# Cloud Computing

Spring 2016

# Reading

- Above the Clouds: A Berkeley View of Cloud Computing. Technical Report
- EMC Lecture Notes — Cloud Infrastructure and Services

- Originally took 4 years
- Required 1 hour per frame
- 3D version in 2009: less than 1/24th of a second per frame.



1 petabytes of required space to render the movie.

Rendering farm:

- 34 racks, 4 chassis each, 32 m/c each
- 40,000 processors, 104TB memory, 10GB networking
- 1.4 million tasks per day
- 24 hours for over a month

# Scaling of Computation

PC

Server

Cluster

- Question!

  What should I do if resources on my computer are not enough?

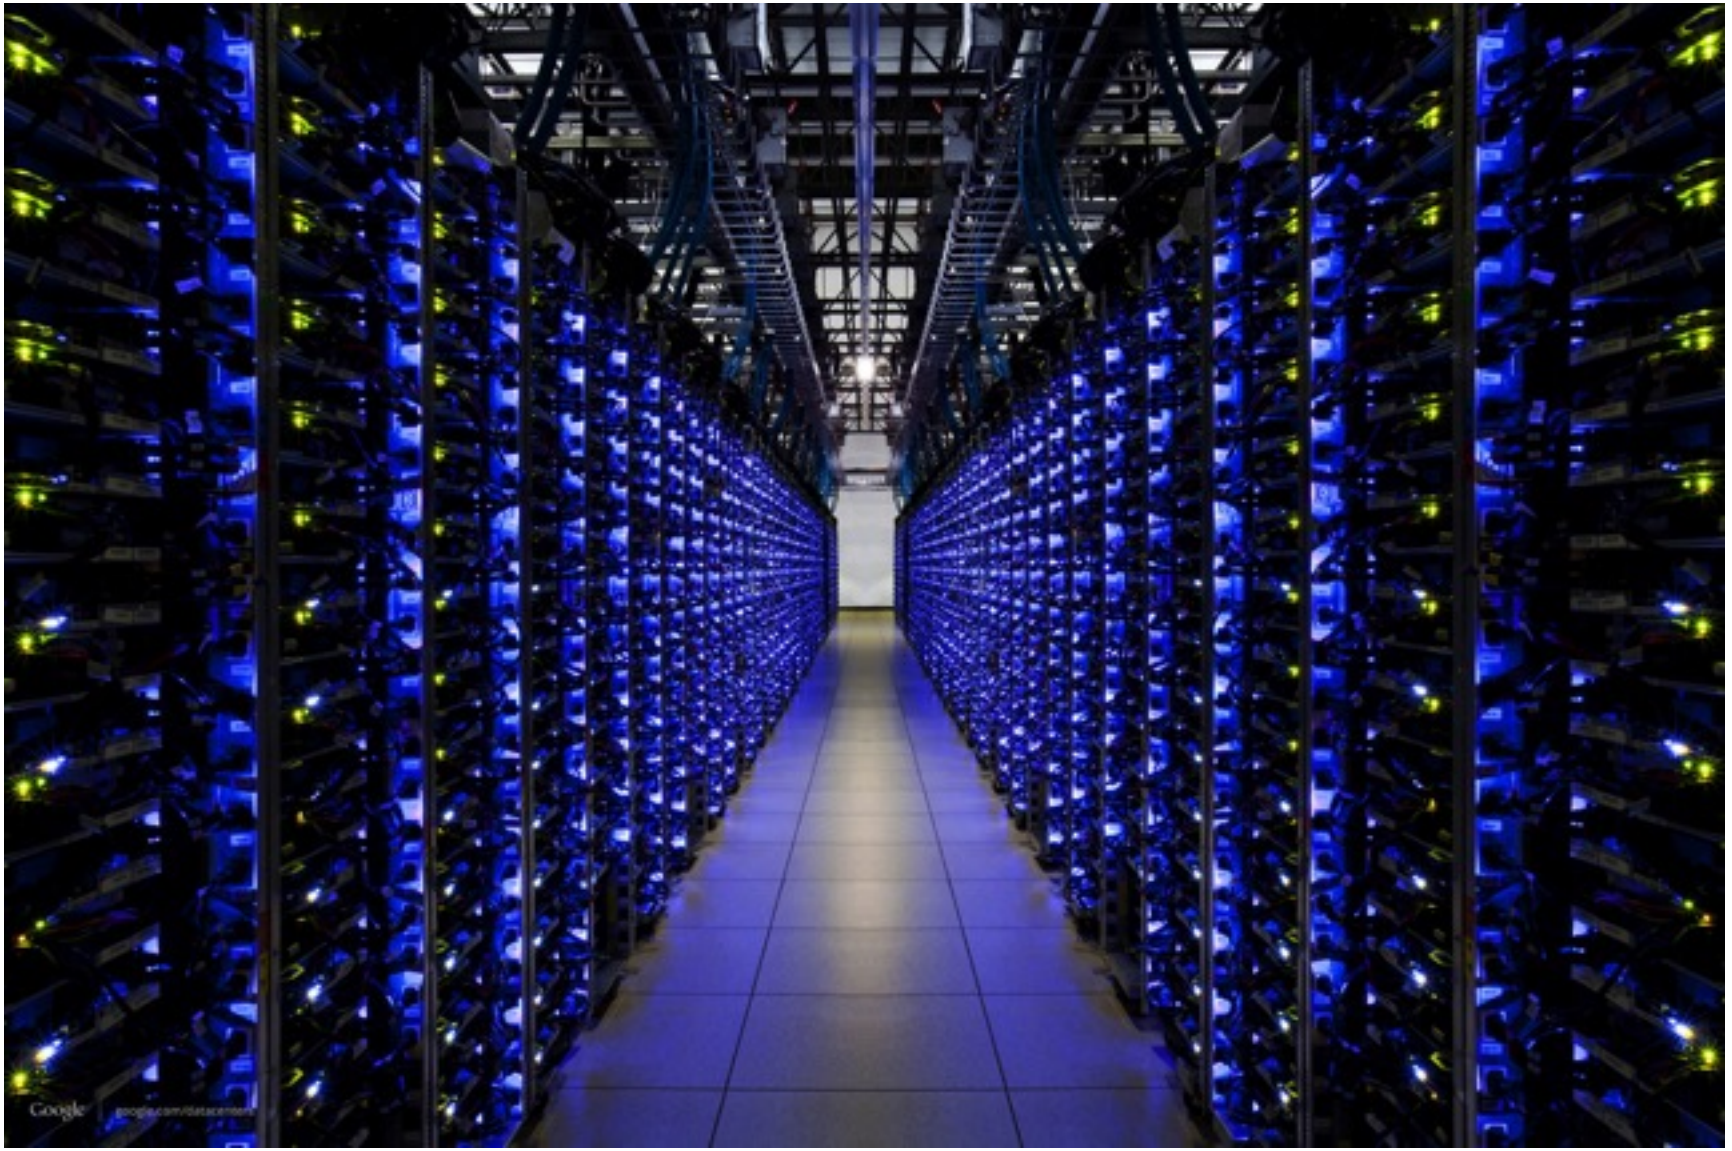


Inside Google Data Center

# Anatomy of a Cluster

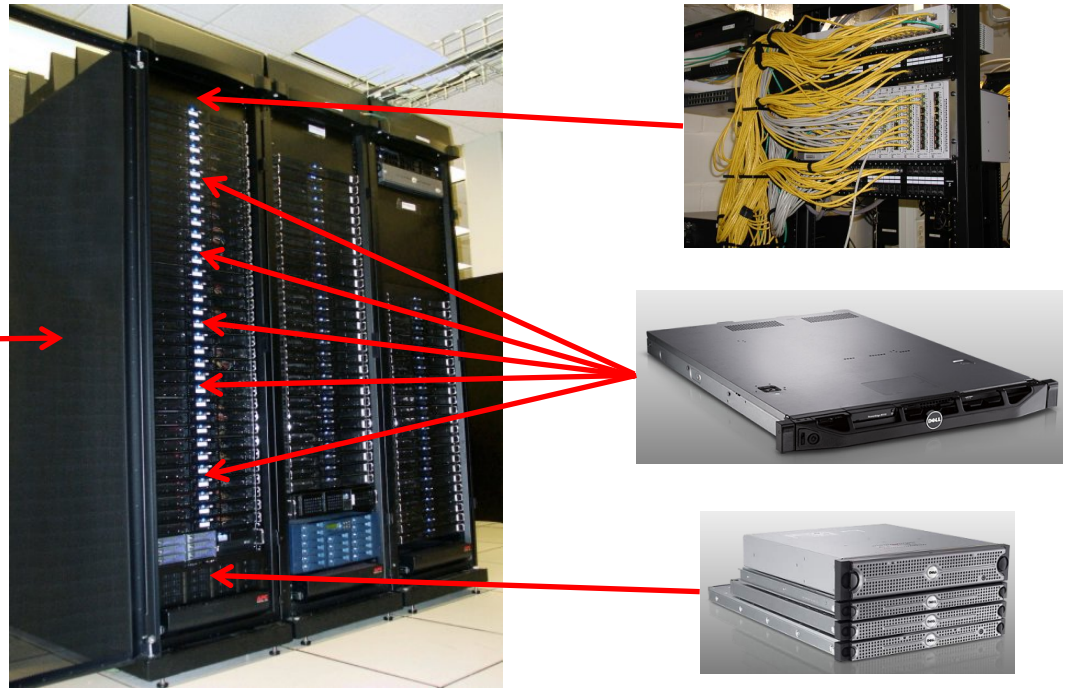

# Rack vs Blade Servers

One rack unit is 1.75 inches or 44.45 mm

Bladed server technology is commonly used to deploy compute systems in a CDC

- Consolidates power- and system-level function into a single, integrated chassis
- Enables the addition of server modules as hot-pluggable components
- Provides increased server performance and availability without increase in size, cost, or complexity

Google Data Center, Oklahoma, USA

# Geo-Distribution of Google Data Centers



- Replication
- Availability
- Fault-tolerance

# Outline

- Introduction
- Cluster Computing, Grid Computing, and Utility Computing
- Cloud Computing Characteristics
- Cloud Service Models
- Cloud Deployment Models
- Cloud Infrastructure
- Challenges

# Cluster Computing

- Definition: A *cluster computer is a set of* interconnected computers that cooperate closely to provide a single, integrated high performance computing capability (CL)
- Building a super computer using relatively simple **computers** connected using a **high speed network**
- Usually used in parallel programming
- Single master node to allocate nodes to a parallel program, maintains a queue of submitted jobs, and provides an interface for users
- Alternative architectures where no single master exists are also used
- "**homogeneous**"

# Grid Computing

- The name ‘Grid’ is used to refer to middleware that is designed to enable the sharing of resources such as files, computers, software, data and sensors on a very large scale
- “**heterogeneous**”: no assumptions are made concerning hardware, operating systems, networks, administrative domains, security policies, etc
- Resources from different organizations are brought together to allow the collaboration of a group of people or institutions
- Virtual organizations

# Which Option Would You Choose?

Option A:

Option B:

# Utility Computing

- Power generator per house vs. power plants and customers pay per their usage
- Advantages:
  - It is more economical to run a big data center than building small ones at each enterprise
  - Pay as you go
  - No need for resource provisioning
  - Easier to scale
  - High utilization

# Outline

- Introduction
- Cluster Computing, Grid Computing, and Utility Computing
- Cloud Computing Characteristics
- Cloud Service Models
- Cloud Deployment Models
- Cloud Infrastructure
- Challenges

# New in Cloud Computing

- The illusion of infinite computing resources available on demand
  - Users do not need to plan far ahead for provisioning
- No up-front commitment by Cloud users
  - Companies can start small and increase hardware resources only when there is an increase in their needs
- Pay for use of computing resources on a short-term basis as needed
  - e.g., processors by the hour and storage by the day
  - release resources as needed, rewarding conservation

"Above the Clouds: A Berkeley View of Cloud Computing", RAD lab, UC Berkeley

# Elasticity: Economics in the Cloud

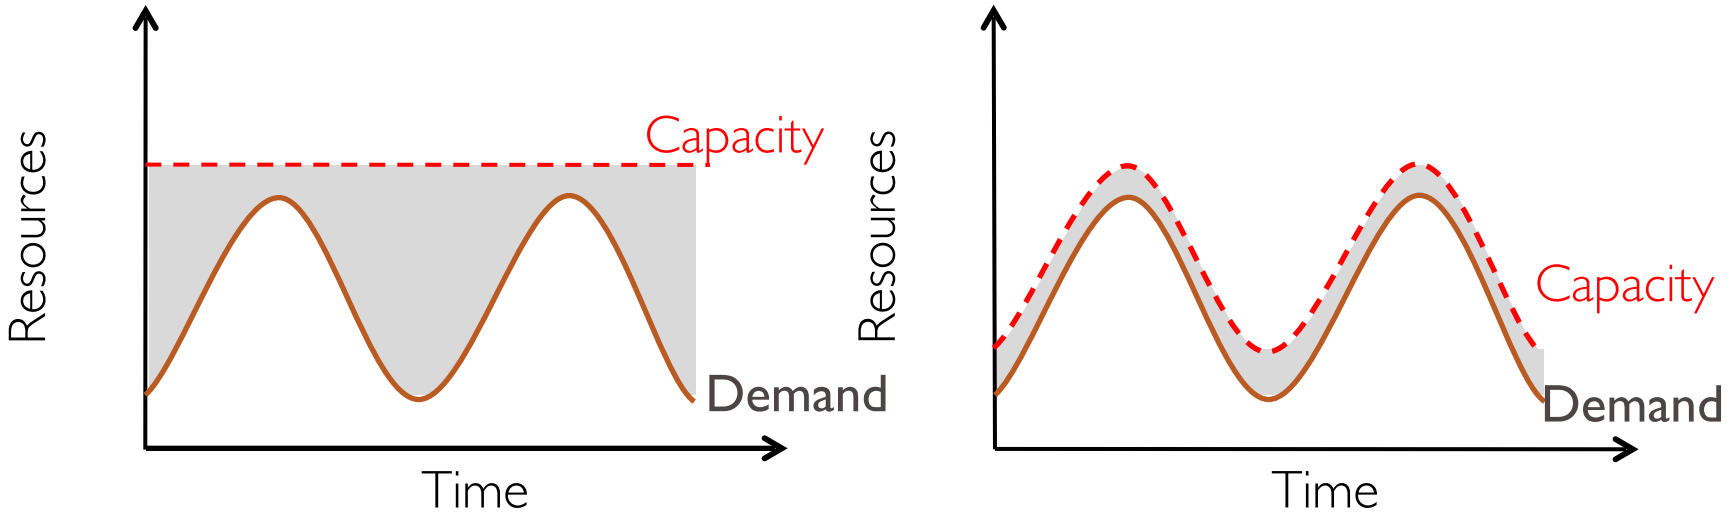


Traditional Infrastructures

Deployment in the Cloud

Unused resources

"Above the Clouds: A Berkeley View of Cloud Computing", RAD lab, UC Berkeley

# Definition?

- The interesting thing about Cloud Computing is that we've redefined Cloud Computing to include everything that we already do.... I don't understand what we would do differently in the light of Cloud Computing other than change the wording of some of our ads.

  Larry Ellison, quoted in the Wall Street Journal, September 26, 2008

- A lot of people are jumping on the [cloud] bandwagon, but I have not heard two people say the same thing about it. There are multiple definitions out there of "the cloud".

  Andy Isherwood, quoted in ZDnet News, December 11, 2008

# Characteristics of Cloud Computing

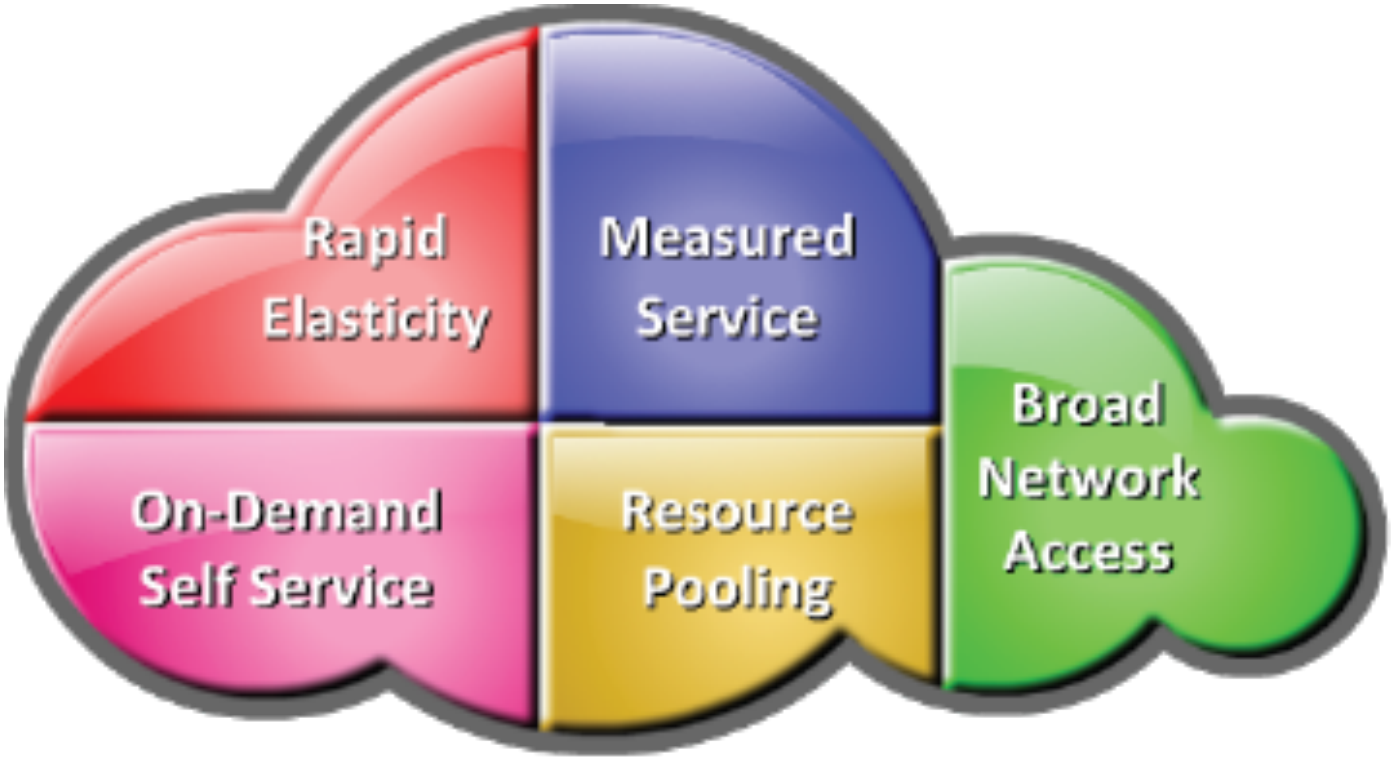


Based on the definition of Cloud Computing by NIST

Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# On-Demand Self-Service


- Enables consumers to get ***computing resources as and when required***, without any human intervention
- Facilitates consumer to leverage ***"ready to use" services*** or, enables to choose required services from the service catalog

Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# Broad Network Access


- Cloud services are accessed via the network, usually the internet
- Enables accessing the services from anywhere across the globe

Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# Resource Pooling


- Resources (compute, storage, network) are pooled to serve multiple consumers
  - Based on multi-tenant model
- No knowledge about the exact location of the resources provided
- Resources are dynamically ***assigned*** and reassigned
  - based on the consumer demand

Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# Rapid Elasticity

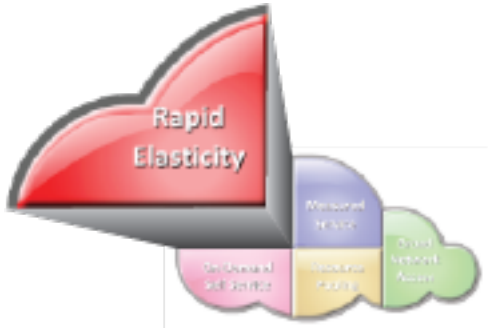


- Resources can be both scaled up and scaled down dynamically
- To the consumer, the Cloud appears to be infinite

Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# Metered Service


- Consumers are billed based on the metered usage of Cloud resources:
  - Cost incurred on a pay-per-use basis
  - Pricing/billing model is tied up with the required service levels
- Resource usage is monitored and reported

Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# Outline

- Introduction
- Cluster Computing, Grid Computing, and Utility Computing
- Cloud Computing Characteristics
- Cloud Service Models
- Cloud Deployment Models
- Cloud Infrastructure
- Challenges

# Cloud Service Models

- Cloud Service can be classified into three categories:
  - Infrastructure-as-a-Service (IaaS)
  - Platform-as-a-Service (PaaS)
  - Software-as-a-Service (SaaS)

# Infrastructure-as-a-Service

- Can hire infrastructure components such as servers, storage, and network
- Can deploy and run software, including OS and applications
- Pays for infrastructure componer usage
  - Storage capacity, CPU usage, etc
- Examples:
  - Amazon Elastic Compute Cloud (EC2)
  - Windows Azure Virtual Machines

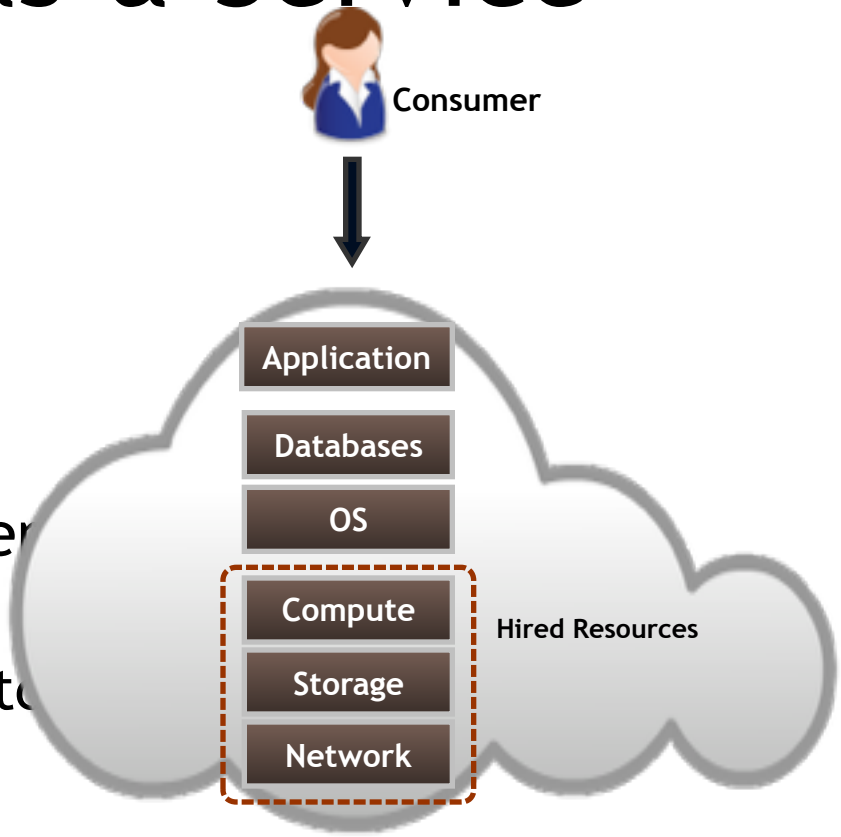


Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# Platform-as-a-Service

- Can deploy consumer-created or acquired applications on the Cloud provider's infrastructure
- Consumer has control over:
  - Deployed applications
  - Possible application hosting environment configurations
- Consumer is billed for platform software components:
  - OS, Database, Middleware
- Examples:
  - Google App Engine
  - Microsoft Azure Platform

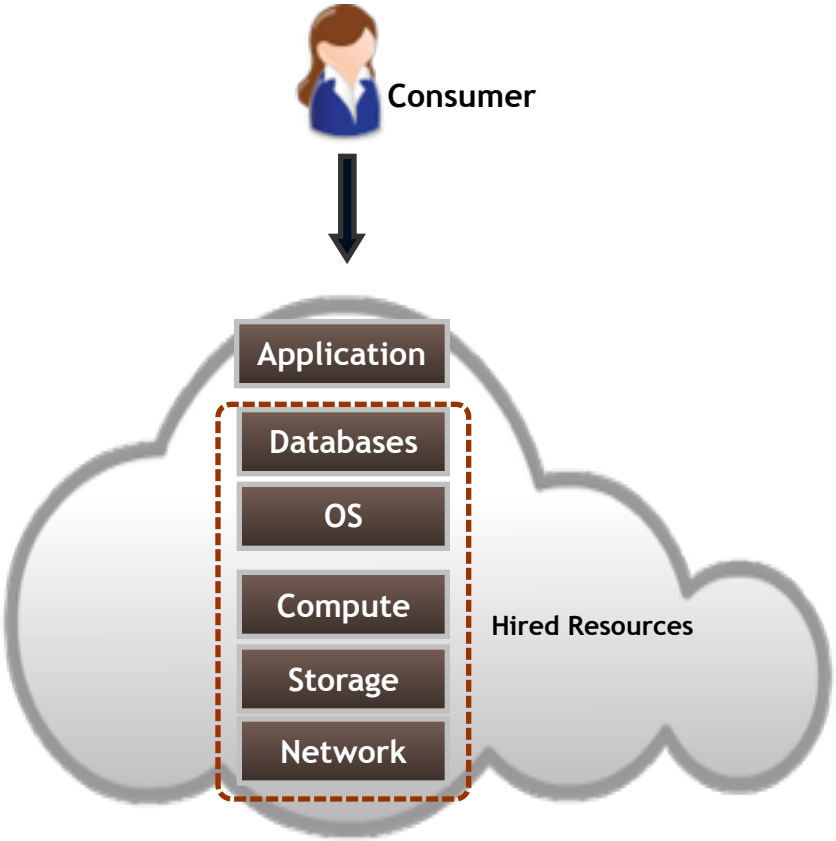


Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# Software-as-a-Service

- Use provider's applications running in a Cloud infrastructure
- Application is accessible from various client devices
- Billing is based on the application usage
- Examples:
  - EMC Mozy for online backup
  - Salesforce.com for CRM application
  - Google Apps

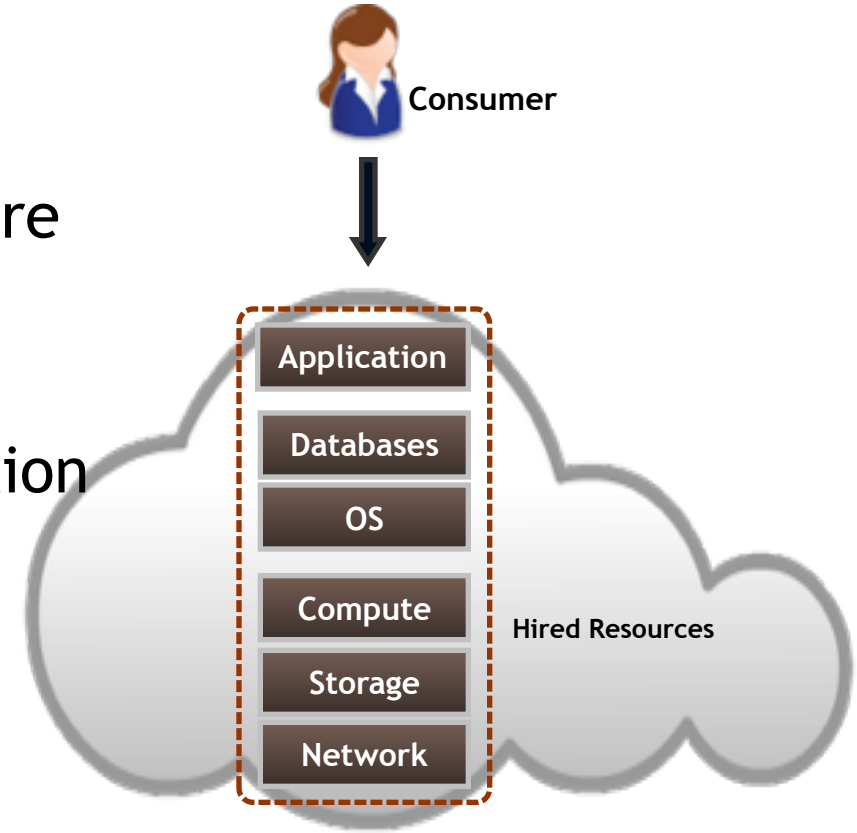


Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# Outline

- Introduction
- Cluster Computing, Grid Computing, and Utility Computing
- Cloud Computing Characteristics
- Cloud Service Models
- Cloud Deployment Models
- Cloud Infrastructure
- Challenges

# Cloud Deployment Models

- Public Cloud:
  - Resources are made available to the general public or organizations and are owned by the Cloud service provider
- Private Cloud:
  - Operated solely for one organization and is not shared with other organizations.
  - Greatest level of security and control
- Hybrid Cloud:
  - Organization consumes resources from both private and public Clouds
- Community Cloud:
  - The Cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g. all universities, government)

# Public Cloud

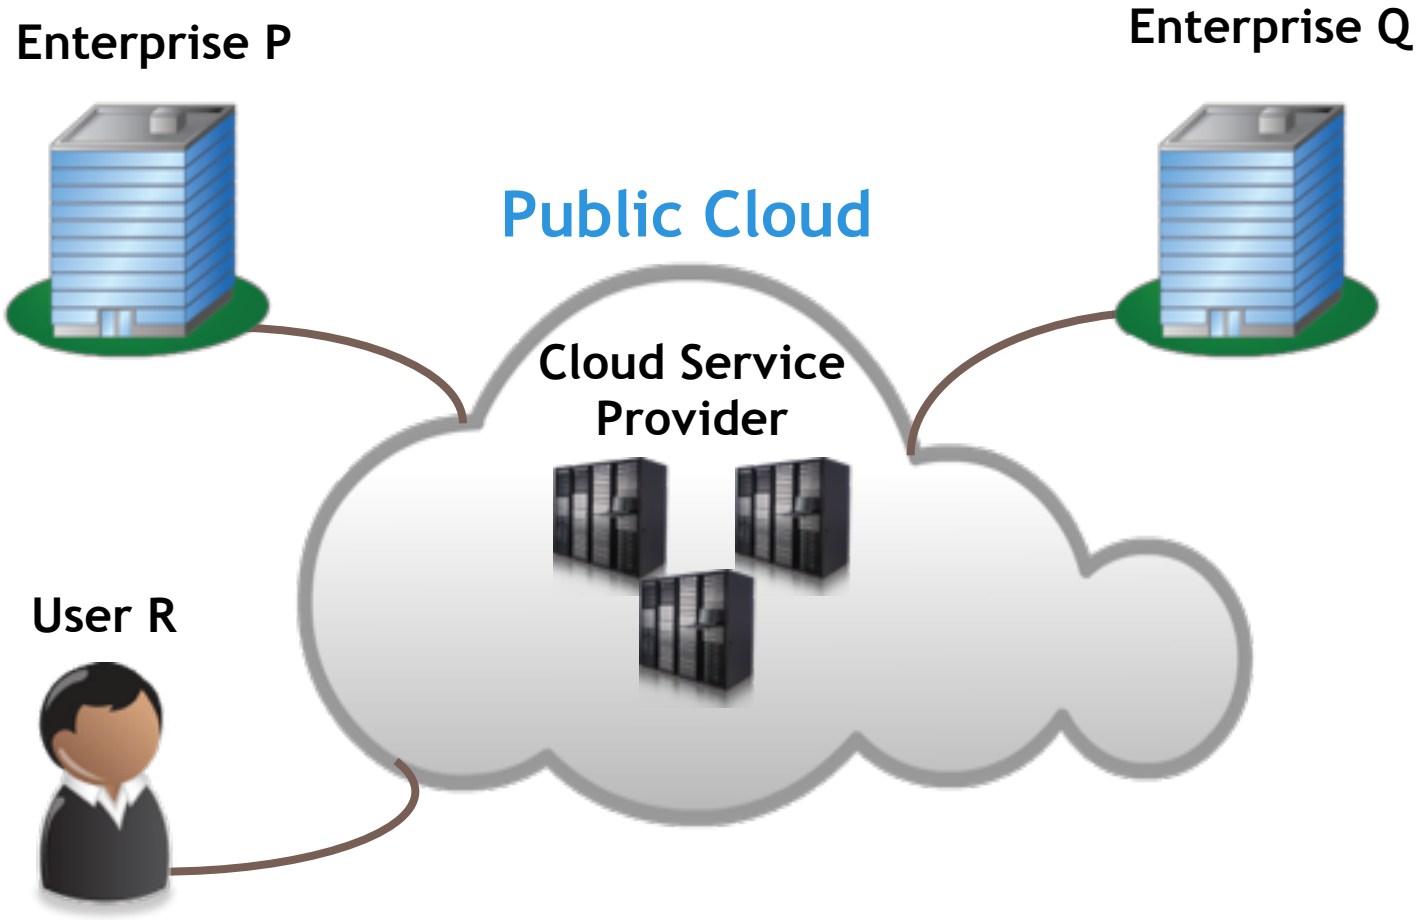


Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# Private Cloud

## On-premise Private Cloud

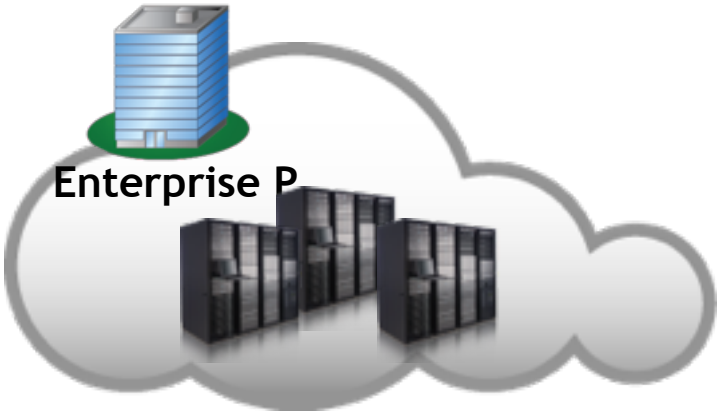


## Externally hosted Private Cloud

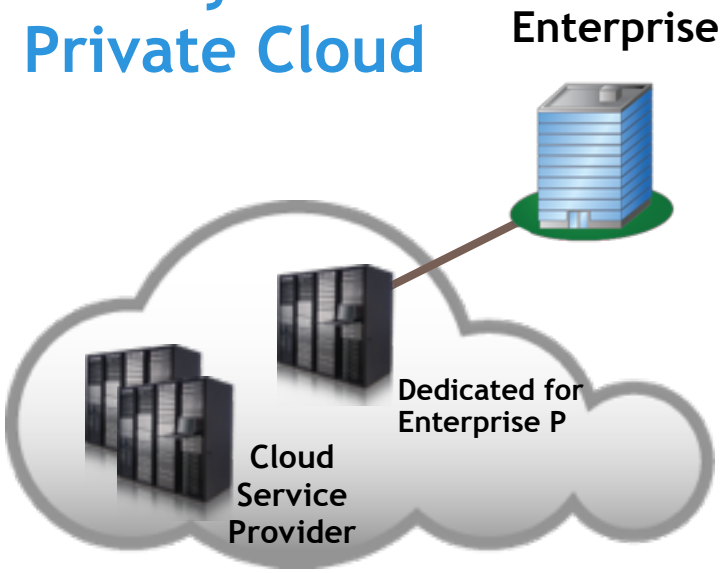


Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# Hybrid Cloud

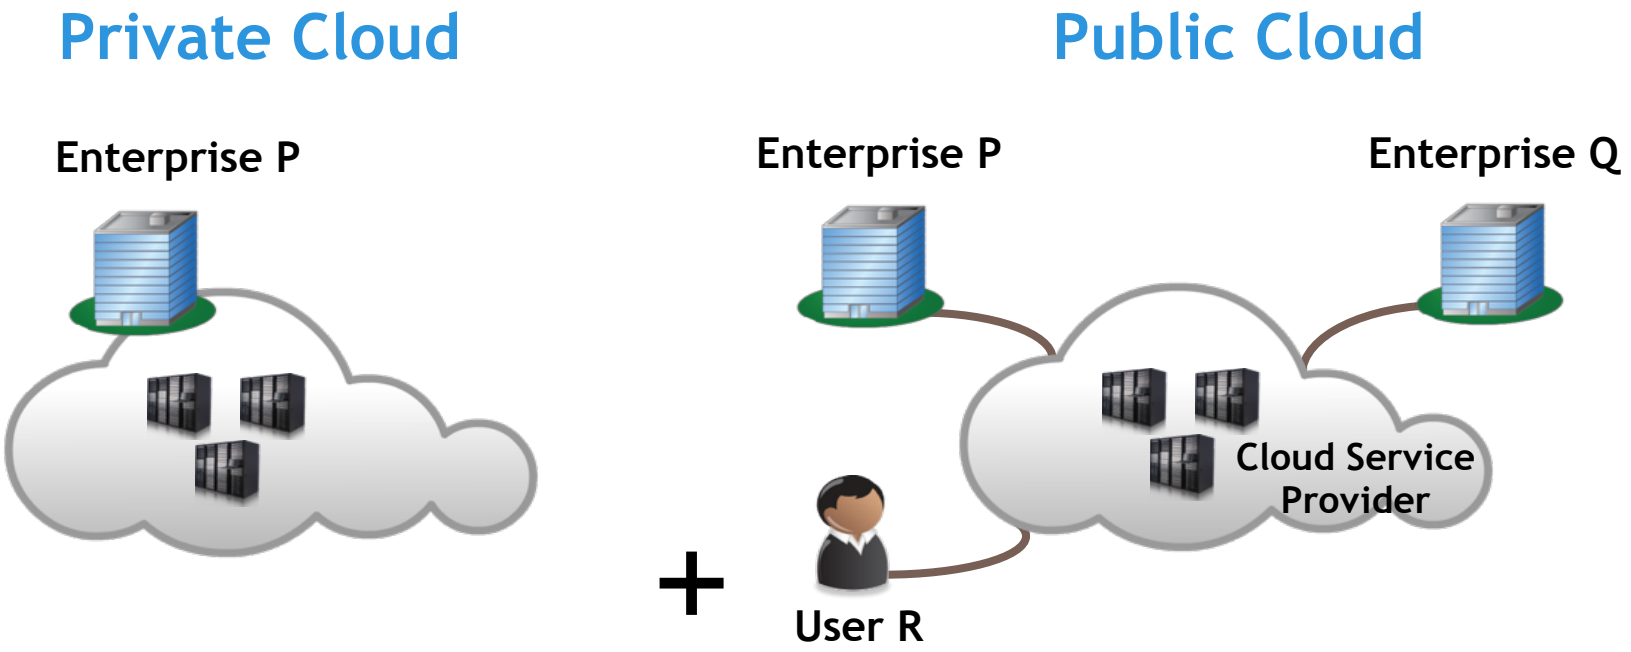


Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# Community Cloud

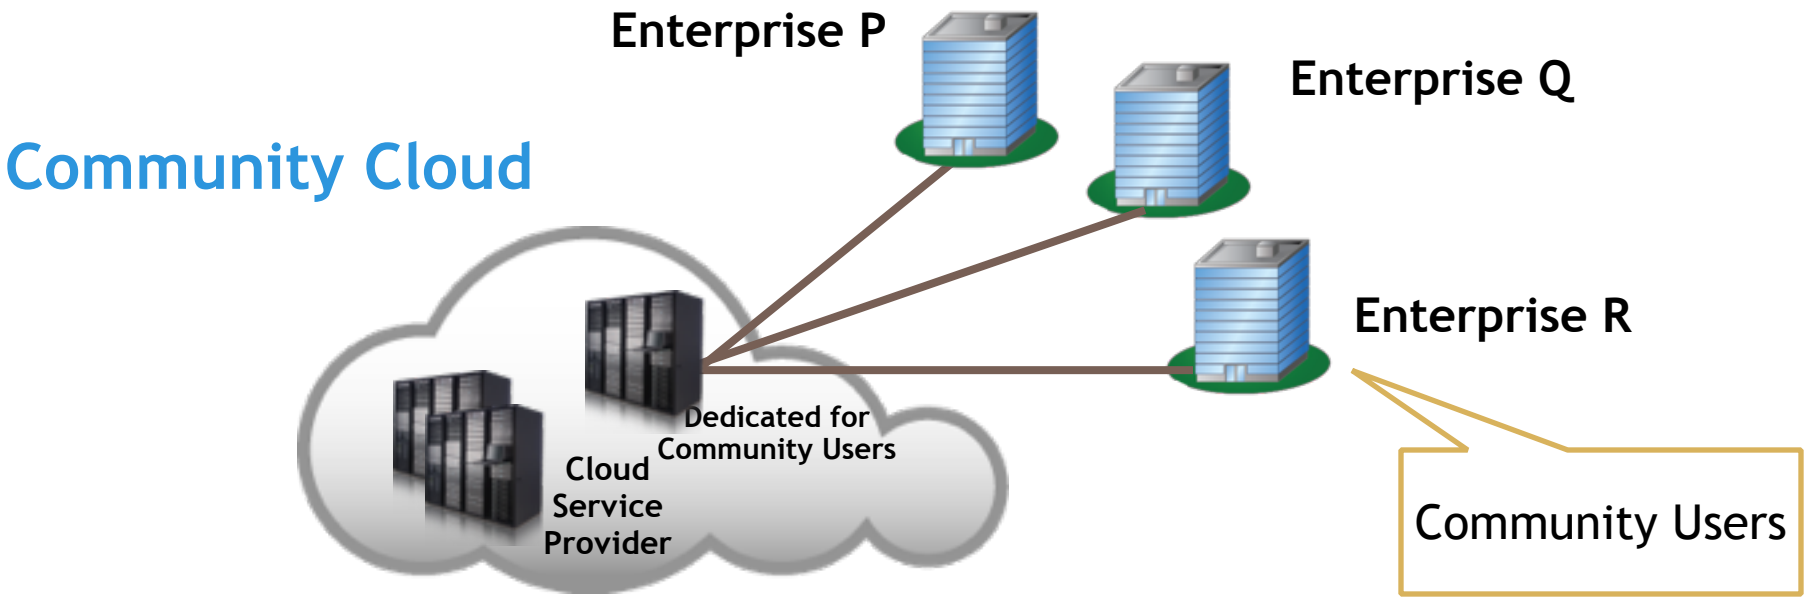


Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# Case Study: Amazon Web Service (AWS)

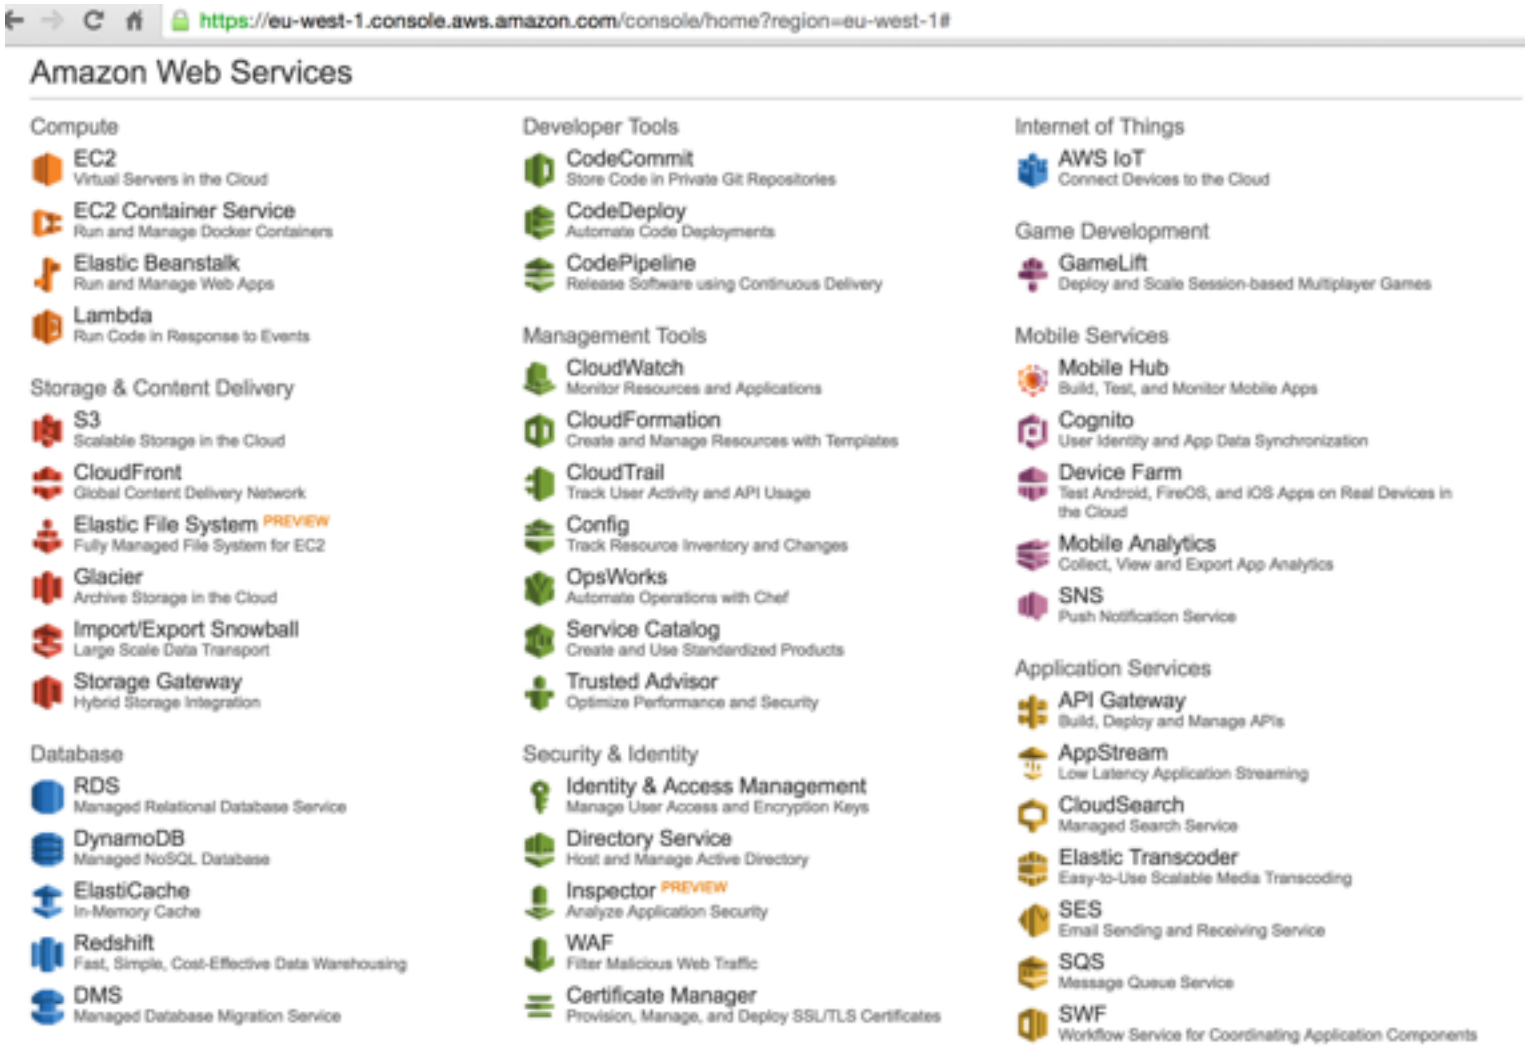

# Case Study: Amazon Web Service (AWS)

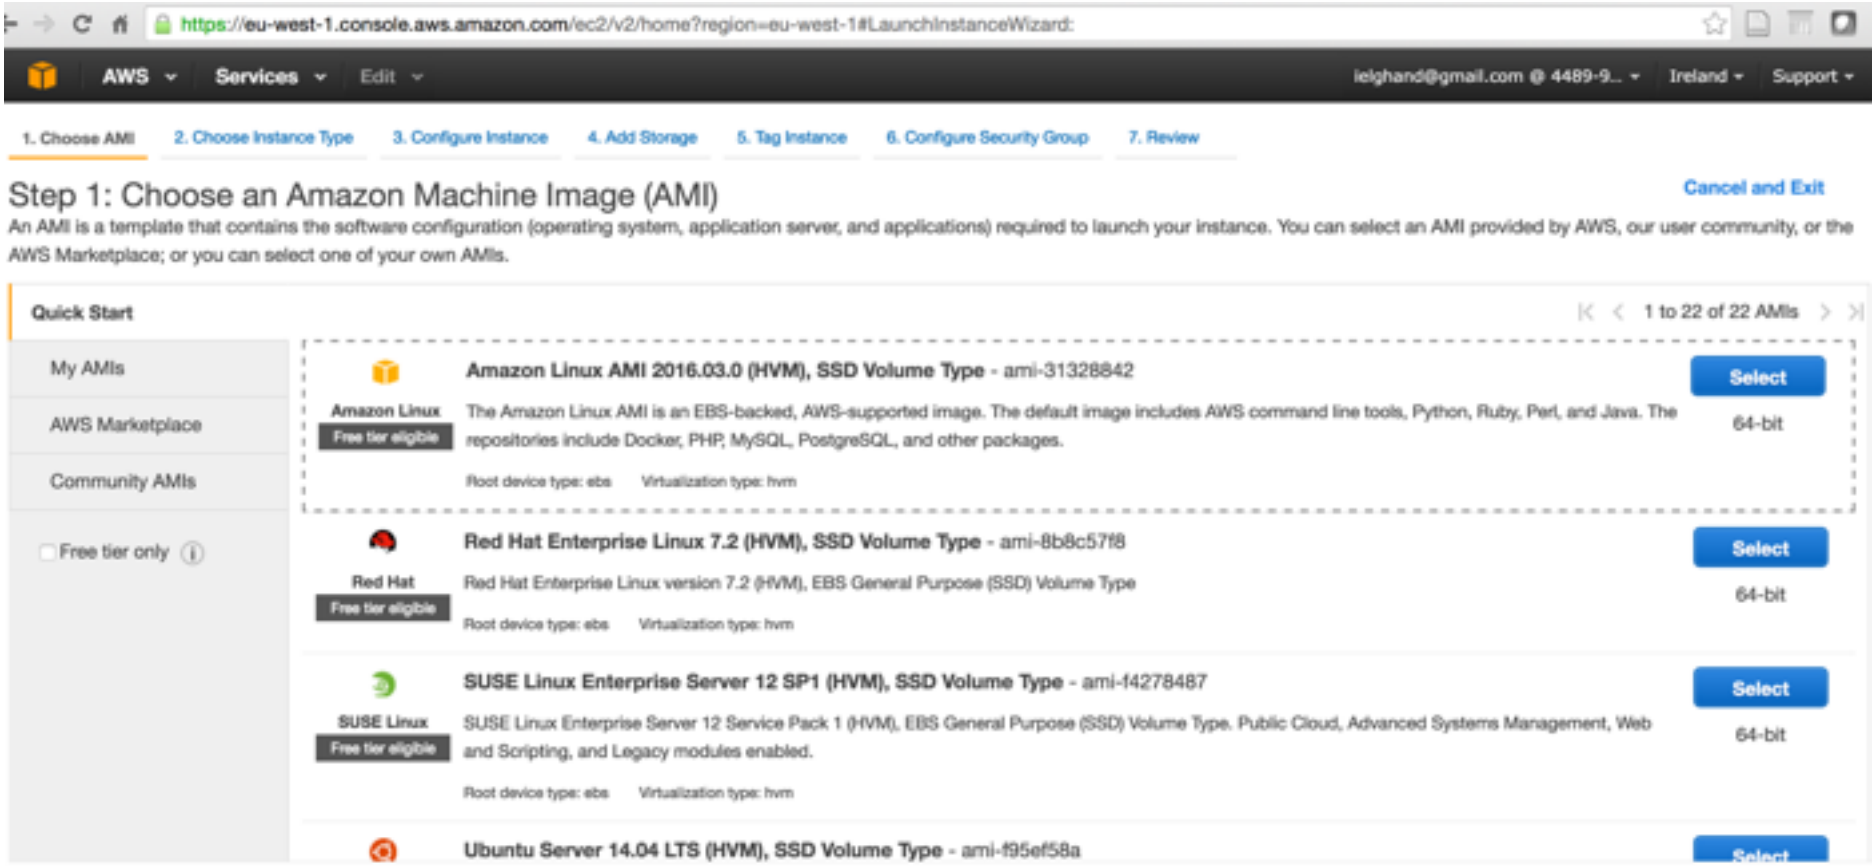

# Case Study: Amazon Web Service (AWS)

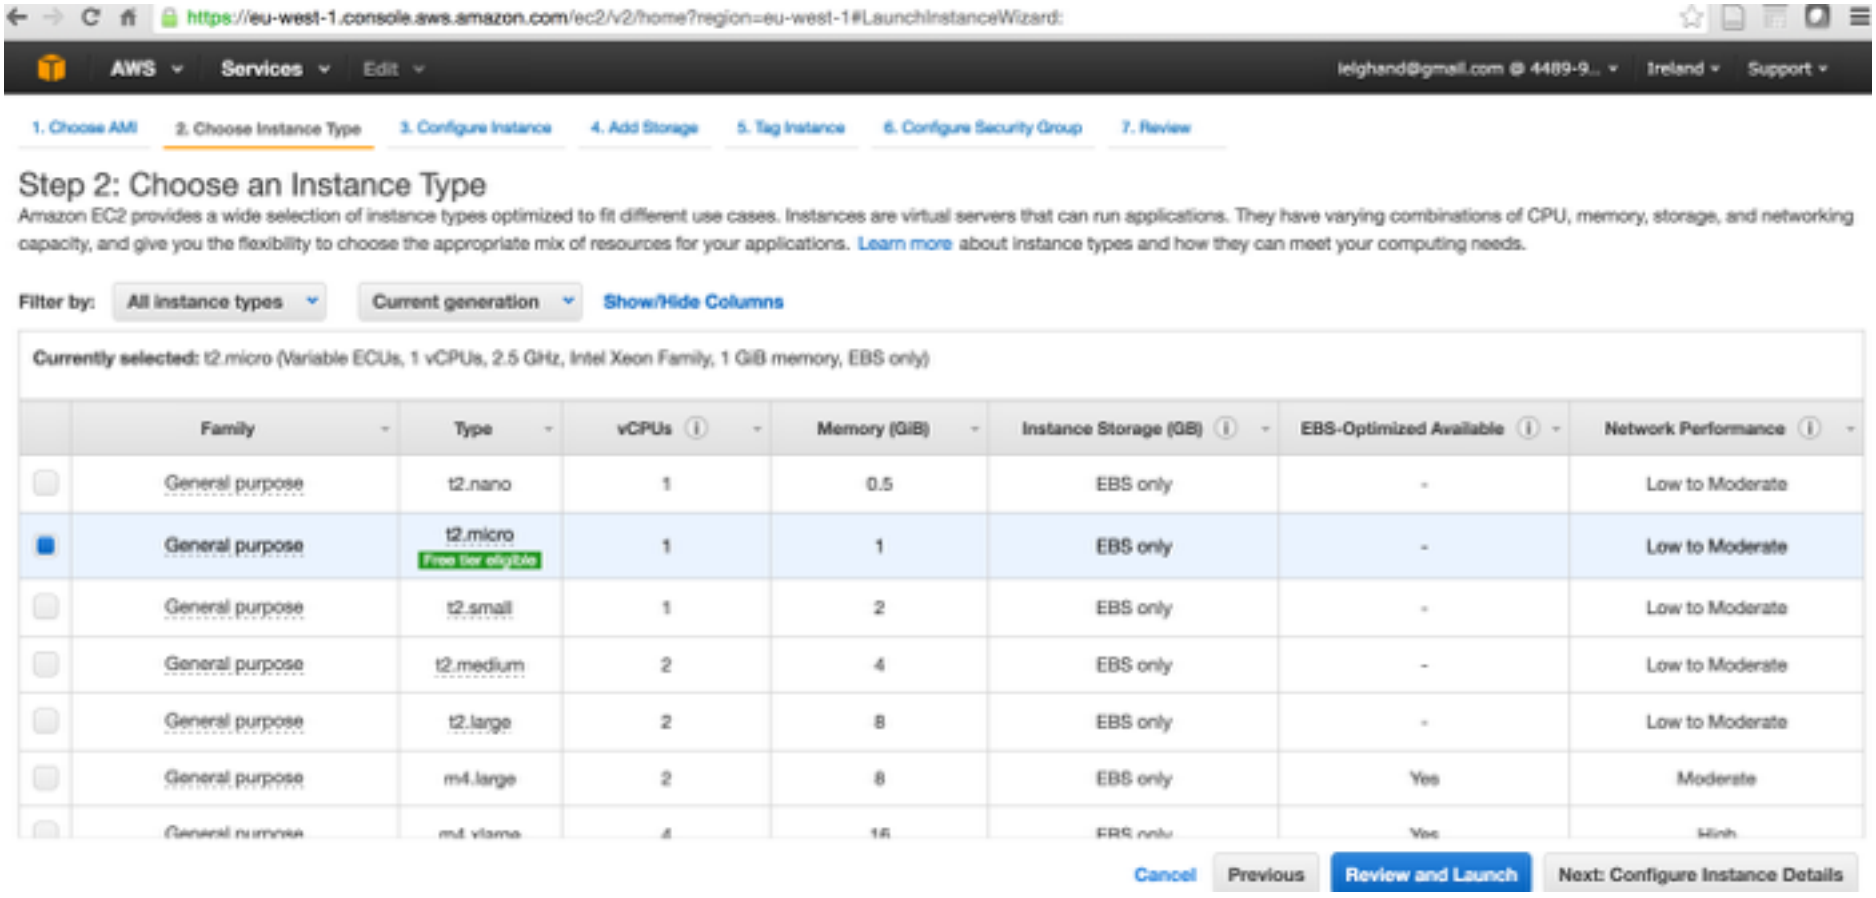

# Outline

- Introduction
- Cluster Computing, Grid Computing, and Utility Computing
- Cloud Computing Characteristics
- Cloud Service Models
- Cloud Deployment Models
- Cloud Infrastructure
- Challenges

# Virtualization

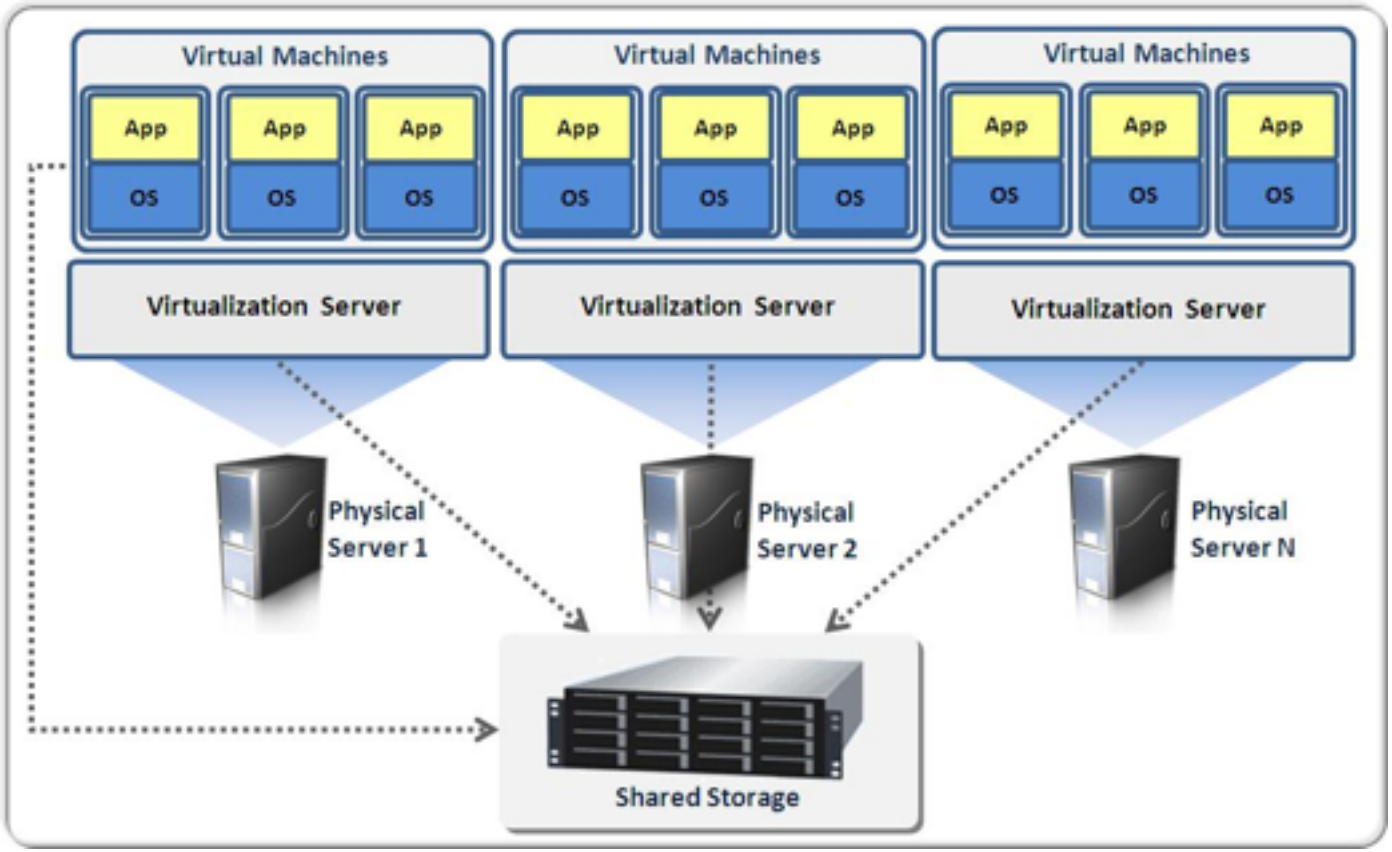


- It is a technique of masking or abstracting the physical compute hardware and enabling multiple operating systems (OSs) to run concurrently on a single or clustered physical machine(s)

# Advantages of Virtualization

- Runs multiple operating systems (OSs) per machine concurrently
- Makes OS and applications h/w independent
- Isolates VM from each other, hence no conflicts
- Improves resource utilization
- Offers flexible infrastructure at low cost

# Cloud Infrastructure Framework

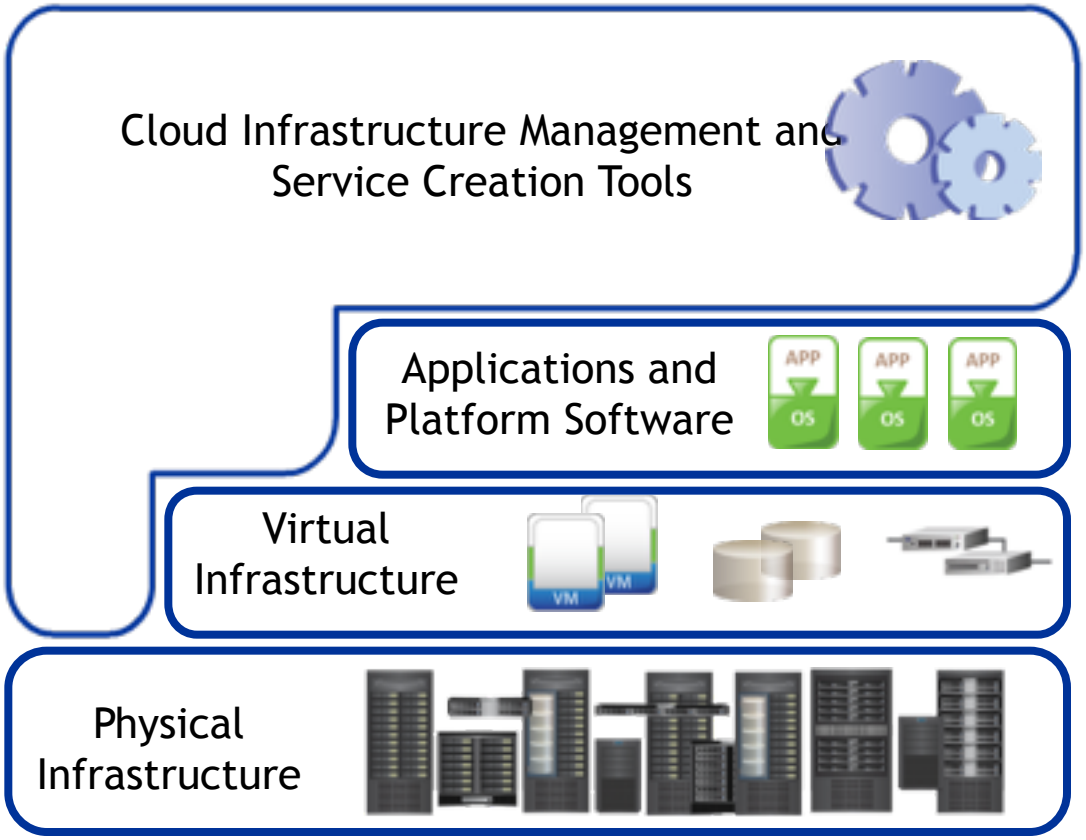


Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# Cloud Infrastructure and Management

- Virtualized infrastructure
- Components:
  - Cloud Services creation processes
  - Cloud Services management processes
- Manage physical and virtual infrastructures
- Handle service requests and provisions Cloud services
- Provide administrators a single management interface to manage resources across a virtual data center

Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# Outline

- Introduction
- Cluster Computing, Grid Computing, and Utility Computing
- Cloud Computing Characteristics
- Cloud Service Models
- Cloud Deployment Models
- Cloud Infrastructure
- Challenges

# Cloud Challenges – Consumer's Perspective

- Security and Regulation
  - Sensitive data?
  - Financial information, medical records
- Network latency
  - Real time applications may suffer due to network latency and limited bandwidth
- Supportability
  - Legacy or Custom applications may not be compatible with Cloud platform
- Interoperability
  - Lack of standardization across Cloud-based platforms

Slide credit: EMC Cloud Infrastructure and Services Student Guide.

# Thank You