



# The Education Cloud MECA

**Danilo Demarchi**

Politecnico di Torino

Department of Electronics and Telecommunications

- Develop Cloud-based European infrastructure
- Provide a range of open educational resources
- Remote access and sharing of educational and professional software
- Remote and practice-based learning facilities.

- Laboratory experiences
- CAD tools
- Project ideas
- A common infrastructure as “educational cloud” on top of the cloud software/hardware infrastructure
- E-Learning materials for at least 16 courses
  - CAD systems
  - Microelectronics technologies
  - Test & Characterisation
  - Application of integrated circuits and systems
  - Open educational resources to strengthen the virtual mobility
  - Each university will provide remote access to facilities, laboratories in a cloud teaching system

- KVM
  - Linux-based Virtualization
- Cloudstack
  - Open Source software designed to deploy and manage large networks of virtual machines
- Cloudify
  - Open Source cloud orchestration framework

- KVM – the Kernel-based Virtual Machine – is a Linux kernel module that turns Linux into a hypervisor
- Supports multiple architectures: x86 (32- and 64-bit) s390 (mainframes), PowerPC, ia64 (Itanium)
- Tightly integrated into Linux
- In the Virtual Machine is possible to install almost any Operating System

- Useful for setting up Virtual Machines on which Students can remotely access and to do Lab works
  - Using their PCs, so from Standard Classrooms
  - In the Campus Area
  - With controlled accesses from Home (not implemented at Polito)

- It is possible using FreeNX Server (the free version of OpenNX), and NoMachine (Open Source Client for accessing Remote Machines)
- The students can enter with his/her privileges and work
- The Session can be interrupted and the work can start as stopped opening a new session, also on another PC

- Apache CloudStack is an Open Source **Infrastructure-as-a-Service (IAAS)** platform that manages and orchestrates pools of storage, network, and computer resources to build a public or private IaaS compute cloud
- With CloudStack you can:
  - Set up an on-demand elastic cloud computing service
  - Allow end-users to provision resources





# 300+

## Enterprise and Education



## Service Providers and Telcos

Large Scale  
Production Clouds  
In Deployment



Production sites with over

# 40,000+

Servers



## Web 2.0



# CloudStack Architecture

4/29/2012

Chiradeep Vittal

Alex Huang

# CloudStack Supports Multiple Cloud Strategies

**MECA**

## Private Clouds

### On-premise Enterprise Cloud



- Dedicated resources
- Security & total control
- Internal network
- Managed by Enterprise or 3<sup>rd</sup> party

## Public Clouds

### Hosted Enterprise Cloud



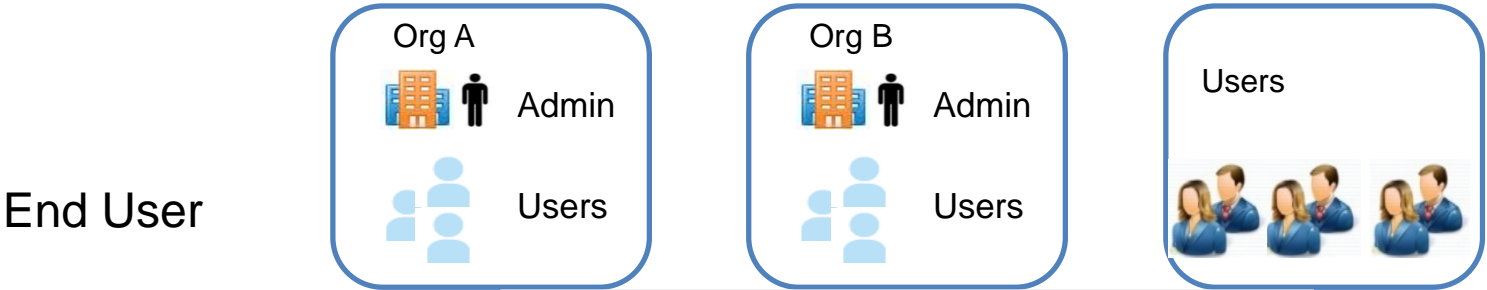
- Dedicated resources
- Security
- SLA bound
- 3<sup>rd</sup> party owned and operated

### Multi-tenant Public Cloud



- Mix of shared and dedicated resources
- Elastic scaling
- Pay as you go
- Public internet, VPN access

# CloudStack Provides On-demand Access to Infrastructure Through a Self-Service Portal



cloudstack



# Open Flexible Platform

## Compute



### Hypervisor

XenServer

VMware

Oracle VM

KVM

Bare metal

## Storage



### Block & Object

Local Disk

iSCSI

Fiber  
Channel

NFS

Swift

Primary Storage

Secondary Storage

## Network



### Network & Network Services

Network  
Type

Isolation

Firewall

Load  
balancer

VPN

# Create Custom Virtual Machines via Service Offerings

Add instance

1 Setup > 2 Select a template > 3 Service Offering > 4 Data Disk Offering > 5 Network > 6 Review


Please select networks for your virtual machine

Networks			
<input type="checkbox"/>	CustomerNet	Direct	<input type="radio"/> Default
<input type="checkbox"/>	CustomerNet	Direct	<input type="radio"/> Default
<input type="checkbox"/>	vlan100	Direct	<input type="radio"/> Default
<input type="checkbox"/>	vlan100	Direct	<input type="radio"/> Default
<input type="checkbox"/>	vlan100	Direct	<input type="radio"/> Default

Add new network

<input type="checkbox"/>	NEW
--------------------------	-----

Previous Cancel Next



Select Operating System

- Windows, Linux

Select Compute Offering

- CPU & RAM

Select Disk Offering

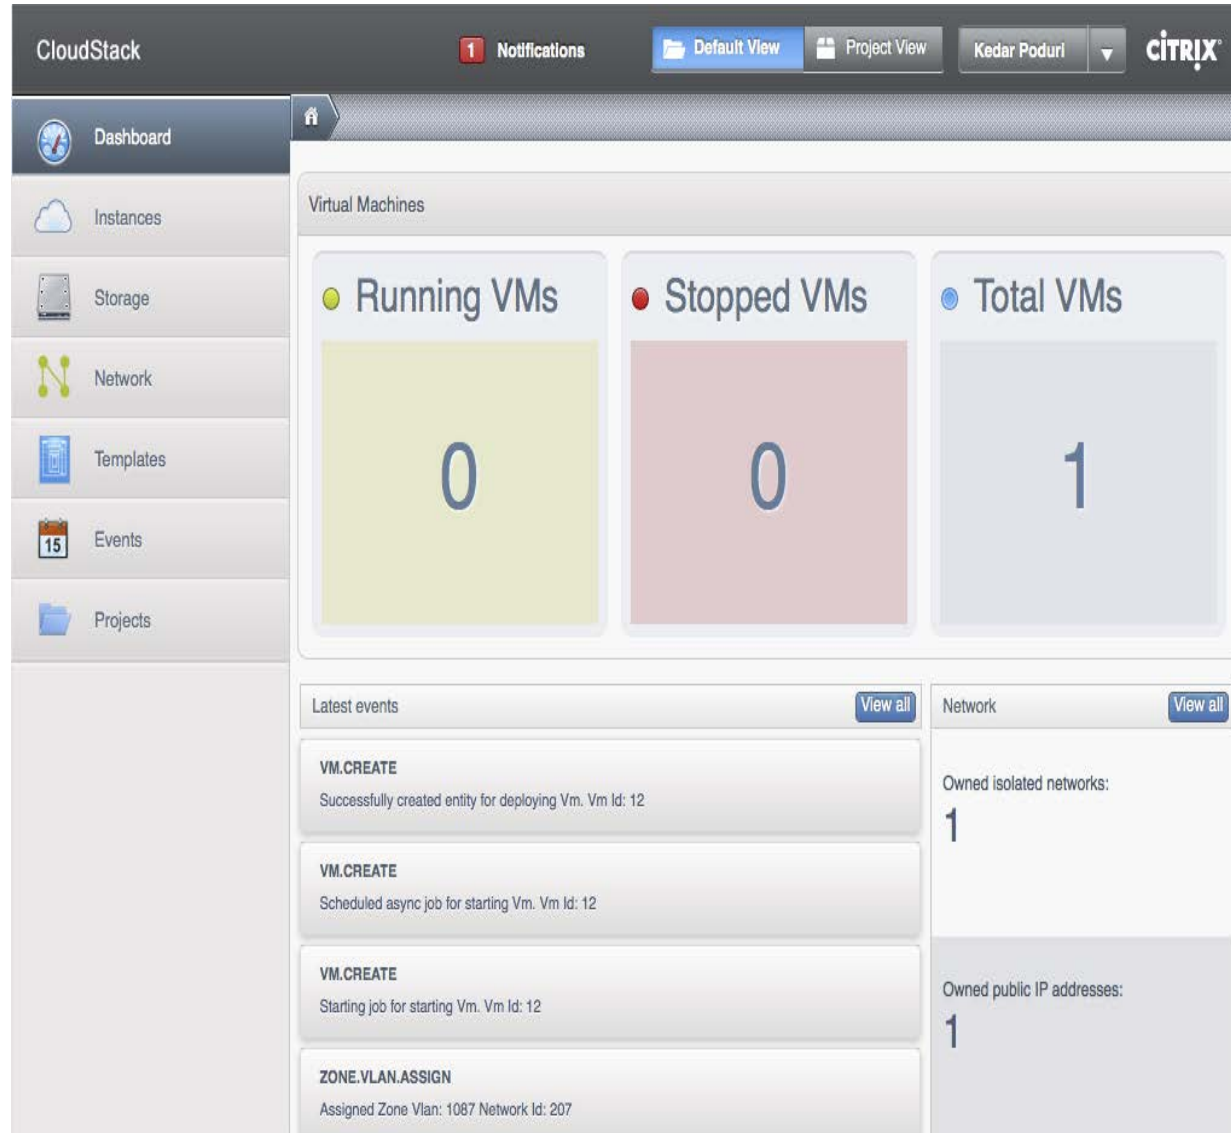
- Volume Size

Select Network Offering

- Network & Services

Create VM

# Dashboard Provides Overview of Consumed Resources



- Running, Stopped & Total VMs
- Public IPs
- Private networks
- Latest Events

# Virtual Machine Management



Users

VM Operations

Console Access

VM Status

Change Service Offering



Start



Stop



Restart



Destroy

```
vmtoolsd --
$ dir
2007/12/27 14:25:00  b cywin1.com*
--
$ ls -ltr
total 320
-rwxrwxrwx 1 carlos 0x0  0 Oct 5 2007 c*
-rwxrwxrwx 3 carlos 0x0  0 Oct 5 2007 b
-rwxrwxrwx 4 carlos 0x0  0 Oct 5 2007
-rwxrwxrwx 2 carlos 0x0  0 Dec 26 14:27
-rwxrwxrwx 1 carlos 0x0  25K Dec 27 14:26 20071227.com*
$
```

- CPU Utilized
- Network Read
- Network Writes

2 CPUs

1 GB  
RAM

20 GB

20 Mbps



4 CPUs

4 GB  
RAM


200 GB

100  
Mbps





# Volume & Snapshot Management


Add / Delete Volumes




Create Templates from Volumes



Schedule Snapshots



View Snapshot History



# Network & Network Services

- Create Networks and attach VMs
- Acquire public IP address for NAT & load balancing
- Control traffic to VM using ingress and egress firewall rules
- Set up rules to load balance traffic between VMs

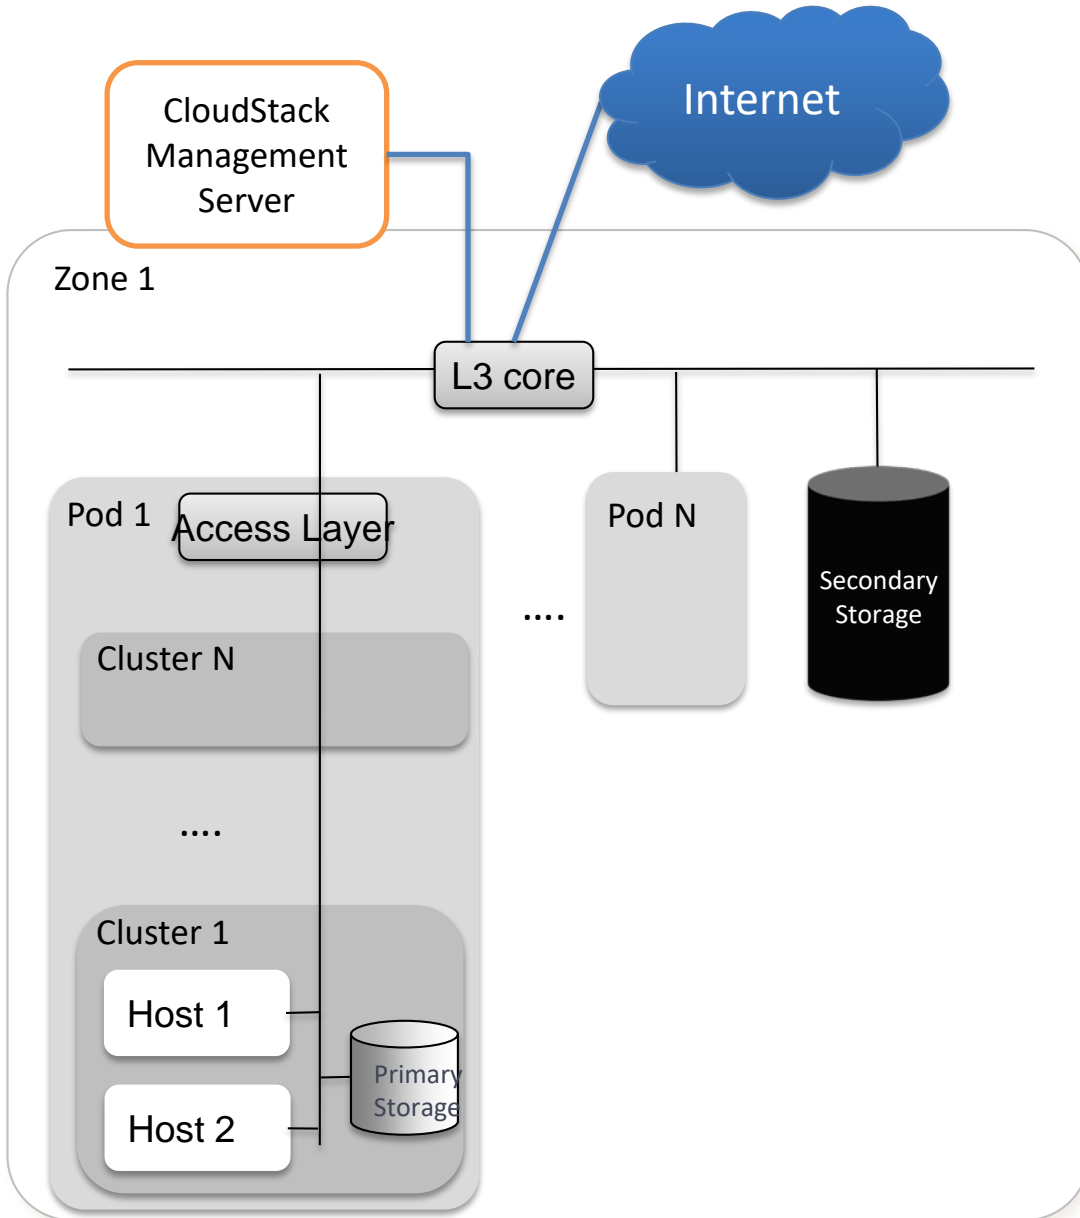
CloudStack 0 Notifications Default View Project View kedar CITRIX

Dashboard  
Instances  
Storage  
**Network**  
Templates  
Accounts  
Domains  
Events  
System  
Projects  
Global Settings  
Configuration

Network  
Select view: IP Addresses Filter By: Allocated

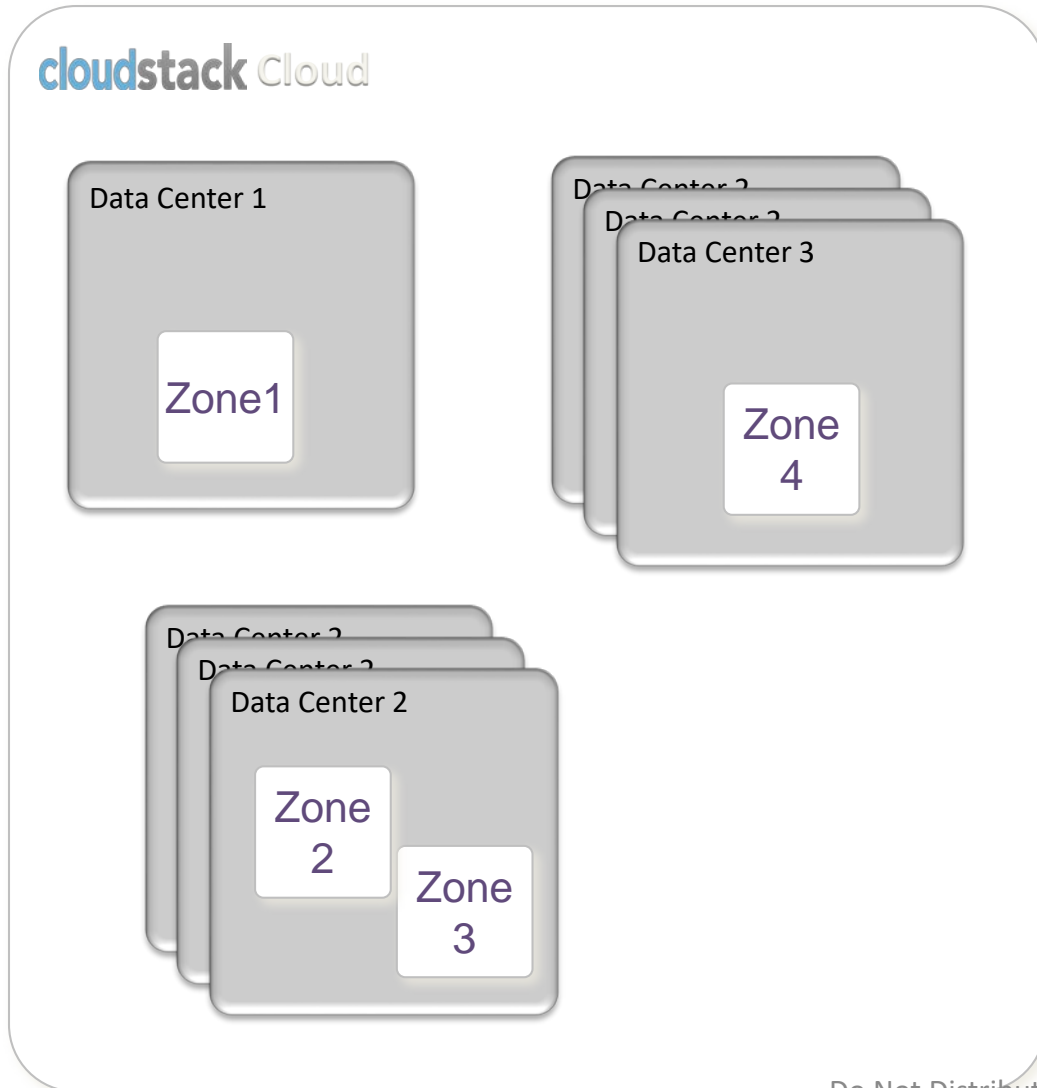
IP	Zone	VLAN	Network Type	State	Actions
72.52.67.77	San Jose	67	200	Allocated	
72.52.67.62 [Source NAT]	San Jose	67	200	Allocated	
72.52.67.52	San Jose	67	200	Allocated	
72.52.67.51	San Jose	67	200	Allocated	
72.52.67.46 [Source NAT]	San Jose	67	200	Allocated	
72.52.67.45 [Source NAT]	San Jose	67	200	Allocated	
72.52.67.42	San Jose	67	200	Allocated	
72.52.67.37	San Jose	67	200	Allocated	
72.52.67.31 [Source NAT]	San Jose	67	200	Allocated	
72.52.67.23	San Jose	67	200	Allocated	
72.52.67.22	San Jose	67	200	Allocated	
72.52.67.20	San Jose	67	200	Allocated	
72.52.126.69 [Source NAT]	San Jose	72	200	Allocated	
72.52.126.68 [Source NAT]	San Jose	72	200	Allocated	

# CloudStack Deployment Architecture



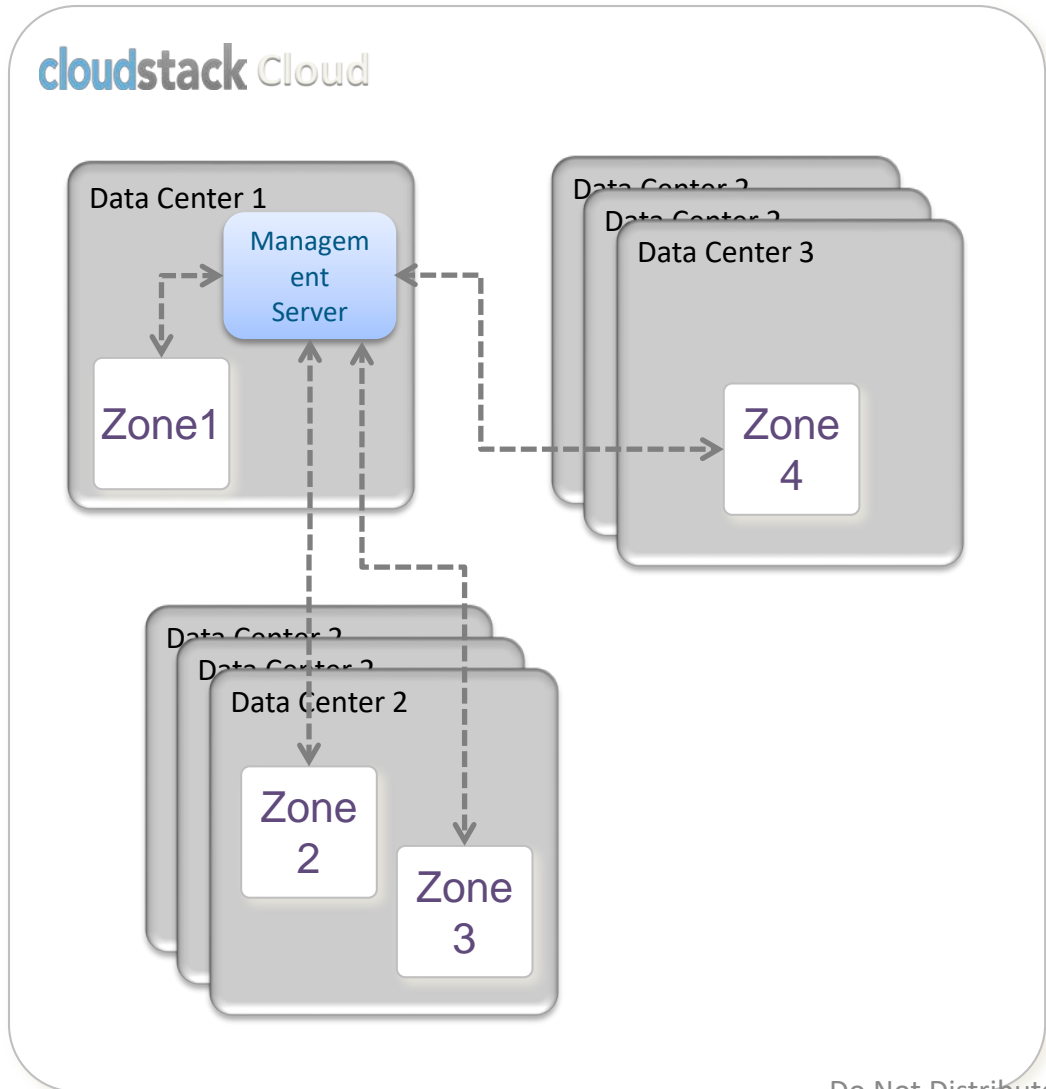
- Hypervisor is the basic unit of scale.
- Cluster consists of one or more hosts of same hypervisor
- All hosts in cluster have access to shared (primary) storage
- Pod is one or more clusters, usually with L2 switches.
- Availability Zone has one or more pods, has access to secondary storage.
- One or more zones represent cloud

# CloudStack Cloud Architecture



CloudStack Cloud can have one or more Availability Zones (AZ).

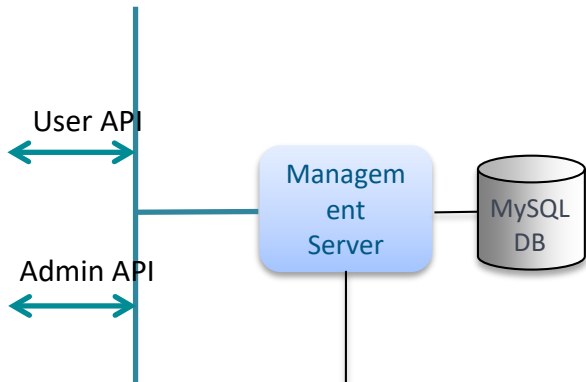
# Management Server Managing Multiple Zones



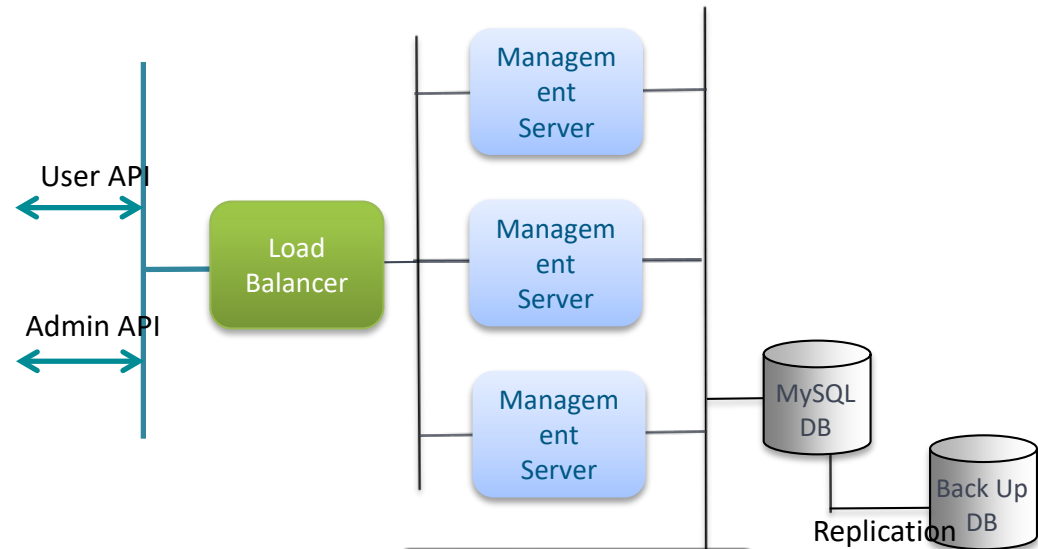
- Single Management Server can manage multiple zones
- Zones can be geographically distributed but low latency links are expected for better performance
- Single MS node can manage up to 5K hosts.
- Multiple MS nodes can be deployed as cluster for scale or redundancy

# Management Server Deployment Architecture

## Single-node Deployment



## Multi-node Deployment



Infrastructure Resources

- MS is stateless. MS can be deployed as physical server or VM
- Single MS node can manage up to 10K hosts. Multiple nodes can be deployed for scale or redundancy
- Commercial: RHEL 5.4+; FOSS: Ubuntu 10.0.4, Fedora 16

Infrastructure Resources

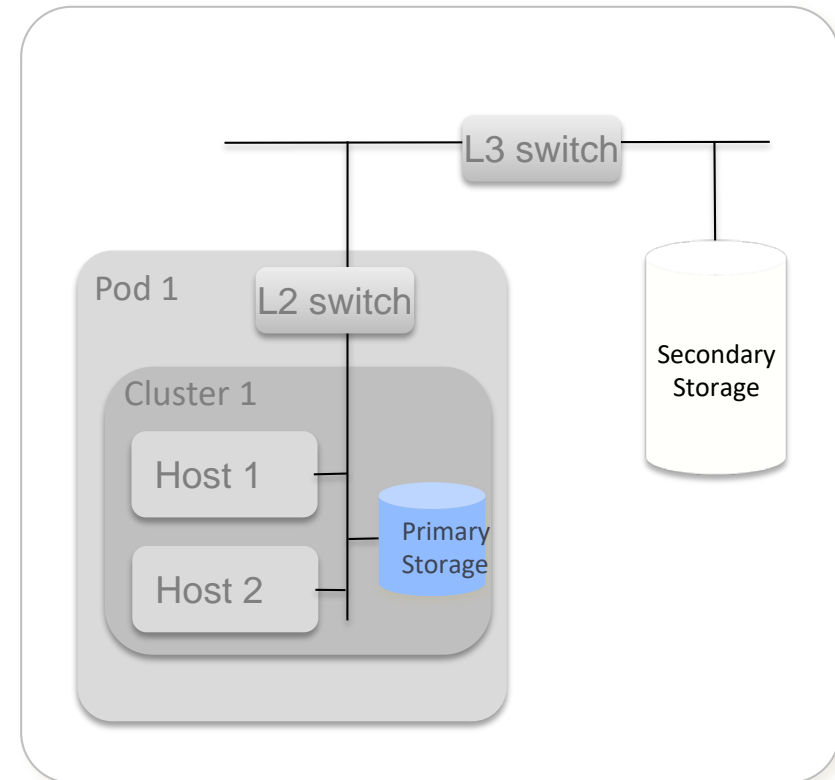
# CloudStack Storage

## Primary Storage

- Configured at Cluster-level. Close to hosts for better performance
- Stores all disk volumes for VMs in a cluster
- Cluster can have one or more primary storages
- Local disk, iSCSI, FC or NFS


## Secondary Storage

- Configured at Zone-level
- Stores all Templates, ISOs and Snapshots
- Zone can have one or more secondary storages
- NFS, OpenStack Swift



# Core CloudStack Components


- **Hosts**
  - Servers onto which services will be provisioned
- **Primary Storage**
  - VM storage
- **Cluster**
  - A grouping of hosts and their associated storage
- **Pod**
  - Collection of clusters
- **Network**
  - Logical network associated with service offerings
- **Secondary Storage**
  - Template, snapshot and ISO storage
- **Zone**
  - Collection of pods, network offerings and secondary storage
- **Management Server Farm**
  - Responsible for all management and provisioning tasks





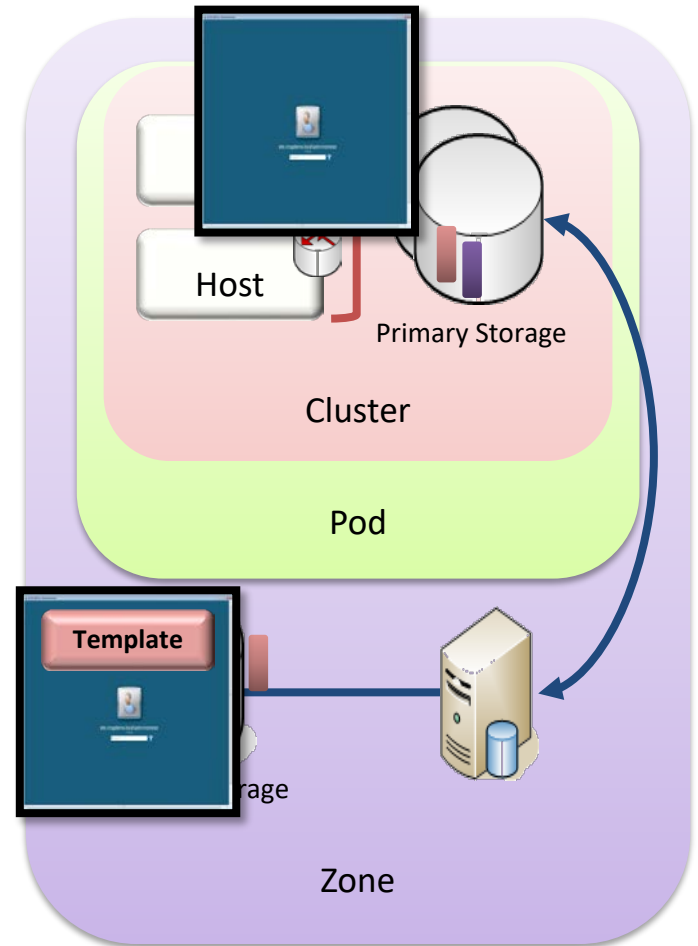
# Understanding the Role of Storage and Templates

- Primary Storage
  - Cluster level storage for VMs
  - Connected directly to hosts
  - NFS, iSCSI, FC and Local
- Secondary Storage
  - Zone level storage for template, ISOs and snapshots
  - NFS or OpenStack Swift via CloudStack System VM
- Templates and ISOs
  - Imported into CloudStack
  - Can be private or public



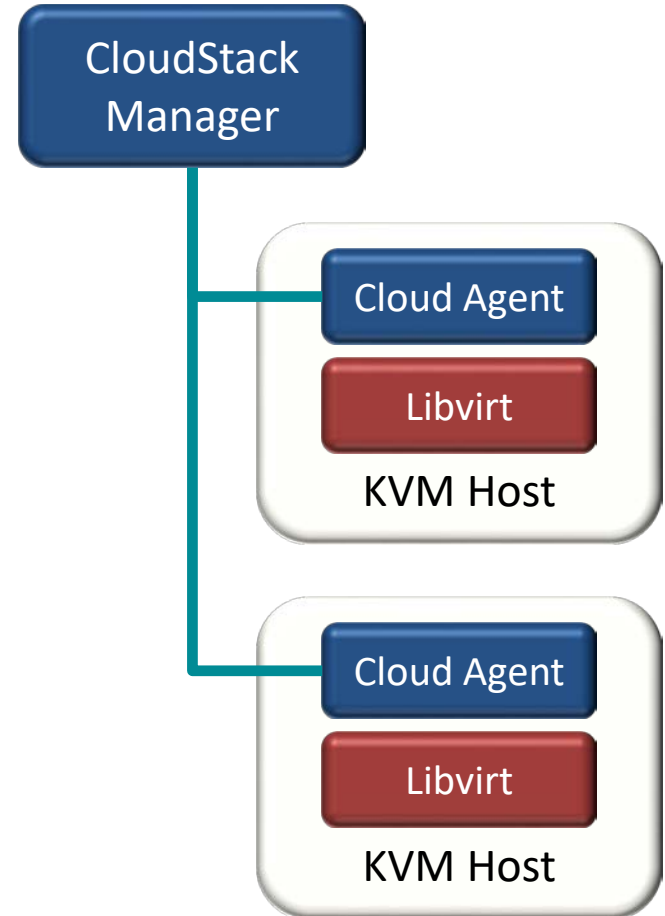
# Provisioning Process

1. User Requests Instance
2. Provision Optional Network Services
3. Copy instance template from secondary storage to primary storage on appropriate cluster
4. Create any requested data volumes on primary storage for the cluster
5. Create instance
6. Start instance

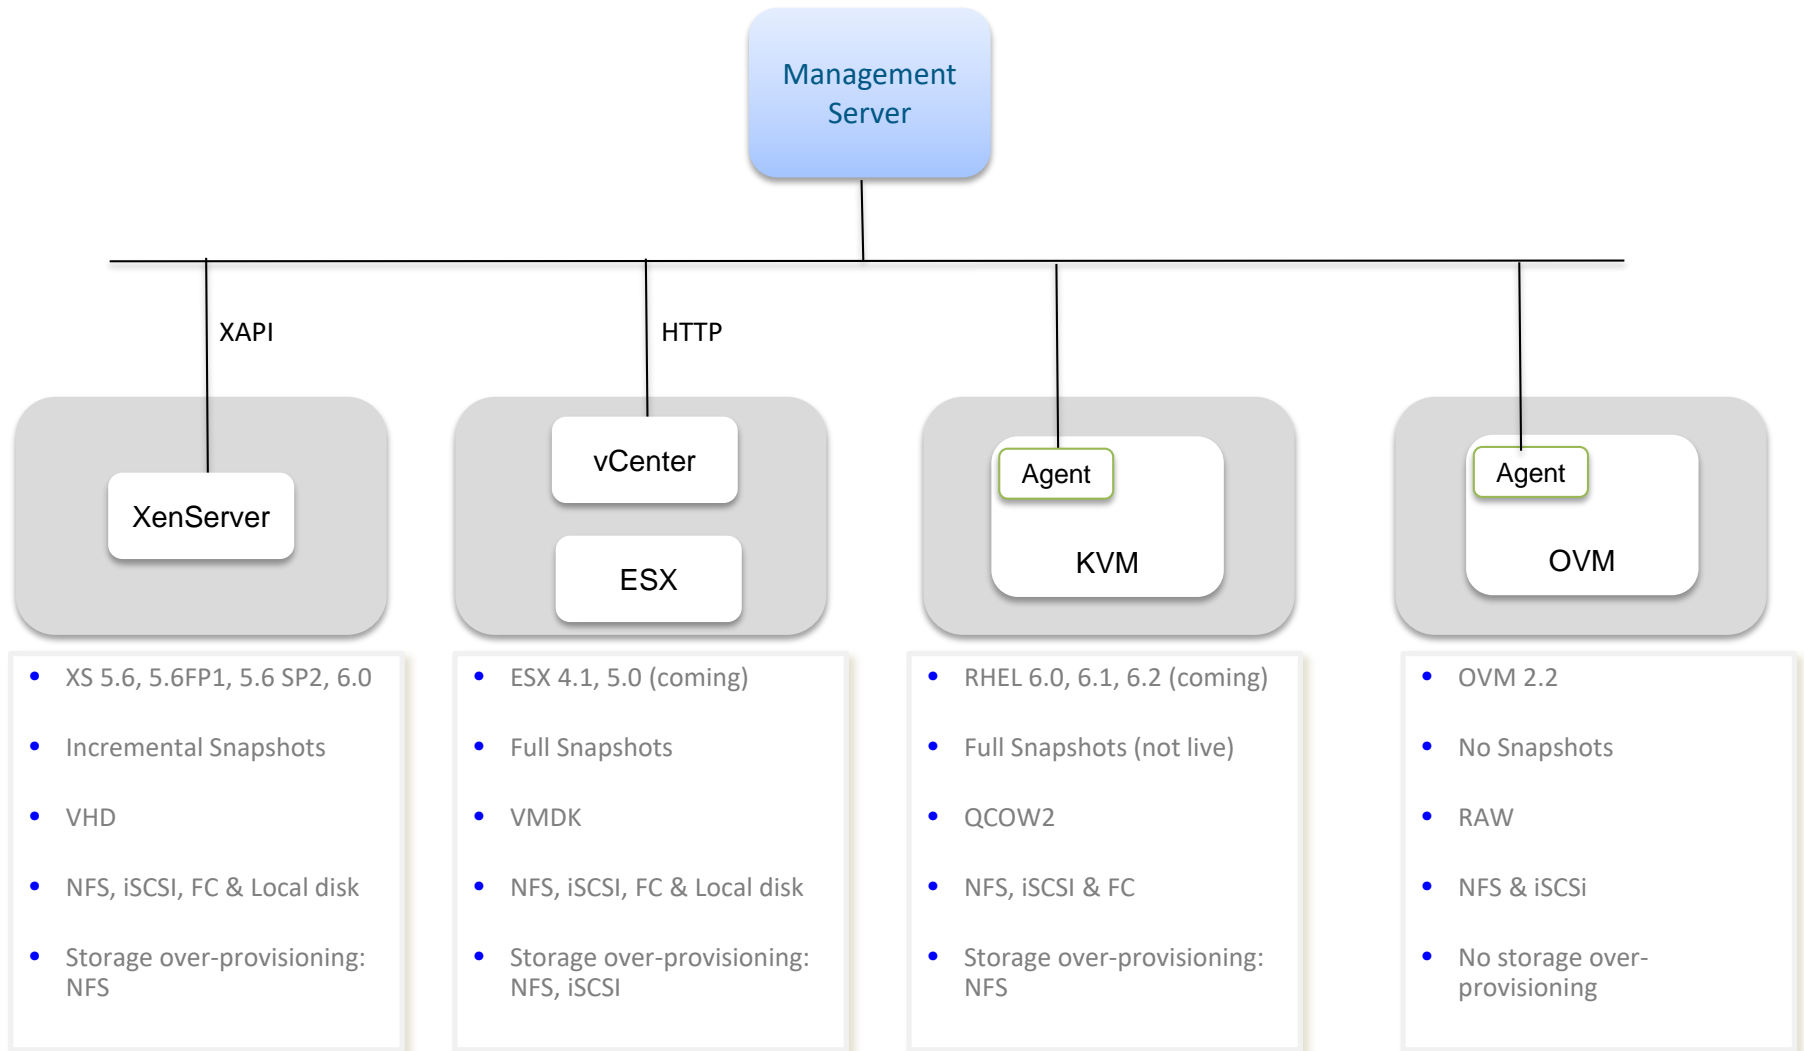


# RedHat Enterprise Linux (KVM)

- Integrates with libvirt using Cloud Agent
- Snapshots at host level
- System VM control channel at host level
- Network management is host level
- Only RHEL 6, not RHEV
  - Also supports Ubuntu 10.04



# Management Server Interaction with Hypervisors



# Multi-tenancy & Account Management

cloudstack Cloud

Resources

VMs, IPs,  
Snapshots...

Domain

Org A



Admin

Domain

Reseller A



Admin

Sub-Domain

Org C



Admin

Resources

VMs, IPs,  
Snapshots...

Account



Group A

Account



Group B



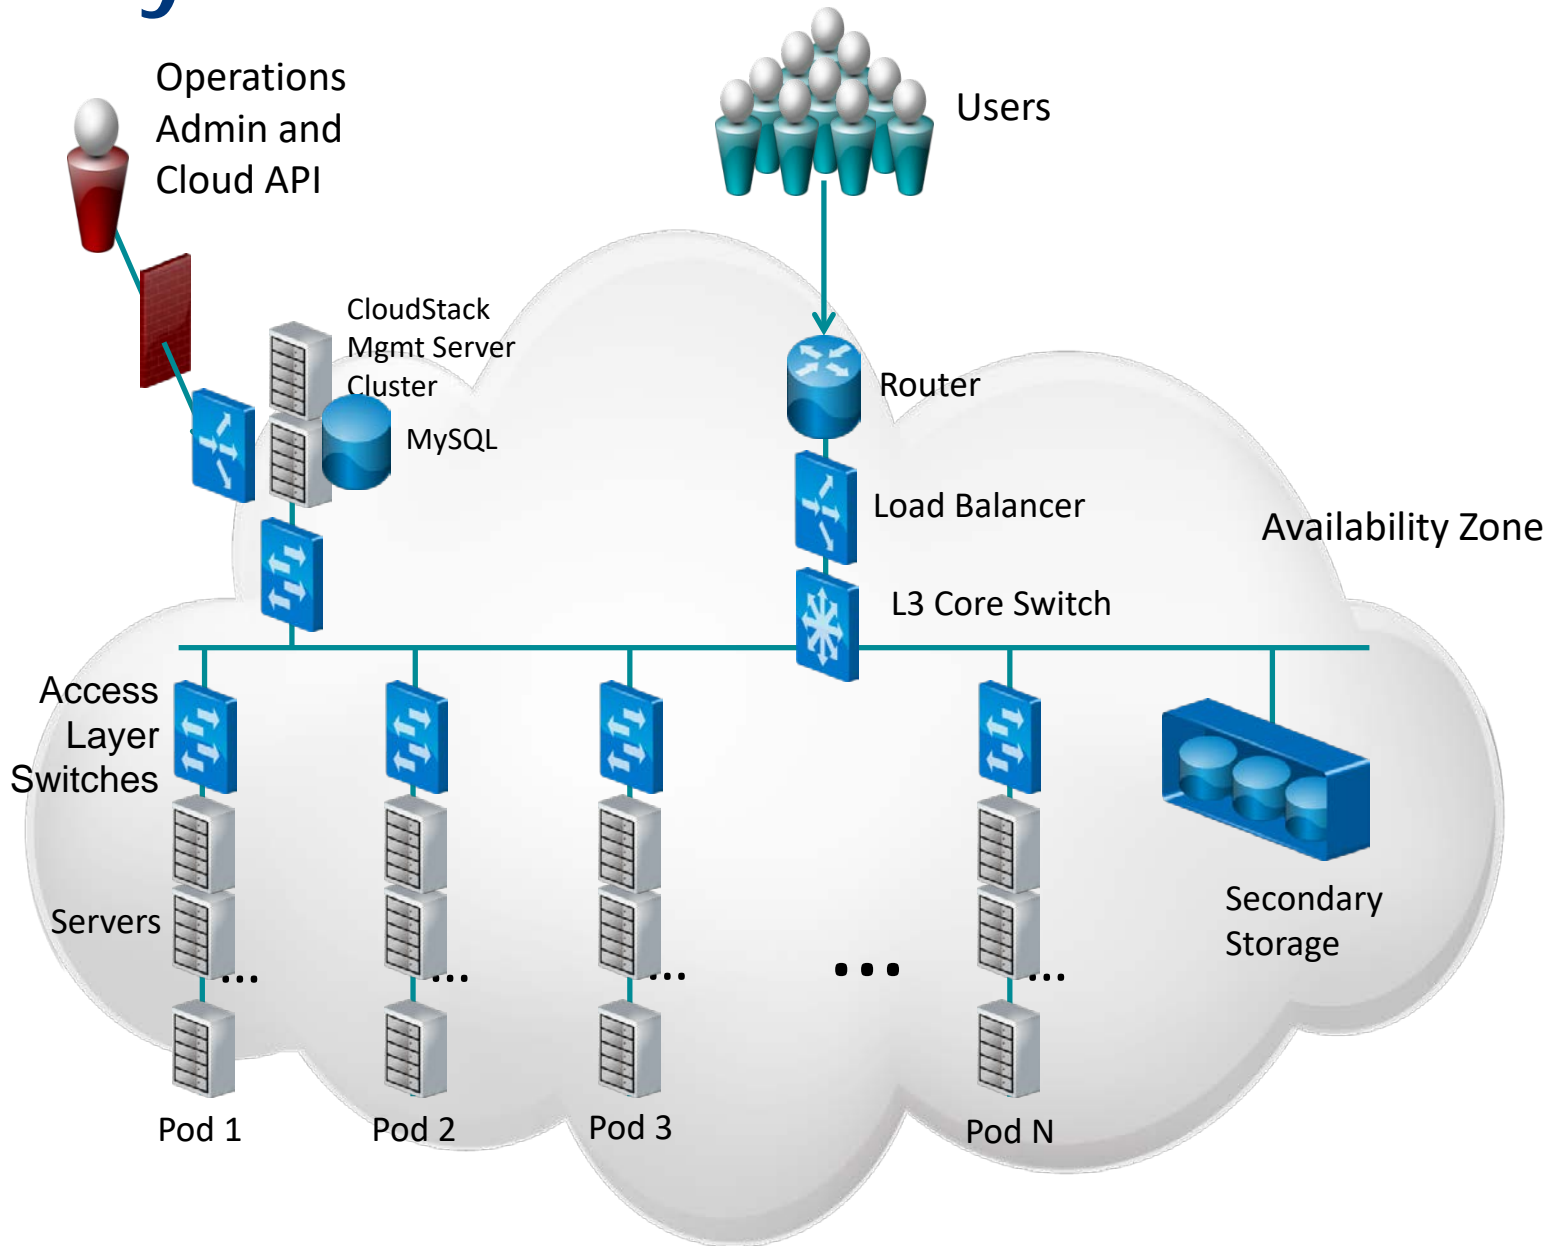
User 1



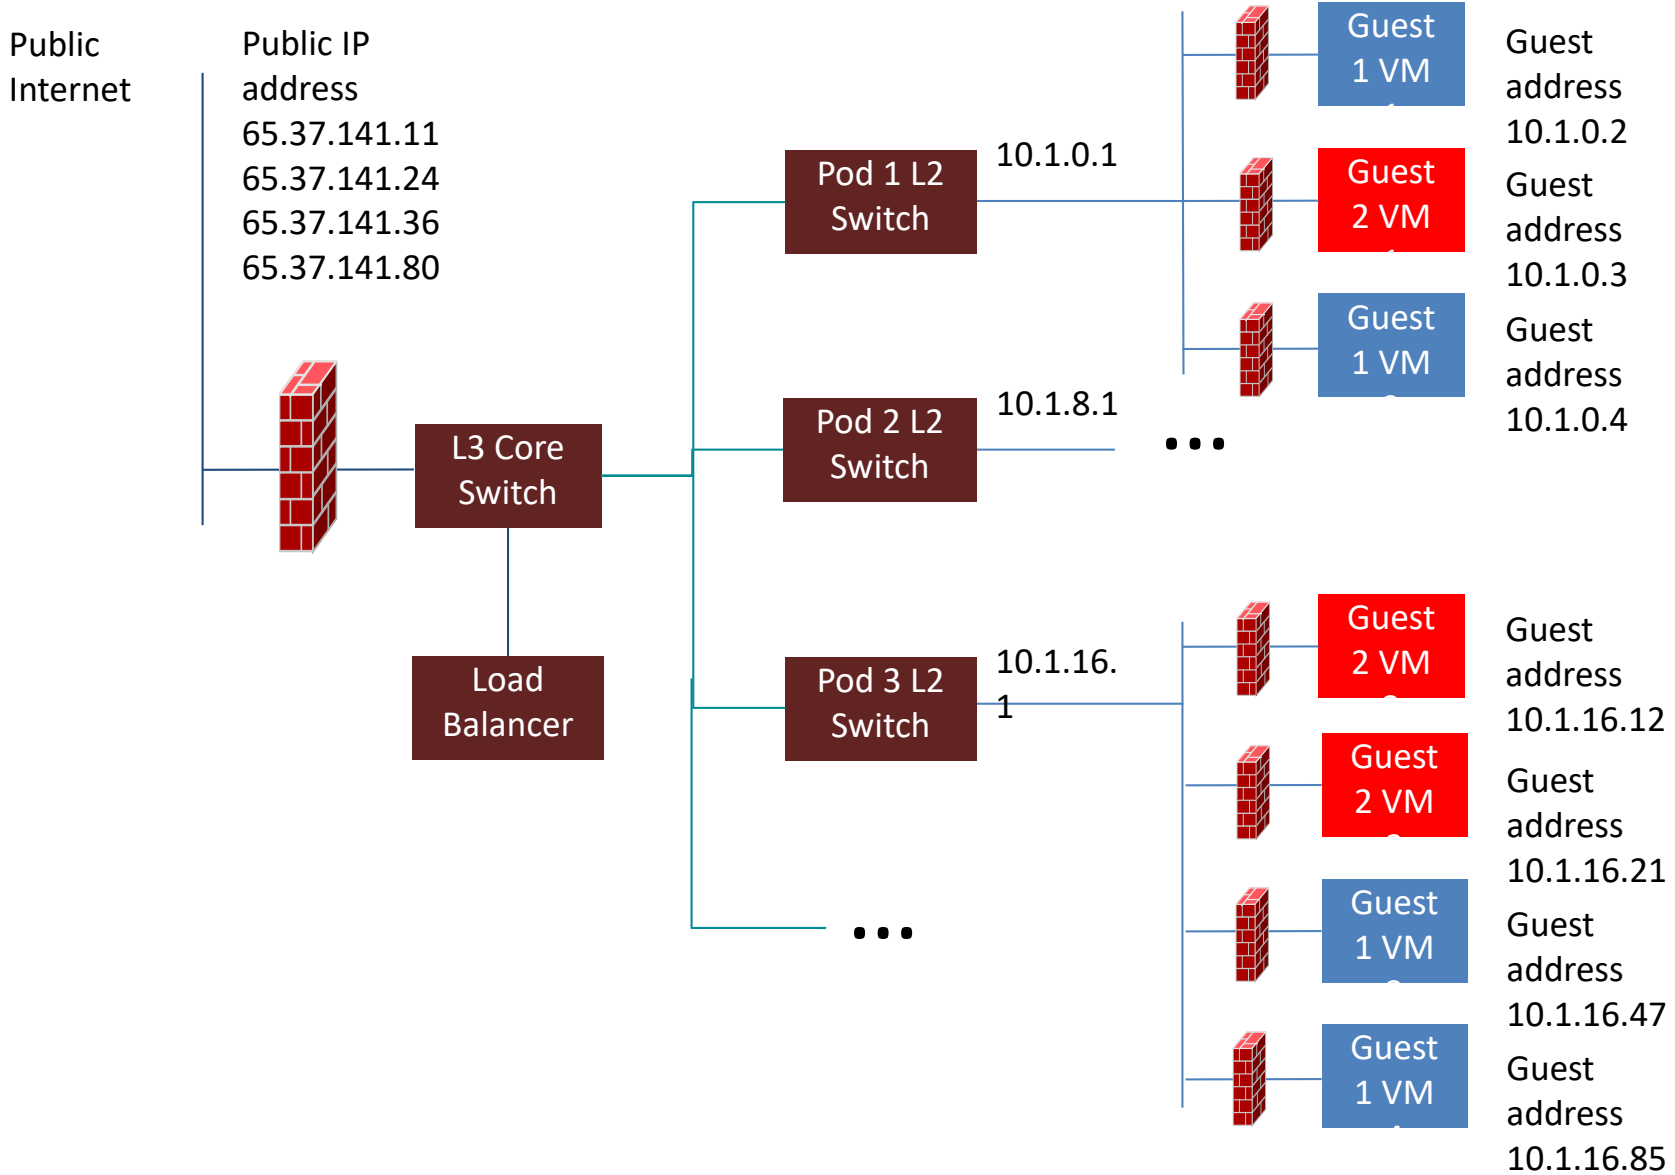
User 2

- Domain is a unit of isolation that represents a customer org, business unit or a reseller
- Domain can have arbitrary levels of sub-domains
- A Domain can have one or more accounts
- An Account represents one or more users and is the basic unit of isolation
- Admin can limit resources at the Account or Domain levels

# Physical Network



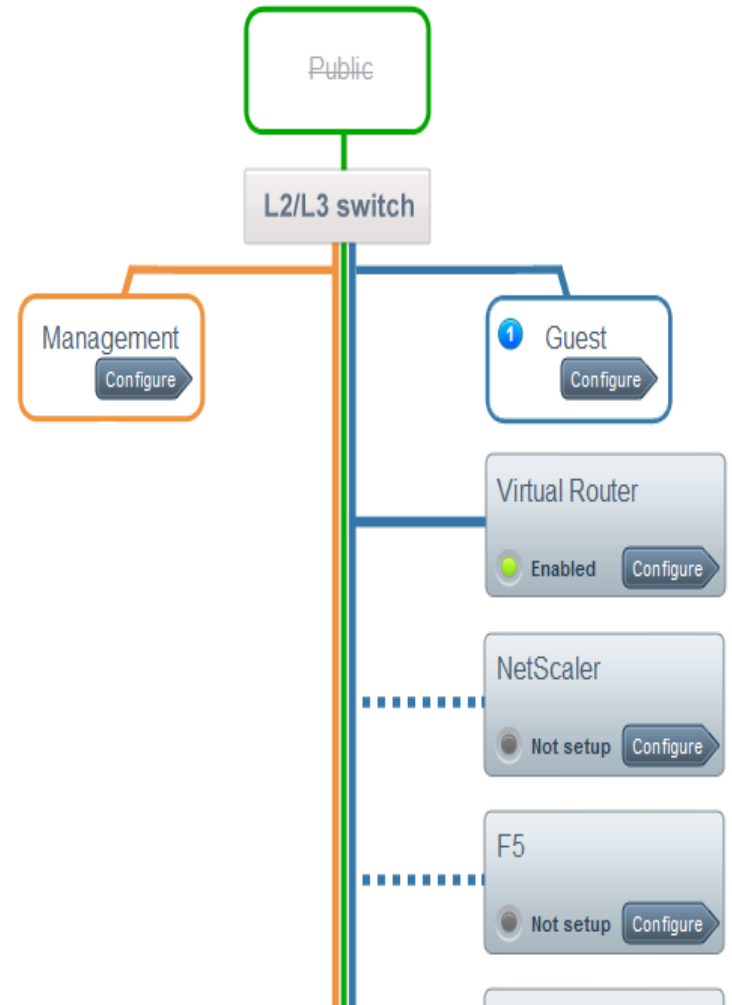
# Guest Networks with L3 isolation



# Network Offerings

- Cloud provider defines the feature set for guest networks
- Toggle features or service levels
  - Security groups on/off
  - Load balancer on/off
  - Load balancer software/hardware
  - VPN, firewall, port forwarding
- User chooses network offering when creating network
- Enables upgrade between network offerings
- Default offerings built-in
  - For classic CloudStack networking

Zone Configuration - Network 1





# Networking Principles in Apache CloudStack

- **Flexibility**
  - Allow various combinations of technology for L2-L7 network services
  - Allow different providers (vendors) for the same network service in a Cloud POP
- **Pluggability**
  - Plugins allow vendors to drop in vendor-specific configuration and lifecycle management code
- **Service scalability**
  - Scale out using virtual appliances when possible
  - Scale up using hardware appliances if needed

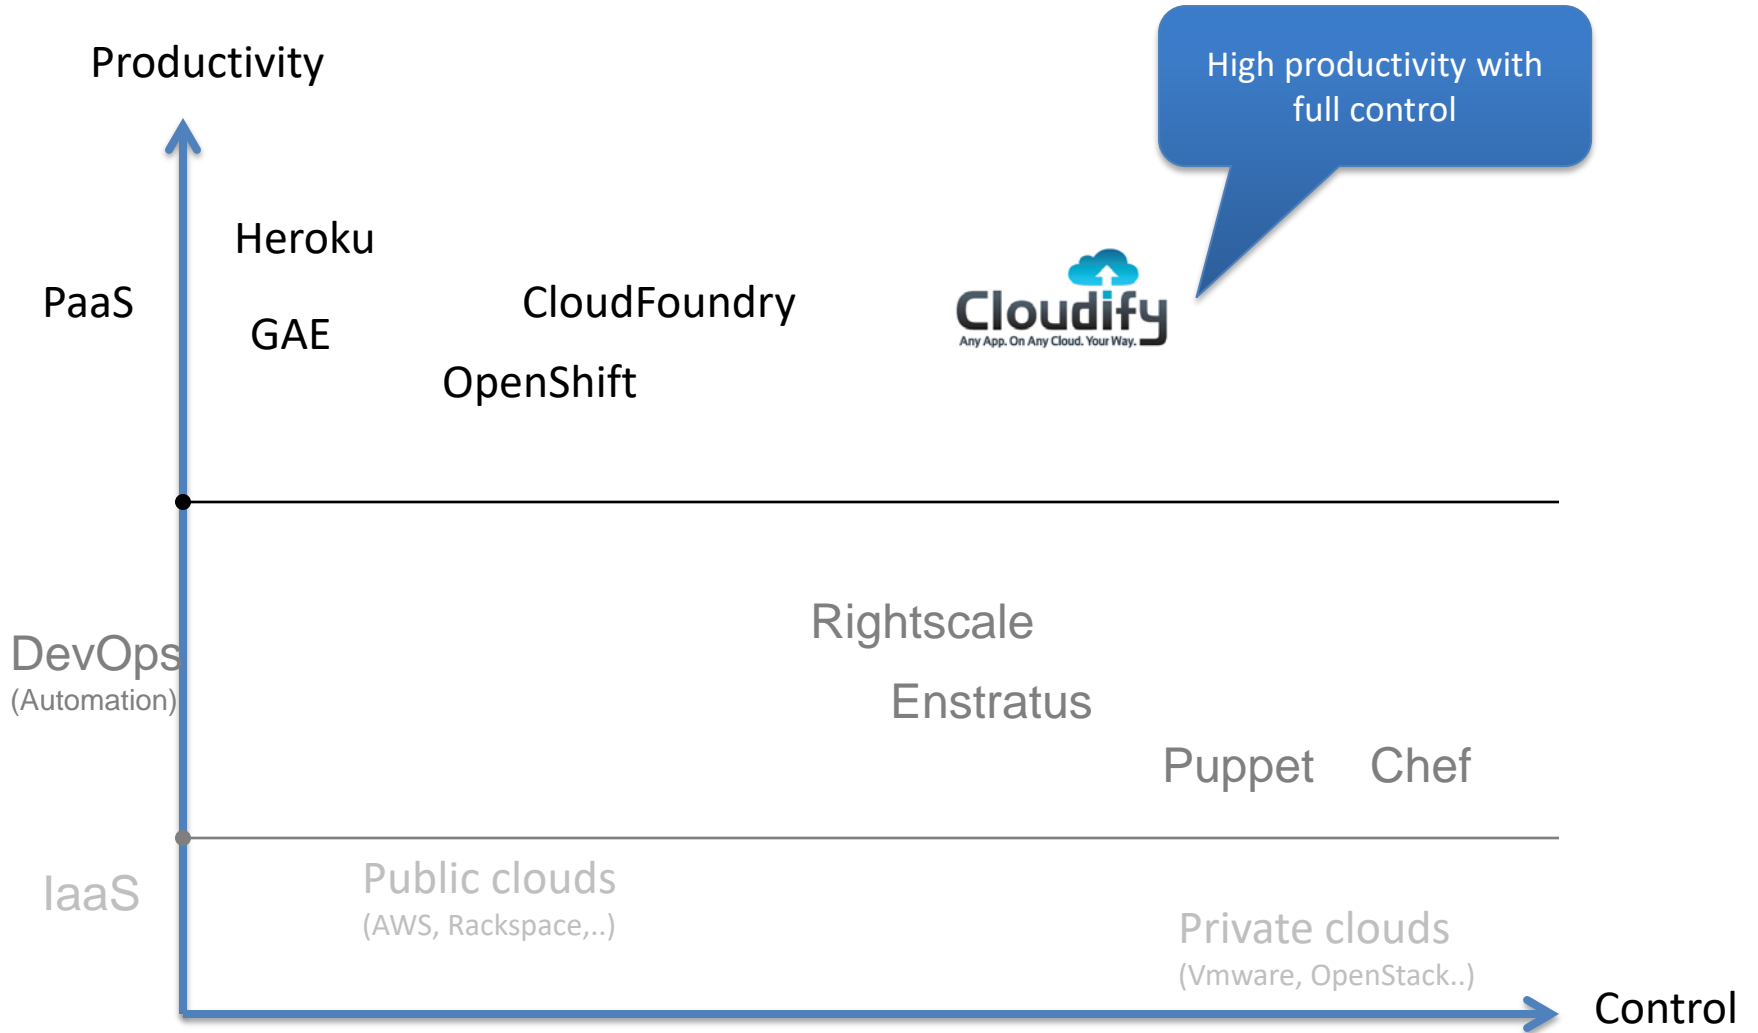


**On-Board and Scale**

**Any App. To Any Cloud. Your Way.**



- Cloudfify is an open source cloud orchestration framework
- It is a **Platform as a Service** (PaaS), a category of cloud computing services that provides a platform allowing customers to develop, run, and manage web applications
- Cloudfify allows to:
  - Model applications and services and automate their entire life cycle
  - Deploy on any cloud or data center environment
  - To monitor all aspects of the deployed application





3

2



4



4

- 1 Upload your recipes
- 2 Cloudfify creates VMs & installs agents
- 3 Agents install and manage your app
- 4 Cloudfify automates the scaling

# Extensive Platform Support



Leverage any Chef Cookbook



Focus on Big Data applications



Dozens of built in recipes GitHub Recipe Repository


# Understand Application Structure & Dependencies



```
application {  
  name="petclinic"  
  service {  
    name = "mongod"  
  }  
  service {  
    name = "mongoConfig"  
  }  
  service {  
    name = "apacheLB"  
  }  
  service {  
    name = "mongos"  
    dependsOn = ["mongoConfig", "mongod"]  
  }  
  service {  
    name = "tomcat"  
    dependsOn = ["mongos", "apacheLB"]  
  }  
}
```

# Auto Scale Your Way

```
scalingRules ([
  scalingRule {
    serviceStatistics {
      metric "Total Requests Count"
      statistics Statistics.maximumThroughput
      movingTimeRangeInSeconds 20
    }
    highThreshold {
      value 1
      instancesIncrease 1
    }
    lowThreshold {
      value 0.2
      instancesDecrease 1
    }
  }
])
```





- LICENSES are not “Cloudable” !!!
- Selection of Cloud Training Tools
  - MOOC
  - Hardware mCloud
  - Software mCloud
  - Lab mCloud
- Selection of Tools and Configurations to be installed in the MECA shared Training Platform
- Selection of Resources by each Partner