

Cloud Computing



What is Cloud Computing?

- Cloud computing is a model for enabling *convenient, on-demand network access* to a *shared pool of configurable computing resources* (e.g., networks, servers, storage, applications, and services) [Mell_2009], [Berkely_2009].
- It can be *rapidly provisioned* and *released* with minimal management effort.
- It provides *high level abstraction* of computation and storage model.
- It has some essential **characteristics, service models, and deployment models.**



Essential Characteristics

- ***On-Demand Self Service:***
 - A consumer can unilaterally provision computing capabilities, automatically **without requiring human interaction with each service's provider.**
- ***Heterogeneous Access:***
 - Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous **thin** or **thick** client platforms.



Essential Characteristics (cont.)

- ***Resource Pooling:***
 - The provider's computing resources are pooled to serve multiple consumers using a ***multi-tenant model***.
 - Different physical and virtual resources dynamically assigned and reassigned according to consumer demand.
- ***Measured Service:***
 - Cloud systems ***automatically control*** and ***optimize*** resources used by leveraging a metering capability at some level of abstraction appropriate to the type of service.
 - ***It will provide analyzable and predictable computing platform.***



Service Models

- ***Cloud Software as a Service (SaaS):***
 - The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure.
 - The applications are accessible from various client devices such as a web browser (e.g., web-based email).
 - The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage,...
 - ***Examples: Caspio, Google Apps, Salesforce, Nivio, Learn.com.***



Service Models (cont.)


- ***Cloud Platform as a Service (PaaS):***
 - The capability provided to the consumer is to deploy onto the cloud infrastructure *consumer-created* or *acquired applications* created using *programming languages and tools* supported by the provider.
 - The consumer does not manage or control the underlying cloud infrastructure.
 - Consumer has control over the deployed applications and possibly application hosting environment configurations.
 - ***Examples: Windows Azure, Google App.***



Service Models (cont.)

- ***Cloud Infrastructure as a Service (IaaS):***
 - The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources.
 - The consumer is able to deploy and run arbitrary software, which can include operating systems and applications.
 - The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).
 - ***Examples: Amazon EC2, GoGrid, iland, Rackspace Cloud Servers, ReliaCloud.***

Service Models (cont.)



Cloud Computing Stack

Service Model at a glance: Picture From http://en.wikipedia.org/wiki/File:Cloud_Computing_Stack.svg



Deployment Models



Private Cloud: .

The cloud is operated **solely** for an organization. It may be managed by the organization or a third party and may exist on premise or off premise. —

Community Cloud: .

The cloud infrastructure is shared by several organizations and supports a specific community that has **shared concerns**. —

It may be managed by the organizations or a third party and may exist on premise or off premise —

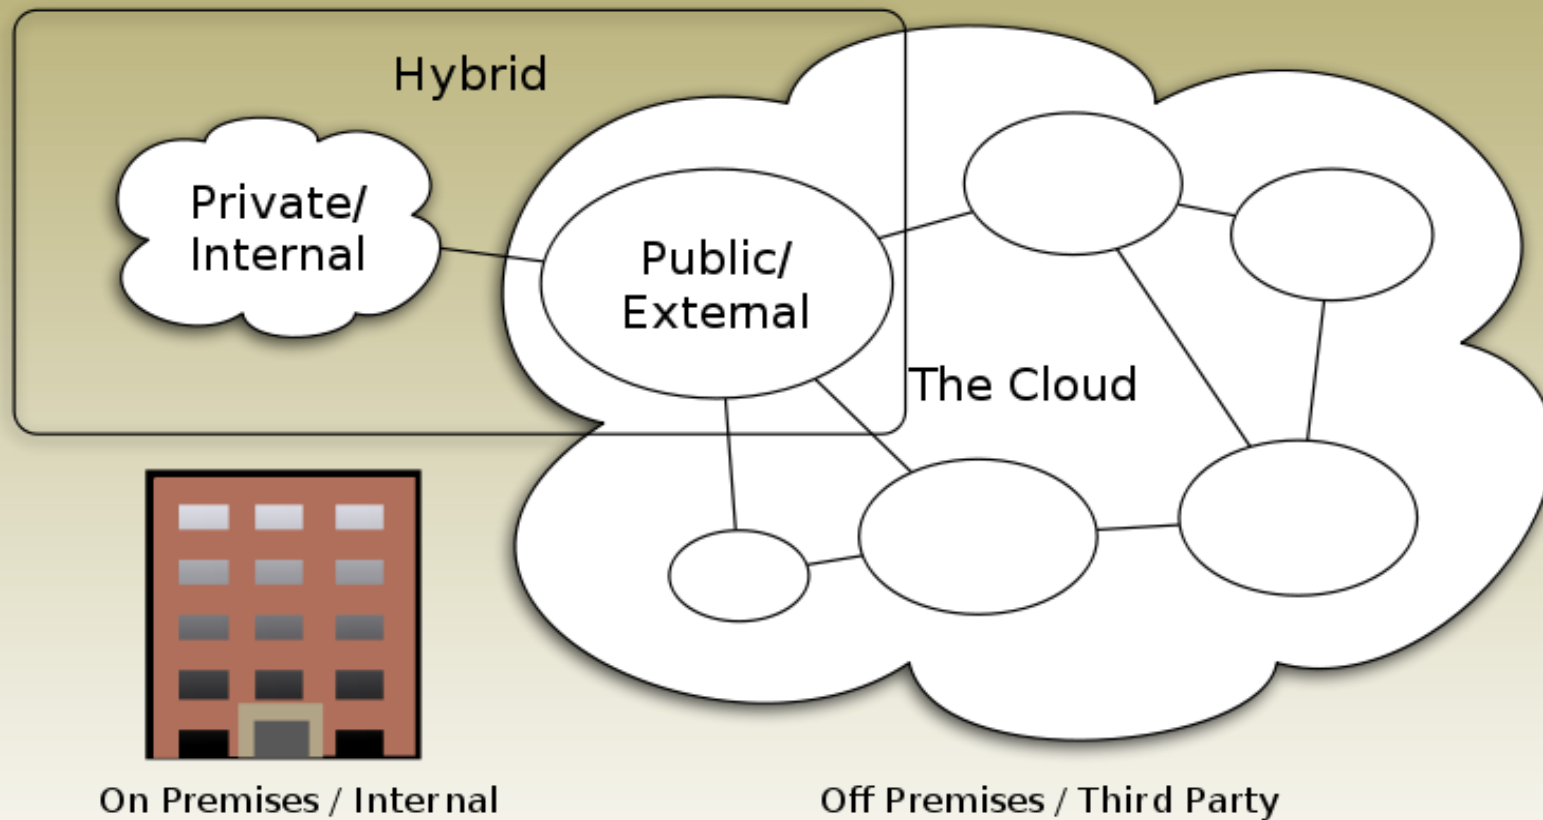


Public Cloud: .

The cloud infrastructure is made available to the general public or a large industry group and it is owned by an organization selling cloud services. —

Hybrid cloud: .

The cloud infrastructure is a composition of two or more clouds (private, community, or public). —



Cloud Computing Types

CC-BY-SA 3.0 by Sam Johnston

Advantages of Cloud Computing

- Cloud computing do not need high quality equipment for user, and it is very easy to use.
- Provides dependable and secure data storage center.
- Reduce run time and response time.
- Cloud is a large resource pool that you can buy on-demand service.
- Scale of cloud can extend dynamically providing nearly infinite possibility for users to use internet.



What is EC₂ ?

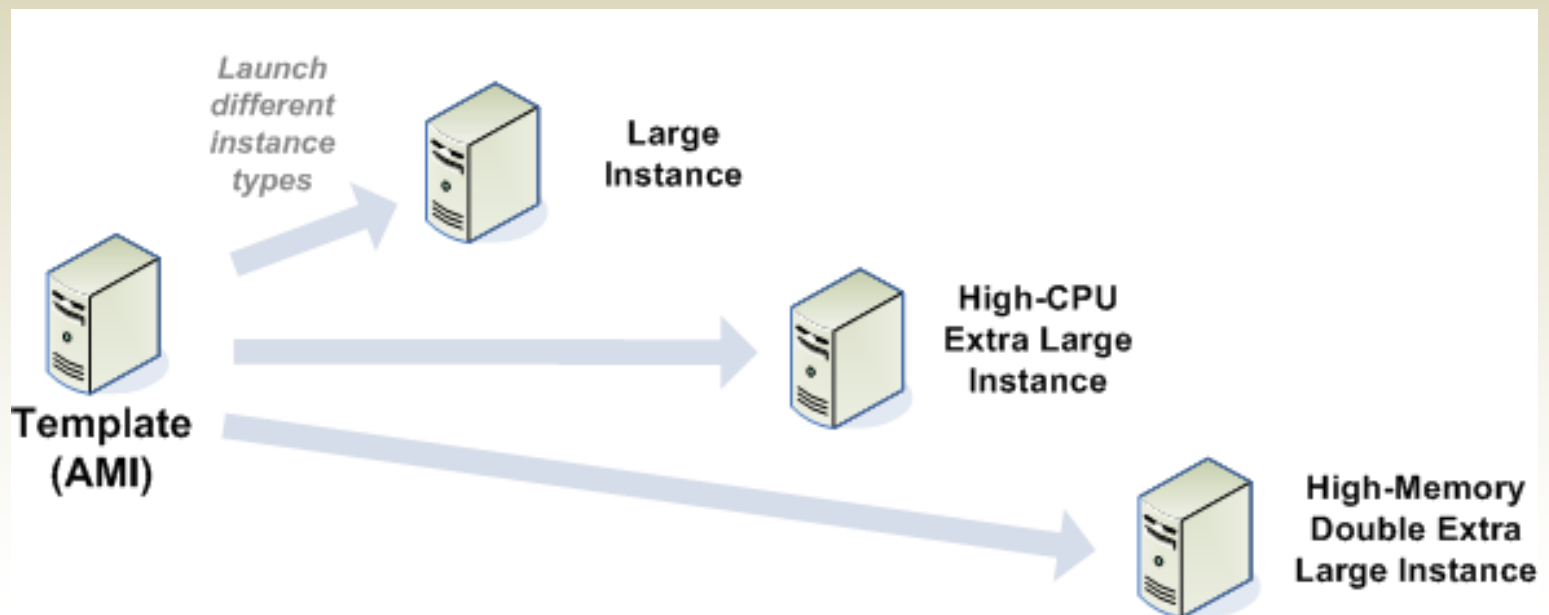
- Amazon Elastic Compute Cloud (EC₂) is a web service that provides resizable computing capacity that one uses to build and host different software systems.
- Designed to make web-scale computing easier for developers.
- A user can create, launch, and terminate server instances as needed, paying by the hour for active servers, hence the term "elastic".



EC2 Infrastructure Concepts

Amazon Machine Images - AMI

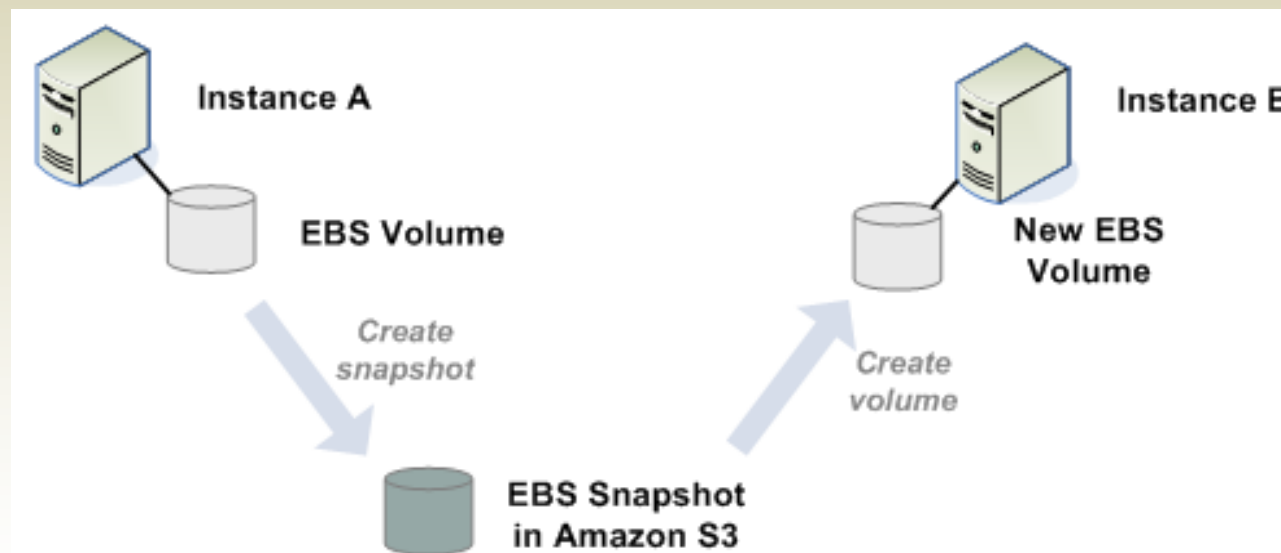
- Is an immutable representation of a set of disks that contain an operating system, user applications and/or data.
- From an AMI, one can launch multiple instances, which are running copies of the AMI.



Type	CPU	Memory	Local Storage	Platform	I/O	Name
Small	1 EC2 Compute Unit (1 virtual core with 1 EC2 Compute Unit)	1.7 GB	160 GB instance storage (150 GB plus 10 GB root partition)	32-bit	Moderate	m1.small
Large	4 EC2 Compute Units (2 virtual cores with 2 EC2 Compute Units each)	7.5 GB	850 GB instance storage (2 x 420 GB plus 10 GB root partition)	64-bit	High	m1.large
Extra Large	8 EC2 Compute Units (4 virtual cores with 2 EC2 Compute Units each)	15 GB	1690 GB instance storage (4 x 420 GB plus 10 GB root partition)	64-bit	High	m1.xlarge
Micro	Up to 2 EC2 Compute Units (for short periodic bursts)	613 MB	None (use Amazon EBS volumes for storage)	32-bit or 64-bit	Low	t1.micro
High-CPU Medium	5 EC2 Compute Units (2 virtual cores with 2.5 EC2 Compute Units each)	1.7 GB	350 GB instance storage (340 GB plus 10 GB root partition)	32-bit	Moderate	c1.medium

Elastic Block Store(EBS) volume

- An EBS volume is a read/write disk that can be created by an AMI and mounted by an instance.
- Volumes are suited for applications that require a database, a file system, or access to raw block-level storage.



Amazon S3

- S3 = Simple storage Service
- A SOA – Service Oriented Architecture which provides online storage using web services.
- Allows read, write and delete permissions on objects.
- Uses REST and SOAP protocols for messaging.

Amazon SimpleDB

- Amazon SimpleDB is a highly available, flexible, and scalable non-relational data store that offloads the work of database administration.
- Creates and manages multiple geographically distributed replicas of your data automatically to enable high availability and data durability.
- The service charges you only for the resources actually consumed in storing your data and serving your requests.

Networking

- Each instance launched is assigned two addresses a private address and a public IP address.
- A replacement instance has a different public IP address.
- Amazon EC2 offers Elastic IP addresses (static IP addresses) for dynamic cloud computing.

S3

EC2

VPC

Elastic MapReduce

CloudFront

RDS

SNS

Navigation

Region: US East

> EC2 Dashboard

INSTANCES

> Instances

> Spot Requests

IMAGES

> AMIs

> Bundle Tasks

ELASTIC BLOCK STORE

> Volumes

> Snapshots

NETWORKING & SECURITY

> Elastic IPs

> Security Groups

> Placement Groups

> Load Balancers

> Key Pairs

My Instances

Launch Instance

Instance Actions

Reserved Instances

Show/Hide

Refresh

Help

Viewing:

All Instances

All Instance Types

	Name	Instance	Type	Status	Lifecycle	Public DNS
<input checked="" type="checkbox"/>	Web Server	i-841948e9	m1.small	running	normal	ec2-67-202-15-66.compute-1.

1 EC2 Instance selected

EC2 Instance: i-841948e9

Description

Monitoring

Tags

AMI ID:	ami-08728661	Zone:	us-east-1b
Security Groups:	sg-22-open	Type:	m1.small
Status:	running	Owner:	043708602122
VPC ID:	-	Subnet ID:	-
Virtualization:	paravirtual	Placement Group:	
Reservation:	r-7de68517	RAM Disk ID:	-
Platform:	-	Key Pair Name:	GSG_Keypair
Kernel ID:	aki-407d9529	Monitoring:	basic
AMI Launch Index:	0	Elastic IP:	-
Root Device:	/dev/sda1	Root Device Type:	ebs

AWS Management Console

Region: <input type="text" value="US East (Virginia)"/>		
	Linux/UNIX Usage	Windows Usage
Standard On-Demand Instances		
Small (Default)	\$0.085 per hour	\$0.12 per hour
Large	\$0.34 per hour	\$0.48 per hour
Extra Large	\$0.68 per hour	\$0.96 per hour
Micro On-Demand Instances		
Micro	\$0.02 per hour	\$0.03 per hour
Hi-Memory On-Demand Instances		
Extra Large	\$0.50 per hour	\$0.62 per hour
Double Extra Large	\$1.00 per hour	\$1.24 per hour
Quadruple Extra Large	\$2.00 per hour	\$2.48 per hour
Hi-CPU On-Demand Instances		
Medium	\$0.17 per hour	\$0.29 per hour
Extra Large	\$0.68 per hour	\$1.16 per hour
Cluster Compute Instances		
Quadruple Extra Large	\$1.60 per hour	N/A*
Cluster GPU Instances		
Quadruple Extra Large	\$2.10 per hour	N/A*
* Windows® is not currently available for Cluster Compute or Cluster GPU Instances		

Pricing is per instance-hour consumed for each instance, from the time an instance is launched until it is terminated. Each partial instance-hour consumed will be billed as a full hour.

References

- Mobile cloud computing: Big Picture by M. Reza Rahimi
- <http://aws.amazon.com/ec2/>
- Amazon Elastic Compute Cloud – User Guide, API Version 2011-02-28.



Next: Amazon DynamoDB

A fully managed proprietary NoSQL database service that supports key-value and document data structures

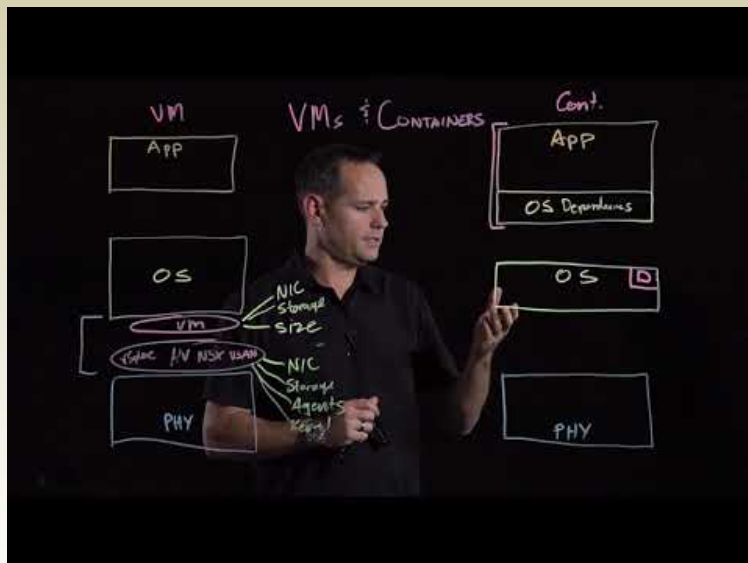


Next: Redis - a highly available in memory key value store.

Comparing DynamoDB and Redis

<https://db-engines.com/en/system/Amazon+DynamoDB%3BRedis>

Virtual Machines vs. Containers



Kubernetes - container orchestration

