Intel® Cloud Builders Guide
Intel® Xeon® Processor Servers
Powerleader PR* Series Servers
Microsoft System Center Virtual Machine Manager Self-Service Portal* 2.0

Intel® Cloud Builders Guide: Cloud Design and Deployment on Intel® Platforms

Powerleader Power Rack Server*

Microsoft System Center Virtual Machine Manager Self-Service Portal* 2.0

Intel® Xeon® Processor 5500 Series

Intel® Xeon® Processor 5600 Series

AUDIENCE AND PURPOSE

For cloud service providers, hosters and enterprise IT organizations who are looking to build their own cloud infrastructure, the decision to use a cloud for the delivery of IT services is best done by starting with the knowledge and experience gained from previous work. This reference architecture, also called "this guide," outlines a private cloud setup using Windows Server, Hyper-V* and the Microsoft System Center Virtual Machine Manager Self-Service Portal* 2.0 (VMMSSP) on the Powerleader Power-Rack* (PR) Series Servers, powered by the Intel® Xeon® processor.

Using the contents of this guide, which includes detailed scripts and screen shots, should significantly reduce the learning curve for building and operating a new cloud computing infrastructure.

Because the creation and operation of a cloud requires integration and customization to existing IT infrastructure and business requirements, it is not expected that this guide can be used "as-is." For example, adaptation to an existing network and identification of management requirements are out of scope for this guide. Therefore, it is expected that the user of this guide will make significant adjustments to the design to meet specific customer requirements. This guide is assumed to be a starting point for that journey.

Table of Contents

Executive Summary

Enterprise IT departments can build private clouds today on the Microsoft* server platform, including Windows Server*, Hyper-V* and the System Center* suite of management software. Microsoft is committed to making it easier for customers to deploy private clouds in their data centers. Enabling private clouds is one component of Microsoft's commitment to deliver IT as a Service. The other components of IT as a Service are the Windows Azure Platform Appliance* and Windows Azure + SQL Azure* in Microsoft's data center.

Private clouds enable agility and efficiency as they provide customers with the benefits of an automated and optimized virtualized data center. The term "private cloud" refers to a style of computing where scalable and elastic IT-enabled capabilities are delivered as a utility-like service to internal customers. Private clouds share the core attributes of cloud computing:

- Scalable
- Elasticity
- Multi-tenancy
- Metered by use
- Self-service

A private cloud differs from a public cloud in that it represents a set of computing resources that are dedicated to one customer (not shared with other customers). Resources may be located on premises or externally hosted by a third-party service provider. If the dedicated resources are hosted, this is a special type of private cloud called a "hosted private cloud".

Private clouds change IT processes and make IT more agile and efficient. For example, instead of waiting two months to get new hardware for a new business unit IT application, in a private cloud environment the business unit requests and provisions new resources via a self-service mechanism and the IT department will be able to immediately service the business. This is possible because the IT staff has a private cloud already deployed, which makes use of a set of pooled resources.

VMMSSP* Overview

The Microsoft System Center Virtual Machine Manager Self-Service Portal* 2.0 (VMMSSP*[1]) is a free, partner-extensible portal that enables private cloud and IT as a Service with Windows Server*, Hyper-V* and System Center Virtual Machine Manager*. With the portal, customers and partners can dynamically pool, allocate, and manage resources to offer Infrastructure as a Service (IaaS). Key benefits of the portal include:

- Allocation of data center resources: The portal pools data center infrastructure resources, such as network, storage, load balancers, virtual machine templates, and domains, and makes them available to business units to meet their infrastructure needs. It also establishes the costs associated with reserving and using infrastructure resources.
- Simplification of business unit on-boarding: The portal simplifies the process of on-boarding a new business unit infrastructure request. It provides a way for data center administrators to register business unit requirements in one centralized location. Business unit administrators can request resources in the organization's infrastructure pool to host their IT services.
- Validation and provision of infrastructure: The portal also simplifies the process for data center administrators to validate and provision a business unit IT administrator's infrastructure requests. Using the portal, data center administrators can provision the requested resources and assign them to a requesting business unit IT administrator.
- Self-service provisioning: The portal provides an end-user self-service feature for virtual machine provisioning. It streamlines the business unit IT user's experience in managing virtual machines. And its extensibility scripts reduce the manual steps associated with provisioning virtual machine resources.
- Partner extensibility: The toolkit includes powerful extensibility features for Independent Hardware Vendors (IHV), Independent Software Vendors (ISV) and System Integrators (SI). Partners can customize different virtual machine actions, such as create, delete, stop, start, shutdown, connect, and pause, to make use of the unique characteristics of their infrastructure.

To provide these benefits, the portal includes the following components:

- Self-service portal to enable consumers of IT to request and provision infrastructure for their applications and services, and to on-board new business unit IT departments
- Tested documentation.
- Dynamic provisioning engine to rapidly provision virtualized infrastructure.
- Guidance to help partners easily extend functionality.

In addition to out-of-box functionality, VMMSSP* makes use of technologies like Windows PowerShell*, Microsoft's command shell for scripting and development, for extensibility purposes. The self-service portal provides an interface in which the default virtual machine actions can be extended; for example, scripts that interact with storage area networks (SANs) or load balancers that support virtual machines can be added.
The self-service portal uses XML to represent the virtual machine actions. Each action, such as "CreateVM", is composed of tasks. Each task, in turn, contains scripts and related parameters. Use the self-service portal to add, remove, or edit tasks.

An action XML segment stores one complete set of actions and tasks. The self-service portal provides a default action XML segment; this action XML segment can be cloned in order to extend the virtual machine actions. In this way, multiple custom sets of virtual machine actions can be created. The data center administrator can configure PowerShell scripts to include specific "cmdlets" to be executed in response to a triggered virtual machine action (for example, to create a virtual machine).

Usage Scenarios

The following are usage scenarios which create private clouds using VMMSSP*:

- Organizations that want to replace traditional IT operations and move to a flexible consumption-based model, but still require data and assets to reside on premises.
- Organizations that are entering a new business domain and want to run the supporting IT infrastructure on a separate, scalable, flexible model.
- Organizations that want to rebuild their IT infrastructure from scratch, implementing new policies, procedures, and business processes.

Enterprise Customer VMMSSP* Scenario

A common scenario is a business unit IT manager who needs an application to be built; a portal would capture the IT requirements for this new workload. These requirements may include compute, storage, and network assets, as well as specific virtual machine profiles that the IT manager wishes to use. The data center administrator (DC Admin) is then notified of the request, checks for the availability of resources in the data center, and assigns them to the requested infrastructure. Once ready, the request is approved and a self-service portal can be accessed, where full control of the custom-built environment for the workload can be taken. Virtual machines can be created, started, stopped, or deleted.

Data center administrators can access the infrastructure and virtual machine reports using a Microsoft System Center Virtual Machine Manager Self-Service Portal* 2.0 dashboard. To summarize:

1. The IT manager on-boards the application capturing requirements for a workload.
2. The data center IT Admin validates the request and assigns appropriate compute, network and storage resources.
3. The IT manager user accesses the portal, with full control of the environment.
4. The DCIT administrator accesses reports regarding workloads, costs, etc.

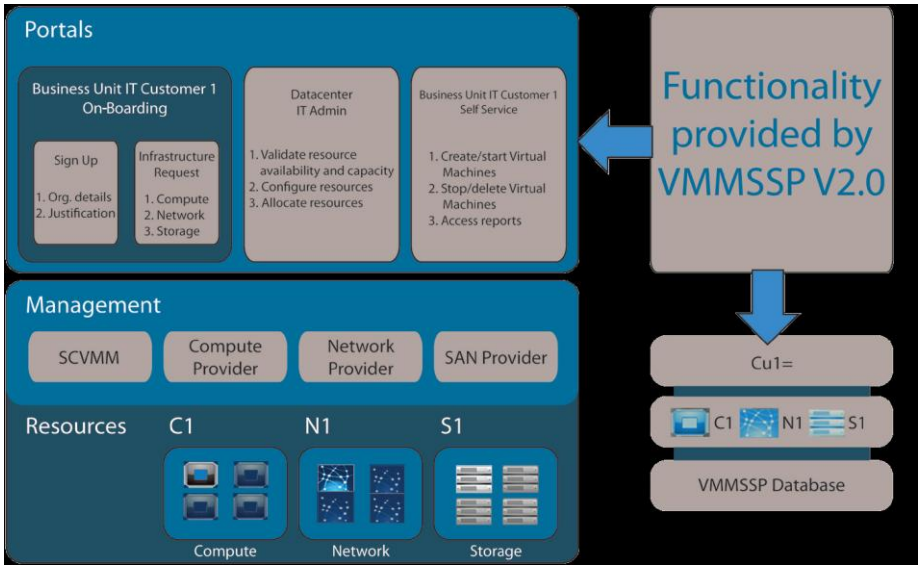

Figure 1 - VMMSSP* V2.0 Overview

Systems Integrator VMMSSP* Scenario

SIs can approach their accounts to help them assess their infrastructure readiness before deploying the VMMSSP*. For this purpose, VMMSSP provides SIs with extensive documentation to design and implement the portal within the data center. VMMSSP also provides guidance for extending virtual machine actions.

Test Bed Blueprint Overview

Test Bed Design Considerations

As the previous usage scenarios suggest, the deployment of a cloud-based infrastructure is often driven by the need to lower operational costs and/or rapidly respond to changes in compute resource demand. To meet those challenges, cloud applications must either directly make use of distributed capabilities in the cloud (using tools such as Apache Hadoop*[2]) or use the flexibility that virtualization provides within the cloud environment to optimize resource utilization and improve resiliency to failures.

The possibility of hardware failures requires that either the application or the cloud management software be able to perform the necessary recovery actions. The state of each compute element must be maintained in shared storage or the application must be responsible for retry/recovery on a failing device. For a typical non-cloud aware enterprise workload, the application is usually wrapped into a virtual machine, and that virtual machine and its data are stored on shared storage. In this way, if a server fails, the cloud management software can simply restart the virtual machine on another server.

The Intel® Cloud Builders Test Bed follows these principles. The diagram shown in Figure 2 represents a physical construction of the cloud environment used for the tests discussed in this reference architecture.

Intel® Xeon® processor 5500 and 5600 series[3] servers are connected using a 1 GbE network and 10 GbE network using "top of rack" switches. The switches were configured to achieve the logical architecture shown in Figure 2. In this configuration, there are no "special" servers and no "special" connections; all servers are configured identically. This uniformity allows for simple replacement or reassignment of workloads. Additional design considerations are:

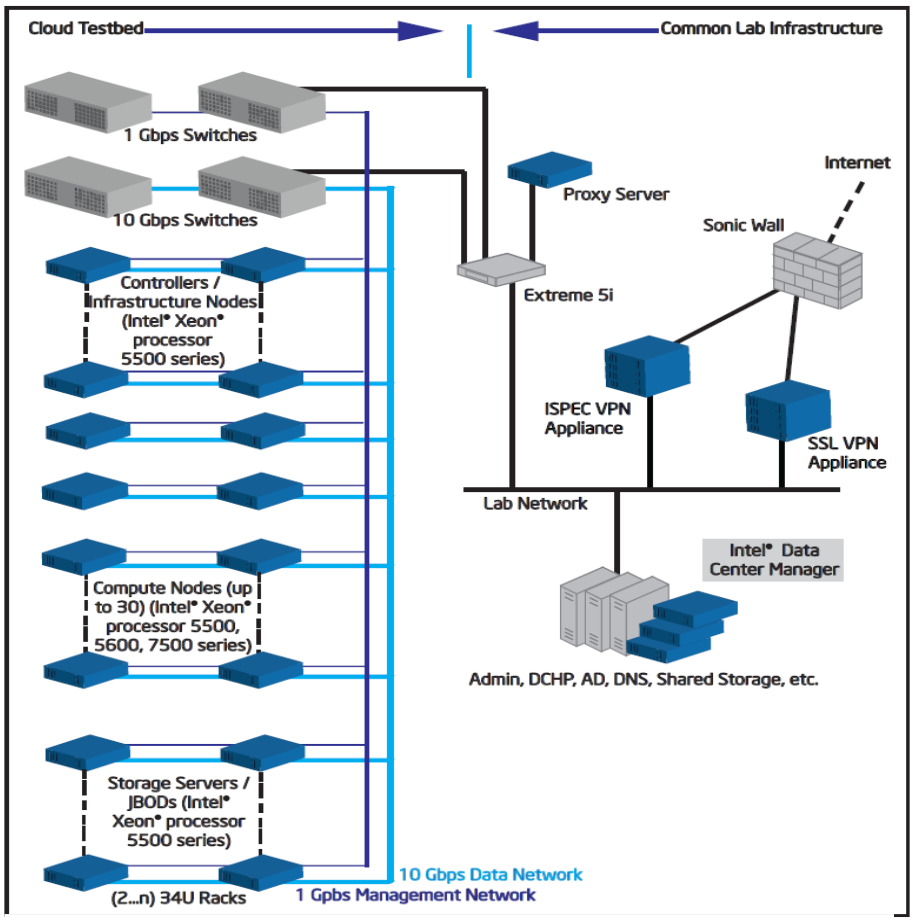

Figure 2 - Physical Representation of the Test Bed Cloud Infrastructure

Note: Figure 2 represents the Test Bed configuration, and while similar, is not the actual Powerleader configuration.

- The 1 GbE network was used for cloud management (provisioning/de-provisioning) of cloud resources and carries all cloud user (administrative and application) traffic.
- The 10 GbE network was available but not utilized.
- A single-tiered domain structure to accommodate the various needs of the Microsoft* components was implemented; this included an Active Directory Server with Domain Naming System (DNS) and Dynamic Host Configuration Protocol (DHCP) roles. The setup provided the capabilities for key infrastructure components to utilize the features of account control, domain name resolution, and dynamic IP addressing for virtual machines.
- The use of off-the-shelf x86 Intel® architecture-based server building blocks means that any server in the rack can be used for any purpose.
- 10 dedicated compute nodes were made available for the VMMSSP* to create and configure virtual machines. Each compute node had either a 150 or 200 GB virtual drive attached through an Internet Small Computer System Interface (iSCSI) located on a dedicated storage server. When attached and configured, the iSCSI drives look like direct attached storage devices.
- Microsoft Windows Storage Server* 2008 R2 was implemented on a dedicated system containing 1.8 Terabytes of storage. The storage was divided into 8 virtual iSCSI targets ranging from 150 Gb to 200 Gb and assigned to specific compute nodes. This was merely a convenient and low-cost way to make a large quantity of storage capacity accessible to the compute nodes.

The above considerations are typical of a cloud data center design: highly efficient servers, single- or multi-tiered networks, high virtualization, and consolidated storage. These attributes allow cost-effective operation and support a highly-automated infrastructure.

Hardware Description

The Intel® Xeon® processor 5500 and 5600 series, Intel's latest innovation in processor technology, which provides a foundation for designing new cloud data centers to achieve greater performance while using less energy and space and dramatically reducing operating costs[4], was used for the purposes of this guide. The Intel® Xeon® processor 5500 and 5600 series combine unprecedented security, performance, and energy efficiency, including:

- Intelligent performance that automatically varies the processor frequency to meet business and application performance requirements[5].
- Automated energy efficiency that scales energy usage to the workload to achieve optimal performance per watt and to reduce operating costs.
- Flexible virtualization that offers best-in-class performance and manageability in virtualized environments to strengthen the infrastructure and to reduce costs[6].

Hardware Topology

The test bed used 8 compute nodes utilizing Intel® Xeon® processor 5600 series. The Active Directory Server was set up as a private network with its own DNS and DHCP. Initial configuration consisted of:

- Two infrastructure nodes.
- Eight compute nodes.
- 1GbE network.

System	Processor Configuration	Additional Information
2 Primary Infrastructure Nodes	Intel® Xeon® Processor See the product page at http://www.intel.com/p/en_US/products/server/processor/xeon5000	Powerleader PR2510N Server Processor: Intel® Xeon® processor E5620 2-way x 4 cores = 8 cores Memory: 24 GB (6*4GB) DDR3 RAM Storage: 2* 1000 GB HDD, RAID1 Network: 2* 1000Mb
8 Compute Nodes	Intel® Xeon® Processor See the product page at http://www.intel.com/p/en_US/products/server/processor/xeon5000	Powerleader PR1510N Server Processor: Intel® Xeon® processor E5620 2-way x 4 cores = 8cores Memory: 24GB (6*4GB) DDR3 RAM Storage: 1000 GB HDD DAS, 2.4TB iSCSI (Shared)
Storage Servers	Intel® Xeon® Processor See the product page at http://www.intel.com/p/en_US/products/server/processor/xeon5000	Powerleader PR2510N Server Processor: Intel® Xeon® processor E5620*2 Memory: 24 GB RAM Storage min: 4x1000 GB HDD configured (RAID10) 2.4TB allocated for iSCSI Targets

Table 1 - Cloud Test Bed Systems Configuration

Software Description

System Component	Software Component
Hypervisor VMM	Microsoft Windows Server* 2008 R2 with Hyper-V* role activated Microsoft System Center Virtual Machine Manager* 2008 R2 Microsoft System Center Virtual Machine Manager Administration Console* Microsoft System Center Virtual Machine Manager Self Service Portal* v. 2.0
Active Directory Server	Microsoft Windows Server* 2008 (64-bit) AD server configured as primary domain controller DNS enabled DHCP enabled
SQL Server	64-bit Microsoft SQL Server* 2008 R2
Web Server	64-bit Microsoft IIS Server 7.5 on Windows* 2008 R2
Compute Nodes	64-bit Microsoft Windows Server* 2008 R2
Storage Node	Microsoft Windows Storage Server* 2008 R2 2.4 TB of available storage space (8) iSCSI targets created

Table 2 - Cloud Test Bed Software Configuration

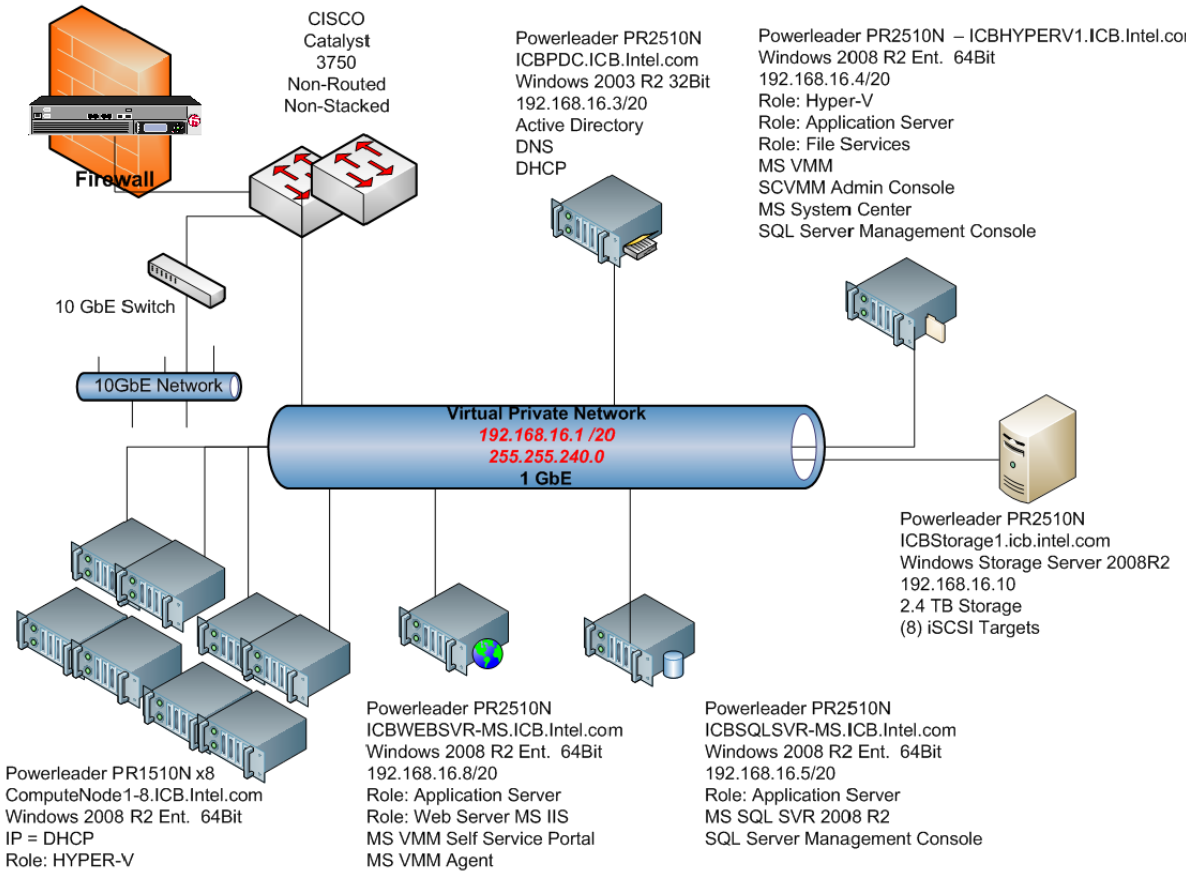

Figure 3 - Intel® Cloud Builders Layout for Microsoft VMMSSP*

Technical Review

Use Case Details

The system described in this guide went through a number of use cases ranging from basic system functionality to advanced cases, including performance and power management. The following use cases were applied:

- Configure the portal for the data center.
- Register a business unit.
- Create users/view users.
- Configure the virtual machine templates.
- Configure the network.
- Create the infrastructure.
- Approve an infrastructure request.
- Create multiple virtual machines.
- Access a virtual machine using the Web portal.
- Delete virtual machines using the Web portal.
- Submit a change request for additional storage.
- Decommission an infrastructure.
- Compare performance between Intel® Xeon® X5570 and Intel® Xeon® X5670 processors.
- Apply a power-capping policy on compute nodes in a cloud environment.
- Dynamic power reconfiguration.

Actors

For the following use cases, *DCITAdmin1* is the data center administrator with access to resources such as compute nodes, network, storage, load balancers, virtual machine templates, and domains. *BUITAdmin1* and *BUITAdmin2* are business unit administrators who can request resources from *DCITAdmin1* to host their IT services.

Preconditions

1. The Admin portal is available.
2. *BUITAdmin1* and *BUITAdmin2* have user accounts in the domain.
3. *BUITAdmin1* and *BUITAdmin2* have requested to register their respective business units and *DCITAdmin* has approved their request. (If the business unit is registered, the admin account is created automatically by the toolkit).
4. Templates are already created and available on the library server.
5. The hypervisor needs to be preconfigured for the hosts.
6. The request for infrastructure creation has already been submitted by *BUITAdmin1*, and *DCITAdmin* has approved it.
7. The virtual machine is already created and running.
8. For the power management use cases, the virtual machines also have the Java workload installed and running.
9. Intel® Data Center Manager (Intel® DCM) is installed and configured on a console system.
10. The managed compute nodes are configured with firmware that is compatible with the applicable version of Intel® DCM, and the nodes are successfully registered with Intel® DCM.

Execution and Results

Configure the Portal for the Data Center

1. Log into the toolkit Web portal as the *DCITAdmin1* set up when the toolkit was installed. Verify that the following tabs are visible:
   - Requests
   - Infrastructure
   - Virtual Machines
   - Jobs
   - User Roles
   - Settings
2. Click the *Settings* tab, then click "Configure Datacenter Resources".

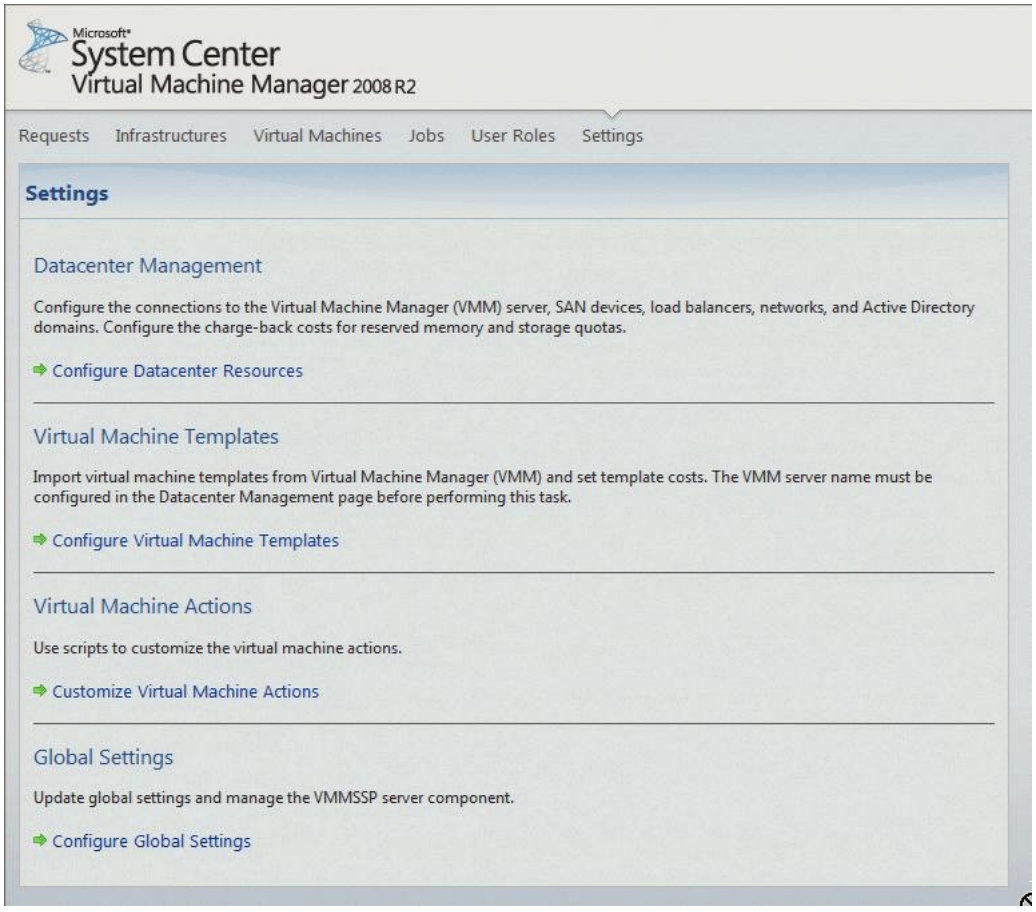

Figure 4 - Microsoft System Center Virtual Machine Manager* Portal

3. In the **Settings: Datacenter Management** page, enter the necessary information, then click "Save and Close".
4. The page closes and the *Settings* tab displays.
5. In the **Settings: Datacenter Management** page, verify all entered data is successfully captured in the system.

Register a Business Unit

1. Log into the portal as *BUITAdmin1*.

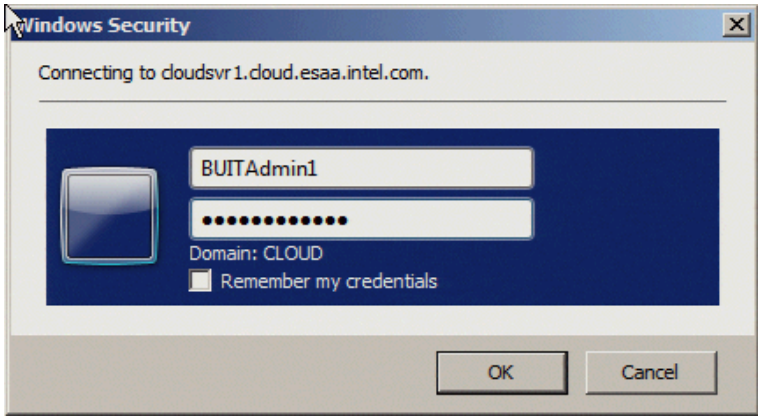

Figure 5 - Login Screen

2. In the toolkit self-service portal, click "Register New Business Unit", enter the necessary information, then click "Submit".
3. Log into the toolkit Web portal as *BUITAdmin2*.
4. In the toolkit self-service portal, click "Register New Business Unit", enter the necessary information, then click "Submit".
5. Log on to the toolkit Web portal as *DCIT Admin*. In the *Requests* tab, click the request, then click "Approve" to approve both business units.
6. Verify the business unit registration is submitted successfully.

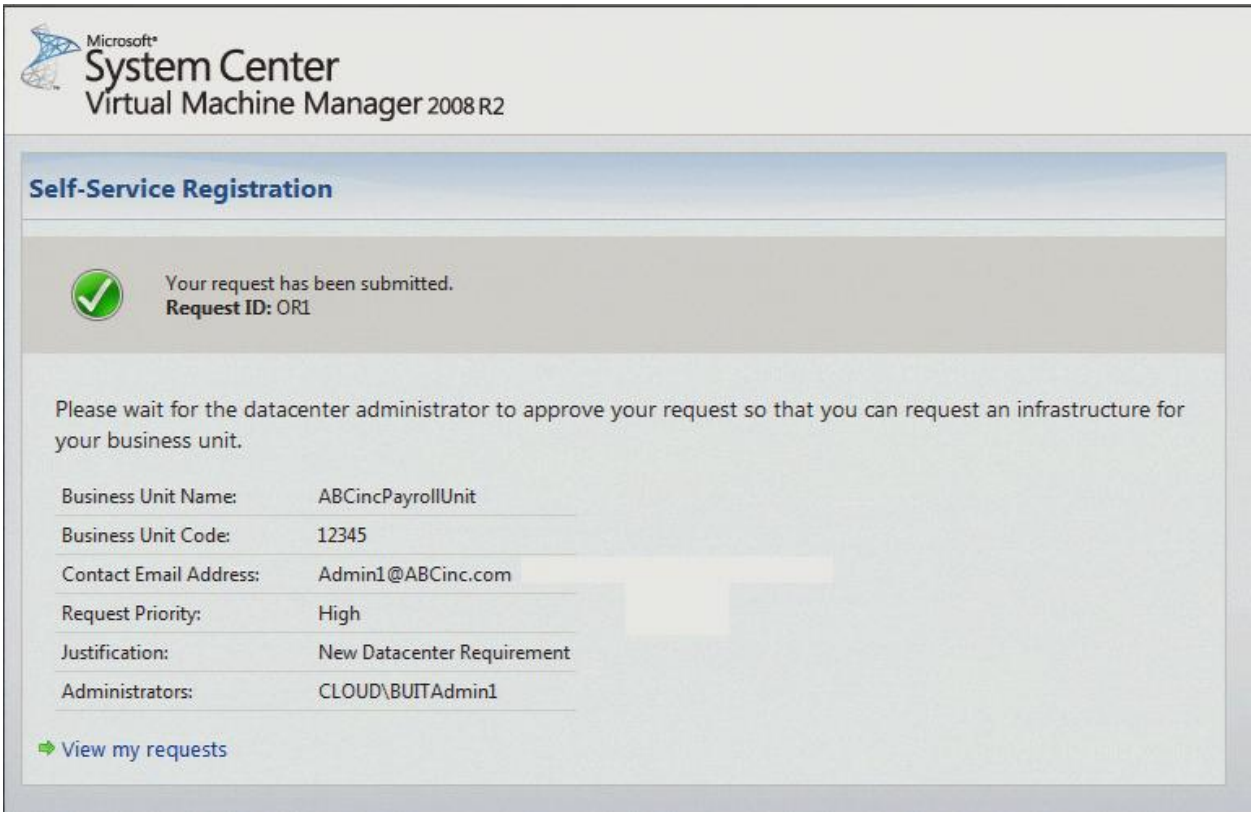

Figure 6 - Self-Service Registration Page

Create Users/View Users

1. Log into the toolkit Web portal as *DCITAdmin1*.
2. Click the *User Roles* tab, then highlight *BUITAdmin*.
3. Click "View/Edit Members" > "Users".
4. Verify that the *BUITAdmin1* and *BUITAdmin2* accounts are created successfully and are listed under Users.

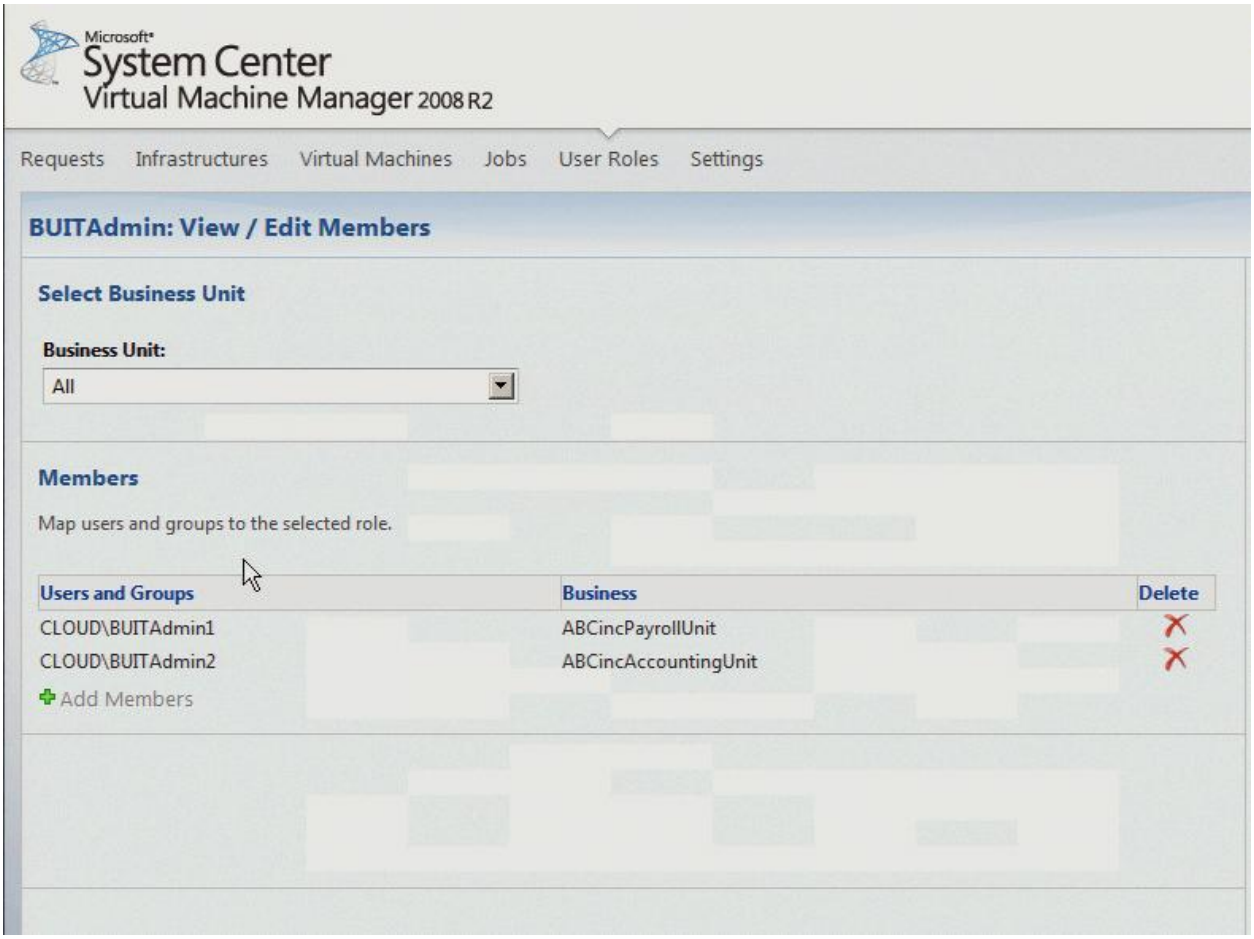

Figure 7 - View/Edit Members Page

Configure the Virtual Machine Templates

1. Log into the toolkit Web portal as *DCITAdmin1*.
2. Click the *Settings* tab, then click "Configure Virtual Machine Templates".
3. In the **Settings: Virtual Machine Templates** page, under "Other Tasks", click "Import Template".
4. In the **Import Virtual Machine Templates** page, select the library server and library shares from the drop-down boxes, then click "Search".
5. All of the templates in the library share are displayed.

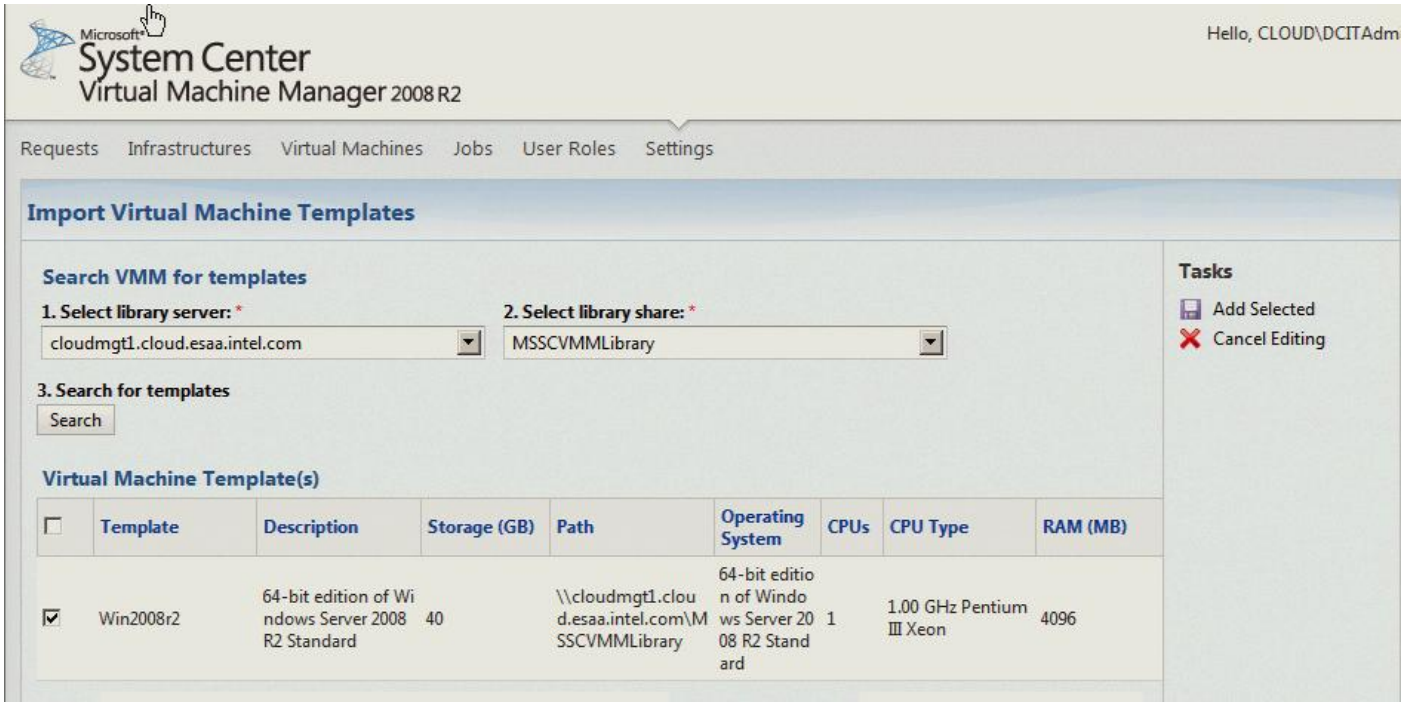

Figure 8 - Import Virtual Machine Templates Page

6. Select one or more templates that need to be made available to the toolkit, then click "Add Selected". The selected templates are added to the list.
7. Configure the cost for the template, then click "Save and Close". The template is saved and made available to the toolkit.

8. Open the **Settings: Virtual Machine Templates** page to verify that the template is listed.

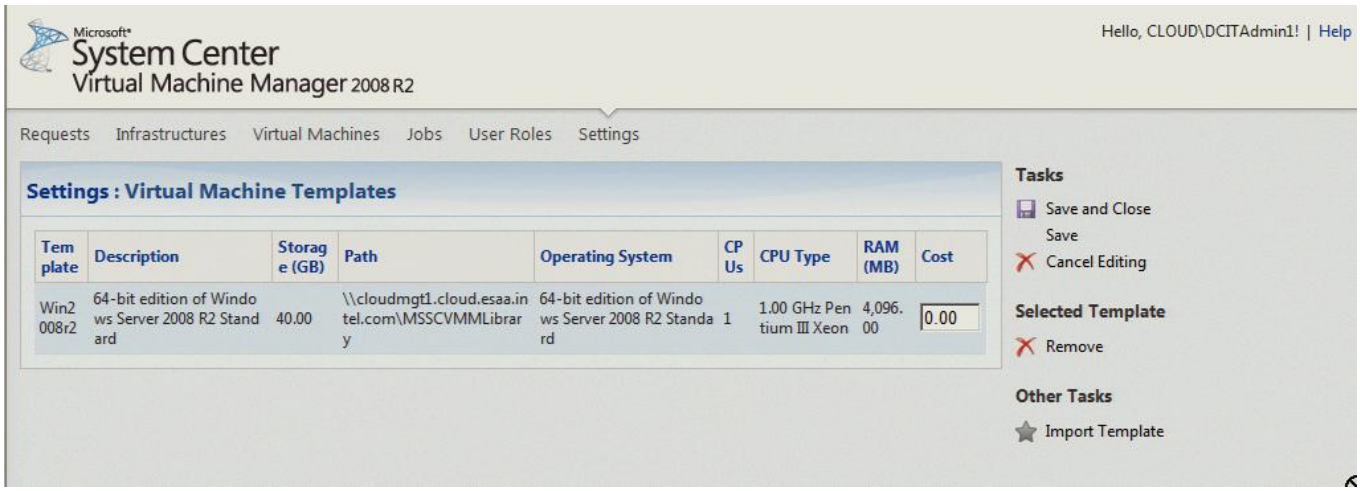

Figure 9 – Settings: Virtual Machine Templates

Configure the Network

1. Log into the toolkit Web portal as *DCITAdmin1*.
2. In the *Settings* tab, under "Datacenter Management", click the "Configure Datacenter Resources" link.
3. In the left column, click "Add Network".
4. In the "Add/Edit" a Network dialog box, fill-in the "Network Name" text box.
5. In the "Network Type" drop-down list, click either "Intranet" or "Private VLAN".
6. If the network uses static IP addresses, select the "Static IP" check box.
   - In the "Start IP Address" text box, specify the lower limit of IP addresses for virtual machines in the environment, then in the "End IP Address" text box, specify the upper limit of IP addresses.
   - Fill in the "Subnet Mask" text box, then click "Done".
7. Click "Submit".
8. In the *Settings* tab, verify the newly created network is present.

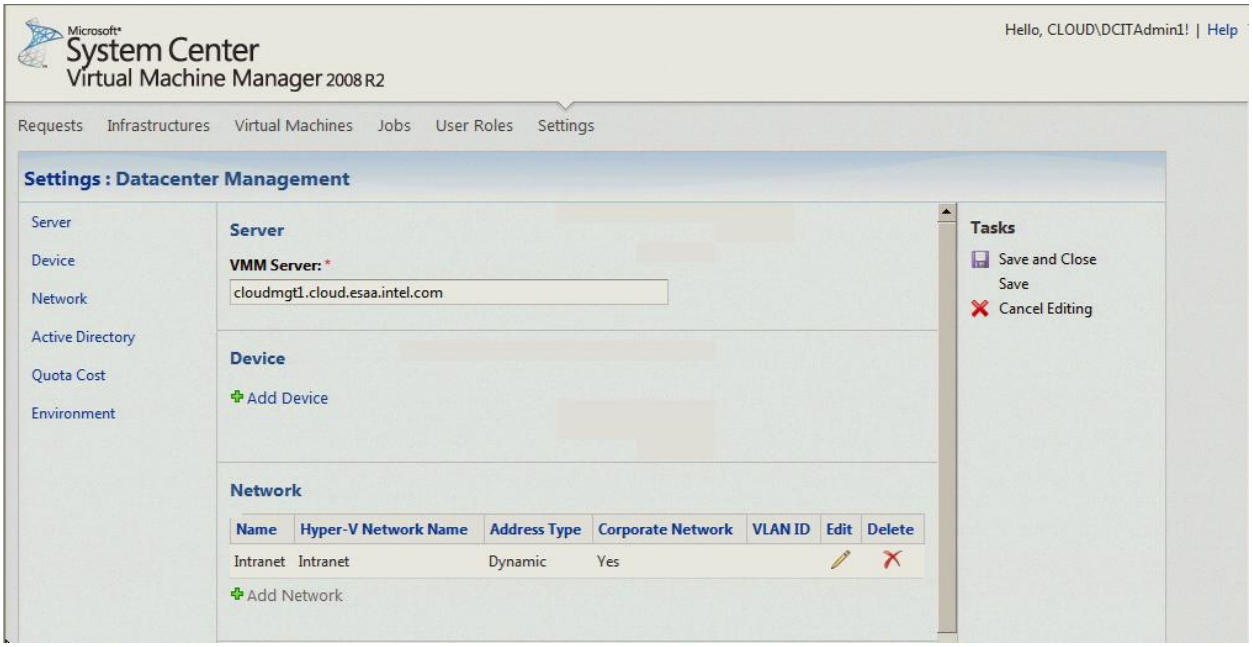

Figure 10 - Settings: Datacenter Management Page

Create the Infrastructure

1. Log into the toolkit Web portal as a *BUIT Admin* of the previously registered business unit.
2. In the *Requests* tab, click the business unit request approved previously, then click "Create infrastructure request".
3. In the **Create Infrastructure Request - Information** page, in the *Information* tab, fill in the boxes with the appropriate information, then click "Next".

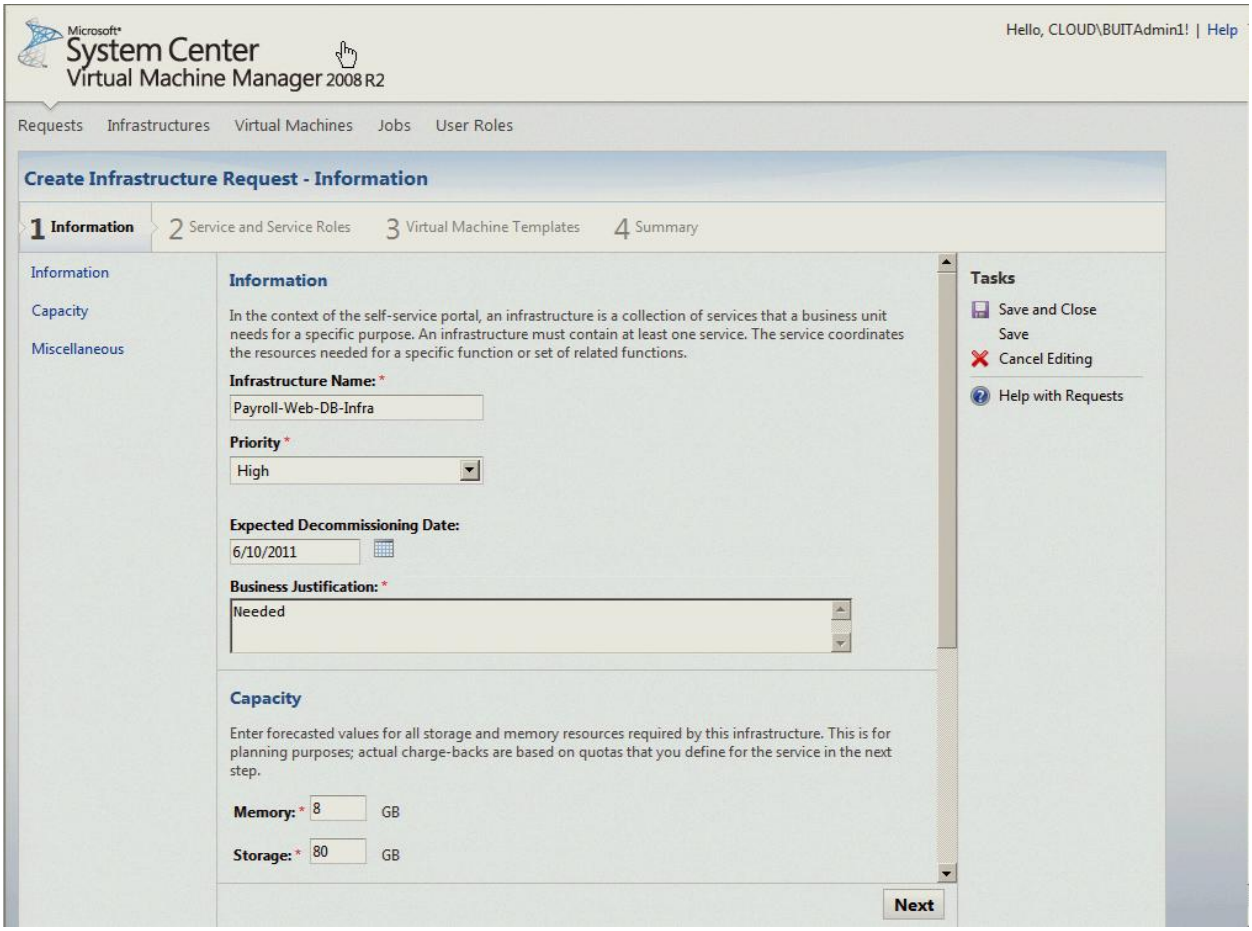

Figure 11 - Create Infrastructure Request - Information Page

4. In the *Service and Service Roles* tab, fill in the boxes with the appropriate information, then click "Add Service Role".
5. In the Add Service Role dialog box, enter the appropriate information, then click "Save and Close".
6. The Add Service Role dialog box closes and the *Service and Service Roles* tab opens. All information entered previously is retained. The service role is added successfully to the service.
7. Click "Next".
8. In the *Virtual Machine Template* tab, select a pre-configured template, then click "Next".
9. Click "Next"..

10. In the **My Requests** page, ensure all entered information is correct, then click "Submit Request".

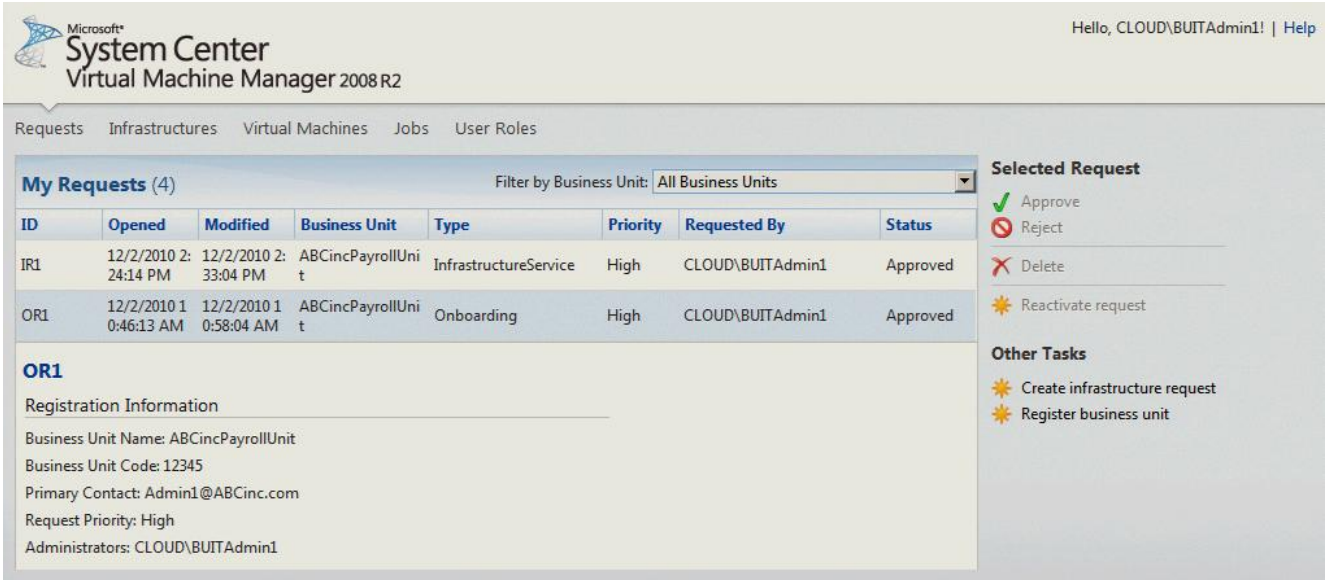

Figure 12 – My Reguests Page

Approve an Infrastructure Request

1. Log into the toolkit Web portal as *DCITAdmin1*.
2. Click the *Requests* tab.
3. Highlight the requested infrastructure service submitted by *BUITAdmin1*.
4. Click the pencil icon to edit the service. Fill out the information, including:
   - Host group
   - Library server location
   - Information to store offline virtual machines.
5. Click the pencil icon next to each service role and edit the service role information.
6. Approve the infrastructure request.
7. In the *Virtual Machine Template* tab, select a pre-configured template and associate a cost to it.
8. In the Summary page, verify that all of the entered information is correct, then click "Save and Close".
9. Click "Approve".
10. Click the *Infrastructures* tab and verify that the approved infrastructure is listed under "My Infrastructures".

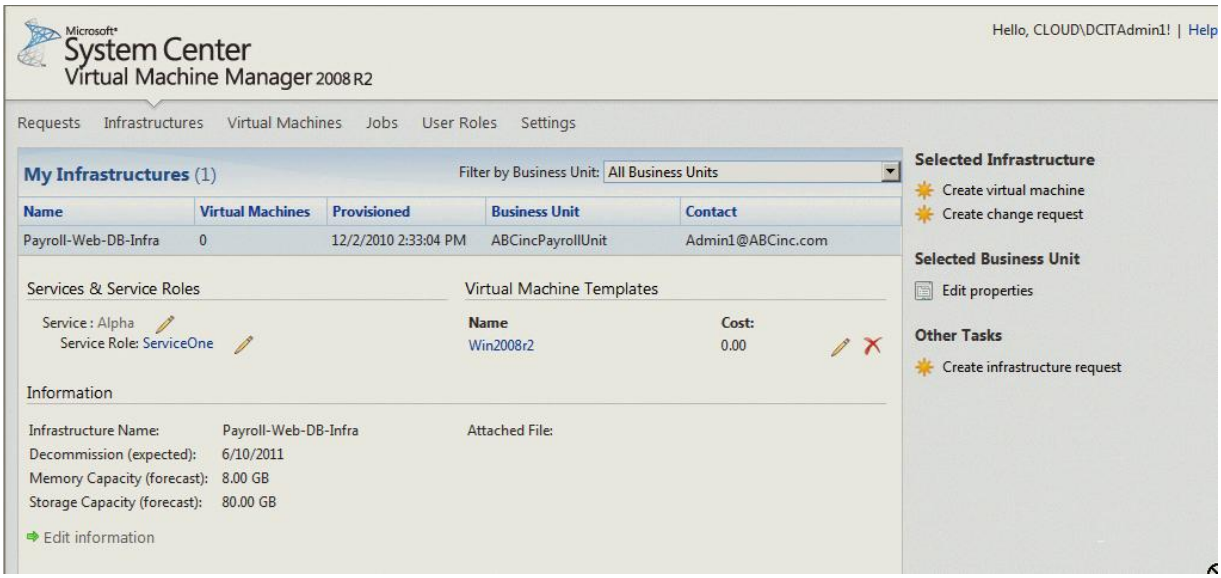

Figure 13 - My Infrastructures Page

Create Multiple Virtual Machines

1. Log into the toolkit Web portal as *BUITAdmin1* for the business unit registered previously.
2. Click the *Virtual Machines* tab.
3. Click "Create virtual machine".
4. In the "Create Virtual Machine" form, enter the following information:
   - **Business Unit:** Select the business unit configured previously
   - **Infrastructure:** Select the infrastructure created previously
   - **Service:** Select the service created previously
   - **Service Role:** Select the service role created previously
   - **Option:** With a single name and sequential number
   - **Number of virtual machines to create:** 16
   - **Computer Name:** DBServer_
   - **Index suffix start:** 001
   - **Template:** Select the name of the virtual machine template associated with the service
   - **Network:** Select the name of the network configured previously
5. Click "Create".
6. The "Create Virtual Machine" request is submitted successfully. The toolkit creates corresponding jobs to carry out the operation.
7. In the *Jobs* tab, monitor the jobs to verify they are created for each virtual machine.
8. Go to the host computer to verify the virtual machines are created successfully.

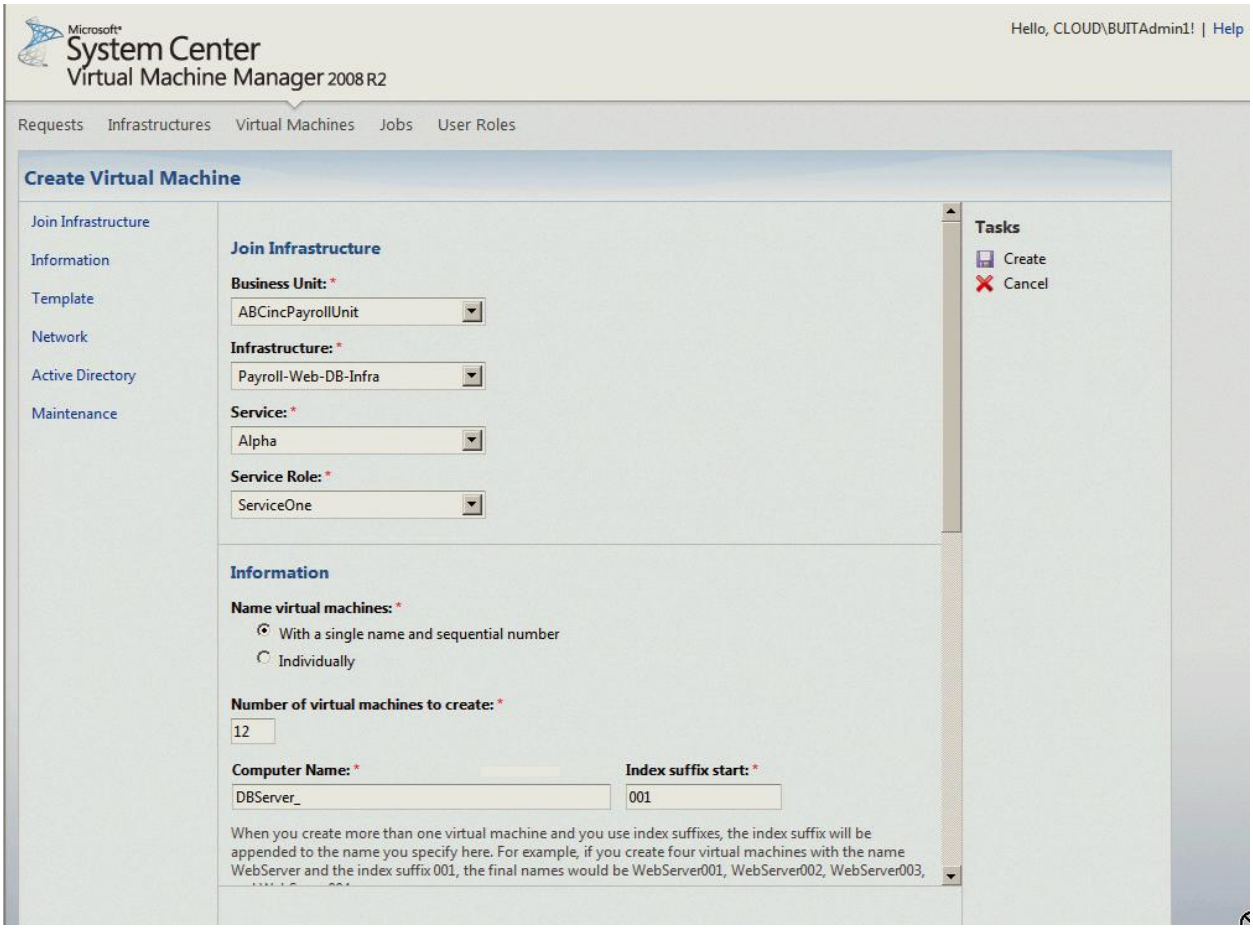

Figure 14 - Create Virtual Machine Page

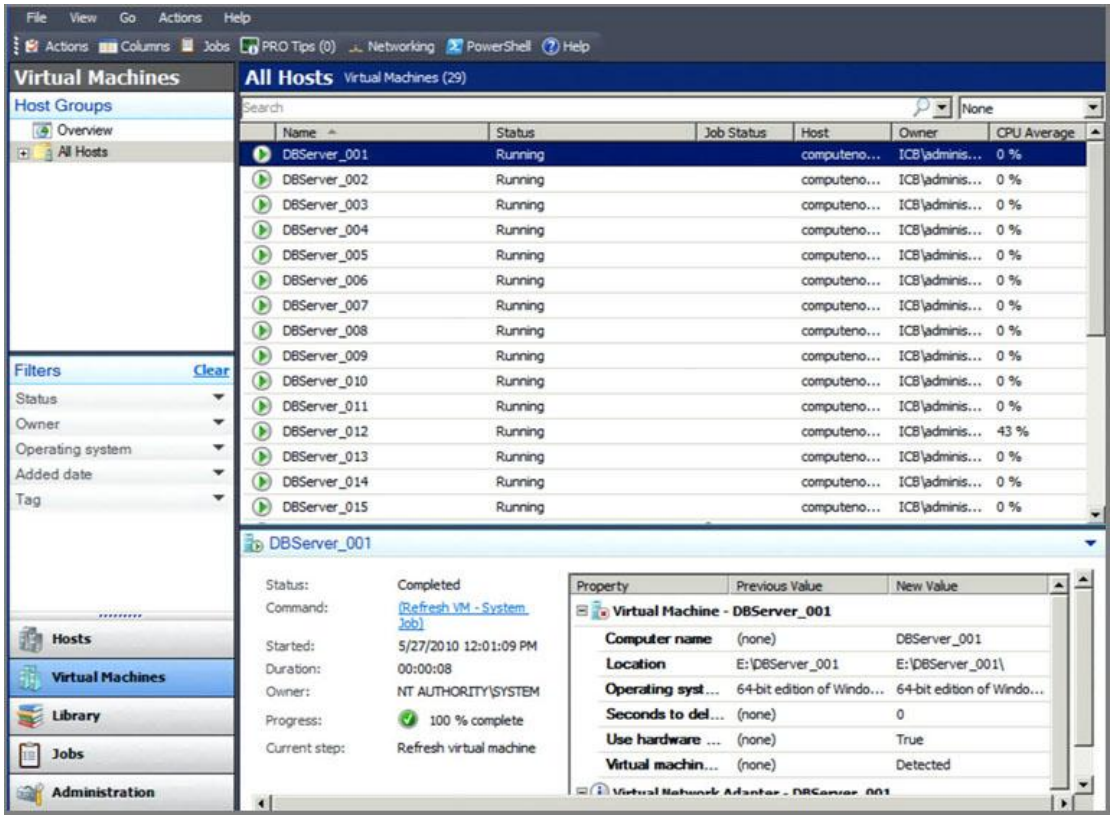

Figure 15 - View Virtual Machines

Access a Virtual Machine Using the Web Portal

1. Log into the toolkit Web portal as *BUITAdmin1* for the business unit registered previously.
2. Click the *Virtual Machines* tab.
3. Highlight a virtual machine to connect to.
4. Click "Connect".
5. Verify that the virtual machine is accessible from the Web portal.

Delete Virtual Machines Using the Web Portal

1. Log into the toolkit Web portal as *BUITAdmin1*.
2. Click the *Virtual Machines* tab.
3. Highlight the virtual machine to be deleted.
4. Click "Shut down" to stop the virtual machine.
5. Wait until the Job Status changes from "Running" to "Stopped".
6. Click "Delete".
7. Verify the selected virtual machine is deleted.

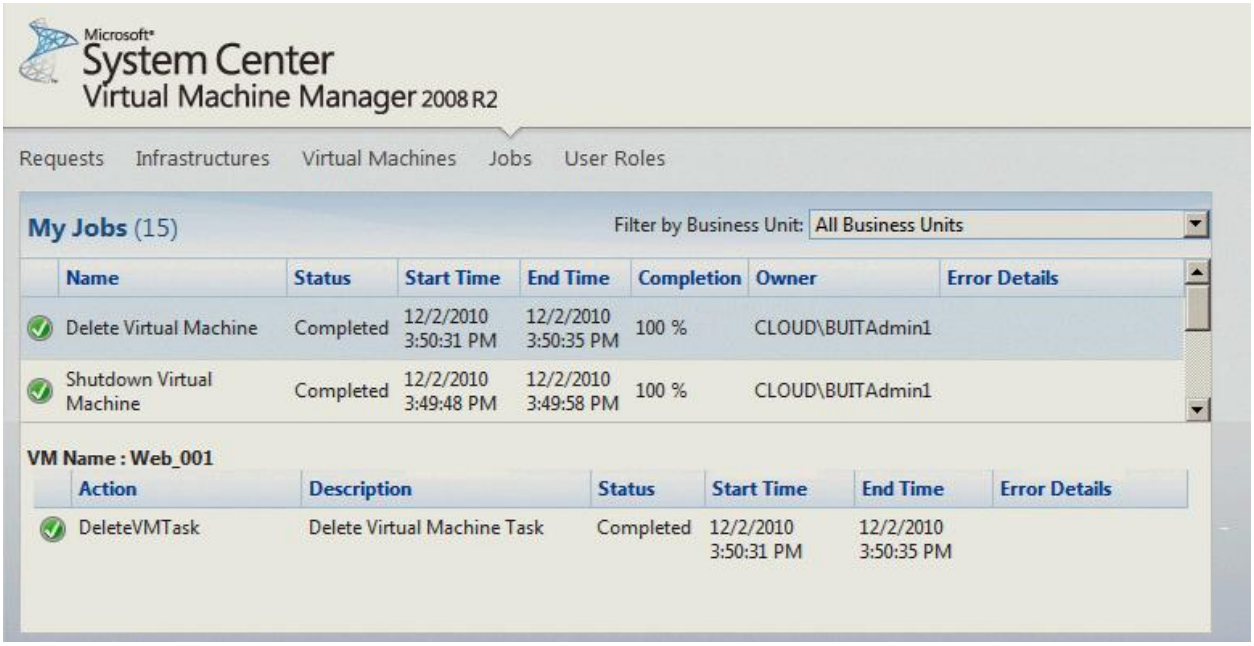

Figure 16 – My Jobs Page

Submit a Change Request for Additional Storage

1. Log into the toolkit Web portal as *BUITAdmin1*.
2. Click the *Infrastructures* tab.
3. Select the infrastructure on which the change request is to be submitted.
4. Click "Create Change Request".
5. Select the Priority from the drop down list.
6. Enter a Justification in the Justification text box.
7. In the "Object to Change" drop down list, select "Infrastructure".
8. In the "Type of Change" drop down list, select "Edit".
9. Click "Next".
10. Enter the amount of storage required, then click "Next".
11. Click "Submit Request".
12. Log into the toolkit Web portal as *DCITAdmin1*.
13. Approve the request to change the amount of storage.
14. Verify the requested storage is available.

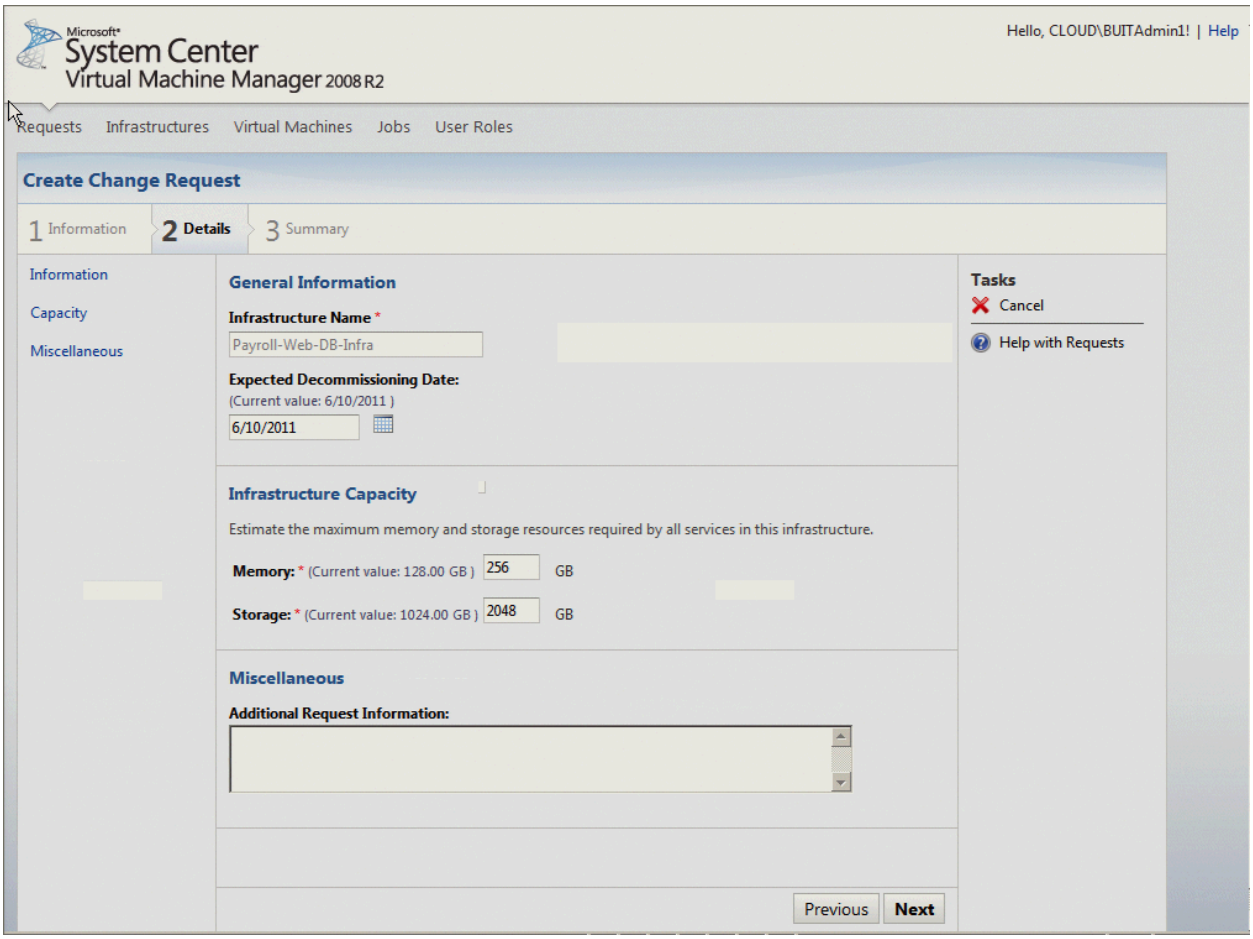

Figure 17 - Create Change Request Page

Decommission an Infrastructure

1. Log into the toolkit Web portal as *BUITAdmin1*.
2. Click the *Infrastructures* tab.
3. Click "Create Change Request".
4. Select the Priority from the drop down list.
5. Enter a Justification in the Justification text box.
6. In the "Object to Change" drop-down list, select "Infrastructure".
7. In the "Type of Change" drop-down list, select "Decommission".
8. Click "Next".
9. Click "Submit Request". The decommission infrastructure request will be created for the selected infrastructure.
10. Log into the toolkit Web portal as *DCITAdmin1*.
11. Approve the request to decommission the infrastructure.
12. Verify the infrastructure is decommissioned.

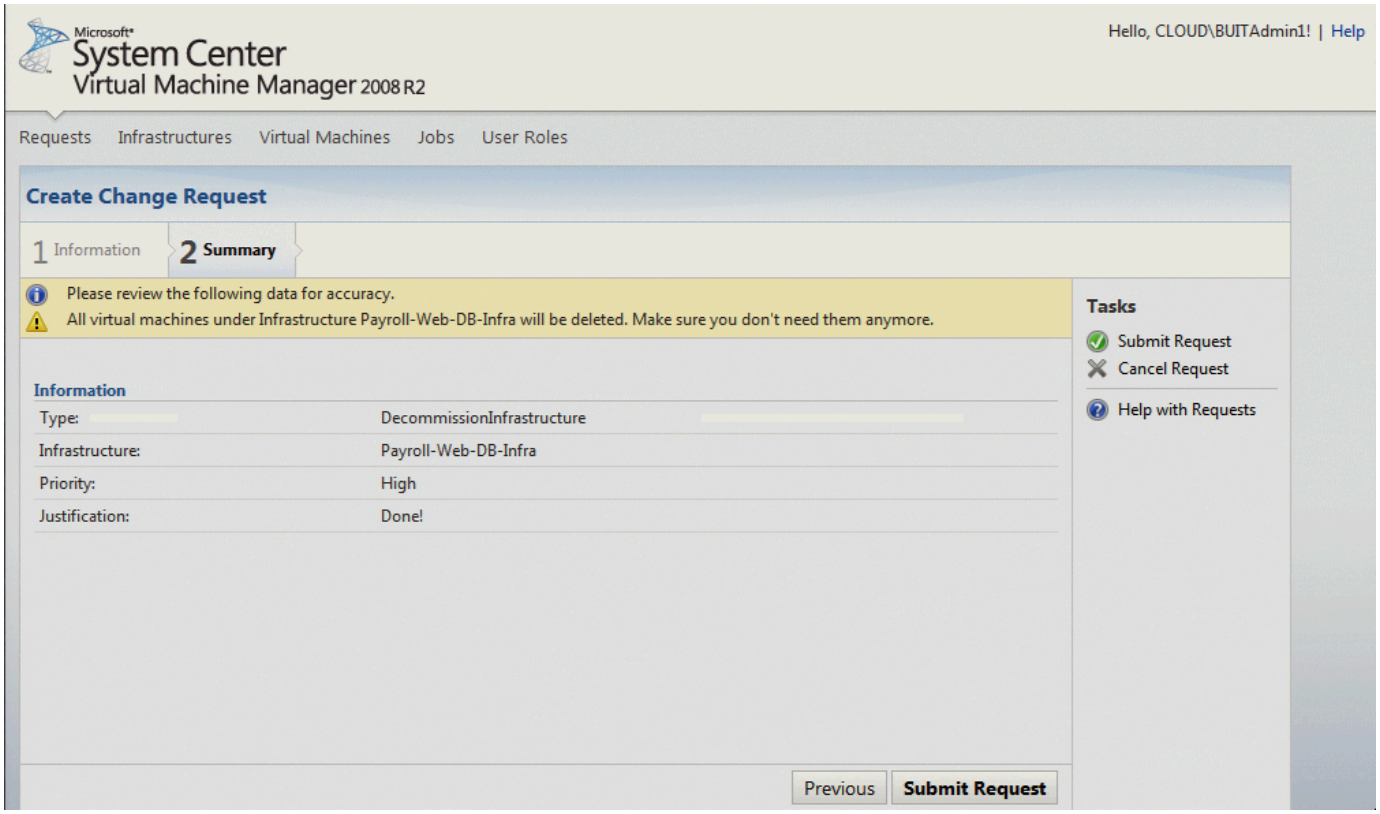

Figure 18 - Create Change Request Page

Things to Consider

Storage

For this guide, a storage iSCSI for the bulk storage requirements in the cloud was used. The storage server could have been connected to a Storage Area Network (SAN). The choice in this case was completely arbitrary and based solely on the capabilities of the lab setup. The more fundamental requirement is for the VMMSSP* to have access to a shared storage device with sufficient capacity and performance. A network attached storage (NAS) device connected to a SAN would have worked just as well. The choice is driven by design factors such as the back-up strategy, performance, and cost.

Scalability

The scalability of the solution is heavily impacted by:

- Network technology (e.g. 10 GbE) and architecture.
- Selected storage architecture.
- Choice of server hardware for compute nodes and storage nodes.

SSD Drives Use

The performance of storage nodes (and compute nodes when local storage is utilized), as well as the overall power consumption of the cloud, may be favorably impacted by the use of solid-state (SSD) drives. This was not specifically tested for this guide.

Networking

VMMSSP* supports several network technology architectures. For this test bed, a very simple network topology

we utilized. The network for this test bed performed well, but more advanced technologies and architectures (e.g. 10 GbE, FCoE) are more suited for production deployments, especially when performance is considered.

Conclusion

In this guide, the components and implementation process for a private cloud is described, a prototype built jointly by Intel and Microsoft* with the goal of demonstrating possible architecture and deployment of a private cloud. The prototype supports a virtual data center on its own VLAN with its own designated storage, physical, and virtual compute resources. The prototype system has the standard functionality expected from any cloud system, such as the ability to provision and destroy virtual machines as needed. The VMMSSP* functionality was illustrated with several real world use cases with defined resource pools, multiple virtual machines were created, and the infrastructure was decommissioned. Also, possible enhancements that could improve performance and scalability of the prototype were described in this guide, which would make the system more suitable for real production environments.

Additional Information

- Intel® Cloud Builders: http://intel.com/cloudbuilders
- Intel® Xeon® processors: http://intel.com/xeon
- Powerleader Servers: http://www.powerleader.com.cn
- Microsoft*: http://www.microsoft.com

Glossary

Intel® Data Center Manager (Intel® DCM): Intel® DCM is an SDK from Intel that provides policy-based tools for managing power in the data center. Used in conjunction with Intel® Intelligent Node Manager and integrated with management consoles, Intel® DCM provides benefits such as increased rack density within space, power and cooling constraints through fine-grained power control, and reduced capital costs by right-sizing power and cooling infrastructure based on actual power consumption trends.

Data Center Manageability Interface (DCMI): DCMI Specifications are derived from Intelligent Platform Management Interface (IPMI) 2.0, which has been widely adopted by the computing industry for server management and system-health monitoring. The DCMI specifications define a uniform set of monitoring, control features, and interfaces that target the common and fundamental hardware management needs of server systems that are used in large deployments within data centers, such as Internet Portal data centers. This includes capabilities such as secure power and reset control, temperature monitoring, event logging, and others.

Hardware profile: Represents a non-overlapping, logical group of nodes that are provisioned with the identical set of kernel, network configuration, disk partitioning, and device modules.

Infrastructure as a Service (IaaS): Infrastructure as a Service is the delivery of technology infrastructure, such as network, storage, and compute, as a service, typically through virtualization. Users subscribe to this virtual infrastructure on demand as opposed to purchasing servers, software, data center space, or network equipment. Billing is typically based on the resources consumed.

Software profile: Defines non-overlapping, logical groups of nodes that are provisioned with identical sets of software.

Power Manager: The function that is responsible for managing the power utilization in the cloud data center. The use cases described in this guide used Intel® DCM for the Power Manager.

VDC: Virtual Data Center.

VM: Virtual Machine.

Intel* Virtualization Technology (Intel® VT): Provides comprehensive hardware assists, which boost virtualization software performance and improve application response times. Intel® VT reduces demands placed on virtualization software so that more applications and heavier workloads per server can be consolidated to get better value from server and software investments.

Endnotes

- [1] Microsoft VMMSSP, http://www.microsoft.com/virtualization/en/us/private-cloud.aspx
- [2] Apache Hadoop, http://hadoop.apache.org/
- [3] Intel® Xeon® Processor 5500 Series Software Industry Testimonials, http://www.intel.com/business/software/testimonials/xeon5500.htm and Intel® Xeon® Processor 5600 Series, http://www.intel.com/itcenter/products/xeon/5600/index.htm
- [4] Why the Intel® Xeon® Processor 5500 Series is the Ideal Foundation for Cloud Computing, http://communities.intel.com/docs/DOC-4213 and Intel in cloud computing Wiki, http://communities.intel.com/docs/DOC-4230
- [5] Intel® Xeon® Processor 5500 Series Software Industry Testimonials, http://www.intel.com/business/software/testimonials/xeon5500.htm and Intel® Xeon® Processor 5600 Series, http://www.intel.com/itcenter/products/xeon/5600/index.htm
- [6] Intel® Virtualization Technology, http://www.intel.com/technology/virtualization/

To learn more about deployment of cloud solutions, visit http://intel.com/cloudbuilders

Disclaimers

Δ Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See www.intel.com/products/processor_number for details.

? Hyper-Threading Technology requires a computer system with an Intel processor supporting Hyper-Threading Technology and an HT Technology enabled chipset, BIOS and operating system. Performance will vary depending on the specific hardware and software you use. See http://www.intel.com/info/hyperthreading/ for more information including details on which processors support HT Technology.

◊ Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM) and, for some uses, certain platform software enabled for it. Functionality, performance or other benefits will vary depending on hardware and software configurations and may require a BIOS update. Software applications may not be compatible with all operating systems. Please check with your application vendor.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order. Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web site at www.intel.com.

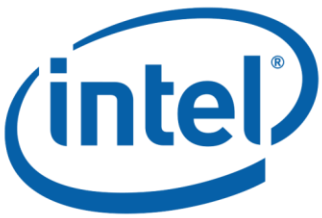

Printed in the USA 0111/IP/PRW/PDF Please Recycle 324948-001US