



Microsoft System Center

Integrated Cloud Platform

David Ziembicki
Mitch Tulloch, Series Editor

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

Copyright © 2014 Microsoft Corporation (All)

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

Library of Congress Control Number: 2014935076
ISBN: 978-0-7356-8314-3

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

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

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

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

Acquisitions Editor: Anne Hamilton

Developmental Editor: Karen Szall

Project Editor: Karen Szall

Editorial Production: Diane Kohnen, S4Carlisle Publishing Services

Copyeditor: Andrew Jones

Cover Illustration: Twist Creative • Seattle

Cover Design: Microsoft Press Brand Team

Contents

<i>Introduction</i>	<i>vii</i>
Chapter 1 Hybrid cloud computing and the Microsoft Cloud OS	1
The Microsoft Cloud OS vision	1
Hybrid cloud architectures	2
Chapter 2 Private cloud	5
Software-defined storage	5
Software-defined storage platform	7
Software-defined storage management	11
Additional storage capabilities	13
Cloud-integrated storage	14
Software-defined networking	15
Software-defined network platform	15
Network architecture	19
Software-defined network management	20
Cloud-integrated networking	21
Software-defined compute	22
Software-defined compute platform	22
Software-defined compute management	25
Cloud-integrated compute	26

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

microsoft.com/learning/booksurvey

Software-defined management	26
SQL Server 2012	26
System Center 2012 R2 Virtual Machine Manager	27
System Center 2012 R2 Operations Manager	28
System Center 2012 R2 Service Manager	29
System Center 2012 R2 Data Protection Manager	29
System Center 2012 R2 Orchestrator	29
System Center 2012 R2 App Controller	30
System Center 2012 R2 Windows Azure Pack	30
System Center 2012 R2 Configuration Manager	31
System Center 2012 R2 fabric management architecture	31
Chapter 3 Public cloud	35
Windows Azure overview	35
Windows Azure compute services	36
Windows Azure storage and data services	37
Windows Azure network services	39
Windows Azure application services	39
Extending the datacenter fabric to Windows Azure.	41
Extending the datacenter network to Windows Azure	41
Extending datacenter storage to Windows Azure	44
Extending datacenter compute to Windows Azure	45
Extending datacenter fabric management to Windows Azure	46
Self-Service	46
Updating and update management	47
Monitoring and alerting	48
Orchestration and automation	50
Backup and disaster recovery	51

Chapter 4	Service provider cloud	53
	Cloud OS Network	53
	Extending the datacenter fabric to a service provider.....	54
	Extending the datacenter network to service providers	54
	Extending datacenter storage to service providers	54
	Extending datacenter compute to service providers	55
	Extending datacenter fabric management to a service provider.....	56
	Service Provider Foundation	56
	Windows Azure Pack	59
	System Center 2012 R2	63
	Hyper-V Replica	63
	Conclusion	65

What do you think of this book? We want to hear from you!

Microsoft is interested in hearing your feedback so we can continually improve our books and learning resources for you. To participate in a brief online survey, please visit:

microsoft.com/learning/booksurvey

Introduction

Microsoft System Center: Integrated Cloud Platform is targeted toward IT executives and architects interested in the big picture of how Microsoft's cloud strategy is delivered using Windows and Microsoft System Center. We provide an all-encompassing approach to understanding and architecting Windows Server 2012 R2, System Center 2012 R2, and Windows Azure based solutions for infrastructure as a service. The combination of Windows, System Center, and Windows Azure is a cloud-integrated platform, delivering what Microsoft calls the "Cloud OS," which is a common platform spanning private cloud, public cloud (Windows Azure), and service provider clouds. This platform enables a single virtualization, identity, data, management, and development platform across all three cloud types.

This book is organized by cloud type and we begin with a short overview of the Cloud OS strategy from Microsoft and a high-level hybrid cloud architecture that will be detailed throughout the book. Next we cover the design and deployment of private cloud solutions using Windows and System Center to deliver the software-defined datacenter where storage, network, compute, and management are all virtualized and delivered by the Microsoft platform. We cover some of the substantial cost savings that can be achieved using the Microsoft storage platform, the multi-tenancy enabled by our network virtualization platform, and the consolidation ratios that can be provided by Hyper-V's scalability and high performance.

With a private cloud foundation in place, we next move to the public cloud and detail how to extend the private cloud datacenter (network, storage, compute, management) to Windows Azure while treating it as a seamless extension to your datacenter. Finally, the third cloud type, service provider clouds, are covered using the same approach—extending your datacenter to service providers. The end result is a robust hybrid cloud architecture where consumers of IT within an organization can choose the optimal location to host their virtual machines and services on any of the three cloud types based on which cloud makes the most sense for their workload.

Acknowledgments

This book summarizes the detailed architecture and design work captured in the Infrastructure as a Service (IaaS) reference architecture guides from Microsoft Services. The architectures represent years of lessons learned from our largest and

most complex customer implementations. Contributors to this body of knowledge include: Joel Yoker, Adam Fazio, Artem Pronichkin, Jeff Baker, Michael Lubanski, Robert Larson, Steve Chadly, Alex Lee, Yuri Diogenes, Carlos Mayol Berral, Ricardo Machado, Sacha Narinx, Thomas Ellermann, Aaron Lightle, Ray Maker, TJ Onishile, Ian Nelson, Shai Ofek, Anders Ravnholt, Ryan Sokolowski, Avery Spates, Andrew Weiss, Yuri Diogenes, Michel Luescher, Robert Heringa, Tiberiu Radu, Elena Kozylkova, and Jim Dial.

Errata, updates, & book support

We've made every effort to ensure the accuracy of this book and its companion content. You can access updates to this book—in the form of a list of submitted errata and their related corrections—at:

<http://aka.ms/SCcloudplat/>

If you discover an error that is not already listed, please submit it to us at the same page.

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

Please note that product support for Microsoft software and hardware is not offered through the previous addresses. For help with Microsoft software or hardware, go to <http://support.microsoft.com>.

We want to hear from you

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

<http://aka.ms/tellpress>

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

Stay in touch

Let's keep the conversation going! We're on Twitter:
<http://twitter.com/MicrosoftPress>.

Hybrid cloud computing and the Microsoft Cloud OS

A number of key trends are driving the evolution of information technology (IT) today. New applications requiring global scale, social integration, and mobile capability are critical in many industries. The proliferation of devices such as smart phones and tablets is driving the need for applications and services delivery to nearly everywhere on the globe. The explosion of data and the insight that can be gained from the exponential growth in data is generating demand for enormous storage and analysis capability. These trends have triggered significant changes to how IT must be delivered, resulting in the evolution of cloud computing.

Cloud computing is delivered in many forms such as private cloud in an organization's datacenter, public cloud in a provider such as Microsoft's datacenter, or a multitude of service provider clouds from a range of different organizations. Each provides a different set of features, capabilities, cost points, and service level agreements (SLA).

Within this environment, organizations have a wide range of options for their cloud computing needs and an increasing challenge of how to manage a distributed, cloud-based infrastructure as well as their various applications and services. As a leading provider of on-premises software solutions and one of the largest global cloud providers, Microsoft has created a single integrated cloud platform to meet customer's needs: the Cloud OS.

The Microsoft Cloud OS vision

The Microsoft Cloud OS strategy can be summarized by the following quote from the white paper "Unified Management for the Cloud OS: System Center 2012 R2" published in October 2013:

"The Microsoft vision for a new era of IT provides one consistent platform for infrastructure, applications, and data: the Cloud OS. The Cloud OS spans your datacenter environments, service provider datacenters, and Windows Azure, enabling you to easily and cost-effectively cloud optimize your business."

This strategy is unique in the industry as Microsoft is the only global provider of leading on-premises software for private cloud, large scale public cloud with Windows Azure, and a global service provider ecosystem.

The Cloud OS strategy provides a common identity, virtualization, management, development, and data platform across private cloud, public cloud, and service-provider cloud as illustrated in Figure 1-1.

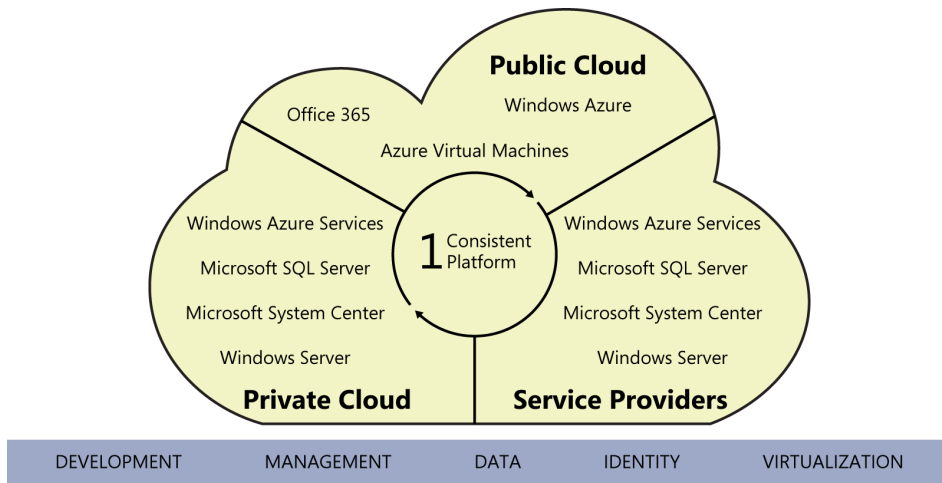


FIGURE 1-1 The Microsoft Cloud OS vision.

The various combinations of private, public, and service provider clouds are commonly referred to as hybrid cloud architectures. The ability to both provide the various types of cloud infrastructure as well as the ability to manage resources across all of them requires an integrated cloud platform such as Microsoft's Cloud OS comprised of Windows Server, Windows Azure, and System Center.

Hybrid cloud architectures

The key attribute of the Cloud OS vision is hybrid cloud architecture, in which customers have the option of leveraging on-premises infrastructure, Windows Azure, or Microsoft hosting-partner infrastructure. The customer IT organization will be both a consumer and provider of services, enabling workload and application development teams to make sourcing selections for services from all three of the possible infrastructures or create solutions that span them.

Starting from the bottom, the diagram in Figure 1-2 illustrates the cloud infrastructure level (public, private, and hosted clouds), the cloud service catalog space, and examples of application scenarios and service-sourcing selections (for example, a workload team determining if it will use virtual machines that are provisioned on-premises, in

Windows Azure, or in a Microsoft hosting partner.) The Cloud OS strategy provides a common identity, virtualization, management, development, and data platform across private cloud, public cloud, and service provider cloud.

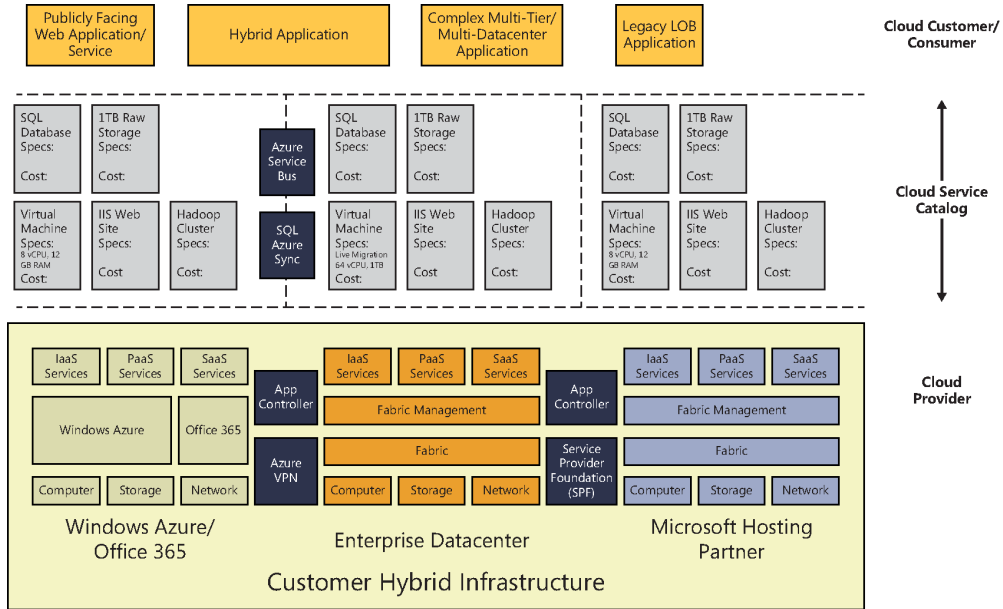


FIGURE 1-2 Hybrid cloud architecture details.

The benefits of this approach are that virtual machines, applications, and services can be hosted on the cloud that makes the most sense for each workload in terms of cost, capability, or SLA. Additionally, the Cloud OS enables “VM Mobility” as all three components (private, public/Azure, service provider) utilize the same underlying Windows Server 2012 R2 and Hyper-V infrastructure meaning that virtual machines can be moved to any of the cloud types without having to convert or modify them. The Cloud OS is an integrated cloud platform where System Center 2012 R2 is able to manage the private cloud as well as virtual machines, applications, and services hosted in Windows Azure or service provider clouds.

In the next several chapters we will outline how to use the Cloud OS to build a software-defined datacenter and private cloud with Windows Server, Hyper-V, and System Center as well as consume Windows Azure and service provider clouds by extending your datacenter and System Center management platform to those clouds. The end result will be a hybrid cloud architecture that enables applications, workloads, and services to be hosted on the cloud that makes the most sense for them while providing an integrated management capability across the hybrid cloud.

Private cloud

In this chapter we'll begin the design of the private cloud portion of the hybrid cloud architecture. The sample design we'll build over the next several chapters is an overview of the detailed architecture provided in the following guides on Microsoft TechNet:

- "Infrastructure as a Service Product Line Architecture - Fabric Architecture Guide" found at <http://aka.ms/iaasfabricarchitecture>
- "Infrastructure as a Service Product Line Architecture - Fabric Management Guide" found at <http://aka.ms/iaasfabricmanagement>
- "Infrastructure as a Service Product Line Architecture - Deployment Guide" found at <http://aka.ms/iaasdeployment>

Software-defined storage

For the purposes of this book, we will build a private cloud architecture consisting of a storage scale-unit, a compute scale-unit, and a network scale-unit which establish a single-rack configuration supporting over 1,000 virtual machines, over half a petabyte of storage, over one million IOPS capacity, and over 40 Gb/s to/from the external LAN. The scale-unit architecture can be expanded with additional racks. The sample architecture is illustrated in Figure 2-1 and in Table 2-1. While this book describes how to build such an architecture, several of the Microsoft OEM partners deliver turn-key solutions using this design approach.

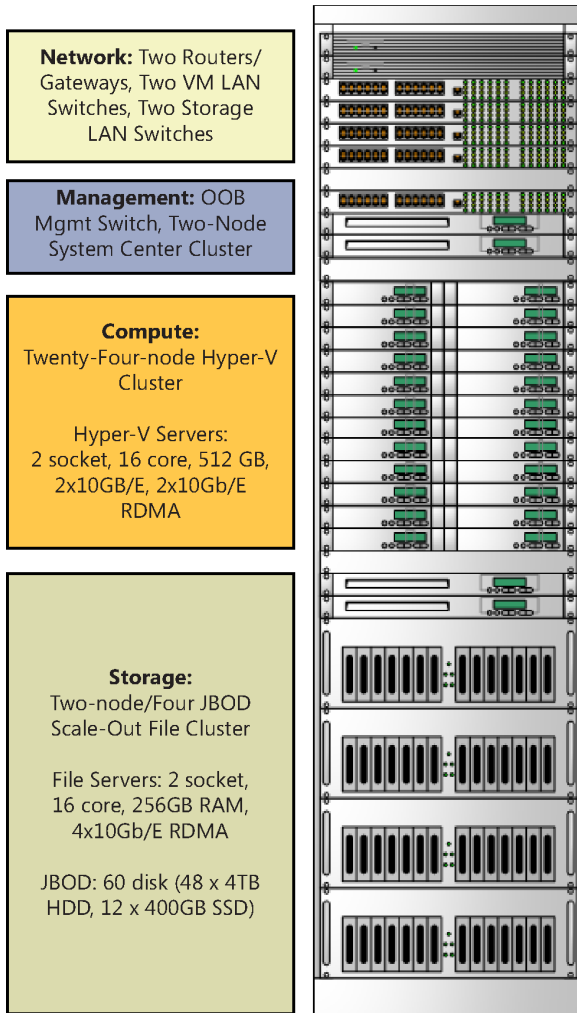


FIGURE 2-1 The sample private cloud architecture used for this book.

TABLE 2-1 Details of Sample Private Cloud Architecture Used for this Book

Functionality	Details
Network	Two Routers/Gateways, Two VM LAN Switches, Two Storage LAN Switches
Management	OOB Mgmt Switch, Two-Node System Center Cluster
Compute	Twenty-Four-node Hyper-V Cluster Hyper-V Servers: 2 socket, 16 core, 512 GB, 2x10GB/E, 2x10Gb/E RDMA
Storage	Two-node / Four JBOD Scale-Out File Server cluster File Servers: 2 socket, 16 core, 256GB RAM, 4x10Gb/E RDMA JBOD: 60 disk (48 x 4TB HDD, 12 x 400Gb SSD)

We'll start with a new approach to enterprise storage called software-defined storage or virtual SAN. In most enterprise datacenters today, storage infrastructure and management is one of the highest cost areas of IT. This is in stark contrast to large cloud providers such as Microsoft which have enormous storage infrastructures which dwarf most enterprises but are far more cost efficient. How is this possible? Through the use of commodity hardware and advanced software where all of the storage "intelligence" is provided not by custom hardware but by software.

Software-defined storage platform

With Windows Server 2012 R2, Microsoft has added substantial software-defined storage capabilities to the platform, enabling customers to establish advanced storage infrastructures at substantially lower costs than traditional hardware-based SAN solutions. Figure 2-2 illustrates the architecture of a software-defined storage solution using Windows Server 2012 R2.




FIGURE 2-2 A sample architecture for a software-defined storage solution based on Windows Server 2012 R2.

While this architecture provides many of the same capabilities as a SAN, it is comprised of the following commodity hardware components:

- **SAS disks** SAS disks provide high performance in throughput and, more importantly, low latency. SAS drives typically have a rotational speed of 10,000 or 15,000 RPM with an average latency of 2 ms to 3 ms and 6 Gbps interfaces. There are also SAS SSDs supporting substantially higher IOPS than spinning disks. SAS disks can support dual interface ports which is required for using clustered storage spaces. The SCSI Trade Association has a range of information about SAS. SAS disks are very common and are available from a wide range of vendors and price points.

- **SAS JBOD** SAS JBOD (“just a bunch of disks”) refers to the disk trays or enclosures where SAS disks are housed. The difference between JBOD and an array or SAN is that a JBOD tray does not have any RAID, storage management, or other intelligence built-in, it is simply a physical component providing SAS connectivity between servers and multiple disks. SAS JBOD typically support 24 or 60 SAS disks in a single enclosure with two to four SAS ports for server connectivity.
- **Windows Server 2012 Scale-out File Servers** In a traditional SAN architecture, most of the functionality and intelligence is provided by the SAN controllers. These are proprietary hardware and software solutions from SAN vendors. In the Microsoft software-defined storage architecture, this functionality is provided by standard server hardware running Windows Server 2012 R2. Just as a SAN controller provides disk resiliency through RAID and advanced features such as tiering and quality of service, the Windows Server 2012 R2 file server infrastructure provides the same capabilities through software combined with commodity server hardware.

With the physical infrastructure in place, the software-defined capabilities of Windows Server 2012 R2 can then be utilized. The Windows Server 2012 R2 platform enables a range of storage virtualization capabilities called Storage Spaces. Storage Spaces enables cost-efficient, highly available, scalable, and flexible storage solutions. Storage Spaces delivers advanced storage virtualization capabilities for single server and scalable multinode cluster deployments.

With Storage Spaces, the Windows storage stack has been enhanced to incorporate two new abstractions:

- **Storage pools** A collection of physical disks that enable you to aggregate disks, expand capacity in a flexible manner, and delegate administration.
- **Storage spaces** Virtual disks created from free space in a storage pool. Storage spaces have such attributes as resiliency level, storage tiers, fixed provisioning, and precise administrative control.

Storage Spaces is manageable through the Windows Storage Management API in Windows Management Instrumentation (WMI), Windows PowerShell, and through the File and Storage Services role in Server Manager. Storage Spaces is completely integrated with failover clustering for high availability, and it is integrated with CSV for scale-out deployments. In addition, System Center 2012 R2 enables full deployment and management of the software-defined storage architecture using Virtual Machine Manager which will be covered in detail later in this chapter.

While the focus of this chapter is the Microsoft software-defined storage architecture using commodity components to achieve extremely cost efficient and high performance virtual machine storage, it is very important to understand that the new Microsoft storage platform is multiprotocol and able to support and enhance heterogeneous storage environments. Windows File server clusters can front-end both Fibre Channel and iSCSI-based SAN environments for customers with existing investments. Additionally, file server clusters based on Windows Server 2012 R2 can present three types of storage: SMB 3.0 file shares, iSCSI targets, and NFS shares as illustrated in Figure 2-3.

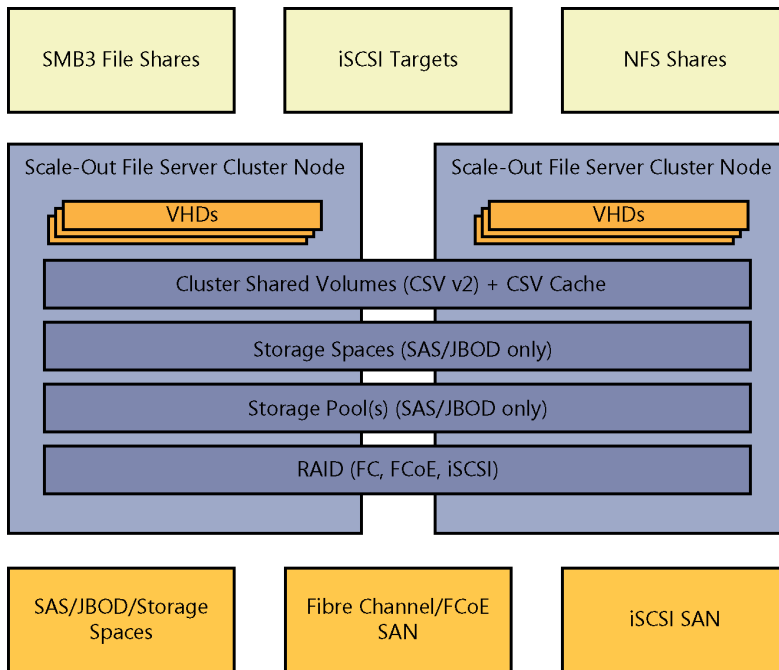


FIGURE 2-3 Supported storage for file server clusters based on Windows Server 2012 R2.

This flexibility allows for a wide range of storage hardware to be utilized and adds significant performance and availability features to each of the supported storage features included in Windows Server 2012 R2, such as:

- Chkdsk enhancements
- CSV v2
- Data deduplication
- Improved NTFS availability
- iSCSI target improvements
- Live storage migration
- NFS improvements
- ODX
- QoS
- ReFS
- SMB application support
- SMB Direct
- SMB multichannel
- SMB scale-out
- SMB transparent failover

- SMB VSS for remote file shares
- Storage spaces
- Storage Tiering
- Thin and trim provisioning
- Virtual fibre channel
- Write Back Cache

An example of the design of the software-defined storage architecture is illustrated in Figure 2-4. Note the SAS disk, JBOD, and Windows file server components. This design details a highlight available architecture using a Scale-out File Server cluster and clustered storage spaces.




FIGURE 2-4 A design for a software-defined storage architecture.

The above design can support a significant number of virtual machines and IOPS using only two file servers and four JBODs. The architecture is a scale-out design meaning additional servers and JBODs can be added in order to support a larger number of virtual machines or applications.

Key factors in sizing and designing the software-defined storage architecture include the number of virtual machines to be hosted, storage capacity, IOPS required, resiliency required, etc. Those requirements then impact the number and types of disk, the ratios of HDD to SSD, how many SAS and Ethernet/RDMA connections per file server, and so on.

The above design approach provides a continuously available infrastructure meaning if you have two or more file servers, three or more JBODs, and redundant network and storage connectivity, any component in the architecture can fail with no downtime of the storage or virtual machines.

While a detailed design is beyond the scope of this book, significant detail is provided in the “Infrastructure as a Service Product Line Architecture” document referred to at the beginning of this chapter. In that document we provide a detailed reference architecture for the software-defined storage approach (as well as designs for non-converged and converged storage architectures).

Software-defined storage management

With a general understanding of the software-defined storage architecture, it becomes clear that there is a significant amount of configuration possibilities as each layer of the architecture such as hardware, operating system, failover clustering, storage spaces, and file server role have a multitude of settings and options available. While all of these are configurable via Windows PowerShell to enable automation, System Center 2012 R2 Virtual Machine Manager (VMM) is able to automate the deployment and management of the software-defined storage architecture.

Using VMM to deploy the software-defined storage architecture begins with ensuring the VMM fabric (library, host groups, network, and storage discovery) is configured. This ensures basic prerequisites such as operating system images and other environment configuration settings are specified. The process for using VMM to deploy Scale-out File Server cluster is documented in detail on Microsoft TechNet (see <http://technet.microsoft.com/en-us/library/gg610634.aspx>) and the following steps are summarized from that article:

1. Perform initial configuration of the physical computers. This includes configuring the basic input/output system (BIOS) to support virtualization, setting the BIOS boot order to boot from a Pre-Boot Execution Environment (PXE)-enabled network adapter as the first device, and configuring the logon credentials and IP address settings for the baseboard management controller on each computer.
2. Create Domain Name System (DNS) entries and Active Directory computer accounts for the computer names that will be provisioned, and allow time for DNS replication to occur. This step is not required, but it is strongly recommended in an environment where you have multiple DNS servers, where DNS replication may take some time.
3. Prepare the PXE server environment, and add the PXE server to VMM management.
4. Add the required resources to the VMM library. These resources include a generalized virtual hard disk with an appropriate operating system that will be used as the base image, and optional driver files to add to the operating system during installation.

5. In the library, create one or more host profiles, or as of Virtual Machine Manager 2012 R2 (VMM), physical computer profile. These profiles include configuration settings, such as the location of the operating system image, and hardware and operating system configuration settings.
6. To create a Hyper-V host, run the Add Resources Wizard to discover the physical computers, to configure settings such as the host group and the host or physical computer profile to use, to configure custom deployment settings, and to start the operating system and Hyper-V deployment.
7. To create a Scale-out File Server cluster (as of System Center 2012 R2 Virtual Machine Manager only), run the Create Clustered File Server Wizard to discover the physical computers, to configure settings such as the cluster name, provisioning type, and discovery scope, and to start the Scale-out File Server cluster deployment.
8. During deployment, the VMM management server restarts the physical computers by issuing "Power Off" and "Power On" commands to the BMC through out-of-band management. When the physical computers restart, the PXE server responds to the boot requests from the physical computers.
9. The physical computers boot from a customized Windows Preinstallation Environment (Windows PE) image on the PXE server. The Windows PE agent prepares the computer, configures the hardware when it is necessary, downloads the operating system image (.vhd or .vhdx file) together with any specified driver files from the library, and applies the drivers to the operating system image.
10. Roles are then enabled as follows:
 - For Hyper-V hosts, the Hyper-V role is enabled.
 - For Scale-out File Servers (as of VMM 2012 R2 only) the Failover Cluster feature and File Server role are enabled. Then, after the cluster is created, the Scale-out File Server role is enabled in the cluster.
11. The computer is then restarted.

To deploy the basic software-defined storage scale unit (the two-node scale-out file cluster illustrated previously), the above procedure would be utilized to configure two bare-metal servers with Windows Server 2012 R2. Those two servers would then be configured by VMM to form a Scale-out File Server cluster using the following steps:

1. Enable the file server role on the computers.
2. Enable the Scale-out File Server role on the cluster.
3. Add the provisioned computers as a Scale-out File Server cluster under VMM management.

The above procedures can also be performed in one process using the Create Clustered File Server Wizard in VMM.

With the above process completed, a new two-node Scale-out File Server cluster is now part of the fabric defined and managed by VMM. VMM will discover all of the physical

storage (SAS JBOD, disks, and so on) attached to the cluster and be able to manage and configure that as well. The process consists of creating a storage pool using some or all of the physical disks available to the cluster. While simple, this part of the setup is critical as your choices of which disks (HDD, SSD, or combination of both) determine the capacity and performance characteristics of the pool you are about to configure.

After the storage pool(s) have been configured, the next step is to create storage spaces, cluster shared volumes, and file shares to present the storage. This also is accomplished in VMM using simple wizards. The Create File Share Wizard will ask you which storage pool you would like to create the share on then it will ask for a critical piece of information, the resilience and redundancy options for the storage space that will be created on the storage pool. The options are:

- **Parity** Allows you to select Single or Dual.
- **Mirror** Allows you to select Two-way or Three-way.

These settings determine the resiliency to disk failure that the storage space can provide. Parity provides better capacity utilization but is not as high performance as mirroring. Dual parity or Three-way mirroring provide higher resiliency as more disks can fail without losing data than Single parity or Two-way mirroring.

With the deployment of SMB 3.0 file shares on the scale-out file cluster, the architecture is now able to present high speed and high availability storage. From bare-metal servers and JBOD in the rack, VMM is able to deploy and configure the complete storage architecture. Advanced features such as tiering, QoS, RDMA, and many others are available. At a cost point far lower than most SANs, this architecture provides an excellent starting point for a virtualized private cloud architecture.

Additional storage capabilities

While we have discussed in some detail software-defined storage and management, System Center also provides robust support for managing SAN and converged storage infrastructures. Many organizations have significant investments in storage that they want to continue to leverage and VMM provides the same management capabilities for physical storage infrastructures as for virtual or software-defined storage infrastructures.

VMM has been enhanced to support managing disparate storage architectures including Fibre Channel and iSCSI SAN. VMM can add and discover external storage arrays that are managed by Storage Management Initiative—Specification (SMI-S) or Store Management Provider (SMP) providers. VMM can also manage virtual Fibre Channel so that an existing Fibre Channel SAN can be utilized by guest virtual machines. Similarly to the SMB 3.0 software-defined storage approach, a significant amount of storage integration and management can be performed in software with VMM in concert with physical storage infrastructure.

Cloud-integrated storage

In addition to on-premises, Windows Server 2012 R2 storage solutions, the Microsoft storage platform also includes cloud-integrated storage using StorSimple.

StorSimple cloud-integrated storage (CiS) provides primary storage, backup, archive, and disaster recovery. Combined with Windows Azure, this hybrid cloud storage solution optimizes total storage costs and data protection for enterprises.

Cloud-integrated storage enables a seamless continuum of storage, comprised of multiple tiers such as local SSD, local HDD, and remote Windows Azure storage with the ability to place data in the most optimal location based on usage and cost. Figure 2-5 illustrates extending the previously described storage architecture comprised of Windows Server 2012 R2 and SAS JBOD storage to include the StorSimple appliance and connectivity to Windows Azure storage for a complete cloud-integrated storage solution with multiple storage tiers.

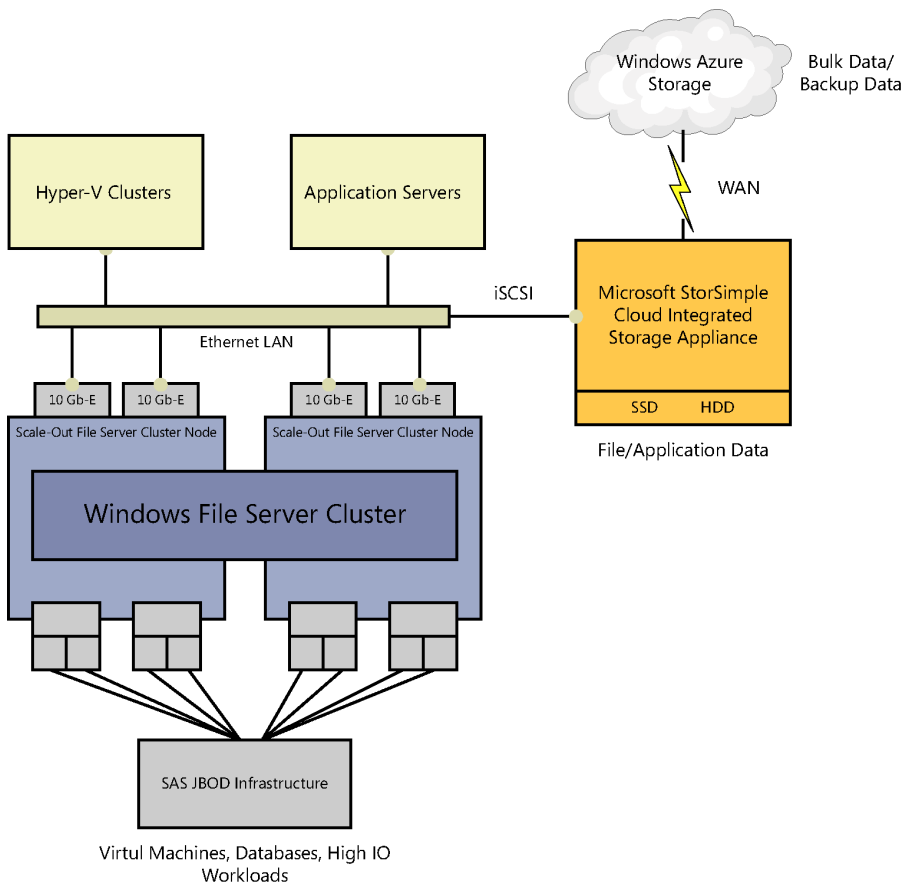


FIGURE 2-5 A storage architecture that includes the StorSimple appliance and connectivity to Windows Azure storage.

A full discussion of hybrid storage is beyond the scope of this book but is the focus of another Microsoft Press e-book titled *Rethinking Enterprise Storage: A Hybrid Cloud Model* (ISBN 9780735679603), by Marc Farley http://blogs.msdn.com/b/microsoft_press/archive/2013/07/26/free-ebook-rethinking-enterprise-storage-a-hybrid-cloud-model.aspx.

Software-defined networking

The concepts of software-defined networking are similar to those of software-defined storage in that the software provides the majority of the intelligence and functionality of the network infrastructure. This can also be described as separating the control plane (how network traffic is routed/processed) from the data plane (the packets and data that flow and traverse the network) and implementing the control plane in software as opposed to hardware (for example, virtual routers instead of physical routers). The benefits are the same as with storage such as increased flexibility and agility in being able to re-configure the network architecture as needs change without having to replace hardware.

Software-defined network platform

As with storage, Windows Server 2012 R2 and System Center 2012 R2 contain large investments in software-defined networking capability. Many of the design requirements were driven by the needs of large enterprises and service providers architecting large scale, multitenant infrastructure as a service (IaaS) solutions. A number of different platform and management capabilities are required to truly deliver a software-defined networking solution.

Hyper-V NIC teaming

From Windows Server 2012 onward, network interface card (NIC) teaming is a built-in feature of the operating system with a simple and easy to use interface for rapidly configuring teaming for highly available network connectivity to hosts and virtual machines. NIC teaming includes several modes and options which can be configured for different design scenarios. Windows Server NIC teaming is the foundation of a software-defined network infrastructure as it ensures that all higher-level networking capabilities are built on a highly available foundation with hosts using two or more network adapters. NIC teaming enables both network high availability as well as bandwidth aggregation.

Hyper-V Virtual Switch

As described on Microsoft TechNet, the Hyper-V Virtual Switch is a software-based layer-2 network switch that is available in Hyper-V Manager when you install the Hyper-V server role. The Hyper-V Virtual Switch includes programmatically managed and extensible capabilities to connect virtual machines to both virtual networks and the physical network. In addition, Hyper-V Virtual Switch provides policy enforcement for security, isolation, and service levels.

With built-in support for Network Device Interface Specification (NDIS) filter drivers and Windows Filtering Platform (WFP) callout drivers, the Hyper-V Virtual Switch enables independent software vendors (ISVs) to create extensible plug-ins (known as Virtual Switch Extensions) that can provide enhanced networking and security capabilities. Virtual Switch Extensions that you add to the Hyper-V Virtual Switch are listed in the Virtual Switch Manager feature of Hyper-V Manager.

Virtual Switch extension types include capturing, filtering, and forwarding extensions which correspond to the types of actions the extensions can take. For example, a capture extension can capture and examine traffic but cannot change it. A filtering extension can make policy decisions such as evaluating firewall rules and determine whether to allow the traffic to pass through the switch or not. Finally, forwarding extensions can forward traffic flow information to an external system such as a virtual appliance for network policy enforcement. An example of a full featured forwarding extension is the Cisco Nexus 1000v solution for Hyper-V.

A diagram of the Hyper-V Virtual Switch architecture, derived from a diagram on Microsoft MSDN ([http://msdn.microsoft.com/en-us/library/windows/hardware/hh582268\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/hardware/hh582268(v=vs.85).aspx)), is illustrated in Figure 2-6.

The Hyper-V Virtual Switch is the key enabling feature for software-defined networking as it exists between the Hyper-V host's physical network connectivity and all of the host's virtual machines. Having a software layer at that point enables the features listed above as well as many others. The extensible design of the switch allows enhancements by Microsoft or partners to add new capabilities.

The features of the Hyper-V Virtual Switch include:

- **ARP/ND Poisoning (spoofing) protection** Provides protection against a malicious VM using Address Resolution Protocol (ARP) spoofing to steal IP addresses from other VMs. Provides protection against attacks that can be launched for IPv6 using Neighbor Discovery (ND) spoofing.
- **DHCP Guard protection** Protects against a malicious VM representing itself as a Dynamic Host Configuration Protocol (DHCP) server for man-in-the-middle attacks.
- **Port ACLs** Provides traffic filtering based on Media Access Control (MAC) or Internet Protocol (IP) addresses/ranges, which enables you to set up virtual network isolation.
- **Trunk mode to a VM** Enables administrators to set up a specific VM as a virtual appliance, and then direct traffic from various VLANs to that VM.
- **Network traffic monitoring** Enables administrators to review traffic that is traversing the network switch.
- **Isolated (private) VLAN** Enables administrators to segregate traffic on multiple VLANs, to more easily establish isolated tenant communities.
- **Bandwidth limit and burst support** Bandwidth minimum guarantees amount of bandwidth reserved. Bandwidth maximum caps the amount of bandwidth a VM can consume.

- **ECN marking support** Explicit Congestion Notification (ECN) marking—also known as Data Center TCP (DCTCP)—enables the physical switch and operating system to regulate traffic flow such that the buffer resources of the switch are not flooded, which results in increased traffic throughput.
- **Diagnostics** Diagnostics allow easy tracing and monitoring of events and packets through the virtual switch.

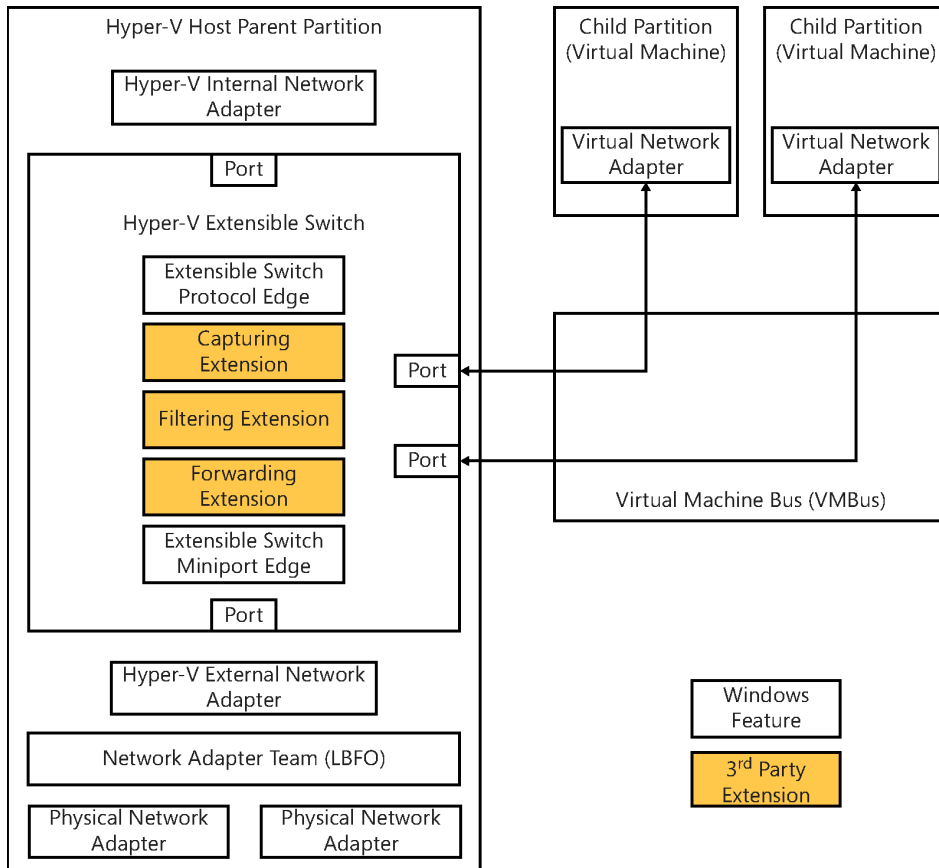


FIGURE 2-6 An example of the Hyper-V Virtual Switch architecture.

The above features can be combined with NIC teaming to enable highly available network access to virtual machines. The security features can be used to ensure that virtual machines that may become compromised are not able to impact other virtual machines through ARP spoofing or DHCP man-in-the-middle attacks. Port ACLs open a wide range of scenarios for protecting virtual machines through access control lists on the virtual switch.

Several of the Hyper-V Virtual Switch features establish the foundation for secure, multitenant environments. Network quality of service (QoS) is enabled through bandwidth limiting and burst support to prevent virtual machines from becoming “noisy neighbors” or

consuming too much host capacity. Private Virtual LANs (PVLANS) enable isolation of virtual machine network traffic.

Hyper-V Network Virtualization

Hyper-V Network Virtualization provides the concept of a virtual network that is independent of the underlying physical network. With this concept of virtual networks, which are composed of one or more virtual subnets, the exact physical location of an IP subnet is decoupled from the virtual network topology. As a result, customers can easily move their subnets to the cloud while preserving their existing IP addresses and topology in the cloud, so that existing services continue to work unaware of the physical location of the subnets. A high-level diagram of a virtualized network environment is illustrated in Figure 2-7.

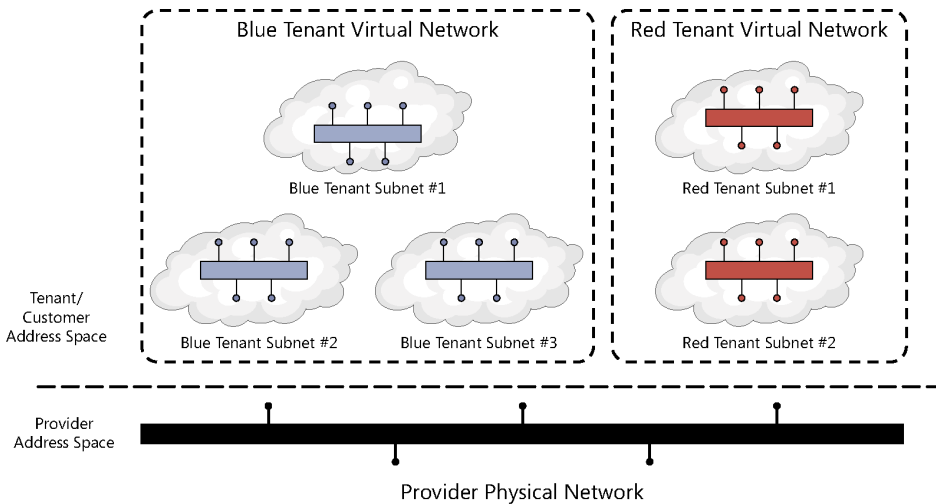


FIGURE 2-7 An example of Hyper-V Network Virtualization.

Hyper-V Network Virtualization in Windows Server 2012 R2 provides policy-based, software-controlled network virtualization that reduces the management overhead. In addition, it provides cloud hosting providers with better flexibility and scalability for managing virtual machines to achieve higher resource utilization.

The details of Hyper-V Network Virtualization are fairly complex and beyond the scope of this book. Several key points for consideration are the separation and isolation of the virtual networks created per tenant, the ability for tenants to “bring their own IP address and subnets,” and the separation of all of the tenant/customer virtual networks from the provider or datacenter physical network infrastructure. The primary value is that changes to the virtual networks, such as creation/modification/deletion, do not require changes to the underlying physical network infrastructure. This capability is in contrast to VLAN-based approaches which often do require changes to physical network infrastructure configuration of switch ports.

Another key consideration with network virtualization is that traffic encapsulation using Network Virtualization for Generic Routing Encapsulation (NVGRE) is the mechanism utilized for virtualizing IP addresses and subnets. The customer or tenant network traffic (from the virtual networks) is encapsulated inside the provider address space packets. This virtualization is what enables the separation between the tenants and the provider. All of the network virtualization functionality works between all Hyper-V hosts in the provider environment, however a key question is how this network virtualization works between Hyper-V hosts and non-virtualized servers or networks outside of the provider datacenter. In those scenarios, the functionality of a network virtualization gateway. The gateway, either a physical or virtual appliance, sits at the edge of the Hyper-V Network Virtualization infrastructure and outside networks, to de-encapsulate outbound virtual network traffic and encapsulate inbound traffic to virtual networks.

Hyper-V Network Virtualization builds on and works with the NIC Teaming and Virtual Switch capabilities described previously to enable the complete software-defined network infrastructure required for today's large scale, multitenant datacenters. These foundational technologies exist in Windows Server 2012 R2, but to enable true software-defined networking, a centralized management capability spanning all participating Hyper-V servers is required.

Network architecture

In the design example illustrated in Figure 2-1 and detailed in Table 2-1, several different physical networks were described. The first is a physical management network for accessing the baseboard management controllers on all of the physical servers as well as management ports on any of the network and/or storage devices. This connectivity is provided by a dedicated physical network switch in the single rack design.

The next two networks are the storage and LAN networks. These can share the same physical infrastructure but in the sample design we illustrate a dedicated pair of switches for the storage network which would host all of the SMB storage traffic between the Hyper-V hosts and the Scale-out File Server clusters.

The second pair of dedicated switches host the LAN traffic between all of the physical servers and virtual machines. These switches can be connected to any other datacenter networks that will be utilizing Hyper-V Network Virtualization.

For connectivity to datacenter networks or external networks not utilizing network virtualization, the Hyper-V Network Virtualization Gateways are utilized. These are illustrated as physical servers or appliances however they can also be deployed as virtual machines by System Center VMM so the two single-rack unit appliances could be dedicated Hyper-V hosts running one or more network virtualization gateways.

You should plan carefully to ensure that there is balance between the storage, LAN, and external network infrastructures. You want to avoid bottlenecks between the external to LAN and the hosted virtual machines and between the virtual machines and the file server-based storage infrastructure. An additional consideration arises in large scale scenarios where

more than one rack scale-unit will be deployed to ensure that there is adequate bandwidth between the physical networks spanning racks.

Software-defined network management

All of the technologies discussed so far in Windows Server 2012 R2 are the basis for a software-defined network and delivered through the Hyper-V hosts in the overall IaaS architecture. System Center 2012 R2 VMM provides the centralized management solution for the IaaS fabric. Similar to software-defined storage where VMM provides deployment and management of storage infrastructure, VMM also provides the management capability for software-defined networking. VMM is used to manage and configure the provider network and establish the tenant virtual networks.

VMM utilizes a relatively complex but very flexible set of abstractions to represent all of the software-defined network elements and constructs. The primary construct is a “logical network” which consists of child objects such as “network sites,” “IP subnet/VLANs,” and “IP Address Pools.” These constructs enable modeling of complex network infrastructures. Logical networks can be assigned to Hyper-V hosts or host groups by VMM so that those hosts and virtual machines are configured to utilize those logical networks. These constructs are utilized to configure the fabric or provider network infrastructure in a VMM fabric.

When utilizing Hyper-V Network Virtualization, VMM is also utilized to configure all of the available virtual networks for each tenant as well as their component subnets and IP address pools. Figure 2-8, adapted from Microsoft TechNet (<http://technet.microsoft.com/en-us/library/jj983727.aspx>), illustrates the relationships between these objects.

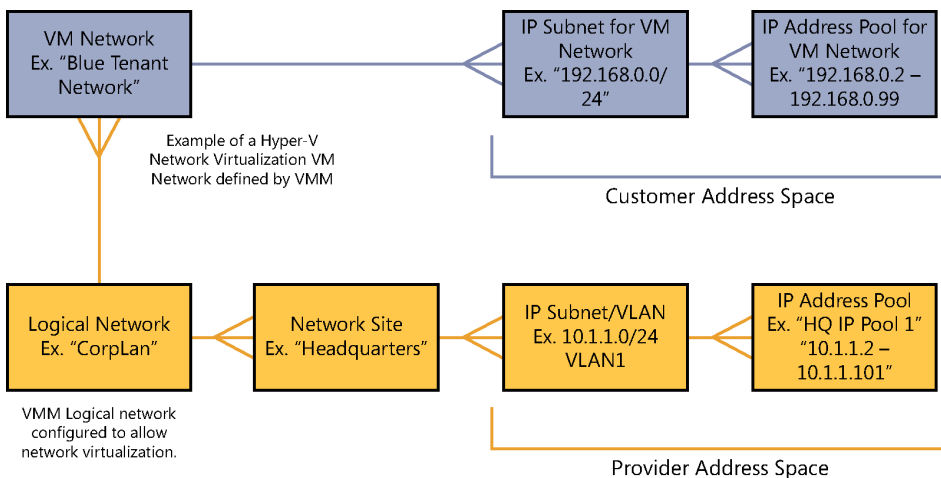


FIGURE 2-8 An example of the VMM networking object model.

The diagram shows the provider address space (the physical infrastructure managed by VMM) and one VM network, which is a tenant or customer address space defined via Hyper-V Network Virtualization. For simplicity only one VM network is shown, in reality there may be

hundreds or more VM networks in scenarios such as a service provider with multiple, isolated VM networks for each of their customers.

In addition to the capabilities discussed so far, VMM is also able to integrate with and manage physical network infrastructure components such as Top of Rack (ToR) switches. VMM can integrate with network equipment that supports Open Management Infrastructure (OMI). In this case, VMM can configure things such as switch port types, trunking, and so on.

A final software-defined networking capability provided by VMM is the automatic deployment and configuration of Hyper-V Network Virtualization gateways. VMM can deploy both stand-alone and highly available pairs of virtual machines acting as network virtualization gateways. This functionality is provided using VMM service templates.

When utilized together, all of the VMM network management capabilities deliver a complete software-defined networking infrastructure, such as the physical fabric network (provider), virtual networks (tenant/customer), and network virtualization gateways that can all be provisioned and managed by VMM.

Cloud-integrated networking

As discussed in Chapter 1, “Hybrid cloud computing and the Microsoft Cloud OS,” the Cloud OS strategy encompasses more than just the private cloud on-premises datacenter by addressing both the public cloud (Windows Azure) and service provider cloud. Cloud-integrated networking refers to extending the on-premises datacenter network to both the public cloud and service providers.

Utilizing a combination of Hyper-V Network Virtualization, network virtualization gateways, and site-to-site VPN between both the private cloud datacenter and the service provider datacenter, an organization can establish a software-defined network that spans both infrastructures. The service provider must enable such functionality using Hyper-V and related components, which is one of the reasons for choosing service providers such as those in the Microsoft Cloud OS Network who utilize the Microsoft platform as the basis of their hosting infrastructure.

In addition to service providers, the private cloud datacenter network can also be extended to Windows Azure utilizing Windows Azure Virtual Network. Using VPN technology, the datacenter network can be extended to Windows Azure using several different methods. The first uses the public Internet as the underlying transport by using VPN gateway devices in the private cloud datacenter configured to connect using VPN to Windows Azure virtual networks. The second method entails working with a Microsoft partner (such as AT&T or Equinix) who enable VPN connectivity to Windows Azure over their private networks (that is, not traversing the public Internet).

Using the above capabilities, a software-defined network spanning private, public, and service provider cloud can be configured using the combination of Windows Server 2012 R2, Windows Azure, and System Center 2012 R2. A high-level view of such an architecture is illustrated in Figure 2-9.

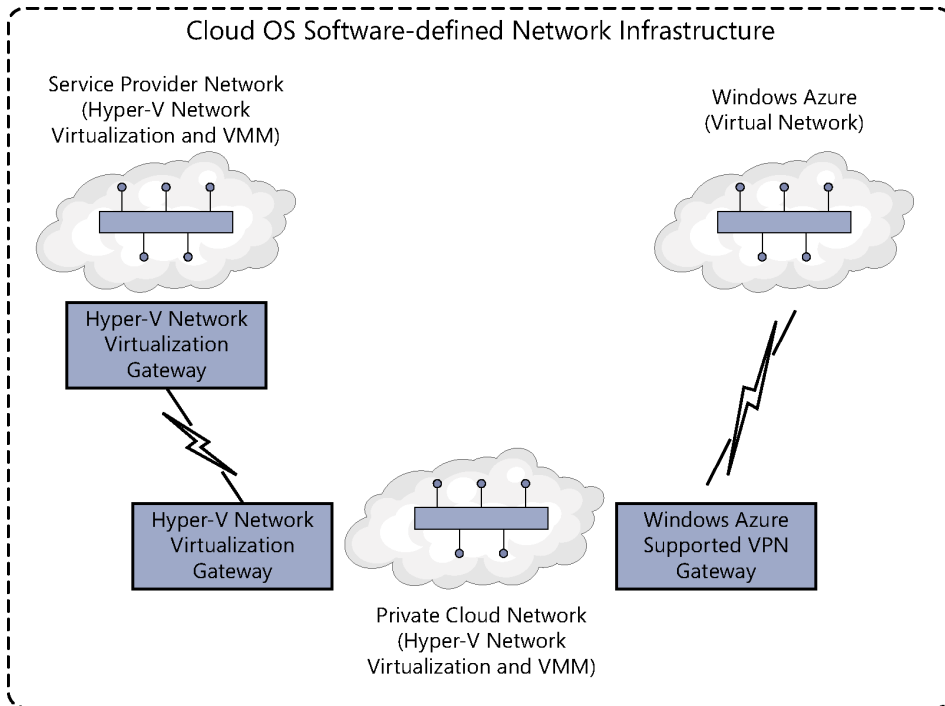


FIGURE 2-9 An example of a Cloud OS software-defined network infrastructure.

The software-defined network infrastructure provides flexibility in where virtual machines and workloads are hosted while enabling connectivity between all workloads regardless of which cloud they are hosted on.

Software-defined compute

Software-defined compute is simply another name for operating system virtualization. As with software-defined networking and storage, the virtualization platform defines the features and capability of the virtualized compute infrastructure in the form of virtual machines. In addition to the consolidation benefits of running multiple virtual machines on a physical server, there are also significant benefits in terms of standardization of host hardware, the ability to live migrate running virtual machines, add resources to running virtual machines, and a number of other capabilities that increase the flexibility and agility of the datacenter.

Software-defined compute platform

Windows Server 2012 R2 and Hyper-V are the software-defined compute platform from Microsoft. From the 2012 wave onward, Hyper-V includes hundreds of new features and capabilities. Hyper-V is the key foundational element of the Microsoft Cloud OS. The Cloud

OS is comprised of private cloud, Windows Azure, and service provider clouds, all of which utilize Hyper-V as the underlying virtualization platform. This enables the concept of software-defined compute to span all three clouds and for virtual machines to be moved from one cloud to another or created on the cloud that is most optimal for the workload.

Several of the largest investments in Hyper-V improvements relate to scalability and availability. Table 2-2 outlines the host, virtual machine, and cluster scalability improvements between Windows Server 2008 R2 and Windows Server 2012 R2.

TABLE 2-2 Hyper-V Scalability

System	Resource	Maximum Number		Improvement Factor
		Windows Server 2008 R2	Windows Server 2012 R2	
Host	Logical processors on hardware	64	320	5×
	Physical memory	1 TB	4 TB	4×
	Virtual processors per host	512	2048	4×
Virtual machine	Virtual processors per virtual machine	4	64	16×
	Memory per virtual machine	64 GB	1 TB	16×
	Virtual disk capacity	2 TB	64 TB	32×
	Active virtual machines	384	1024	4×
Cluster	Nodes	16	64	4×
	Virtual machines	1000	8000	8×

NOTE System Center 2012 R2 Virtual Machine Manager can manage up to 1000 hosts and 25,000 virtual machines.

Those improvement have massive implications for datacenter design and the levels of consolidation now possible. Most organizations that are virtualized still typically run between 15 and 30 virtual machines per physical server. This table shows that should your hardware be capable, Hyper-V will support up to 1,024 virtual machines on a physical host. Effectively, Hyper-V has leapt ahead of the capability of mainstream server hardware. That means that currently the physical server and its cost are the limiting factor, but these improvement open up the distinct possibility of running hundreds of virtual machines per host, offering another round of significant server consolidation potential.

As it relates to the software-defined datacenter, the ability to use larger hosts and larger clusters enables large deployment scale units and resource pools. This provides efficiency of management and capacity as fewer spare nodes or space capacity is required as compared to smaller clusters.

Recalling the reference design illustrated in Figure 2-10, note the Hyper-V cluster in addition to the Scale-out File Server cluster detailed in the section titled “Software-defined storage.”

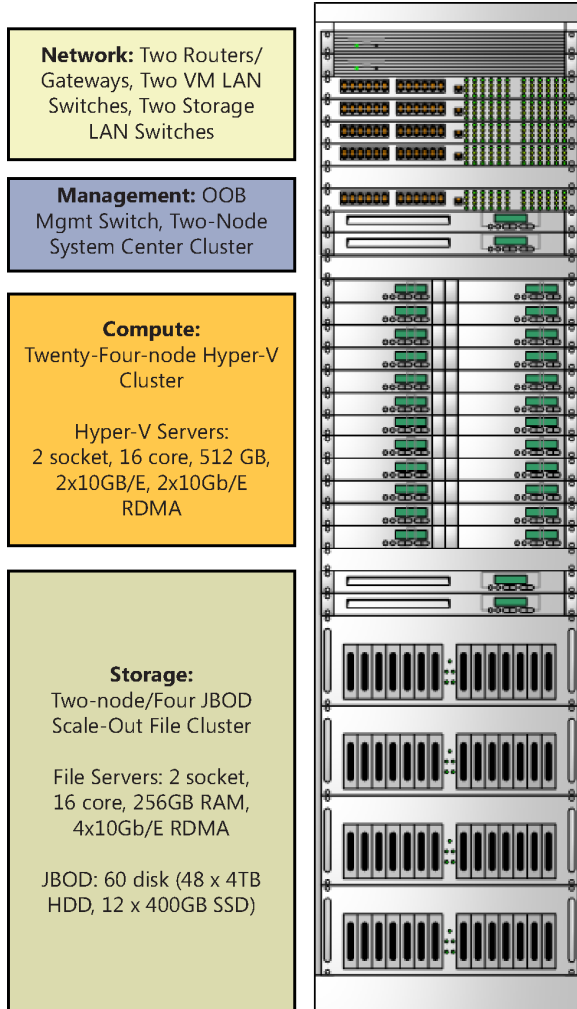


FIGURE 2-10 The software-defined datacenter reference architecture.

In this particular example, a 24-node Hyper-V cluster is illustrated. In reality, the sizing of the Hyper-V cluster depends on a number of different, critical variables such as the desired number and type of virtual machines being hosted, the physical attributes of the host servers (two servers per one rack unit in this example), and the ratio of Hyper-V hosts/virtual machines to scale-out file cluster IO capacity. In this example, we use a 24-node Hyper-V cluster paired with a 4-node scale-out file cluster with four SAS JBOD trays. This is a typical design pattern for a high scale and low cost software-defined datacenter scale unit.

This design pattern should be able to run well in excess of 1,000 virtual machines in a single rack footprint with the ability to scale out as many racks as needed. Further, if you price out such a solution using commodity components as compared to many of the “converged” architectures on the market, this approach can yield substantial cost savings.

A more detailed diagram of the integration between the Hyper-V and file server clusters is illustrated in Figure 2-11.

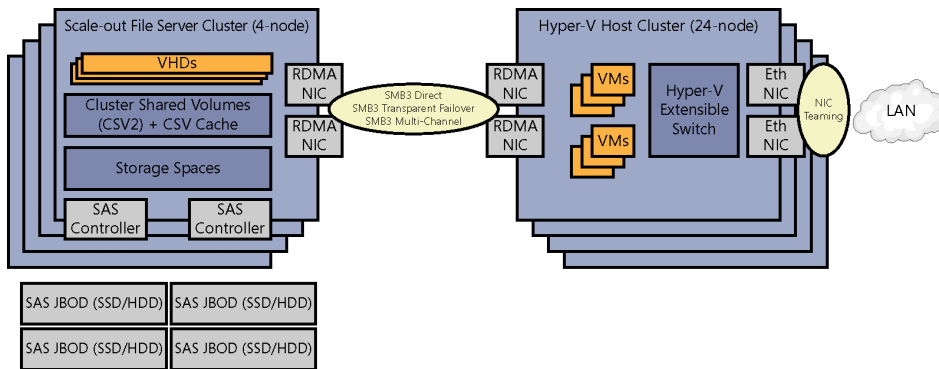


FIGURE 2-11 An example of Scale-out File Server and Hyper-V cluster design.

In this design, the Scale-out File Server cluster and associated SAS JBOD are the software-defined storage infrastructure. The Hyper-V cluster accesses this storage using the SMB3 protocol and a number of associated hardware (RDMA) and software (SMB3 Direct, Multichannel, and Transparent Failover) for very high speed and low latency connectivity to storage. Note on the Hyper-V clusters, the reference architecture utilizes four network adapters, two supporting RDMA for accessing the file cluster and two without RDMA which are teamed for host, cluster, and virtual machine LAN traffic. As mentioned previously, this is an overview of the detailed architecture provided in the “Infrastructure as a Service Product Line Architecture” document referred to at the beginning of this chapter.

Software-defined compute management

System Center 2012 R2, in particular VMM, is the software-defined datacenter management tool from Microsoft. It is complemented by the other components of System Center, all of which are discussed in subsequent sections. A single VMM server is capable of managing up to 1,000 Hyper-V hosts and 25,000 virtual machines.

VMM is able to discover, inventory, and deploy the Windows Server 2012 R2 operating system to physical servers and then create Hyper-V clusters out of them. The process is similar to that outlined previously for the Scale-out File clusters so it will not be repeated here. The end result is that by using VMM the entire software-defined datacenter from storage, to network, to compute can be provisioned and configured using WMM which enables rapid deployment of new physical infrastructure and scale-out capability. The single rack reference architecture illustrated above can all be deployed using VMM.

A key improvement with the 2012 wave of Windows and System Center is that both products are now developed on the same schedule, eliminating the previous months of lag time between a release of Windows and Hyper-V and the corresponding release of System Center to can manage it. In addition, a key design goal of VMM was the ability to manage all of the features delivered in Hyper-V using VMM.

Cloud-integrated compute

Cloud-integrated compute refers to the ability to choose the most appropriate cloud for a unit of compute such as a virtual machine. Any given virtual machine in the Cloud OS concept can be hosted in Hyper-V on-premises using a design, such as the one presented in this book, in your organization's datacenter or it could be hosted in a Cloud OS network service provider's datacenter on Hyper-V or it could be hosted in Windows Azure IaaS, which is also built on Hyper-V. Using System Center 2012 R2 – App Controller, authorized users can provision virtual machines to any connected cloud. In later chapters we'll discuss additional cloud integration capabilities such as extending the datacenter network to Windows Azure using VPN and extending the datacenter network to service providers using network virtualization. These capabilities enable the datacenter to span all three clouds and enable compute, storage, and networking to be consumed from any cloud.

Software-defined management

To this point we have covered the concept of the Cloud OS and detailed software-defined storage, networking, and compute. We briefly discussed the management of those capabilities using System Center, however, in this section we will deal with the topic of managing the private cloud infrastructure in more depth. System Center 2012 R2 is comprised of a suite of components, each focused on part of the infrastructure management lifecycle such as provisioning, monitoring, backup, and disaster recovery.

The Microsoft Press book *Introducing Microsoft System Center 2012 R2* provides a deeper dive on all of the System Center components, so we will primarily cover the software-defined management and cloud integration features of each component, then present an architecture for deploying System Center.

SQL Server 2012

When discussing System Center, we begin with the required Microsoft SQL Server infrastructure underpinning it. Delivering a highly available and well performing System Center is heavily dependent on an associated highly available and high performance SQL infrastructure. Later in this section we'll detail what a best practices implementation looks like, for now realize that SQL server is a key component of the management infrastructure.

System Center 2012 R2 Virtual Machine Manager

System Center VMM is the solution for software-defined and cloud-integrated datacenter management from Microsoft. VMM can establish the datacenter foundation from bare-metal deployment of Scale-out File Server and Hyper-V clusters to applying software updates to those clusters. VMM can integrate with and manage a variety of storage and network infrastructure components. For heterogeneous environments, VMM can manage both VMware and Citrix XenServer environments in addition to Hyper-V.

Virtual Machine Manager can be used to deploy and manage the software-defined datacenter from the datacenter fabric (physical storage, network, and host resource) to the virtual machines and clouds, to the deployment and management of applications and services running in the virtual machines.

Software-defined storage deployment

VMM can be utilized to deploy bare-metal physical servers including pre-boot settings, operating system deployment, and post-deployment configuration. This includes configuration of the File server role which is the pre-cursor to establishing the software-defined storage infrastructure. Once the file servers are provisioned, VMM can create a Scale-out File Server cluster from them and begin the process of configuring storage pools from the attached SAS JBOD storage, then storage spaces, cluster shared volumes, and associated settings to deploy the complete software-defined storage infrastructure.

Software-defined compute deployment

With the storage infrastructure in place, VMM can be utilized to deploy bare-metal servers and configure them to be Hyper-V hosts, then form Hyper-V host failover clusters from the deployed servers. During this process, the Hyper-V clusters are configured to utilize the VMM deployed and managed storage infrastructure as the highly available storage for the Hyper-V clusters leveraging the full set of SMB3 capabilities discussed in the : “Software-defined storage” section. Using VMM for deploying both the Scale-out File Server and Hyper-V clusters significantly reduces the time required and increases the consistency of the deployment when compared to the long list of configuration steps that would have to be performed identically on all the nodes if done manually. VMM also enables rapid scale-out by adding additional clusters or nodes using the same automation when needed.

Software-defined network deployment

With the software-defined storage and compute foundation in place, the final part of fabric deployment can be performed which is establishing the software-defined network infrastructure. This entails creating the appropriate port profiles, logical switches, and virtual networks as described in the “Software-defined networking” section. This step might also include adding third-party extensions to the Hyper-V virtual switches in the host

infrastructure or configuring any number of capabilities such as NIC teaming, QoS, port ACLs, and other settings.

Another critical deployment step for the software-defined network is the deployment of network virtualization gateways, also fully automated by VMM, to enable connectivity to and from the isolated virtual networks created in the infrastructure. VMM includes service templates which assist in automatically deploying virtual machines to perform the network virtualization gateway functionality.

VMM also can manage IP addressing, both static and dynamic, or it can integrate with the IP Address Management (IPAM) capability of Windows Server 2012 R2.

Software-defined management

With the storage, network, and compute fabric deployed, VMM provides a number of additional capabilities. From a fabric perspective, VMM supports on-demand compliance scanning and updating of the fabric. VMM can monitor the update status of the fabric servers, scan for compliance, and apply updates for selected servers.

VMM supports automated updates of Hyper-V host clusters. When VMM performs update remediation on a host cluster, VMM places one cluster node at a time in maintenance mode and then installs updates. If the cluster supports live migration, intelligent placement is used to migrate virtual machines off of the cluster node.

One of the primary benefits of a software-defined datacenter is the ability to optimize the usage of infrastructure from a capacity and power perspective dynamically. An example of this is the Dynamic Optimization and Power Optimization features in VMM. With Dynamic Optimization, VMM live migrates virtual machines within a host cluster to improve load balancing among hosts and to correct any placement constraints for virtual machines to optimize the cluster based on policies configured by the administrator. With Power Optimization, VMM helps manage energy efficiency by turning off hosts in a cluster that are not needed to meet resource requirements and turns the hosts back on when they are needed again.

Beyond managing the fabric, VMM is also the foundation of application and service deployment, including complex multi-tier services consisting of many virtual machines. VMM can deploy individual virtual machines, VM roles which as single tier, scale-out constructs of one or more identical VMs, such as a web farm, and service templates which are n-tier models of complex applications or services.

System Center 2012 R2 Operations Manager

System Center Operations Manager is the monitoring and alerting component of System Center covering physical, virtual, and applications/service resources. Operations Manager is a key component of software-defined datacenter management as it provides a view of the entire physical and virtual infrastructure. In recent versions, Operations Manager has expanded to support monitoring Linux systems as well as network and storage devices.

Operations Manager continues to be extended by a wide range of partners through management packs. From an IT process automation perspective, Operations Manager is frequently the source of alerts and events which are the triggers for process automation or Orchestrator runbooks. Examples include a performance alert triggering a runbook to scale out a web farm, or a hardware fault triggering a runbook to place a Hyper-V host into maintenance mode.

Operations Manager also delivers cloud-integrated management capability as it includes robust support for monitoring resources deployed in the public cloud including both Windows Azure and Amazon Web Services. The heterogeneous monitoring capability spanning both private and public clouds is a key differentiator and pre-requisite for the software-defined datacenter.

System Center 2012 R2 Service Manager

System Center Service Manager deals with the ITIL-based service management and human workflow side of process automation. Until Service Manager was released, System Center had long been missing a centralized configuration management database (CMDB) consolidating all of the discovered inventory and configuration information from the entire System Center suite—from devices inventoried by Configuration Manager to users from Active Directory to virtual resources from VMM. Service Manager implements ITIL-based service management processes, such as Incident and Change Management, by enabling a human workflow engine for topics such as help desk ticketing, approvals, and routing. Service Manager includes a customizable self-service portal and extensible service catalog.

System Center 2012 R2 Data Protection Manager

System Center Data Protection Manager (DPM) provides backup and disaster recovery functionality for Microsoft applications and services. From backing up data or Microsoft applications such as SharePoint or SQL Server to recovery services in an alternate site, DPM is designed to provide a cost-efficient solution for backup and disaster recovery. DPM is also evolving to be a cloud-integrated backup solution through the ability to utilize Windows Azure storage as the target for backups.

System Center 2012 R2 Orchestrator

System Center Orchestrator adds a workflow engine, authoring experience, and execution infrastructure for runbooks, which are instances of IT process automation. While each System Center component discussed in this chapter includes automation of certain processes, they typically deal with only part of the management lifecycle. For processes which need to span the lifecycle, or which need to integrate with multiple System Center or third-party systems, Orchestrator is an excellent solution.

System Center 2012 R2 App Controller

App Controller provides a common self-service experience to configure, deploy, and manage virtual machines and services across private and public clouds. App Controller has the ability to connect to VMM-based private clouds (consisting of Hyper-V, VMware, or Xen), Windows Azure, and service provider clouds through Service Provider Foundation (SPF) and VMM running at the service provider (which will be described in later in this book). App Controller provides valuable functionality for certain use cases but the clear direction of the System Center suite from a self-service point of view is the Windows Azure Pack.

System Center 2012 R2 Windows Azure Pack

The Windows Azure Pack integrates with System Center and Windows Server to help provide a self-service portal for managing services such as websites, virtual machines, and service bus. Windows Azure Pack also provides a portal for administrators to manage resource clouds, scalable web hosting, and more. The diagram in Figure 2-12 illustrates the high-level Windows Azure Pack conceptual architecture.

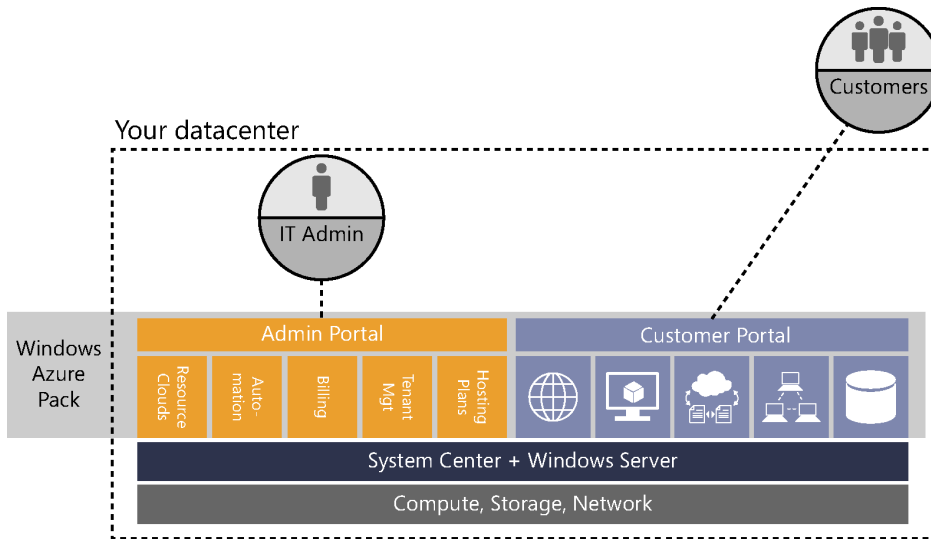


FIGURE 2-12 The Windows Azure Pack conceptual architecture.

Windows Azure Pack is a critical piece of the Cloud OS as it provides user interface and API consistency between Windows Azure (public cloud) and private cloud or service provider clouds. This provides a common user experience for consumers of the Cloud OS regardless of where their virtual machines, websites, and services are deployed.

The Windows Azure Pack can be utilized by either enterprises wishing to deploy a robust self-service capability for their private cloud infrastructure or by service providers looking to enable self-service for their commercially hosted services.

In addition to the self-service portal and SPF APIs, Windows Azure Pack also includes Service Management Automation (SMA). Service Management Automation is a set of tools that is integrated as the Automation extension in Windows Azure Pack. Administrators and developers can use SMA to construct, run, and manage runbooks to integrate, orchestrate, and automate IT business processes. SMA runbooks utilize the Windows PowerShell workflow engine.

A frequently asked question is the relationship and seeming overlap between Orchestrator and SMA. Both have appropriate uses in the R2 wave (Orchestrator for integration across disparate management systems and SMA for all other automation using Windows PowerShell) and is part of the overall evolution of the orchestration capability from Microsoft which will continue to be expanded in SMA. So for all new automation activities that can be performed with SMA, that is the recommended path. For those that cannot be achieved with SMA, Orchestrator is the recommended solution.

System Center 2012 R2 Configuration Manager

System Center Configuration Manager provides client device and application management. From deployment of desktops and devices to managing application delivery and virtualization, Configuration Manager is a key component of an enterprise management infrastructure. Configuration Manager is primarily a device and application management platform but still provides functionality in terms of software and operating system deployment that may be required in some datacenter scenarios.

System Center 2012 R2 fabric management architecture

A software-defined datacenter fabric is comprised of storage, network, and compute as we have seen in previous sections. Fabric management is provided by System Center and therefore a robust architecture for both SQL Server and System Center is required for a highly available fabric management capability. In an IaaS design, either for private cloud or service provider cloud, high availability of the management infrastructure is required as that management infrastructure is what the self-service capability utilizes to provide services to consumers.

The previous section just scratched the surface on the capability of System Center. As the suite has grown in capability, it has also grown in complexity of deployment. Just as an enterprise resource planning (ERP) suite may run an organization's entire set of business processes and therefore requires significant design and implementation planning, System Center is capable of running an organization's entire Cloud OS and software-defined datacenter infrastructure, therefore it also warrants significant design and implementation planning.

Recalling the reference architecture rack diagram illustrated earlier in Figure 2-1, a third cluster (in addition to the Scale-out File Server and Hyper-V host cluster) is included as part of the design. This cluster is the fabric management cluster, a set of Hyper-V hosts (typically

two to four nodes) dedicate to running the SQL Server and System Center infrastructure required for software-defined datacenter management. A frequent question is why the recommendation of having a dedicated fabric management cluster, why not run the SQL and System Center virtual machines alongside the workload virtual machines on the fabric Hyper-V cluster(s)? There are several reasons for the separation, the primary reason being predictable performance and high availability. Having the separation of fabric management from fabric ensures that there is dedicated capacity for the management infrastructure so that it remains available and high performance regardless of the amount of utilization of the fabric cluster. This ensures that should workloads on the fabric start consuming all of the available fabric capacity, the fabric management infrastructure, with its own dedicated capacity, will be able to monitor and react to the surge in usage. If fabric management was co-located with the fabric, performance degradation or competition might occur. Given these and other considerations, our strong recommendation and reference architecture specify a dedicated Hyper-V host cluster for the fabric management deployment.

Using the dedicated fabric management cluster, the reference architecture utilizes a fully virtualized SQL Server guest cluster as the basis for all of the required System Center database functionality. As mentioned at the beginning of this section, a high performance and high availability SQL Server foundation is absolutely critical for a robust deployment of System Center for fabric management.

A key aspect of the reference architecture detailed below is full virtualization and scale-out design. What that means is that fabric management is deployed initially in the smallest footprint possible (though still quite a large set of virtual machines) based on the expected capacity of the fabric to be managed. Since elasticity is a key cloud attribute, the fully virtualized design of the fabric management infrastructure enables each major part (SQL guest cluster, System Center components, and Windows Azure Pack) to be scale-out independently by adding additional virtual machines. Similarly, should the fabric management cluster require more than two nodes to achieve the appropriate performance, it also is just matter of adding additional nodes to the fabric management cluster. In all cases, this scale out can be performed with minimal downtime to any of the fabric management components.

Fabric management SQL Server design

In this section, we present an overview of the SQL Server 2012 design for fabric management which is captured in detail in the “IaaS Product Line Architecture (PLA) Fabric Management Architecture Guide” on Microsoft TechNet at <http://aka.ms/iaasfabricmanagement>. The SQL design is the output of all of the recommended and best practices for both SQL cluster design and each System Center component’s requirements. The design assumes full implementation of all System Center features except for Configuration Manager and Data Protection Manager which are optional components.

The design leverages a SQL Server 2012 guest cluster and multiple SQL Server instances within the guest cluster to follow either best practices for separation (such as database engine from analysis service) or constraints (such as scale-out and scale-up) where instances can be managed individually and be distributed between the guest cluster nodes. The end result is a very complex design, however, one which has been validated across both the relevant product groups and Microsoft Consulting Services, and Premier Support as our standard reference architecture for deploying SQL and System Center to provide highly available IaaS capability. The required and optional SQL instances are illustrated in Figure 2-13.

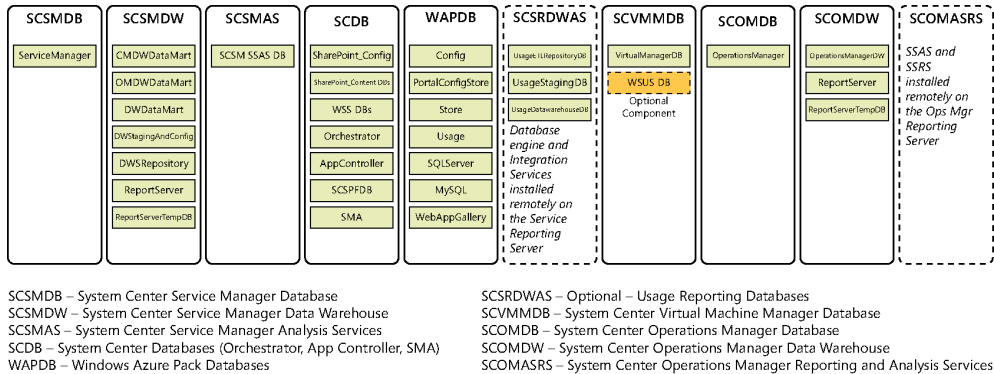


FIGURE 2-13 An example of fabric management SQL design.

As mentioned, the SQL instances are all hosted in a SQL guest cluster running on the fabric management host cluster. Each database instance minimally requires two LUNs, Shared VHDX, or SMB3 file shares for database and log storage. More advanced designs for larger scale may use three or more for each instance. The detailed design of the SQL guest cluster is included in the “Infrastructure as a Service Product Line Architecture - Fabric Management Guide” mentioned previously. There are several options for the shared storage required for the SQL guest cluster including iSCSI, virtual fiber channel, Shared VHDX, and SMB3 file shares. A detailed discussion of these options is beyond the scope of this book but covered in detail in the PLA.

Fabric Management System Center Design

With the underlying SQL Server architecture defined, the deployment of the System Center 2012 R2 suite can be designed. Like the SQL guest cluster, each of the System Center components deployed are deployed using a high availability design using either highly available VMs, guest clustering, or redundant/load-balanced virtual machines using application level high availability. The diagram in Figure 2-14 illustrates the smallest footprint design of the fabric management architecture. A second design pattern supporting larger scale is detailed in the “Infrastructure as a Service Product Line Architecture - Fabric Architecture Guide” found on Microsoft TechNet at <http://aka.ms/iaasfabricarchitecture>.

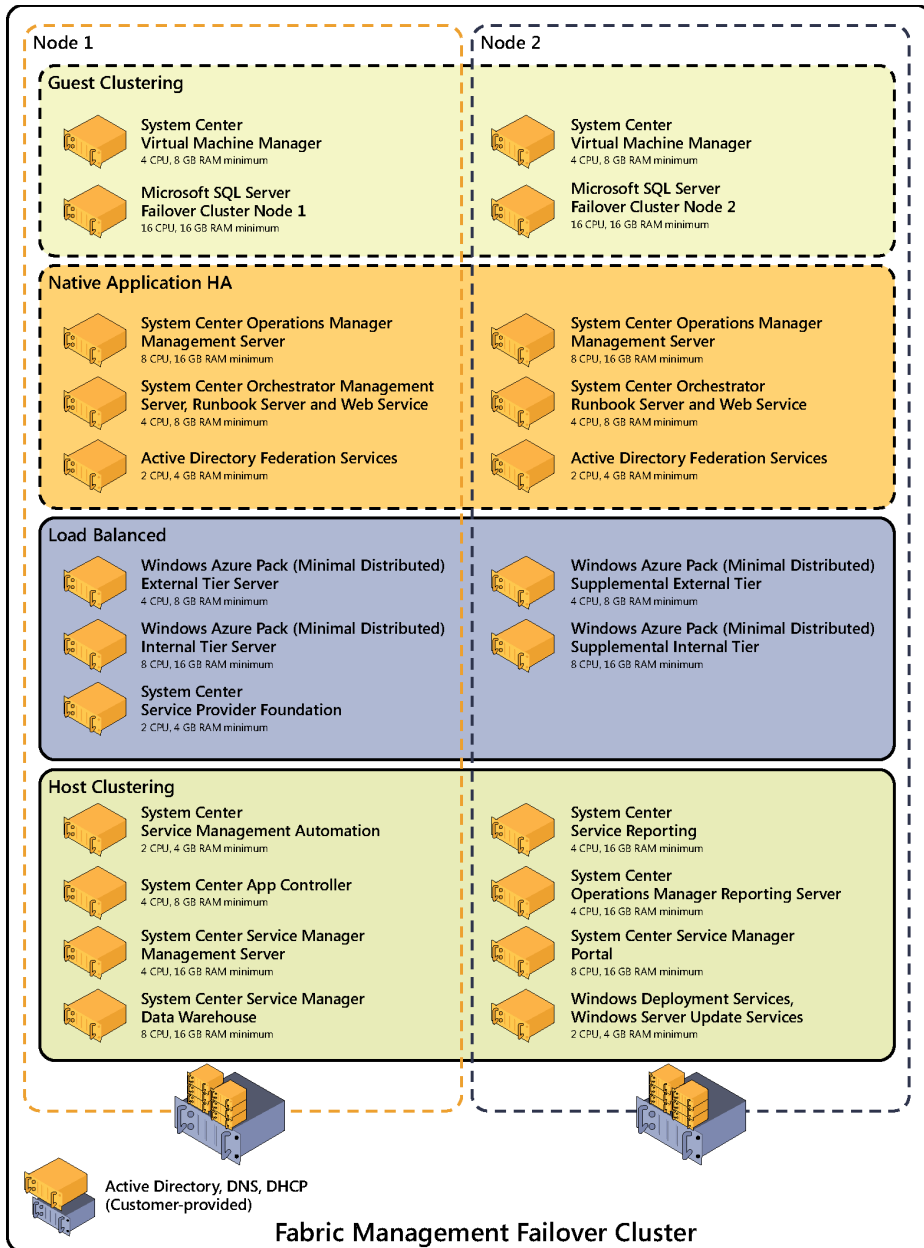


FIGURE 2-14 The Fabric Management Failover Cluster and System Center design.

The fabric management architecture is complicated, however, when considering the wide range of management capability provided by System Center, the inclusion of Windows Azure Pack, and the ability of a two to four physical node fabric management cluster to manage thousands of virtual machines, the complexity is put in some perspective.

Public cloud

*P*ublic cloud refers to large scale cloud services delivered by an organization such as Microsoft. Unlike private cloud, where organizations own and operate the infrastructure, public cloud enables organizations to consume cloud services and capacity on-demand. Microsoft provides a large portfolio of both consumer and enterprise cloud services such as Office 365 and Windows Azure. In a hybrid cloud infrastructure, Windows Azure delivers the public cloud infrastructure and Platform as a Service (PaaS) capabilities required for a robust hybrid infrastructure.

Windows Azure overview

Windows Azure is the public cloud solution from Microsoft for Infrastructure and Platform as a Service. Windows Azure is one of the largest investments in the history of Microsoft considering the massive datacenter, compute, storage, and network capacity in addition to research and development of the various Windows Azure services.

Windows Azure is a global service hosted in a Microsoft world class datacenter infrastructure. Many of the Windows Azure services provide a financially backed service level agreement (SLA) and all Windows Azure services use a pay for consumption model where the user is billed based on how much capacity they utilize. Windows Azure enables an organization to host their workloads and applications in the cloud while also connecting to on-premises resources in a hybrid cloud model.

Windows Azure is built using the same Windows Server and Hyper-V foundation as the Microsoft private cloud solution described in the previous chapter. This foundation enables virtual machine portability between the private cloud and the public cloud. Adding Windows Azure to the hybrid cloud infrastructure provides an effectively unlimited amount of capacity distributed across geographically separated datacenters. For both large and small organizations, this can be a significant benefit.

This section will provide a brief overview of all of the major Windows Azure services so you have awareness of the large and growing set of capabilities in Windows Azure. Additional information can be found at <http://www.windowsazure.com>.

Windows Azure compute services

Windows Azure currently includes the following compute services:

- Virtual Machines
- Web Sites
- Mobile Services
- Cloud Services

A summary of each of these is derived from Windows Azure documentation (<http://www.windowsazure.com/en-us/documentation/>).

Virtual Machines

Windows Azure provides a wide range of Infrastructure as a Service (IaaS) features such as virtual machines, storage, and network resources. Creating a new virtual machine (or many virtual machines) typically takes no longer than five minutes and is performed via the Windows Azure portal or through REST APIs or Windows PowerShell. Windows Azure IaaS virtual machines are offered with the specifications listed in Table 3-1 with a correspondingly higher price for virtual machines with more cores or RAM.

TABLE 3-1 Windows Azure virtual machine sizes

Compute Instance Name	Virtual Cores	RAM
Extra Small (A0)	Shared	768 MB
Small (A1)	1	1.75 GB
Medium (A2)	2	3.5 GB
Large (A3)	4	7 GB
Extra Large (A4)	8	14 GB
A5	2	14 GB
A6	4	28 GB
A7	8	56 GB

Windows Azure supports a wide range of virtual machine operating systems including Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2. Windows Azure also supports Linux virtual machines including Ubuntu, CentOS, Suse, and Oracle Linux. In addition to virtual machines with just an operating system, Windows Azure also provides virtual machines with applications such as SQL, Microsoft SharePoint, as well as Oracle database and other applications. Finally, Windows Azure also allows users to upload their own custom virtual machine images (such as a reference virtual machine created in Hyper-V and uploaded to Azure).

Web Sites

Windows Azure Web Sites allow rapid deployment of web applications and integration with various Microsoft and third-party or open source development frameworks. Windows Azure Web Sites are elastic and scalable with the ability to scale out a web application to additional virtual machines on demand or automatically based on autoscaling policies. Windows Azure includes a number of different web applications (such as blog/CMS platforms, development frameworks) which can be deployed into Windows Azure Web Sites from the gallery.

Mobile Services

Mobile Services enables mobile application development by providing features to structure storage, authenticate users, and send push notifications. Mobile Services provides SDKs for Windows, Android, iOS, and HTML as well as a flexible REST API. Mobile Services lets you to build connected applications for any platform and deliver a consistent experience across devices.

Cloud Services

Windows Azure Cloud Services enables rapid deployment of highly available web applications. Rather than creating and uploading virtual machines, with Cloud Services you upload your application and Windows Azure executes the deployment details such as provisioning, load balancing, and health monitoring. Cloud Services are key to the Windows Azure availability model that underpin several of the Windows Azure SLAs.

Windows Azure storage and data services

Windows Azure currently includes the following storage and data services:

- Storage
- SQL Database
- HDInsight
- Cache
- Backup
- Recovery Manager

A summary of each of these is derived from Windows Azure documentation (<http://www.windowsazure.com/en-us/documentation/>)

Storage

Windows Azure Storage provides a robust, distributed storage architecture for data and virtual machine storage. Windows Azure Storage provides three storage constructs: blobs, queues, and tables. Blobs store unstructured binary and text data. Queues store messages that a client can access. Tables store nonrelational structured data. For Windows Azure virtual

machines, unlike on-premises virtual machines where the virtual machine's VHD file is stored on a disk or LUN, a virtual machine's VHD file is stored in Windows Azure blob storage which is an extremely high availability service where each blob is replicated to three locations within one datacenter and three locations in a geographically separate datacenter by default.

SQL Database

Windows Azure SQL Database is a fully managed relational database service that delivers flexible manageability, includes built-in high availability, offers predictable performance, and supports massive scale-out. With Windows Azure SQL Database, developers have direct access to a managed SQL capability without have to create and maintain virtual machines running SQL server.

HDInsight

HDInsight is a Hadoop-based service from Microsoft that brings a 100 percent Apache Hadoop solution to the cloud. This platform manages data of any type, whether structured or unstructured, and of any size. With HDInsight you can seamlessly process data of all types through the Microsoft data platform, which provides simplicity and ease of management. You can analyze Hadoop data with PowerPivot, Power View, and other Microsoft Business Intelligence (BI) tools through integration with Microsoft data platform.

Cache

Windows Azure Cache is a distributed, in-memory, scalable solution that enables developers to build highly scalable and responsive applications by providing super-fast access to data.

Backup

Windows Azure Backup helps you protect important server data off-site with automated backup to Windows Azure. Backups are encrypted before transmission and stored encrypted in Windows Azure. These backups are off-site protected by reliable Windows Azure storage, reducing the need to secure and protect on-site backup media. Cloud backups can be managed from the backup tools in Windows Server, Windows Server Essentials, or System Center Data Protection Manager.

Recovery Manager

Windows Azure Hyper-V Recovery Manager can help protect important services by coordinating the replication and recovery of Hyper-V and System Center 2012 R2 private clouds at a secondary location.

System Center 2012 Virtual Machine Manager (VMM) clouds can be protected through automating the replication of the virtual machines that compose them at a secondary location. The ongoing asynchronous replication of each VM is provided by Windows Server 2012 Hyper-V Replica and is monitored and coordinated by Hyper-V Recovery Manager.

Windows Azure network services

Windows Azure currently includes the following network services:

- Virtual Network
- Traffic Manager

A summary of each of these is derived from Windows Azure documentation (<http://www.windowsazure.com/en-us/documentation/>).

Virtual Network

Windows Azure Virtual Network enables you to create a logically isolated section in Windows Azure and securely connect it to an on-premises datacenter or a single client machine using an IPsec connection. Virtual Network makes it easy for you to take advantage of scalable, on-demand infrastructure of Windows Azure while providing connectivity to data and applications on-premises.

Windows Azure virtual machines can take advantage of a number of advanced networking capabilities such as isolated virtual networks per subscription, virtual private network (VPN) connectivity between an on-premises datacenter network and Windows Azure, as well as a number of other features such as load balancing, DHCP, port ACLs, and many others.

Windows Azure IaaS provides an easy on ramp to public cloud by supporting a wide range of virtual machines and workloads that can be moved from on-premises hosting to Windows Azure.

Traffic Manager

Traffic Manager allows you to load balance incoming traffic across multiple hosted Windows Azure services whether they're running in the same datacenter or across different datacenters around the world.

Windows Azure application services

Windows Azure currently includes the following application services:

- Active Directory
- Media Services
- Content Delivery Network
- Service Bus
- Multi-Factor Authentication
- Scheduler
- Notification Hubs
- Visual Studio Online
- BizTalk Services

A summary of each of these is derived from Windows Azure documentation (<http://www.windowsazure.com/en-us/documentation/>).

Active Directory

Windows Azure Active Directory is a comprehensive identity and access management cloud solution. It combines core directory services, advanced identity governance, security, and application access management. Windows Azure Active Directory also offers developers an identity management platform to deliver access control to their applications, based on centralized policy and rules. For enterprises with more demanding needs, an advanced offering, Windows Azure Active Directory Premium, helps complete the set of capabilities that this identity and access management solution delivers.

Media Services

Media Services offer the flexibility, scalability, and reliability of a cloud platform to handle high quality media experiences for a global audience. Media Services include cloud-based versions of many existing technologies from the Microsoft Media Platform and our media partners, including ingest, encoding, format conversion, content protection, and both on-demand and live streaming capabilities.

Content Delivery Network

The Windows Azure Content Delivery Network (CDN) offers developers a global solution for delivering high-bandwidth content by caching blobs and static content of compute instances at physical nodes in the United States, Europe, Asia, Australia, and South America.

Service Bus

Windows Azure Service Bus provides the messaging channel for connecting your cloud applications to your on-premises applications, services, and systems.

Multi-Factor Authentication

Windows Azure Multi-Factor Authentication reduces organizational risk and helps enable regulatory compliance by providing an extra layer of authentication, in addition to a user's account credentials, to secure employee, customer, and partner access. Windows Azure Multi-Factor Authentication can be used for both on-premises and cloud applications.

Scheduler

Windows Azure Scheduler allows you to invoke actions—such as calling HTTP/S endpoints or posting a message to a storage queue—on any schedule. With Scheduler, you create jobs in the cloud that reliably call services both inside and outside of Windows Azure and run those jobs on demand, on a regularly recurring schedule, or designate them for a future date.

Notification Hubs

Notification Hubs provide a highly scalable, cross-platform push notification infrastructure that enables you to either broadcast push notifications to millions of users at once or tailor notifications to individual users.

Visual Studio Online

Host code, plan and track projects, and collaborate with team members to ship better software with Visual Studio Online. With Visual Studio Online, you get an end-to-end, cloud-based Application Lifecycle Management (ALM) solution that handles everything from hosted code repos and issue tracking to load testing and automated builds.

BizTalk Services

Windows Azure BizTalk Services is a simple, powerful, and extensible cloud-based integration service that provides Business-to-Business (B2B) and Enterprise Application Integration (EAI) capabilities for delivering cloud and hybrid integration solutions. The service runs in a secure, dedicated, per-tenant environment that you can provision on demand.

The next two sections discuss how to extend the on-premises datacenter fabric to include Windows Azure infrastructure services as well as extending fabric management to include managing Windows Azure.

Extending the datacenter fabric to Windows Azure

A key attribute of the Cloud OS strategy is delivering a hybrid infrastructure spanning private cloud, Windows Azure, and service provider clouds. This section covers extending the private cloud fabric (compute, storage, and network) to Windows Azure. Subsequently, we'll cover extending fabric management to resources hosted in Windows Azure.

Extending the datacenter network to Windows Azure

The first step in extending the fabric to Windows Azure is establishing secure network connectivity between the private cloud datacenter and Windows Azure. Windows Azure provides several methods for establishing secure VPN connectivity between a private cloud datacenter and Windows Azure.

Windows Azure Virtual Network

The “Windows Azure overview” section provided a brief overview of Windows Azure Virtual Network. With Virtual Network, you can create private networks in Windows Azure and specify your own private IP address ranges to be used in your virtual network. Resources placed in a virtual network, such as virtual machines, can only be accessed from other

resources within the virtual network or over administrator specified publically accessible endpoints which can be configured with access control lists (ACLs).

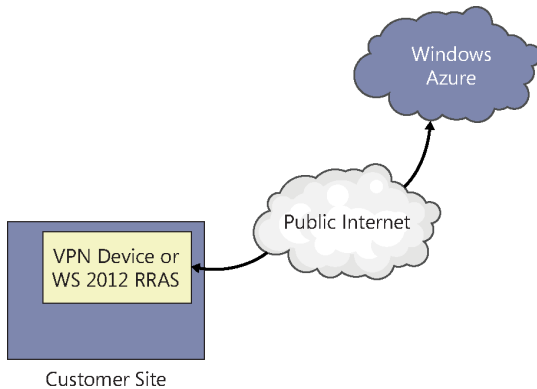
A second key feature of Virtual Network is that it can be used to create a secure, cross-premises VPN connection between Windows Azure and your datacenter. This is what we refer to as extending your datacenter fabric to Windows Azure as you can extend your network to include Virtual Networks you've established in Windows Azure, use a common IP addressing scheme (e.g. 10.x.x.x or 192.x.x.x) across the private and public cloud resources and even set up your own DNS servers either in your Virtual Network or on-premises.

Within a Windows Azure Virtual Network, you can establish multiple virtual machines and cloud services which can all communicate using that network. Again, you can determine if you want to allow any external connectivity from outside the virtual network. The Windows Azure Network Security (<http://go.microsoft.com/fwlink/p/?linkid=389558&clcid=0x409>) whitepaper provides depth on some of the topology options and security considerations.

When establishing VPN connectivity to Windows Azure, there are two primary options. The first is targeted toward individual users such as developers who may need to connect to your Windows Azure Virtual Networks from the Internet or arbitrary networks (such as a developer who moves between locations). In this case, the Windows Azure "point-to-site" capability can be utilized which consists of downloading a VPN connection profile from your Windows Azure Virtual Network that the developer installs on their workstation enabling them to use the VPN client built into Windows to connect to the Windows Azure Virtual Network. The point-to-site capability does not require a VPN device or special hardware. Point-to-site connectivity utilizes Secure Sockets Tunneling Protocol (SSTP).

The second method of establishing connectivity is the site-to-site VPN capability of Windows Azure Virtual Network. The site-to-site capability requires the installation and configuration of a VPN device (or Windows Server 2012 R2 Routing and Remote Access Server) in your datacenter to connect to a Windows Azure Virtual Network Gateway you configure on your Virtual Network. Microsoft provides a list of VPN devices that have been tested for compatibility with the site-to-site VPN capability described at <http://msdn.microsoft.com/en-us/library/windowsazure/jj156075.aspx>. Figure 3-1 illustrates the site-to-site VPN connectivity for Windows Azure Virtual Network adapted from this diagram on Microsoft TechNet (<http://www.windowsazure.com/en-us/services/expressroute/>).

Currently, Windows Azure Virtual Network only supports connecting your virtual network to one on-premises site or VPN device. You can have multiple Virtual Networks and connect them back to one or more sites, but there currently can only be a one-to-one relationship between a given Virtual Network and on-premises site. Site-to-site connectivity uses IPSec and IKEv2.



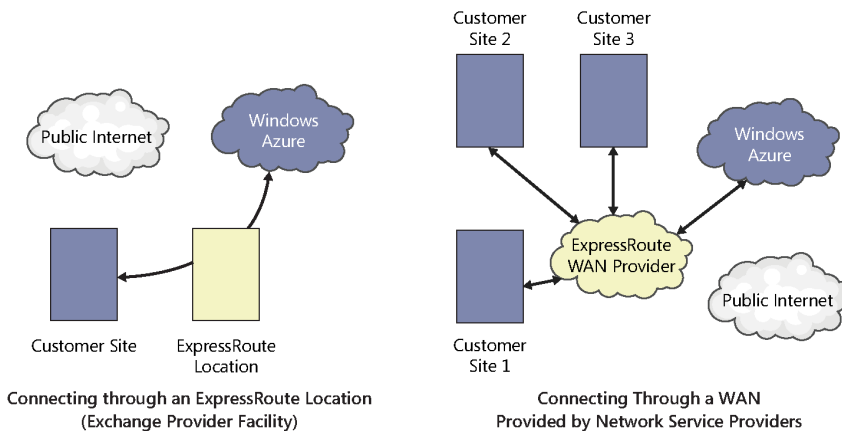
Connecting through Windows Azure Virtual Network Site-to-Site VPN

FIGURE 3-1 A connection through Windows Azure Virtual Network Site-to-Site VPN.

Windows Azure ExpressRoute

Windows Azure ExpressRoute enables you to create private connections between Azure datacenters and infrastructure that's on your premises or in a colocation environment. ExpressRoute connections do not go over the public Internet, and offer more reliability, faster speeds, lower latencies, and higher security than typical connections over the Internet. In some cases, using ExpressRoute connections to transfer data between on-premises and Windows Azure can also yield significant cost benefits.

With ExpressRoute, you can establish connections to Windows Azure at an ExpressRoute location (Exchange Provider facility) or directly connect to Windows Azure from your existing WAN network (such as a MPLS VPN) provided by a network service provider. The diagram in Figure 3-2 illustrates the two options adapted from this diagram on Microsoft TechNet (<http://www.windowsazure.com/en-us/services/expressroute/>).



Connecting through an ExpressRoute Location (Exchange Provider Facility)

Connecting Through a WAN Provided by Network Service Providers

FIGURE 3-2 An example of Windows Azure ExpressRoute connectivity options.

Note that in both cases, traffic between the customer site and Windows Azure does not traverse the public Internet. Unlike Windows Azure Virtual Network, which can be configured by an on-premises administrator, ExpressRoute requires collaboration with a service provider. At the time of publication these included AT&T, Equinix, and Level(3). More information on ExpressRoute can be found here: <http://www.windowsazure.com/en-us/services/expressroute/>.

These two options enable the ability to extend your private cloud datacenter network to Windows Azure. This opens a variety of scenarios such as extending your Active Directory into Windows Azure or using your on-premises System Center infrastructure to manage and monitor your resources in Windows Azure.

Extending datacenter storage to Windows Azure

There are several approaches to extending your private cloud storage infrastructure to Windows Azure for effectively unlimited storage capacity. As discussed in the “Windows Azure overview” section, Windows Azure provides highly available storage through three foundational storage constructs: blobs, tables, and queues. A wide variety of storage solutions can be built using these constructs such as applications exposing blob storage as shares, drives, or other common storage access scenarios. Solutions exist from both an IaaS and a PaaS perspective.

StorSimple

Cloud-integrated storage from Microsoft StorSimple provides primary storage, backup, archive, and disaster recovery, combined with Windows Azure. As discussed briefly in Chapter 2, StorSimple couples an on-premises storage appliance with Windows Azure blob storage. The on-premises appliance can provide two tiers of storage: hard disks (HDD) and solid-state disks (SSD). Windows Azure storage is a logical third tier of storage. Policies configured by the administrator determine when and what type of data is kept on SSD, which is put on HDD, and which is moved to Windows Azure. This is commonly referred to as *storage tiering*, where the most frequently accessed or important data is kept on the highest speed (but typically more costly) storage while less frequently accessed or important data is moved to less expensive storage such as Windows Azure. StorSimple also enables interesting backup and disaster recovery scenarios because StorSimple devices in different datacenters can be used to access snapshots and data in Windows Azure for rapid recovery.

Figure 3-3 illustrates a multiple-tier storage infrastructure for extending the datacenter storage fabric to Windows Azure that uses all of the elements described in this book. The highest performance tier is the Windows Scale-out File Server cluster infrastructure on-premises using SAS JBOD with SSD/HDD. While this infrastructure itself can be configured with multiple tiers, for simplicity it is illustrated as a single tier. This tier is ideal for virtual machine storage, high IO databases, etc. The Microsoft StorSimple appliance provides the point of access for the next three tiers. Tiers two and three represent the HDD and SSD tiers local to the StorSimple appliance also on-premises. These intermediate tiers are optimal for

application and file data. The final tier is also accessible through the StorSimple appliance, however, the data is stored in Windows Azure. All of the StorSimple tiers, including the Windows Azure tier, are presented on-premises as iSCSI targets meaning nearly any storage client can access them because iSCSI is widely supported in all operating systems.

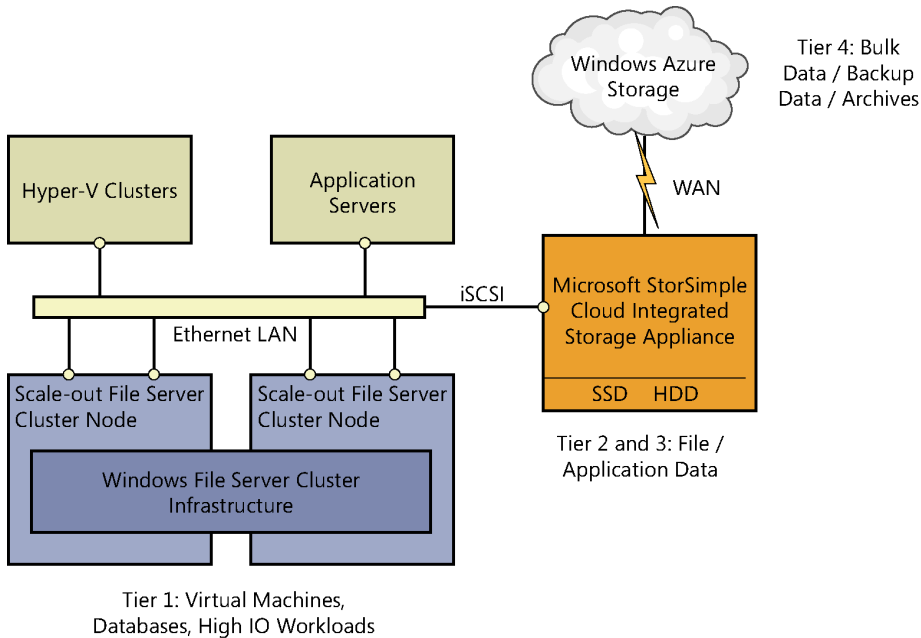


FIGURE 3-3 An example of Microsoft storage fabric spanning on-premises and Windows Azure.

PaaS storage

In addition to the IaaS approaches to utilizing Windows Azure storage, there are a variety of mechanisms from a PaaS perspective to store data in Windows Azure. Windows Azure storage can be accessed via the Windows Azure APIs from any application with connectivity to the Internet. A wide range of third-party applications and solutions can also be utilized to access Windows Azure storage. While the PaaS scenarios are beyond the scope of this book, it is important to realize the flexibility that the PaaS methods provide to applications and developers.

Extending datacenter compute to Windows Azure

Extending the datacenter compute fabric to Windows Azure entails using services such as virtual machines and HDInsight to augment your compute capacity with the effectively unlimited capacity of Windows Azure.

Windows Azure Virtual Machines

Windows Azure Virtual Machines, both IaaS and PaaS, enable you to deploy your workloads in Azure, burst to Windows Azure for extra capacity, or use Windows Azure as a backup or disaster recovery capability. Virtual Machines provide a wide range of scenarios for augmenting your on-premises compute capacity and over time Windows Azure will likely become the primary option for many, if not all, workload deployments.

HDInsight

A second scenario for extending a compute fabric to Windows Azure is in the area of big data, analysis, and high performance computing. Most organizations can benefit from the advances in big data, business intelligence, and related capabilities but purchasing, implementing, and managing large scale data solutions on-premises is cost prohibitive for many organizations. Windows Azure enables on-demand solutions for these topics through HDInsight which provides Apache Hadoop capability. In addition, Windows Azure provides the capability of Windows Server 2012 R2 High Performance Computing (HPC) clusters to be built in Windows Azure. The key feature is that very large clusters can be created by any organization and they only incur cost during the time they are utilized, then they can be easily de-commissioned.

The ability to extend the datacenter network, storage, and compute fabric to Windows Azure affords any organization access to world class datacenters and associated cloud services. Utilization-based pricing provides an easy on-ramp to capabilities that many organizations would otherwise be unable to utilize. The next section discusses how to enable a seamless management capability across the private cloud and Windows Azure.

Extending datacenter fabric management to Windows Azure

Once the network, storage, and compute fabric has been extended to Windows Azure, the next step is extending the fabric management capability of Microsoft System Center to encompass all of the resources hosted in Windows Azure. In addition, Microsoft has introduced new cloud-based management services that are hosted in Windows Azure called Windows Intune and System Center Advisor, which are management systems that are operated by Microsoft but can manage customer devices and infrastructure.

Self-Service

Microsoft provides two solutions for IaaS Self-Services. The first is System Center 2012 R2 App Controller. The second is the combination of the Windows Azure management portal and the Windows Azure Pack.

System Center 2012 R2 App Controller

System Center 2012 R2 App Controller provides a single self-service experience to configure, deploy, and manage virtual machines and services. App Controller enables a single self-service portal to span VMM-based private clouds, Windows Azure, and Microsoft service-provider partner clouds. App Controller provides an example of the design goal of the Cloud OS providing capabilities which span the three clouds.

Windows Azure Pack

The Windows Azure Pack integrates with System Center and Windows Server to help provide a self-service portal for managing services such as websites, virtual machines, and service bus. Windows Azure Pack also provides a portal for administrators to manage resource clouds, scalable web hosting, and more. Windows Azure Pack effectively provides a copy of the Windows Azure management portal which can be run in a private cloud datacenter or a service provider datacenter. Unlike App Controller, which is a single portal able to connect to all three clouds, Windows Azure Pack provides the same user interface as Windows Azure but is a separate portal. In the Cloud OS, the same user interface is provided across all three clouds but through three distinct portals. Over time continued convergence and commonality between Windows Azure, private cloud, and service provider cloud is expected.

Updating and update management

For scenarios in which granular update management is not required (where the update policy in effect is to utilize the built-in policies, such as download and apply all updates), the standard Windows or Microsoft Update that is available over the Internet can be utilized by Windows Azure virtual machines. An example of where this might apply is development or test in the cloud, or other cases in which granular management via WSUS or Configuration Manager is not required.

Windows Server Update Services

Windows Server Update Services (WSUS) enables IT administrators to deploy the latest Microsoft product updates. By using WSUS, administrators can fully manage the distribution of updates that are released through Microsoft Update to computers in their network. The WSUS server provides the features that administrators need to manage and distribute updates through a management console. In addition, a WSUS server can be the update source for other WSUS servers within the organization. The WSUS server that acts as an update source is called an upstream server. In a WSUS implementation, at least one WSUS server in the network must connect to Microsoft Update to get available update information. WSUS can be deployed either on-premises or in Windows Azure. Keep in mind that network traffic egressing Windows Azure incurs cost while traffic into Windows Azure does not.

System Center 2012 R2 Configuration Manager

System Center 2012 R2 Configuration Manager supports two scenarios for managing hybrid environments. The two supported scenarios are described in <http://support.microsoft.com/kb/2889321>. The first is using an on-premises deployment of Configuration Manager to manage virtual machines both locally and in Windows Azure over the site-to-site VPN capability described in previous sections. The second is a single-server, primary site deployment of Configuration Manager in a Windows Azure virtual machine to manage the other Windows Azure-hosted virtual machines.

Another use case for System Center 2012 R2 Configuration Manager and Windows Azure is cloud-based distribution points. You can use a cloud service in Windows Azure to host a distribution point. When you use a cloud-based distribution, you configure client settings to enable users and devices to access the content, and you specify a primary site to manage the transfer of content to the distribution point. Additionally, you specify thresholds for the amount of content that you want to store on the distribution point and the amount of content that you want to allow clients to transfer from the distribution point. Based on these thresholds, Configuration Manager can raise alerts that warn you when the combined amount of content that you have stored on the distribution point is near the specified storage amount or when transfers of data by clients are close to the thresholds that you defined.

Monitoring and alerting

In hybrid cloud architectures, there is a choice in terms of where management servers and services are hosted (on-premises or in the cloud) and whether the same management infrastructure is utilized for both on-premises and cloud-hosted resources.

System Center 2012 R2 Operations Manager

As with Configuration Manager, System Center 2012 R2 Operations Manager can be deployed in several scenarios for monitoring both on-premises and Windows Azure resources. The Operations Manager infrastructure can be deployed on-premises and is used to monitor both on-premises servers and virtual machines, as well as the Windows Azure IaaS-hosted resources, such as virtual machines and storage, by extending the datacenter network to Windows Azure using site-to-site VPN. Alternatively, a dedicated deployment of Operations Manager can be deployed in Windows Azure for managing and monitoring the Windows Azure-hosted resources.

When utilizing Operations Manager to manage Windows Azure-based resources, there are two levels of management capability. The first is using the Windows Azure management APIs in conjunction with the Windows Azure Fabric Management Pack for Operations Manager.

The Management Pack for Windows Azure Fabric enables you to monitor the availability and performance of Windows Azure fabric resources that are running on Windows Azure. The management pack runs on a specified proxy agent and then uses various Windows Azure APIs to discover and collect instrumentation information remotely about a specified Windows Azure resource, such as a cloud service, storage, or virtual machine. The Management Pack for Windows Azure Fabric offers the following functionality:

- Discovers Windows Azure Cloud Services.
- Provides status of each role instance.
- Collects and monitors performance information per role instance.
- Collects and monitors Windows events per role instance.
- Collects and monitors the .NET Framework trace messages from each role instance.
- Grooms performance, event, and the .NET Framework trace data from Windows Azure Storage.
- Changes the number of role instances.
- Discovers Windows Azure Virtual Machines.
- Provides status of each role instance of the Virtual Machines.
- Discovers Windows Azure Storage.
- Monitors availability and size of each Storage and optionally alerts.
- Discovers relationships between discovered Windows Azure resources, to see which other resources a particular Windows Azure resource uses. This information is then displayed in a topology dashboard.
- Monitors management and cloud service certificates and alerts if the certificates are about to expire.
- Includes a new Distributed Application template that lets you create distributed applications that span Windows Azure as well as on-premises resources, for hybrid monitoring scenarios.
- Includes a set of dashboards for the hybrid monitoring scenarios.

This first level of management capability provided by Operations Manager and the Windows Azure Fabric Management Pack does not require the deployment of agents or code into the virtual machines and thus can function against any supported Windows Azure resources.

The second level of management capability using Operations Manager entails deploying operations manager agents into the Windows Azure-hosted virtual machines in the same fashion as on-premises hosted resources. This requires implementation of the site-to-site VPN capability.

The combination of the two approaches enables seamless management and monitoring of both private cloud and Windows Azure resources using Operations Manager.

System Center Global Service Monitor

System Center Global Service Monitor is a cloud service that provides a simplified way to monitor the availability of external web-based applications from multiple locations around the world. More importantly, Global Service Monitor monitors applications from the perspective of the customers who use them. Because Global Service Monitor monitors from locations that are correlated to customer geographies, application owners can gain insight into customer experiences in addition to the separate problems that relate to external factors—such as Internet or network problems—from application or service problems.

Global Service Monitor integrates with the Operations Manager console, so that you can monitor external and internal-facing web applications in the same place that you monitor other applications. Using Global Service Monitor, the Operations Manager console integration lets you monitor web applications from both internal and external locations. In Global Service Monitor, you can use your management group and obtain access to agents in the cloud that are provided by Microsoft. This lets you monitor web applications from 15 locations and then report to your management group. You can also use your own agents as watcher nodes to monitor internal locations and applications.

Windows Azure Diagnostics

Primarily utilized in PaaS scenarios, Windows Azure Diagnostics (<http://msdn.microsoft.com/en-us/library/gg433048.aspx>) enables you to collect diagnostic data from an application that is running in Windows Azure. You can use diagnostic data for debugging and troubleshooting, measuring performance, monitoring resource usage, traffic analysis and capacity planning, and auditing. After the diagnostic data has been collected, it can be transferred to a Windows Azure storage account for persistence.

Orchestration and automation

Microsoft provides two solutions for orchestration and automation. The first is Windows Azure PowerShell and the second System Center 2012 R2 Orchestrator.

Windows Azure PowerShell

Windows Azure PowerShell is a powerful automation capability that you can use to control and automate the deployment and management of your workloads in Windows Azure. Windows Azure PowerShell can be used for provisioning virtual machines, setting up virtual networks and cross-premises networks, and managing cloud services in Windows Azure. Virtually all Windows Azure services can be managed using Windows Azure PowerShell.

System Center 2012 R2 Orchestrator

Using System Center Orchestrator 2012 R2, you can automate and orchestrate a wide range of activities. These activities can include direct Windows Azure management tasks, such as working with storage or virtual machines, but can also include scenarios such as orchestrating activities within virtual machines and services that are deployed in Windows Azure.

Orchestration in a hybrid cloud environment requires careful planning: there is a wide range of requirements and options. In hybrid cloud architectures, there is choice in terms of where Orchestrator management servers and services are hosted (on-premises or in the cloud) and whether the same management infrastructure is utilized for both on-premises and cloud-hosted resources.

Microsoft provides integration packs for each of the System Center components. This enables Orchestrator runbooks to automate a wide range of management tasks across physical, virtual, and application resources. The Integration Pack for Windows Azure is an add-on for Orchestrator in System Center 2012 R2 that enables you to automate Windows Azure operations that relate to certificates, deployments, cloud services, storage, and virtual machines by using the Windows Azure Service Management REST API.

Backup and disaster recovery

Microsoft provides two solutions for backup and disaster recovery. The first is Windows Azure Backup and the second is Hyper-V Recovery Manager.

Windows Azure Backup

Windows Azure Backup is a new feature in Windows Azure that seamlessly enables off-site file and folder backups from the on-premises Windows Server, Windows Server Essentials, or System Center Data Protection Manager to Windows Azure.

Using incremental backups, only changes to files are transferred to the cloud. This helps ensure efficient use of storage, reduced bandwidth consumption, and point-in-time recovery of multiple versions of the data. Configurable data-retention policies, data compression, and data-transfer throttling also offer you added flexibility and help boost efficiency. Backups are stored in Windows Azure and are “off-site,” reducing the need to secure and protect on-site backup media.

The backup data is encrypted prior to being stored in Windows Azure. The customer is responsible for managing encryption keys and backup of those keys. Customer data is never decrypted in Windows Azure; for restores, the data is decrypted on the on-premises client side by the customer.

Hyper-V Recovery Manager

Windows Azure Hyper-V Recovery Manager can help you protect important services by coordinating the replication and recovery of System Center 2012 private clouds at a secondary location.

System Center 2012 R2 VMM private clouds can be protected through automation of the replication of the virtual machines that compose them at a secondary location. The ongoing asynchronous replication of each virtual machine is provided by Windows Server 2012 R2 Hyper-V Replica and is monitored and coordinated by Hyper-V Recovery Manager.

The service helps automate the orderly recovery in the event of a site outage at the primary datacenter. Virtual machines can be brought up in an orchestrated fashion to help restore service quickly. This process can also be used for testing recovery or transferring services temporarily. Windows Azure Hyper-V Recovery Manager provides the following functionality:

- Windows Azure–based portal and service that orchestrates DR operations:
 - Across two Virtual Machine Manager–managed data centers or private clouds
 - For Hyper-V virtual machines that are running on Windows Server 2012 and above
- Leverages Hyper-V Replica technology for replication
- Provides single-click at scale configuration of settings across sites
- Provides “Recovery Plan” feature to enable grouping, prioritizing, and sequencing of disaster recovery operations across a large number of virtual machines
- Leverages Windows Azure Portal to provide multisite DR operations from anywhere

Using System Center 2012 R2, a single solution for fabric management can be utilized both for the private cloud and Windows Azure hosted resources. In later chapters, we’ll see that this also can be extended to service provider clouds for a single management solution spanning the three clouds in the Cloud OS.

Service provider cloud

The third cloud type in the Cloud OS vision is the service provider cloud. These are clouds hosted by a provider other than Microsoft or the end customer. As discussed previously, service provider clouds are ideal for use cases that either Windows Azure can't support or where Windows Azure isn't available. Service provider clouds are also ideal for use cases that are too expensive for hosting on-premises or where on-premises staff is not trained or capable of managing the particular solution (consider a hosted enterprise resource planning instance such as SAP).

Cloud OS Network

The Cloud OS vision depends on a robust service provider ecosystem running the Microsoft platform to enable the common virtualization, identity, data, management, and development capabilities as the private cloud and Windows Azure. To achieve this goal of a robust ecosystem, Microsoft established the Cloud OS Network.

The Cloud OS Network is a worldwide consortium of cloud service providers who have embraced the Cloud OS vision. These organizations offer solutions based on the Microsoft Cloud Platform designed to meet customer business needs. Members of this network combine Microsoft technology with their hosting and geographic expertise to provide flexibility and choice for hybrid infrastructure solutions. A current list of the Cloud OS Network partners can be found at: <http://www.microsoft.com/en-us/server-cloud/cloud-os-network.aspx#fbid=tKWR1hKoghK>.

Members of the Cloud OS Network have built hosting solutions using the same products and architectures as described in Chapter 2, "Private cloud." The same Infrastructure as a Service (IaaS) Product Line Architecture referenced in those chapters and used in private cloud deployments is utilized in the service provider deployments by the Cloud OS Network partners. This provides not just commonality of product utilization, but also commonality of architectures and solutions.

Many of the Cloud OS Network partners are utilizing the Microsoft software-defined storage, network, and compute architectures described in this book because many of those capabilities were specifically engineered to support service provider needs.

By choosing Cloud OS Network partners, you can extend your datacenter fabric and fabric management to service providers in addition to Windows Azure and achieve the Cloud OS vision of a unified virtualization, identity, data, management, and development platform across the entire hybrid infrastructure.

Extending the datacenter fabric to a service provider

Many of the same concepts from extending the datacenter fabric to Windows Azure also apply to extending the fabric to service providers. In some cases, there are different features or capabilities utilized and those are what we will cover in this chapter.

Extending the datacenter network to service providers

Similar to extending the datacenter network to Windows Azure, extending to service providers also entails using VPN capability, however, in this case the combination of VPN and Hyper-V Network Virtualization can be utilized. Cloud OS Network partners that enable Hyper-V Network Virtualization capability do so through a combination of Hyper-V, Virtual Machine Manager (VMM), and Windows Azure Pack capabilities. The end result is that you can extend your datacenter network and bring your own IP address spaces to the service provider datacenter. As described in the section titled “Software-defined networking” in Chapter 2, Hyper-V Network Virtualization allows the service provider to run multiple isolated tenant networks side by side allowing each tenant to bring their own IP address ranges.

If the service provider utilizes Service Provider Foundation and Windows Azure Pack (each described later in this chapter) they can expose a similar network configuration and VPN connectivity self-service interface as Windows Azure, again providing a common experience across the three cloud types even though the underlying implementation is different.

Establishing network connectivity to the service provider cloud opens a variety of application and management scenarios. Similar to the Windows Azure scenarios, establishing network connectivity enables you to either deploy Microsoft System Center 2012 for fabric management in the private cloud datacenter and managing both private cloud and service provider hosted resources or alternatively, placing the System Center implementation at the service provider and managing on-premises resources from that implementation. When network connectivity is established between clouds using VPN and network virtualization, nearly all System Center management scenarios become possible—the most important being the ability to deploy agents into the running virtual machines for management and monitoring.

Extending datacenter storage to service providers

The combination of VPN and network connectivity to service providers enable several storage scenarios, such as hosting file servers and storage at the provider, replicating storage between the private cloud and the service provider using technologies such as Distributed File System (DFS) in Windows.

Similarly to Windows Azure, service providers are also able to leverage economies of scale and offer raw capacity such as storage at prices that may be less than what enterprises are able to achieve on-premises. Additionally, since most service providers are not creating global scale services like Windows Azure where mass standardization is required, they often have the flexibility to fill niches and needs that may not be able to be profitably delivered by Windows Azure. Examples might be hosting particular types of regulated data, or hosting data in specific regions where Windows Azure does not maintain a datacenter. For these and many other possible reasons, extending storage to the service provider cloud is a valuable option.

Many of the Windows features and capabilities, such as DFS, DFS-R, and File Services, are well known as they have been part of Windows Server for many years and therefore they will not be covered in detail here. We simply want to emphasize that many of the architectures you might have used between your on-premises datacenters can also be used when a service provider running the Cloud OS and allowing network connectivity is selected.

Extending datacenter compute to service providers

Extending the datacenter compute infrastructure to service providers is the same as extending to Windows Azure. Cloud OS Network partners enable Hyper-V based hosting of virtual machines. Cloud OS Network partners also have the option of using Windows Azure Pack as the self-service user interface to their hosted solutions and providing the same user interface being utilized by Windows Azure and their customer's private cloud infrastructures.

An additional capability service providers can provide, which is potentially highly valuable to customers and not currently provided by Windows Azure, is being a replication target for Hyper-V Replica. Windows Server 2012 R2 enhances Hyper-V Replica to support three replicas of a virtual machine: the primary or source virtual machine, a secondary replica, and a tertiary replica. For source virtual machines in a private cloud, the secondary replica could be hosted either on-premises or at a service provider as can the tertiary replica. This enables two interesting scenarios. The first is a customer maintaining both the primary replica in the private cloud and the secondary in a second datacenter in their private cloud with the tertiary being hosted at a service provider. The second scenario is where the customer maintains the primary virtual machine with both the secondary and tertiary replicas being hosted by a service provider. This can relieve an organization from the expense of maintaining a backup or disaster recovery datacenter while opening up opportunities for service providers to bring significant value to their customers.

The third cloud type in the Cloud OS, the service provider cloud, enables a number of scenarios for extending the hybrid cloud fabric to the Microsoft Cloud OS Network partners. With the fabric extended to service providers, fabric management must also be extended to encompass the provider hosted resources.

Extending datacenter fabric management to a service provider

The final step in building the Cloud OS hybrid infrastructure is extending fabric management to encompass the resources hosted at Cloud OS Network service providers. This section will be brief because nearly all of the approaches that were utilized for managing Windows Azure hosted resources also apply to managing service provider hosted resources. Two enabling technologies, Service Provider Foundation (SPF) and Windows Azure Pack (WAP) provide API and UI commonality respectively between private cloud, Windows Azure, and service provider cloud.

Service Provider Foundation

Service Provider Foundation is provided with System Center 2012 R2 Orchestrator. Service Provider Foundation exposes an extensible OData web service that interacts with System Center 2012 R2 VMM. This enables service providers and hosters to design and implement multi-tenant self-service portals that integrate with the IaaS capabilities available in a Windows Server 2012 R2 and System Center 2012 R2 cloud environment.

In many cases, service providers want to enable a robust self-service capability for their customers, typically in the form of a self-service portal providing secure access for the customer to provision and manage resources in the service providers shared infrastructure. Creating such a portal requires a robust set of web services and APIs for the portal to use and interact with the infrastructure. For the Microsoft platform, this function is provided by SPF.

Figure 4-1, adapted from Microsoft TechNet diagram (<http://technet.microsoft.com/en-us/library/jj642897.aspx>), illustrates the high-level architecture enabled by SPF.

The tenant represents a service provider's customer, and the tenant has assets on the service provider's infrastructure. Each tenant has their own administrators, applications, scripts, and other tools. The service provider could be an enterprise IT organization providing services to business units or it could be a commercial service provider or hoster.

The service provider provides tenants an environment, which can include virtual machines or other resources. The service provider in this case is assumed to have an existing self-service portal, which all tenants can use (later we'll discuss the Windows Azure Pack portal provided by Microsoft). On the back end, the service provider has a set of resources (compute, storage, network), which is called the fabric. The service provider allocates those resources into discrete groups according to the service provider's needs in terms of performance, isolation, etc. Each of these groups is known as a stamp. The service provider assigns the tenant's access

to stamps in whatever manner is appropriate. The tenant's resources may be provisioned across several stamps, according to the service provider's policies and business model. SPF makes it possible for the service provider to present an aggregated view to the tenant of all their resources regardless of which stamp they are hosted on. SPF also enables a set of application programming interfaces (APIs) tenants can utilize to manage their resources.

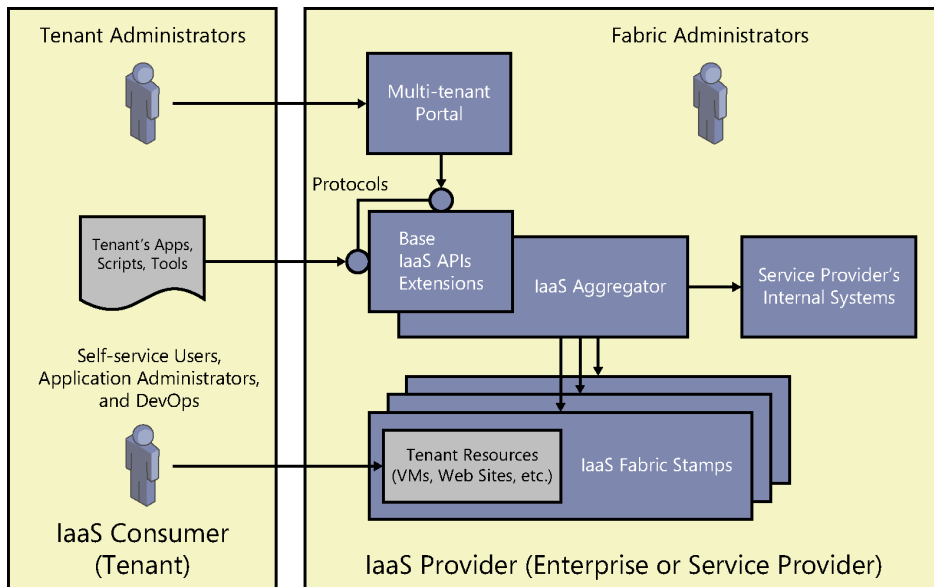


FIGURE 4-1 An example of the Service Provider Foundation architecture.

A stamp in Service Provider Foundation is a logical scale unit of compute, storage, and network designed for scalability that provides pre-determined amount of capacity. An example of a stamp is the single rack architecture described in previous chapters where a balanced mix of compute, storage, and network capacity is designed to support a specified number of virtual machines. As tenant demand increases, the service provider deploys additional stamps to meet demand. As described previously, these stamps can be deployed from bare-metal by VMM.

Figure 4-2, also adapted from Microsoft TechNet (<http://technet.microsoft.com/en-us/library/jj642897.aspx>) provides a view of how both the service provider's custom portal and System Center 2012 R2 App Controller can serve as the front end to the service provider's hosted IaaS stamps.

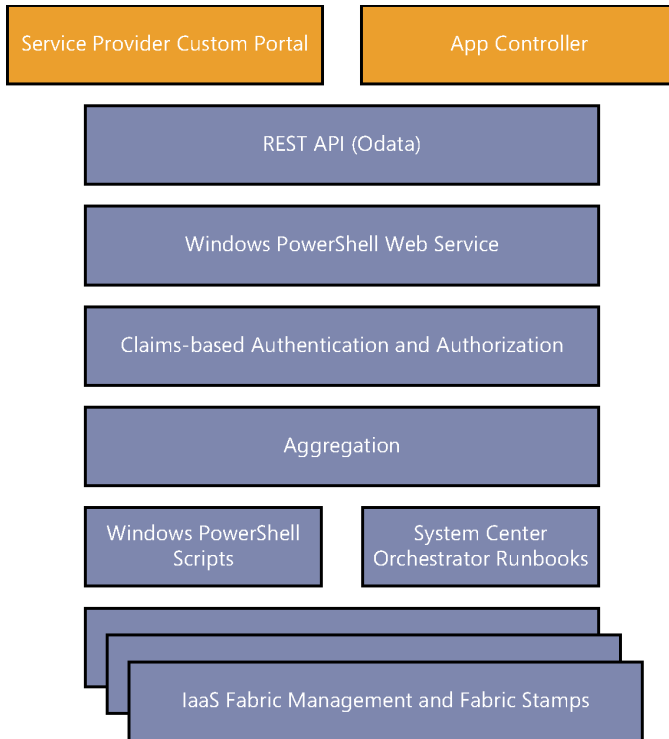


FIGURE 4-2 An example of the Service Provider Foundation integration architecture.

Service Provider Foundation uses a SQL Server database to aggregate the tenant resources, which are managed with Windows PowerShell scripts and Orchestrator runbooks. The service provider can distribute tenant resources among management stamps depending on their own policies while to the tenant their resources are aggregated and appear to be hosted on a single infrastructure.

Another key service provided by SPF is usage metering. SPF provides usage metering that enables service providers to:

- Obtain metrics for tenant usage consumption for virtual machines, CPU, memory, network, and disk.
- Determine capacity utilization.
- Bill tenants for their usage according to their plans.

Usage metering is a critical requirement for service providers as it enables them to monitor and track various metrics that they use to bill their customers for services provided.

The Usage Service captures the tenant-specific resource allocation and consumption information in a uniform manner across the hosted services (VMs, web sites, etc.). The Usage Service treats all services uniformly and collects information across these services and stores them for a limited period of time in a SQL Server database. The information is designed to be

used by billing or financial systems for chargeback and monetization of the provided services. The information captured consists of actions performed with billing impact on the self-service tenant portals or at the Service Management API layer, meaning regardless of whether a tenant performs an action through the portal, PowerShell, or API, the Usage Metering service will capture those actions.

The Usage Metering service does not provide a billing system but is designed to enable third-party billing systems by capturing the required data and making that data available via a REST API.

Windows Azure Pack

Windows Azure Pack (WAP) has been described previously as providing an Azure-consistent self-service user interface for private and service provider clouds. Windows Azure Pack is a collection of Windows Azure technologies, available to Microsoft customers at no additional cost, for installation into private cloud or service provider data centers. It runs on top of Windows Server 2012 R2 and System Center 2012 R2.

Windows Azure Pack includes the following capabilities as documented on Microsoft TechNet (<http://technet.microsoft.com/en-us/library/dn296435.aspx>):

- **Management portal for tenants** A customizable self-service portal for provisioning, monitoring, and managing services such as Web Site Clouds, Virtual Machine Clouds, and Service Bus Clouds.
- **Management portal for administrators** A portal for administrators to configure and manage resource clouds, user accounts, and tenant offers, quotas, and pricing.
- **Service management API** A REST API that helps enable a range of integration scenarios including custom portal and billing systems.
- **Web Site Clouds** A service that helps provide a high-density, scalable shared web hosting platform for ASP.NET, PHP, and Node.js web applications. The Web Site Clouds service includes a customizable web application gallery of open source web applications and integration with source control systems for custom-developed web sites and applications.
- **Virtual Machine Clouds** A service that provides IaaS capabilities for Windows and Linux virtual machines. The Virtual Machine Clouds service includes a VM template gallery, scaling options, and virtual networking capabilities.
- **Service Bus Clouds** A service that provides reliable messaging services between distributed applications. The Service Bus Clouds service includes queued and topic-based publish/subscribe capabilities.
- **SQL and MySQL** Services that provide database instances. These databases can be used in conjunction with the Web Sites service.
- **Automation** The capability to automate and integrate additional custom services into the services framework, including a runbook editor and execution environment.

Windows Azure Pack also provides APIs and builds on the SPF APIs. In addition to virtual machines, Windows Azure Pack also enables web sites, databases, and service bus services similar to Windows Azure but hosted in the private cloud or service provider cloud. Figure 4-3 illustrates the architecture that a service provider would deploy.

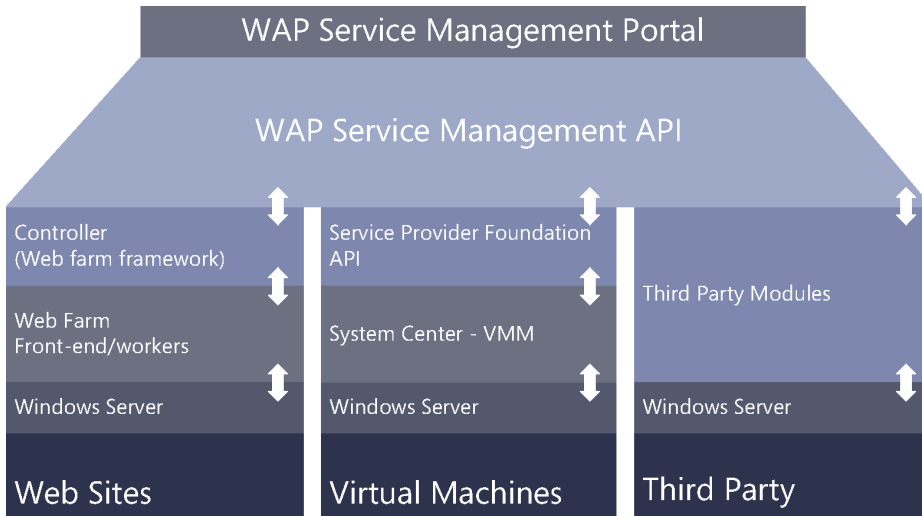


FIGURE 4-3 An example of the Windows Azure Pack and SPF architecture.

The references to third-party modules indicate the extensibility model enabled by both SPF and the Windows Azure Pack for the advantage of Microsoft's large partner ecosystem. Multiple partners have created extensions for connecting to billing systems, providing hosted services beyond those provided by WAP natively, and many other scenarios.

Windows Azure Pack enables significant new capabilities to Windows and System Center. Given its pedigree from Windows Azure and its target use cases with large enterprises and service providers, Windows Azure Pack is delivered through a highly available set of web services and capabilities requiring a relatively complex architecture.

Windows Azure Pack is comprised of several required and optional components. Each component can be deployed in physical or virtual machines and each can be deployed as scale-out, load-balanced tiers. This section provides the suggested machine topologies for these components.

There are four defined patterns for Windows Azure Pack deployment:

- Express Deployment (single server)
- Basic Distributed Deployment
- Minimal Distributed Deployment
- Scaled Distributed Deployment

The latter two deployment patterns are recommended for production environments. The Minimal Distributed Deployment, illustrated in Figure 4-4, is appropriate for enterprise or small service provider deployments.

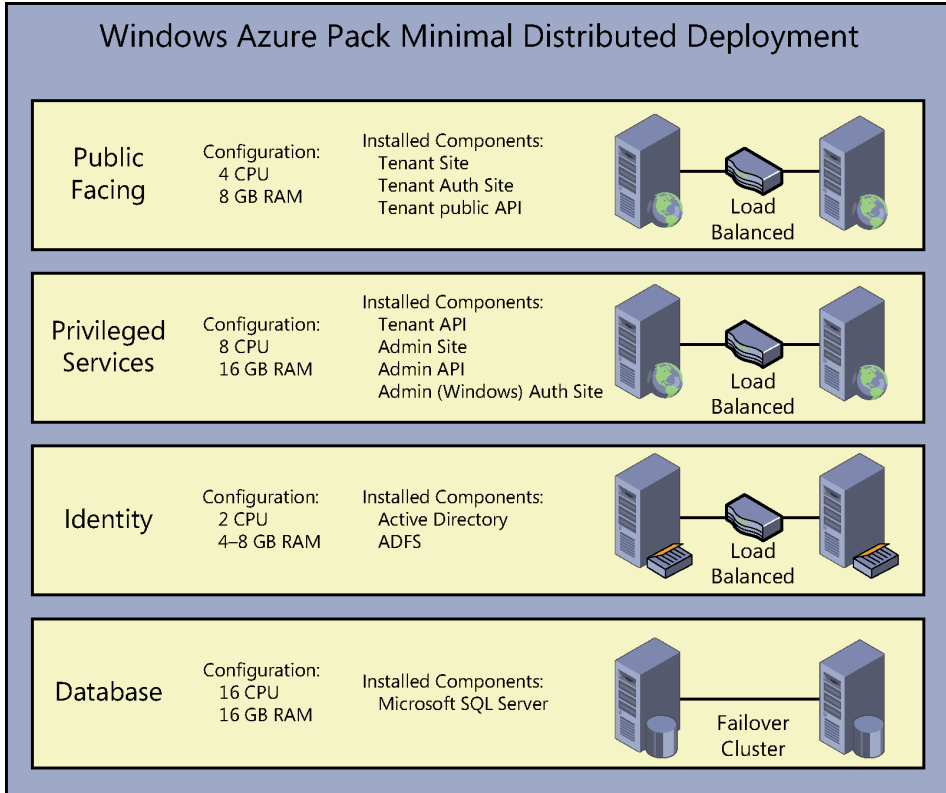


FIGURE 4-4 An example of the Windows Azure Pack Minimal Distributed Deployment.

For large enterprises or service providers requiring higher scale, the Windows Azure Pack Scaled Distributed Deployment pattern can be utilized. This pattern further separates the layers of the Windows Azure Pack architecture into their own sets of load-balanced servers (or virtual machines). Figure 4-5 illustrates this deployment pattern.

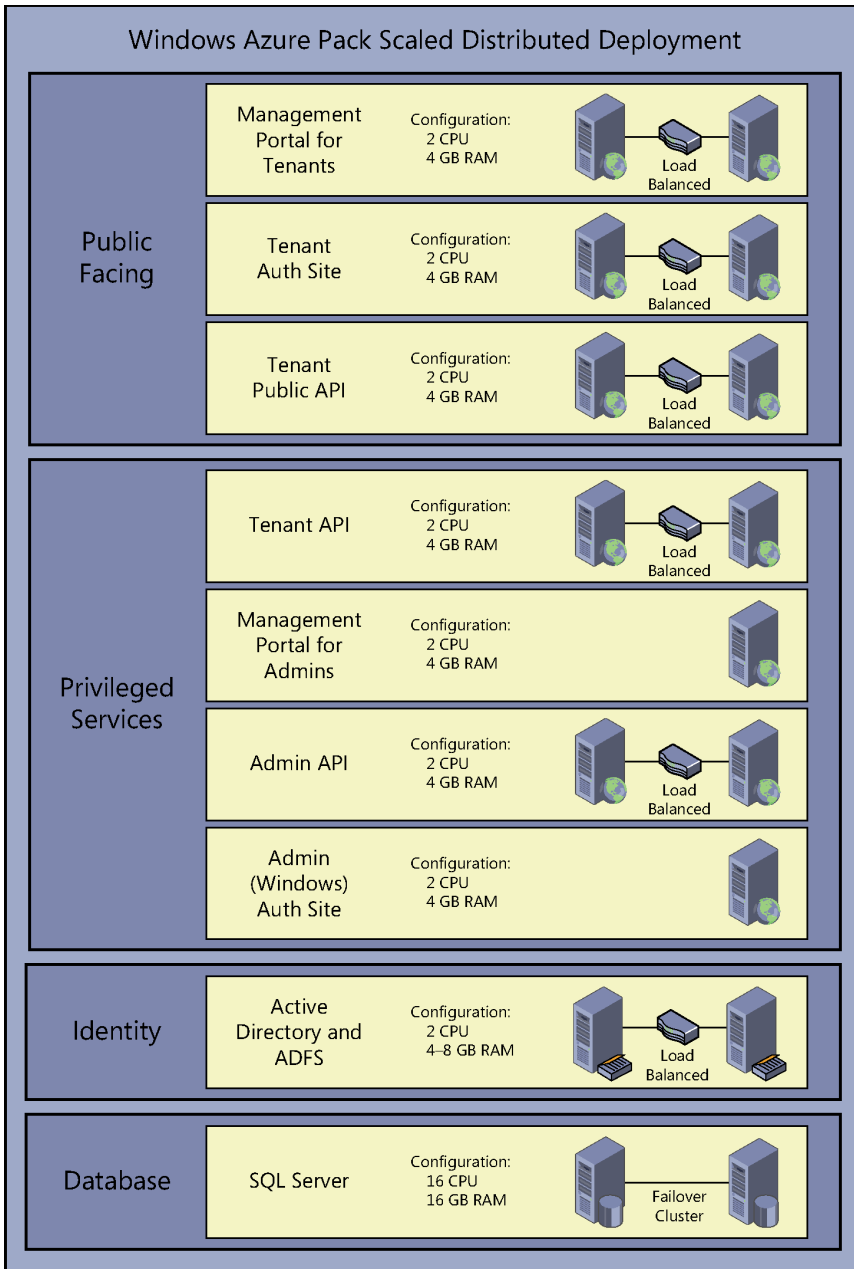


FIGURE 4-5 An example of the Windows Azure Pack Minimal Distributed Deployment.

For both enterprises and service providers utilizing Windows Azure Pack, our reference architecture for IaaS referred to in previous sections considers WAP to be a part of fabric management, meaning it is deployed (using either of the two patterns illustrated above) on

the fabric management cluster as virtual machines. In our reference architecture, WAP also utilizes the fabric management SQL Server guest cluster that the rest of the System Center components utilize and the WAP database requirements are illustrated alongside those of System Center.

System Center 2012 R2

As with extending fabric management to Windows Azure, System Center 2012 R2 can also be utilized to manage resources hosted at service providers. Many of the same deployment options apply such as using an on-premises deployment of System Center to manage service provider hosted resources over a VPN connection to the provider or deploying System Center itself at the service provider to manage all of your resources hosted there. Given the similarity in approach, those options won't be repeated here.

Utilizing System Center 2012 R2 – App Controller, and enterprise can connect App Controller to any service provider cloud that exposes SPF endpoints to them. What this means is that from within App Controller, it can be configured to provision virtual machines to connected service provider clouds in addition to VMM-based private clouds and Windows Azure.

Hyper-V Replica

As discussed previously, Hyper-V Replica provides asynchronous replication of Hyper-V virtual machines between two (or three) hosting servers. It is simple to configure and does not require either shared storage or any particular storage hardware. Replication works over any ordinary IP-based network, and the replicated data can be encrypted during transmission. Hyper-V Replica works with standalone servers, failover clusters, or a mixture of both. The servers can be physically colocated or widely separated geographically. The physical servers do not need to be in the same domain, or even joined to any domain at all.

When replication is enabled, changes in the primary virtual machines are transmitted over the network periodically to the Hyper-V Replica virtual machines. The exact frequency varies depending on how long a replication cycle takes to finish (depending in turn on the network throughput, among other things), but generally, replication data is sent to the Hyper-V Replica server every 5 minutes in Windows Server 2012. In Windows Server 2012 R2, you can configure the replication frequency, so that the changes are sent every 30 seconds, every 5 minutes, or every 15 minutes.

A primary scenario that service providers can enable (which is currently not supported in Windows Azure) is the Hyper-V Replica hosting possibility that was described earlier in this chapter in the section titled "Extending datacenter compute to service providers." This scenario is where the service provider serves as a replication target for the secondary or tertiary replicas of on-premises virtual machines to achieve a disaster recovery capability without the expense of a second or third datacenter being incurred by the customer.

The typical scenario for Hyper-V Replica is replicating virtual machines from your primary datacenter to a secondary datacenter. With Windows Server 2012 R2, the ability to replicate

to a third or tertiary datacenter was introduced. Figure 4-6 illustrates this capability from a private cloud perspective.

Hyper-V Replica Between Primary, Secondary, and Tertiary Datacenter





FIGURE 4-6 A Hyper-V Replica between three private cloud datacenters.

With the addition of a service provider cloud and assuming the service provider enables the capability, the service provider could host either the secondary or the secondary and tertiary replicas. Figure 4-7 illustrates both design options.

Hyper-V Replica Between Primary, Secondary, and Service Provider (Tertiary) Datacenter



Hyper-V Replica Between Primary, Service Provider Secondary, and Tertiary Datacenter




FIGURE 4-7 An example of Hyper-V Replica between three private cloud datacenters.

The above scenario can be highly valuable both for enterprises and service providers because maintaining secondary or disaster recovery datacenters is a significant cost for enterprises that can be reduced by leveraging a service provider. For service providers, this is a net new service they can offer to customers. This is another example of the new approaches to IT challenges enabled by the Cloud OS.

Conclusion

In this book we've described the vision of the Cloud OS and detailed the architectures and capabilities of the Windows Server 2012 R2 and System Center 2012 R2 private and service provider clouds as well as the Windows Azure public cloud. The combination of the three, as illustrated in Figure 4-8, comprises the Cloud OS hybrid infrastructure. System Center delivers the integrated cloud platform management suite required to utilize all three cloud types as a single platform.

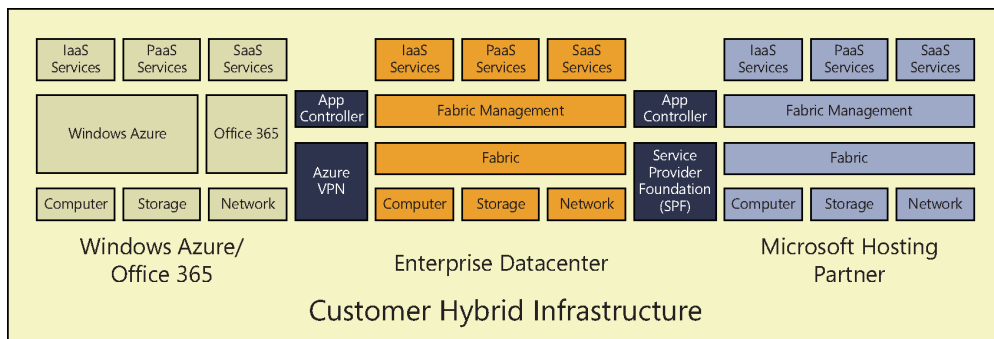


FIGURE 4-8 An example of the Cloud OS hybrid infrastructure.

The hybrid infrastructure enables an IT organization to provide a robust cloud service catalog consisting of infrastructure and platform capabilities such as virtual machines, web sites, and storage. The cloud service catalog might have multiple items of the same type (such as virtual machines), each with different cost, performance, SLA, and other characteristics depending on the cloud type they are hosted in. This enables IT consumers to choose the cloud service and cloud type or location which is most optimal for their use case. The Microsoft Cloud OS hybrid infrastructure provides a common virtualization, identity, data, management, and development platform across all three cloud types while delivering a common user interface and experience for both administrators and consumers.

Additional resources

The following resources, which are referenced in this book, can be found on Microsoft TechNet and in other Microsoft properties. They have been collected here for ease of reference.

- Infrastructure-as-a-Service Product Line Architecture-Deployment Guide:
<http://aka.ms/iaasdeployment>
- Infrastructure as a Service Product Line Architecture-Fabric Architecture Guide:
<http://aka.ms/iaasfabricarchitecture>
- Infrastructure as a Service Product Line Architecture - Fabric Management Guide:
<http://aka.ms/iaasfabricmanagement>
- Microsoft Cloud OS Vision:
<http://www.microsoft.com/en-us/server-cloud/cloud-os/>
- Microsoft Cloud OS Network:
<http://www.microsoft.com/en-us/server-cloud/cloud-os-network.aspx>
- Windows Azure:
<http://www.windowsazure.com>
- Building Clouds Blog:
<http://blogs.technet.com/b/privatecloud/>
- Server & Tools Blog:
<http://blogs.technet.com/b/serverandtools/>

About the author



DAVID ZIEMBICKI is a Senior Architect in Microsoft Services' Americas Office of the CTO. David's areas of expertise include private and hybrid cloud, virtualization, and datacenter automation. He has been a leading infrastructure architect across hundreds of strategic projects with public sector and Fortune 500 customers in multiple industries throughout his IT career. David is a lead architect for Microsoft's Datacenter Services Portfolio and the Microsoft Private Cloud Fast Track program. He is

a course instructor, published author, and regular speaker on Microsoft Cloud, Datacenter, and Infrastructure solutions.

David's blog can be found at <http://davidzi.com/blog> and he is on Twitter at <http://www.twitter.com/davidzi>.

About the series editor



MITCH TULLOCH is a well-known expert on Windows Server administration and virtualization. He has published hundreds of articles on a wide variety of technology sites and has written or contributed to over two dozen books, including *Windows 7 Resource Kit* (Microsoft Press, 2009), for which he was lead author; *Understanding Microsoft Virtualization Solutions: From the Desktop to the Datacenter* (Microsoft Press, 2010); and *Introducing Windows Server 2012* (Microsoft Press, 2012), a free e-book that has been downloaded almost three quarters of a million times.

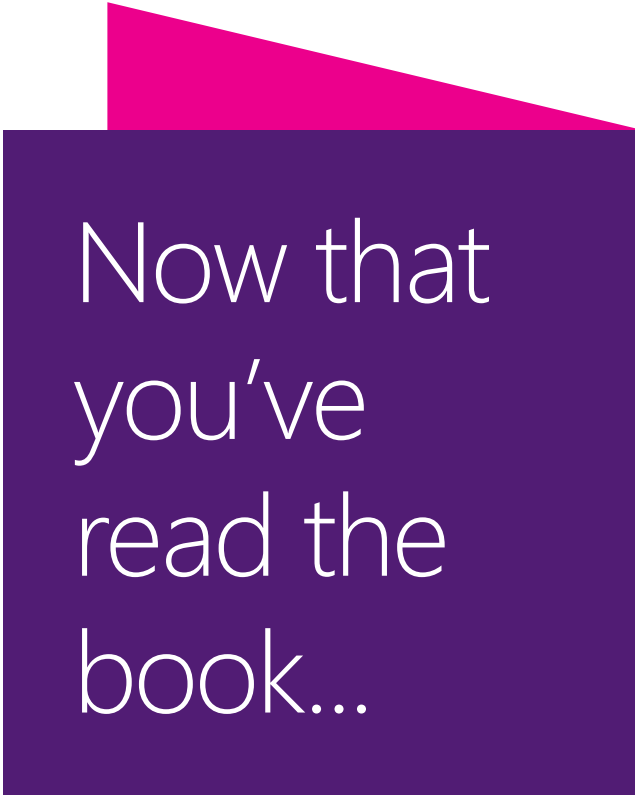
Mitch has been repeatedly awarded Most Valuable Professional (MVP) status by Microsoft for his outstanding contributions to supporting the global IT community. He is a nine-time MVP in the technology area of Windows Server Software Packaging, Deployment & Servicing. You can find his MVP Profile page at <http://mvp.microsoft.com/en-us/mvp/Mitch%20Tulloch-21182>.

Mitch is also Senior Editor of WServerNews (<http://www.wservernews.com>), a weekly newsletter focused on system administration and security issues for the Windows Server platform. With more than 100,000 IT pro subscribers worldwide, WServerNews is the largest Windows Server–focused newsletter in the world.

Mitch runs an IT content development business based in Winnipeg, Canada, that produces white papers and other collateral for the business decision maker (BDM) and technical decision maker (TDM) audiences. His published content ranges from white papers about Microsoft cloud technologies to reviews of third-party products designed for the Windows Server platform. Before starting his own business in 1998, Mitch worked as a Microsoft Certified Trainer (MCT) for Productivity Point.

For more information about Mitch, visit his website at <http://www.mtit.com>.

You can also follow Mitch on Twitter at <http://twitter.com/mitchtulloch> or like him on Facebook at <http://www.facebook.com/mitchtulloch>.



Now that
you've
read the
book...

Tell us what you think!

Was it useful?

Did it teach you what you wanted to learn?

Was there room for improvement?

Let us know at <http://aka.ms/tellpress>

Your feedback goes directly to the staff at Microsoft Press,
and we read every one of your responses. Thanks in advance!

