



Dr. Selvakumar Manickam

selva@usm.my



This is what a 500 MB hard drive is like in 1956 (no)

Just one
MicroSD card
stores more than
the rest combined...


25 years
of storage




Many Cloud Providers

- AWS: Amazon Web Services
 - EC2: Elastic Compute Cloud
 - S3: Simple Storage Service
 - EBS: Elastic Block Storage
- Microsoft Azure
- Ali Cloud
- Google Compute Engine
- Rightscale, Salesforce, EMC, Gigaspaces, 10gen, Datastax, Oracle, VMWare, Yahoo, Cloudera
- And many many more!

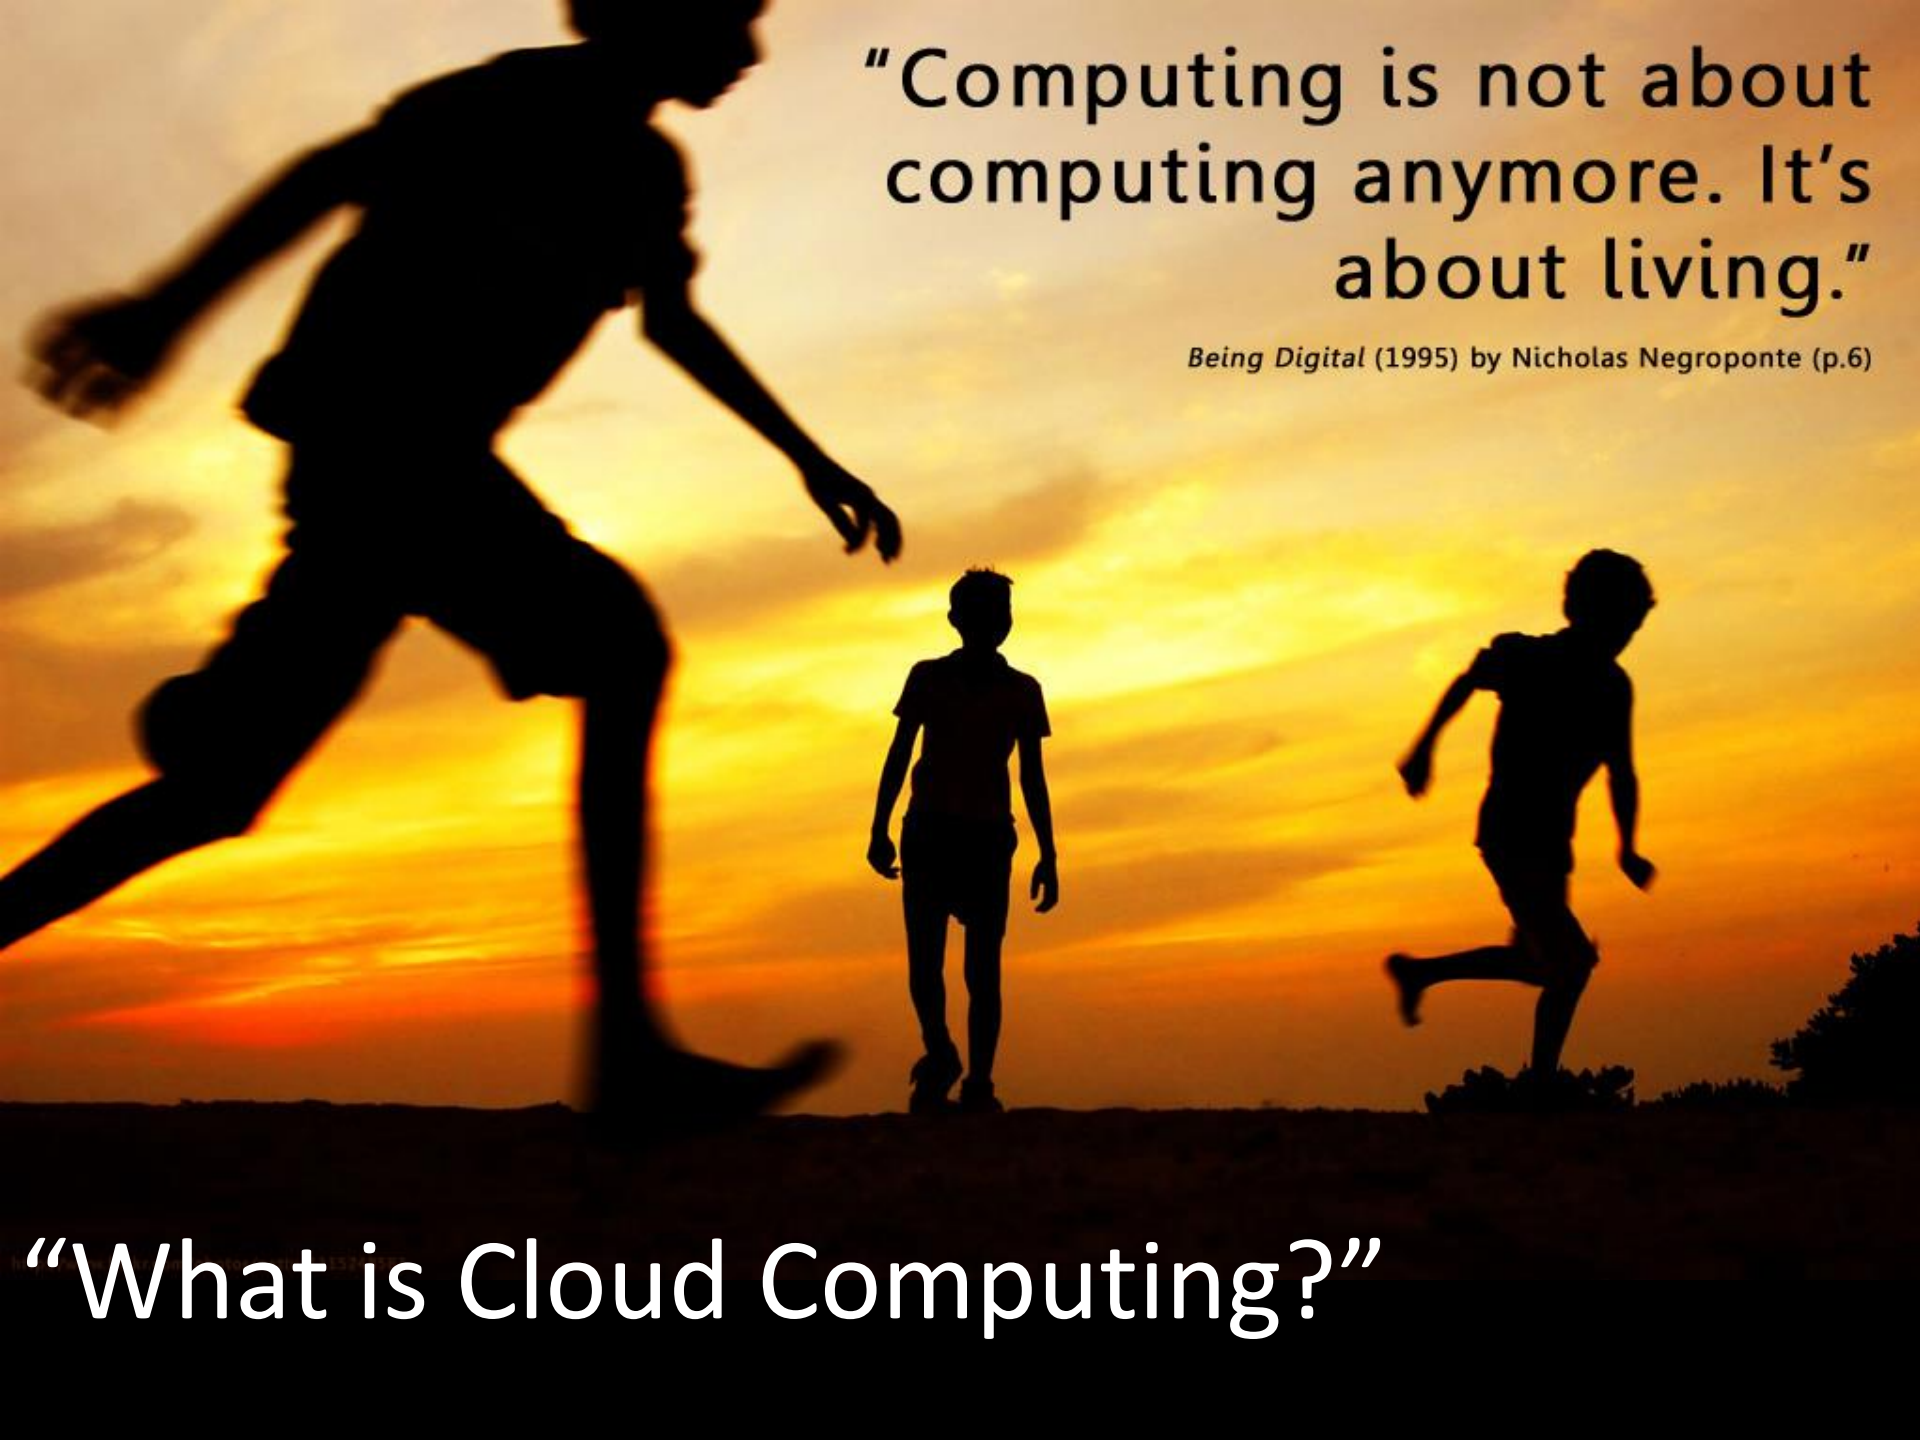
The Hype



Cluster Computing
 Cloud Computing
 Grid Computing



Google news

The background of the slide features a vibrant sunset with a gradient of orange, yellow, and red. In the foreground, the silhouettes of three people are captured in motion. On the left, a large silhouette of a person is running towards the right. In the center, a smaller silhouette of a person stands still, facing forward. On the right, another silhouette of a person is running towards the left. The overall scene conveys a sense of movement and transition.

“Computing is not about
computing anymore. It’s
about living.”

Being Digital (1995) by Nicholas Negroponte (p.6)

“What is Cloud Computing?”



The cloud is one of the largest shifts in IT over the past decade, fueling innovation across the globe.

Today, it powers many businesses and mission-critical applications, allowing organizations to take advantage of its inherent flexibility.






It's Time to Move to the Cloud



\$200B

Gartner estimates cloud-related spend to exceed by 2016.




90%

of new spending in the next six years will be cloud-based.



87%

of businesses are using public cloud.



Common implies multi-tenancy, not single or isolated tenancy

Location-independent


Online

Utility implies pay-for-use pricing

Demand implies ~infinite, ~immediate, ~invisible scalability



Cloud Computing Infrastructure Models

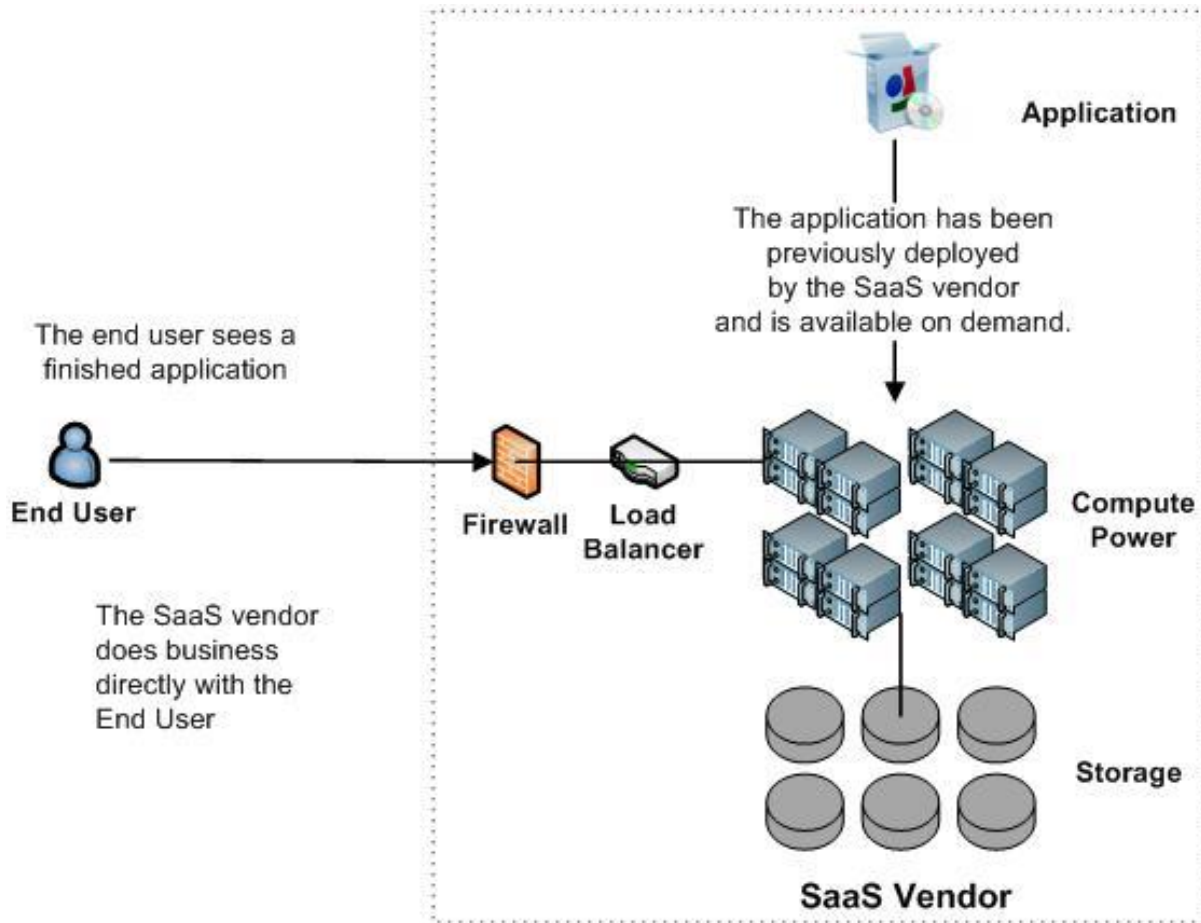




Architectural Layers of Cloud Computing

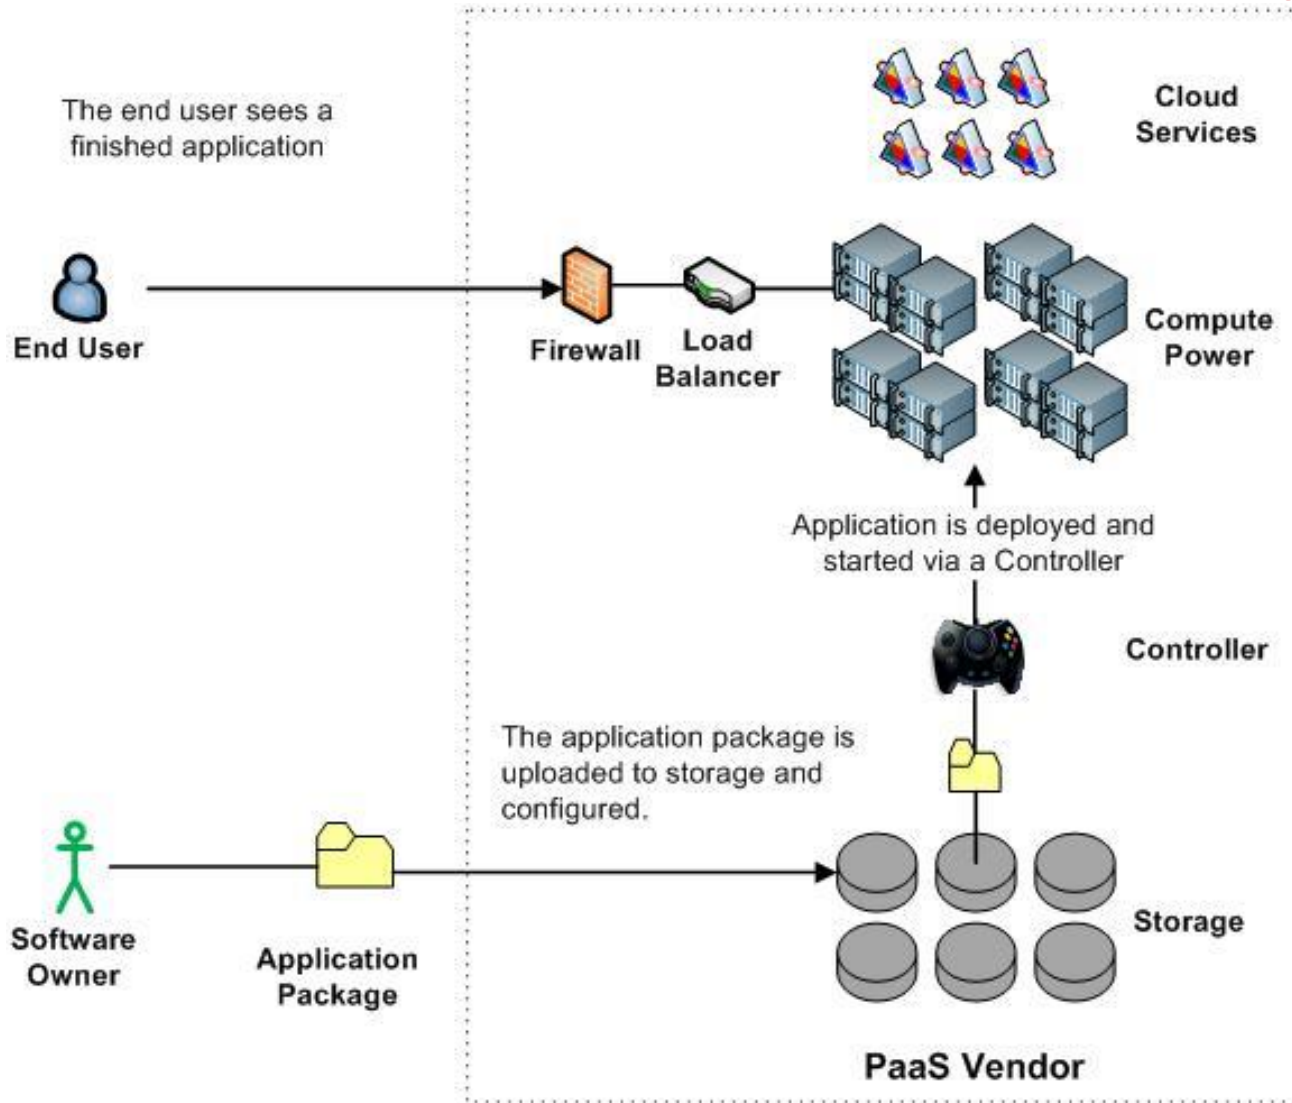


Software as a Service (SaaS)





Platform as a Service (PaaS)






Windows Azure





Infrastructure as a Service (IaaS)

The end user sees a finished application






Software as a Service (SaaS) Architectures

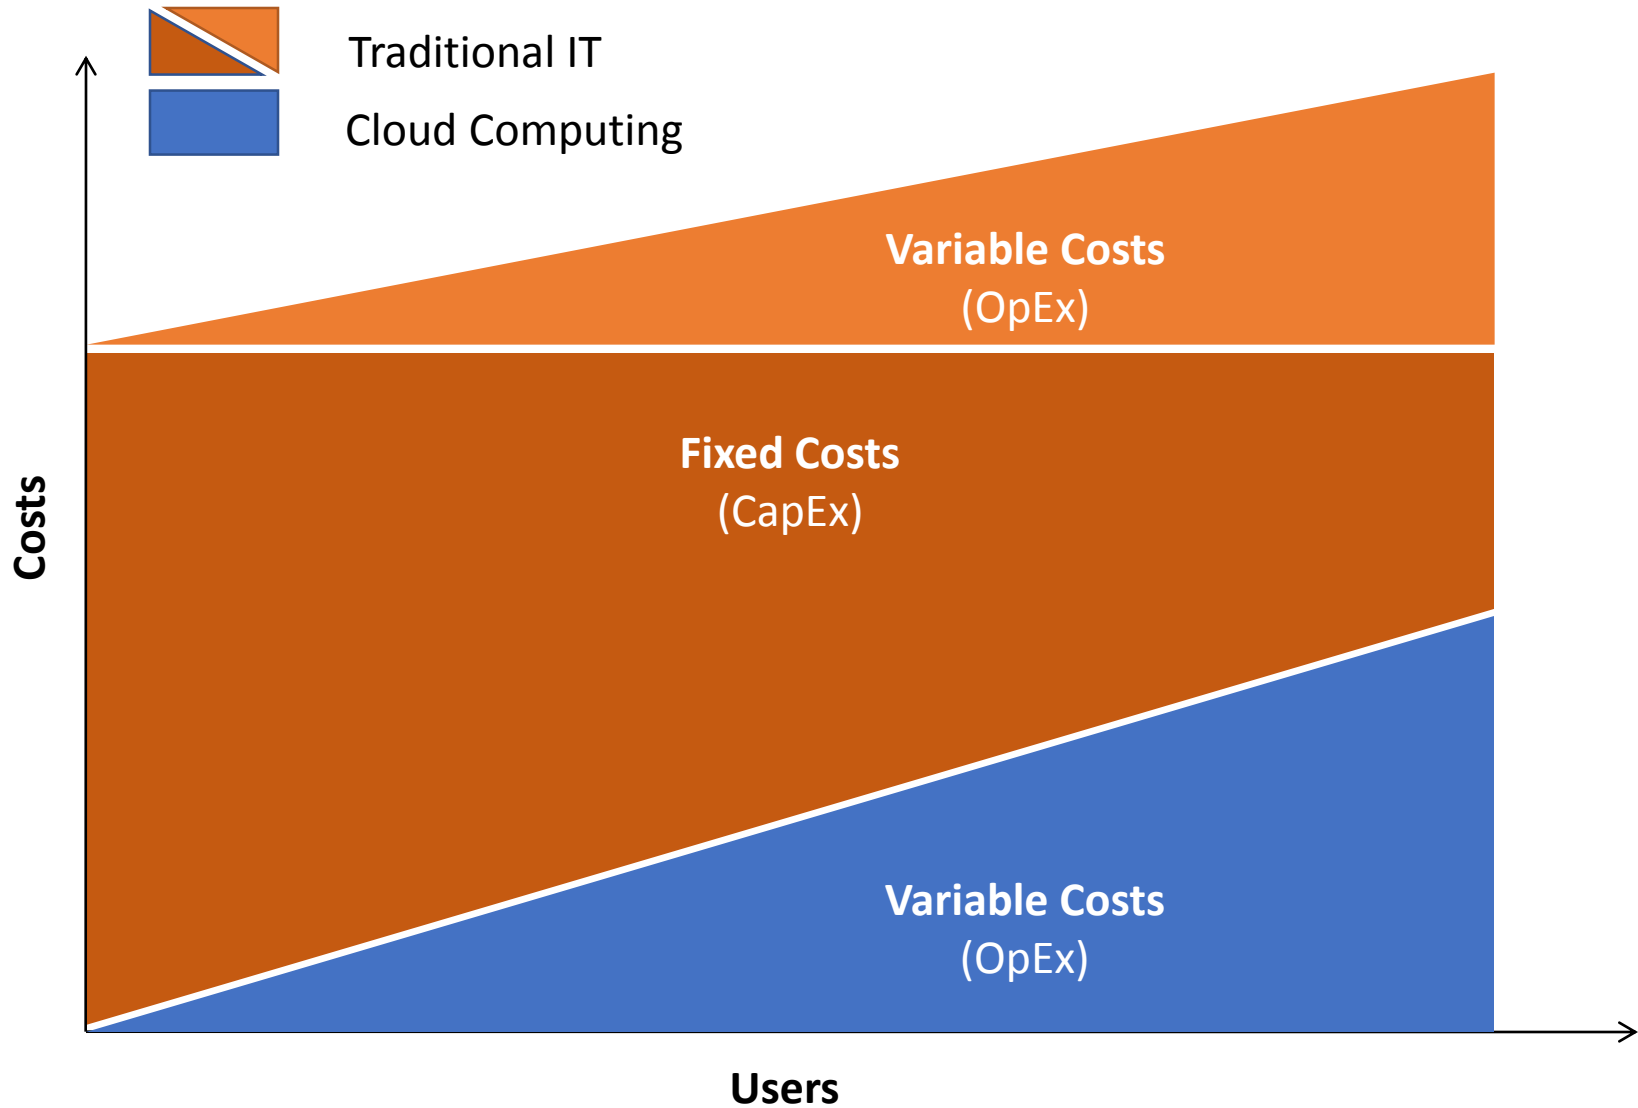


Platform as a Service (PaaS) Architectures



Infrastructure as a Service (IaaS) Architectures

Cloud Computing Economics



Pros and Cons





Cloud Computing Security



MOBILE

AUTHENTICATION

AUTHORIZATION

WEB APPLICATION

SECURITY

NETWORK

IDENTITY MANAGEMENT

DATA LEAK

COMPLIANCE

VIRUS

REMOTE ACCESS

MALWARE


ACCESS

END POINT



Cloud Security Statistics

What is Your No.1 Issue Slowing Adoption of Public Cloud Computing?





Problems Associated with Cloud Computing

- Most security problems stem from:
 - Loss of control
 - Lack of trust (mechanisms)
 - Multi-tenancy
- These problems exist mainly in 3rd party management models
 - Self-managed clouds still have security issues, but not related to above

Loss of Control in the Cloud

- Consumer's loss of control
 - Data, applications, resources are located with provider
 - User identity management is handled by the cloud
 - User access control rules, security policies and enforcement are managed by the cloud provider
 - Consumer relies on provider to ensure
 - Data security and privacy
 - Resource availability
 - Monitoring and repairing of services/resources



Cloud Computing Privacy Issues

- Location and qualification of roles (controllers, processors - chain of (sub)providers and...affected individuals)
- Location and transborder flows (requirements)
- (Un) Certainty and (Mis) trust about the use of personal data (unlawful secondary use/ disclosure to LEAs ?)
- Concerns about security (security measures/ data breaches)
- Transparency



Lack of Trust in the Cloud

- A brief deviation from the talk
 - (But still related)
 - Trusting a third party requires taking risks
- Defining trust and risk
 - Opposite sides of the same coin (J. Camp)
 - People only trust when it pays (Economist's view)
 - Need for trust arises only in risky situations
- Defunct third party management schemes
 - Hard to balance trust and risk
 - e.g. Key Escrow (Clipper chip)
 - Is the cloud headed toward the same path?



Multi-tenancy Issues in the Cloud

- Conflict between tenants' opposing goals
 - Tenants share a pool of resources and have opposing goals
- How does multi-tenancy deal with conflict of interest?
 - Can tenants get along together and 'play nicely' ?
 - If they can't, can we isolate them?
- How to provide separation between tenants?

Security Issues in the Cloud

- In theory, minimizing any of the issues would help:
 - Loss of Control
 - Take back control
 - Data and apps may still need to be on the cloud
 - But can they be managed in some way by the consumer?
 - Lack of trust
 - Increase trust (mechanisms)
 - Technology
 - Policy, regulation
 - Contracts (incentives): topic of a future talk
 - Multi-tenancy
 - Private cloud
 - Takes away the reasons to use a cloud in the first place
 - VPC: its still not a separate system
 - Strong separation

The Notorious Nine Threats

1. Data Breaches
2. Data Loss
3. Account Hijacking
4. Insecure APIs
5. Denial of Service
6. Malicious Insiders
7. Abuse of Cloud Services
8. Insufficient Due Diligence
9. Shared Technology Issues
10. BYOC

Open Source Advantages

- Leverage the work of a growing community of developers
- Works across multiple hardware infrastructure
- Possible to deploy at service providers and on-premise
- Customized to fit individual needs or to add additional services

What is OpenStack

OpenStack is an open source cloud IaaS platform which provides compute, storage, and networking resources with service components.

OpenStack is open source software for creating private and public clouds, built and disseminated by a large and democratic community of developers, in collaboration with users.

OpenStack is managed by the OpenStack Foundation, a non-profit that oversees both development and community-building around the project.

Introduction

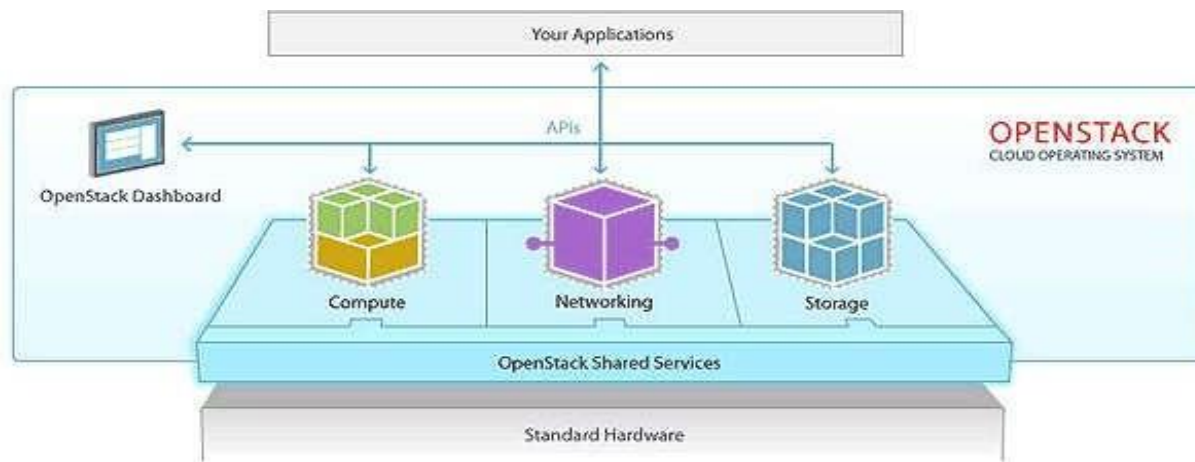
OpenStack lets users deploy virtual machines and other instances that handle different tasks for managing a cloud environment on the fly.

It makes horizontal scaling easy, which means that tasks that benefit from running concurrently can easily serve more or fewer users on the fly by just spinning up more instances.

For example, a mobile application that needs to communicate with a remote server might be able to divide the work of communicating with each user across many different instances, all communicating with one another but scaling quickly and easily as the application gains more users.



OpenStack is open source software, which means that anyone who chooses to can access the source code, make any changes or modifications they need, and freely share these changes back out to the community at large. It also means that OpenStack has the benefit of thousands of developers all over the world working in tandem to develop the strongest, most robust, and most secure product that they can.



Major Components of OpenStack

Nova (Compute)

Swift (Object Storage)

Cinder (Block Storage)


Glance (Image Service)

Neutron (Networking)

Horizon (Dashboard)

Ceilometer (Telemetry)

Heat (Orchestration)





Nova (Compute)

Nova compute or the king service provides a platform on which we are going to run our guest machines; It's the virtual machine provisioning and management module that defines drivers that interact with underlying virtualization. It provides a Control plane for an underlying hypervisors. Each hypervisor requires a separate Nova Instance. Nova supports almost all hypervisors known to man.

Glance (Image Service)

In simple words glance is the Image Registry, it stores and Manage our guest (VM) images, Disk Images, snap shots etc. It also contains prebuilt VM templates so that you can try it on the fly. Instances are booted from our glance image registry. User can create custom images and upload them to Glance later reuse. A feature of Glance is to store images remotely so to save local disk space.



Swift (Object Storage)

Swift Offers cloud storage software, Look at it as Dropbox or Google drive, as they are not attached to servers, they are individual addressable objects. It's built for scale and optimized for durability, availability, and concurrency across the entire data set. Swift is ideal for storing unstructured data that can grow without bound. Swift provides redundancy checksum for files, Files are stored as segments and a manifest file tracks them.

Ceilometer (Telemetry)

This module is responsible for metering Information. It can be used generate bills and based on the statistics of usage. Its API can be used with external billing systems.

Administrators can create certain alarms that are triggered based on performance statistics



Neutron (Networking)

Neutron, the networking component which was formally called Quantum. This component provides the software defined Networking Stack for Openstack. It Provides networking as a service. It gives the cloud tenants an API to build rich networking topologies, and configure advanced network policies in the cloud. It enables innovation plugins (open and closed source) that introduce advanced network capabilities which let anyone build advanced network services (open and closed source) that plug into Openstack tenant networks means that you can create advance managed switches and routers. You can even create an intelligent switch from a PC (Yes you can use it as a standalone component) and use it to replace your managed switch or at least make it act as a backup Switch.

Horizon (Dashboard)

Horizon is the Dashboard to Openstack, your eyes and ears. It provides a web based user interface to OpenStack services including Nova, Swift, Keystone etc.



Cinder (Block Storage)

Cinder is also one of the storage modules of Openstack; Think of it as an external hard drive or like a USB device. It has the performance characteristics of a Hard drive but much slower than Swift and has low latency. Block Volumes are created in Swift and attached to running Volumes for which you want to attach an extra partition or for copying data to it. It survives the termination of an Instance. It is used to keep persistence storage. Cinder Images are mostly stored on our shared storage environment for readily availability. These Images can be clones and snapshot which can be turned in to bootable images. Its similar to Amazon Elastic block Storage.

Heat (Orchestration)

It creates a human and machine-accessible service for managing the entire lifecycle of infrastructure and applications within Openstack clouds. It contains human readable templates with simple instruction that is read by the Heat Engine. Heat along with ceilometer (explained below) can create an auto-scaling the cloud.

Why to use OpenStack?

- Enables rapid innovation
- Cuts down time-to-market
- Boosts scalability and resource utilization
- Eases regulatory compliance
- Devoid of vendor lock-in

Who is using OpenStack?

List of top companies using OpenStack.

Company	Website	Country	Revenue	Company Size
Red Hat Inc	redhat.com	United States	>1000M	5000-10000
ViaSat, Inc.	viasat.com	United States	>1000M	1000-5000
World Wide Technology	wwt.com	United States	>1000M	1000-5000
Mirantis, Inc.	mirantis.com	United States	50M-100M	500-1000
Canonical Ltd.	canonical.com	United Kingdom	50M-100M	500-1000

Companies using OpenStack, by industry

Industry	Number of companies
Computer Software	1957
Information Technology and Services	1103
Staffing and Recruiting	449
Higher Education	286
Telecommunications	256
Computer Hardware	194
Internet	171
Hospital & Health Care	125
Financial Services	107
Retail	94

OpenStack Community – 60+ companies



Review Security Strategies

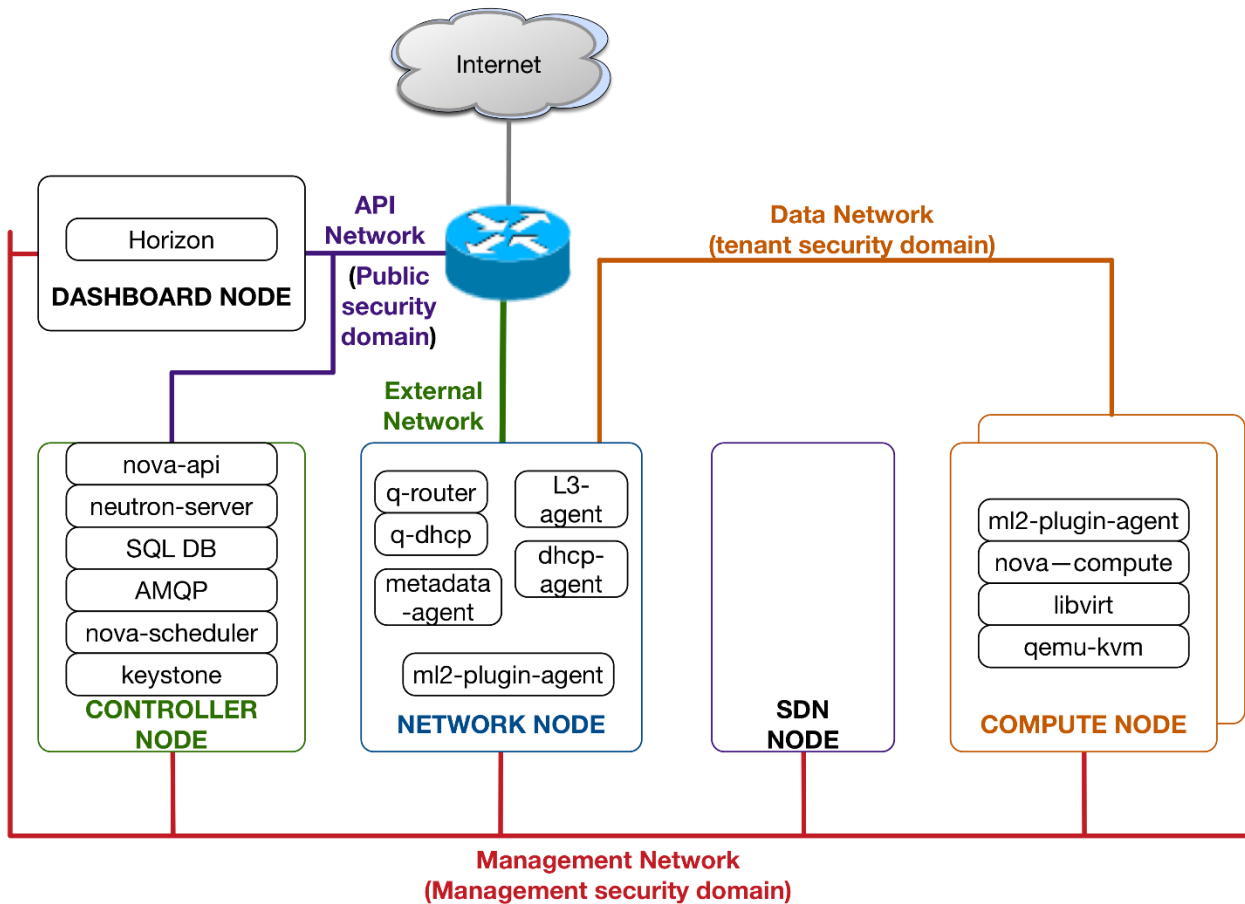
- “Defense in depth” as the Primary Strategy
 - Use a series of defensive mechanisms
 - Protect each and every service component, communication channel and API access
 - Perimeter protection is one of the best defenses but not the only one
- Role Based Access Control (RBAC) & Domains
 - Restrict access to project resources based on Roles
 - Leverage Keystone “domains”
- Harden Neutron Service Components
 - Patch any reported neutron security vulnerabilities
 - e.g. https://www.cvedetails.com/vulnerability-list/vendor_id-11727/Openstack.html
 - Proper service configuration (Service owner, file permissions etc.)



Review Security Strategies (Continued..)

- Evaluate the security impact of using a tenant network type
 - Flat/VLAN/VXLAN/GRE etc.
- Limit Resources available to Projects
 - Self-Service network users have the ability to create, update, and destroy network resources
 - Set Quotas to limit network resources
- Isolate Network access
 - Security domains to control interaction to networking services
 - Network segmentation for various types of traffic (e.g. Management, API, External, Data)
 - Use Neutron Security Groups & Anti-Spoofing filters to protect project instances

Network Security Domains



Networking Security Checklist

- Are all interactions with the networking service isolated into security domains?
- Does the ML2 mechanism driver mitigate ARP spoofing?
- Considered the pros and cons of supporting various tenant network types?
- Have you hardened all neutron service components?
- Are you using neutron security-groups and enabled port-security?
- Are all communications using SSL encryption?
- Has RBAC been implemented using the concept of least privilege?
- Have you investigated the maturity and security features of the various pluggable neutron components you are using?
- Are you using quotas to limit project resources?



USM

UNIVERSITI
SAINS
MALAYSIA



Thank You

Dr. Selvakumar Manickam

selva@usm.my