Intel® Cloud Builders Guide
Intel® Xeon® Processor-based Servers
HP ProLiant SL* Scalable Systems
Enomaly Elastic Computing Platform*

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

HP ProLiant SL* Scalable Systems
Enomaly Elastic Computing Platform*, Service Provider Edition

Intel® Xeon® Processor 5500 Series
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. The establishment of 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 as it supplies a basic hardware blueprint and available cloud software management solutions such as the Enomaly Elastic Computing Platform*, Service Provider Edition on HP ProLiant SL* Scalable System servers. The use cases described in this reference architecture, also called "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 reference architecture is cloud service providers, cloud hosters, and enterprise IT administrators who are looking to realize revenue potential as they expand their existing data center infrastructure to offer cloud computing services to their customers or internal users.

Table Of Contents

Executive Summary

The ideas and principles of the cloud are compelling because of the agility and cost savings. In recent years, many companies have set up cloud infrastructures that can be accessed over the public Internet, and these companies are now offering services that customers can utilize 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).

For preexisting workloads, workloads unique to an organization, and workloads that need to remain under close control of an organization, the use of an IaaS structure provides the best choice. The support and hosting of an IaaS cloud service, whether a revenue-producing, customer-facing service (a public cloud), or an internal (private cloud) behind a corporate firewall, raises the question of which steps should be used to build the cloud.

Given the many choices, trade-offs, and decisions, it is clear that the ability to build on a known and understood starting point is key. This paper summarizes a body of work about the construction of a cloud using HP ProLiant SL Scalable System servers[1] running Intel® Xeon® processor 5600 series-based microarchitecture,[2] and Enomaly Elastic Computing Platform* (ECP), Service Provider Edition (SPE).

Our cloud implementation consists of five HP ProLiant SL Scalable System servers running Intel Xeon processor 5600 series exposed as a compute service. Enomaly ECP SPE does not differentiate between compute and storage pools, and it offers a very simple mechanism to build and manage the cloud. In this paper, we focus on how Enomaly ECP SPE provides the ability to support multi-tenancy, which enables our infrastructure to host virtual machines (VMs) belonging to different "customers," along with isolation of storage and network. Enomaly ECP SPE provides a very powerful yet simple portal interface through which we demonstrate how easy it is to establish an IaaS with connected storage.

Our design consists of one cluster of compute capacity and a single storage pool. In theory, there is no upper limit on the number of virtual machines that can be provisioned using Enomaly ECP SPE, which can support very large clouds with many thousands of servers in more complex architectures than the one implemented here. Depending on the end user resource requirements, the number of virtual machines allowed per server can be defined as needed. The cloud setup has a total of 160 GB RAM across five nodes. Since the VMMs (hypervisors) occupy approximately 500 MB, we can support up to 5x31=155 VMs with the five server nodes in our environment, with 1 GB RAM provided to each VM. When cloud configurations for lighter weight computing resources are designed, each virtual machine can be assigned less than 1 GB of RAM, and the specific allotment can be customized.

Using Enomaly ECP SPE's portal interface (accessible through any modern browser), we were able to provision virtual machines for multiple "customers," and connect to these machines using remote connection protocols such as Remote Desktop Protocol (RDP) and virtual network computing (VNC). The traffic for each "customer" is isolated from all of the other "customers" using VLAN packet tagging. In our case-specific cloud configuration, the storage for the virtual image repository was provided using the network file system (NFS), and the virtual machine instances resided on the local disk of the node on which the virtual machine was running. NFS was chosen because of its ease of use and its quick deployment. For production purposes, and in parallel or cluster configurations, we recommend more robust storage architecture be used. Some options to consider include cluster file system solutions such as Global File System (GFS-available with major Linux* distributions, and supported by Red Hat Inc.*), which may be combined with internet small computer system interface (iSCSI), such as the HP StorageWorks* P4000 G2 SAN solutions, or Fibre Channel storage area network (SAN) storage such as the HP StorageWorks P2000 G3 MSA Array Systems. Lighter-weight solutions can also be considered, depending on the scale of the cloud deployment. These solutions might include network-attached storage (NAS) technologies, such as the HP StorageWorks X9000 Network Storage System, as well as simpler server-based storage solutions supplemented with high-performance or bonded network links can also be considered, depending on the scale of the cloud deployment.

Network traffic is managed by the virtualized Enomaly ECP SPE node. One of the components that contributes to system bottlenecks is network traffic throughput. In future testing, we will use 10 GB Ethernet to help address the growing bandwidth needs of a cloud computing environment. We implemented Enomaly ECP Manager Agents and Enomaly ECP Node Agents on HP ProLiant SL170z* servers running Intel Xeon processor 5600 series.

These servers have excellent performance/watt per sq. ft. characteristics and are ideal for highly dense cloud architectures. They deliver up to a 60 percent performance boost over Intel Xeon processor 5500 series, use up to 30 percent less power, and offer new security features including Intel® Advanced Encryption Standard–New Instructions (AES-NI), which speeds data encryption, and Intel® Trusted Execution Technology (Intel® TXT), which prevents the insertion of malicious software prior to virtual machine launch in virtualized environments.

Introduction

Design Principles

The basic design principle we used provides an implementation of cloud architecture that addresses scalability and flexibility, multi-tenancy with isolation of data, and reliability, all while optimizing capital and operational expenses.

To achieve the lowest operational costs for the infrastructure, we used a homogeneous pool of compute and storage with uniform connectivity because it is the simplest to manage, to troubleshoot, and from which to add and/or remove capacity. This is not to say that a cloud contains only one kind of server or storage element. Rather, we strive to create pools of uniform compute, storage, and networking, which allows new pools to be added over time, for example, as new generations of server technology become available. This design approach allows for the automation of workload placement, the connection of external networks and storage, and the delegation of workload management within the cloud.

Virtualization offers the ability to scale, and provides the flexibility required to address unexpected or sudden changes of workload without the addition of new hardware. In a public cloud environment, extreme care must be taken in the design to ensure that data and workload isolation exist between users. By contrast, in an enterprise, reasonable controls to ensure isolation are required, but some access constraints might be relaxed.

We needed to make the assumption that all components of the cloud can potentially fail, such as hardware, network, or storage. Either the application or the cloud management software needs to be able to perform the necessary recovery actions in the event of failure. The state of each compute element must be maintained in shared storage or the application must accept responsibility for retrying the actions on a failing device. For a typical non-cloud-aware enterprise workload, this would require that the application be wrapped in a virtual machine and all the corresponding data stored in a shared storage. In this manner, whenever a server fails, the cloud management software can simply migrate or restart the virtual machine on another server.

Some design considerations used to build this implementation include:

- The use of standard HP Proliant server building blocks means that any server in the rack can be used for any purpose. In this configuration, there are no "special" servers and no special connections—all servers are physically configured the same way. This allows for simple replacement or reassignment of workloads, and the simple automation of workload configuration at the time servers are provisioned.
- A flat layer-2 network enables us to reduce the implementation costs and retain the flexibility to assign any workload to any server. This approach is flexible for small clouds, but may create bottlenecks for a larger number of servers as this design is aggregated at a larger data center level. Enomaly ECP SPE supports much more sophisticated network topologies, but they were not required for the purpose of this test bed.
- The KVM* hypervisor was used to provide isolation for the workloads on the servers. KVM is one of several hypervisors supported by Enomaly ECP SPE, in addition to Xen* and VMware*.
- An NFS server was used in combination with a direct attached storage (DAS) as a convenient and low-cost way to provide storage accessible to the clusters. Enomaly ECP SPE also supports the use of sophisticated storage architectures such as cluster file systems and SAN, and can avoid the use of node-local storage.

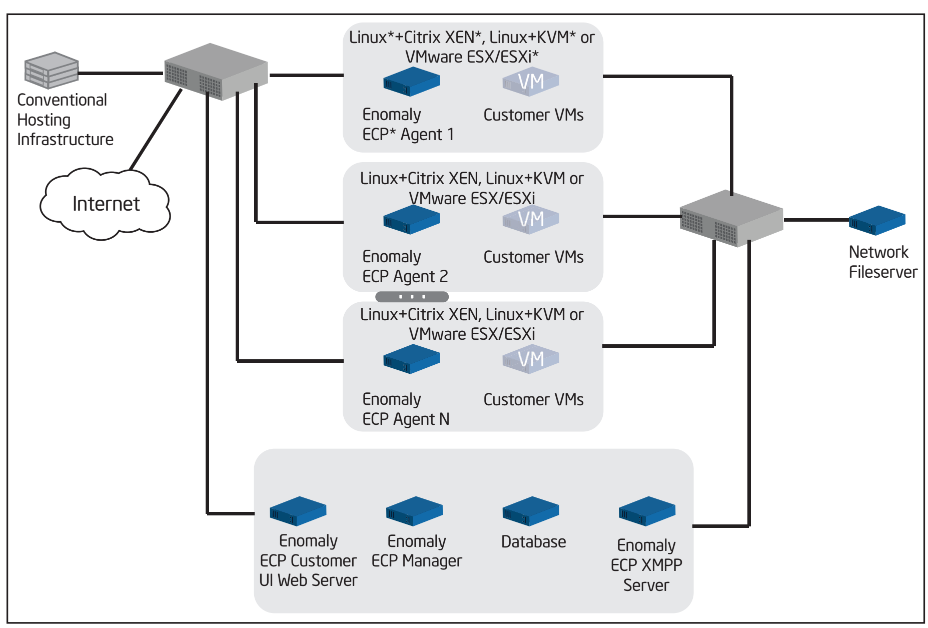

Figure 1: Basic Architecture Deployment for Enomaly ECP* SPE

Enomaly ECP Logical Architecture Configurations

Figure 1 shows the logical architecture of an Enomaly ECP SPE[3] deployment that was used in this hardware design. This is a basic architecture for an Enomaly ECP SPE deployment, suitable for small-scale test beds. The approach used in this paper is to create a compute pool with access to a common storage pool and access point.

Figure 2 shows more sophisticated logical architectures also supported by Enomaly ECP SPE. This is provided for reference only, and was not employed in our test bed.

Enomaly ECP SPE Implementation Overview

Enomaly's ECP Offering

Enomaly ECP is a software product comprising a cloud computing platform for service providers or enterprises. Enomaly's ECP, Service Provider Edition, is the answer for service providers who want to offer the powerful, flexible, and compelling economics of cloud computing to their customers.

Enomaly ECP, Service Provider Edition

Enomaly's SPE offering is a complete "Cloud in a Box" solution that is specifically designed to meet the requirements of carriers, service providers, and hosting providers who want to offer revenue-generating Infrastructure on Demand cloud computing services—with compelling and highly differentiated feature sets—to their customers quickly and easily.

Enomaly ECP SPE features include:

- Fully multi-tenant, carrier-class cloud service platform
- Detailed resource metering and accounting

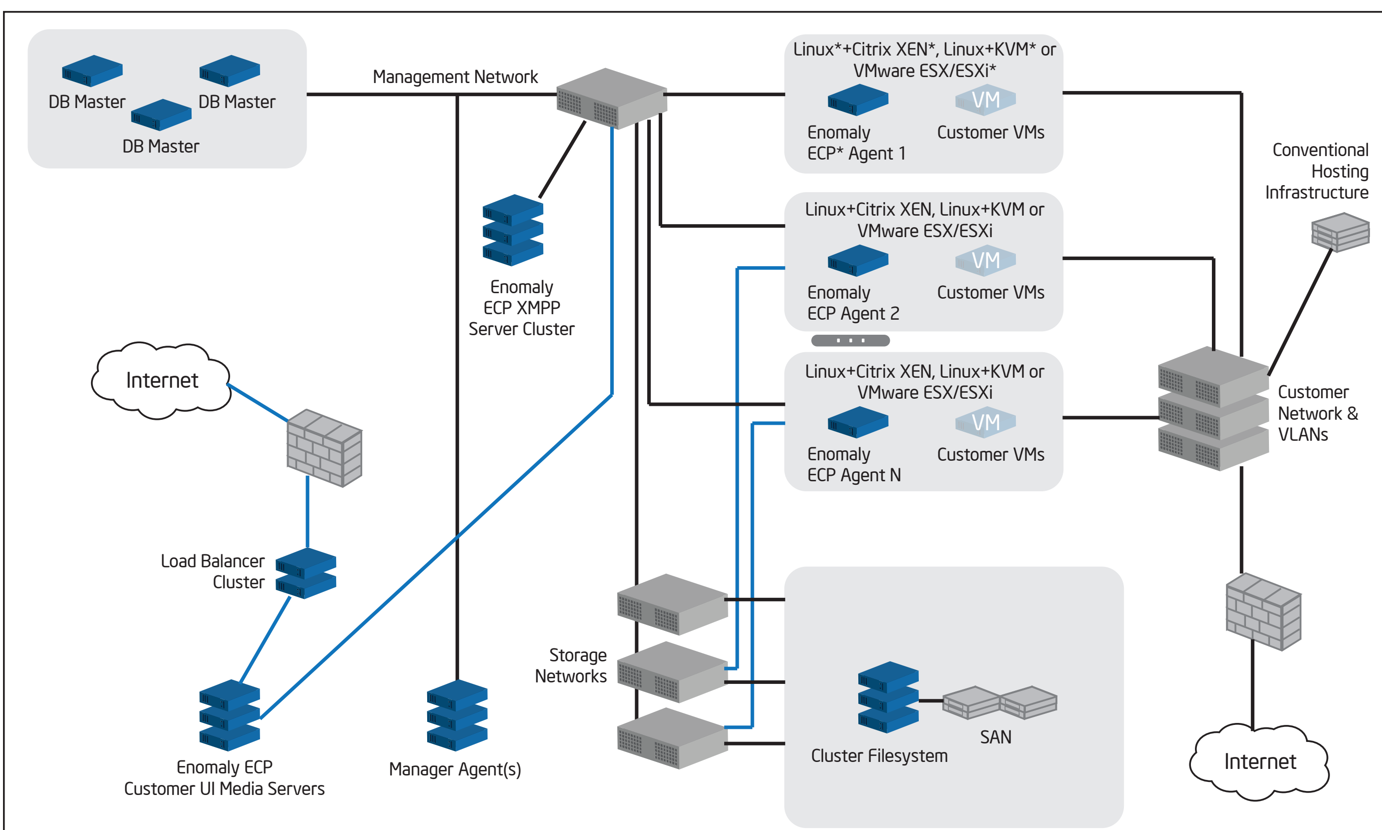

Figure 2: Carrier-Grade Deployment Architecture for Enomaly ECP* SPE

- Simple but powerful customer self-service portal (fully brandable and internationalized)
- HA capability for customer workloads
- Strong security capabilities (plus an additional High Assurance feature set available in Enomaly ECP, High Assurance Edition)
- Scale-up as well as scale-out capability
- Instant startup of hot spare servers
- Public and private (per-customer) VLAN security
- Integrated App Store: Through the Enomaly ECP App Center, a service provider can publish pre-built cloud applications directly to customers, which makes them available either free of charge or for a fee. Customers can directly provision VMs on the cloud from this library of pre-existing system images, including both business applications and infrastructure components, such as load balancers and firewalls.[4]
- Powerful automation APIs (both customer-facing API and administrative API)
- Intelligent and dynamic workload provisioning and orchestration
- Flexible integration with billing, provisioning, and operational monitoring systems
- Ability to utilize Enomaly's robust provisioning and orchestration engines, the earlier open source versions of which are already deployed by 15,000 organizations around the world

Enomaly ECP Architectural Principles and Design Philosophy

The Enomaly ECP[5] is a complete end-to-end platform that enables a telco or a service provider, such as a hosting company, to deliver IaaS to its customers. Enomaly ECP is based on a robust distributed architecture. Enomaly ECP applies the key attributes of the Internet cloud to the service provider data center; like the Internet, Enomaly ECP uses a decentralized command and control structure model and can therefore continue to operate in the event of critical failures. The Enomaly ECP system architecture provides an optimal degree of fault-tolerance, backward and forward application compatibility, extensibility, reliability, maintainability, availability, serviceability, usability, and security.

Scalable Design

Modular and Composable

Enomaly ECP comprises a large number of modular and composable components. Once a user designs a virtual architecture for a specific application, that virtual architecture can be put in place as you run and interconnect the modules requested by the application. This enables Enomaly ECP to operate efficiently in a wide range of application deployments.

Decentralization

Enomaly ECP achieves high levels of reliability and reaches massive scale through an architecture patterned on that of the Internet itself, in which multiple redundant components interoperate without a single central point of coordination.

In order to build scalable cloud platforms that can operate effectively at 1,000-node, and even 100,000-node and million-node scales, it is essential to use loosely coupled and decentralized architectures. The cloud must run like a decentralized organism, without a single person or organization managing it. Like the Internet, it should allow 99 percent of its day-to-day operations to be coordinated without a central authority. Applications deployed on a cloud managed in this fashion are more adaptive and fault tolerant because single points of failure are reduced or eliminated.

Oversubscription

To be cost-effective, a carrier-class IaaS platform must be able to efficiently provision resources for both peak and low-volume usage periods and provide an agreed upon minimum service standard.

Enomaly's ECP approach to this challenge is based on two essential principles—oversubscription and Class of Service—based resource quota management, which together enable powerful capacity control. In a public cloud infrastructure, an oversubscription model depends on the ratio of allocated resources to the maximum peak usage levels, the frequency and volume of peak usage, and the minimum service level agreement. The key is to manage resources around the standard deviation from the normal usage benchmarks, while simultaneously guaranteeing a particular quality of service for each customer.

Enomaly's ECP quota system manages predetermined levels of deviation across a specified resource pool that serves a population of customers. Service providers can oversubscribe their environments while remaining protected from (intentional or inadvertent) overuse and misuse. At the same time, they can allow for a variety of pricing and costing schemes to be implemented using a model that incorporates usage tiers, quality-of-service tiers, and the ability to provision additional resources dynamically as desired.

Unified Access to Heterogeneous Resources

Enomaly ECP offers unified access mechanisms, which provide common management points for heterogeneous collections of resources, including multiple virtual server platforms, multiple hypervisors, and large numbers of compute nodes potentially spanning multiple data centers. All of Enomaly

ECP's interfaces—the administrative user interface, the fully brandable customer self-service interface, and the end-user and administrative representational state transfer (ReST) APIs—can simultaneously manage diverse resources distributed across multiple geographies.

Command and Control Architecture of Enomaly ECP

Enomaly ECP provides a lightweight publish/subscribe messaging bus for command and control operations between the Enomaly ECP Agent software that controls each compute node and the Enomaly ECP Manager cluster. This scalable message bus is based on the Extensible Messaging and Presence Protocol (XMPP). XMPP is an open technology for real-time communication which is well-suited to the lightweight middleware role. XMPP runs on standard HTTP (tcp/80) and HTTPS (tcp/443) ports, which allows connections across most firewalls.

Our implementation supports all Enomaly ECP SPE features (e.g., provisioning, monitoring, virtual machine management, high availability, etc.) required to deliver IaaS functionality. The following sections will provide details on the software implementation, the hardware test bed, the test cases executed, and things to consider when you design a cloud.

Test Bed Blueprint Overview

Hardware Description

A cloud hardware infrastructure requires various hardware components for customers to host their cloud-based services. Typical hardware components include server platforms to host the compute loads, storage platforms to host data, network components to manage internal and external traffic, and other data center infrastructure components.

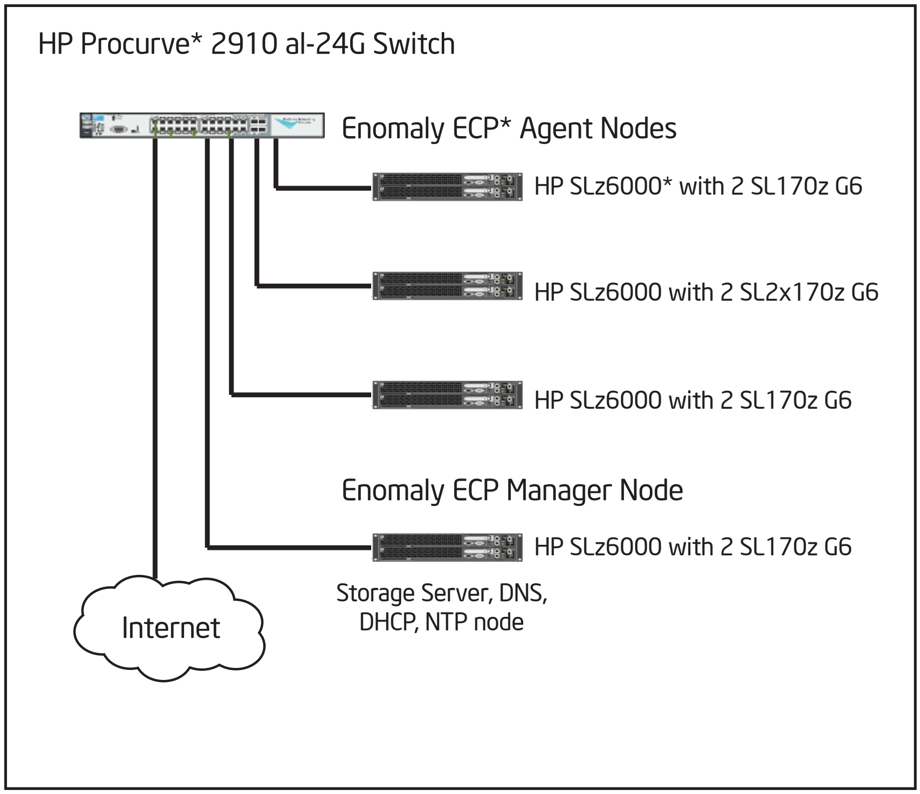

Figure 3: Current HP Test Environment

In our test bed, we used HP ProLiant SL Scalable System servers running Intel Xeon processor 5600 series[6] which provides a foundation to design new cloud data centers to achieve greater performance while using less energy and space, and dramatically reducing operating costs. Figure 3 shows the test bed physical architecture.

The HP ProLiant SL Scalable System servers running Intel Xeon processor 5600 series offers several features that help make it the best performing server for the cloud.[7] Some of these features include:

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

Physical Architecture

A small scale implementation was put in place at an HP cloud lab based on the Basic Deployment Architecture for Enomaly ECP SPE (see Figure 1) according to the following specifications.

System Design

Following the basic tenet of cloud design, we partitioned the cloud design into four functional blocks: infrastructure control for managing the cloud, and storage, network, and compute elements for hosting cloud-based services. Table 1 describes in detail how we distributed the servers to handle the various functional blocks. For this preliminary testing, we used a simple network configuration with a single HP ProCurve 2910al-24G* 1 GbE switch (Figure 4).

Further testing with performance evaluation will be conducted in the near future and it will utilize a combination of 10 GbE and 1 GbE switches with isolated storage, and public and administrative networks (Figure 4).

	Hostname	Network
Storage Server	XNODE1	172.17.86.1
Enomaly ECP* Master Node	XNODE3	172.17.86.3
Enomaly ECP Node 01	XNODE2	172.17.86.2
Enomaly ECP Node 02	XNODE4	172.17.86.4
Enomaly ECP Node 03	XNODE5	172.17.86.5

Table 2. Network IP Configuration (Network: 172.17.80.0/20 (255.255.240.0))

Our design reserves static IP addresses for nodes that participate in hosting compute and storage: XNODE3, the Enomaly ECP master, XNODE2, XNODE4, XNODE5, the Enomaly ECP agent nodes, and XNODE1, the Enomaly ECP storage node.

System	Processor Configuration	Detail Configuration
1 Enomaly ECP* Master, 2 Enomaly ECP Agent Nodes and Storage Server	HP ProLiant SL170z* G6 server[8] Intel® Xeon® processor X5670	▪ Form Factor: 1U Rack Mount Server ▪ Processor: Intel® Xeon® processor 5600 series 2.93 GHz; 2-way x 4 cores = 8 cores ▪ Memory: 32 GB RAM ▪ Storage: 4 x 500 GB HDD SATA2
1 Enomaly ECP Agent Nodes	HP ProLiant SL2x170z* G6 Server with Intel® Xeon® processor X5570	▪ Form Factor: 1U Rack Mount Server (half-width, 2 per 1U) ▪ Processor: Intel® Xeon® 5570 processor 2.8 GHz; 2-way ▪ Memory: 32 GB RAM ▪ Storage: 2x1 TB GB HDD SATA
Network	HP ProCurve 2910 al-24G	▪ Ports: 24x 1000BASE-T ports ▪ VLAN: 802.1Q; Port-based and tag-based VLAN

Table 1: Initial HP ProLiant SL Scalable Systems Cloud Test Bed System Configuration

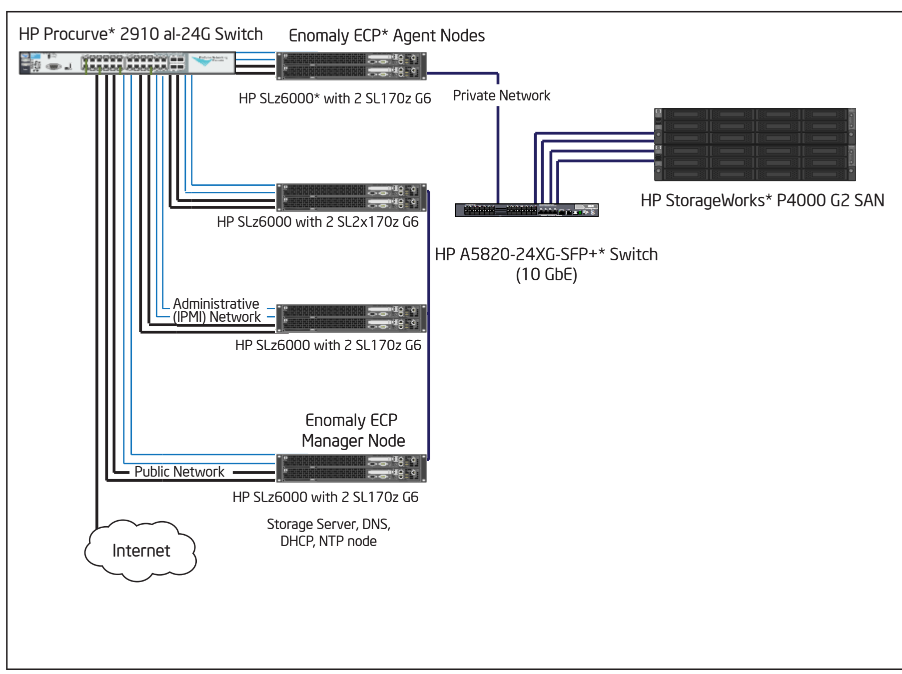

Figure 4: Forthcoming High Performance HP Networking Test Environment.

Technical Review

Installation Overview

This section discusses the installation process for the cloud infrastructure we created at the HP cloud lab. To set up the cloud infrastructure, we confirmed that there were sufficient hardware, software, memory, network, and storage resources, and we checked the overall network and dynamic host configuration protocol (DHCP) connectivity. We then configured the storage elements, switch and VLAN settings, installed the Enomaly ECP master node and four agent nodes, and verified the network/database connectivity. To test the newly configured Enomaly ECP environment, we created VMs and validated VM provisioning and management. For more specific details on each step, including specific device and server configuration parameters, please refer to the Enomaly ECP SPE and Intel test bed Installation Details document.[9]

Use Case Details

Actors

- Service Provider (SP)
- Service Consumers (SC1, SC2)

Use Case Overview

The cloud test bed App Center was provisioned with an instance of CentOS 5.4 (64-bit), and during the testing phase of the cloud implementation, we ran the following 14 use cases successfully.

Pre-conditions:

1. The cloud management software, Enomaly ECP SCE, is installed and ready to go.
2. Compute and storage nodes are installed and registered.
3. SC1 and SC2 (Service 1 and Service 2) application services are packaged in VMs to be deployed on the compute nodes.

Use Cases:

1. Package Applications Virtual Machines: Deploy application service packages.
2. Create Users: Ensure that admin can create two users (SC1, SC2) using the admin portal.
3. SC1and SC2: Create instance of the Service 1. Service 2: Instantiate VMs that make up Service 1 and Service 2, including IP address and links to storage.
4. SC1 and SC2: Monitor state of Service 1. Service 2: Observe the state of the newly created VMs.
5. Clone a VM: Ensure that two users (SC1, SC2) can clone VMs.
6. SC1: Scale-out Service 1. Service 2: Add an app front-end VM and add to load-balance pool. Test for application scalability.
7. Network Isolation: Verify network isolation among the VMs.
8. SC1: Terminate an app front-end VM. Remove from load-balance pool. Terminate a VM and observe for results.
9. VM Configuration Management: Verify VM configurations by logging into the VMs belonging to SC1 and SC2.
10. Generate Billing Report: Verify that consumers of the cloud, SC1 and SC2 can examine billing records.
11. SP: Add bare-metal capacity to existing cluster. Add a new server to an existing cluster.
12. SC1 and SC2: Scale-out Service 1. Service 2: Test for vertical application scalability on both Service 1 and Service 2.
13. SP: Fail a server. Ensure that the cluster is still operational.
14. SC1 and SC2: Shutdown Service 1. Service 2: End the application service and remote access to users SC1 and SC2.

Execution and Results

All the test cases described above were successfully executed at an HP lab facility, by HP staff, with support assistance from Enomaly, Inc. Detailed documentation of each of the test case inputs, as well as system responses, including screen shots of each step, are shown in the following subsections. All use cases are performed through the self-service web UI portal unless otherwise stated.

Setup and Software Verification

1. Login as "admin" through the Web UI portal (for our setup: http://172.17.86.3).

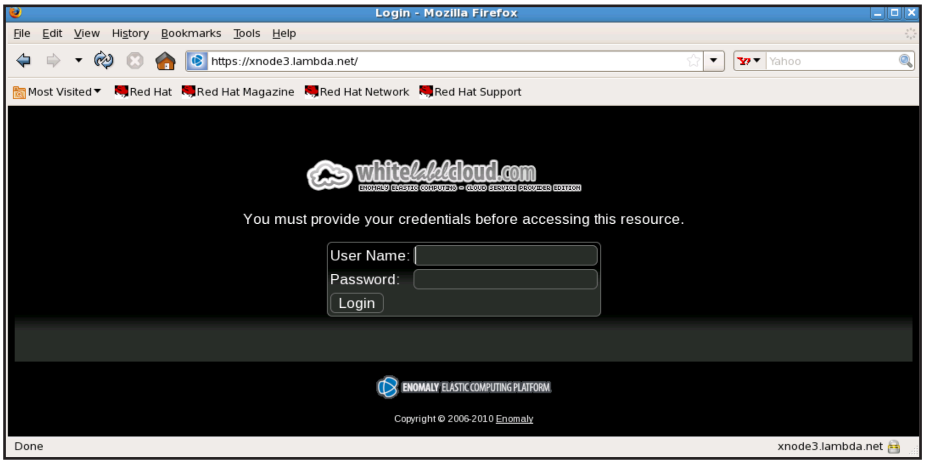

Figure 5

2. Verify that the cluster is operational by monitoring the dashboard that shows the CPU usage with aggregate cluster load and active transactions.

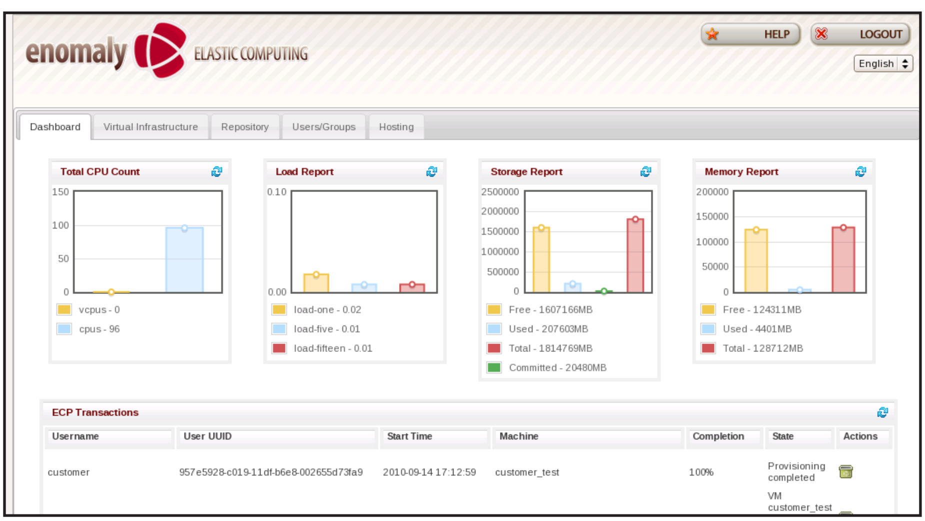

Figure 6

3. Logout as "admin" and login as "customer" to access the "App Center" home page.

Figure 7

Install a VM with the "CentOS 5.4 Gold" VM template provided by Enomaly ECP SPE 3.3.

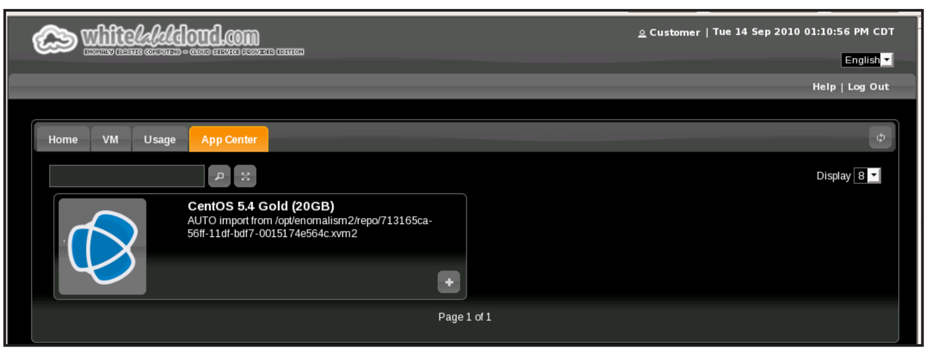

Figure 8

Our test used these values: VM Name: "test," Hardware Template: "Medium," and default settings for other options. The "Home" page can be used to monitor VM installation progress.

4. The "VM" page should be used to verify that the VM "test" was installed and powered on automatically.

Figure 9

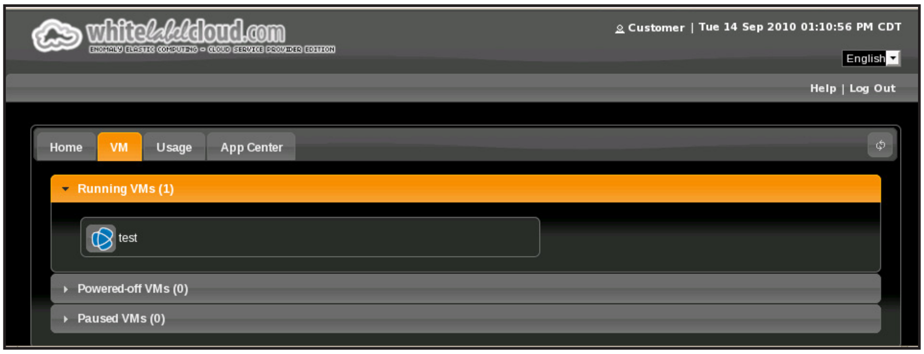

Figure 10

5. Open a VNC connection to the newly created VM with the information provided on the "Details" page and log into VM with the account/password.

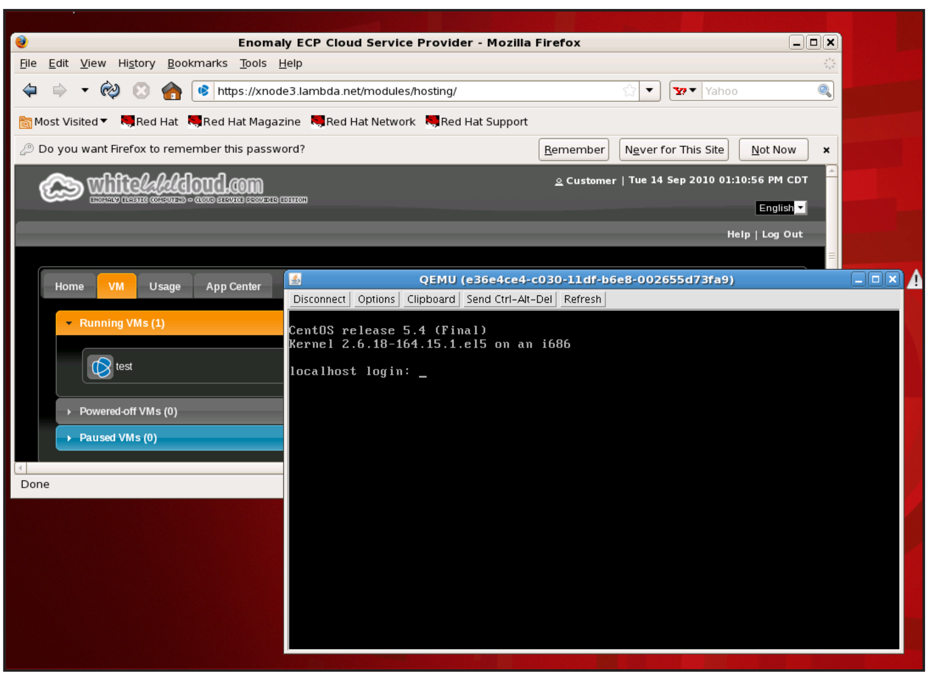

Figure 11

6. Power off VM "test." The newly created VM, "test," should be moved from group "Running VMs" to group "Powered-off VMs" on "VM" page.

Create a Virtual Machine Image for Use

1. Copy a bootable installation ISO image file for your operating system to "/opt/enomalism2/iso/" directory of Enomaly ECP master node. For this test case, we used Red Hat Enterprise* Linux* 5.5 installation media, with an ISO file named "RHEL5.5-Server-20100322.0-x86_64-DVD.iso."

2. Login as "admin" and connect to the Web UI portal. Go to "Repository/VM Creator" page and create a VM appliance. Choose File on Server: "RHEL5.5-Server-20100322.0-x86_64-DVD.iso," VM Type: "KVM Machine," RAM: 2048 and default settings for other options.

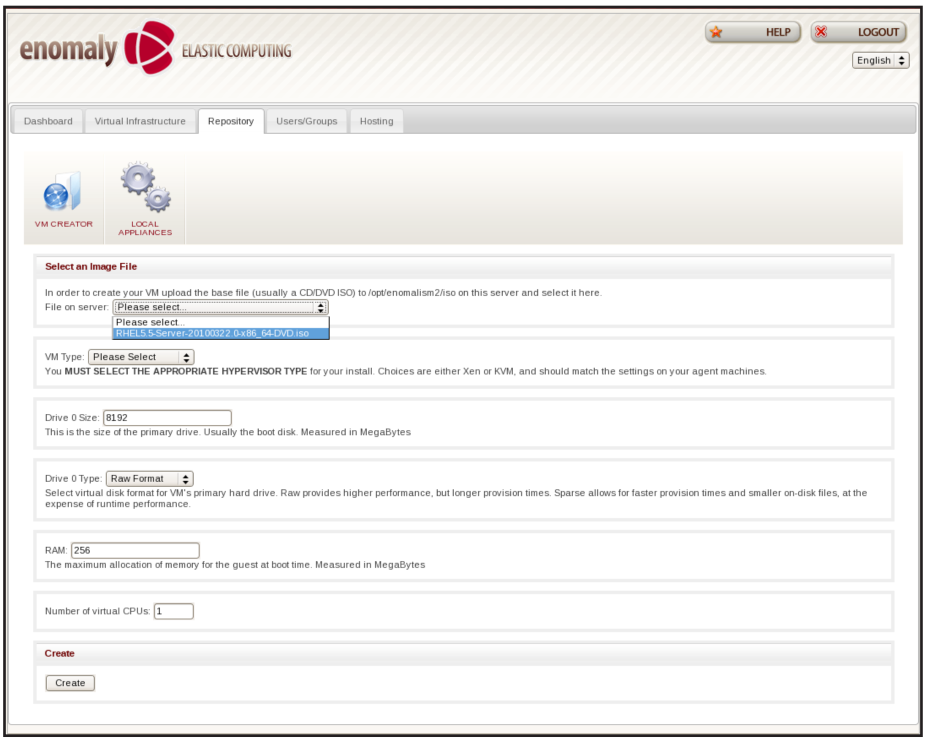

Figure 12

3. After the VM appliance is created successfully, it will be listed in "Repository/LOCAL APPLIANCES." Rename it to "RHEL5.5 ISO."
4. Go to "Virtual Infrastructure/ELASTIC VALET" page and provision a new VM; package to provision: "RHEL5.5 ISO" and default settings for other options.
5. Go to "Virtual Infrastructure/INFRASTRUCTURE" page and see the newly provisioned VM. Power on the VM and open a VNC console connection to the VM. Enter the Red Hat Linux installation screen. Install the OS within the VM and customize it as desired, obtaining a freshly installed Red Hat Enterprise Linux 5.5 system.

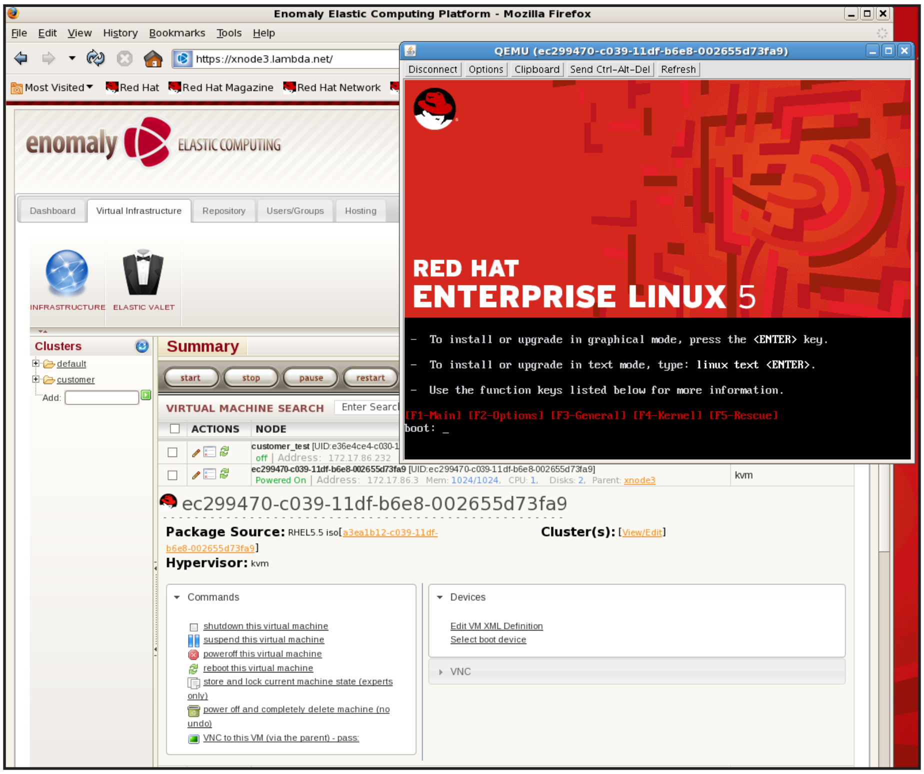

Figure 13

After installation and prior to rebooting, select "Edit the VM XML Definition" and remove the <disk device="cdrom" ..> element so that the VM will boot from the virtual hard drive.

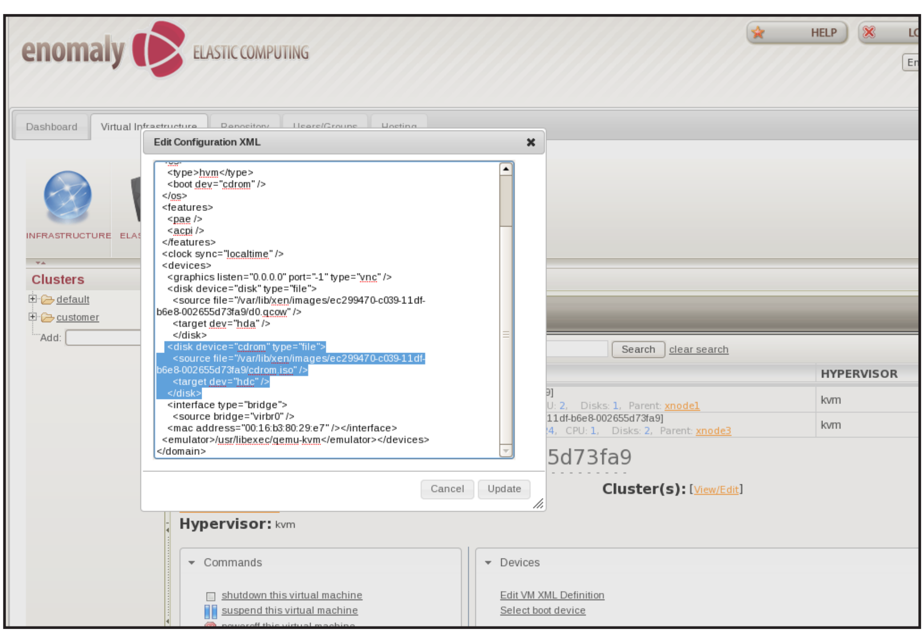

Figure 14

Power off the VM and then select "Package the VM." After the packaging operation completes successfully, the new VM will be listed on the "Repository/LOCAL APPLIANCES" page renamed as "RHEL5.5 Installed."

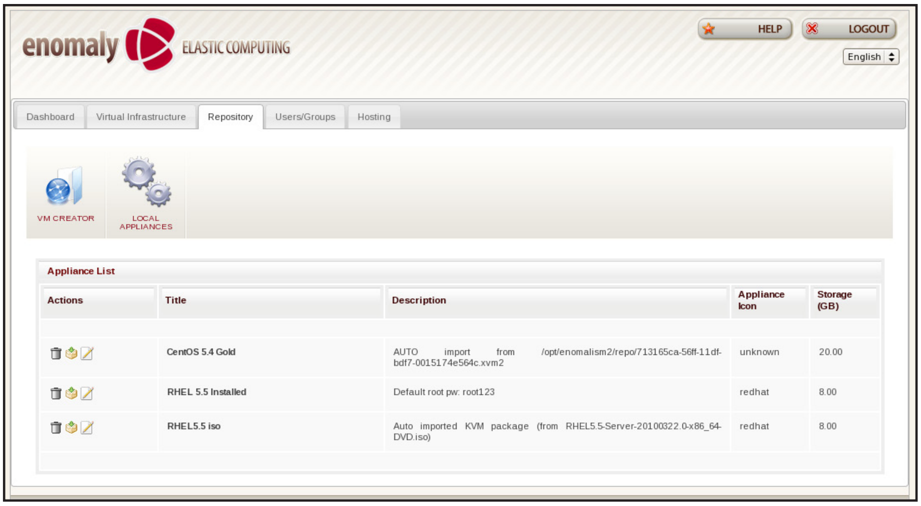

Figure 15

6. Go to the "Repository/LOCAL APPLIANCES" page and install the "RHEL5.5 Installed" appliance. After the operation completes successfully, it will be listed in the "Repository/LOCAL APPLIANCES" page, and will now be available to end users according the access controls defined on the Permissions page.

Create Users

1. From the web UI portal, go to "Hosting/QUOTA" page and create two Quotas.

Quota Q1:

Max VMs:	10
Max CPUs:	20
Max Mem:	32768
Max HD:	500000

Quota Q2:

Max VMs:	20
Max CPUs:	40
Max Mem:	65536
Max HD:	1000000

Table 3

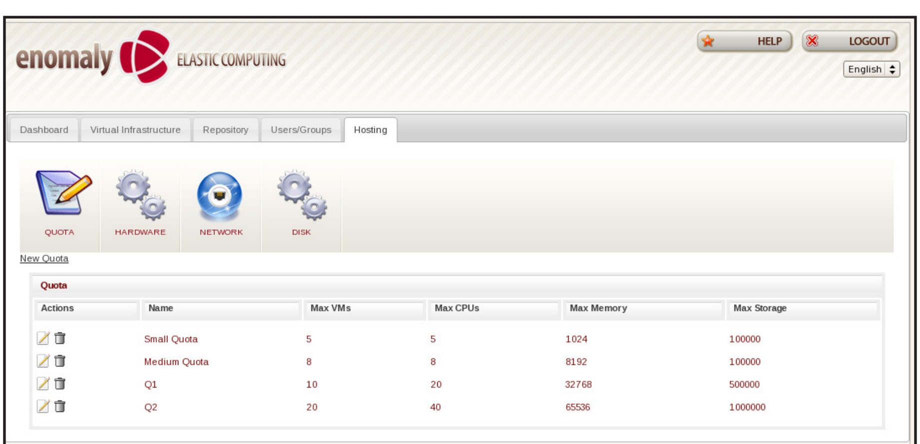

Figure 16

2. Go to “Users/Groups” page, and create two Groups and two Users, and assign Quotas to Groups.

Create Group G1	Group Name: G1	Display Name: G1
Create Group G2	Group Name: G2	Display Name: G2
Create User SC1	Username: SC1	Display Name: SC1

Table 4

Select Groups: G1 default settings for other options Assign Quota Q1

Create User SC2; Username: SC2; Display Name: SC2

Select Groups: G2 default settings for other options Assign Quota Q2

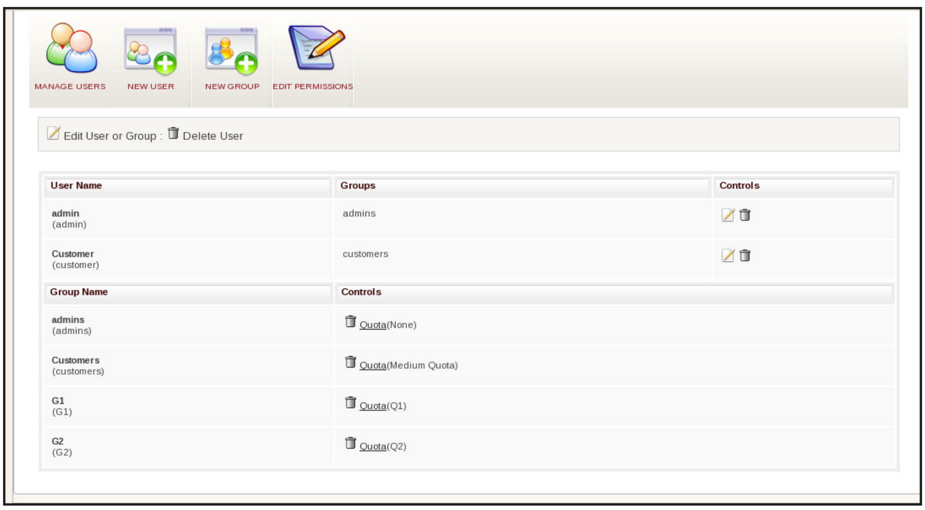

Figure 17

3. Go to “Hosting/HARDWARE” and create three hardware profiles.

Hardware Profile	CPU Architecture	Hypervisor Type	RAM (MB)	CPUs
Small	x86_64	kvm-hvm	2024	1
Medium	x86_64	kvm-hvm	2048	2
Large	x86_64	kvm-hvm	8192	8

Table 5

4. Go to “Hosting/NETWORK” page and create four VLANs.

VLAM Name	VLAN ID
VLAN100	100
VLAN200	200
VLAN300	300
VLAN400	400

Table 6

5. Go to “Users/Groups/EDIT PERMISSIONS” page and grant “read” Permissions to groups G1 and G2.

6. Logout and login as "SC1" and check "Usage" page to verify current resource usage and quota status.

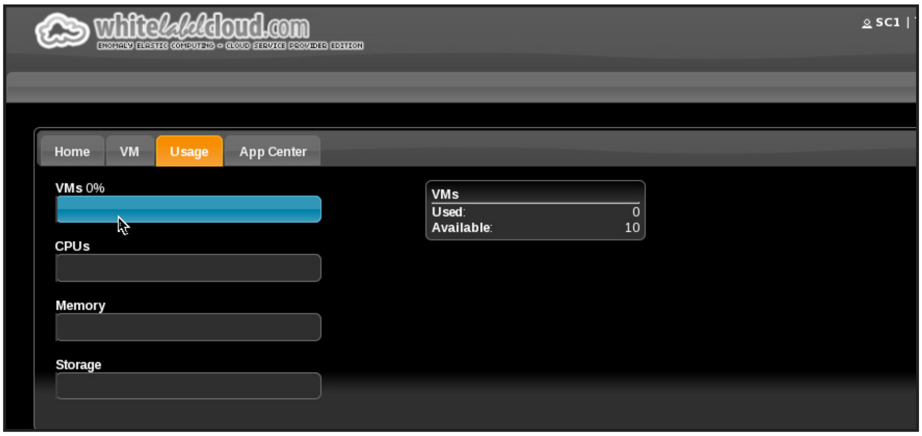

Figure 18

7. For the same user (SC1), verify the approved templates by visiting the "App Center" page.

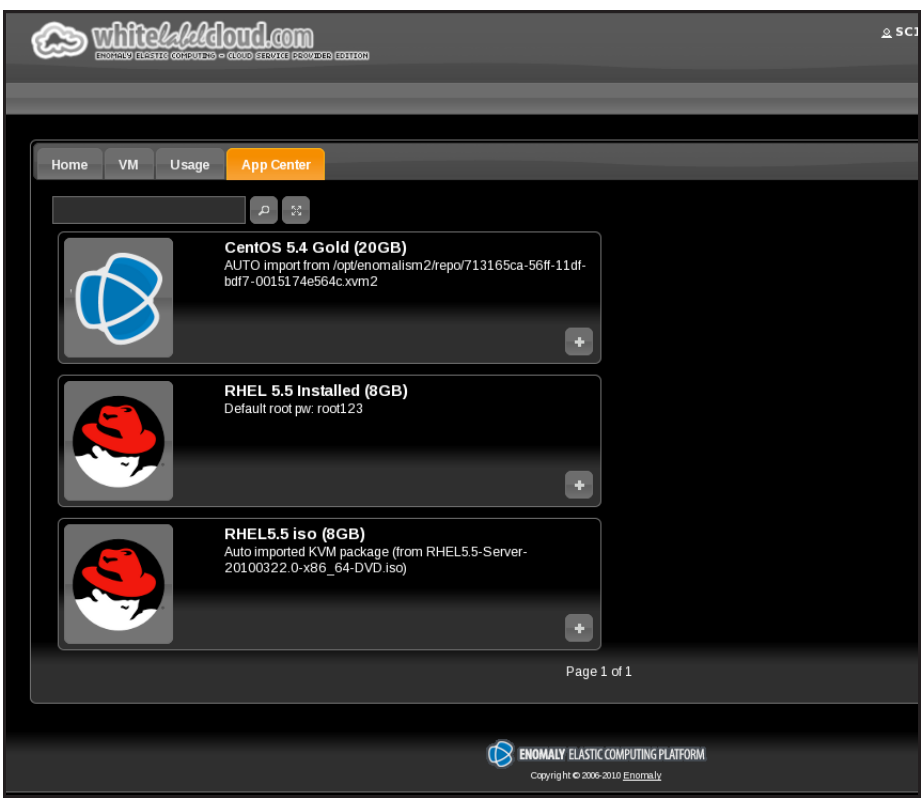

Figure 19

8. Logout and repeat the above two steps for user SC2 to verify quota and VM settings.

Create Service Instances by Multiple Users

1. Login in as user "SC1" and go to the "App Center" page and install a VM with the "CentOS54 Gold" VM template provided by Enomaly ECP SPE 3.3 using the "small" hardware template and VM name: "centos54-1." Verify that "centos54-1" was created successfully by accessing the VM page. By visiting the "Usage" page, we can verify that the resource usage/quota has been updated. Please refer to section 5.3.1 for the VM creation details.
2. Wait for about 1 minute, and then open the "Details" page of VM "centos54-1," go to the "Network" tab, and verify that the VM's IP address was provisioned successfully.

Figure 20

3. Open a VNC connection to VM "centos54-1." Log into the VM using the account information, and use ping to verify that there is a working network connection:

```
$ ping xnode3.lamba.net.
```

The ping should be successful.

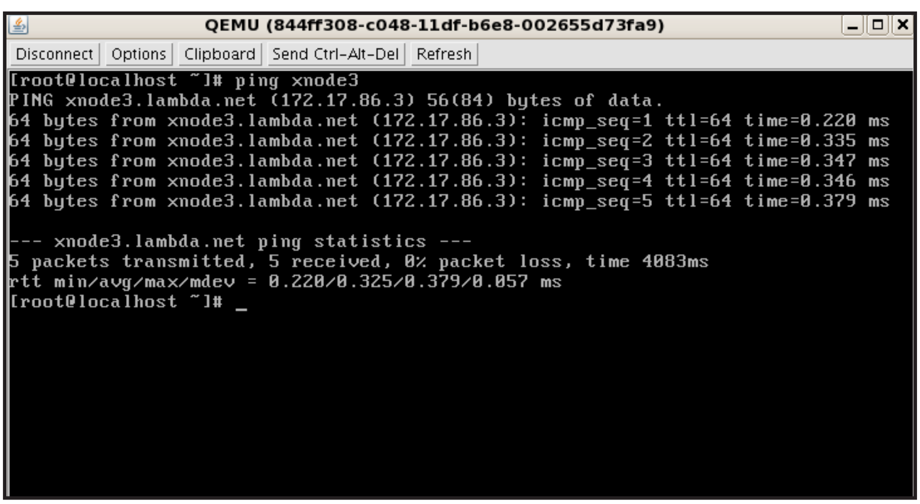

Figure 21

Monitor Services

1. Install additional VM as stated below and verify that the VMs are created successfully.

VM Name	Hardware Template
centos54 - 2	Medium
centos54 - 3	Large
centos54 - 4	Large

Table 7

2. Use the "App Center" page to install another VM with the "CentOS54 Gold" VM template (VM Name: "exceed," hardware template: "Large"), but get failure notification due to exceeding quota limits.

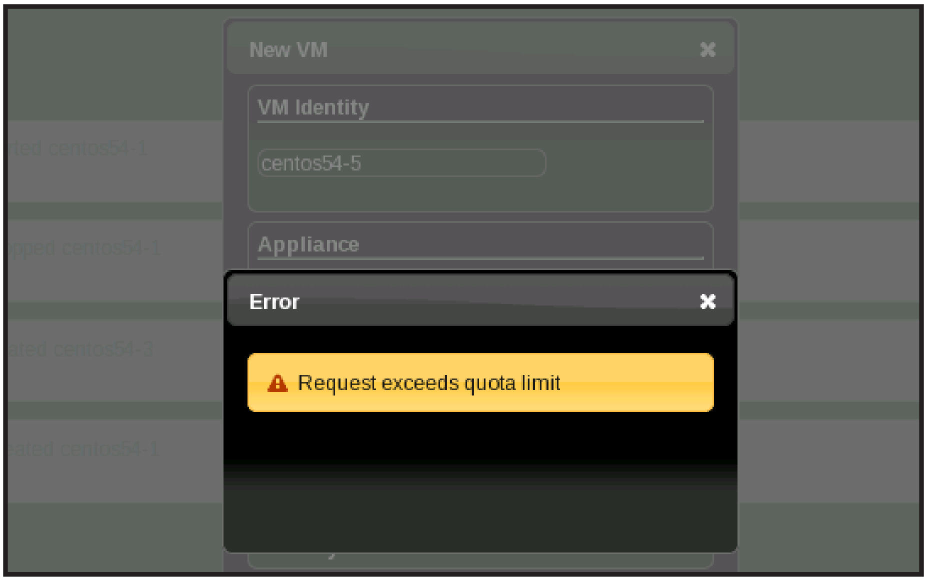

Figure 22

3. Logout as "SC1" and login as the user "SC2."
4. Go to "App Center" page and install a VM with the "rhel5.5 installed" VM template (VM Name: "rhel55-1," hardware template: "Small") and verify that the VM has been successfully created.
5. Open a VNC connection to VM "rhel55-1." Log into the VM using the account information, and use ping to verify that there is a working network connection:

```
$ ping xnode3.lamba.net.
```

The ping should be successful.

6. Repeat above steps to install more VMs:

VM Name	HardwareTemplate
rhel55-2	Medium
rhel55-3	Large
rhel55-4	Large
rhel55-5	Large
rhel55-6	Large

Table 8

7. Logout and login as "admin" and then check resource reports updated in the Dashboard tab. View VMware vSphere* VMs and templates to see organizations created.

Figure 23

Clone a VM by Multiple Users

1. Go to "VM" page and power off VM "centos54-1." After the VM is moved to group "Powered-off VMs," select to clone this VM. The cloned VM, which is named after the original VM name with suffix of current date/time, will be powered on automatically.

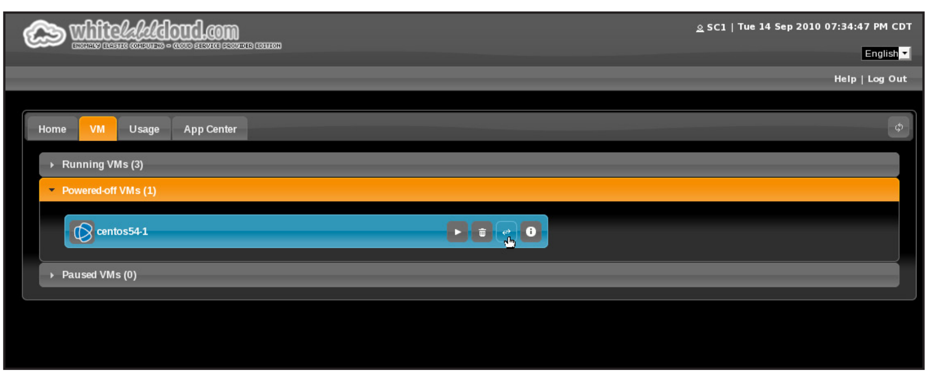

Figure 24

2. After successful cloning and powering the "centos54-1" and its cloned component, the VM page should reflect the cloned VMs.

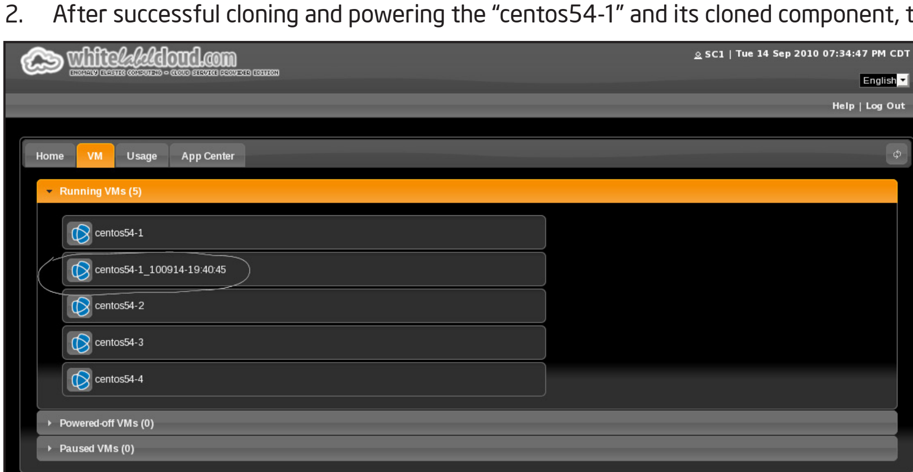

Figure 25

Network Isolation

1. Login as "SC1" user and power down all the VMs. Reconfigure the network settings of the VM "centos54-1." Access the Network tab. Reconfigure from network "default" to "VLAN100." Repeat the same for "centos54-2." Access the Network tab. Reconfigure from network "default" to "VLAN200."
2. Power on VMs: "centos54-1," "centos54-2," and the cloned VM of "centos54-1." Check that only the cloned VM of "centos54-1" can still get an IP address from DHCP server, as the other two VMs are isolated.
3. Open VNC connection to VM "centos54-1," login to the VM and reconfigure the network IP address using following command:

```
# ifconfig eth0 192.168.190.101 netmask 255.255.255.0
```

4. Open VNC connection to VM "centos54-2," login to the VM and reconfigure the network IP address using following command:

```
# ifconfig eth0 192.168.190.102 netmask 255.255.255.0
```

5. Logout and login as "SC2." Power down all the VMs. Open "Details" page of the VM "rhel55-1" and go to "Network" tab. Reconfigure it from network "default" to "VLAN100" and VM "rhel55-2" from network "default" to "VLAN200."
6. Power on VMs: "rhel55-1," "rhel55-2," and the cloned VM of "rhel55-1." Check that only the cloned VM of "rhel55-1" can still get IP from DHCP server, as other two VMs are isolated.
7. Open a VNC connection to VM "rhel55-1" and "rhel55-2," login to the VM, and reconfigure the network IP address to be unique.
8. Logout and Login as "SC1" user. Open a VNC connection to the VM, "centos54-1." Login to the VM and ping the other VMs as follows:

`$ ping centos54-2`	No replies since the VM is not in the same VLAN
`$ ping rhel55-1`	Get replies since this VM is in the same VLAN
`$ ping rhel55-2`	No replies since the VM is not in the same VLAN

Table 9

VM Configuration Management

1. Go to "Virtual Infrastructure/INFRASTRUCTURE" page, open the page for VM "SC2_rhel55-2" under "SC2/XNODE3" (your node name will probably be different), and change the VM's CPU/Memory/HD configuration. Change the VLAN from "default" to "VLAN200."

Billing for Multiple Users

1. Billing data is provided for programmatic access, but can be retrieved manually with a specially written URL containing the start date/time and ending date/time of interest. For example, access http://172.18.1.105:8080/modules/billingreports/?start_date=2010:02:09:11:00:00&end_date=2010:02:11:11:00:00 to get billing information for the period from 11:00.00am on Feb. 9, 2010 until 11:00.00am on Feb. 11, 2010.

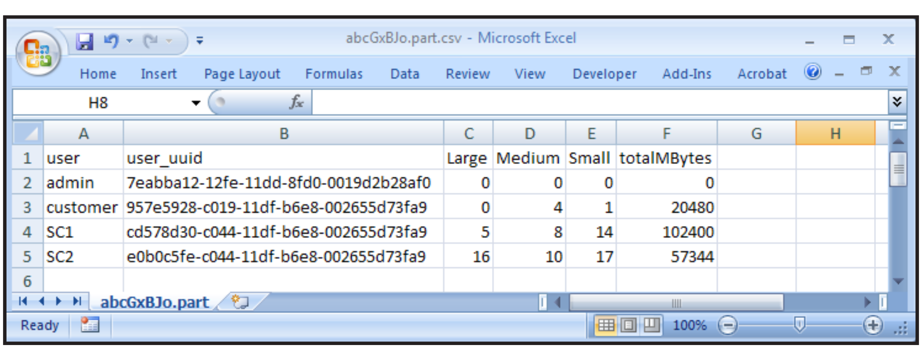

user	user_uuid	Large	Medium	Small	totalMBytes
admin	7eabba12-12fe-11dd-8fd0-0019d2b28af0	0	0	0	0
customer	957e5928-c019-11df-b6e8-002655d73fa9	0	4	1	20480
SC1	cd578d30-c044-11df-b6e8-002655d73fa9	5	8	14	102400
SC2	e0b0c5fe-c044-11df-b6e8-002655d73fa9	16	10	17	57344

Figure 26

2. Save the output as a comma separated value file (.csv).

Add New Servers

1. Access “Virtual Infrastructure/INFRASTRUCTURE” page, and check existing Enomaly ECP nodes. Four Enomaly ECP nodes, XNODE1, XNODE2, XNODE3, and XNODE4 should be visible.

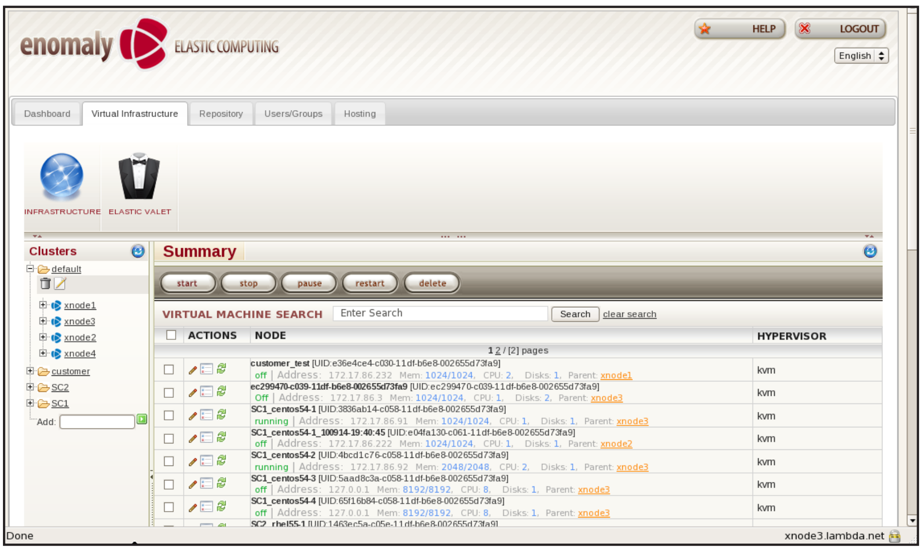

Figure 27

2. Login to xnode5 and start to install an Enomaly ECP SPE 3.3 agent node installation (Refer to 5.3.1.2 for more details). Access the “Virtual Infrastructure/INFRASTRUCTURE” page on the Web UI portal. XNODE5 should have been added to the existing Enomaly ECP nodes.

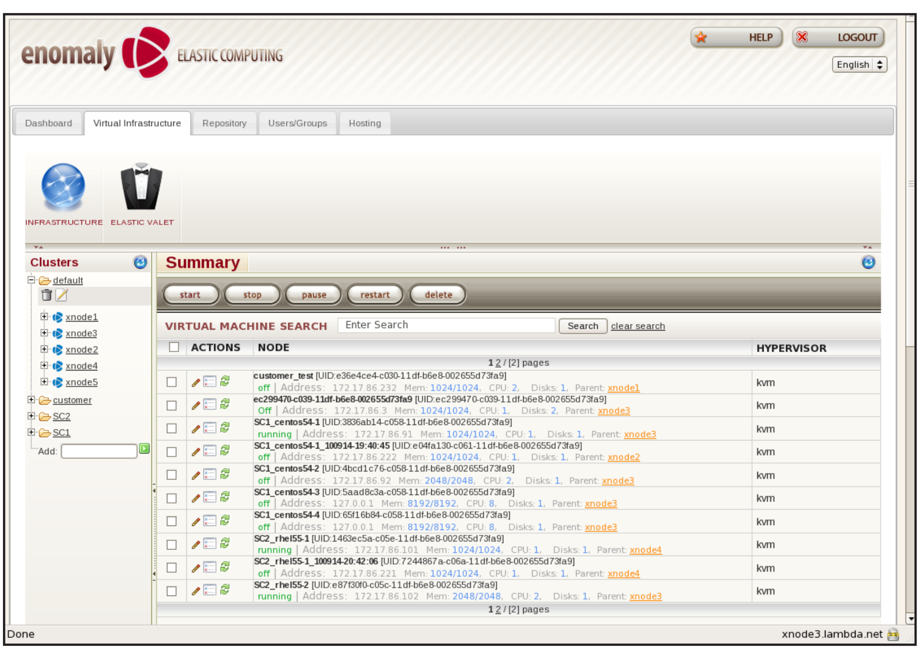

Figure 28

Scale Up and Down Services

1. Login as “SC1” through the Web UI. Access the “VM” page, power on VM “centos54-2,” open its “Details” page, and go to the “Usage” tab to check the VM’s current CPU/Memory/HD configuration.

Figure 29

2. Power off VM “centos54-2,” open its “Details” page, and change its hardware template from “Medium” to “Large.”

Hardware Profile	CPU Architecture	Hypervisor Type	RAM (MB)	CPUs
Small	x86_64	kvm-hvm	2024	1
Medium	x86_64	kvm-hvm	2048	2
Large	x86_64	kvm-hvm	8192	8

Table 10

3. Power on VM “centos54-2,” open its “Details” page and go to the “Usage” tab to check the VM’s current CPU/Memory/HD configuration.

Figure 30

Log into the VM and verify that the virtualized hardware configuration corresponds to what you specified in Enomaly ECP:

```
# cat /proc/cpuinfo > cpuinfo.info
```

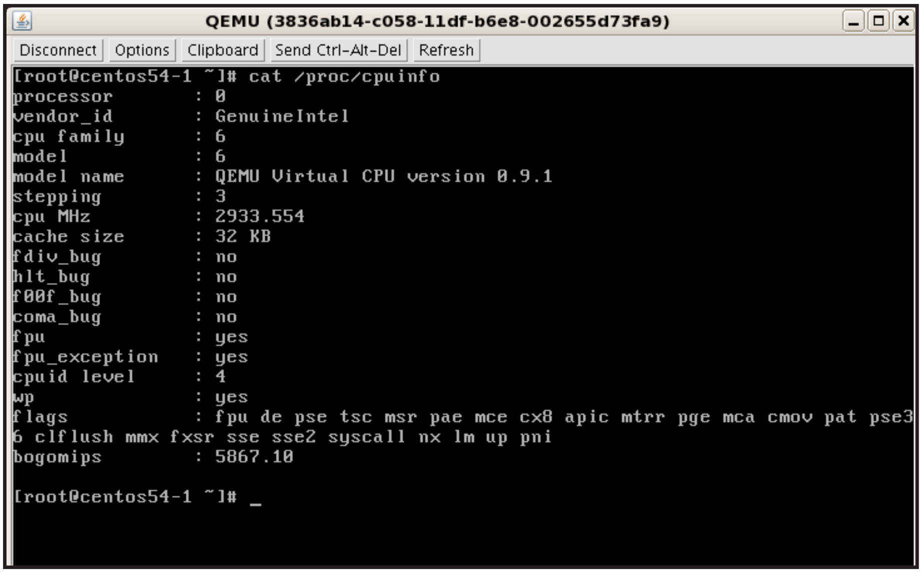

Figure 31

```
# cat /proc/meminfo > meminfo.info
```

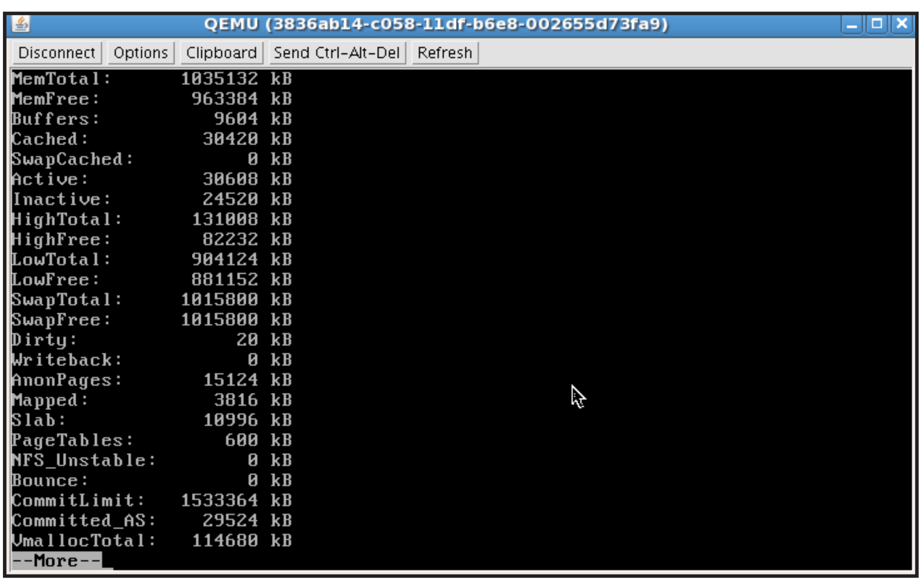

Figure 32: Partial Output from meminfo.info

Failure Scenario

1. Open a SSH tunnel to the Enomaly ECP master node with command:

```
# ssh -L 8081:127.0.0.1:8081 root@172.17.86.3
```

(The equivalent Putty (http://www.chiark.greenend.org.uk/~sgtatham/putty/) command in MS Windows environments is: plink -ssh -L 8081:localhost:8081 root@172.17.86.3)

2. Connect to Enomaly ECP Manager Status Report https://localhost:8081, and check existing Enomaly ECP nodes. All should be active.

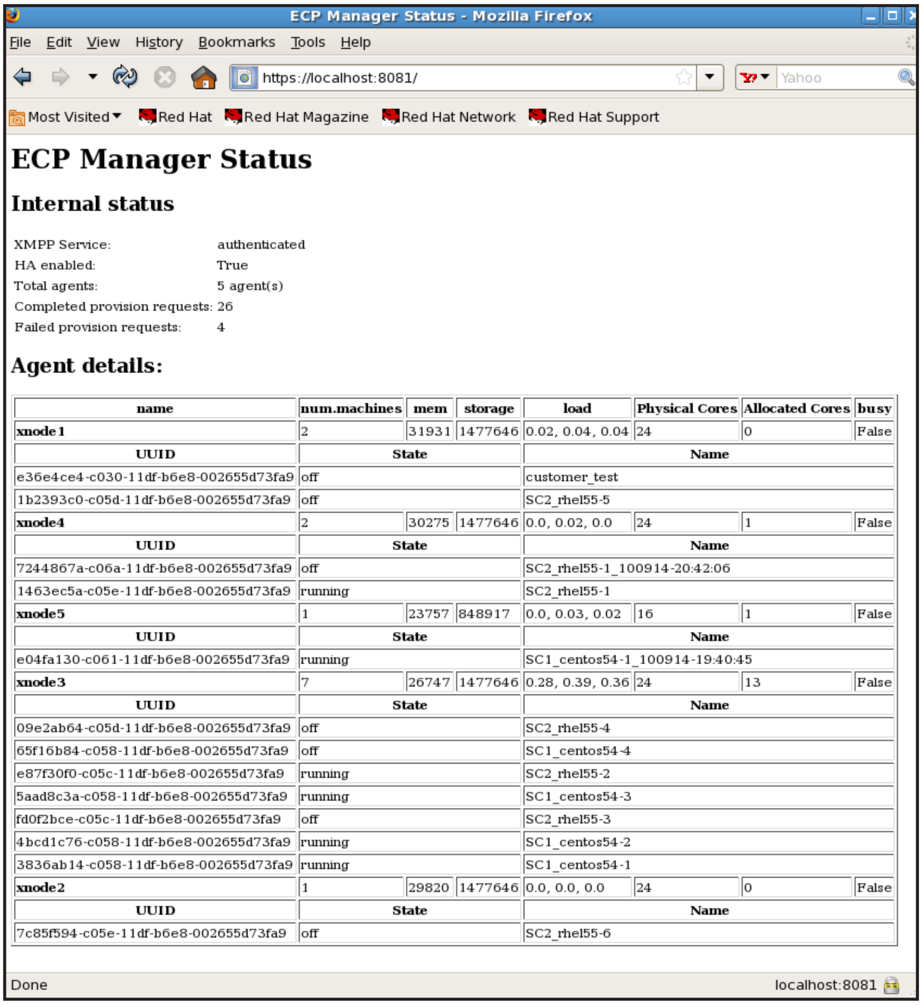

ECP Manager Status

Internal status

XMPP Service:	authenticated
HA enabled:	True
Total agents:	5 agent(s)
Completed provision requests:	26
Failed provision requests:	4

Agent details:

name	num.machines	mem	storage	load	Physical Cores	Allocated Cores	busy
xnode1	2	31931	1477646	0.02, 0.04, 0.04	24	0	False
UUID	**State**			**Name**			
e36e4ce4-c030-11df-b6e8-002655d73fa9	off			customer_test			
1b2393c0-c05d-11df-b6e8-002655d73fa9	off			SC2_rhel55-5			
xnode4	2	30275	1477646	0.0, 0.02, 0.0	24	1	False
UUID	**State**			**Name**			
7244867a-c06a-11df-b6e8-002655d73fa9	off			SC2_rhel55-1_100914-20:42:06			
1463ec5a-c05e-11df-b6e8-002655d73fa9	running			SC2_rhel55-1			
xnode5	1	23757	848917	0.0, 0.03, 0.02	16	1	False
UUID	**State**			**Name**			
e04fa130-c061-11df-b6e8-002655d73fa9	running			SC1_centos54-1_100914-19:40:45			
xnode3	7	26747	1477646	0.28, 0.39, 0.36	24	13	False
UUID	**State**			**Name**			
09e2ab64-c05d-11df-b6e8-002655d73fa9	off			SC2_rhel55-4			
65f16b84-c058-11df-b6e8-002655d73fa9	off			SC1_centos54-4			
e87f30f0-c05c-11df-b6e8-002655d73fa9	running			SC2_rhel55-2			
5aad8c3a-c058-11df-b6e8-002655d73fa9	running			SC1_centos54-3			
fd0f2bce-c05c-11df-b6e8-002655d73fa9	off			SC2_rhel55-3			
4bcd1c76-c058-11df-b6e8-002655d73fa9	running			SC1_centos54-2			
3836ab14-c058-11df-b6e8-002655d73fa9	running			SC1_centos54-1			
xnode2	1	29820	1477646	0.0, 0.0, 0.0	24	0	False
UUID	**State**			**Name**			
7c85f594-c05e-11df-b6e8-002655d73fa9	off			SC2_rhel55-6			

Figure 33

3. Reboot XNODE4. XNODE4 should be marked with strikeout and the VMs running on XNODE4 changed from running state to unknown state in about 10 seconds.
4. As XNODE4 comes back to operational state in about 5 minutes, the strikeout mark on XNODE4 should be removed, and all VMs running on XNODE4 moved to another node and changed to the running state.

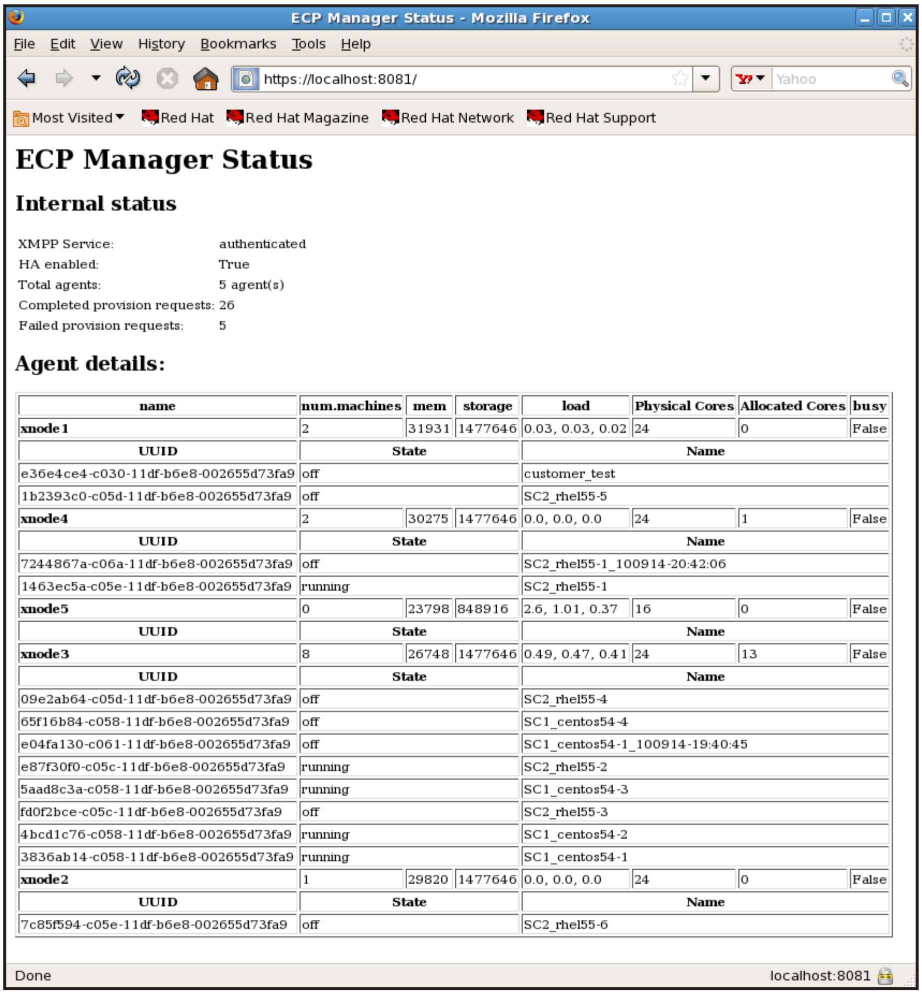

name	num.machines	mem	storage	load	Physical Cores	Allocated Cores	busy
xnode1	2	31931	1477646	0.03, 0.03, 0.02	24	0	False
UUID	State			Name			
e36e4ce4-c030-11df-b6e8-002655d73fa9	off			customer_test			
1b2393c0-c05d-11df-b6e8-002655d73fa9	off			SC2_rhel55-5			
xnode4	2	30275	1477646	0.0, 0.0, 0.0	24	1	False
UUID	State			Name			
7244867a-c06a-11df-b6e8-002655d73fa9	off			SC2_rhel55-1_100914-20:42:06			
1463ec5a-c05e-11df-b6e8-002655d73fa9	running			SC2_rhel55-1			
xnode5	0	23798	848916	2.6, 1.01, 0.37	16	0	False
UUID	State			Name			
xnode3	8	26748	1477646	0.49, 0.47, 0.41	24	13	False
UUID	State			Name			
09e2ab64-c05d-11df-b6e8-002655d73fa9	off			SC2_rhel55-4			
65f16b84-c058-11df-b6e8-002655d73fa9	off			SC1_centos54-4			
e04fa130-c061-11df-b6e8-002655d73fa9	off			SC1_centos54-1_100914-19:40:45			
e87f30f0-c05c-11df-b6e8-002655d73fa9	running			SC2_rhel55-2			
5aad8c3a-c058-11df-b6e8-002655d73fa9	running			SC1_centos54-3			
fd0f2bce-c05c-11df-b6e8-002655d73fa9	off			SC2_rhel55-3			
4bcd1c76-c058-11df-b6e8-002655d73fa9	running			SC1_centos54-2			
3836ab14-c058-11df-b6e8-002655d73fa9	running			SC1_centos54-1			
xnode2	1	29820	1477646	0.0, 0.0, 0.0	24	0	False
UUID	State			Name			
7c85f594-c05e-11df-b6e8-002655d73fa9	off			SC2_rhel55-6			

Figure 34

5. Reboot Enomaly ECP master with the following observations:
   - The Web UI portal and Enomaly ECP Manager Status Report were unavailable until the Enomaly ECP master came back to operational state.
   - Enomaly ECP Agent nodes should be still running. Verify by pinging VMs running on nodes.

Shutdown Services

1. Login as “SC1” through the Web UI portal http://172.18.1.105:8080.
2. Power off all VMs except “centos54-1” (Small) and “centos54-2” (Medium). The snapshot gives the information that only 15 percent of the CPU quota is currently being used.

Figure 35

3. Power off VMs "centos54-1" (Small) and "centos54-2" (Medium).
4. Go to the "Usage" page and check resource usage/quota. The CPU quota will drop down to zero percent.

Observations and Analysis

Enomaly ECP SPE provides a solution that enables a service provider (either a customer-facing, revenue-generating business unit, or an IT service organization within an enterprise) to offer a self-service Infrastructure on Demand solution to its customers or users. The test cases above illustrate the core capabilities of the platform for this purpose.

Stress Testing the Environment

Verifying Failure Modes

While using the HA module of Enomaly ECP SPE, various failure modes result in automatic recovery of VMs and hosts. It is necessary to thoroughly stress test any hosting environment prior to moving into production. As a minimum, our recommendation is to at least test the following scenarios:

a. Reboot Host Machine (Loss of Host < 10 minutes):
   1. Observe the Manager Status report at https://<primary host name>:8081.
   2. Notice that host machine has strikeout through name and VMs move to unknown state.
   3. When host returns, notice that the agent will remove strikeout on name, and VMs will show "Off" state.
   4. Observe that all VMs will restart after a short period of time.

b. Power Failure on Host Machine (Loss of Host > 10 minutes):
   1. Observe the Manager Status report at https://<primary host name>:8081.
   2. Notice that after 10 minutes, host machine has strikeout through name and VMs move to unknown state.
   3. Note that all VMs were moved to other hosts and put into "running" state.
   4. If host returns, notice that the strikeout will be removed and agent will show no running VMs, but will be ready to provision jobs.

c. Host Intermittent Network Loss (Loss of Network < 10 minutes):
   1. Observe the Manager Status report at https://<primary host name>:8081.
   2. Notice that agent will not strikeout, as failure has not been detected.
   3. Note that all VMs cannot be pinged from this host until the VMs resume operation after network is restored.

d. Host Permanent Network Loss (Loss of Network > 10 minutes):
   1. Observe the Manager Status report at https://<primary host name>:8081.
   2. Notice that after 10 minutes, host machine has strikeout through name and VMs move to unknown state.
   3. Note that all VMs moved to other hosts and put into "running" state.

4. If network returns, notice that strikeout will be removed and agent will show no running VMs, but will be ready to provision jobs.

e. Primary Host Loss:

1. Observe that no web UIs will respond in this state, however VMs will continue to run on all hosts.
2. When primary host returns to service, notice that all command/control will return to administrators and users.

Verifying Performance

A major performance criterion to test in an Enomaly ECP cluster is the Disk I/O available to the VMs. This can be tested by doing the following:

1. Install Bonnie++ on each host to test the disk I/O to the NFS share. You must have a baseline value of performance from your /var/lib/xen shared directory. If any host is showing wildly varying values, NFS tuning may be required. This test should be performed sequentially and not in parallel.
2. Install Bonnie++ inside a VM. Determine if the values returned are within 60 to 80 percent of the values off the host machine itself. Overhead loss is to be expected, but lower than 60 percent will indicate a configuration issue.

Bonnie++ is an open source benchmark suite that is aimed at performing a number of simple tests of hard drive and file system performance. The benchmark can be downloaded from www.coker.com.au/bonnie++/.

Next Steps

The previous design implements a multi-tenant cloud. However, this design has only limited means to ensure that the software and BIOS on each node are in fact the exact versions that are approved. With SL170Z Proliant servers running the

Intel Xeon processor 5600 series, the trusted multi-tenancy technology is in place to be able to measure the key components of the BIOS and the hypervisor prior to 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:

1. Intel® Trusted Execution Technology (Intel® TXT)[10]: 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.
2. Intel® Advanced Encryption Standard – New Instructions (AES-NI)[11]: Intel AES-NI are a new set of instructions available on Intel Xeon processor 5600 series based on the 32nm Intel® micro-architecture. 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.
3. Enomaly ECP High Assurance Edition: The High Assurance Edition extends the feature set of Enomaly ECP, Service Provider Edition, with a unique set of high-security capabilities well suited to meet the needs of customers who require a higher level of security than that offered by any of the commodity cloud computing services available in the marketplace. Enomaly provides a true Trusted Cloud platform, with continuous security assurance by means of unique hardware-assisted mechanisms. The customer can be assured, for example, that:
   - Hardware has not been modified to duplicate data to some storage medium of which the customer is unaware
   - The hypervisor has not been modified to copy memory state or VM images
   - No hostile kernel modules have been injected into the guest OS
4. Intel® Intelligent Power Node Manager and Intel® Data Center Manager are used for power management and coordinate with policies to optimize power usage. These tools can allow for power capping, which avoids circuit breaker trips or hot spots in the data center. Please refer to http://software.intel.com/sites/datacentermanager/intel_node_manager_v2e.pdf to learn about Baidu’s case study utilizing the Intel® Intelligent Power Node Manager.

5. Test with improved HP shared storage products, including the HP StorageWorks X9000 Network Storage[12] and HP P4000 G2 SAN[13] storage solutions.

Things to Consider

Scalability

The scalability of the cloud solution could be impacted by:

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

Another interesting option to consider is the use of solid-state drives (SSDs) as hard disk replacements, since SSDs can significantly improve performance in a cloud. In addition, when you plan cloud implementations, security should be of primary concern. These topics are each discussed briefly below.

Network Technology Architecture

Enomaly ECP SPE 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, channel bonding technologies, and highly segmented network architectures) will be more suitable for production deployments.

Storage Architecture

Enomaly ECP SPE supports several storage architectures. For this test bed we selected the simplest option, a single NFS. This performed acceptably for our purposes, but more advanced architectures are recommended for production deployments.

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 factors of processor architecture, and specific feature sets available in the processor such as Intel® VT-d. 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 technologies, please refer to www.intel.com/technology/virtualization/ and download.intel.com/business/resources/briefs/xeon5500/xeon_5500_virtualization.pdf.

Solid-State Drives

Security is a key consideration in the selection and management of IaaS. A complete discussion of best practices for cloud security, from the perspective of both the Service Provider and the end-user organization, is beyond the scope of this document. However, the following points should be considered:

- Established best practices for host security in a conventional physical-host context (e.g., password management, patch management, server hardening, anti-malware, etc.) should be applied equally to virtual hosts operating on an IaaS platform.
- IaaS platforms such as Enomaly ECP provide full isolation between virtual servers by employing full hardware-assisted virtualization. This provides each virtual server with its own virtual hardware, its own private operating system instance, etc. This contrasts with cloud platforms based on “domains” or “containers,” in which some virtual hardware and some operating system components are shared between virtual hosts, which creates additional avenues of attack.
- Enomaly ECP can further segregate the virtual hosts belonging to different customers at the level of the network, which isolates the network traffic of each customer into one or more private VLANs.

Enomaly ECP, High Assurance Edition, uses the Intel® Trusted Execution Technology (Intel® TXT) capabilities of SL170Z Proliant servers running Intel® Xeon® processor 5600 series to deliver an IaaS environment that is strongly protected against hacking, tampering, and unauthorized administrative changes.

Additionally, to meet the highest security requirements, Enomaly provides a High Assurance Edition (HAE) of the Enomaly ECP platform.

Additional Info

Intel Cloud Builders: http://www.intel.com/cloudbuilders

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

Enomaly: www.enomaly.com

HP ProLiant Servers: www.hp.com/go/proliant

Glossary

Cloud Computing: An evolution in IT consumption and delivery made available on a self-service basis through the Internet with a flexible, pay-as-you-go business model.

Domain: An instance of an operating system (or subsystem in the case of container virtualization) that runs on a virtualized machine provided by the hypervisor.

Enomaly Elastic Computing Platform (ECP): A software product comprised of a cloud computing platform for service providers or enterprises.

HA: High Availability

Hypervisor: A layer of software that allows the virtualization of a node in a set of virtual machines, which supports one or more virtual machines with possibly different configurations than the node itself.

LVM: Logical Volume Manager

Node or Host: A single physical machine.

Pool: Provides a means to take a chunk of storage and carve it up into volumes. A pool can be used to manage things such as a physical disk, a NFS, an iSCSI target, a host adapter, or an LVM group.

VM: Virtual Machine

Volume: A single storage resource which can be assigned to a guest, or used to create further pools. A volume could be a block device, a raw file, or a special format file.

Endnotes

1. HP ProLiant Servers, http://h10010.www1.hp.com/wwpc/us/en/sm/WF02a/15351-15351-3896136.html

2. Intel Xeon 5600 series product information: http://ark.intel.com/ProductCollection.aspx?series=47915

3. Enomaly ECP SPE Datasheet, www.enomaly.com/Cloud-Hosting-Pr.465.0.html; Enomaly ECP Architecture Whitepaper, www.enomaly.com/Cloud-Hosting-Pr.466.0.html

4. For more information and screen shots visit our Web site at www.enomaly.com/Cloud-Service-Pr.cloudhosting.0.html

5. Enomaly ECP HAE Datasheet, www.enomaly.com/Cloud-Hosting-Pr.465.0.html

6. Why the Intel® Xeon® Processor 5600 Series is the Ideal Foundation for Cloud Computing, communities.intel.com/docs/DOC-4213; Intel Xeon 5500 Series Software Industry Testimonials, www.intel.com/business/software/testimonials/xeon5500.htm

7. Intel in cloud computing Wiki, communities.intel.com/docs/DOC-4230

8. HP ProLiant SL Scalable System, http://h10010.www1.hp.com/wwpc/us/en/sm/WF02a/15351-15351-3896136.html

9. Enomaly ECP Intel Test bed Installation Details, enomaly.com/fileadmin/ICB/Enomaly_ECP_Intel_Test bed_Installation_Details.pdf

10. Intel® AESNI and Intel® Trusted Execution Technology, http://www.intel.com/technology/dataprotection/index.htm

11. See number 8.

12. HP StorageWorks X9000 Network Storage Systems, http://h10010.www1.hp.com/wwpc/us/en/sm/WF0a/12169-3798502-4059049-4059049-4059049-4058820.html

13. HP StorageWorks P4000 G2 SAN Solutions, http://h10010.www1.hp.com/wwpc/us/en/sm/WF2a/12169-304616-3930449-3930449-3930449-4118659.html

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

To learn more about HP ProLiant servers, visit http://www.hp.com/go/proliant

To learn more about Enomaly ECP SPE, visit http://www.enomaly.com/Cloud-Service-Pr.cloudhosting.0.html

Disclaimers

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

Internal Intel measurements for Intel® Xeon® processor X5680 vs. Intel® Xeon® processor X5570 on BlackScholes*

Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, Go to: http://www.intel.com/performance/resources/benchmark_limitations.htm

No computer system can provide absolute security under all conditions. Intel® Trusted Execution Technology is a security technology under development by Intel and requires for operation a computer system with Intel® Virtualization Technology, an Intel Trusted Execution Technology-enabled processor, chipset, BIOS, Authenticated Code Modules, and an Intel or other compatible measured virtual machine monitor. In addition, Intel Trusted Execution Technology requires the system to contain a TPMv1.2 as defined by the Trusted Computing Group and specific software for some uses. See http://www.intel.com/technology/security/ for more information.

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. Go to: http://www.intel.com/products/processor%5Fnumber/

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

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

Printed in USA 1010/SK/PRW/PDF Please Recycle 324422-001 US