Dell Private Cloud Solution

Based on the Microsoft Hyper-V Cloud Reference Architecture

A Dell Technical White Paper

For the latest information, see www.microsoft.com/privatecloud.

April 2011

Contents

Tables

Figures

1. Introduction

The Microsoft Hyper-V Cloud Fast Track Program is a joint effort between Microsoft and its partner original equipment manufacturers (OEMs) to help organizations quickly develop and implement private clouds, while reducing both the cost and the risk.

This particular reference architecture combines Microsoft software, consolidated guidance, and validated configurations with Dell partner technology—such as computing power, network and storage architectures, and value-added software components. It is designed to provide an overview of Dell and Microsoft guiding principles and design criteria for this solution, and to illustrate how this Hyper-V technology cloud solution conforms to these principles.

The private cloud model provides much of the efficiency and agility of cloud computing, along with the increased control and customization achieved through dedicated private resources. With the Microsoft Hyper-V Cloud Fast Track program, Microsoft and its hardware partners provide organizations both the control and the flexibility required to reap the full benefits of the private cloud.

For more information about how to configure, deploy, and operate the Dell private cloud solution based on the Microsoft Hyper-V cloud reference architecture, see the Business-Ready Configuration for Microsoft Hyper-V R2 on Dell PowerEdge R-Series Servers with EqualLogic Storage solutions guide, or contact a Dell or Microsoft Fast Track representative for guidance.

2. Hyper-V Cloud Fast Track Program Description

2.1 The Program

The Microsoft Hyper-V Cloud Fast Track Program is a joint reference architecture for building private clouds that combines Microsoft software, consolidated guidance, and validated configurations with OEM partner technology, including computing power, network and storage architectures, and value-added software components.

Hyper-V Cloud Fast Track solutions provide a turnkey approach for delivering preconfigured and validated implementations of the private cloud. With local control over data and operations, IT professionals can dynamically pool, allocate, and manage resources for agile infrastructures-as-a-service. In addition, business unit managers can deploy line-of-business applications with speed and consistency using self-provisioning (and decommissioning) and automated data center services in a virtualized environment.

2.2 Business Benefits

The Microsoft Hyper-V Cloud Fast Track Program provides a reference architecture for building private clouds on each organization’s unique terms. Each fast-track solution helps organizations implement private clouds with increased ease and confidence. Among the benefits of the Microsoft Hyper-V Cloud Fast Track Program are faster deployment, reduced risk, and a lower cost of ownership.

Faster deployment:

- End-to-end architectural and deployment guidance
- Streamlined infrastructure planning due to predefined capacity
- Enhanced functionality and automation through deep knowledge of infrastructure
- Integrated management for virtual machine (VM) and infrastructure deployment
- Self-service portal for rapid and simplified provisioning of resources

Reduced risk:

- Tested, end-to-end interoperability of compute, storage, and network
- Predefined, out-of-box solutions based on a common cloud architecture that has already been tested and validated
- High degree of service availability through automated load balancing

Lower cost of ownership:

- A cost-optimized, platform- and software-independent solution for rack system integration
- High performance and scalability with Windows Server 2008 R2 operating system advanced platform editions of Hyper-V technology
- Minimized backup times and fulfilled recovery time objectives for each business-critical environment

3. Technical Overview

3.1 Hyper-V Cloud Architecture Principles

Microsoft Hyper-V cloud architecture principles conform to the cloud attributes outlined by the National Institute of Standards and Technology (NIST) definition of cloud computing version 15: on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service. Similarly, the Microsoft Hyper-V cloud architecture is based on seven principles: resource pooling, elasticity and the perception of infinite capacity, perception of continuous availability, predictability, metering/chargeback, multitenancy, and security and identity:

3.1.1 Resource Pooling

Resource optimization is a principle that drives efficiency and cost reduction. It is primarily achieved through resource pooling. Abstracting the platform from the physical infrastructure enables optimization of resources through shared use. Allowing multiple consumers to share resources results in higher resource utilization and a more efficient use of the infrastructure. Optimization through abstraction is a key factor behind many of the Hyper-V cloud principles, ultimately helping to improve agility and drive down costs.

3.1.2 Elasticity and Perception of Infinite Capacity

From a consumer's perspective, cloud services appear to have infinite capacity. Using the "electric utility provider" as a metaphor, consumers use as much or as little of the service as they need. This utility approach to computing requires that capacity planning be proactive so that requests can be satisfied on demand. Applying the principle of elasticity reactively and in isolation often leads to inefficient use of resources and unnecessary costs. But when an organization encourages desired

consumer behavior, it can use this principle to balance the desire for agility with the cost of unused capacity.

3.1.3 Perception of Continuous Availability

From the consumer's perspective, cloud services always appear to be available when needed. The consumer should never experience an interruption of service, even if failures occur within the Hyper-V cloud environment. To achieve this perception, organizations must take a mature service management approach that combines inherent application resiliency with infrastructure redundancies in a highly automated environment. As with the perception of infinite capacity, this principle can only be achieved in conjunction with the other Hyper-V cloud principles.

3.1.4 Predictability

Whether you're a consumer or a provider, predictability is a fundamental cloud principle. From the vantage point of the consumer, cloud services should be consistent; they should have the same quality and functionality any time they are used. To achieve predictability, a provider must deliver an underlying infrastructure that assures a consistent experience to the hosted workloads. This consistency is achieved through the homogenization of underlying physical servers, network devices, and storage systems.

From the provider's service management perspective, this predictability is driven through the standardization of service offerings and processes. The principle of predictability is needed to ensure service quality.

3.1.5 Metering and Chargeback (Service Providers' Approach to Delivering IT)

When IT professionals are asked to deliver a service to the business, they typically purchase the necessary components and then build an infrastructure specific to the service requirements. This typically results in longer time to market and increased costs due to duplicate infrastructure. In addition, the service often fails to meet business expectations of agility and cost control. The problem is often compounded when an existing service needs to be expanded or upgraded.

Taking a service provider's perspective toward delivering infrastructure transforms the IT approach. If infrastructure is provided as a service, IT can use a shared resource model that makes it possible to achieve economies of scale. This, combined with the other principles, helps the organization to realize greater agility at lower cost.

3.1.6 Multitenancy

Multitenancy refers to the ability of the infrastructure to be logically subdivided and provisioned to different organizations or organizational units. The traditional example is a hosting company that provides servers to multiple customer organizations. Increasingly, this is also a model being utilized by a centralized IT organization that provides services to multiple business units within a single organization, treating each as a customer or tenant.

3.1.7 Security and Identity

Security for the Hyper-V cloud is founded on two paradigms: protected infrastructure and network access.

Protected infrastructure takes advantage of security and identity technologies to ensure that hosts, information, and applications are secured across all scenarios in the data center, including the physical (on-premises) and virtual (on-premises and cloud) environments.

Application access helps ensure that IT managers can extend vital applications to internal users as well as to important business partners and cloud users.

Network access uses an identity-centric approach to ensure that users—whether they're based in the central office or in remote locations—have more secure access no matter what device they're using. This helps ensure that productivity is maintained and that business gets done the way it should.

Most important from a security standpoint, the secure data center makes use of a common integrated technology to assist users in gaining simple access using a common identity. Management is integrated across physical, virtual, and cloud environments so that businesses can take advantage of all capabilities without the need for significant additional financial investments.

3.2 Conceptual Architecture

One of the key drivers of the layered approach to infrastructure architecture presented in this white paper is to allow complex workflow and automation to be developed over time. This is accomplished by creating a collection of simple automation tasks, assembling them into procedures that are managed by the management layer, and then creating workflows and process automation that are controlled by the orchestration layer.

In a modular architecture, the concept of a *scale unit* refers to the extent to which a module can scale before another module is required. For example, an individual server can be considered a scale unit. In this case, the single server can be expanded to a certain point in terms of central processing units (CPU) and random access memory (RAM). However, beyond these limits an additional server is required to continue scaling. Each scale unit also has associated amounts of labor, such as physical installation labor and configuration labor. With large scale units such as a preconfigured full rack of servers, the labor overhead can be minimized.

It is critical to know the scale limits of all components, both hardware and software, to determine the optimum scale units for input to the overall architecture. Scale units enable the documentation of all the requirements (such as space, power, HVAC, and connectivity) that are needed for implementation.

3.3 Servers

The hardware architecture choices that are available to data center architects are constantly evolving. Choices range from rack-mounted servers, to tightly integrated, highly-redundant blade systems, to container models. The same spectrum exists for storage and networking equipment.

Server scale limits are well published and include factors such as the number and speed of CPU cores, maximum amount and speed of RAM, and the number and type of expansion slots. Particularly important are the number and type of onboard input-output (I/O) ports, as well as the number and type of supported I/O cards. Both Ethernet and Fibre Channel expansion cards often provide multiport options where a single card can have four ports. Additionally, in blade server architectures, there are often limitations on the number of I/O cards and supported combinations. It is important to be aware of these limitations, as well as the oversubscription ratio between blade I/O ports and any blade chassis switch modules.

A single server is not typically a good scale unit for a Hyper-V cloud solution, due to the amount of overhead and cost required to install and configure an individual server, as well as the lack of high availability.

3.4 Storage

Storage architecture is a critical design consideration for Hyper-V cloud solutions. The topic is challenging, because it is rapidly evolving in terms of new standards, protocols, and implementations. Storage and the support of storage networking are critical to the overall performance of the environment; however, they also can play a large role in the overall cost, because storage tends to be one of the more expensive items in a Hyper-V cloud solution.

Storage architectures today have several layers including the storage arrays, the storage network, the storage protocol, and for virtualization, the clustered volume manager that utilizes the physical storage.

One of the primary objectives of the private cloud is to enable rapid provisioning and deprovisioning of VMs. Doing so at large scale requires tight integration with the storage architecture and robust automation. Provisioning a new VM on an already existing logical unit number (LUN) is a simple operation. However, provisioning a new LUN, adding it to a host cluster, and then provisioning a new VM on that LUN involves relatively complex tasks that also greatly benefit from automation.

3.5 Networking

Many network architectures include a tiered design with three or more layers such as core, distribution, and access. Designs are driven by the port bandwidth and quantity required at the edge, as well as the ability of the distribution and core layers to provide higher-speed uplinks to aggregate traffic. Additional considerations include Ethernet broadcast boundaries and limitations, and spanning tree and or other loop avoidance technologies.

A dedicated management network is a common feature of advanced data center virtualization solutions. Most virtualization vendors recommend that hosts be managed via a dedicated network to avoid competition with guest traffic needs and to provide a degree of separation for security and ease of management purposes. This typically implies dedicating one network interface card (NIC) per host and one port per network device to the management network.

With advanced data center virtualization, a frequent use case is to provide isolated networks where different “owners” such as particular departments or applications are provided their own dedicated networks. Multitenant networking refers to the use of technologies such as virtual local area networks (VLANs) or Internet Protocol security (IPSec) isolation techniques to provide dedicated networks that utilize a single network infrastructure or wire.

Managing the network environment in an advanced data center virtualization solution can present challenges that must be addressed. Ideally, network settings and policies are defined centrally and applied universally by the management solution. In the case of IPSec-based isolation, this can be accomplished using the Active Directory service and Group Policy to control firewall settings across the hosts and guest as well as the IPSec policies controlling network communication.

For VLAN-based network segmentation, several components—including the host servers, host clusters, Microsoft System Center Virtual Machine Manager, and the network switches—must be configured correctly to enable both rapid provisioning and network segmentation. With Hyper-V and host clusters, identical virtual networks must be defined on all nodes so that a VM can fail over to any node and maintain its connection to the network. At large scale, this can be accomplished via scripting with the Windows PowerShell command-line interface.

3.6 Virtualization

The virtualization layer is one of the primary enablers in mature IT environments. The decoupling of hardware, operating systems, data, applications, and user state opens up a wide range of options for better management and distribution of workloads across the physical infrastructure. The ability of the virtualization layer to migrate running VMs from one server to another with zero downtime, as well as many other features that are provided by hypervisor-based virtualization technologies, provides a rich set of capabilities. These capabilities can be utilized by the automation, management, and orchestration layers to maintain desired states (such as load distribution) or to proactively address decaying hardware (such as prefailure detection) or other issues that would otherwise cause faults or service disruptions.

As with the hardware layer, the virtualization layer must be able to be managed by the automation, management, and orchestration layers. The abstraction of software from hardware that virtualization provides moves the majority of management and automation into the software space, instead of requiring people to perform manual operations on physical hardware.

3.7 Automation

The ability to automate all expected operations over the lifetime of a hardware or software component is critical. Without this capability deeply embedded across all layers of the infrastructure, dynamic processes will grind to a halt as soon as user intervention or other manual processing is required.

Windows PowerShell and several other foundational technologies, including Windows Management Instrumentation (WMI) and Web Services for Management (WS-Management), provide a robust automation layer across nearly all Microsoft products, as well as a variety of non-Microsoft hardware and software. This layer provides a single automation framework and scripting language to be used across the entire infrastructure.

The automation layer is made up of foundational automation technologies plus a series of single-purpose commands and scripts that perform operations such as starting or stopping a VM, rebooting a server, and applying a software update. These atomic units of automation are combined and executed by higher-level management systems. The modularity of this layered approach dramatically simplifies development, debugging, and maintenance.

3.8 Management

The management layer consists of the tools and systems that are utilized to deploy and operate the infrastructure. In most cases, this consists of a variety of toolsets for managing hardware, software, and applications. Ideally, all components of the management system would use the automation layer and not introduce their own protocols, scripting languages, or other technologies, because this increases complexity and may require additional staff expertise.

The management layer is utilized to perform activities such as provisioning the storage area network (SAN), deploying an operating system, or monitoring an application. A key attribute of the management layer is its ability to manage and monitor every single component of the infrastructure remotely and to capture the dependencies among all of the infrastructure components.

3.9 Orchestration

The orchestration layer uses the management and automation layers. In much the same way that an enterprise resource planning (ERP) system manages a business process such as order fulfillment and handles exceptions such as inventory shortages, the orchestration layer provides an engine for

IT-process automation and workflow. The orchestration layer is the critical interface between the IT organization and its infrastructure and transforms intent into workflow and automation.

Ideally, the orchestration layer provides a graphical user interface in which complex workflows that consist of events and activities across multiple management-system components can be combined, to form an end-to-end IT business process such as automated patch management or automatic power management. The orchestration layer must provide the ability to design, test, implement, and monitor these IT workflows. Microsoft System Center Opalis is an automation platform for orchestrating and integrating IT tools that provides this functionality.

3.10 Service Management

The service management layer provides the means for automating and adapting IT service management best practices, such as those found in Microsoft Operations Framework (MOF) and the Information Technology Infrastructure Library (ITIL), to provide built-in processes for incident resolution, problem resolution, and change control. By providing an integrated service management platform, Microsoft System Center Service Manager can reduce costly downtime and improve service in the data center.

3.11 Tenant/User Self-Service

The tenant/user self-service layer provides an interface for Hyper-V cloud tenants or authorized users to request, manage, and access services, such as VMs, that are provided by the Hyper-V cloud architecture. Using role-based access control and authorization, the self-service layer provides the ability to delegate certain aspects of administration (such as starting/stopping VMs) to designated “tenant administrators.”

4. Reference Architecture

4.1 Workload Categories

4.1.1 Server Virtualization and Consolidation

Server virtualization is based on the abstraction of physical system resources so that multiple logical partitions can be created and can host a heterogeneous set of operating systems that run simultaneously on a single physical server.

Rather than paying for many under-utilized servers, each dedicated to a specific workload, server virtualization allows those workloads to be consolidated onto a smaller number of more efficiently utilized physical systems. Server virtualization provides the following benefits:

- Consolidates multiple, underutilized physical servers on a single host, running VMs
- Reduces workforce, space, and kilowatts by taking advantage of virtualization for server consolidation and agility
- Helps save money because less management, less space, and fewer kilowatt hours are needed

Virtualization can also help to simplify and accelerate provisioning, one of the tenets of the private cloud. The acquisition of workload resources and hardware can be decoupled. Adding the capability that is required for a particular business process (such as a web commerce engine) becomes streamlined and immediate. In a more advanced virtualized environment, workload requirements can be self-provisioning, which results in dynamic resource allocation.

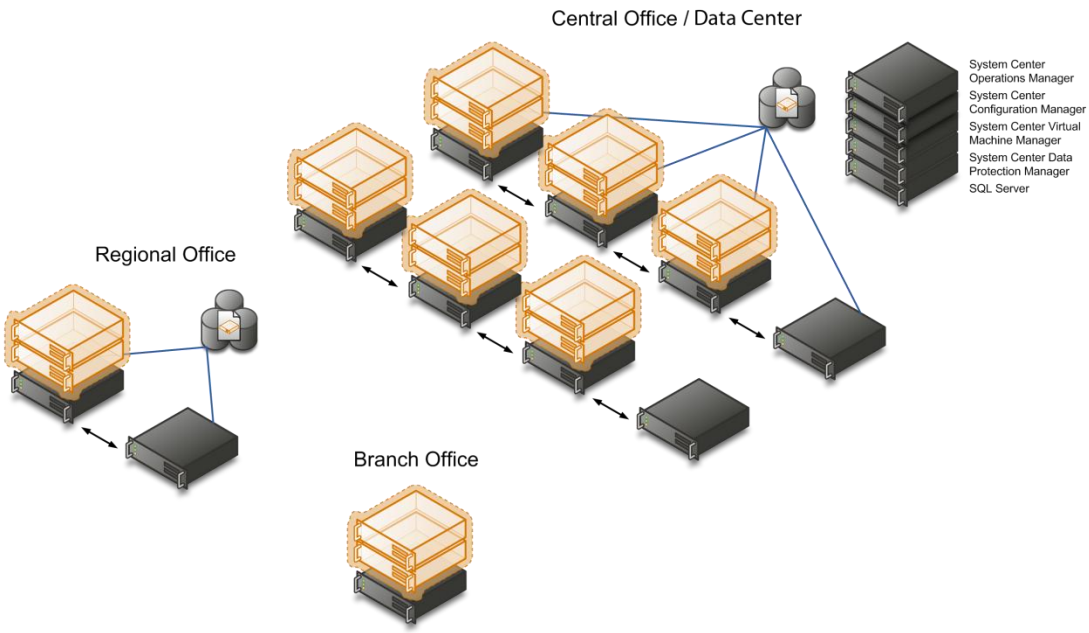

Figure 1. Server-based Consolidation Through Virtualization

Although virtualization-based server consolidation can provide many benefits, it can also add complexity if the environment is not managed properly. The savings from hardware consolidation could be offset by increases to IT management overhead. Because creating VMs is so easy, an unintentional and unnecessary virtual sprawl can result that far exceeds physical server capacity and that outpaces

the tools used to manage VMs. A properly managed virtual infrastructure, however, automatically determines which servers are the best candidates for virtualization, allowing administrators to initiate automated processes to convert those servers to VMs and provision them to the right hosts in minutes—compared to the weeks or months it would take to procure and configure physical servers manually.

4.1.2 Virtual Desktop Infrastructure

Virtual desktop infrastructure (VDI) allows IT managers to deploy desktops in VMs on secure and centralized hardware. A centralized and optimized virtual desktop enables IT staff to build a more agile and efficient IT infrastructure, while allowing users to access and run their desktop applications wherever they may be. Flexible desktop scenarios using the Windows operating system give organizations the ability to choose the client computing scenarios that best meet the unique needs of their business.

When an organization is managing its virtual infrastructure with the same tools it uses to manage its physical assets, this can reduce system complexity streamline and changes made to the overall infrastructure. By using some or all of these technologies together, organizations can provide very flexible solutions to support many user scenarios.

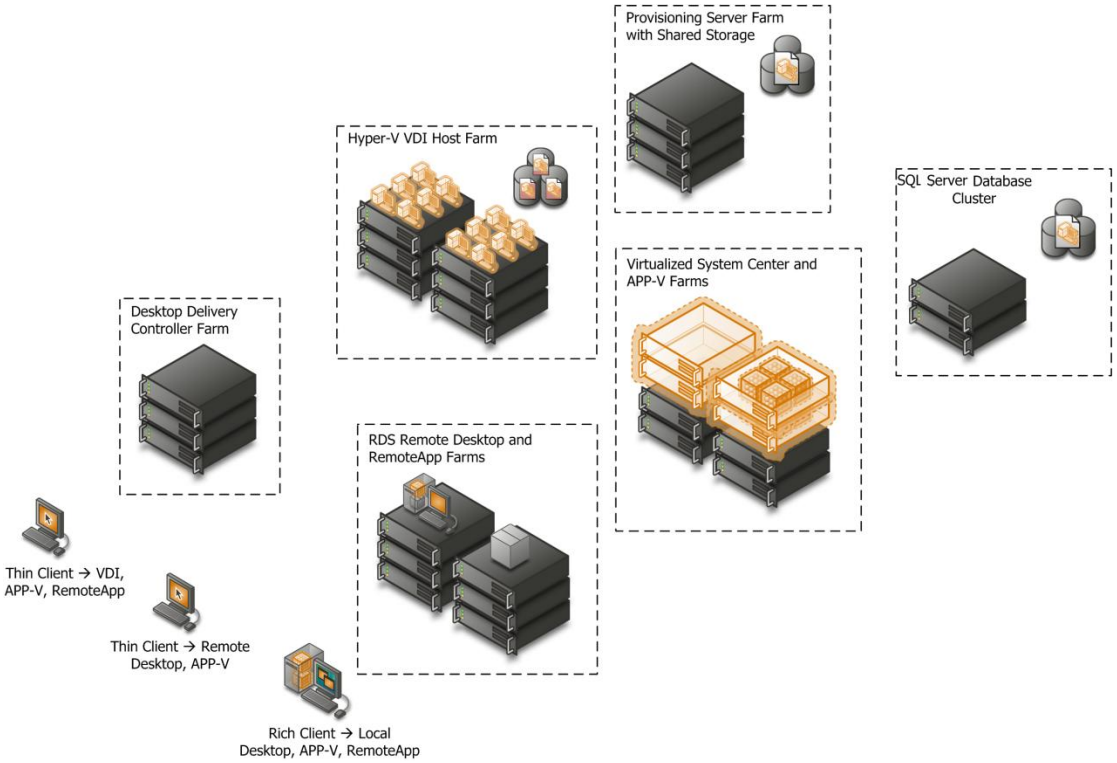

Figure 2. Virtualization Desktop Infrastructure

4.2 Logical Architecture

A private cloud is far more than a highly available infrastructure that provides computing resources to higher-level applications. With cloud computing, a fundamental shift is that of IT moving from server *operator* to service *provider*. This requires a set of services to accompany the infrastructure, such as reporting, usage metering, and self-service provisioning. If these services are unavailable, then the cloud "service layer" is unavailable, and IT is little more than a traditional data center. For this reason, high availability must also be provided to the management systems.

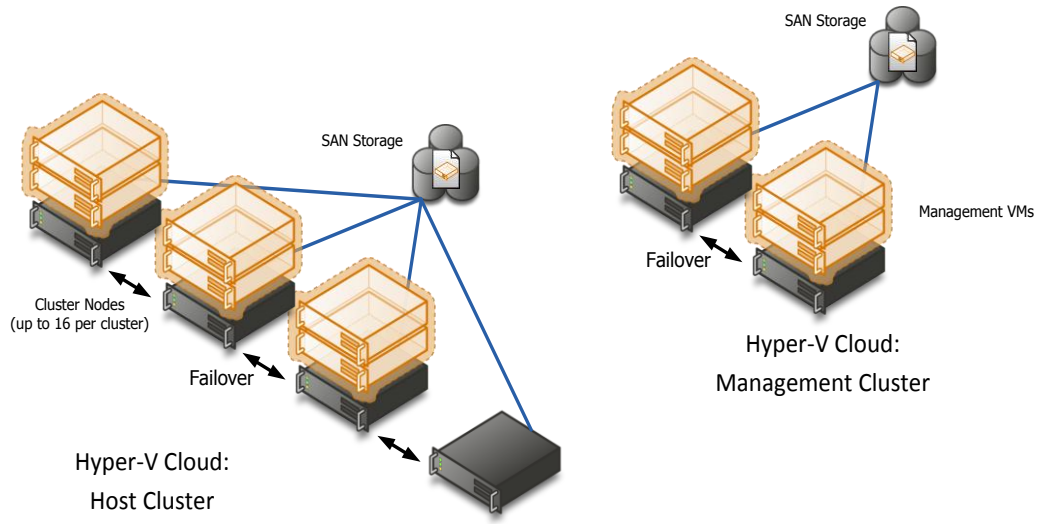

Figure 3. Hyper-V Cloud Logical Architecture

4.3 Server Architecture

The host server architecture is a critical component of the virtualized infrastructure, as well as a key variable in the consolidation ratio and cost analysis. The ability of the host server to handle the workload of a large number of consolidation candidates increases the consolidation ratio and helps provide the desired cost benefit.

The system architecture of the host server refers to the general category of the server hardware itself. The primary factor to keep in mind when selecting system architectures is that each Hyper-V host will contain multiple guests with multiple workloads. Processor, RAM, storage, and network capacity are critical, as well as high I/O capacity and low latency. The host server must be able to provide the required capacity in each of these categories.

Dell has validated an assortment of configurations for Fast Track architecture, consisting of PowerEdge servers, Dell EqualLogic SANs, and PowerConnect networking. These Dell systems are all available with enterprise features that enable high performance and highly available systems. These features are the core enablers of enterprise cloud computing. Dell Fast Track architecture configurations use

- Microsoft Hyper-V technology
- Dell PowerEdge R710 and R610 rack servers
- Dell EqualLogic PS6010XV storage
- 10-gigabit network that utilizes Dell PowerConnect 8024F series switches

This section presents an overview of the Dell components that have been validated for the Fast Track architecture. For more information about Dell configurations, see Business-Ready Configuration for Microsoft Hyper-V R2 on Dell PowerEdge R-Series Servers with EqualLogic Storage.

4.3.1 Rack Server Design

Dell has tested and validated rack mount servers in the eleventh-generation PowerEdge portfolio. These servers are available with ultra-efficient and Energy Star–compliant power supplies. They are combined with dynamic, power-efficient fans with optimized airflow design, to provide enterprise performance in an environmentally friendly footprint. The R710 chassis design provides redundant power connectivity through multiple hot-swappable power supplies. For more PowerEdge server specifications, see www.dell.com/poweredge.

4.3.2 Server Design

The Dell R710 servers validated in this Fast Track solution contain Xeon 5500 series processors with 4-core configurations. The Dell R710 servers can include either Xeon 5500 or 5600 series Xeon processors with up to 6-core configurations. The R710 also supports up to 288 gigabytes (GB) of double data rate 3 (DDR3) memory running at 1333 megahertz (MHz).

Local storage redundancy is provided through a RAID (Redundant Array of Independent Disks) controller: the Dell PowerEdge Expandable RAID Controller (PERC) 6/I adapter, which supports SAS (Serial Attached SCSI [small computer system interface]) devices. This is paired with two 146-GB 10K RPM SAS drives in a RAID 1 configuration.

4.3.3 Server Storage Connectivity

The Dell validated solution utilizes redundant 10-gigabit network interfaces for storage connectivity. The Dell solution uses two Intel X520 DA dual port network adapters in an enhanced small form-factor pluggable (SFP+) form factor. The adapters use Microsoft Multipath I/O (MPIO) for fault tolerance.

4.3.4 Server Network Connectivity

At the heart of the solution's network configuration are the Dell PowerConnect switches. These managed Layer 3 Ethernet switches offer the enterprise-level performance required for the configuration. These switches are capable of supporting all configuration networks with the use of VLANs for traffic isolation, and network teaming for redundancy. For more Dell PowerConnect information, see www.dell.com/powerconnect.

The Dell R710 servers again use the Intel X520 DA network adapter for network connectivity. Network teaming is utilized to provide VLANs and fault tolerance. The VLAN configuration of each host provides a VLAN for host management, cluster private communication, and live migration, as well as a VLAN dedicated to virtual machines. One port from each of the two adapters is used for the server connectivity, and the second port from each adapter is used for Internet SCSI (iSCSI) traffic.

4.3.5 Server High Availability (HA) and Redundancy

All Dell PowerEdge servers, PowerConnect network devices, and Dell EqualLogic storage hardware have redundant power supplies and multiple cooling fans. Additionally, the PowerEdge servers all are available with RAID arrays for fault-tolerant disk drives.

4.4 Storage Architecture

The storage design for any virtualization-based solution is a critical element that is typically responsible for a large percentage of the solution's overall cost, performance, and agility.

4.4.1 Storage Options

Although many storage options exist, organizations should choose their storage devices based on their specific data management needs. Storage devices are typically modular and flexible midrange and high-end SANs. Modular midrange SANs are procured independently and can be chained together to provide greater capacity. They are efficient, can grow with the environment as needed, and require less up-front investment than high-end SANs. Large enterprises may have larger storage demands and may need to serve a larger set of customers and workloads. In this case, high-end SANs can provide the highest performance and capacity. High-end SANs typically include more advanced features such as continuous data availability through technologies like dispersed cluster support.

4.4.2 Cluster Shared Volumes

Windows Server 2008 R2 includes the first version of Windows failover clustering to offer a distributed file access solution. Cluster Shared Volumes (CSV) in Windows Server 2008 R2 is exclusively for use with the Hyper-V role and enables all nodes in the cluster to access the same cluster storage volumes at the same time. This enhancement eliminates the one VM-per-LUN requirement of previous Hyper-V versions without using a third-party file system. CSV uses standard New Technology File System (NTFS) and has no special hardware requirements. From a functional standpoint, if the storage is suitable for failover clustering, it is suitable for CSV.

CSV provides not only shared access to the disk, but also storage path I/O fault tolerance (dynamic I/O redirection). In the event that the storage path on one node becomes unavailable, the I/O for that node is rerouted via Server Message Block (SMB) through another node.

CSV maintains metadata information about the volume access and requires that some I/O operations take place over the cluster communications network. One node in the cluster is designated as the coordinator node and is responsible for these disk operations. However, all nodes in the cluster can read/write directly and concurrently to the same volume (not the same file) through the dedicated storage paths for disk I/O, unless a failure scenario occurs as described above.

1. CSV Characteristics

Table 1 below shows the characteristics that are defined by the New Technology File System (NTFS) and are inherited by CSV.

Table 1. CSV Parameter Characteristics

CSV Parameter	Characteristic
Maximum Volume Size	256 terabytes (TB)
Maximum # Partitions	128
Directory Structure	Unrestricted
Maximum Files per CSV	4+ Billion
Maximum VMs per CSV	Unlimited

2. CSV Volume Sizing

Because all cluster nodes can access all CSV volumes simultaneously, IT managers can now use standard LUN allocation methodologies based on performance and capacity requirements of the expected workloads. Generally speaking, isolating the VM operating system I/O from the application data I/O is a good start, in addition to application-specific considerations such as segregating database I/O from logging I/O and creating SAN volumes and storage pools that factor in the I/O profile itself (that is, random read and write operations versus sequential write operations).

The architecture of CSV differs from traditional clustered file systems, which frees it from common scalability limitations. As a result, there is no special guidance for scaling the number of Hyper-V nodes or VMs on a CSV volume. The important thing to keep in mind is that all VM virtual disks running on a particular CSV will contend for storage I/O.

Also worth noting is that individual SAN LUNs do not necessarily equate to dedicated disk spindles. A SAN storage pool or RAID array may contain many LUNs. A LUN is simply a logic representation of a disk provisioned from a pool of disks. Therefore, if an enterprise application requires specific storage I/O operations per second (IOPS) or disk response times, IT managers must consider all the LUNs in use on that storage pool. An application that would require dedicated physical disks were it not virtualized may require dedicate storage pools and CSV volumes running within a VM.

Consider the following when setting up your CSV infrastructure:

- At least four CSVs per host cluster are recommended for segregating operating system I/O, random read/write I/O, sequential I/O, and other VM-specific data.
- Create a standard size and IOPS profile for each type of CSV LUN to utilize for capacity planning. When additional capacity is needed, provision additional standard CSV LUNs.
- Consider prioritizing the network used for CSV traffic. For more information, see Designating a Preferred Network for Cluster Shared Volumes Communication in the Microsoft TechNet Library.

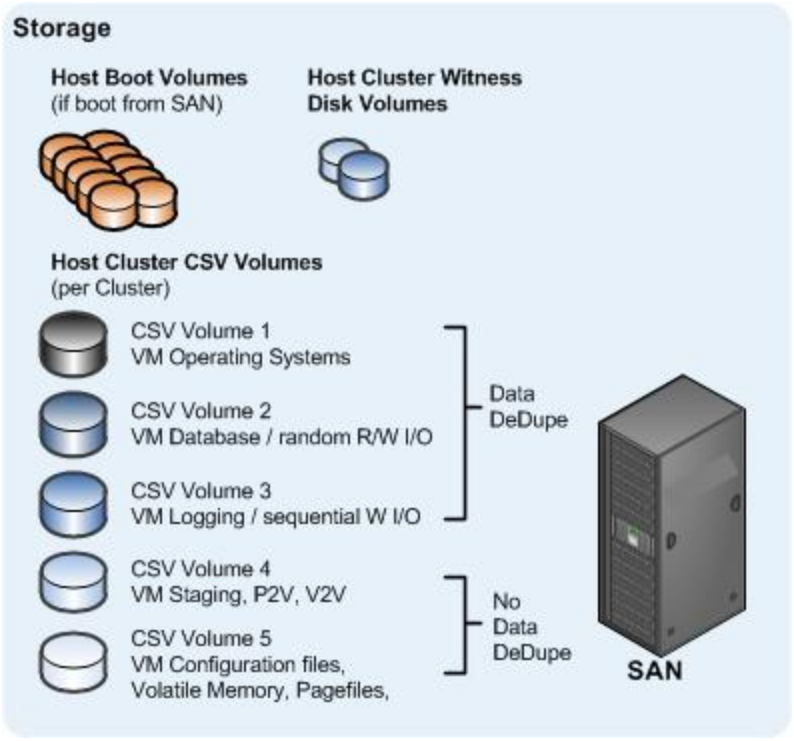

Figure 4. Example of a Common CSV Design for a Large Hyper-V Cluster

The Dell solution provides support for CSV, offering the capability to host multiple virtual machines on a single volume and to migrate those virtual machines independently among the Dell R710 servers in the cluster. Each R710 server can simultaneously read and write directly (using its 10-gigabit iSCSI adapters) to the volume on the storage array.

4.4.3 SAN Design

Dell has validated an EqualLogic PS6010XV iSCSI SAN solution for this configuration. The EqualLogic iSCSI system provides a high-performance and highly available storage solution.

The Dell EqualLogic PS Series of iSCSI storage arrays offers high performance, reliability, intelligent automation, and seamless virtualization of a single pool of storage to enable simplified enterprise storage deployment and management. Table 2 lists the EqualLogic PS 6000 series storage array features.

Table 2. Features of the EqualLogic PS 6000 Series Storage Array

Feature	Limit
Maximum members in a group	16
Maximum members in a pool	8
Maximum volumes in a group	1024
Maximum number of pools in a group	4
Snapshots in a group	10000
Snapshots of a volume	512
Replicas of a volume	512
Replication partners per group	16
Replication partners per volume	1
Volumes in a collection	8
Collections in a group (snapshot and replication)	100
Volume connections per pool and per group (each time an iSCSI initiator connects to a volume counts as a connection)	4096 with 4 pools, 1024 per pool
Access control records per volume and its snapshots	16
Maximum volumes enabled for replication	128
Simultaneous management sessions (any combination of graphical user interface, Telnet, or scripting sessions)	7

For more information about Dell storage solutions, visit www.dell.com/equallogic.

1. High Availability

The Dell validated storage configuration provides high availability throughout the physical architecture.

- Redundant switches, switch paths, cables, and network interfaces are used.
- All systems are cabled and powered to ensure that there is no single point of failure.
- Multipathing is used on the host.

The usages of VLANs and their paths are illustrated in Figure 5.

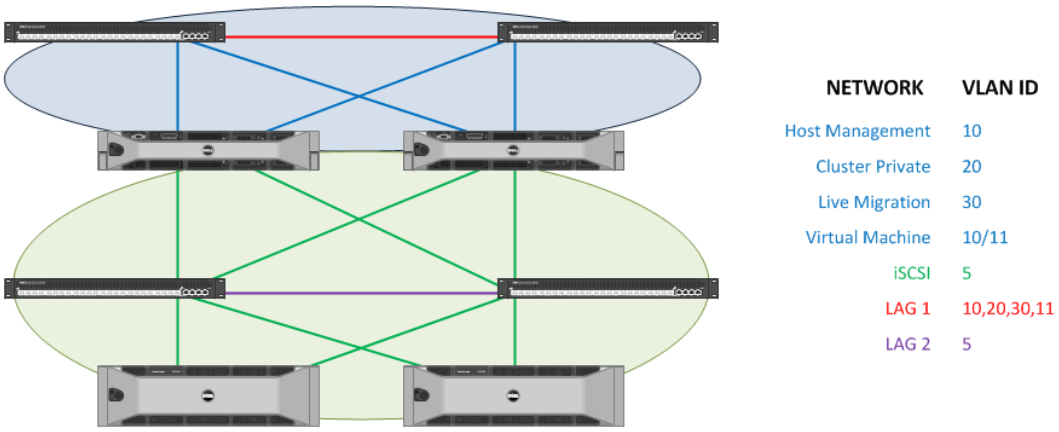

Figure 5. Physically Autonomous iSCSI with VLANs

The SAN uses the multipathing solution from the Microsoft MPIO driver and is enhanced with the device-specific module (DSM) from Dell EqualLogic to provide fault tolerance at the storage level. The EqualLogic SAN is made highly available through redundant power from independent power distribution units (PDUs), redundant storage controllers, redundant target ports of NICs per controller, redundant network switches, cables and interfaces, and RAIDs with hot spare drives. Data storage redundancy is also a capability of the EqualLogic PS6000 series SAN through volume mirroring and synchronous or asynchronous replication.

2. Performance

Storage performance is a complex mix of drive, interface, controller, cache, protocol, SAN, HBA, driver, and operating system considerations. The overall performance of the storage architecture is typically measured in terms of maximum throughput, maximum IOPS, and latency or response time. Although each of these performance measurements is important, IOPS and latency are the most relevant to server virtualization.

Most modern SANs use a combination of high-speed disks, slower-speed disks, and large memory caches. A storage controller cache can improve performance during burst transfers or when the same data is accessed frequently by storing it in the cache memory, which is typically several orders of magnitude faster than the physical disk I/O. However, it is not a substitute for adequate disk spindles because caches are ineffective during heavy write operations.

For storage at the hardware level, the Dell EqualLogic SAN provides high-speed drives and large caches, as well as 10-gigabit interfaces for increased throughput. The EqualLogic SAN is also available with solid-state drives that can further increase performance. It is also capable of a variety of RAID levels that can increase performance; when the SAN is joined by additional members, these RAID levels can be tiered to provide the best performance given specific size constraints.

3. Drive Types

The type of hard drive utilized in the host server or the storage array has the most significant impact on the overall storage architecture performance. As with the storage connectivity, high IOPS and low latency are more critical than maximum sustained throughput when it comes to host server sizing and guest performance. When selecting drives, this translates into selecting those with the highest rotational speed and lowest latency possible. Utilizing 15-KB RPM drives over 10-KB RPM drives can result in up to 35 percent more IOPS per drive.

The EqualLogic PS6010XV uses 600-GB 15-KB RPM SAS drives. For more information, see Business-Ready Configuration for Microsoft Hyper-V R2 on Dell PowerEdge R-Series Servers with EqualLogic Storage (page 15).

4. RAID Array Design

The RAID type should provide both high availability and high performance, even in the event of disk failures and RAID parity rebuilds. In general, RAID level 10 (0+1), 5, or 6 is recommended for virtual machine volumes. RAID 1 is also acceptable for host boot volumes. A fault-tolerant RAID configuration is required. Workloads will dictate the specific RAID design.

The Dell configuration is validated with RAID 6. The EqualLogic SAN does, however, support RAID levels 1, 5, 6, 10, and 50.

5. Multipathing

In all cases, multipathing should be used. Generally, storage vendors build a DSM on top of Windows Server 2008 R2 MPIO software. Each DSM and HBA has its own unique multipathing options, recommended number of connections, and other particulars.

The DSM driver for EqualLogic is installed via the EqualLogic Host Integration Toolkit (HIT).The HIT includes software for configuring the MPIO settings, so that users can select a load balance policy (round robin, least queue depth, and failover only), connections per volume, connections per member, networks available for MPIO, and minimum adapter speed.

6. iSCSI

The Dell validated configuration uses a dedicated switch fabric for iSCSI networks. The traffic is also removed from the default VLAN 1 and placed on to VLAN 5.

- **Encryption and authentication.** The Dell configuration is isolated from all other traffic. However, if multiple clusters or systems are used on the same SAN, proper segregation or device isolation must be provided. In other words, the storage used by cluster A must be visible only to cluster A, and not to any other cluster, nor to a node from a different cluster. To achieve this isolation, the EqualLogic SAN provides authentication via Challenge-Handshake Authentication Protocol (CHAP), iSCSI initiator, or Internet Protocol (IP) address. The IP address is used to restrict access in the validated configuration.
- **Jumbo frames.** If supported at all points in the entire path of the iSCSI network, jumbo frames can increase throughput by up to 20 percent. Jumbo frames are supported in Hyper-V at the host and guest levels. The Dell EqualLogic iSCSI SAN utilizes jumbo frames by default. No further configuration is necessary at the SAN level. Hosts and intermediate I/O fabrics still require configuration.

7. Thin Provisioning

Particularly in virtualization environments, thin provisioning is a common practice. This allows for efficient use of the available storage capacity. The LUN and corresponding CSV can grow as needed, typically in an automated fashion (auto-grow), to ensure availability of the LUN. However, storage can become overprovisioned in this scenario, so careful management and capacity planning are critical.

The Dell EqualLogic SAN allows for thin provisioning both during and after volume creation. Volumes that are created as a thin-provisioned volume consume a minimum of 10 percent of the volume's logical size. With the EqualLogic the administrator can also configure notification thresholds on thin-provisioned volumes to alert the administrator that additional storage should be provisioned.

8. Volume Cloning

Volume cloning is another common practice in virtualization environments. This can be used for both host and VM volumes to dramatically decrease host installation times and VM provisioning times.

Volume cloning is available in the EqualLogic PS 6010XV arrays. The clones are identical copies of the source, including content and size, but have unique iSCSI qualified names (IQNs) and names and are available immediately. Clones in the PS 6010XV can be created from another volume in the pool, from a snapshot in the pool, or from an inbound replica from a secondary group.

9. Volume Snapshots

SAN volume snapshots offer a common method of providing a point-in-time, instantaneous backup of a SAN volume or LUN. These snapshots are typically block-level and only utilize storage capacity as blocks change on the originating volume.

The Dell EqualLogic PS 6010XV provides snapshot capabilities. The required snapshot space is tied directly to a volume and defaults to 100 percent of the size of the volume. This allows for the data of an entire volume to change and makes space available for the snapshot. The EqualLogic also allows snapshots to be assigned their own unique IQNs, which allows snapshots to be mounted just like any other iSCSI volume. Data can be extracted from the snapshot and restored if necessary.

10. Storage Tiering

Tiering storage is the practice of physically partitioning data into multiple distinct classes based on price, performance, or other attributes. Data can be dynamically moved among classes in a tiered storage implementation based on access activity or other considerations. This is normally achieved through a combination of varying types of disks that are used for different data types such as production, non-production, or backups.

Storage tiering is provided dynamically in EqualLogic SANs. As additional members are added to a group and configured to use different RAID levels, volumes can move to different pools without any interruption in service.

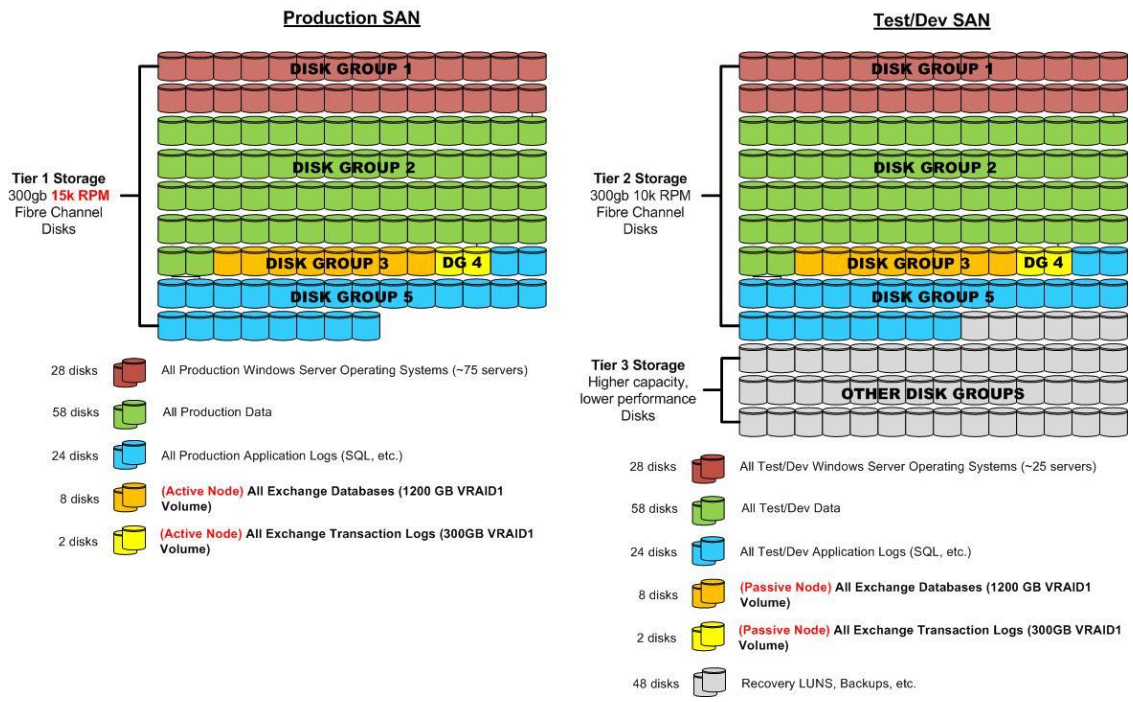

Figure 6. Tiered Storage Design

4.5 Network Architecture

4.5.1 Core, Distribution, and Access Network Design Tiers

Many network architectures include a tiered design with three or more tiers such as core, distribution, and access. Designs are driven by the port bandwidth and quantity required at the edge, as well as the ability of the distribution and core tiers to provide higher-speed uplinks to aggregate traffic. Additional considerations include Ethernet broadcast boundaries and limitations, and Spanning Tree Protocol (STP) and other loop-avoidance technologies.

Figure 7 illustrates the Dell tiered network design solution. The End of Row and Core switches are undefined, but the concept remains the same.

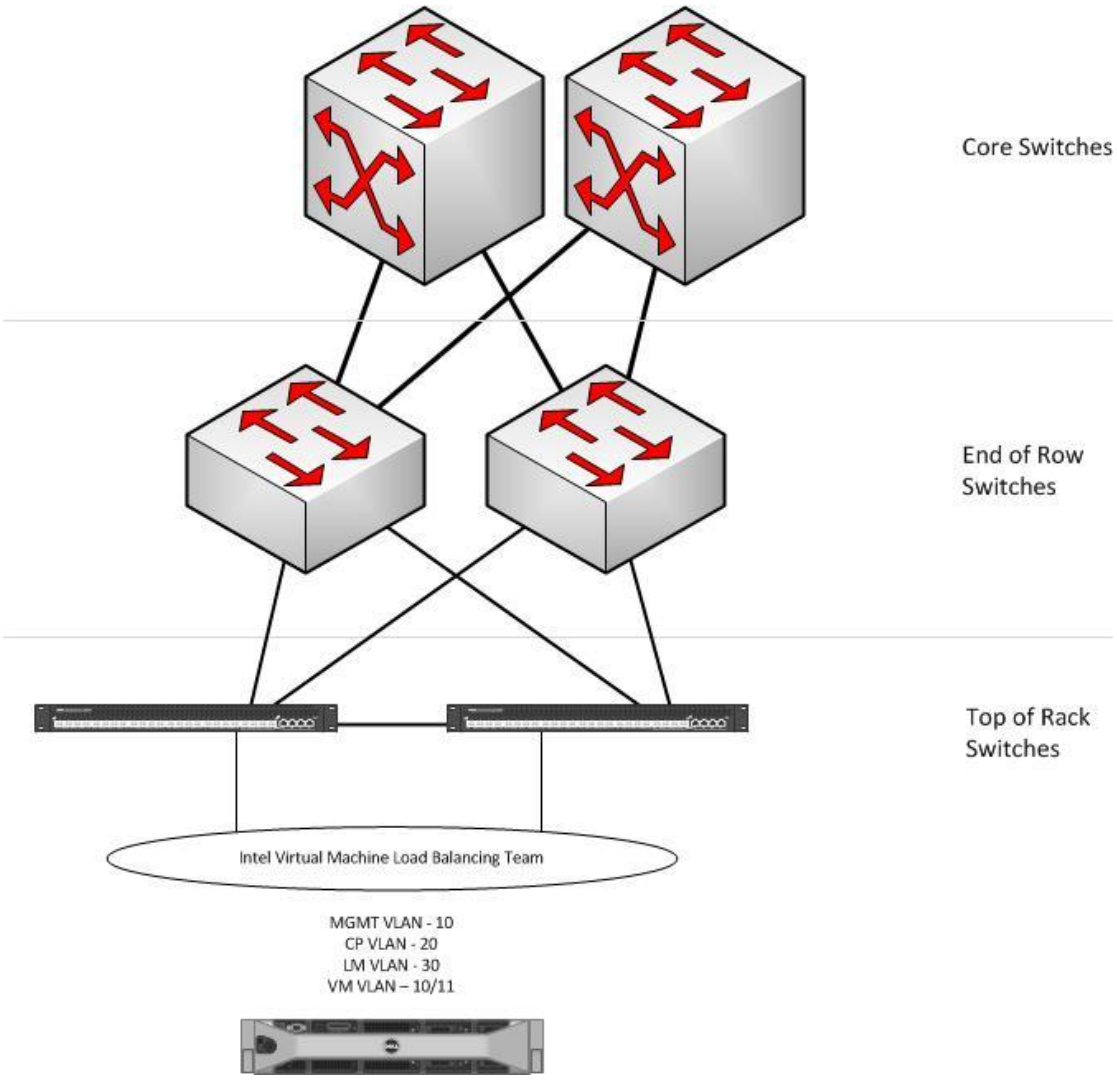

Figure 7. Dell Tiered Network Architecture

Note the following configuration details:

- VLANs 10, 11, 20, and 30 are defined on the PowerConnect 8024F switches.
- Link aggregation groups (LAGs) are configured for inter-switch and uplink traffic.
- The network design allows for the loss of any switch without dropping host server connectivity.

4.5.2 High Availability and Redundancy

Dell PowerConnect switches utilize multiple power supplies and multiple cooling fans for redundancy within each switch. Each switch is cabled to provide a path between the switches. Each inter-switch link is connected using a LAG, which increases the amount of bandwidth and provides fault tolerance at the switchport and cable levels. The switches are also configured to use STP to allow for multiple switching paths. More complex STP designs using MSTP are also possible with the PowerConnect switches, giving the network administrator the capability of segmenting traffic such as live migration or cluster private/CSV traffic from uplinks.

PowerEdge servers provide network fault tolerance through the use of software teaming. Dell has validated teaming software Intel network adapters in this configuration. The teaming drivers from Intel also provide VLAN capabilities.

The guidelines for fault-tolerant teaming and VLAN configurations are available in the Dell Business Ready Configuration Reference Architecture and are specific to each network environment. The network design must allow for the loss of any switch module or switch without dropping host server connectivity. For more information, see Business-Ready Configuration for Microsoft Hyper-V R2 on Dell PowerEdge R-Series Servers with EqualLogic Storage.

4.6 Virtualization Architecture

4.6.1 Windows Server 2008 R2 and Hyper-V Host Design

1. Operating System Configuration

The following list outlines the general considerations for the Hyper-V host operating system. Note that these are not meant to be installation instructions but rather the process requirements.

- Use Windows Server 2008 R2 Datacenter Edition, with either the full or server core installation option.
- Use the latest hardware device drivers.
- The Hyper-V parent partition operating system is domain-joined.
- Hyper-V server role and failover clustering features are required.
- Apply relevant Windows updates, and all relevant out-of-band updates that are not offered on Microsoft Update. For more information, see Hyper-V Update List for Windows Server 2008 R2 in the Microsoft TechNet Library.
- All nodes, networks, and storage must pass the Cluster Validation Wizard in Windows Server 2008 R2.

2. Performance Settings

The following Windows Server 2008 R2 Hyper-V network performance improvements should be considered for production use:

- **TCP checksum offload** benefits both CPU and overall network throughput performance, and is fully supported by live migration.
- **Jumbo frames** capability is extended to VMS with Hyper-V in Windows Server 2008 R2. Just as in physical network scenarios, jumbo frames add the same basic performance enhancements to virtual networking. That includes up to six times larger payloads per packet, which improves overall throughput and also reduces CPU utilization for large file transfers.
- **Virtual machine queue** (VMQ) allows the host's single NIC card to appear as multiple NICs to the VMs by allowing the host's NIC to direct memory access (DMA) packets directly into individual VM memory stacks. Each VM device buffer is assigned a VMQ, which avoids needless packet copies and route lookups in the virtual switch. The result is less data in the host's buffers and an overall performance improvement to I/O operations.

The Dell configuration uses jumbo frames and Transmission Control Protocol (TCP) offloading and is configured in the advanced properties of the Intel driver. The maximum frame size should be set. (This value also must be set on the switches in the iSCSI Fabric.) The result is less data in the host's buffers and an overall performance improvement to I/O operations.

3. IP Networks Configuration

Ensure that the following rules are followed when setting up the IP network:

- The cluster heartbeat network must be on a distinctly separate subnet from the host management network.
- The VM network adapter should not be shared with the host operating system and therefore should not have the TCP/IP protocols (IPv4 and IPv6) bound to it.
- The iSCSI network must be on a distinctly separate and isolated network, with a dedicated IP range used only for storage.

4. MPIO Configuration

Microsoft Multipath I/O (MPIO) architecture supports iSCSI, Fibre Channel, and Serial Attached Storage SAN connectivity by establishing multiple sessions or connections to the storage array.

Multipathing solutions use redundant physical path components—adapters, cables, and switches—to create logical paths between the server and the storage device. In the event that one or more of these components fails, causing the path to fail, multipathing logic uses an alternate path for I/O so that applications can still access their data. Each network interface card (in the iSCSI case) or HBA should be connected using redundant switch infrastructures to provide continued access to storage in the event of a failure in a storage fabric component.

Failover times vary by storage vendor, and can be configured by using timers in the Microsoft iSCSI Software Initiator driver. Default MPIO settings (two connections per slice, six connections per member, and least queue depth) are used for Dell EqualLogic DSM settings. The number of usable paths may be configured if necessary to fully utilize all available storage paths.

MPIO must be used on storage adapters. Follow MPIO best practices, as documented in Appendix B of the MPIO white paper Windows Server High Availability with Microsoft MPIO.

5. NIC Teaming Configuration

NIC teaming or link aggregation (IEEE 802.3ad) enables network maintenance to occur at all points within the data center network topology without affecting applications. This technology bonds physical NICs together to form one or more logical network team that sends traffic to all NICs in the team. This allows a single NIC, cable, or switch to sustain a planned or unplanned outage without disrupting the host's Ethernet traffic. The NIC manufacturer is also the software provider for the NIC teaming software. Each NIC teaming software application has its own unique set of requirements, features, teaming modes, and configuration recommendations. NIC teaming should be used to provide high availability to the VM networks.

Note: MPIO, not teaming, should be used for storage traffic in conjunction with iSCSI.

The Dell-validated solution uses NIC teaming software from Intel. The teaming mode used is Virtual Machine Load Balancing, which enables transmit load balancing while still providing fault tolerance in the event of network card, cable, or network device interruption or failure.

4.6.2 Hyper-V Host Cluster Design

A Hyper-V host cluster is a group of independent servers that work together to increase the availability of applications and services. The clustered servers (called nodes) are connected by physical cables and by software. In the case of unplanned downtime—if one of the nodes fails—another node begins to provide service (a process known as failover), which means the VMs on the failing node are

automatically restarted on another node in the cluster. In case of a live migration, one or more VMs are moved from one node to another node in the cluster, and users experience no perceptible service interruption.

The host servers are one of the critical components of a dynamic virtual infrastructure. Consolidation of multiple workloads onto the host servers requires that those servers be highly available. Windows Server 2008 R2 provides advances in failover clustering that enable high availability and live migration of VMs between physical nodes.

1. Server Topology

A Hyper-V cloud consists of at least two Hyper-V host clusters. The first consists of at least two nodes and is referred to as the management cluster. The second and any additional clusters are referred to as host clusters. Each host cluster can contain up to 16 nodes. Host clusters require some form of shared storage such as a Fibre Channel or iSCSI SAN.

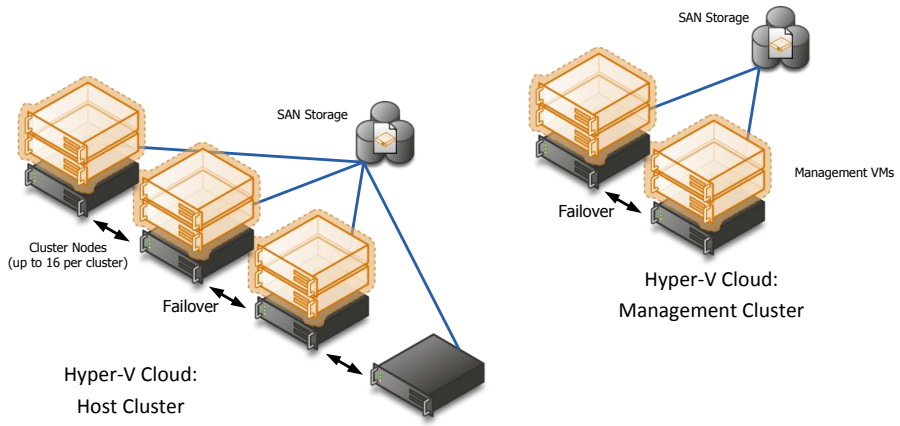

Figure 8. Example of the Topology of a Hyper-V Cloud

The Hyper-V Cloud Fast Track configurations are made up of the following elements:

- **Management network.** A dedicated management network is required so hosts can be managed via a dedicated network, avoiding competition with guest traffic needs. A dedicated network provides a degree of separation for security and ease of management purposes. This typically implies dedicating one network adapter per host and one port per network device to the management network. This network is used for remote administration of the host, communication to management systems (System Center agents), and other administrative tasks.

  Additionally, most server manufacturers provide a separate out-of-band management capability that enables remote management of server hardware outside of the host operating system. For more information, see Hyper-V: Live Migration Network Configuration Guide in the Microsoft TechNet Library.

- **iSCSI network.** If using iSCSI, a dedicated iSCSI network is required so that storage traffic is not in contention with any other traffic. This typically implies dedicating two network adapters per host and two ports per network device to the storage network. For all iSCSI storage connections, an MPIO configuration with two independent physical ports is required.

- **CSV/cluster communication network.** Usually, when the cluster node that owns a virtual hard disk (VHD) file in CSV performs disk I/O, the node communicates directly with the storage devices, for example through a SAN. However, storage connectivity failures sometimes prevent

a given node from communicating directly with the storage device. To maintain functionality until the failure is corrected, the node redirects the disk I/O through a cluster network (the preferred network for CSV) to the node where the disk is currently mounted. This is called CSV redirected I/O mode. For all CSV network connections, a teamed configuration is required.

- **Live migration network.** During live migration, the contents of the memory of the VM running on the source node needs to be transferred to the destination node over a LAN connection. To ensure high-speed transfer, a dedicated, redundant, or teamed 1-Gbps (or better) live migration network is required. For best performance, utilize a dedicated or shared 10-gigabit Ethernet connection for the live migration network. This significantly reduces the time required to evacuate the VMs off a host with zero downtime during maintenance or Windows updates.
- **VM networks.** The VM network or networks are dedicated to VM LAN traffic. The VM network can be based on two or more 1-GB Ethernet networks, one or more networks created via NIC teaming, or virtual networks created from shared 10-gigabit Ethernet NICs. Implement one or more dedicated VM networks. If using 1-GB Ethernet NICs, ensure that all Hyper-V hosts have two or more dedicated network adapters connected to the VM network for exclusive use by the guest VMs. If using 10-gigabit NICs, ensure that a teamed, virtual NIC is presented to the guest VMs to ensure redundancy.

Table 3 shows the network VLAN configuration that the Dell validated solution recommends.

Table 3. Recommended Dell Network VLAN Configuration

Description	VLAN	IP Range /subnet	Routable?	Tagging Method
Host Management	10	172.10.0.0/16	yes	Tagged at host
Cluster Private / CSV	20	172.20.0.0/16	no	Tagged at host
Live Migration	30	172.30.0.0/16	no	Tagged at host
iSCSI Network	5	172.5.0.0/16	yes	Untagged, primary VLAN set on switch
Virtual Machine	10,11	172.10.0.0/16 172.11.0.0/16	yes	Tagged at VM

2. Storage Topology

Cluster Shared Volumes (CSV) is a feature that simplifies the configuration and management of Hyper-V VMs in failover clusters. With CSV on a failover cluster that runs Hyper-V, multiple VMs can use the same LUN (disk) yet fail over (or move from node to node) independently of one another. CSV provides increased flexibility for volumes in clustered storage—for example; it allows IT managers to keep system files separate from data to optimize disk performance, even if the system files and the data are contained within virtual hard disk (VHD) files. If you choose to use live migration for your clustered VMs, CSV can also provide performance improvements for the live migration process. CSV is available in Windows Server 2008 R2 Enterprise and Datacenter Editions and Microsoft Hyper-V Server 2008 R2.

Figure 9 illustrates a design where a Hyper-V host cluster is utilizing three CSV LUNs to store different data types. One stores the guest VM operating system partitions, another stores the guest VM data partition, and the third stores guest VM application data such as database files.

CSV must be enabled and able to be utilized for storing multiple VMs on a single LUN.

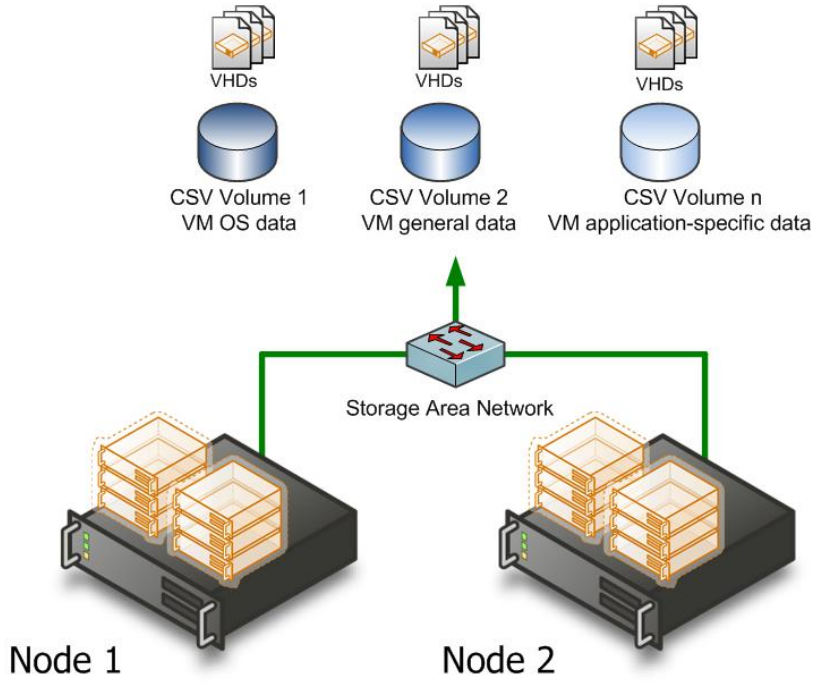

Figure 9. Example of a Common CSV Design for a Large Hyper-V cluster

4.6.3 Hyper-V VM Design

Standardization is a key tenet of private cloud architectures that helps to drive performance predictability and optimize IT usage. This also applies to VMs. A standardized collection of VM templates can both drive predictable performance and greatly improve capacity planning capabilities.

To make the most efficient use of available resources, use documented, standardized VM configurations for all VMs, both management and tenants. Table 4 shows an example for Windows Server 2003 and Windows Server 2008 R2 guests.

Table 4. Sample VM Configuration

Template	Specs	Network	Operating System	Unit Cost
Template 1 – Small	1 vCPU, 2 GB memory, 50-GB disk	VLAN x	Windows Server 2003 R2	1
Template 2 – Med	2 vCPU, 4 GB memory, 100-GB disk	VLAN x	Windows Server 2003 R2	2
Template 3 – Large	4 vCPU, 8 GB memory, 200-GB disk	VLAN x	Windows Server 2003 R2	4
Template 4 – Small	1 vCPU, 2 GB memory, 50-GB disk	VLAN x	Windows Server 2008 R2	1
Template 5 – Med	2 vCPU, 4 GB memory, 100-GB disk	VLAN x	Windows Server 2008 R2	2
Template 6 – Large	4 vCPU, 8 GB memory, 200-GB disk	VLAN x	Windows Server 2008 R2	4

Considerations when setting up the virtual machines include:

- VM storage
- VM networking
- Virtual processors

1. VM Storage

a. Dynamically Expanding Disks

Dynamically expanding VHDs provide storage capacity as needed to store data. The size of the VHD file is small when the disk is created and grows as data is added to the disk. The size of the VHD file does not shrink automatically when data is deleted from the virtual hard disk. However, you can compact the disk to decrease the file size after data is deleted by using the Edit Virtual Hard Disk Wizard in Windows Server 2008.

b. Fixed-Size Disks

Fixed virtual hard disks provide storage capacity by using a VHD file that is expanded to the size specified for the virtual hard disk when the disk is created. The size of the VHD file remains fixed regardless of the amount of data stored. However, you can use the Edit Virtual Hard Disk Wizard to increase the size of the virtual hard disk, which increases the size of the VHD file. By allocating the full capacity at the time of creation, fragmentation at the host level is typically not an issue (fragmentation inside the VHD itself must be managed within the guest).

c. Differencing Disks

Differencing virtual hard disks provide storage that enables you to preserve the content of a VHD and record all changes in a separate VHD file linked to it. In that configuration, the parent disk/parent VHD is read-only.

d. Pass-Through Disks
Hyper-V enables VM guests to directly access local disks or SAN LUNs that are attached to the physical server, without requiring the volume to be presented to the host server. The VM guest accesses the disk directly (utilizing the disk's globally unique identifier or GUID) without having to utilize the host's file system. Given that the performance difference between fixed-disk and pass-through disks is now negligible, the decision should be mainly based on manageability. For instance, if the data on the volume will be very large (hundreds of gigabytes), a VHD is hardly portable at that size given the extreme amount of time it takes to copy. Hence, a pass-through disk would be a reasonable alternative. Also bear in mind the backup scheme. With pass-through disks, the data can only be backed up from within the guest. When utilizing pass-through disks, there is no VHD file created; the LUN is used directly by the guest.

e. In-Guest iSCSI Initiator
Hyper-V can also utilize iSCSI storage by directly connecting to iSCSI LUNs via the guest's virtual network adapters. This is mainly used for access to large volumes, for access to volumes on SANs that the Hyper-V host itself is not connected to, or for guest clustering.

Utilize fixed disks for production environments, which provide better performance than the other options and ease the monitoring of storage availability. Utilizing fixed disks allocates the full size of the disk upon creation.

Dynamically expanding virtual hard disks are also a viable option for production use. However, they carry other risks such as storage oversubscription and fragmentation, so use these with caution.

Differencing disks are not recommended for production server workloads, because they introduce additional challenges with respect to storage management and storage subsystem efficiency. In cases where maximum size and extreme performance are required and portability (in the VHD sense) is not a requirement, use pass-through disks.

For in-guest iSCSI, ensure that a separate virtual network is utilized for access to the iSCSI storage to obtain acceptable performance. If the VM iSCSI network is shared with Ethernet traffic, utilize quality of service (QoS) to provide performance guarantees to the different networks. Consider using jumbo frames within the guest to improve iSCSI performance.

2. VM Networking
Hyper-V guests support two types of virtual network adapters: synthetic and emulated. The faster performing of the two, synthetic, makes use of the Hyper-V VMBus architecture and is the high-performance, native device in the VM. Synthetic devices require that the Hyper-V integration components be installed within the guest. Emulated adapters are available to all guests even if integration components are not available.

Always use synthetic virtual network adapters when possible. Because there are integration services for all supported Hyper-V guest operating systems, the primary reason to use the emulated network adapter is for pre-boot execution environment (PXE) booting.

You can create many virtual networks on the server running Hyper-V to provide a variety of communications channels. For example, you can create networks to provide the following:

- Communications between VMs only. This type of virtual network is called a private network.

- Communications between the host server and VMs. This type of virtual network is called an internal network.
- Communications between a VM and a physical network by creating an association to a physical network adapter on the host server. This type of virtual network is called an external network.

For the private cloud scenario, use one or more external networks per VM, and segregate the networks with VLANs and other network security infrastructure as needed.

3. Virtual Processors

A logical processor (LP) is defined as a processing core seen by the host operating system or parent partition. For example, in the case of Intel Hyper-Threading Technology, each thread is considered an LP. Hyper-V supports a maximum ratio of 8 virtual processors (VPs) per logical processor for server workloads, and a maximum ratio of 12 virtual processors for VDI workloads.

Therefore, looking at a server workload scenario, a 16-LP server supports a maximum of 128 VPs. That would in turn equate to 128 single-processor VMs, 64 dual-processor VMs, or 32 quad-processor VMs. The VP/LP ratio is a maximum supported limit. Different workloads will require different ratios.

For a table showing the supported number of VPs in a Hyper-V guest, see About Virtual Machines and Guest Operating Systems in the Microsoft TechNet Library.

4.7 Management Architecture

4.7.1 Management Scenarios

1. Infrastructure Deployment

When deploying several to hundreds of hosts, choosing the deployment model is a critical decision. Choices range from manual installation, which is highly inefficient, to varying degrees of automation, up to enterprise-class management systems. To achieve the architecture principle of predictability, all infrastructure components should be able to be deployed and configured in a repeatable and automated fashion. Examples include configuring network infrastructure, storage infrastructure, provisioning servers, and creating clusters.

The key components for successful deployment are the Microsoft Deployment Toolkit 2010 (MDT) and Windows Deployment Services (WDS). These are complemented by standard Windows Server roles such as Active Directory Domain Services, Domain Name System (DNS), and Dynamic Host Configuration Protocol (DHCP).

Using these technologies, it is possible to provide a robust deployment infrastructure using standard in-box solutions and toolkits.

The process and automation detailed in this document are based on the following choices:

- Windows Server 2008 R2 Datacenter Edition is installed with the core installation option for all Hyper-V hosts.
- Windows PowerShell 2.0, part of the Windows Management Framework, is available on all hosts.
- Windows PowerShell remoting through Windows PowerShell Remote is enabled on all hosts.
- All hosts are domain-joined.
- All hosts can be joined to up to16-node host clusters using CSVs.

- DHCP and DHCP reservations are to be used for all IP addressing.
- Microsoft Deployment Toolkit 2010 is to be used to create master images.
- Windows Deployment Services (WDS) is to be used to deploy all hosts.
- A network infrastructure supporting PXE boot and infrastructure deployment is available.
- A hardware management system is available that can remotely power off, power on, and reboot servers.
- The media access control (MAC) address of every NIC in every server is documented and associated to the physical network it is connected to. (That is, if iSCSI is used, the MAC address of each NIC connecting to the iSCSI network and the server it belongs to is documented.)

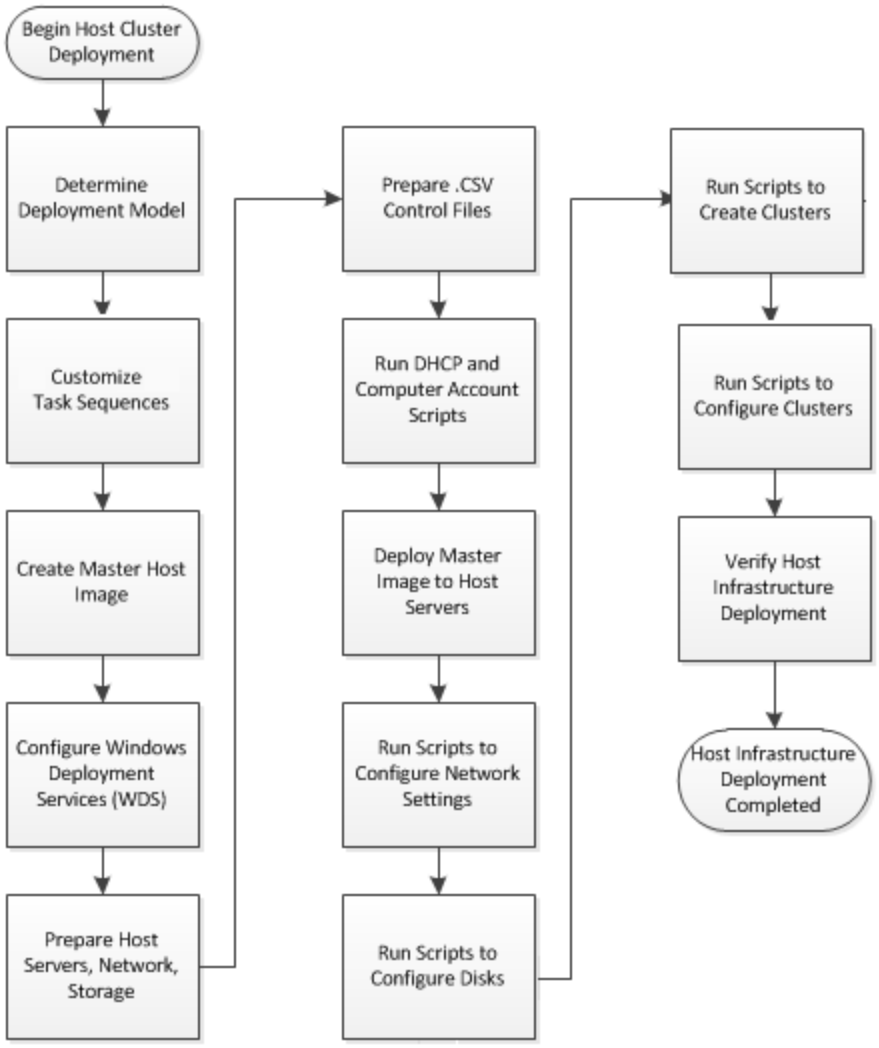

Figure 10. Host Cluster Deployment Process

2. VM Provisioning and Deprovisioning

One of the primary cloud attributes is user self-service, or providing the consumers of a service with the ability to request that service and have it be automatically provisioned for them. In the Hyper-V cloud solution, this refers to the ability of users to request one or more VMs or to delete one or more of their existing VMs. The infrastructure scenario supporting this capability is the VM provisioning and

deprovisioning process. This process is initiated from the self-service portal or tenant user interface and triggers an automated process or workflow in the infrastructure through System Center Virtual Machine Manager to either create or delete a VM based on the authorized settings input by the user or tenant. Provisioning could be template-based, such as requesting a small, medium, or large VM template, or it could be a series of selections made by the user (such as vCPUs and RAM). If authorized, the provisioning process creates a new VM per-user request, adds the VM to any relevant management products in the Hyper-V cloud (such as Microsoft System Center), and enables access to the VM by the requestor.

3. Infrastructure Monitoring

The Hyper-V cloud creates the ability to monitor every major component of the solution and generate alerts based on performance, capacity, and availability metrics. Examples include monitoring server availability, CPU, and storage utilization, which are achieved by taking advantage of integration through management packs and through Microsoft System Center Operations Manager.

4. Infrastructure Maintenance

The Hyper-V cloud makes it possible to perform maintenance on any component of the solution without affecting the availability of the solution or infrastructure to the consumer. Examples include the need to update or patch a host server, and adding additional storage to the SAN. For instance, during maintenance mode of a host System Center Virtual Machine Manager, the Hyper-V cloud ensures that VMs are evacuated via live migration, and no new workloads are placed on that host until maintenance is complete and the host is returned to normal operation.

5. Resource Optimization

Elasticity, perception of infinite capacity, and perception of continuous availability are Hyper-V cloud architecture principles that relate to resource optimization. Resources are optimized by dynamically moving workloads around the infrastructure based on performance, capacity, and availability metrics. One example is System Center Operations Manager and System Center Virtual Machine Manager working in coordination to distribute workloads across the infrastructure for maximum performance. Another example is System Center Operations Manager and System Center Virtual Machine Manager working in coordination to consolidate as many workloads as possible to the smallest number of hosts for a higher consolidation ratio.

6. Backup and Disaster Recovery

The Hyper-V cloud solution provides a means of data backup and disaster recovery for both the VMs and the host infrastructure. For example, Microsoft System Center Data Protection Manager 2010 delivers unified data protection for Windows server-based programs such as Microsoft SQL Server, Microsoft Exchange Server, Microsoft SharePoint Server, as well as virtualization servers and file servers.

7. Reporting

The Hyper-V cloud solution offers a centralized reporting capability that provides standard reports detailing capacity, utilization, and other system metrics. The reporting functionality serves as the foundation for capacity-based or utilization-based billing and chargeback to tenants. System Center Virtual Machine Manager Self Service Portal 2.0 provides comprehensive reporting capabilities.

4.7.2 Automation

The automation layer is made up of the foundational automation technology plus a series of single-purpose commands and scripts that perform operations such as starting or stopping a VM, rebooting a server, or applying a software update. These atomic units of automation are combined and executed by

higher-level management systems. The modularity of this layered approach dramatically simplifies development, debugging, and maintenance.

The Windows Management Framework is a core set of infrastructure technologies that combine to provide advanced local and remote automation and system management (see Figure 11). Key underpinning technologies such as Windows Management Instrumentation (WMI), Web Services for Management (WS-Management), and Background Intelligent Transfer Service (BITS) are used by higher-level layers such as Windows PowerShell, which itself is utilized by higher-level layers such as scripts, user interfaces, and suites such as Microsoft System Center.

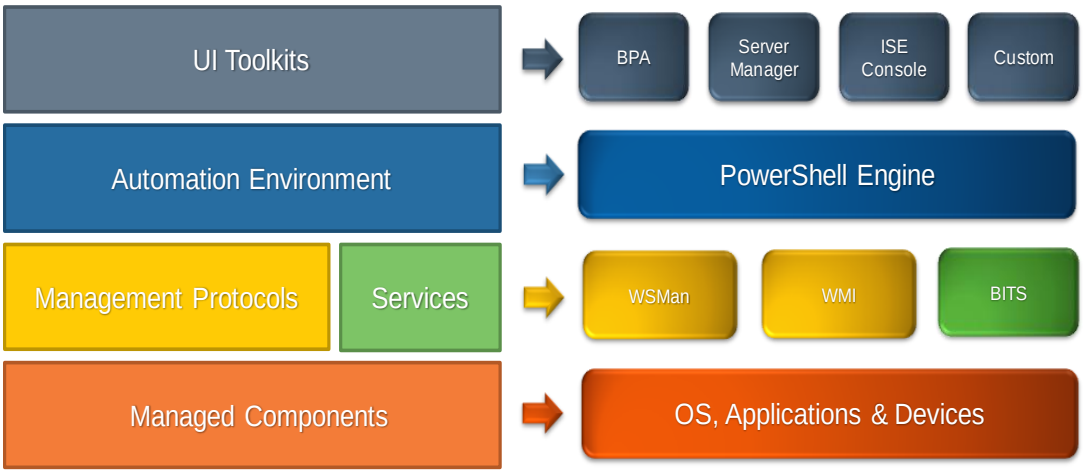

Figure 11. Windows Management Framework

The Windows Management Framework Core package provides updated management functionality and includes Windows PowerShell 2.0 and Windows Remote Management (WinRM) 2.0.

Windows Management Instrumentation (WMI) is the Microsoft implementation of Web-based Enterprise Management (WBEM), which is an industry initiative to develop a standard technology for accessing management information in an enterprise environment (see Figure 12). WMI uses the Common Information Model (CIM) industry standard to represent systems, applications, networks, devices, and other managed components. CIM is developed and maintained by the Distributed Management Task Force (DMTF).

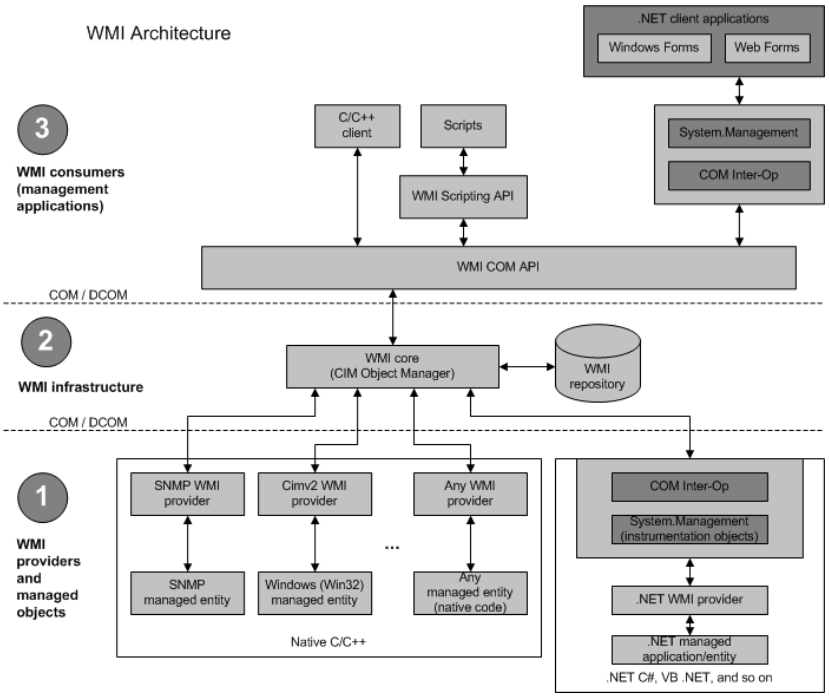

Figure 12. WMI Architecture

The WS-Management protocol specification provides a common way for systems to access and exchange management information across an IT infrastructure. WinRM and Intelligent Platform Management Interface (IPMI), along with the Event Collector, are components of Hardware Management in Windows Server. For more information, see Hardware Management Introduction in the Microsoft TechNet Library.

4.7.3 Private Cloud Management

The management layer is made up of tools that are utilized to deploy and operate the infrastructure. In most cases, this consists of a variety of toolsets for managing hardware, software, and applications that make use of the underpinnings described in Sections 4.7.1 through 4.7.3 of this white paper. The Microsoft System Center family includes a comprehensive set of capabilities for management of the cloud fabric. This section describes how individual products fulfill their roles within the private cloud management stack.

1. SQL Server 2008 SP1

Microsoft System Center components are database-driven applications. This makes for a highly available and well-performing database platform that is critical to the overall management of the environment.

SQL Server Configuration

- 1 non-high-availability (non-HA) VMs on different Hyper-V hosts (clustered from within the VM-guest cluster)
- Windows Server 2008 R2 Enterprise Edition
- 4 vCPUs
- 8 GB of memory
- 3 vNICs (1 client connections, 1 cluster communications, 1 iSCSI)
- Storage: 1 operating system VHD, 3 x dedicated iSCSI LUNs

Table 5. SQL Server Data Locations

LUN	Purpose	Size
LUN 1, CSV Volume	VM operating system	30-GB VHD
LUN 2, iSCSI	SQL Server databases	varies
LUN 3, iSCSI	SQL Server logging	varies
LUN 4, iSCSI	SQL Server cluster quorum	1 GB

Table 6. Databases

DB Client	Instance Name	DB name	Authentication
VMM SSP	<Instance 1>	<SCVMMSSP>	Win Auth
WSUS	<Instance 1>	<WSUS_DB>	Win Auth
Ops Mgr	<Instance 1>	<Ops Mgr_DB>	Win Auth
Ops Mgr	<Instance 2>	<Ops Mgr_DW_DB>	Win Auth
VMM	<Instance 1>	<VMM_DB>	Win Auth

2. Microsoft System Center Virtual Machine Manager 2008 R2

System Center Virtual Machine Manager 2008 R2 helps to enable the centralized management of physical and virtual IT infrastructure, increased server utilization, and dynamic resource optimization across multiple virtualization platforms. It includes end-to-end capabilities such as planning, deploying, managing, and optimizing the virtual infrastructure.

Scope

System Center Virtual Machine Manager is used to manage only Hyper-V Cloud Fast Track hosts and guests in a single data center. No virtualization infrastructure outside of the solution should be managed by System Center Virtual Machine Manager. The System Center Virtual Machine Manager configuration is designed only within the scope of this architecture. Therefore, for managing elements outside this scope, other factors might lead to a different approach or design.

Servers

- 1 HA VM
- Windows Server 2008 R2
- 2 vCPUs
- 4 GB of memory
- 1 vNIC
- Storage: 1 operating system VHD, 1 x data VHD or pass-through volume

Roles

The following roles are required by System Center Virtual Machine Manager:

- SCVMM Server
- Administrator Console
- Command Shell
- SCVMM Library
- SQL Server Database (remote)

Any role that is not listed here will not be installed.

Operations Manager Integration

In addition to the built-in roles, System Center Virtual Machine Manager is configured to integrate with System Center Operations Manager. System Center Virtual Machine Manager uses System Center Operations Manager 2007 to monitor the health and availability of the VMs and VM hosts that System Center Virtual Machine Manager is managing. System Center Virtual Machine Manager also uses System Center Operations Manager to monitor the health and availability of the System Center Virtual Machine Manager server, database server, library servers, and self-service web servers, and to provide diagram views of the virtualized environment in the System Center Virtual Machine Manager Administrator Console. Integration with System Center Operations Manager is also a prerequisite for enabling Performance and Resource Optimization (PRO), a feature of System Center Virtual Machine Manager, and for configuring reporting in System Center Virtual Machine Manager.

VMM Library Placement

Libraries are the repository for VM templates, VHDs, and ISOs, and therefore serve a very important role. The library share itself resides on the System Center Virtual Machine Manager server. However,

the share is placed on its own logical partition and corresponding VHD or pass-through disk whose underlying disk subsystem has sufficient performance to service the provisioning demands.

Performance and Resource Optimization (PRO)
PRO is a feature of System Center Virtual Machine Manager 2008 R2 that enables dynamic management of virtualized infrastructure. The host-level PRO actions in the System Center Virtual Machine Manager 2008 Management Pack recommend migrating the VM with the highest resource usage on the host whenever the CPU or memory usage on the host exceeds the threshold defined by a PRO monitor. The VM is migrated to another host in the host group or host cluster that is running the same virtualization software. If an IT organization has a workload that is not suitable for migration when running in a VM, it can exclude that VM from host-level PRO actions. The VM will not be migrated even if it has the highest usage of the elevated resource on the host.

Dell PRO-enabled Management Packs
PRO provides an open and extensible framework for the creation of management packs for virtualized applications or associated hardware. In building these management packs enabled for PRO, partners can create a customized solution combining their offerings with the comprehensive monitoring and issue-resolution capabilities of PRO.

Partners incorporate deep product and process awareness of their solutions in their PRO-enabled management packs, as realized in rules and policies that PRO will act on in the event of poor performance or a pending failure. With these predetermined watch points and resolution steps, PRO can react dynamically to adverse situations.

The Dell Server PRO Management Pack v2.0 is a Dell-created PRO pack that integrates with Dell OpenManage to monitor events such as loss of power supply redundancy, temperature threshold, server storage battery error, and internal disk failures. The PRO Tips generates support actions such as the live migration of all virtual machines off of the alerting host.

For more information, see both Dell Server PRO Management Pack 2.0 for Microsoft System Center Virtual Machine Manager and the User's Guide for the Dell Server PRO Management Pack, available at support.us.dell.com.

3. Microsoft System Center Operations Manager 2007 R2

System Center Operations Manager agents are deployed to the fabric management hosts and VMs and to scale unit hosts and VMs. These in-guest agents are used to provide performance and health status of the operating system only. The scope of the System Center Operations Manager instance is for Hyper-V cloud infrastructure monitoring only. Application-level monitoring is out of scope for this System Center Operations Manager instance to keep the architecture workload neutral.

Servers

- 1 HA VM
- Windows Server 2008 R2
- 2 vCPUs
- 4 GB of memory
- 1 vNIC
- Storage: 1 operating system VHD

Roles

The following roles are required by System Center Operations Manager:

- Root Management Server
- Reporting Server (data will reside in SQL Server)
- Data Warehouse (data will reside in SQL Server)
- Operator Console
- Command Shell

Any role that is not listed here will not be installed.

Management Packs

The following System Center Operations Manager Management Packs are highly recommended:

- System Center Virtual Machine Manager 2008 Management Pack for Operations Manager 2007
- Windows Server Operating System Management Pack for Operations Manager 2007
- Windows Server Failover Clustering Management Pack for Operations Manager 2007
- Windows Server Hyper-V Management Pack for Operations Manager 2007
- SQL Server Monitoring Management Pack
- Windows Server Internet Information Services 7 Management Pack for Operations Manager 2007

To provide deeper integration of the hardware layer into the management stack, two additional recommended management packs for this described configuration are:

- *Dell EqualLogic Monitoring*: The Dell EqualLogic Storage Management Pack Suite v4.0 enables System Center Operations Manager (2007 R2/SP1) to discover, monitor, and accurately depict the status of Dell EqualLogic PS arrays on a defined network segment. For the Dell EqualLogic Storage Management Pack Suite v4.0, visit support.us.dell.com.
- *Dell Server Monitoring*: The Dell Server Management Pack Suite is also a required monitor and alert on Dell PowerEdge servers, Dell Remote Access Cards (DRAC), Chassis Management Controllers (CMC) and more. To obtain the Dell Server Management Pack Suite, along with its documentation, visit support.us.dell.com.

To find additional management packs for specific System Center or OEM products, see the Management Pack Catalog in Microsoft Pinpoint.

Reporting
Integrating System Center Virtual Machine Manager and System Center Operations Manager provides the capability to generate the following reports:

- *Virtualization Candidates*. Identifies physical computers that are good candidates for conversion to VMs. Uses server performance metrics available in System Center Operations Manager.
- *VM Utilization*. Reports resource utilization by your VMs. Reports underutilized or overutilized VMs.
- *VM Allocation*. Calculates chargeback to cost centers. Reports CPU, memory, disk, and network usage.
- *Host Utilization*. Reports the number of VMs running on each host. Reports average usage and total or maximum values for processors, memory, and disk space.
- *Host Utilization Growth*. Reports the percentage of change in resource usage and number of running VMs. System Center Operations Manager must be on a dedicated VM and use the remote SQL Server instance. In addition, System Center Operations Manager must be integrated with System Center Virtual Machine Manager.

4. Maintenance and Patch Management
In a cloud infrastructure, it is important to ensure a high level of consistency across all VMs as well as hosts. By using Windows Server Update Services (WSUS), administrators fully manage and automate the distribution of updates to the host and VMs in a consistent manner.

This ability is augmented by System Center Configuration Manager 2007 R2, which comprehensively assesses and deploys servers across physical and virtual environments. In addition, System Center Configuration Manager makes it possible to assess variation from desired configurations, take hardware and software inventory, and remotely administer computers.

Patch management of offline VMs and VM libraries within a cloud is another important aspect that needs to be addressed. Microsoft Virtual Machine Servicing Tool addresses this need by providing the capability to update offline virtual hard disks as well as templates with the latest operating system and application patches, without introducing vulnerabilities into the infrastructure.

5. Backup and Disaster Recovery

In a virtualized data center, there are three commonly used backup types: host-based, guest-based, and SAN-based. Table 7 contrasts these backup types.

Table 7. Comparison of Common Data Center Backup Types

Capability	Host Based	Guest Based	SAN Snapshot
Protection of VM configuration	X		X*
Protection of host and cluster configuration	X		X*
Protection of virtualization-specific data such as VM snapshots	X		X
Protection of data inside the VM	X	X	X
Protection of data inside the VM stored on pass-through disks		X	X
Support for VSS-based backups for supported operating systems and applications	X	X	X*
Support for Continuous Data Protection	X	X	
Ability to granularly recover specific files or applications inside the VM		X	

***Depends on storage vendor's level of Hyper-V Integration**

Data security and availability are paramount. Accordingly, a Hyper-V Cloud Fast Track implementation is designed to provide:

- The capability to support the Hyper-V Volume Shadow Copy Service (VSS) writer for host-side backup.
- Backup storage separate from the SAN (SAN snap technology can be used in conjunction with a mechanism to move the backup off the production SAN).
- The capability to restore individual files from the VM backup.
- Application awareness in backup solutions.

System Center Data Protection Manager 2010

System Center Data Protection Manager 2010 is a comprehensive backup solution that can be used in the Hyper-V Cloud Fast Track implementations and provides continuous data protection for VMs hosted on servers running Hyper-V. This protection includes online backup of supported guest VMs hosted on clustered or standalone systems, protection of VMs during the live migration process, and item-level recovery from host-level backup. In addition, System Center Data Protection Manager 2010 offers disk-to-disk, disk-to-tape, and disk-to-disk-to-tape technologies, all of which maintain the business value of a virtualized infrastructure by ensuring that the infrastructure is well protected and always available.

Dell EqualLogic Auto-Snapshot Manager/Microsoft Edition (ASM/ME)
Dell provides a VSS-aware quick recovery application called Auto-Snapshot Manager/Microsoft Edition (ASM/ME) that is included with the EqualLogic Host Integration Tools for Windows. ASM/ME is a VSS requestor application that enables the creation of data- and application-consistent Smart Copies (point-in-time copies) of NTFS volumes, Exchange databases, Hyper-V VMs, and SQL Server databases utilizing the built-in hardware volume snapshot capabilities of PS Series arrays. For more information, see Microsoft Hyper-V Virtual Machine Protection Using EqualLogic Auto-Snapshot Manager / Microsoft Edition, available at the EqualLogic Resource Center.

6. Tenant/User Self-Service Portal
Microsoft System Center Virtual Machine Manager Self-Service Portal 2.0 is the Microsoft self-service solution for Hyper-V cloud. The tenant/user self-service layer provides an interface for Hyper-V cloud tenants or authorized users to request, manage, and access the services, such as VMs, provided by the Hyper-V cloud architecture. Using role-based access control and authorization, the self-service layer provides the ability to delegate certain aspects of administration (such as starting/stopping VMs) to designated "tenant administrators." In addition, the self-service portal integrates and makes use of the underlying technologies discussed in the previous sections.

The self-service portal has three components:

- **Website.** The website is a web-based component that provides a user interface to the self-service portal. Through the website, users can perform various tasks such as pooling infrastructure assets in the self-service portal, extending VM actions, creating business unit and infrastructure requests, validating and approving requests, and provisioning VMs (using the self-service VM provisioning feature). Users can also use the website to view information related to these tasks.
- **Database.** The database is a SQL Server database that stores information about configured assets, data related to business units and requests, and information about what has been provisioned to various business units.
- **Server.** The server is a Windows service that runs default and customized VM actions that the user requests through the website.

System Center Virtual Machine Manager Self-Service Portal 2.0 also provides self-service access that enables on-demand provisioning, deprovisioning, reporting, monitoring, and metering/billing. For more information, see Section 4.7.5 Security, later in this white paper.

7. Dell Storage Management
Dell provides tools for the management of its storage solutions. Although initial setup is done via a serial console, there are other, more user-friendly tools like web portals to use as well. All tools provide an interface to provision and deprovision storage LUNS, manage security, monitor hardware health, create and delete snapshots, and more. For more information, see the Dell Storage Solutions Guide for Microsoft Hyper-V.

8. Dell PowerConnect Network Management
The Dell PowerConnect Series switches are managed through a serial console, Telnet, Secure Shell (SSH), or a web portal. For configuration documentation, visit the Network Switch Manuals section of the Dell support website.

9. Dell Server Management Utilities

Dell PowerEdge Server Out-of-Band Management

The Integrated Dell Remote Access Controller 6 (iDRAC6) is a systems management hardware and software solution that provides remote management capabilities, crashed system recovery, and power control functions for Dell PowerEdge systems. The iDRAC6 uses an integrated system-on-chip microprocessor to remotely monitor and control systems. The iDRAC6 is also manageable through the Windows Remote Management framework, allowing close integration with Microsoft System Center components.

For additional information about the iDRAC6, see Integrated Dell Remote Access Controller 6 (iDRAC6) Version 1.5 User Guide. For information about Windows Remote Management integration, see Remote Management Using Microsoft Windows PowerShell and the Dell Lifecycle Controller.

Dell OpenManage Server Administrator

Dell OpenManage Server Administrator (OMSA), a security-rich web tool, can help customers manage individual servers and their internal storage from virtually anywhere at any time. OMSA is available with a Dell server at no additional charge. Dell OMSA:

- Helps to simplify single-server monitoring, with a secure command-line or web-based management graphical user interface
- Provides views of system configuration, health, and performance
- Provides online diagnostics to help isolate problems, or shut down and restart the server

For more information about Dell OMSA, visit www.dell.com/openmanage

10. Storage, Network, and Server Management

The underlying hardware of the Hyper-V Cloud Fast Track configurations builds the backbone of the infrastructure. Therefore, it is important to make sure that the management of the hardware layer is tightly integrated into the higher-level management layers. There are management packs for storage, network, and server that bridge the virtual/physical divide and allow System Center Operations Manager and System Center Virtual Machine Manager to understand the health of the hardware layer and implement automated actions accordingly.

11. Server Out-of-Band Management Configuration

The out-of-band management feature in System Center Configuration Manager 2007 SP2 allows administrators to connect to a computer's management controller when the computer is turned off or otherwise unresponsive through the operating system. Out-of-band management supplements in-band management.

4.7.4 Orchestration

The orchestration layer provides the ability to design, test, implement, and monitor IT workflows within the cloud environment. System Center Opalis is an automation platform that makes it possible to automate best practices such as those found in Microsoft Operations Framework (MOF) and Information Technology Infrastructure Library (ITIL) in the cloud environment.

Through its workflow designer, System Center Opalis automatically shares data and initiates tasks in System Center Operations Manager, System Center Configuration Manager, System Center Service Manager, System Center Virtual Machine Manager, Active Directory, and third-party tools. This ensures

repeatable, consistent results by removing the latency associated with manual coordination service delivery.

System Center Opalis fosters integration, efficiency, and business alignment of data center IT services by:

- Reducing unanticipated errors and service delivery time by automating tasks across vendor and organization silos.
- Integrating System Center with non-Microsoft tools to enable interoperability across the data center.
- Orchestrating tasks across systems for consistent, documented, compliant activity.

4.7.5 Security

The three pillars of IT security are confidentiality, integrity, and availability (CIA). These tenets apply directly to a private cloud as well. The Hyper-V Cloud Fast Track Program was created from the ground up to address security threats. Generally, threat modeling assumes the following conditions:

- Organizations have resources (in this case, IT components) that they want to protect.
- All resources are likely to exhibit some vulnerabilities.
- People might exploit these vulnerabilities to cause damage or gain unauthorized access to information.
- Properly applied security counter-measures help to mitigate threats that exist because of vulnerabilities.

Threat modeling is addressed through Microsoft Operations Framework (MOF) 4.0, a framework that provides practical guidance for managing IT practices and activities throughout the entire IT life cycle.

Security for the Hyper-V cloud is founded on two paradigms: protected infrastructure and network access. These paradigms are discussed in the next two sections.

1. Protected Infrastructure

A defense-in-depth strategy is utilized at each layer of the Hyper-V cloud architecture. Security technologies and controls are implemented in a coordinated fashion.

An entry point represents data or process flow that traverses a trust boundary. Any portions of an IT infrastructure in which data or processes traverse from a less-trusted zone into a more-trusted zone should have a higher review priority.

Users, processes, and IT components all operate at specific trust levels that vary between fully trusted and fully distrusted. Typically, parity exists between the level of trust assigned to a user, process, or IT component and the level of trust associated with the zone in which the user, process, or component resides.

A defense–in-depth strategy, with overlapping layers of security, is the best way to counter threats, and the least-privileged user account (LUA) approach is an important part of that strategy. The Hyper-V cloud uses the LUA approach, which helps ensure that users follow the principle of least privilege and always log on with limited user accounts. This strategy also aims to limit the use of administrative credentials to administrators, and then only for administrative tasks.

Several of the elements of defense-in-depth strategy are provided by the following systems:

- **Delegation, access control, and Data Management with Active Directory.** The Hyper-V cloud is based on a Windows environment that uses Active Directory to manage access control. Access control is the means by which administrators can control or delegate the ability of other users to manipulate objects in Active Directory and also to perform actions on domain controllers and file servers. Understanding the access control model in Active Directory is essential to being able to delegate administration.
- **Role-based security for System Center Virtual Machine Manager.** System Center Virtual Machine Manager implements role-based security to provide finer control over who can do what within the virtualized environment. This security model supports delegated administration. A user role defines a set of operations (grouped in a profile) that can be performed on a selected set of objects (defined by the user role's scope). Within that framework, an organization can create delegated administrator roles that allow, for example, a high-level administrator to manage all operations in a New York office, a specialized administrator to manage all library servers, or an advanced user to set up complex virtual environments within a single lab. An organization also can create self-service user roles that allow users to perform a specified set of operations on their own VMs.
- **System Center Virtual Machine Manager Self-Service Portal 2.0 access control.** The self-service portal utilizes a role-based access control authorization model based on standard ITIL/MOF best practices. This model lets users securely allow tenant access to the private cloud infrastructure resources. The self-service portal provides four default user roles: data center administrator (DCIT Admin), business unit administrator (BUIT Admin), advanced operator, and business unit user. Data center administrators can also create custom user roles with any combination of permissions for VM actions.

2. Network Access

Windows Firewall with Advanced Security combines a host firewall and IPSec. Network Access Protection (NAP) is a platform that allows network administrators to define specific levels of network access based on a client's identity, the groups to which the client belongs, and the degree to which the client complies with corporate governance policies. If a client is not compliant, NAP provides a mechanism for automatically bringing the client into compliance (a process known as remediation) and then dynamically increasing its level of network access.

- **End-point protection (antivirus and anti-malware).** The antivirus solution on the host is designed to support Hyper-V and can be configured according to Knowledge Base article 961804. The host-based firewall must be enabled and configured.
- **Microsoft Forefront.** Microsoft Forefront Client Security delivers end-to-end security and access to information through an integrated line of protection, access, and identity management products. The Forefront-based solutions help to deliver a contextual and user-centric security solution aligned to the needs of cloud customers:
  - *Multilayered protection*. Forefront delivers leading malware protection solutions across end points, messaging and collaboration application servers, and the network.
  - *Identity-based access*. Microsoft identity-based access technologies and Forefront solutions build upon the Active Directory infrastructure to enable policy-based user access to applications, devices, and information.

4.7.6 Service Management

The service management layer makes it possible to automate and adapt IT service management best practices, such as those found in MOF and ITIL, thereby providing built-in processes for incident resolution, problem resolution, and change control.

Conclusion

The Microsoft Hyper-V Cloud Fast Track program is a comprehensive approach to cloud computing developed jointly by Microsoft and its OEM partners. This private cloud solution combines Microsoft software, consolidated guidance, and validated configurations with Dell hardware technology—including computing power, network and storage architectures, and value-added software components. By tapping into the Fast Track program, organizations can rapidly deploy private clouds while reducing their risk and their total cost of ownership. The ultimate goal is to give organizations both the control and flexibility they need to harness the full power of the private cloud.