

Intel® Cloud Builder Guide to Cloud Design and Deployment on Intel® Xeon® Processor-based Platforms

Enomaly Elastic Computing Platform,* Service Provider Edition

White Paper

Intel® Cloud
Builder Guide

Enomaly Elastic
Computing Platform*



Audience and Purpose of This Paper

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 Builder program to help any customer design, deploy, and manage a cloud infrastructure.

The Intel Cloud Builder program provides a starting point by supplying a basic hardware blueprint and available cloud software management solutions such as the Enomaly Elastic Computing Platform,* Service Provider Edition. 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 who are looking to realize the revenue potential of leveraging their existing data center infrastructure by offering cloud computing services to their customers or internal users.



Table of Contents

Executive Summary	3	Technical Review, continued	
Introduction	4	Network Isolation	12
Design Principles	4	VM Configuration Management	12
ECP Logical Architecture Configurations	5	Billing for Multiple Users	12
Enomaly ECP SPE Implementation Overview	5	Add New Servers	13
Enomaly's ECP Offering	5	Scale Up and Down Services	13
Scalable Design	6	Failure Scenario	14
Command and Control Architecture of Enomaly ECP	6	Shutdown Services	14
Tested Blueprint Overview	7	Observations and Analysis	14
Hardware Description	7	Stress Testing the Environment	14
Physical Architecture	7	Verifying Failure Modes	14
System Design	7	Verifying Performance	15
Technical Review	8	Next Steps	15
Installation Overview	8	Things to Consider	16
Use Case Details	8	Scalability	16
Actors	8	Network Technology Architecture	16
Use Case Overview	8	Storage Architecture	16
Execution and Results	9	Hardware Considerations	16
Setup and Software Verification	9	Solid-State Drives	16
Create a Virtual Machine Image for Use	10	Security Considerations	16
Create Users	10	Additional Info	16
Create Service Instances by Multiple Users	11	Glossary	17
Monitor Services	11	References	17
Clone a VM by Multiple Users	12		

Executive Summary

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 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, for workloads unique to an organization, and those workloads that need to remain under close control of an organization, using an IaaS structure provides the best choice. To support and host an IaaS cloud service, whether as a revenue-producing customer-facing service (a Public Cloud) or internal to an organization behind corporate firewall (a Private Cloud), the natural question would be what are the steps to build this cloud.

Given the many choices, trade-offs, and decisions, it is clear that being able to build on a known and understood starting point is key. This report summarizes a body of work consisting of building a cloud leveraging Intel® Xeon® processor 5500 series-based microarchitecture and Enomaly Elastic Computing Platform* (ECP), Service Provider Edition (SPE).

Our cloud implementation consists of five Intel Xeon processor 5500 series-based servers exposed as a compute service. Enomaly ECP SPE does not differentiate between compute and storage pools and 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, enabling 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 the ease of establishing 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 $5 \times 31 = 155$ VMs with the five server nodes in our environment with 1 GB RAM provided to each VM. When designing cloud configurations for lighter weight computing resources, 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 allow for the provisioning of 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 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 GFS (available with major Linux* distributions, and supported by Red Hat Inc.), which may be combined with iSCSI or Fibre Channel SAN storage from any of a large number of vendors. Lighter-weight solutions based on NAS technologies, 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 Enomaly ECP SPE node that is virtualized. One of the components that attribute to system bottlenecks is network traffic throughput. In our testbed, we used 10 GB Ethernet to help address the growing bandwidth needs of a cloud computing environment. Enomaly ECP Manager Agents and Enomaly ECP Node Agents were implemented on Intel Xeon processor 5500 series-based servers.

The Intel® Xeon® processor 5600 series-based servers have excellent performance/watt/sq. ft. characteristics and are ideal for highly dense cloud architectures. Intel Xeon processor 5600-based platforms also provide power management features such as Intel® Intelligent Power Node Manager that help further reduce the power footprint.

Intel Xeon processor 5600 series delivers up to 60% performance boost over Intel Xeon processor 5500 series,¹ up to 30% lower power,² and new security features – Intel® Advanced Encryption Standard–New Instructions (AES-NI), which speeds data encryption, and Intel® Trusted Execution Technology³ (Intel® TXT), which prevents the insertion of malicious software prior to virtual machine launch in virtualized environments.

Introduction

Design Principles

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

To achieve lowest operational costs for the infrastructure, a homogeneous pool of compute and storage with uniform connectivity was used as it is the simplest to manage, easiest to troubleshoot, and easiest on 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, allowing new pools to be added over time, for example as new generations of server technology becomes 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 exists between users. By contrast, in an enterprise, reasonable controls to ensure isolation are required but some access constraints might be relaxed.

Lastly, we needed to be able to make the assumption that all components of the cloud can potentially fail – hardware, network, storage, etc. Either the application or the cloud management software need to be able to perform the necessary recovery actions in the event of failure. It is required that the state of each compute element be maintained in

shared storage or that the application 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 for building this implementation include:

- The use of standard Intel® 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 of provisioning.
- 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 testbed.
- The Xen* Hypervisor was used to provide isolation for the workloads on the servers, and network traffic was segregated by Enomaly ECP SPE using network packet tagging (VLAN) and routing. Xen is one of several hypervisors supported by Enomaly ECP SPE, in addition to KVM 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.



Source: Enomaly Inc. Reprinted with permission from Enomaly Inc.

Figure 1. Basic Architecture Deployment for Enomaly ECP* SPE.

ECP Logical Architecture Configurations

Figure 1 shows the logical architecture of an Enomaly ECP SPE [1] deployment that was used in this hardware design. This is a basic architecture for Enomaly ECP SPE deployment, suitable for small-scale testbeds. 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 testbed.

Enomaly ECP SPE Implementation Overview

Enomaly’s ECP Offering

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.

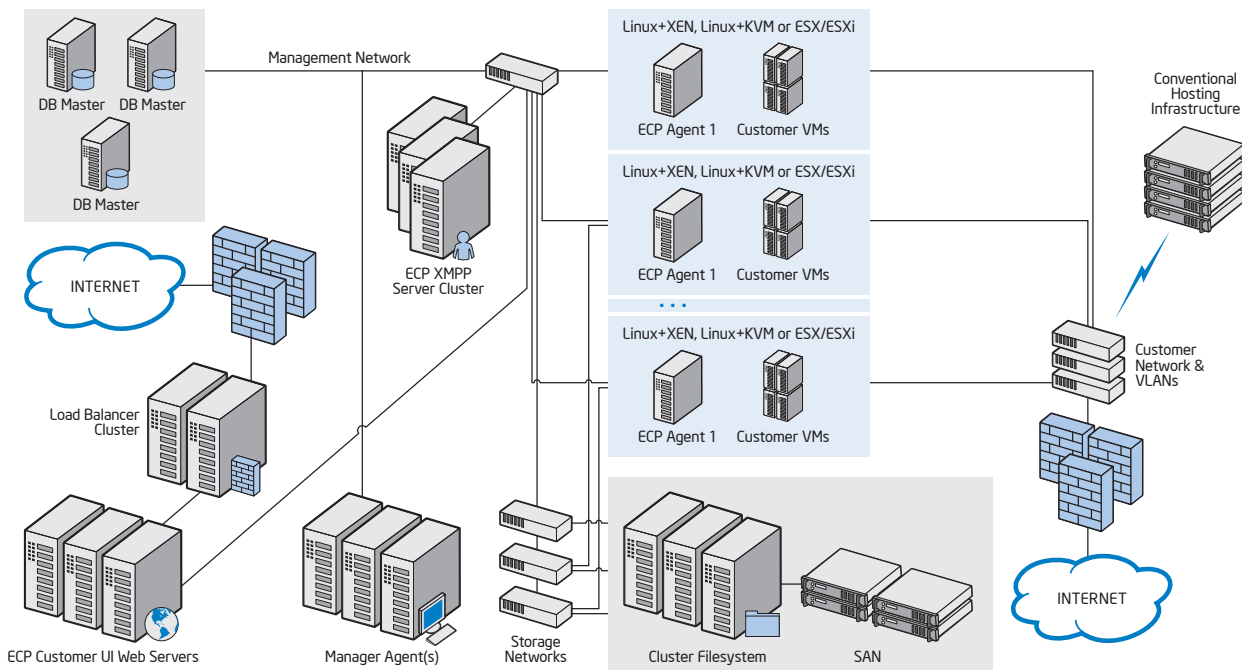
Enomaly ECP, Service Provider Edition

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

ECP SPE Features:

- Fully multi-tenant, carrier-class cloud service platform
- Detailed resource metering and accounting
- 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 ECP App Center, a service provider can publish pre-built cloud applications directly to customers, making 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 as well as infrastructure components such as load balancers and firewalls. More information and screen shots are on our Web site at www.enomaly.com/Cloud-Service-Pr.cloudhosting.0.html.

Enomaly Elastic Computing Platform* (ECP) is a software product comprising a cloud computing platform for service providers or enterprises.



Source: Enomaly Inc. Reprinted with permission from Enomaly Inc.

Figure 2. Carrier-Grade Deployment Architecture for Enomaly ECP* Service Provider Edition.

- 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 leverage Enomaly's robust provisioning and orchestration engines, the earlier open source versions of which are already deployed by 15,000 organizations around the world.

ECP Architectural Principles and Design Philosophy

The Enomaly ECP [2] [3] is a complete end-to-end platform enabling a Telco or a service provider such as a hosting company to deliver an IaaS offering to its customers. Enomaly ECP is based on a robust distributed architecture. 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 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

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 by running and interconnecting the modules requested by the application. This enables ECP to operate efficiently in a wide range of application deployments.

Decentralization

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, capable of operating effectively at 1,000-node and even 100,000-node and million-node scale, it is essential to leverage 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.

Over-Subscription

To be cost-effective, a carrier-class IaaS platform must be capable of efficiently provisioning resources for both peak and low-volume usage periods while providing 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 serving a population of customers. Service providers can oversubscribe their environments, while remaining protected from (intentional or inadvertent) overuse and misuse, and while allowing 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 provides 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 ECP's interfaces – the administrative user interface, the fully brandable customer self-service interface, and the end-user and administrative REST APIs – can simultaneously manage diverse resources distributed across multiple geographies.

Command and Control Architecture of Enomaly ECP

ECP provides a lightweight publish/subscribe messaging bus for command and control operations between the ECP Agent software controlling each compute node and the 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 allowing 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 testbed, the test cases executed, and things to consider when designing a cloud.

Testbed 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.

In our testbed, we used Intel Xeon processor 5500 series [4],[5], which provides a foundation for designing new cloud data centers to achieve greater performance while using less energy and space, and dramatically reducing operating costs. Figure 3 shows the testbed physical architecture.

The Intel Xeon processor 5500 series offers several features that help it make the best performing server for the cloud [6]. 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 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

To meet the requirements of the Intel-Enomaly Testbed, a small-scale implementation was put in place based on the Basic Deployment Architecture for Enomaly ECP SPE (Figure 1, on page 4), according to the following specifications.

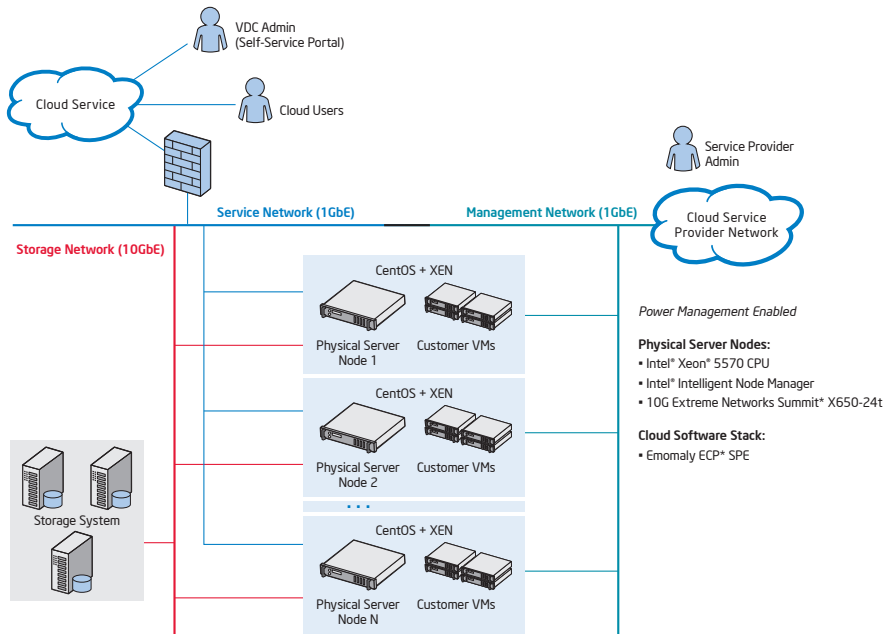
System Design

Following the basic tenet of cloud design, we partitioned the cloud design into four functional blocks: Infrastructure control for managing the cloud, storage, network and compute elements for hosting cloud-based services. Table 1 describes in detail on how we distributed the servers to handle the various functional blocks.

Furthermore, the network was partitioned for isolation, security and QoS requirements:

- An external customer network for end customer accessing the cloud via a self service portal
- An internal management network over which all command and control functions are managed
- An internal private network for managing data from the virtual machines hosted on the compute nodes

We used two Cisco 3750G* (1G) switches and one Extreme Networks Summit* X650-24t (10G) for the storage network. All switches were "out of the box" (i.e., no VPN configuration with all the 48 ports configured onto a single subnet) with a total of three subnets.



Source: Intel® Cloud Builder Testbed

Figure 3. Tested Physical Architecture.

Table 1. Cloud Testbed System Configuration

System	Processor Configuration	Other Info
1 ECP* Master and 4 ECP Agent Nodes	Intel® Xeon® Processor X5570 ^A	<ul style="list-style-type: none"> ▪ Form Factor: 1U Rack Mount Server ▪ Processor: Intel® Xeon® processor 5500-based series 2.93 GHz; 2-way x 4 cores = 8 cores ▪ Memory: 32 GB RAM ▪ Storage: 4 x 500 GB HDD SATA2 ▪ Service and Management Network: Intel® 82576 Gigabit Network Connection ▪ Storage Network: Intel® 82598EB 10-Gigabit AT2 Server Adapter
1 Storage Server	Intel® Xeon® Processor E5345 ^A	<ul style="list-style-type: none"> ▪ Form Factor: 2U Tower Server ▪ Processor: Intel® Xeon® processor 5300-based series 2.33 GHz; 2-way x 4 cores = 8 cores ▪ Memory: 16 GB RAM ▪ Local Storage: 1x160 GB HDD SATA ▪ Expansion Storage: 10 x 160 GB SDD
1 Proxy Server	Intel® Xeon® Processor	<ul style="list-style-type: none"> ▪ Form Factor: 1U Rack Mount Server ▪ Processor: Intel® Xeon® processor 2.8 GHz; 2-way ▪ Memory: 2 GB RAM ▪ Storage: 2x160 GB HDD SATA
Network	Summit* X650-24t	<ul style="list-style-type: none"> ▪ Ports: 24x 10GBASE-T ports (port 1# ~ 24#); 4x 1000BASE-X (SFP) ports (port 25# ~ 28#); 2x SummitStack ports (shared with last 2x 10GBASE-T ports/dedicated 10 Gbps without auto-negotiation) ▪ VLAN: 802.1Q; Port-based and tag-based VLAN

Table 2. Network IP Configuration

	Hostname	Storage Network	Private Network	Service Network
Storage Server	ECPstorage	192.168.182.240	192.168.181.240	172.18.1.112
ECP Master Node	ECP-master	192.168.182.100	192.168.181.100	172.18.1.105
ECP Node 01	ECPO1	192.168.182.101	192.168.181.101	172.18.1.106
ECP Node 02	ECPO2	192.168.182.102	192.168.181.102	172.18.1.107
ECP Node 03	ECPO3	192.168.182.103	192.168.181.103	172.18.1.108
ECP Node 04	ECPO4	192.168.182.104	192.168.181.104	172.18.1.109

Our design reserves static IP addresses for nodes that participate in hosting compute and storage: ECP-master, ECPO1, ECPO2, ECPO3, ECPO4 and ECPstorage.

An additional server was configured to act as a caching proxy server. This was done specifically to speed the process by which the servers were able to get up and running. Since each server is installed with a clean copy of Enomaly ECP SPE, each server will need to get updates and packages from the Internet. The caching Proxy server is configured to cache these packages to speed up the process.

Technical Review

Installation Overview

This section discusses the installation process for Intel-Enomaly cloud infrastructure. To set up the cloud infrastructure, we validated that there was sufficient hardware, software, memory, network and storage resources and the overall network and DHCP connectivity. We then configured the storage elements, switch and VLAN settings, installed the ECP master node and four agent nodes and verified the network, database connectivity. To test the newly configured 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 Intel Testbed Installation Details document.[7]

Use Case Details

Actors

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

Use Case Overview

The cloud testbed App Center was provisioned with an instance of CentOS 5.4 (64-bit), and during the testing phase of the cloud implementation, the following 14 use cases were successfully run:

Pre-conditions:

1. The cloud management software, Enomaly ECP SCE software is installed and ready to go.
2. Compute and storage nodes are installed and registered.
3. SC1 and SC2 application services (Service 1 and Service 2) 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) via admin portal.
3. SC1, SC2: Create instance of the Service 1, Service 2: Instantiate VMs that make up the Service 1, Service 2 including IP address and links to storage.
4. SC1, 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 VMs 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 and 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, SC2 can examine billing records.
11. SP: Add bare-metal capacity to existing cluster: Add a new server to an existing cluster.
12. SC1, 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, SC2: Shutdown Service 1, Service 2: End the application service and remote access to users SC1 and SC2.

Execution and Results

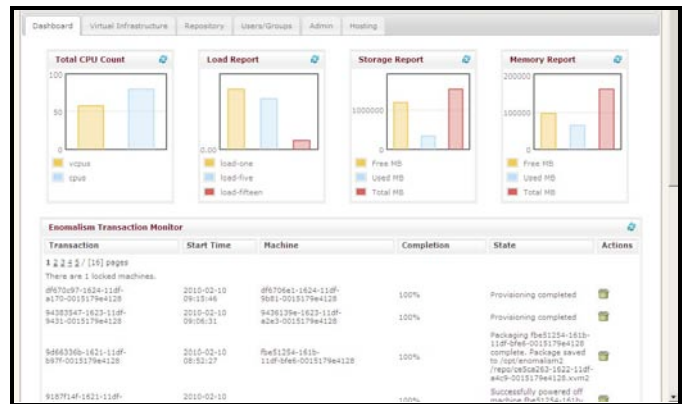
All the test cases described above were successfully executed. Detailed documentation of each of the test case inputs as well as system responses, including screen shots of each step, are shown in the following subsections. All use cases are performed through the self-service web UI portal unless otherwise stated.

Setup and Software Verification

1. A successful login as “admin” through the Web UI portal (for our setup: <http://172.18.1.105:8080>)



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



3. Access the “App Center” Home page to install a VM with the “centos54” VM template provided by ECP SPE 3.0.4. 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 if the VM “test” was installed and powered on automatically.





- Open VNC connection to the newly created VM with the information provided on the "Details" page and login into VM with account/password.
- Power off VM "test" and see it is moved from group "Running VMs" to group "Powered-off" VMs on "VM" page.



Create a Virtual Machine Image for Use

- Copy a bootable installation ISO image file for your operating system to "/opt/enomalism2/iso/" directory of ECP master node. For this test case, we use Windows Server* 2003 installation media, with an ISO file named "WindowsServer2003EntR2SP2_x86.iso".
- 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: "WindowsServer2003EntR2SP2_x86.iso", VM Type: "XEN KVM Machine", RAM: 1024 and default settings for other options.
- After the VM appliance is created successfully, it will be listed in "Repository/LOCAL APPLIANCES" page, rename it to "install: windows server 2003".
- Go to "Virtual Infrastructure/ELASTIC VALET" page and provision a new VM; Package to provision: "install: windows server 2003" and default settings for other options.
- 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 and enter the Windows OS installation screen. Install the Windows OS within the VM and customize it as desired, obtaining a freshly installed Windows system. 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 and renamed to "windows server 2003".
- Go to the "Repository/REMOTE APPLIANCES" page and install the "windows server 2003" 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.

Actions	Title	Description	Type	OS	Storage (GB)
	2P win 350	Auto imported XEN package (from windows-server-0303040812SP2_x86.iso) e24e34-1f1-110f-8a79-0515179e4128	Libvirt	windows	8.00
	2P win03	PKG_RuS1284_1433-110f-8a79-0515179e4128_x86packed on host Feb 10 17:00:09 2008	Libvirt	windows	8.00
	centos4	ALTO Import from Ap/monalim2 /repo/4407468_c94a_110f-8a79-0515179e4128.vmd	Libvirt	unknown	20.00
	installwindows server 2003	Auto imported XEN package (from windows-server-0303040812SP2_x86.iso)	Libvirt	unknown	8.00
	XEN S.P	VM SrvL S.P. This appliance contains a single pre-installed VM (3008-001) on the main hardware	Libvirt	unknown	8.00
	windows server 2003	Mj4799-154d-110f-0591-0015179e4128.vmd	Libvirt	windows	8.00

Create Users

- 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

- Go to "Users/Groups" page, 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

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
    
```

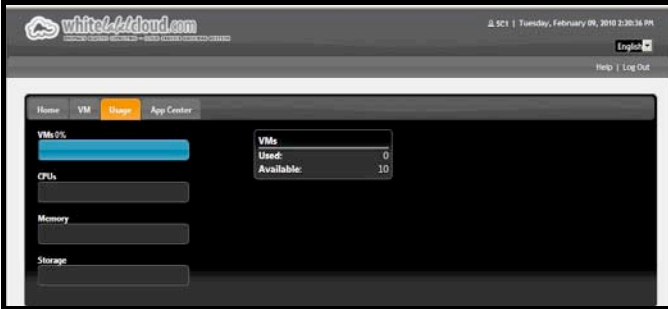
Small:i686 CPU Arch 1024 Memory	xen-hvm Hypervisor 1 CPUs
Medium:i686 CPU Arch 2048 Memory	xen-hvm Hypervisor 2 CPUs
Large:x86_64 CPU Arch 8192 Memory	xen-hvm Hypervisor 8 CPUs

- Go to "Hosting/NETWORK" page and create four VLANs.

```

VLAN100: 100 VLAN ID
VLAN200: 200 VLAN ID
VLAN300: 300 VLAN ID
VLAN400: 400 VLAN ID
    
```

- Go to "Users/Groups"/"EDIT PERMISSIONS" page and grant "read" Permissions to Groups Group G1 and G2.
- Logout and login as "SC1" and verify if "Usage" Page to check current resource usage and quota status.




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



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

Create Service Instances by Multiple Users

- Go to "App Center" page and install a VM with the "centos54" VM template provided by ECP SPE 3.0.4, using the "small" Hardware Template and a VM name: "centos54-1". Verify if "centos54-1" created successfully by accessing the VM page. By visiting the "usage" page, we can verify if the resource usage/quota has been updated. Please refer to section 5.3.1 for the VM creation details.
- Wait for about 1 minute, and then open "Details" page of VM "centos54-1" and go to "Network" tab, check if the VM's IP address was provisioned successfully.



- Open VNC connection to VM "centos54-1", log into the VM using account information.

```
$ ping www.intel.com
```

The ping should be successful.

Monitor Services


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

VM Name: "centos54 - 2" / Hardware Template: "Medium"

VM Name: "centos54 - 3" / Hardware Template: "Large"

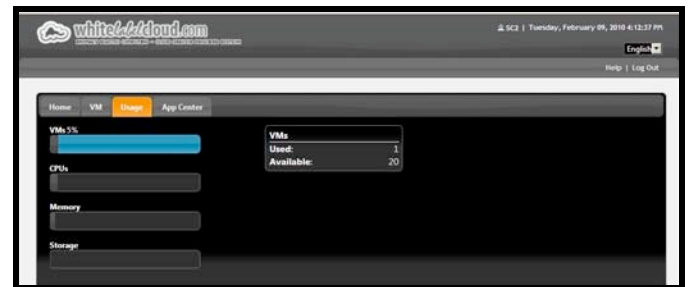
VM Name: "centos54 - 4" / Hardware Template: "Large"

- Use the "App Center" page to install another VM with the "centos54" VM template (VM Name: "exceed", Hardware Template: "Large"), but get failure notification due to exceeding quota limits.



- Logout as "admin" and login as the user "SC2".

- Go to "App Center" page and install a VM with the "windows server 2003" VM template (VM Name: "win2003-1", Hardware Template: "Small") and verify that the VM is successfully created.



- Open VNC connection to VM "win2003-1", log into the VM using account information.

a. \$ ping www.intel.com

b. The ping should be successful.

- Repeat above steps to install more VMs:

VM Name: "win2003 - 2" with Hardware Template: "Medium"

VM Name: "win2003 - 3" with Hardware Template: "Large"

VM Name: "win2003 - 4" with Hardware Template: "Large"

VM Name: "win2003 - 5" with Hardware Template: "Large"


VM Name: "win2003 - 6" with Hardware Template: "Large"

- Logout and login as "admin" and then check resource reports updated in dashboard.




Clone a VM by Multiple Users

- 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.



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



Network Isolation

- 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".
- 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 IP from DHCP server, as other two VMs are isolated.

- Open VNC connection to VM "centos54-1", login to the VM and reconfigure network IP address using following command:

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

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

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

- Logout and login as "SC2"; Power down all the VMs. Open "Details" page of the VM "win2003-1" and go to "Network" tab, reconfigure it from network "default" to "VLAN100" and VM "win2003-2" from network "default" to "VLAN300".

- Power on VMs: "win2003-1", "win2003-3", and the cloned VM of "win2003-1". Check that only the cloned VM of "win2003-1" can still get IP from DHCP server, as other two VMs are isolated.

- Open VNC connection to VM "win2003-1" and "win2003-3" and login to the VM and reconfigure network IP address to be unique.

- Logout and Login as "SC1" user. Open VNC connection to the VM and login to the VM and ping other VMs as follows:

```
$ ping centos54-2 // No replies as the VM is not in the same VLAN
```

```
$ ping win2003-1 // Get replies as this VM is in the same VLAN
```

```
$ ping win2003-3 // No replies as the VM is not in the same VLAN
```

VM Configuration Management

- Go to "Virtual Infrastructure/INFRASTRUCTURE" page, open the page for VM "SC2_win2003-2" under "SC2/ECP01", and check VM's CPU/Memory/HD configuration "default" to "VLAN200".



Billing for Multiple Users

- 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_uid	Small	Medium	Large	Total MBytes
admin	7eabba12-12fe-11dd-8fd0-0019d2b28af0	0	0	0	0
customer	b0e229cc-115f-11df-83ba-0015179e4128	5	31	8	9729
SC1	25576b1c-1524-11df-9d72-0015179e4128	88	49	88	358400
SC2	52db6d35-1524-11df-ab85-0015179e4128	133	44	175	81920

2. Save the output as "billing.csv".

Add New Servers

1. Access "Virtual Infrastructure/INFRASTRUCTURE" page, and check existing ECP nodes. Four ECP nodes ECP-master, ECPO1, ECPO2 and ECPO3 should be visible.



2. Login ECPO4 and start to install ECP SPE 3.0.4 for Agent Node Installation (Refer to 5.3.1.2 for more details). Access "Virtual Infrastructure/INFRASTRUCTURE" page on Web UI portal. ECPO4 should be added to the existing ECP nodes.

Scale Up and Down Services

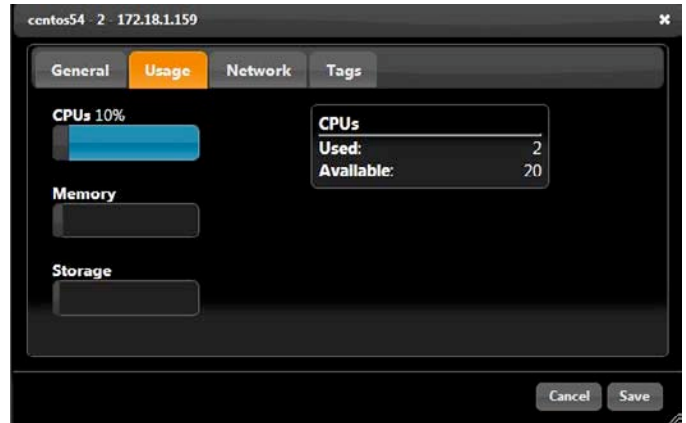
1. Login as "SC1" through the Web UI portal <http://172.18.1.105:8080>. Access the "VM" page, power on VM "centos54-2", open its "Details" page and go to "Usage" tab to check the VM's current CPU/Memory/HD configuration.

2. Power off VM "centos54-2", open its "Details" page and change Hardware Templates from "Medium" to "Large".


```
Small:i686 CPU Arch xen-hvm Hypervisor
1024 Memory 1 CPUs

Medium:i686 CPU Arch xen-hvm Hypervisor
2048 Memory 2 CPUs

Large:x86_64 CPU Arch xen-hvm Hypervisor
8192 Memory 8 CPUs
```



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



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

```
(Partial output from cpuinfo.info)
processor: 0
vendor_id: GenuineIntel
model name: Intel® Xeon® CPU X5570 @ 2.93GHz
processor: 1
vendor_id: GenuineIntel
model name: Intel® Xeon® CPU X5570 @ 2.93GHz
processor: 2
vendor_id: GenuineIntel
model name: Intel® Xeon® CPU X5570 @ 2.93GHz
processor: 3
vendor_id: GenuineIntel
model name: Intel® Xeon® CPU X5570 @ 2.93GHz
```

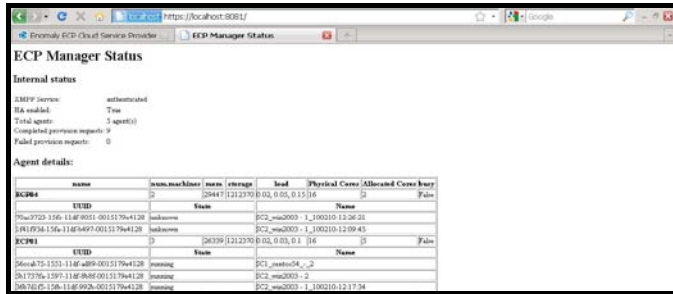
```
# cat /proc/meminfo > meminfo.info
```

```
(Partial output from meminfo.info)
MemTotal: 3893932 kB
MemFree: 3813588 kB
```

Failure Scenario

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

```
# ssh -L 8081:127.0.0.1:8081 root@172.18.1.105
```
2. Connect to ECP Manager Status Report <https://localhost:8081>, and check existing ECP nodes. All should be active.
3. Reboot ECP04. ECP04 marked with strikeout and the VMs running on ECP04 changed from running state to unknown state in about 10 seconds.



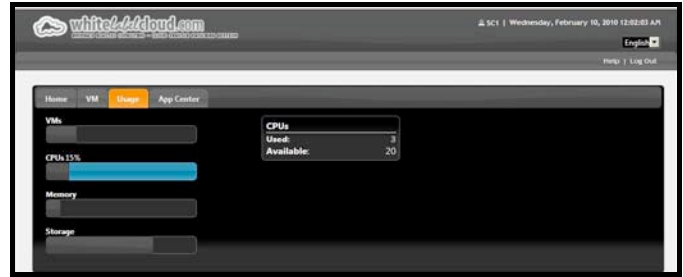
4. ECP04 comes back to operational state in about 5 minutes, the strikeout mark on ECP04 removed, and all VMs running on ECP04 moved to ECP01 and changed to running state.



5. Reboot ECP-master with the following observations:
 - Web UI portal and ECP Manager Status Report unavailable until ECP-master came back to operational state.
 - 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% of the CPU quota is currently being used.



3. Power off VMs "centos54-1" (Small) and "centos54-2" (Medium).
4. Go to "Usage" page and check resource usage/quota. The CPU quota will drop down to 0%.

Observations and Analysis

Enomaly ECP SPE provides a solution enabling 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 validate 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 for 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 for 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 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 60-80% of the values off the host machine itself. Overhead loss is to be expected, but lower than 60% 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 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) [8]:** 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) [8]:** 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. **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 not aware.
 - 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. **Power Management:** Utilize Intel® Intelligent Power Node Manager and Intel® Data Center Manager [9] along with policies to optimize power usage. Use cases could be implemented to achieve power capping to avoid circuit breaker trips or to avoid hot spots in the data center. Please refer to [10] to learn about Baidu's case study utilizing the Intel® Intelligent Power Node Manager.

Things to Consider

Scalability

The scalability of the cloud solution could be impacted by:

- Network technology (e.g., 10GigE) 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 the performance in a cloud. In addition, when planning 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 testbed we selected a very simple network topology. The selection made for this testbed performed well for our purposes, but more advanced technologies and architectures (e.g., 10-GigE, 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 testbed 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 IO support feature sets, such as the Intel® Xeon® processor 5500 series, and the Intel Xeon processor 5600 series, is strongly recommended. For more details on Intel® Virtualization technologies please refer to [11].

Solid-State Drives

The performance of storage nodes (and compute nodes when local storage is utilized), as well as the overall power consumption of the cloud, may be favorably impacted by the use of SSDs. This was not specifically tested within our exercise.

Security Considerations

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, isolating the network traffic of each customer into one or more private VLANs.

Enomaly ECP* High Assurance Edition, leverages the Intel® Trusted Execution Technology (Intel® TXT) capabilities of 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 Builder: intel.com/software/cloudbuilder

Intel Xeon processors: intel.com/xeon

Enomaly: www.enomaly.com

Glossary

To avoid ambiguity about the terms used, here are the definitions for some of the specific concepts used in this paper:

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

A **Domain** is an instance of an operating system (or subsystem in the case of container virtualization) running on a virtualized machine provided by the Hypervisor.

Enomaly Elastic Computing Platform (ECP) is a software product comprising a cloud computing platform for service providers or enterprises.

HA – High Availability

A **Hypervisor** is a layer of software allowing the virtualization of a node in a set of virtual machines, supporting one or more virtual machines with possibly different configurations than the node itself.

LVM – Logical Volume Manager

A **Node or Host** is a single physical machine.

A **Pool** provides a means for taking a chunk of storage and carving 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, an LVM group.


VM – Virtual Machine

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

References

- [1] Enomaly ECP Architecture Whitepaper, www.enomaly.com/Cloud-Hosting-Pr.466.0.html
- [2] Enomaly ECP SPE Datasheet, www.enomaly.com/Cloud-Hosting-Pr.465.0.html
- [3] Enomaly ECP HAE Datasheet, www.enomaly.com/Cloud-Hosting-Pr.465.0.html
- [4] Why the Intel® Xeon® Processor 5500 Series is the Ideal Foundation for Cloud Computing, communities.intel.com/docs/DOC-4213
- [5] Intel Xeon 5500 Series Software Industry Testimonials, www.intel.com/business/software/testimonials/xeon5500.htm
- [6] Intel in cloud computing Wiki, communities.intel.com/docs/DOC-4230
- [7] Enomaly ECP Intel Testbed Installation Details, enomaly.com/fileadmin/ICB/Enomaly_ECP_Intel_Testbed_Installation_Details.pdf
- [8] Intel® AESNI and Intel® Trusted Execution Technology, www.intel.com/technology/security/
- [9] Intel® Intelligent Power Node Manager and Intel Data Center Manager, software.intel.com/sites/datacentermanager/
- [10] Baidu POC, software.intel.com/sites/datacentermanager/intel_node_manager_v2e.pdf
- [11] Intel Virtualization Technology, www.intel.com/technology/virtualization/ and download.intel.com/business/resources/briefs/xeon5500/xeon_5500_virtualization.pdf

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



⁴ 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.

¹ Source: Internal Intel measurements for Intel® Xeon® processor X5680 vs. Intel® Xeon® processor X5570 on BlackScholes*.

² Source: Fujitsu Performance measurements comparing Xeon L5650 vs X5570 SKUs using SPECint_rate_base2006. See <http://docs.ts.fujitsu.com/dl.aspx?id=0140b19d-56e3-4b24-a01e-26b8a80cfe53> and <http://docs.ts.fujitsu.com/dl.aspx?id=4af74e10-24b1-4cf8-bb3b9c4f5f177389>.

³ 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.

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

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

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

Copyright © 2010 Intel Corporation. All rights reserved. Intel, the Intel logo, Xeon, and Xeon inside are trademarks of Intel Corporation in the U.S. and other countries.

Copyright © 2010 Enomaly Inc. Enomaly and the Enomaly logo are registered trademarks of Enomaly Inc.

*Other names and brands may be claimed as the property of others.