Intel® Cloud Builder Guide
Intel® Xeon® Processor-based Servers
Microsoft System Center Virtual Machine
Manager Self-Service Portal 2.0*

# Intel® Cloud Builder Guide: Cloud Design and Deployment on Intel® Platforms

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 paper," 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 Intel® Xeon® processor series-based servers. Using the contents of this paper, which includes detailed scripts and screen shots, should significantly reduce the learning curve for building and operating your first 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 paper can be used "as-is." For example, adaptation to an existing network and identification of management requirements are out of scope for this paper. Therefore, it is expected that the user of this paper will make significant adjustments to the design to meet specific customer requirements. This paper 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 as they are:

- 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.

## Introduction

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 you can extend the default virtual machine actions; for example, you can add scripts that interact with storage area networks (SANs) or load balancers that support virtual machines.

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. You 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; you can clone this action XML segment in order to extend the virtual machine actions. In this way, if needed, you can create multiple custom sets of virtual machine actions.

The data center administrator can configure PowerShell scripts to include specific cmdlets they would like to execute 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.

### How Would an Enterprise Customer Use VMMSSP?

A common scenario is if you were a business unit IT manager who needs an application to be built, so you go to a portal where you 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 you, the IT manager, wish 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 you can access a self-service portal, where you can now take full control of the environment you custom built for your workload. Here you can create virtual machines, start them, stop them, or delete them.

Datacenter 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, where you have full control of the environment.
4. The DCIT administrator accesses reports regarding your workloads, costs, etc.

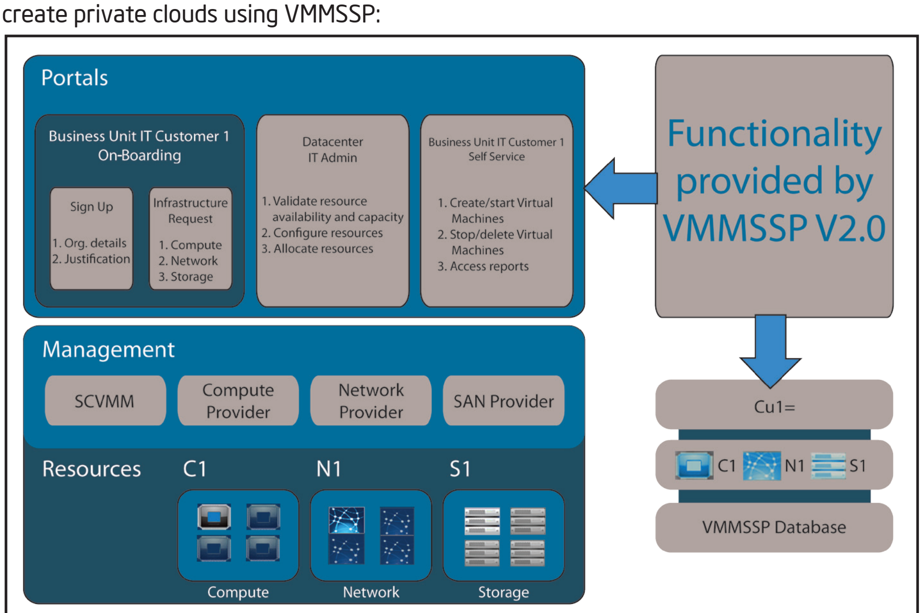


**Figure 1: Overview of VMMSSP V2.0**

### How Would a Systems Integrator (SI) Use VMMSSP?

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 datacenter. VMMSSP also provides guidance for extending virtual machine actions.

## Test Bed Blueprint Overview

### Test Bed Design Considerations

As the previously discussed 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 Builder Test Bed follows these principles. The diagram shown in Figure 2 represents a physical construction of the cloud environment

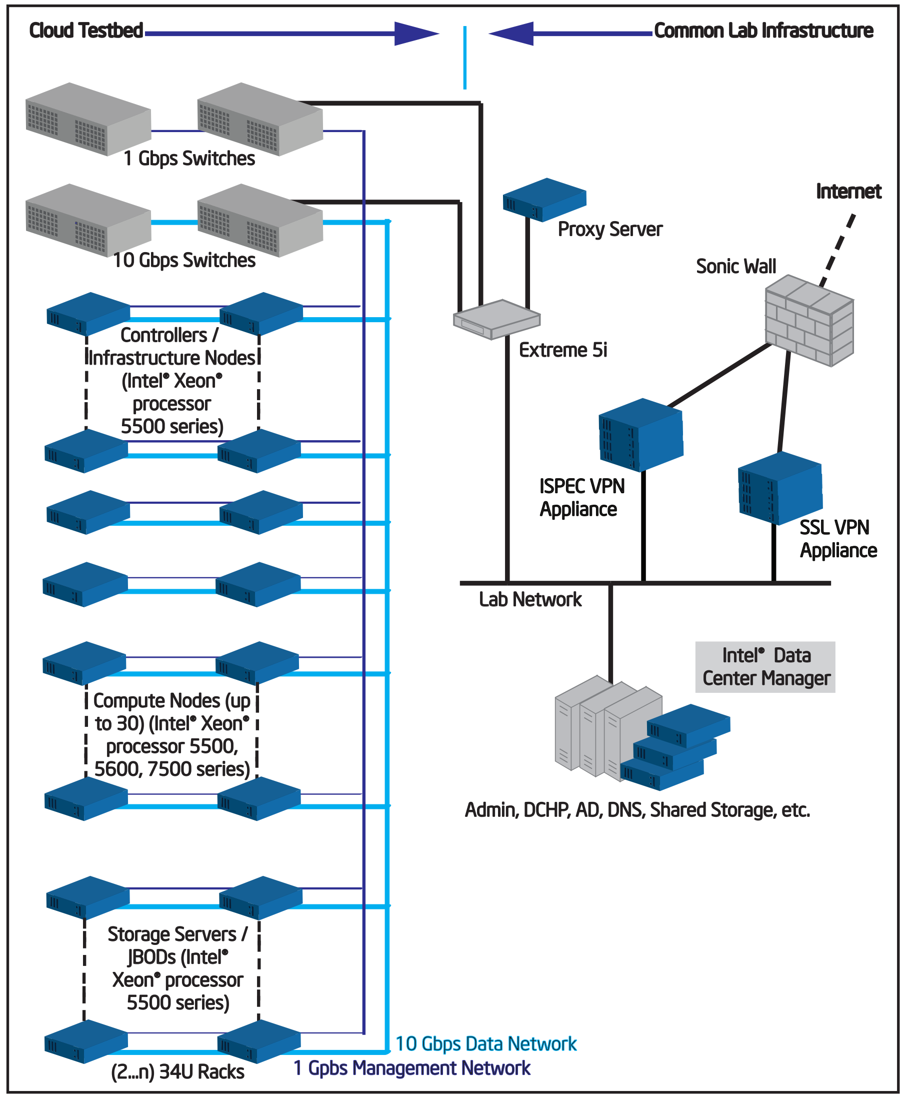


**Figure 2: Physical representation of the cloud infrastructure**

used for the tests discussed in this reference architecture.

Intel® Xeon® processor 5500 and 5600[3] series 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 given below:

- The 1 GbE network was used for cloud management (provisioning/de-provisioning) of cloud resources and carries all the cloud user (administrative and application) traffic.
- The 10 GbE network was available but not utilized.
- We chose to implement a single-tiered domain structure to accommodate the various needs of the Microsoft components; 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

We used Intel's latest innovation in processor technology, the Intel Xeon processor 5500 and 5600 series, 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]

The Intel Xeon processor 5500 and 5600 series combine unprecedented security, performance, and energy efficiency. Some of these features include:

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

### Topology

The test bed used up to eight compute nodes utilizing Intel Xeon processor 5500 series and two compute nodes utilizing Intel Xeon processor 5600 series-based servers. The Active Directory Server was set up as a private network with its own DNS and DHCP.

Initial configuration consisted of:

- Four infrastructure nodes
- Ten compute nodes
- 1GbE network

| System | Processor Configuration | Additional Information |
|---|---|---|
| **Primary Infrastructure Nodes** | **Intel® Xeon® Processor** X5570 | **Form Factor: 2U Rack Mount Server**<br>**Processor:** Intel Xeon processor 5500 series-based; **2.93GHz; 2-way x 4 cores = 8 cores**<br>**Memory: 24GB RAM**<br>**Storage: 300GB HDD**<br>**Intel Xeon Processor 5500 Series Product Information:**<br>http://ark.intel.com/ProductCollection.aspx?codeName=33163<br>Intel® Xeon® processor 5000 series product support:<br>http://www.intel.com/support/processors/xeon5k/ |
| **8 Compute Nodes** | **Intel Xeon Processor** X5570 | **Form Factor: 1U Rack Mount Server**<br>**Processor:** Intel Xeon processor 5500 series-based; **2.93GHz; 2-way x 4 cores = 8 cores**<br>**Memory: 24GB RAM**<br>**Storage: 136GB HDD DAS + 150 or 200 Gb iSCSI** |
| **2 Compute Nodes** | **Intel® Xeon® Processor** X5667 | Form Factor: 1U Rack Mount Server<br>**Processor:** Intel Xeon processor 5600 series-based; **3.06GHz; 2-way x 6 cores = 12 cores**<br>**Memory: 24GB RAM**<br>**Storage: 136GB HDD + 150GB or 200GB iSCSI attached** |
| **Storage Servers** | **Intel® Xeon® Processor** X5667 | **Form Factor: 5U Tower**<br>**Processor:** Intel Xeon processor 5500- series-based **2.93GHz ; 2-way x 4 cores = 8 cores**<br>**Memory: 24GB RAM**<br>**Storage: 6x300GB HDD configured as 1.4 TB RAID**<br>iSCSI Targets: **(4) 200GB, (4) 150 GB** |

**Table 1. Cloud Test Bed Systems Configuration**

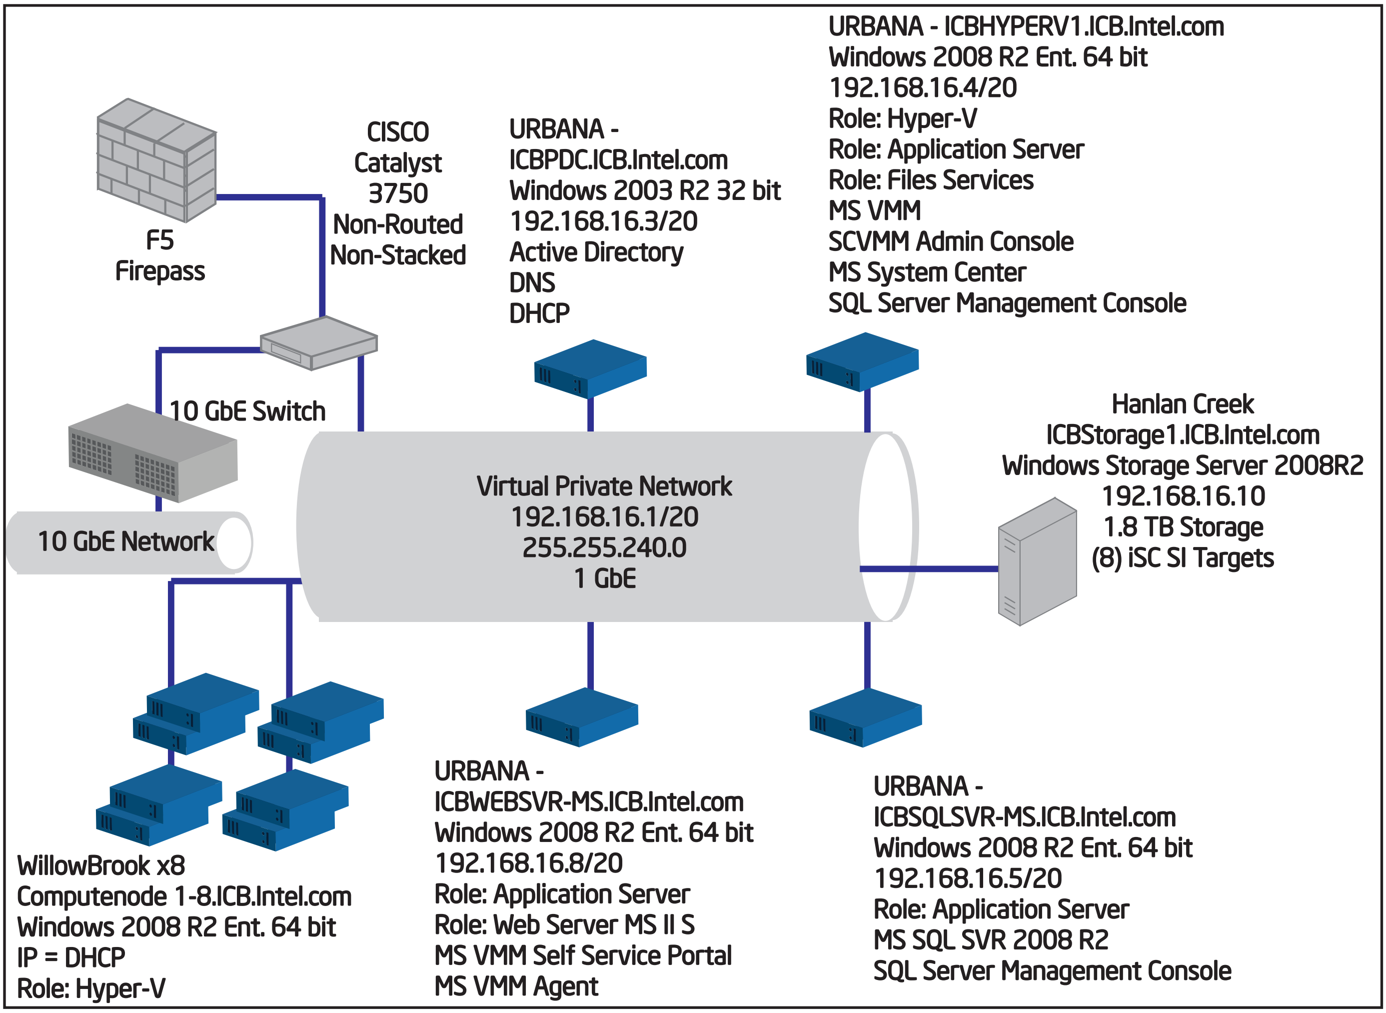


Figure 3: Intel Cloud Builder Layout for Microsoft VMMSSP

**Software Description**

| System Component | Software Component |
|---|---|
| **Hypervisor**<br>**VMM** | ▪ Microsoft Windows 2008 R2 with Hyper-V role activated<br>▪ System Center Virtual Machine Manager 2008 R2<br>▪ System Center Virtual Machine Manager Administration Console<br>▪ System Center Virtual Machine Manager Self Service Portal v. 2.0 |
| **Active Directory Server** | ▪ Windows Server 2003 (32-bit or 64-bit)<br>▪ AD server configured as primary domain controller<br>▪ DNS enabled<br>▪ 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 2008 R2 Server |
| **Storage Node** | ▪ Microsoft Windows Storage Server 2008 R2<br>▪ 1.4 TB of available storage space<br>▪ (8) iSCSI targets created |

Table 2: Cloud Test Bed Software Configuration

## Technical Review

### Use Case Details

We exercised the system through a number of use cases ranging from those covering basic system functionality to advanced cases, including performance and power management. We applied the following use cases:

- 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 all of the following use cases, DCITAdmin1 is the data center administrator who has 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*

### 1. Configure the Portal for the Data Center

1. Log on to the toolkit Web portal as the DCITAdmin1 that you set up when you installed the toolkit. Verify that the following tabs are visible: **Requests**, **Infrastructure**, **Virtual Machines**, **Jobs**, **Users**, and **Settings**.

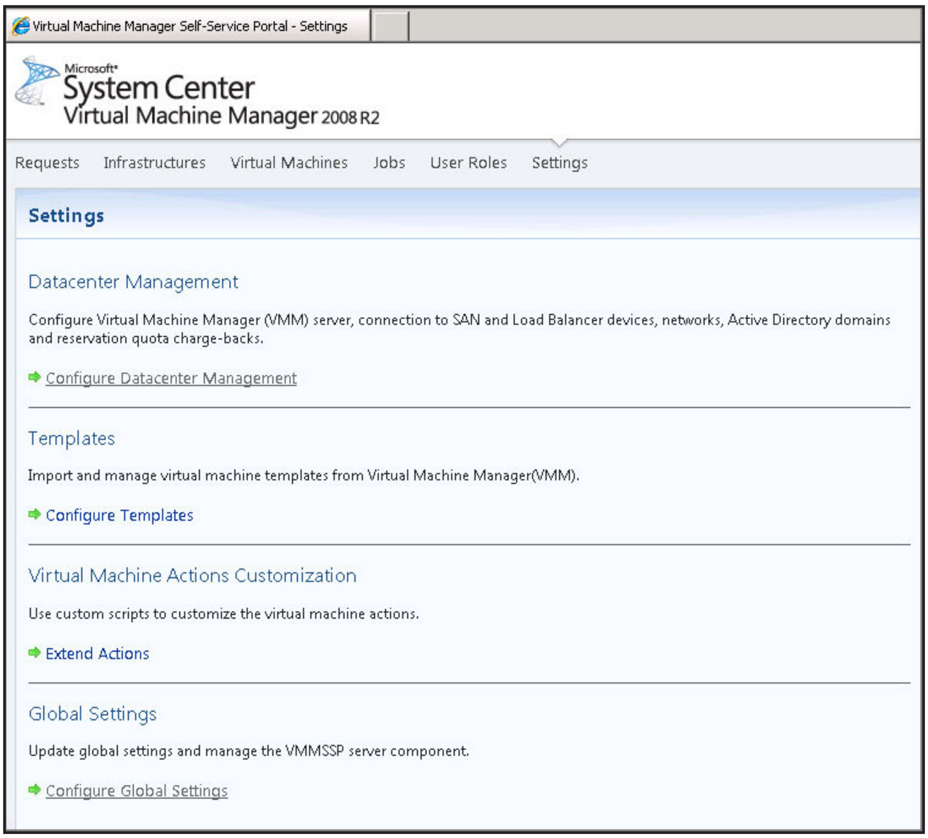


**Figure 4: Microsoft System Center Virtual Machine Manager Portal**

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

## 2. Register a Business Unit

1. Log on to the portal as BUITAdmin1.

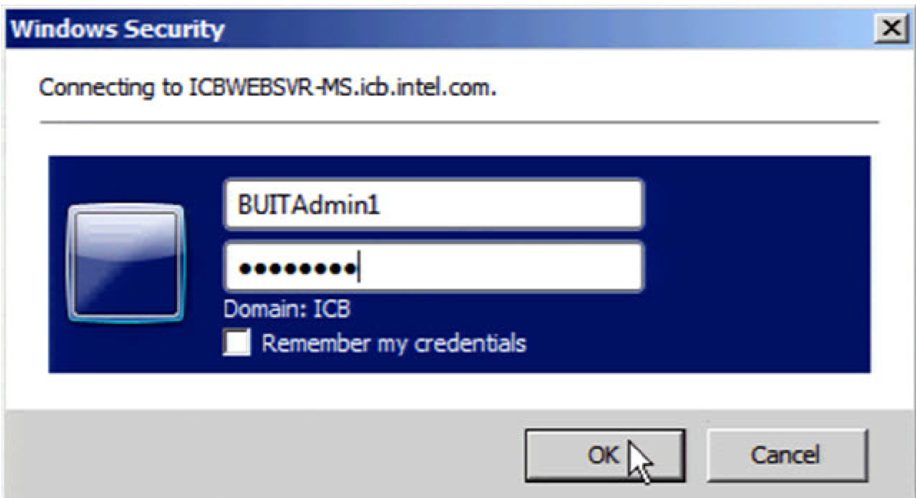


**Figure 5: Login Screen**

2. In the toolkit self-service portal, click **Register New Business Unit**, enter the necessary information, and then click **Submit**.

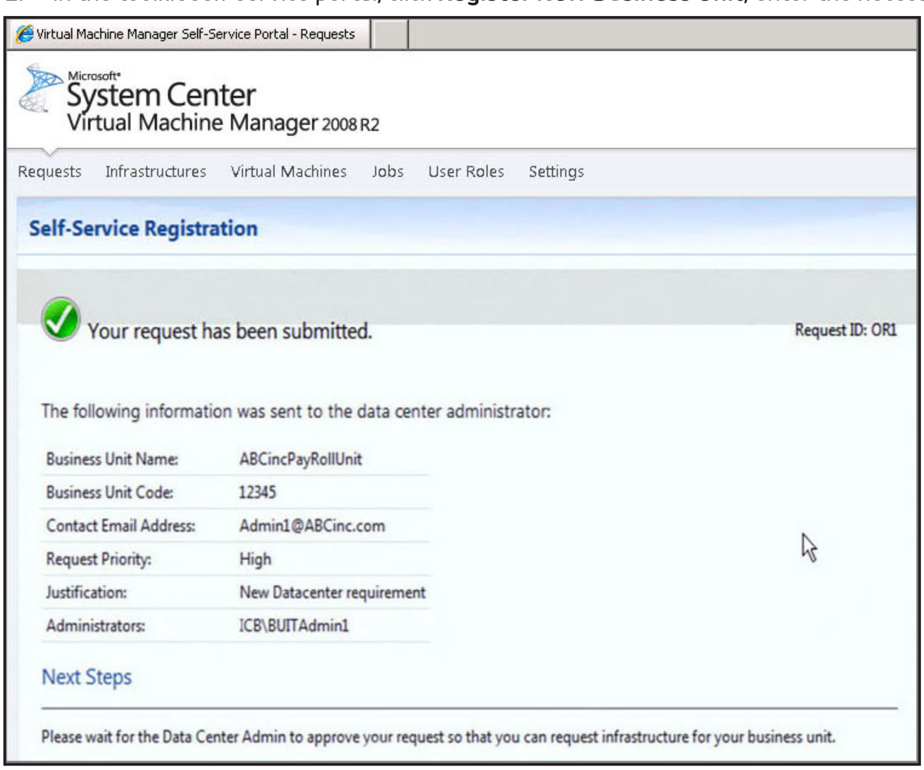


**Figure 6: Self-Service Registration**

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

### 3. Create Users/View Users

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

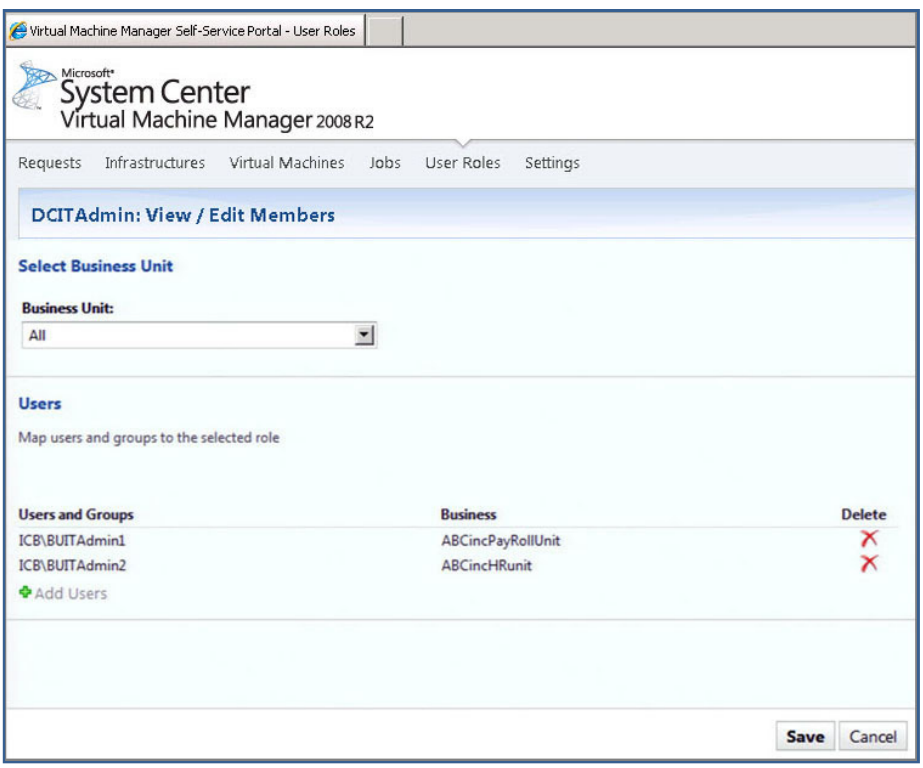


Figure 7: Add/Edit Users

## 4. Configure the Virtual Machine Templates

1. Log on to the toolkit Web portal as DCITAdmin1.
2. Click the **Settings** tab, and then click **Configure Templates**.
3. On the **Settings: Virtual Machine Template** page, under **Other Tasks**, click **Add Templates**.
4. On the **Add Virtual Machine Templates** page, select the library server and library shares from the drop-down boxes, and then click **Search**.

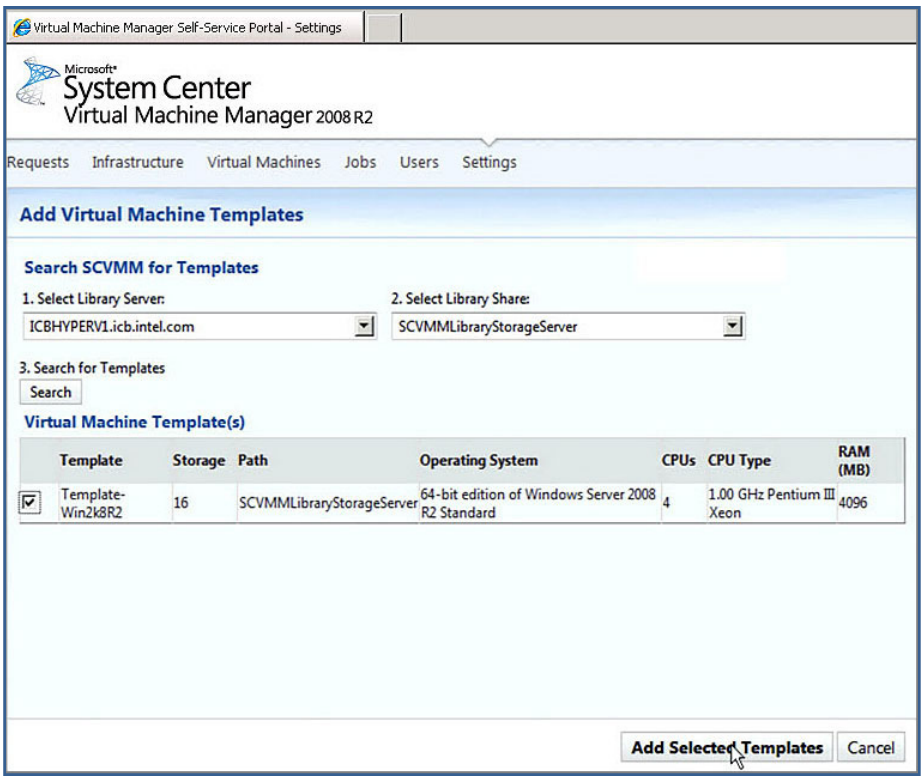


Figure 8: Locate the Virtual Machine Template

5. Select one or more templates that need to be made available to the toolkit and then click **Add Selected Templates**.

   The selected templates are added to the list.
6. Configure the cost for the template and then click **Save and Close**.

   The template is saved and made available to the toolkit.

7. Open the **Settings: Virtual Machine Template** page to verify that the template is listed.

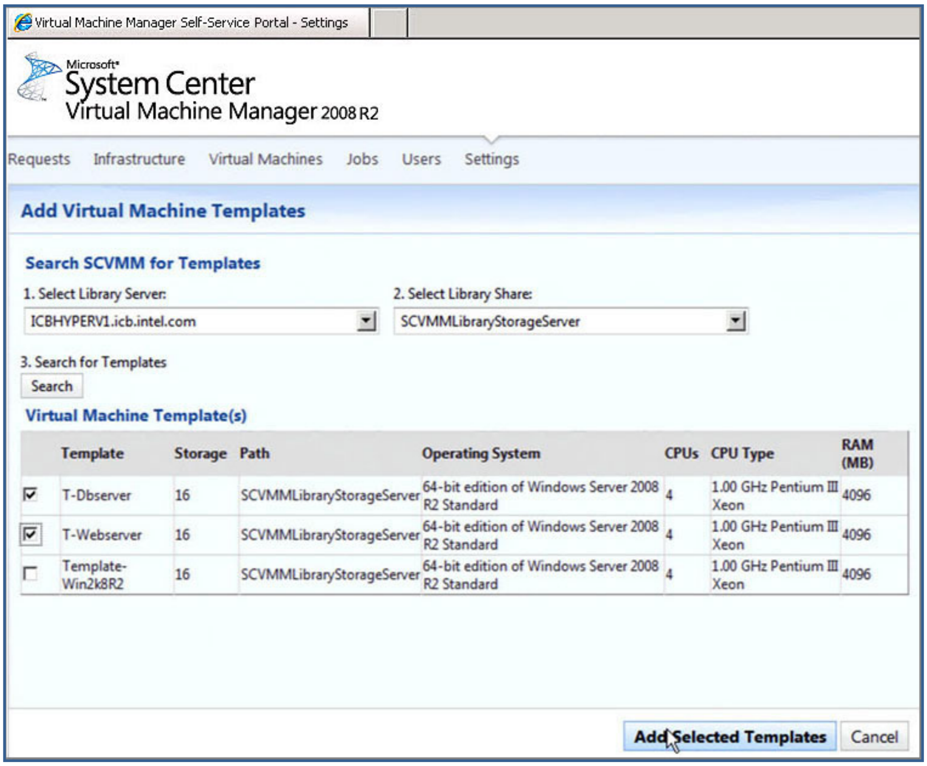


**Figure 9: Verify that the template is added**

## 5. Configure the Network

1. Log on to the toolkit Web portal as DCITAdmin1.
2. On the **Settings** tab, under **Datacenter Management**, click **Configure Datacenter Management** link.
3. In the left column, click **Network**, and then 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.
7. In the **Start IP Address** text box, specify the lower limit of IP addresses for virtual machines in the environment, and then in the **End IP Address** text box, specify the upper limit of IP addresses.
8. Fill in the **Subnet Mask** text box, and then click **Done**.
9. Click **Save and Close**.

10. On the **Settings** tab, verify that the newly created network is present.

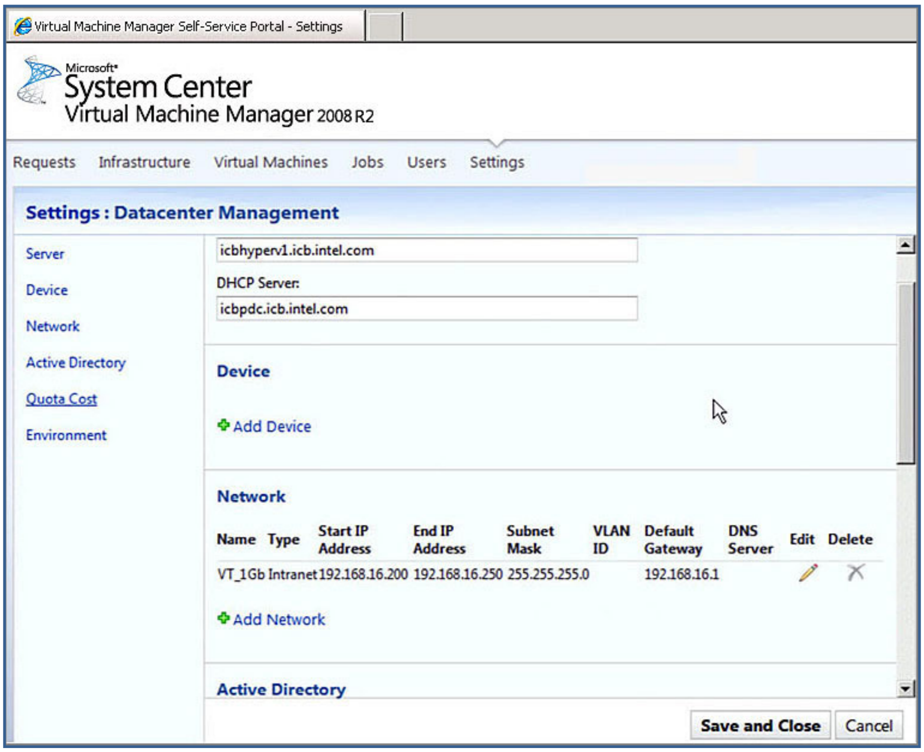


**Figure 10: Configure the Network**

## 6. Create the Infrastructure

1. Log on to the toolkit Web portal as a BUIT Admin of the business unit registered previously.
2. On the **Requests** tab, click the business unit request approved previously, and then click **Create infrastructure request**.

3. On the **New Infrastructure Request** page, on the **Basic Information** tab, fill in the boxes with the appropriate information, and then click **Next**.

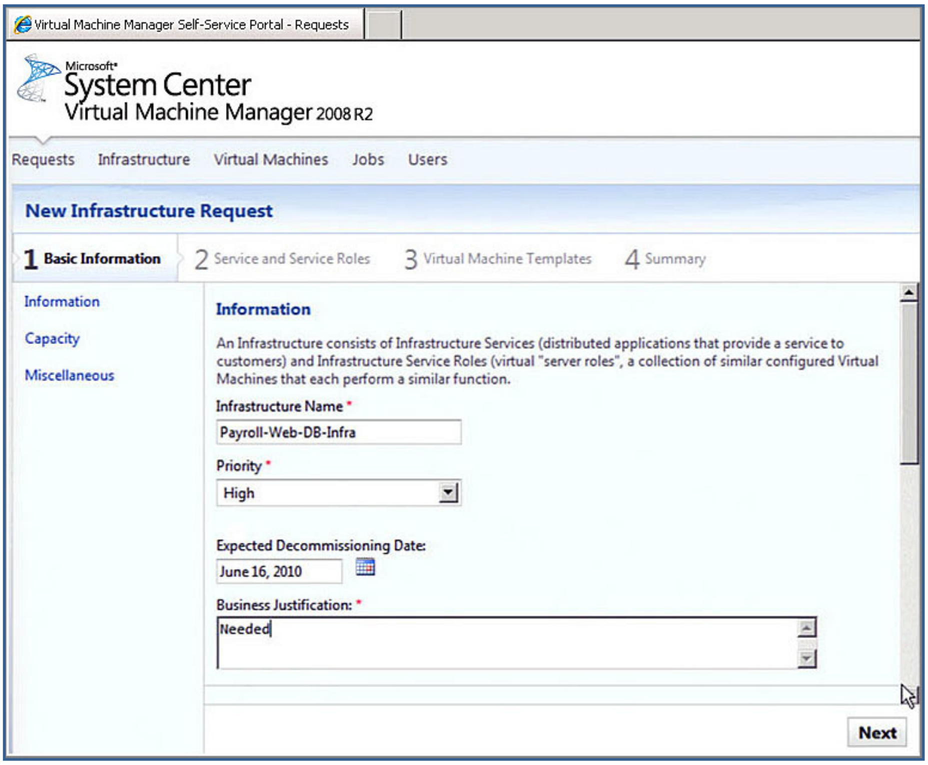


Figure 11: Create New Infrastructure Request

4. On the **Service and Service Roles** tab, fill in the boxes with the appropriate information, and then click **Add Service Role**.
5. In the **Add Service Role** dialog box, enter the appropriate information, and then click **Save**.
   The **Add Service Role** dialog box closes and the **Service and Service Roles** tab opens. All of the information entered previously is retained. The service role is added successfully to the service.
6. Click **Next**.

7. On the **Virtual Machine Template** tab, select a pre-configured template and associate a cost to it.

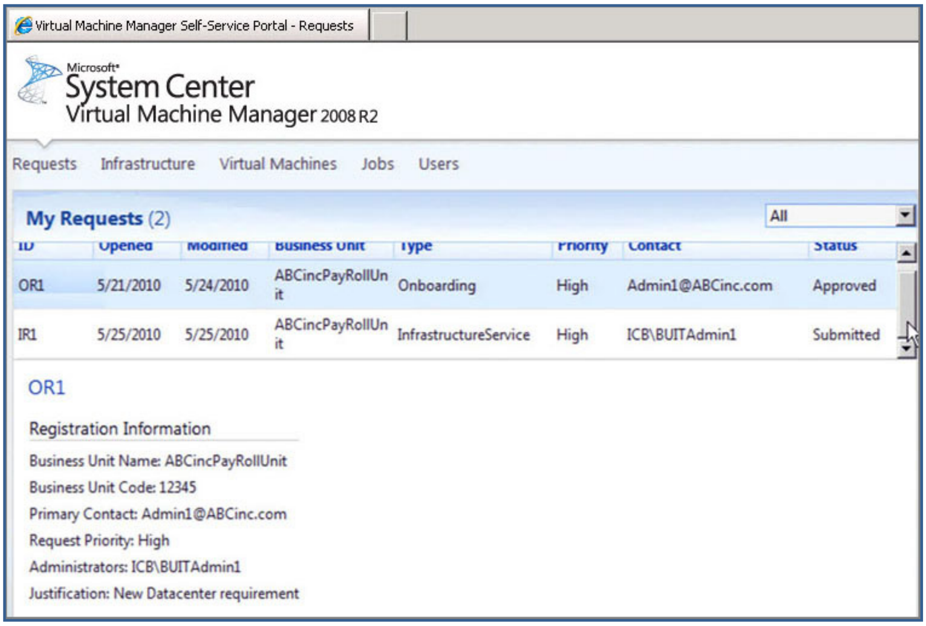


**Figure 12: Summary Page**

8. On the **Summary** page, make sure that all of the entered information is correct, and then click **Submit**.
9. Log on to portal as DCITAdmin1.
10. Approve the submitted request to create infrastructure.
11. Verify that the infrastructure is created successfully.

## 7. Approve an Infrastructure Request

1. Log on to 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 such as host group, library server location, and information to store offline virtual machines.
5. Click the pencil icon next to each service role and edit the service role information.
   You can now approve the infrastructure request.
6. On the **Virtual Machine Template** tab, select a pre-configured template and associate a cost to it.
7. On the **Summary** page, verify that all of the entered information is correct and then click **Submit**.

8. Click the **Infrastructure** tab and verify that the approved infrastructure is listed under **My Infrastructure**.

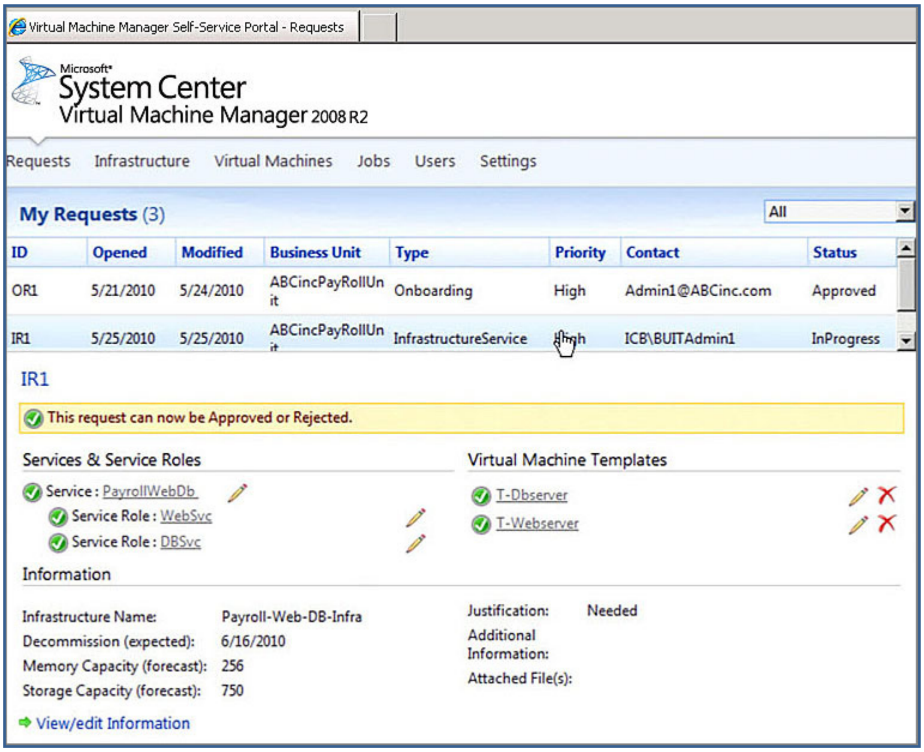


**Figure 13: Approve/Reject Requests**

## 8. Create Multiple Virtual Machines

1. Log on to 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:

- **Option**: With a single name and sequential number
- **Number of virtual machines to create**: 16
- **Computer Name**: DBServer_
- **Index suffix start**: 001
- **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
- **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**.
   The Create Virtual Machine request is submitted successfully. The toolkit creates corresponding jobs to carry out the operation.
6. On the **Jobs** tab, monitor the jobs to verify that they are created for each virtual machine.

7. Go to the host computer to verify that the virtual machines have been created successfully.

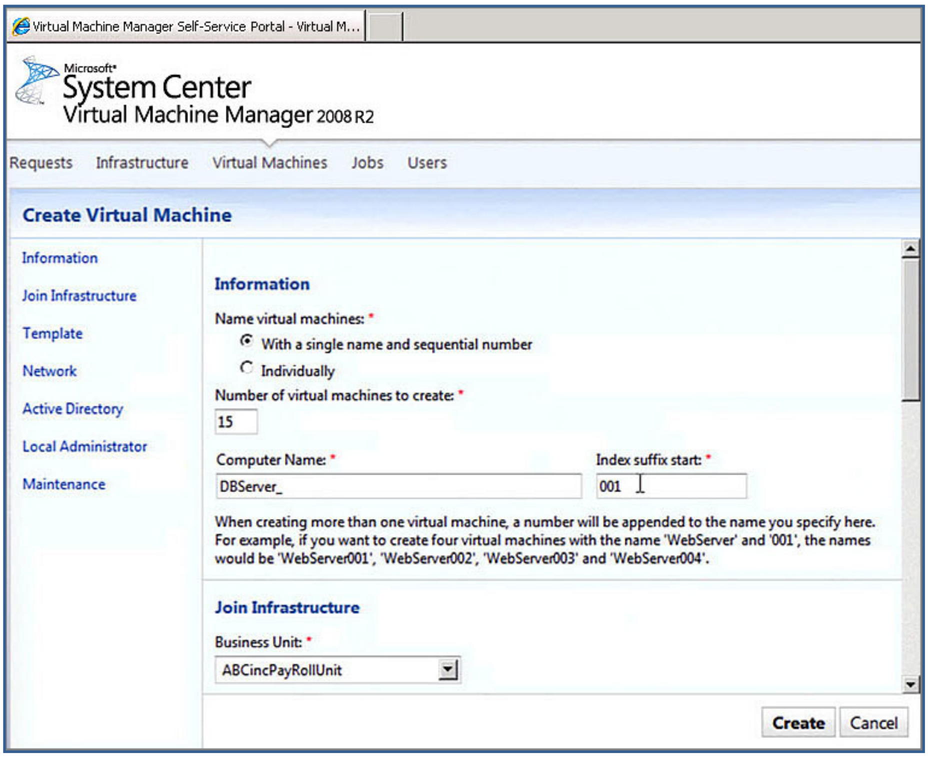


Figure 14: Create Virtual Machines

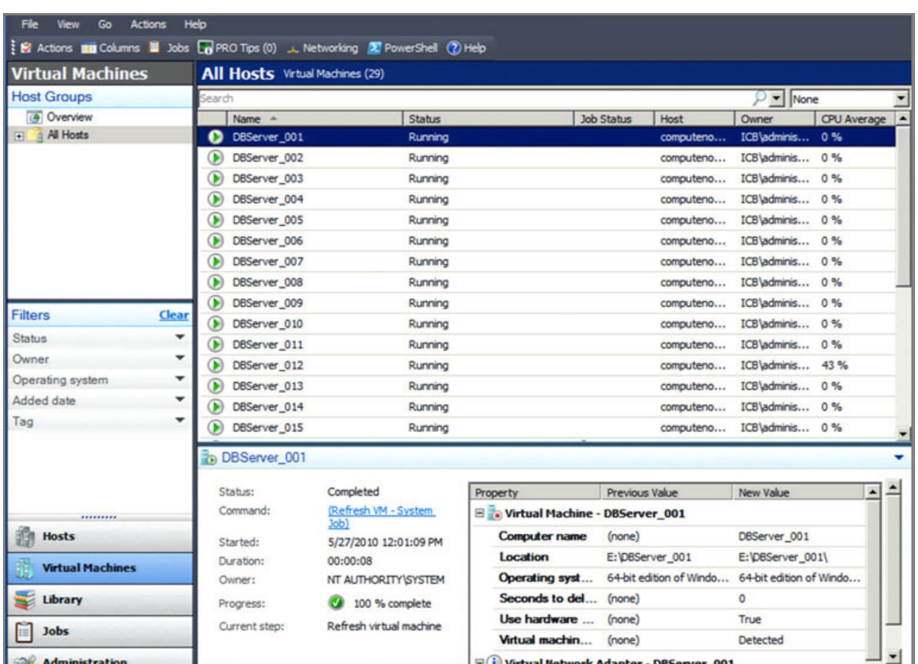


Figure 15: View Virtual Machines

## 9. Access a Virtual Machine Using the Web Portal

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

## 10. Delete Virtual Machines Using the Web Portal

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

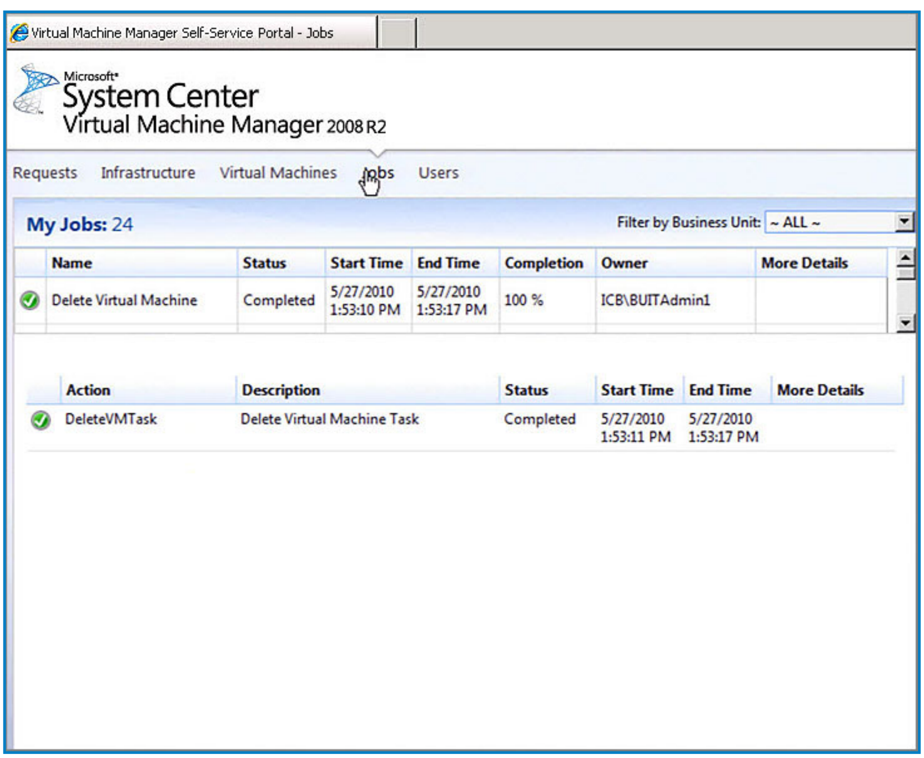


Figure 16: Delete Virtual Machines

## 11. Submit a Change Request for Additional Storage

1. Log on to the toolkit Web portal as BUITAdmin1.
2. Select the infrastructure on which the change request is to be submitted.

3. Click **New Change Request**.
4. In the **Object to Change** drop down list, select **Infrastructure**.
5. In the **Type of Change** drop down list, select **Edit**.
6. Enter the amount of storage required and click **Next**.
7. Click **Submit Request**.
8. Log on to the toolkit Web portal as DCITAdmin1.
9. Approve the request to decommission the infrastructure.
10. Verify that requested storage is available.

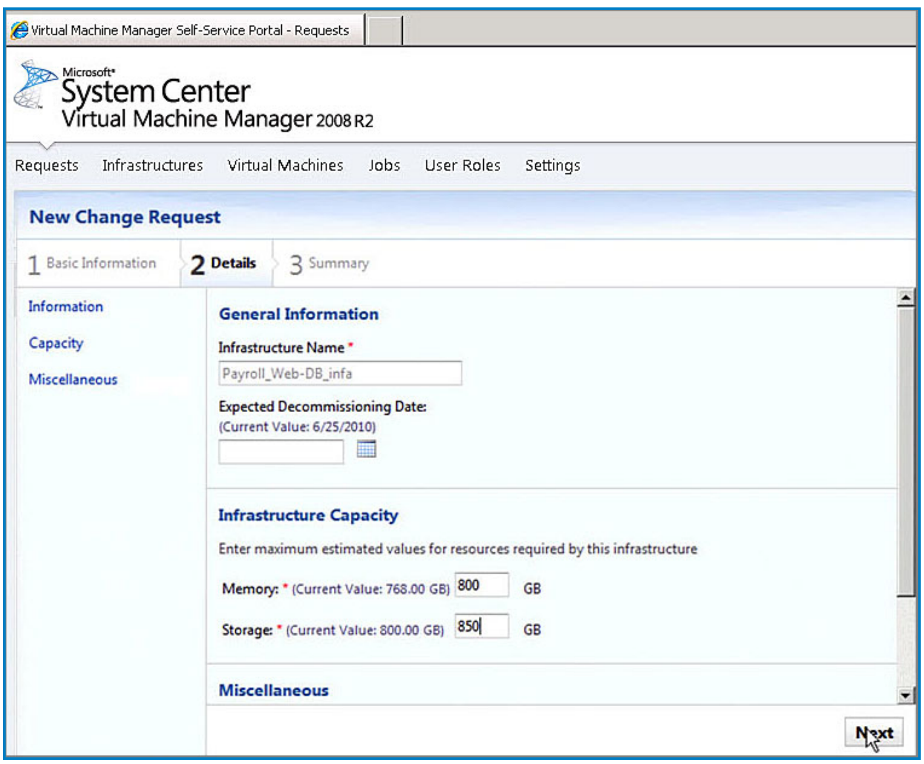


**Figure 17: Change Request**

## 12. Decommission an Infrastructure

1. Log on to the toolkit Web portal as BUITAdmin1.
2. Click the **Infrastructure** tab.
3. Click **New Change Request**.
4. In the **Object to Change** drop-down list, select **Infrastructure**.
5. In the **Type of Change** drop-down list, select **Decommission**.
6. Click **Next**.
7. Click **Submit Request**.
   The decommission infrastructure request will be created for the selected infrastructure.
8. Log on to the toolkit Web portal as DCITAdmin1.

9. Approve the request to decommission the infrastructure.
10. Verify that the infrastructure has been decommissioned.

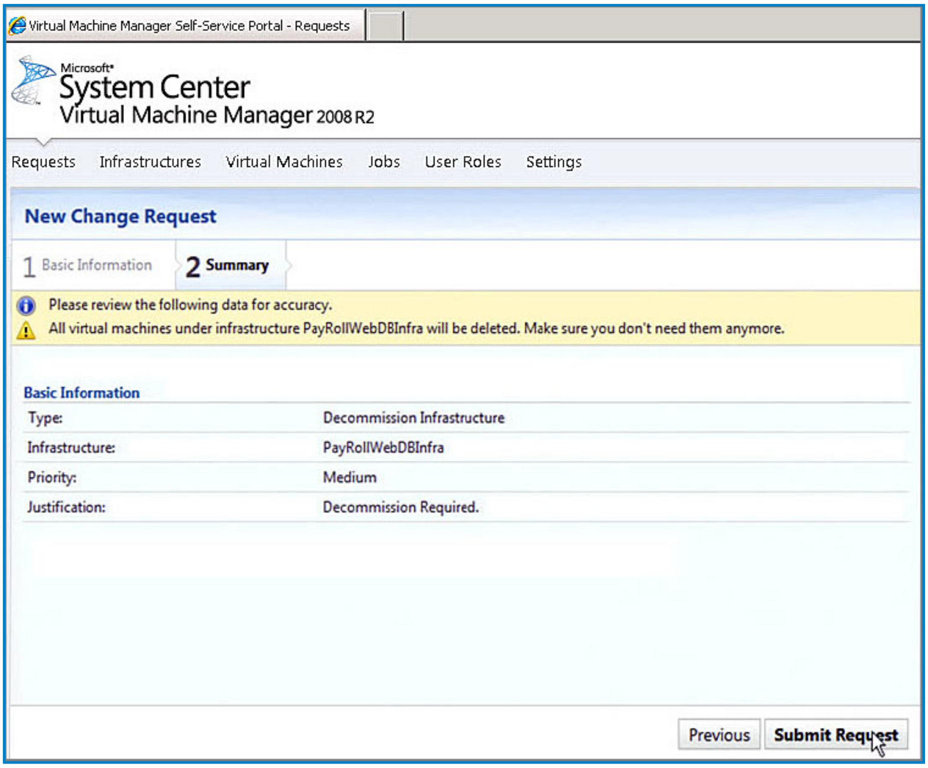


Figure 18: Decommission Infrastructure

## 13. Compare Performance between Intel Xeon X5570 and Intel Xeon X5670 processors

1. Connect to VMMSSP and ensure that the virtual machines created have the Java workload installed and running.
2. Click **Hosts** in the navigation pane.
3. Click the hostgroup and make sure that it includes only hosts with Xeon X5570 processors.
4. Log on to the Self-Service Portal of VMMSSP.
5. Start the Java workload benchmark on all of the virtual machines.
6. Record the number of transactions completed by the workload.
7. Connect to System Center Virtual Machine Manager.
8. In the navigation pane, click **Hosts**.
9. Click the hostgroup and make sure that the host group only includes hosts with Xeon X5670 processors.
10. Log on to the toolkit Web portal.
11. Start the Java workload benchmark on all virtual machines.
12. Record the number of transactions completed by the workload.

Table 3 shows the Java workload performance for the two processors:

| Compute Node1 | Xeon 5570 | Xeon 5670 |
|---|---|---|
| VM1 | 87,175 | 91,535 |
| VM2 | 78,887 | 106,869 |
| VM3 | 83,363 | 106,135 |
| VM4 | 66,478 | 68,628 |
| Total Business Ops/s | 315,903 | 373,167 |

| | |
|---|---|
| Performance Gain Xeon 5670/5570 | 18.13% |

**Table 3: Java Workload Performance**

## 14. Apply a Power Capping Policy on Compute Nodes in a Cloud Environment

1. Start the Java workload on all four virtual machines on each of the compute nodes.
2. Note the aggregate power consumption in each of the compute nodes using Intel DCM UI.
3. Apply the power capping policy at 200W through the Intel DCM UI.
4. Note the power consumption at maximum capping.

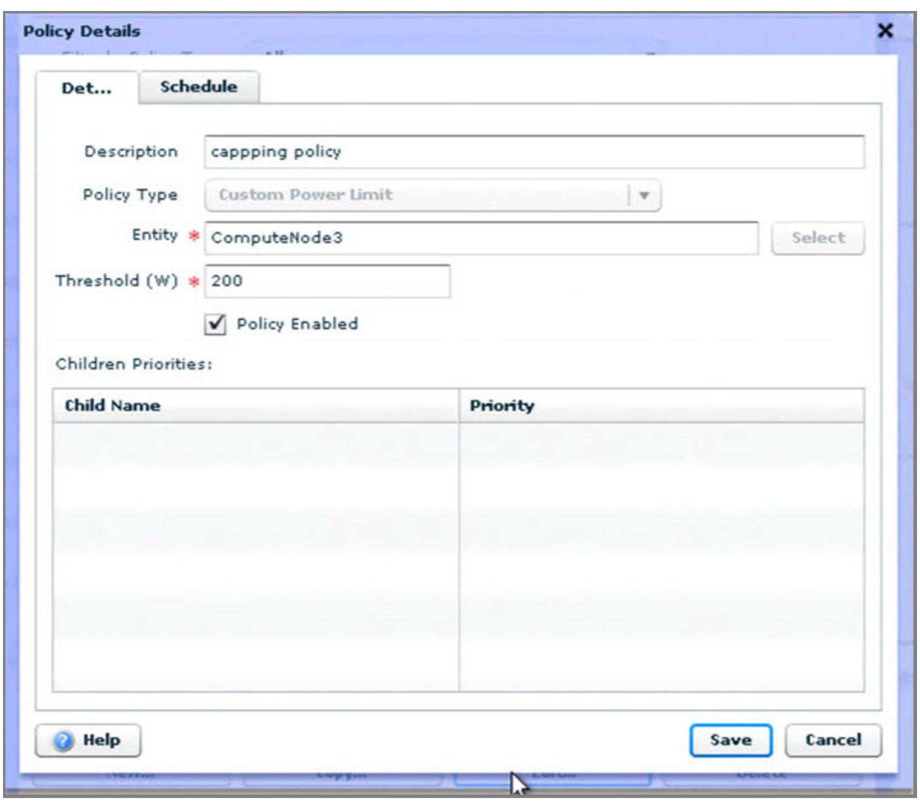


**Figure 19: Power Capping Policy**

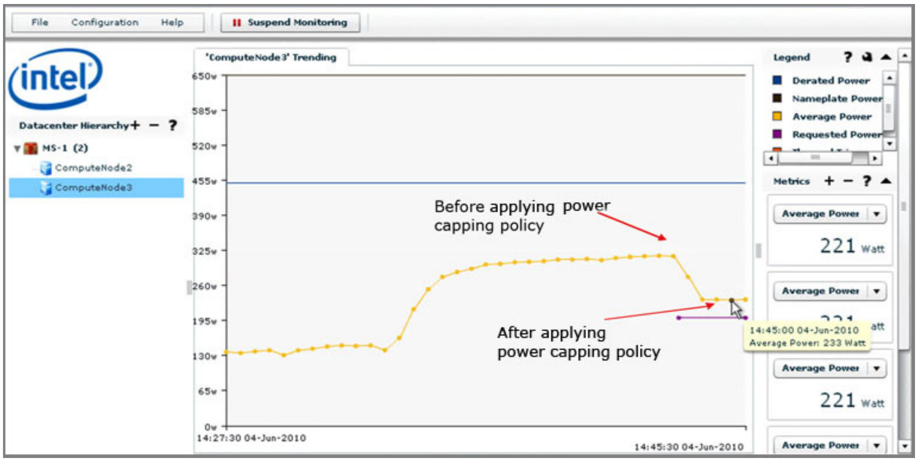

Figure 20: Monitor Power Consumption

## 15. Note Dynamic Power Reconfiguration

Note the power consumption at different stages from power up to full load and measure the time it takes for a compute node:

- To go from hibernation to full load state
- To go from full load to full stop

1. Start all of the compute nodes.
2. Start a Java workload on all virtual machines for each of the compute nodes.
3. Record the power consumption graph from the DCM UI.
4. Record the time it takes to bring the cluster to full load.
5. Note the range of power consumption for the cluster from zero workload to maximum workload.
6. Terminate all workloads, sign off each tenant as quickly as possible, and shut down all of the nodes.

7. Record the time it takes to bring the cluster to full stop.

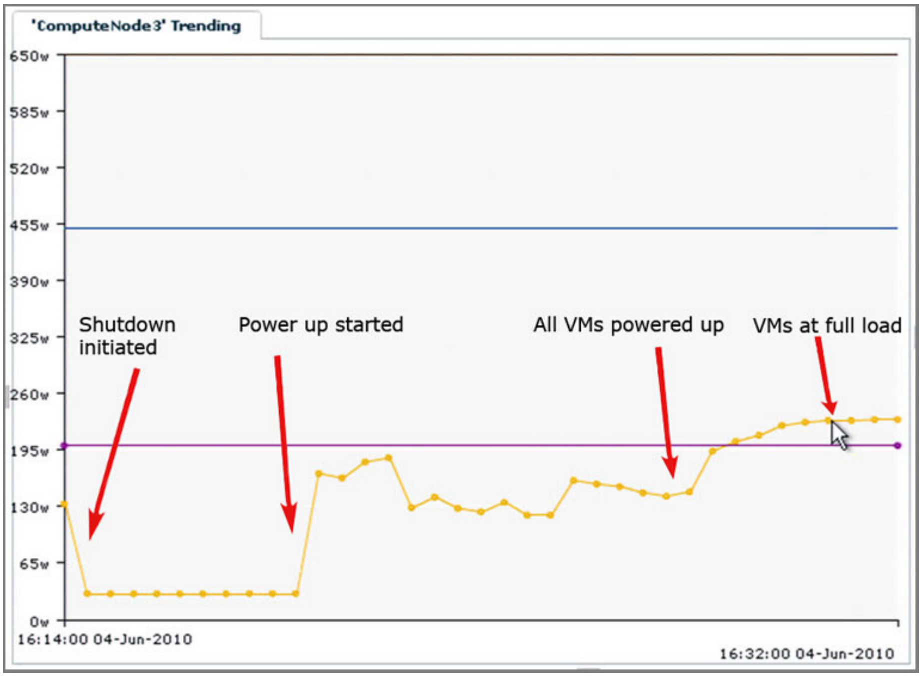


**Figure 21: Dynamic Reconfiguration Results**

## Things to Consider

### Storage

We chose to use a storage iSCSI for the bulk storage requirements in our cloud. We could just as well have connected the storage server to a Storage Area Network (SAN). The choice in this case was completely arbitrary and based solely on the capabilities in our 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

### Use of SSD Drives

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 within our exercise.

### Networking

VMMSSP supports several network technology architectures. For this test bed, we selected a very simple network topology. The selection made for this test bed performed well for our purposes, but more advanced technologies and architectures (e.g. 10 GbE, FCoE) will be more suitable for production deployments, especially when performance is under consideration.

## Conclusion

In this paper, we have described the components and implementation process for a private cloud, 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 we created has the standard functionality expected from any cloud system, such as the ability to provision and destroy virtual machines as needed. We illustrated the VMMSSP functionality with several real world use cases when we defined the resource pools, we created multiple virtual machines, and we decommissioned the infrastructure. We also describe possible enhancements that could improve performance and scalability of the prototype described in this paper, which would make the system more suitable for real production environments.

## Additional Information

Intel Cloud Builder: http://intel.com/software/cloudbuilder

Intel Xeon processors: http://intel.com/xeon

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 paper 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 you can consolidate more applications and heavier workloads per server to get better value from your 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® 5500 Series Software Industry Testimonials, http://www.intel.com/business/software/testimonials/xeon5500.htm and Intel Xeon 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 5500 Series Software Industry Testimonials, http://www.intel.com/business/software/testimonials/xeon5500.htm and Intel Xeon 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 www.intel.com/software/cloudbuilder

## Disclaimers



Printed in USA 0910/TC/PRW/PDF Please Recycle 324327-001 US