules to permit replication between the primary and Replica servers and sites.
- Failover Clustering feature, if you want to use Hyper-V Replica on a clustered virtual machine.

Summary

You can use Hyper-V Replica to provide a virtual machine-level replication solution that efficiently replicates data over a network link to a remote site with little reliance on storage arrays or other software replication technologies. Hyper-V Replica provides a storage-agnostic and workload-agnostic solution that replicates more efficiently, periodically, and asynchronously over IP-based networks, typically to a remote site. It also lets you easily test the Replica virtual machine with little disruption to the ongoing replication. If a failure occurs at the primary site, you can restore business operations by bringing up the replicated virtual machine at the Replica site. Hyper-V Replica provides a virtual machine-level, more affordable, more reliable, and more manageable replication solution that is integrated with Hyper-V Manager and the Failover Clustering feature in Windows Server 2012.

Server Message Block 3.0

Server Message Block 3.0 in Windows Server 2012 provides support for server and storage that's continuously available, with minimal gaps in availability to customers. It has these advantages:

- Storage and systems are resilient to planned, and unplanned, downtimes.
- Data losses are greatly reduced.
- Lower acquisition and management costs.
- Includes safeguards against site failures and disasters.
- Files can be transferred faster, more reliably, and more efficiently.

Lower costs, simplify deployments, and increase availability

In previous versions of Microsoft Windows Server, Storage Area Networks (SANs) were expensive to acquire and operate, didn't allow for straightforward virtual machine or database mobility, and required vendor-specific tools and specialized knowledge to set up. Also, file share failovers were not transparent and required considerable intervention by an administrator to recover.

Customers expect storage to be continuously available, with little or minimal downtime. In previous versions of Windows Server, the time from a failed storage connection to the application failing was often 60 seconds or more. The goal of Windows Server 2012 is to limit this to 25 seconds or less.

Note

By default, SMB 3.0 is enabled in Windows Server 2012. This replaces SMB2, which was introduced with Windows Server 2008 R2 and Windows Vista, and it also replaces SMB2.1, which was introduced with Windows Server 2008 R2 and Windows 7.

Technical description

In SMB 3.0, 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 are used to store data on file shares.

SMB 3.0 for Windows Server 2012 includes these improvements:

- **SMB 3.0 Transparent Failover:** You can now more easily perform hardware or software maintenance of nodes in a file server cluster 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, SMB 3.0 Transparent Failover enables file shares to fail over to another cluster node without interrupting server applications that are storing data on these file shares.
- **SMB 3.0 Multichannel**: This allows aggregation of network bandwidth and network fault tolerance if multiple paths are available between the SMB 3.0 client and the SMB 3.0 server. Server applications can then take full advantage of all available network bandwidth and are resilient to a network failure.

  To use SMB 3.0 Multichannel, one computer should be configured as the file server (SMB 3.0 server) and the other as the file client (SMB 3.0 client). SMB 3.0 automatically detects and uses multiple network connections using any of the configurations illustrated in the following diagram:

Figure 4: SMB 3.0 Multichannel configuration

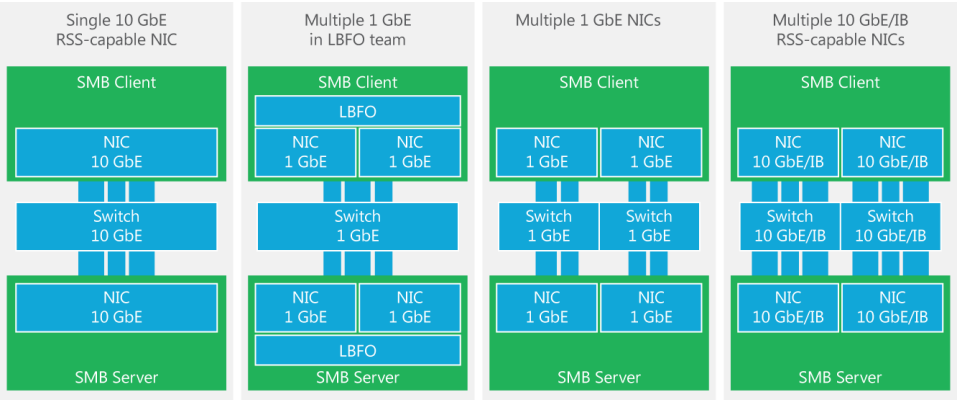

SMB 3.0 automatically detects and uses multiple network connections if these configurations are used:

- **Single 10– Gigabit Ethernet (GbE) network adapters:** Each of the two computers is configured with a single 10–GbE network interface.
- **Dual 1–GbE network adapters in a team:** Each of the two computers is configured with two 1–GbE network interfaces configured as a Load Balancing and Failover (LBFO) team. Each SMB 3.0 client network adapter talks to an SMB 3.0 server network adapter using its teamed interfaces.

- **Dual 1-GbE network adapters:** Each of the two computers is configured with two 1–GbE network interfaces. Each SMB 3.0 client network adapter talks to an SMB 3.0 server network adapter using a different subnet.
- **Dual 10–GbE network adapters:** Each of the two computers is configured with two 10–GbE network interfaces. Each SMB 3.0 client network adapter talks to an SMB 3.0 server network adapter using a different sub.
- **Dual Infiniband network adapters:** Each of the two computers is configured with two IB network interfaces. Each SMB 3.0 client network adapter talks to an SMB 3.0 server network adapter using a different subnet.
- **SMB 3.0 Direct.** This uses a special type of network adapter that has RDMA capability and can function at higher speeds with very low latency, while using very little CPU. For workloads such as Hyper-V or SQL Server, this allows a remote file server to have performance comparable to local storage.

  RDMA provides a more secure way to enable a DMA engine to transfer buffers between two machines across the network, as shown in the following figure.

Figure 5: SMB 3.0 Direct using a direct buffer transfer between two RDMA-capable network cards

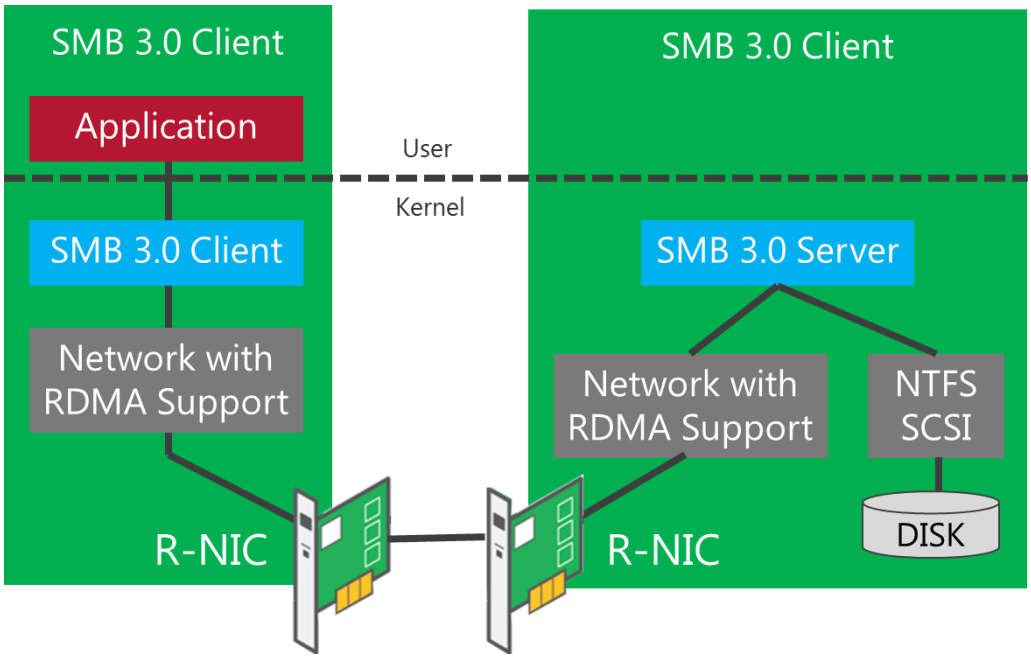

- **SMB 3.0 performance counters for server applications:** Performance counters provide detailed information about I/O size, I/O latency, IOPS, and so on. This enables an SQL Server database or Hyper-V administrator to analyze the performance of the SMB 3.0 file shares where their data is stored.
- **SMB 3.0 performance optimizations:** The SMB 3.0 client and SMB 3.0 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, the large maximum transmission unit (MTU) is turned on by default, which significantly enhances performance in large sequential transfers, such as SQL Server data warehouse, database backup or restore, and deploying or copying virtual hard disks.

- **SMB 3.0 management with Windows PowerShell:** With Windows PowerShell, you can manage SMB 3.0 on the file server, end-to-end, from the command line.
- **SMB 3.0 remote file storage:** Hyper-V can now store virtual machine files (including configuration, virtual hard disk files, and snapshots) in shared folders that use the SMB 3.0 protocol. Support for storing database files in shared folders that use the SMB protocol was introduced in SQL Server 2008 R2.
- **Volume Shadow Copy Service (VSS) for SMB 3.0 File Servers:** Disk snapshots are the primary means of nonintrusive data protection. It is nonintrusive because VSS snapshots take seconds and data I/O can be maintained without interruptions. Snapshots can be used for recovery, backup, and other purposes, as shown in the following diagram. In previous versions of Windows Server, VSS worked only on local disks. Windows Server 2012 allows VSS to be used on external file shares.

Figure 6: Diagram of VSS snapshots

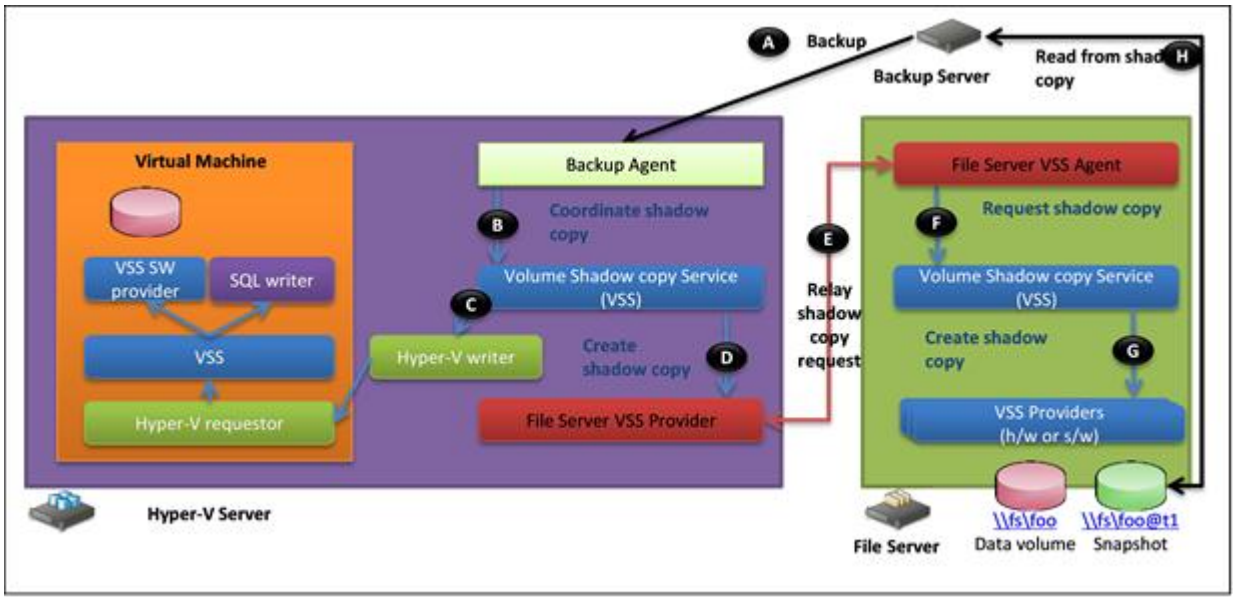

- **SMB 3.0 scale-out file servers:** Previous versions of Windows Server provided an active/passive solution. A file share, and its storage, could only be online on one cluster node at a time. To use multiple cluster nodes, customers often used multiple file server resource groups. Access to a particular share was provided only through one cluster node at a time, restricting bandwidth to the available bandwidth of the specific cluster node of the share. With Windows Server 2012, share data can be accessed through any cluster node.
- **SMB 3.0 storage availability and performance.** For virtual environments and SQL Server running on shared storage, availability of a connection to the storage and the performance of the network are critical components to overall application availability and performance. Because of availability and performance issues, Windows-based file storage has had limited adoption in these scenarios in the past. In Windows Server 2012, the SMB 3.0 layer automatically recognizes all network adapters with available network paths to the storage and routes traffic through all of them. This helps provides better performance through load balancing and against network adapter failures.

Requirements

SMB 3.0 Transparent Failover requires:

- **A failover cluster running Windows Server 2012 with at least two nodes:** The cluster must pass the cluster validation tests in the validation wizard.
- **"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.
- **Client computers running Windows 8 or Windows Server 2012:** Both computers must include the updated SMB 3.0 client that supports continuous availability.

Note

Down-level clients can connect to file shares that have the continuous availability property, but SMB 3.0 Transparent Failover is not supported for these clients.

SMB 3.0 Multichannel requires:

- **At least two computers running Windows Server 2012:** No extra features need to be installed because 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–GbE network interface.
  - **Dual 1–GbE network adapters:** Each computer must be configured with two 1–GbE network interfaces. Each SMB 3.0 client network adapter talks to an SMB 3.0 server network adapter 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 an LBFO team. Each SMB 3.0 client network adapter and an SMB 3.0 server network adapter communicate with each other by using their teamed interfaces.
  - **Dual 10-GbE network adapters:** Each computer must be configured with two 10–GbE network interfaces. Each SMB 3.0 client network adapter communicates with an SMB 3.0 server network adapter by using a different subnet.
  - **Dual Infiniband network adapters:** Each computer must be configured with two Infiniband network interfaces. Each SMB 3.0 client network adapter communicates with an SMB2 server network adapter by using a different subnet.

SMB 3.0 Direct requires:

- **At least two computers running Windows Server 2012:** No extra features need to be installed; the technology is available by default.
- **Network adapters with RDMA capability:** Currently, these network adapters come in three different types: iWARP, Infiniband, and RDMA over Converged Ethernet (RoCE).

Scenario

Fernando is an IT administrator for a large financial services company. His company has a large number of file server clusters in the datacenter to support numerous applications and files for thousands of users.

Fernando wants to provide continuous availability of applications and files for users when performing maintenance on nodes in file server clusters, or in the event of hardware failure on a cluster node. He also would like improved network and CPU performance and increased resiliency to network failures.

Fernando upgrades his servers to Windows Server 2012 to use SMB 3.0 support. SMB 3.0 Transparent Failover allows him to move file shares between cluster nodes without interrupting server applications storing data on those shares. The SMB 3.0 Multichannel feature allows his applications to take full advantage of all available network bandwidth while adding resiliency to potential network failures.

These and other enhancements meet his needs for better network and CPU performance, along with continuous availability, convenience, and lower cost.

Summary

The SMB 3.0 layer automatically recognizes all network adapters with available network paths to the storage, and routes traffic through all of them. This provides better performance through load balancing and also delivers insurance against network adapter failures.

With SMB 3.0 Transparent Failover, you can move file shares between cluster nodes without interrupting server applications storing data on those shares. The SMB 3.0 Multichannel feature allows your applications to take full advantage of all available network bandwidth while adding resiliency to potential network failures.

These and other enhancements provide you with better network and CPU performance, along with continuous availability, convenience, and lower cost.

DHCP Server Failover

Windows Server 2012 supports the DHCP Failover protocol as described in the Internet Engineering Task Force (IETF) Internet draft. Through this protocol, the DHCP Server Failover feature allows two DHCP servers to synchronize lease information almost instantaneously and to provide high availability of the DHCP service. If one server becomes unavailable, the other server assumes responsibility for servicing clients for the same subnet. You can also configure failover with load-balancing, with client requests distributed between the two DHCP servers.

Technical description

DHCP failover in Windows Server 2012 provides support for two DHCPv4 servers.

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

Hot standby mode

In this mode, two servers operate in a failover relationship in which an active server leases IP addresses and configuration information to all clients in a scope or subnet. The two DHCP servers in a failover relationship don't need to be on the same subnet. DHCP service can be provided to the subnet by DHCP Relay. The secondary server assumes this responsibility if the primary server becomes unavailable. A server is primary or secondary in the context of a subnet, so that a server that is primary for one subnet could be secondary for another.

Hot standby mode of operation is recommended for deployments in which a central office or datacenter server acts as a standby backup server to a server at a remote site that is local to the DHCP clients. In this hub-and-spoke deployment, shown in the following figure, it is undesirable to have a remote standby server service any clients unless the local DHCP server becomes unavailable.

Figure 7: Hot standby DHCP failover in a hub-and-spoke deployment

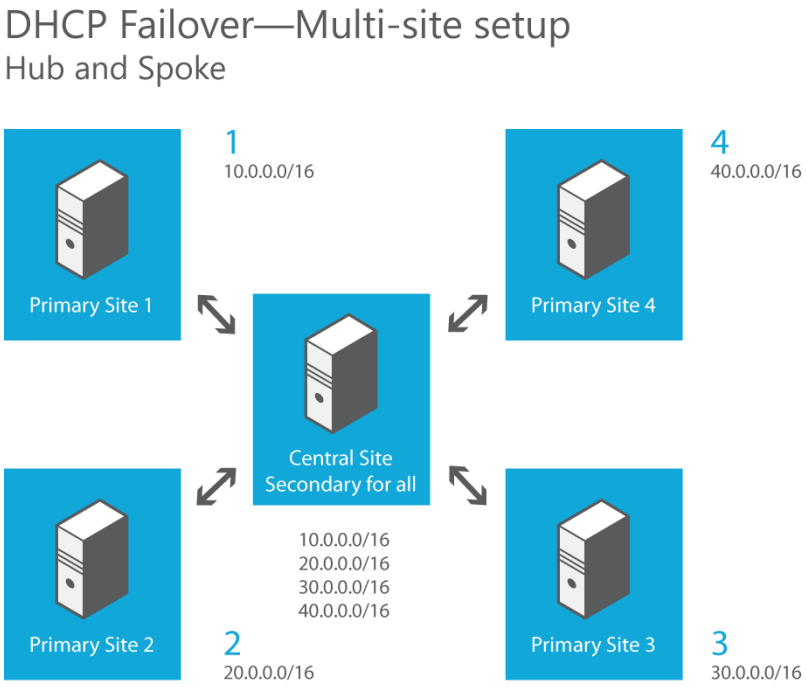

Load-sharing mode

In load sharing mode, which is the default mode, the two servers simultaneously serve IP addresses and options to clients on a particular subnet. The two DHCP servers in a failover relationship don't need to be on the same subnet) The client requests are load-balanced and shared between the two servers. This mode is recommended for deployments in which both servers in a failover relationship are located at the same physical site, as shown in the following two figures. Both servers respond to DHCP client requests based on the load distribution ratio configured by the administrator.

Figure 8: Load-sharing DHCP failover in a single site with a single subnet

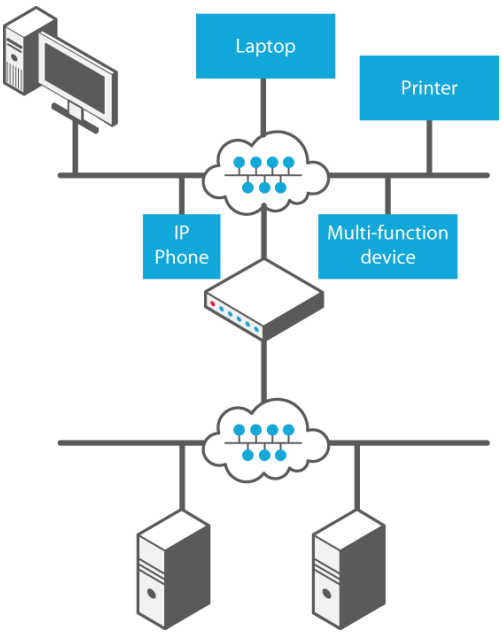

Figure 9: Load-sharing DHCP failover in a single site with multiple subnets

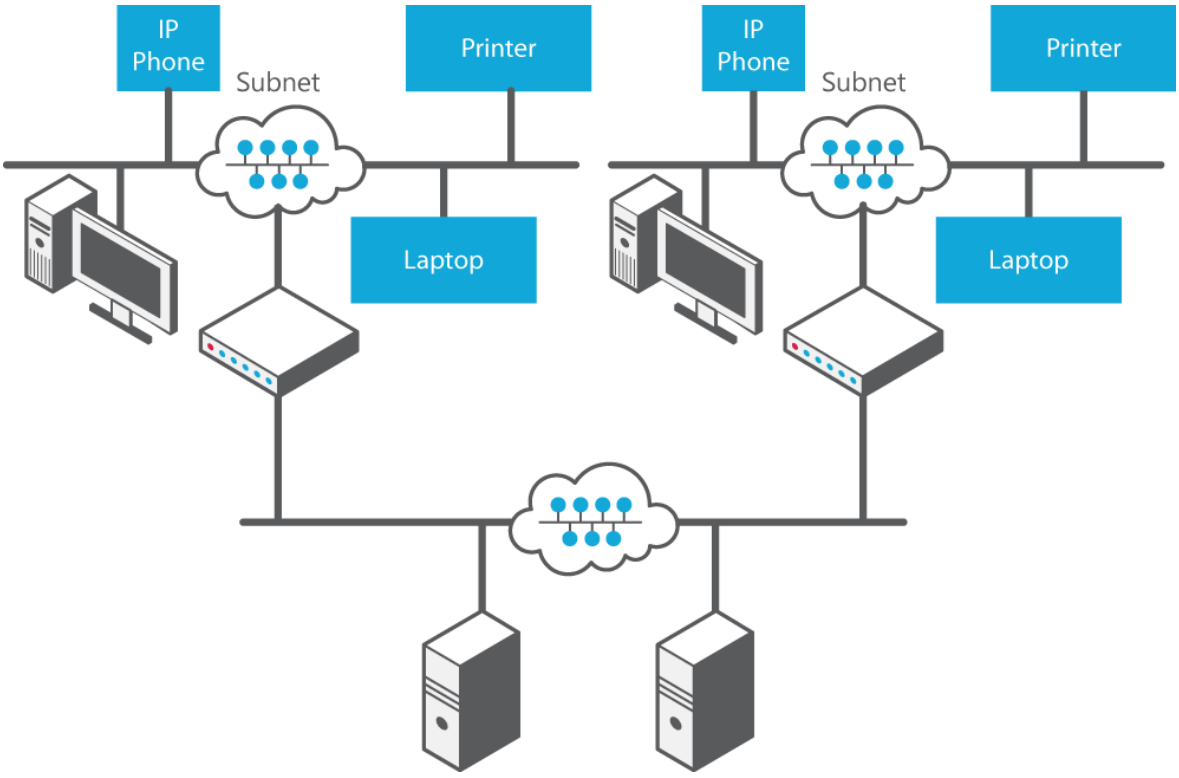

Requirements

DHCP failover requires:

- **Software:** DHCP failover requires two DHCP servers running Windows Server 2012.
- **Number of servers:** DHCP failover is deployed between two DHCP servers participating in the failover relationship.
- **Time synchronization:** For DHCP failover to function correctly, time must be kept synchronized between the two servers in a failover relationship. Time synchronization can be maintained by deploying the Network Time Protocol (NTP) or any alternative mechanism. Running the Failover Configuration Wizard compares the current time on the two servers. If the time difference is greater than 1 minute, the failover setup process halts with a critical error that instructs the administrator to synchronize the time.

Scenario

Grant is the network administrator at a midsized manufacturing company. His company has a centralized corporate office with multiple subnets and dozens of client computers.

Grant needs reliability for network services and access. As part of his plan to make sure there's high availability of all resources, Grant deploys two DHCP servers running Windows Server 2012 with DHCP Failover in a load-sharing relationship. Client lease requests are load balanced between the two servers, and each acts as a backup for the other in case of a hardware failure.

Summary

With Windows Server 2012, you can deploy DHCP Failover in a load-sharing relationship. Client lease requests will be load-balanced between the servers, and each will act as a backup for the other in case of a hardware failure.

Windows Server 2012 offers a DHCP failover solution that allows straightforward deployment of a highly available DHCP service with low capital expenditure and low maintenance overhead.

Performance

Support for SR-IOV networking devices

For virtual networking, a primary goal is near native I/O. Windows Server 2012 adds the ability to assign SR-IOV functionality from physical devices directly to virtual machines. This gives virtual machines the ability to bypass the software-based Hyper-V Virtual Switch, and directly address the network adapter. As a result, CPU overhead and latency is reduced, with a corresponding rise in throughput.

Technical description

The SR-IOV standard was introduced by the PCI-SIG, the special interest group that owns and manages PCI specifications as open industry standards. SR-IOV works in conjunction with system chipset support for virtualization technologies that provide remapping of interrupts and Direct Memory Access (DMA) and lets SR-IOV-capable devices be assigned directly to a virtual machine. Hyper-V in Windows Server 2012 enables support for SR-IOV–capable network devices and lets an SR-IOV virtual function of a physical network adapter be assigned directly to a virtual machine. This increases network throughput and reduces network latency while also reducing the host CPU overhead that is required for processing network traffic. The following figure shows the architecture of SR-IOV support in Hyper-V.

Figure 10: SR-IOV support in Hyper-V

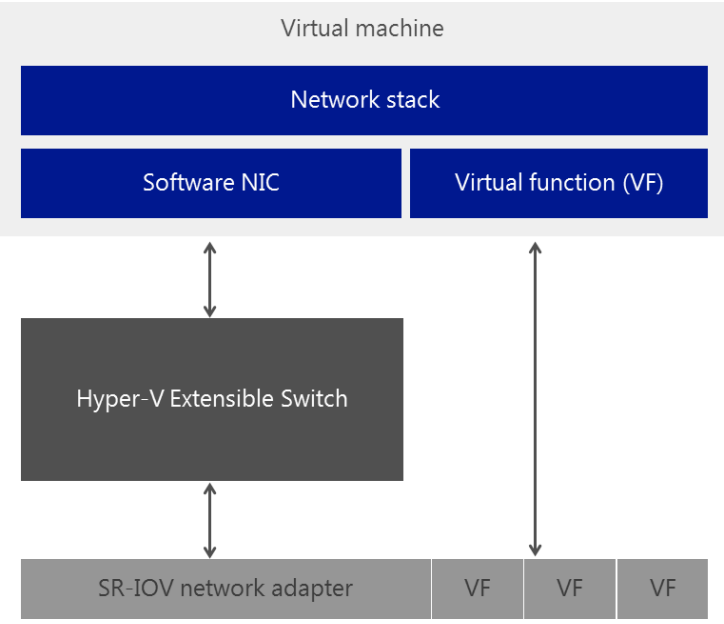

This configuration increases network throughput and reduces network latency while also reducing the host CPU overhead required to process network traffic.

Requirements

SR-IOV networking requires:

- A host system that supports SR-IOV, such as Intel VT-d2, including chipset support for interrupt and DMA remapping and correct 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 the management operating system (which runs the Hyper-V role) and each virtual machine where a virtual function is assigned.

SR-IOV causes the virtual machine's traffic to bypass the Hyper-V Virtual Switch. If any switch port policies are set, SRIOV functionality is revoked for that virtual machine.

Summary

Support for SR-IOV–capable systems and network devices in Windows Server 2012 allows you to assign SR-IOV–capable network adapters directly to a virtual machine. This helps maximize your network throughput, while minimizing network latency and the CPU overhead required to process network traffic.

Dynamic Virtual Machine Queue (D-VMQ)

Virtual Machine Queue (VMQ) allows the host's network adapter to pass DMA packets directly into individual virtual machine memory stacks. Each virtual machine device buffer is assigned a VMQ, which avoids needless packet copies and route lookups in the virtual switch. Essentially, VMQ allows the host's single network adapter to appear as multiple network adapters to the virtual machines, allowing each virtual machine its own dedicated network adapter. The result is less data in the host's buffers and an overall performance improvement to I/O operations.

Technical Description

Windows Server 2008 R2: Offload routing and filtering of network packets to the network adapter (enabled by hardware-based receive queues) to reduce host overhead.

Figure 11: Network I/O path without VMQ

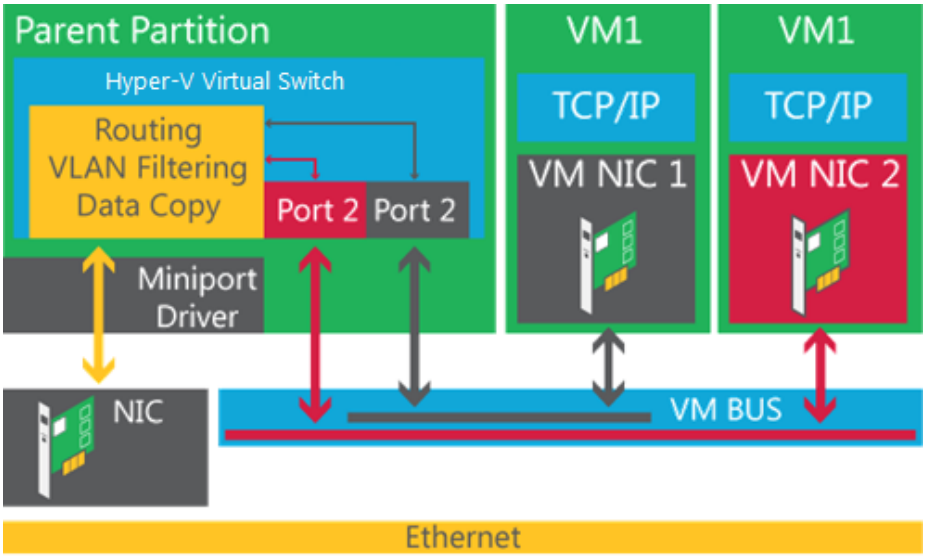

The VMQ is a hardware virtualization technology for the efficient transfer of network traffic to a virtualized host operating system. A VMQ-capable network adaptor classifies incoming frames to be routed to a receive queue based on filters which associate the queue with a virtual machine's virtual network adaptor. These hardware queues may be affinitized to different CPUs thus allowing for receive scaling on a per-virtual machine network adaptor basis. Windows Server 2008 R2 allowed administrators to statically configure the number of processors available to process interrupts for VMQ. Without VMQ, CPU 0 would run hot with increased network traffic. With VMQ, the interrupts were spread across more processors. However, network load may vary over time. A fixed number of processors may not be suitable in all traffic regimes.

New in Windows Server 2012: Windows Server 2012 dynamically distributes incoming network traffic processing to host processors, based on processor use and network load. In times of heavy network load, Dynamic VMQ (D-VMQ) automatically uses more processors. In times of light network load, Dynamic VMQ relinquishes those same processors.

Figure 12: Network I/O path with VMQ

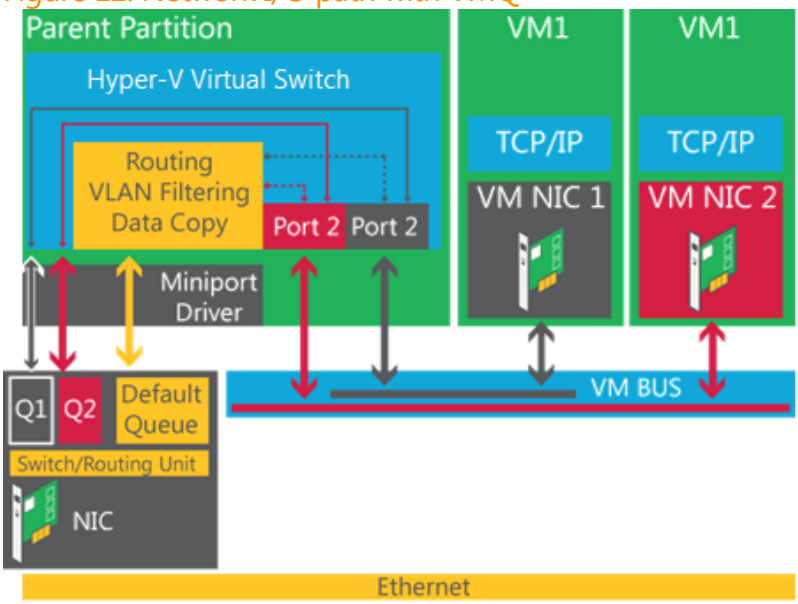

Requirements

Dynamic VMQ requires hardware network adaptors and drivers that support Network Device Interface Specification (NDIS) 6.30.

Summary

VMQ spreads interrupts for network traffic across available processors. In Windows Server 2012, the Dynamic VMQ capability allows an adaptive algorithm to modify the CPU affinity of queues without removing or re-creating queues. This results in a better match of network load to processor use, resulting in increased network performance.

Accelerating Network I/O

Windows Server 2012 helps meet the needs of the most demanding enterprises with lower costs for infrastructure and operations. It provides the necessary higher scalability and improved performance to support low latency applications in both physical and virtual environments.

Technical description

Windows Server 2012 automatically optimizes performance for critical applications by providing these benefits:

- **Predictability:** Responds in a predictable manner to changing I/O.
- **Scalability:** Easily reduces latency to handle applications with a high IOPS.
- **Absolute low latency:** Decreases the amount of end-to-end transaction processing necessary for critical applications.

Windows Server 2012 also enhances network performance with receive-side scaling (RSS) and Receive Segment Coalescing (RSC) improvements.

Receive-side scaling

RSS spreads monitoring interrupts over multiple processors, so a single processor isn't required to handle all I/O interrupts, which was common with earlier versions of Windows Server. Active load balancing between the processors tracks the load on the different CPUs and then transfers the interrupts as needed.

You can select which processors will be used for handling RSS requests, including processors that are beyond 64 KB, which allows you to take advantage of very high-end computers that have a large number of logical processors.

The following figure illustrates RSS on a network adapter with eight RSS queues.

Figure 13: RSS with four nodes and eight queues

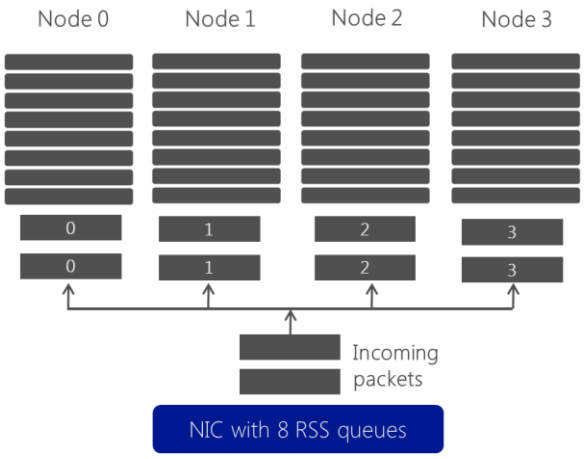

RSS works with inbox NIC Teaming or LBFO to remove a limitation in earlier versions of Windows Server where a choice had to be made between using hardware drivers or RSS. RSS will also work for User Datagram Protocol (UDP) traffic and can manage and debug applications that use WMI and Windows PowerShell.

Receive Segment Coalescing

RSC improves the scalability of the servers by reducing the overhead for processing a large amount of network I/O traffic. It accomplishes this by coalescing multiple inbound packets into a large buffer, as shown in the following diagram.

Figure 14: Receive Segment Coalescing

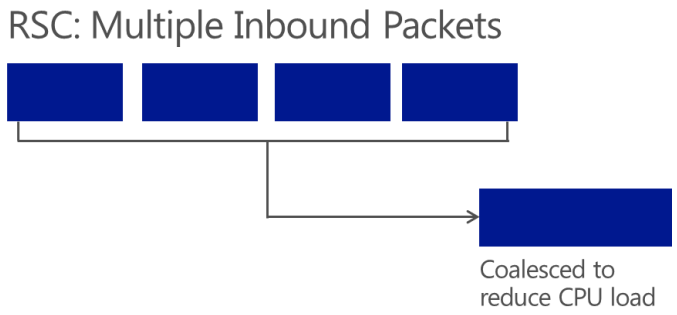

In early testing, RSC reduced CPU use up to 20 percent.

Requirements

The Accelerating Network I/O feature requires:

- Windows Server 2012.
- Network adapters that support the features described earlier.

Summary

By using Windows Server 2012, your servers can take advantage of enhancements for reducing latency through higher I/O per second, reduced end-to-end transaction processing, active load balancing between processors, and many other RSS improvements. These enhancements help meet the needs of the most advanced applications, which means better performance for users accessing servers and applications in physical or virtualized environments, over local or wide-area networks.

IPsec Task Offload for Virtual Machines

IPsec protects network communication by authenticating and encrypting some or all of the contents of network packets. IPsec Task Offload in Windows Server 2012 leverages the hardware capabilities of server network adaptors to offload IPsec processing. This reduces the CPU overhead of IPsec encryption and decryption significantly.

In Windows Server 2012, IPsec Task Offload is extended to virtual machines as well. Customers using virtual machines who want to protect their network traffic with IPsec can take advantage of the IPsec hardware offload capability available in server network adaptors, which frees up CPU cycles to perform more application-level work and leaves the per packet encryption/decryption to hardware.

Technical Description

IPsec processing is computationally expensive relative to network communication without IPsec. In general, encryption and decryption in software consumes CPU cycles. Hardware can efficiently perform per packet encryption and decryption. The IPsec information necessary to perform per packet operations is contained in a Security Association (SA). An SA is required for each IPsec connection. A pair of SAs is required for each connection, a transmit SA and a receive SA. In Windows Server 2008 R2 IPsec Task Offload version 2 (IPsecTOv2) was enhanced to support intelligent offloading of SAs used for per packet encryption/decryption in a hardware network adaptor. Using the network adaptor to perform the per packet cryptographic functions results in lower CPU use and higher network throughput.

Requirements

IPsec Task Offload requires the support of hardware network adaptors and drivers with this capability in NDIS 6.30.

Summary

Windows Server 2012 extends IPsec Task offload to virtual machines. Hosting providers who want to offer more secure hosted cloud environments to address security and compliance concerns, and customers using virtual machines who want to protect their network traffic with IPsec can take advantage of the IPsec hardware offload capability available in server network adaptors. Windows Server 2012 provides a better performing IPsec solution for virtual machines.

Manageability

Resource Metering in Hyper-V

Your computing resources are limited. You need to know how different workloads draw upon these resources, even when they're virtual. In Windows Server 2012, Hyper-V introduces Resource Metering, a technology that helps you track historical data of the use of virtual machines. With Resource Metering, you gain insight into the resource use of specific servers. You can use this data to perform capacity planning, to monitor consumption by different business units or customers, and to capture data needed to help redistribute the costs of running a workload. You could also use the information this feature provides to help create a billing solution, so customers of your hosting services could be charged proportionately for their resource use.

Metrics for resource use

Windows Server 2012 offers two ways to obtain historical data on the customer's use of virtual machine resources: Hyper-V cmdlets in Windows PowerShell and the new APIs in the Virtualization WMI Provider.

Hyper-V exposes the metrics in the following table for resource use.

Table 1: Metrics exposed by

Metric	Units	Description
Average CPU use	Megahertz (MHz)	The average amount of CPU used by a virtual machine over a period of time.
Average memory use	Megabytes (MB)	The average amount of physical memory used by a virtual machine over a period of time.
Minimum memory use	MB	The lowest amount of physical memory assigned to a virtual machine over a period of time.
Maximum memory use	MB	The highest amount of physical memory assigned to a virtual machine over a period of time.
Maximum disk allocation	MB	The highest amount of disk space capacity allocated to a virtual machine over a period of time.
Incoming network traffic	MB	Total incoming network traffic, for a virtual network adapter over a period of time.
Outgoing network traffic	MB	Total outgoing network traffic for a virtual network adapter over a period of time.

Use of Network Metering Port access control lists (ACLs)

Enterprises pay for the Internet traffic in and out of their datacenters but not 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.

Metering virtual machine use in a multitenant environment

Hyper-V in Windows Server 2012 lets providers build a multitenant environment, in which virtual machines can be served to multiple clients in a more isolated and secure way, as shown in the following figure. Because a single client may have many virtual machines, aggregation of resource use data can be a challenging task. However, Windows Server 2012 simplifies this task by using resource pools, a feature available in Hyper-V. Resource pools are logical containers that collect resources of the virtual machines that belong to one client, permitting single-point querying of the client's overall resource use.

The following figure is an example of Resource Metering in a two-tenant environment that is built with Hyper-V in Windows Server 2012.

Figure 15: A two-tenant environment built with Hyper-V in Windows Server 2012

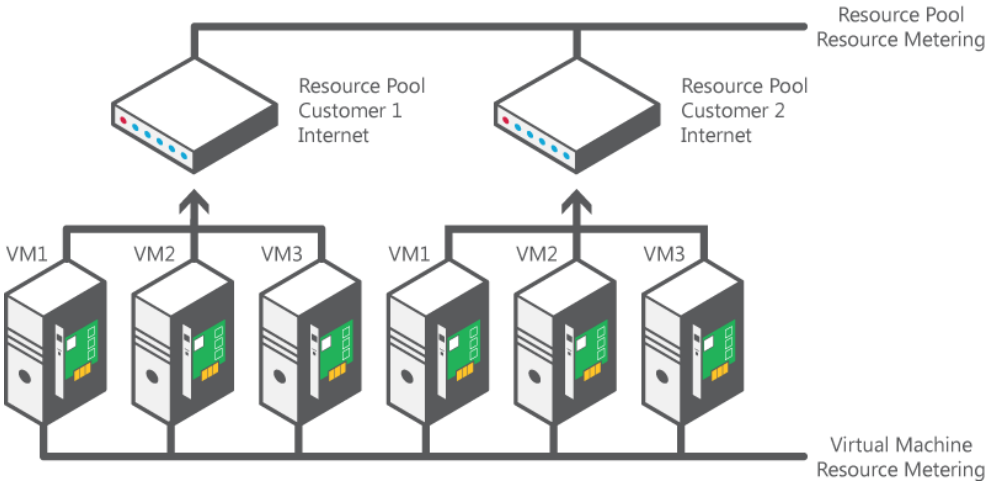

The following figure shows a basic model of Resource Metering.

Figure 16: Basic model of metering resource use

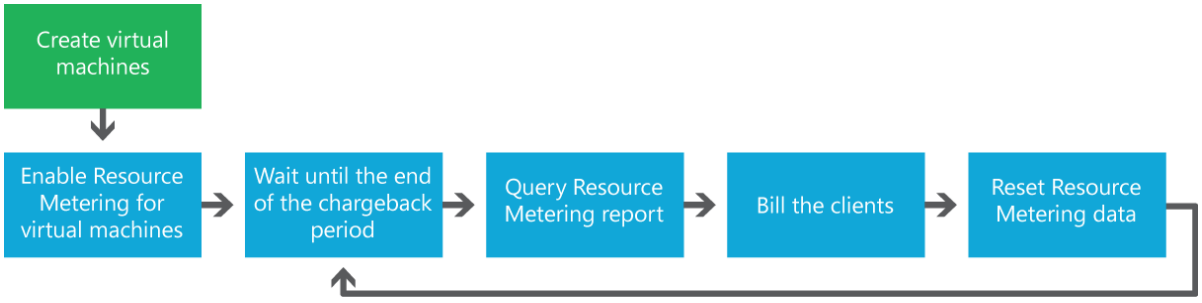

In this model, a hosting provider:

1. Creates virtual machines for a customer and enables Resource Metering once for the virtual machines. In a multitenant environment, the provider would enable metering on each resource pool. Hyper-V then tracks resource use for each virtual machine until that virtual machine is deleted.
2. Queries resource use data at the end of each chargeback period, and uses the data to bill clients as needed.
3. Resets the data at the end of each chargeback period, so that Hyper-V can begin tracking resource use for the new chargeback period.

Resource Metering works with all Hyper-V operations. Movement of virtual machines between Hyper-V hosts (such as through live, offline, or storage migration) does not affect the collected data.

Requirements

To use the Resource Metering feature You need:

- Windows Server 2012
- Hyper-V server role

Resource Metering is not supported for:

- Storage accessed through a virtual Fibre Channel adapter.
- Physical disks attached directly to a virtual machine.
- Network adapters configured with OffloadWeight.

Note

Network offload weight helps make sure that limited hardware resources are dynamically assigned to the right virtual machines. As virtual machines move around inside the datacenter, the network offload weight is used to prioritize which virtual machines obtain access to network hardware offloads that have finite limits (such as SR-IOV).

Summary

The Resource Metering feature in Windows Server 2012 Hyper-V makes it easier for you to track historical data about each customer's use of virtual machines. Through resource pools, which are part of this technology, Hyper-V lets providers aggregate use data in a multitenant environment in which different customers or business units may have many virtual machines. With this feature, you can perform capacity planning or monitor resource consumption by various business units or customers. Third-party independent software vendors (ISVs) can use data that this feature provides to build more reliable, end-to-end chargeback solutions.

IP Address Management (IPAM)

Windows Server 2012 introduces IPAM, a framework for discovering, monitoring, auditing, and managing the IP address space and the associated infrastructure servers on a corporate network. IPAM provides:

- Automatic IP address infrastructure discovery.
- Migration of IP address data from spreadsheets or other tools.
- Custom IP address space display, reporting, and management.
- Audit of server configuration changes and tracking of IP address use.
- Monitoring and specific scenario-based management of DHCP and Domain Name System (DNS) services.

The other salient aspects and features of IPAM are:

- Agentless architecture.
- Distributed deployment.
- Remote management (the IPAM console can be remote and manage any instance of IPAM server in the network).
- Infrastructure servers of Windows Server 2008, Windows Server 2008 R2, Windows Server 2008 R2 SP1, and Windows Server 2012.
- Backup and restore and disaster discovery.
- Migration.
- Active Directory integration.
- In-box availability.

Technical description

IPAM gives you a choice of two architectures:

- **Distributed:** An IPAM server deployed at all site in an enterprise. This mode of deployment is largely preferred to reduce network latency in managing infrastructure servers from a centralized IPAM server.
- **Centralized:** One IPAM server in an enterprise. This is deployed even in the case of the distributed mode. Administrators have one single console to visualize, monitor, and manage the complete IP address space of the network and the associated infrastructure servers.

An example of the distributed IPAM deployment method is shown in the following figure, with one IPAM server located at the corporate headquarters and others 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 17: An example distributed IPAM architecture

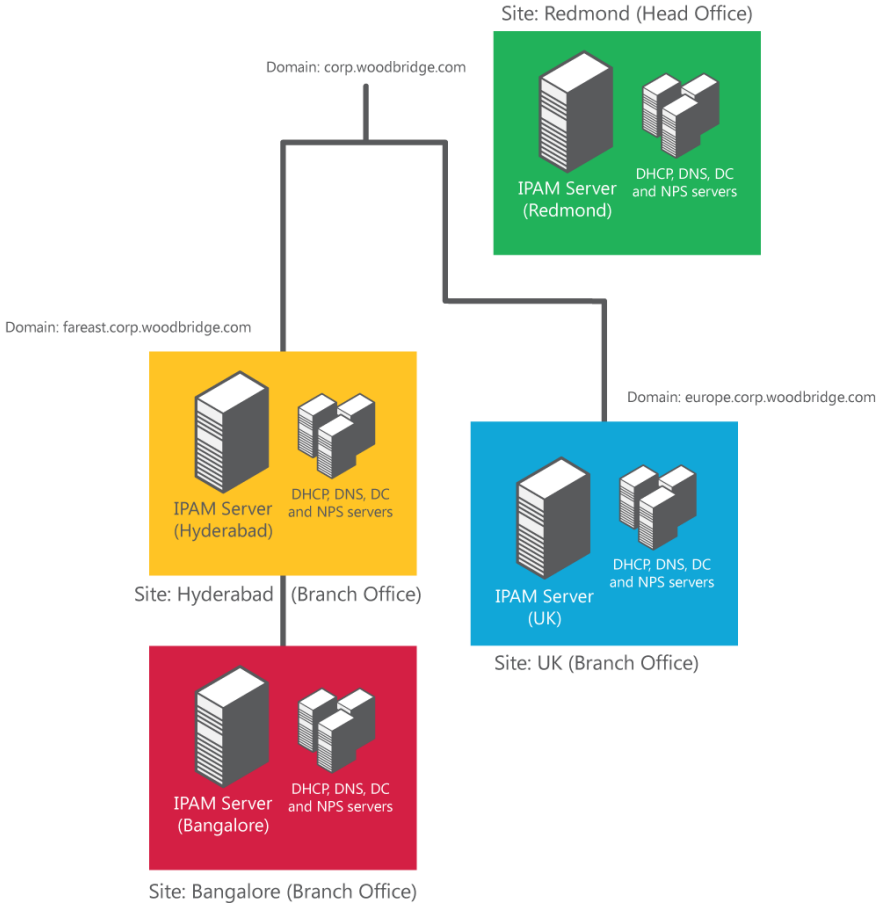

IPAM monitoring

IPAM periodically attempts to locate the domain controller, DNS, and DHCP servers on the network that are within the scope of discovery specified and allow manual addition of Network Policy Server (NPS). You must choose whether these servers are managed by IPAM or are unmanaged. To be managed by IPAM, server security settings and firewall ports must be configured to allow the IPAM server access to perform required monitoring and configuration functions. You can manually configure these settings or use Group Policy objects (GPOs) to automatically configure them. If you choose the automatic method, settings are applied when a server is marked as managed, and settings are removed when it is marked as unmanaged.

The IPAM server communicates with managed servers by using a remote procedure call (RPC) or WMI interface, as shown in the following figure. 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 using IPAM. Zone status monitoring and a limited set of configuration functions are also available for DNS servers.

Figure 18: IPAM server communications

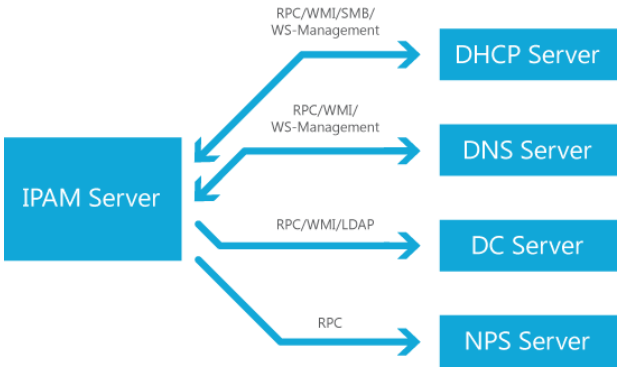

Server discovery

IPAM supports Active Directory–based auto-discovery of DNS and DHCP servers on the network. Discovery is based on the domains and server roles selected during configuration of the scope of discovery. IPAM discovers the domain controller, DNS servers, and DHCP servers in the network and confirms their availability based on role-specific protocol transactions. In addition to automatic discovery, IPAM also supports manual addition of a server to the list of servers in the IPAM system.

Figure 19: IPAM Server Discovery view

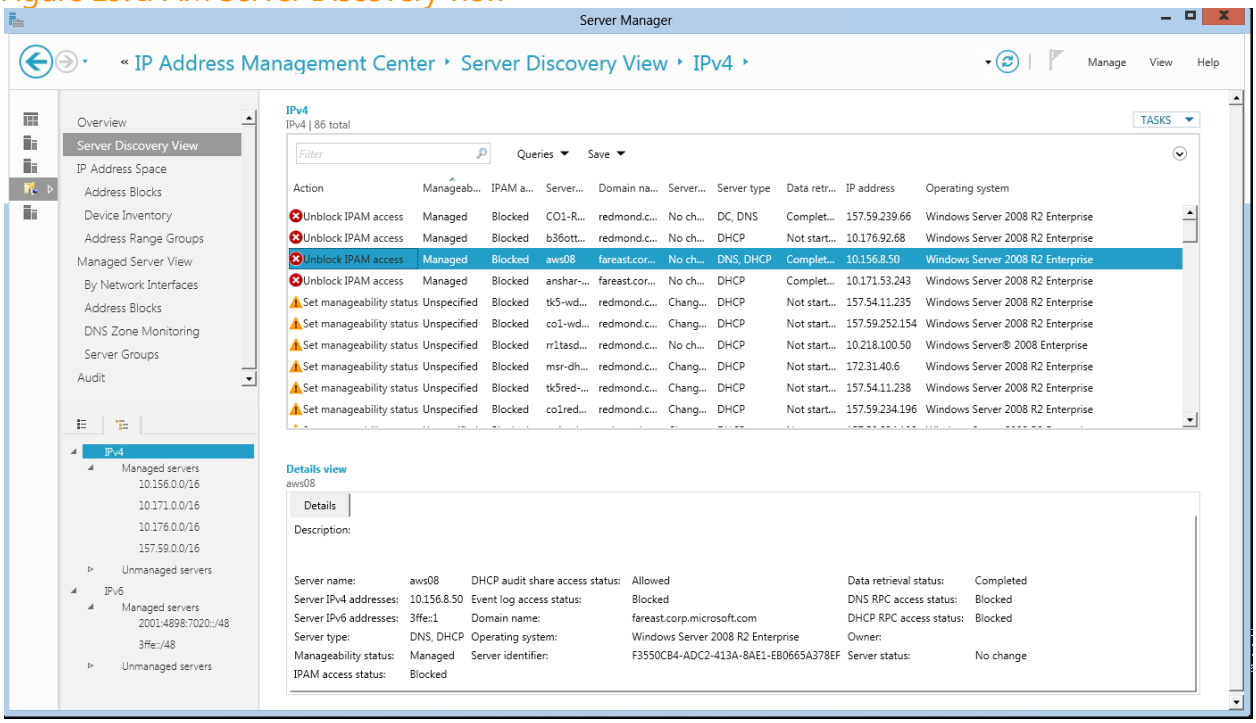

Managed servers

Configuring the manageability status of a server as “Managed” indicates that it is part of the IPAM server’s managed environment. 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

Configuring the manageability status of a server as Unmanaged indicates that the server is considered to be outside the IPAM server's managed environment. IPAM doesn't collect data from these servers.

IPAM security groups

IPAM creates security groups with unique permission settings that let you control the access level of specific users and groups. These are the available IPAM security groups:

- **IPAM users:** Members have read-only permission and can use all views in the IPAM interface, with the exception of the IP address tracking view.
- **IPAM MSM administrators:** Members of the multiserver management (MSM) administrators group have read and write permissions to manage infrastructure servers (DHCP) and other common IPAM 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 permission to manage IP address space and other IPAM common tasks. Network administrators are typically members of this group.
- **IPAM IP Tracking administrators:** This is a special group with permission to view IP address tracking data on the network. You can use this group to protect privacy information that might be contained in IP address tracking data.
- **IPAM administrators:** Members can 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:

- **Server Discovery:** Automatically discovers domain controllers, DHCP servers, and DNS servers in the selected domains.
- **Server Configuration:** Collects configuration information from DHCP and DNS servers for display in IP address space and server management functions.
- **Address Utilization:** Collects IP address space use data from DHCP servers for display of current and historical use.
- **Event Collection:** Collects DHCP and IPAM server operational events. Also collects 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.

You can use Task Scheduler to modify these tasks.

Requirements

This feature requires:

- Windows Server 2012 (on the server running IPAM).
- Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012 (on the domain controllers, DHCP servers, DNS servers, and Network Policy Server [NPS] servers).
- A single Active Directory forest.
- A domain member computer to act as the IPAM server. You can't install the IPAM feature on an Active Directory domain controller.

Summary

IPAM in Windows Server 2012 reduces the time and expense needed to manage IP address space. For example, by using the centralized architecture feature in IPAM, you can more easily obtain updated information about IP address ranges used in branch offices, and the particular addresses assigned to infrastructure servers. This strategy lets IT personnel track how IP addresses are assigned throughout your organization with a minimum of administrative effort.

Security

DNSSEC

DNSSEC is a suite of additions to DNS that helps protect DNS traffic from attack. By validating a digital signature attached to each DNS response, the resolver verifies the authenticity of DNS data, even from an untrusted DNS server. Specifically, DNSSEC provides origin authority, data integrity, and authenticated denial of existence.

Windows Server 2012 extends and simplifies your implementation of DNSSEC by providing:

- Support for the latest standards.
- Integration with Active Directory.
- Simple deployment.
- A "sign and forget" operation experience.

Technical Description

Windows Server 2012 adds several new features to the management and implementation of DNSSEC.

Support for the latest standards

Windows Server 2012 supports the latest DNSSEC standards, including:

- Next Secure3 (NSEC3) signing per RFC 5155, including support for authenticated denial of existence.
- RSA/SHA-2 per RFC 5702.
- Automated Trust Anchor Rollover per RFC 5011.

Simple deployment

DNSSEC deployment is a phased process that begins with signing DNS zones. After the zones are signed and the DNS system is stable, validation of the DNS responses is enabled on caching resolvers. Finally, Windows 7 and Windows 8 clients have a set of options for establishing trust of the responses from the caching resolver.

Windows Server 2012 allows signing and management of DNSSEC through the standard DNS management tools including DNS Manager and Windows PowerShell.

Figure 20: Signing a zone with DNS Manager

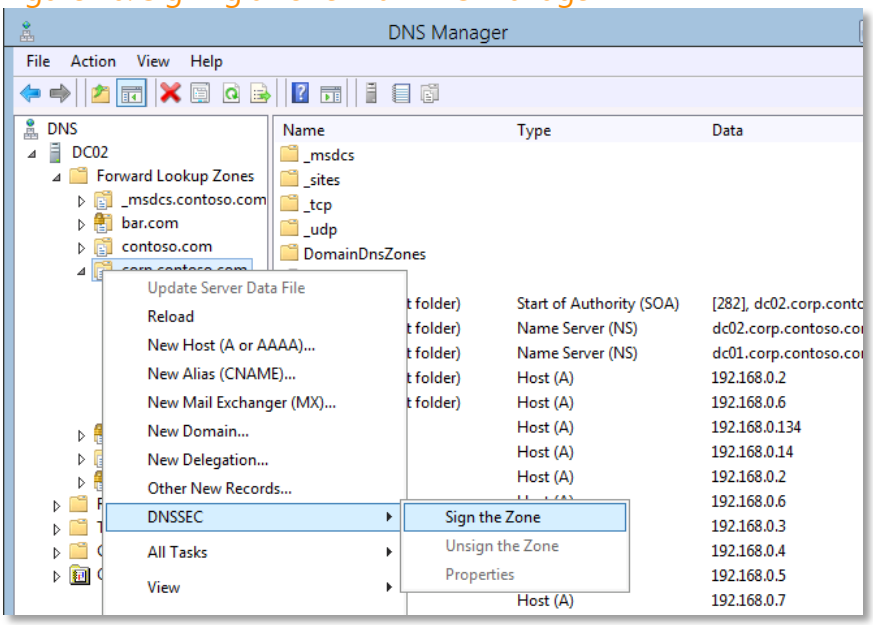

When you sign a zone using DNS Manager, a wizard appears that includes a "recommended settings" option, enabling one-step signing of a zone.

Figure 21: Configuring DNSSEC signing for a zone

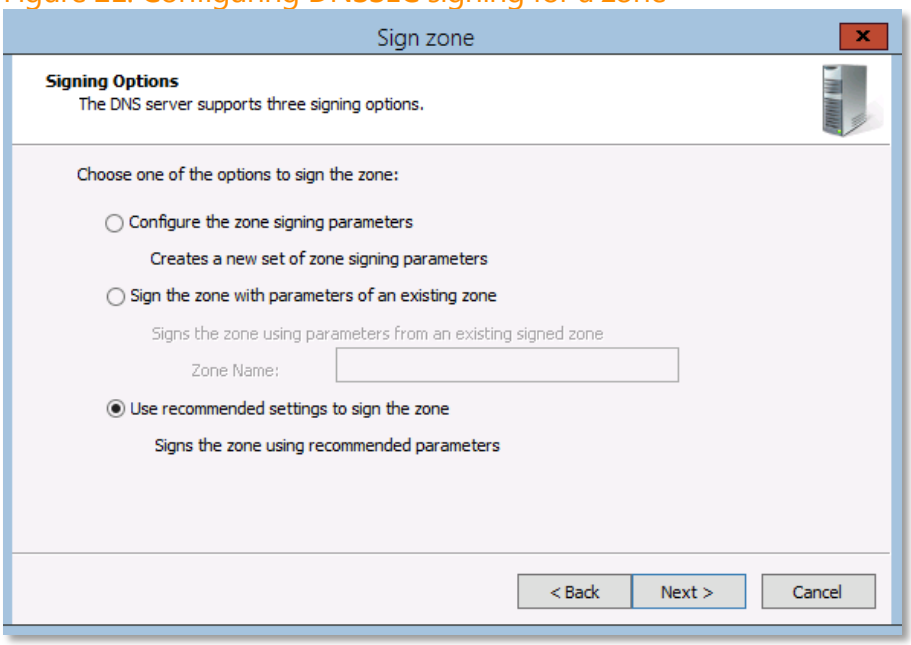

Active Directory Integration

Windows Server 2012 includes added support for Active Directory integration with DNSSEC. This includes:

- Key generation and replication.
- DNS dynamic updates in DNSSEC signed zones.
- Automated trust anchor distribution through Active Directory.

Figure 22: Diagram of Active Directory Integration features

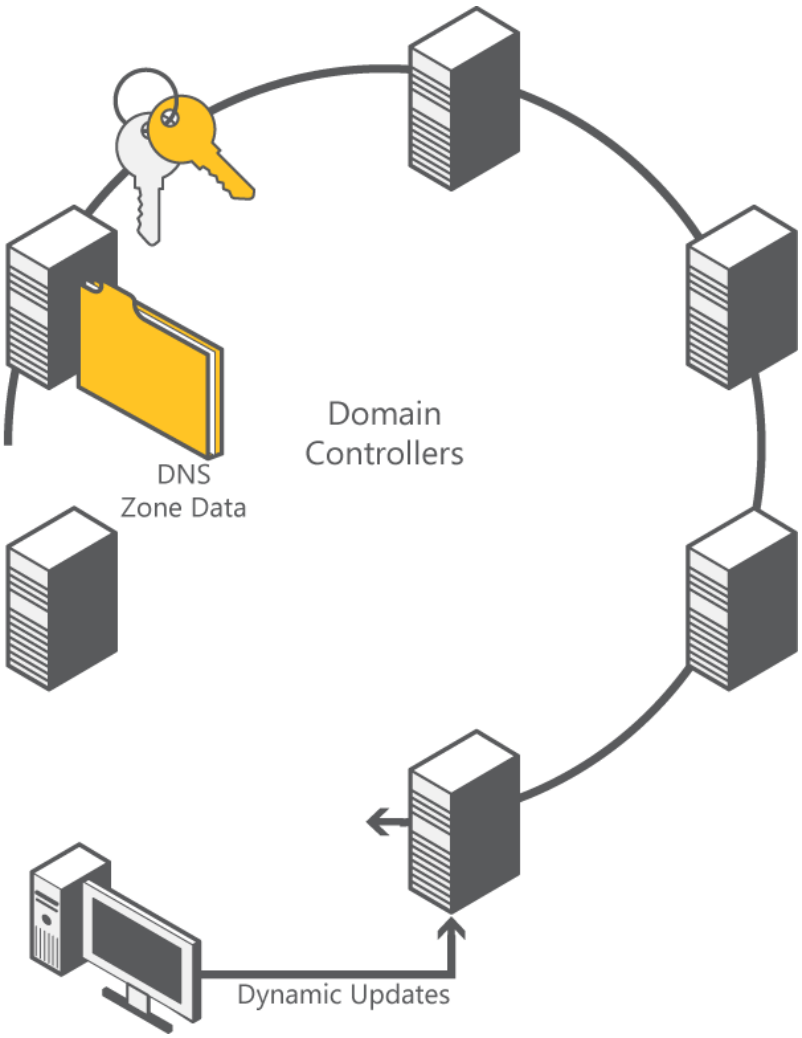

The following sections discuss some key aspects of the DNSSEC support in Windows Server 2012, which enable the signing of DNS zones in Active Directory environments and validation of DNS responses at the caching resolver.

Key operations

In a multimaster DNS deployment, which is typical in enterprise Active Directory environments, each master server is identical to all the other master servers. However, for DNSSEC, a single master server performs key generation and key management for a given zone. Windows Server 2012 introduces the concept of a key master for DNSSEC. Any authoritative DNS server that hosts a primary copy of the zone can be designated as the key master, and any server can be the key master for multiple zones if it hosts a primary copy of the zone.

An administrator selects the key master to be the DNS server in charge of key management for a particular zone. Different DNS servers can also function as key masters for different zones. The key master is in the context of a given 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 key master for that zone. If desired, the key master role can be moved to a different DNS server, as long as it also hosts a primary copy of the zone. DNSSEC is configured on the key master and is

responsible for key generation and key management for the zone. The key master tracks signature expiration and automatically refreshes signatures that are about to expire, ensuring continuous availability of zone data. When the key master updates keys, Active Directory automatically replicates these updates to all other domain controllers.

The key master is responsible for distribution of private keys and zone-signing information. Only the key master can perform any operations with or on the key signing key (KSK). The key master 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 key master must be online and highly available for uninterrupted service for key signing operations.

Trust anchors

In cryptographic terms, a trust anchor is an authoritative entity represented by a public key and associated data. In DNSSEC terms, a trust anchor is a configured DNSKEY resource record or a DNS 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. A trust anchor 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 hosting the zone is a domain controller, you can decide to distribute trust anchors automatically to other DNS servers in the domain. This includes file-backed zones hosted on domain controllers. If the key master is a stand-alone server that is not integrated with Active Directory, for example, a member server, this option isn't available.

Dynamic zone signing

With Windows Server 2012, the DNS server can automatically perform post-signing tasks that could only be performed manually in earlier versions of Windows. This significantly reduces the cost of ownership of a DNSSEC-signed zone.

The following post-signing tasks are handled automatically by Windows Server 2012:

- **Automatic signing of zone data:** After keys have been replicated to a DNS server through Active Directory, zone data is automatically signed on each DNS server. Updates are sent and received in an unsigned state through Active Directory replication. After it receives the unsigned update, each authoritative server generates the necessary DNSSEC signatures to secure the updated zone data and then adds it to its own copy of the zone.
- **Background zone signing:** A DNSSEC-signed zone remains online and available for queries or updates at all times during the zone-signing process.
- **Dynamic updates:** These updates can be enabled on a signed zone, and any DNS server that is authoritative for the zone can accept dynamic updates.
- **Scavenging:** With Windows Server 2012, scavenging of stale records in a signed zone works exactly as it does in an unsigned zone.

Key rollovers

DNSSEC keys don't have a permanent lifetime, and require periodic replacement. The longer a key is in use, the greater the risk of compromise. The key replacement process, or key rollover, is a vital part of operating a DNSSEC-signed zone.

Windows Server 2012 DNS provides support for automated key rollover management, including provisioning and configuring the supported methods for ZSK and KSK rollover and the actual process of performing a key rollover.

Automated updating of DNSSEC trust anchors

To validate the DNSSEC-protected data, DNSSEC-aware resolvers must have knowledge of the trust anchor for that zone. DNS in Windows Server 2012 provides for DNSSEC-aware resolvers to poll for trust anchor changes and automatically retain their copy of the trust anchor updated per RFC 5011.

Requirements

This feature requires:

- **Online zone signing:** An authoritative DNS server running Windows Server 2012.
- **Trust anchor distribution:** An Active Directory–integrated authoritative DNS server running Windows Server 2012 is required for automatic trust anchor distribution. Manual trust anchor distribution on DNS servers that are not authoritative and not integrated with Active Directory also requires Windows Server 2012 if the zone is signed with NSEC3 or uses a signing algorithm other than RSA/SHA-1.
- **Hosting signed zones:** If a zone is signed with NSEC3 or uses a signing algorithm other than RSA/SHA-1, authoritative DNS servers must be running Windows Server 2012. In Windows Server 2012, a read-only domain controller (RODC) can host a secondary copy of a signed zone.
- **Validation of signed zones:** If a zone is signed with NSEC3 or uses a signing algorithm other than RSA/SHA-1, resolving DNS servers must be running Windows Server 2012 to validate DNSSEC-protected responses.
- **Client computers:** Windows 7 or Windows 8.

Summary

Windows Server 2012 simplifies configuration and management of DNSSEC by adding support for online signing and automated key management. These added features make it easier for you to protect your internal DNS infrastructure and your external DNS communication, while reducing management overhead and lowering total cost of ownership.

Linking Private Clouds with Public Cloud Services

Windows Server 2012 includes these new or expanded features to provide more security, convenience, and adaptability in multitenant environments, or when linking private and public clouds:

- Hyper-V Network Virtualization, including Generic Routing Encapsulation (GRE) and IP Address Rewrite.
- Hyper-V Extensible Switch, including Port ACLs, private virtual LAN (PVLAN), ARP/ND spoofing protection, DHCP Guard, Trunk mode to virtual machines, and monitoring.
- Quality of Service (QoS).
- RDP WAN optimizations.
- WebSocket protocol.
- Server Name Indicator (SNI).

Hyper-V Network Virtualization

Isolating virtual machines of different departments or customers can be a challenge on a shared network. When departments or customers need to isolate complete networks of virtual machines, the challenge becomes even greater. Traditionally, VLANs are used to isolate networks, but they can become very complex to manage on a large scale. Hyper-V Network Virtualization solves this problem. With it, you can isolate network traffic from different business units or customers on a shared infrastructure without using VLANs. Hyper-V Network Virtualization also lets you move virtual machines as needed within your virtual infrastructure while preserving their virtual network assignments. Finally, you can use Hyper-V Network Virtualization to transparently integrate these private networks into a preexisting infrastructure on another site.

With Hyper-V Network Virtualization, it's straightforward to place any virtual machine on any node, regardless of its IP address. This allows multiple virtual machines to have their own IP addresses, which makes the transition to private clouds easier for customers with existing on-premises virtual machines or physical servers.

Windows Server 2012 is a more complete virtualization platform, extending beyond the server and into the network. Windows Server 2012 offers these benefits:

- **Flexibility:** Move virtual machines between host computers, networks, and datacenters with few scheduling downtime, affecting users, or reconfiguring the server. Create private networks by using your existing network infrastructure, and rearrange them for maximum efficiency without affecting services. Connect virtual machines from around the world as if they were in the same rack. Create pools of storage from disks, and distribute the storage anywhere in your network.
- **Stability:** Achieve the efficiency of cloud computing while you support service level agreements (SLAs) and the QoS of dedicated server and network hardware. Eliminate data loss and downtime by continuously backing up and replicating virtual machines, across your datacenter or around the world.

- **Power:** With support for 320 logical processors in a server and 64 virtual processors per virtual machine, you can reach higher levels of performance and consolidate more servers onto fewer hosts.
- **Manageability:** Log on to any service, on your local network or in the cloud, with your Active Directory credentials. Move your domain controllers into the cloud. Monitor and bill network use for individual virtual machines.

New and enhanced functionality

Windows Server 2012 adds significant functionality to Hyper-V Network Virtualization to enhance performance, while simplifying configuration and management. These are some of the new features and benefits:

- **Onboarding workloads to a shared infrastructure as a service (IaaS) is easier:** Customers can keep their internal IP addresses while they move workloads onto shared IaaS clouds, minimizing configuration changes needed for IP addresses, DNS names, security policies, and virtual machine configurations.
- **Server and network administration are decoupled;** 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.
- **Tenant isolation no longer depends on VLANs:** In software-defined, policy-based datacenter networks, network traffic isolation no longer depends on VLANs; it's enforced within Hyper-V hosts based on multitenant isolation policy. Network administrators can use VLANs for traffic management of the physical infrastructure if the topology is primarily static.
- **Flexible workload placement and cross-subnet live migration are possible:** Services and workloads can be placed on or migrated to any server in the datacenter while keeping their IP addresses, without being limited to physical IP subnet hierarchy or VLAN configurations.
- **The network is simplified and server/network resource use is improved**: The rigidity of VLANs and dependency of virtual machine placement on physical network infrastructure result in overprovisioning and underuse. The increased flexibility of virtual machine workload placement helps simplify network management and improves server and network resource use.
- **Network Virtualization is compatible with existing infrastructure and emerging technology:** Hyper-V Network Virtualization can be deployed in today's datacenter, and is compatible with emerging datacenter "flat network" technologies, such as Transparent Interconnection of Lots of Links (TRILL), an IETF standard, architecture intended to expand Ethernet topologies.
- **Standards-based approach is used:** Hyper-V Network Virtualization uses existing standard IP protocol and GRE header formats.
- **Permits use of Windows PowerShell/ WMI:** Use Windows PowerShell to easily script and automate administrative tasks.

Technical description

Hyper-V Network Virtualization extends the concept of server virtualization to allow multiple virtual networks, potentially with overlapping IP addresses, to be deployed on the same physical network. With Hyper-V Network Virtualization, you can set policies that isolate traffic in your dedicated virtual network independently of the physical infrastructure. The following figure illustrates how you can use Hyper-V Network Virtualization to isolate network traffic belonging to two different customers. In the figure, Blue and Red virtual machines are hosted on a single physical network, or even on the same physical server.

However, because they belong to separate Blue and Red virtual networks, the virtual machines can't communicate with each other, even if the customers assigns them IP addresses from the same address space.

Figure 23: Using Hyper-V Network Virtualization to isolate network traffic belonging to two different customers

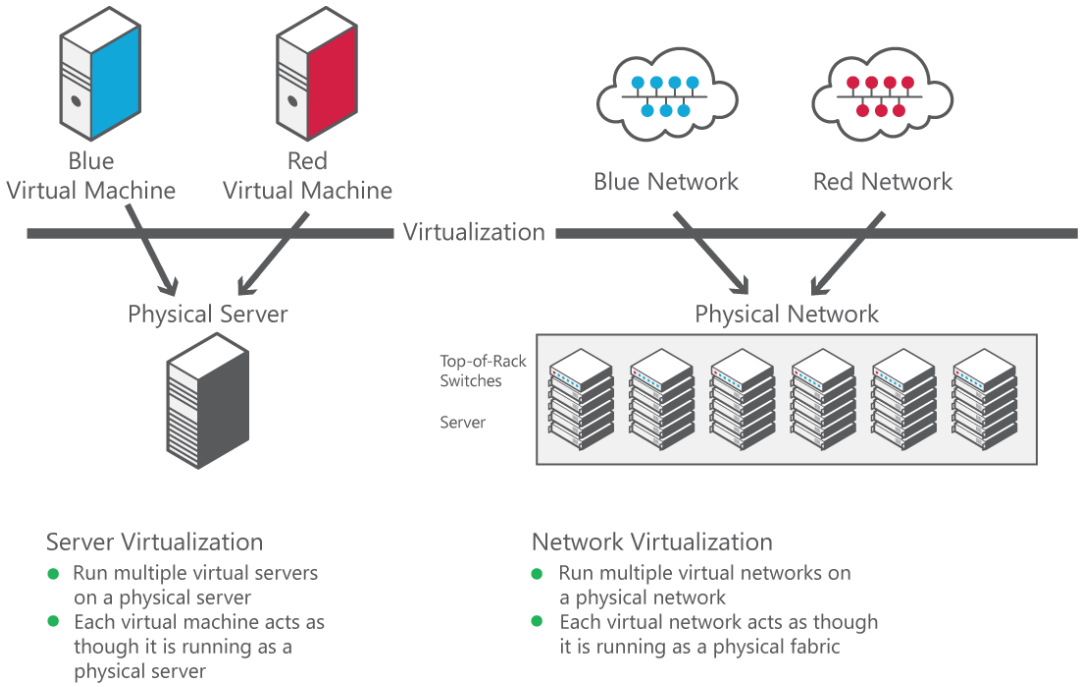

To virtualize the network with Hyper-V Network Virtualization, each virtual machine is assigned two IP addresses:

- **Customer address:** The IP address that customers assign based on their own intranet infrastructure. It allows customers 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 customers.
- **Provider address:** The IP address that hosts assign based on their physical network infrastructure. The provider address appears in the packets on the wire exchanged with the server running Hyper-V hosting the virtual machine. It is visible on the physical network, but it is not visible to the virtual machine.

The layer of customer addresses is consistent with customers' network topology, which is virtualized and decoupled from the underlying physical network addresses, as implemented by the layer of provider addresses. With Hyper-V Network Virtualization, any virtual machine workload can execute unmodified on any Windows Server 2012 server running Hyper-V within any physical subnet if the servers Hyper –V have the appropriate policy settings that can map between the two addresses. This approach has many benefits, including cross-subnet live migration, customer virtual machines running IPv4 while the host provider is running an IPv6 datacenter or vice versa, and using IP address ranges that overlap between customers. The biggest advantage of having separate customer and provider addresses is that it allows customers to move their virtual machines to the cloud without reconfiguring them in any way.

In the following figure, Blue Corp and Red Corp are both moving infrastructure to a hosted cloud environment. The figure shows the virtual and physical infrastructure environment.

Figure 24: Physical and virtual networking infrastructure for Blue Corp and Red Corp

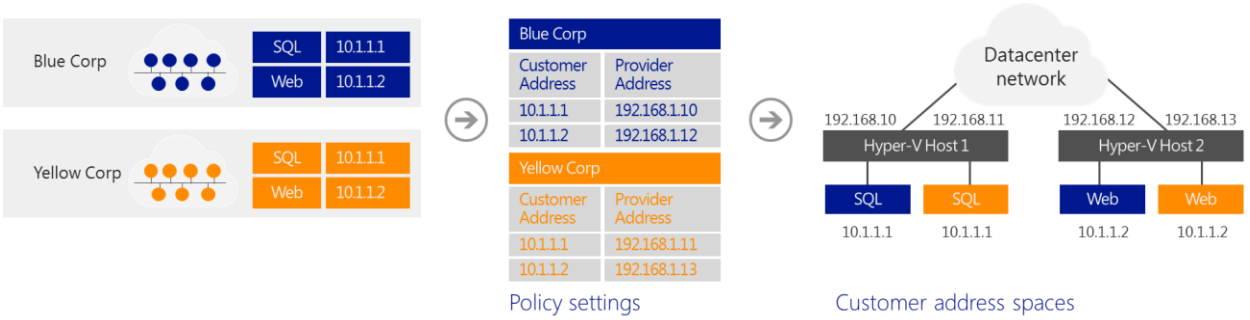

Before moving to the hosting provider's shared cloud service:

- Blue Corp ran a SQL Server instance (named SQL) at the IP address 10.1.1.1 and a web server (named WEB) at the IP address 10.1.1.2, which uses its SQL Server for database transactions.
- Red Corp ran a SQL Server instance (also named SQL) and assigned the IP address 10.1.1.1, and a web server (also named WEB) at the IP address 10.1.1.2, which uses its SQL Server for database transactions.

Both Blue Corp and Red Corp move their respective SQL Servers and WEB servers to the same hosting provider's shared IaaS 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):

- Customer addresses of Blue Corp's virtual machines: SQL is 10.1.1.1; WEB is 10.1.1.2.
- Customer addresses of Red Corp's virtual machines: SQL is 10.1.1.1; WEB is 10.1.1.2.

Both companies are assigned the following provider addresses by their hosting provider when the virtual machines are provisioned:

- Provider addresses of Blue Corp's virtual machines: SQL is 192.168.1.10; WEB is 192.168.1.12.
- Provider addresses of Red Corp's virtual machines: SQL is 192.168.1.11; WEB is 192.168.1.13.

The hosting provider creates policy settings, consisting 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 both 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:

- Hyper-V Host 2, based on its policy settings, translates the addresses in the packet from the following:
  - Source: 10.1.1.2 (customer address of Blue Corp WEB).
  - Destination: 10.1.1.1 (customer address of Blue Corp SQL).

  It translates the preceding addresses to the following:
  - Source: 192.168.1.12 (provider address for Blue Corp WEB).
  - Destination: 192.168.1.10 (provider address for Blue Corp SQL).
- When the packet is received at Hyper-V Host 1, based on its policy settings, Hyper-V Network Virtualization translates the addresses in the packet from the following:
  - Source: 192.168.1.12 (provider address for Blue Corp WEB).

- Destination: 192.168.1.10 (provider address for Blue Corp SQL).

Before Hyper-V delivers the packet to the Blue Corp SQL virtual machine, it translates the preceding addresses back to the following:

- Source: 10.1.1.2 (customer address of Blue Corp WEB).
- Destination: 10.1.1.1 (customer address of Blue Corp SQL).

When the Blue Corp SQL virtual machine on Hyper-V Host 1 responds to the query, the following happens:

- Hyper-V Host 1, based on its policy settings, translates the addresses in the packet from the following:
  - Source: 10.1.1.1 (customer address of Blue Corp SQL).
  - Destination: 10.1.1.2 (customer address of Blue Corp WEB).

  It translates the preceding addresses to the following:
  - Source: 192.168.1.10 (provider address for Blue Corp SQL).
  - Destination: 192.168.1.12 (provider address for Blue Corp WEB).
- When the packet is received at Hyper-V Host 2, based on its policy settings, Hyper-V Network Virtualization translates the addresses in the packet from the following:
  - Source: 192.168.1.10 (provider address for Blue Corp SQL).
  - Destination: 192.168.1.12 (provider address for Blue Corp WEB).

  Before it delivers the packet to the Blue Corp WEB virtual machine, it translates the preceding addresses to the following:
  - Source: 10.1.1.1 (customer address of Blue Corp SQL).
  - Destination: 10.1.1.2 (customer address of Blue Corp WEB).

A similar process for traffic between the Red Corp WEB and SQL virtual machines uses the settings in the Red Corp isolation group. With Hyper-V Network Virtualization, Red Corp and Blue Corp virtual machines interact in the same way as if they were on their original intranets, but they are never in communication with each other, even though they're using the same IP addresses. These two sets of servers are isolated from each other by the separate customer and provider addresses, the policy settings of the Hyper-V hosts, and the address translation between customer and provider addresses for inbound and outbound virtual machine traffic.

Setting and maintaining the Hyper-V Network Virtualization capabilities requires the use of a policy management server, which can be integrated into tools that manage virtual machines.

IP Address Rewrite

IP Address Rewrite modifies the customer IP addresses of the virtual machine's packets before they're transferred on the physical network, as shown in the following diagram.

Figure 25: Example of IP Address Rewrite

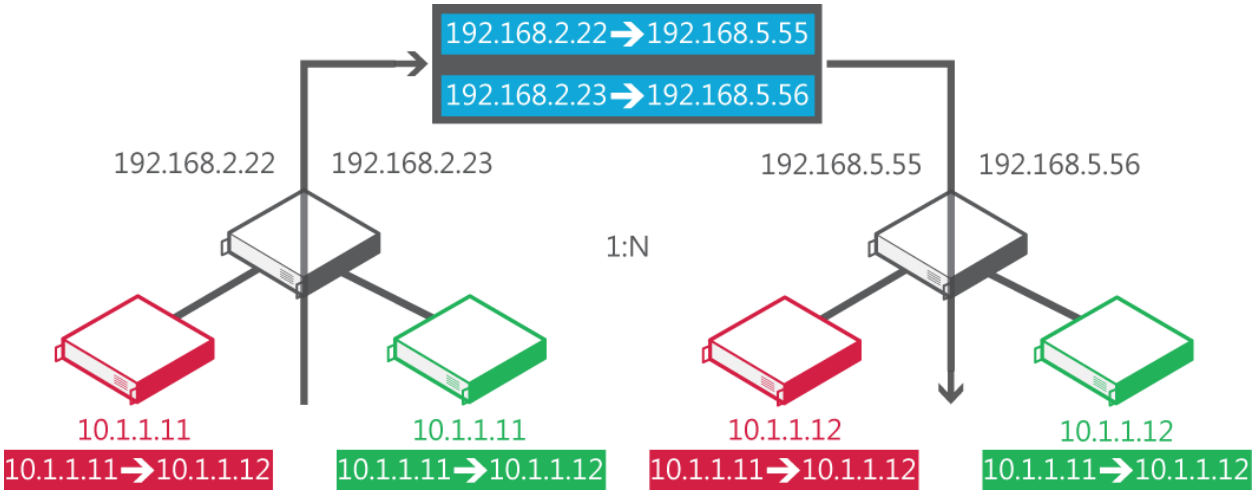

IP Address Rewrite can provide better performance because it's compatible with the existing Windows networking offload technologies, such as VMQ. It also means that there is no need to upgrade existing network adapters, switches, or other network appliances.

Generic Routing Encapsulation (GRE)

With GRE, all the virtual machine's packets are encapsulated with a new header before they're sent on the wire. GRE provides better network scalability because all virtual machines on a specific host can share the same provider IP address, as shown in the following diagram.

Figure 26: Example of Generic Routing Encapsulation

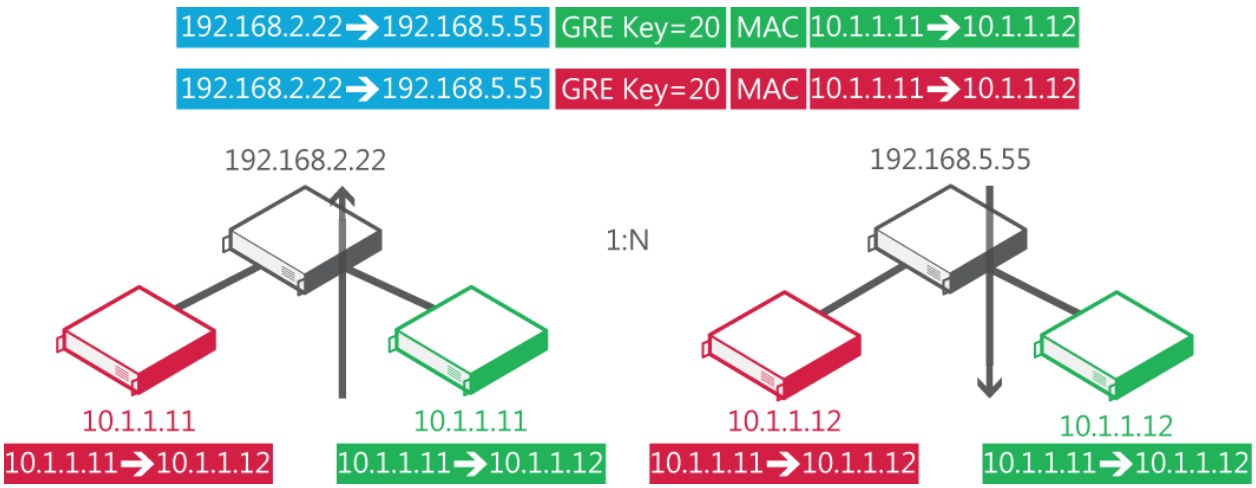

Requirements

Hyper-V Network Virtualization requires Windows Server 2012 and the Hyper-V server role.

Summary

Hyper-V Network Virtualization lets you customize policies to shape network traffic in your private cloud, without relying on complex VLANs. With Network Virtualization, different customers can share physical network infrastructures more securely and can choose IP addresses for their hosted services that preserve integration with their private networks. Finally, Hyper-V Network Virtualization lets you use your infrastructure more efficiently because you're free to move customer workloads anywhere in your datacenter so you can pursue more aggressive consolidation strategies.

Hyper-V Extensible Switch

Windows Server 2012 provides improved multitenant security for customers on a shared IaaS cloud through the new Hyper-V Extensible Switch. The Hyper-V Extensible Switch is a layer-2 virtual interface that provides programmatically managed and extensible capabilities to connect virtual machines to the physical network, as shown in the following diagram.

Figure 27: Hyper-V Extensible Switch

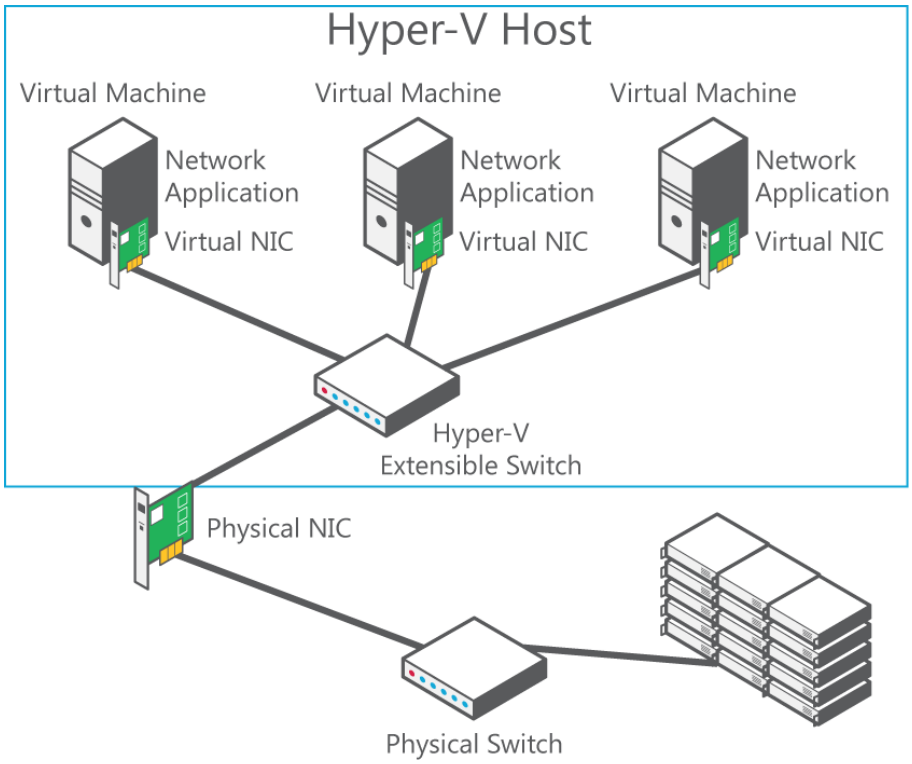

Technical description

Windows Server 2012 provides the isolation and security capabilities for multitenancy by offering these new features:

- Multitenant virtual machine isolation through PVLANs
- Protection from Address Resolution Protocol/Neighbor Discovery (ARP/ND) poisoning (also called spoofing)
- Protection against DHCP snooping and DHCP Guard
- Virtual port ACLs
- The capability to trunk traditional VLANs to virtual machines.
- Monitoring
- Windows PowerShell/Windows Management Instrumentation (WMI).

Virtual machine isolation with PVLANs

Windows Server 2012 introduces support for PVLANs, which provides isolation between two virtual machines on the same VLAN.

When a virtual machine doesn't need to communicate with other virtual machines, you can use PVLANs to isolate it from other virtual machines in your datacenter. To do this, assign every virtual machine in a PVLAN one primary VLAN ID and one or more secondary VLAN IDs. You can put the secondary PVLANs into one of three modes, as shown in the following table.

Figure 28: Supported modes for secondary PVLANs

PVLAN mode	Description
Isolated	Isolated ports cannot exchange packets with each other at layer 2.
Promiscuous	Promiscuous ports can exchange packets with any other port on the same primary VLAN ID.
Community	Community ports on the same VLAN ID can exchange packets with each other at layer 2.

These PVLAN modes determine which virtual machines on the PVLAN a virtual machine can communicate with. If you want to isolate a virtual machine, you put it into isolated mode.

ARP/ND spoofing protection

Hyper-V Extensible Switch provides protection against a malicious virtual machine stealing IP addresses from other virtual machines through ARP spoofing (known as ARP poisoning in IPv4). With this type of man-in-the-middle attack, a malicious virtual machine sends a fake ARP message, which associates its own MAC address to an IP address it doesn't own. Unsuspecting virtual machines send the network traffic targeted to that IP address to the MAC address of the malicious virtual machine instead of the intended destination.

For IPv6, Windows Server 2012 provides equivalent protection for Neighbor Discovery (ND) spoofing.

DHCP Guard protection

In a DHCP environment, a rogue DHCP server could intercept client DHCP requests and provide incorrect address information. The rogue DHCP server could cause traffic to be routed to a malicious intermediary that sniffs all traffic before forwarding it to the legitimate destination. To protect against this type of man-in-the-middle attack, the Hyper-V administrator designates which Hyper-V Extensible switch ports can have DHCP servers connected to them. DHCP server traffic from other Hyper-V Extensible switch ports is automatically dropped. The Hyper-V Extensible Switch now protects against a rogue DHCP server attempting to provide IP addresses that would cause traffic to be rerouted.

Virtual port ACLs

Port ACLs provide a mechanism for network isolation and metering network traffic for a virtual port on the Hyper-V Extensible Switch. By using port ACLs, you can meter the IP addresses or MAC addresses that can, or can't, communicate with a virtual machine. For example, you can use port ACLs to enforce isolation of a virtual machine by allowing it to communicate with only the Internet or a predefined set of addresses.

By using the metering capability, you can measure network traffic going to or from a specific IP address or MAC address, which allows you to report on traffic sent or received from the Internet or from network storage arrays.

You can configure multiple port ACLs for a virtual port. Each port ACL consists of a source or destination network address and a permit to deny or meter action. The metering capability also supplies information about the number of instances where traffic was attempted, either to or from a virtual machine, from a restricted ("deny") address.

Trunk mode to virtual machines

With the Hyper-V Extensible Switch trunk mode, traffic from multiple VLANs can be directed to a single network adapter in a virtual machine that could previously receive traffic from only one VLAN. As a result, traffic from different VLANs is consolidated, and a virtual machine can listen in on multiple VLANs. This feature helps you shape network traffic and enforce multitenant security in your datacenter.

Monitoring

Many physical switches can monitor the traffic from specific ports flowing through specific virtual machines on the switch. Hyper-V Extensible Switch provides port mirroring, which allows an administrator to designate which virtual ports should be monitored and to which virtual port the monitored traffic should be delivered for further processing. For example, a security monitoring virtual machine can look for anomalous patterns in the traffic flowing through other specific virtual machines on the switch. In addition, an administrator can diagnose network connectivity issues by monitoring traffic bound for a particular virtual switch port.

Windows PowerShell/WMI

As with all features in Windows Server 2012, you can use Windows PowerShell cmdlets for the Hyper-V Extensible Switch to build command-line tools or automated scripts for setup, configuration, monitoring, and troubleshooting. These cmdlets can be run remotely. Windows PowerShell also allows third parties to build their own tools to manage the Hyper-V Extensible Switch.

Requirements

Hyper-V Extensible Switch extensibility is built into the Hyper-V server role and requires Windows Server 2012.

Summary

The Hyper-V Extensible Switch in Windows Server 2012 offers several new features to provide isolation and security capabilities for multitenancy. For example, hosting providers can use port ACLs to fully isolate customers' networks from one another without the requirement to set up and maintain VLANs.

Hyper-V Extensible Switch also provides protection against ARP poisoning and spoofing, protection against DHCP snooping, VLAN trunk mode support, and port mirroring for monitoring virtual machine network traffic. Like most of the new or enhanced features in Windows Server 2012, Hyper-V Extensible Switch offers improved manageability by supporting Windows PowerShell and WMI for command-line or automated scripting, and full event logging.

Extending the Hyper-V Extensible Switch for new capabilities

Many enterprises need the ability to extend virtual switch features with their own plug-ins to suit their virtual environment. When IT pros install virtual switches, they look for the same kind of functionality that they can achieve on physical networks, such as adding firewalls, intrusion detection systems, and network traffic monitoring tools. The challenge has been finding straightforward ways to add virtualized appliances, extensions, and other features and functions to virtual switches. Most virtual switch technology offerings are built around closed systems that make it difficult for enterprise developers and third-party vendors to build solutions and to quickly and easily install new functionality into their virtual switches.

The Hyper-V Extensible Switch overcomes that challenge. With the Hyper-V Extensible Switch, IT pros can easily add more functionality to their virtual machines and networks. At the same time, internal enterprise developers and third-party providers have an open platform for creating solutions that extend the basic functionality of the switch. If you're in charge of making IT purchasing decisions at your company, you want a virtualization platform does not lock you in to a small set of compatible features, devices, or technologies.

Technical description

The Hyper-V Extensible Switch in Windows Server 2012 is a layer-2 virtual network switch that provides programmatically managed and extensible capabilities to connect virtual machines to the physical network. The Hyper-V Extensible Switch is an open platform that lets multiple vendors provide extensions that are written to standard Windows API frameworks. The reliability of extensions is strengthened through the Windows standard framework and reduction of required third-party code for functions and is backed by the Windows Hardware Quality Labs (WHQL) certification program. You can manage the Hyper-V Extensible Switch and its extensions by using Windows PowerShell, programmatically with WMI or the Hyper-V Manager UI.

Extensibility

The Hyper-V Extensible Switch architecture in Windows Server 2012 is an open framework that lets third parties add new functionality such as monitoring, forwarding, and filtering into the virtual switch. Extensions are implemented by using NDIS filter drivers and Windows Filtering Platform (WFP) callout drivers. These two public Windows platforms for extending Windows networking functionality are used as follows:

- **NDIS filter drivers:** Used to monitor or modify network packets in Windows. NDIS filters were introduced with the NDIS 6.0 specification.
- **WFP callout drivers**: Introduced in Windows Vista and Windows Server 2008, let ISVs create drivers to filter and modify TCP/IP packets, monitor or authorize connections, filter IPsec-protected traffic, and filter RPCs. Filtering and modifying TCP/IP packets provides unprecedented access to the TCP/IP packet processing path. In this path, you can examine or modify outgoing and incoming packets before additional processing occurs. By accessing the TCP/IP processing path at different layers, you can more easily create firewalls, antivirus software, diagnostic software, and other types of applications and services. For more information, see the Windows Filtering Platform.

Extensions can extend or replace three aspects of the switching process:

- Ingress filtering
- Destination lookup and forwarding
- Egress filtering

In addition, by monitoring extensions you can gather statistical data by monitoring traffic at different layers of the switch. Multiple monitoring and filtering extensions can be supported at the ingress and egress portions of the Hyper-V Extensible Switch. Only one instance of the forwarding extension may be used per switch instance, and it overrides the default switching of the Hyper-V Extensible Switch.

The following table lists the various types of Hyper-V Extensible Switch extensions.

Table 2: Types of Hyper-V Extensible Switch extensions

Extension	Purpose	Potential examples	Extensibility component
Network Packet Inspection	Inspecting network packets, but not altering them.	sFlow and network monitoring	NDIS filter driver
Network Packet Filter	Injecting, modifying, and dropping network packets.	Security	NDIS filter driver
Network Forwarding	Third-party forwarding that bypasses default forwarding.	OpenFlow, Virtual Ethernet Port Aggregator (VEPA), and proprietary network fabrics	NDIS filter driver
Firewall/ Intrusion Detection	Filtering and modifying TCP/IP packets, monitoring or authorizing connections, filtering IPsec-protected traffic, and filtering RPCs.	Virtual firewall and connection monitoring	WFP callout driver

The Hyper-V Extensible Switch provides an open switch API that lets enhanced switch and management products work with Hyper-V.

The following figure shows the architecture of the Hyper-V Extensible Switch and the extensibility model.

Figure 29: Architecture of the Hyper-V Extensible Switch

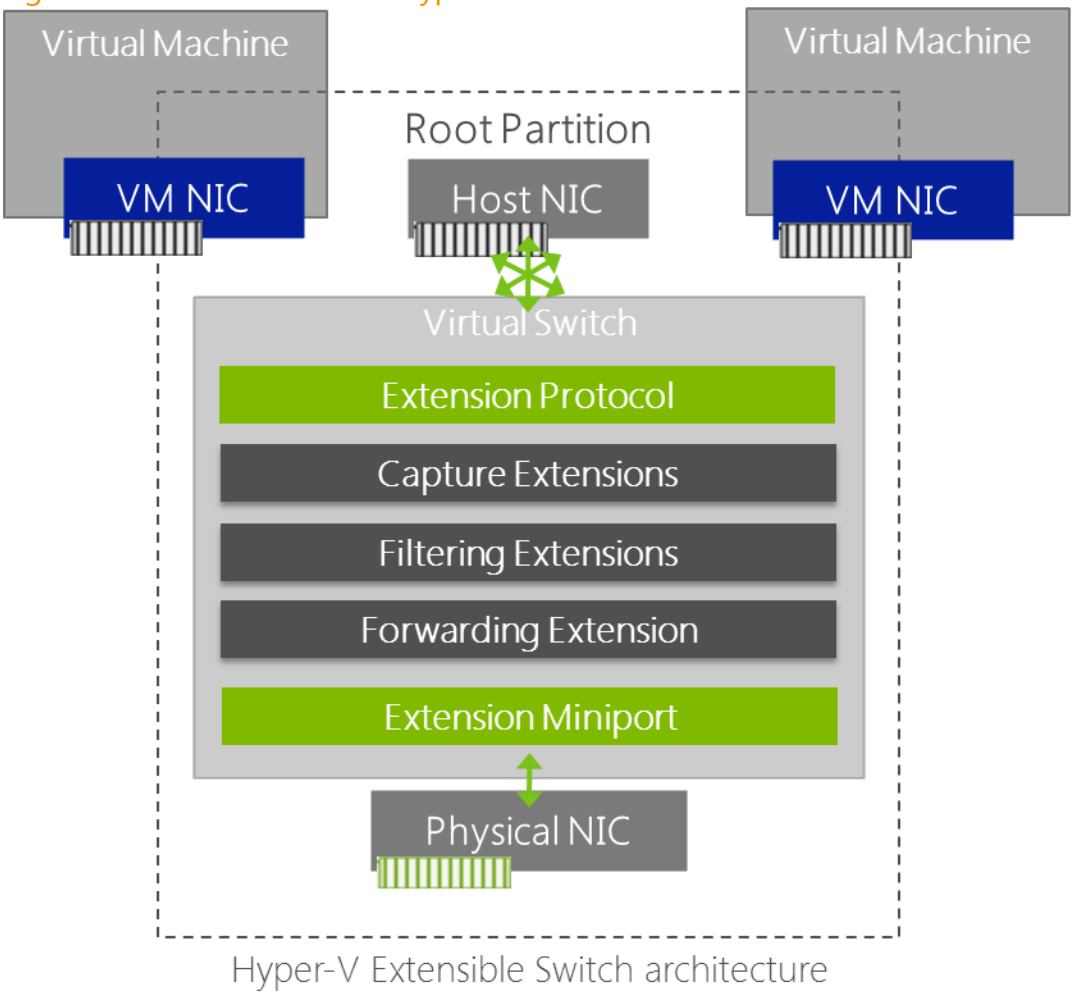

38

Some other features of Hyper-V Extensible Switch extensibility are:

- **Extension uniqueness:** Extension state/configuration is unique to each instance of an Extensible Switch on a machine.
- **Extensions learn virtual machine life cycle:** Virtual machine activity cycle is similar to that of physical servers, having peak times during various parts of the day or night based on their core workloads. Extensions can learn the flow of network traffic based on the workload cycle of your virtual machines and optimize your virtual network for greater performance.
- **Extensions can veto state changes:** Extensions can implement monitoring, security, and other features to further improve the performance, management, and diagnostic enhancements of the Hyper-V Extensible Switch. Extensions can help make the system more secure and reliable by identifying harmful state changes, and stopping them from being implemented.
- **Multiple extensions on same switch:** Multiple extensions can coexist on the same Hyper-V Extensible Switch.

Manageability

You can use these management features, built into the Hyper-V Extensible Switch, to troubleshoot and resolve problems on Hyper-V Extensible Switch networks:

- **Windows PowerShell and scripting support:** Windows Server 2012 provides Windows PowerShell cmdlets for the Hyper-V Extensible Switch that let you build command-line tools or automated scripts for setup, configuration, monitoring, and troubleshooting. Windows PowerShell also lets third parties to build their own Windows PowerShell–based tools to manage the Hyper-V Extensible Switch.
- **Unified tracing and enhanced diagnostics:** The Hyper-V Extensible Switch includes unified tracing to provide two levels of troubleshooting. At the first level, the Event Tracing for Windows (ETW) provider for the Hyper-V Extensible Switch permits tracing packet events through the Hyper-V Extensible Switch and extensions, making it easier to pinpoint where an issue occurred. The second level permits capturing packets for a full trace of events and traffic packets.

Requirements

Hyper-V Extensible Switch extensibility is built into the Hyper-V server role and requires Windows Server 2012.

Summary

The Hyper-V Extensible Switch is an open platform, so that third-party vendors can provide plug-ins that supply additional functionality such as traffic monitoring, firewall filters, and switch forwarding. The management of these plug-ins is unified through Windows PowerShell cmdlets and WMI scripting.

The Hyper-V Extensible Switch permits easier implementation and management of virtualized datacenters by providing:

- **Open platform to fuel plug-ins:** The Hyper-V Extensible Switch is an open platform that lets plug-ins sit in the virtual switch between all traffic, including virtual machine–to–virtual machine traffic. Extensions can provide traffic monitoring, firewall filters, and switch forwarding. To jump-start the ecosystem, several partners will announce extensions when the Hyper-V Extensible Switch is released. No "one-switch-only" solution for Hyper-V will occur.
- **Free core services:** Core services are provided for extensions. For example, all extensions have live migration support by default, and no special coding for services is required.
- **Windows reliability and quality:** Extensions experience a high level of reliability and quality from the strength of the Windows platform and Windows logo certification program that sets a high bar for extension quality.
- **Unified management:** Managing extensions is integrated into Windows management through Windows PowerShell cmdlets and WMI scripting.
- **Easier support:** Unified tracing means it's quicker and easier to diagnose issues when they arise. Less down time increases availability of services.
- **Live migration support:** The Hyper-V Extensible Switch provides capabilities so extensions can participate in Hyper-V live migration.

The Hyper-V Extensible Switch gives third-party vendors the opportunity to develop custom solutions for handling network traffic in a Windows Server 2012 virtual network. For example, these solutions can be used to emulate a vendor's physical switch and its policies or to monitor and analyze traffic.

Quality of Service (QoS)

QoS for networks is an industry-wide set of standards and mechanisms for ensuring high-quality performance for critical applications. By employing QoS mechanisms, you can use existing resources efficiently and provide the required level of service without expanding your network or infrastructure. The goal of QoS is to provide preferential delivery service for the applications that need it by ensuring sufficient bandwidth, controlling latency, and reducing data loss.

Windows Server 2012 expands the power of QoS with the ability to guarantee a minimum bandwidth to a virtual machine or a service. This is important for hosting companies that need to honor SLA clauses that promise a minimum network bandwidth to customers. It's equally important to enterprises that need to have predictable network performance when they run virtualized server workloads on shared hardware.

Addressing the needs of enterprises and public cloud hosting providers

Public cloud hosting providers want to host customers on a server running Hyper-V and be able to have a certain level of performance based on SLAs. They don't want customers impacted or compromised by other customers on their shared infrastructure, which includes CPU, storage, and network resources. Enterprises have similar requirements. They want to run multiple application servers on a server running Hyper-V and have each application server deliver predictable performance. Lack of performance predictability has traditionally driven administrators to put fewer virtual machines on a capable server or to simply shy away from virtualization; because of this, they spend more money on physical equipment and infrastructure.

Currently, most hosting providers and enterprises use a dedicated network adapter and a dedicated network for a specific type of workload, such as storage or live migration, to achieve network performance isolation on a server running Hyper-V. Although this deployment strategy works for those using 1–GbE network adapters, it is impractical for those using, or planning to use, 10–GbE network adapters. One 10-GbE network adapter, or two for high availability, provides sufficient bandwidth for all of the workloads on a server running Hyper-V in most deployments, but 10–GbE network adapters and switches are considerably more expensive than their 1–GbE counterparts. To best use 10–GbE hardware, a server running Hyper-V requires new capabilities to manage bandwidth.

Technical description

Windows Server 2012 includes new QoS bandwidth management features that enable hosting providers and enterprises to provide services with predictable network performance to virtual machines on a server running Hyper-V. Windows Server 2012 supports bandwidth floors, as well as bandwidth caps. Windows Server 2012 also takes advantage of Data Center Bridging (DCB)–capable hardware to converge multiple types of network traffic on a single network adapter, with a guaranteed level of service to each type. With Windows PowerShell, you can configure all of these features manually or automate them in a script to manage a group of servers, regardless of whether they're domain-joined or stand-alone, with minimal dependencies.

Rate limiting

In Windows Server 2008 R2, QoS supported the enforcement of maximum bandwidth. This is known as rate limiting. Consider a typical server running Hyper-V in which there are four types of network traffic sharing a single 10–GbE network adapter:

- **Virtual machine:** Traffic between virtual machines and resources on other servers.
- **Storage:** Traffic to and from storage.
- **Live migration:** Traffic for live migration of virtual machines between servers running Hyper-V.
- **Cluster heartbeat:** Traffic to and from a Cluster Shared Volume (CSV), intercommunication between nodes in a cluster.

If virtual machine data is rate-limited to 3 gigabits per second (Gbps), the sum of the virtual machine data transfers cannot exceed 3 Gbps at any time, even if the other network traffic types don't use the remaining 7 Gbps of bandwidth. However, this also means that the other types of traffic can reduce the actual amount of bandwidth available for virtual machine data to unacceptable levels, depending on how their maximum bandwidths are defined.

Minimum bandwidth

QoS in Windows Server 2012 introduces a new bandwidth management feature: minimum bandwidth. Unlike maximum bandwidth, which is a bandwidth cap, minimum bandwidth is a bandwidth floor. It guarantees a certain amount of bandwidth to a specific type of traffic. The following figure shows how minimum bandwidth works for each of the four types of network traffic flows in three different time periods: T1, T2, and T3.

Figure 30: Assigning minimum bandwidth to services

Service	Reservation	Runtime bandwidth demand (gigabits per second) T1	T2	T3
Virtual machine	30%	4	4	2
Storage	40%	5	5	6
Live migration	20%	0	3	2
Cluster Shared Volume	10%	0.5	1	0

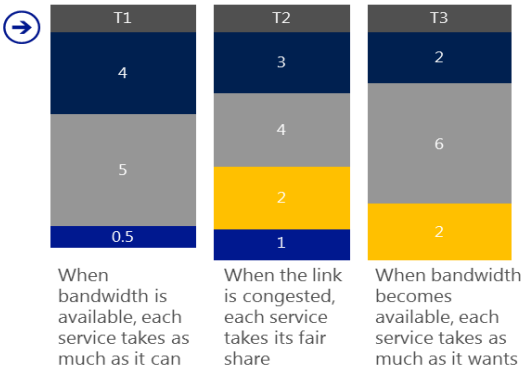

In the previous figure, the table on the left shows the configuration of the minimum amount of required bandwidth that a given type of network traffic flow needs. For example, storage is configured to have at least 40%of the bandwidth (4 Gbps of a 10–GbE network adapter) at any time. The chart on the right shows the actual amount of bandwidth each type of network traffic has in T1, T2, and T3. In this example, storage is actually 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 desired network bandwidth exceeds the available bandwidth (such as in the T2 period in the figure), minimum bandwidth makes sure that each type of network traffic receives 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's congestion—there's 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 using available bandwidth.

Windows Server 2012 offers two different mechanisms to enforce minimum bandwidth: through the enhanced packet scheduler in Windows and through the network adapters that support Data Center Bridging DCB. In both cases, network traffic first must be classified. Windows classifies a packet itself or gives instructions to a network adapter to classify it. The result of classification is a number of traffic flows in Windows, and a given packet can belong to only one of them.

For example, a traffic flow could be a live migration connection, a file transfer between a server and a client, or a Remote Desktop Connection. Based on how the bandwidth policies are configured, either the packet scheduler in Windows or the network adapter dispatches the packets at a rate equal to, or higher than, the minimum bandwidth configured for the traffic flow.

Each of the two mechanisms has its advantages and disadvantages:

- The software solution, which is built on the new packet scheduler in Windows Server 2012, provides a fine granularity of classification. It's the only viable choice if many traffic flows require minimum bandwidth enforcement. A typical example is a server running Hyper-V hosting many virtual machines, where each virtual machine is classified as a traffic flow.
- The hardware solution, which depends on DCB support on the network adapter, supports far fewer traffic flows. However, it can classify network traffic that doesn't originate from the networking stack. A typical scenario involves a Converged Network Adapter (CNA) that supports Internet SCSI (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 doesn't process this offloaded traffic, DCB is the only viable choice to enforce minimum bandwidth.

You can employ both mechanisms on the same server. For example, a server running Hyper-V has two physical network adapters: one that binds to a virtual switch and serves virtual machine data, and the other that serves the rest of the traffic of the host server. You can enable the software-based minimum bandwidth in Hyper-V to make sure bandwidth is shared fairly among virtual machines and enable the hardware-based minimum bandwidth on the second network adapter to make sure bandwidth is shared fairly among various types of network traffic from the host server.

Many configuration options use network adapters on servers running Hyper-V. The following figure shows an example.

Figure 31: Assigning physical network adapters to services and virtual machines in Hyper-V

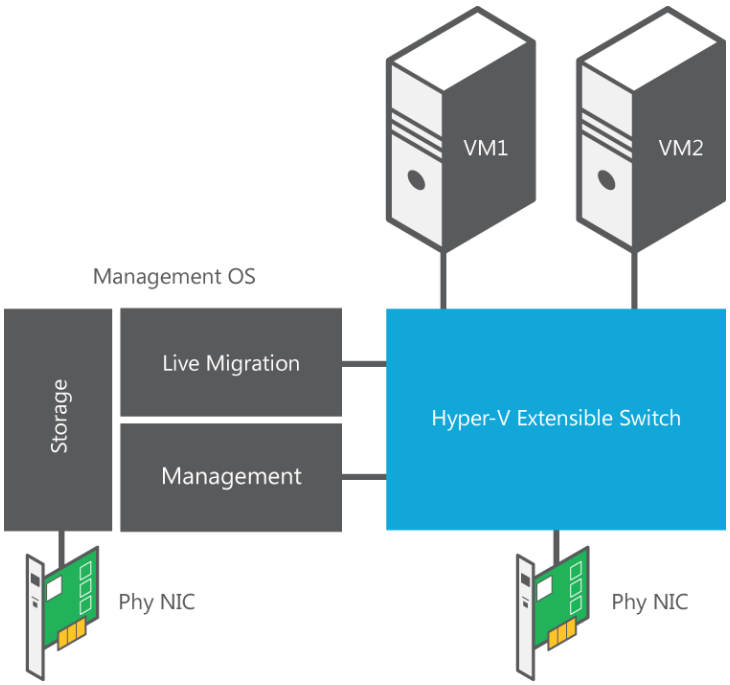

In this example, Hyper-V simply uses a dedicated network adapter for storage and routes both live migration and management traffic through the Hyper-V Extensible Switch.

We don't recommend that you enable both mechanisms at the same time for a specific type of network traffic. For example, consider live migration and storage traffic that are configured to use the second network adapter on the server running Hyper-V. If you've already configured the network adapter to allocate bandwidth for live migration and storage traffic using DCB, you shouldn't also configure the packet scheduler in Windows to do the same, and vice versa. Enabling both mechanisms at the same time for the same types of network traffic compromises the intended results.

Configuring and managing QoS

In Windows Server 2012, you can dynamically manage QoS policies and settings with Windows PowerShell. The new QoS cmdlets support both the QoS functionalities available in Windows Server 2008 R2, such as maximum bandwidth and priority tagging, and the new features in Windows Server 2012, such as minimum bandwidth.

Requirements

Minimum QoS can be enforced through two methods:

- The first method relies on software built into Windows Server 2012 and has no additional requirements.
- The second method, which is hardware-assisted, requires a network adapter that supports DCB.

Summary

By taking advantage of the QoS enhancements in Windows Server 2012, you can configure minimum and maximum bandwidth levels for virtual machines. You can take advantage of the increased throughput of 10–GbE network adapters, while also configuring different bandwidth settings based on workloads for those adapters.

With Windows Server 2012, you can better control bandwidth, latency, and data loss. You can also use Windows PowerShell to automate most of the management of QoS features to reduce time and effort.

QoS allows you to guarantee minimum levels of service for hosted customers, without incurring higher costs from expanding your network or infrastructure. You can also eliminate much of the complexity that used to be required to set up and manage QoS bandwidth support.

Remote Desktop Protocol (RDP) WAN Optimizations

RDP improvements in Windows Server 2012 provide reliable network connections over a WAN to users of devices, such as PCs, laptops, tablets, and phones. You can access your desktop and applications from anywhere, using any device, without the latency and jitters of a typical WAN connection.

Technical description

With today's modern workforce, clients frequently need to connect from branch offices, homes, or hotels over low-bandwidth connections. To support VDI, Remote Desktop Services sessions, or Windows Server 2008 R2 RemoteApp sessions over WANs, Remote Desktop Services must adapt to different network conditions and be quick and responsive.

RDP in Windows Server 2012 and Windows 8 responds to this challenge by including optimizations for low-bandwidth, high-latency connections. To help achieve this, RDP adds these improvements:

- **User Datagram Protocol (UDP):** RDP chooses between TCP and UDP transports, depending on the content type and connection quality. When Remote Desktop is enabled on a computer, UDP for port 3389 is automatically enabled in Windows Firewall. For enhanced performance, verify that this port is enabled on your network.
- **Forward error correction (FEC):** RDP also supports FEC. The server sends redundant data across the network to recover quickly from packet loss without requiring retransmission, even over lossy networks.
- **Network auto-detect:** RDP detects end-to-end network speed by measuring latency, maximum bandwidth, and packet loss and then adjusts the type of connection and the data transfer based on the available bandwidth.
- **Dynamic transport detection:** RDP uses dynamic transport detection to select the appropriate transport to communicate with the client. The system first tries using UDP as the transport mechanism. If that fails, it automatically switches to TCP, using the most appropriate transport to achieve the recommended user experience.
- **Congestion control:** RDP employs congestion control to prevent loss of packets, recover quickly from transmission gaps, and avoid further delays. This helps maintain the necessary flow of data to the client to provide a seamless experience over an RDP connection.

RDP uses the information from network auto-detect and congestion control to adjust the outbound data to improve the overall user experience.

Requirements

To take advantage of RDP WAN optimizations, you'll need Windows Server 2012.

Summary

The RDP enhancements in Windows Server 2012 optimize WAN connections to provide a more seamless experience for users. The RDP WAN optimizations employ FEC, network auto-detection, transport detection, congestion control, and other techniques to compensate for variable network conditions and provide a positive working experience for remote users.

WebSocket Protocol

Web applications have become more data-intensive and interactive, with data constantly flowing from server applications to Asynchronous JavaScript and XML (AJAX) client applications in the browser. Stock quotes, news updates, interactive charts, and collaboration applications are some examples of real-time data flows.

To the user, real-time data seems to be constantly available and automatically flowing into the browser. To the developer, the mechanics of delivering this type of data have been constrained due to the inflexible request/response nature of Hypertext Transfer Protocol (HTTP).

To address this issue, Windows Server 2012 introduces support for WebSocket Protocol, an open standards-based network protocol that defines how to transfer data using encrypted, real-time bidirectional communications.

With WebSocket Protocol, bidirectional communications between a client and the server are improved, helping to enhance the overall performance of data-intensive and interactive web applications.

Technical description

WebSocket Protocol provides encrypted, real-time, bidirectional communications between a client and a server. The client could be a browser running a JavaScript application, an immersive Windows-based application, or an application running on a different platform, such as a mobile device. The WebSocket Protocol feature in Windows Server 2012 is supported in Microsoft Internet Information Services (IIS) 8.0, ASP.NET 4.5, and Windows Communication Foundation (WCF), using either native or managed programming APIs for writing server-side WebSocket applications.

Establishing a WebSocket Connection

As shown in the following figure, a WebSocket connection is established using an HTTP handshake. Data is sent over the underlying TCP connection using binary or UTF-8 messages. The connection can traverse firewalls and proxies, but for better results, it should be tunneled over Secure Sockets Layer (SSL).

Figure 32: Establishing a WebSocket connection

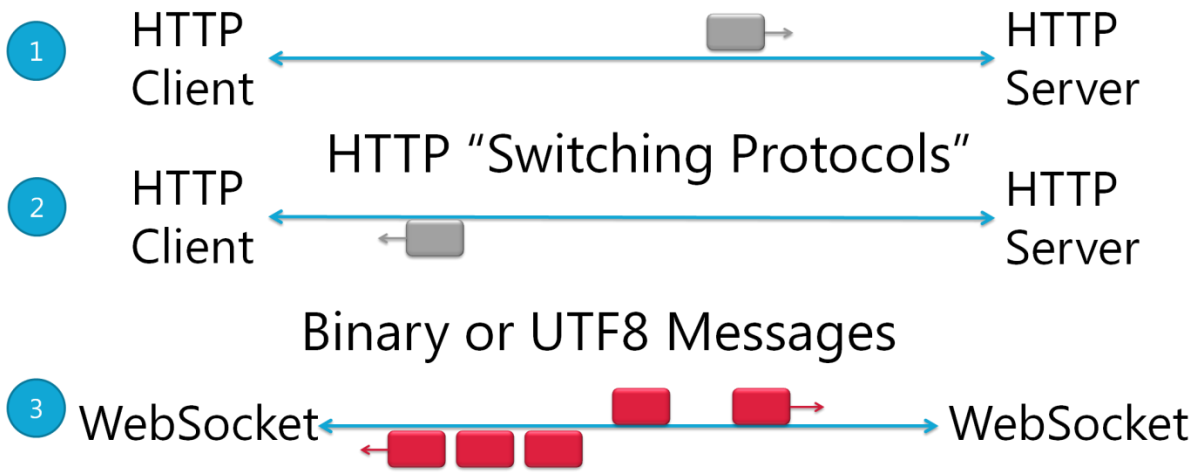

For more information, read "The WebSocket Protocol" at http://tools.ietf.org/html/draft-ietf-hybi-thewebsocketprotocol-10.

Requirements

WebSocket Protocol requires:

- Windows Server 2012
- IIS 8.0
- ASP.NET 4.5

Summary

With the new WebSocket Protocol feature in Windows Server 2012, you can add encrypted, authenticated, and real-time, bidirectional client-server communications to your ASP.NET 4.5 web applications, which helps to enhance the overall performance of data-intensive and interactive web applications.

Server Name Indicator (SNI)

Administrators can more easily create and administer secure websites using a new feature in Windows Server 2012 called Server Name Indicator (SNI). With SNI, a host name is used as a third piece of information to uniquely identify the network endpoint, eliminating the need for a dedicated IP address for each secure site.

Simplify management and improve SSL scalability

IIS 8.0 support for the SNI standard allows multiple SSL-protected websites to share a single public IP address. SNI-capable browsers (including most modern browsers, for example, Windows Internet Explorer 7, 8, and 9 running on Windows Vista or Windows 7) can send the host name as a part of the SSL negotiation process, eliminating the need to have a dedicated IP address for each secure site. To provide greater SSL scalability, Windows Server 2012 provides a new certificate store designed to support thousands of SSL certificates on a single server. Combined, these two features allow web hosting providers and enterprises with web farms to host more SSL-protected websites using fewer servers and IP addresses.

Technical description

SNI is an open-standards extension to the SSL and Transport Layer Security (TLS) protocols that browsers and web servers use to authenticate servers and encrypt communications. With SNI, browsers can use the host name to identify the web server that they want to communicate with, instead of just the IP address and port. Most modern browsers support the SNI standard.

Windows Server 2012 provides a new certificate store named Web Hosting that you can use to associate SSL certificates with websites and applications. The Web Hosting certificate store is designed to scale to thousands of certificates.

Requirements

SNI requires:

- Windows Server 2012
- SSL certificates

Summary

With SNI, Windows Server 2012 allows you to share a single IP address and server between thousands of SSL-protected websites. And, with the new Web Hosting certificate store, you can benefit from the increased scalability of storing thousands of SSL certificates on a single server. This provides increased secure site density and reduced hosting costs. As an added benefit, you can conserve your IPv4 address space.

Connecting Users Easily to IT Resources

Windows Server 2012 gives IT pros the tools and features needed to connect remote users to resources and services more easily and efficiently. These remote access features are discussed in this paper:

- DirectAccess and virtual private network (VPN) integration, including Routing and Remote Access service (RRAS)
- Microsoft BranchCache

DirectAccess and VPN Integration

Windows Server 2012 provides an integrated remote access solution that's easier to deploy. Employees can access corporate network resources when working remotely, and IT administrators can manage corporate computers located outside the internal network. Windows Server 2012 accomplishes this by integrating two existing remote access technologies: DirectAccess for automatic, transparent connectivity and traditional VPNs for compatibility.

Secure and efficient remote access

Increasing numbers of employees work remotely, but they're expected to maintain a high level of productivity away from the office. This expectation increases the need for remote users to have secure remote access to corporate networks.

DirectAccess

DirectAccess lets remote users more securely access internal resources without the need for a VPN to provide connectivity. DirectAccess improvements in Windows Server 2012 include simplified deployment steps, support for new deployment scenarios, a streamlined management experience, and improved scalability and performance.

DirectAccess transparently connects client computers to the internal network every time the computer connects to the Internet, even before the user logs on, as shown in the following figure.

Figure 33: DirectAccess connection architecture

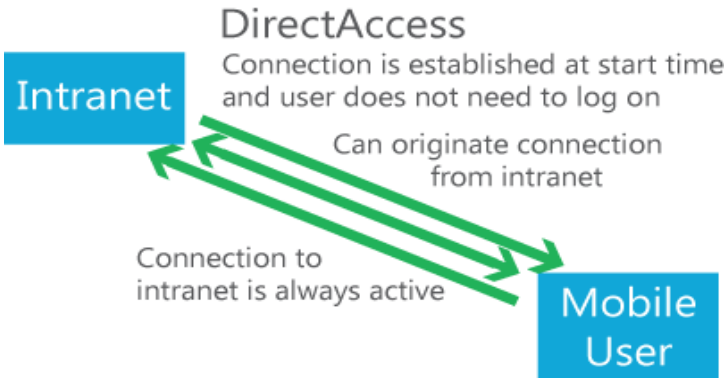

DirectAccess also lets administrators easily monitor connections and remotely manage DirectAccess client computers on the Internet.

Routing and Remote Access service (RRAS)

RRAS provides traditional client VPN connectivity for unmanaged client computers and for computers running older operating systems that can't connect through DirectAccess. In addition, RRAS site-to-site VPN provides connectivity between VPN servers, as shown in the following figure.

Figure 34: VPN connection architecture

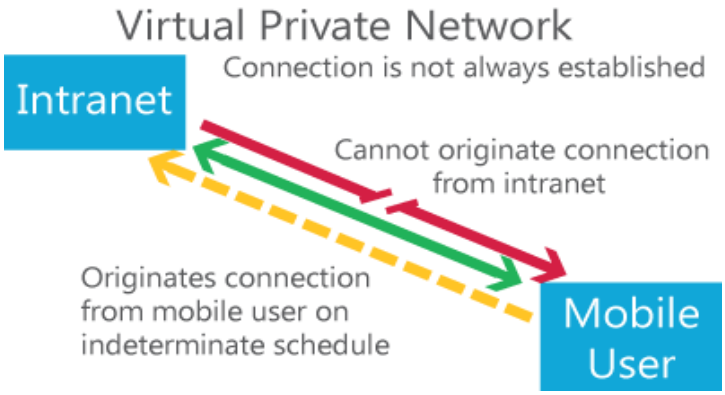

Integrated remote access

Remote access in Windows Server 2012 integrates DirectAccess and RRAS VPN. Administrators can configure DirectAccess and VPNs together in the VDI (Virtualization Desktop Infrastructure) management console by using a single set of wizards. They can also configure other RRAS features by using the legacy RRAS management console. The new role allows easier migration of Windows 7 RRAS and DirectAccess deployments, and it provides new features and improvements.

Cross-premises connectivity

Windows Server 2012 provides a highly cloud-optimized operating system. VPN site-to-site functionality in remote access provides cross-premises connectivity between enterprises and hosting service providers. Cross-premises connectivity enables organizations to connect to private subnets in a hosted cloud network. It also enables connectivity between geographically separate enterprise locations. With cross-premises connectivity, you can use existing networking equipment to connect to hosting providers by using the industry standard Internet Key Exchange version 2 (IKEv2) and IPsec protocol.

The following figure demonstrates how two organizations achieve cross-premises deployments by using Windows Server 2012.

Figure 35: Example of a cross-premises deployment

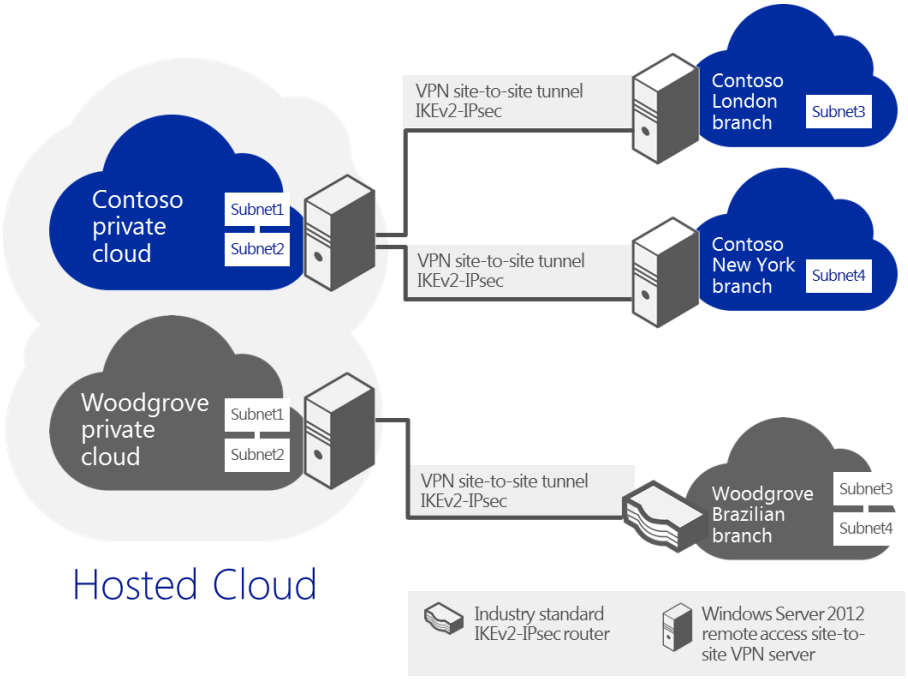

These steps describe the procedures used by Contoso and Woodgrove for the cross-premises deployment shown in the preceding figure:

1. Contoso.com and Woodgrove.com offload some of their enterprise infrastructure in a hosted cloud.
2. The hosting provider provides private clouds for each organization.
3. In the hosted cloud, virtual machines running Windows Server 2012 are configured as VDI (Virtualization Desktop Infrastructure) servers running site-to-site VPN.
4. In each hosted private cloud, a cluster of two or more VDI (Virtualization Desktop Infrastructure) servers is deployed to provide high availability and failover.
5. Contoso.com has two branch office locations. In each location, a Windows Server 2012 VDI (Virtualization Desktop Infrastructure) server is deployed to provide a cross-premises connectivity solution to the hosted cloud and between the branch offices.
6. Woodgrove.com can use existing routers to connect to the hosted cloud because cross-premises functionality in Windows Server 2012 complies with IKEv2 and IPsec standards.

Improved management experience

By using the new VDI (Virtualization Desktop Infrastructure) management console, you can configure, manage, and monitor multiple DirectAccess and VPN remote access servers in a single location. The console provides a dashboard that provides you with information about server and client activity. For more granular information, you can generate detailed reports. Operations status provides comprehensive monitoring information about specific server components. Event logs and tracing help diagnose specific issues. By using client monitoring, you can see detailed views of connected users and computers, and you can monitor which resources the clients are accessing. Accounting data can be logged to a local database or a Remote Authentication Dial-In User Service (RADIUS) server.

In addition to the VDI (Virtualization Desktop Infrastructure) management console, you can use Windows PowerShell command-line interface tools and automated scripts for remote access setup, configuration, management, monitoring, and troubleshooting.

On client computers, users can access the Network Connectivity Assistant (NCA) application, integrated with Windows Network Connection Manager, to see a concise view of the DirectAccess connection status and links to corporate help resources, diagnostics tools, and troubleshooting information. Users can also enter one-time password (OTP) credentials if OTP authentication for DirectAccess is configured.

Ease of deployment

The enhanced installation and configuration design in Windows Server 2012 allows you to set up a working deployment without changing your internal networking infrastructure. In simple deployments, you can configure DirectAccess without setting up a certificate infrastructure. DirectAccess clients can now authenticate themselves by using only Active Directory credentials; no computer certificate is required. In addition, you can select to use a self-signed certificate created automatically by DirectAccess for IP-HTTPS and for authentication of the network location server.

To further simplify deployment, DirectAccess in Windows Server 2012 supports access to internal servers that are running IPv4 only. An IPv6 infrastructure isn't required for DirectAccess deployment.

New and improved deployment scenarios

Remote access in Windows Server 2012 includes additional enhancements, including integrated deployment for several scenarios that required manual configuration in Windows Server 2008 R2. These include force tunneling (which sends all traffic through the DirectAccess connection), Network Access Protection (NAP) compliance, support for locating the nearest remote access server from DirectAccess clients in different geographical locations, and deploying DirectAccess for only remote management.

With Windows Server 2012, you can now configure a DirectAccess server with two network adapters at the network edge or behind an edge device, or with a single network adapter running behind a firewall or network address translation (NAT) device. The ability to use a single adapter removes the requirement to have dedicated public IPv4 addresses for DirectAccess deployment. With this configuration, clients connect to the DirectAccess server by using IP-HTTPS.

In Windows Server 2012, you can configure remote access servers in a multisite deployment that allows users in dispersed geographical locations to connect to the multisite entry point closest to them. You can distribute and balance traffic across the multisite deployment by using an external global load balancer.

DirectAccess in Windows Server 2008 R2 provides standard client IPsec authentication and two-factor authentication by using smart cards. DirectAccess in Windows Server 2012 adds support for two-factor authentication using an OTP, if third-party vendors provide the ability to use OTP solutions.

For two-factor smart card authentication, Windows Server 2012 supports the use of Trusted Platform Module (TPM)–based virtual smart card capabilities that are available in Windows 8. The TPM of client computers can act as a virtual smart card for two-factor authentication, which reduces the overhead and costs incurred in a smart card deployment.

Windows Server 2012 introduces the ability of computers to join a domain and receive domain settings remotely via the Internet. Using this capability, deployment of new computers in remote offices and provisioning of client settings to DirectAccess clients is easier.

You can configure client computers running Windows 8, Windows 7, and Windows Server 2008 R2 as DirectAccess clients. Clients running Windows 8 have access to all DirectAccess features, and they have an improved experience when connecting from behind a proxy server that requires authentication. Clients not running Windows 8 have these limitations:

- Need to download and install the DirectAccess Connectivity Assistant tool.
- Requires a computer certificate for authentication.
- In a multisite deployment, they'll be configured to always connect through the same entry point.

Scalability improvements

Remote access offers several scalability improvements, including support for more users with better performance and lower costs:

- You can cluster multiple remote access servers for load balancing, high availability, and failover. Cluster traffic can be load-balanced using Windows Network Load Balancing (NLB) or a third-party load balancer. Servers can be added to or removed from the cluster with few interruptions to the connections in progress.
- The VDI (Virtualization Desktop Infrastructure) server role takes advantage of SR-IOV for improved I/O performance when running 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.
- Optimization improvements in IP-HTTPS use the encryption that IPsec provides. This optimization, combined with the removal of the SSL encryption requirement, increases scalability and performance.

Requirements

This feature requires:

- Windows Server 2012, joined to an Active Directory Domain Services (AD DS) domain (for DirectAccess).
- The VDI (Virtualization Desktop Infrastructure) role.
- Client computers running Windows 7 or Windows 8 (for DirectAccess).

Summary

With the new DirectAccess and VDI (Virtualization Desktop Infrastructure) enhancements, you can easily provide more secure remote access connections for your users, and log reports for monitoring and troubleshooting those connections. The new features in Windows Server 2012 support deployments in dispersed geographical locations, improved scalability with high availability, and improved performance in virtualized environments.

BranchCache Enhancements

In Windows Server 2012, BranchCache has a streamlined deployment process and an improved ability to optimize bandwidth over WAN connections between BranchCache–enabled content servers and remote client computers. This functionality lets remote client computers access data and run applications in a more secure, efficient, and scalable way.

With BranchCache, users access files and applications over a WAN with the performance and experience of accessing local resources over a LAN.

Work productively across WANs

The recent trend of moving content servers to off-premises locations means that these servers frequently deliver content to remote users over WAN connections. The additional pressure on WAN connections can increase networking costs and reduce productivity by slowing application response times and content delivery speeds.

BranchCache provides a solution to these IT and business needs. It lets branch office clients identify, download, and share data efficiently, even over WAN connections.

BranchCache is improved in Windows 8 and Windows Server 2012, with a streamlined deployment process and an improved ability to optimize bandwidth over WAN connections between BranchCache–enabled content servers and remote client computers. This functionality lets remote client computers access data and run applications in a more secure, efficient, and scalable way.

Improvements in BranchCache

With these new BranchCache features in Windows Server 2012, you can deploy BranchCache in larger branch offices:

- **Deployment of multiple hosted cache servers:** 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 needed at a location.
- **Improved database performance:** BranchCache now uses the Extensible Storage Engine (ESE) database technology that powers Microsoft Exchange Server. This allows a single hosted cache server to meet the demands of more people while using the same hardware. 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 large organizations.

New tools and a simplified deployment model make BranchCache more effective, easier to implement, and less expensive to operate for these reasons:

- **BranchCache no longer requires office-by-office configuration:** Deployment is streamlined because there's no requirement for a separate GPO for each location. Only a single GPO that contains a small group of settings is required to deploy BranchCache in any size organization, from a small business to a large enterprise.
- **Client computer configuration is automatic:** Clients can be configured through Group Policy as distributed cache–mode clients by default. However, they search for a hosted cache server, and if one is discovered, clients automatically self-configure as hosted cache–mode clients.

- **Cache data is kept encrypted, and hosted cache servers do not require server certificates:** BranchCache security provides improved data encryption and other technologies, providing data security without requiring a public key infrastructure or additional drive encryption.
- **BranchCache provides tools to manipulate data and preload the content at remote locations:** You can push content to branch offices so it's immediately available when the first user requests it. This allows you to distribute content during periods of low WAN use.
- **BranchCache is deeply integrated with the Windows File Server:** BranchCache uses the Windows File Server state-of-the-art technology to divide files into small pieces and eliminate duplicates. This greatly 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 division calculations are performed once and may be done offline:** When a client computer that's running Windows 8 downloads content from a file server or web server running Windows Server 2012 and using new disk deduplication technology, there's no need for BranchCache to calculate how to divide the content, because the file server and web server have already made these calculations. Content information is calculated offline, before a BranchCache client 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 calculations have already been performed.
- **BranchCache is now manageable with Windows PowerShell and WMI:** This enables scripting and remote management of BranchCache content servers, hosted cache servers, and client computers.

Technical description

To optimize WAN bandwidth, BranchCache downloads content from your content servers and caches it at branch office locations, letting client computers at the branch office locations access the content locally.

After a client downloads content once, other client computers that request the same content don't download it from the content servers over the WAN connection. Instead, they retrieve small IDs, 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 Windows server or other client computers, depending on the mode in which you've deployed BranchCache.

The following figure illustrates this process.

Figure 36: Clients and content servers using content information structure to identify and transmit content

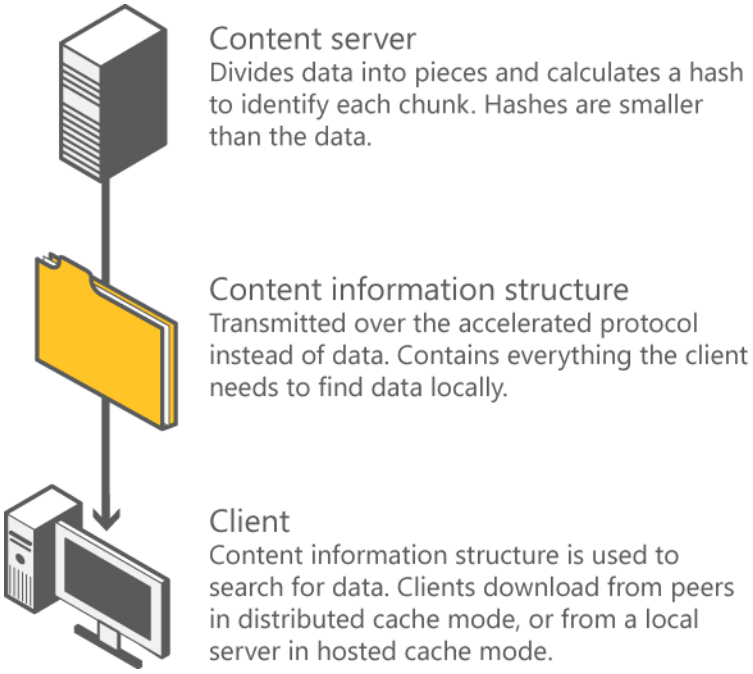

BranchCache modes

BranchCache can be configured in one of two modes: hosted cache mode or distributed cache mode.

If the branch office location has a server running Windows, you can configure BranchCache client computers in hosted cache mode. The branch server, called a hosted cache server, can store data locally after it's requested and transferred from the home office servers.

For offices that don't have an available server to deploy as a hosted cache server, you can configure BranchCache clients in distributed cache mode. In this mode, BranchCache-enabled client computers cache downloaded content and share it with other clients in the office.

BranchCache content servers

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

- **File server:** Supported file servers include computers running Windows Server 2012 or Windows Server 2008 R2 and have the File Services server role and the BranchCache for Network Files role service installed. These file servers use the SMB protocol to exchange information. After you install the file server, you must also share folders and enable hash generation for shared folders by using Group Policy or Local Computer Policy to enable BranchCache.
- **Web server:** Supported web servers include computers running Windows Server 2012 or Windows Server 2008 R2, those that have the Web Server (IIS 8.0) server role installed, and those that use HTTP or HTTP Secure (HTTPS). In addition, the web server must have the BranchCache feature installed.
- **Application server:** Supported application servers include computers 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 implements a secure-by-design approach that easily works alongside your existing network security architectures, without requiring additional equipment or additional, complex security configurations. BranchCache keeps cached content and content information encrypted and doesn't allow unauthorized access to files in the cache. BranchCache can speed up encrypted communication (HTTPS or IPsec) while helping preserve security.

BranchCache is noninvasive and doesn't alter any Windows authentication or authorization processes. After you deploy BranchCache, authentication is still performed by using domain credentials, and authorization with ACLs is unchanged. In addition, other configurations continue to function just as they did before BranchCache deployment.

The BranchCache security model is based on the creation of content information on the content servers. This metadata, which is much smaller than the size of the actual content it represents, takes the form of a series of hashes.

After content information is created, it's used in BranchCache message exchanges instead of the actual content, and it's exchanged by using the supported protocols (HTTP, HTTPS, and SMB).

Requirements

BranchCache requires:

- One or more content servers running Windows Server 2012
- Client computers running Windows 8
- Optionally, one or more hosted cache servers running Windows Server 2012
- Network connectivity, such as a VPN or DirectAccess connection, between content server and office locations

Scenario

Julia is the IT administrator for a large insurance company headquartered in Columbus, Ohio. The company has five branch offices across North America—three larger offices that have a large, rapidly expanding workforce and two smaller offices that host a handful of remote workers. Julia wants to consolidate most resources at the main office, while still providing remote users with quick, seamless access to files over a WAN. Using Windows Server 2012, she deploys BranchCache with multiple cache servers in hosted cache mode at the three large sites. Because the two smaller offices don't have servers, remote users at those locations use distributed cache mode.

Julia easily accommodates new users because of the scalable performance and single, easy-to-configure GPO for all locations. In addition, users at all branch offices have convenient access to files at the corporate headquarters, all with optimized performance and higher security.

Summary

BranchCache provides remote users with convenient access to files and data over WANs, while eliminating the need to duplicate expensive storage systems in branch offices.

Windows Server 2012 allows you to deploy BranchCache with multiple cache servers in hosted cache mode at your larger remote locations. For smaller offices, you can enable distributed cache mode, which helps optimize access to files for the remote users at those locations. In addition, you can more easily accommodate new users because of the scalable performance and single, easy-to-configure GPO for all locations. Branch office users maintain productivity from the unique caching infrastructure and bandwidth optimizations, while administrators benefit from enhanced scalability and simplified deployment and management.

Conclusion

Windows Server 2012 makes it as straightforward to manage a network as a single server, giving you the reliability and scalability of multiple servers without the costs. Windows Server 2012 automatically reroutes around storage, server, and network failures, keeping file services online with minimal noticeable downtime. It also helps remote users to more easily connect to corporate resources by providing highly available servers and network storage, and by compensating for high latency, low bandwidth, and network congestion.

In summary, the new and enhanced features in Windows Server 2012 help you manage your private clouds efficiently with flexibility, and with easier consolidation capabilities. They also help you link private clouds with public cloud services for efficient multitenancy and more seamless communication. And they help you to connect users more easily to IT resources, regardless of location.

List of charts, tables, and figures