Xen and CloudStack

Ewan Mellor

Director, Engineering, Open-source Cloud Platforms

Citrix Systems

August 27-28, 2012
San Diego, CA, USA

Agenda

- What is CloudStack?
- Move to the Apache Foundation
- CloudStack architecture on Xen
- The future for CloudStack
- Developing on CloudStack

CloudStack is…

- Infrastructure-as-a-Service software
- For people to build Amazon-style clouds
  - either as a public service,
  - or as a private cloud within your own organization.

CloudStack is…

Compute

Hypervisor

- XenServer
- Oracle VM
- vSphere
- KVM
- Bare metal

Storage

Block & Object

- Local Disk
- iSCSI
- Fiber Channel
- NFS
- Swift

Network

Network & Network Services

- Network Type
- Isolation
- Firewall
- Load balancer
- VPN

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			
☐	CustomerNet	Direct	◯ Default
☐	CustomerNet	Direct	◯ Default
☐	vlan100	Direct	◯ Default
☐	vlan100	Direct	◯ Default
☐	vlan100	Direct	◯ Default

Add new network	
☐	NEW

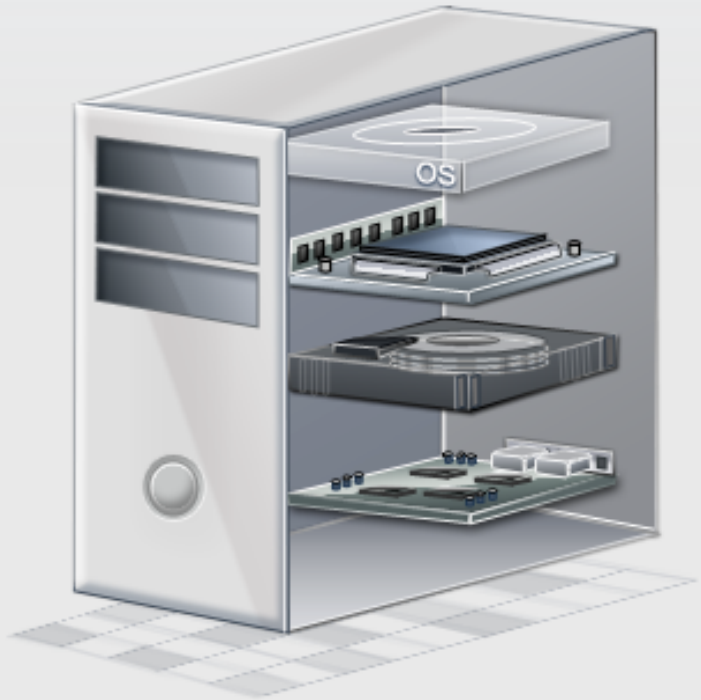

Previous

Cancel

Next

CloudStack versus OpenStack

CloudStack	**OpenStack**
Apache Foundation (incubating)	OpenStack Foundation (founding in progress)
2008-present	2010-present
Java	Python
Fairly centralized	Aggressively distributed
Complete solution	More of a framework

Apache incubation

- Citrix donated the CloudStack code to the Apache Software Foundation and Apache Incubator
- The project now needs to prove that it can operate openly and independent of Citrix
- If it can do that, it can be accepted as a top-level Apache project.

CloudStack and Xen

Most CloudStack users choose Xen or XenServer

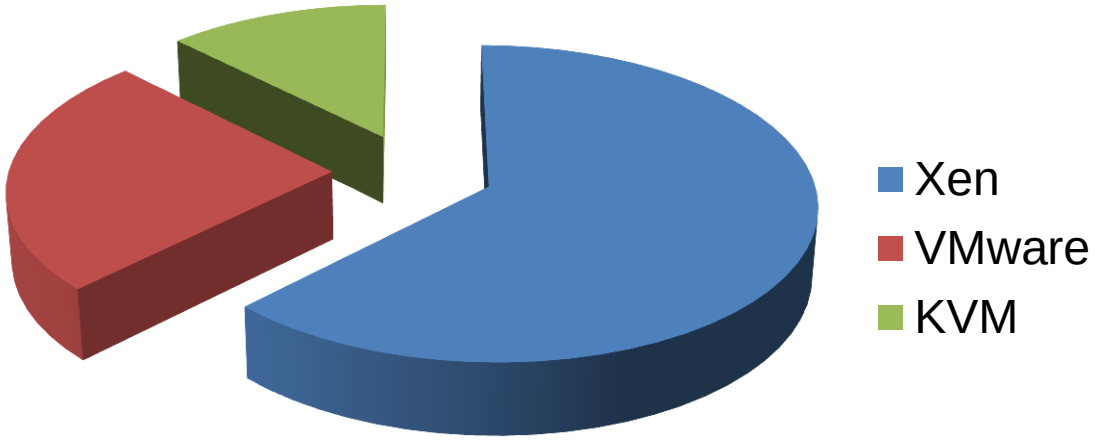

It is the combination that works the best!

CloudStack is…

- A cluster of management servers
- An orchestration engine
- State and user management
- Region, zone, pod, cluster model
- A large collection of effectors
- System VMs
- Usage reporting
- User interface

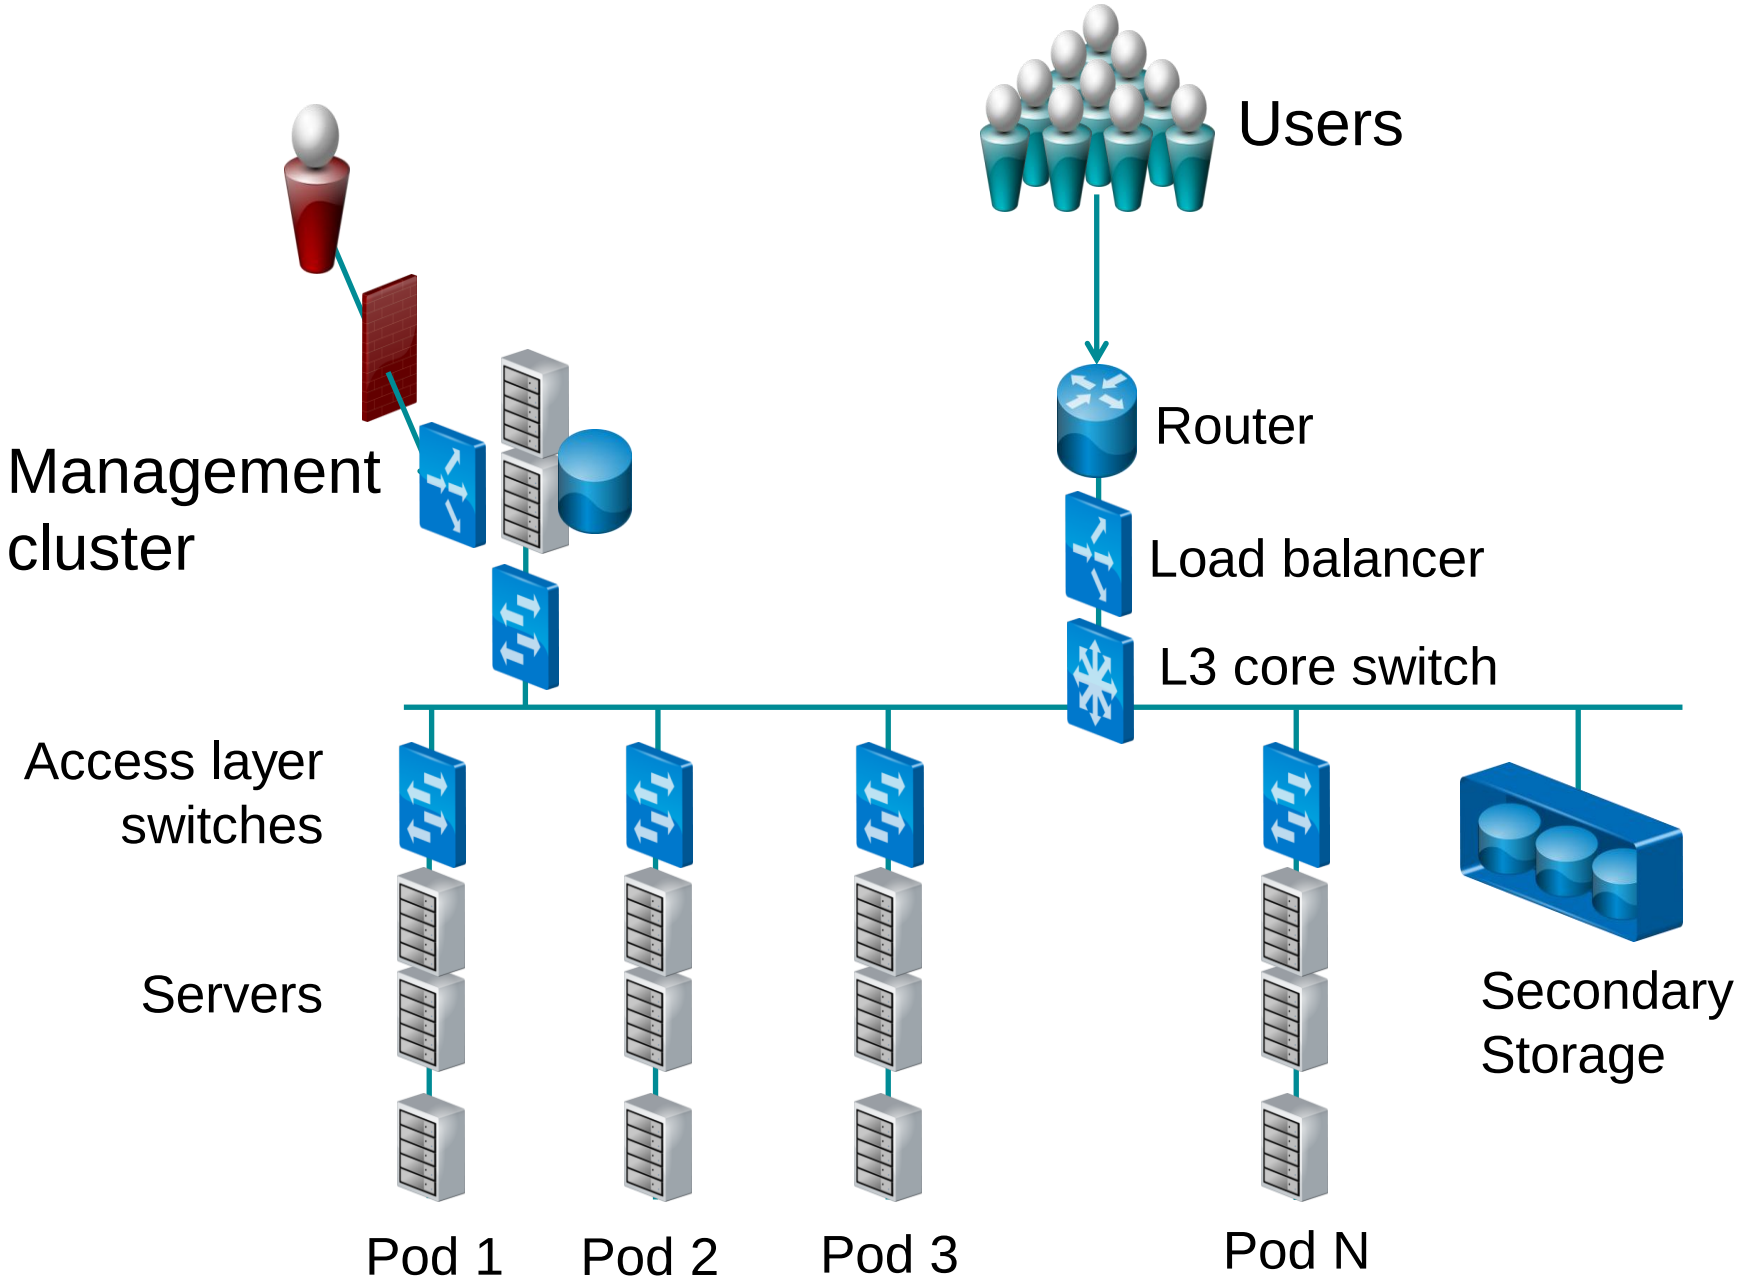
Users
Router
Load balancer
L3 core switch
Management cluster
Access layer switches
Servers
Secondary Storage
Pod 1
Pod 2
Pod 3
Pod N

Two types of storage

Primary storage

- Stores virtual disks
- Configured at cluster level
- Close to hosts for better performance
- Requires high IOPS

Secondary storage

- Stores templates, ISOs and snapshots
- Configured at zone level
- Zone can have one or more
- High capacity, low cost

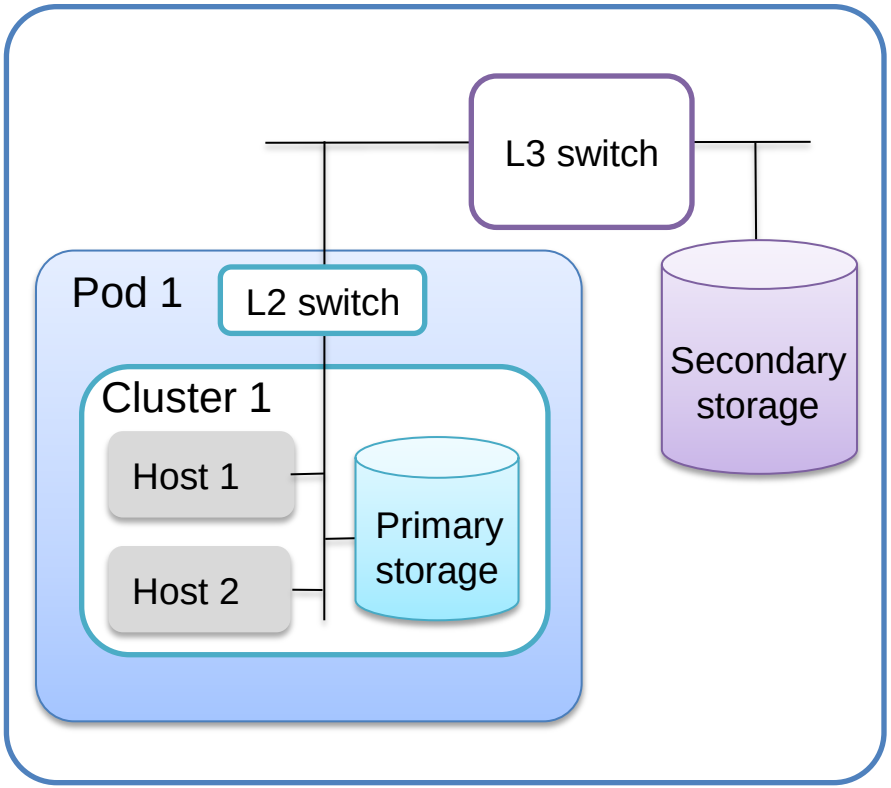

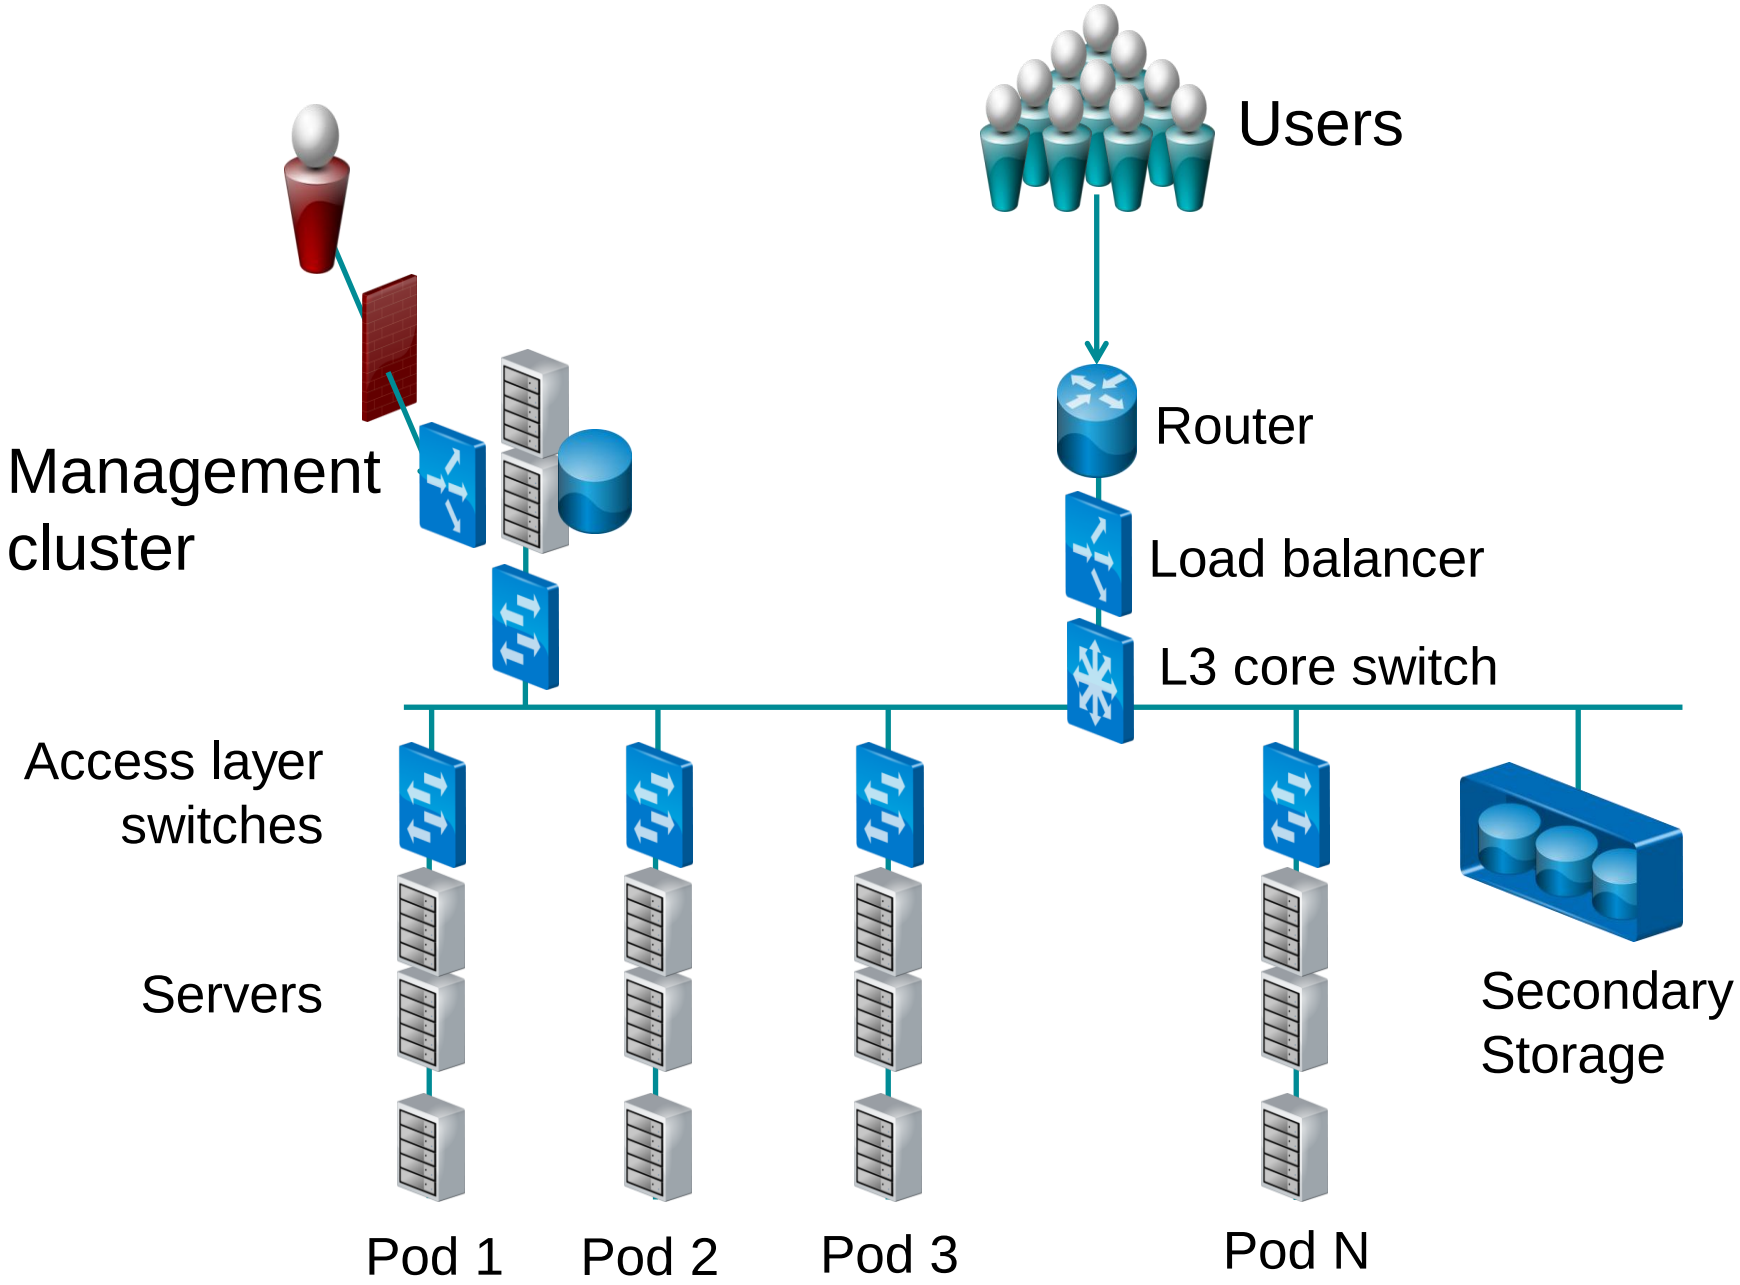
Users
Router
Load balancer
L3 core switch
Management cluster
Access layer switches
Servers
Secondary Storage
Pod 1
Pod 2
Pod 3
Pod N

Networking features

- Choice of network isolation
  - Physical, VLAN, L3 (anti-spoof), overlay (GRE)
- Multiple networks
  - Shared networks, project networks
- IPAM / DHCP
- Gateway (inc VRRP)
- VPN, stateful firewall, NAT, port forwarding
- Traffic monitoring
- Load balancing
- User-data / password-change services

Virtual network topologies

CS Virtual Router provides network services

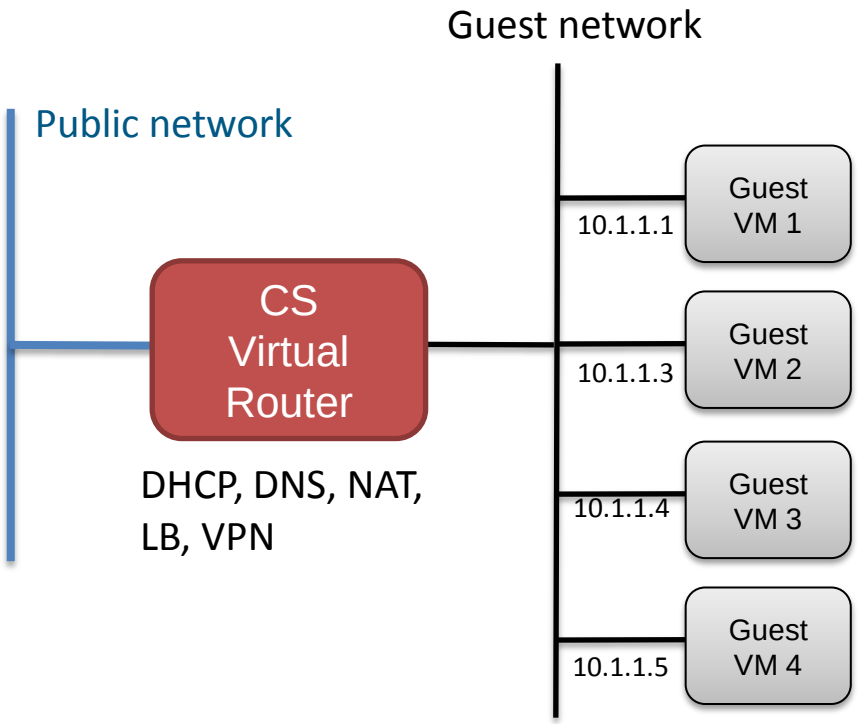

External devices provide network services

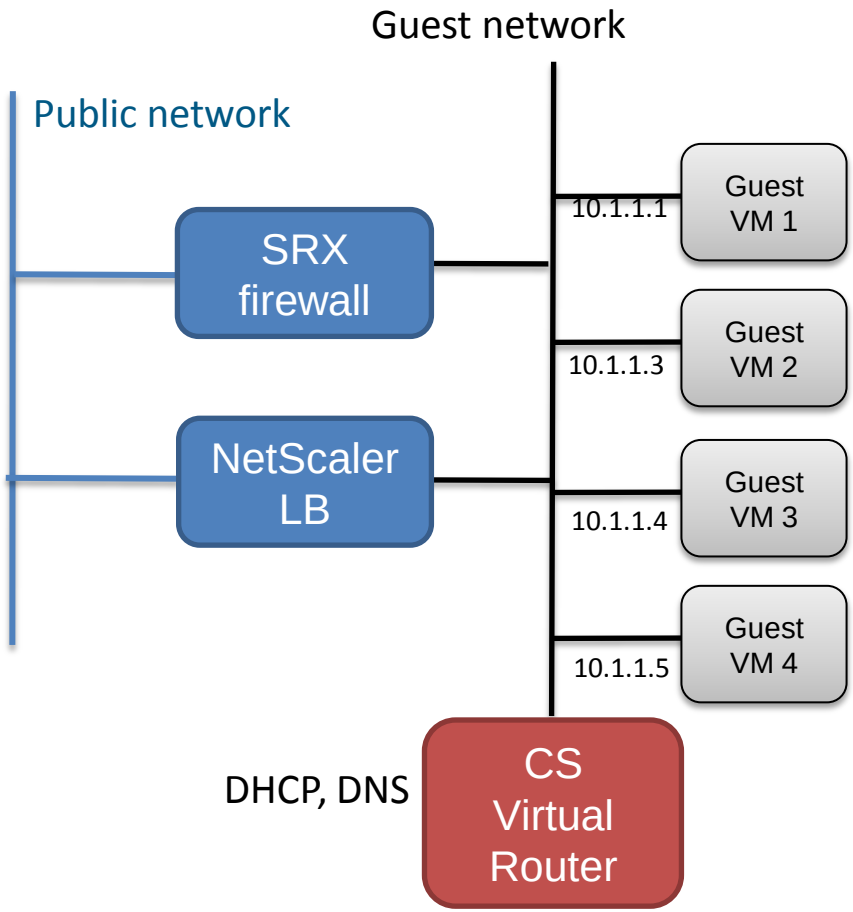

Multi-tier virtual network topology

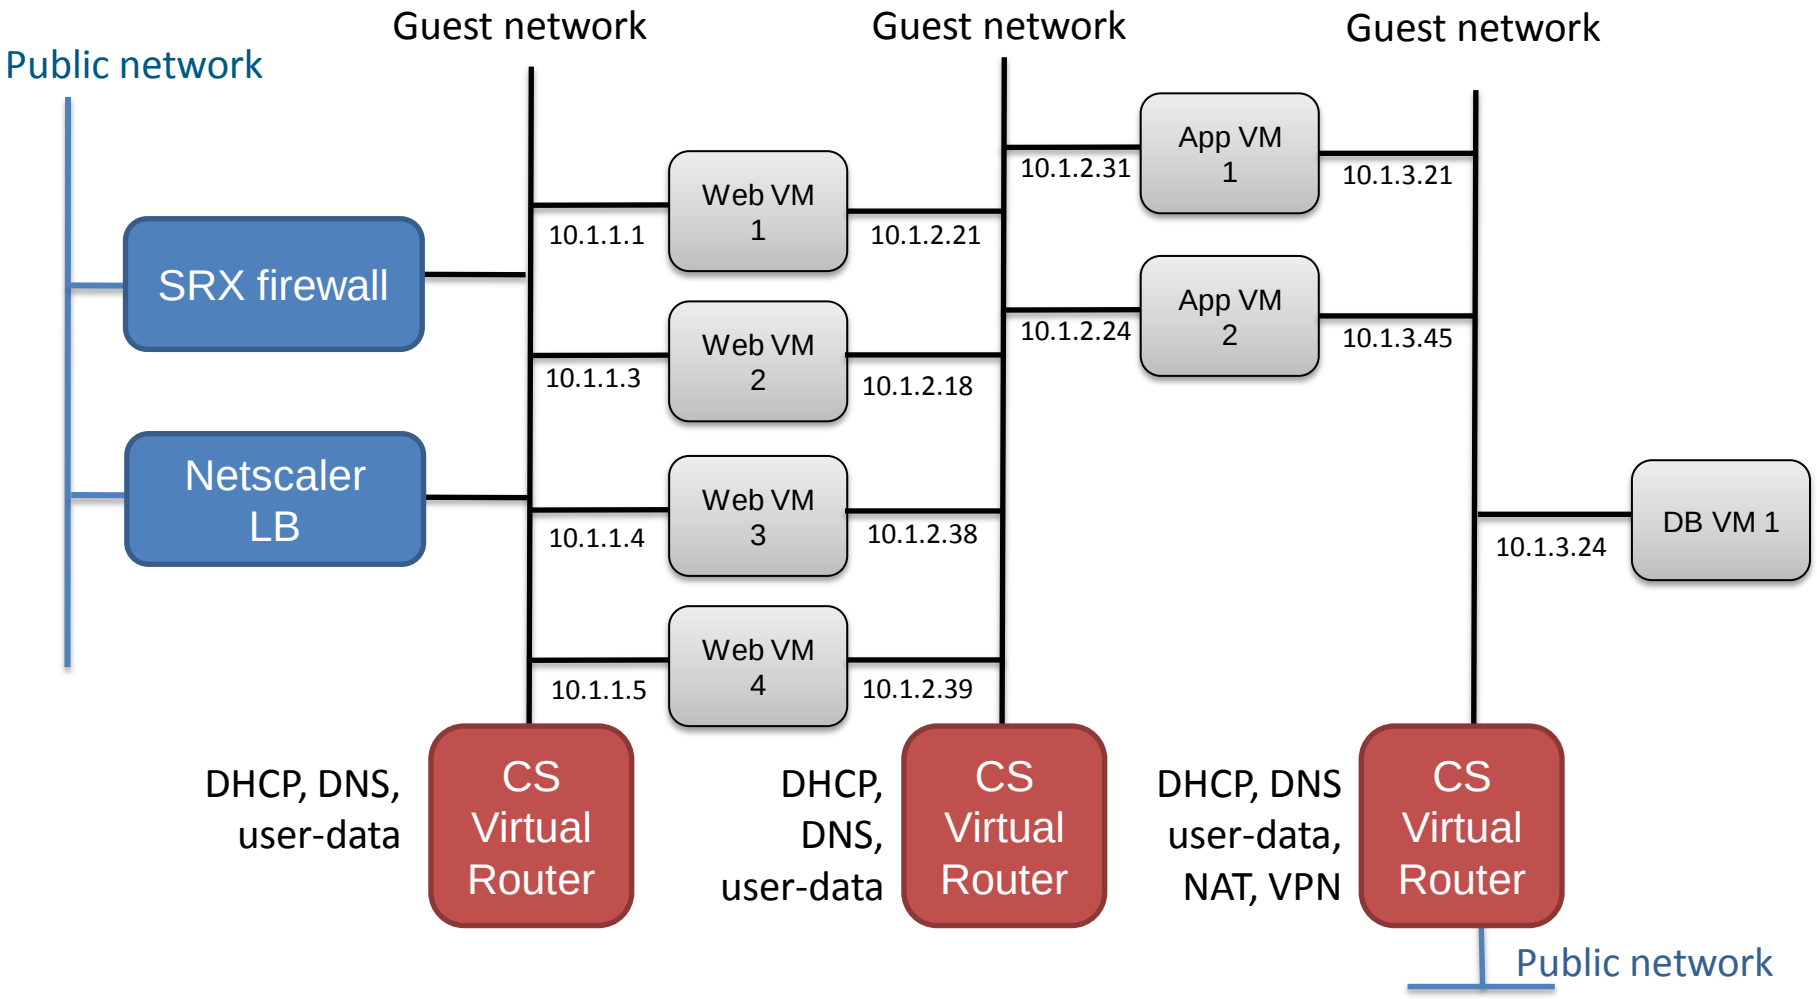

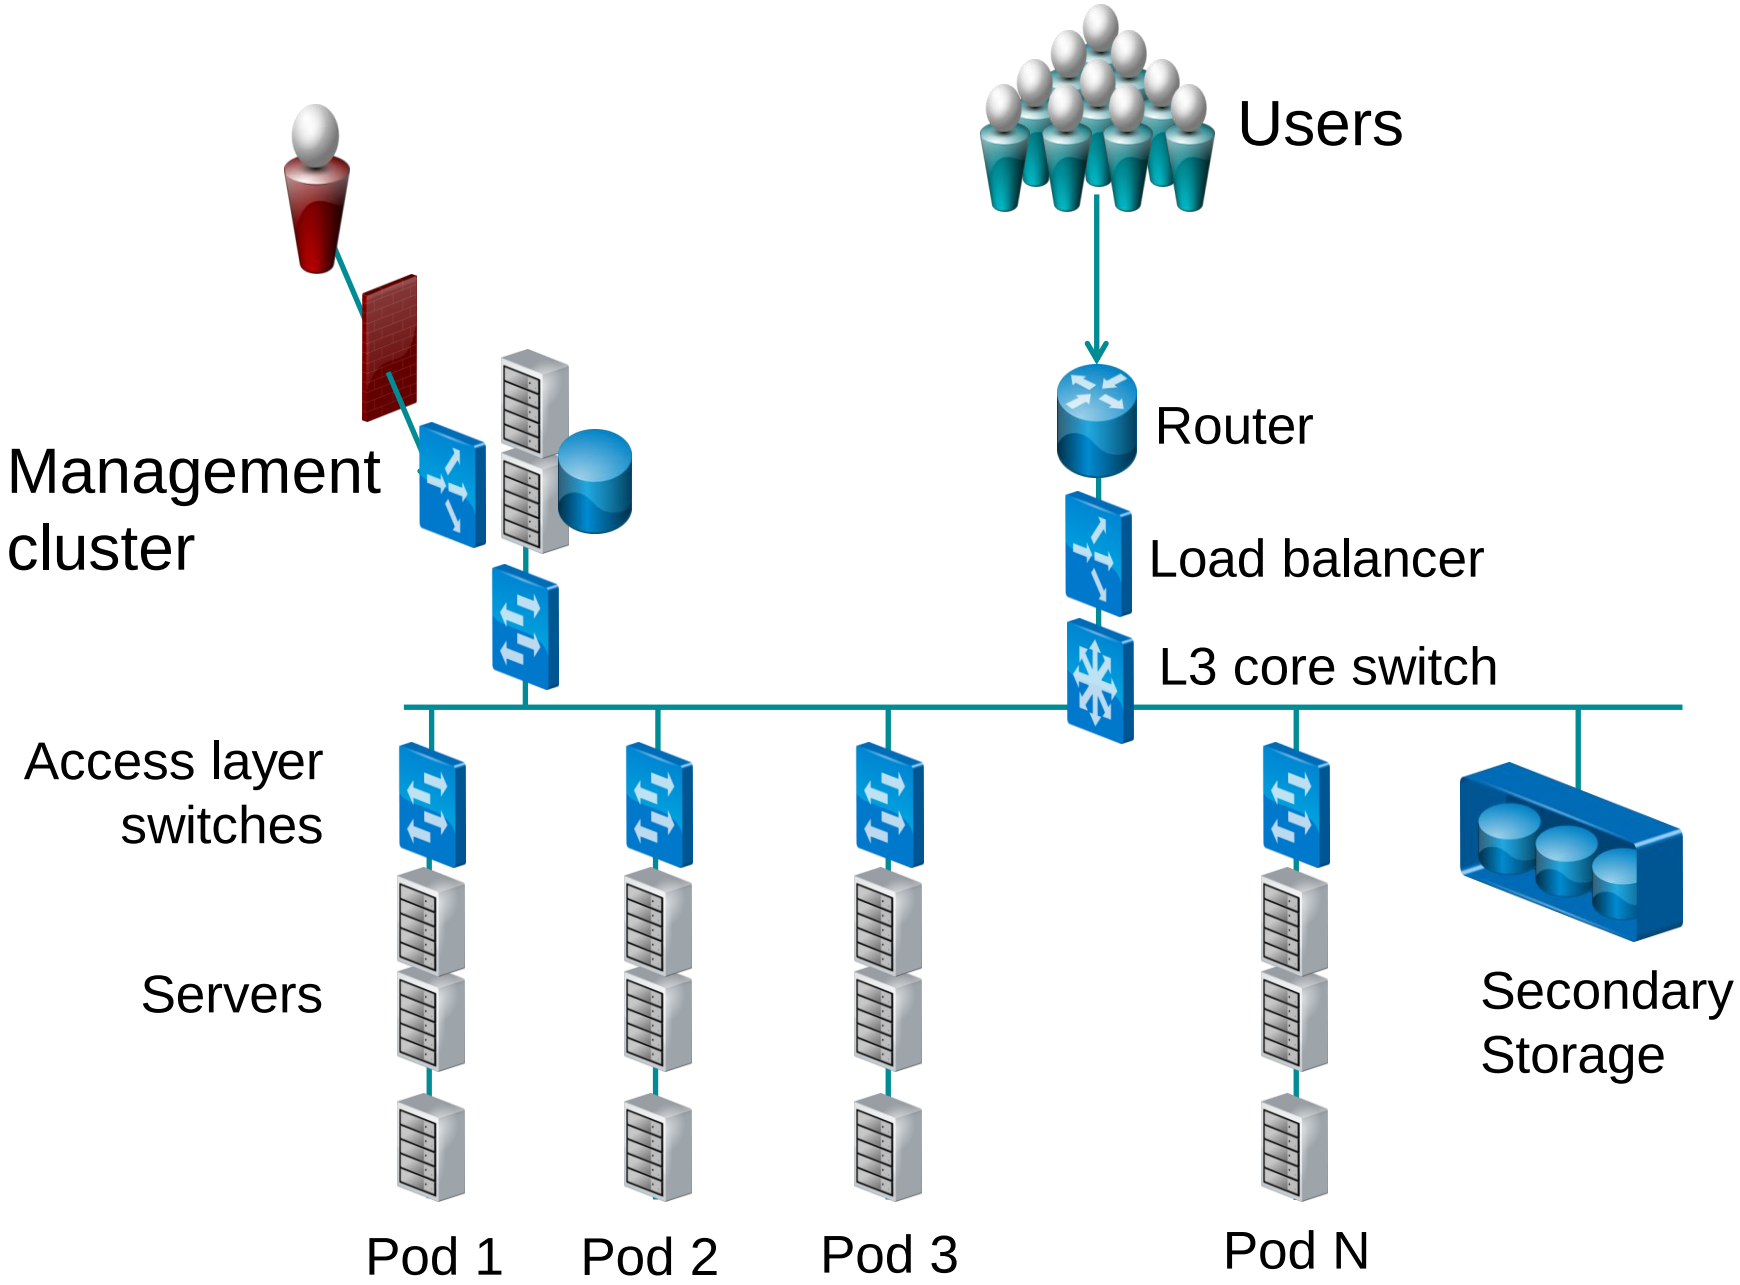
Users
Router
Load balancer
L3 core switch
Management cluster
Access layer switches
Servers
Secondary Storage
Pod 1
Pod 2
Pod 3
Pod N

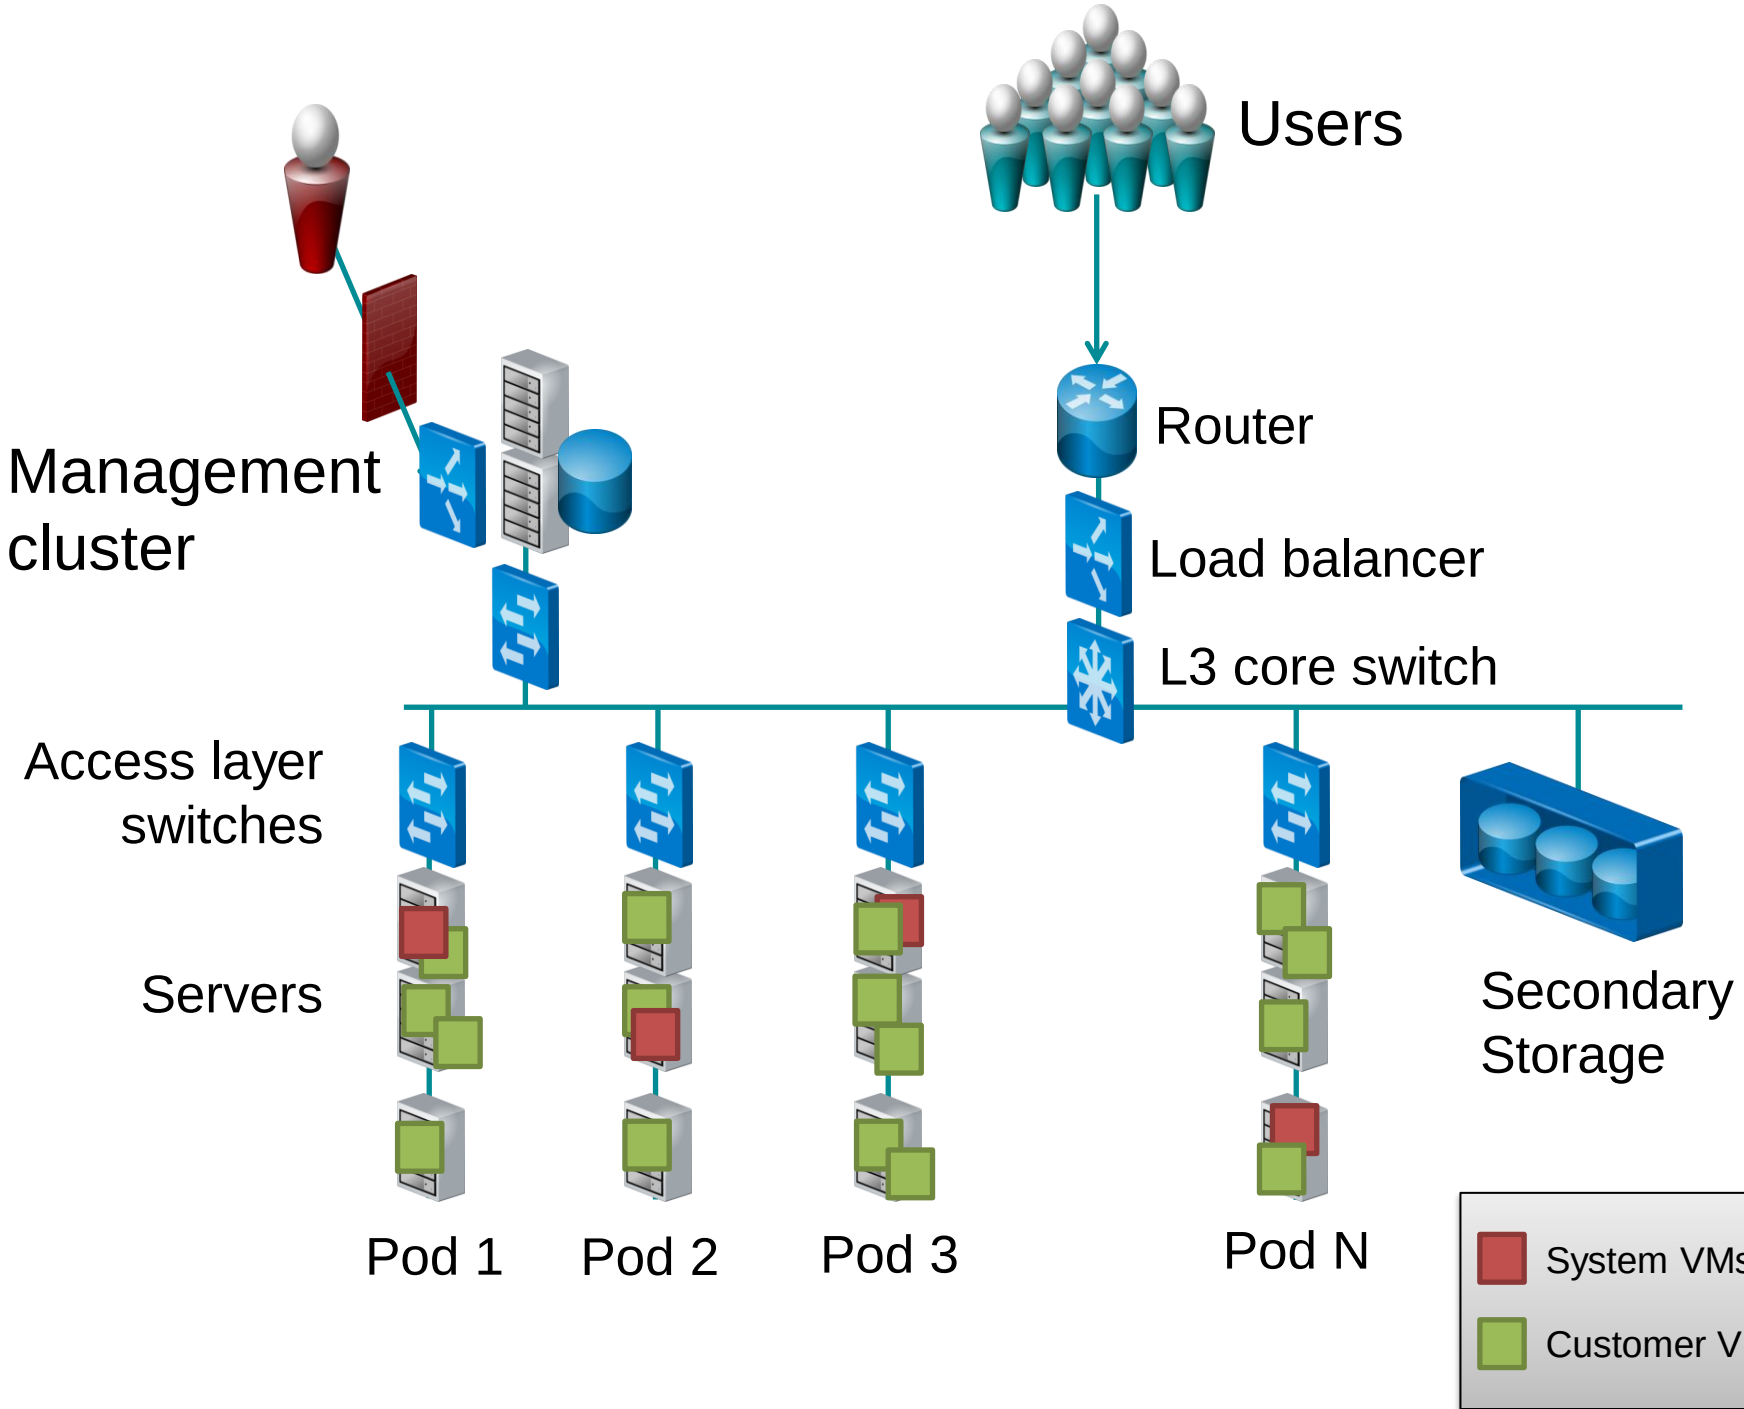
Users
Router
Load balancer
L3 core switch
Management cluster
Access layer switches
Servers
Secondary Storage
Pod 1
Pod 2
Pod 3
Pod N
System VMs
Customer VMs

Driver domains

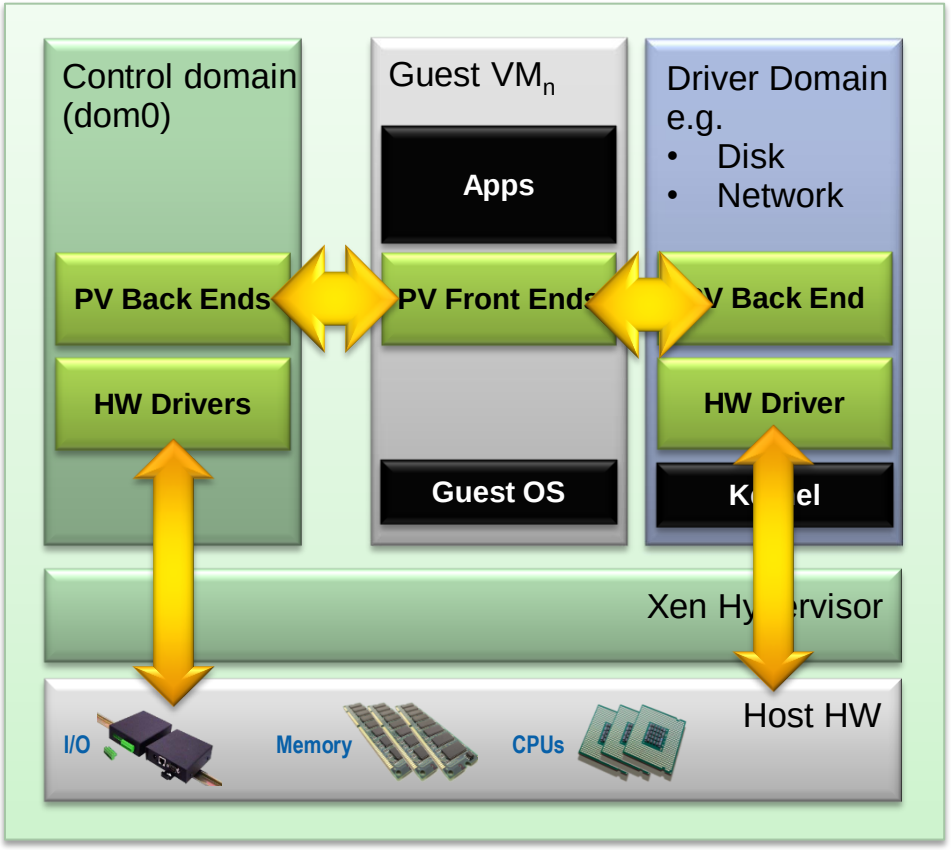

The road ahead

The future for CloudStack

- Disentangle some core components
  - Make it easier for people to work on subsystems
  - Make it easier to scale the API layers

S3 API system VM

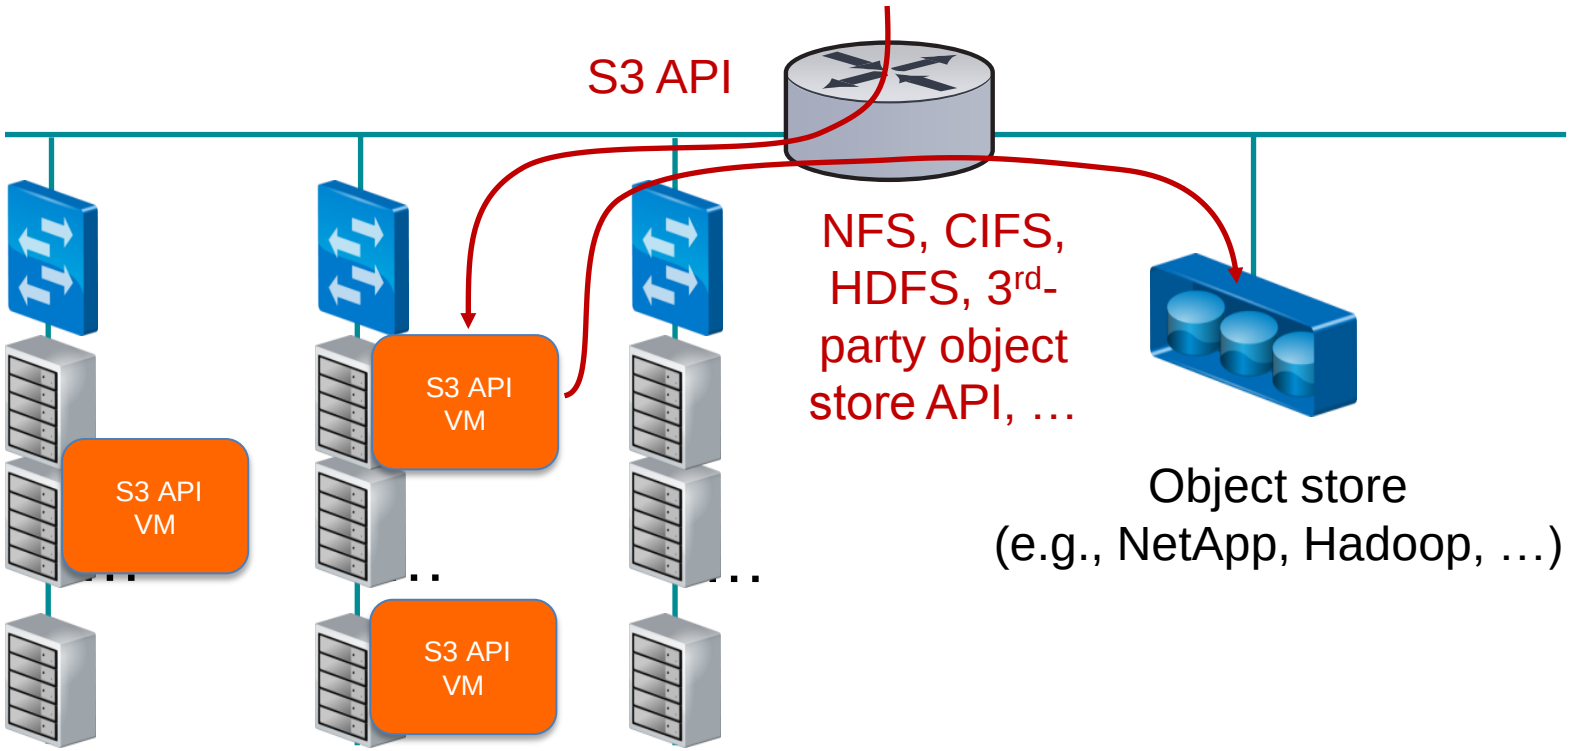

The future for CloudStack

- Absorb new changes from Xen and XS
  - Storage migration
  - The Windsor architecture
  - I/O subsystem changes

The future for CloudStack

- New object storage systems
  - Caringo CAStor
  - HDFS
  - Riak CS
- New block storage systems
  - GlusterFS?
  - Ceph?
- Open vSwitch / OpenFlow / VXLAN

DevCloud

DevCloud – What is inside?

- VirtualBox VM
  - Xen Cloud Platform
  - Ubuntu 12.04 dom0 (aka Kronos)
    - JVM, Tomcat, developer tools
    - NFS server
    - MySQL
  - Working, pre-configured CloudStack
  - System VMs
  - Tiny Linux template

Why?

- Eliminates
  - need for another machine / hypervisor
  - Primary and secondary storage
  - Networking, VLAN, etc
  - Install and configuration of environment
- Disposable changes
  - Revert to snapshot

Run

- Start the DevCloud VM from VirtualBox UI
- Access the CloudStack UI
  - http://localhost:8080/client
- Access the dom0
  - ssh root@localhost –p 2222
- From the UI, enable the zone
  - System VMs will start up

cloudstack

open source cloud computing

Xen®

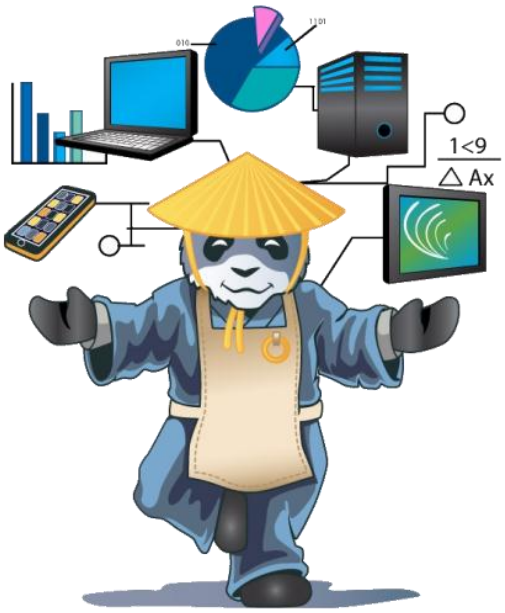

CloudStack architecture

Management Server Cluster

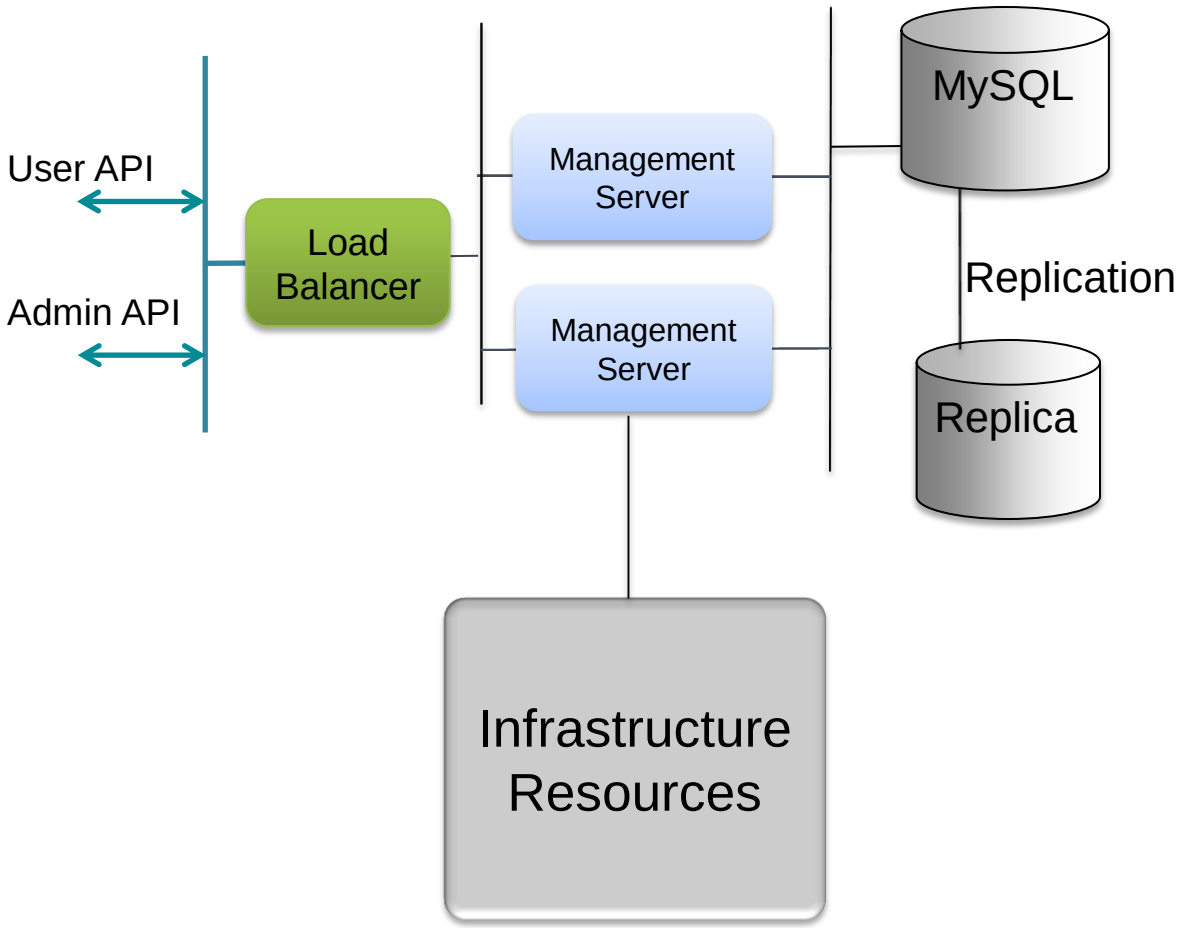

- MS is stateless. MS can be deployed as physical server or VM
- Single MS node can manage up to 10K hosts. Multiple nodes can be deployed for scale or redundancy

Components

- Hosts
  - Servers onto which services will be provisioned
- Primary Storage
  - VM disk storage
- Cluster
  - A grouping of hosts and their associated storage
- Pod
  - Collection of clusters in the same failure boundary
- Network
  - Logical network associated with service offerings
- Secondary Storage
  - Template, snapshot and ISO storage
- Zone
  - Collection of pods, network offerings and secondary storage

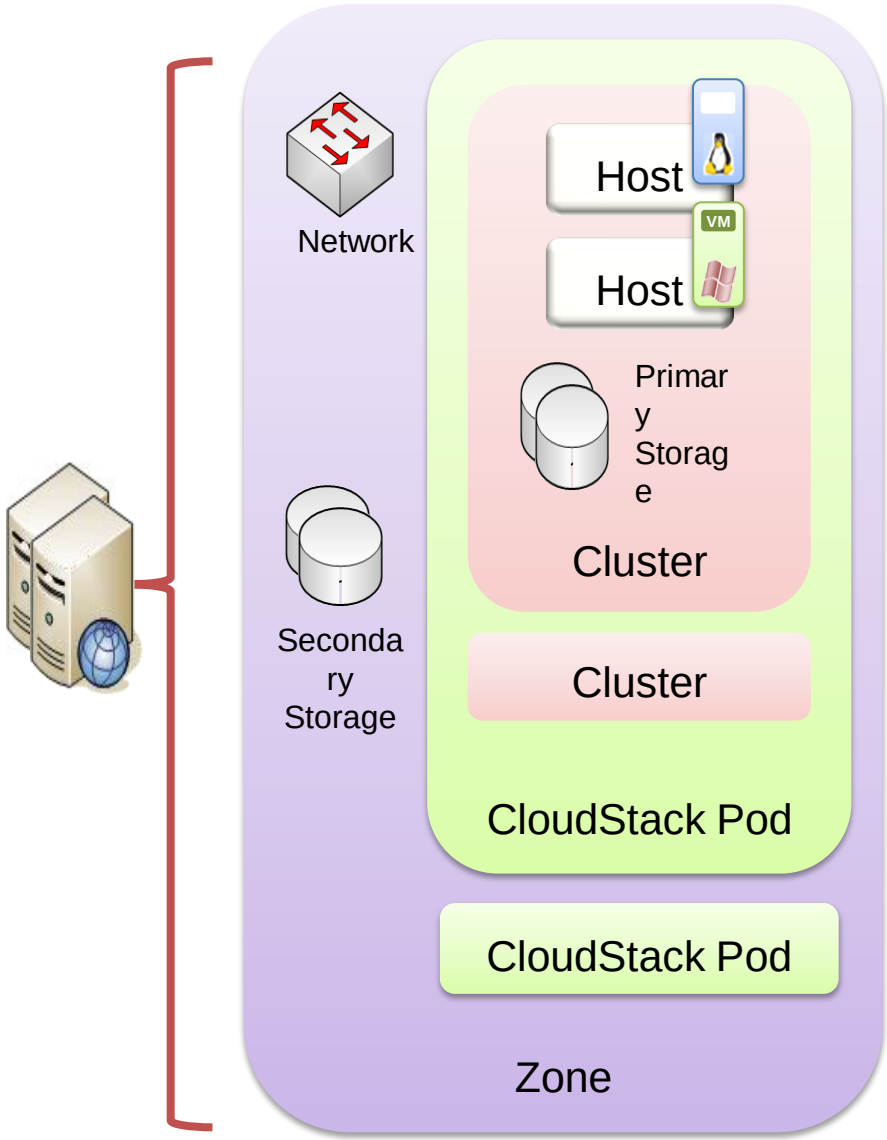

Deployment Architecture

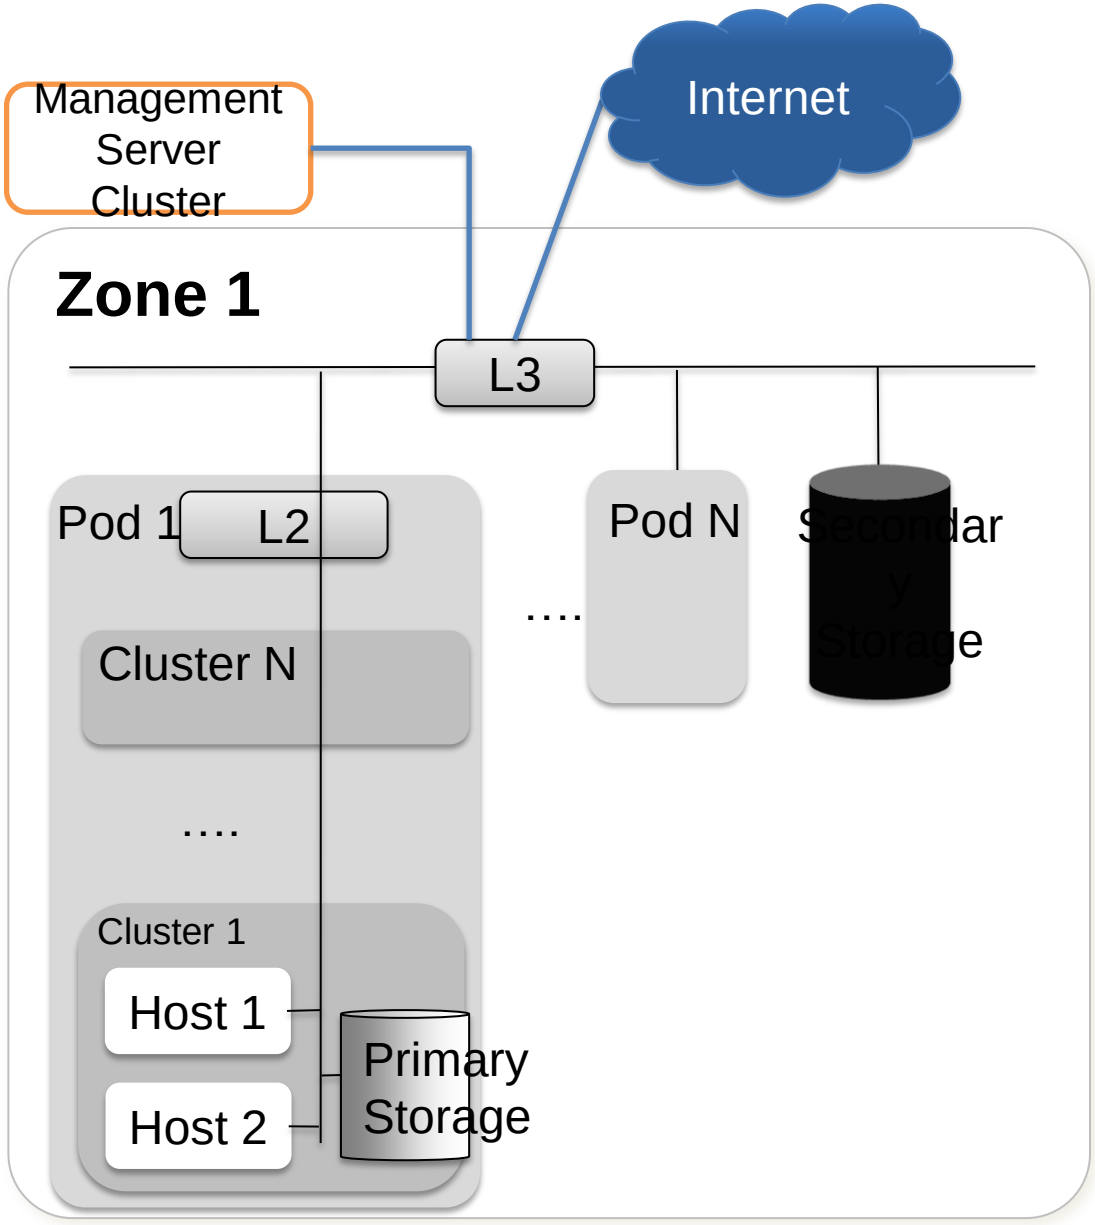

- Hypervisor is the basic unit of scale.
- Cluster consists of one ore more hosts of same hypervisor
- All hosts in cluster have access to shared (primary) storage
- Pod is one or more clusters, usually with L2 switches.
- Availability Zone has one or more pods, has access to secondary storage.
- One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

DevCloud step-by-step

Overview

- Install Virtual Box
- Download DevCloud virtual machine
- Install DevCloud VM
- Run the VM
- Login to the CloudStack GUI
- Start a CloudStack VM.

Virtual Box

- Download / install from Oracle
- Ensure that VT-d and virtualization extensions are enabled in your laptop CPU
- Recommend at least 2 GB for DevCloud VM

DevCloud VM

- Download from http://download.cloud.com/templates/devcloud/DevCloud.ova
- Latest install instructions here:

  http://wiki.cloudstack.org/display/COMM/DevCloud
- Import DevCloud into VirtualBox
  - Change memory allocation to 2 GB
- Snapshot for known good state