Intel® Cloud Builders Guide
Intel® Xeon® Processor-based Servers
IBM CloudBurst* Solution

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

IBM CloudBurst* Solution

Intel® Xeon® Processor 5600 Series

AUDIENCE AND PURPOSE

Cloud computing offers a path to greater scalability and lower costs for service providers, infrastructure hosting companies, and large enterprises. Establishing an infrastructure that can provide such capabilities requires experience. Intel has teamed up with leading cloud vendors through the Intel® Cloud Builders program to help any customer design, deploy, and manage a cloud infrastructure.

The Intel Cloud Builders program provides a starting point by supplying a basic hardware blueprint and available cloud software management solutions such as the IBM CloudBurst*. The use cases described in this paper can be used as a baseline to build more complex usage and deployment models to suit specific customer needs. The audience for this paper is cloud service providers, cloud hosters and enterprise IT that are looking to realize the revenue potential of their existing data center infrastructure by offering cloud computing services to their customers or internal users.

Table of Contents

Executive Summary

A traditional enterprise tends to pull together resources and deploy them in support of a business function workload, one project at a time, or in silos. The resources are dedicated to the workload and are unable to support other workloads where they could be leveraged as added support. Cloud computing, on the other hand, leverages a "pooled resources" environment that uses virtualization so the physical assets can support multiple workloads. The assets (hardware, software, and delivery) must be standardized and automated so they can be delivered efficiently, and so they can enable the self-service, self-management of cloud computing. This is what delivers a responsive end-user experience. For the end-user, a cloud is elastic in scalability and is accessible from any device, anywhere, anytime. And if users are charged, they are required to pay only for what they use when they are using it. From a provider's perspective, it's about an environment of highly virtualized resources that are location-independent and have automated service management to handle provisioning, de-provisioning, change management, security, and overall environment controls.

The ideas and principles of the cloud are of great interest because of its agility benefits and cost savings. In recent years, many companies have set up cloud infrastructures that can be accessed over the public Internet and are offering services that customers can use to host their applications. Typical cloud services can be classified as Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS) and Software-as-a-Service (SaaS).

A cloud is based on three fundamentals: standardization, virtualization, and automation. These factors enable the cloud to provide scalable and efficient services and help consumers drive down costs and improve service. Clients are using the cloud to achieve real, measurable business results.

An IaaS structure is the best choice for pre-existing workloads, for workloads unique to an organization, and for workloads that need to remain under close control of an organization. Starting from a proven blueprint is key to effectively supporting and hosting an IaaS cloud service, whether as a revenue-producing, customer-facing service (a public cloud), or an internal one behind the corporate firewall (a private cloud).

This reference architecture summarizes the process of building a cloud using servers based on the Intel® Xeon® processor 5600 series (codenamed Westmere-EP) and IBM CloudBurst software, to provide a starting point to building a cloud. This paper provides a high-level overview of the IBM CloudBurst solution and supplements any technical documentation that may be available at the IBM CloudBurst website [2].

Introduction

Acquiring IT as a form of service can follow one of three models:

- Monitoring servers
- Hosting and managed services
- Outsourcing

Cloud computing is a new delivery and consumption model for each of these. A company can build out its own private cloud environment, purchase a pre-integrated system designed to deliver a cloud-delivered workload, or acquire a cloud-delivered workload as a service from a third party.

IBM offers technologies—hardware and software—as well as key services to help you and your business benefit from cloud computing. The best way to make use of cloud computing as a viable delivery model is to lay out a clear path based on integrating the cloud platform with your overall IT delivery strategy.

IBM CloudBurst*: An Integrated Cloud Building-Block Platform

IBM CloudBurst is a self-service provisioning purpose-built system that combines IBM System x* hardware, with its management software (e.g., IBM Systems Director*); IBM Storage; VMware vSphere* 4.1; and IBM Tivoli* software packages (see Figure 1). IBM CloudBurst offers a set of unique capabilities for organizations to quick-start their cloud implementation:

- **Self-contained on-premise cloud** includes prepackaged hardware, software, and services based on the IBM System x BladeCenter* platform and Tivoli Service Management products.
- **Web 2.0 self-service portal** provides automated request and (de-)provisioning of production or development/test workloads using virtualization technologies across the server, network, and storage, including reservation of compute and storage resources.
- **Pre-packaged automation templates and workflows** for most common resource types, including VMware virtual machines (provisioned to capabilities).
- **Integrated core service management capabilities** include real-time monitoring of virtualized resources, energy management, (de-)provisioning, patch management and remediation, security, usage and accounting, reusable library for rapid deployment, and pre-built reports (BIRT).
- **Modular/plug and play**: Incrementally, automatically expandable and scalable.
- **Multi-tenant**: Management of multi-customer, multi-project collections of virtual systems.
- **Quick-start implementation** services to get the cloud platform up and running in days.
- **Extensibility** across the data center with Tivoli System Automation Manager (TSAM) integration.

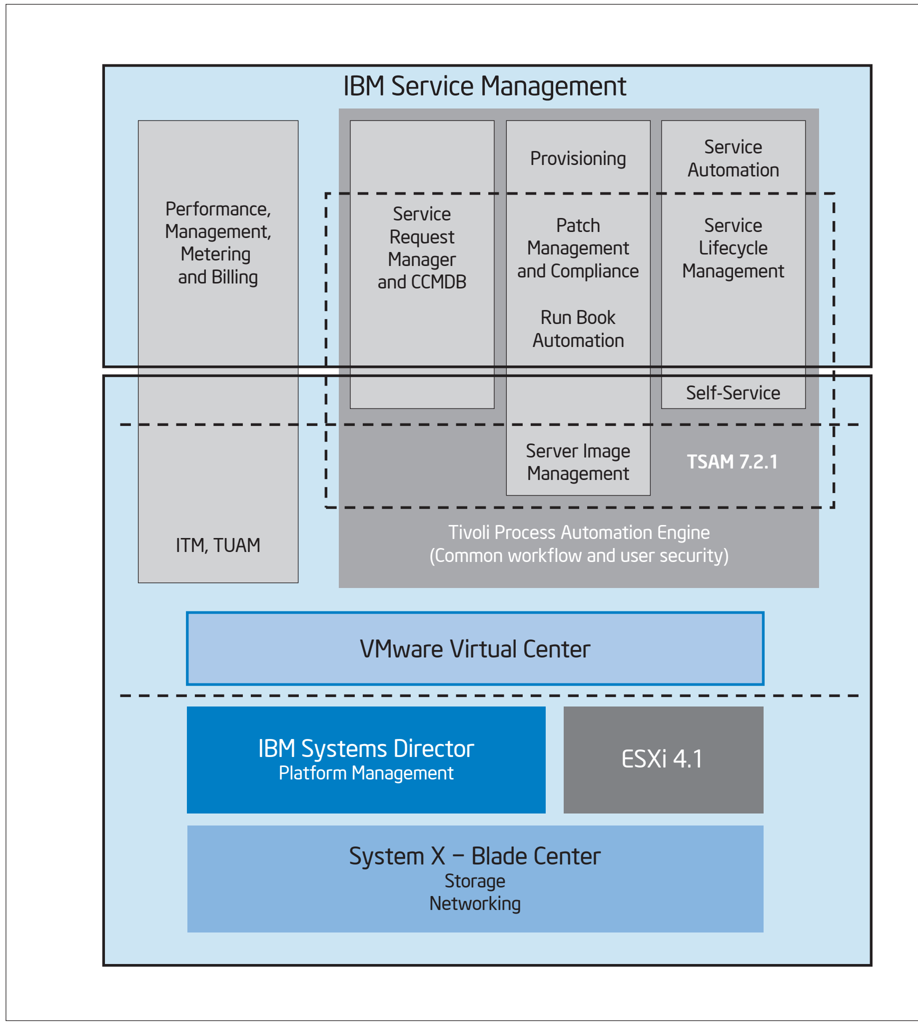

Figure 1. IBM CloudBurst Reference Architecture

- **Enablement for integrating IBM WebSphere* CloudBurst** outside-the-box.
- **Usage and accounting** provides metering and accounting for cloud services, and integration to billing systems.
- **Virtualized hardware management** provides enhanced management of the virtual environment.
- **Energy monitoring** of the hardware infrastructure.
- **Tivoli Service Automation Manager (TSAM) includes:**
  - Orchestration of cloud operations
  - Integration point for service-management capabilities
  - Service catalog and templates
  - Automated provisioning of virtual systems
- **Monitoring** of both physical and virtual server environments.
- **High Availability:** Make management system DB highly available.
- **Server, storage, and network hardware are** preinstalled and configured on IBM hardware.

IBM CloudBurst Solution Overview

IBM CloudBurst v2.1 is a prepackaged and self-contained service delivery platform that can be easily and quickly implemented in a data-center environment. It allows the data center to accelerate the creation of service platforms for multiple workload types that are highly integrated, offer flexibility, and optimize resources. These platforms provide an enhanced request-driven user experience, while helping drive down costs and accelerating time to value for the business. As a cloud computing quick-start, IBM CloudBurst allows organizations to prove the benefits of this delivery model in a defined portion of their data center or for a specific internal project.

IBM CloudBurst combines the necessary hardware, software, and services components to rapidly implement cloud computing. As a single solution, IBM CloudBurst de-mystifies the intricacies of implementing a cloud computing model and can help enable organizations to quickly realize the benefits and business potential of a dynamic infrastructure (see Figure 2). Cloud computing is a services acquisition and delivery model for IT resources that can help improve business performance and help control the costs of delivering IT resources to an organization.

Built on the IBM System x BladeCenter platform, IBM CloudBurst v2.1 focuses on expanding existing core service management capabilities into production clouds. IBM CloudBurst includes:

- IBM Tivoli Service Automation Manager (TSAM) v7.2 for greater ease of use:
  - Enhanced Web 2.0 image management interface
  - Resource reservation capabilities, allowing environments to be scheduled
  - Delivery of integrated IBM Tivoli Usage and Accounting Manager to enable chargeback for cloud services, to optimize system usage
  - Enables the creation of resource usage and accounting data to feed into Tivoli Usage and Accounting Manager, which allows for tracking, planning, budgeting, and chargeback of system resource usage
- High-availability option using Tivoli Systems Automation Manager and VMware High Availability that provides protection against unplanned management system outages and that can simplify virtual machine mobility during planned changes.
- Integration with Tivoli Monitoring for Energy Management to monitor and manage energy usage of IT and facility resources. This helps optimize energy consumption for a more efficient use of resources, to help lower operating cost.
- The ability to manage other heterogeneous resources outside of the IBM CloudBurst environment (requires purchase of additional licenses).

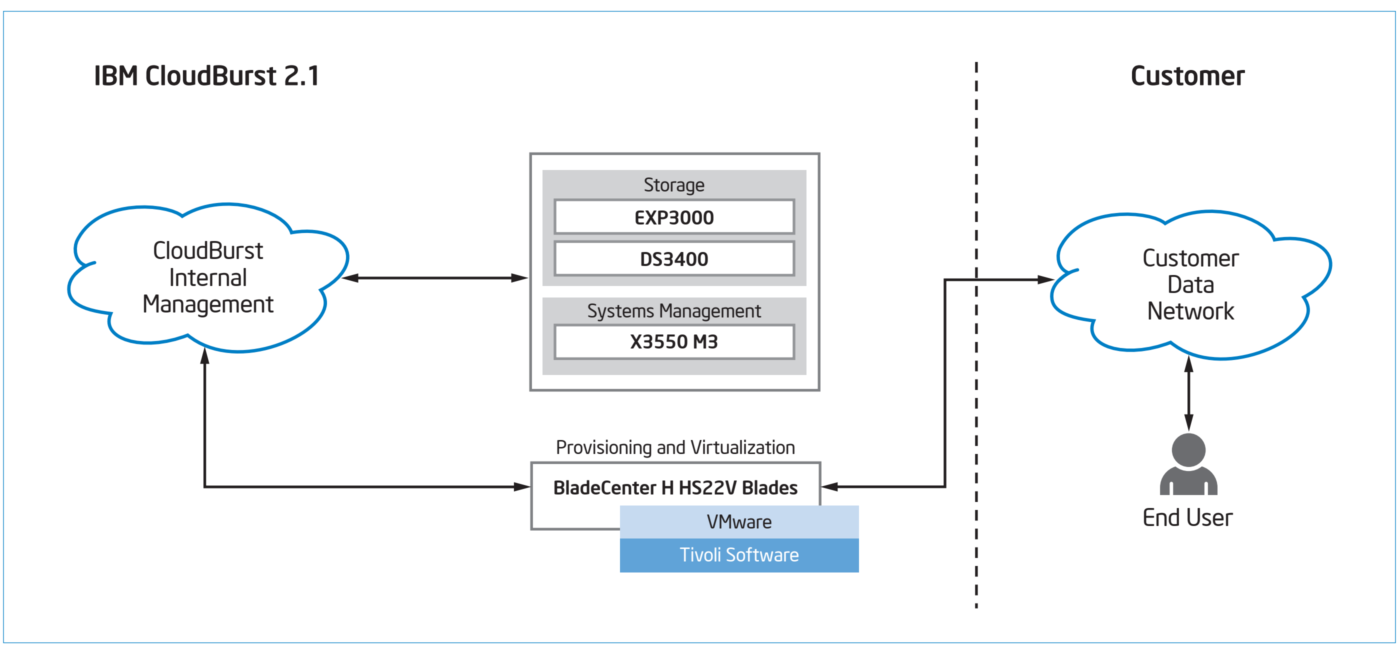

Figure 2. IBM CloudBurst* Conceptual View

Hardware Configuration

The IBM CloudBurst 2.1 appliance is orderable in four configuration sizes: small, medium, large and x-large. Each configuration size has a unique preconfigured template.

- The small configuration is a **single BladeCenter H** chassis with a maximum of **four** blades and one **DS3400** storage controller. This is the minimum configuration size.
- The medium configuration is also a **single BladeCenter H** chassis, but with **five** to **14** blades. As the blade quantity increases, additional **EXP3000** storage expansions are added to provide additional VM storage.
- The large configuration includes **two BladeCenter H** chassis, with **15** to **28** blades, and two DS3400 storage controllers. Additional **EXP3000** storage expansions are added to accommodate the additional blades. **SAN24B-4** Express Fibre Channel switches are also added.
- The xlarge configuration includes up to four **BladeCenter H** chassis, with **29** to **56** blades, and four DS3400 storage controllers. Additional **EXP3000** storage expansions are added to accommodate the additional compute capacity. In addition to the **SAN24B-4** Fibre Channel switches, BNT G8124 10Gb Ethernet switches are also added.

Detailed Hardware Configuration

The IBM CloudBurst configurations outlined above are based on the following hardware components:

Rack-level Components

- One 42U rack
- One System x BladeCenter H chassis
- Redundant 8G FC networking
- Redundant 1G Ethernet networking
- One DS3400 dual controller
- Twelve 600 GB 15K drives
- One x3550 M3 management server
- One to two HS22V (for HA) blade(s) for IBM CloudBurst software
- Three HS22V provisioning blades

HS22V Blade

- Dual Six-Core Intel® Xeon® processor X5660: 6C, 2.8 GHz, 12 MB L2 cache, 95 W
- 72 GB memory
- Diskless HDD
- VMware ESXi* 4.1 Hypervisor on embedded USB key
- Two 10 Gbps network interface controllers (NICs)
- Dual-port, 8 Gbps Fibre Channel (FC)

x3550 M3 Management Node

- Dual Quad-Core Intel Xeon processor X5620: 2.4 GHz, 12 MB L2 cache, 80 W
- 24 GB memory
- HDD: Six 300 GB SAS 10K (arranged as three RAID1 pairs)
- Four 1Gbps network interface controllers (NICs)
- Dual port, 4Gbps Fibre Channel (FC)
- Remote presence capability

Processor Description

IBM CloudBurst blade servers support Intel® Xeon® processor 5600 series (4 or 6 cores per processor) [6]. The following sections describe key Intel® processor technologies that enable implementation of virtualized, multi-tenant data centers:

Intel® Turbo Boost Technology§ allows processors to deliver higher-speed execution on demand by using available power to run at a higher frequency.

Intel® Hyper-Threading Technology†: Many server applications require parallel, multi-threaded execution. Intel® Hyper-Threading Technology enables simultaneous multi-threading within each processor core, up to two threads per core. Hyper-threading reduces computational latency, making optimal use of every clock cycle. For example, while one thread is waiting for a result or event, another thread is executing in that core, to maximize the work from each clock cycle.

Intel® QuickPath Technology is a scalable, shared-memory architecture that delivers a high-memory bandwidth to enable top performance for bandwidth-intensive applications. It provides high-speed point-to-point connections between processors, and between processors and the I/O hub. Each processor has its own dedicated memory that it accesses directly through an integrated memory controller. If a processor needs to access the dedicated memory of another processor, it can do so through a high-speed Intel QuickPath Interconnect (Intel QPI) that links all the processors.

Intel® Intelligent Power Technology: Within a single server, Intel® Intelligent Power Technology minimizes power consumption when server components are not fully utilized. Integrated power gates allow individual idling cores to be reduced to near-zero power independent of other operating cores. Automated low-power states automatically put processor and memory into the lowest available power states that will meet the requirements of the current workload. Processors are enhanced with more and lower CPU power states, and the memory and I/O controllers have new power management features.

Intel® Virtualization Technology◊ **(Intel® VT)** enhances virtualization performance with new hardware-assist capabilities across the following elements of the server:

- **Processor (Intel® VT-x):** Intel VT-x provides hardware-assisted page-table management, allowing a guest OS more direct access to the hardware and reducing compute-intensive software translation from the VMM. Intel VT-x also includes Intel® VT FlexMigration and Intel® VT FlexPriority, which are capabilities for flexible workload migration and performance optimization across the full range of 32-bit and 64-bit operating environments.

- **Chipset (Intel® VT-d):** Intel VT-d helps speed data movement and eliminates much of the performance overhead by giving designated virtual machines their own dedicated I/O devices, thus reducing the overhead of the VMM in managing I/O traffic.

Intel® Advanced Encryption Standard (AES) New Instructions (Intel® AES-NI) enables robust encryption without additional appliances and increased performance overhead. This technology improves CPU performance for encryption by as much as up to 52 percent for secure internet transactions and allows broader use of encryption throughout the data center. [4]

Software Configuration

The IBM CloudBurst solution consists of the following software components on various servers in the solution (see Figure 3).

Management Stack on the x3550 M3

The following software utilities are included for IBM CloudBurst 2.1 hardware management:

- Microsoft Windows* Server 2008 R2 Standard Edition, 64-bit
- IBM Systems Director* v6.2
  - Active Energy Manager v4.2
  - Network Control v1.2
- LSI "Eagle 2" SMI-S
- vSphere Client 4 64-bit
- IBM DS Storage Manager
- ToolsCenter
  - Bootable Media Creator v2.0
  - Advanced Setup Utility v3.2
  - Dynamic Systems Analysis v3.0
  - UpdateXpress System Pack Installer v4.0
- IBM DB2* v9.5 FP5 64-bit for IBM Systems Director
- IBM Tivoli Monitoring (ITM) for Energy Management 6.2.1.01
- IBM Java 6.0
- Advanced Management Module (BladeCenter Open Fabric Manager (BOFM) Basic is included)
- As an option, IBM System Storage SAN Volume Controller (SVC) with Console v5.1.0

CloudBurst Stack on HS22V Management Blade

The following software utilities are installed on the blades:

- Tivoli Service Automation Manager (TSAM) 7.2.1 (including TPM, service catalog, TPAE)
- IBM Tivoli Monitoring 6.2.2 [9], TEMS plus OS agents, Green Energy Manager
- IBM Tivoli Usage and Accounting Manager (ITUAM) 7.1.2
- TSA
- Middleware: DB2 ESE 9.5 FP3; WebSphere Application Server ND 6.1.0
- x86 OS: SUSE Enterprise Linux* 10, IBM Java 6.0
- VMware ESXi 4.0u1 hypervisor on each blade
- VMware VirtualCenter* 4.0 running in a VM on the Management Blade (Windows Server 2008 R2 Standard, 64-bit)

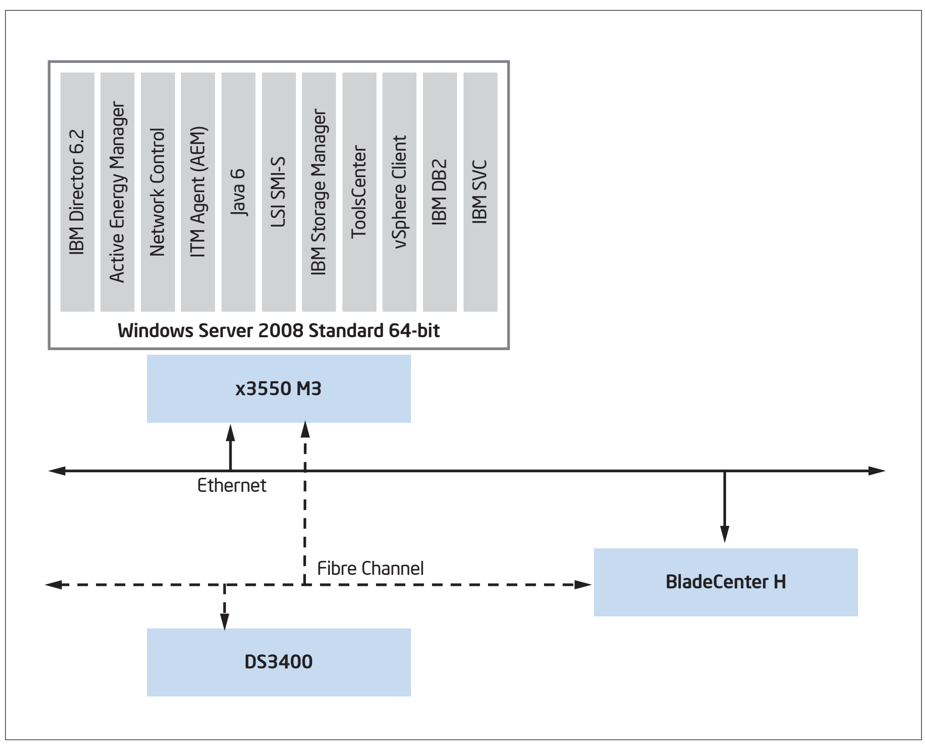

Figure 3. IBM CloudBurst* Management Software Stack

CloudBurst Network Architecture

The following sections describe the network architecture of IBM CloudBurst.

Management Network

- The management network is tied together by redundant SMC 8126 rack switches. These two SMC switches are connected by a redundant link.
- Each component with a 1GbE management interface is redundantly connected to the SMC switches as shown in Figure 4.
- All management network traffic is untagged (no VLAN), except for the IBM Tivoli Monitoring component (the ITM virtual machine) of the Tivoli software stack, which requires access to the x3550 M3 to communicate with Active Energy Manager (a plug-in of IBM Systems Director). Therefore, the management network must also be accessible to the HS22V blades in the Management Cluster. To maintain a more secure environment, management traffic that enters the BladeCenter is tagged with a VLAN 70 as it enters the chassis. This VLAN is only accessible to the virtual machines in the Management Cluster (Tivoli software stack VMs and VMware vCenter* VMs).
- The SAN24B-4 Express Fibre Channel switches are only included in the large template.
- The SVC (SAN Volume Controller) may also connect to the management network; SVC is a field installable option providing access to the client's data storage.
- Figure 4 also highlights the available connections from the x3550 M3 to a customer's 1GbE management network.

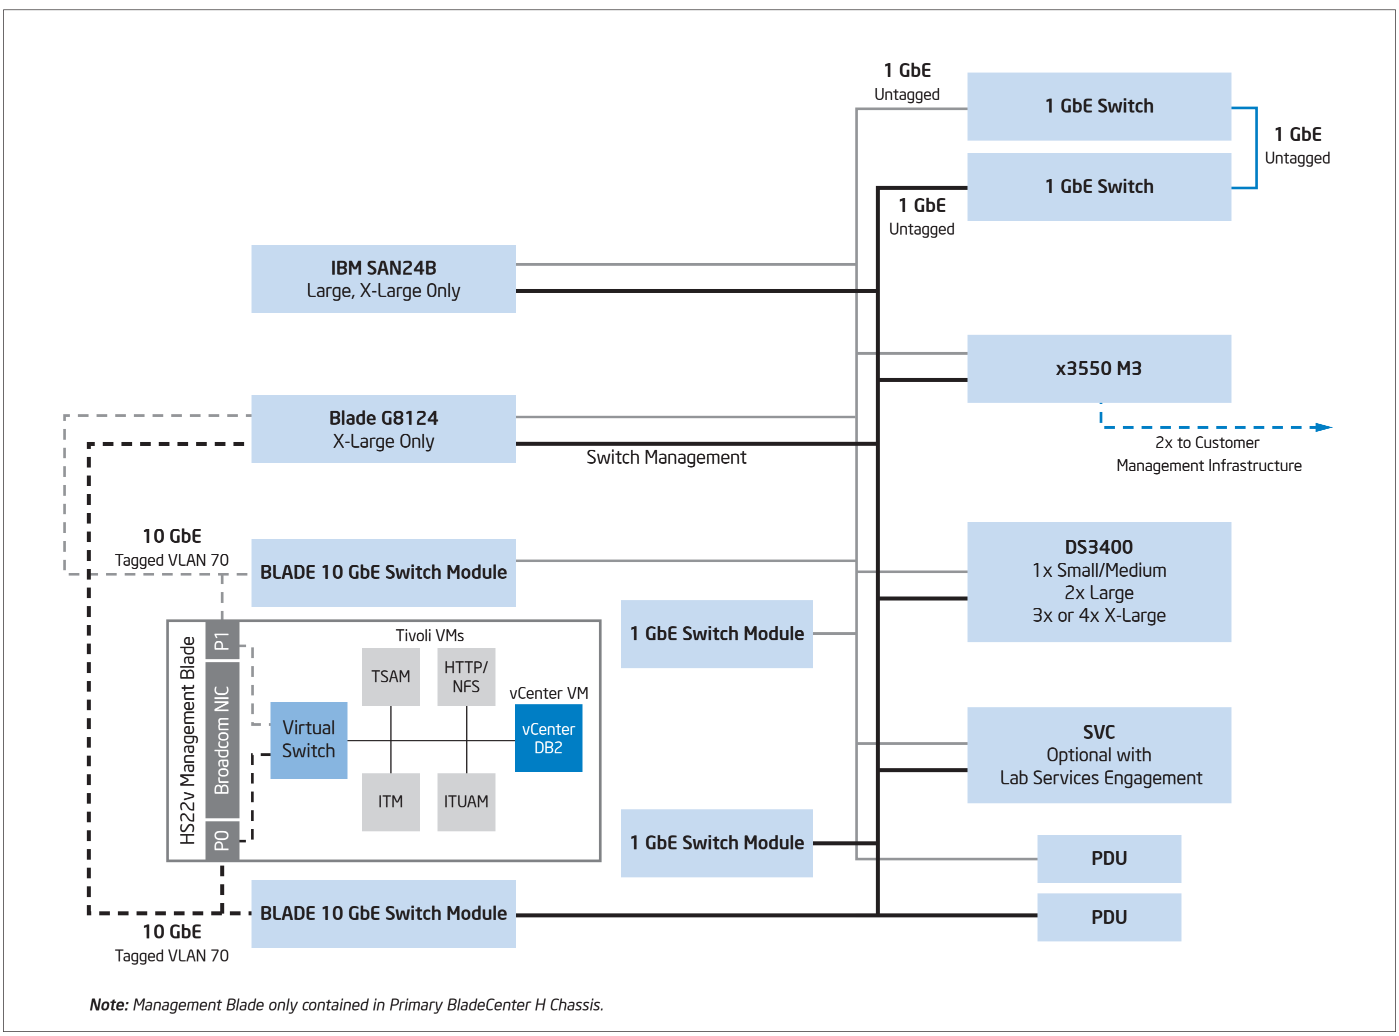

Figure 4. IBM CloudBurst* Management Network Architecture

VM Network

The VM network is a separate and private logical network that is used for inter-communication between the VMs of the Tivoli software stack and that provides a monitoring function for the provisioned VMs (see Figure 5). VLAN 90 is used for inter-communication between the Tivoli software VMs, and VLANs 100-129 are used for agent monitoring. Each switch in the solution is pre-configured with the VLANs required for the VM network. For all three configurations, this includes the BNT Virtual Fabric 10Gb Switch Modules. All of the VLANs associated with the VM network are not exposed to any of the interfaces connecting to the customer data network.

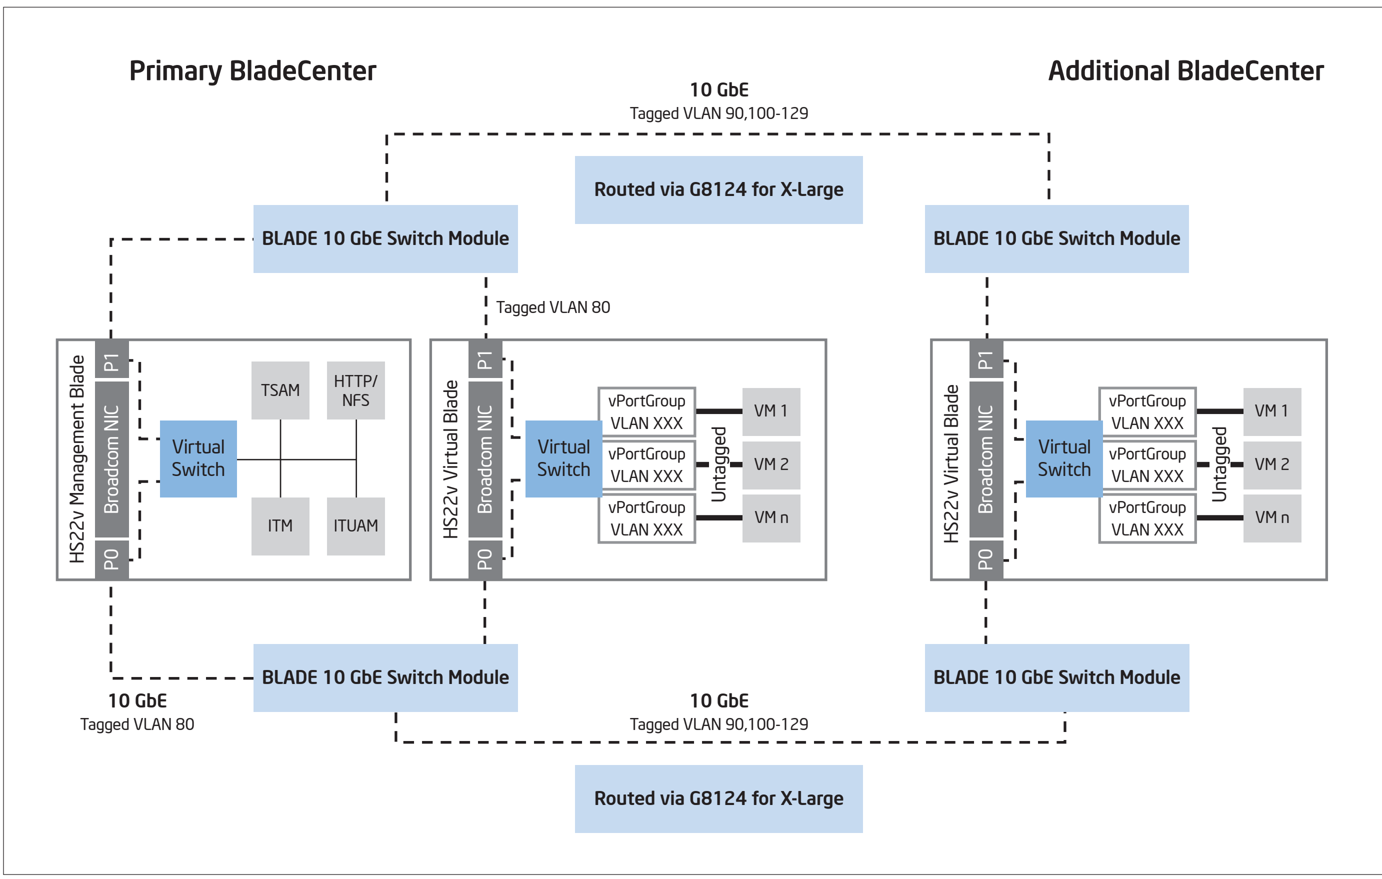

Figure 5. IBM CloudBurst* VM Network

VMotion* Network

The VMotion* network is a logical segment of the 10GbE network that the VMware Infrastructure uses to migrate VMs between the hosts in a given cluster (see Figure 6). All of the switches in the solution are pre-configured with VLAN 80 to enable VMotion network traffic.

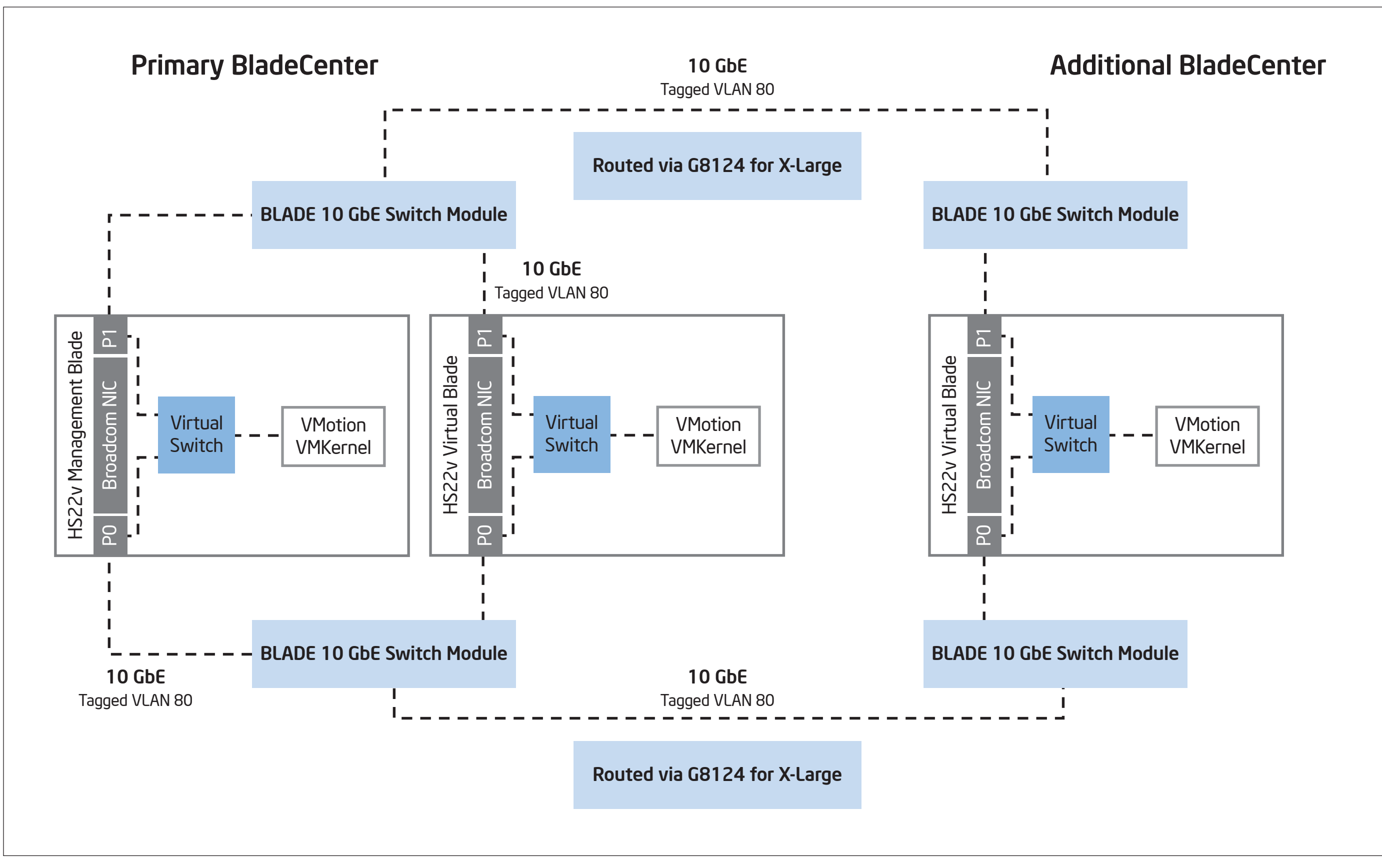

Figure 6. IBM CloudBurst* VMotion* Network

Use Case Details

To better understand the IBM CloudBurst solution [1], a series of use cases were developed involving either a cloud administrator (enterprise IT or service provider), or an end-user of the cloud. These scenarios include:

1. Accessing the private cloud via a self-service portal
2. Managing users in Tivoli Service Automation Manager
3. Creating and deleting virtual machines (virtual machine management)
4. Billing and metering
5. High availability
6. Energy management
7. External virtual infrastructure management

After completing these tasks, you will have a functional private cloud solution.

Note: The information in the following sections is obtained from the official IBM documentation of various products. Please consult the respective product documentation [1] [7] [8] for more details on the specific use cases.

Use Case 1: Accessing the Private Cloud via a Self-Service Portal

The IBM Tivoli Service Automation Manager (TSAM) assists in the automated provisioning, management, and de-provisioning of IT landscapes, comprised of hardware servers, networks, operating systems, middleware, and applications. Tivoli Service Automation Manager uses IBM Service Management and the Tivoli process automation engine as an integration platform.

Tivoli Service Automation Manger includes both self-service and administrative user interfaces. The self-service user interface, introduced with Tivoli Service Automation Manager v7.2, is tailored to the users of self-service offerings. Its design is based on the Web 2.0 standard, which allows context-sensitive, real-time display updating based on the user's current entry or selection. This allows faster access to the necessary information without having to go through a sequence of clicks, dialogs, and panels.

Tivoli Service Automation Manager User Roles

User authorizations for Tivoli Service Automation Manager are managed with self-service and administrative roles. There are four roles for the self-service user interface, and an administrator role for the administrative user interface.

The self-service user interface identifies four user roles:

Cloud Administrator

The following tasks are limited to the cloud administrator:

- Define new teams, user accounts and their associated roles
- Modify their own information
- Register and unregister software images
- Allow resource allocations and changes
- Check the status of projects and monitor the servers for all users
- Approve or deny provisioning requests made by team administrators

Cloud Manager

Users in this role are the read-only administrators of the cloud. They can:

- Modify their own information (except for role and team membership)
- Check the status of projects and monitor the virtual servers for any team

Team Administrator

Users in this role can:

- Create user accounts for other users
- Modify their own information (except for role and team membership)
- Place requests for provisioning servers and check the status of projects
- Monitor the servers
- Change server status or password
- Log in and use the provisioned servers and applications

Team User

Users in this role can perform the following tasks:

- Modify their own information (except for role and team membership)
- View projects available for their team
- Check the status of the servers provisioned for their team
- Log in and use the provisioned servers and applications

Important: One user can be assigned only one role; for example, a team administrator in one team cannot be a team user in another team.

The administrative user interface identifies one user role:

Tivoli Service Automation Manager Administrator

Users in this role can perform tasks available only in the administrative interface:

- Run discoveries and workflows
- Display usage and accounting reports
- Configure, register, and unregister software stacks

These tasks are described in the *Tivoli Service Automation Manager Installation and Administration Guide*. [10]

Self-Service Image Management

Software and server images can be maintained in the Tivoli Provisioning Manager image library for selection at provisioning time. New server image templates can be created and imported to the library. Once the images are in the library, they need to be registered so that they can be used to provision new virtual servers. In addition, a snapshot-like image of an entire provisioned server can be saved and then restored in the current project, thus initializing the server at the state represented by the image.

Self-Service Virtual Server Management

Tivoli Service Automation Manager supports end-user-initiated provisioning and management of virtual servers. The functionality addresses a long-standing need by data centers to efficiently manage the self-service deployment of virtual servers and associated software. Using a set of simple, point-and-click tools, an end user can select a soft-ware stack and have the software automatically installed or uninstalled in a virtual host that is automatically provisioned.

This functionality ensures the integrity of fulfillment operations that involve a wide range of resource actions:

- Creating virtual servers within a new deployment project or adding new virtual servers to an existing project, with optional scheduling for implementation at some future time.
- For each virtual server created, installing a software image that includes an operating system and other applications that are associated with the image.
- Installing additional software on the provisioned virtual machines.
- Deleting a virtual server when its usefulness is at an end, freeing up its resources for use by other servers.
- Saving virtual server images and restoring the servers to their previous state.
- Saving and restoring images of servers within the project.
- Deleting individual servers.
- Canceling a project and deleting all of the associated virtual servers.
- Starting, stopping, and restarting virtual servers.
- Resetting the administrator password on a virtual server.
- Adding, removing, and modifying users and user teams.

Before users in the data center can create and provision virtual servers, administrators perform a set of setup tasks, including configuring the integration; setting up the virtualization environments managed by the various hypervisors, and running a Tivoli Provisioning Manager discovery to discover servers and images across the data center.

After this initial setup, the administrator associates the virtual server offerings with Tivoli Provisioning Manager virtual server templates. In addition, the image library is used as the source for software images to be used in provisioning the virtual servers.

Use Case 2: Managing Users in Tivoli Service Automation Manager

A team is a logical group of users accessing the same resources. Each team is separately tracked for the usage and accounting data. A user can belong to more than one team. Each project can be assigned to one team only, but the team can have access to more than one project. When you create a user you need to assign a role to them. Each role has a different set of privileges.

Creating New Teams

To create a new team:

1. Log in to the self-service interface as a cloud administrator.
2. From the Home panel, select **Request a New Service > Virtual Server Management > Manage Users and Teams > Create Team**. The Create Team window is shown in Figure 7.
3. Provide the name for the team. It must be no more than 8 characters and unique, for example: **TestTeam**.
4. You are going to use the accounting function in the next tutorial, so specify the project account. It must be no more than 20 characters; for example, **TestTeamAccount.**
5. Provide a team description (optional).
6. Select the default **PMRDPCAUSR** user from the list, and click > to add them to the team.

   Note: Up to 15 users can be added in one request. To add more users to the team, you can create a new request to modify the team that has been created.
7. Click **OK** to submit the request.

Creating New Users

To use a project, the user must be a member of the team that has access to this project.

Cloud and team administrators can create new users and add them to the teams. When creating a new user, cloud administrators can assign all roles and pick from all teams available in the cloud. Team administrators can only assign the Team Administrator and Team User roles, and can add the user only to the team they themselves belong to.

In one request, a user can be added to up to five teams. To add a user to more teams, use the Modify User Information request after submitting this one.

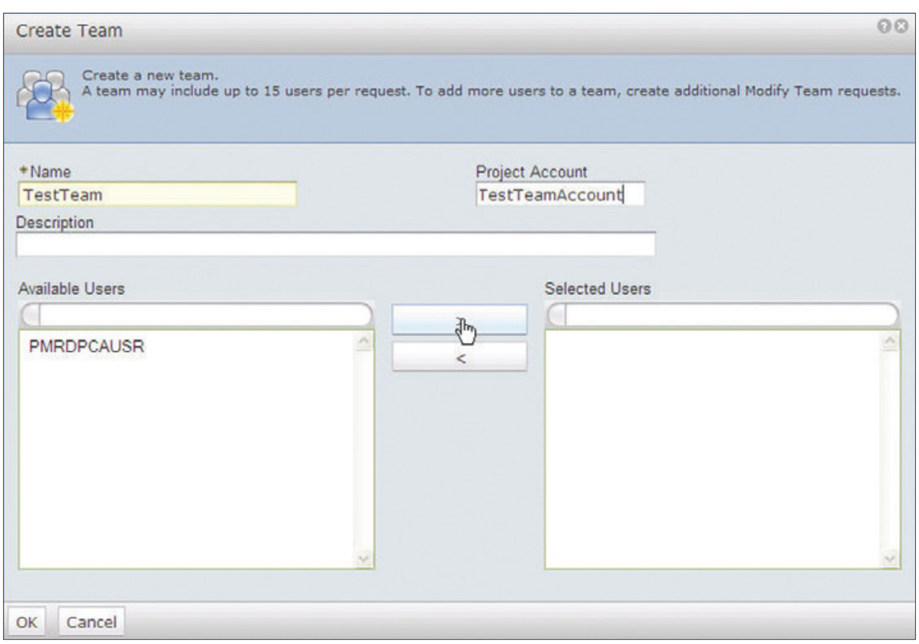

Figure 7. Create New Teams

To create a new user:

1. Log in to the self-service user interface.
2. From the Home panel, select **Request a New Service > Virtual Server Management > Manage Users and Teams> Create User**. The Create User window opens (see Figure 8).
3. Provide a User ID and a Display Name for the user. The User ID must be no more than 30 characters and unique (not in use or deleted), for example **TeamAdmin**.
4. Specify and confirm the Password. The Password must be at least six characters.
5. Select a Role for the user; for example, **Team Administrator**.

   Important: One user can be assigned only one role, so, a team administrator in one team cannot be a team user in another team.
6. Click to view the list of available teams and select the team to add the user to it. If you created **TestTeam** in the previous module, you can now add the user to it.

   Tip: If there are no teams available, after finishing this task, use the Create Team request (when creating a team, you can specify which users to add to it). You can also use the Modify User request to add a user to a team at any time. If the user does not belong to any team, they are not able to use any of the provisioned virtual machines.
7. Select **Make the Account Active** to enable the user to log in after the new user registration is completed.
8. Fill in the personal information details as required.
9. Set the Regional Settings:
   - **Locale:** Select the language in which the user is to see the information in the self-service interface.
   - **Timezone:** Select the timezone in which the user is to see the information in the self-service interface.
10. Click **OK** to submit the request.

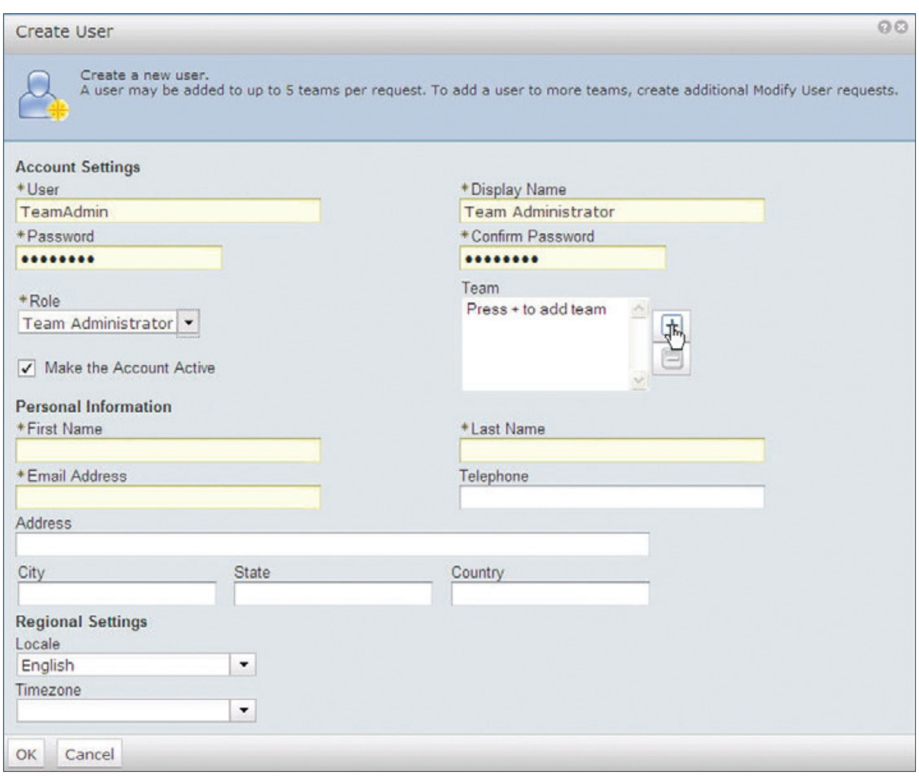

Figure 8. Create User

Use Case 3: Creating/Deleting Virtual Machines (Virtual Machine Management)

Creating a Project

A project is a group of virtual resources of the same type that can be used by a team. Each project is assigned to one team, and all team members can access the virtual machines provisioned for that project and monitor the state of the project. The resources provisioned in a project can be modified—you can add and remove servers, install additional software on the servers, or modify server capacity.

Only cloud administrators can create new projects. Team administrators can request new projects, and, depending on the approval workflow configuration, the request will be auto-approved (default), or sent for cloud administrator's approval.

To create a project:

1. Log in to the self-service interface as a cloud administrator.
2. In the Home panel, click **Request a New Service > Virtual Server Management**.
3. Depending on the hypervisor (virtual technology) that is configured in your environment, select **Create Project** with the configured server type. The Create Project window opens. (see Figure 9).
4. Enter a Project Name, for example: **TestProject**.
5. Select the team that should have access to the project. Only one team can use a project. If you created **TestTeam** in the previous lesson, select that team.
6. Optional: Type a description for the project.
7. Set the Start/End Date and Time for the project.

   You can select a future start date to create a reservation. If you set the End Date, at that time all of the servers that have been provisioned within that project are de-provisioned. The host names that were automatically assigned to these servers are freed up for use by other virtual servers that are created in the data center. Any image saved for a server that participated in the project is deleted.
8. From the Resource Group Used to Reserve Resources drop-down menu, select **Monitoring Agent to be Installed**. This agent performs scans to identify server resource utilization (memory, CPU, and disk space). The usage values can then be viewed in the Manage Servers panel.

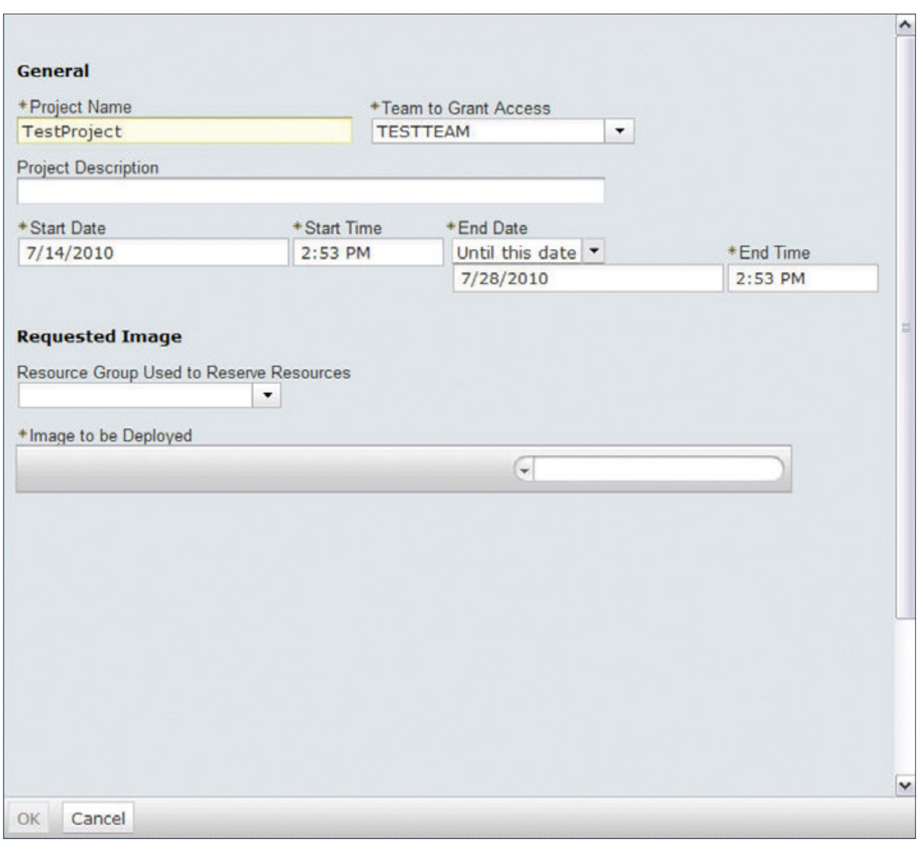

Figure 9. Create Project

9. Select the image to be deployed. Only registered images are available.
10. Use the arrows to select additional software to be installed on the provisioned servers.

If the software item requirements can be fulfilled, its name will appear in the Selected Software box. Otherwise an error message will be displayed. The software items will be installed in the same order they are added in the box.

11. To specify any additional parameters for the software items, select an item and click **Configure**. Fill in the fields as required.
12. In the Resources section, adjust the settings of the requested resources as required.
13. Click **OK** to submit the request.

Saving a Virtual Server Image

You can save snapshot-like images of the existing provisioned virtual servers in their current state. The saved image can then be used in a subsequent request to restore a virtual server, but not to create new servers.

To save an image of a virtual server:

1. In the Home panel, click **Request a New Service > Virtual Server Management > Backup and Restore Server Image > Create Server Image**. The Create Server Image opens (see Figure 10).
2. In the corresponding fields, type Virtual Server Image Name and Description.
3. Select the project in which the server is provisioned (Project Name). Project Details will display.
4. From the Select a Server to Save an Image display list, select a server that you want to create an image of. You can view the list of any previously saved images for this server.
5. Click **OK** to submit the request.

When the My Requests section on the right side of the Home panel shows this request as "Resolved", the image has been saved. You can now restore a server to its previous state at any time. You can also use the Remove Server Type Saved Image request to remove the images that are no longer needed.

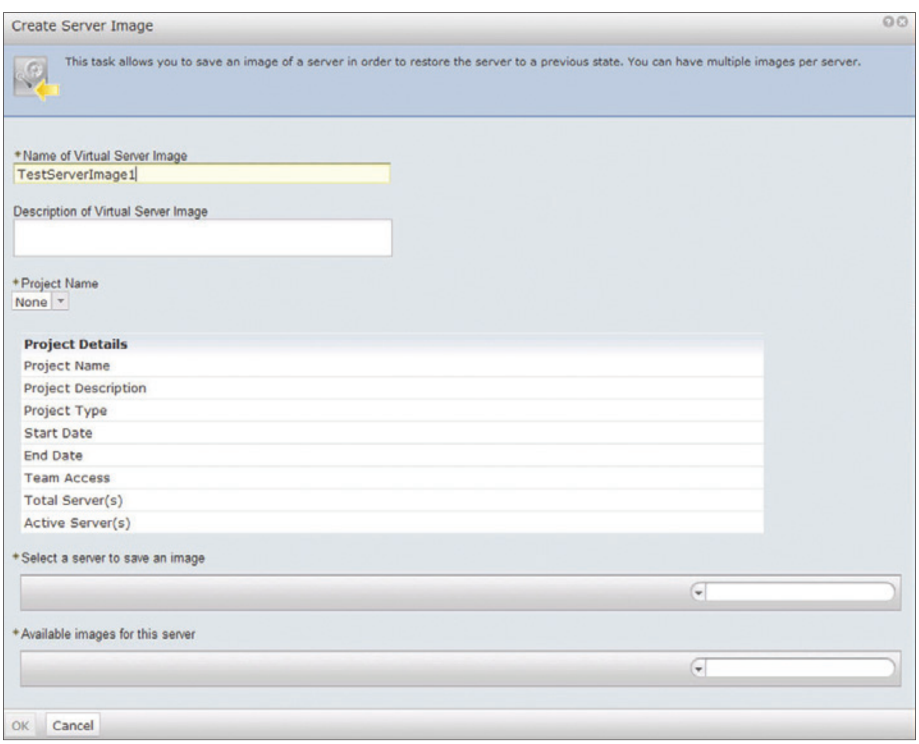

Figure 10. Create Server Image

Restoring a Virtual Server from an Image

You can restore a virtual machine to its previous state using the image saved in the previous lesson.

To restore a server from an image:

1. In the Home panel, click **Request a New Service > Virtual Server Management > Backup and Restore Server Image > Restore Server from Image**. The Restore Server from Image window opens (see Figure 11).
2. From the Project Name list, select a project for which you want to restore a server.
3. Select a server that you want to restore.
4. From the list that displays, select the image that is to be used for restoring the server.

   Note: The list is limited to the available images for servers that are currently provisioned.
5. Click **OK** to submit the request.

When the My Requests section on the right side of the Home panel shows this request as "Resolved", the server has been restored.

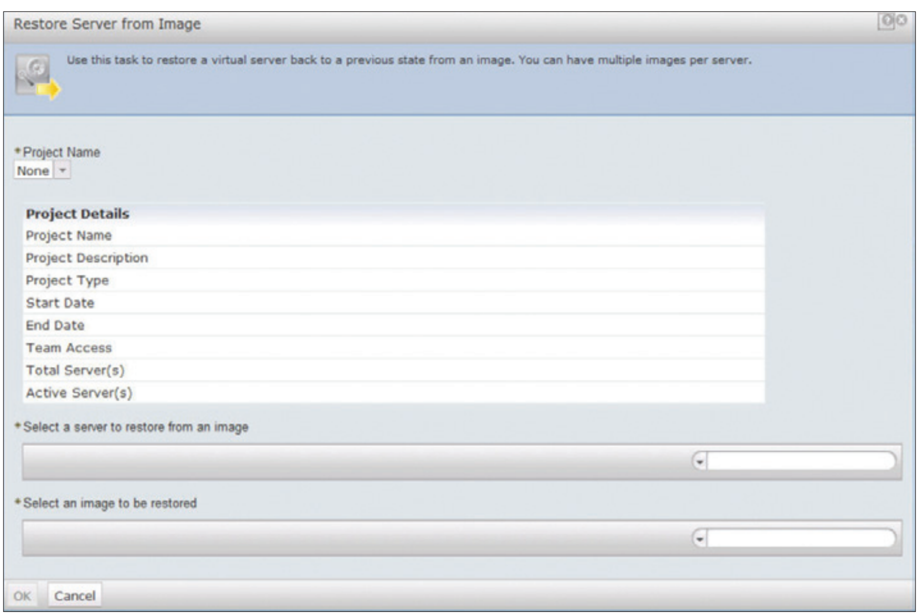

Figure 11. Restore a Virtual Server from an Image

Use Case 4: Billing and Metering

Managing Reports

The following IBM CloudBurst software components provide various reporting functions:

- Tivoli Service Automation Manager: Reports on service definitions and service instances
- Tivoli Usage and Accounting Manager: Reports including accounting and invoicing information
- Tivoli Monitoring: Reports including performance and availability of monitored resources

Running the Tivoli Service Automation Manager Reports

There are seven Service Automation Manager reports. Depending on the services that you have configured, the name of the report will vary:

- **All Service Definitions**
  AllSD<service>.rptdesign A list of all service definitions and their instances
- **Service Definition Summary**
  SDSummary<service>.rptdesign
  Basic service definition information and list of service definitions and deployments
- **Service Definition Details**
  SDDetail<service>.rptdesign
  Service-dependent content with identifying information, status, topology, monitoring and resource assignments, and service definition history
- **Service Deployment Summary**
  SISummary<service>.rptdesign
  Service-dependent graphics and lists showing CPU and memory levels, servers by owner, servers by platform architecture, monitoring agents by type, and instances
- **Service Deployment Details**
  SIDetail<service>.rptdesign
  Service-dependent content including graphics showing CPU and memory types, servers by platform architecture, monitoring agents, and service history
- **Co-Located Nodes**
  coLocation<service>.rptdesign
  Information concerning components of the service instance that reside on the same physical server
- **Service Deployment History Summary**
  SISummaryHistory<service>.rptdesign
  Service-dependent graphics showing CPU, memory, servers (over time and for all deployments)

Each report can be exported from the administrative interface in various formats, or sent by e-mail to the specified address.

To run the reports:

1. Log in to the administrative interface.
2. Click **Go To > Administration > Reporting > Report Administration**
3. In the Application field, type **PMZHB** and press **Enter** to view the list of all Service Automation Manager reports.
4. Select a report that you want to view. The report details are displayed in the Report tab.
5. If you are running the reports for the first time, click **Generate Request Page**, and wait until the process is finished. This step needs to be performed only once.
6. Click **Preview**. The request page is displayed.

   You can use the BIRT engine to view the report in HTML format, or have the report e-mailed in PDF format. If you do not specify the e-mail address the report will be generated in HTML format and presented online in BIRT Viewer. You can also schedule the reports in the future. In this lesson, you will view the report immediately as HTML.
7. In the Schedule section, select **Immediate** to run the report without scheduling, and click **Submit**.

Once the report is generated you can export it in the following ways:

- To export the report data in a .csv format, click **Export Data** on the Report toolbar, specify any required options, and click **OK**.
- To export the report in .ppt, .pdf, .doc, .postscript, or .xsl formats, click **Export Report** on the Report toolbar, select Format and any other required options, and click **OK**.

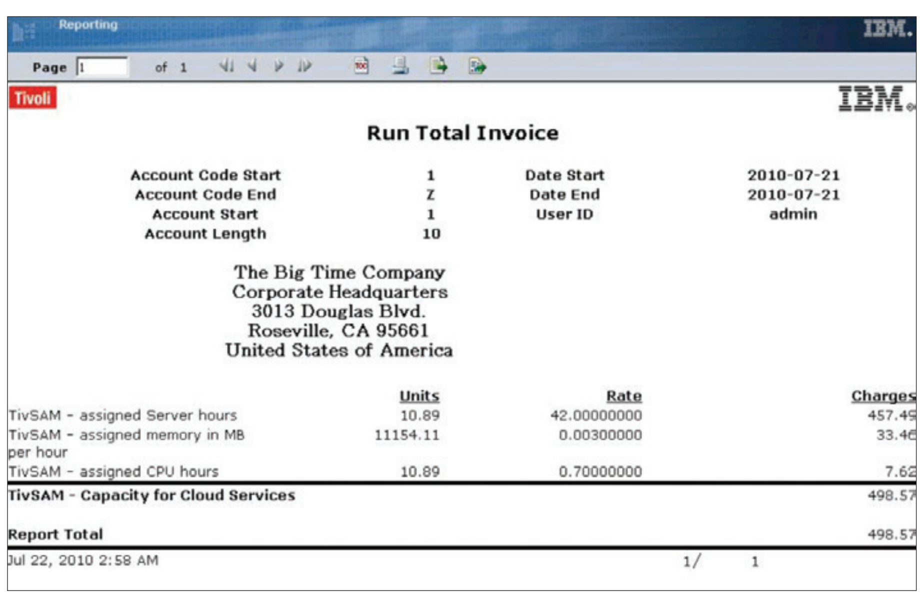

Figure 12. Example Tivoli* Usage Report

Use Case 5: Providing High Availability

IBM CloudBurst provides high availability to protect the cloud environment if the host management blade fails. The VMware High Availability feature provides high availability for the provisioned virtual machines, including failover protection against hardware and operating system failures. It monitors virtual machines, restarting them on other physical servers in the resource pool when a failure is detected. It also protects applications against operating system failures, automatically restarting the virtual machines when such failures occur.

As an option, high availability for the management blade can be provided using both the VMware High Availability feature and Tivoli System Automation for Multiplatforms:

- VMware provides high availability for the management services provided by the IBM CloudBurst software stack (apart from HTTP/NFS/Samba services and the services running on icb-tivsam). If a blade fails, the system restarts the IBM Tivoli Monitoring, and Usage and Accounting Manager virtual servers on a different blade in the IBM CloudBurst box.
- Tivoli System Automation for Multiplatforms provides an improved high availability solution for the NFS virtual server (HTTP/NFS/Samba services) and the Tivoli Service Automation Manager virtual server. It detects not only hardware and operating system but also software failures (for example, service crashes), and provides a fully automated recovery.

The Tivoli System Automation for Multiplatforms solution for high availability offers significant improvements over the high availability of the hypervisor. For instance, only the Tivoli System Automation for Multiplatforms can detect and recover from software failures by restarting or failing over the software components. For hardware, Tivoli System Automation for Multiplatforms can detect and recover from blade, network, and disk failures.

The hypervisor, by contrast, detects only the outage of a complete server. **Tivoli System Automation for Multiplatforms** also provides application start-and-stop automation, since restarting the operating system does not necessarily guarantee that the service is back up and running.

Use Case 6: Energy Management

The cloud administrator can manage the energy use of the cloud environment using IBM Tivoli Monitoring for Energy Management and IBM Systems Director Active Energy Manager.

With the energy management software, the cloud administrator can:

- Monitor real-time energy data
- Define automation that detects power-related errors based on events and monitored data
- Collect and analyze historical information about power use in the infrastructure

The Active Energy Manager agent provides energy status and data to the IBM Tivoli Monitoring software so common problems can be identified and resolved. With the Tivoli Enterprise Portal, the cloud administrator can view the energy status and data that the agent provides. The Active Energy Manager agent also provides predefined workspaces in the Tivoli Enterprise Portal for monitoring common energy data, and customized basic workspaces for managing the IBM CloudBurst appliance environment.

Use Case 7: External Virtual Infrastructure Management

The cloud administrator can use the IBM CloudBurst service automation capabilities to manage not only the internal resources provided with the product but also external resources accessible through the customer network. The administrator can control the automated provisioning of the virtual systems using the internal IBM CloudBurst resources or other external pools of virtual resources, including additional types of hypervisors not provided with IBM CloudBurst—IBM WebSphere CloudBurst Appliance, VMware, KVM, Xen, and others.

Use Case Summary

IBM CloudBurst is an integrated, pre-packaged solution that combines IBM System x hardware (with its respective management software, e.g., IBM Systems Director), VMware vSphere 4.1, and Tivoli software packages to create a private cloud solution. These use cases demonstrate that the private cloud is ready to use via the self-service portal.

Next Steps

The next step is to evaluate IBM Service Delivery Manager on Intel processor-based IBM hardware as a potential multi-tenant Intel® Cloud Builders reference architecture.

IBM Service Delivery Manager

For customers needing flexible hardware choices, IBM also provides IBM Service Delivery Manager (ISDM) [5], a pre-integrated software stack deployed as a set of virtual images that automate IT service deployment and provide resource monitoring and cost management in a cloud. ISDM is a cloud management platform that enables the data center to accelerate the creation of service platforms for a variety of workload types with a high degree of integration, flexibility, and resource optimization. ISDM includes these core service management capabilities:

- A self-service portal interface for in-advance computing reservations of virtualized environments, including storage and networking resources
- Automated provisioning and de-provisioning of resources
- Real-time monitoring of physical and virtual cloud resources
- Integrated usage and accounting chargeback capabilities that can help system administrators track and optimize system usage
- Built-in high availability of the cloud management platform
- Pre-packaged automation templates and workflows for most common resource types
- Ability to leverage existing hardware for the cloud infrastructure

Intel® Xeon® Processor 5600 Series

With the help of Intel Xeon processor 5600 series, the trusted multi-tenancy technology is in place to measure the key components of the BIOS and hypervisor before allowing the server to join a pool of resources in the cloud. Some of the features in the next-generation Intel Xeon processor 5600 series include:

- **Intel® Trusted Execution Technology (Intel® TXT) [4]:** Using capabilities in the processor, chipset, BIOS, and a Trusted Platform Module (TPM), Intel TXT provides a mechanism for enabling a very small atomic level of "assumed trust" while allowing a robust basis for verification of platform components such as BIOS, option ROMs, etc. up to a hypervisor or operating system. With Intel TXT, the assumed trust (root of trust) is pushed down into the processor itself—perhaps the best-protected component of any platform.
- **Intel® Advanced Encryption Standard: New Instructions (Intel® AES-NI) [4]:** Intel AES-NI are a new set of instructions available on Intel Xeon processor 5600 series based on the 32nm Intel® Nehalem microarchitecture. These instructions enable fast and secure data encryption and decryption, using the Advanced Encryption Standard (AES) which is defined by FIPS Publication number 197. Since AES is currently the dominant block cipher, and it is used in various protocols, the new instructions are valuable for a wide range of applications.

Things to Consider

Networking

Cloud data center networks are complex systems. Traditionally, data center servers are connected to several different networks for functions such as production traffic, backups, storage, management, VM migration, etc. While Ethernet is the predominant networking technology, many data centers use Fibre Channel technology for their storage network.

If the traffic load demands it, individual network ports are aggregated to provide more capacity for a particular function. This results in many network connections for each server, plus the supporting cables and switch ports. Up to 12 network connections are common for a virtualized server today, which increases system complexity and raises costs.

The transition to 10Gb Ethernet allows consolidation of multiple separate Ethernet ports into fewer 10GbE ports, simplifying the data center network and reducing costs and while providing greater overall platform networking bandwidth.

Power Management

To save energy, intelligent power regulation is needed. Some energy savings are possible with permanent capping, but this is limited by the capping range, or by the need to remove the capping policy to meet performance goals.

An alternative to capping is dynamic power management policies that take advantage of virtualized cloud data centers. This means that power capping levels vary over time and become control variables. With dynamic power management, equipment can be shut down to actually reduce energy consumption, not just provide power management. These shutdowns alter the profile and topology of the cloud infrastructure, and require a mechanism to control them.

Dynamic power management is especially helpful in a virtualized environment, where applications are not bound to a physical host and can be moved around within a pool of servers to optimize the overall power and thermal performance of the pool. The loose binding between applications and hosts allows administrators to treat a group of hosts as a pooled resource, for optimizations that are not possible with individual machines, such as powering down equipment during low demand.

Scalability and Sizing Guidelines

As with any virtualized infrastructure, it is important to properly plan and size the virtual environment to meet current and future requirements. The following guidelines are based on the **IBM BladeCenter Virtualization Sizing Guide**. In addition to these general guidelines, please consult your IBM technical sales representative to determine the right sizing for your IBM CloudBurst solution.

Server Memory and CPU Headroom

It is important to plan headroom, or reserve server capacity, for peak periods when some consolidated workloads require greater server resources than average or steady-state periods. To calculate the recommended amount of memory and CPU headroom, IBM considered the variation between average and peak server resource requirements for each consolidated workload and statistically analyzed the probability that workloads would simultaneously reach peak utilization. The amount of CPU headroom varies based on the number of CPU cores and also on whether Intel Hyper-Threading Technology (Intel HT) is enabled. Less CPU headroom is needed as the number of CPU cores increases.

Memory Capacity versus CPU Utilization Constraints

Memory capacity, not CPU utilization, is often the primary constraint to adding more virtual machines on a consolidated host server. Workloads that are consolidated based on peak utilization characteristics tend to be CPU-constrained, and those based on average utilization characteristics tend to be memory-capacity-constrained or constrained by a hypervisor limit such as the maximum supported number of virtual machines or virtual CPUs.

Memory Speed

The speed at which the memory DIMMs operate is affected by factors such as the number of DIMMs populated on the processor memory channels or the Intel processor Quick Path Interconnect (QPI) frequency. For Intel EP-class processors, populating three memory DIMMs on a memory channel generally results in memory being clocked at a slower speed than if only two memory DIMMs were populated on the memory channel. For Intel EP- and EX-class processors, the QPI frequency can range from 6.4 GT/s to 4.8 GT/s. The maximum memory speed is a function of the QPI speed. For example, current Intel processors that support a maximum QPI speed of 4.8 GT/s limit the maximum memory speed to 800 MHz.

Effect of Intel Hyper-Threading Technology on CPU Headroom

Enabling Intel Hyper-Threading Technology improves server consolidation performance for most commercial applications. To reflect this performance benefit in the core-based CPU headroom, a host server with Intel Hyper-Threading Technology enabled should use the CPU headroom corresponding to a host server with 25 percent more CPU cores. For example, a host server with 8 CPU cores and Intel Hyper-Threading Technology enabled should use the CPU headroom metric corresponding to a server with 10 CPU cores. Less CPU headroom is needed when Intel Hyper-Threading Technology is enabled. See the IBM BladeCenter Virtualization Sizing Guide for more details.

Peak Utilization versus Average Utilization

To correctly size the number of virtual machines suitable for a specific host server, it is important to understand the characteristics of the workloads being consolidated. Consider the average and peak utilization of the processor, memory, disk I/O, and network I/O required by the pre-consolidation workloads.

Hypervisor Limits

In some cases, the number of recommended virtual machines can exceed a hypervisor's software limit. VMware vSphere and Microsoft Hyper-V* currently have maximum virtual machine limits of 320 and 384 VMs per server, respectively. It's important to be aware of all hypervisor software limitations when reviewing the sizing guidelines.

Virtualization Consolidation Zones

Besides the tangible average and peak hardware utilization characteristics of typical server consolidation workloads, IBM has defined three Consolidation Zones that also consider the intangible workload characteristics such as tolerance for performance degradation and workload predictability.

When determining the Consolidation Zone to use for sizing the number of virtual machines a host server can support, consider the following business and workload characteristics of the customer environment:

- **Business Criticality**: Consider the role or usage model of the consolidated host server.
  - Development: A server used for development, testing, or QA
  - Production-regular: A production server that can tolerate hours of down time
  - Production-important: A production server that can tolerate up to an hour of down time
  - Production-mission critical: A production server that can only tolerate minimal down time (minutes)
  - Production-clustered: A mission-critical server that is currently clustered
- **Age of Servers:** Consider the age of the pre-consolidation server when evaluating the measured CPU utilization statistics. Older servers may have "Moderate to High" CPU utilization. However, their processing capacity may be a small fraction of more modern servers.
- **Type of Workload**: Heavyweight workloads run optimally on physical servers that can scale-up processors and memory capacity, whereas lightweight workloads can run well on servers with fewer processors and less memory capacity, and where the servers can be scaled out.
- **CPU Utilization**: Consider the CPU utilization of the pre-consolidation server relative to its age. For servers less than three years old, CPU utilization less than 10 percent is considered low. For servers more than five years old, discount the measured CPU utilization since it will equate to a small fraction of a modern server's CPU utilization.
- **Memory Footprint**: Consider the actual memory used by a workload and not the physical memory installed on the server.
- **Disk I/O**: Consider the average amount of disk I/O that a workload requires.
- **Network I/O**: Consider the average amount of network I/O that a workload requires.
- **Workload Predictability:** Consider the variability of a workload's requirement for server resources.
- **Tolerance for Performance Degradation:** Consider Quality of Service Agreements for response-time criteria, number of concurrent users, etc. Determine if the workload can tolerate increased latency during periods of peak utilization.
- **Desired Consolidation Ratio:** Consider the customer preference for the number of virtual machines to consolidate on a host server.
- **Virtualization Experience:** Consider the customer's experience with virtualization technology to determine their readiness to pursue more aggressive consolidation ratios. Customers without prior virtualization consolidation experience may prefer to limit the number of workloads operating on a single host server. Experienced customers may have the internal business processes in place to support higher consolidation ratios.

Use of Solid-State Drives (SSDs) to Improve Performance

Using SSDs may improve the performance of storage nodes (and compute nodes when local storage is utilized), as well as reduce the overall power consumption of the cloud. While not specifically tested in this reference architecture, you may want to consider using SSDs when implementing a large-scale cloud, to improve performance.

Hardware Considerations

A full discussion of processor and overall server performance considerations is beyond the scope of this paper. However, it is important to note that the performance of virtual machines running on the cloud platform is heavily influenced by the processor architecture, and specific feature sets available in the processor such as Intel Virtualization Technology (Intel VT). The use of high-performance server processors equipped with virtualization and I/O support feature sets, such as the Intel Xeon processor 5600 series, is strongly recommended. For more details on Intel Virtualization Technology, please refer to [3].

Glossary

BIRT: Business Intelligence and Reporting Tools (BIRT) is an open-source Eclipse*-based reporting system that integrates with your Java/J2EE application to produce compelling reports.

BladeCenter: One of the leading open-blade architecture products in the IT market, with a focus on processor, memory, and I/O configurability delivered with multi-generational forward- and backward-compatibility through collaboration with major IT players.

Data Network: The data network is network that is designed to allow for the transmission of data between the various computing resources and the enterprise storage solution, including shared storage, backup, archiving, and disaster recovery.

Fibre Channel (FC): FC is a gigabit-speed network technology primarily used for storage networking.

High Availability (HA): HA delivers the availability needed by many applications running in virtual machines, independent of the operating system and application running in it. HA provides uniform, cost-effective failover protection against hardware and operating system failures within a virtualized IT environment.

Infrastructure as a Service (IaaS): The delivery of computer infrastructure (typically a platform virtualization environment) as a service.

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 you can consolidate more applications and heavier workloads per server to get better value from your server and software investments.

Kernel-based Virtual Machine (KVM): Is a full virtualization solution for Linux on x86 hardware containing virtualization extensions (Intel VT). It consists of a loadable kernel module, kvm.ko, that provides the core virtualization infrastructure and a processor-specific module. Using KVM, one can run multiple virtual machines running unmodified Linux or Windows images. Each virtual machine has private virtualized hardware: a network card, disk, graphics adapter, etc.

Platform as a Service (PaaS): Enables the consumer to deploy consumer-created or acquired applications onto the cloud infrastructure.

Provisioning: The process of providing customers or clients with accounts, the appropriate access to those accounts, all the rights associated with those accounts, and all of the resources necessary to manage the accounts.

Storage Area Network (SAN): A SAN is a dedicated network separate from LANs and WANs that connects all the storage resources in a cloud platform. It consists of a collection of SAN hardware and that typically has high inter-connection rates between the various storage devices, and software that manages, monitors, and configures the SAN.

Software as a Service (SaaS): Enables the consumer to use the provider's applications that are running on a cloud infrastructure. The applications are accessible from various client devices through a thin-client interface such as a web browser.

Virtual LAN (VLAN): A VLAN is a group of hosts that communicate as if they were connected to the same broadcast domain, regardless of their physical location.

Virtual Machine (VM): A VM is a software implementation of a computer that executes programs like a physical machine.

VMotion: VMware VMotion technology, unique to VMware, leverages the complete virtualization of servers, storage, and networking to move an entire running virtual machine instantaneously from one server to another.

Additional Information

For more information on the Intel Cloud Builders program, visit **www.intel.com/cloudbuilders**

For more information on Intel Xeon processors, visit **www.intel.com/xeon**

For more information on IBM CloudBurst, visit **https://www.ibm.com/developerworks/wikis/display/ibmcloudburst/Home**

References

[1] IBM CloudBurst v2.1. Software User's Guide: **http://publib.boulder.ibm.com/infocenter/tivihelp/v10r1/topic/com.ibm.cb.doc_2.1/tsam.cb.book.pdf**

[2] IBM CloudBurst Wiki: **https://www.ibm.com/developerworks/wikis/display/ibmcloudburst/Home**

[3] Intel Virtualization Technology: **www.intel.com/technology/virtualization** and **http://download.intel.com/business/resources/briefs/xeon5500/xeon_5500_virtualization.pdf**

[4] Intel AES-NI and Intel Trusted Execution Technology (Intel TXT]: **www.intel.com/technology/security/**

[5] IBM Service Delivery Manager: **http://www-01.ibm.com/software/tivoli/products/service-delivery-manager/features.html?S_CMP=rnav**

[6] Intel Xeon Processors: **http://www.intel.com/xeon**

[7] Administering IBM CloudBurst: **http://publib.boulder.ibm.com/infocenter/tivihelp/v10r1/index.jsp?topic=%2Fcom.ibm.cb.doc_2.1%2Ft_cb_config_multiple_network.html**

[8] IBM Tivoli Service Automation Manager: **http://publib.boulder.ibm.com/infocenter/tivihelp/v10r1/topic/com.ibm.cb.doc_2.1/c_functionality.html**

[9] IBM Tivoli Monitoring: **http://publib.boulder.ibm.com/infocenter/tivihelp/v10r1/topic/com.ibm.cb.doc_2.1/c_cb_monitoringoverview.html**

[10] Tivoli Service Automation Manager Installation and Administration Guide: **https://www.ibm.com/developerworks/wikis/display/tivolidoccentral/Tivoli+Service+Automation+Manager**

Printed in USA 1010/RR/OCG/XX/PDF Please Recycle 324485-001US